



ENRIC RIBERA BORRELL¹, JANNES QUER, LORENZ RICHTER², CHRISTOF
SCHÜTTE

**Improving control based importance
sampling strategies for metastable diffusions
via adapted metadynamics**

¹  0000-0002-3897-5984

²  0000-0001-5028-5639

Zuse Institute Berlin
Takustr. 7
14195 Berlin
Germany

Telephone: +49 30 84185-0
Telefax: +49 30 84185-125

E-mail: bibliothek@zib.de
URL: <http://www.zib.de>

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Improving control based importance sampling strategies for metastable diffusions via adapted metadynamics

Enric Ribera Borrell^{1, 2}, Janes Quer¹, Lorenz Richter^{1,3,4} and Christof Schütte^{1,2}

¹*Institute of Mathematics, Free University Berlin, 14195 Berlin, Germany*

²*Zuse Institute Berlin, 14195 Berlin, Germany*

³*Institute of Mathematics, BTU Cottbus-Senftenberg, 03046 Cottbus, Germany*

⁴*dida Datenschmiede GmbH, 10827 Berlin, Germany*

June 13, 2022

Abstract

Sampling rare events in metastable dynamical systems is often a computationally expensive task and one needs to resort to enhanced sampling methods such as importance sampling. Since we can formulate the problem of finding optimal importance sampling controls as a stochastic optimization problem, this then brings additional numerical challenges and the convergence of corresponding algorithms might as well suffer from metastability. In this article, we address this issue by combining systematic control approaches with the heuristic adaptive metadynamics method. Crucially, we approximate the importance sampling control by a neural network, which makes the algorithm in principle feasible for high dimensional applications. We can numerically demonstrate in relevant metastable problems that our algorithm is more effective than previous attempts and that only the combination of the two approaches leads to a satisfying convergence and therefore to an efficient sampling in certain metastable settings.

1 Introduction

The accurate computation of rare events is of great importance in multiple applications, relating to fields such as molecular dynamics, epidemiology, engineering or finance, to name just a few. One is typically interested in events that happen only very rarely, but are still relevant for certain phenomena of interest. Since analytical computations are mostly infeasible in practice, one usually relies on Monte Carlo approximations for the desired quantities. The related sampling problem, however, can be very challenging mainly for two reasons: a potentially high dimension of the problem at hand as well as large statistical errors of corresponding estimators, which are rooted in the characteristic of the events being rare. Loosely speaking, the difficulty of sampling rare events is based on its very definition: it is hard to observe an event (frequently) if it almost never appears (at least in relation to the typical timescales for which a simulation is feasible). In fact, the characteristic exponential divergence of the relative error with the parameter that controls the rarity of the quantity of interest poses great computational challenges.

In this article we shall focus on rare events in stochastic processes, where one is interested in sampling regions of the state space which are unlikely to be visited. In particular, we are interested in processes that exhibit some sort of metastability, where particles that follow the dynamics stay in certain regions of the space for a very long time. In fact, the average waiting of switching between metastable events is orders of magnitude longer than the timescale of the process itself. This is for instance typical in molecular simulations with particles following the Langevin dynamics in which a potential function governs the evolution of the stochastic process, see e.g. [14]. Here, metastable regions correspond to local minima of the potential, which are separated by so-called energy barriers, and transitions between those regions are of interest since they correspond to macroscopic properties of corresponding molecules. These are for instance reaction rates or conformation changes, such as the folding of a protein or a phase transition. However, those transitions happen only very rarely so that a simulation of transition trajectories can be extremely difficult from a computational point of view. On the one hand, the time to overcome energy barriers might be extremely large (in fact, it scales exponentially with the height of the energy barrier [2]), on the other hand, variances

of estimators related to those rare transitions might be large¹.

One idea to overcome those challenges is to apply importance sampling. Abstractly speaking, the basic idea is to sample from another probability distribution and weight the resulting random variables back in order to still get an unbiased estimator for the quantity of interest. Since we are interested in path-dependent quantities we consider importance sampling in the space of continuous trajectories. This corresponds to adding a function to the drift of original dynamics. One can think of the additional function as a control function or force that pushes trajectories into desired regions of the state space and thereby allows for overcoming possible energy barriers. Equivalently, one can think of modifying the original physical potential such that it appears less “rugged” and particles are no longer trapped in local minima. In principle, it is possible to design modifications of the potential rather freely. However, one has to keep in mind that these modifications influence the quality of the importance sampling estimator significantly, see e.g. [16]. A systematic approach for finding good control functions that aim to minimize the variance of the estimator is related to a stochastic optimal control problem [18] (for further variational perspectives we refer to [33]). This perspective then allows for numerical strategies such as iterative stochastic optimization methods that aim to find efficient controls in practice. At the same time, especially in metastable situations, those approaches hold two additional challenges that might make corresponding algorithms infeasible in applications:

In order to compute a first iteration in the stochastic optimization procedure the rare event of interest must at least be simulated once. If this does not happen, one can usually not proceed.

Even if one manages to simulate rare events with great computational effort, the estimated objectives in the optimization routines as well as their gradients might suffer from high variances, which might make convergence of the method very slow.

In this article we develop an algorithm that shall address these two aspects and improve importance sampling based estimation in metastable scenarios. In particular, we will combine systematic control-based approaches with heuristic adaptive methods that are related to the so-called metadynamics algorithm.

1.1 Previous work

We have mentioned before that rare event sampling occurs in multiple different fields of application, where each field adds a different perspective. In the sequel, let us review some of those perspectives and relate to works that are relevant for our endeavor.

Adaptive biasing techniques Methods that aim to modify the potential on the fly in order to remove metastable features of the dynamics depending on the particles in the simulation are often subsumed under the term *adaptive biasing techniques*. A well known method is called metadynamics [28], which was developed in order to improve the sampling related to stationary distributions of complex molecular systems. Many extensions and applications have been published throughout the last years. For a good review on recent developments we refer to [4, 41] and the references therein. Convergence results and high dimensional adaptations can be found in [12, 25]. An extension to importance sampling for path dependent properties of interest has been proposed in [37] and similar ideas based on the adaptive biasing force technique has been suggested by [43]. This method has been used in many applications and different extensions have for instance been proposed in [3, 21]. To our knowledge, for the adaptive biasing force methods no extension for path dependent quantities has been considered yet. Let us also note that related non-equilibrium methods have been addressed, see e.g. [27] or [47], but as before, an extension to path dependent problems is usually not covered.

Rare event sampling in an asymptotic regime Many methods for rare event estimation have been developed in an asymptotic regime, relying usually on large deviation arguments. Those strategies are often connected to the associated Hamilton-Jacobi-Bellman equations and can for instance be found in [5, 6, 8]. For variance reduction strategies in a zero noise limit we refer to [42], which relies on optimal control strategies of the corresponding deterministic problem. The special situation of attractors with resting points has been addressed in [7], and in [40] the variance of importance sampling based on asymptotic arguments applied in a nonasymptotic regime has been analyzed. Even though all of these method can be applied to path dependent quantities, an application to high dimensional applications is usually not addressed.

¹Note that those two aspects usually interact.

Nonasymptotic importance sampling Importance sampling in a nonasymptotic regime, which targets at sampling path dependent quantities, corresponds to a controlled stochastic process, see e.g. [31]. A strategy that aims to identify optimal importance sampling controls has been suggested in [18]. Numerically, the approach rests on the approximation of the control by a linear combination of ansatz functions. In [17] the corresponding method is analyzed from the perspective of path space measures and variational formulations of the problem are considered. Further variational perspectives have been suggested in [33], putting additional emphasis on certain numerical robustness properties and allowing for high-dimensional applications by modeling the control with neural networks. For strategies that are based on backward stochastic differential equations we refer for instance to [15]. The optimal control attempt has also been combined with model reduction techniques in [19, 20] and [45], noting that one of the main drawbacks of this approach is the placing of ansatz functions over the domain of interest. For a statistical analysis of importance sampling in path space that highlights its non-robustness in particular in high dimensions we refer to [16]. We also refer to [38] for a comprehensive introduction to nonasymptotic importance sampling for path functionals.

Optimal control problems Due to the connections of (optimal) importance sampling and optimal control theory we may as well refer to numerical strategies that allow to solve (high-dimensional) control problems. One strategy is to solve the related Hamilton-Jacobi-Bellman equation, which can for instance be tackled with deep learning based strategies in high-dimensions, see e.g. [13, 33, 44, 46]. Let us in particular refer to [32], where elliptic partial differential equations (PDEs) are considered, which are relevant for the problems we focus on in this article. Let us further highlight [33] where robustness properties of loss functions have been analyzed, leading in particular to the novel log-variance divergence, which exhibits favorable numerical properties. For approximating control functions with tensor trains we for instance refer to [10].

1.2 Outline of the article

The article is structured as follows. In Section 2 we state the rare event problem and discuss issues appearing in naive Monte Carlo estimations. In Section 2.1 we introduce importance sampling as a strategy to overcome those issues and in Section 2.2 we subsequently show how one can aim for optimal importance sampling strategies by deriving an equivalent optimal control problem via PDE arguments. In Section 3 we then address computational aspects of solving this control problem via an optimization approach. In particular, Section 3.1 is devoted to the computation of gradients that are needed in iterative optimization methods, Section 3.2 introduces the metadynamics based initialization method and Section 3.3 discusses the approximation of the control functions via neural networks, which then allows us to pose our final algorithm. In Section 4 we subsequently demonstrate in different numerical examples that the algorithm can significantly improve sampling performance in high-dimensional metastable scenarios. Finally, Section 5 provides a conclusion and an outlook for further research questions. For the proofs and additional statements we refer to Appendix A.

2 Sampling metastable dynamics

In this article we focus on stochastic dynamical systems which exhibit metastable features. To be precise, we consider the overdamped Langevin equation

$$dX_s = -\nabla V(X_s)ds + \sigma(X_s)dW_s, \quad X_0 = x \in \mathbb{R}^d, \quad (1)$$

on a bounded domain $D \subset \mathbb{R}^d$, where $(W_s)_{s \geq 0}$ is a d -dimensional Brownian motion. The function $V : \mathbb{R}^d \rightarrow \mathbb{R}$ shall be understood as a potential that for instance governs the dynamics of multiple atoms in a physical system, and for the diffusion coefficient we usually choose $\sigma(x) = \sqrt{2\beta^{-1}} \text{Id}$, where $\beta > 0$ denotes the inverse temperature². We assume that there exists a unique strong solution to SDE (1) and that the resulting process X is ergodic such that we can guarantee convergence to a unique equilibrium distribution, see e.g. [29] or [36] for details on these assumptions.

Given a target set $T \subset D$, let us define the first hitting time of the process X as

$$\tau := \inf\{s > 0 \mid X_s \in T\},$$

²Note that in principle both the potential and the diffusion coefficient could be made time-dependent as well.

and note that it is a.s. finite³. We can now define our quantity of interest $I : C([0, 1], \mathbb{R}^d) \rightarrow \mathbb{R}$ as

$$I(X) := \exp(-W(X)), \quad (2)$$

with the path functional $W : C([0, 1], \mathbb{R}^d) \rightarrow \mathbb{R}$ given as

$$W(X) := \int_0^1 f(X_s) ds + g(X), \quad (3)$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ are such that W is integrable. Our goal is to compute the expectation value of the quantity of interest I ,

$$I(x) := \mathbb{E}^x[I(X)], \quad (4)$$

which we can view as a function of the initial value x , where we introduce the shorthand notation $\mathbb{E}^x[I(X)] := \mathbb{E}[I(X) | X_0 = x]$. Let us recall that a stochastic process, such as the one defined in (1), is metastable if its dynamic behavior is characterized by unlikely transition events between the so-called metastable regions. In the particular case of an overdamped Langevin process, one can distinguish between two types of metastability, coming either from energetic or entropic barriers [29]. In this article we focus on the former. Let us recall that in this case both the temperature β^{-1} and the height of the energetic barriers determine the strength of the metastability [2]. In particular, by Kramer's law the mean hitting time satisfies the large deviations asymptotics

$$\mathbb{E}[\tau] \sim \exp\left(\frac{2V}{\beta}\right) \quad \text{as } \beta \rightarrow \infty, \quad (5)$$

where V is the energy barrier that the dynamics has to overcome in order to reach the target set T . An illustration of this exponential dependency is provided in Example 2.1.

Example 2.1 (Double well potential). *For an illustration, let us consider the one-dimensional double well potential*

$$V(x) = \alpha(x^2 - 1)^2, \quad (6)$$

where $\alpha > 0$ modulates the height of the energetic barrier and thereby influences the strength of the metastability, see the left panel of Figure 1. Let us consider the initial value $x = -1$ in the left well of the potential. We choose the target set $T = [1, 3]$ to be supported in the right well so that the particles need to cross the potential barrier. In the right panel of Figure 1 we plot the expected hitting time of reaching T for different values of α and β when using naive Monte Carlo estimations. Indeed we observe the exponential dependence as indicated by (5). As mentioned above, one can aim to speed up sampling by reducing the trajectory lengths when applying an importance sampling based sampling scheme. It turns out that there is an optimal way to design such a scheme, leading to substantially reduced mean hitting times which do not scale exponentially with the energy barrier anymore. We will show how to design this optimal scheme in the upcoming sections.

