

Diffusion Models for Constrained Domains

Nic Fishman

Department of Statistics, University of Oxford

njwfish@gmail.com

Leo Klarner

Department of Statistics, University of Oxford

leo.klarner@stats.ox.ac.uk

Valentin De Bortoli

CNRS, ENS Ulm

valentin.debortoli@gmail.com

Emile Mathieu

Department of Engineering, University of Cambridge

ebm32@cam.ac.uk

Michael Hutchinson

Department of Statistics, University of Oxford

michael.hutchinson@stats.ox.ac.uk

Reviewed on OpenReview: <https://openreview.net/forum?id=xuWTFQ4VGO>

Abstract

Denosing diffusion models are a novel class of generative algorithms that achieve state-of-the-art performance across a range of domains, including image generation and text-to-image tasks. Building on this success, diffusion models have recently been extended to the Riemannian manifold setting, broadening their applicability to a range of problems from the natural and engineering sciences. However, these Riemannian diffusion models are built on the assumption that their forward and backward processes are well-defined for all times, preventing them from being applied to an important set of tasks that consider manifolds defined via a set of inequality constraints. In this work, we introduce a principled framework to bridge this gap. We present two distinct noising processes based on (i) the *logarithmic barrier* metric and (ii) the *reflected* Brownian motion induced by the constraints. As existing diffusion model techniques cannot be applied in this setting, we derive new tools to define such models in our framework. We then demonstrate the practical utility of our methods on a number of synthetic and real-world tasks, including applications from robotics and protein design.

1 Introduction

Diffusion models (Sohl-Dickstein et al., 2015; Song & Ermon, 2019; Song et al., 2021; Ho et al., 2020) have recently been introduced as a powerful new paradigm for generative modelling. They work as follows: noise is progressively added to data following a Stochastic Differential Equation (SDE)—the forward *noising* process—until it is approximately Gaussian. The generative model is given by an approximation of the associated *time-reversed* process called the backward *denoising* process. This is also an SDE whose drift depends on the gradient of the logarithmic densities of the forward process, referred to as the Stein score. This score is approximated by leveraging techniques from deep learning and score matching (Hyvärinen, 2005; Vincent, 2011). Building on the success of diffusion models in domains such as images and text, this framework has recently been extended to a wide range of Riemannian manifolds (De Bortoli et al., 2022; Huang et al., 2022), broadening their applicability to various important modelling domains from the natural and engineering sciences—including Lie groups such as the group of rotations $SO(3)$, the group of rigid body motions $SE(3)$, and many others (see e.g. Trippe et al., 2022; Corso et al., 2022; Watson et al., 2022; Leach et al., 2022; Urain et al., 2022; Yim et al., 2023).




Figure 1: The behaviour of different types of noising processes considered in this work defined on the unit interval. \mathbf{B}_t : Unconstrained (Euclidean) Brownian motion. $\hat{\mathbf{B}}_t$: Log-barrier forward noising process. $\bar{\mathbf{B}}_t$: Reflected Brownian motion. All sampled with the same initial point and driving noise. Black line indicates the boundary.

However, a key assumption of the Riemannian diffusion models introduced in De Bortoli et al. (2022) and Huang et al. (2022) is that the stochastic processes they consider are defined *for all times*. While this holds for a large class of stochastic processes, it is not the case for most manifolds defined via a set of inequality constraints. For instance, in the case of the hypercube $(-1, 1)^d$ equipped with the Euclidean metric, the Riemannian Brownian motion coincides with the Euclidean d -dimensional Brownian motion $(\mathbf{B}_t)_{t \in [0, T]}$ as long as $\mathbf{B}_t \in (-1, 1)$. With probability one, $(\mathbf{B}_t)_{t \geq 0}$ escapes from $(-1, 1)$, meaning that the Riemannian Brownian motion is not defined for all times and the frameworks introduced in De Bortoli et al. (2022) and Huang et al. (2022) do not apply. Such constrained manifolds comprise a wide variety of settings—including polytopes and convex sets of Euclidean spaces—and are studied across a large number of disciplines, ranging from computational statistics (Morris, 2002), over robotics (Han & Rudolph, 2006) and quantum physics (Lukens et al., 2020), to computational biology (Thiele et al., 2013). Deriving principled diffusion models that are able to operate directly on these manifolds is thus of significant practical importance, as they enable generative modelling in data-scarce and safety-critical settings in which constraints on the modelled domain may reduce the number of degrees of freedom or prevent unwanted behaviour.

As sampling problems on such manifolds are important (Kook et al., 2022; Heirendt et al., 2019), a flurry of Markov chain based methods have been developed to sample from unnormalised densities. Successful algorithms include the reflected Brownian motion (Williams, 1987; Petit, 1997; Shkolnikov & Karatzas, 2013), log-barrier methods (Kannan & Narayanan, 2009; Lee & Vempala, 2017; Noble et al., 2022; Kook et al., 2022; Gatmiry & Vempala, 2022; Lee & Vempala, 2018) and hit-and-run approaches in the case of polytopes (Smith, 1984; Lovász & Vempala, 2006). In this work, we study the generative modelling counterparts of these algorithms through the lens of diffusion models. Among existing methods for statistical sampling on constrained manifolds, the geodesic Brownian motion (Lee & Vempala, 2017) and the reflected Brownian motion (Williams, 1987) are continuous stochastic processes, and thus well suited for extending the continuous Riemannian diffusion framework developed by De Bortoli et al. (2022) and Huang et al. (2022). In particular, we introduce two principled diffusion models for generative modelling on constrained domains based on (i) the geodesic Brownian motion, leveraging tools from the log-barrier methods, and (ii) the reflected Brownian motion. In both cases, we show how one can extend the ideas of time-reversal and score matching to these settings. We demonstrate the practical utility of these methods on a range of tasks defined on convex polytopes and the space of symmetric positive definite matrices, including the constrained conformational modelling of proteins and robotic arms. The code for all of our experiments is available [here](#).

2 Background

Riemannian manifolds. A Riemannian manifold is a tuple $(\mathcal{M}, \mathbf{g})$ with \mathcal{M} a smooth manifold and \mathbf{g} a metric which defines an inner product on tangent spaces. The metric \mathbf{g} induces key quantities on the manifold, such as an exponential map $\exp_x : T_x \mathcal{M} \rightarrow \mathcal{M}$, defining the notion of following straight lines on manifolds, a gradient operator ∇^1 and a divergence operator div^2 . It also induces the Laplace-Beltrami operator Δ and

¹The (Riemannian) gradient ∇ is defined s.t. for any smooth $f \in C^\infty(\mathcal{M})$, $x \in \mathcal{M}$, $v \in T_x \mathcal{M}$, $\mathbf{g}(x)(\nabla f(x), v) = \text{d}f(x)(v)$.

²The Riemannian divergence acts vector fields and can be defined using the volume form of \mathcal{M} .




Figure 2: A convex polytope defined by six constraints $\{f_i\}_{i \in \mathcal{I}}$, along with the log barrier potential, and ‘straight trajectories’ under the log-barrier metric and under the Euclidean metric with and without reflection at the boundary.

consequently a Brownian motion with density (w.r.t. the volume form³), whose density is given by the heat equation $\partial_t p_t = \Delta p_t$. We refer the reader to Appendix A for a brief introduction to differential geometry, to Lee (2013) for a thorough treatment and to Hsu (2002) for details on stochastic analysis on manifolds.

