CONTROLLED SEQUENTIAL MONTE CARLO

BY JEREMY HENG[1], ADRIAN N. BISHOP[2], GEORGE DELIGIANNIDIS[3,*] AND
ARNAUD DOUCET[3,**]

[1]*ESSEC Business School, heng@essec.edu*
[2]*CSIRO, adrian.bishop@uts.edu.au*
[3]*University of Oxford, *deligian@stats.ox.ac.uk; **doucet@stats.ox.ac.uk*

Sequential Monte Carlo methods, also known as particle methods, are a popular set of techniques for approximating high-dimensional probability distributions and their normalizing constants. These methods have found numerous applications in statistics and related fields; for example, for inference in nonlinear non-Gaussian state space models, and in complex static models. Like many Monte Carlo sampling schemes, they rely on proposal distributions which crucially impact their performance. We introduce here a class of controlled sequential Monte Carlo algorithms, where the proposal distributions are determined by approximating the solution to an associated optimal control problem using an iterative scheme. This method builds upon a number of existing algorithms in econometrics, physics and statistics for inference in state space models, and generalizes these methods so as to accommodate complex static models. We provide a theoretical analysis concerning the fluctuation and stability of this methodology that also provides insight into the properties of related algorithms. We demonstrate significant gains over state-of-the-art methods at a fixed computational complexity on a variety of applications.

1. Introduction. Sequential Monte Carlo (SMC) methods have found a wide range of applications in many areas of statistics as they can be used, among others things, to perform inference for dynamic nonlinear non-Gaussian state space models [32, 33, 41] but also for complex static models [8, 12, 39]; see [16, 30, 31] for recent reviews of this active area. Although these methods are supported by theoretical guarantees [11], the number of samples required to achieve a desired level of precision of the corresponding Monte Carlo estimators can be prohibitively large in practice, especially so for high-dimensional problems.

The present work is one means to address the computational difficulties with SMC in offline inference settings. In particular, we leverage ideas from optimal control theory and we seek novel SMC methods that achieve a desired level of precision at a fraction of the computational cost of state-of-the-art algorithms. We introduce a class of algorithms that will be referred to as controlled SMC, under which the sequence of SMC proposal distributions are related naturally with an associated optimal control problem. The cost functional is the Kullback–Leibler divergence from the sought after proposals to the target distributions and we may account for an arbitrary current proposal estimate. With this formulation, the optimal proposal distributions are specified by the optimal control policy of a related dynamic programming recursion. In general, this dynamic programming recursion is intractable. However, by making this connection, we can then exploit an array of methods and procedures for so-called approximate dynamic programming (ADP). Broadly speaking, a single iteration of our proposed methodology involves: (1) based on a current sequence of proposal distributions, running a SMC method to obtain a collection of samples that approximate the sequence

of SMC target distributions; (2) using these samples as support points, we approximate intractable backward recursions using regression to compute a new policy that specifies a new sequence of approximately optimal proposal distributions. Continuing in this manner allows us to further refine the proposal distributions and improve our approximation of the target distributions, via a novel iteration of SMC and ADP.

Prior influential work in [42] proposed a motivating method in the context of importance sampling, for computing the marginal likelihood in state space models. In this contribution, the sequential structure which defines the marginal likelihood is exploited and proposal distributions are defined by a sequence of parameterized Markov transition kernels. A criterion based on the variance of importance weights is introduced to optimize these parameters and an iterative procedure with fixed random numbers is proposed. The work of [47] extends [42] by employing these optimized proposal distributions within a SMC methodology. In particular, [47] identified the appropriate importance weights one should use for resampling, which is crucial to ensure that the variance of the marginal likelihood estimator remains controlled. Moreover, [47] also recommends relaxing the use of common random variables across iterations. Recent work in [22], again with a focus on discrete time state space models, may be viewed as an extension of [47] where resampling is performed at every iteration, instead of just the last. The resulting algorithm is numerically much more stable than [42, 47]. Although the iterative procedures in [22, 42, 47] are similar in spirit to our proposed methodology, the main and important difference is that all of these works employ an optimality criterion, to learn proposal distributions, that is not adjusted across iterations to account for any improvements made in prior iterations. Finally, we highlight related ideas in [28, 45] where the focus is partially observed diffusion models and the algorithms proposed therein are based on other strategies to learn a parameterized additive control directly. Such ideas have also been exploited in physics to perform rare event simulation for diffusions [40].

Our work extends these contributions in the following ways. First, these preceding works [22, 42, 47] consider only state space models. In contrast, the methodology proposed here allows us to perform inference for static models; a direct extension of these prior methods [22, 42, 47] to static models is infeasible, as it leads to algorithms which are not implementable.

Second, in contrast to the methodology in [22, 42, 47], the Kullback–Leibler optimality criterion at each iteration in our approach, is dependent on the approximately optimal proposal distributions computed at the preceding iteration; that is, we seek to minimize the residual discrepancy between any previously estimated proposals and the target distributions. This difference allows us to elucidate the effect each iteration in our method has on refining proposal distributions and improves algorithmic performance as illustrated in Section 6.2. See also [49, 50] for related iterative procedures in continuous-time optimal control approximation.

The controlled SMC methodology is one of the main contributions of this work. Another contribution is to provide a detailed theoretical analysis of various aspects of our methodology. In Proposition 5.1, we provide a backward recursion that characterizes the error of policies estimated using our ADP procedure. This error is given naturally in terms of function approximation errors with finite samples and the stability properties of the dynamic programming recursion defining the optimal policy, which is addressed in Proposition 5.2. These results show that we can obtain good approximations of the optimal policy, and hence the optimal proposal distributions, if the function classes employed are "rich" enough and the number of samples used to learn policies is sufficiently large. In Theorem 5.1, we then establish a central limit theorem for our ADP algorithm as the number of samples used in the policy learning goes to infinity. This reveals that the algorithm concentrates around an idealized ADP algorithm and provides a precise characterization of how Monte Carlo errors correlate over time. These preceding results concern a single iteration of our proposed

method and may be applied to the existing algorithms discussed above, for example, [22, 42, 47]. Using the notion of iterated random functions, we introduce a novel framework in Theorem 5.2 to understand the asymptotic behaviour of our algorithm as the number of iterations converges to infinity. This elucidates the need for iterating the ADP procedure and provides insight into the number of iterations required in practice. The discussion surrounding Theorem 5.2 also emphasizes a key difference between the newly proposed method and existing work in [22, 42, 47]. After the first version of this work appeared, a similar approach was developed for generic stochastic control problems in [23]. Our results hold under strong assumptions but appear to capture our experimental results remarkably well.

The rest of this paper is organized as follows. We introduce SMC methods in the framework of Feynman–Kac models [11] in Section 2 and twisted variants in Section 3, as this affords us the generality to cover both state space models and static models. We then identify the optimal policy that induces an optimal SMC method in Section 4.1. We describe general methods to approximate the optimal policy in Section 4.2 and develop an iterative scheme to refine policies in Section 4.3. The proposed methodology is illustrated on a neuroscience application in Section 4.4. We present the results of our analysis in Section 5 and conclude with applications in Sections 6–7. All proofs are given in the Supplementary Material [24] which also includes three additional applications. MATLAB code to reproduce all numerical results is available online.[1]

2. Motivating models and sequential Monte Carlo.

2.1. *Notation.* We first introduce notation used throughout the article. Given integers $n \leq m$ and a sequence $(x_t)_{t \in \mathbb{N}}$, we define the set $[n : m] = \{n, \ldots, m\}$ and write the subsequence $x_{n:m} = (x_n, \ldots, x_m)$. When $n < m$, we use the convention $\prod_{t=m}^{n} x_t = 1$. Let $(\mathsf{E}, \mathcal{E})$ be an arbitrary measurable space. We denote the set of all finite signed measures by $\mathcal{S}(\mathsf{E})$, the set of all probability measures by $\mathcal{P}(\mathsf{E}) \subset \mathcal{S}(\mathsf{E})$ and the set of all Markov transition kernels on $(\mathsf{E}, \mathcal{E})$ by $\mathcal{M}(\mathsf{E})$. Given $\mu, \nu \in \mathcal{P}(\mathsf{E})$, we write $\mu \ll \nu$ if μ is absolutely continuous w.r.t. ν and denote the corresponding Radon–Nikodym derivative as $d\mu/d\nu$. For any $x \in \mathsf{E}$, δ_x denotes the Dirac measure at x. The set of all real-valued, $\mathcal{E}$-measurable, lower bounded, bounded or continuous functions on E are denoted by $\mathcal{L}(\mathsf{E})$, $\mathcal{B}(\mathsf{E})$ and $\mathcal{C}(\mathsf{E})$, respectively. Given $\gamma \in \mathcal{S}(\mathsf{E})$ and $M \in \mathcal{M}(\mathsf{E})$, we define $(\gamma \otimes M)(dx, dy) = \gamma(dx) M(x, dy)$ and $(M \otimes \gamma)(dx, dy) = M(y, dx)\gamma(dy)$ as finite signed measures on the product space $\mathsf{E} \times \mathsf{E}$, equipped with the product σ-algebra $\mathcal{E} \times \mathcal{E}$. Given $\gamma \in \mathcal{S}(\mathsf{E})$, $M \in \mathcal{M}(\mathsf{E})$, $\varphi \in \mathcal{B}(\mathsf{E})$, $\xi \in \mathcal{B}(\mathsf{E} \times \mathsf{E})$, we define the integral $\gamma(\varphi) = \int_{\mathsf{E}} \varphi(x)\gamma(dx)$, the signed measure $\gamma M(\cdot) = \int_{\mathsf{E}} \gamma(dx) M(x, \cdot) \in \mathcal{S}(\mathsf{E})$ and functions $M(\varphi)(\cdot) = \int_{\mathsf{E}} \varphi(y) M(\cdot, dy) \in \mathcal{B}(\mathsf{E})$, $M(\xi)(\cdot) = \int_{\mathsf{E}} \xi(\cdot, y) M(\cdot, dy) \in \mathcal{B}(\mathsf{E})$.

2.2. *Feynman–Kac models.* We begin by introducing Feynman–Kac models [11] and defer a detailed discussion of their applications to Sections 2.3–2.4. Consider a nonhomogeneous Markov chain of length $T + 1 \in \mathbb{N}$ on a measurable space $(\mathsf{X}, \mathcal{X})$, associated with an initial distribution $\mu \in \mathcal{P}(\mathsf{X})$, and a collection of Markov transition kernels $M_t \in \mathcal{M}(\mathsf{X})$ for $t \in [1 : T]$. We denote the law of the Markov chain on path space X^{T+1}, equipped with the product σ-algebra $\mathcal{X}^{T+1}$, with

$$(1) \qquad \mathbb{Q}(dx_{0:T}) = \mu(dx_0) \prod_{t=1}^{T} M_t(x_{t-1}, dx_t)$$

---

[1]Link: https://github.com/jeremyhengjm/controlledSMC.

and denote expectations w.r.t. $\mathbb{Q}$ by $\mathbb{E}_{\mathbb{Q}}$, whereas we write $\mathbb{E}_{\mathbb{Q}}^{t,x}$ for conditional expectations on the event $X_t = x \in \mathsf{X}$. Given a sequence of strictly positive functions $G_0 \in \mathcal{B}(\mathsf{X})$ and $G_t \in \mathcal{B}(\mathsf{X} \times \mathsf{X})$ for $t \in [1 : T]$, we define the Feynman–Kac path measure

$$(2) \qquad \mathbb{P}(\mathrm{d}x_{0:T}) = Z^{-1} G_0(x_0) \prod_{t=1}^{T} G_t(x_{t-1}, x_t) \mathbb{Q}(\mathrm{d}x_{0:T}),$$

where $Z := \mathbb{E}_{\mathbb{Q}}[G_0(X_0) \prod_{t=1}^{T} G_t(X_{t-1}, X_t)]$ denotes the normalizing constant. Equation (2) can be understood as the probability measure obtained by repartitioning the probability mass of $\mathbb{Q}$ with the potential functions $(G_t)_{t \in [0:T]}$.

To examine the time evolution of (2), we define the following sequence of positive signed measures $\gamma_t \in \mathcal{S}(\mathsf{X})$ for $t \in [0 : T]$ by

$$(3) \qquad \gamma_t(\varphi) = \mathbb{E}_{\mathbb{Q}}\left[\varphi(X_t) G_0(X_0) \prod_{s=1}^{t} G_s(X_{s-1}, X_s) \right]$$

and their normalized counterparts $\eta_t \in \mathcal{P}(\mathsf{X})$ by

$$(4) \qquad \eta_t(\varphi) = \gamma_t(\varphi)/Z_t$$

for $\varphi \in \mathcal{B}(\mathsf{X})$, $t \in [0 : T]$, where $Z_t := \gamma_t(\mathsf{X})$. Equations (3) and (4) are known as the unnormalized and normalized (updated) Feynman–Kac models respectively [11], Definition 2.3.2. These models are determined by the triple $\{\mu, (M_t)_{t \in [1:T]}, (G_t)_{t \in [0:T]}\}$, which depends on the specific application of interest. The measure η_T is the terminal time marginal distribution of $\mathbb{P}$ and $Z = Z_T = \mu(G_0) \prod_{t=1}^{T} \eta_{t-1}(M_t(G_t))$.

