

APPLICATIONS OF THE CROSS-ENTROPY METHOD TO IMPORTANCE SAMPLING AND OPTIMAL CONTROL OF DIFFUSIONS

WEI ZHANG², HAN WANG¹, CARSTEN HARTMANN^{1,*}, MARCUS WEBER², AND CHRISTOF SCHÜTTE^{1,2}

Abstract. We study the cross-entropy method for diffusions. One of the results is a versatile cross-entropy algorithm that can be used to design efficient importance sampling strategies for rare events or to solve optimal control problems. The approach is based on the minimization of a suitable cross-entropy functional, with a parametric family of exponentially tilted probability distributions. We illustrate the new algorithm with several numerical examples and discuss algorithmic issues and possible extensions of the method.

Key words. importance sampling, optimal control, cross-entropy method, rare events, change of measure.

AMS subject classifications.

1 Introduction This article deals with the application of the cross-entropy method to diffusion processes, specifically, with the application to importance sampling for rare events and optimal control. Generally, the cross-entropy method is a Monte-Carlo method that was originally developed for the efficient simulation of rare events in queuing models and that has been extended to, e.g., combinatorial optimization or analysis of networks in the meantime [18, 4]. To our knowledge, however, the cross-entropy approach has not been analyzed or used in combination with diffusion processes, even though there have been significant research activities in the direction of efficient algorithms for importance sampling and optimal control of high dimensional multiscale diffusions; see, e.g., [5, 23, 24] for some ideas related to importance sampling of rare events or [20, 25] for problems in optimal control.

We will exploit the fundamental duality between importance sampling and optimal control, which arises due to the fact that both problems admit a variational formulation that boils down to finding an optimal transformation of the underlying path space measure [6]. Algorithms for computing an optimal change of measure usually seek an approximation of the optimal measure with respect to some distance on the space of (probability) measures. Here, we will use the Kullback-Leibler divergence, which, although not a metric, is a numerically convenient and widely used similarity measure for probability measures. The cross-entropy method then provides a general algorithm to find the minimizer of the Kullback-Leibler divergence among a family of parameterized probability measures, and the main purpose of this paper is to formulate the method in the context of diffusions and to discuss its application to importance sampling and optimal control.

This paper is organized as follows. The cross-entropy method in path space is outlined in Section 2 and discussed in the context of importance sampling and the dual optimal control problem. Section 3 is devoted to the formulation of the cross-entropy algorithm for diffusions. Several numerical examples are studied in Section 4. We summarize our findings in Section 5.

¹Institute of Mathematics, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany

²Zuse Institute Berlin, Takustrasse 7, 14195 Berlin, Germany

*Corresponding author. Email: chartman@mi.fu-berlin.de

2 The cross-entropy method in path space In this section, we discuss how to use the cross-entropy method for stochastic differential equations. The mathematical set-up will be largely based on the application of the method to importance sampling, which is the most standard application of the cross-entropy method in the literature. The associated (dual) optimal control problem will be briefly discussed at the end of this section.

2.1 Problem set-up Consider $z_s \in \mathbb{R}^n$ satisfying

$$\begin{aligned} dz_s &= b(z_s)ds + \sqrt{2\epsilon} dw_s, \quad 0 \leq s \leq T \\ z_0 &= x \end{aligned} \tag{2.1}$$

where $\epsilon > 0$ is constant, $b: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a smooth vector field, and w is n -dimensional Brownian motion. Further let $O \subset \mathbb{R}^n$ be open and bounded and call

$$\tau = \inf\{s > 0: z_s \notin O\} \tag{2.2}$$

the first exit time of the set $O \subset \mathbb{R}^n$. In the following we will use Z to denote a path (trajectory) $\{z_s: 0 \leq s \leq T\}$ and use the notation $z_s \in \mathbb{R}^n$ for the state at time s . Accordingly we denote by $F(Z)$ a path functional that, throughout this paper, is assumed to be of the form

$$F(Z) = \exp\left(-\frac{1}{\epsilon} \int_0^{\tau \wedge T} G(z_s) ds - \frac{1}{\epsilon} H(z_{\tau \wedge T})\right). \tag{2.3}$$

for some continuous and bounded functions $G, H: \mathbb{R}^n \rightarrow \mathbb{R}$ and with $\tau \wedge T = \min\{\tau, T\}$. We consider a Monte Carlo method to compute the quantity

$$\ell(x) = \mathbf{E}(F(Z)), \tag{2.4}$$

with $\mathbf{E}(\cdot) = \mathbf{E}(\cdot | z_0 = x)$ denoting the conditional expectation over all realizations of (2.1) starting at $z_0 = x$. A special and interesting case is when $G = 0$ and $H = -\epsilon \log \mathbf{1}_{\partial O}$ with ∂O denoting the smooth boundary of the set O , in which case

$$\mathbf{E}(\mathbf{1}_{\partial O}(z_{\tau \wedge T})) = \mathbf{P}(\tau \leq T), \tag{2.5}$$

is the probability of trajectories starting at $z_0 = x$ to reach the boundary of O before time T . The exit time distribution of a set provides details about, e.g., transition mechanisms and is a common quantity to analyze metastable dynamics.

2.2 Importance sampling We now formulate the cross-entropy method for diffusions. In doing so, we follow the relevant literature [4, 18] and first introduce the general concept of importance sampling in path space, before we establish the link with optimal control. Suppose we are able to generate path samples from a family of probability measures $\{\mu_\lambda\}_{\lambda \in \mathcal{F}}$ on the space of continuous functions $C([0, T], \mathbb{R}^n)$ that are parametrized by $\lambda \in \mathcal{F} \subset \mathbb{R}^m$ where the dynamics (2.1) corresponds to $\lambda = \mathbf{0}$; for the sake of simplicity, we set $\mathcal{F} = \mathbb{R}^m$. We use the shorthand $\mu = \mu_{\mathbf{0}}$ and refer to $\mu_{\lambda \neq \mathbf{0}}$ as the *tilted probability measure*. We further assume that every μ_λ has a probability density $f(\cdot; \lambda)$ with respect to the scaled Wiener measure ν_ϵ .*

*The scaled measure ν_ϵ is the probability measure induced by the scaled Brownian motion $\sqrt{2\epsilon}w_s$ on the space $C([0, T], \mathbb{R}^n)$. It is related to the standard Wiener measure underlying the standard Brownian motion w_s by rescaling of time, which follows from the fact that w_s and $\sqrt{2\epsilon}w_{s/(2\epsilon)}$ have the same law.

The idea of importance sampling is—instead of drawing samples from the measure μ —to generate samples from an alternative probability measure $\eta = g(\cdot)\nu_\epsilon$ that is absolutely continuous with respect to μ , but that yields Monte-Carlo estimators that have, e.g., smaller variance or bounded relative error as the probability of the rare goes to zero [12]. Using independent draws from η an unbiased Monte Carlo estimator of (2.4) is given by

$$\ell_N = \frac{1}{N} \sum_{i=1}^N \frac{F(\tilde{Z}_i) f(\tilde{Z}_i; \mathbf{0})}{g(\tilde{Z}_i)}, \quad (2.6)$$

where the trajectories $\tilde{Z}_i, i = 1, \dots, N$ are independent realizations from η . It is well known that the optimal measure η^* that minimizes the variance of the estimator has the density

$$g^*(Z) = \frac{F(Z) f(Z; \mathbf{0})}{\ell}. \quad (2.7)$$

It is easy to see that the thus defined η^* yields a zero variance estimator. Note however that it depends on $\ell = \mathbf{E}(F(Z))$, which is the quantity that we want to compute.