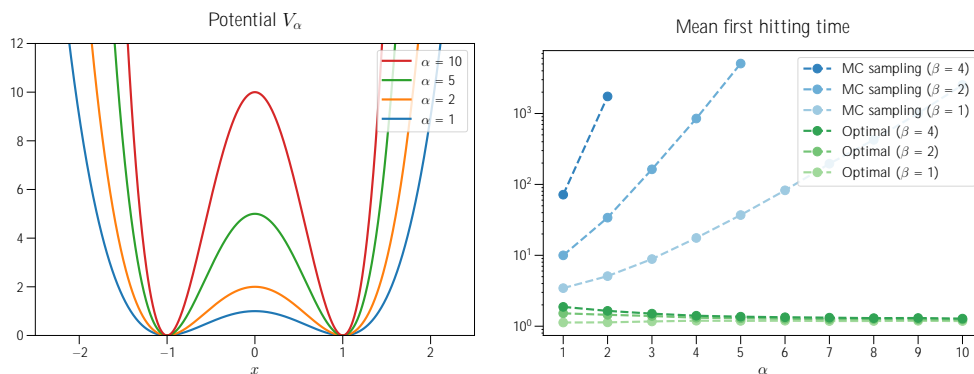


Figure 1: Left panel: The double well potential for different values of α . Right panel: Mean first hitting times for different values of α and β .

³Let $\tilde{\gamma}(y; t|x)$ be the probability distribution of the particles which have not arrived in T yet, given the position x and assuming the time evolution t . Then, the homogeneous (absorbing) boundary condition for the mean first hitting time implies that $\lim_{t \rightarrow \infty} \tilde{\gamma}(y; t|x) = 0$ for all $y \in \mathbb{R}^d$, which means that the particles will eventually leave the domain and therefore $\tau < \infty$ holds almost surely [36].

2.1 Monte Carlo approximations and importance sampling

Since there is no closed-form formula available for the computation of the expectation value (4), we must rely on its Monte Carlo estimator

$$\hat{\cdot}(x) := \frac{1}{K} \sum_{k=1}^K I(X^{(k)}), \quad (7)$$

where $fX^{(k)}g_{k=1}^K$ are independent realizations of the process X , all starting at x . For a finite sample size K the estimator is unbiased and the usual behaviors of the variance and the relative error hold, i.e.

$$\text{Var}(\hat{\cdot}(x)) = \frac{\text{Var}(I(X))}{K}, \quad \text{RE}(\hat{\cdot}) := \frac{\sqrt{\text{Var}(\hat{\cdot}(x))}}{\mathbb{E}^x[I(X)]} = \frac{\text{RE}(I(X))}{\sqrt{K}} \quad (8)$$

for any $x \geq \mathbb{R}^d$. In a metastable system the intrinsic relative error of the quantity of interest, $\text{RE}(I(X))$, can be very large. As a consequence, reducing the relative error of the Monte Carlo estimator underneath a prescribed positive value $\varepsilon > 0$, i.e. $\text{RE}(\hat{\cdot}) \leq \varepsilon$, might imply that one needs a very large number of trajectories, namely $K \geq (\text{RE}(I(X))/\varepsilon)^2$. Thus, in order to make numerical estimations feasible, one often needs to rely on methods that reduce the inherent variance of the corresponding stochastic quantities. One such method is importance sampling, on which we shall focus in the sequel.

The general idea of importance sampling is to draw random variables from another probability measure and subsequently weight them back in order to still have an unbiased estimator of the desired quantity of interest [35]. In the case of stochastic processes this change of measure corresponds to adding a control to the original process (1), yielding the controlled dynamics

$$dX_s^u = (-rV(X_s^u) + \sigma(X_s^u)u(X_s^u))ds + \sigma(X_s^u)dW_s, \quad X_0^u = x, \quad (9)$$

where the control u is an Itô integrable function that satisfies a linear growth condition, i.e. $u \geq U$ with

$$U = \tilde{f}u \geq C^1(\mathbb{R}^d, \mathbb{R}^d) : u \text{ grows at least linearly in } xg.$$

Further details can be found in e.g. [16, 17, 18, 33]. The controlled dynamics (9) can now be related to the original one (1) via a change of measure in path space, which can be made explicit via Girsanov's formula (see Appendix A.2 for details). To be precise, it holds

$$\mathbb{E}^x[I(X)] = \mathbb{E}^x[I(X^u)M^u], \quad (10)$$

where the exponential martingale

$$M^u := \exp\left(\int_0^u u(X_s^u) dW_s - \frac{1}{2} \int_0^u |ju(X_s^u)|^2 ds\right) \quad (11)$$

corrects for the induced bias.

Relating to Example 2.1, the control $u \geq U$ can intuitively be understood as an external force aiming to push particles over the energy barrier such that they can escape from metastable regions and reach desired target sets. In principle, the importance sampling relation (10) stays intact for any $u \geq U$, however, it turns out that the variance of corresponding estimators significantly depends on an appropriate choice of u , see [16]. In particular, it does not suffice to somehow push particles over existing barriers – instead, the specific control protocol needs to be chosen very carefully. Clearly, a natural goal for designing an optimal control $u \geq U$ is to aim for minimizing the variance of the importance sampling estimator, i.e.

$$\text{Var}(I(X^u)M^u) = \inf_{u \geq U} \text{Var}(I(X^u)M^u)g. \quad (12)$$

In the next section we shall discuss how this objective can in fact be linked to a classical optimal control problem, which will subsequently lead to feasible numerical strategies.

2.2 Optimal control characterizations and associated boundary value problems

In order to derive the connection between variance minimization as stated in (12) and a classical optimal control problem, we will essentially argue via PDEs that are associated to our estimation problem⁴. Let us first recall via the Feynman-Kac theorem [29, Proposition 6.1] that the expectation $\mathbb{E}[\phi(X_T) | \mathcal{F}_t^X]$ (considered as a function of the initial value), as defined in (4), fulfills the elliptic boundary value problem

$$(L - f(x)) \phi(x) = 0 \quad x \in S, \quad (13a)$$

$$\phi(x) = \exp(-\int_t^T g(x_s) ds) \quad x \in \partial S, \quad (13b)$$

where L is the infinitesimal generator of the process X , defined as

$$L = \frac{1}{2} \sum_{i,j=1}^d (\sigma \sigma^\top)_{ij}(x) \frac{\partial^2}{\partial x_i \partial x_j} + \sum_{i=1}^d \frac{\partial}{\partial x_i} V(x) \frac{\partial}{\partial x_i}.$$

The domain $S := D \setminus T^c$ is assumed to be bounded and the functions $f \in C(\mathbb{R}^d, \mathbb{R})$, $g \in C^2(\mathbb{R}^d, \mathbb{R})$ are the same as in (3).

The connection of our estimation problem to an optimal control problem can be revealed when applying the Hopf-Cole transformation (see e.g. Section 4.4.1 in [9], cf. [18]) to the solution of the PDE (13), namely

$$\phi(x) = \exp(-\int_t^T \psi(x_s) ds). \quad (14)$$

One can readily show that ψ now fulfills the non-linear boundary value problem

$$L \psi(x) - \frac{1}{2} \|\sigma^\top \psi(x)\|^2 + f(x) = 0 \quad x \in S, \quad (15a)$$

$$\psi(x) = g(x) \quad x \in \partial S. \quad (15b)$$

The PDE (15) is known as Hamilton-Jacobi-Bellman (HJB) equation, which is a key equation in optimal control theory allowing for a characterization of optimal control strategies. In fact, we can now identify the control problem that corresponds to the above PDE and thus to our estimation problem, by stating the cost functional

$$J(u; x) := \mathbb{E}^x \left[W(X^u) + \frac{1}{2} \int_0^u \|\sigma^\top(X_s^u)\|^2 ds \right], \quad (16)$$

where X^u follows the controlled dynamics as defined in (9) and f and g can be interpreted as running and terminal costs, respectively. The solution to PDE (15) is sometimes called value function in the sense that it offers the optimal *costs-to-go*, depending on the initial value x , i.e.

$$\psi(x) = \inf_{u \in U} J(u; x). \quad (17)$$

Let us make the above observations precise.

Proposition 2.2 (Variance minimization as control problem). *Let us assume there exist solutions $\psi \in C_b^2(\mathbb{R}^d, \mathbb{R})$ and $\phi \in C_b^2(\mathbb{R}^d, \mathbb{R})$ to the elliptic boundary value problems (13) and (15), respectively, and set*

$$u = \sigma^\top \psi - \sigma^\top \psi \log \phi. \quad (18)$$

Then the following are equivalent:

- (i) $u \in U$ minimizes the control costs as defined in (16).
- (ii) $u \in U$ minimizes the variance of the importance sampling estimator as defined in (12).

In fact it holds

$$\text{Var} \left(I(X^u) M^u \right) = 0. \quad (19)$$

Proof. See e.g. [17, Theorem 2] or [33, Theorem 2.2]. □

⁴Note that an alternative derivation can be achieved via certain divergences between path space measures, see [33].

Proposition 2.2 shows that u , which minimizes either (12) or (16), can be recovered from the solution of the HJB equation. This reveals that the optimal control is in fact of gradient form, just as the drift of our original stochastic process (1). We can therefore express the overall drift of the optimally controlled process as

$$r(V + V_{\text{bias}}), \quad (20)$$

where $V_{\text{bias}} = \sigma \sigma^>$ is sometimes called the optimal bias potential, which can be interpreted as being the optimal correction of the original potential V in terms of variance reduction. As stated in Proposition 2.2, one can show that it is optimal in the sense that it drives the variance of the importance sample estimator to zero, thereby yielding a perfect sampling scheme. Let us illustrate this by referring again to Example 2.1, where we have considered a stereotypical double well potential with different energetic barriers depending on the parameter $\alpha > 0$. In Figure 2 we display the optimal control functions and optimal bias potentials, respectively, that allow the trajectories to cross the barrier – note that the control is particularly large in regions where the particles get trapped when not applying the control. The optimal solutions are calculated via a finite difference discretization of the corresponding PDE (13). In Figure 3 we display the resulting Monte Carlo estimators and corresponding relative errors when using either naive Monte Carlo or the optimal importance sampling estimator. Note that the estimators are indeed much more accurate when relying on the optimal importance sampling control. We do not observe a zero relative error⁵ due to the discretization of the process with different step-sizes $\Delta t > 0$, see also Section 4.

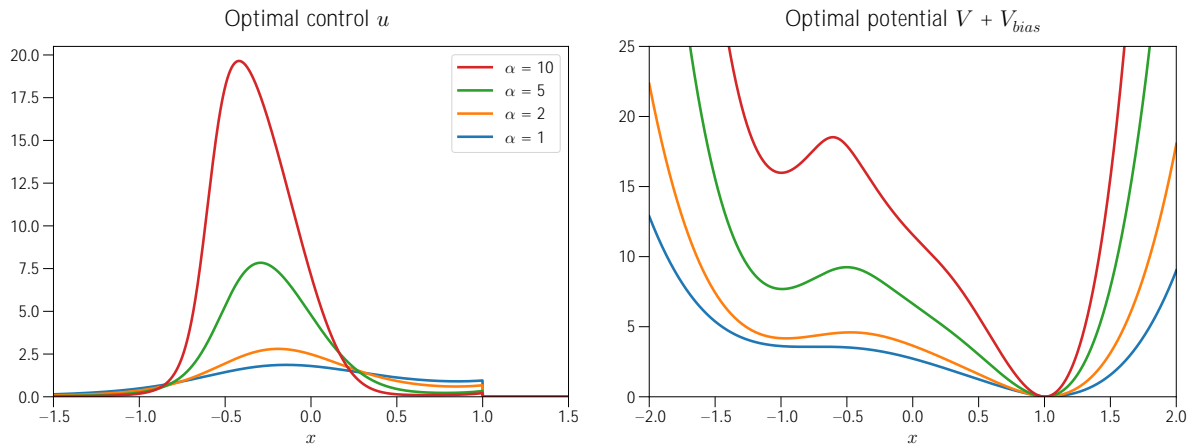


Figure 2: Optimal control functions and optimal potentials for different values of α and inverse temperature $\beta = 1$, see also Example 2.1.

⁵The observation that the relative error for the naive Monte Carlo estimator seems to increase with decreasing Δt is misleading and seems to be due to the fact that hitting times can be simulated more accurately, which leads to smaller values of $\hat{\tau}$ and therefore larger relative errors, see also the left panel of Figure 3.