2.3. *State space models.* Consider an X-valued hidden Markov chain $(X_t)_{t \in [0:T]}$, whose law on $(\mathsf{X}^{T+1}, \mathcal{X}^{T+1})$ is given by

$$\mathbb{H}(\mathrm{d}x_{0:T}) = \nu(\mathrm{d}x_0) \prod_{t=1}^{T} f_t(x_{t-1}, \mathrm{d}x_t),$$

where $\nu \in \mathcal{P}(\mathsf{X})$ and $f_t \in \mathcal{M}(\mathsf{X})$ for $t \in [1 : T]$. The Y-valued observations $(Y_t)_{t \in [0:T]}$ are assumed to be conditionally independent given $(X_t)_{t \in [0:T]}$ and the conditional distribution of Y_t has a strictly positive density $g_t(X_t, \cdot)$ with $g_t \in \mathcal{B}(\mathsf{X} \times \mathsf{Y})$ for $t \in [0 : T]$. Here, $\{\nu, (f_t)_{t \in [1:T]}, (g_t)_{t \in [0:T]}\}$ can potentially depend on unknown static parameters $\theta \in \Theta$, but this is notationally omitted for simplicity. Given access to a realization $y_{0:T} \in \mathsf{Y}^{T+1}$ of the observation process, statistical inference for these models relies on the marginal likelihood of $y_{0:T}$ given θ,

$$Z(y_{0:T}) = \mathbb{E}_{\mathbb{H}}\left[\prod_{t=0}^{T} g_t(X_t, y_t) \right],$$

and/or the smoothing distribution, that is, the conditional distribution of $X_{0:T}$ given $Y_{0:T} = y_{0:T}$ and θ

$$(5) \qquad \mathbb{P}(\mathrm{d}x_{0:T} | y_{0:T}) = Z(y_{0:T})^{-1} \prod_{t=0}^{T} g_t(x_t, y_t) \mathbb{H}(\mathrm{d}x_{0:T}).$$

If we set $\mathbb{Q} \in \mathcal{P}(\mathsf{X}^{T+1})$ defined in (1) equal to $\mathbb{H}$, we recover the Feynman–Kac path measure representation (2) by defining $G_t(x_{t-1}, x_t) = g_t(x_t, y_t)$ for all $t \in [0 : T]$. However, this

FIG. 1.   *Number of neuron activations over $M = 50$ repeated experiments with time* (left) *and effective sample size of bootstrap particle filter with $N = 1024$ particles* (right) *for the neuroscience model with parameters $\alpha = 0.99$ and $\sigma^2 = 0.11$.*

representation is not unique. Indeed any $\mathbb{Q}$ satisfying $\mathbb{H} \ll \mathbb{Q}$ provides a Feynman–Kac path measure representation of (2) by defining the potentials

$$G_0(x_0) = \frac{d(\nu \cdot g_0)}{d\mu}(x_0),$$

$$G_t(x_{t-1}, x_t) = \frac{d(f_t \cdot g_t)(x_{t-1}, \cdot)}{dM_t(x_{t-1}, \cdot)}(x_t), \quad t \in [1:T].$$

As outlined in [16], most SMC algorithms available at present correspond to the same basic mechanism applied to different Feynman–Kac representations of a given target probability measure. The bootstrap particle filter (BPF) presented in [21] corresponds to $\mathbb{Q} = \mathbb{H}$, that is, $M_t(x_{t-1}, dx_t) = f_t(x_{t-1}, dx_t)$ for $t \in [1:T]$, while the popular 'fully adapted' auxiliary particle filter (APF) of [41] uses $M_t(x_{t-1}, dx_t) = \mathbb{P}(dx_t|x_{t-1}, y_t) \propto f_t(x_{t-1}, dx_t)g_t(x_t, y_t)$.

As a motivating example, we consider a model for $T + 1 = 3000$ measurements collected from a neuroscience experiment [48]. The observation $y_t \in \mathsf{Y} = [0:M]$ at each time instance $t \in [0:T]$, shown in the left panel of Figure 1, represents the number of neuron activations over $M = 50$ repeated experiments and is modelled as a binomial distribution with probability of success $p_t \in [0, 1]$. We will write its probability mass function as $y_t \mapsto \mathrm{Bin}(y_t; M, p_t)$. To model the time varying behaviour of activation probabilities, it is assumed that $p_t = \kappa(X_t)$ where $\kappa(u) := (1 + \exp(-u))^{-1}$ for $u \in \mathbb{R}$ is the logistic link function and $(X_t)_{t \in [0:T]}$ is a real-valued first-order autoregressive process. This corresponds to a time homogeneous state space model on $\mathsf{X} = \mathbb{R}$, equipped with its Borel σ-algebra $\mathcal{X} = \mathfrak{B}(\mathbb{R})$, with $\nu = \mathcal{N}(0, 1)$, $f(x_{t-1}, dx_t) = \mathcal{N}(x_t; \alpha x_{t-1}, \sigma^2)\, dx_t$, and $g(x_t, y_t) = \mathrm{Bin}(y_t; M, \kappa(x_t))$ for $t \in [1:T]$, where we denote the Gaussian distribution on $\mathbb{R}^d$ with mean vector $\xi \in \mathbb{R}^d$ and covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ by $\mathcal{N}(\xi, \Sigma)$ and its Lebesgue density by $x \mapsto \mathcal{N}(x; \xi, \Sigma)$. The parameters of this model to be inferred from data are $\theta = (\alpha, \sigma^2) \in [0, 1] \times \mathbb{R}_+$.

2.4. *Static models.*   Suppose we are interested in sampling from a target distribution $\eta(dx) = Z^{-1}\gamma(dx) \in \mathcal{P}(\mathsf{X})$ and/or estimating its normalizing constant $Z = \gamma(\mathsf{X})$. To facilitate inference, we introduce a sequence of probability measures $(\eta_t)_{t \in [0:T]}$ in $\mathcal{P}(\mathsf{X})$ that bridges a simple distribution $\eta_0 = \mu$ to the target distribution $\eta_T = \eta$ with $\eta \ll \mu$. Our implementation in Section 7 adopts the geometric path [12, 17, 39]

$$(6) \qquad \gamma_t(dx) := \mu(dx)\left(\frac{d\gamma}{d\mu}(x)\right)^{\lambda_t}, \qquad \eta_t(dx) := \gamma_t(dx)/Z_t, \quad t \in [0:T],$$

where $Z_t := \gamma_t(\mathsf{X})$ and $(\lambda_t)_{t \in [0:T]} \in [0, 1]^{T+1}$ is an increasing sequence satisfying $\lambda_0 = 0$ and $\lambda_T = 1$; see [12], Section 2.3.1, for choices in other inference settings. In order to define $\mathbb{Q}$,

we introduce a sequence of 'forward' Markov transition kernels $M_t \in \mathcal{M}(\mathsf{X})$ for $t = [1:T]$ where $\eta_{t-1} M_t$ approximates η_t. One expects the distribution $\hat{\eta} = \eta_0 M_1 \cdots M_T$ of samples drawn from a nonhomogeneous Markov chain with initial distribution η_0 and transition kernels $(M_t)_{t \in [1:T]}$ to be close to $\eta_T = \eta$. However, importance sampling cannot be employed to correct for the discrepancy between $\hat{\eta}$ and η, as $\hat{\eta}$ is typically analytically intractable.

SMC samplers described in [12] circumvent this difficulty by performing importance sampling on path space $(\mathsf{X}^{T+1}, \mathcal{X}^{T+1})$ using an artificial extended target distribution of the form

$$\mathbb{P}(\mathrm{d}x_{0:T}) = \eta(\mathrm{d}x_T) \prod_{t=1}^{T} L_{t-1}(x_t, \mathrm{d}x_{t-1}),$$

where $L_t \in \mathcal{M}(\mathsf{X})$ for $t \in [0:T-1]$ is a sequence of auxiliary 'backward' Markov transition kernels. Assuming that we have $L_{t-1} \otimes \gamma_t \ll \gamma_{t-1} \otimes M_t$ with strictly positive and bounded Radon–Nikodym derivative for all $t \in [1:T]$, the Feynman–Kac path measure representation (2) can be recovered by defining

$$(7) \qquad G_0(x_0) = 1, \qquad G_t(x_{t-1}, x_t) = \frac{\mathrm{d}(L_{t-1} \otimes \gamma_t)}{\mathrm{d}(\gamma_{t-1} \otimes M_t)}(x_{t-1}, x_t), \quad t \in [1:T].$$

Under these potentials, the normalized Feynman–Kac models (4) act as the sequence of bridging distributions $(\eta_t)_{t \in [0:T]}$ in this setting. In annealed importance sampling (AIS) [39] and the sequential sampler proposed in [8], one selects $M_t \in \mathcal{M}(\mathsf{X})$ as a Markov chain Monte Carlo (MCMC) kernel that is η_t-invariant and $L_{t-1} \in \mathcal{M}(\mathsf{X})$ as its time reversal, that is, $L_{t-1} \otimes \eta_t = \eta_t \otimes M_t$, so the potentials in (7) simplify to

$$(8) \qquad G_0(x_0) = 1, \qquad G_t(x_{t-1}) = \frac{\gamma_t(x_{t-1})}{\gamma_{t-1}(x_{t-1})}, \quad t \in [1:T].$$

3. Twisted models and sequential Monte Carlo.

3.1. *Twisted Feynman–Kac models.* SMC methods can perform poorly when the discrepancy between $\mathbb{P}$ and $\mathbb{Q}$ is large. The right panel of Figure 1 illustrates that this is the case when we employ BPF on the neuroscience application in Section 2.3: the effective sample size (ESS), a common criterion used to assess the quality of a particle approximation ([32], pp. 34–35) falls below 20% when the data change abruptly. This is because the kernel $M_t(x_{t-1}, \mathrm{d}x_t) = f_t(x_{t-1}, \mathrm{d}x_t)$ used to sample particles at time t does not take the observations into account. Better performance could be obtained using observation-dependent kernels. Indeed, in the context of state space models, the smoothing distribution (5) can be written as $\mathbb{P}(\mathrm{d}x_{0:T} | y_{0:T}) = \mathbb{P}(\mathrm{d}x_0 | y_{0:T}) \prod_{t=1}^{T} \mathbb{P}(\mathrm{d}x_t | x_{t-1}, y_{t:T})$ with

$$(9) \qquad \begin{aligned} \mathbb{P}(\mathrm{d}x_0 | y_{0:T}) &= \frac{\nu(\mathrm{d}x_0) \psi_0^*(x_0)}{\nu(\psi_0^*)}, \\ \mathbb{P}(\mathrm{d}x_t | x_{t-1}, y_{t:T}) &= \frac{f_t(x_{t-1}, \mathrm{d}x_t) \psi_t^*(x_t)}{f_t(\psi_t^*)(x_{t-1})}, \quad t \in [1:T], \end{aligned}$$

where the kernel $f_t(x_{t-1}, \cdot)$ is *twisted* using the so-called backward information filter [5, 6], given by $\psi_t^*(x_t) = \mathbb{P}(y_{t:T} | x_t)$, for $t \in [0:T]$.

The backward information filter can also be defined using the backward recursion

$$(10) \qquad \begin{aligned} \psi_T^*(x_T) &= g_T(x_T, y_T), \\ \psi_t^*(x_t) &= g_t(x_t, y_t) f_{t+1}(\psi_{t+1}^*)(x_t), \quad t \in [0:T-1]. \end{aligned}$$

We can exploit this to obtain an approximation $\hat{\psi}_t(x_t), t \in [0 : T]$ using regression [22, 42, 47]. We can then sample particles at time t using a proposal $M_t^{\hat{\psi}}(x_{t-1}, dx_t) \propto f_t(x_{t-1}, dx_t)\hat{\psi}_t(x_t)$ that approximates $\mathbb{P}(dx_t | x_{t-1}, y_{t:T})$.

Abstracting the above discussion from state space models to general Feynman–Kac models, where the potential G_t might depend on both x_{t-1} and x_t, motivates the following definitions.

DEFINITION 3.1 (Admissible policies).   A sequence of functions $\psi = (\psi_t)_{t \in [0:T]}$ is an admissible policy if these functions are strictly positive and satisfy $\psi_0 \in \mathcal{B}(\mathsf{X})$, $\psi_t \in \mathcal{B}(\mathsf{X} \times \mathsf{X})$ for all $t \in [1 : T]$. The set of all admissible policies will be denoted as Ψ.