The idea of the cross-entropy method is to find the best approximation of η^* among the family $\mu_\lambda, \lambda \in \mathcal{F}$ of tilted probability measures. The approach is based on minimizing the Kullback-Leibler divergence, which in our case can be defined as follows: given two probability measures $\mu_1 = g_1\nu_\epsilon, \mu_2 = g_2\nu_\epsilon$ that are absolutely continuous with respect to the scaled Wiener measure, the *Kullback-Leibler divergence* or *relative entropy* between μ_1 and μ_2 is defined as

$$D(\mu_1, \mu_2) = \mathbf{E}_{\mu_1} \left(\log \frac{d\mu_1}{d\mu_2} \right) \quad (2.8)$$

where the expectation with respect to the measure μ_1 is defined as

$$\mathbf{E}_{\mu_1} \left(\log \frac{d\mu_1}{d\mu_2} \right) = \int \log \frac{d\mu_1}{d\mu_2} d\mu_1 = \int g_1 \log g_1 d\nu_\epsilon - \int g_1 \log g_2 d\nu_\epsilon. \quad (2.9)$$

Cross-entropy method I. The cross entropy method now seeks an optimal change of measure by solving the minimization task

$$\min_{\lambda \in \mathbb{R}^m} D(\eta^*, \mu_\lambda) \quad (2.10)$$

for the tilt parameter $\lambda \in \mathbb{R}^m$. Not knowing what η^* is, this still sounds like an infeasible minimization problem. It turns out, however, that we need to know η^* only up to a constant prefactor, which in our case, since x is fixed, eliminates the unknown quantity $\ell = \ell(x)$. Using (2.7) and (2.8), the minimization problem is equivalent to the following maximization problem

$$\max_{\lambda \in \mathbb{R}^m} \mathbf{E}_\mu (F(Z) \log f(Z; \lambda)). \quad (2.11)$$

For the efficient numerical solution of (2.11) it is often convenient to allow for drawing samples from a probability measure that somehow “in between” μ and μ_{λ^*} . This will give us some extra freedom to use, e.g., an iterative solver for the maximization problem (2.11). Letting $\mathbf{v} \in \mathbb{R}^m$ denote an arbitrary family parameter, our maximization problem can be recast as

$$\max_{\lambda \in \mathbb{R}^m} \mathbf{E}_{\mu_{\mathbf{v}}} (F(Z) h(Z; \mathbf{v}) \log f(Z; \lambda)), \quad (2.12)$$

where

$$h(Z; \mathbf{v}) = \frac{f(Z; \mathbf{0})}{f(Z; \mathbf{v})}. \quad (2.13)$$

An unbiased estimator of (2.12) is

$$\max_{\boldsymbol{\lambda} \in \mathbb{R}^m} \frac{1}{N} \sum_{i=1}^N F(\tilde{Z}_i) h(\tilde{Z}_i; \mathbf{v}) \log f(\tilde{Z}_i; \boldsymbol{\lambda}) \quad (2.14)$$

where $\tilde{Z}_i, i = 1, \dots, N$ are independent realizations generated from $\mu_{\mathbf{v}}$. The necessary condition for $\boldsymbol{\lambda}^*$ being a maximizer of (2.14) is obtained by taking the gradient with respect to $\boldsymbol{\lambda}$:

$$\sum_{i=1}^N F(\tilde{Z}_i) h(\tilde{Z}_i; \mathbf{v}) \nabla_{\boldsymbol{\lambda}} \log f(\tilde{Z}_i; \boldsymbol{\lambda}) = 0. \quad (2.15)$$

The degree of difficulty when solving (2.15) numerically of course depends on the parameterization of the tilted family of distributions. In Section 3, we will introduce a family of tilted distributions that turns (2.15) into a linear system of equations for the unknown $\boldsymbol{\lambda}$.

2.3 Optimal control In this section, we consider the cross-entropy method for the solution of certain optimal control problems for diffusion processes. To this end consider the optimal control problem with cost function [20, 25]

$$J(u) = \mathbf{E}_{\mu_u} \left(\int_0^{\tau \wedge T} G(z_s) + \frac{1}{4} |u_s|^2 ds + H(z_{\tau \wedge T}) \right), \quad (2.16)$$

with bounded continuous function $G, H: \mathbb{R}^n \rightarrow \mathbb{R}$ and $u_s \in \mathbb{R}^n$ being a measurable control that will be specified below. The expectation \mathbf{E}_{μ_u} with respect to the probability measure μ_u is understood as the expectation over all realizations of the controlled dynamics

$$\begin{aligned} dz_s &= (b(z_s) + u_s) ds + \sqrt{2\epsilon} dw_s, \quad 0 \leq s \leq T \\ z_0 &= x \end{aligned} \quad (2.17)$$

We suppose that $G \geq 0$ and, for the ease of notation, we set $H = 0$. We wish to minimize (2.16) under the constraint (2.17) and over all measurable controls u that are adapted to the filtration generated by the Brownian motion driving the dynamics (2.17).

Control problems of the above form are called *linear-quadratic* as the control appears linearly in the equations of motion and quadratically in the cost functional; the dependence on the states z may be nonlinear though. An example is the alignment of a molecule in solution in a laser field (modelled by u), where the cost measures the deviation of the molecule from a given reference configuration and the energy exerted by the laser [21]. Other applications of stochastic control problems of the form (2.16)–(2.17) involve molecular dynamics [20], photochemistry [1], material science [22], or mechanical engineering [26], to mention just a few.

Under suitable conditions on the vector field b , it is known that the optimal control problem (2.16)–(2.17) has a unique viscosity solution in form of a Markovian feedback control [9]. Specifically, there exists a continuous and bounded function $c: [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}^n$, such that the minimizer $\hat{u} = \operatorname{argmin} J(u)$ of the cost functional J is of the form

$$\hat{u}_s = c(s, z_s), \quad (2.18)$$

with $c(t, x) = -2\nabla v(t, x)$ and v being the value function or *optimal cost-to-go*:

$$v(t, x) = \min_u \mathbf{E}_{\mu_u} \left(\int_0^{\tau \wedge T} G(z_s) + \frac{1}{4} |u_s|^2 ds + H(z_{\tau \wedge T}) \Big| z_t = x \right). \quad (2.19)$$

Now call $\mu, \hat{\mu}$ the probability measures on the path space $C([0, T], \mathbb{R}^n)$ corresponding to (2.17) with $u = 0$ and \hat{u} . Then, using the Legendre-type dual relation, we have [3, 10]

$$J(\hat{u}) = -\epsilon \log \mathbf{E}_{\mu} \left(\exp \left(-\frac{1}{\epsilon} \int_0^{\tau \wedge T} G(z_s) ds \right) \right), \quad (2.20)$$

where, by Jensen's inequality (see [9, Sec. VI.2]), we know that $\hat{\mu}$ -a.s.

$$\exp \left(-\frac{1}{\epsilon} \int_0^{\tau \wedge T} G(z_s) ds \right) \frac{d\mu}{d\hat{\mu}} = \mathbf{E}_{\mu} \left(\exp \left(-\frac{1}{\epsilon} \int_0^{\tau \wedge T} G(z_s) ds \right) \right) \quad (2.21)$$

From the above we conclude that (see [9, Sec. VI.3] for details)

$$\begin{aligned} J(u) &= \mathbf{E}_{\hat{\mu}} \left(\int_0^{\tau \wedge T} \left(G(z_s) + \frac{1}{4} |u_s|^2 \right) ds \right) \frac{d\mu_u}{d\hat{\mu}} \\ &= J(\hat{u}) + \mathbf{E}_{\hat{\mu}} \left(\left(\epsilon \log \frac{d\mu}{d\hat{\mu}} + \frac{1}{4} \int_0^{\tau \wedge T} |u_s|^2 ds \right) \frac{d\mu_u}{d\hat{\mu}} \right) \end{aligned} \quad (2.22)$$

After some rearrangement and simplification, we obtain the following simple relationship:

$$J(u) = J(\hat{u}) + \epsilon D(\mu_u, \hat{\mu}). \quad (2.23)$$

where $D(\cdot, \cdot)$ is the Kullback-Leibler divergence as defined in (2.8).

