

RESAMPLING STRATEGY IN SEQUENTIAL MONTE CARLO FOR CONSTRAINED SAMPLING PROBLEMS

Chencheng Cai*, Rong Chen and Ming Lin

Washington State University, Rutgers University and Xiamen University

Abstract: Monte Carlo sample paths of a dynamic system are useful for studying the underlying system and making statistical inferences related to the system. In many applications, the dynamic system being studied requires various types of constraints or observable features. In this study, we use a sequential Monte Carlo framework to investigate efficient methods for generating sample paths (with importance weights) from dynamic systems with rare and strong constraints. Specifically, we present a general formulation of the constrained sampling problem. Under such a formulation, we propose a flexible resampling strategy based on a potentially time-varying lookahead timescale, and identify the corresponding optimal resampling priority scores based on an ensemble of forward or backward pilots. Several examples illustrate the performance of the proposed methods.

Key words and phrases: Constrained sampling, pilot, priority score, resampling, sequential Monte Carlo.

1. Introduction

Stochastic dynamic systems are widely used in fields such as physics, finance, and engineering, among others. One of the important tools used to study a complex dynamic system is to obtain Monte Carlo sample paths of the underlying stochastic process. Such samples can be used for statistical inferences under the Monte Carlo framework, and provide us with a better understanding of the behavior of the system. The sequential Monte Carlo (SMC) method is a class of efficient sampling methods that use the sequential nature of the underlying dynamic process (Gordon, Salmond and Smith (1993); Kong, Liu and Wong (1994); Avitzour (1995); Liu and Chen (1995); Kitagawa (1996); Kim, Shephard and Chib (1998); Pitt and Shephard (1999); Chen, Wang and Liu (2000); Godsill, Doucet and West (2004); Doucet and Johansen (2011)). Although the SMC method is often used to estimate the marginal distribution of the underlying state at each time point (either filtering or smoothing), it also naturally provides sample paths (with importance weights) of the joint distribution of the entire state sequence. Here, we focus on the problem of efficiently generating such sample paths in an SMC framework.

*Corresponding author.

In practice, a stochastic system often comes with external observable information, including direct/indirect measurements and constraints. For example, many physics and financial applications are interested in the distribution of all possible paths of a diffusion process with fixed starting and ending points (a diffusion bridge) (Pedersen (1995); Durham and Gallant (2002); Lin, Chen and Mykland (2010)). In RNA and protein structure studies, the properties of self-loops (a single strand of RNA or protein that forms a loop in space by an internal chemical bond) are often studied using samples from the distribution of self-avoiding walks on a lattice with the same starting and ending points (Zhang et al. (2009); Lin, Chen and Liang (2008)). For example, computing a long-run marginal expected shortfall in financial risk management requires generating sample paths of a bivariate GARCH process, under a crisis constraint enforced on one of the processes (the one representing the market) to end below a threshold, say, a 40% loss (Acharya, Engle and Richardson (2012)).

Here, we provide a general formulation that includes many constrained path simulation problems. The formulation enables a discussion of general guidance for designing efficient SMC implementations for such problems. Because the standard SMC approach sometimes struggles with constrained systems, we propose a flexible resampling strategy and identify the optimal resampling priority scores. We develop two efficient approaches for estimating optimal priority scores, using forward pilots and backward pilots, respectively.

The remainder of this paper is organized as follows. In Section 2, we formally state the constrained sampling problem. Here, we also propose a general framework of an SMC with constraints (SMCc) method, with a flexible resampling strategy and optimal resampling priority scores. Section 3 presents two efficient pilot methods for estimating the optimal priority scores. Three examples with different types of constraints are used to demonstrate the performance of the proposed methods in Section 4. Section 5 concludes the paper.

2. SMC with Constraints

2.1. Stochastic dynamic system with constraints

The class of stochastic dynamic systems we consider contains a sequence of latent random states variables $x_{0:T} = \{x_0, x_1, \dots, x_T\}$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The dynamics of the system are governed by an initial state distribution $p(x_0)$ and a known forward-propagation distribution $p(x_t | x_{0:t-1})$, where $p(\cdot)$ denotes the density function/probability. In addition, external constraints are imposed on the latent states. Here, we consider the new constraints available at time t as an event I_t . Let $C_t = I_0 \cap I_1 \cap \dots \cap I_t$ be the cumulative constraint event up to time t . Then, $C_0 \supset C_1 \supset \dots \supset C_T$ forms a sequence of monotonically nonincreasing events. When there is no additional constraint at time t , we have $I_t = \Omega$ and $C_t = C_{t-1}$.

We focus on an effective SMC method under the importance sampling framework to obtain properly weighted samples of the entire path $x_{0:T}$, given the full constraint set C_T . The target posterior distribution of $x_{0:T}$ is

$$p(x_{0:T} | C_T) \propto p(x_{0:T}, C_T) = p(x_0, C_0) \prod_{t=1}^T p(x_t, C_t | x_{0:t-1}, C_{t-1}), \quad (2.1)$$

where $p(x_0, C_0)$ and $p(x_t, C_t | x_{0:t-1}, C_{t-1}) = p(x_t | x_{0:t-1}, C_{t-1})p(C_t | x_{0:t}, C_{t-1})$, for $t = 1, \dots, T$, are often specified by the system.

Many dynamic systems can be reformulated as a constrained sampling problem, including state-space models. We provide five examples in Appendix A, showing several classical problems, as well as several new classes of problems.

We can measure the *strength* of a constraint as the difference between the joint distributions of the latent states, with and without such a constraint. Specifically, we define the strength of I_t , the constraint imposed at time t , as the following χ^2 -divergence:

$$G(t) := \chi^2 \left\{ p(x_{0:t} | C_{t-1}) \parallel p(x_{0:t} | C_t) \right\} = \text{Var}_{p(x_{0:t} | C_{t-1})} \left\{ \frac{p(x_{0:t} | C_t)}{p(x_{0:t} | C_{t-1})} \right\}. \quad (2.2)$$

This is closely related to importance sampling, because it is the variance of the importance weight $w(x_{0:t}) = p(x_{0:t} | C_t) / p(x_{0:t} | C_{t-1})$ when we generate samples from the proposal distribution $p(x_{0:t} | C_{t-1})$ to make inferences with respect to the target distribution $p(x_{0:t} | C_t)$. A “strong” constraint would alter the distribution of the underlying states significantly.

2.2. Intermediate distributions under constraints

Here, we generate a *properly weighted* sample set $\{(x_{0:T}, w_T^{(i)})\}_{i=1, \dots, n}$, with respect to $p(x_{0:T} | C_T)$ under a general SMC framework, that satisfies

$$\frac{\sum_{i=1}^n w_T^{(i)} h(x_{0:T}^{(i)})}{\sum_{i=1}^n w_T^{(i)}} \xrightarrow{a.s.} \mathbb{E}_{p(x_{0:T} | C_T)} [h(x_{0:T})], \quad (2.3)$$

as $n \rightarrow \infty$, for any function $h(\cdot)$ with finite expectation under $p(x_{0:T} | C_T)$.

The SMC method has been studied extensively in the literature, for example, Liu and Chen (1998) and Doucet and Johansen (2011), and the references therein. We use the following notation for clarity. Consider a sequence of forward intermediate target propagation distributions with densities $\pi_0(x_0), \pi_1(x_{0:1}), \dots, \pi_T(x_{0:T})$. The SMC approach proposes generating samples $x_{0:T}^{(i)} = (x_0^{(i)}, x_1^{(i)}, \dots, x_T^{(i)})$, for $i = 1, \dots, n$, sequentially from a series of proposal conditional distributions $q(x_t | x_{0:t-1})$, for $t = 0, 1, \dots$, and then updating the corresponding *importance weights* using

$$w_t^{(i)} = \frac{\pi_t(x_{0:t}^{(i)})}{q(x_0^{(i)}) \prod_{s=1}^t q(x_s^{(i)} | x_{0:s-1}^{(i)})} = w_{t-1}^{(i)} \frac{\pi_t(x_{0:t}^{(i)})}{\pi_{t-1}(x_{0:t-1}^{(i)}) q(x_t^{(i)} | x_{0:t-1}^{(i)})},$$

with $w_0^{(i)} = \pi_0(x_0^{(i)})/q(x_0^{(i)})$. The distribution $q(x_0) \prod_{s=1}^T q(x_s | x_{0:s-1})$, from which the samples $x_{0:T}^{(i)}$, for $i = 1, \dots, n$, are generated, is called the *sampling distribution* or *proposal distribution*. At each time t , if the intermediate target distribution $\pi_t(x_{0:t})$ is absolutely continuous with respect to the proposal distribution $q(x_0) \prod_{s=1}^t q(x_s | x_{0:s-1})$, the sample set $\{(x_{0:t}^{(i)}, w_t^{(i)})\}_{i=1, \dots, n}$ is properly weighted with respect to $\pi_t(x_{0:t})$.

If we set $\pi_T(x_{0:T})$ to be the joint posterior distribution $p(x_{0:T} | C_T)$ in (2.1), then at the end, we can use the weighted sample set $\{(x_{0:T}^{(i)}, w_T^{(i)})\}_{i=1, \dots, n}$ for statistical inferences of $p(x_{0:T} | C_T)$, as in (2.3). We consider the following three distributions as potential intermediate target distributions $\pi_t(x_{0:t})$: for $t = 0, 1, \dots, T$,

$$\bar{p}_t(x_{0:t}) := p(x_{0:t} | C_T), \quad (2.4)$$

$$p_t^+(x_{0:t}) := p(x_{0:t} | C_{t_+}), \quad (2.5)$$

$$\tilde{p}_t(x_{0:t}) := p(x_{0:t} | C_t), \quad (2.6)$$

where $t_+ \geq t$ is the next time a “strong” constraint is imposed after time t .

The marginal posterior distribution $\bar{p}_t(x_{0:t})$ in (2.4) is naturally induced by the joint posterior distribution $p(x_{0:T} | C_T)$, but is usually infeasible to use as the intermediate target distribution $\pi_t(x_{0:t})$, because it involves the high-dimensional integral

$$p(x_{0:t} | C_T) \propto \int \cdots \int \prod_{s=1}^T p(x_s, C_s | x_{0:s-1}, C_{s-1}) dx_{t+1} \cdots dx_T. \quad (2.7)$$

Conventional SMC approaches (Gordon, Salmond and Smith (1993); Liu and Chen (1998)) use the forward propagation distributions $\tilde{p}_t(x_{0:t})$ in (2.6), using only information up to time t . Often, this is not efficient, especially when future constraints have a strong effect on the distribution of the current state. To overcome this drawback, we propose using $p_t^+(x_{0:t})$ in (2.5) as $\pi_t(x_{0:t})$, which uses part of the future events and constraints to correct the Monte Carlo samples proactively. Recall that $t_+ \geq t$ is the next time when a strong constraint is imposed after time t . This is essentially a partial lookahead scheme, with a variable lookahead timescale.