DEFINITION 3.2 (Twisted path measures).   Given a policy $\psi \in \Psi$ and a path measure $\mathbb{F} \in \mathcal{P}(\mathsf{X}^{T+1})$ of the form $\mathbb{F}(dx_{0:T}) = \nu(dx_0) \prod_{t=1}^T K_t(x_{t-1}, dx_t)$ for some $\nu \in \mathcal{P}(\mathsf{X})$ and $K_t \in \mathcal{M}(\mathsf{X})$ for $t \in [1 : T]$, the ψ-twisted path measure of $\mathbb{F}$ is defined as $\mathbb{F}^\psi(dx_{0:T}) = \nu^\psi(dx_0) \prod_{t=1}^T K_t^\psi(x_{t-1}, dx_t)$ where

(11)
$$\nu^\psi(dx_0) := \frac{\nu(dx_0)\psi_0(x_0)}{\nu(\psi_0)},$$
$$K_t^\psi(x_{t-1}, dx_t) := \frac{K_t(x_{t-1}, dx_t)\psi_t(x_{t-1}, x_t)}{K_t(\psi_t)(x_{t-1})}, \quad t \in [1 : T].$$

For any policy $\psi \in \Psi$, since $\mathbb{P} \ll \mathbb{Q} \ll \mathbb{Q}^\psi$ by positivity of ψ, we can rewrite the measure $\mathbb{P}$ defined in (2) as

(12)
$$\mathbb{P}(dx_{0:T}) = Z^{-1} G_0^\psi(x_0) \prod_{t=1}^T G_t^\psi(x_{t-1}, x_t) \mathbb{Q}^\psi(dx_{0:T}),$$

where the twisted potentials associated with the twisted path measure $\mathbb{Q}^\psi$ are given by

(13)
$$G_0^\psi(x_0) := \frac{\mu(\psi_0)G_0(x_0)M_1(\psi_1)(x_0)}{\psi_0(x_0)},$$
$$G_t^\psi(x_{t-1}, x_t) := \frac{G_t(x_{t-1}, x_t)M_{t+1}(\psi_{t+1})(x_t)}{\psi_t(x_{t-1}, x_t)}, \quad t \in [1 : T - 1],$$
$$G_T^\psi(x_{T-1}, x_T) := \frac{G_T(x_{T-1}, x_T)}{\psi_T(x_{T-1}, x_T)}.$$

Note from (12) that $Z = \mathbb{E}_{\mathbb{Q}^\psi}[G_0^\psi(X_0) \prod_{t=1}^T G_t^\psi(X_{t-1}, X_t)]$ by construction, whereas the triple $\{\mu^\psi, (M_t^\psi)_{t \in [1:T]}, (G_t^\psi)_{t \in [0:T]}\}$ induces the ψ-twisted Feynman–Kac models given by

(14)
$$\gamma_t^\psi(\varphi) = \mathbb{E}_{\mathbb{Q}^\psi}\left[\varphi(X_t)G_0^\psi(X_0) \prod_{s=1}^t G_s^\psi(X_{s-1}, X_s)\right],$$
$$\eta_t^\psi(\varphi) = \gamma_t^\psi(\varphi)/Z_t^\psi,$$

for $\varphi \in \mathcal{B}(\mathsf{X})$, $t \in [0 : T]$, where $Z_t^\psi := \gamma_t^\psi(\mathsf{X})$. For $t \in [0 : T - 1]$, the marginal distributions of the twisted model are given by

(15)
$$\eta_t^\psi(dx_t) = \eta_t(dx_t)M_{t+1}(\psi_{t+1})(x_t)Z_t/Z_t^\psi$$

and do not generally coincide with the ones of the original model (4). However, we stress that they coincide at time T as

$$(16) \qquad Z = Z_T^{\psi} = \mu^{\psi}(G_0^{\psi}) \prod_{t=1}^{T} \eta_{t-1}^{\psi}(M_t^{\psi}(G_t^{\psi})).$$

To illustrate the effect of twisting models in the static setting of Section 2.4, rewriting the twisted potentials (13) using (15) as

$$G_0^{\psi}(x_0) = \frac{\mathrm{d}\eta_0^{\psi}}{\mathrm{d}\mu^{\psi}}(x_0), \qquad G_t^{\psi}(x_{t-1}, x_t) = \frac{\mathrm{d}(L_{t-1} \otimes \gamma_t^{\psi})}{\mathrm{d}(\gamma_{t-1}^{\psi} \otimes M_t^{\psi})}(x_{t-1}, x_t), \quad t \in [1:T],$$

shows that this corresponds to employing the same backward kernels $(L_t)_{t \in [0:T-1]}$, but altered bridging distributions $(\eta_t^{\psi})_{t \in [0:T]}$, initial distribution μ^{ψ} and forward kernels $(M_t^{\psi})_{t \in [1:T]}$.

3.2. *Twisted sequential Monte Carlo.* Consider a policy $\psi \in \Psi$ such that sampling from the initial distribution $\mu^{\psi} \in \mathcal{P}(\mathsf{X})$ and the transition kernels $(M_t^{\psi})_{t \in [1:T]}$ in $\mathcal{M}(\mathsf{X})$ is feasible and evaluation of the twisted potentials (13) is tractable. We can now construct the ψ-twisted SMC method as simply the standard sampling-resampling SMC algorithm applied to ψ-twisted Feynman–Kac models [16]. The resulting algorithm provides approximations of the probability measures $(\eta_t^{\psi})_{t \in [0:T]}$, normalizing constant Z and path measure $\mathbb{P}$, by simulating an interacting particle system of size $N \in \mathbb{N}$. An algorithmic description is detailed in Algorithm 1, where $\mathcal{R}(w_1, \ldots, w_N)$ refers to a resampling operation based on a vector of unnormalized weights $(w_n)_{n \in [1:N]} \in \mathbb{R}_+^N$. For example, this is the categorical distribution on $[1:N]$ with probabilities $(w_n / \sum_{m=1}^{N} w_m)_{n \in [1:N]}$, when multinomial resampling is employed; other lower variance and adaptive resampling schemes can also be considered [18]. All simulations presented in this article employ the systematic resampling scheme.

Given the output of the algorithm, that is, an array of X-valued position variables $(X_t^n)_{(t,n) \in [0:T] \times [1:N]}$ and an array of $[1:N]$-valued ancestor variables $(A_t^n)_{(t,n) \in [0:T] \times [1:N]}$, we have a particle approximation of η_t^{ψ} given by the weighted random measure

$$\eta_t^{\psi,N} = \sum_{n=1}^{N} W_t^{\psi,n} \delta_{X_t^n}, \qquad W_t^{\psi,n} := \frac{G_t^{\psi}(X_{t-1}^{A_{t-1}^n}, X_t^n)}{\sum_{m=1}^{N} G_t^{\psi}(X_{t-1}^{A_{t-1}^m}, X_t^m)},$$

Algorithm 1: ψ-twisted sequential Monte Carlo

Input: number of particles $N \in \mathbb{N}$ and policy $\psi \in \Psi$.

1. At time $t = 0$ and particle $n \in [1:N]$:

    (a) sample $X_0^n \sim \mu^{\psi}$;
    (b) sample ancestor index $A_0^n \sim \mathcal{R}(G_0^{\psi}(X_0^1), \ldots, G_0^{\psi}(X_0^N))$.

2. For time $t \in [1:T]$ and particle $n \in [1:N]$:

    (a) sample $X_t^n \sim M_t^{\psi}(X_{t-1}^{A_{t-1}^n}, \cdot)$;
    (b) sample ancestor index $A_t^n \sim \mathcal{R}(G_t^{\psi}(X_{t-1}^{A_{t-1}^1}, X_t^1), \ldots, G_t^{\psi}(X_{t-1}^{A_{t-1}^N}, X_t^N))$.

Output: trajectories $(X_t^n)_{(t,n) \in [0:T] \times [1:N]}$ and ancestries $(A_t^n)_{(t,n) \in [0:T] \times [1:N]}$.

---

for $t \in [1:T]$ (similar expression for $t = 0$) and an unbiased estimator of Z resembling the form of (16)

$$(17) \qquad Z^{\psi,N} = \left\{ \frac{1}{N} \sum_{n=1}^{N} G_0^{\psi}(X_0^n) \right\} \prod_{t=1}^{T} \left\{ \frac{1}{N} \sum_{n=1}^{N} G_t^{\psi}(X_{t-1}^{A_{t-1}^n}, X_t^n) \right\}.$$

With stored trajectories [26], we can also form a particle approximation of $\mathbb{P}$ with $\mathbb{P}^{\psi,N} = N^{-1} \sum_{n=1}^{N} \delta_{X_{0:T}^n}$, where $X_{0:T}^n$ denotes the path obtained by tracing the ancestral lineage of particle X_T^n, that is, $X_{0:T}^n := (X_t^{B_t^n})_{t \in [0:T]}$ with $B_T^n := A_T^n$ and $B_t^n := A_t^{B_{t+1}^n}$ for $t \in [0:T-1]$. Many convergence results are available for these approximations as the size of the particle system N increases [11]. However, depending on the choice of $\psi \in \Psi$, the quality of these approximations may be inadequate for practical values of N; for example, the large variance of (17) often hinders its use within particle MCMC schemes [1] and the approximation $\mathbb{P}^{\psi,N}$ could degenerate quickly with T. The choice of an optimal policy is addressed in the following section.

4. Controlled sequential Monte Carlo.

4.1. *Optimal policies.* Suppose we have an arbitrary current policy $\psi \in \Psi$, initially given by a sequence of constant functions. We would like to twist the path measure $\mathbb{Q}^{\psi} \in \mathcal{P}(\mathsf{X}^{T+1})$ further with a policy $\phi \in \Psi$, so that the resulting twisted path measure $(\mathbb{Q}^{\psi})^{\phi} \in \mathcal{P}(\mathsf{X}^{T+1})$ is in some sense 'closer' to the target Feynman–Kac measure $\mathbb{P}$. Note from Definition 3.2 that $(\mathbb{Q}^{\psi})^{\phi} = \mathbb{Q}^{\psi \cdot \phi}$, where $\psi \cdot \phi = (\psi_t \cdot \phi_t)_{t \in [0:T]}$ denotes element-wise multiplication, is simply the $(\psi \cdot \phi)$-twisted path measure of $\mathbb{Q}$. From (13), the corresponding twisted potentials are given by

$$G_0^{\psi \cdot \phi}(x_0) = \frac{\mu^{\psi}(\phi_0) G_0^{\psi}(x_0) M_1^{\psi}(\phi_1)(x_0)}{\phi_0(x_0)},$$

$$(18) \qquad G_t^{\psi \cdot \phi}(x_{t-1}, x_t) = \frac{G_t^{\psi}(x_{t-1}, x_t) M_{t+1}^{\psi}(\phi_{t+1})(x_t)}{\phi_t(x_{t-1}, x_t)}, \quad t \in [1:T-1],$$

$$G_T^{\psi \cdot \phi}(x_{T-1}, x_T) = \frac{G_T^{\psi}(x_{T-1}, x_T)}{\phi_T(x_{T-1}, x_T)}.$$

The choice of ϕ that optimally refines an arbitrary policy ψ is given by the following optimality result.

PROPOSITION 4.1. *For any $\psi \in \Psi$, under the policy $\phi^* = (\phi_t^*)_{t \in [0:T]}$ defined recursively as*

$$\phi_T^*(x_{T-1}, x_T) = G_T^{\psi}(x_{T-1}, x_T),$$

$$(19) \qquad \phi_t^*(x_{t-1}, x_t) = G_t^{\psi}(x_{t-1}, x_t) M_{t+1}^{\psi}(\phi_{t+1}^*)(x_t), \quad t \in [1:T-1],$$

$$\phi_0^*(x_0) = G_0^{\psi}(x_0) M_1^{\psi}(\phi_1^*)(x_0),$$

the refined policy $\psi^* := \psi \cdot \phi^*$ satisfies the following properties:

1. *the twisted path measure $\mathbb{Q}^{\psi^*}$ coincides with the Feynman–Kac path measure $\mathbb{P}$;*

2. *the normalized Feynman–Kac model $\eta_t^{\psi^*}$ is the time t-marginal distribution of $\mathbb{P}$ and its normalizing constant $Z_t^{\psi^*} = Z$ for all $t \in [0:T]$;*

3. *the normalizing constant estimator $Z^{\psi^*,N} = Z$ almost surely for any $N \in \mathbb{N}$.*

Moreover, if $G_0^{\psi} \in \mathcal{B}(\mathsf{X})$ and $G_t^{\psi} \in \mathcal{B}(\mathsf{X} \times \mathsf{X})$ for $t \in [1:T]$ then $\phi^* \in \Psi$.

This proposition implies that SMC sampling with the optimal ψ^*-twisted version of Algorithm 1 ensures that the normalizing constant estimator is constant over the entire time horizon, and is equal to the desired normalizing constant. This follows because the SMC weights themselves are almost surely constant; one can see this by substituting the optimal choice (19) into (18). The variance of the SMC weights and the constancy of the normalizing constant estimator can both be used (as described later) as measures of performance evaluation or adaptive tuning.

In a state space context, (19) corresponds to the recursion satisfied by the backward information filter introduced in (10) when $\psi \in \Psi$ are constant functions, that is, $\mu^\psi = \mu = \nu$ and $M_t^\psi = M_t = f_t$, $t \in [1:T]$; see, for example, [5, 6].

As it can be shown that ϕ^* is the optimal policy of an associated Kullback–Leibler optimal control problem (Supplementary Material, Section 5), we shall refer to it as the optimal policy w.r.t. $\mathbb{Q}^\psi$, although the optimality properties in Proposition 4.1 only identify a policy up to normalization factors. An application of this result gives us the optimal policy $\psi^* = \psi \cdot \phi^*$ w.r.t. $\mathbb{Q}$, which is admissible if the original potentials $(G_t)_{t \in [0:T]}$ are bounded.[2]

4.2. *Approximate dynamic programming.* Equation (19) may be viewed as a dynamic programming backward recursion. The optimal policy ϕ^* w.r.t. $\mathbb{Q}^\psi$ will give rise to an optimally controlled SMC algorithm via a $\psi^* = (\psi \cdot \phi^*)$-twisted version of Algorithm 1. In all but simple cases, the recursion (19) defining ϕ^* is intractable. We now exploit the connection to optimal control by adapting numerical methods (i.e., approximate dynamic programming) for finite horizon control problems ([2], pp. 329–331) to our setup. The resulting methodology approximates ϕ^* by combining function approximation and iterating the backward recursion (19).

In the following, we will approximate $V_t^* := -\log \phi_t^*$, $t \in [0:T]$ as this corresponds to learning the optimal value functions of the associated control problem. Compared to learning optimal policies directly, as considered in [22], the latter choice is often more desirable as computing in logarithmic scale offers more numerical stability and the minimization is additionally analytically tractable in important scenarios. Moreover, this allows us to relate regression errors to performance properties of the resulting twisted SMC method in the next section.