Cross-entropy method II. Computing the optimal control \hat{u} can be tedious, or is infeasible if the dynamics are high dimensional. As a remedy we suggest again to find the best approximation of $\hat{\mu} = \hat{g}(\cdot) \nu_\epsilon$ among a suitably defined family $\mu_\lambda, \lambda \in \mathcal{F}$ of tilted probability measures that are absolutely continuous with respect to ν_ϵ . Sticking to the notation from Section 2, it readily follows that the minimizer of (2.23) has the density

$$\hat{g}(Z) \propto \exp \left(-\frac{1}{\epsilon} \int_0^{\tau \wedge T} G(z_s) ds \right) f(Z; 0) \quad (2.24)$$

with respect to ν_ϵ , which is should be compared to the corresponding expression (2.7) for the optimal importance sampling distribution. By Girsanov's theorem there is a one-to-one correspondence between the control force $u = u(\lambda)$ and a certain family of *exponentially tilted* probability measures μ_λ . Instead of minimizing (2.16) subject to the dynamics (2.17), we solve the constrained optimization problem

$$\min_{\lambda \in \mathbb{R}^m} J(u(\lambda)). \quad (2.25)$$

subject to the dynamics (2.17). From (2.23), we know that solving (2.25) is equivalent to minimizing the Kullback-Leibler divergence D between μ_u and $\hat{\mu}$, which, however, is still not easy. On the other hand, inspired by the discussions in Section 2, we can apply cross-entropy method

to minimize the relaxed entropy functional $D(\hat{\mu}, \mu_u)$ rather than $D(\mu_u, \hat{\mu})$. As a consequence, the problem is reduced to the case in Section 2.

Note that the relaxed problem solved by the cross-entropy method is different from (2.25), as the Kullback-Leibler divergence is not symmetric. Yet both (2.25) and its relaxed version agree at the minimum, therefore the hope is that the latter yields a reasonable approximation of the optimal control problem—given that the family of tilted distributions is cleverly chosen.

3 Cross-entropy algorithm In this section, we will specify the family $\{\mu_\lambda\}_{\lambda \in \mathcal{F}}$ of tilted probability measures that we are going to use for the procedure introduced above and formulate the cross-entropy algorithm. As a first step, let μ denote the probability measure on $C([0, T], \mathbb{R}^n)$ that is induced by the dynamics (2.1) and let ν_ϵ be scaled Wiener measure associated with

$$\begin{aligned} dx_s &= \sqrt{2\epsilon} dw_s, \quad 0 \leq s \leq T \\ x_0 &= x. \end{aligned} \tag{3.1}$$

By Girsanov's theorem [15], we have

$$d\mu = \exp\left(-\frac{1}{\epsilon} S(Z)\right) d\nu_\epsilon, \tag{3.2}$$

with the action

$$\begin{aligned} S(Z) &= -\sqrt{\frac{\epsilon}{2}} \int_0^{\tau \wedge T} b(z_s) \cdot dw_s - \frac{1}{4} \int_0^{\tau \wedge T} |b(z_s)|^2 ds \\ &= -\frac{1}{2} \int_0^{\tau \wedge T} b(z_s) \cdot dz_s + \frac{1}{4} \int_0^{\tau \wedge T} |b(z_s)|^2 ds \end{aligned} \tag{3.3}$$

where the stochastic integral is interpreted in the sense of Itô [15]. A remark is in order.

REMARK 1. We could rewrite (3.2)–(3.3) as a Stratonovich integral using the relationship

$$\int_0^{\tau \wedge T} b(z_s) \cdot dz_s = \int_0^{\tau \wedge T} b(z_s) \circ dz_s - \epsilon \int_0^{\tau \wedge T} \operatorname{div} b(z_s) ds, \tag{3.4}$$

by which we obtain

$$d\mu = \exp\left(\frac{1}{2\epsilon} \int_0^{\tau \wedge T} b(z_s) \circ dz_s - \frac{1}{4\epsilon} \int_0^{\tau \wedge T} (|b(z_s)|^2 + 2\epsilon \operatorname{div} b(z_s)) ds\right) d\nu_\epsilon \tag{3.5}$$

The associated action functional

$$S(Z) = -\frac{1}{2} \int_0^{\tau \wedge T} b(z_s) \cdot dz_s + \frac{1}{4} \int_0^{\tau \wedge T} |b(z_s)|^2 ds + \frac{\epsilon}{2} \int_0^{\tau \wedge T} \operatorname{div} b(z_s) ds \tag{3.6}$$

is closely related to what is known as the Onsager-Machlup functional in the physical literature. See [7] for details. We will stick to Itô interpretation of (3.3) in the following.

3.1 Choosing a family of path space measures We will confine our attention to a special class of tilted probability distributions that is suggested by the the optimal control problem from in Section 2.3. To this end, let $\{\phi_i\}_{1 \leq i \leq m}$ denote a set of continuously differentiable basis functions $\phi_i: [0, T] \times \mathbb{R}^n \rightarrow \mathbb{R}$. The cross-entropy method will be based on realizations of

$$\begin{aligned} dz_s &= (b(z_s) + c(s, z_s; \lambda)) ds + \sqrt{2\epsilon} dw_s, \quad 0 \leq s \leq T \\ z_0 &= x, \end{aligned} \tag{3.7}$$

with

$$c(s, x; \boldsymbol{\lambda}) = 2 \sum_{i=1}^m \lambda_i \nabla \phi_i(s, x). \quad (3.8)$$

(The scaling factor 2 is merely conventional.) It follows from Girsanov's theorem that the associated path probability measure $\mu_{\boldsymbol{\lambda}}$ has a density $f(\cdot; \boldsymbol{\lambda})$ with respect to the scaled Wiener measure ν_{ϵ} . It is given by the usual exponential expression

$$f = \exp \left(\frac{1}{2\epsilon} \int_0^{\tau \wedge T} (b(z_s) + c(s, z_s; \boldsymbol{\lambda})) \cdot dz_s - \frac{1}{4\epsilon} \int_0^{\tau \wedge T} |b(z_s) + c(s, z_s; \boldsymbol{\lambda})|^2 ds \right). \quad (3.9)$$

As a consequence, we can generate independent samples from $\mu_{\boldsymbol{\lambda}}$ by repeatedly simulating the controlled dynamics (3.7). Since the tilting parameter $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_m)$ enters linearly, the associated cross-entropy maximization problem (2.14) is strictly convex and thus has a unique solution. Note that, indeed, $\mu_{\mathbf{0}} = \mu$ is the probability measure corresponding (2.1).

