



HAL
open science

Adaptive Sampling of Large Deviations

Grégoire Ferré, Hugo Touchette

► **To cite this version:**

Grégoire Ferré, Hugo Touchette. Adaptive Sampling of Large Deviations. Journal of Statistical Physics, 2018, 172 (6), pp.1525 - 1544. 10.1007/s10955-018-2108-8 . hal-01938825

HAL Id: hal-01938825

<https://hal.science/hal-01938825v1>

Submitted on 29 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Adaptive sampling of large deviations

Grégoire Ferré*

CERMICS, École des Ponts ParisTech, Université Paris-Est, F-77455 Marne-la-Vallée, France

Hugo Touchette†

*National Institute for Theoretical Physics (NITheP), Stellenbosch 7600, South Africa and
Institute of Theoretical Physics and Department of Physics,
University of Stellenbosch, Stellenbosch 7600, South Africa*

(Dated: July 4, 2018)

We introduce and test an algorithm that adaptively estimates large deviation functions characterizing the fluctuations of additive functionals of Markov processes in the long-time limit. These functions play an important role for predicting the probability and pathways of rare events in stochastic processes, as well as for understanding the physics of nonequilibrium systems driven in steady states by external forces and reservoirs. The algorithm uses methods from risk-sensitive and feedback control to estimate from a single trajectory a new process, called the driven process, known to be efficient for importance sampling. Its advantages compared to other simulation techniques, such as splitting or cloning, are discussed and illustrated with simple equilibrium and nonequilibrium diffusion models.

I. INTRODUCTION

We consider in this paper the problem of estimating large deviation functions characterizing the fluctuations of time-integrated functionals of Markov processes in the long-time limit. These functions have a wide range of applications in engineering and physical sciences, where they are used to predict the probability of rare events [1–3] and to understand how these events arise via transition paths or modified processes [4–6]. Large deviation theory also underlies now much of the research on nonequilibrium systems driven in steady states by non-conservative forces or boundary reservoirs [7–9]. In this context, large deviation functions play the role of nonequilibrium potentials, similar to the entropy or free energy, that characterize the steady state and fluctuations of physical quantities, such as energy or particle currents exchanged with reservoirs, as well as the presence of symmetries [10] and phase transitions in fluctuations [11–14].

Recently, many efforts have been devoted to developing numerical algorithms for estimating large deviation functions that go beyond the direct sampling of probabilities, which require prohibitively large samples. The most popular algorithms are variations of two basic approaches used in rare event simulations, namely: 1) splitting [15–17] or cloning [18–20] algorithms, which use population dynamics to estimate probabilities or generating functions that have a multiplicative structure in time, and 2) importance sampling [21–23] (including transition path sampling [24]), which modifies the process to be simulated, so as to transform rare events into typical events that can be simulated efficiently. Deterministic methods not based a priori on sampling can also be used, including spatial discretizations of various spectral and optimal control representations of large deviation functions, which work well for low-dimensional systems, in addition to action minimization methods, which can be applied in the low-noise or low-temperature limit [25–27].

In this paper, we propose an algorithm that combines spectral methods with importance sampling to efficiently estimate large deviation functions in an adaptive way. The core of the algorithm comes from recent works on learning algorithms for risk-sensitive control of Markov chains [28–31], which we adapt to continuous-time diffusion processes and to the problem of estimating large deviation functions. The algorithm works by estimating or learning “on the fly” a modified process, called the auxiliary or *driven process*, which corresponds to the process that is asymptotically

equivalent to the original process conditioned on the rare event of interest [6] or, alternatively, to the exponential tilting of the original process, known to be efficient for importance sampling [32]. This modified process is given by a principal eigenvalue problem related to the Feynman–Kac equation or, equivalently, by a stochastic optimal control problem [32] that we solve iteratively using stochastic approximation and feedback control methods.

The main advantage of this algorithm, compared to splitting or cloning, is that it does not require the simulation of many copies of the process considered – it runs on one long trajectory of that process, modified with a feedback-reinforcement rule, to adaptively learn the driven process, thereby reducing significantly the complexity of estimating large deviation functions. The calculation of error bars for the estimated quantities is also simplified compared to other techniques, as the algorithm is based on simple time averages and stochastic approximations [33–35]. Finally, the errors incurred by discretizing continuous-time processes and functionals can be analysed in a precise way, in principle, via Feynman–Kac semi-groups [36].

We discuss these advantages and test the algorithm in Sec. IV with simple equilibrium and nonequilibrium diffusions, after introducing the general model and notations in Sec. II and the algorithm in Sec. III. The results at this point are preliminary and are presented as a proof of concept of the algorithm. More detailed results about the time-discretization and sampling errors will be addressed in future works, together with more complex applications involving interacting particle systems and higher-dimensional diffusions.

II. FRAMEWORK

A. Model and notations

We consider an ergodic diffusion $(X_t)_{t \geq 0}$ evolving in a state space $\mathcal{X} \subset \mathbb{R}^d$ according to the following stochastic differential equation (SDE):

$$dX_t = b(X_t) dt + \sigma dW_t, \quad (1)$$

where $b : \mathcal{X} \rightarrow \mathbb{R}^d$ is a smooth function, called the *drift*, W_t is an m -dimensional Brownian motion, and σ is a $d \times m$ matrix, assumed to be constant for simplicity.¹ The generator of this diffusion reads

$$L = b \cdot \nabla + \frac{1}{2} \nabla \cdot D \nabla, \quad (2)$$

where \cdot denotes the scalar product and $D = \sigma \sigma^\top$, with \top as the transpose, is the diffusion matrix, assumed to be positive definite. This is the generator of the evolution semi-group P_T , defined by

$$P_T \varphi(x) = \mathbb{E}[\varphi(X_T) | X_0 = x], \quad (3)$$

for all time $T \geq 0$ and any smooth test function φ . The dual L^\dagger of L in the space $L^2(dx)$ of square-integrable functions with respect to the Lebesgue measure is the generator of the Fokker-Planck equation

$$\partial_t \rho(x, t) = L^\dagger \rho(x, t), \quad (4)$$

which gives the evolution of the probability density $\rho(\cdot, t)$ of X_t starting from some initial density $\rho(\cdot, 0)$ for X_0 .

¹ See [6] for a treatment of diffusions with multiplicative noise.

Our goal here is to study the fluctuations of time-integrated functionals of X_t , called *observables*, having the general form

$$A_T = \frac{1}{T} \int_0^T f(X_t) dt + \frac{1}{T} \int_0^T g(X_t) \circ dX_t, \quad (5)$$

where $f : \mathcal{X} \rightarrow \mathbb{R}$ and $g : \mathcal{X} \rightarrow \mathbb{R}^d$ are reasonably smooth functions (*e.g.* continuous) and \circ denotes the Stratonovich product [37]. Such a functional defined over the time horizon $[0, T]$ can represent, for example, a control cost associated with the state X_t and its increments [38] or a physical quantity integrated in time, such as the work performed on a particle by external forces or the heat exchanged by a particle with its environment [39].

Assuming that the process is ergodic with respect to an invariant measure $\mu(dx) = \rho^*(x) dx$ with smooth density ρ^* , we have

$$A_T \xrightarrow{T \rightarrow \infty} \int_{\mathcal{X}} f(x) \rho^*(x) dx + \int_{\mathcal{X}} g(x) \cdot J^*(x) dx = a^*, \quad (6)$$

almost surely, where

$$J^*(x) = b(x) \rho^*(x) - \frac{D}{2} \nabla \rho^*(x) \quad (7)$$

is the stationary current field associated with ρ^* [32]. The theory of large deviations [2] refines this ergodic theorem, generalized here with the additional g term, by providing estimates for the rate at which the probability distribution of A_T concentrates on its ergodic value a^* . Such estimates can be derived under general conditions (see [2]) and take, in the simplest case, the form

$$\lim_{T \rightarrow \infty} -\frac{1}{T} \log \mathbb{P}(A_T \in B) = \min_{a \in B} I(a), \quad (8)$$

for any Borel subset B of \mathbb{R} , where $I : \mathbb{R} \rightarrow \mathbb{R}_+$ is a positive function such that $I(a^*) = 0$. When this limit exists, A_T is said to satisfy the *large deviation principle* (LDP) with rate function I . Formally, this means that

$$\mathbb{P}(A_T \in da) = e^{-TI(a) + o(T)} da, \quad (9)$$

where $o(T)$ denotes corrections in the exponential that grow slower than linearly in T . Thus, we see that the rate function provides useful information about the fluctuations of A_T : the likelihood that $A_T = a$ decays exponentially with time for all $a \neq a^*$, since $I(a) > 0$ in this case, and converges otherwise to 1 as $T \rightarrow \infty$, since $I(a^*) = 0$. Moreover, the rate function is in general not a parabola, meaning that it describes fluctuations that are generally not Gaussian. In this sense, large deviation theory is often seen as an extension of both the ergodic theorem, which describes the concentration of A_T towards its mean, and the central limit theorem, which describes the local Gaussian fluctuations of A_T around its mean [2].