The distribution $p_t^+(x_{0:t}) = p(x_{0:t} | C_{t_+})$ in (2.5) is a compromise between $\bar{p}_t(x_{0:t}) = p(x_{0:t} | C_T)$ in (2.4) and $\tilde{p}_t(x_{0:t}) = p(x_{0:t} | C_t)$ in (2.6), where the former considers the whole constraint set, and the latter ignores all future constraints. At time t , the distribution $p_t^+(x_{0:t})$ seeks the next available strong constraint for guidance, but not the entire future constraint set, as $\bar{p}_t(x_{0:t})$ would require.

In most cases, the next available strong constraint plays an important role in shaping the path distribution. The compromised distribution $p_t^+(x_{0:t})$ balances the use of future constraints and computational efficiency.

To use $p_t^+(x_{1:t})$ in (2.5) as the intermediate target distribution $\pi_t(x_{1:t})$, the “optimal” proposal distribution is $q(x_t|x_{0:t-1}) = p_t^+(x_t|x_{0:t-1}) = p(x_t|x_{0:t-1}, C_{t+})$, which incorporates the cumulative constraints up to time t_+ , as suggested by Kong, Liu and Wong (1994) and Liu and Chen (1998). However, this is often difficult, especially when t_+ is far away from t , because it involves a high-dimensional integral similar to that in (2.7). On the other hand, $\tilde{p}_t(\cdot)$ in (2.6) is often easy to work with, using the proposal distributions equal or close to $\tilde{p}_t(x_t|x_{0:t-1}) = p(x_t|x_{0:t-1}, C_t)$. Specifically, we propose generating samples under the distribution $\tilde{p}_t(\cdot)$, but using a resampling step so that the resulting samples (before weighting) follow $p_t^+(\cdot)$.

2.3. Optimal resampling scores under constraints

An important component of an SMC implementation is the resampling step, in which each sample $x_{0:t}^{(i)}$ is assigned a priority score $\beta_t^{(i)} > 0$ that reflects the algorithm’s preference on this sample. Then a new set of samples $\{x_{0:t}^{*(i)}\}_{i=1,\dots,n}$ is drawn from the current set of samples $\{x_{0:t}^{(i)}\}_{i=1,\dots,n}$, with replacement, with probabilities proportional to the priority scores. Next, the sample weights are adjusted to $w_t^{*(i)} = w_t^{(i)} / \beta_t^{(i)}$. The resampling step tends to remove samples with low priority scores, and duplicate those with high priority scores. The resulting weighted sample set $\{(x_{0:t}^{*(i)}, w_t^{*(i)})\}_{i=1,\dots,n}$ remains to be properly weighted with respect to $\pi_t(x_{0:t})$ (Douc, Moulines and Stoffer (2014) Chap. 10.3; Tsay and Chen (2019), Lemma 8.7).

In conventional SMC approaches, the resampling priority scores are often chosen as the sample weights, that is, $\beta_t^{(i)} = w_t^{(i)}$, such that the samples have equal weights after resampling, and hence can be viewed as samples of the intermediate target distribution $\pi_t(x_{0:t})$ (Gordon, Salmond and Smith (1993); Kong, Liu and Wong (1994); Liu and Chen (1998)). This is a natural choice for filtering problems. However, it is not necessarily a good choice if we focus on the final target distribution $\pi_T(x_{0:T})$. There is a great flexibility in the selection of the priority scores, as long as it does not assign a zero score to any of the samples. With a careful construction, it may improve the sampling efficiency significantly (Pitt and Shephard (1999); Zhang et al. (2007)). The design is the key to developing an efficient SMC for the constrained systems presented here.

In the constraint problem with intermediate target distribution $p_t^+(x_{0:t})$, we have that

$$p_t^+(x_{0:t}) \propto \tilde{p}_t(x_{0:t})p(C_{t+}|x_{0:t}, C_t).$$

A convenient way of drawing samples from the distribution $p_t^+(\cdot)$ uses this fact to conduct a resampling step of the samples generated from $\tilde{p}_t(\cdot)$. Specifically,

Algorithm 1: Sequential Monte Carlo with Constraints (SMCc).

-
- At times $t = 0, 1, \dots, T$:
 - Propagation: For $i = 1, \dots, n$,
 - * Draw $x_t^{(i)}$ from distribution $q(x_t | x_{0:t-1}^{(i)})$ and let $x_{0:t}^{(i)} = (x_{0:t-1}^{(i)}, x_t^{(i)})$.
 - * Update weights by setting

$$w_t^{(i)} \leftarrow w_{t-1}^{(i)} \cdot \frac{p(x_{0:t}^{(i)} | C_t)}{p(x_{0:t-1}^{(i)} | C_{t-1})q(x_t^{(i)} | x_{0:t-1}^{(i)})}$$

$$\propto w_{t-1}^{(i)} \cdot \frac{p(x_t^{(i)}, C_t | x_{0:t-1}^{(i)}, C_{t-1})}{q(x_t^{(i)} | x_{0:t-1}^{(i)})}.$$
 - Set priority scores $\beta_t^{(i)} = w_t^{(i)} \hat{p}(C_{t+} | x_{0:t}^{(i)}, C_t)$, $i = 1, \dots, n$.
 - Resampling:
 - * Draw samples $\{J_1, \dots, J_n\}$ from the set $\{1, \dots, n\}$ with replacement, with probabilities proportional to $\{\beta_t^{(i)}\}_{i=1, \dots, n}$.
 - * Let $x_{0:t}^{*(i)} = x_{0:t}^{(J_i)}$ and $w_t^{*(i)} = w_t^{(J_i)} / \beta_t^{(J_i)}$ for $i = 1, \dots, n$.
 - * Return the new set $\{(x_{0:t}^{(i)}, w_t^{(i)})\}_{i=1, \dots, n} \leftarrow \{(x_{0:t}^{*(i)}, w_t^{*(i)})\}_{i=1, \dots, n}$.
 - Return the weighted sample set $\{(x_{0:T}^{(i)}, w_T^{(i)})\}_{i=1, \dots, n}$.
-

we propose tracking the exact weight $w_t^{(i)}$ under the distribution $\tilde{p}_t(\cdot)$, but use a resampling step with the priority score

$$\beta_t^{(i)} := w_t^{(i)} \frac{p_t^+(x_{0:t}^{(i)})}{\tilde{p}_t(x_{0:t}^{(i)})} \propto w_t^{(i)} p(C_{t+} | x_{0:t}^{(i)}, C_t), \quad (2.8)$$

so that the resulting samples $\{x_{0:t}^{*(i)}\}_{i=1, \dots, n}$ follow $p_t^+(\cdot)$. We call $\beta_t^{(i)} = w_t^{(i)} p(C_{t+} | x_{0:t}^{(i)}, C_t)$ the *optimal priority score* with respect to the constraint set C_{t+} .

We use the resampling approach to incorporate information on future constraints, because it is easy to conduct, and refer to this method as the SMCc method. The exact value of $p(C_{t+} | x_{0:t}^{(i)}, C_t)$ is often difficult to obtain. We can use an approximated value, $\hat{p}(C_{t+} | x_{0:t}^{(i)}, C_t)$, to construct $\beta_t^{(i)}$ in (2.8). The method is presented in Algorithm 1. We discuss how to approximate $p(C_{t+} | x_{0:t}^{(i)}, C_t)$ in Section 3.

Note that, in the SMCc algorithm in Algorithm 1, the sample set $\{(x_{0:t}^{(i)}, w_t^{(i)})\}_{i=1, \dots, n}$ obtained at each time $t < T$ is properly weighted with respect to $\tilde{p}_t(x_{0:t}) = p(x_{0:t} | C_t)$, not $p_t^+(x_{0:t})$. However, we are only interested in the entire sample path $x_{0:T}$, which follows the desired distribution $\tilde{p}_T(x_{0:T}) = p_t^+(x_{0:T}) = p(x_{0:T} | C_T)$ at time T .

The simple random sampling with replacement in Algorithm 1 is not the most efficient method for resampling. Better resampling schemes, such as the residual resampling method (Liu and Chen (1998)) and the systematic resampling method (Carpenter, Clifford and Fearnhead (1999)), reduce the variation introduced by the resampling step.

2.4. Discussion

Determination of t_+ : In practice, the selection of t_+ depends on specific problems and can be defined by the user. For example, when the constraints are rare and strong, as in the diffusion bridge sampling problem (Lin, Chen and Mykland (2010)), one may use $t_+ = T$, the end constraint. When the constraints are frequent and weak, as in the state-space model, where the observation y_t serves as a weak constraint, we may use $t_+ = \min\{t + d, T\}$, taking into account the information of the next d observations $y_{t+1}, \dots, y_{t+d}$, as in Lin, Chen and Liu (2013). This is a fixed lookahead approach. We can also use the constraint strength measure $G(t)$ defined in (2.2) as a general guide for determining t_+ . For example, we may set $t_+ = \min_{s>t}\{G(s) > c\}$, for some threshold value c . In practice, the exact value of $G(t)$ may not be easy to compute, but can be estimated using a trial run of an SMC with a small sample size, because $G(t)$ can be estimated using the variance of the importance weights of the target distribution $p(x_{0:t} | C_t)$ to the proposal distribution $p(x_{0:t} | C_{t-1})$. Specifically, using $q(x_t | x_{0:t-1}^{(i)}) = p(x_t | x_{0:t-1}^{(i)}, C_{t-1})$ as the proposal distribution, and resampling with the importance weights as the priority score at every step, $G(t)$ can be estimated by $\hat{G}(t) = \widehat{\text{var}}(w_t) / \bar{w}_t^2$, where $\bar{w}_t$ and $\widehat{\text{var}}(w_t)$ are the sample mean and sample variance, respectively, of the weights $\{w_t^{(i)}\}_{i=1, \dots, n}$. Then, the time t is identified as a time point with a strong constraint when $\hat{G}(t)$ exceeds a certain threshold value or, equivalently, when the effective sample size, defined as $n / [1 + \hat{G}(t)]$ in the SMC literature (Kong, Liu and Wong (1994)), is less than a certain value. The procedure only needs to be carried out once to identify all strong constraints.