3.2 Iterative cross-entropy optimization of control policies

Defining

$$f(Z; \boldsymbol{\lambda}) = \exp \left(-\frac{1}{\epsilon} S(Z; \boldsymbol{\lambda}) \right), \quad (3.10)$$

with the action

$$S = -\frac{1}{2} \int_0^{\tau \wedge T} (b(z_s) + c(s, z_s; \boldsymbol{\lambda})) \cdot dz_s + \frac{1}{4} \int_0^{\tau \wedge T} |b(z_s) + c(s, z_s; \boldsymbol{\lambda})|^2 ds \quad (3.11)$$

and noting that

$$\nabla_{\boldsymbol{\lambda}} \log f(Z; \boldsymbol{\lambda}) = -\frac{1}{\epsilon} \nabla_{\boldsymbol{\lambda}} S(Z; \boldsymbol{\lambda}), \quad (3.12)$$

the necessary optimality condition (2.15) can be recast as a linear system of equations:

$$A\boldsymbol{\lambda} = \mathbf{r}, \quad (3.13)$$

where $A = (A_{ij})_{1 \leq i, j \leq m}$ and $\mathbf{r} = (r_i)_{1 \leq i \leq m}$ with

$$\begin{aligned} A_{ij} &= 2 \sum_{k=1}^N F(Z_k) h(Z_k; \mathbf{v}) \int_0^{\tau \wedge T} \nabla \phi_i(s, z_s^k) \nabla \phi_j(s, z_s^k) ds, \\ r_i &= \sum_{k=1}^N F(Z_k) h(Z_k; \mathbf{v}) \left(\int_0^{\tau \wedge T} \nabla \phi_i(s, z_s^k) \cdot dz_s^k - \int_0^{\tau \wedge T} \nabla \phi_i(s, z_s^k) \cdot b(z_s^k) ds \right). \end{aligned} \quad (3.14)$$

Here $\mathbf{v} \in \mathbb{R}^m$ is an arbitrary vector, and $Z_k = (z_s^k)_{0 \leq s \leq T}$ denotes the sample paths of (3.7) with control $u_s = c(s, z_s^k; \mathbf{v})$. Note that the realizations are generated from $\mu_{\mathbf{v}}$ and therefore do not depend on $\boldsymbol{\lambda}$. Further notice that the matrix A is positive definite if the basis functions ϕ_i are linearly independent, which implies that (3.14) has a unique solution.

It thus seems that the solution of the discrete maximization problem (2.14) can be obtained by just solving the linear equation (3.14). However, in real applications the expectation value ℓ of $F(Z)$ in (2.4) is very small so that it is difficult to estimate the coefficients (3.14) accurately enough, which renders the solution of (2.14) inaccurate. In many applications, the reason for this

is metastability of the dynamics when $\epsilon \ll 1$. In this case, the trajectories are long and ℓ is small, which means computing (3.14) is both time-consuming and inaccurate. Inspired by the original cross-entropy method [4], we may overcome this problem by starting from a higher temperature (here: ϵ) and compute (3.14) while decreasing the temperature. The proposed iterative method to solve (2.14) is illustrated in Algorithm 1.

Algorithm 1 Cross-entropy algorithm

- 1: Define $\epsilon_0 > \epsilon_1 > \dots > \epsilon_k = \epsilon$, set $\mathbf{v}^{(0)} = 0$.
 - 2: **for** $j = 0$ to k **do**
 - 3: generate N_j trajectories $z_i, i = 1, 2, \dots, N$ from dynamics (3.7), with $\boldsymbol{\lambda} = \mathbf{v}^{(j)}$, $\epsilon = \epsilon_j$.
 - 4: compute the coefficients of $A^{(j)}, \mathbf{r}^{(j)}$ from (3.14) with $\mathbf{v} = \mathbf{v}^{(j)}$, and solve the linear equations $A^{(j)}\mathbf{v}^{(j+1)} = \mathbf{r}^{(j)}$.
 - 5: **end for**
-

We conclude this subsection with a few remarks on possible extensions of the method.

REMARK 2. *It is straightforward to relax the restriction of the fixed initial condition and consider distributed initial conditions instead, i.e. $z_0 = x \in \mathbb{R}^n$ following some probability distribution π on \mathbb{R}^n . All considerations and the cross-entropy method remain valid without alterations, if the sum over the N realizations of the dynamics is replaced by the sum over all realization and the sum over sufficiently many independent initial conditions $x \sim \pi$.*

REMARK 3. *We briefly mention two possible generalization of the above algorithm. The first generalization concerns dynamics with multiplicative noise:*

$$dz_s = b(z_s)ds + \sqrt{2\epsilon}\sigma(z_s)dw_s, \quad (3.15)$$

where the $n \times n$ matrices $a(\cdot) = \sigma(\cdot)\sigma(\cdot)^T$ are positive definite with bounded inverses. Defining the weighted scalar product $\langle u, v \rangle = u^T(a(z))^{-1}v$, then all considerations remain valid, with the dot product $u \cdot v$ being replaced by $\langle u, v \rangle$. In particular, (3.3) must be replaced by

$$S(Z) = -\frac{1}{2} \int_0^{\tau \wedge T} \langle b(z_s), dz_s \rangle + \frac{1}{4} \int_0^{\tau \wedge T} \|b(z_s)\|^2 ds, \quad (3.16)$$

where $\|v\| = \sqrt{\langle v, v \rangle}$ is the norm induced by the scalar product $\langle \cdot, \cdot \rangle$. Another important class of systems are Langevin-type diffusions with degenerate noise:

$$\begin{aligned} dx_s &= y_s ds \\ dy_s &= -(\nabla V(x_s) + y_s) ds + \sqrt{2\epsilon}dw_s, \end{aligned} \quad (3.17)$$

Here $(x_s, y_s) \in \mathbb{R}^{2n}$ are the state variables and $V: \mathbb{R}^n \rightarrow \mathbb{R}$ is a smooth potential energy that is bounded below and sufficiently growing at infinity; more general variants of (3.17) can be considered too, but we refrain from discussing the most general scenario here. Langevin diffusions have the property that, even though the noise is degenerate and hence the tilting of the distribution can be only in the direction of some variables, one has control over the full path space distribution.[†]

[†]The reason for this lies in the fact that the Langevin equation satisfies a condition known as *complete controllability*, which ensures that noise drives all degrees of freedom in the system [13].

REMARK 4. *If the terminal time T is large, it is possible to suppress the time dependence of the basis functions and consider only functions $\phi_j = \phi_i(x)$. In this case the optimal tilting is not explicitly time-dependent as is the case in optimal stopping problems (see, e.g., [16]).*

4 Numerical examples In this section, we will study the cross-entropy method with some concrete dynamics and illustrate some numerical results.

4.1 Optimal transition of a one-dimensional bistable system We begin by studying the following optimal control problem: minimize

$$J(u) = \mathbf{E}_{\mu_u} \left(\tau + \frac{1}{4} \int_0^\tau u_s^2 ds \right), \quad (4.1)$$

under the tilted dynamics

$$dz_s = (u_s - V'(z_s)) ds + \sqrt{2\epsilon} dw_s. \quad (4.2)$$

with w_s standard one-dimensional Brownian motion and the bistable potential

$$V(x) = \frac{1}{2}(x^2 - 1)^2. \quad (4.3)$$

The double well potential V has two minima at $x_0 = -1$ and $x_2 = 1$ and a local maximum at $x_1 = 0$, and we define

$$\tau = \inf\{s > 0: |z_s - x_2| \leq 1\} \quad (4.4)$$

to be the first hitting time of the right well. The temperature (noise level) is set to $\epsilon = 0.2$. Throughout this section we choose a fixed initial condition $z_0 = x_0$.