Constrained manifolds. In this work, we are concerned with *constrained* manifolds. More precisely, given a Riemannian manifold $(\mathcal{N}, \mathfrak{h})$, we consider a family of real functions $\{f_i : \mathcal{N} \rightarrow \mathbb{R}\}_{i \in \mathcal{I}}$ indexed by \mathcal{I} . We then define

$$\mathcal{M} = \{x \in \mathcal{N} : f_i(x) < 0, i \in \mathcal{I}\}. \quad (1)$$

In this scenario, $\{f_i\}_{i \in \mathcal{I}}$ is interpreted as a set of constraints on \mathcal{N} . For example, choosing $\mathcal{N} = \mathbb{R}^d$ and affine constraints $f_i(x) = \langle a_i, x \rangle - b_i$, $x \in \mathbb{R}^d$, we get that \mathcal{M} is an open polytope as illustrated in Figure 2. This setting naturally appears in many areas of engineering, biology, and physics (Boyd et al., 2004; Han & Rudolph, 2006; Lukens et al., 2020).

While the two methods we introduce in Section 3 can be applied to the general framework (1), in our applications, we focus on two specific settings: (a) Polytopes— $\mathcal{N} = \mathbb{R}^d$ and (b) symmetric positive-definite (SPD) matrices under trace constraints— $\mathcal{N} = \mathcal{S}_{++}^d$.

Continuous diffusion models. We briefly recall the framework for constructing continuous diffusion processes introduced by Song et al. (2021) in the context of generative modelling over \mathbb{R}^d . At minimum diffusion models need four things: (i) A forward noising process converging to an invariant distribution. (ii) A time reversal for the reverse process. (iii) A discretization of the continuous-time process for the forward/reverse process. (iv) A score matching loss — in this paper we will focus on the implicit score matching loss. Song et al. (2021) consider a forward *noising process* $(\mathbf{X}_t)_{t \in [0, T]}$ which progressively noises a data distribution p_0 into a Gaussian $\mathcal{N}(0, \text{Id})$. More precisely $(\mathbf{X}_t)_{t \in [0, T]}$ is an Ornstein–Uhlenbeck (OU) process which is given by the following stochastic differential equation (SDE)

$$d\mathbf{X}_t = -\frac{1}{2}\mathbf{X}_t dt + d\mathbf{B}_t, \quad \mathbf{X}_0 \sim p_0.$$

Under mild conditions on p_0 , the time-reversed process $(\overleftarrow{\mathbf{X}}_t)_{t \in [0, T]} = (\mathbf{X}_{T-t})_{t \in [0, T]}$ also satisfies an SDE (Cattiaux et al., 2021; Haussmann & Pardoux, 1986) given by

$$d\overleftarrow{\mathbf{X}}_t = \left\{ \frac{1}{2}\overleftarrow{\mathbf{X}}_t + \nabla \log p_{T-t}(\overleftarrow{\mathbf{X}}_t) \right\} dt + d\mathbf{B}_t, \quad \overleftarrow{\mathbf{X}}_0 \sim p_T, \quad (2)$$

where p_t denotes the density of \mathbf{X}_t . This construction allows direct sampling of the forward process and leverages the Euler-Maruyama discretisation to facilitate sampling of the reverse process. Finally, the quantity $\nabla \log p_t$ is referred as the Stein score and is unavailable in practice. It can be approximated with a score network $s_\theta(t, \cdot)$ trained by minimising a denoising score matching (dsm) loss

$$\mathcal{L}(\theta) = \mathbb{E}[\lambda_t \|\nabla \log p_{t|0}(\mathbf{X}_t | \mathbf{X}_0) - s_\theta(t, \mathbf{X}_t)\|^2], \quad (3)$$

³We assume that \mathcal{M} is orientable and therefore that a volume form exists.

or an equivalent implicit score matching (ism) loss


$$\mathcal{L}(\theta) = \mathbb{E}[\lambda_t \{ \frac{1}{2} \|s_\theta(t, \mathbf{X}_t)\|^2 + \text{div}(s_\theta)(t, \mathbf{X}_t) \}] + C, \quad (4)$$

where $C \geq 0$ and $\lambda_t > 0$ is a weighting function, and the expectation is taken over $t \sim \mathcal{U}([0, T])$ and $(\mathbf{X}_0, \mathbf{X}_t)$. For an arbitrarily flexible score network, the global minimiser $\theta^* = \text{argmin}_\theta \mathcal{L}(\theta)$ satisfies $s_{\theta^*}(t, \cdot) = \nabla \log p_t$.

3 Inequality-Constrained Diffusion Models

We are now ready to introduce our methodology to deal with manifolds defined via *inequality constraints* (1). In Section 3.1, we propose a Riemannian diffusion model endowed with a metric induced by a log-barrier potential. In Section 3.2, we introduce a *reflected* diffusion model. While both models extend classical diffusion models to inequality-constrained settings, they exhibit very different behaviours. We discuss their practical differences in Section 5.1.

3.1 Log-barrier diffusion models



$$d\mathbf{X}_t = \frac{1}{2} \text{div}(\mathbf{g}^{-1})(\mathbf{X}_t) dt + \mathbf{g}^{-1/2}(\mathbf{X}_t) d\mathbf{B}_t$$

Figure 3: Convergence of the Barrier Langevin dynamics on the unit interval to the uniform distribution.




Figure 4: Illustrative diagram of the barrier method and the change of metric.

Barrier Langevin dynamics. Barrier methods work by constructing a smooth potential $\phi : \mathcal{M} \rightarrow \mathbb{R}$ such that it blows up on the boundary of a desired set, see Nesterov et al. (2018). Such potentials form the basis of interior point methods in convex optimisation (Boyd et al., 2004). Of these functions, the *logarithmic barrier* is the most popular among practitioners (Lee & Vempala, 2017). For a convex polytope \mathcal{M} defined by the constraints $Ax < b$, with $A \in \mathbb{R}^{m \times d}$ and $b \in \mathbb{R}^m$, the logarithmic barrier $\phi : \mathcal{M} \rightarrow \mathbb{R}_+$ is given for any $x \in \mathcal{M}$ by

$$\phi(x) = - \sum_{i=1}^m \log(\langle A_i, x \rangle - b_i). \quad (5)$$

Assuming that $\|A_i\| = 1$, we have that for any $x \in \mathcal{M}$, $\phi(x) = - \sum_{i=1}^m \log(d(x, \partial\mathcal{M}_i))$, where $\partial\mathcal{M}_i = \{x \in \mathbb{R}^d : \langle A_i, x \rangle = b_i\}$. More generally we can define for any $x \in \mathcal{M}$

$$\phi(x) = - \sum_{i=1}^m \log(d(x, \partial\mathcal{M}_i))$$

where $d(x, \partial\mathcal{M}_i)$ computes the minimum geodesic distance from x to the boundary $\partial\mathcal{M}$. In general this is a highly non-trivial optimization problem, contrary to the polytope case which admits a simple closed form.