Relation to other particle filters: The idea of using future constraints to estimate the current state under an SMC framework has been studied in Pitt and Shephard (1999), Doucet, Briers and Sénécal (2006), Lin, Chen and Liu (2013), Whiteley and Lee (2014), and others. The auxiliary particle filter (Pitt and Shephard (1999)) suggests first conducting resampling according to the priority score $\beta_t = w_t p(C_{t+\Delta} | x_{0:t}, C_t)$ for a certain number of lookahead steps $\Delta > 0$ at time t (usually $\Delta = 1$), then drawing samples of x_{t+1} from $q(x_{t+1} | x_{0:t}) = p(x_{t+1} | x_{0:t}, C_{t+\Delta})$. The twisted particle filter in Whiteley and Lee (2014) introduces a special sample to incorporate information on future constraints in the SMC implementation. Whiteley and Lee (2014) show the theoretical properties of using such a procedure. The block sampling method

in Doucet, Briers and Sénécal (2006) also uses future constraints to update x_t . These methods all need to evaluate $p(C_{t+\Delta} | x_{0:t}, C_t)$. When $p(C_{t+\Delta} | x_{0:t}, C_t)$ does not have a closed form, the extended Kalman filter is often used to find an approximation (Jazwinski (1970)), but this incurs a high computational cost and can be poor when Δ is large. In Section 3, we develop computationally efficient methods to approximate $p(C_{t_+} | x_{0:t}, C_t)$ in the optimal priority score $\beta_t = w_t p(C_{t_+} | x_{0:t}, C_t)$.

3. Approximation of the Optimal Priority Score Using Pilots

We evaluate the term $p(C_{t_+} | x_{0:t}, C_t)$ in the optimal priority score (2.8), which is

$$p(C_{t_+} | x_{0:t}, C_t) = \int \cdots \int \prod_{s=t+1}^{t_+} p(x_s, C_s | x_{0:s-1}, C_{s-1}) dx_{t+1} \cdots dx_{t_+}. \quad (3.1)$$

The integrand is often well defined by the system, but the integral does not have a closed-form solution, in most cases.

We can sometimes assume a parametric form for $p(C_{t_+} | x_{0:t}, C_t)$, based on some prior knowledge. Zhang et al. (2007) and Lin et al. (2008) use the SMCc approach with $t_+ = T$ to generate protein conformation samples that satisfy certain distance constraints between the molecules of the conformation. The parametric functions they use to approximate $p(C_T | x_{0:t}, C_t)$ are based on the distance between the current location x_t and the final destination x_T . The particle efficient importance sampling (PEIS) method of Scharth and Kohn (2016) uses $t_+ = T$, and approximates the optimal conditional proposal distributions $p(x_t | x_{0:t-1}, C_T)$ and resampling priority scores $\beta_t = w_t p(C_T | x_{0:t}, C_t)$ within some parametric families. The method tunes the parameters in the functions using an iterative local optimization routine. However, the performance of these parametric approximation methods depends greatly on the choice of the parametric family.

Here, we propose two *ensemble* pilot approaches for estimating $p(C_{t_+} | x_{0:t}^{(j)}, C_t)$ nonparametrically. Compared with existing “individual” pilot approaches in the existing literature (Wang, Chen and Guo (2002); Zhang and Liu (2002); Lin, Chen and Liu (2013)), the proposed “ensemble” pilot approaches pool all pilot samples together, and use a nonparametric smoothing technique to improve the estimation accuracy and reduce the computational cost.

For ease of presentation, we consider approximating $p(C_{t_+} | x_{0:t}^{(j)}, C_t)$, for all t between two predetermined time points $t_1 < t_2$, both with strong constraints. Under this setting, we have $t_+ = t_2$, for every $t_1 < t \leq t_2$.

3.1. Approximation based on forward pilots

Suppose there exists a low-dimensional statistic $S(x_{0:t})$ that summarizes $x_{0:t}$ such that

$$p(x_{t+1:t+d}, C_{t+d} | x_{0:t}, C_t) = p(x_{t+1:t+d}, C_{t+d} | S(x_{0:t}), C_t), \quad (3.2)$$

for all t and $d = 0, 1, \dots$. Under this assumption, $p(C_{t+} | x_{0:t}, C_t) = p(C_{t+} | S(x_{0:t}), C_t)$ is a function of $S(x_{0:t})$. We further assume that there exists a function $\phi(\cdot)$ such that $S(x_{0:t+1}) = \phi(S(x_{0:t}), x_{t+1})$. When the system is Markovian, then $S(x_{0:t}) = x_t$.

In the forward pilot approach, we apply the SMC to generate pilot samples sequentially from time t_1 to t_2 using a proposal distribution that encourages the pilot samples to satisfy the constraints C_{t_2} . The pilots and their corresponding weights contain information on C_{t_2} , and can be used to estimate $p(C_{t+} | x_{0:t}^{(j)}, C_t)$, for $t_1 < t < t_2 = t_+$. The detailed implementation is presented in Algorithm 2. Note that for $U_t^{(j)} = \prod_{s=t+1}^{t_2} \tilde{u}_s^{(j)}$ defined in Algorithm 2, we have

$$\mathbb{E}(U_t^{(j)} | S_t^{(j)} = S) = p(C_{t+} | S(x_{0:t}) = S, C_t),$$

for all $t_1 < t < t_2$. Therefore, we can use $\{(U_t^{(j)}, S_t^{(j)})\}_{j=1, \dots, m}$ to estimate $p(C_{t+} | S(x_{0:t}), C_t)$ using the nonparametric histogram function (3.3) in Algorithm 2. We choose not to use the kernel smoothing method here to control the computational cost, because $\hat{p}(C_{t+} | S(x_{0:t}), C_t)$ needs to be evaluated for all $x_{0:t}^{(j)}$, for $j = 1, \dots, n$, and at each time t .

The pilot sample idea has been proposed by Wang, Chen and Guo (2002) and Zhang and Liu (2002), and is used for delayed estimation in Lin, Chen and Liu (2013). They consider an individual pilot algorithm in the sense that the pilot samples are generated for each sample path $x_{0:t}^{(i)}$ in Algorithm 1, and are only used to compute $\beta_t^{(i)}$ for that particular sample path. The computational cost of such an individual pilot algorithm is very high, because it requires generating pilot samples for every $x_{0:t}^{(i)}$, for $i = 1, \dots, n$, at each time t .

Instead, in the ensemble pilot approaches, we use all pilot samples $\tilde{x}_{0:t}^{(j)}$, for $j = 1, \dots, m$, with $\tilde{S}_t^{(j)} = S(\tilde{x}_{0:t}^{(j)})$ close to $S(x_{0:t}^{(i)})$, to estimate $p(C_{t+} | x_{0:t}^{(i)}, C_t)$. Because this is a global operation, we can afford to use a large number of pilots, and thus obtain accurate estimates of $\hat{p}(C_{t+} | x_{0:t}^{(i)}, C_t) = \hat{p}(C_{t+} | S(x_{0:t}^{(i)}), C_t)$. Additionally, this algorithm needs to be conducted only once to obtain $\hat{p}(C_{t+} | S(x_{0:t}), C_t)$, for all $t_1 < t < t_2$, yielding a significant reduction in the computational cost.

The accuracy of $\hat{p}(C_{t+} | S(x_{0:t}), C_t)$ depends on the choice of the proposal distribution $\varphi(x_s | S(x_{0:s-1}))$ in Algorithm 2 used to generate the pilots. Because there is a strong constraint I_{t_2} at time t_2 , it is important to incorporate I_{t_2} in the proposal distribution $\varphi(\cdot)$ when generating pilot samples from t_1 to t_2 . This ensures that the pilot samples have a reasonably large probability of satisfying the constraint at time t_+ .

Algorithm 2: Forward Pilot Algorithm.

- Initialization: For $j = 1, \dots, m$, draw samples $\tilde{S}_{t_1}^{(j)}$ from a proposal distribution $\varphi(S)$ that covers the support of $S(x_{0:t_1})$.
- For $t = t_1 + 1, \dots, t_2$, draw pilot samples forwardly as follows.
 - Generate samples $\tilde{x}_t^{(j)}$ from a proposal distribution $\varphi(\tilde{x}_t | \tilde{S}_{t-1}^{(j)})$, and calculate $\tilde{S}_t^{(j)} = \phi(\tilde{S}_{t-1}^{(j)}, \tilde{x}_t^{(j)})$ for $j = 1, \dots, m$.
 - Calculate the incremental weights

$$\tilde{u}_t^{(j)} = \frac{p(\tilde{x}_t^{(j)}, C_t | S(\tilde{x}_{0:t-1}) = \tilde{S}_{t-1}^{(j)}, C_{t-1})}{\varphi(\tilde{x}_t^{(j)} | \tilde{S}_{t-1}^{(j)})}, \quad j = 1, \dots, m.$$

- For $t = t_2 - 1, t_2 - 2, \dots, t_1 + 1$:
 - Compute $U_t^{(j)} = \prod_{s=t+1}^{t_2} \tilde{u}_s^{(j)}$ for $j = 1, \dots, m$.
 - Let $\mathcal{S}_1 \cup \dots \cup \mathcal{S}_D$ be a partition of the support of $S(x_{0:t})$. Estimate $p(C_{t+} | x_{0:t}, C_t) = p(C_{t+} | S(x_{0:t}), C_t)$ by

$$f_t(S(x_{0:t})) = \sum_{d=1}^D \gamma_{t,d} \mathbb{I}(S(x_{0:t}) \in \mathcal{S}_d) \quad (3.3)$$

with $\gamma_{t,d} = \sum_{j=1}^m U_t^{(j)} \mathbb{I}(\tilde{S}_t^{(j)} \in \mathcal{S}_d) / \sum_{j=1}^m \mathbb{I}(\tilde{S}_t^{(j)} \in \mathcal{S}_d)$, where $\mathbb{I}(\cdot)$ is the indicator function.

- Return the estimated functions $\{\hat{p}(C_{t+} | x_{0:t}, C_t) = f_t(S(x_{0:t}))\}_{t=t_1+1, \dots, t_2-1}$ to compute the priority scores for Algorithm 1.