B. Large deviation functions and driven process

In practice, the rate function $I(a)$ can be calculated in many different ways other than by direct sampling, which requires exponentially large samples with T [21]. The most common method proceeds from the *scaled cumulant generating function* (SCGF), defined for $k \in \mathbb{R}$ by

$$\lambda(k) = \lim_{T \rightarrow \infty} \frac{1}{T} \log \mathbb{E} \left[e^{k A_T} \right]. \quad (10)$$

By the Gärtner-Ellis Theorem [2], the Legendre–Fenchel transform of this function yields the rate function:

$$I(a) = \sup_{k \in \mathbb{R}} \{ka - \lambda(k)\}, \quad (11)$$

provided, essentially, that $\lambda(k)$ is differentiable; see [2] for more precise conditions.

For the SDE (1) and the additive functional (5), the SCGF is known to be given by the principal eigenvalue of the operator

$$\mathcal{L}_k = b \cdot (\nabla + kg) + \frac{1}{2}(\nabla + kg) \cdot D(\nabla + kg) + kf, \quad (12)$$

which is the generator of the *Feynman–Kac semi-group* P_T^k , defined by

$$P_T^k \varphi(x) = \mathbb{E} \left[\varphi(X_T) e^{TkA_T} \mid X_0 = x \right], \quad (13)$$

for any smooth function φ ; see Appendix A.2 of [6]. In the end, the rate function can thus be computed by solving the spectral problem

$$\mathcal{L}_k h_k = \lambda(k) h_k, \quad (14)$$

where $\lambda(k)$ is the principal eigenvalue of \mathcal{L}_k and h_k its corresponding eigenvector. This holds provided that this operator has reasonable spectral properties, made precise in the following assumption.

Assumption 1. *The operator \mathcal{L}_k defined in (12) acting on $L^2(\mu)$ has an isolated largest eigenvalue $\lambda(k)$. Its multiplicity is one and it is associated with a regular eigenvector $h_k \in L^2(\mu)$ such that for all $x \in \mathcal{X}$, $h_k(x) > 0$.*

This assumption holds for many systems, in particular when \mathcal{X} is bounded [36, 40] or when b and g are gradient fields with appropriate growth conditions; see [41, Sec. 2.5]. In practice, the spectral problem (14) can be solved numerically using standard projection or discretization (Galerkin) methods, which work well for low-dimensional systems [42], or more involved real space renormalization methods when dealing with higher-dimensional systems [43]. Note that, for $g = 0$, \mathcal{L}_k is the usual Feynman–Kac generator with source term kf . Moreover, for $k = 0$, $\mathcal{L}_0 = L$ is the generator of the SDE, so that $\lambda(0) = 0$ and $h_0 = 1$, the constant unit function.

The numerical method that we propose in the next section attempts to estimate the spectral elements $\lambda(k)$ and h_k in a different way using the fact that they are solutions of the family of eigenproblems

$$P_T^k h_k = e^{T\lambda(k)} h_k, \quad \forall T > 0, \quad (15)$$

which can be approximated stochastically. The method also exploits a connection between large deviations and control theory showing that $\lambda(k)$ is the ergodic limit of an optimal control cost satisfying a stochastic Hamilton–Jacobi–Bellman equation [44], and that h_k determines the controlled diffusion achieving the optimal cost, which is the driven process mentioned in the introduction. These results are explained in detail in [32] (see also references therein); for the purpose of this paper, we only state them without proofs.

The controlled diffusion, denoted by $(\tilde{X}_t)_{t \geq 0}$, satisfies the SDE

$$d\tilde{X}_t = b_k(\tilde{X}_t) dt + \sigma dW_t, \quad (16)$$

where

$$b_k(x) = b(x) + D[kg(x) + \nabla \log h_k(x)] \quad (17)$$

is the optimal control drift defined for all $x \in \mathcal{X}$. Under Assumption 1, this modified diffusion is ergodic with respect to a new invariant measure μ_k , whose density is

$$\rho_k(x) = h_k(x)l_k(x), \quad (18)$$

where l_k is the dual of h_k in $L^2(dx)$ satisfying $\mathcal{L}_k^\dagger l_k = \lambda(k)l_k$ [6, Sec. 3.3]. These two functions are normalized such that

$$\int_{\mathcal{X}} h_k(x)l_k(x) dx = 1, \quad \int_{\mathcal{X}} l_k(x) dx = 1. \quad (19)$$

Because of the change of process, the observable A_T must concentrate on a new value, which can be shown to be given by $\lambda'(k)$, that is,

$$A_T \xrightarrow{T \rightarrow \infty} \lambda'(k) = a_k, \quad (20)$$

almost surely with respect to $(\tilde{X}_t)_{t \geq 0}$. Similarly, the control cost

$$C_T = kA_T - R_T, \quad (21)$$

where

$$R_T = \frac{1}{2T} \int_0^T [b(\tilde{X}_t) - b_k(\tilde{X}_t)] \cdot D^{-1}[b(\tilde{X}_t) - b_k(\tilde{X}_t)] dt, \quad (22)$$

reaches in the ergodic limit the value $\lambda(k)$, so that

$$C_T \xrightarrow{T \rightarrow \infty} \lambda(k) \quad (23)$$

almost surely under $(\tilde{X}_t)_{t \geq 0}$ [32, Sec. 4]. Finally, it can be shown by Legendre duality that the rate function at the value $a_k = \lambda'(k)$ is given by the ergodic limit of R_T above, leading to

$$R_T \xrightarrow{T \rightarrow \infty} I(a_k), \quad (24)$$

almost surely with respect to $(\tilde{X}_t)_{t \geq 0}$ [32, Sec. 4]. Note that the diffusion is not modified for $k = 0$, so that $a_0 = a^*$, $\lambda'(0) = a^*$, and $I(a^*) = 0$.

These ergodic limits provide direct estimators of the SCGF and the rate function, based on a single trajectory of the driven process, which can be simulated for different values of the parameter $k \in \mathbb{R}$. For the SCGF, there are in fact three possible estimators:

1. From (20): The value of A_T , integrated numerically with the condition $\lambda(0) = 0$.
2. From (23): The value of C_T .
3. The eigenvalue returned by the algorithm proposed in Sec. III.

In practice, we find that the last estimator is more stable, although the first and second are more adapted to obtain error bars.

For the rate function, we have two possible estimators:

1. The Legendre transform (11) of the SCGF, given in parametric form by

$$I(a_k) = ka_k - \lambda(k), \quad (25)$$

where a_k is either estimated from A_T or by taking the numerical derivative of $\lambda(k)$.

2. From (24): The value of R_T obtained for the value of A_T , giving the couple (A_T, R_T) .

In practice, we find that the first estimator based on the Legendre transform is more reliable. By comparison, the computation of $I(a)$ based on (22) involves the optimal drift b_k and, therefore, the logarithmic derivative of h_k , which is more difficult to estimate in a stable way.

In all cases, error bars can be constructed from the same trajectory by estimating, in principle, the variance of A_T , C_T and R_T using covariance techniques for Markov processes [33–35]. This is an advantage over splitting and cloning algorithms, for which the calculation of errors bars is difficult, as they involve correlated copies or “clones” of the simulated process [45–47].

In closing, it is interesting to note that the driven process can also be interpreted as the change of process in the importance sampling of the probability $\mathbb{P}(A_T \in B)$ that is optimal in the sense of logarithmic or asymptotic efficiency [21]. Therefore, it can be used not only to estimate the SCGF and rate function, but also to estimate the actual probability $\mathbb{P}(A_T \in B)$ in an efficient way [28]. The optimal change of process in this case is known to correspond to the exponential tilting of the original process [21], which is a time-dependent process in general; see Appendix D of [6]. In the ergodic limit, this process converges to a homogeneous process corresponding exactly to the driven process (16). We refer to [32] for more details about these results.

III. ADAPTIVE ALGORITHM

We are now ready to present the algorithm for estimating the SCGF and the rate function of A_T for continuous-time diffusions. The algorithm is based, as mentioned, on prior algorithms proposed in [28–31] for Markov chains and exploits the fact that

$$e^{-T\lambda(k)}(P_T^k \varphi) \xrightarrow{T \rightarrow \infty} h_k \int_{\mathcal{X}} \varphi(x) l_k(x) dx, \quad (26)$$

for any smooth test function φ . This result follows under Assumption 1; see [6, Sec. III.B] or [36, Sec. 6.1] for a proof. The algorithm that we propose works from this limit by approximating the action of the Feynman–Kac semi-group P_T^k in a stochastic way as a time average computed over a long trajectory of the diffusion. Moreover, it continuously modifies the diffusion as we estimate h_k to construct the driven process $(\tilde{X}_t)_{t \geq 0}$, which underlies the estimators of the large deviation functions.