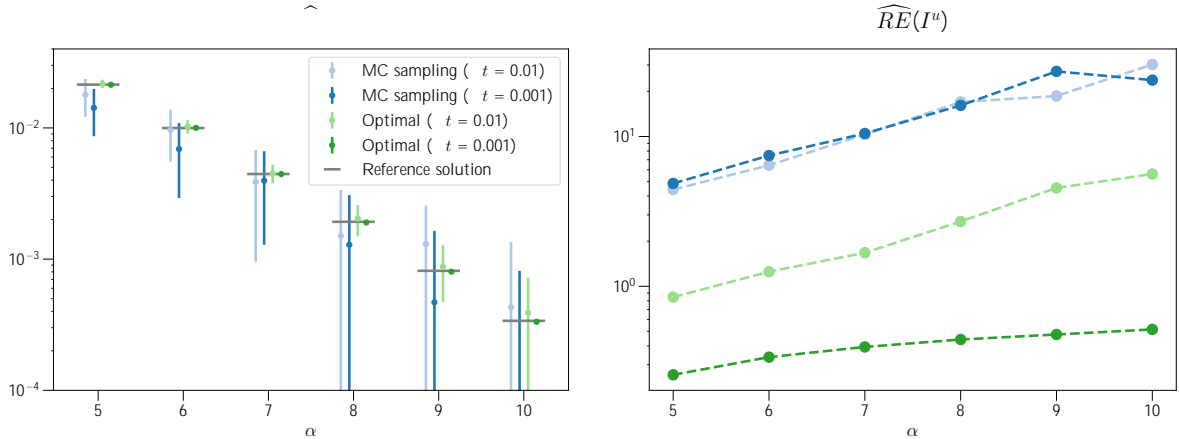


Figure 3: We compare the naive Monte Carlo estimator with optimal importance sampling relying on a discretization of PDE (13) for different values of α , inverse temperature $\beta = 1$ and different time steps t with sample size $K = 10^3$. The confidence intervals are computed according to (40).

3 Numerical strategies for solving the optimal control problem

Numerically solving optimal control problems as the one in (13) or (almost equivalently) solving high-dimensional PDEs such as the one stated in (15) can be challenging. In particular in high dimensional settings, this task seems hopeless when relying on classical grid based methods such as finite differences or finite elements since these methods suffer from the curse of dimensionality [44]. We will therefore work with an approach that relies on an optimization procedure aiming to iteratively minimize the cost functional (16) over a prescribed function class in the spirit of machine learning (cf. [33]). Our envisioned procedure can be described as follows:

- (i) Initialize the control u with an appropriate choice u^{init} .
- (ii) Simulate realizations of the controlled process X^u as defined in (9).
- (iii) Compute an estimator of the cost functional $J(u; x)$ as well as its derivative with respect to u .
- (iv) Update the control u by gradient descent.
- (v) Repeat steps (ii)–(iv) until convergence.

We argue that two aspects are crucial in order to implement the above scheme in practice. On the one hand, we need to design gradient estimators that are feasible in the sense that they can cope with random stopping times and at the same time exhibit sufficiently low variances. On the other hand, in particular in metastable settings, an appropriate initialization u^{init} is important. Otherwise initial gradient information might turn out to be useless for instance due to increased variances or due to long trajectory simulations. As a remedy, we suggest a novel simulation algorithm which tries to combine ideas from existing approaches and thereby overcome these known issues. In particular, we will suggest to identify feasible control initializations coming from an adapted version of the heuristic metadynamics algorithm.

3.1 Gradient computations

Let us first address the issue of computing gradients of the cost functional

$$J(u; x) = \mathbb{E}^x \left[W(X^u) + \frac{1}{2} \int_0^u j u(X_s^u)^2 ds \right], \quad (21)$$

as already defined in (16), with respect to the control. An inherent difficulty is that both the running costs and the process depend on the control u , and the latter implies that also the hitting time τ^u depends on u . We will approach this difficulty by an appropriate change of the path space measure and first compute a functional derivative in the Gâteaux sense. Subsequently we can relate the rather abstract result to implementable gradients by considering controls $u = u_\theta$ that are parametrized by a parameter vector $\theta \in \mathbb{R}^p$.

The Gâteaux derivative can then be identified with a gradient with respect to θ by considering a special Gâteaux derivative direction. Notably, this strategy will lead to a gradient estimator that can cope with the fact that the random hitting time τ^u appears both in the integral limit as well as in the terminal costs and depends on the function u , with respect to which we differentiate. Eventually, we can numerically compute our gradient estimator by a Monte Carlo approximation.

Let us start by recalling the definition of the Gâteaux derivative.

Definition 3.1 (Gâteaux derivative). We say that $J : U \rightarrow \mathbb{R}$ is *Gâteaux differentiable* at $u \in U$ if for all $\phi \in U$ the mapping

$$\varepsilon \mapsto J(u + \varepsilon\phi; x) \quad (22)$$

is differentiable at $\varepsilon = 0$. The Gâteaux derivative of J in direction ϕ is then defined as

$$\frac{\delta}{\delta u} J(u; x; \phi) := \left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} J(u + \varepsilon\phi; x). \quad (23)$$

We can now compute the functional derivative of $J(u; x)$.

Proposition 3.2 (Gâteaux derivative of cost functional). *The Gâteaux derivative of the cost functional defined in (16) in the direction $\phi \in U$ is given by*

$$\frac{\delta}{\delta u} J(u; x; \phi) = \mathbb{E}^x \left[\int_0^u (u - \phi)(X_s^u) ds + \left(W(X^u) + \frac{1}{2} \int_0^u j u (X_s^u)^2 ds \right) \int_0^u \phi(X_s^u) dW_s \right]. \quad (24)$$

Proof. See Appendix A.3. □

Proposition 3.2 is valid for any direction $\phi \in U$. Let us note that we are particularly interested in the directions⁶ $\phi = \frac{\partial}{\partial \theta_i} u$ for all $i \in \{1, \dots, pg\}$. This choice is motivated by the chain rule of the Gâteaux derivative, which, under suitable assumptions, states that

$$\frac{\partial}{\partial \theta_i} J(u; x) = \left. \frac{\delta}{\delta u} \right|_{u=u_\theta} J\left(u; x; \frac{\partial}{\partial \theta_i} u\right). \quad (25)$$

We therefore readily get the following formula for the gradient of J with respect to the parameter θ .

Corollary 3.3 (Gradient of cost functional). *Let $u = u_\theta$ be parametrized by the parameter vector $\theta \in \mathbb{R}^p$, then the partial derivatives of the control functional (16) with respect to the parameters are given by*

$$\begin{aligned} \frac{\partial}{\partial \theta_i} J(u; x) = \mathbb{E}^x & \left[\int_0^u \left(u - \frac{\partial}{\partial \theta_i} u \right) (X_s^{u_\theta}) ds \right. \\ & \left. + \left(W(X^{u_\theta}) + \frac{1}{2} \int_0^u j u (X_s^{u_\theta})^2 ds \right) \int_0^u \left(\frac{\partial}{\partial \theta_i} u \right) (X_s^{u_\theta}) dW_s \right] \end{aligned} \quad (26)$$

for any $i \in \{1, \dots, pg\}$.

Remark 3.4. *Note that the gradient given by (26) is equivalent to the one derived in [18] up to discretization and up to a more general approximating function. For convenience, we repeat the related derivation for general parametrized functions u_θ in Appendix A.1. Further note that – contrary to the statement in [18] – our analysis shows that the gradient is in fact exact, even though it involves the random hitting time τ^u , which depends on u .*

In principle, the gradient from Corollary 3.3 can be implemented straightforwardly by Monte Carlo approximation. However, even when relying on automatic differential tools, the repeated computation of, for

⁶We assume that the functions $u_\theta, \theta \in \mathbb{R}^p$, as well as all partial derivatives $\frac{\partial}{\partial \theta_i} u_\theta$ lie in U .

instance, $\frac{\partial}{\partial \vartheta} u$ might be costly. Let us therefore state a loss functional that is more convenient from a computational point of view, namely

$$J_e(u, u_{\#}; x) = \mathbb{E}^x \left[\frac{1}{2} \int_0^u j_{u_{\#}}(X_s^{u_{\theta}})^2 ds + \left(W(X^{u_{\theta}}) + \frac{1}{2} \int_0^u j_u(X_s^{u_{\theta}})^2 ds \right) \int_0^u u_{\#}(X_s^{u_{\theta}}) dW_s \right], \quad (27)$$

which now depends on two parameter vectors $\theta, \vartheta \in \mathbb{R}^p$. It is straightforward to see that the gradient of the actual cost functional (16), stated in Corollary 3.3, can then be recovered via

$$\left. \nabla_{\#} J_e(u, u_{\#}; x) \right|_{\vartheta = \theta} = \nabla J(u; x). \quad (28)$$

In practice, setting $\vartheta = \theta$ only after the differentiation is achieved by removing the parameter θ from the computational graph of automatic differentiation. Note that with this trick only one backward pass is needed.

Both $\frac{\partial}{\partial \vartheta} J$ and J_e , as defined in (26) and (27), respectively, can now be approximated by Monte Carlo, yielding for instance the estimator

$$\begin{aligned} \widehat{J}_e(u, u_{\#}; x) = & \frac{1}{K} \sum_{k=1}^K \left(\frac{1}{2} \int_0^u j_{u_{\#}}(X_s^{u_{\theta};(k)})^2 ds \right. \\ & \left. + \left(W(X^{u_{\theta};(k)}) + \frac{1}{2} \int_0^u j_u(X_s^{u_{\theta};(k)})^2 ds \right) \int_0^u u_{\#}(X_s^{u_{\theta};(k)}) dW_s^{(k)} \right), \end{aligned} \quad (29)$$

where $X^{u_{\theta};(k)}$ and $W^{(k)}$ are i.i.d. realizations of the controlled process (9) and of Brownian motion, respectively.

3.2 Efficient initializations of stochastic optimization via metadynamics

We have so far computed a gradient estimator that allows for stochastic optimization in the spirit of reinforcement learning. To be precise, we can run gradient descent like algorithms that iteratively minimize a suitable objective function with the aim to improve the control which is applied to the dynamics. In this section we shall address the question of how to initialize the approximating function $u \in \mathcal{U}$ in such an iterative optimization procedure. This is in particular important in problems with random hitting times which depend crucially on the applied control u – as we have already seen in Figures 1-3.

Aiming for reasonable initializations of $u \in \mathcal{U}$ that can in particular cope with strong metastabilities of the dynamics, we suggest to rely on ideas from the so-called *metadynamics* algorithm. In its original form *metadynamics* is an adaptive method for sampling the free energy profile of high-dimensional molecular systems [28, 41]. Its main intention is related to sampling from corresponding stationary distributions, in particular in systems that exhibit high metastabilities. The approach can be described quite vividly: one iteratively “fills up” regions with low potential energy in order to be able to eventually escape local minima that account for metastable behaviors. Those regions can be identified as the ones where the trajectories spend a sufficient amount of time. The filling is specifically achieved by adding Gaussian functions every fixed time interval until the trajectories are eventually able to escape the local minima. Analog to (20), the idea is therefore to modify the potential via

$$V + V_{\text{bias}}, \quad (30)$$

where now the bias potential is given by a sum of unnormalized Gaussian functions, i.e.

$$V_{\text{bias}}(x; \eta, \mu, \Sigma) = \sum_{m=1}^M \eta_m \widetilde{N}(x; \mu_m, \Sigma_m), \quad (31)$$

where $\widetilde{N}(x; \mu_m, \Sigma_m)$ is an unnormalized density of a multivariate normal distribution, i.e.

$$\widetilde{N}(x; \mu_m, \Sigma_m) = \exp \left(-\frac{1}{2} (x - \mu_m)^T \Sigma_m^{-1} (x - \mu_m) \right), \quad (32)$$

with mean $\mu_m \in \mathbb{R}^d$, covariance matrix $\Sigma_m \in \mathbb{R}^{d \times d}$ and $\eta_m \in \mathbb{R}$ being an appropriate weight. The intuition is that the added bias potential prevents the trajectory from going back to the already visited states. This

bias potential can also be interpreted as a control function. The control resulting from the bias potential is then given by

$$u^{\text{init}}(x) = \sigma^{-1}(x) \nabla V_{\text{bias}}(x; \eta, \mu, \cdot) = \sigma^{-1}(x) \sum_{m=1}^M \eta_m \nabla \tilde{N}(x; \mu_m, \Sigma_m). \quad (33)$$

We use the underlying idea of the metadynamics algorithm in order to compute a reasonable initialization for our optimization procedure, however, we adapt the stopping criterion used in the original version: we add unnormalized Gaussian functions in the simulation until a certain predefined set (usually the target set) is hit by the perturbed trajectory. The intuition behind this strategy is that we essentially want to “fill up” those local minima of the potential which influence the estimation of the quantity of interest. Note that this has already been considered in [37], however, not in combination with optimal control ideas. We then take the resulting control u^{init} as a rough initial guess of the optimal control, which is likely to push trajectories over the energy barrier such that trajectories are not trapped in metastable regions anymore. The expected advantages of such a control initialization are twofold. On the one hand the trajectories (in particular at the beginning of the control optimization) are expected to be much shorter, which reduces the runtime significantly. On the other hand we expect the variances of the gradient estimators, as for instance defined in Corollary 3.3, to be smaller, which will result in faster convergence of gradient descent algorithms. Crucially, we should note that without our metadynamics based initialization method the optimization might not even converge at all. We refer to Section 4 for illustrative examples of those aspects. The basic principle of our metadynamics based initialization method is illustrated in Figure 4.