While developed and most commonly used in optimisation, barrier methods can also be used for sampling (Lee & Vempala, 2017). The core idea of barrier methods is to ‘warp the geometry’ of the constrained space, stretching it as the process approaches the boundary so that it never hits it, hence bypassing the need to explicitly deal with the boundary. Assuming ϕ to be strictly convex and smooth, its Hessian $\nabla^2 \phi$ is positive definite and thus defines a valid Riemannian metric on \mathcal{M} . The formal approach to ‘warping the geometry’ of the convex space with the boundary is to endow \mathcal{M} with the Hessian as a Riemannian metric $\mathbf{g} = \nabla^2 \phi$, making it into a Hessian Manifold, see Shima & Yagi (1997). In the special case where the barrier is given by (5), we get for any $x \in \mathcal{M}$

$$\mathbf{g}(x) = A^\top S^{-2}(x) A \text{ with } S(x) = \text{diag}(b_i - \langle A_i, x \rangle)_i.$$

Equipped with this Riemannian metric, we consider the following Langevin dynamics as a forward process

$$d\mathbf{X}_t = \frac{1}{2}\text{div}(\mathbf{g}^{-1})(\mathbf{X}_t)dt + \mathbf{g}(\mathbf{X}_t)^{-\frac{1}{2}}d\mathbf{B}_t, \quad (6)$$

with $\text{div}(F)(x) \triangleq (\text{div}(F_1)(x), \dots, \text{div}(F_d)(x))^\top$, for any smooth $F: \mathbb{R}^d \rightarrow \mathbb{R}^d$. Under mild assumptions on \mathcal{M} , $\text{div}(\mathbf{g}^{-1})$ and \mathbf{g}^{-1} we get that $(\mathbf{X}_t)_{t \geq 0}$ is well-defined and for any $t \geq 0$, $\mathbf{X}_t \in \mathcal{M}$. In particular, for any $t \geq 0$, \mathbf{X}_t does not reach the boundary. This stochastic process was first proposed by Lee & Vempala (2017) in the context of efficient sampling from the uniform distribution over a polytope. Under similar conditions, \mathbf{X}_t admits a density p_t w.r.t. the Lebesgue measure and we have that $\partial_t p_t = \frac{1}{2}\text{Tr}(\mathbf{g}^{-1}\nabla^2 p_t)$. In addition, $(\mathbf{X}_t)_{t \geq 0}$ is irreducible. Hence, assuming that \mathcal{M} is compact, the uniform distribution on \mathcal{M} is the unique invariant measure of the process $(\mathbf{X}_t)_{t \geq 0}$ and $(\mathbf{X}_t)_{t \geq 0}$ converges to the uniform distribution in some sense. We refer the reader to Appendix C for a proof of these results.

Time-reversal. Assuming that \mathbf{g}^{-1} and its derivative are bounded on \mathcal{M} , the time-reversal of (6) is given by Cattiaux et al. (2021), in particular we have

$$\begin{aligned} d\overleftarrow{\mathbf{X}}_t &= [-\frac{1}{2}\text{div}(\mathbf{g}^{-1}) + \text{div}(\mathbf{g}^{-1}) + \mathbf{g}^{-1}\nabla \log p_{T-t}](\overleftarrow{\mathbf{X}}_t)dt + \mathbf{g}(\overleftarrow{\mathbf{X}}_t)^{-\frac{1}{2}}d\mathbf{B}_t, \\ &= [\frac{1}{2}\text{div}(\mathbf{g}^{-1}) + \mathbf{g}^{-1}\nabla \log p_{T-t}](\overleftarrow{\mathbf{X}}_t)dt + \mathbf{g}(\overleftarrow{\mathbf{X}}_t)^{-\frac{1}{2}}d\mathbf{B}_t. \end{aligned} \quad (7)$$

$\overleftarrow{\mathbf{X}}_0$ is initialised with the uniform distribution on \mathcal{M} (which is close to the one of \mathbf{X}_T for large T).

The estimation of the score term $\nabla \log p_t$ is done by minimising the ism loss function (4). We refer to Section 3.3 for details on the training and parameterisation.

Sampling. Sampling from the forward (6) and backward (7) processes, once the score is learnt, requires a discretisation scheme. We use Geodesic Random Walks (GRW) (Jørgensen, 1975) for this purpose, see Algorithm 1. This discretisation is a generalisation of the Euler-Maruyama discretisation of SDE in Euclidean spaces, where the + operator is replaced by the exponential mapping on the manifold, computing the geodesics.

Algorithm 1 *Geodesic Random Walk.* Discretisation of the SDE $d\mathbf{X}_t = d(t, \mathbf{X}_t)dt + d\mathbf{B}_t$.

Require: T (simulation time), N (number of steps), X_0^γ (initial point), d (drift function)

$\gamma = T/N$

for $k \in \{0, \dots, N-1\}$ **do**

$Z_{k+1} \sim \mathcal{N}(0, \text{Id})$

$W_{k+1} = \gamma d(k\gamma, X_k) + \sqrt{\gamma}Z_{k+1}$

$X_{k+1}^\gamma = \exp_{X_k}[W_{k+1}] \approx \text{proj}_{\mathcal{M}}(X_k + W_{k+1})$

return $\{X_k\}_{k=0}^N$

However, contrary to De Bortoli et al. (2022), we will not have access explicitly to the exponential mapping of the Hessian manifold. Instead, we rely on an approximation, using a *retraction* (see Absil & Malick (2012); Boumal (2023) for a definition and alternative schemes).

3.2 Reflected diffusion models

Another approach to deal with the geometry of \mathcal{M} is to use the standard metric \mathfrak{h} and forward dynamics of \mathcal{N} and constraining it to \mathcal{M} by *reflecting* the process whenever it would encounter a boundary. We will first assume that \mathcal{M} is compact and convex. To simplify the presentation, we focus on the Euclidean case $\mathcal{N} = \mathbb{R}^d$ with smooth boundary ∂M^4 .

The key difference between this approach and the barrier approach is that in the reflected case we leave the geometry unchanged, so all we need to do is show that the dynamics induced by reflecting the forward process whenever it hits the boundary leads to an invariant distribution and admits a time-reversal.

⁴We refer to Appendix G for a definition of smooth boundary.




Figure 5: *Left*: Convergence of the reflected Brownian motion on the unit interval to the uniform distribution. *Right*: Value of $|\mathbf{k}|_t$ for the trajectory samples on the left through time.

It is worth noting that while barrier approaches have received considerable theoretical attention in the sampling literature (Lee & Vempala, 2017; Noble et al., 2022), reflected methods have remained comparatively undeveloped from the methodological and practical point of view. In the next section, we recall the basics of reflected stochastic processes.