The algorithm is presented next. Our contribution compared to [28–31] is to consider general time-continuous processes, time-additive functionals of these processes that depend on both their state and increments, and, more importantly, to construct the driven process explicitly so as to estimate the SCGF and the rate function adaptively using the estimators introduced in the previous section.

A. Time discretization

The first step required in the algorithm is to discretize in time the SDE (1) and its associated Feynman–Kac semi-group (13) by transforming the Markov diffusion $(X_t)_{t \geq 0}$ into a Markov chain $(x_n)_{n \in \mathbb{N}}$ with a small time step. Many discretization schemes can be used for the SDE; see, for

instance, [48]. Here we use the standard Euler–Maruyama scheme with constant time step Δt , given by

$$x_{n+1} = x_n + b(x_n)\Delta t + \sigma\sqrt{\Delta t}\xi_n, \quad (27)$$

where $(\xi_n)_{n \geq 0}$ is a sequence of independent standard d -dimensional Gaussian random variables. The corresponding discretization of the evolution semi-group P_T over a time step Δt is denoted by $Q_{\Delta t}$, so that

$$Q_{\Delta t}\varphi(x) = \mathbb{E}[\varphi(x_{n+1})|x_n = x], \quad (28)$$

for any test function φ and $x \in \mathcal{X}$. We refer to [41, 48] for more information about the discretization of SDEs and their weak error analysis.

Many discretizations also exist for the Feynman–Kac semi-group P_T^k . Here, we use the natural scheme where the diffusion is discretized as above and the integral of A_T is discretized as a Riemann sum with the left-point rule for the integral involving f and the mid-point rule for the Stratonovich integral involving g [36]. The action of P_T^k is thus replaced by

$$Q_{\Delta t}^k\varphi(x) = \mathbb{E} \left[e^{k \left[f(x_n)\Delta t + g\left(\frac{x_{n+1}+x_n}{2}\right) \cdot (x_{n+1}-x_n) \right]} \varphi(x_{n+1}) \mid x_n = x \right]. \quad (29)$$

For our purposes, h_k will be approximated by recursive applications of $Q_{\Delta t}^k$, based on the following assumption.

Assumption 2. *There exist a time step $\Delta t^* > 0$ and $p > 0$ such that, for $0 < \Delta t \leq \Delta t^*$,*

$$Q_{\Delta t}^k h_{k,\Delta t} = e^{\Delta t \lambda_{\Delta t}(k)} h_{k,\Delta t}, \quad (30)$$

where

$$h_{k,\Delta t} = h_k + O(\Delta t^p), \quad \lambda_{\Delta t}(k) = \lambda(k) + O(\Delta t^p). \quad (31)$$

This assumption means that the time-discretized operator $Q_{\Delta t}^k$ admits h_k as an approximate eigenvector with approximate eigenvalue $\lambda(k)$. This applies, for example, when \mathcal{X} is compact. In this case, precise estimates for the errors in (31) are obtained for $g = 0$ in [36] and can be extended to $g \neq 0$.

In the following, we will drop the subscript Δt on $h_{k,\Delta t}$ and $\lambda_{\Delta t}(k)$ to simplify the notations, and will present the algorithm essentially as if h_k were an exact eigenvector of $Q_{\Delta t}^k$ with exact eigenvalue $\lambda(k)$. However, we should keep in mind that this is only approximately true due to the errors in Δt . We will comment on this in Sec. IV with specific numerical examples.

B. Stochastic approximation and annealing

The main ingredient of the algorithm is the limit (26) of the Feynman–Kac semi-group, which shows that h_k and $\lambda(k)$ can be computed by successively applying $Q_{\Delta t}^k$ to an initial guess φ , so as to obtain

$$(Q_{\Delta t}^k)^n \varphi \sim e^{n\Delta t \lambda(k)} h_k \int_{\mathcal{X}} \varphi(x) l_k(x) dx \quad (32)$$

as $n \rightarrow \infty$. To perform this iteration, which is a functional version of the well-known *power method* for matrices [49], we apply a stochastic approximation [30, 35, 50] whereby the expectation appearing in the action of $Q_{\Delta t}^k$ is replaced by the iterates of the Markov chain:

$$\begin{aligned} Q_{\Delta t}^k \varphi(x_n) &= \mathbb{E} \left[e^{k \left[f(x_n) \Delta t + g \left(\frac{x_{n+1} + x_n}{2} \right) \cdot (x_{n+1} - x_n) \right]} \varphi(x_{n+1}) \mid x_n \right] \\ &\approx e^{k \left[f(x_n) \Delta t + g \left(\frac{x_{n+1} + x_n}{2} \right) \cdot (x_{n+1} - x_n) \right]} \varphi(x_{n+1}), \end{aligned} \quad (33)$$

where x_{n+1} is a random variable distributed according to $Q_{\Delta t}(x_n, \cdot)$. This approximation is known to reproduce the expectation as a statistical average in the ergodic limit $n \rightarrow \infty$ [29, 30, 35].

In our case, we simulate not the Markov chain x_n , but a modified chain, corresponding to the discretization of the driven process (16), which we express as

$$\tilde{x}_{n+1} = \tilde{x}_n + [b(\tilde{x}_n) + F_n(\tilde{x}_n)] \Delta t + \sigma \sqrt{\Delta t} \xi_n, \quad (34)$$

where

$$F_n = D(kg + \nabla \log h_k^n) \quad (35)$$

is the extra biasing force derived, according to (17), from the estimate h_k^n of h_k at time n . In this case, the evolution (33) is modified by Girsanov formula [37] to

$$\begin{aligned} Q_{\Delta t}^k \varphi(\tilde{x}_n) &= \mathbb{E} \left[e^{k \left[f(\tilde{x}_n) \Delta t + g \left(\frac{\tilde{x}_{n+1} + \tilde{x}_n}{2} \right) \cdot (\tilde{x}_{n+1} - \tilde{x}_n) \right]} \varphi(\tilde{x}_{n+1}) R_n(\tilde{x}_n, \tilde{x}_{n+1}) \mid \tilde{x}_n = x \right] \\ &\approx e^{k \left[f(\tilde{x}_n) \Delta t + g \left(\frac{\tilde{x}_{n+1} + \tilde{x}_n}{2} \right) \cdot (\tilde{x}_{n+1} - \tilde{x}_n) \right]} \varphi(\tilde{x}_{n+1}) R_n(\tilde{x}_n, \tilde{x}_{n+1}), \end{aligned} \quad (36)$$

where

$$R_n(\tilde{x}_n, \tilde{x}_{n+1}) = \exp \left(-\frac{1}{2\sigma^2} F_n^2(\tilde{x}_n) \Delta t - \sqrt{\frac{\Delta t}{\sigma^2}} F_n(\tilde{x}_n) \cdot \xi_n \right) \quad (37)$$

is the Radon–Nikodym derivative of the transition kernel of x_n with respect to that of \tilde{x}_n .

In the end, we also apply an annealing scheme, commonly used in stochastic approximations, which consists in replacing the update rule (32), defined by $\varphi^0 = \varphi$ and $\varphi^{n+1} = Q_{\Delta t}^k \varphi^n$, by the scheme

$$\varphi^{n+1} = \varphi^n + a_n \left(Q_{\Delta t}^k \varphi^n - \varphi^n \right), \quad (38)$$

where $(a_n)_{n \in \mathbb{N}}$ is a decreasing sequence, often called the adaptation or learning sequence, which acts as a smoothing parameter, filtering here the noisy update of the eigenfunction. This sequence is usually chosen in such a way that

$$\sum_{n \geq 0} a_n = \infty, \quad \sum_{n \geq 0} a_n^2 < \infty, \quad (39)$$

with the understanding that a_n should not be decreased too slowly, so as to limit noise, nor too fast, so as to reach the “correct” fixed point. These conditions can be relaxed under stability assumptions or by an averaging procedure [50]; see [35] for more details.

C. Spatial projection

The iteration just described for approximating h_k can be performed numerically by discretizing the state space \mathcal{X} into small cells (grid discretization). For high-dimensional systems, however, it is more convenient to use a Galerkin-type approximation of the eigenvalue and eigenfunction [30], obtained by projecting the problem onto a set of basis functions $\{\phi_j\}_{j=1}^M$ with $\phi_j : \mathcal{X} \rightarrow \mathbb{R}$ [42]. Let us denote by $\mathcal{H}_M = \text{Span}\{\phi_j\}_{j=1}^M$ the space spanned by these functions. Then the M -dimensional eigenproblem that we need to solve is

$$Q_{\Delta t}^k h = e^{\Delta t \lambda} h, \quad h \in \mathcal{H}_M, \quad (40)$$

using, for notational convenience, the same symbols h and λ for the exact and the projected spectral elements. We also drop from now on the parameter k , which will be implicit.