The control task thus consists in minimizing the transition time from the left to the right well by tilting the potential energy landscape, while penalizing too strong tilting. The cost functional considered here is a variant of (2.16) for $T \rightarrow \infty$ with running cost $G = 1$ and terminal cost $H = 0$. In all numerical computations, however, $T = \infty$ is replaced by a large but finite terminal time $T < \infty$, so that $\tau \wedge T \approx \tau < \infty$; the latter is to make sure that Girsanov's theorem is applied to a finite stopping time.

Representation and optimization of control policies. As basis functions for representing the feedback controls we use three (unnormalized) Gaussians of the form

$$\phi_i(x) = \exp\left(-\frac{(x - x_i)^2}{2r^2}\right), \quad i = 0, 1, 2, \quad (4.5)$$

with $r = 0.5$ (see Fig. 4.1). Note that the basis functions are independent of time, which is due to the fact that the time dependence of the optimal tilting is relatively weak in our case.

We generate trajectories using the Euler-Maruyama scheme with time step $dt = 1.0 \times 10^{-3}$. The number of realizations used in Algorithm 1 is set to $N_j = 10^4$ for all temperature steps $\epsilon_j = (2j + 1)^{-1}$, $j = 0, 1, 2$. The algorithm is initialized with $\mathbf{v}^{(0)} = \mathbf{0}$, from which $\mathbf{v}^{(j+1)}$ is obtained in the j th step with $j = 0, 1, 2$. Note that applying a control force $u_s^j = c(s, z_s; \mathbf{v}^{(j)})$ in the j -th iteration is equivalent to modifying the potential by

$$V^{I,j}(x) = V(x) - 2 \sum_{i \in I} v_i^{(j)} \phi_i(x) \quad (4.6)$$




Fig. 4.1: One-dimensional basis functions ϕ_0, ϕ_1, ϕ_2 .

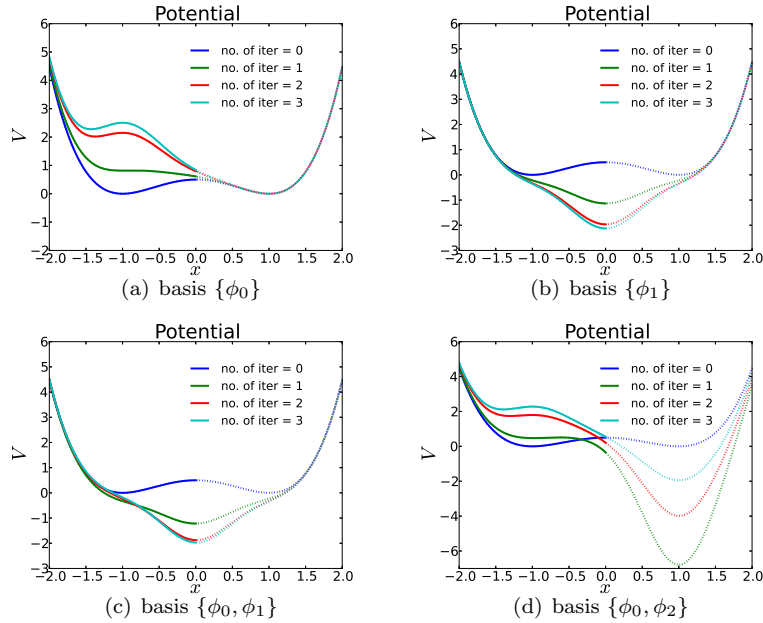


Fig. 4.2: Effective potentials for one-dimensional dynamics. The modified potentials $V^{I,j}(x)$, $j = 0, 1, 2, 3$, obtained by performing Algorithm 1, with different sets of basis functions as explained in the text.

where $I \subset \mathbb{N}$ is the index set of the basis functions. We denote the optimized potential by $V^I(x) = V^{I,3}(x)$. The numerical results for index sets $I \subset \{0, 1, 2\}$ are presented in Figure 4.2 and Figure 4.3. Figure 4.2 shows the modified potentials using four different index sets

$$I \in \{\{0\}, \{1\}, \{0, 1\}, \{0, 2\}\} . \quad (4.7)$$

It can be seen that the solution is relatively sensitive to basis functions that either do not capture the relevant region of state space (here: the transition region around the maximum at $x_1 = 0$) or that are supported in regions that are not sampled (here: rightmost energy well).

Comparison with reference solution. To analyze the accuracy of the cross-entropy algorithm in more detail, we computed the solution of the optimal control problem (4.1) by solving a Feymann-Kac type elliptic boundary value problem using a highly accurate finite element discretization (see [20] for details); this is our reference solution. We then apply the cross-entropy method with 17 Gaussian basis functions with centres $a_k = -1.5 + 0.1k$, $k = 0, 1, \dots, 16$ and variance $r = 0.1$ and compute the modified potential via Algorithm 1. From Fig. 4.3, we see that cross-entropy solutions with 17 basis functions can approximate the reference solution quite well. We also observe that a similarly good approximation can be obtained with a *single* well-chosen basis function ϕ_1 . This indicates the possibility to solve high-dimensional problems with few basis functions.

With the optimized potentials, we then generate $N = 10^6$ samples from the controlled dynamics and compute the value of the associated cost function (4.1). The results are presented in Table 4.1: We clearly observe that the best results are obtained when the basis functions capture the relevant transition region (here: ϕ_1), since with the index sets 2 and 3, the cost value 1.31 is closer to the exact solution 1.25 than elsewhere.




Fig. 4.3: Optimal control problem. Optimized potential $V^I(x)$ with a single basis function ϕ_1 compared to the cross-entropy optimized potential with 17 basis functions and the (“exact”) reference solution. The potentials are vertically shifted for better presentation.

4.2 Exit time distribution of a one-dimensional bistable system We continue our study with the computation of the exit time distribution $\mathbf{P}(\tau \leq T)$ of the bistable dynamics from the last subsection; the dynamics and the parameters are the same as in Section 4.1. Our aim is to compute a discrete approximation of the exit time distribution using $T \in \{2.0, 1.0, 0.5, 0.3, 0.2, 0.1\}$ using importance sampling with two sets of basis functions, $\{\phi_0\}$ and $\{\phi_1\}$ and compare it to standard Monte Carlo.

Acceleration index. In order to measure the speed-up gained by importance sampling (IS) compared to standard Monte Carlo (MC) we define the acceleration index

$$\mathcal{I} = \frac{\text{Var}_{MC}}{\text{Var}_{IS}} \quad (4.8)$$

as the ratio of MC and IS sample variances. According to the central limit theorem, the variances of the two methods will decrease as N^{-1} with the number N of trajectories, and hence the

basis set	coefficients	cost	mean of τ
$\{\phi_0\}$	$(-1.252, 0, 0)$	5.14	2.02
$\{\phi_1\}$	$(0, 1.313, 0)$	1.31	0.57
$\{\phi_0, \phi_1\}$	$(-0.078, 1.246, 0)$	1.31	0.57
$\{\phi_0, \phi_2\}$	$(-1.139, 0, 0.975)$	3.88	1.73
17 Gaussians /reference		1.27/1.25	0.52/0.52

Table 4.1: One-dimensional optimal control problem with different sets of basis functions. The cost is given by the expectation (4.1) that is computed from $N = 10^6$ samples. The last column shows the mean of the stopping time τ under the modified dynamics. The last row displays the results for 17 Gaussian basis functions and the reference solution.