Let $(X_t^n)_{(t,n) \in [0:T] \times [1:N]}$ and $(A_t^n)_{(t,n) \in [0:T] \times [1:N]}$ denote the trajectories and ancestries, obtained by running a ψ-twisted SMC. At time T, to approximate $V_T^* := -\log \phi_T^* = -\log G_T^\psi$, we consider the least squares problem

$$(20) \qquad \hat{V}_T = \arg \min_{\varphi \in \mathsf{F}_T} \sum_{n=1}^{N} (\varphi(X_{T-1}^{A_{T-1}^n}, X_T^n) + \log G_T^\psi(X_{T-1}^{A_{T-1}^n}, X_T^n))^2,$$

where F_T is a pre-specified function class. An approximation of ϕ_T^* can then be obtained by taking $\hat{\phi}_T := \exp(-\hat{V}_T)$. To iterate the backward recursion $\phi_{T-1}^* = G_{T-1}^\psi M_T^\psi(\phi_T^*)$, we set $\xi_{T-1} := G_{T-1}^\psi M_T^\psi(\hat{\phi}_T)$ by plugging in the approximation $\hat{\phi}_T \approx \phi_T^*$ and consider the least squares problem

$$(21) \qquad \hat{V}_{T-1} = \arg \min_{\varphi \in \mathsf{F}_{T-1}} \sum_{n=1}^{N} (\varphi(X_{T-2}^{A_{T-2}^n}, X_{T-1}^n) + \log \xi_{T-1}(X_{T-2}^{A_{T-2}^n}, X_{T-1}^n))^2,$$

where F_{T-1} is another function class to be specified. As before, we form the approximation $\hat{\phi}_{T-1} := \exp(-\hat{V}_{T-1})$. Continuing in this manner until time 0 gives us an approximation $\hat{\phi} =$

---

[2]For ease of presentation, the notion of admissibility adopted in Definition 3.1 is more stringent than necessary as nonadmissible optimal policies can still lead to valid optimal SMC methods.

Algorithm 2: Approximate dynamic programming

Input: policy $\psi \in \Psi$ and output of ψ-twisted SMC method (Algorithm 1).

1. Initialization: set $M_{T+1}^{\psi}(\hat{\phi}_{T+1})(X_T^n) = 1$ for $n \in [1:N]$.
2. For time $t \in [1:T]$:

    (a) set $\xi_t(X_{t-1}^{A_{t-1}^n}, X_t^n) = G_t^{\psi}(X_{t-1}^{A_{t-1}^n}, X_t^n) M_{t+1}^{\psi}(\hat{\phi}_{t+1})(X_t^n)$ for $n \in [1:N]$;

    (b) fit $\hat{V}_t = \arg\min_{\varphi \in \mathsf{F}_t} \sum_{n=1}^N (\varphi(X_{t-1}^{A_{t-1}^n}, X_t^n) + \log \xi_t(X_{t-1}^{A_{t-1}^n}, X_t^n))^2$;

    (c) set $\hat{\phi}_t = \exp(-\hat{V}_t)$.

3. At time $t = 0$:

    (a) set $\xi_0(X_0^n) = G_0^{\psi}(X_0^n) M_1^{\psi}(\hat{\phi}_1)(X_0^n)$ for $n \in [1:N]$;

    (b) fit $\hat{V}_0 = \arg\min_{\varphi \in \mathsf{F}_0} \sum_{n=1}^N (\varphi(X_0^n) + \log \xi_0(X_0^n))^2$;

    (c) set $\hat{\phi}_0 = \exp(-\hat{V}_0)$.

Output: policy $\hat{\phi} = (\hat{\phi}_t)_{t \in [0:T]} \in \Psi$.

$(\hat{\phi}_t)_{t \in [0:T]}$ of ϕ^*. We shall refer to this procedure as the approximate dynamic programming (ADP) algorithm and provide a detailed description in Algorithm 2.

Restricting the function classes $(\mathsf{F}_t)_{t \in [0:T]}$ to contain only lower bounded functions ensures that the estimated policy $\hat{\phi}$ lies in Ψ, hence the refined policy $\psi \cdot \hat{\phi}$ also lies in Ψ. We defer a detailed discussion on the choice of function classes and shall assume for now this is such that under the refined policy $\psi \cdot \hat{\phi} \in \Psi$, sampling from initial distribution $\mu^{\psi \cdot \hat{\phi}} \in \mathcal{P}(\mathsf{X})$, transition kernels $(M_t^{\psi \cdot \hat{\phi}})_{t \in [1:T]}$ in $\mathcal{M}(\mathsf{X})$ is feasible and evaluation of twisted potentials $(G_t^{\psi \cdot \hat{\phi}})_{t \in [0:T]}$ is tractable.

As the size of the particle system N increases, it is natural to expect $\hat{\phi}$ to converge (in a suitable sense) to a policy defined by an idealized algorithm that performs the least squares approximations in (20)–(21) using L^2-projections. This will be established in Section 5.2 for a common choice of function class. It follows that the quality of $\hat{\phi}$, as an approximation of the optimal policy ϕ^*, will depend on the number of particles N and the 'richness' of chosen function classes $(\mathsf{F}_t)_{t \in [0:T]}$. A more precise characterization of the ADP error in terms of approximate projection errors will be given in Section 5.1.

4.3. *Policy refinement.* If the recursion (19) could be performed exactly, no policy refinement would be necessary as we would initialize ψ as a policy of constant functions and obtain the optimal policy $\psi^* = \phi^*$ w.r.t. $\mathbb{Q}$. This will not be possible in practical scenarios. Given a current policy $\psi \in \Psi$, we employ ADP and obtain an approximation $\hat{\phi}$ of the optimal policy ϕ^* w.r.t. $\mathbb{Q}^{\psi}$. The residuals from the corresponding least squares approximations (20)–(21) are given by

$$\varepsilon_T^{\psi} := \log \hat{\phi}_T - \log G_T^{\psi}, \qquad \varepsilon_t^{\psi} := \log \hat{\phi}_t - \log G_t^{\psi} - \log M_{t+1}^{\psi}(\hat{\phi}_{t+1}), \quad t \in [0:T-1].$$

From (18), these residuals are related to twisted potentials of the refined policy $\psi \cdot \hat{\phi}$ via

$$(22) \qquad \log G_0^{\psi \cdot \hat{\phi}} = \log \mu^{\psi}(\hat{\phi}_0) - \varepsilon_0^{\psi}, \qquad \log G_t^{\psi \cdot \hat{\phi}} = -\varepsilon_t^{\psi}, \quad t \in [1:T].$$

Using this relation, we can monitor the efficiency of ADP via the variance of SMC weights in the $(\psi \cdot \hat{\phi})$-twisted version of Algorithm 1. It follows from (22) that the Kullback–Leibler

divergence from $(\mathbb{Q}^\psi)^{\hat{\phi}}$ to $\mathbb{P}$ is at most

$$(23) \qquad |\log \mu^\psi(\hat{\phi}_0) - \log Z| + \|\varepsilon_0^\psi\|_{L^1(\mathbb{P}_0)} + \sum_{t=1}^T \|\varepsilon_t^\psi\|_{L^1(\mathbb{P}_{t-1,t})},$$

where $\|\cdot\|_{L^1}$ denotes the L^1-norm w.r.t. the one time $(\mathbb{P}_t)_{t\in[0:T]}$ and two time $(\mathbb{P}_{t,s})_{(t,s)\in[0:T-1]\times[t+1:T]}$ marginal distributions of $\mathbb{P}$. This shows how performance of $(\psi \cdot \hat{\phi})$-twisted SMC depends on the quality of the ADP approximation of the optimal policy w.r.t. $\mathbb{Q}^\psi$.

If we further twist the path measure $\mathbb{Q}^{\psi\cdot\hat{\phi}}$ by a policy $\hat{\zeta} \in \Psi$, the subsequent ADP procedure defining $\hat{\zeta}$ would consider the least squares problem

$$(24) \qquad -\log\hat{\zeta}_T := \arg\min_{\varphi\in\mathsf{F}_T} \sum_{n=1}^N (\varphi(X_{T-1}^{A_{T-1}^n}, X_T^n) - \varepsilon_T^\psi(X_{T-1}^{A_{T-1}^n}, X_T^n))^2,$$

at time T, and for $t \in [1:T-1]$

$$(25) \qquad -\log\hat{\zeta}_t := \arg\min_{\varphi\in\mathsf{F}_t} \sum_{n=1}^N (\varphi(X_{t-1}^{A_{t-1}^n}, X_t^n) - (\varepsilon_t^\psi - \log M_{t+1}^{\psi\cdot\hat{\phi}}(\hat{\zeta}_{t+1}))(X_{t-1}^{A_{t-1}^n}, X_t^n))^2,$$

where $(X_t^n)_{(t,n)\in[0:T]\times[1:N]}$ and $(A_t^n)_{(t,n)\in[0:T]\times[1:N]}$ denote the output of $(\psi \cdot \hat{\phi})$-twisted SMC. Equations (24)–(25) reveal that it might be beneficial to have an iterative scheme to refine policies as this allows repeated least squares fitting of residuals, in the spirit of L^2-boosting methods [7]. Moreover, it follows from (22)–(23) that errors would not accumulate over iterations. The resulting iterative algorithm, summarized in Algorithm 3, will be referred to as the controlled SMC method (cSMC). The overall computational complexity is of order $I \times T \times (NC_{\text{sample}} C_{\text{evaluate}} + C_{\text{approx}})$, where $C_{\text{sample}}(d)$ is the cost of sampling from each initial distribution or transition kernel in (11), $C_{\text{evaluate}}(d)$ is the cost of evaluating each twisted potential in (13), and $C_{\text{approx}}(N, d)$ is the cost of each least squares approximation.[3] The first iteration of the algorithm would coincide with that of [22] for state space models, if regressions were computed on the natural scale; subsequent iterations differ in policy refinement

---

Algorithm 3: Controlled sequential Monte Carlo

Input: number of particles $N \in \mathbb{N}$ and iterations $I \in \mathbb{N}$.

1. Initialization: set $\psi^{(0)}$ as constant one functions.
2. For iterations $i \in [0:I-1]$:

    (a) run $\psi^{(i)}$-twisted SMC method (Algorithm 1);
    (b) perform ADP (Algorithm 2) with SMC output to obtain policy $\hat{\phi}^{(i+1)}$;
    (c) construct refined policy $\psi^{(i+1)} = \psi^{(i)} \cdot \hat{\phi}^{(i+1)}$.

3. At iteration $i = I$:

    (a) run $\psi^{(I)}$-twisted SMC method (Algorithm 1).

Output: trajectories $(X_t^n)_{(t,n)\in[0:T]\times[1:N]}$ and ancestries $(A_t^n)_{(t,n)\in[0:T]\times[1:N]}$ from $\psi^{(I)}$-twisted SMC method.

---

[3]The dependence of these costs on their arguments will depend on the specific problem of interest and the choice of function classes. As an example, C_{approx} will depend linearly on N in the case of linear least squares regression.

strategy. To maintain a coherent terminology, we will refer to the standard SMC method and ψ^*-twisted SMC method as the uncontrolled and optimally controlled SMC methods, respectively. From the output of the algorithm, we can estimate $\mathbb{P}$ with $\mathbb{P}^{\psi^{(I)},N}$ and its normalizing constant Z with $Z^{\psi^{(I)},N}$ as explained in Section 3.2.

It is possible to consider performance monitoring and adaptive tuning for the SMC sampling and the iterative policy refinement. Recalling the relationship between residuals and twisted potentials (22), we note that monitoring the variance of the SMC weights using, for example, the ESS, allows us to evaluate the effectiveness of the ADP algorithm and to identify time instances when the approximation is inadequate. We can also deduce if the estimated policy is far from optimal by comparing the behaviour of the normalizing constant estimates across time with those when the optimal policy is applied, as detailed in Proposition 4.1. When implementing Algorithm 3, the number of iterations $I \in \mathbb{N}$ can be pre-determined using preliminary runs or chosen adaptively until successive policy refinement yields no improvement in performance. For example, one can iterate policy refinement until the ESS across time achieves a desired minimum threshold and/or there is no improvement in ESS across iterations; see Section 6.2 for a numerical illustration. In Section 5.3, under appropriate regularity assumptions, we show that this iterative scheme generates a geometrically ergodic Markov chain on Ψ and characterize its unique invariant distribution. For all numerical examples considered in this article, we observe that convergence happens very rapidly, so only a small number of iterations is necessary.

4.4. *Illustration on neuroscience model.* We now apply our proposed methodology on the neuroscience model introduced in Section 2.3. We take BPF as the uncontrolled SMC method, that is, we set $\mu = \nu$ and $M_t = f$ for $t \in [1:T]$. Under the following choice of function classes,

$$(26) \qquad \mathsf{F}_t = \{\varphi(x_t) = a_t x_t^2 + b_t x_t + c_t : (a_t, b_t, c_t) \in \mathbb{R}^3\}, \quad t \in [0:T],$$

the policy $\psi^{(i)} = (\psi_t^{(i)})_{t \in [0:T]}$ at iteration $i \in [1:I]$ of Algorithm 3 has the form

$$\psi_t^{(i)}(x_t) = \exp(-a_t^{(i)} x_t^2 - b_t^{(i)} x_t - c_t^{(i)}), \quad t \in [0:T],$$

where $a_t^{(i)} := \sum_{j=1}^i a_t^j$, $b_t^{(i)} := \sum_{j=1}^i b_t^j$, $c_t^{(i)} := \sum_{j=1}^i c_t^j$ for $t \in [0:T]$ and $(a_t^{j+1}, b_t^{j+1}, c_t^{j+1})_{t \in [0:T]}$ denotes the coefficients estimated using linear least squares at iteration $j \in [0: I-1]$. Exact expressions of the twisted initial distribution, transition kernels and potentials, required to implement cSMC are given in Section 9.3 of Supplementary Material.