In the following let us suggest two versions of a metadynamics based initialization algorithm. The first one builds the bias potential V_{bias} by sampling just one trajectory, as already considered in [37]. This trajectory follows the dynamics of the controlled stochastic process (9) until it hits the target set T . In particular, the control u is modified on the fly after each specified time interval δ by adding another unnormalized Gaussian function to the potential according to (31) with μ_m being the averaged position of the particle over the last time interval. This method ensures that metastable regions get “filled up” when visited by the trajectory. The time interval δ , the covariance matrix of the Gaussians Σ_m and the weights η_m should be chosen such that the original potential is not perturbed too much, however still allowing for a significant reduction of the hitting time τ^u . It is possible that using different versions of the metadynamics algorithm like *well-tempered metadynamics* can lead to further improvements of the initialization procedure. For simplicity, we however choose constant weights and covariance matrices. Let us summarize our first method in Algorithm 1.

Algorithm 1 One trajectory adapted metadynamics

Consider specified target set T and dynamics X^u as in (9).
Choose time interval δ , weight $\eta > 0$ and covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$.
Initialize $X_0^u = x$, $V_{\text{bias}}^{(0)} = 0$, $m = 0$.
while trajectory has not arrived in T do
 Run dynamics X^u with control $u = \sigma^{-1} \nabla V_{\text{bias}}^{(m)}$ for time interval δ .
 Choose $\mu_m = \frac{1}{\delta} \int_m^{m+\delta} X_s^u ds$.
 Adapt bias potential $V_{\text{bias}}^{(m+1)} = V_{\text{bias}}^{(m)} + \eta \tilde{N}(\cdot; \mu_m, \Sigma)$.
 $m = m + 1$.
end while
Set $V_{\text{bias}} = V_{\text{bias}}^{(m)}$.

In high dimensional settings the considered dynamics often exhibits multiple metastable regions. In this case, a bias potential provided by only one trajectory might not be sufficient since we cannot guarantee the trajectory to visit all metastable regions before hitting the target set. Therefore, only some of those regions might be “filled up”. To deal with this issue, we suggest to sample multiple trajectories and execute corresponding bias potential modifications cumulatively. To be precise, each trajectory starts at the same initial position, one after the other. The first trajectory starts with the zero bias initialization, the second one considers the bias potential that came out after running the first trajectory, and so on. Ideally one would like to stop sampling trajectories once the bias potential is already properly “filled up”, but this can be difficult to determine. In any case, it is unlikely that we perturb the potential much more than needed since otherwise trajectories would hit the target set even before the time interval δ elapsed.

Furthermore, we introduce the scaling factor $r \in (0, 1)$ that shall stabilize the procedure by reducing the effect of adding further Gaussians for each trajectory. The cumulative method is summarized in Algorithm 2.

Algorithm 2 Cumulative adapted metadynamics

Consider specified target set T and dynamics X^u as in (9).
Choose time interval δ , number of trajectories K^{meta} , weight $\eta > 0$, scaling factor $r \in (0, 1)$ and covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$.
Initialize $V_{\text{bias}}^{(0)} = 0$.
for $k \in \{1, \dots, K^{\text{meta}}\}$ do
 Set $X_0^{u:(k)} = x, m = 0$.
 while trajectory has not arrived in T do
 Run dynamics $X^{u:(k)}$ with control $u = -\Sigma^{-1} \nabla V_{\text{bias}}^{(m)}$ for time interval δ .
 Choose $\mu_m = \frac{1}{\delta} \int_m^{m+\delta} X_s^{u:(k)} ds$.
 Adapt bias potential $V_{\text{bias}}^{(m+1)} = V_{\text{bias}}^{(m)} + r^k \eta \tilde{N}(x; \mu_m, \Sigma)$.
 $m = m + 1$.
 end while
end for
Set $V_{\text{bias}} = V_{\text{bias}}^{(m)}$.

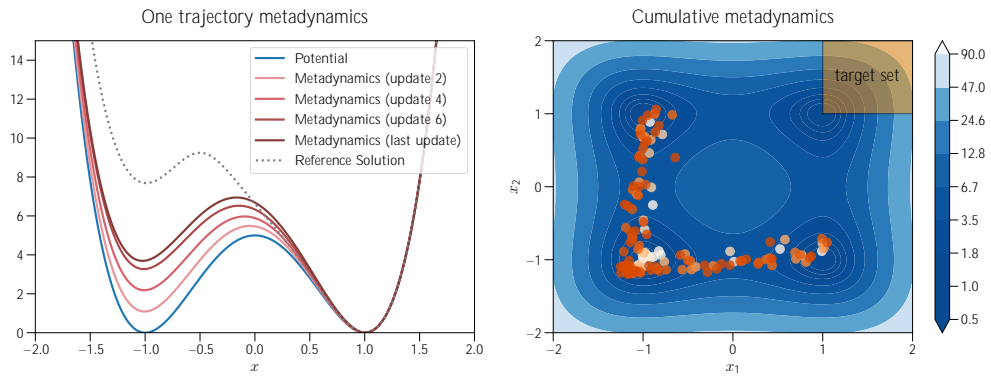


Figure 4: Left panel: perturbed potential for different updates of $V_{\text{bias}}^{(m)}$ during Algorithm 1 for the double well potential Example 2.1 with $\alpha = 5$ and $\beta = 1$. Right panel: centers of the added unnormalized Gaussian functions after Algorithm 2 for the two-dimensional extension of the double well potential. The color of the centers represent the weights of the added functions.

3.3 Control function approximations

Finally, we need to specify how to approximate the control function $u \in U$. The general idea is to rely on parametrized functions u , specified by the parameter vector $\theta \in \mathbb{R}^p$. In particular, we consider a linear combination of ansatz functions (Galerkin approach) as well as neural network approximations. The former match well with the structure of the metadynamics based initialization algorithm that we have introduced before, but suffer from the curse of dimensionality. The latter seem well suited for high-dimensional problems, but need an additional step in order to benefit from our initialization strategy. Note that either function space needs to be sufficiently large in order to approximate the optimal control u well enough.

In the Galerkin approach the control u is projected onto a space consisting of finitely many ansatz functions. A clever choice of ansatz functions depends on the problem at hand and one might for instance consider radial symmetric functions, polynomials or piecewise linear functions with Chebyshev coefficients, see e.g. [20]. As a related method let us mention tensor train approximations and refer to [10, 39] for further details. In this work we rely on Gaussian ansatz functions since they match well with the aforementioned initialization strategy. To be precise, let us choose the control approximation

$$u(x) = \sum_{i=1}^p \theta_i r \mathcal{N}(x; \mu_i, \Sigma), \quad (34)$$

where $N(\cdot; \mu_j, \Sigma_j)$ is the density of a multivariate normal distribution with mean $\mu_j \in \mathbb{R}^d$ and covariance matrix $\Sigma_j \in \mathbb{R}^{d \times d}$, as in (33).

Feed-forward neural networks, on the other hand, are nonlinear functions that exhibit remarkable approximation properties [1, 24]. They essentially consist of compositions of affine-linear maps and nonlinear activation functions. In particular, we define a *feed-forward neural network* $u : \mathbb{R}^{d_0} \rightarrow \mathbb{R}^{d_L}$ with L layers by

$$u(x) = A_L \rho(A_{L-1} \rho(\dots \rho(A_1 x + b_1) \dots) + b_{L-1}) + b_L, \quad (35)$$

with matrices $A_l \in \mathbb{R}^{d_l \times d_{l-1}}$, vectors $b_l \in \mathbb{R}^{d_l}$, $1 \leq l \leq L$, and a nonlinear activation function $\rho : \mathbb{R} \rightarrow \mathbb{R}$ that is applied componentwise. The collection of matrices A_l and vectors b_l comprises the learnable parameters θ . For our control approximations, we can now choose $d_0 = d_L = d$. Note that the choice of the so-called architecture of neural networks, i.e. the number of parameters in each layer, is not always straight-forward and requires some finetuning.

For initializing u with the control u^{init} obtained by one of the two adapted metadynamics algorithms, which have been suggested in Section 3.2, we can consider a least squares minimization on a given domain D . That means we minimize the loss

$$L_{\text{init}}(\theta) := \mathbb{E} \left[|u(X) - u^{\text{init}}(X)|^2 \right], \quad (36)$$

where $X \sim \nu$ is sampled from a prescribed measure ν that has full support on the domain D , e.g. the uniform measure. For the parametric approximation we can solve the minimization of the loss explicitly by solving a least squares problem. When considering neural networks we have to minimize L_{init} by some variant of gradient descent where the different parameters, such as batch size and stopping criterion, have to be chosen depending on the problem. Further details on the applied minimization method are provided in Section 4. For this minimization and for the computation of the gradient of the control cost, we rely on automatic differentiation tools such as PyTorch. For convenience let us state our final algorithm.

Algorithm 3 Efficient importance sampling

Choose a metadynamics time interval δ , weight $\eta > 0$ (potentially a scaling factor $r \in (0, 1)$), covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$, a function approximation u , a gradient based optimization algorithm, a corresponding learning rate $\lambda > 0$, a sample size K , a step size ϵ and a stopping criterion.

Compute u^{init} by either Algorithm 1 or 2.

Initialize $u = u^{\text{init}}$ by minimizing $L_{\text{init}}(\theta)$.

repeat

Simulate K samples of $X^{u:(k)}$.

Compute the gradient of the control cost estimator \hat{J}_e defined in (29) via automatic differentiation.

Update the parameters based on the optimization algorithm.

until stopping criterion is fulfilled.

Do importance sampling according to (10) with the control u .

Result: Low-variance estimate of J_e as defined in (4).

4 Numerical examples

In this section we demonstrate that our proposed Algorithm 3 can indeed lead to low-variance estimators of observables that involve random stopping times. In particular, we will show in both low and high-dimensional metastable examples that the combination of control based importance sampling together with reasonable initializations leads to improved estimators. Throughout we will consider the overdamped Langevin equation as stated in (9) with $\sigma(x) = \sqrt{2\beta^{-1}} \text{Id}$ on the domain $D = [-3, 3]^d$. We consider a multi-dimensional extension of the double well potential⁷

$$V(x) = \sum_{i=1}^d \alpha_i (x_i^2 - 1)^2, \quad (37)$$

where the parameter $\alpha \in \mathbb{R}^d$ as well as the inverse temperature β encode the strength of the metastability. We aim to compute the quantity

$$\mathbb{E} [e^{-V(x)}], \quad (38)$$

⁷Notice that even though the potential is symmetric in all dimensions we cannot decouple the estimation or control problem, respectively.

by choosing $f = 1$ and $g = 0$ in the observable (3). If not stated differently we set the initial value of the process to $x = (1, \dots, 1)^>$ for all examples. The inverse temperature is set to $\beta = 1$ so that the metastability is mainly influenced by the choice of α . If not otherwise stated the target set is chosen to be $T = [1, 3]^d$.