The eigenfunction h is expressed in that basis as

$$h(x) = \sum_{j=1}^M \alpha_j \phi_j(x) = \alpha^\top \phi(x), \quad (41)$$

where $\alpha = [\alpha_1, \dots, \alpha_M]^\top$ and $\phi(x) = [\phi_1(x), \dots, \phi_M(x)]^\top$. Multiplying (40) by ϕ_i for $i \in \{1, \dots, M\}$ and integrating over any measure η on \mathcal{X} yields

$$\sum_{j=1}^M \alpha_j \int_{\mathcal{X}} \phi_i(Q_{\Delta t}^k \phi_j) d\eta = e^{\Delta t \lambda} \sum_{j=1}^M \alpha_j \int_{\mathcal{X}} \phi_i \phi_j d\eta, \quad i \in \{1, \dots, M\}. \quad (42)$$

As a result, we see that the vector of coefficients $\alpha \in \mathbb{R}^M$ is the principal solution of the eigenproblem

$$A\alpha = \Lambda B\alpha, \quad \lambda = \frac{1}{\Delta t} \log \Lambda, \quad (43)$$

where

$$A = \int_{\mathcal{X}} \phi(Q_{\Delta t}^k \phi^\top) d\eta, \quad B = \int_{\mathcal{X}} \phi \phi^\top d\eta. \quad (44)$$

Note that the matrix B is invertible as soon as the $\{\phi_j\}_{j=1}^M$ form a linearly independent family.

D. Algorithm

We are now ready to describe all the steps of the algorithm that estimates the principal eigenvalue of \mathcal{L}_k and its corresponding eigenfunction. For a fixed $k \in \mathbb{R}$, we initiate the process at a position $x_0 \in \mathcal{X}$ and define a first approximation h^0 of h as

$$h^0 = (\alpha^0)^\top \phi, \quad (45)$$

where ϕ is the vector of basis functions and α^0 is the initial vector of coefficients, chosen such that h^0 is constant. At each iteration, we then perform the following steps:

1. Draw a new position \tilde{x}_{n+1} according to the Markov chain (34).
2. Compute the extra bias F_n according to (35), which in the function basis takes the form

$$F_n = D \left(kg + \frac{\sum_{j=1}^M \alpha_j^n \nabla \phi_j}{\sum_{j=1}^M \alpha_j^n \phi_j} \right). \quad (46)$$

3. Compute the Girsanov weight R_n according to (37).
4. Compute the matrices A_{n+1} and B_{n+1} using the formulae

$$\begin{aligned} A_{n+1} &= \frac{1}{n+1} \sum_{m=0}^n e^{k[f(\tilde{x}_m)\Delta t + g\left(\frac{\tilde{x}_{m+1} + \tilde{x}_m}{2}\right)(\tilde{x}_{m+1} - \tilde{x}_m)]} \phi(\tilde{x}_m) \phi(\tilde{x}_{m+1})^\top R_m(\tilde{x}_m, \tilde{x}_{m+1}), \\ B_{n+1} &= \frac{1}{n+1} \sum_{m=0}^n \phi(\tilde{x}_m) \phi(\tilde{x}_m)^\top, \end{aligned} \quad (47)$$

which follow by projecting (36) and (44), respectively.

5. Update the coefficient vector α^n , giving the decomposition of the iterate h^n , as

$$\alpha^{n+1} = \alpha^n + a_n \left(\frac{B_n^{-1} A_n}{h^n(x_0)} - \text{Id} \right) \alpha^n, \quad (48)$$

where Id is the identity matrix of size M [30].

6. Estimate the eigenvalue as

$$\lambda^{n+1} = \frac{n}{n+1} \lambda^n + \frac{1}{(n+1)\Delta t} \log(h^n(\tilde{x}_0)), \quad (49)$$

which follows from (43).

Repeating these steps, it can be proved that the iterates h^n and λ^n converge to the solution of the spectral problem (40), following the analysis found in [28–30]; see also [35]. Moreover, because $h^n \rightarrow h$, the Markov chain $(\tilde{x}_n)_{n \in \mathbb{N}}$ samples in the long run the Euler–Maruyama discretization of the driven process (16).

E. Remarks

The following are technical remarks worth noting about the algorithm:

1. The matrix A_n can be updated at each step using

$$A_{n+1} = \frac{n}{n+1} A_n + \frac{1}{n+1} e^{k[f(\tilde{x}_n)\Delta t + g\left(\frac{\tilde{x}_{n+1} + \tilde{x}_n}{2}\right)(\tilde{x}_{n+1} - \tilde{x}_n)]} \phi(\tilde{x}_n) \phi(\tilde{x}_{n+1})^\top R_n, \quad (50)$$

instead of the sum shown in (47). Similarly, the matrix B_n can be updated, following [30], using the Sherman–Morrisson–Woodbury formula, which leads here to

$$B_{n+1}^{-1} = \frac{n+1}{n} B_n^{-1} - \frac{n+1}{n} \frac{B_n^{-1} \phi(\tilde{x}_{n+1}) \phi(\tilde{x}_{n+1})^\top B_n^{-1} R_{n+1}}{\frac{n}{n+1} + \phi(\tilde{x}_{n+1})^\top B_n^{-1} \phi(\tilde{x}_{n+1}) R_{n+1}}. \quad (51)$$

The advantage of this formula is that the computation required for updating the coefficient α^n scales with the number M of basis vectors as M^2 , as is common in the power method [49], whereas the typical cost of inverting the matrix B scales as M^3 .

2. We normalize the eigenfunction at every iteration by setting $h(\tilde{x}_0) = 1$ at an arbitrary location, taken here to be the initial state \tilde{x}_0 [28]. This prevents the norm of the eigenfunction from diverging or decaying to zero, as is common in the power method, and provides an estimate of $\lambda(k)$ through (49). This normalization step can be based on other norms, at the expense of computing integrals.
3. The algorithm can be run with the original (unbiased) process X_t , but the estimation of the matrices A and B in this case typically suffers from large statistical errors due to the high variance of the underlying estimator. This is a known problem related to the estimation of Feynman–Kac functionals and exponential integrals in general [51–54]. Biasing the dynamics with the driven process \tilde{X}_t reduces this variance in an optimal way (in the sense of asymptotic or logarithmic efficiency) by forcing the exploration of the process in important regions of the state space where the integrand of the generating function $\mathbb{E}[e^{kTA_T}]$ is largest [32].
4. As the dynamics is biased towards the driven process \tilde{X}_t , the vector α^n of basis coefficients representing h_k converges towards the solution of the eigenproblem (43), with the matrices A and B computed as ergodic averages under \tilde{X}_t , so that $\eta = \mu_k$ in (44). Even for a small number of basis functions, it would be a priori impossible to compute these matrices by numerical quadrature. This shows that the algorithm can be used to obtain good approximations of h_k even for high-dimensional systems, provided that enough basis functions are used to represent the support of μ_k , which is typically concentrated on a subset of \mathcal{X} .
5. The algorithm is stable despite the fact that it includes the Girsanov reweighting factor, which is exponential in time. The reason for this stability, already noted in [29], is that the Girsanov weight is computed and accumulated incrementally over single time steps.
6. The learning sequence $(a_n)_{n \in \mathbb{N}}$ is chosen here in the following way:
 - (a) For $0 \leq n \leq N_1$, we take $a_n = 0$, so there is no adaption at the level of h and λ , although the matrices A and B are evolved. This “burn-in” period allows for a better initial guess of the various functions estimated through their ergodic averages.
 - (b) For $N_1 < n \leq N_1 + N_{\text{iter}} - N_2$, we take $a_n = 1$, that is, the full information of the process is taken into account.
 - (c) For $N_1 + N_{\text{iter}} - N_2 < n \leq N_1 + N_{\text{iter}}$, we take $a_n = C/(n - (N_1 + N_{\text{iter}} - N_2))$ with constant $C > 0$, so the process learns less with time, smoothing the noise in the long run. In the following, we choose $C = 1$.

The times N_1 and N_2 can be fixed or can be chosen dynamically according to some stopping rule.

7. In practice, we can perform independent simulations with different values of k to obtain an interpolation of the SCGF over some range, say, $[k_{\min}, k_{\max}]$, which can then be used to obtain the rate function by Legendre transform. Alternatively, we can do a simulation in which k is slowly increased from $k = 0$ in a “quasi-static” way, so as to adaptively update the biasing force over a range of values for k [55] or to reach some prescribed value of the SCGF for an unknown k [28].

IV. APPLICATIONS

We apply in this section the algorithm to two simple test cases involving one-dimensional diffusions. The first is the Ornstein-Uhlenbeck process, for which the large deviation functions of the area per unit time are known exactly [6], while the second is a driven diffusion on the circle, often used in physics as a model of nonequilibrium systems, including Josephson junctions perturbed by thermal noise and Brownian particles controlled by external forces [56–58]. We discuss for both the convergence and efficiency of the algorithm.