Skorokhod problem. The reflected Brownian motion is defined as the solution to the Skorokhod problem. Roughly speaking a solution to the Skorokhod problem consists of two coupled processes, $(\bar{\mathbf{B}}_t, \mathbf{k}_t)_{t \geq 0}$, such that $(\mathbf{B}_t)_{t \geq 0}$ acts *locally* as a Euclidean Brownian motion $(\mathbf{B}_t)_{t \geq 0}$ and \mathbf{k}_t compensates for the excursion of $(\mathbf{B}_t)_{t \geq 0}$ so that $(\bar{\mathbf{B}}_t)_{t \geq 0}$ remains in \mathcal{M} . We say that $(\bar{\mathbf{B}}_t, \mathbf{k}_t)_{t \geq 0}$ is a solution to the *Skorokhod problem* (Skorokhod, 1961) if $(\mathbf{k}_t)_{t \geq 0}$ and $(\bar{\mathbf{B}}_t)_{t \geq 0}$ are two processes satisfying mild conditions, see Appendix D for a rigorous introduction, such that for any $t \geq 0$,

$$\bar{\mathbf{B}}_t = \bar{\mathbf{B}}_0 + \mathbf{B}_t - \mathbf{k}_t \in \mathcal{M}, \quad (8)$$

and $|\mathbf{k}|_t = \int_0^t \mathbf{1}_{\bar{\mathbf{B}}_s \in \partial \mathcal{M}} d|\mathbf{k}|_s$, $\mathbf{k}_t = \int_0^t \mathbf{n}(\bar{\mathbf{B}}_s) d|\mathbf{k}|_s$, where $(|\mathbf{k}|_t)_{t \geq 0}$ is the total variation of $(\mathbf{k}_t)_{t \geq 0}$ and we recall that \mathbf{n} is the outward normal to \mathcal{M} ⁵. When $(\bar{\mathbf{B}}_t)_{t \geq 0}$ hits the boundary, the condition $\mathbf{k}_t = \int_0^t \mathbf{n}(\bar{\mathbf{B}}_s) d|\mathbf{k}|_s$, tells us that $-\mathbf{k}_t$ “compensates” for $\bar{\mathbf{B}}_t$ by pushing the process back into \mathcal{M} along the inward normal $-\mathbf{n}$, while the condition $|\mathbf{k}|_t = \int_0^t \mathbf{1}_{\bar{\mathbf{B}}_s \in \partial \mathcal{M}} d|\mathbf{k}|_s$ can be interpreted as \mathbf{k}_t being constant when $(\bar{\mathbf{B}}_t)_{t \geq 0}$ does not hit the boundary. As a result $(\bar{\mathbf{B}}_t)_{t \geq 0}$ can be understood as the continuous-time counterpart to the reflected Gaussian random walk. The process $(\mathbf{k}_t)_{t \geq 0}$ can be related to the notion of *local time* (Revuz & Yor, 2013) and quantifies the amount of time $(\bar{\mathbf{B}}_t)_{t \geq 0}$ spends at the boundary $\partial \mathcal{M}$. Lions & Sznitman (1984, Theorem 2.1) ensure the existence and uniqueness of a solution to the Skorokhod problem. One key observation is that the event $\{\bar{\mathbf{B}}_t \in \partial \mathcal{M}\}$ has probability zero (Harrison & Williams, 1987, Section 7, Lemma 7). As in the *unconstrained* setting, one can describe the dynamics of the density of $\bar{\mathbf{B}}_t$.

Proposition 3.1 (Burdzy et al. (2004)). *For any $t > 0$, the distribution of $\bar{\mathbf{B}}_t$ admits a density w.r.t. the Lebesgue measure denoted p_t . In addition, we have for any $x \in \text{int}(\mathcal{M})$ and $x_0 \in \partial \mathcal{M}$*

$$\partial_t p_t(x) = \frac{1}{2} \Delta p_t(x), \quad \partial_{\mathbf{n}} p_t(x_0) = 0, \quad (9)$$

where we recall that \mathbf{n} is the outward normal to \mathcal{M} .

Note that contrary to the unconstrained setting, the heat equation has *Neumann* boundary conditions. Similarly to the compact Riemannian setting (Saloff-Coste, 1994) it can be shown that the reflected Brownian motion converges to the uniform distribution on \mathcal{M} exponentially fast (Loper, 2020; Burdzy et al., 2006), see Section 3.2. Hence, $(\bar{\mathbf{B}}_t)_{t \geq 0}$ is a candidate for a forward noising process in the context of diffusion models.

Time-reversal. In order to extend the diffusion model approach to the reflected setting, we need to derive a *time-reversal* for $(\bar{\mathbf{B}}_t)_{t \in [0, T]}$. Namely, we need to characterise the evolution of $(\bar{\mathbf{X}}_t)_{t \in [0, T]} = (\bar{\mathbf{B}}_{T-t})_{t \in [0, T]}$. It can be shown that the time-reversal of $(\bar{\mathbf{B}}_t)_{t \in [0, T]}$ is also the solution to a Skorokhod problem.

⁵We extend the normal \mathbf{n} to the whole space by letting $\mathbf{n}(x) = 0$ if $x \notin \partial \mathcal{M}$.

Algorithm 2 *Reflected step algorithm.* The algorithm operates by repeatedly taking geodesic steps until one of the constraints is violated, or the step is fully taken. Upon hitting the boundary we parallel transport the tangent vector to the boundary and then reflect it against the boundary. We then start a new geodesic from this point in the new direction. The arg intersect_t function computes the distance one must travel along a geodesic in direction \mathbf{s} til constraint f_i is intersected. For a discussion of paralleltransport , exp_g and reflect please see Appendix A.

Input: $x \in \mathcal{M}$, $\mathbf{v} \in \mathbb{T}_x \mathcal{M}$, $\{f_i\}_{i \in \mathcal{I}}$

```

 $\ell \leftarrow \|\mathbf{v}\|_g$ 
 $\mathbf{s} \leftarrow \mathbf{v} / \|\mathbf{v}\|_g$ 
while  $\ell \geq 0$  do
   $d_i = \text{arg intersect}_t[\text{exp}_g(x, t\mathbf{s}), f_i]$ 
   $i \leftarrow \text{arg min}_i d_i \text{ s.t. } d_i > 0$ 
   $\alpha \leftarrow \min(d_i, \ell)$ 
   $x' \leftarrow \text{exp}_g(x, \alpha\mathbf{s})$ 
   $\mathbf{s} \leftarrow \text{paralleltransport}_g(x, \mathbf{s}, x')$ 
   $\mathbf{s} \leftarrow \text{reflect}(\mathbf{s}, f_i)$ 
   $\ell \leftarrow \ell - \alpha$ 
   $x \leftarrow x'$ 
return  $x$ 

```

Algorithm 3 *Reflected Random Walk.* Discretisation of the SDE $d\mathbf{X}_t = b(t, \mathbf{X}_t)dt + d\mathbf{B}_t - d\mathbf{k}_t$.