In our numerical simulations we discretize the controlled stochastic process in time $0 < t_1 < \dots < T_{\tilde{N}} = \tau$ using the Euler-Maruyama scheme

$$\widehat{X}_{n+1}^u = \widehat{X}_n^u + (\tau V(\widehat{X}_n^u) + \sqrt{2\beta^{-1}} u(\widehat{X}_n^u)) \tau + \sqrt{2\beta^{-1}} \xi_{n+1} \tau, \quad X_0^u = x, \quad (39)$$

where $\tau = t_{n+1} - t_n$ is a time-step and $\xi_{n+1} \sim N(0, \text{Id})$ is a standard normally distributed random variable [22]. Note that the length of each discrete trajectory is random according to $\tilde{N} = b\tau^u / \tau c$. For each experiment we monitor the importance sampling mean as the Monte Carlo estimator of (10) and its variance and relative error accordingly. For the Monte Carlo estimators we compute confidence intervals by

$$\widehat{\sigma}(x) = 1.96 \frac{\sqrt{\widehat{\text{Var}}(I^u)}}{\sqrt{K}}, \quad (40)$$

where $\widehat{\text{Var}}$ is the estimated variance computed with a sample size $K_{\text{var}} = 10^5$, and K the sample size of the Monte Carlo estimator. We also keep track of mean first hitting time $\mathbb{E}[\tau^u]$ and the time needed for the last trajectory of an ensemble to reach the target set T .

$$L^2(u) := \mathbb{E} \left[\int_0^u \int u - u_{\text{ref}} f^2(X_s^u) ds \right]. \quad (41)$$

The control approximation with Gaussian ansatz functions is done according to (34) with p Gaussians uniformly distributed over the domain D . The covariance matrix is constant, $\Sigma_i = 0.5 \text{Id}$, for all $i \in \{1, \dots, p\}$ and the number of ansatz functions changes depending on the example. The initialization of u with the control u^{init} is achieved by solving a least squares problem. For the neural network representation we consider a feed-forward neural network according to (35) with two hidden layers, $d_1 = d_2 = 30$ and activation function $\rho(x) = \tanh(x)$. The initialization of u with the control u^{init} is achieved after minimizing the mean squared error loss (36) by using the Adam algorithm with learning rate $\lambda = 0.01$ [26]. If not otherwise specified the training data points for this approximation problem have been uniformly sampled from the domain just one time and have been used for all gradient steps. A total of 10^3 gradient steps suffices to obtain a good approximation. In order to have fair comparisons we set the control $u^{(0)}$ to be the zero function when not considering a metadynamics based initialization.

Moreover, the control optimization in Algorithm 3 is implemented using the Adam algorithm with learning rate $\lambda = 0.01$. If not otherwise stated, the batch-size is set to be $K = 10^3$ and the time step $\tau = 0.005$. We repeat all of our experiments multiple times with different random seeds and different time intervals δ in order to guarantee generalizability. Each experiment requires just one CPU core and the maximum value of allocated memory is set to 100GB.

4.1 Metastable double well potential

Let us start with a one-dimensional metastable example for which $\alpha = 5$. We approximate the control with neural networks or Gaussians ansatz functions, where both are either initialized with the zero function or an initialization given by Algorithm 1, for which we choose $\delta = 0.2$, $\eta = 1.0$ and $\sigma^2 = 0.5$. For this example we consider a finer time step $\tau = 0.001$ and a batch size of $K = 1000$. The resulting modified potential consist of $M = 7$ unnormalized Gaussian functions. Figure 5 shows the L^2 approximation error compared to a reference solution as well as the relative error of the importance sampling estimator as a function of the gradient steps. We can see that the neural network performs better and that the control initializations speed up the convergence significantly. Note that the learnt importance sampling control leads to a similar relative error compared to a reference optimal control.

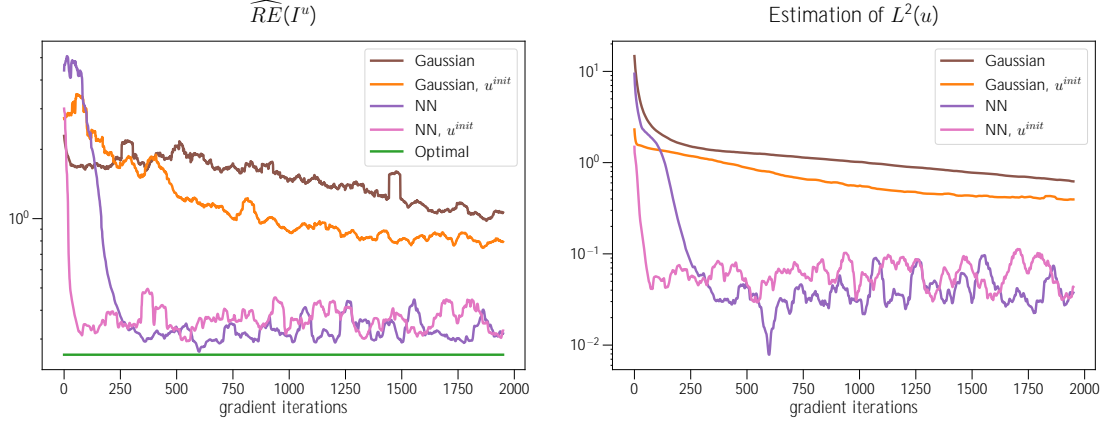


Figure 5: Importance sampling relative error and estimation of $L^2(u)$ at each gradient step.

In Figure 6 we display the approximated functions, once as the control and once as the perturbed potential. We can see that in particular the neural network approximation agrees well with the reference solution, whereas both the Gaussian approximation and the metadynamics attempt without control optimizations based control are off.

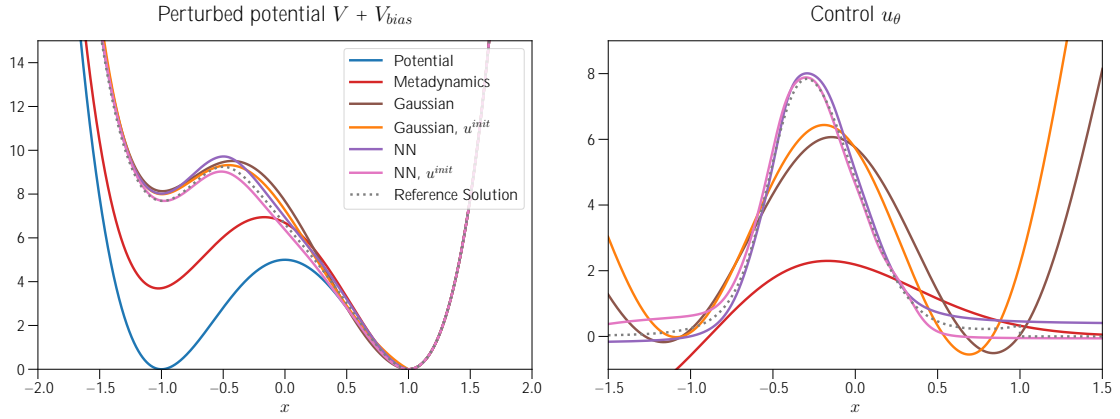


Figure 6: Perturbed potential and control after convergence.

We have stated before that the metastability depends on the parameter α . Let us therefore vary this parameter and compare performances of the following schemes against each other: naive Monte Carlo (MC sampling), a bias constructed by an adapted metadynamics algorithm without optimization as proposed in [37] (Metadynamics), a Gaussian control representation without initialization (Gaussians) as well as with the metadynamics based initialization (Gaussians, u^{init}), a neural network representation without initialization (NN) as well as with initialization (NN, u^{init}), a sampling with the discretized optimal control calculated with a finite difference method (Optimal) and the reference solution from the PDE (Reference Solution).

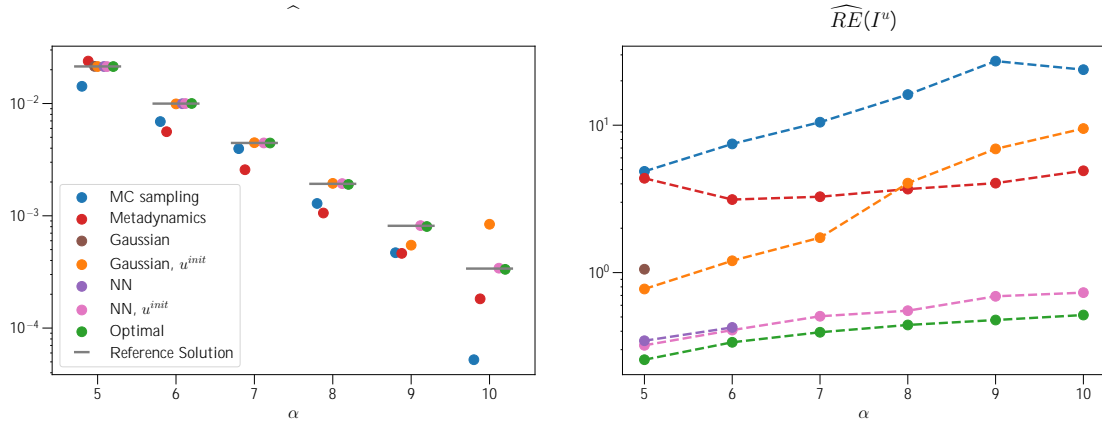


Figure 7: Estimation $\hat{\cdot}$ and importance sampling relative error of the different methods described above.

In Figure 7 we display the estimator as well as the relative errors. We can see that the estimation of the expectation value gets worse with increasing value of α , in particular when relying on naive Monte Carlo estimation, the Gaussian control approximations or the metadynamics algorithm only. We should highlight that without control initialization we are not able to get results for $\alpha > 6$ for the importance sampling estimators since corresponding control based optimization algorithms exceed the memory constraints. The reason for this is that the first sampling of the gradient estimator takes very long and thus the allocated memory capacity is exceeded. With the adapted metadynamics based initializations, on the other hand, we can observe that the optimal control importance sampling strategies yield valid estimators with low relative error even for large metastabilities. However, the neural network representation with initialization results in a much smaller relative error than the Gaussian representation also with initialization. Moreover, we stress the fact that doing importance sampling right after metadynamics does not guarantee satisfactory results. This is probably due to u^{init} still being off from u^* , noting that there is an exponential dependency on the variance in the distance of the used control to the optimal control [16].

4.2 Multi-dimensional extension of the double well potential

Let us now repeat the above experiment for multi-dimensional problems. We start with $d = 2$, for which we can still compute a reference solution by approximating PDE (13) via the finite difference method. This example is followed by an example in $d = 4$ for which the PDE can not be discretized anymore due to the curse of dimensionality. Here we compute a reference value of $\mathbb{E}[I^u]$ by Monte Carlo estimation using a very large batch size.

Example in $d = 2$

For the $2d$ example we choose $\alpha = (5, 5)^T$. As before we compute the metadynamics based initialization of the control according to Algorithm 1 with time interval $\delta = 1.0$, weight $\eta = 1.0$ and covariance matrix $\Sigma = 0.5 \text{Id}$. The resulting modified potential consists of $M = 10$ unnormalized Gaussian functions. For the linear combination of ansatz functions we choose $p = 100$ Gaussians, again placed on an equidistant grid in the domain D .

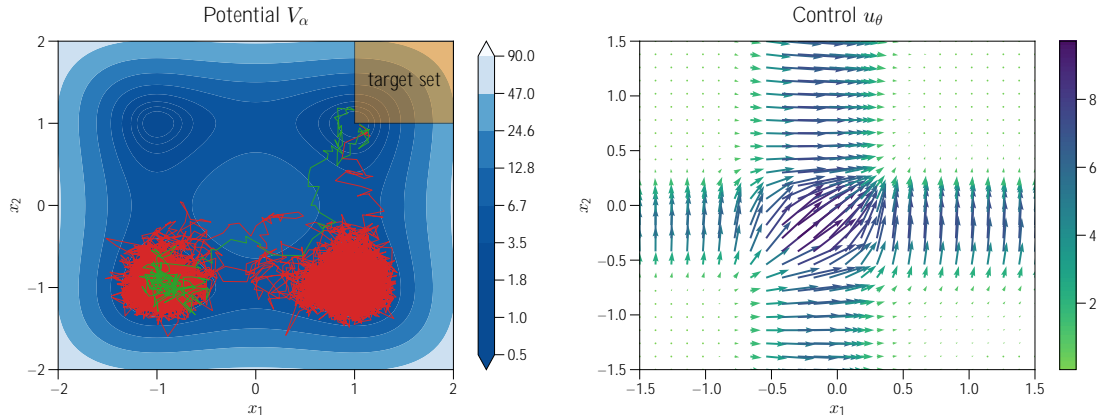


Figure 8: Left panel: The double well potential for $d = 2$ has four minima at $x = (\pm 1, \pm 1)$ and a local maximum at $x = (0, 0)$. The uncontrolled trajectory (in red) is long and gets trapped in two of the minima for a significant amount of time. The optimal importance sampling trajectory (in green) on the other hand is much shorter. Right panel: control approximated by a neural network using Algorithm 3.