A. Ornstein-Uhlenbeck process

The first example that we consider is the mean area or mean position

$$A_T = \frac{1}{T} \int_0^T X_t dt \quad (52)$$

of the Ornstein-Uhlenbeck process on \mathbb{R} satisfying the SDE

$$dX_t = -2\theta X_t dt + \sqrt{2} dW_t, \quad (53)$$

where $\theta > 0$. In the notations of Sec. II, we thus have $b(x) = -2\theta x$, $\sigma = \sqrt{2}$, $f(x) = x$ and $g(x) = 0$, so that

$$L = -2\theta x \frac{d}{dx} + \frac{d^2}{dx^2}, \quad (54)$$

and $\mathcal{L}_k = L + kx$.

For this process and observable, it can be checked [6] that the SCGF, corresponding to the dominant eigenvalue of \mathcal{L}_k , is

$$\lambda(k) = \frac{k^2}{4\theta^2}, \quad (55)$$

so that $I(a) = \theta^2 a^2$ from (25). This is expected, since the integral of a Gaussian process is also Gaussian. Moreover, the associated dominant eigenfunction is

$$h_k(x) = e^{\frac{k}{2\theta}x}, \quad (56)$$

leading with (17) to the optimal drift

$$b_k(x) = -2\theta x + 2(\log h_k(x))' = -2\theta x + \frac{k}{\theta}. \quad (57)$$

This shows that a fluctuation of A_T is created in an optimal way by adding a constant to the drift, which “moves” the Gaussian stationary density of the Ornstein-Uhlenbeck process

$$\rho^*(x) = \sqrt{\frac{\theta}{\pi}} e^{-\theta x^2} \quad (58)$$

to

$$\rho_k(x) = \sqrt{\frac{\theta}{\pi}} e^{-\theta(x-m_k)^2}, \quad m_k = \frac{k}{2\theta^2}, \quad (59)$$

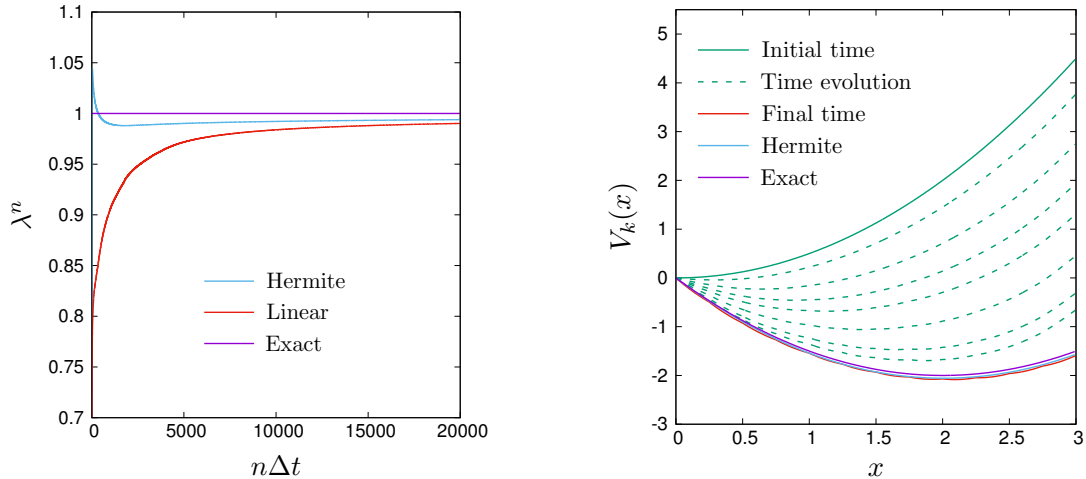


FIG. 1. Large deviations of the mean area for the Ornstein-Uhlenbeck process. Left: Evolution of the estimated SCGF in time. Right: Evolution of the biasing potential $V_k(x)$ in time for the linear basis (curves from top to bottom), compared with the final estimate for the Hermite basis and the exact result.

leading to

$$A_T \xrightarrow{T \rightarrow \infty} m_k = a_k \quad (60)$$

almost surely with respect to \tilde{X}_t , in agreement with (20).

We compare the algorithm against these exact results using a simple mesh discretization (first-order finite elements), defined by the basis functions

$$\phi_j(x) = \begin{cases} \frac{x-x_j}{\delta} + 1, & x \in [x_j - \delta, x_j], \\ -\frac{x-x_j}{\delta} + 1, & x \in [x_j, x_j + \delta], \\ 0, & \text{otherwise,} \end{cases} \quad (61)$$

where the points x_j define the centers of each “cell” of width $\delta > 0$. In the following, we refer to this basis simply as the *linear basis*. To illustrate the flexibility of the algorithm, we also perform simulations using a Hermite polynomial basis, which forms a complete orthonormal basis in $L^2(\rho^*)$. We run the algorithm with $k = 1$, $\theta = 1/2$, $x_0 = 0$, $T = 2 \times 10^4$, $T_1 = N_1 \Delta t = T_2 = N_2 \Delta t = 2 \times 10^3$ with $\Delta t = 5 \times 10^{-3}$, which is a standard time step used in simulations relative to the basic timescale of the dynamics, corresponding here to $1/(2\theta) = 1$. For the linear basis, we use $M = 61$ equally spaced cells with $\delta = 0.25$ around $x = 0$, whereas for the Hermite basis we use only $M = 10$ basis functions.

Figure 1 illustrates the results of a typical simulation, starting on the left with the evolution of the estimate of the SCGF, given by (49), as time increases. We observe a very good agreement in the long run with the exact value, which for the parameters used is equal to 1, with a faster convergence observed for the Hermite basis compared to the linear basis. As mentioned before, we can also recover the SCGF by recording the stationary value of A_T , which corresponds as above to $m_k = a_k = \lambda'(k)$, and numerically integrate the result in k from $\lambda(0) = 0$. The results obtained are similar to those obtained from the eigenvalue estimate (49), and are not shown for this reason.

To understand the convergence of the eigenvalue at the process level, we show in the right plot of Fig. 1 the evolution of the effective potential $V_k(x)$ associated with the modified drift $b_k(x)$ according to

$$b_k(x) = -V_k'(x). \quad (62)$$

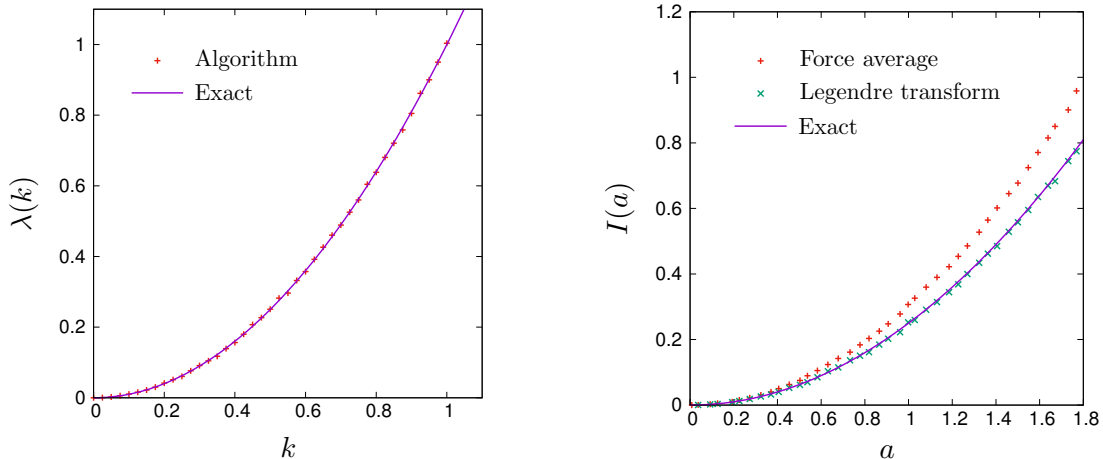


FIG. 2. Large deviations of the mean area for the Ornstein-Uhlenbeck process. Left: Estimated SCGF compared with the exact result. Right: Estimated rate function based on the Legendre transform of the SCGF and on the averaging of the biasing force (using the Hermite basis), compared with the exact result.

We show the results for both the linear basis and the Hermite basis, with the zero of the potential arbitrarily set at $V_k(0) = 0$. We see from these that the tail of the potential takes longer to be estimated correctly, as the process starts to explore values away from $m_0 = 0$. After the convergence time, there is a good agreement with the parabola

$$V_k(x) = \theta(x - m_k)^2, \quad (63)$$

which is the exact solution predicted by (59). The small errors are due to the finite time step Δt used and the finite basis function set. The small “wiggles” seen in the potential obtained with the linear basis come from the fact that this basis is piecewise linear, so that $V_k(x)$ is piecewise logarithmic. The potential obtained with the Hermite basis is smoother, as expected.