Figure 2 illustrates that the parameterization (26) provides a good approximation of the optimal policy. We note (left panel) the significant improvement of ESS across iterations, and see how this may be used as a measure of performance evaluation. In the right panel, we can also deduce how far the estimated policy is from optimality by observing the behaviour of normalizing constant estimates as discussed previously. Indeed, while the uncontrolled SMC approximates $Z_t = Z_t^{\psi^{(0)}} = p(y_{0:t})$, the controlled SMC scheme approximates $Z_t^{\psi^*} = p(y_{0:T})$ for all $t \in [0:T]$.

Moreover, we see from the left panel of Figure 3 that the improvement in performance is reflected in the estimated policy's ability to capture abrupt changes in the data. This plot also demonstrates the effect of policy refinement: by refitting residuals from previous iterations (24)–(25), the magnitude of estimated coefficients decreases with iterations as the residuals can be adequately approximated by simpler functions. Lastly, in the right panel of Figure 3, we illustrate the invariant distribution of coefficients estimated by cSMC using a long run of $I = 1000$ iterations, with the first 10 iterations discarded as burn-in. These plots show that the distribution concentrates as the size of the particle system N increases, which is consistent with our findings presented in Section 5.3.

FIG. 2. *Comparison of uncontrolled and controlled SMC methods in terms of effective sample size* (left) *and normalizing constant estimation* (right) *on the neuroscience model introduced in Section* 2.3. *The parameters are* $\alpha = 0.99$, $\sigma^2 = 0.11$ *and the algorithmic settings of cSMC are* $I = 3$, $N = 128$.

5. Analysis. This section considers several theoretical aspects of the proposed methodology, and may be skipped without affecting the methodological developments thus far and the experimental results that follow.

5.1. *Policy learning.* The goal of this section is to characterize the error of ADP (Algorithm 2) for learning the optimal policy (19) in terms of regression errors. We first define, for any $\mu \in \mathcal{P}(\mathsf{E})$, the set $\mathfrak{L}^2(\mu)$ of $\mathcal{E}$-measurable functions $\varphi : \mathsf{E} \to \mathbb{R}^d$ such that $\|\varphi\|_{L^2(\mu)} := (\int_{\mathsf{E}} |\varphi(x)|^2 \mu(dx))^{1/2} < \infty$, and $L^2(\mu)$ as the set of equivalence classes of functions in $\mathfrak{L}^2(\mu)$ that agree μ-almost everywhere. To simplify notation, we introduce some operators.

DEFINITION 5.1 (Bellman operators). Given $\psi \in \Psi$ such that $G_0^\psi \in \mathcal{B}(\mathsf{X})$ and $G_t^\psi \in \mathcal{B}(\mathsf{X} \times \mathsf{X})$ for $t \in [1 : T]$, we define the operators $Q_t^\psi : L^2(\nu_{t+1}^\psi) \to L^2(\nu_t^\psi)$ for $t \in [0 : T-1]$ as

$$Q_0^\psi(\varphi)(x) = G_0^\psi(x) M_1^\psi(\varphi)(x), \quad \varphi \in L^2(\nu_1^\psi),$$

$$Q_t^\psi(\varphi)(x, y) = G_t^\psi(x, y) M_{t+1}^\psi(\varphi)(y), \quad \varphi \in L^2(\nu_{t+1}^\psi),$$


FIG. 3. *Applying controlled SMC method on the neuroscience model introduced in Section* 2.3: *coefficients estimated at each iteration with* $N = 128$ *particles* (left) *and invariant distribution of coefficients with various number of particles* (right).

where $v_0^\psi := \mu^\psi \in \mathcal{P}(\mathsf{X})$ and $v_t^\psi := \eta_{t-1}^\psi \otimes M_t^\psi \in \mathcal{P}(\mathsf{X} \times \mathsf{X})$ for $t \in [1 : T]$. For notational convenience, define $Q_T^\psi(\varphi)(x, y) = G_T^\psi(x, y)$ for any φ (take v_{T+1}^ψ as an arbitrary element in $\mathcal{P}(\mathsf{X} \times \mathsf{X})$).

Although these operators are typically used to define unnormalized predictive Feynman–Kac models [11], Proposition 2.5.1, we shall adopt terminology from control literature and refer to them as Bellman operators. It can be shown that these Bellman operators are well defined and are in fact bounded linear operators; see Proposition 5.2. In this notation, we can rewrite (19) more succinctly as

$$(27) \qquad \phi_T^* = G_T^\psi, \qquad \phi_t^* = Q_t^\psi \phi_{t+1}^*, \quad t \in [0 : T - 1].$$

To understand how regression errors propagate in time, for $-1 \le s \le t \le T$, we define the Feynman–Kac semigroup $Q_{s,t}^\psi : L^2(v_{t+1}^\psi) \to L^2(v_{s+1}^\psi)$ associated to a policy $\psi \in \Psi$ as

$$(28) \qquad Q_{s,t}^\psi(\varphi) = \begin{cases} \varphi & s = t, \\ Q_{s+1}^\psi \circ \cdots \circ Q_t^\psi(\varphi) & s < t, \end{cases}$$

for $\varphi \in L^2(v_{t+1}^\psi)$. To describe regression steps taken to approximate the intractable recursion (27), we introduce the following operations.

DEFINITION 5.2 (Logarithmic projection). On a measurable space $(\mathsf{E}, \mathcal{E})$, let $v \in \mathcal{P}(\mathsf{E})$, $\xi : \mathsf{E} \to \mathbb{R}_+$ be a $\mathcal{E}$-measurable function such that $-\log \xi \in L^2(v) \cap \mathcal{L}(\mathsf{E})$, and $\mathsf{F} \subset \mathcal{L}(\mathsf{E})$ be a closed linear subspace of $L^2(v)$. We define the (F, v)-projection operator $P^v : \mathcal{B}(\mathsf{E}) \to \mathcal{B}(\mathsf{E})$ as

$$(29) \qquad P^v \xi = \exp\Big(-\arg\min_{\varphi \in \mathsf{F}} \|\varphi + \log \xi\|_{L^2(v)}^2\Big).$$

The projection theorem gives existence of a unique $P^v \xi$. We have chosen to define $-\log P^v \xi$ as the orthogonal projection of $-\log \xi$ onto F, as this corresponds to learning the optimal value functions of the associated control problem. Since projections are typically intractable, a practical implementation will involve a Monte Carlo approximation of (29).

DEFINITION 5.3 (Approximate projection). Following notation in Definition 5.2, given a consistent approximation v^N of v, that is, $v^N(\varphi) \to v(\varphi)$ almost surely for any $\varphi \in L^1(v)$, we define the approximate (F, v)-projection operator $P^{v,N} : \mathcal{B}(\mathsf{E}) \to \mathcal{B}(\mathsf{E})$ as the (F, v^N)-projection operator. We additionally assume that the function class F is such that $P^{v,N} \xi$ is a random function for all $\xi \in \mathcal{B}(\mathsf{E})$.

If $\psi \in \Psi$ is the current policy, we use the output of ψ-twisted SMC (Algorithm 1) to learn the optimal policy ϕ^*, through the empirical measures

$$(30) \qquad v_0^{\psi,N} = \frac{1}{N} \sum_{n=1}^N \delta_{X_0^n}, \qquad v_t^{\psi,N} = \frac{1}{N} \sum_{n=1}^N \delta_{(X_{t-1}^{A_{t-1}^n}, X_t^n)}, \quad t \in [1 : T],$$

which are consistent approximations of $(v_t^\psi)_{t \in [0:T]}$ [11], defined in Definition 5.1. Given pre-specified closed and linear function classes $\mathsf{F}_0 \subset L^2(v_0^\psi) \cap \mathcal{L}(\mathsf{X})$, $\mathsf{F}_t \subset L^2(v_t^\psi) \cap \mathcal{L}(\mathsf{X}^2)$, $t \in [1 : T]$, we denote the approximate (F_t, v_t^ψ)-projection operator by $P_t^{\psi,N}$ for $t \in [0 : T]$. We can now write our ADP algorithm detailed in Algorithm 2 succinctly as

$$(31) \qquad \hat{\phi}_T = P_T^{\psi,N} G_T^\psi, \qquad \hat{\phi}_t = P_t^{\psi,N} Q_t^\psi \hat{\phi}_{t+1}, \quad t \in [0 : T - 1].$$

The following result characterizes how well (31) can approximate (27).

PROPOSITION 5.1. *Suppose that we have a policy $\psi \in \Psi$, number of particles N and closed, linear function classes $\mathsf{F}_0 \subset L^2(v_0^\psi) \cap \mathcal{L}(\mathsf{X})$, $\mathsf{F}_t \subset L^2(v_t^\psi) \cap \mathcal{L}(\mathsf{X}^2)$, $t \in [1 : T]$ such that*:

[A1] the Feynman–Kac semigroup defined in (28) satisfies

$$\tag{32} \left\| Q_{s,t}^\psi(\varphi) \right\|_{L^2(v_{s+1}^\psi)} \leq C_{s,t}^\psi \|\varphi\|_{L^2(v_{t+1}^\psi)}, \quad -1 \leq s < t \leq T - 1,$$

for some $C_{s,t}^\psi \in [0, \infty)$ and all $\varphi \in L^2(v_{t+1}^\psi)$;

[A2] the approximate (F_t, v_t^ψ)-projection operator satisfies

$$\sup_{\xi \in \mathsf{S}_t^\psi} \mathbb{E}^{\psi,N} \left\| P_t^{\psi,N} \xi - \xi \right\|_{L^2(v_t^\psi)} \leq e_t^{\psi,N} < \infty,$$

where $\mathsf{S}_t^\psi := \{ Q_t^\psi \exp(-\varphi) : \varphi \in \mathsf{F}_{t+1} \}$ for $t \in [0 : T - 1]$ and $\mathsf{S}_T^\psi := \{ G_T^\psi \}$. Then the policy $\hat{\phi} \in \Psi$ generated by ADP algorithm (31) satisfies

$$\tag{33} \mathbb{E}^{\psi,N} \| \hat{\phi}_t - \phi_t^* \|_{L^2(v_t^\psi)} \leq \sum_{u=t}^T C_{t-1,u-1}^\psi e_u^{\psi,N}, \quad t \in [0 : T],$$

where $C_{t-1,t-1}^\psi = 1$ and $\mathbb{E}^{\psi,N}$ denotes expectation w.r.t. the law of the ψ-twisted SMC method (Algorithm 1).

Equation (33) reveals how function approximation errors propagate backwards in time. If the choice of function class is 'rich' enough and the number of particles is sufficiently large, then these errors can be kept small and ADP provides a good approximation of the optimal policy. If the number of particles is taken to infinity, the projection errors are driven solely by the choice of function class (as the latter dictates $e_t^{\psi,\infty}$). Moreover, observe that these errors are also modulated by stability constants of the Feynman–Kac semigroup in (32). We now establish the inequality (32). For $\varphi \in \mathcal{B}(\mathsf{E})$, we write its supremum norm as $\|\varphi\|_\infty = \sup_{x \in \mathsf{E}} |\varphi(x)|$.

PROPOSITION 5.2. *Suppose $\psi \in \Psi$ is such that $G_0^\psi \in \mathcal{B}(\mathsf{X})$, $G_t^\psi \in \mathcal{B}(\mathsf{X} \times \mathsf{X})$ for $t \in [1 : T]$ and let $\delta := \max_{t \in [0:T]} \|G_t^\psi\|_\infty$ (and $Z_{-1}^\psi := 1$). Then (32) holds with*

$$\tag{34} \begin{aligned} C_{s,t}^\psi &= \left(Z_t^\psi / Z_s^\psi \prod_{u=s+1}^t \|G_u^\psi\|_\infty \right)^{1/2} \\ &\leq (Z_t^\psi / Z_s^\psi)^{1/2} \delta^{(t-s)/2}, \quad -1 \leq s < t \leq T - 1. \end{aligned}$$

For the case $G_t^\psi(x, y) = G_t^\psi(y)$ for all $x, y \in \mathsf{X}$ and $t \in [1 : T]$, if we assume additionally for each $t \in [1 : T]$ that:

[A3] there exist $\sigma_t^\psi \in \mathcal{P}(\mathsf{X})$ and $\kappa_t^\psi \in (0, \infty)$ such that for all $x \in \mathsf{X}$ we have

$$\tag{35} M_t^\psi(x, \mathrm{d}y) \leq \kappa_t^\psi \sigma_t^\psi(\mathrm{d}y).$$

Then inequality (32) holds with

$$\tag{36} C_{s,t}^\psi = \left[\kappa_{s+2}^\psi \|G_{s+1}^\psi\|_\infty \sigma_{s+2}^\psi (Q_{s+1,t}^\psi(1)) \frac{Z_t^\psi}{Z_s^\psi} \right]^{1/2}, \quad -1 \leq s < t \leq T - 1.$$

The assumption of bounded potentials is typical in similar analyses of ADP errors [20], Section 8.3.3, and stability of SMC methods [11]. The second part of Proposition 5.2 shows that it is possible to exploit regularity properties of the transition kernels to obtain better constants $C_{s,t}^{\psi}$. Conditions such as (35) are common in the filtering literature; see, for example, [13], equation (9), and [11], Chapter 4.