T	coefficients	$\mathbf{P}(\tau \leq T)$	Var	Accel. \mathcal{I}	Traj. Usage
2.0	$(-0.592, 0, 0)$	9.22×10^{-2}	1.8×10^{-2}	4.7	61%
1.0	$(-0.984, 0, 0)$	3.23×10^{-2}	2.2×10^{-3}	13.7	48%
0.5	$(-1.570, 0, 0)$	6.45×10^{-3}	1.3×10^{-4}	50.3	39%
0.3	$(-2.321, 0, 0)$	9.51×10^{-4}	4.8×10^{-6}	198.6	33%
0.2	$(-3.219, 0, 0)$	9.53×10^{-5}	8.7×10^{-8}	1091.8	27%
0.1	$(-5.830, 0, 0)$	1.22×10^{-7}	7.0×10^{-13}	very large	16%

Table 4.2: Computation of $\mathbf{P}(\tau \leq T)$ by importance sampling, based on $N = 10^6$ independent realizations a single basis function ϕ_0 . “Accel.” \mathcal{I} (acceleration) measures the computational speed-up of importance sampling relative to standard MC. “Traj. Usage” denotes the portion of trajectories satisfying $\tau \leq T$ under the modified dynamics.

acceleration index \mathcal{I} has the following interpretation: If N is sufficiently large and IS reaches a certain error with N trajectories, MC requires about $\mathcal{I}N$ trajectories to achieve the same error. Thus, \mathcal{I} is the speed-up factor of IS relative to MC, provided that we can ignore the computational overhead associated with importance sampling (which is the case here).

Comparison with standard Monte-Carlo. Tables 4.2–4.4 and Figure 4.4. Tables 4.2 and 4.3 show the results of importance sampling with basis function $\{\phi_0\}$ and $\{\phi_1\}$, respectively, Table 4.4 records the result of vanilla Monte Carlo. We observe that, for each value of T , the variances of the importance samplers are largely reduced compared to those of standard Monte Carlo. The difference increases when T decreases as is to be expected, for the event $\{\tau \leq T\}$ is rarer the smaller T is. Note that for $T = 0.1$ standard Monte Carlo cannot be used at all because not a single realization is generated that hits the set boundary, while importance sampling still gives a reasonable estimate; see Table 4.4.

In Figure 4.4, the optimized potentials resulting from IS are plotted for different values of T and different sets of basis functions. As expected the optimized potentials become increasingly different from the original potential the smaller the value of T , indicating that larger forces are needed when the event is rarer.

T	coefficients	$\mathbf{P}(\tau \leq T)$	Var	Accel. \mathcal{I}	Traj. Usage
2.0	(0, 0.680, 0)	9.23×10^{-2}	9.2×10^{-3}	9.1	85%
1.0	(0, 1.059, 0)	3.23×10^{-2}	1.0×10^{-3}	30.0	81%
0.5	(0, 1.636, 0)	6.46×10^{-3}	7.4×10^{-5}	86.9	68%
0.3	(0, 2.360, 0)	9.49×10^{-4}	3.0×10^{-6}	310.9	56%
0.2	(0, 3.237, 0)	9.56×10^{-5}	5.9×10^{-8}	1621.2	46%
0.1	(0, 5.821, 0)	1.21×10^{-7}	4.1×10^{-13}	very large	26%

Table 4.3: Computation of $\mathbf{P}(\tau \leq T)$ for the one-dimensional dynamics, based on $N = 10^6$ independent realizations a single basis function $\{\phi_1\}$; see Tab. 4.2 for comparison.

T	$\mathbf{P}(\tau \leq T)$	Var	Accel.	Traj. Usage
2.0	9.23×10^{-2}	8.4×10^{-2}	1.0	9.2%
1.0	3.22×10^{-2}	3.1×10^{-2}	1.0	3.2%
0.5	6.28×10^{-3}	6.2×10^{-3}	1.0	0.6%
0.3	1.00×10^{-3}	1.0×10^{-3}	1.0	0.1%
0.2	9.30×10^{-5}	9.3×10^{-5}	1.0	0.009%
0.1	0.00	–	1.0	0.0%

Table 4.4: Computation of $\mathbf{P}(\tau \leq T)$ for the one-dimensional dynamics by standard Monte-Carlo, based on $N = 10^6$ independent realizations; compare Tabs. 4.2 and 4.3.

4.3 Conformational transition of solvated butane As a nontrivial test case, we now apply the cross-entropy algorithm to the conformational dynamics of solvated butane. Specifically, we compute the cumulative distribution function of the *gauche-trans* transition time.

The butane is simulated in a $3.0 \times 3.0 \times 3.0 \text{ nm}^3$ box with periodic boundary conditions, using the GROMOS 45a3 force field [19] with a modified GROMACS 4.5 [17]. The simulation box contains 900 SPC/E [2] water molecules. The dynamics are governed by the underdamped Langevin equation

$$\begin{aligned} dr_i &= m_i^{-1} p_i dt, \\ dp_i &= [f_i(r) + u_i] dt - \gamma p_i dt + \sigma_i dw_t, \end{aligned} \quad (4.9)$$

where $r = (r_1, \dots, r_N) \in \mathbb{R}^{3N}$ and $p = (p_1, \dots, p_N) \in \mathbb{R}^{3N}$, with r_i denoting the Cartesian coordinate of the i -th atom and with $p_i \in \mathbb{R}^3$ being the conjugate momentum. Friction constant γ and noise coefficients σ_i are coupled by the fluctuation-dissipation relation $\sigma_i^2 = 2\gamma m_i k_B \Theta$ where k_B is Boltzmann’s constant, Θ is the temperature and m_i is the mass of the i -th atom (here: $\epsilon = k_B \Theta$). Here $f_i(r)$ denotes the force resulting from the GROMOS 45a3 force field and u_i is the additional biasing force on the i -th atom. The force term reads

$$f_i(r) = -\nabla_{r_i} V_{\text{bonded}}(r) - \nabla_{r_i} V_{\text{non-bonded}}(r). \quad (4.10)$$




Fig. 4.4: The modified potentials for $\mathbf{P}(\tau \leq T)$ with different sets of basis functions. The potentials become increasingly different from the original one when T decreases. The rightmost panel shows the cumulative distribution function of τ for $T \in [0.1, 2.0]$.




Fig. 4.5: The butane molecule.

It involves bonded and non-bonded interactions. The bonded interactions are

$$\begin{aligned}
 V_{\text{bonded}}(r) &= V_b(r) + V_\theta(r) + V_\phi(r) \\
 &= \sum_{k=1}^{N_b} \frac{1}{4} k_b (b_k^2(r) - b_{k,0}^2)^2 + \sum_{k=1}^{N_\theta} \frac{1}{2} k_\theta (\cos(\theta_k(r)) - \cos \theta_{k,0})^2 \\
 &\quad + \sum_{k=1}^{N_\phi} k_\phi [1 + \cos(\delta_k) \cos(m_k \phi_k(r))] \tag{4.11}
 \end{aligned}$$

where the bond potential $V_b(r)$ is the energy due to the covalent bonds in the system, with $b_k(r)$ being the instantaneous length of the k -th bond, $b_{k,0}$ the constant equilibrium length, and k_b the force constant, $V_\theta(r)$ is the energy of the bond angle, with $\theta_k(r)$, $\theta_{k,0}$ and k_θ denoting instantaneous angle, equilibrium angle and force constant; the third term is the energy of the torsional dihedral angle, with $\phi_k(r)$, δ_k , m_k and k_ϕ denoting the instantaneous dihedral angle, phase shift, multiplicity and the force constant. For a butane molecule as illustrated in Figure 4.5 covalent bonds are between the atoms 1–2, 2–3 and 3–4, bond angle interactions are between atom triples 1–2–3 and 2–3–4, and dihedral angle interactions involve the atom quadruple 1–2–3–4. The orders of magnitude of k_b ($= 7.15 \times 10^6$ kJ/(mol \times nm 4) for all bonds) and k_θ ($= 530$ kJ/mol for all bond angles) are much larger than the dihedral angle constant k_ϕ ($= 5.92$ kJ/mol).