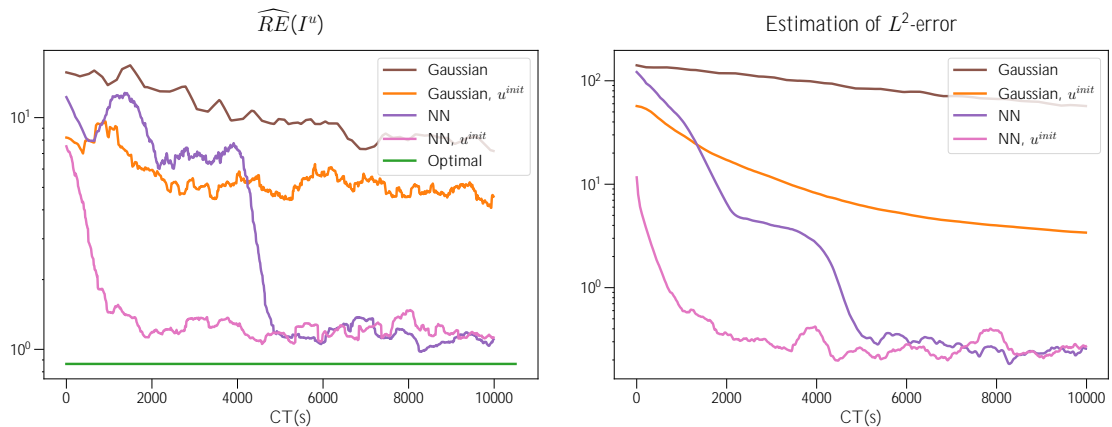


Figure 9: Importance sampling relative error and estimation of $L^2(u)$ as a function of the computation time.

Let us highlight two important aspects of the experiment. First, we see in Figure 9 that even after a runtime of 10^4 seconds both Gaussian ansatz approximation attempts in contrast neural networks have not converged. We have observed in our experiments that the Gaussian approximation is sensible with respect to the number of ansatz functions, their placing in the considered domain and the choice of the covariance matrix. Note that such hyperparameters do not have to be tuned for neural networks. Second, we can observe that the optimization initialized with the adapted metadynamics algorithm results in a faster convergence. In Figure 9 we see that our suggested approach needs only half of computational time to converge in comparison with the simulation using zero initialization. Note that the applied control yields shorter trajectories and thus reduced overall computational costs – this can for instance be seen in Figure 8.

Example in $d = 4$

Let us now consider an example in $d = 4$. We again use the potential stated in (37) and set $\alpha = (5, 5, 5, 5)^>$. Note that the potential now has 16 minima. This time we rely on Algorithm 2 for our metadynamics based control initialization, for which we choose $\delta = 5$, $K^{\text{meta}} = 100$, $\eta = 1$, $r = 0.95$ and $\beta = 0.5 \text{Id}$. The main advantage of the cumulative version of the adapted metadynamics algorithm is that we now rely on multiple trajectories for finding a good control initialization. This method is more robust and the trajectories explore a larger part of the domain. For the chosen parameters the resulting modified potential consists of $M = 284$ unnormalized Gaussian functions. For $d = 4$ a relevant step is the initialization of the control with u^{init} by minimizing the loss stated in (36). In such a case the support of the bias potential can be small in comparison with the considered domain and sampling the training data uniformly might not be feasible. Instead, we sample new training data for each gradient step following a normal distribution centered in the different unnormalized Gaussian which the chosen adapted version of the metadynamics algorithm has

added. A total of 10^4 gradient steps and 10^3 sampled points for each gradient step suffice to provide a good initialization. Figure 10 shows the estimation of $\hat{\mu}$ provided by the NN approximation with the above mentioned metadynamics based initialization as well as the relative error of the importance sampling estimator as a function of the gradient steps. As a reference value for μ we take a Monte Carlo estimator that relies on $K = 10^8$ trajectories.

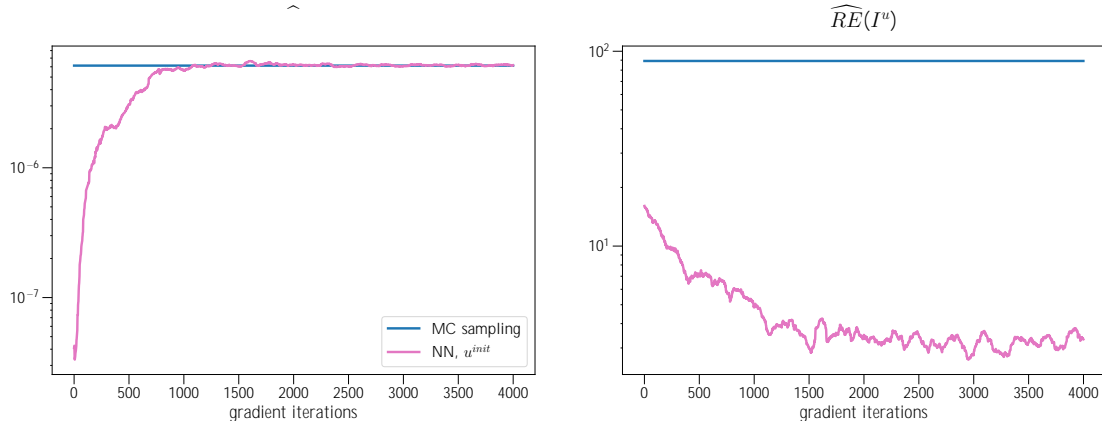


Figure 10: Estimation of $\hat{\mu}$ and the importance sampling relative error as a function of the gradient iterations.

In our experiments, it turns out that we indeed need to rely on the metadynamics based control initialization, since the zero initialization does not work due to long trajectories and memory issues. Note that the target set is much smaller in comparison to the rest of the domain, which, together with the intrinsic metastability of the system, implies very long trajectories. In Figure 10 we notice that the suggested metadynamics based initialization converges and gives an accurate estimator with a smaller relative error. In Figure 11 we observe that the estimation via naive Monte Carlo sampling requires the simulation of a huge amount of trajectories and is less accurate as the estimation via our suggested procedure.

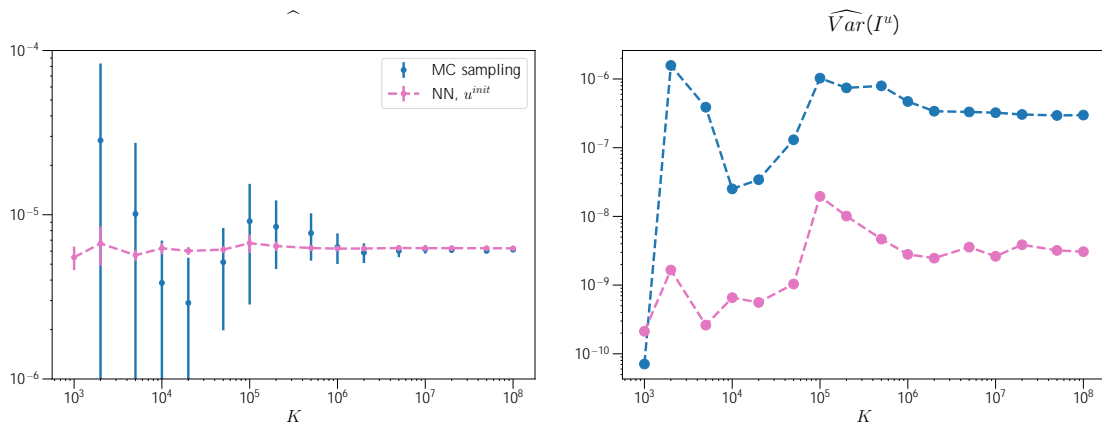


Figure 11: Estimation of $\hat{\mu}$ as well as the corresponding sampling variance when using either naive Monte Carlo or the learnt importance sampling estimating according to Algorithm 3 as a function of different batch sizes K . Note that the Monte Carlo estimator corresponds to a non controlled sampling, i.e. $u = 0$.

4.3 High-dimensional example with metastable features

In our final example we consider a high-dimensional example in $d = 20$, where one dimension is particularly metastable. We set $\alpha_1 = 5$ and $\alpha_j = 0.5$ for $j \geq 2, \dots, 20$. Here, we set the target set to $T = [1, 3] \times [3, 3]^{d-1}$, so that the trajectories stop after overcoming the potential barrier in the metastable coordinate. Here we again do not have a reference solution due to the curse of dimensionality.

As in the previous example we use the cumulative version of the adapted metadynamics algorithm to generate a good control initialization u^{init} . Algorithm 2 is implemented with $\delta = 2$, $K^{\text{meta}} = 100$, $\eta = 1$, $r = 0.95$ and $\beta = 0.5 \text{Id}$. The biasing potential obtained after the metadynamics procedure consists of $M = 295$ unnormalized Gaussian functions. We compare our results with a naive Monte Carlo estimator with a batch

size of $K = 10^8$.

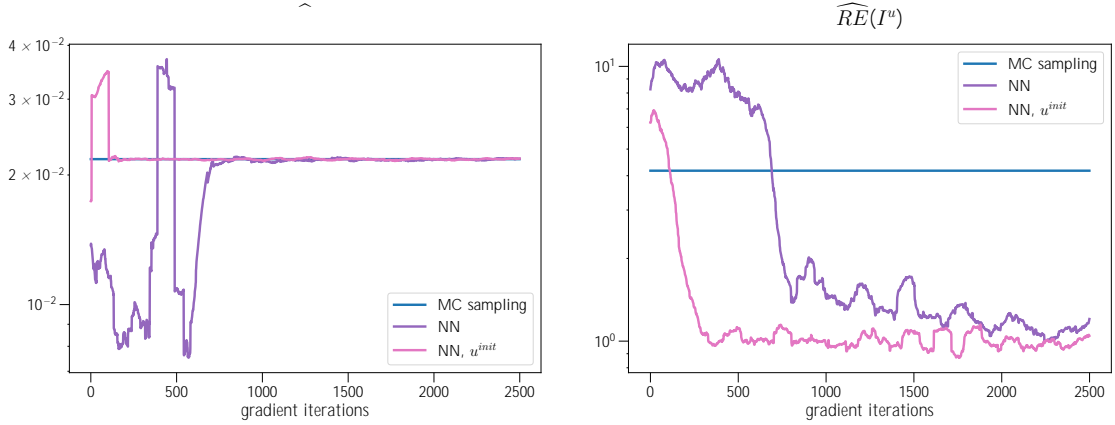


Figure 12: Estimation of $\hat{\mu}$ and importance sampling relative error at each gradient iteration.

In our experiments we can see that the optimization procedure benefits from the metadynamics based initialization. While here both methods eventually find the same minimum, the initialized version converges much faster. In Figure 12 we can observe that the not initialized approximation reaches a similar relative error estimator after 2500 gradient iterations. In Figure 13 we can see the projection of the control in the i -th coordinate as a function of x_i for a fixed value of x_j , $j \in \{2, \dots, 20\}$ after convergence. We compare the metastable direction $i = 1$ with the others, e.g. $i = 2$.

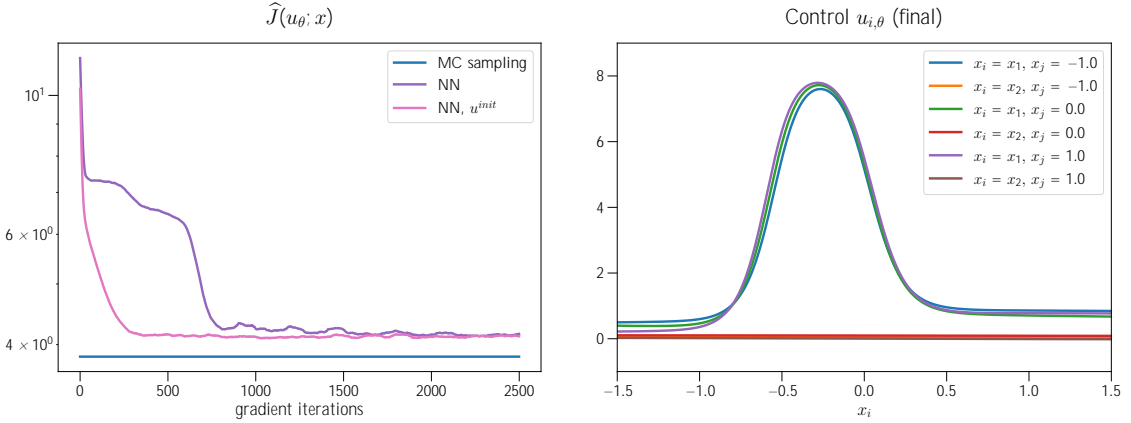


Figure 13: Left panel: minimization of the objective function during the approximation procedure. Right panel: different components of the control function after convergence.

5 Conclusion and outlook

In this paper we have presented a novel method that improves the sampling of metastable diffusions. To be precise, this method is able to cope with two particular challenges, namely 1) the simulation of certain rare events of interest and 2) high variances in the gradient computation of importance sampling optimization algorithms. To overcome those issues, we have suggested to combine optimal control based importance sampling with the metadynamics algorithm. In fact, we are able to combine the best of those two algorithms: offering reasonable initializations, but still striving for systematic variance reductions of Monte Carlo estimators. For the control function approximation we rely on neural networks, which allow for high-dimensional applications and offer further advantages in contrast to a linear combination of ansatz functions, since those would have to be placed explicitly in the domain of interest and would require the tuning of additional hyperparameters. We have further derived a gradient estimator that can easily be computed with automatic differentiation tools and thereby allows for efficient computations. We could demonstrate in multiple numerical experiments that our method improves the convergence of optimizing control based importance sampling significantly. In particular, we note that often our methods work well when alternative strategies do not produce reasonable results anymore. Overall, we are thus able to design low variance importance sampling estimators even in very metastable scenarios.