5.2. *Limit theorems.* We now study the asymptotic behaviour of the ADP algorithm (31), with a current policy $\psi \in \Psi$, as the size of the particle system N grows to infinity. For a common choice of function class, we will establish convergence to a policy $\tilde{\phi} = (\tilde{\phi}_t)_{t \in [0:T]}$, defined by the idealized algorithm

$$(37) \qquad \tilde{\phi}_T = P_T^{\psi} G_T^{\psi}, \qquad \tilde{\phi}_t = P_t^{\psi} Q_t^{\psi} \tilde{\phi}_{t+1}, \quad t \in [0:T-1],$$

where P_t^{ψ} denotes the $(\mathsf{F}_t, v_t^{\psi})$-projection operator for $t \in [0:T]$. In particular, we consider logarithmic projections that are defined by linear least squares approximations; this corresponds to function classes of the form

$$(38) \qquad \mathsf{F}_t := \{\Phi_t^T \beta : \beta \in \mathbb{R}^M\}, \quad t \in [0:T],$$

where $\Phi_0 \subset L^2(v_0^{\psi}) \cap \mathcal{L}(\mathsf{X})$, $\Phi_t \subset L^2(v_t^{\psi}) \cap \mathcal{L}(\mathsf{X}^2)$, $t \in [1:T]$ are vectors of $M \in \mathbb{N}$ pre-specified basis functions. We will treat M as fixed in our analysis and refer to [20], Theorem 8.2.4, for results on how M should increase with N to balance the tradeoff between enriching (38) and the need for more samples to achieve the same estimation precision. We denote by $\tilde{\phi} := (\tilde{\phi}_t)_{t \in [0:T]}$ the policy generated by the idealized algorithm (37) where $\tilde{\phi}_t := \exp(-\Phi_t^T \beta_t^{\psi})$, β_t^{ψ} being the corresponding least squares estimate. This result builds upon the central limit theorem for particle methods established in [9, 11, 29].

THEOREM 5.1. *Consider the ADP algorithm* (31) *with current policy* $\psi \in \Psi$, *under linear least squares approximations* (38). *Under appropriate regularity conditions, for all* $x \in \mathsf{X}^{2T+1}$, *the estimated policy* $\hat{\phi}(x)$ *converges in probability to the policy* $\tilde{\phi}(x)$ *as* $N \to \infty$. *Moreover, for all* $x \in \mathsf{X}^{2T+1}$,

$$(39) \qquad \sqrt{N}(\hat{\phi}(x) - \tilde{\phi}(x)) \overset{d}{\longrightarrow} \mathcal{N}(0_{(T+1)}, \Omega^{\psi}(x))$$

for some $\Omega^{\psi} : \mathsf{X}^{2T+1} \to \mathbb{R}^{(T+1) \times (T+1)}$, *where* $\overset{d}{\longrightarrow}$ *denotes convergence in distribution and* $0_p = (0, \ldots, 0)^T \in \mathbb{R}^p$ *is the zero vector.*

A precise mathematical statement of this result and its proof are given in Section 3 of Supplementary Material. Note that the proof relies on a technical central limit theorem on path space that can be deduced in the case of multinomial resampling from [11], Theorem 9.7.1. The exact form of Ω^{ψ} reveals how errors correlate over time and suggests that we may expect the variance of the estimated policy to be larger at earlier times, due to the inherent backward nature of the ADP approximation.

5.3. *Iterated approximate dynamic programming.* We provide here a theoretical framework to understand the qualitative behaviour of policy $\psi^{(I)}$, estimated by Algorithm 3, as the number of iterations I grows to infinity. This offers a novel perspective of iterative algorithms for finite horizon optimal control problems and may be of general interest.

To do so, we require the set of all admissible policies to be a complete separable metric space. This follows if we impose that X is a compact metric space and work with

$\Psi := \mathcal{C}(\mathsf{X}) \prod_{t=1}^{T} \mathcal{C}(\mathsf{X} \times \mathsf{X})$, equipped with the metric $\rho(\varphi, \xi) := \sum_{t=0}^{T} \|\varphi_t - \xi_t\|_\infty$ for $\varphi = (\varphi_t)_{t \in [0:T]}$, $\xi = (\xi_t)_{t \in [0:T]} \in \Psi$; noncompact state spaces can also be accommodated with a judicious choice of metric (see, e.g., [4], p. 380).

We begin by writing the iterative algorithm with $N \in \mathbb{N}$ particles as an iterated random function $F^N : \mathsf{U} \times \Psi \to \Psi$, defined by $F_U^N(\psi) = \psi \cdot \hat{\phi}$, where $\hat{\phi}$ is the output of ADP algorithm (31) and $U \in \mathsf{U}$ encodes all uniform random variables needed to simulate a ψ-twisted SMC method (Algorithm 1). As the uniform variables $(U^{(I)})_{I \in \mathbb{N}}$ used at every iteration are independent and identically distributed, iterating F^N defines a Markov chain $(\psi^{(I)})_{I \in \mathbb{N}}$ on Ψ. We will write $\mathbb{E}$ to denote expectation w.r.t. the law of $(U^{(I)})_{I \in \mathbb{N}}$ and $\pi^{(I)} \in \mathcal{P}(\Psi)$ to denote the law of $\psi^{(I)}$. Similarly, we denote the iterative scheme with exact projections by $F : \Psi \to \Psi$, defined as $F(\psi) = \psi \cdot \tilde{\phi}$, where $\tilde{\phi}$ is the output of the idealized ADP algorithm (37). We denote by $\varphi^* \in \Psi$ a fixed point (if it exists) of F, that is, $F(\varphi^*) = \varphi^*$. The following is based on results developed in [14].

THEOREM 5.2. *Assume that the iterated random function F^N satisfies*:

[A4] $\mathbb{E}[\rho(F_U^N(\varphi_0), \varphi_0)] < \infty$ for some $\varphi_0 \in \Psi$,

[A5] there exists a measurable function $L^N : \mathsf{U} \to \mathbb{R}_+$ with $\mathbb{E}[L_U^N] < \alpha$ for some $\alpha \in [0, 1)$ such that $\rho(F_U^N(\varphi), F_U^N(\xi)) \leq L_U^N \rho(\varphi, \xi)$ for all $\varphi, \xi \in \Psi$.

Then the Ψ-valued Markov chain $(\psi^{(I)})_{I \in \mathbb{N}}$ generated by Algorithm 3 admits a unique invariant distribution $\pi \in \mathcal{P}(\Psi)$ and

$$(40) \qquad \varrho(\pi^{(I)}, \pi) \leq C(\psi^{(0)}) r^I, \quad I \in \mathbb{N},$$

for some $C : \Psi \to \mathbb{R}_+$ and $r \in (0, 1)$, where ϱ denotes the Prohorov metric on $\mathcal{P}(\Psi)$ induced by the metric ρ. If we suppose in addition that:

[A6] for each $\psi \in \Psi$, $\rho(F_U^N(\psi), F(\psi)) \leq N^{-1/2} E_U^{\psi, N}$ where $(E_U^{\psi, N})_{N \in \mathbb{N}}$ is a uniformly integrable sequence of nonnegative random variables with finite mean that converges in distribution to a limiting distribution with support on $\mathbb{R}_+$, then we also have that

$$(41) \qquad \mathbb{E}_\pi[\rho(\psi, \varphi^*)] \leq N^{-1/2} \mathbb{E}[E_U^{\varphi^*, N}](1 - \alpha)^{-1},$$

where φ^* is a fixed point of F and $\mathbb{E}_\pi$ denotes expectation w.r.t. $\psi \sim \pi$.

Assumption A5 requires the ADP procedure to be sufficiently regular; that is, for two policies $\varphi, \xi \in \Psi$ that are close, given the same uniform random variables U to simulate a φ-twisted and ξ-twisted SMC method, the policies $\hat{\varphi}$ (w.r.t. $\mathbb{Q}^\varphi$) and $\hat{\xi}$ (w.r.t. $\mathbb{Q}^\xi$) estimated by (31) should also be close enough to keep the Lipschitz constant L_U^N small. Assumption A6 is necessary to quantify the Monte Carlo error involved when employing approximate projections and can be deduced for example using the central limit theorem in (39). See Section 4 of the Supplementary Material for a discussion on when and why contraction occurs, and a simple example where Assumptions A4–A6 are verified.

The first part of Theorem 5.2, which establishes existence of a unique invariant distribution and geometric convergence to the latter, follows from standard theory on iterated random functions; see, for example, [14]. The second conclusion of Theorem 5.2, which provides a characterization of the limiting distribution, is to the best of our knowledge novel. A fixed point φ^* can be interpreted as a policy for which subsequent refinement using exact (i.e., with $N \to \infty$) projections onto the same function classes yields no change.

FIG. 4. *Assessing performance on the neuroscience model introduced in Section* 2.3 *based on* 100 *independent repetitions of each algorithm*: *sample relative variance of smoothing expectation* (left) *and log-marginal likelihood estimates* (right).

6. Application to state space models.

6.1. *Neuroscience model.* We return to the neuroscience model introduced in Section 2.3 and explore cSMC's utility as a smoother, with algorithmic settings described in Section 4.4, in comparison to the forward filtering backward smoothing (FFBS) procedure of [15, 30]. We consider an approximation of the maximum likelihood estimate (MLE) $(\alpha, \sigma^2) = (0.99, 0.11)$ as parameter value and the smoothing functional $x_{0:T} \mapsto M(\kappa(x_0), \ldots, \kappa(x_T))$ whose expectation represents the expected number of neuron activations at each time. Although BPF's particle approximation of the smoothing distribution degenerates quickly in time, cSMC with $I = 3$ iterations offers a marked improvement; for example, the number of distinct ancestors at the initial time is on average 63 times that of BPF. We use $N = 1024$ particles for cSMC and select the number of particles in FFBS to match compute time. The results, displayed in the left panel of Figure 4, show some gains over FFBS and especially so at later times.

We then investigate the relative variance of the log-marginal likelihood estimates obtained using cSMC and BPF in a neighbourhood of the approximate MLE. As the marginal likelihood surface is rather flat in α, we fix $\alpha = 0.99$ and vary $\sigma^2 \in \{0.01, 0.02, \ldots, 0.2\}$. We use $I = 3$ iterations, $N = 128$ particles for cSMC and $N = 5529$ particles for BPF to match computational cost. The results, reported in the right panel of Figure 4, demonstrate that while the relative variance of estimates produced by BPF increases exponentially as σ^2 decreases, that of cSMC is stable across the values of σ^2 considered.

Lastly, we perform Bayesian inference on the unknown parameters $\theta = (\alpha, \sigma^2)$ and compare the efficiency of cSMC and BPF within a particle marginal Metropolis–Hastings (PMMH) algorithm [1]. We specify a uniform prior on $[0, 1]$ for α and an independent inverse-Gamma prior distribution $\mathcal{IG}(1, 0.1)$ for σ^2. Initializing at $\theta = (0.99, 0.11)$, we run two PMMH chains $(\theta_k^{cSMC})_{k \in [0:K]}$, $(\theta_k^{BPF})_{k \in [0:K]}$ of length $K = 100,000$. Both chains are updated using an independent Gaussian random walk proposal with standard deviation $(0.002, 0.01)$, but rely on cSMC or BPF to produce unbiased estimates of the marginal likelihood when computing acceptance probabilities. To ensure a fair comparison, we use $I = 3$ iterations and $N = 128$ particles for cSMC which matches the compute time taken by BPF with $N = 5529$ particles, so that both PMMH chains require very similar computational cost. The autocorrelation functions of each PMMH chain, shown in Figure 5, reveal that the $(\theta_k^{BPF})_{k \in [0:K]}$ chain has poorer mixing properties. These differences can be summarized by the effective sample size, computed as the length of the chain K divided by the estimated integrated autocorrelation time for each parameter of interest, which was found to be $(4356, 2442)$ for $(\theta_k^{BPF})_{k \in [0:K]}$ and $(20,973, 13,235)$ for $(\theta_k^{cSMC})_{k \in [0:K]}$.

FIG. 5. *Autocorrelation functions of PMMH chains, with marginal likelihood estimates produced by cSMC or BPF, for parameters of the neuroscience model introduced in Section* 2.3.

6.2. *The Lorenz-96 model.* Following [38], we consider the Lorenz-96 model [34] in a low noise regime, that is, the Itô process $\xi(s) = (\xi_i(s))_{i \in [1:d]}$, $s \geq 0$ defined as the weak solution of the stochastic differential equation

$$(42) \qquad d\xi_i = (-\xi_{i-1}\xi_{i-2} + \xi_{i-1}\xi_{i+1} - \xi_i + \alpha)\,dt + \sigma_f dB_i, \quad i \in [1:d],$$

where indices should be understood modulo d, $\alpha \in \mathbb{R}$ is a forcing parameter, $\sigma_f^2 \in \mathbb{R}_+$ is a noise parameter and $B(s) = (B_i(s))_{i \in [1:d]}$, $s \geq 0$ is a d-dimensional standard Brownian motion. The initial condition is taken as $\xi(0) \sim \mathcal{N}(0_d, \sigma_f^2 I_d)$. We assume that the process is observed at a regular time grid of size $h > 0$ according to $Y_t \sim \mathcal{N}(H\xi(s_t), R)$, $s_t = th$, $t \in [0:T]$, and consider the partially observed case where $H_{ii} = 1$ for $i = 1, \ldots, p$ and 0 otherwise with $p = d - 2$.