Require: T (simulation time), N (number of steps), X_0^γ (initial point), $\{f_i\}_{i \in \mathcal{I}}$ (boundary functions)

```

 $\gamma = T/N$ 
for  $k \in \{0, \dots, N-1\}$  do
   $Z_{k+1} \sim \mathcal{N}(0, \text{Id})$ 
   $X_{k+1}^\gamma = \text{ReflectedStep}[X_k^\gamma, \sqrt{\gamma}Z_{k+1}, \{f_i\}_{i \in \mathcal{I}}]$ 
return  $\{X_k^\gamma\}_{k=0}^N$ 

```

Theorem 3.2. *There exist $(\overleftarrow{\mathbf{k}}_t)_{t \geq 0}$ a bounded variation process and a Brownian motion $(\mathbf{B}_t)_{t \geq 0}$ such that*

$$\overleftarrow{\mathbf{X}}_t = \overleftarrow{\mathbf{X}}_0 + \mathbf{B}_t + \int_0^t \nabla \log p_{T-s}(\overleftarrow{\mathbf{X}}_s) ds - \overleftarrow{\mathbf{k}}_t.$$

In addition, for any $t \in [0, T]$ we have

$$|\overleftarrow{\mathbf{k}}|_t = \int_0^t \mathbf{1}_{\overleftarrow{\mathbf{X}}_s \in \partial \mathcal{M}} d|\overleftarrow{\mathbf{k}}|_s, \quad \overleftarrow{\mathbf{k}}_t = \int_0^t \mathbf{n}(\overleftarrow{\mathbf{X}}_s) d|\overleftarrow{\mathbf{k}}|_s.$$

The proof, see Appendix G, follows Petit (1997) which provides a time-reversal in the case where \mathcal{M} is the positive orthant. It is based on an extension of Haussmann & Pardoux (1986) to the reflected setting, with a careful handling of the boundary conditions. In particular, contrary to Petit (1997), we do not rely on an explicit expression of p_t but instead use the intrinsic properties of $(\mathbf{k}_t)_{t \geq 0}$. Informally, Theorem 3.2 means that the process $(\overleftarrow{\mathbf{X}}_t)_{t \in [0, T]}$ satisfies

$$d\overleftarrow{\mathbf{X}}_t = \nabla \log p_{T-t}(\overleftarrow{\mathbf{X}}_t) dt + d\mathbf{B}_t - d\overleftarrow{\mathbf{k}}_t, \quad (10)$$

which echoes the usual time-reversal formula (2). In practice, in order to sample from $(\overleftarrow{\mathbf{X}}_t)_{t \in [0, T]}$, one needs to consider the *reflected* version of the *unconstrained* dynamics $d\overleftarrow{\mathbf{X}}_t = \nabla \log p_{T-t}(\overleftarrow{\mathbf{X}}_t) dt + d\mathbf{B}_t$.

Sampling. In practice, we approximately sample the reflected dynamics by considering the Markov chain given by Algorithm 3. We refer to Pacchiarotti et al. (1998) and Bossy et al. (2004) for weak convergence results on this numerical scheme in the Euclidean setting.




Figure 6: Reflection against a linear boundary. For a step s with magnitude $|s|$ and direction \hat{s} the distance to the boundary described by the normal \hat{n} and offset b is $d = \frac{\langle \hat{s}, x_t \rangle - b}{\langle \hat{s}, \hat{n} \rangle}$. The reflected direction is given by $\hat{s}' = \hat{s} - 2\langle \hat{s}, \hat{n} \rangle \hat{n}$.




Figure 7: Reflection against a spherical boundary. For a sphere of radius r , the distance to the boundary from x_t in direction s is given by $d = \frac{1}{2}(\langle \hat{s}, x_t \rangle^2 + 4(r^2 - \|x_t\|^2))^{1/2} - \frac{1}{2}\langle \hat{s}, x_t \rangle$. The normal at the intersection can be computed as the unit vector in the direction $-2(d\hat{s} + x_t)$, and then \hat{s}' as above.

Likelihood evaluation. In the case of a reflected Brownian motion, it is possible to compute an equivalent ODE in order to perform likelihood evaluation. The associated ODE was first derived in Lou & Ermon (2023). The form of the ODE and the proof that it remains in \mathcal{M} are postponed to Appendix F.

3.3 Loss and score network parameterisation

In order to train log-barrier and reflected diffusion models we prove that we can use a *tractable* score matching loss in constrained manifolds. We will prove that the implicit score-matching loss leads to the recovery of the correct score when we have a boundary, so long as we enforce that the score is zero on the boundary (see Appendix E.1). This proof holds for both the log-barrier and the reflected process.

Proposition 3.3. *Let $s \in C^\infty([0, T] \times \mathbb{R}^d, \mathbb{R}^d)$ such that for any $x \in \partial\mathcal{M}$ and $t \geq 0$, $s_t(x) = 0$. Then, there exists $C > 0$ such that*

$$\mathbb{E}[\|\nabla \log p_t - s_t\|^2] = \mathbb{E}[\|s_t\|^2 + 2 \operatorname{div}(s_t)] + C,$$

where \mathbb{E} is taken over $\mathbf{X}_t \sim p_t$ and $t \sim \mathcal{U}([0, T])$.

This result immediately implies we can optimise the score network using the ism loss function so long as we enforce a Neumann boundary condition. The estimation of the score term $\nabla \log p_t$ is also done by minimising the ism loss function (4).

The score term $\nabla \log p_{T-t}$ appearing in both time-reversal processes (7) and (10) is intractable. It is thus approximated with a *score network* $s_\theta(t, \cdot) \approx \nabla \log p_t$. We use a multi-layer perceptron architecture with sin activation functions, see Appendix J for more details on the experimental setup. Due to boundary condition (9), the normal component of the score is zero at the boundary in the reflected case. A similar result holds in the log-barrier setting. This is additionally required for our proof of the ism loss.

Following Liu et al. (2022), we can accommodate this in our score parameterisation by additionally scaling the score output of the neural network by a monotone function $h(d(x, \partial\mathcal{M}))$ where d is the distance from x to the boundary and $h(0) = 0$. In particular, we use a clipped ReLU function: $s_\theta(t, x) = \min(1, \operatorname{ReLU}(d(x, \partial\mathcal{M}) - \delta)) \cdot \operatorname{NN}_\theta(t, x)$ with $\delta > 0$ a threshold so the model is forced to be zero “close” to the boundary as well as exactly on the boundary, see Appendix J for an illustration. The inclusion of this scaling function is necessary to produce reasonable results as we show in Appendix E.2. The weighting function in (4) is set to $\lambda_t = (t + 1)$.

The forward processes (6) and (8) for the barrier and reflected methods cannot be sampled in closed form, so at training time samples from the conditional marginals $p(\mathbf{X}_t | \mathbf{X}_0)$ are obtained by discretising these processes. As to take the most of this computational overhead, we use several samples from the discretised forward trajectory $(\mathbf{X}_{t_1}, \dots, \mathbf{X}_{t_k} | \mathbf{X}_0)$ instead of only using the last sample.

4 Related Work

Sampling on constrained manifolds. Sampling from a distribution on a space defined by a set of constraints is an important ingredient in several computational tasks, such as computing the volume of a polytope (Lee & Vempala, 2017). Incorporating such constraints within MCMC algorithms while preserving fast convergence properties is an active field of research (Kook et al., 2022; Lee & Vempala, 2017; Noble et al., 2022). In this work, we are interested in sampling from the uniform distribution defined on the constrained set in order to define a proper *forward process* for our diffusion model. Log-barrier methods such as the Dikin walk or Riemannian Hamiltonian Monte Carlo (Kannan & Narayanan, 2009; Lee & Vempala, 2017; Noble et al., 2022) change the geometry of the underlying space and define stochastic processes which never violate the constraints, see Kannan & Narayanan (2009) and Noble et al. (2022) for instance. If we keep the Euclidean metric, then *unconstrained* stochastic processes might not be well-defined for all times. To counter this effect, it has been proposed to *reflect* the Brownian motion (Williams, 1987; Petit, 1997; Shkolnikov & Karatzas, 2013). Finally, we also highlight hit-and-run approaches (Smith, 1984; Lovász & Vempala, 2006), which generalise Gibbs’ algorithm and enjoy fast convergence properties provided that one knows how to sample from the one-dimensional marginals.