Repeating the simulations for other values of k , we can recover $\lambda(k)$, as shown in the left plot of Fig. 2. The result is in good agreement with the exact solution (55). Estimating $\lambda'(k)$ with (20) leads to an estimated rate function, by the Legendre transform (25), which also agrees well with the exact rate function, shown in the right plot of Fig. 2. By comparison, the rate function obtained from the time average (24) of the extra biasing force (with the Hermite basis) is not as good: it lies above the exact rate function and shows a larger offset or error as k increases, which does not decrease by reducing Δt or increasing the number of basis functions.

To understand this error, we show in the left plot of Fig. 3 the evolution of the extra biasing force F_n for $k = 1$, estimated at the current location of the system, which should approach the constant 2, following (57). The evolution is noisy, as can be seen, which is expected, since F_n is estimated by the logarithmic derivative

$$F_n(x) = 2 \frac{(h^n(x))'}{h^n(x)}, \quad (64)$$

computed in the function basis from (46). The derivative amplifies the Monte Carlo errors inherent in the estimate h^n . In addition, the denominator often takes small values away from the mean position m_k , which makes the estimation of the optimal force still more difficult. The right plot of Fig. 3 shows that the noise on F_n is considerably filtered out by the time average underlying the estimator R_T of the rate function, although a bias remains even after the convergence time of the SCGF, which leads to the offset seen in Fig. 2. The results in both cases are more noisy for the linear basis because the derivative is not continuous across the different cells.

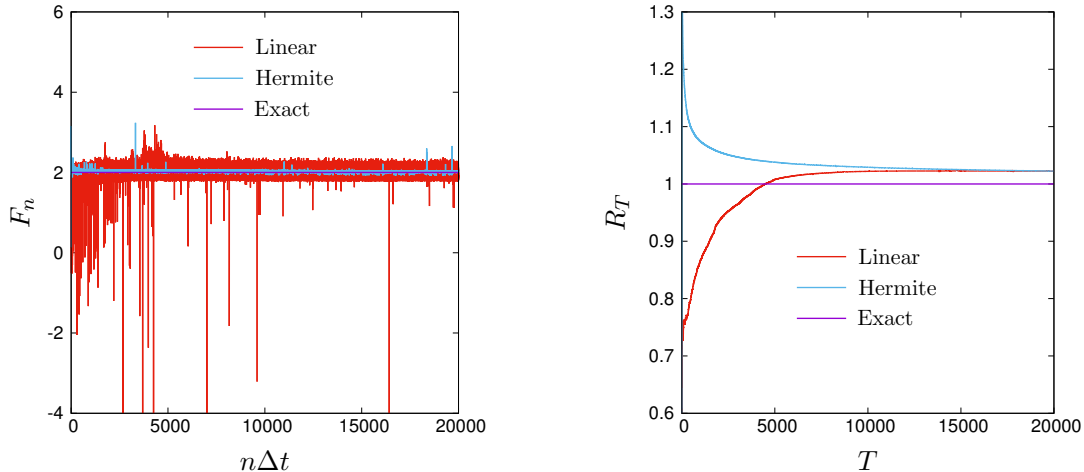


FIG. 3. Large deviations of the mean area of the Ornstein-Uhlenbeck process. Left: Biasing force F_n at the current state \tilde{x}_n of the process at time n . Right: Evolution of the time average R_T , which gives the rate function according to (24).

The offset on R_T remains more or less constant by running independent simulations (error bars on 30 simulations are too small to show), and so appears to be a systematic error or bias. Many factors can account for this bias. First, the rate function estimated from the limit shown in (24) is known to be an upper bound on the true rate function [32], which is tight if and only if the modified drift estimated in the simulation is the optimal control drift (17). Here, F_n is not constant, as predicted from (57), so we expect the estimate R_T to lie above its expected value. Second, we have noticed in simulations that A_T underestimates a_k for large k , which has the effect of further “pushing” the estimate of the rate function above $I(a)$. This is most likely due again to F_n being non-constant. Finally, the time average of F_n involves the ratio of two functions, according to (64), which fluctuate in time via the updating of α^n in (46). As a result, we expect this additional randomness to artificially increase the second moment of F_n , leading to a further bias.

It is difficult to isolate these factors, and all, in fact, seem to play a role. In future works, it would be interesting to study the bias observed in R_T by computing, for example, the time average of F_n^2 using the final estimate of h_k rather than the time-evolved estimate h^n . Different annealing sequences or averaging techniques could also be used to “filter” the time average of F_n^2 and mitigate the bias on R_T . For now, the most efficient and reliable way to compute the rate function is to use the estimator based on the SCGF and its Legendre transform.

B. Periodic diffusion

We consider for the second test a diffusion on the unit circle satisfying the SDE

$$dX_t = (-V'(X_t) + \gamma)dt + \sqrt{2}dW_t, \quad (65)$$

with the periodic potential $V(x) = \cos(2\pi x)$ and $\gamma \in \mathbb{R}$ a constant drive. For $\gamma \neq 0$, the total drift cannot be expressed as the gradient of a smooth periodic function, so X_t is a nonequilibrium process violating detailed balance [56]. The observable studied for this process is the winding number

$$A_T = \frac{1}{T} \int_0^T dX_t, \quad (66)$$

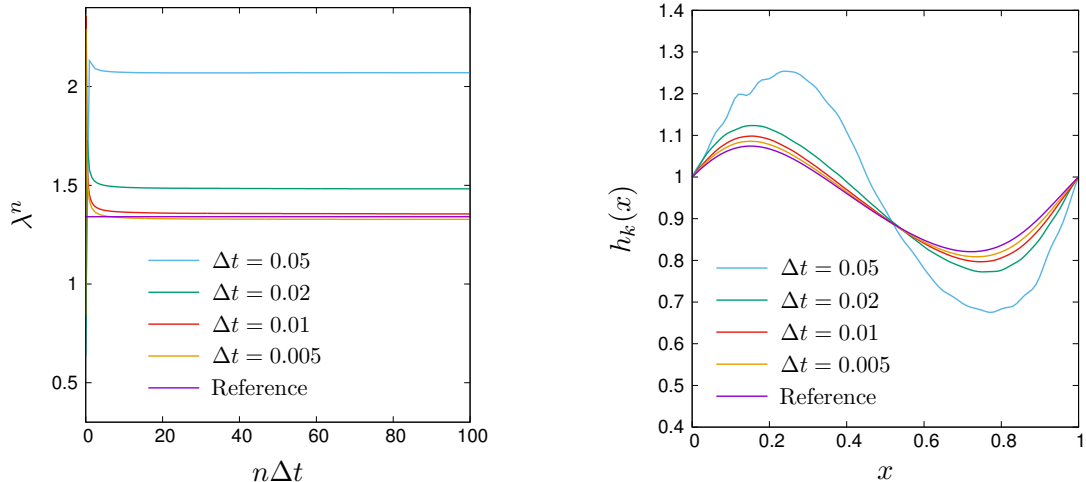


FIG. 4. Current large deviations of the driven periodic diffusion. Left: Evolution of the estimated SCGF with time for different time steps Δt . Right: Estimated eigenfunction h_k for different time steps Δt compared with the reference Fourier solution.

calculated with the real rather than periodic state, which can be interpreted physically as the mean velocity or current of a Brownian particle moving around the circle [59]. In the notation of Sec. II, we have $b(x) = -V'(x) + \gamma$, $\sigma = \sqrt{2}$, $f(x) = 0$, and $g(x) = 1$.

The SCGF and the rate function of this observable are not known exactly, but Galerkin approximations can easily be found by Fourier series, using the basis functions $\phi_j(x) = e^{j2\pi ix}$ with $j \in \mathbb{Z}$ [59–61]. This basis is used to compute reference values for the SCGF and the rate function, devoid of Δt errors, by projecting the spectral problem (14) in Fourier space and by ensuring that enough basis functions are used. We find in practice that $M = 41$ Fourier modes are sufficient. We use the same Fourier basis for the algorithm, also with $M = 41$ modes, in addition to $T = 2 \times 10^4$ and $T_1 = T_2 = 2 \times 10^3$ for the integration times, as in the first test.

We show in Fig. 4 the results of a typical simulation for $\gamma = 1$ and $k = 1$. The left plot in this figure shows the evolution in time of the estimated SCGF from the eigenvalue iteration (49), while the right plot shows the final estimated eigenfunction h_k for different time steps Δt , compared with the reference Fourier solution. We see that the results for the eigenvalue and the eigenvector significantly depart from the reference solutions for $\Delta t = 0.05$, but converge to them as Δt is decreased, in accordance with Assumption 2 and the theoretical results of [36]. For the same time step used for the Ornstein-Uhlenbeck process, namely, $\Delta t = 5 \times 10^{-3}$, we see no notable difference between the estimated eigenfunction and the reference values, leading to a precise estimation of the SCGF. The convergence of the eigenvalue here is much faster than for the Ornstein-Uhlenbeck process because the space explored is compact, being limited to $[0, 1]$ with periodic boundary conditions.