As discussed in [38], an efficient discretization scheme in this low noise regime [36], Chapter 3, is given by adding Brownian increments to the output of a high-order numerical integration scheme on the drift of (42). Incorporating time discretization gives a time homogenous state space model on $(\mathsf{X}, \mathcal{X}) = (\mathbb{R}^d, \mathfrak{B}(\mathbb{R}^d))$ with $\nu = \mathcal{N}(0_d, \sigma_f^2 I_d)$, $f(x_{t-1}, dx_t) = \mathcal{N}(x_t; q(x_{t-1}), \sigma_f^2 h I_d)\,dx_t$ and $g(x_t, y_t) = \mathcal{N}(y_t; Hx_t, R)$ for $t \in [1:T]$, where $y_{0:T} \in \mathsf{Y}^{T+1} = (\mathbb{R}^p)^{T+1}$ is a realization of the observation process and $q : \mathsf{X} \to \mathsf{X}$ denotes the mapping induced by a fourth-order Runge–Kutta (RK4) method on $[0, h]$. We will take noise parameters as $\sigma_f^2 = 10^{-2}$, $R = \sigma_g^2 I_p$, observe the process for 10 time units, that is, set $h = 0.1$, $T = 100$ and implement RK4 with a step size of 10^{-2}. For this application, we can employ the fully adapted APF as uncontrolled SMC method [41], that is, set $\mu = \nu^\psi$ and $M_t = f^\psi$ for $t \in [1:T]$ with policy $\psi_t = g$, $t \in [0:T]$.

Our ADP approximation will utilize the function classes

$$(43) \qquad \mathsf{F}_t = \{\varphi(x_t) = x_t^T A_t x_t + x_t^T b_t + c_t : (A_t, b_t, c_t) \in \mathbb{S}_d \times \mathbb{R}^d \times \mathbb{R}\}, \quad t \in [0:T],$$

where $\mathbb{S}_d = \{A \in \mathbb{R}^{d \times d} : A = A^T\}$. Under this parameterization, the policy $\psi^{(i)} = (\psi_t^{(i)})_{t \in [0:T]}$ at iteration $i \in [1:I]$ of Algorithm 3 is given by

$$(44) \qquad -\log \psi_t^{(i)}(x_t) = x_t^T A_t^{(i)} x_t + x_t^T b_t^{(i)} + c_t^{(i)}, \quad t \in [0:T],$$

where $A_t^{(i)} := \sum_{j=1}^i A_t^j$, $b_t^{(i)} := \sum_{j=1}^i b_t^j$, $c_t^{(i)} := \sum_{j=1}^i c_t^j$ for $t \in [0:T]$ and $(A_t^{j+1}, b_t^{j+1}, c_t^{j+1})_{t \in [0:T]}$ denotes coefficients estimated using linear least squares at iteration $j \in [0 : I - 1]$. Having APF as uncontrolled SMC is also equivalent to taking BPF as uncontrolled with an initial policy $\psi^{(0)} = (\psi_t^{(0)})_{t \in [0:T]}$ of the form (44) with $A_t^{(0)} := \frac{1}{2}\sigma_g^{-2} H^T H$, $b_t^{(0)} :=$

$-\sigma_g^{-2} H^T y_t$ and $c_t^{(0)} := \frac{1}{2}\sigma_g^{-2} y_t^T y_t + \frac{1}{2} p \log(2\pi) + \frac{1}{2} p \log(\sigma_g^2)$ for $t \in [1:T]$. For $A \in \mathbb{S}_d$, the notation $A \succ 0$ refers to A being positive definite. If the constraints $(\sigma_f^{-2} I_d + 2A_0^{(i)})^{-1} \succ 0$, $(\sigma_f^{-2} h^{-1} I_d + 2A_t^{(i)})^{-1} \succ 0, t \in [1:T]$ are satisfied or imposed,[4] then sampling from the twisted initial distribution and transition kernels is feasible and evaluation of the corresponding potentials is also tractable; see Section 9.2 of Supplementary Material for exact expressions. The diagnostics discussed in Section 4.4 indicate that (44) provides an adequate approximation of the optimal policy by adapting to the chaotic behaviour of the Lorenz system.

We begin by comparing the relative variance of the log-marginal likelihood estimates obtained by cSMC and APF, as α takes values in a regular grid between 2.5 to 8.5. We consider $d = 8$ and simulate observations under the model with $\alpha = 4.8801$, $\sigma_g^2 = 10^{-4}$. We employ $N = 512$ particles and the following adaptive strategy within cSMC: perform policy refinement until the minimum ESS over time is at least 90%, terminating at a maximum of 4 iterations. To ensure a fair comparison, the number of particles used in APF is chosen to match computation time. The results, plotted in the left panel of Figure 6, show that cSMC offers significant variance reduction across all values of α considered. Moreover, we see from the right panel of Figure 6 that the adaptive criterion allows us to adaptively increase the number of iterations as we move away from the data generating parameter. We then compare cSMC against the iterated APF [22], Algorithm 4, when function approximations are performed in the logarithmic scale (43). Using $N = 512$ particles and $I = 3$ iterations with the fully adapted APF as initialization for both algorithms, the sample variance of cSMC log-marginal likelihood estimates at $\alpha = 4.8801$, based on 1000 independent repetitions, was smaller than iterated APF at each iteration $i \in [1:3]$, with a relative ratio of $\{0.99, 0.94, 0.92\}$, respectively.

Next, we consider configurations $(d, \sigma_g^2) \in \{8, 16, 32, 64\} \times \{10^{-4}, 10^{-3}, 10^{-2}\}$ with $\alpha = 4.8801$ and generate observations under the model. We use $I = 1$ iteration for cSMC in all configurations and increase the number of particles N with d for both algorithms. As before, N is chosen so that both methods require the same compute time to ensure a fair comparison. The relative variance of both methods are reported in Table 1. These results indicate several order of magnitude gains over APF in all configurations considered.


FIG. 6. *Lorenz*-96 *model of Section* 6.2 *with data generating parameter* $\alpha = 4.8801$: *sample relative variance of log-marginal likelihood estimates based on* 100 *independent repetitions of each algorithm* (left), *average number of iterations taken by cSMC with adaptation* (right).

---

[4]In our numerical implementation, we find that these constraints are already satisfied when the step size h is sufficiently small. Otherwise, they can be imposed by projecting onto the set of real symmetric positive definite matrices using the algorithm in [25].

TABLE 1
Algorithmic settings and performance of APF and cSMC for each dimension d and observation noise σ_g^2 considered. Notationally, N refers to the number of particles and RVAR is the sample relative variance of log-marginal likelihood estimates over 100 independent repetitions of each method

				Observation noise		
				$\sigma_g^2 = 10^{-4}$	$\sigma_g^2 = 10^{-3}$	$\sigma_g^2 = 10^{-2}$
			N	$\log_{10}(\text{RVAR})$	$\log_{10}(\text{RVAR})$	$\log_{10}(\text{RVAR})$
Algorithm	APF	$d = 8$	1382	-6.7263	-5.6823	-4.4061
		$d = 16$	2027	-7.4056	-5.9009	-4.4719
		$d = 32$	4034	-7.5943	-5.4901	-4.1039
		$d = 64$	11,468	-7.5173	-5.3765	-3.1057
	cSMC	$d = 8$	512	-11.1252	-10.4173	-8.66563
		$d = 16$	512	-11.8899	-11.1011	-9.29596
		$d = 32$	1024	-12.5804	-11.8622	-9.6577
		$d = 64$	4096	-13.5959	-12.7691	-9.74631

7. Application to static models. We now detail how the proposed methodology can be applied to static models described in Section 2.4. The framework introduced in [12] generalizes the AIS method of [39] and the sequential sampler of [8] by allowing arbitrary forward and backward kernels instead of being restricted to MCMC kernels. This degree of freedom is useful here as sampling from twisted MCMC kernels and computing integrals w.r.t. these kernels is typically impossible.

7.1. *Setup.* We consider the Bayesian framework where the target distribution of interest is a posterior distribution $\eta(dx) = Z^{-1}\mu(dx)\ell(x, y)$ defined on $(\mathsf{X}, \mathcal{X}) = (\mathbb{R}^d, \mathfrak{B}(\mathbb{R}^d))$, given by a Bayes update with a prior distribution $\mu \in \mathcal{P}(\mathsf{X})$ and a likelihood function $\ell : \mathsf{X} \times \mathsf{Y} \to \mathbb{R}_+$. In applications, the marginal likelihood $Z(y) := \int_{\mathsf{X}} \mu(dx)\ell(x, y)$ of observations $y \in \mathsf{Y}$ is often also a quantity of interest. Assuming η has a strictly positive and continuously differentiable density $x \mapsto \eta(x)$ w.r.t. Lebesgue measure on $\mathbb{R}^d$, we select the forward kernel M_t related to the transition kernel of an unadjusted Langevin algorithm (ULA) [43, 44] targeting η_t defined in (6), for example, we will define $M_t(x_{t-1}, dx_t) = \mathcal{N}(x_t; x_{t-1} + \frac{1}{2}h\Gamma\nabla\log\eta_t(x_{t-1}), h\Gamma)\,dx_t$ where $h > 0$ denotes the step size, and Γ is a positive definite pre-conditioning matrix (which in the simplest case may be the identity $\Gamma = I_d$).

Under appropriate regularity conditions, for sufficiently small h, M_t admits an invariant distribution that is close to η_t [35]. Moreover, as the corresponding Langevin diffusion is η_t-reversible, this suggests that M_t will also be approximately η_t-reversible for small h. This prompts the choice of backward kernel $L_{t-1}(x_t, dx_{t-1}) = M_t(x_t, dx_{t-1})$, in which case, we expect the potentials (7) to be close to (8) when the step size is small. We have limited the scope of this article to overdamped Langevin dynamics; future work could consider the use of generalized Langevin dynamics and other nonreversible dynamics.

7.2. *Log-Gaussian Cox point process.* We end with a challenging high dimensional application of Bayesian inference for log-Gaussian Cox point processes on a dataset[5] concerning the locations of 126 Scots pine saplings in a natural forest in Finland [10, 19, 37]. The actual square plot of 10×10 square metres is standardized to the unit square and locations are plotted in the left panel of Figure 7. We then discretize $[0, 1]^2$

[5]The dataset can be found in the R package spatstat as finpines.

FIG. 7. *Locations of* 126 *Scots pine saplings in a natural forest in Finland* (left) *and log-marginal likelihood estimates obtained with* 100 *independent repetitions of cSMC, standard AIS and adaptive AIS* (right).

into a 30×30 regular grid. Given a latent intensity process $\Lambda = (\Lambda_m)_{m \in [1:30]^2}$, the number of points in each grid cell $Y = (Y_m)_{m \in [1:30]^2} \in \mathbb{N}^{30^2}$ are modelled as conditionally independent and Poisson distributed with means $a\Lambda_m$, where $a = 1/30^2$ is the area of each grid cell. The prior distribution for Λ is specified by $\Lambda_m = \exp(X_m)$, $m \in [1:30]^2$, where $X = (X_m)_{m \in [1:30]^2}$ is a Gaussian process with constant mean $\mu_0 \in \mathbb{R}$ and exponential covariance function $\Sigma_0(m, n) = \sigma^2 \exp(-|m - n|/(30\beta))$ for $m, n \in [1:30]^2$. We will adopt the parameter values $\sigma^2 = 1.91$, $\beta = 1/33$ and $\mu_0 = \log(126) - \sigma^2/2$ estimated by [37]. This application corresponds to dimension $d = 900$, a prior distribution $\mu = \mathcal{N}(\mu_0 1_d, \Sigma_0)$ with $1_d = (1, \ldots, 1)^T \in \mathbb{R}^d$ and likelihood function $\ell(x, y) = \prod_{m \in [1:30]^2} \exp(x_m y_m - a \exp(x_m))$, where $y = (y_m)_{m \in [1:30]^2} \in \mathsf{Y} = \mathbb{N}^d$ is the given dataset.

For this application, cSMC relies on pre-conditioned ULA moves with the choice of $\Gamma^{-1} = \Sigma_0^{-1} + a \exp(\mu_0 + \sigma^2/2) I_d$ considered in [19]. As the above choice of pre-conditioning captures the curvature of the posterior distribution, we adopt the following function classes:

$$\mathsf{F}_0 = \{\varphi(x_0) = x_0^T A_0 x_0 + x_0^T b_0 + c_0 : (A_0, b_0, c_0) \in \mathbb{S}_d \times \mathbb{R}^d \times \mathbb{R}\},$$

(45)
$$\mathsf{F}_t = \{\varphi(x_{t-1}, x_t) = x_t^T A_t x_t + x_t^T b_t + c_t - (\lambda_t - \lambda_{t-1}) \log \ell(x_{t-1}, y)$$

$$: (A_t, b_t, c_t) \in \mathbb{S}_d \times \mathbb{R}^d \times \mathbb{R}\}, \quad t \in [1:T],$$

where $(A_t)_{t \in [0:T]}$ are restricted to diagonal matrices to reduce the computational overhead involved in estimating large number of coefficients for a problem of this scale. The rationale for approximating the x_{t-1} dependency in $\psi_t^*(x_{t-1}, x_t)$, $t \in [1:T]$ is based on the argument that the potentials (7) would be close to that of AIS (8) for sufficiently small step size h. We refer to Section 8.1 of Supplementary Material for exact expressions required to implement cSMC. As before, the diagnostics discussed in Section 4.4 reveal that such a parameterization offers an adequate approximation of the optimal policy.