Therefore the vibrations of the bonds and bond angles can be viewed as fast motions in the system. The butane dihedral angle potential has multiplicity $m = 3$ and phase shift $\delta = 0$, which implies that the dihedral angle is essentially populated around three angles -60° , 60° and 180° . The first two are called *gauche conformations*, the other one is called *trans conformation*; transitions between these conformations are the slow motions in the system. We constrain the bonds and bond angles of the water molecules using the SETTLE algorithm [14]. The non-bonded interaction potential is given by

$$V_{\text{non-bonded}}(r) = \sum_{i,j} \left[\frac{C_{12,ij}}{r_{ij}^{12}} - \frac{C_{6,ij}}{r_{ij}^6} \right] + \sum_{i,j} \frac{1}{4\pi\epsilon_0} \frac{q_i q_j}{r_{ij}}, \quad (4.12)$$

where r_{ij} is the relative distance between atom i and j . The first term of (4.12) is the van der Waals interaction and the second term is the Coulomb interaction, where q_i denotes the partial charge due to the atom i . During the simulation, the van der Waals potential is evaluated numerically by the cut-off method with cut-off radius 9 nm, while for the Coulomb energy the smooth particle mesh Ewald method (SPME) is used [8].

Importance sampling and parametrization of control policies. In order to compute $\mathbf{P}(\tau \leq T)$, we first generate an ensemble of initial conditions and run an equilibrium simulation (i.e. with $u = 0$ in (4.9)) at 300 K, with friction constant 1.0 ps^{-1} . The dihedral angle ϕ of the butane molecule is monitored for every interval of 10 ps. If ϕ is in the range $[40^\circ, 80^\circ]$ (corresponding to the *gauche* conformation), then the system state (including all water coordinates) is recorded, and used for the later study of *gauche-trans* transition. In our simulation, we only record the first 5000 of these system states as an equilibrium sample of the *gauche* conformation.

As a next step, we run stopped simulations of (4.9) at 300 K with feedback control ($u \neq 0$) and friction constant 10 ps^{-1} ; the process is stopped at $\tau \wedge T$ where τ is the time of first exit from the *gauche* conformation. We define the exit from the *gauche* conformation as the entrance to the *trans* conformation, which happens when the dihedral angle ϕ is equal or larger than 150° .

The candidate feedback control $u_t = c(r_t; \boldsymbol{\lambda})$ with $c = (c_1, \dots, c_N)$ used in the simulation is assumed to depend only on the dihedral angle (that is a function of atomic coordinates):

$$c_i(r; \boldsymbol{\lambda}) = -\nabla_{r_i} V_{\text{ctrl}}(\phi(r)) = -\nabla_{r_i} \left[\sum_{k=1}^{N_c} \lambda_k \cos(k\phi(r)) \right]. \quad (4.13)$$

The assumption that the control depends only on the dihedral angle is justified by the following three observations: (1) the dihedral angle fully describes the conformational transition of the butane molecule, (2) it is the slowest degree of freedom in the system relative to bond lengths, bond angles and the motions of the water molecules, and (3) the explicit time dependence of the control force is negligible; the influence of the potentially slow overall rotations of the butane molecule that can couple to the internal conformational degrees of freedom on the optimal change of measure is ruled out by numerical tests as is described below. As a consequence the optimal control will be a function of the dihedral angle only—at least to a good approximation [25]. In this simulation, we choose the control potential $V_{\text{ctrl}}(\phi(r))$ to be a sum of cosine functions, because the system is symmetric around $\phi = 0^\circ$; the number of control functions is kept fixed at $N_c = 8$ throughout the simulation. Numerical results confirm that the coefficient λ_8 is already very small compared to the dominant coefficients, which means that the number of basis functions




Fig. 4.6: The dihedral potential and the dihedral potential with added control potential.

is large enough to capture the essential structure of the control. In all simulations the Langevin equation (4.9) is discretized by the BAOAB scheme [11] with constant time step of 5×10^{-4} ps.

We apply the cross-entropy algorithm to butane with $T = 1.0, 0.5, 0.2$ and 0.1 ps. From each of the 5,000 equilibrium system states from the *gauche* conformation, we have launched 4 independent realizations of length $\tau \wedge T$ for $T = 1.0, 0.5,$ and 0.2 , which gives $M_{IS} = 20,000$ independent trajectories; for the smallest value $T = 0.1$ ps, we have simulated 12 independent trajectories from each initial condition, resulting in 60,000 trajectories in total.

In order to compute the optimal control force u^* efficiently we employ a further simplification and remove all the water molecules. This is done because the vacuum simulation is much cheaper than the simulation with water and the control forces calculated in vacuum are an accurate approximation of the controls for the solvated molecule; the latter was verified numerically by applying the cross-entropy method to solvated butane with the vacuum solution as initial guess; no further iteration of the control forces was needed in this case.

Simulation results. The results of the importance sampling (IS) computations are summarized in Figure 4.6, Table 4.5 and Table 4.6. In vacuum, we find probabilities $\mathbf{P}(\tau \leq T) = 2.16 \times 10^{-2}, 8.66 \times 10^{-3}, 1.48 \times 10^{-3}$ and 6.13×10^{-5} for $T = 1.0, 0.5, 0.2$ and 0.1 ps, respectively. These values do not significantly differ from those of the solvated system as is shown in the second column of Tab. 4.5. For comparison, Table 4.6 shows the result reference estimates of $\mathbf{P}(\tau \leq T)$, based on brute-force Monte-Carlo (MC) simulation with $M_{MC} = 100,000$ independent realizations. The small absolute error of the IS scheme indicates that the assumption that the control can be expressed solely in terms of the dihedral angle is reasonable. Only for $T = 0.1$ ps the reference MC estimate cannot be trusted, because only 9 of the 100,000 trajectories hit the *trans* conformation before time T ; hence for $T = 0.1$ ps, neither the estimate of $\mathbf{P}(\tau \leq T)$ nor of the error are reliable.

The fourth column of Table 4.5 shows the variance of the IS estimator that is drastically reduced as compared with the brute force MC simulation. As before we have computed the total acceleration index \mathcal{I} , as a result of both variance reduction and the speed up of the rare transition events (see the fifth column of Table 4.5). The column "trajectory usage" presents

T [ps]	$\mathbf{P}(\tau \leq T)$	Error	Var	Accel. \mathcal{I}	Traj. Usage
0.1	4.30×10^{-5}	0.77×10^{-5}	3.53×10^{-6}	12.2	0.4%
0.2	1.21×10^{-3}	0.11×10^{-3}	2.50×10^{-4}	4.8	5.4%
0.5	6.85×10^{-3}	0.38×10^{-3}	2.88×10^{-3}	2.4	8.3%
1.0	1.74×10^{-2}	0.08×10^{-2}	1.21×10^{-2}	1.4	12.3%

Table 4.5: Results for solvated butane, with controls acting on the dihedral angle only. “Error” denotes the expected error $\sqrt{\text{Var}/M_{IS}}$ of the IS estimator, where M_{IS} is the number of trajectories used. The meaning of the other columns is the same as in Tab. 4.2.