As for the Ornstein-Uhlenbeck process, we can repeat these simulations over a range of values for k to obtain the SCGF and the rate function. The left plot of Fig. 5 shows that the SCGF is in good agreement with the Fourier solution for $\Delta t = 5 \times 10^{-3}$, and so is the rate function estimated by Legendre transform. However, as seen before, the rate function estimated with the time average R_T of the biasing force shows an offset, although smaller this time, which comes from the noisy estimation of F_n . As before, the estimator of the rate function that should be used is the one based on the Legendre transform of the SCGF, with a_k estimated by A_T .

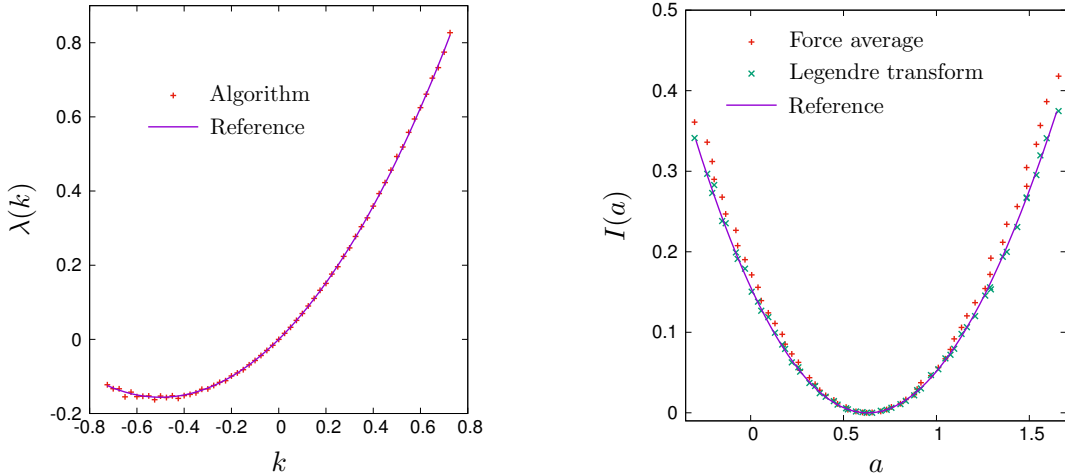


FIG. 5. Current large deviations for the driven periodic diffusion. Left: Estimated SCGF compared with the reference Fourier solution. Right: Estimated rate function based on the Legendre transform and on the averaging of the biasing force, compared with the reference Fourier solution.

V. CONCLUSION

We have presented a new algorithm for estimating the large deviation functions of time-integrated observables of Markov processes, which characterize the likelihood of their fluctuations in the long-time limit. The algorithm draws on earlier results on stochastic control [28–31], and works by adaptively estimating the principal eigenvalue and eigenfunction of a spectral problem related to the large deviation problem. The adaptive part consists in modifying the process considered, using feedback and reinforcement learning, so as to reach the so-called driven process, which is known to be optimal for the purpose of estimating large deviation functions using importance sampling [32]. In this sense, the algorithm relates to many adaptive importance sampling methods that have been proposed recently for rare event simulations; see, for instance, [62–65]. It is also closely related to diffusion Monte Carlo methods [45, 66, 67], which attempt to estimate the ground-state wavefunction of a many-body quantum system by simulating a related stochastic process.

The proposed algorithm can be applied to diffusions, as illustrated here, but also to Markov chains in discrete or continuous time, and works for both equilibrium and nonequilibrium systems, that is, reversible and non-reversible processes, respectively. Moreover, although the test cases that we have presented are simple, they clearly show that the algorithm has the potential to improve upon other simulation methods, especially cloning methods, since it runs on a single simulation of the process and provides information about how fluctuations are created in time by constructing the driven process in a non-parametric form and with no prior information. A modification of the original cloning algorithm was proposed recently [46] to construct the driven process, but it is based on a different feedback rule that compares two time-dependent histograms, whose estimation is noisy and requires a large number of clones. The results presented here show that a single clone, evolving over a long-enough time, is sufficient. This obviously cuts the computational complexity of estimating large deviations, but also simplifies, as mentioned, the error and convergence analyses of the algorithm.

As with any new proposal, more work is needed to understand the benefits and limitations of the algorithm, to test its applicability to realistic systems, and to benchmark it against other numerical methods. Of particular importance is to derive precise error estimates associated with the space and time discretizations of the spectral elements, large deviation functions, and the driven process.

Recent results about the discretization errors in Δt for the SCGF can be found in [36]. These discretization errors are also present in cloning algorithms when applied to diffusions, so they are not specific to the algorithm presented here. The use of Galerkin discretizations is also common to both algorithms and requires further investigations, particularly in the low-noise limit and for processes involving many interacting particles.

ACKNOWLEDGMENTS

We are grateful to Florian Angeletti for useful discussions in the initial phase of this work. G.F. is supported by the Labex Bezout. H.T. was supported by the National Research Foundation of South Africa (Grant nos 90322 and 96199) and Stellenbosch University (Project Funding for New Appointee). This research was also supported in part by the International Centre for Theoretical Sciences (ICTS) during a visit for participating in the program “Large deviation theory in statistical physics: Recent advances and future challenges” (Code: ICTS/Prog-ldt/2017/8).

* gregoire.ferre@enpc.fr

† htouchette@sun.ac.za, htouchet@alum.mit.edu

- [1] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis*, Stochastic Modeling Series (Chapman and Hall, London, 1995).
- [2] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed. (Springer, New York, 1998).
- [3] F. den Hollander, *Large Deviations*, Fields Institute Monograph (AMS, Providence, 2000).
- [4] R. L. Jack and P. Sollich, “Large deviations and ensembles of trajectories in stochastic models,” *Prog. Theoret. Phys. Suppl.* **184**, 304–317 (2010).
- [5] R. Chetrite and H. Touchette, “Nonequilibrium microcanonical and canonical ensembles and their equivalence,” *Phys. Rev. Lett.* **111**, 120601 (2013).
- [6] R. Chetrite and H. Touchette, “Nonequilibrium Markov processes conditioned on large deviations,” *Ann. Henri Poincaré* **16**, 2005–2057 (2015).
- [7] H. Touchette, “The large deviation approach to statistical mechanics,” *Phys. Rep.* **478**, 1–69 (2009).
- [8] B. Derrida, “Non-equilibrium steady states: Fluctuations and large deviations of the density and of the current,” *J. Stat. Mech.* **2007**, P07023 (2007).
- [9] R. J. Harris and H. Touchette, “Large deviation approach to nonequilibrium systems,” in *Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond*, Reviews of Nonlinear Dynamics and Complexity, Vol. 6, edited by R. Klages, W. Just, and C. Jarzynski (Wiley-VCH, Weinheim, 2013) pp. 335–360.
- [10] R. J. Harris and G. M. Schütz, “Fluctuation theorems for stochastic dynamics,” *J. Stat. Mech.* **2007**, P07020 (2007).
- [11] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, “Dynamical first-order phase transition in kinetically constrained models of glasses,” *Phys. Rev. Lett.* **98**, 195702 (2007).
- [12] L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler, “Dynamic order-disorder in atomistic models of structural glass formers,” *Science* **323**, 1309–1313 (2009).
- [13] C. P. Espigares, P. L. Garrido, and P. I. Hurtado, “Dynamical phase transition for current statistics in a simple driven diffusive system,” *Phys. Rev. E* **87**, 032115 (2013).
- [14] A. Aminov, G. Bunin, and Y. Kafri, “Singularities in large deviation functionals of bulk-driven transport models,” *J. Stat. Mech.* **2014**, P08017 (2014).
- [15] T. Dean and P. Dupuis, “Splitting for rare event simulation: A large deviation approach to design and analysis,” *Stoch. Proc. Appl.* **119**, 562–587 (2009).
- [16] F. Cérou and A. Guyader, “Adaptive multilevel splitting for rare event analysis,” *Stoch. Anal. Appl.* **25**, 417–443 (2007).