We select as competing algorithms: (1) standard AIS with pre-conditioned Metropolis-adjusted Langevin algorithm (MALA) moves and (2) an adaptive (pre-conditioned) AIS. For both cSMC and standard AIS, we adopt the geometric path (6) with $\lambda_t = t/T$ and fix the number of time steps as $T = 20$. We use $N = 4096$ particles, $I = 3$ iterations for cSMC and 5 times more particles for standard AIS to ensure that our comparison is performed at a fixed computational complexity. Using pilot runs, we chose a step size of 0.4 for MALA to achieve suitable acceptance probabilities, and a smaller step size of 0.05 for ULA as this improves the approximation in (45). For the adaptive AIS algorithm, we also adopt (6) with λ_t adapted so that the ESS% is maintained above 80% [27, 46, 51] and with an adaptive step size chosen to ensure an acceptance probability within the range of 30% to 50% at each time step [3, 27].

Since the runtime of adaptive AIS is random, we choose the number of particles to ensure the averaged computational cost matches that of cSMC and standard AIS; this is typically on the order of 2 times as many particles as cSMC.

The results obtained show that standard AIS performs poorly in this scenario, providing high variance estimates of the log-marginal likelihood compared to each iteration of cSMC, as displayed in the right panel of Figure 7. Adaptive AIS performs better than standard AIS but it is still outperformed by cSMC. The sample variance of log-marginal likelihood estimates is 573 times smaller for the last iteration of cSMC compared to standard AIS, and it is 200 times smaller compared to adaptive AIS. The mean squared error[6] of adaptive AIS algorithm is 920 times larger than that of cSMC.

SUPPLEMENTARY MATERIAL

Supplementary material for "Controlled sequential Monte Carlo" (DOI: 10.1214/19-AOS1914SUPP; .pdf). The supplement contains proofs of all results, a detailed description of the connection to Kullback–Leibler control, three more applications employing other flexible function classes and some model specific expressions.

REFERENCES

[1] ANDRIEU, C., DOUCET, A. and HOLENSTEIN, R. (2010). Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **72** 269–342. MR2758115 https://doi.org/10.1111/j.1467-9868.2009.00736.x

[2] BERTSEKAS, D. P. and TSITSIKLIS, J. N. (1996). *Neuro-dynamic Programming*. Athena Scientific, Belmont, MA.

[3] BESKOS, A., JASRA, A., KANTAS, N. and THIERY, A. (2016). On the convergence of adaptive sequential Monte Carlo methods. *Ann. Appl. Probab.* **26** 1111–1146. MR3476634 https://doi.org/10.1214/15-AAP1113

[4] BICHTELER, K. (2002). *Stochastic Integration with Jumps. Encyclopedia of Mathematics and Its Applications* **89**. Cambridge Univ. Press, Cambridge. MR1906715 https://doi.org/10.1017/CBO9780511549878

[5] BRESLER, Y. (1986). Two-filter formula for discrete-time non-linear Bayesian smoothing. *Internat. J. Control* **43** 629–641.

[6] BRIERS, M., DOUCET, A. and MASKELL, S. (2010). Smoothing algorithms for state-space models. *Ann. Inst. Statist. Math.* **62** 61–89. MR2577439 https://doi.org/10.1007/s10463-009-0236-2

[7] BÜHLMANN, P. and YU, B. (2003). Boosting with the L_2 loss: Regression and classification. *J. Amer. Statist. Assoc.* **98** 324–339. MR1995709 https://doi.org/10.1198/016214503000125

[8] CHOPIN, N. (2002). A sequential particle filter method for static models. *Biometrika* **89** 539–552. MR1929161 https://doi.org/10.1093/biomet/89.3.539

[9] CHOPIN, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Ann. Statist.* **32** 2385–2411. MR2153989 https://doi.org/10.1214/009053604000000698

[10] CHRISTENSEN, O. F., ROBERTS, G. O. and ROSENTHAL, J. S. (2005). Scaling limits for the transient phase of local Metropolis–Hastings algorithms. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **67** 253–268. MR2137324 https://doi.org/10.1111/j.1467-9868.2005.00500.x

[11] DEL MORAL, P. (2004). *Feynman–Kac Formulae. Probability and Its Applications* (*New York*). Springer, New York. MR2044973 https://doi.org/10.1007/978-1-4684-9393-1

[12] DEL MORAL, P., DOUCET, A. and JASRA, A. (2006). Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **68** 411–436. MR2278333 https://doi.org/10.1111/j.1467-9868.2006.00553.x

[13] DEL MORAL, P. and GUIONNET, A. (1999). Central limit theorem for nonlinear filtering and interacting particle systems. *Ann. Appl. Probab.* **9** 275–297. MR1687359 https://doi.org/10.1214/aoap/1029962742

[14] DIACONIS, P. and FREEDMAN, D. (1999). Iterated random functions. *SIAM Rev.* **41** 45–76. MR1669737 https://doi.org/10.1137/S0036144598338446

---

[6]Computed by taking reference to an estimate obtained using many repetitions of a SMC sampler with a large number of particles.

[15] DOUCET, A., GODSILL, S. J. and ANDRIEU, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat*. *Comput*. **10** 197–208.

[16] DOUCET, A. and JOHANSEN, A. M. (2011). A tutorial on particle filtering and smoothing: Fifteen years later. In *The Oxford Handbook of Nonlinear Filtering* (D. Crisan and B. L. Rozovsky, eds.) 656–704. Oxford Univ. Press, Oxford. MR2884612

[17] GELMAN, A. and MENG, X.-L. (1998). Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statist*. *Sci*. **13** 163–185. MR1647507 https://doi.org/10.1214/ss/1028905934

[18] GERBER, M., CHOPIN, N. and WHITELEY, N. (2019). Negative association, ordering and convergence of resampling methods. *Ann*. *Statist*. **47** 2236–2260. MR3953450 https://doi.org/10.1214/18-AOS1746

[19] GIROLAMI, M. and CALDERHEAD, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat*. *Soc*. *Ser. B. Stat*. *Methodol*. **73** 123–214. MR2814492 https://doi.org/10.1111/j.1467-9868.2010.00765.x

[20] GOBET, E. (2016). *Monte-Carlo Methods and Stochastic Processes*: *From Linear to Non-linear*. CRC Press, Boca Raton, FL. MR3558839

[21] GORDON, N. J., SALMOND, D. and SMITH, A. F. M. (1993). A novel approach to non-linear/non-Gaussian Bayesian state estimation. *IEE Proc*. *Radar Signal Process*. **140** 107–113.

[22] GUARNIERO, P., JOHANSEN, A. M. and LEE, A. (2017). The iterated auxiliary particle filter. *J. Amer*. *Statist*. *Assoc*. **112** 1636–1647. MR3750887 https://doi.org/10.1080/01621459.2016.1222291

[23] GUPTA, A., JAIN, R. and GLYNN, P. (2018). A fixed point theorem for iterative random contraction operators over Banach spaces. Available at arXiv:1804.01195.

[24] HENG, J., BISHOP, A. N, DELIGIANNIDIS, G. and DOUCET, A. (2020). Supplement to "Controlled sequential Monte Carlo." https://doi.org/10.1214/19-AOS1914SUPP.

[25] HIGHAM, N. J. (1988). Computing a nearest symmetric positive semidefinite matrix. *Linear Algebra Appl*. **103** 103–118. MR0943997 https://doi.org/10.1016/0024-3795(88)90223-6

[26] JACOB, P. E., MURRAY, L. M. and RUBENTHALER, S. (2015). Path storage in the particle filter. *Stat*. *Comput*. **25** 487–496. MR3306720 https://doi.org/10.1007/s11222-013-9445-x

[27] JASRA, A., STEPHENS, D. A., DOUCET, A. and TSAGARIS, T. (2011). Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scand*. *J. Stat*. **38** 1–22. MR2760137 https://doi.org/10.1111/j.1467-9469.2010.00723.x

[28] KAPPEN, H. J. and RUIZ, H.-C. (2016). Adaptive importance sampling for control and inference. *J. Stat*. *Phys*. **162** 1244–1266. MR3462019 https://doi.org/10.1007/s10955-016-1446-7

[29] KÜNSCH, H. R. (2005). Recursive Monte Carlo filters: Algorithms and theoretical analysis. *Ann*. *Statist*. **33** 1983–2021. MR2211077 https://doi.org/10.1214/009053605000000426

[30] KÜNSCH, H. R. (2013). Particle filters. *Bernoulli* **19** 1391–1403. MR3102556 https://doi.org/10.3150/12-BEJSP07

[31] LIN, M., CHEN, R. and LIU, J. S. (2013). Lookahead strategies for sequential Monte Carlo. *Statist*. *Sci*. **28** 69–94. MR3075339 https://doi.org/10.1214/12-STS401

[32] LIU, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. *Springer Series in Statistics*. Springer, New York. MR1842342

[33] LIU, J. S. and CHEN, R. (1998). Sequential Monte Carlo methods for dynamic systems. *J. Amer. Statist*. *Assoc*. **93** 1032–1044. MR1649198 https://doi.org/10.2307/2669847

[34] LORENZ, E. N. (1996). Predictability: A problem partly solved. In *Proc*. *Seminar on Predictability* **1**.

[35] MATTINGLY, J. C., STUART, A. M. and HIGHAM, D. J. (2002). Ergodicity for SDEs and approximations: Locally Lipschitz vector fields and degenerate noise. *Stochastic Process*. *Appl*. **101** 185–232. MR1931266 https://doi.org/10.1016/S0304-4149(02)00150-3

[36] MILSTEIN, G. N. and TRETYAKOV, M. V. (2004). *Stochastic Numerics for Mathematical Physics*. *Scientific Computation*. Springer, Berlin. MR2069903 https://doi.org/10.1007/978-3-662-10063-9

[37] MØLLER, J., SYVERSVEEN, A. R. and WAAGEPETERSEN, R. P. (1998). Log Gaussian Cox processes. *Scand*. *J. Stat*. **25** 451–482. MR1650019 https://doi.org/10.1111/1467-9469.00115

[38] MURRAY, L. M., SINGH, S., JACOB, P. E. and LEE, A. (2016). Anytime Monte Carlo. Available at arXiv:612.03319.

[39] NEAL, R. M. (2001). Annealed importance sampling. *Stat*. *Comput*. **11** 125–139. MR1837132 https://doi.org/10.1023/A:1008923215028

[40] NEMOTO, T., BOUCHET, F., JACK, R. L. and LECOMTE, V. (2016). Population-dynamics method with a multicanonical feedback control. *Phys*. *Rev. E* **93** 062123. MR3713187 https://doi.org/10.1103/physreve.93.062123

[41] PITT, M. K. and SHEPHARD, N. (1999). Filtering via simulation: Auxiliary particle filters. *J. Amer. Statist*. *Assoc*. **94** 590–599. MR1702328 https://doi.org/10.2307/2670179

[42] RICHARD, J.-F. and ZHANG, W. (2007). Efficient high-dimensional importance sampling. *J. Econometrics* **141** 1385–1411. MR2413506 https://doi.org/10.1016/j.jeconom.2007.02.007

[43] ROBERTS, G. O. and STRAMER, O. (2002). Langevin diffusions and Metropolis–Hastings algorithms. *Methodol. Comput. Appl. Probab.* **4** 337–357. MR2002247 https://doi.org/10.1023/A:1023562417138

[44] ROBERTS, G. O. and TWEEDIE, R. L. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2** 341–363. MR1440273 https://doi.org/10.2307/3318418

[45] RUIZ, H.-C. and KAPPEN, H. J. (2017). Particle smoothing for hidden diffusion processes: Adaptive path integral smoother. *IEEE Trans. Signal Process.* **65** 3191–3203. MR3650283 https://doi.org/10.1109/TSP.2017.2686340

[46] SCHÄFER, C. and CHOPIN, N. (2013). Sequential Monte Carlo on large binary sampling spaces. *Stat. Comput.* **23** 163–184. MR3016936 https://doi.org/10.1007/s11222-011-9299-z

[47] SCHARTH, M. and KOHN, R. (2016). Particle efficient importance sampling. *J. Econometrics* **190** 133–147. MR3425276 https://doi.org/10.1016/j.jeconom.2015.03.047

[48] TEMEREANCA, S., BROWN, E. N. and SIMONS, D. J. (2008). Rapid changes in thalamic firing synchrony during repetitive whisker stimulation. *J. Neurosci.* **28** 11153–11164.

[49] THEODOROU, E. A. and TODOROV, E. (2012). Relative entropy and free energy dualities: Connections to path integral and KL control. In *Proceedings* 51*st IEEE Conference on Decision and Control* (*CDC*), 1466–1473. Oxford Univ. Press, Oxford.

[50] THIJSSEN, S. and KAPPEN, H. J. (2015). Path integral control and state-dependent feedback. *Phys. Rev. E* **91** 032104. MR3460947 https://doi.org/10.1103/PhysRevE.91.032104

[51] ZHOU, Y., JOHANSEN, A. M. and ASTON, J. A. D. (2016). Toward automatic model comparison: An adaptive sequential Monte Carlo approach. *J. Comput. Graph. Statist.* **25** 701–726. MR3533634 https://doi.org/10.1080/10618600.2015.1060885