Diffusion models on manifolds. De Bortoli et al. (2022) extended the work of Song et al. (2021) to Riemannian manifolds by defining forward and backward stochastic processes in this setting. Concurrently, a similar framework was introduced by Huang et al. (2022), extending the maximum likelihood approach of Huang et al. (2021). Existing applications of denoising diffusion models on Riemannian manifolds have been focused on well-known manifolds for which one can find metrics so that the framework of De Bortoli et al. (2022) applies. In particular, on compact Lie groups, geodesics and Brownian motions can be defined in a canonical manner. Their specific structure can be leveraged to define efficient diffusion models (Yim et al., 2023). Leach et al. (2022) define diffusion models on $SO(3)$ for rotational alignment, while Jing et al. (2022) consider the product of tori for molecular conformer generation. Corso et al. (2022) use diffusion models on $\mathbb{R}^3 \times SO(3) \times SO(2)$ for protein docking applications. RFDiffusion (Watson et al., 2022) and FrameDiff also incorporate $SE(3)$ diffusion models. Finally, Urain et al. (2022) introduce a methodology for $SE(3)$ diffusion models with applications to robotics.

Comparison with Lou & Ermon (2023). We now discuss a few key differences between our work and the reflected diffusion models presented in Lou & Ermon (2023). First, in the hypercube setting, our methodologies are identical from a theoretical viewpoint. The main difference is that Lou & Ermon (2023) use an approximated version of the DSM loss (3), whereas we rely on the ISM loss (4). Lou & Ermon (2023) exploit the specific factorised structure of the hypercube to make the DSM loss tractable, leading to significant practical advantages. These can however be directly employed in our framework for reflected models, and also in the log-barrier setting. Second, in our work, we target scientific applications where the underlying geometry is not Euclidean, whereas Lou & Ermon (2023) focus on the case where the constrained domain of interest is the hypercube (a subset of Euclidean space) with image applications, or where the domain can be easily projected into the hypercube, such as the simplex. Our approach is designed to handle a wider range of settings, such as non-convex polytopes in Euclidean space, or non-Euclidean geometries.

5 Experimental results

To demonstrate the practical utility of the constrained diffusion models introduced in Section 3, we evaluate them on a series of increasingly difficult synthetic tasks on convex polytopes, including the hypercube and the simplex, in Section 5.1. We then highlight their applicability to real-world settings by considering two problems from robotics and protein design. In particular, we show that our models are able to learn distributions over the space of $d \times d$ symmetric positive definite (SPD) matrices S_{++}^d under trace constraints in Section 5.2—a setting that is essential to describing and controlling the motions and exerted forces of robotic platforms (Jaquier et al., 2021). In Section 5.3, we use the parametrisation introduced in Han & Rudolph (2006) to map the problem of modelling the conformational ensembles of proteins under positional constraints on their endpoints to the product manifold of a convex polytope and a torus. The code is available [here](#). We refer the reader to Appendix J for more details on the experimental setup.

Table 1: MMD metrics between samples from synthetic distributions and trained constrained and unconstrained (Euclidean) diffusion models. Means and confidence intervals are computed over 5 different runs.

| Space | d | Log-barrier | | Reflected | | Euclidean | |
|-------------|-----|-------------|--------------------|-----------|--------------------|-----------|--------------------|
| | | MMD | % in \mathcal{M} | MMD | % in \mathcal{M} | MMD | % in \mathcal{M} |
| $[-1, 1]^d$ | 2 | .066±.006 | 100.0 | .055±.015 | 100.0 | .062±.011 | 98.8 |
| | 3 | .209±.077 | 100.0 | .080±.004 | 100.0 | .076±.004 | 98.5 |
| | 10 | .330±.004 | 100.0 | .313±.048 | 100.0 | .081±.005 | 96.4 |
| Δ^d | 2 | .050±.012 | 100.0 | .043±.002 | 100.0 | .055±.013 | 96.4 |
| | 3 | .238±.009 | 100.0 | .181±.007 | 100.0 | .068±.014 | 96.3 |
| | 10 | .275±.015 | 100.0 | .290±.009 | 100.0 | .060±.003 | 92.6 |




Figure 8: Histograms of samples from the data distribution and from trained constrained diffusion models.

5.1 Method characterisation on convex polytopes

First, we aim to assess the empirical performance of our methods on constrained manifolds of increasing dimensionality. For this, we focus on polytopes and construct synthetic datasets that represent bimodal distributions. In particular, we investigate two specific instances of polytopes: hypercubes and simplices. In Appendix J.1 we also present results on the Birkhoff polytope. We quantify the performance of each model via the Maximum Mean Discrepancy (MMD) (Gretton et al., 2012), which is a kernel-based metric between distributions. We present a qualitative comparison of the logarithmic barrier and reflected Brownian motion models in Figure 8, and observe that both methods recover the data distribution on the two-dimensional hypercube and simplex, although the reflected method seems to produce a better fit. In Table 1, we report the MMD between the data distribution and the learnt diffusion models. Here, we similarly observe that the reflected method consistently yields better results than the log-barrier one.

We additionally compare both of these models to a set of unconstrained Euclidean diffusion models, noting that they are outperformed by the constrained models in lower dimensions, but generate better results in higher dimensions. There are a number of potential explanations for this: First, we note that the mixture of Normal distributions we use as a synthetic data-generating process places significantly less probability mass near the boundary as its dimensionality increases, more closely resembling an unconstrained mixture distribution that is easier for the Euclidean diffusion models to learn, while posing a challenge to the log-barrier and reflected diffusion models that initialise at the uniform distribution within the constraints. Additionally, we note that the design space and hyperparameters used for all experiments were informed by best practices for Euclidean models that may be suboptimal for the more complex dynamics of constrained diffusion models.