3.2. Approximation based on backward pilots

Assume that the stochastic dynamic system is Markovian, that is,

$$p(x_t, I_t | x_{0:t-1}, C_{t-1}) = p(x_t, I_t | x_{t-1}),$$

for all t . Then, $p(C_{t+} | x_{0:t}, C_t) = p(I_{t+1:t+} | x_t)$ is a function of x_t , and does not depend on the past constraints before time t . Here, $I_{t+1:t+}$ denotes the cumulative constraints imposed between time $t + 1$ and t_+ (inclusive). For such a Markovian system, we can extend the backward pilot sampling method proposed in Lin, Chen and Mykland (2010) to a more general SMCc setting. In this approach, the pilot samples are generated in the reverse time direction, starting from $t = t_+$, the time point with a strong constraint, and propagating backward. The method is presented in Algorithm 3.

In Algorithm 3, the weight for the backward pilot $\tilde{x}_{t:t_+}$ is

$$\tilde{w}_t = \frac{p(\tilde{x}_{t+1:t_+}, I_{t+1:t_+} | \tilde{x}_t)}{r(\tilde{x}_{t:t_+})},$$

where $r(\tilde{x}_{t:t_+})$ is the proposal distribution used to generate the backward pilots. Taking the expectation conditional on $\tilde{x}_t$, we have

$$\begin{aligned} \mathbb{E}(\tilde{w}_t | \tilde{x}_t) &= \int \cdots \int \frac{p(\tilde{x}_{t+1:t_+}, I_{t+1:t_+} | \tilde{x}_t)}{r(\tilde{x}_t, \tilde{x}_{t+1:t_+})} r(\tilde{x}_{t+1:t_+} | \tilde{x}_t) d\tilde{x}_{t+1:t_+} \\ &= \frac{p(I_{t+1:t_+} | \tilde{x}_t)}{r(\tilde{x}_t)}, \end{aligned}$$

where $r(\tilde{x}_{t+1:t_+} | \tilde{x}_t)$ and $r(\tilde{x}_t)$ are the conditional distribution and the marginal distribution, respectively, induced from $r(\tilde{x}_{t:t_+})$. Therefore,

$$p(C_{t_+} | \tilde{x}_{0:t}, C_t) = p(I_{t+1:t_+} | \tilde{x}_t) = r(\tilde{x}_t) \mathbb{E}(\tilde{w}_t | \tilde{x}_t).$$

Again, we use the pilot samples $\{(\tilde{x}_t^{(j)}, \tilde{w}_t^{(j)})\}_{j=1, \dots, m}$ to estimate $r(\tilde{x}_t)$ and $\mathbb{E}(\tilde{w}_t | \tilde{x}_t)$ by nonparametric smoothing. A histogram estimator is

$$\begin{aligned} &\hat{p}(I_{t+1:t_+} | x_t) \\ &= \hat{r}(x_t) \hat{\mathbb{E}}(\tilde{w}_t | x_t) \\ &= \sum_{d=1}^D \frac{\sum_{j=1}^m \mathbb{I}(\tilde{x}_t^{(j)} \in \mathcal{X}_d)}{m|\mathcal{X}_d|} \mathbb{I}(x_t \in \mathcal{X}_d) \cdot \sum_{d=1}^D \frac{\sum_{j=1}^m \tilde{w}_t^{(j)} \mathbb{I}(\tilde{x}_t^{(j)} \in \mathcal{X}_d)}{\sum_{j=1}^m \mathbb{I}(\tilde{x}_t^{(j)} \in \mathcal{X}_d)} \mathbb{I}(x_t \in \mathcal{X}_d) \\ &= \sum_{d=1}^D \frac{\sum_{j=1}^m \tilde{w}_t^{(j)} \mathbb{I}(\tilde{x}_t^{(j)} \in \mathcal{X}_d)}{m|\mathcal{X}_d|} \mathbb{I}(x_t \in \mathcal{X}_d), \end{aligned}$$

where $\mathbb{I}(\cdot)$ is the indicator function, $\mathcal{X}_1 \cup \mathcal{X}_2 \cup \cdots \cup \mathcal{X}_D$ is a partition of the support of x_t , and $|\mathcal{X}_d|$ denotes the volume of $\mathcal{X}_d$.

Compared with the forward pilot method in Algorithm 2, the backward pilots here are generated backward, starting from the constrained time point t_+ . The strong constraint I_{t_+} is automatically incorporated in the proposal distribution to generate $\tilde{x}_{t_+}$ at the beginning. Hence, it is often expected to have a more accurate estimate of $p(C_{t_+} | x_{0:t}, C_t) = p(I_{t+1:t_+} | x_t)$. However, this method requires that the system be Markovian.

This algorithm extends the backward pilot algorithm in Lin, Chen and Mykland (2010) for generating samples of diffusion bridges, allowing more general constraint problems, including those in which frequent (but weak) constraints exist between two rare and strong constraints. It also allows for more flexible constraints than the fixed-point constraints considered in Lin, Chen and Mykland (2010).

Algorithm 3: Backward Pilot Algorithm.

- Initialization: For $j = 1, \dots, m$, draw samples $\tilde{x}_{t_2}^{(j)}$ from a proposal distribution $r(x_{t_2})$ approximately proportional to $p(I_{t_2} | x_{t_2})$ and set $\tilde{w}_{t_2}^{(j)} = 1/r(\tilde{x}_{t_2}^{(j)})$.
- For $t = t_2 - 1, \dots, t_1 + 1$, draw pilot samples backward as follows.
 - Generate samples $\tilde{x}_t^{(j)}$, $j = 1, \dots, m$, from a proposal distribution $r(\tilde{x}_t | \tilde{x}_{t+1}^{(j)})$.
 - Update weights by

$$\tilde{w}_t^{(j)} = \tilde{w}_{t+1}^{(j)} \frac{p(\tilde{x}_{t+1}^{(j)}, I_{t+1} | \tilde{x}_t^{(j)})}{r(\tilde{x}_t^{(j)} | \tilde{x}_{t+1}^{(j)})}, \quad j = 1, \dots, m.$$

- Let $\mathcal{X}_1 \cup \dots \cup \mathcal{X}_D$ be a partition of the support of x_t . Estimate $p(C_{t+} | x_{0:t}, C_t) = p(\mathcal{I}_{t+1:t+} | x_t)$ by

$$f_t(x_t) = \sum_{d=1}^D \eta_{t,d} \mathbb{I}(x_t \in \mathcal{X}_d), \quad (4.1)$$

where $\eta_{t,d} = (1/m|\mathcal{X}_d|) \sum_{j=1}^m \tilde{w}_t^{(j)} \mathbb{I}(\tilde{x}_t^{(j)} \in \mathcal{X}_d)$, and $|\mathcal{X}_d|$ denotes the volume of the subset $\mathcal{X}_d$.

- Return the estimated functions $\{\hat{p}(C_{t+} | x_{0:t}, C_t) = f_t(x_t)\}_{t=t_1+1, \dots, t_2-1}$ to compute the priority scores for Algorithm 1.

4. Examples

In this section, we apply the SMCc method to several problems with different types of constraints; see Appendix A. The first example (expected shortfall) belongs to the rare and strong constraint case, with a strong constraint at the end; the second example (discretized diffusion process) belongs to the intermediate constraint case; and the last example (robotic control) is a stopping time problem.

For each problem, we compare feasible SMCc approaches with other SMC algorithms, without using the proposed special design of priority scores. The acceptance-rejection sampling is referred to simply Rejection. We denote the conventional SMC algorithm with a manipulated drift as SMC-drift, and the SMCc approaches that use a parametric function, forward pilots, or backward pilots to approximate $p(C_{t+} | x_{0:t}, C_t)$ as SMCc-PA, SMCc-FP, and SMCc-BP, respectively. The system in the first example is not Markovian, and so we only use SMCc-FP. For the other examples, it is more convenient to use SMCc-BP, because it is difficult to construct an effective proposal distribution for generating forward pilots that meets the end-point constraints, whereas backward pilot generation is relatively straightforward. Next, provide the details of specific implementations.

4.1. Long-run marginal expected shortfall

Measuring the systemic risk of a firm is important for risk management. Acharya, Engle and Richardson (2012) proposed using the long-run marginal expected shortfall (LRMES) as a systemic risk index, which is defined as the expected capital shortfall of a firm during a financial crisis. In particular, a financial crisis is defined as a decrease in the market index by 40% within six months (126 trading days). Let $x_{m,t}$ and $x_{f,t}$ be the daily logarithmic prices of the market and the firm, respectively, at time t . The LRMES of the firm is defined as (with $T = 126$)

$$\text{LRMES} = \mathbb{E}(1 - e^{x_{f,T} - x_{f,0}} \mid e^{x_{m,T} - x_{m,0}} < 60\%).$$

Following Brownlees and Engle (2012) and Duan and Zhang (2016), we assume that the bivariate process $(x_{m,t}, x_{f,t})$ follows a GJR-GARCH model, as in Glosten, Jagannathan and Runkle (1993). In particular, the $(x_{m,t}, x_{f,t})$ process follows

$$\begin{bmatrix} x_{m,t} \\ x_{f,t} \end{bmatrix} = \begin{bmatrix} x_{m,t-1} \\ x_{f,t-1} \end{bmatrix} + \begin{bmatrix} \sigma_{m,t}^2 & \rho_t \sigma_{m,t} \sigma_{f,t} \\ \rho_t \sigma_{m,t} \sigma_{f,t} & \sigma_{f,t}^2 \end{bmatrix}^{1/2} \begin{bmatrix} \varepsilon_{m,t} \\ \varepsilon_{f,t} \end{bmatrix}, \quad (4.2)$$

where $[\cdot]^{1/2}$ denotes the matrix square root, and $\varepsilon_{m,t}$ and $\varepsilon_{f,t}$ are independent $N(0, 1)$ innovations. The evolution law of the time-varying covariance matrix in (4.2) is given in detail in Appendix B, along with the coefficient used in the simulation.