- [17] D. Aristoff, T. Lelièvre, C. G. Mayne, and I. Teo, “Adaptive multilevel splitting in molecular dynamics simulations,” *ESAIM: Proc.* **48**, 215–225 (2015).
- [18] P. Grassberger, “Go with the winners: A general Monte Carlo strategy,” *Comp. Phys. Comm.* **147**, 64–70 (2002).
- [19] C. Giardinà, J. Kurchan, and L. Peliti, “Direct evaluation of large-deviation functions,” *Phys. Rev. Lett.* **96**, 120603 (2006).
- [20] V. Lecomte and J. Tailleur, “A numerical approach to large deviations in continuous time,” *J. Stat. Mech.* **2007**, P03004 (2007).
- [21] J. A. Bucklew, *Introduction to Rare Event Simulation* (Springer, New York, 2004).
- [22] S. Juneja and P. Shahabuddin, “Rare-event simulation techniques: An introduction and recent advances,” in *Handbooks in Operations Research and Management Science*, Vol. Volume 13, edited by S. G. Henderson and B. L. Nelson (Elsevier, Amsterdam, 2006) Chap. 11, pp. 291–350.
- [23] S. Asmussen and P. W. Glynn, *Stochastic Simulation: Algorithms and Analysis*, Stochastic Modelling and Applied Probability (Springer, New York, 2007).
- [24] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, “Transition path sampling: Throwing ropes over rough mountain passes, in the dark,” *Ann. Rev. Phys. Chem.* **53**, 291–318 (2002).
- [25] M. Heymann and E. Vanden Eijnden, “Pathways of maximum likelihood for rare events in nonequilibrium systems: Application to nucleation in the presence of shear,” *Phys. Rev. Lett.* **100**, 140601 (2008).
- [26] E. Vanden-Eijnden and J. Weare, “Rare event simulation of small noise diffusions,” *Comm. Pure Appl. Math.* **65**, 1770–1803 (2012).
- [27] T. Grafke, R. Grauer, and T. Schäfer, “The instanton method and its numerical implementation in fluid mechanics,” *J. Phys. A: Math. Theor.* **48**, 333001 (2015).
- [28] V. S. Borkar, S. Juneja, and A. A. Kherani, “Performance analysis conditioned on rare events: An adaptive simulation scheme,” *Commun. Info. Syst.* **3**, 259–278 (2004).
- [29] T. P. I. Ahamed, V. S. Borkar, and S. Juneja, “Adaptive importance sampling technique for Markov chains using stochastic approximation,” *Oper. Res.* **54**, 489–504 (2006).
- [30] A. Basu, T. Bhattacharyya, and V. S. Borkar, “A learning algorithm for risk-sensitive cost,” *Math. Oper. Res.* **33**, 880–898 (2008).
- [31] V. S. Borkar, “Learning algorithms for risk-sensitive control,” in *Proc. 19th Int. Symp. Math. Theory Networks and Systems* (2010) pp. 1327–1332.
- [32] R. Chetrite and H. Touchette, “Variational and optimal control representations of conditioned and driven processes,” *J. Stat. Mech.* **2015**, P12001 (2015).
- [33] D. Chauveau and J. Diebolt, “Estimation of the asymptotic variance in the CLT for Markov chains,” *Stoch. Models* **19**, 449–465 (2003).
- [34] G. O. Roberts and J. S. Rosenthal, “General state space Markov chains and MCMC algorithms,” *Prob. Surveys* **1**, 20–71 (2004).
- [35] A. Benveniste, M. Métivier, and P. Priouret, *Adaptive Algorithms and Stochastic Approximations*, Stochastic Modelling and Applied Probability, Vol. 22 (Springer, 2012).
- [36] G. Ferré and G. Stoltz, “Error estimates on ergodic properties of discretized Feynman–Kac semigroups,” (2017), [arXiv:1712.04013](https://arxiv.org/abs/1712.04013).
- [37] G. A. Pavliotis, *Stochastic Processes and Applications* (Springer, New York, 2014).
- [38] V. Y. Chernyak, M. Chertkov, J. Bierkens, and H. J. Kappen, “Stochastic optimal control as non-equilibrium statistical mechanics: Calculus of variations over density and current,” *J. Phys. A: Math. Theor.* **47**, 022001 (2014).
- [39] K. Sekimoto, *Stochastic Energetics*, Lect. Notes. Phys., Vol. 799 (Springer, New York, 2010).
- [40] J. Bierkens, V. Y. Chernyak, M. Chertkov, and H. J. Kappen, “Linear PDEs and eigenvalue problems corresponding to ergodic stochastic optimization problems on compact manifolds,” *J. Stat. Mech.* **2016**, 013206 (2016).
- [41] T. Lelièvre and G. Stoltz, “Partial differential equations and stochastic methods in molecular dynamics,” *Acta Numer.* **25**, 681–880 (2016).
- [42] F. Chatelin, *Spectral Approximation of Linear Operators*, Classics in Applied Mathematics (SIAM, Philadelphia, 2011).
- [43] M. Gorissen and C. Vanderzande, “Finite size scaling of current fluctuations in the totally asymmetric exclusion process,” *J. Phys. A: Math. Theor.* **44**, 115005 (2011).
- [44] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions*, Stochastic

- Modelling and Applied Probability, Vol. 25 (Springer, New York, 2006).
- [45] M. Rousset, “On the control of an interacting particle estimation of Schrödinger ground states,” *SIAM J. Math. Anal.* **38**, 824–844 (2006).
 - [46] T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte, “Population-dynamics method with a multicanonical feedback control,” *Phys. Rev. E* **93**, 062123 (2016).
 - [47] T. Nemoto, E. Guevara Hidalgo, and V. Lecomte, “Finite-time and finite-size scalings in the evaluation of large-deviation functions: Analytical study using a birth-death process,” *Phys. Rev. E* **95**, 012102 (2017).
 - [48] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations* (Springer, Berlin, 1992).
 - [49] J. Demmel, *Applied Numerical Linear Algebra* (SIAM, Philadelphia, 1997).
 - [50] B. T. Polyak and A. B. Juditsky, “Acceleration of stochastic approximation by averaging,” *SIAM J. Cont. Opt.* **30**, 838–855 (1992).
 - [51] C. Hartmann and C. Schütte, “Efficient rare event simulation by optimal nonequilibrium forcing,” *J. Stat. Mech.* **2012**, P11004 (2012).
 - [52] C. Hartmann, R. Banisch, M. Sarich, T. Badowski, and C. Schütte, “Characterization of rare events in molecular dynamics,” *Entropy* **16**, 350 (2014).
 - [53] W. Zhang, H. Wang, C. Hartmann, M. Weber, and C. Schütte, “Applications of the cross-entropy method to importance sampling and optimal control of diffusions,” *SIAM J. Sci. Comp.* **36**, A2654–A2672 (2014).
 - [54] C. M. Rohwer, F. Angeletti, and H. Touchette, “Convergence of large deviation estimators,” *Phys. Rev. E* **92**, 052104 (2015).
 - [55] T. Nemoto and S.-I. Sasa, “Computation of large deviation statistics via iterative measurement-and-feedback procedure,” *Phys. Rev. Lett.* **112**, 090602 (2014).
 - [56] H. Risken, *The Fokker-Planck Equation: Methods of Solution and Applications*, 3rd ed. (Springer, Berlin, 1996).
 - [57] P. Reimann, “Brownian motors: Noisy transport far from equilibrium,” *Phys. Rep.* **361**, 57–265 (2002).
 - [58] S. Ciliberto, S. Joubaud, and A. Petrosyan, “Fluctuations in out-of-equilibrium systems: From theory to experiment,” *J. Stat. Mech.* **2010**, P12003 (2010).
 - [59] P. Tsobgni Nyawo and H. Touchette, “Large deviations of the current for driven periodic diffusions,” *Phys. Rev. E* **94**, 032101 (2016).
 - [60] J. Mehl, T. Speck, and U. Seifert, “Large deviation function for entropy production in driven one-dimensional systems,” *Phys. Rev. E* **78**, 011123 (2008).
 - [61] T. Nemoto and S.-I. Sasa, “Variational formula for experimental determination of high-order correlations of current fluctuations in driven systems,” *Phys. Rev. E* **83**, 030105 (2011).
 - [62] P. Dupuis and H. Wang, “Dynamic importance sampling for uniformly recurrent Markov chains,” *Ann. Appl. Prob.* **15**, 1–38 (2005).
 - [63] T. Nemoto, R. L. Jack, and V. Lecomte, “Finite-size scaling of a first-order dynamical phase transition: Adaptive population dynamics and an effective model,” *Phys. Rev. Lett.* **118**, 115702 (2017).
 - [64] J. Blanchet and H. Lam, “State-dependent importance sampling for rare-event simulation: An overview and recent advances,” *Surv. Oper. Res. Manag. Sci.* **17**, 38–59 (2012).
 - [65] H. J. Kappen and H. C. Ruiz, “Adaptive importance sampling for control and inference,” *J. Stat. Phys.* **162**, 1244–1266 (2016).
 - [66] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, “Quantum Monte Carlo simulations of solids,” *Rev. Mod. Phys.* **73**, 33 (2001).
 - [67] L.-H. Lim and J. Weare, “Fast randomized iteration: Diffusion Monte Carlo through the lens of numerical linear algebra,” *SIAM Rev.* **59**, 547–587 (2017).