5.2 Modelling robotic arms under force and velocity constraints

Accurately determining and controlling the movement and exerted forces of robotic platforms is a fundamental problem in many real-world robotics applications. A kinetostatic descriptor that is commonly used to quantify the ability of a robotic arm to move and apply forces along certain dimensions is the so-called manipulability ellipsoid $E \in \mathbb{R}^d$ (Yoshikawa, 1985), which is naturally described as a symmetric positive definite (SPD)

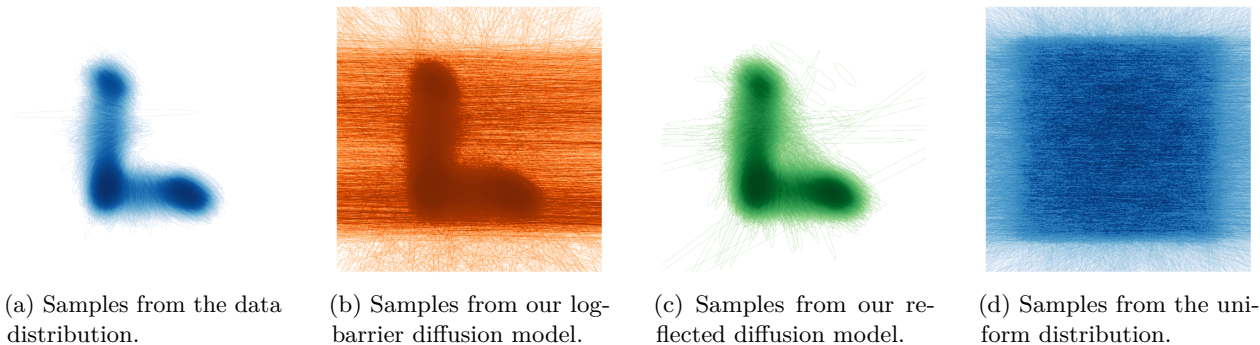


Figure 9: Samples in $S_{++}^2 \times \mathbb{R}^2$ from (a) the data distribution, (b) our log-barrier diffusion model, (c) our reflected diffusion model and (d) the uniform distribution. Each sample is visualised as the manipulability ellipsoid encoded by the SPD matrix $M \in S_{++}^2$ placed at the corresponding location in \mathbb{R}^2 . Additional results and full correlation plots are postponed to Appendix J.2.

matrix $M \in \mathbb{R}^{d \times d}$ (Jaquier et al., 2021). The manifold of such $d \times d$ SPD matrices, denoted as S_{++}^d , is defined as the set of matrices $\{x^\top M x \geq 0, x \in \mathbb{R}^d : M \in \mathbb{R}^{d \times d}\}$. In many practical settings, it may be desirable to constrain the volume of E to retain flexibility or limit the magnitude of an exerted force, which can be expressed as an upper bound on the trace of M , i.e. as the inequality constraint $\sum_{i=1}^d M_{ii} < C$ with $C > 0$. Constraining the rest of the entries of the matrix to ensure it is SPD is non-trivial. Alternatively, we can parameterise the SPD matrices via their Cholesky decomposition. Each SPD matrix has a unique decomposition of the form $M = LL^\top$, where L is a lower triangular matrix with strictly positive diagonal (Golub & Van Loan, 2013, p.143). Constraining the entries of this matrix simply requires ensuring the diagonal is positive. The trace of the SPD matrix is given by $\text{Tr}(M) = \text{Tr}(LL^\top) = \sum_{i,j} L_{i,j}^2$, and results in the constraint on the entries of L to live in a ball of radius C . We additionally model the two-dimensional position of the arm. In summary, the space over which we parameterise the diffusion models is defined as $\{L \in \mathbb{R}^{d(d+1)/2} : L_{i,i} > 0, \sum_{i,j} L_{i,j}^2 < C\} \times \mathbb{R}^2$. Under the Euclidean metric, we can apply both our log-barrier and reflected approaches. The positive diagonal (linear) constraint is handled similarly to the polytope setting. The reflection on the ball boundary is defined and illustrated on Figures 6 and 7.

Using this framework, we model the datasets presented in Jaquier et al. (2021) (see Appendix H for full experimental details). The joint distribution over the SPD matrices (represented as ellipsoids) and their positions is presented in Figure 9. We qualitatively observe that the reflected method is able to model the joint data distribution better than the log-barrier one. This is reflected by an MMD of 0.161 and 0.247, respectively.

5.3 Modelling protein loops with anchored endpoints

Modelling the conformational ensembles of proteins is an important task in the field of molecular biology, particularly in the context of bioengineering and drug discovery. In many data-scarce practical settings such as antibody or enzyme design, it is often unnecessary or even undesirable to model the structure of an entire protein, as researchers are primarily interested in specific functional sites with distinct biochemical properties. However, generating conformational ensembles for such substructural elements necessitates positional constraints on their endpoints to ensure that they can be accommodated by the remaining scaffold.


This problem can be reformulated as modelling a spatial chain with spherical joints and fixed end points. Following the framework outlined in Han & Rudolph (2006), we parametrise the conformations of such chains with d fixed-length links and arbitrary end-point constraints as the product of a convex polytope \mathbb{P} and torus \mathbb{T} : $\mathbb{P}^{(d-3)} \times \mathbb{T}^{(d-2)}$. The essential idea of this parameterisation is to fix one end point as an ‘‘anchor’’ and model the chain as the series of triangles formed by the anchor and each pair of adjacent joints in the chain. A point in the polytope corresponds to the lengths of the diagonals of these triangles, and a point in the torus to the angles between each pair of subsequent triangles. See Figure 10 for an illustration and Appendix I.1 for a full description.



(a) An illustrative diagram of the parameterisation used for the conformational modelling of the C_α trace of a cyclic peptide, introduced in Han & Rudolph (2006).

(b) The convex polytope constraining the diagonals of the triangles for the given bond lengths in the illustrated molecule. The total design space is the product of this polytope with the 4D flat torus.

Figure 10: Illustrations of the parameterisation used to model distribution of polypeptide backbone conformations under anchor point distance constraints as the product of a convex polytope \mathbb{P} and torus \mathbb{T} : $\mathbb{P}^3 \times \mathbb{T}^4$.



(a) Samples from the data distribution.

(b) Samples from our log-barrier diffusion model.

(c) Samples from our reflected diffusion model.

(d) Samples from the uniform distribution.

Figure 11: Planar projection of the modelled C_α chains from (a) the training dataset, (b) our log-barrier diffusion model, (c) our reflected diffusion model and (d) the uniform distribution. Additional results and full correlation plots are postponed to Appendix J.3.

Using this framework, we model the conformational landscape of the cyclic peptide c-AAGAGG, consisting of a circular polypeptide chain with coinciding endpoints. We generate 10^6 backbone conformations using tools from molecular dynamics (Eastman et al., 2017; Hornak et al., 2006) and divide them into training and evaluation datasets, (see Appendix I.2 for full experimental details). Drawing on the definition above, the space on which we effectively parameterise our constrained diffusion models for a circular polypeptide chain of length $d = 6$ is given by the product manifold $\mathbb{P}^3 \times \mathbb{T}^4$.

To learn a distribution over this space, we leverage the methodology introduced in Section 3 for the polytope component \mathbb{P}^3 and in De Bortoli et al. (2022) for the torus component \mathbb{T}^4 . A qualitative comparison of samples from the data distribution, our log-barrier and reflected diffusion models, and the uniform distribution is presented in Figures 11a to 11d. For enhanced visual clarity, we project the modelled spatial chain onto the 2D plane by removing the (unconstrained) torus component of the product manifold and only plotting the planar chains encoded by the (constrained) polytope component (a correlation plot of the full product manifold is presented in Figure 25).

It is apparent that the data distribution is highly multimodal, encompassing a large number of locally optimal conformational clusters. Nevertheless, our reflected and log-barrier diffusion models are able to robustly approximate this energetic landscape, producing samples that reflect key conformational states and producing comparable MMD metrics of 0.032 ± 0.021 and 0.032 ± 0.001 , respectively. As a point of comparison, the uniform distribution on the polytope-torus product has an MMD of 0.112 ± 0.001 .