Without loss of generality, we set $x_{m,0} = x_{f,0} = 0$. In the following, we use $p(\cdot)$ to denote the distribution law under model (4.2), with the parameters outlined in Appendix B. If we draw the sample paths $\{(x_{m,0:T}^{(i)}, x_{f,0:T}^{(i)}, w_T^{(i)})\}_{i=1,\dots,n}$ properly weighted with respect to the distribution $p(x_{m,1:T}, x_{f,1:T} \mid x_{m,0} = 0, x_{f,0} = 0, x_{m,T} < c)$, with $c = \log 0.6$, then the LRMES can be estimated by $\sum_{i=1}^n w_T^{(i)} (1 - e^{x_{f,T}^{(i)}}) / \sum_{i=1}^n w_T^{(i)}$. Note that

$$\begin{aligned} & p(x_{m,1:T}, x_{f,1:T} \mid x_{m,0}, x_{f,0}, x_{m,T} < c) \\ & \propto \mathbb{I}(x_{m,T} < c) p(x_{m,1:T}, x_{f,1:T} \mid x_{m,0}, x_{f,0}) \\ & = \mathbb{I}(x_{m,T} < c) \prod_{t=1}^T p(x_{m,t} \mid x_{m,0:t-1}) \prod_{t=1}^T p(x_{f,t} \mid x_{m,0:t}, x_{f,0:t-1}). \end{aligned}$$

Once we obtain a set of sample paths $\{(x_{m,0:T}^{(i)}, w_T^{(i)})\}_{i=1,\dots,n}$ properly weighted with respect to the distribution $p(x_{m,1:T} \mid x_{m,0}, x_{m,T} < c) \propto \mathbb{I}(x_{m,T} < c) \prod_{t=1}^T p(x_{m,t} \mid x_{m,0:t-1})$, the sample paths $\{x_{f,1:T}^{(i)}\}_{i=1,\dots,n}$ can be drawn easily from $p(x_{f,1:T} \mid x_{m,0}, x_{m,1:T}, x_{f,0}) = \prod_{t=1}^T p(x_{f,t} \mid x_{m,0:t}, x_{f,0:t-1})$ using (4.2). Hence, we focus on sampling $x_{m,0:T}$. We use Rejection, SMC-drift, SMCc-PA, and SMCc-




Figure 1. Box plots of 100 independent estimates of the long-run marginal expected shortfall (LRMES) using different methods. The horizontal line is the “true” LRMES estimated using 100,000 accepted samples generated using the Rejection method.

FP to generate samples from the distribution $p(x_{m,1:T}, x_{f,1:T} | x_{m,0}, x_{m,T} < c)$. We also compare their performance in terms of estimating the LRMES. The implementation details are listed in Appendix B.

For fair comparison, the numbers of Monte Carlo samples in the different methods are adjusted so that each method takes approximately the same CPU time. Specifically, we set the accepted sample sizes to $n = 5$ for the Rejection method, because the acceptance rate is about 0.0001. We use $n = 15000, 12000$, and $10,000$ for SMC-drift, SMCc-PA, and SMCc-FP, respectively. SMCc-FP uses $m = 1000$ forward pilots. In SMCc-PA and SMCc-FP, we perform resampling every five steps. Once $\{(x_{m,0:T}^{(i)}, w_T^{(i)})\}_{i=1,\dots,n}$ is obtained, $x_{f,1:T}^{(i)}$ is sampled accordingly, and the corresponding LRMES is estimated. Figure 1 shows box plots of 100 independent estimates of the LRMES using the varies. The horizontal line is the “true” LRMES estimated using 100,000 accepted samples generated using the Rejection method. The results show that SMCc-FP performs best under a fixed computational cost. Figure 2 plots 50 sample paths of $x_{m,0:T}$ generated using different methods (without weight adjustment). Note that the sample paths generated by Rejection exactly follow the true target distribution. The figures show that SMCc-FP can generate samples close to the true target distribution, with far lower computational cost. The samples generated by SMC-drift and SMC-PA tend to have less diversity and move more aggressively toward the constraint region.

4.2. A diffusion system with intermediate noisy observations

In this section, we conduct a simulation study on a system with periodic and intermediate constraints, corresponding to Case 3 in Appendix A. Consider a




Figure 2. Sample paths of $X_{m,0:T}$ generated by different methods before weight adjustment. The horizontal line denotes a 40% price decrease.

discretized diffusion process x_t , governed by

$$x_t = x_{t-1} + \delta \sin(x_{t-1} - \pi) + \varepsilon_t, \quad (4.3)$$

where $\varepsilon_t \sim N(0, \delta)$, as in Beskos et al. (2006). This is a discretized version of the diffusion process $dX_\lambda = \sin(X_\lambda - \pi)d\lambda + dW_\lambda$, where W_λ is a standard Brownian motion, with step size δ . We take $\delta = 0.1$ in this example.

In this simulation study, two noisy observations of x_t are made at real times $T = 30$ and $T = 60$ (the indices for x_t are $t = 30/\delta = 300$ and $60/\delta = 600$, respectively), with

$$Y_{30} \sim N(x_{30/\delta}, \sigma^2) \quad \text{and} \quad Y_{60} \sim N(x_{60/\delta}, \sigma^2).$$

We also fix the two endpoints at $x_0 = a$ and $x_{90/\delta} = b$. The discretized time points $T_0 = 0$, $T_1 = 30$, $T_2 = 60$, and $T_3 = 90$ are considered to be strong constraints. We apply the SMCc-BP method to generate sample paths of $x_{0:T}$, conditional on the constraints $(x_0, Y_{30}, Y_{60}, x_{90/\delta})$. We take equation (4.3) as the proposal distribution when generating forward paths. The backward pilots are generated from the proposal distribution $r(\tilde{x}_t | \tilde{x}_{t+1}) \sim N(\tilde{x}_{t+1} - \delta \sin(\tilde{x}_{t+1} - \pi), \delta)$.

Note that this process shows a jump behavior among the stable levels at $x = 2k\pi$, for $k = 0, \pm 1, \pm 2, \dots$ (Lin, Chen and Mykland (2010)). In this experiment, we set $x_0 = 0$, $Y_{30} = 6.49$, $Y_{60} = -5.91$, and $x_{90/\delta} = -1.17$, corresponding to the stable levels $0, 2\pi, -2\pi$, and 0 , respectively. Because Y_{30} and Y_{60} differ by a gap of two stable levels, this is a very rare event.

We investigate three levels of measurement errors for the observations Y_{30} and Y_{60} : $\sigma = 0.01$ for very accurate observations, $\sigma = 1.0$ for moderately accurate observations, and $\sigma = 2.0$ for noisy observations. Note that in this experiment, we fix the observations Y_{30} and Y_{60} , but change the underlying assumption of their distributions to reflect the strength of their respective constraints. In each setting, a total of 5,000 sample paths are generated, with 300 backward pilots used to estimate the resampling priority scores. Figure 3 plots the generated sample paths before weight adjustment for each level of error. Figure 4 shows a histogram of the marginal samples of $x_{60/\delta}$ before weight adjustment, which is obtained from the generated sample set $\{x_{0:T}^{(i)}\}_{i=1, \dots, n}$, without considering the




Figure 3. Sampled paths before weight adjustment for $\sigma = 0.01$ (top panel), $\sigma = 1.0$ (middle panel), and $\sigma = 2.0$ (bottom panel) when $x_0 = 0$, $Y_{30} = 6.49$, $Y_{60} = -5.91$, and $x_{90/\delta} = -1.17$.

weights. When the observations are accurate ($\sigma = 0.01$), the two observations act like fixed-point constraints that force all sample paths to pass through the observations. When the observation error is large ($\sigma = 2$), a high proportion of sample paths remain at the original stable level, and only a small proportion are drawn toward the observations. The marginal distributions of $x_{60/\delta}$ show clear differences in the above three cases. Samples from all three levels of error retain the jumping nature of the underlying process, and the SMCc-BP approach is capable of dealing with different levels of observational errors.

4.3. Robotic control

In this example, we consider a robotic control problem in a well-known mechanical system called “Acrobot” (Murray and Hauser (1991)), which consists of two arms of identical inertia and mass, as shown in the left panel of Figure 5. The two arms are allowed to move only in the vertical plane, and the state of the system is determined using a four-dimensional vector $\boldsymbol{\theta} = (\theta_1, \theta_2, \dot{\theta}_1, \dot{\theta}_2)$, where θ_1 and θ_2 are the two angle positions, as marked in Figure 5, and $\dot{\theta}_1$ and $\dot{\theta}_2$ are their respective velocities. The system is controlled only through κ , that is, the torque of the actuator at the joint of the two arms.

A common control task in robotics is to find a control sequence $\kappa_0, \kappa_1, \dots$, following which, a system starting at state $\boldsymbol{\theta}_0$ will reach the desired target state $\boldsymbol{\theta}_*$ as fast as possible (Spong (1995); Perez et al. (2012); Duan et al. (2016)). In our experiment, we set the starting state $\boldsymbol{\theta}_0 = (0, \pi/2, 0, 0)$ and the target state $\boldsymbol{\theta}_* = (0, -\pi/2, 0, 0)$, as shown in the middle and right panels, respectively, of Figure 5.




Figure 4. Histogram of the marginal samples of x_{60}/δ before weight adjustment for $\sigma = 0.01$ (top panel), $\sigma = 1.0$ (middle panel), and $\sigma = 2.0$ (bottom panel) when $x_0 = 0$, $Y_{30} = 6.49$, $Y_{60} = -5.91$, and $x_{90}/\delta = -1.17$.




Figure 5. Acrobot with two arms (left panel), starting position θ_0 at $(0, \pi/2, 0, 0)$ (middle panel), and target position θ_* at $(0, -\pi/2, 0, 0)$ (right panel).

The detailed four-dimensional system of the Acrobot and its discretization to differential equations with a fixed time interval ($\delta = 0.02$ seconds) are shown in Appendix C. If we treat the control sequence κ_t as random innovations, then the problem of finding a sequence of controls to move the system from an initial state θ_0 to a fixed ending state θ_* becomes one of finding the paths of a dynamic system with end-point constraints. Here, we assume that κ_t is independent over time, and follows a uniform $[-5, 5]$ distribution. Note that κ_t can only provide a one-dimensional nonlinear control, and takes values in a limited range. It is often a challenging problem to find the sequence of controls to reach the target state, because the system is nonlinear, especially when the target position has zero velocity.

Let $x_0 = \theta_0$ and define the target region as $\Gamma = \{x : \|x - \theta_*\|_\infty \leq 0.01\}$, where $\|\cdot\|_\infty$ denotes the sup norm of a vector. The problem of finding a control




Figure 6. Left panel: Sample paths generated using the SMCc-BP method (in gray), with the “optimal” path (in black) that reaches the target state θ_* at $t = 68$. The control sequence for the “optimal” path is shown in the top panel. Right panel: Sample paths generated using the random search method for $t = 0, 1, \dots, 68$.

sequence $\kappa_0, \kappa_1, \dots$ that moves the system from θ_0 to θ_* can be formulated as a constrained sampling problem, as outlined in Case 5 of Appendix A.