The stochastic optimization algorithms that we have considered in this work can be understood as some sort of reinforcement learning. For future work, it might therefore be interesting to consider tricks that have been developed in this fruitful field of machine learning in recent years. In particular, it might be promising to alter optimization objectives by for instance adding additional terms that strive for a minimization of the hitting times, while still keeping the variance of the estimators low. Such strategies might become particular relevant in realistic examples in ever larger dimensions with even higher metastabilities. Furthermore, we believe that a connection of our algorithms to model reduction attempts might be fruitful such that (as in the original version of the algorithms) metadynamics based control initialization would only be executed in the relevant coordinates. With that we are optimistic that our proposed method will lead to more efficient sampling of real physical systems.

Acknowledgement

The research of E.R.B and L.R. has been funded by Deutsche Forschungsgemeinschaft (DFG) through grant CRC 1114 “Scaling Cascades in Complex Systems”, A05 (Probing scales in equilibrated systems by optimal nonequilibrium forcing, project number 235221301). The research of J.Q. has been funded by the Einstein Foundation Berlin.

Code availability

The code used for the numerical examples is available on GitHub at www.github.com/riberaborrell/sde-importance-sampling.

A Appendix

A.1 Alternative gradient computations

In this section we present an alternative way of computing the gradient of the control functional (16), now relying on a discrete version of the controlled stochastic process. This gradient estimator has already been suggested in [18], however, we generalize it to more general function classes. The strategy is to consider the controlled process in discrete time, first with deterministic time horizons. Then, by relying on transition probabilities of the discretized process, we can compute the gradient of the discrete cost functional. Eventually, we can change to random stopping times, which yields a gradient that can be estimated by Monte Carlo.

Let us start by stating the discrete version of our controlled process (9) on a time grid $0 < t_1 < \dots < t_N$, for a fixed N , namely

$$\hat{X}_{n+1}^u = \hat{X}_n^u + \left(r V(\hat{X}_n^u) + \sigma(\hat{X}_n^u)u(\hat{X}_n^u) \right) \Delta t + \sigma(\hat{X}_n^u)\xi_{n+1} \sqrt{\Delta t}, \quad (42)$$

where $t = t_{n+1} - t_n$ is the time increment and $\xi_n \sim \mathcal{N}(0, \text{Id})$ are standard normally distributed random variables. Moreover, the discrete version of the control functional (16) reads

$$\widehat{J}(u; x) := \mathbb{E}^\times[h(\widehat{X}^u)], \quad h(\widehat{X}^u) := \sum_{n=0}^{N-1} \left(f(\widehat{X}_n^u) + \frac{1}{2}ju(\widehat{X}_n^u)^2 \right) t + g(\widehat{X}_N^u). \quad (43)$$

The process \widehat{X}^u is a discrete Markov process and by using the Chapman-Kolmogorov equation (see e.g. Section 2.2 in [36]) we can express its joint probability density function conditional on $\widehat{X}_0^u = x$ in terms of the transition densities $G: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_0$

$$\rho(\widehat{X}^u) = \prod_{n=0}^{N-1} G(\widehat{X}_{n+1}^u | \widehat{X}_n^u). \quad (44)$$

From (42) we know that for any discrete steps $x, y \geq \mathbb{R}^d$, G is of multivariate normal form, namely

$$G(y|x) = N(y; x + (rV(x) + \sigma(x)u(x))t, \sigma(x)\sigma^\top(x)t). \quad (45)$$

Let us for computational convenience assume $\sigma(x) = \sqrt{2\beta^{-1}} \text{Id}$, then

$$G(y|x) = \frac{1}{(4\pi\beta^{-1}t)^{d/2}} \exp\left(-\frac{\beta}{4}t \left| \frac{y-x}{t} + rV(x) - \sqrt{2\beta^{-1}}u(x) \right|^2\right). \quad (46)$$

By combining the transition densities in the above expression (44), we get

$$\rho(\widehat{X}^u) = \frac{1}{Z} \exp(-S(\widehat{X}^u)), \quad (47)$$

where the so called discrete action and its normalization factor are given by

$$S(\widehat{X}^u) = \frac{\beta}{4}t \sum_{n=0}^{N-1} \left| \frac{\widehat{X}_{n+1}^u - \widehat{X}_n^u}{t} + rV(\widehat{X}_n^u) - \sqrt{2\beta^{-1}}u(\widehat{X}_n^u) \right|^2, \quad Z = (4\pi\beta^{-1}t)^{Nd/2}. \quad (48)$$

With the help of *infinitesimal perturbation analysis* we can now obtain an estimator for the gradient of the above discretized loss function with respect to the parameter vector θ .

Proposition A.1 (Derivative of discrete cost functional). *Consider h as defined in (43). We can compute the derivative of the discrete cost functional with respect to a parameter θ_i as*

$$\frac{\partial}{\partial \theta_i} \widehat{J}(u; x) = \mathbb{E}^\times \left[\frac{\partial}{\partial \theta_i} h(\widehat{X}^u) - h(\widehat{X}^u) \frac{\partial}{\partial \theta_i} S(\widehat{X}^u) \right], \quad (49a)$$

$$= \mathbb{E}^\times \left[\sum_{n=0}^{N-1} u(\widehat{X}_n^u) \frac{\partial}{\partial \theta_i} u(\widehat{X}_n^u) t + \left(\sum_{n=0}^{N-1} \left(f(\widehat{X}_n^u) + \frac{1}{2}ju(\widehat{X}_n^u)^2 \right) t + g(\widehat{X}_N^u) \right) \sum_{n=0}^{N-1} \xi_{k+1} \frac{\partial}{\partial \theta_i} u(\widehat{X}_n^u)^{\rho} t \right] \quad (49b)$$

where the expectation is to be understood with respect to the discrete path measure with density ρ as defined in (47) and the discrete action S is defined in (48).

Proof. See Appendix A.3. □

As mentioned before, Proposition A.1 holds for the case of a fixed time horizon. Let us now replace fixed times by random stopping times of the controlled dynamics, namely $\tau^u := \inf\{s > 0 | X_s^u \geq T\}$. Note that this stopping time now depends on the control u (and therefore on the parameter θ) and one could be tempted to incorporate this dependency in the gradient computations. However, we have seen in Corollary 3.3 that the derived gradient is in fact exact.

A.2 Girsanov's theorem

Girsanov's theorem [34] provides a formula for changes of measures in path space, which are relevant for our importance sampling computations. Let us therefore provide a brief summary of the theorem.

Let $\tilde{\omega} = C([0, T], \mathbb{R}^d)$ be the space of continuous paths equipped with the supremum norm and let $\mathcal{F} = \mathcal{B}(\tilde{\omega})$ denote the corresponding σ -algebra. First, we define $(M_t^u)_{0 \leq t \leq T}$ by

$$M_t^u := \exp \left(\int_0^t u(X_s^u) \, dW_s - \frac{1}{2} \int_0^t |ju(X_s^u)|^2 \, ds \right), \quad M_0^u := 0, \quad (50)$$

where X^u is the controlled process following (9) and $(\tilde{W}_t)_{0 \leq t \leq T}$ is a Brownian Motion with additional drift, determined for all $t \geq [0, T]$ by

$$\tilde{W}_t := W_t + \int_0^t u(X_s^u) \, ds.$$

If $(M_t^u)_{0 \leq t \leq T}$ is a martingale w.r.t. the canonical filtration of the Brownian motion $(W_t)_{0 \leq t \leq T}$ then the Girsanov Theorem [34, Thm 8.6.8] states that there exists a probability measure \mathbb{Q} absolutely continuous w.r.t. the original probability measure \mathbb{P} characterized by $M_T^u = \frac{d\mathbb{Q}}{d\mathbb{P}}$, i.e. for all $A \in \mathcal{F}$

$$\mathbb{Q}(A) = \mathbb{E}_{\mathbb{Q}}[\mathbb{1}_A] = \mathbb{E}_{\mathbb{P}}[M_T^u \mathbb{1}_A],$$

such that $(\tilde{W}_t)_{0 \leq t \leq T}$ is a Brownian motion with respect to \mathbb{Q} and (X_t^u, \tilde{W}_t) is a weak solution of (1), i.e.

$$\mathbb{Q}\text{-law of } (X_t^u)_{0 \leq t \leq T} = \mathbb{P}\text{-law of } (X_t)_{0 \leq t \leq T}.$$

Notice that for applying Girsanov's theorem, one has to assume that the process (50) is a martingale. Novikov's condition provides us with a sufficient requirement for stochastic processes of the form (50) to be a martingale, see [23]. Namely, it suffices that for all $t \geq [0, T]$

$$\mathbb{E}^x \left[\exp \left(\frac{1}{2} \int_0^t |ju(X_s^u)|^2 \, ds \right) \right] < 1.$$

Girsanov's theorem can be extended to bounded stopping times (see [37, Prop 1]). If the stopping time is bounded the fulfillment of Novikov's condition has already been discussed in [29], [37]. In this case, it holds that

$$\mathbb{E}_{\tilde{\mathbb{P}}}^x[\exp(-W(X))] = \mathbb{E}_{\mathbb{Q}}^x[\exp(-W(X^u))] = \mathbb{E}_{\tilde{\mathbb{P}}}^x[\exp(-W(X^u))M_u^u], \quad (51)$$

where \mathbb{Q} is the Wiener measure of the Brownian motion with drift, \tilde{W} . This implies that the random variable $I^u : \tilde{\omega} \rightarrow \mathbb{R}$ given by

$$I^u = \exp(-W(X^u))M_u^u \quad (52)$$

$$= \exp \left(\int_0^u f(X_s^u) \, ds - g(X_u^u) - \int_0^u u(X_s^u) \, dW_s - \frac{1}{2} \int_0^u |ju(X_s^u)|^2 \, ds \right) \quad (53)$$

is equivalent to our quantity of interest I , as defined in (2). We call the quantity I^u the (re-weighted) importance sampling quantity of interest.

A.3 Proofs

Proof of Proposition 3.2. The proof is adapted from [33] and we refer to a similar computation in [11] and to further technical details in [30].

For $\varepsilon \geq \mathbb{R}$ and $\phi \in C_b^1(\mathbb{R}^d, \mathbb{R}^d)$, let us define the change of measure

$$(\varepsilon, \phi) = \exp \left(\varepsilon \int_0^u \phi(X_s^u) \, dW_s - \frac{\varepsilon^2}{2} \int_0^u |j\phi(X_s^u)|^2 \, ds \right), \quad \frac{d\mathbb{Q}}{d\mathbb{P}} = (\varepsilon, \phi). \quad (54)$$

According to Girsanov's theorem, the process $(\tilde{W}_t)_{0 \leq t \leq T}$, defined as

$$\tilde{W}_t = W_t + \varepsilon \int_0^t \phi(X_s^u) \, ds, \quad (55)$$

is a Brownian motion under \mathbb{Q} . We therefore obtain

$$J(u + \varepsilon\phi; x) = \mathbb{E}^x \left[\left(\frac{1}{2} \int_0^{u+\varepsilon\phi} j(u + \varepsilon\phi)(X_s^{u+\varepsilon\phi})^2 ds + \int_0^{u+\varepsilon\phi} f(X_s^{u+\varepsilon\phi}) ds + g(X_{u+\varepsilon\phi}^{u+\varepsilon\phi}) \right) \right] \quad (56a)$$

$$= \mathbb{E}^x \left[\left(\frac{1}{2} \int_0^u j(u + \varepsilon\phi)(X_s^u)^2 ds + \int_0^u f(X_s^u) ds + g(X_u^u) \right) \mathbb{1}(\varepsilon, \phi) \right]. \quad (56b)$$

Using dominated convergence, we can interchange derivatives and integrals (for technical details, we refer to [30]) and compute

$$\frac{d}{d\varepsilon} \Big|_{\varepsilon=0} J(u + \varepsilon\phi; x) = \mathbb{E}^x \left[\int_0^u (u - \phi)(X_s^u) ds + \left(\frac{1}{2} \int_0^u j u (X_s^u)^2 ds + \int_0^u f(X_s^u) ds + g(X_u^u) \right) \int_0^u \phi(X_s^u) dW_s \right]. \quad (57)$$