T [ps]	$\mathbf{P}(\tau \leq T)$	Error	Var	Accel.	Traj. Usage
0.1	9.00×10^{-5}	3.00×10^{-5}	9.00×10^{-5}	1.0	0.009%
0.2	1.29×10^{-3}	0.11×10^{-3}	1.29×10^{-3}	1.0	0.1%
0.5	7.41×10^{-3}	0.27×10^{-3}	7.36×10^{-3}	1.0	0.7%
1.0	1.78×10^{-2}	0.04×10^{-2}	1.75×10^{-2}	1.0	1.8%

Table 4.6: Results for solvated butane: Brute force Monte Carlo estimate of $\mathbf{P}(\tau \leq T)$.

the percentage of the trajectories that makes it to the *trans* conformation within time interval $[0, T]$. Figure 4.6 shows the effective dihedral angle potential (i.e. the original dihedral potential $V_\phi(\phi)$ plus the control potential $V_{\text{ctrl}}(\phi)$) where we only show the effective energy in the range $[40^\circ, 150^\circ]$, because the initial states of the trajectories are located in the range $[40^\circ, 80^\circ]$, and the trajectories are stopped when they reach $\phi = 150^\circ$. For an easy comparison, all effective energies are shifted by a constant, so that they all coincide at $\phi = 150^\circ$. As expected the resulting control forces are stronger the smaller T is.

5 Discussions As a continuation of our works [25, 24], we propose a cross-entropy algorithm for diffusion processes and study its application to importance sampling and optimal control. For instance, in our previous work [24], we have analyzed the effect of the use of sub-optimal controls in multiscale systems with explicit scale separation, e.g., slow-fast systems or diffusions in the small-noise limit. Here the situation is different, in that no such small parameter or detailed information regarding the relevant degrees of freedom is used. On the other hand, in the cross-entropy method, the approximation of the target measure and, consequently, efficient importance sampling or control strategies crucially depend on a sensible choice of a function basis. A good choice can be often based on prior knowledge about the dynamical system, such as metastable states or reactive coordinates; it is easy to imagine that there is no way to obtain satisfactory results when the basis functions used are not supported along the relevant degrees of freedom. The relation between the cross-entropy method and methods for multiscale dynamical systems is an interesting and yet open question that will be addressed in future work.

Acknowledgement The authors acknowledge financial support by the DFG Research Center MATHEON and the Einstein Center of Mathematics (ECMath).

REFERENCES

- [1] E. ASPLUND AND T. KLÜNER, *Optimal control of open quantum systems applied to the photochemistry of surfaces*, Phys. Rev. Lett., 106 (2011), p. 140404.
- [2] H. BERENDSEN, J. GRIGERA, AND T. STRAATSMA, *The missing term in effective pair potentials*, Journal of Physical Chemistry, 91 (1987), pp. 6269–6271.
- [3] P. DAI PRA., L. MENEGHINI, AND W. RUNGALDIER, *Connections between stochastic control and dynamic games*, Mathematics of Control, Signals and Systems, 9 (1996), pp. 303–326.
- [4] P.-T. DE BOER, D. KROESE, S. MANNOR, AND R. RUBINSTEIN, *A Tutorial on the Cross-Entropy Method*, Annals of Operations Research, 134 (2005), pp. 19–67.
- [5] P. DUPUIS, K. SPILIOPOULOS, AND H. WANG, *Importance sampling for multiscale diffusions*, Multiscale Model. Simul., 10 (2012), pp. 1–27.
- [6] P. DUPUIS AND H. WANG, *Importance sampling, large deviations, and differential games*, Stochastics and Stochastic Reports, 76 (2004), pp. 481–508.
- [7] D. DÜRR AND A. BACH, *The Onsager-Machlup function as Lagrangian for the most probable path of a diffusion process*, Comm. Math. Phys., 60 (1978), pp. 153–170.
- [8] U. ESSMANN, L. PERERA, M. BERKOWITZ, T. DARDEN, H. LEE, AND L. PEDERSEN, *A smooth particle mesh ewald method*, The Journal of Chemical Physics, 103 (1995), p. 8577.
- [9] W. H. FLEMING AND H. M. SONER, *Controlled Markov Processes and Viscosity Solutions*, Springer, 2006.
- [10] C. HARTMANN, R. BANISCH, M. SARICH, T. BADOWSKI, AND C. SCHÜTTE, *Characterization of rare events in molecular dynamics*, Entropy, 16 (2014), pp. 350–376.
- [11] B. LEIMKUHLE AND C. MATTHEWS, *Rational construction of stochastic numerical methods for molecular sampling*, Applied Mathematics Research eXpress, 2013 (2013), pp. 34–56.
- [12] J. S. LIU, *Monte Carlo Strategies in Scientific Computing*, Springer, first ed., 2008.
- [13] J. C. MATTINGLY AND A. M. STUART, *Geometric ergodicity of some hypo-elliptic diffusions for particle motions*, Markov Process. Related Fields, 8 (2002), pp. 199–214.
- [14] S. MIYAMOTO AND P. KOLLMAN, *SETTLE: an analytical version of the SHAKE and RATTLE algorithm for rigid water models*, Journal of Computational Chemistry, 13 (2004), pp. 952–962.
- [15] B. ØKSENDAL, *Stochastic Differential Equations: An Introduction with Applications*, Springer, 6th ed., 2010.
- [16] G. PESKIR AND A. SHIRYAEV, *Optimal Stopping and Free-Boundary Problems*, Birkhäuser, Basel, 2006.
- [17] S. PRONK, S. PÁLL, R. SCHULZ, P. LARSSON, P. BJELKMAR, R. APOSTOLOV, M. SHIRTS, J. SMITH, P. KASSON, D. VAN DER SPOEL, AND H. B. L. E., *Gromacs 4.5: a high-throughput and highly parallel open source molecular simulation toolkit*, Bioinformatics, (2013), pp. 1–10.
- [18] R. Y. RUBINSTEIN AND D. P. KROESE, *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning (Information Science and Statistics)*, Springer, 1 ed., 2004.
- [19] L. D. SCHULER, X. DAURA, AND W. F. VAN GUNSTEREN, *An improved gromos96 force field for aliphatic hydrocarbons in the condensed phase*, Journal of Computational Chemistry, 22 (2001), pp. 1205–1218.
- [20] C. SCHÜTTE, S. WINKELMANN, AND C. HARTMANN, *Optimal control of molecular dynamics using markov state models*, Mathematical Programming, 134 (2012), pp. 259–282.
- [21] H. STAPELFELDT, *Laser aligned molecules: Applications in physics and chemistry*, Physica Scripta, 2004 (2004), pp. 132–136.
- [22] A. STEINBRECHER, *Optimal control of robot guided laser material treatment*, in Progress in Industrial Mathematics at ECMI 2008, A. D. Fitt, J. Norbury, H. Ockendon, and E. Wilson, eds., Springer Berlin Heidelberg, 2010, pp. 505–511.
- [23] E. VANDEN-EIJNDEN AND J. WEARE, *Rare event simulation of small noise diffusions*, Comm. Pure Appl. Math., 65 (2012), pp. 1770–1803.
- [24] W. ZHANG, C. HARTMANN, M. WEBER, AND C. SCHÜTTE, *Importance sampling in path space for diffusion processes*, submitted, (2014). Preprint available under publications.mi.fu-berlin.de/1321/.
- [25] W. ZHANG, J. LATORRE, , G. PAVLIOTIS, AND C. HARTMANN, *Optimal control of multiscale systems using reduced-order models*, arXiv:1406.3458v1, (2014).
- [26] M. ZYSKOWSKI, M. SAIN, AND R. DIERSING, *State-feedback, finite-horizon, cost density-shaping control for the linear quadratic gaussian framework*, J. Optim. Theory Appl., 150 (2011), pp. 251–274.