We propose using the (modified) SMCc-BP method outlined in Appendix D, where (C.2) is used for forward propagation, and a resampling step is performed with the priority scores $\beta_t^{(i)}$ in (D.2), which is estimated using backward pilots. The backward pilots $\{\tilde{x}_{1:T}^{(j)}\}_{j=1}^m$ are generated backward from $\tilde{x}_T^{(j)} = \theta_*$, using

$$\tilde{x}_{t-1}^{(j)} = \tilde{x}_t^{(j)} - \tilde{v}(\tilde{x}_t^{(j)}, \tilde{\kappa}_t)\delta,$$

where $\tilde{\kappa}_t$ is drawn from a uniform $[-5, 5]$ distribution, and the function $\tilde{v}(\cdot)$ is similar to $v(\cdot)$ in (C.3), but with δ replaced with $-\delta$.

In the simulation, we use $n = 10000$ forward samples and $m = 5000$ backward pilots for the SMCc-BP method, with $a = 0.1\delta$ in (D.1) and the maximum duration $T = 100$. The results are compared with those of the random search method using $n = 15000$ samples. The sample trajectories from the two methods are plotted in Figure 6. The SMCc-BP method found an “optimal” path that reaches the target state θ_* at $t = 68$, corresponding to a total time of 1.36 seconds, shown by the black line in the left panel. On the other hand, the random search method was not able to find any paths close to θ_* before time $T = 100$. Note that at $t = 50$, the “optimal” trajectory has a large velocity $\dot{\theta}_2 \approx -2\pi$ (the black line in the bottom figure in the left panel) to swing the second arm up to the target position. This state is quite far away from the target position, but was found and guided by the backward pilots. On the other hand, the simple random search does not have such information to use.




Figure 7. The minimum distance of the samples generated using the SMCc-BP method and the random search method (SMC) to the target state against time.

Figure 7 plots the minimum distance $\|x_t^{(j)} - \theta_*\|_\infty$ of all the generated samples to the target state against time under the SMCc-BP and simple random search methods. Note that the minimum distance at $t = 50$ obtained using the SMCc-BP method is much larger than that of the random search method. However, it seems to be necessary to allow the arms to gain velocity in order for θ_2 to move to the target quickly. Guided by the backward pilots, the SMCc-BP method was able to find such a path. Thus at $t = 68$, some of the SMCc-BP samples reach the target state.

5. Conclusion

We focus on the problem of generating (weighted) sample paths of a dynamic system with rare and strong constraints, under an SMC framework. There are numerous SMC implementations for various problems. Some algorithms and approaches can be used for constrained systems, though mostly in the case of frequent and weak constraints, such as the delayed approaches and fixed-lag smoothing algorithms in state-space models, including the auxiliary particle filter (APF) of Pitt and Shephard (1999). We focus on generating the entire sample path from the joint distribution $x_{0:T}$, in contrast to the typical filtering and smoothing problems that often focus on marginal distributions. We also consider the problem in which the objective is to reach a fixed target (constraint region) with a variable time duration, that is, the stopping time problem. A systematic formulation of the problem is proposed by introducing a sequence of constraint events and their constraint strength measures. Under the guidance of three closely related distributions (the target, compromised, and propagation) induced under the system and a variable lookahead timescale, we introduce a resampling operation with optimal priority scores (resampling probabilities) for efficient operations. Two efficient approaches for approximating the needed optimal priority scores are developed, based on nonparametric smoothing techniques that

use an ensemble of forward or backward pilots, respectively. The framework is general enough to encompass earlier studies, and lays a foundation of further development of efficient implementations of the SMC method with constrained dynamic systems. We show the effectiveness of the approach using several examples.

Note that our approach can be viewed as an extension of the APF proposed by Pitt and Shephard (1999), the lookahead strategies of Chen, Wang and Liu (2000) and Lin, Chen and Liu (2013), and the iterated APF of Guarniero, Johansen and Lee (2016). These special SMC algorithms seek higher efficiency by incorporating future events and constraints when generating samples. For example, the APF conducts resampling based on the priority score $\beta_t^{(i)} = w_t^{(i)} p(y_{t+1} | x_t^{(i)})$ (or its approximation). This corresponds to using $t_+ = t + 1$ under the SMCc framework for state-space models, using the observation y_{t+1} . We focus on sampling problems with rare and strong constraints.

The primary goal of this study is to generate the entire sample path, which is essentially a smoothing problem. Because of the sequential nature of the SMC method, the inherited degeneracy problem cannot be avoided completely, especially because we rely on a resampling scheme to deal with the constraints. The procedures for approximating the optimal priority scores at the very beginning of the sequential process help to improve the sample quality at the early stage, because these samples are designed to follow an approximation of the target smoothing distribution, instead of the filtering distribution. Experiments show that the proposed SMCc approaches alleviate the degeneracy problem much better than simple implementations of the SMC smoother do.

The notation of conditional probability used throughout this paper can be simplified significantly if the underlying dynamic system is Markovian, such that $p(x_t, C_t | x_{0:t-1}, C_{t-1}) = p(x_t, C_t | x_{t-1}, C_{t-1})$. The proposed method in Section 3.1 also becomes simpler, because the statistic $S(x_{0:t})$ can simply be chosen as $S(x_{0:t}) = x_t$.

Acknowledgments

The authors thank the editor, associate editor, and two anonymous referees for their helpful comments and suggestions. Lin's research was supported by the National Natural Science Foundation of China #72033008 and Science Foundation of Ministry of Education of China #19YJA910003.

Appendix

A. Some Examples of the Constrained Problems

Case 1. Frequent constraints: In this case new constraints are imposed frequently on the system. For example, consider a state-space model with state




Figure 8. Segmentation of a stochastic process with intermediate observations.

equation $x_t = f(x_{t-1}, \varepsilon_t)$ and observation equation $y_t = g(x_t, \epsilon_t)$, where ε_t and ϵ_t are random noises with known distributions. The observations y_t can be viewed as constraints imposed on the system. In this case we have a constraint at each time t , in the form $I_t = \{g(x_t, \epsilon_t) = y_t\}$ and the cumulative constraints $C_t = \{g(x_s, \epsilon_s) = y_s, s = 0, \dots, t\}$. Using SMC to study state-space models has been extensively studied with a vast literature, including the cases with small observation noises (ϵ_t with small variance), such as using the full information proposal distribution (Liu and Chen (1998)), the auxiliary particle filter (Pitt and Shephard (1999)), the independent particle filter (Lin et al. (2005)) or certain short-range lookahead method (Lin, Chen and Liu (2013)). In this paper we do not focus on this frequent constraints case.

Case 2. Rare and strong constraints: In this case the constraints occur rarely, but with strong effect. For example, in the problem of generating diffusion bridge paths that connect two fixed endpoints $x_0 = a$ and $x_T = b$, the cumulative constraint events are $C_0 = \dots = C_{T-1} = \{x_0 = a\}$ and $C_T = \{x_0 = a, x_T = b\}$. Then $p(x_{0:T} | C_{T-1}) = p(x_{0:T} | x_0 = a)$ and $p(x_{0:T} | C_T) = \delta_D(x_T - b)p(x_{0:T} | x_0 = a)/p(x_T = b | x_0 = a)$, where $\delta_D(\cdot)$ is the Dirac delta function. Hence $G(T) = \infty$, which indicates that $I_T = \{x_T = b\}$ is a strong constraint. If the constraint at T is not a fixed point, but a noisy measurement of x_T , then the strength of the constraint would depend on T and variance of the measurement error.

Case 3. Periodic and intermediate constraints: In certain systems we have noisy measurements of the unobservable states $x_{0:T}$ periodically. For example, let y_k be a sequence of noisy measurements at time $t_1, \dots, t_K$. The intermediate observations split the whole path into K segments as shown in Figure 8. Again, the constraints strength depend on how frequent one observes the intermediate observations and how strong (or accurate) y_k 's are.

Case 4. Multilevel constraints: In some applications, there may exist multiple levels of constraints, including those with a hierarchical structure, such as one level of weak but frequent constraints and another level of strong but rare constraints. A special case is a standard state-space model with two fixed endpoint constraints. Specifically, suppose a state-space model is governed by the state dynamics $x_t = f(x_{t-1}, \varepsilon_t)$ and the observation equation $y_t = g(x_t, \epsilon_t)$. In addition, two fixed endpoint constraints are imposed on $x_{0:T}$ with $x_0 = a$

and $x_T = b$. The routine observations $y_1, \dots, y_{T-1}$ can be viewed as a layer of weak constraints and the fixed point constraints are viewed as a layer of rare and strong constraints. Although $G(T)$ is still infinity here, this case is “easier” than that without the observations $y_1, \dots, y_{T-1}$, if the observation sequence is “faithful” in the sense that it is based on one realization of the bridge $x_0 = a, x_1, \dots, x_{T-1}, x_T = b$ (e.g. $y_t = x_t + \epsilon_t$). The observations $y_1, \dots, y_{T-1}$ provide general guidance for the system to meet the endpoint constraint $x_T = b$.

Case 5. Constraints in stopping time problems: In the stopping time problems, we are often interested in generating sample paths that eventually reach a constrained set within a maximum duration T . Let

$$\tau = \min \{t \geq 0 : x_t \in \Gamma\} \quad (\text{A.1})$$

be a stopping time, which is the first time that the process $\{x_0, x_1, \dots\}$ reaches the constrained set Γ . The constraint of such problems is in the form of $C = \{x_{0:\infty} : \tau \leq T\}$. Note that C is a “global” constraint, which is different from the local constraints imposed on specific times as we discussed before. In the optimal stopping time problems, one would be interested in finding the “optimal path” that satisfies the constraint with the smallest stopping time τ . We show such an application in Section 4.3.