□

Proof of Proposition A.1. We follow essentially the arguments of [18] without restricting the choice of the space of possible controls to linear combination of vector fields related to Gaussian ansatz functions. First, let us compute the partial derivatives of the discrete control functional (43) and the discrete action (48), namely

$$\frac{\partial}{\partial \theta_i} h(\hat{X}^u) = t \sum_{n=0}^{N-1} u(\hat{X}_n^u) \frac{\partial}{\partial \theta_i} u(\hat{X}_n^u) \quad (58)$$

and

$$\frac{\partial}{\partial \theta_i} S(\hat{X}^u) = \frac{\rho - \bar{\beta}}{2} t \sum_{n=0}^{N-1} \left(\frac{\hat{X}_{n+1}^u - \hat{X}_n^u}{t} + rV(\hat{X}_n^u) - \sqrt{2\beta^{-1}} u(\hat{X}_n^u) \right) \frac{\partial}{\partial \theta_i} u(\hat{X}_n^u) \quad (59a)$$

$$= \frac{\rho - \bar{\beta}}{t} \sum_{n=0}^{N-1} \xi_{n+1} \frac{\partial}{\partial \theta_i} u(\hat{X}_n^u). \quad (59b)$$

Let us write the expectation of the loss function as an integral over the state space

$$\hat{J}(u; x) = \mathbb{E}^x[h(\hat{X}^u)] = \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} h(\rho) dx_1 \dots dx_N.$$

Then, the partial derivative of the loss function with respect to θ_i can be computed like

$$\frac{\partial}{\partial \theta_i} \hat{J}(u; x) = \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \left(\left(\frac{\partial}{\partial \theta_i} h \right) \rho + h \left(\frac{\partial}{\partial \theta_i} \rho \right) \right) dx_1 \dots dx_N. \quad (60)$$

By using the fact that the normalization factor does not depend on the parameters the partial derivative of the probability density function with respect to θ_i simplifies to

$$\frac{\partial}{\partial \theta_i} \rho = \frac{\partial}{\partial \theta_i} \left(\frac{1}{Z} \exp(-S) \right) = \frac{1}{Z} \frac{\partial}{\partial \theta_i} \exp(-S) = - \left(\frac{\partial}{\partial \theta_i} S \right) \rho,$$

and the expression (60) finally reads

$$\begin{aligned} \frac{\partial}{\partial \theta_i} \hat{J}(u; x) &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \left(\left(\frac{\partial}{\partial \theta_i} h - h \left(\frac{\partial}{\partial \theta_i} S \right) \right) \rho \right) dx_1 \dots dx_N \\ &= \mathbb{E}^x \left[\frac{\partial}{\partial \theta_i} h - h \left(\frac{\partial}{\partial \theta_i} S \right) \right]. \end{aligned}$$

□

References

- [1] F. Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- [2] N. Berglund. Kramers’ law: validity, derivations and generalisations. *Markov Processes and Related fields*, 19(3):459–490, 2013.
- [3] G. Bussi, A. Laio, and M. Parrinello. Equilibrium free energies from nonequilibrium metadynamics. *Phys. Rev. Lett.*, 96:090601, 9, 2006. doi: 10.1103/PhysRevLett.96.090601. url: <http://link.aps.org/doi/10.1103/PhysRevLett.96.090601>.
- [4] J. Comer, J. C. Gumbart, J. Hénin, T. Lelièvre, A. Pohorille, and C. Chipot. The adaptive biasing force method: everything you always wanted to know but were afraid to ask. *The Journal of Physical Chemistry B*, 119(3):1129–1151, 2015. doi: 10.1021/jp506633n. eprint: <https://doi.org/10.1021/jp506633n>. url: <https://doi.org/10.1021/jp506633n>. PMID: 25247823.
- [5] P. Dupuis and H. Wang. Importance sampling, large deviations, and differential games. *Stochastics and Stochastic Reports*, 76(6):481–508, 2004.
- [6] P. Dupuis and H. Wang. Subolutions of an Isaacs equation and efficient schemes for importance sampling. *Mathematics of Operations Research*, 32(3):723–757, 2007.
- [7] P. Dupuis, K. Spiliopoulos, and X. Zhou. Escaping from an attractor: Importance sampling and rest points I. *The Annals of Applied Probability*, 25(5):2909–2958, 2015. doi: 10.1214/14-AAP1064. url: <https://doi.org/10.1214/14-AAP1064>.
- [8] P. Dupuis, H. Wang, and K. Spiliopoulos. Importance sampling for multiscale diffusions. *SIAM Multiscale Modeling and Simulation*, 12:1–27, 2012.
- [9] L. C. Evans. *Partial differential equations*. American Mathematical Society, Providence, R.I., 2010. isbn: 9780821849743 0821849743.
- [10] K. Fackeldey, M. Oster, L. Sallandt, and R. Schneider. Approximative policy iteration for exit time feedback control problems driven by stochastic differential equations using tensor train format. *arXiv preprint arXiv:2010.04465*, 2020.
- [11] E. Fournié, J.-M. Lasry, J. Lebuchoux, P.-L. Lions, and N. Touzi. Applications of Malliavin calculus to Monte Carlo methods in finance. *Finance and Stochastics*, 3(4):391–412, 1999.
- [12] R. Galvelis and Y. Sugita. Neural network and nearest neighbor algorithms for enhancing sampling of molecular dynamics. *Journal of Chemical Theory and Computation*, 13(6):2489–2500, 2017. doi: 10.1021/acs.jctc.7b00188. eprint: <https://doi.org/10.1021/acs.jctc.7b00188>. url: <https://doi.org/10.1021/acs.jctc.7b00188>. PMID: 28437616.
- [13] J. Han, A. Jentzen, et al. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
- [14] C. Hartmann, R. Banisch, M. Sarich, T. Badowski, and C. Schütte. Characterization of rare events in molecular dynamics. *Entropy*, 16(1):350–376, 2014.
- [15] C. Hartmann, O. Kebiri, L. Neureither, and L. Richter. Variational approach to rare event simulation using least-squares regression. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(6):063107, 2019.
- [16] C. Hartmann and L. Richter. Nonasymptotic bounds for suboptimal importance sampling. *arXiv preprint arXiv:2102.09606*, 2021.
- [17] C. Hartmann, L. Richter, C. Schütte, and W. Zhang. Variational characterization of free energy: theory and algorithms. *Entropy*, 19:626, Nov. 2017. doi: 10.3390/e19110626.
- [18] C. Hartmann and C. Schütte. Efficient rare event simulation by optimal nonequilibrium forcing. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(11):P11004, Nov. 2012. doi: 10.1088/1742-5468/2012/11/p11004.
- [19] C. Hartmann, C. Schütte, M. Weber, and W. Zhang. Importance sampling in path space for diffusion processes with slow-fast variables. *Probability Theory and Related Fields*, 170:177–228, 2018. doi: 10.1007/s00440-017-0755-3.
- [20] C. Hartmann, C. Schütte, and W. Zhang. Model reduction algorithms for optimal control and importance sampling of diffusions, 2016. doi: 10.1088/0951-7715/29/8/2298. url: <https://doi.org/10.1088>.

- [21] J. Hénin and C Chipot. Overcoming free energy barriers using unconstrained molecular dynamics simulations. *The Journal of Chemical Physics*, 121(7):2904–2914, 2004. doi: 10.1063/1.1773132.
- [22] D. J. Higham. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review*, 43(3):525–546, 2001. doi: 10.1137/S0036144500378302. eprint: <https://doi.org/10.1137/S0036144500378302>. url: <https://doi.org/10.1137/S0036144500378302>.
- [23] N. Ikeda and S. Watanabe. *Stochastic differential equations and diffusion processes*, volume 24 of *North-Holland Mathematical Library*. North-Holland Publishing Co., Amsterdam, second edition, 1989, pages xvi+555. isbn: 0-444-87378-3.
- [24] A. Jentzen, D. Salimova, and T. Welti. A proof that deep artificial neural networks overcome the curse of dimensionality in the numerical approximation of Kolmogorov partial differential equations with constant diffusion and nonlinear drift coefficients. *arXiv preprint arXiv:1809.07321*, 2018.
- [25] B. Jourdain, T. Lelièvre, and P.-A. Zitt. Convergence of metadynamics: Discussion of the adiabatic hypothesis. *The Annals of Applied Probability*, 31(5):2441–2477, 2021. doi: 10.1214/20-AAP1652. url: <https://doi.org/10.1214/20-AAP1652>.
- [26] D. P. Kingma and J. Ba. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [27] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman. The weighted histogram analysis method for free-energy calculations on biomolecules. I. the method. *Journal of Computational Chemistry*, 13(8), 1992.
- [28] A. Laio and M. Parrinello. Escaping free-energy minima. *PNAS*, 20(10):12562–12566, 2002.
- [29] T. Lelièvre and G. Stoltz. Partial differential equations and stochastic methods in molecular dynamics. *Acta Numerica*, 25:681–880, 2016. doi: 10.1017/S0962492916000039.
- [30] H. C. Lie. *On a strongly convex approximation of a stochastic optimal control problem for importance sampling of metastable diffusions*. PhD thesis, 2016. url: <http://dx.doi.org/10.17169/refubium-8010>.
- [31] G. Milstein. *Numerical integration of stochastic differential equations*. In Springer Berlin Heidelberg, Berlin, Heidelberg, 1995. isbn: 978-94-015-8455-5. doi: <https://doi.org/10.1007/978-94-015-8455-5>.
- [32] N. Nüsken and L. Richter. Interpolating between BSDEs and PINNs—deep learning for elliptic and parabolic boundary value problems. *arXiv preprint arXiv:2112.03749*, 2021.
- [33] N. Nüsken and L. Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial Differential Equations and Applications*, 2(4):1–48, 2021.
- [34] B. Øksendal. *Stochastic differential equations*. In *Stochastic Differential Equations: An Introduction with Applications*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003, pages 65–84. isbn: 978-3-642-14394-6. doi: 10.1007/978-3-642-14394-6_5. url: https://doi.org/10.1007/978-3-642-14394-6_5.
- [35] A. B. Owen. *Monte Carlo theory, methods and examples*. Self-published, 2013.
- [36] G. A. Pavliotis. *Stochastic processes and applications*, volume 60 of *Texts in Applied Mathematics*. Springer, New York, 2014, pages xiv+339. isbn: 978-1-4939-1322-0; 978-1-4939-1323-7. doi: 10.1007/978-1-4939-1323-7. url: <https://doi.org/10.1007/978-1-4939-1323-7>. Diffusion processes, the Fokker-Planck and Langevin equations.
- [37] J. Quer, L. Donati, B. Keller, and M. Weber. An automatic adaptive importance sampling algorithm for molecular dynamics in reaction coordinates. *SIAM Journal on Scientific Computing*, 40:A653–A670, Jan. 2018. doi: 10.1137/17M1124772.
- [38] L. Richter. *Solving high-dimensional PDEs, approximation of path space measures and importance sampling of diffusions*. PhD thesis, BTU Cottbus-Senftenberg, 2021.
- [39] L. Richter, L. Sallandt, and N. Nüsken. Solving high-dimensional parabolic PDEs using the tensor train format. In *International Conference on Machine Learning*, pages 8998–9009. PMLR, 2021.
- [40] K. Spiliopoulos. Non-asymptotic performance analysis of importance sampling schemes for small noise diffusions. *Journal of Applied Probability*, 52:1–14, 2015.
- [41] O. Valsson, P. Tiwary, and M. Parrinello. Enhancing important fluctuations: rare events and metadynamics from a conceptual viewpoint. *Annual Review of Physical Chemistry*, 67(1):159–184, 2016. doi: 10.1146/annurev-physchem-040215-112229. eprint: <https://doi.org/10.1146/annurev-physchem-040215-112229>. url: <https://doi.org/10.1146/annurev-physchem-040215-112229>. PMID: 26980304.

- [42] E. Vanden-Eijden and J. Weare. Rare event simulation of small noise diffusions. *Communications on Pure and Applied Mathematics*, 65:1770–1803, 2012.
- [43] F. Wang and D. P. Landau. Determining the density of states for classical statistical models: a random walk algorithm to produce a flat histogram. *Phys. Rev. E*, 64:056101, 5, Oct. 2001. doi: 10.1103/PhysRevE.64.056101. url: <http://link.aps.org/doi/10.1103/PhysRevE.64.056101>.
- [44] E. Weinan, J. Han, and A. Jentzen. Algorithms for solving high dimensional PDEs: from nonlinear Monte Carlo to machine learning. *Nonlinearity*, 35(1):278, 2021.
- [45] W. Zhang, C. Hartmann, and C. Schütte. Effective dynamics along given reaction coordinates, and reaction rate theory. *Faraday Discuss.*, 2016.
- [46] M. Zhou, J. Han, and J. Lu. Actor-critic method for high dimensional static Hamilton–Jacobi–Bellman partial differential equations based on neural networks. *SIAM Journal on Scientific Computing*, 43(6):A4043–A4066, 2021.
- [47] R. Zwanzig. High temperature equation of state by a perturbation method. I. nonpolar gases. *The Journal of Chemical Physics*, 22(8):1420–1426, 1954.