B. Detailed Information for the LRMES Eexample in Section 4.1

The evolution of the covariance matrix in (4.2) follows two processes, one for the volatilities $(\sigma_{m,t}^2, \sigma_{f,t}^2)$, and a different one for the correlation ρ_t . Specifically, given the initial values of $(\sigma_{m,1}, \sigma_{f,1}, \rho_1)$, and a realization of the innovation process $\{(\varepsilon_{m,t}, \varepsilon_{f,t})\}_{t=1,\dots,T}$, the volatility process $(\sigma_{m,t}^2, \sigma_{f,t}^2)$ used in (4.2) iteratively follows

$$\begin{aligned} \sigma_{m,t}^2 &= \omega_m + [\alpha_m + \gamma_m \mathbb{I}(\varepsilon_{m,t-1} < 0)](\sigma_{m,t-1} \varepsilon_{m,t-1})^2 + \beta_m \sigma_{m,t-1}^2, \\ \sigma_{f,t}^2 &= \omega_f + [\alpha_f + \gamma_f \mathbb{I}(\varepsilon_{f,t-1} < 0)](\sigma_{f,t-1} \varepsilon_{f,t-1})^2 + \beta_f \sigma_{f,t-1}^2, \end{aligned} \quad (\text{B.1})$$

for $t = 2, \dots, T$. This is an asymmetric GARCH process taking into the account of the difference in positive or negative innovations (or shocks).

The time-varying correlation coefficient series $\{\rho_t\}$ follows a separate dynamic conditional correlation (DCC) model. Specifically, let

$$Q_t = \begin{bmatrix} \sigma_{m,t}^{*2} & \rho_t \sigma_{m,t}^* \sigma_{f,t}^* \\ \rho_t \sigma_{m,t}^* \sigma_{f,t}^* & \sigma_{f,t}^{*2} \end{bmatrix}$$

be a sequence of 2×2 covariance matrices. Using the same initial values of $\sigma_{m,1}, \sigma_{f,1}, \rho_1$, and the same realization of the innovation process $\{(\varepsilon_{m,t}, \varepsilon_{f,t})\}_{t=1,\dots,T}$ as above, we let

$$Q_1 = \begin{bmatrix} \sigma_{m,1}^2 & \rho_1 \sigma_{m,1} \sigma_{f,1} \\ \rho_1 \sigma_{m,1} \sigma_{f,1} & \sigma_{f,1}^2 \end{bmatrix}$$

and

$$Q_t = (1 - \alpha_C - \beta_C)Q_1 + \alpha_C \begin{bmatrix} \sigma_{m,t-1}^* \varepsilon_{m,t-1} \\ \sigma_{f,t-1}^* \varepsilon_{f,t-1} \end{bmatrix} \begin{bmatrix} \sigma_{m,t-1}^* \varepsilon_{m,t-1} \\ \sigma_{f,t-1}^* \varepsilon_{f,t-1} \end{bmatrix}' + \beta_C Q_{t-1},$$

for $t = 2, \dots, T$. Then we extract ρ_t from Q_t and use it in (4.2). We note that $\sigma_{m,t}^*$ and $\sigma_{f,t}^*$ in Q_t are different from $\sigma_{m,t}$ and $\sigma_{f,t}$ in (B.1) for $t > 1$. The simulation of the process (4.2) starts with setting the initial values and a realization of the innovations, then construct the covariance process and the $(x_{m,t}, x_{f,t})$ process.

To set parameters in (4.2) for our simulation, we apply the model to S&P500 index and the stock prices of Citigroup from January 2, 2012 to December 31, 2017. The maximum likelihood estimates of the parameters are $\omega_m = 3.35 \times 10^{-6}$, $\alpha_m = 3.35 \times 10^{-6}$, $\gamma_m = 0.152$, $\beta_m = 0.858$, $\omega_f = 4.22 \times 10^{-6}$, $\alpha_f = 0.0148$, $\gamma_f = 0.0542$, $\beta_f = 0.935$, $\alpha_C = 0.0755$, $\beta_C = 0.862$, $\sigma_{m,1} = 0.0113$, $\sigma_{f,1} = 0.03$, and $\rho_1 = 0.705$.

The detailed implementations of the Monte Carlo methods compared in this example are as follows:

Rejection: Samples are generated forward, following the distribution $p(x_{m,1:T} | x_{m,0})$ without considering the constraint. At the end, a sample is accepted if $x_{m,T} < c$.

SMC-drift: Samples are generated with a standard SMC implementation with a proposal distribution that includes a drift term, similar to that proposed by Durham and Gallant (2002). Specifically, the proposal distribution $q(x_{m,t} | x_{m,1:t-1})$ used follows

$$x_{m,t} = x_{m,t-1} + \frac{c}{T} + \sigma_{m,t} \varepsilon_{m,t}, \quad (\text{B.2})$$

and the same evolution law of $\sigma_{m,t}$, with $\varepsilon_{m,t} \sim N(0, 1)$ and $t = 1, \dots, T-1$. Note that $c < 0$, we included a negative drift term c/T in the proposal distribution to force $x_{m,t}$ towards the constraint region. At the end, the samples are weighted by $w_T^{(i)} = \mathbb{I}(x_{m,T}^{(i)} < c) \prod_{t=1}^T p(x_{m,t}^{(i)} | x_{m,0:t-1}^{(i)}) / \prod_{t=1}^T q(x_{m,t}^{(i)} | x_{m,1:t-1}^{(i)})$.

SMCc-PA: Samples are generated using the SMCc method in Algorithm 1, with a parametric priority score function, with $t_+ = T$ for all t . The propagation equation (4.2) is used as the proposal distribution. The priority score used is $\beta_t^{(i)} = w_t^{(i)} \hat{p}(x_{m,T} < c | x_{0:t}^{(i)})$ with $\hat{p}(x_{m,T} < c | x_{0:t}^{(i)}) = \Phi(c; x_t^{(i)}, (T-t)\bar{\sigma}_m^2)$, where $\Phi(c; \mu, \sigma^2)$ is the CDF of $N(\mu, \sigma^2)$ evaluated at the value c and $\bar{\sigma}_m^2$ is the long-term average of $\sigma_{m,t}^2$.

SMCc-FP: Samples are generated using the SMCc method with priority scores estimate based on forward pilots. All settings are similar to SMCc-PA, except that $p(C_T | x_{m,0:t}^{(i)}, C_t) = p(x_{m,T} < c | x_{m,0:t}^{(i)})$ is estimated by forward pilots. The model is not Markovian, but the statistic $S(x_{m,0:t}) = (x_{m,t}, \sigma_{m,t+1})$ can be used and satisfies (3.2). Furthermore, since $p(x_{m,T} < c | x_{m,t}, \sigma_{m,t+1}) = p(x_{m,T} - x_{m,t} < c - x_{m,t} | \sigma_{m,t+1})$, we estimate the conditional cumulative distribution function $p(x_{m,T} - x_{m,t} < \Delta | \sigma_{m,t+1})$, using a histogram estimator with partition $\mathcal{S}_\sigma = \cup_d \{0.005(d-1) < \sigma_{m,t+1} \leq 0.005d\}$.

In the above approaches (except Rejection), we force the samples to satisfy the constraint $x_{m,T} < c$ in the last step, by letting $x_{m,T}^{(i)} \sim N(x_{m,T-1}^{(i)}, \sigma_T^{2(i)})$ truncated on $(-\infty, c)$, and updated their importance weight accordingly.

C. Dynamic system of Acrobot used in Section 4.3

The dynamics of Acrobot can be described by the following differential equation.

$$\begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \\ \ddot{\theta}_1 \\ \ddot{\theta}_2 \end{bmatrix} = \xi(\theta, \kappa) := \begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \\ \frac{d_{22}(u_1 - c_1 - g_1) - d_{12}(u_2 - c_2 - g_2)}{d_{11}d_{22} - d_{12}d_{21}} \\ \frac{d_{11}(u_2 - c_2 - g_2) - d_{21}(u_1 - c_1 - g_1)}{d_{11}d_{22} - d_{12}d_{21}} \end{bmatrix}, \quad (\text{C.1})$$

with

$$\begin{aligned} d_{11} &= I_1 + I_2 + 2m_0l_c^2 + m_0l^2 + 2m_0ll_c \cos \theta_2, \\ d_{22} &= I_2 + m_0l_c^2, \\ d_{12} &= d_{21} = I_2 + m_0l_c^2 + m_0ll_c \cos \theta_2, \\ c_1 &= -m_0ll_c\dot{\theta}_2^2 \sin \theta_2 - 2m_0ll_c\dot{\theta}_1\dot{\theta}_2 \sin \theta_2, \\ c_2 &= m_0ll_c\dot{\theta}_1^2 \sin \theta_2, \\ g_1 &= m_0(l + l_c)g \cos \theta_1 + m_0l_cg \cos(\theta_1 + \theta_2), \\ g_2 &= m_0l_cg \cos(\theta_1 + \theta_2), \\ u_1 &= -\mu\dot{\theta}_1, \\ u_2 &= \kappa - \mu\dot{\theta}_2, \end{aligned}$$

where m_0 , l , l_c are the mass, length and distance between the center of mass and pivot for both arms correspondingly, I_1 and I_2 are their moments of inertia, μ is the friction coefficient, and g is the acceleration of gravity. We set $m_0 = 1.0$, $l = 1.0$, $l_c = 0.5$, $I_1 = I_2 = 1/12$ and $g = 9.8$ as in the international standard of units.

We discretize the time with interval length $\delta = 0.02$ seconds and let $x_t = \theta_{t\delta} = (\theta_{t\delta,1}, \theta_{t\delta,2}, \dot{\theta}_{t\delta,1}, \dot{\theta}_{t\delta,2})$. Using the fourth order Runge-Kutta method (DeVries and Wolf (1994)), we approximate the equation (C.1) by a discrete-time model as follows.

$$x_t = x_{t-1} + v(x_{t-1}, \kappa_{t-1})\delta, \quad (\text{C.2})$$

where $v(x_{t-1}, \kappa_{t-1}) = (v^{(1)} + 2v^{(2)} + 2v^{(3)} + v^{(4)})/6$ with

$$\begin{aligned} v^{(1)} &= \xi(x_{t-1}, \kappa_{t-1}), \\ v^{(2)} &= \xi\left(x_{t-1} + \frac{v^{(1)}\delta}{2}, \kappa_{t-1}\right), \\ v^{(3)} &= \xi\left(x_{t-1} + \frac{v^{(2)}\delta}{2}, \kappa_{t-1}\right), \\ v^{(4)} &= \xi(x_{t-1} + v^{(3)}\delta, \kappa_{t-1}), \end{aligned} \quad (\text{C.3})$$

where $\xi(\cdot)$ is in (C.1) and κ_t is the torque imposed at time $t\delta$.

D. Formulation of the Acrobot example as a constrained problem

The Acrobot problem is a typical example of Case 5 in Appendix A. As a stopping time problem, a standard rejection sampling procedure would generate $x_0, x_1, \dots$ using the forward propagation equation until the sample path reaches the target region Γ before the maximum duration time T . This approach may not be efficient as the probability to reach Γ before T could be extremely small. To find the shortest sample path to reach Γ , we consider SMCc in Algorithm 1 with certain modification as follows. To find the shortest sample path to reach Γ , we consider generating sample paths properly weighted with respect to the target distribution

$$\pi_T(x_{0:T}) \propto p(x_{0:T} | C) e^{-a\tau} \propto p(x_{0:T}) \mathbb{I}(\tau \leq T) e^{-a\tau}, \quad (\text{D.1})$$

where τ is the stopping time defined in (A.1), $C = \{\tau \leq T\}$ and $\mathbb{I}(\cdot)$ is an indicator function. Here an additional term $e^{-a\tau}$, $a > 0$, is added to the target distribution to encourage faster stopping. Set the intermediate target distribution as

$$\tilde{p}_t(x_{0:t}) = p(x_{0:t}), \quad t < T,$$

and let $p_t^+(x_{0:t}) = \pi_T(x_{0:T})$. Similar to (2.8), to effectively generated properly weighted samples with respect to $\pi_T(x_{0:T})$, the optimal priority score in SMCc is

$$\begin{aligned} \beta_t &= w_t \frac{p_t^+(x_{0:t})}{\tilde{p}_t(x_{0:t})} \propto w_t \frac{\int p(x_{0:T}) \mathbb{I}(\tau \leq T) e^{-a\tau} dx_{t+1:T}}{p(x_{0:t})} \\ &= w_t \mathbb{E}[\mathbb{I}(\tau \leq T) e^{-a\tau} | x_{0:t}] \end{aligned}$$

In implementation, when a sample path $x_{0:t}^{(i)}$ reaches the target region Γ , it is “accepted” without further propagation.

Now we discuss how to approximate the term $E[\mathbb{I}(\tau \leq T) e^{-a\tau} | x_{0:t}]$. Since the state sequence is a homogeneous Markovian process, that is, $p(x_t | x_{0:t-1}) = p(x_t | x_{t-1}) = p(x_1 | x_0)$, β_t can be estimated using the backward pilots as in Algorithm 3, with $t_1 = 0$, $t_2 = T$, $I_T = \{x_T \in \Gamma\}$ and $I_t = \Omega$ for $t < T$. Particularly, we use

$$\beta_t^{(i)} = w_t^{(i)} f_t(x_{0:t}^{(i)}) \quad (\text{D.2})$$

with

$$f_t(x_{0:t}) = \begin{cases} e^{-\tau}, & \text{if } \tau(x_{0:t}) \leq t; \\ \sum_{d=1}^D \eta_{t,d} \mathbb{I}(x_t \in \mathcal{X}_d), & \text{if } \tau(x_{0:t}) > t, \end{cases}$$

where

$$\eta_{t,d} = \sum_{s=t}^{T-1} \left[\frac{1}{m|\mathcal{X}_d|} \sum_{j=1}^m e^{-a(t+T-s)} \tilde{w}_s^{(j)} \mathbb{I}(\tilde{x}_s^{(j)} \in \mathcal{X}_d) \right].$$

Here $(1/m|\mathcal{X}_d|) \sum_{j=1}^m e^{-a(t+T-s)} \tilde{w}_s^{(j)} \mathbb{I}(\tilde{x}_s^{(j)} \in \mathcal{X}_d)$ is an approximation of $\mathbb{E}[\mathbb{I}(\tau = t + T - s) e^{-a\tau} | x_{0:t}]$ if $\tau(x_{0:t}) > t$.

References

- Acharya, V., Engle, R. and Richardson, M. (2012). Capital shortfall: A new approach to ranking and regulating systemic risks. *American Economic Review: Papers & Proceedings* **102**, 59–64.
- Avitzour, D. (1995). Stochastic simulation Bayesian approach to multitarget tracking. *IEEE Proceedings-Radar, Sonar and Navigation* **142**, 41–44.
- Beskos, A., Papaspiliopoulos, O., Roberts, G. O. and Fearnhead, P. (2006). Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **68**, 333–382.
- Brownlees, C. and Engle, R. (2012). Volatility, correlation and tails for systemic risk measurement. Technical report, SSRN 1611229.
- Carpenter, J., Clifford, P. and Fearnhead, P. (1999). An improved particle for non-linear problems. *IEEE Proceedings on Radar, Sonar, and Navigation* **146**, 2–7.
- Chen, R., Wang, X. and Liu, J. S. (2000). Adaptive joint detection and decoding in flat-fading channels via mixture Kalman filtering. *IEEE Transactions on Information Theory* **46**, 2079–2094.
- DeVries, P. L. and Wolf, R. P. (1994). A first course in computational physics. *Computers in Physics* **8**, 178–179.
- Douc, R., Moulines, E. and Stoffer, D. (2014). *Nonlinear Time Series: Theory, Methods and Applications with R examples*. Taylor & Francis, CRC Press.
- Doucet, A., Briers, M. and S  n  cal, S. (2006). Efficient block sampling strategies for sequential Monte Carlo methods. *Journal of Computational and Graphical Statistics* **15**, 693–711.
- Doucet, A. and Johansen, A. (2011). A tutorial on particle filtering and smoothing: Fifteen years later. In *The Oxford Handbook of Nonlinear Filtering* (Edited by D. Crisan and B. Rozovsky). Oxford University Press.

- Duan, J. and Zhang, C. (2016). Non-Gaussian bridge sampling with an application. Technical report. National University of Singapore.
- Duan, Y., Chen, X., Houthooft, R., Schulman, J. and Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, 1329–1338. PMLR.
- Durham, G. B. and Gallant, A. R. (2002). Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business & Economic Statistics* **20**, 297–338.
- Glosten, L., Jagannathan, R. and Runkle, D. (1993). Relationship between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* **48**, 1779–1801.
- Godsill, S. J., Doucet, A. and West, M. (2004). Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association* **99**, 156–168.
- Gordon, N. J., Salmond, D. J. and Smith, A. F. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEEE Proceedings F (Radar and Signal Processing)* **140**, 107–113.
- Guarniero, R., Johansen, A. and Lee, A. (2016). The iterated auxiliary particle filter. *Journal of the American Statistical Association* **112**, 1636–1647.
- Jazwinski, A. (1970). *Stochastic Processes and Filtering Theory*. Academic Press.
- Kim, S., Shephard, N. and Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *The Review of Economic Studies* **65**, 361–393.
- Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics* **5**, 1–25.
- Kong, A., Liu, J. S. and Wong, W. H. (1994). Sequential imputations and Bayesian missing data problems. *Journal of the American Statistical Association* **89**, 278–288.
- Lin, M., Chen, R. and Liang, J. (2008). Statistical geometry of lattice chain polymers with voids of defined shapes: Sampling with strong constraints. *The Journal of Chemical Physics* **128**, 1–12.
- Lin, M., Chen, R. and Liu, J. S. (2013). Lookahead strategies for sequential Monte Carlo. *Statistical Science* **28**, 69–94.
- Lin, M., Chen, R. and Mykland, P. (2010). On generating Monte Carlo samples of continuous diffusion bridges. *Journal of the American Statistical Association* **105**, 820–838.
- Lin, M., Lu, H.-M., Chen, R. and Liang, J. (2008). Generating properly weighted ensemble of conformations of proteins from sparse or indirect distance constraints. *The Journal of Chemical Physics* **129**, 1–13.
- Lin, M. T., Zhang, J. L., Cheng, Q. and Chen, R. (2005). Independent particle filters. *Journal of the American Statistical Association* **100**, 1412–1421.
- Liu, J. S. and Chen, R. (1995). Blind deconvolution via sequential imputations. *Journal of the American Statistical Association* **90**, 567–576.
- Liu, J. S. and Chen, R. (1998). Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association* **93**, 1032–1044.
- Murray, R. M. and Hauser, J. E. (1991). *A Case Study in Approximate Linearization: The Acrobat Example*. University of California.
- Pedersen, A. R. (1995). Consistency and asymptotic normality of an approximate maximum likelihood estimator for discretely observed diffusion processes. *Bernoulli* **1**, 257–279.
- Perez, A., Platt, R., Konidaris, G., Kaelbling, L. and Lozano-Perez, T. (2012). LQR-RRT*: Optimal sampling-based motion planning with automatically derived extension heuristics. In *2012 IEEE International Conference on Robotics and Automation*, 2537–2542.

- Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association* **94**, 590–599.
- Scharth, M. and Kohn, R. (2016). Particle efficient importance sampling. *Journal of Econometrics* **190**, 133–147.
- Spong, M. W. (1995). The swing up control problem for the acrobat. *IEEE Control Systems Magazine* **15**, 49–55.
- Tsay, R. and Chen, R. (2019). *Nonlinear Time Series Analysis*. John Wiley & Sons.
- Wang, X., Chen, R. and Guo, D. (2002). Delayed-pilot sampling for mixture Kalman filter with application in fading channels. *IEEE Transactions on Signal Processing* **50**, 241–254.
- Whiteley, N. and Lee, A. (2014). Twisted particle filters. *The Annals of Statistics* **42**, 115–141.
- Zhang, J., Dundas, J., Lin, M., Chen, R., Wang, W. and Liang, J. (2009). Prediction of geometrically feasible three dimensional structures of Pseudoknotted RNA through free energy. *RNA* **15**, 2248–2263.
- Zhang, J., Lin, M., Chen, R., Liang, J. and Liu, J. S. (2007). Monte Carlo sampling of near-native structures of proteins with applications. *Proteins: Structure, Function, and Bioinformatics* **66**, 61–68.
- Zhang, J. L. and Liu, J. S. (2002). A new sequential importance sampling method and its application to the two-dimensional hydrophobic–hydrophilic model. *The Journal of Chemical Physics* **117**, 3492–3498.

Chencheng Cai

Department of Mathematics and Statistics, Washington State University, Pullman, WA 99164, USA.

E-mail: chencheng.cai@wsu.edu

Rong Chen

Department of Statistics, Rutgers University, Piscataway, NJ 08854, USA.

E-mail: rongchen@stat.rutgers.edu

Ming Lin

Wang Yanan Institute for Studies in Economics and Fujian Key Laboratory of Statistical Science, Xiamen University, Xiamen, Fujian 361005, China.

E-mail: linming50@xmu.edu.cn

(Received May 2022; accepted January 2023)