6 Discussion

Learning complex distributions whose support is confined to constrained spaces is a crucial task with many applications in the natural and engineering sciences, including computational statistics (Morris, 2002), robotics (Han & Rudolph, 2006), quantum physics (Lukens et al., 2020) and computational biology (Thiele et al., 2013). In this work we extend continuous diffusion models to this setting, proposing two complementary approaches—one based on log-barrier methods and the other on the reflected Brownian motion. For both methods, we derive the time-reversal formula, propose discretisation schemes and extend the score-matching toolbox. We demonstrate the utility of our methods on a range of synthetic and real-world tasks, including the constrained conformational modelling of proteins and robotic arms, and find that reflected methods, while enjoying fewer theoretical guarantees than their log-barrier counterparts, often yield preferable results.

We conclude by highlighting important directions of future research. First, the computational cost of performing the reflection when discretising the reflected Brownian motion is high. Finding numerically efficient approximations of the reflected process is therefore necessary to extend this methodology to very high dimensional settings. Second, the retraction used in place of the exponential map for the barrier method leads to a high number of discretisation steps to ensure a good approximation. Designing a faster forward process for the log-barrier method is key to targeting more complex distributions.

References

- P-A Absil and Jérôme Malick. Projection-like retractions on matrix manifolds. *SIAM Journal on Optimization*, 22(1):135–158, 2012.
- RN Bhattacharya. Criteria for recurrence and existence of invariant measures for multidimensional diffusions. *The Annals of Probability*, pp. 541–553, 1978.
- Mireille Bossy, Emmanuel Gobet, and Denis Talay. A symmetrized euler scheme for an efficient approximation of reflected diffusions. *Journal of applied probability*, 41(3):877–889, 2004.
- Nicolas Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.
- Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Krzysztof Burdzy, Zhen-Qing Chen, and John Sylvester. The heat equation and reflected brownian motion in time-dependent domains. *The Annals of Probability*, 32(1B):775–804, 2004.
- Krzysztof Burdzy, Zhen-Qing Chen, and Donald E Marshall. Traps for reflected brownian motion. *Mathematische Zeitschrift*, 252:103–132, 2006.
- Patrick Cattiaux, Giovanni Conforti, Ivan Gentil, and Christian Léonard. Time reversal of diffusion processes under a finite entropy condition. *arXiv preprint arXiv:2104.07708*, 2021.
- Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. DiffDock: Diffusion Steps, Twists, and Turns for Molecular Docking, October 2022. URL <http://arxiv.org/abs/2210.01776>.
- Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modeling, 2022.
- Patrick G Dougherty, Ashweta Sahni, and Dehua Pei. Understanding cell penetration of cyclic peptides. *Chemical Reviews*, 119(17):10241–10287, 2019.
- Peter Eastman, Jason Swails, John D Chodera, Robert T McGibbon, Yutong Zhao, Kyle A Beauchamp, Lee-Ping Wang, Andrew C Simmonett, Matthew P Harrigan, Chaya D Stern, et al. Openmm 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS computational biology*, 13(7): e1005659, 2017.
- Lawrence Craig Evans and Ronald F Gariepy. *Measure theory and fine properties of functions*. CRC press, 2015.
- Khashayar Gatmiry and Santosh S Vempala. Convergence of the riemannian langevin algorithm. *arXiv preprint arXiv:2204.10818*, 2022.
- Gene H Golub and Charles F Van Loan. *Matrix computations*. JHU press, 2013.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(null):723–773, March 2012. ISSN 1532-4435.
- Li Han and Lee Rudolph. Inverse kinematics for a serial chain with joints under distance constraints. In *Robotics: Science and systems*, 2006.
- J Michael Harrison and Ruth J Williams. Brownian models of open queueing networks with homogeneous customer populations. *Stochastics: An International Journal of Probability and Stochastic Processes*, 22(2): 77–115, 1987.

- Ulrich G Haussmann and Etienne Pardoux. Time reversal of diffusions. *The Annals of Probability*, pp. 1188–1205, 1986.
- Laurent Heirendt, Sylvain Arreckx, Thomas Pfau, Sebastián N Mendoza, Anne Richelle, Almut Heinken, Hulda S Haraldsdóttir, Jacek Wachowiak, Sarah M Keating, Vanja Vlasov, et al. Creation and analysis of biochemical constraint-based models using the cobra toolbox v. 3.0. *Nature protocols*, 14(3):639–702, 2019.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 2020.
- Viktor Hornak, Robert Abel, Asim Okur, Bentley Strockbine, Adrian Roitberg, and Carlos Simmerling. Comparison of multiple amber force fields and development of improved protein backbone parameters. *Proteins: Structure, Function, and Bioinformatics*, 65(3):712–725, 2006.
- Elton P Hsu. *Stochastic analysis on manifolds*. Number 38. American Mathematical Soc., 2002.
- Chin-Wei Huang, Jae Hyun Lim, and Aaron C Courville. A variational perspective on diffusion-based generative models and score matching. *Advances in Neural Information Processing Systems*, 34:22863–22876, 2021.
- Chin-Wei Huang, Milad Aghajohari, Avishek Joey Bose, Prakash Panangaden, and Aaron Courville. Riemannian Diffusion Models, August 2022.
- Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Nobuyuki Ikeda and Shinzo Watanabe. *Stochastic differential equations and diffusion processes*. Elsevier, 2014.
- Noémie Jaquier, Leonel Rozo, Darwin G Caldwell, and Sylvain Calinon. Geometry-aware manipulability learning, tracking, and transfer. *The International Journal of Robotics Research*, 40(2-3):624–650, 2021.
- Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Torsional Diffusion for Molecular Conformer Generation, June 2022. URL <http://arxiv.org/abs/2206.01729>.
- Erik Jørgensen. The central limit problem for geodesic random walks. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 32(1-2):1–64, 1975.
- Weining Kang and Kavita Ramanan. On the submartingale problem for reflected diffusions in domains with piecewise smooth boundaries. 2017.
- Ravi Kannan and Hariharan Narayanan. Random walks on polytopes and an affine interior point method for linear programming. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, pp. 561–570, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585062. doi: 10.1145/1536414.1536491. URL <https://doi.org/10.1145/1536414.1536491>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. URL <http://arxiv.org/abs/1412.6980>.
- Yunbum Kook, Yin Tat Lee, Ruoqi Shen, and Santosh S Vempala. Sampling with riemannian hamiltonian monte carlo in a constrained space. *arXiv preprint arXiv:2202.01908*, 2022.
- Thomas J Lane. Protein structure prediction has reached the single-structure frontier. *Nature Methods*, pp. 1–4, 2023.
- Adam Leach, Sebastian M Schmon, Matteo T Degiacomi, and Chris G Willcocks. Denoising Diffusion Probabilistic Models on $SO(3)$ for Rotational Alignment. *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*, pp. 8, 2022.
- John M Lee. *Riemannian manifolds: an introduction to curvature*, volume 176. Springer Science & Business Media, 2006.

- John M Lee. Smooth manifolds. In *Introduction to Smooth Manifolds*, pp. 1–31. Springer, 2013.
- Yin Tat Lee and Santosh S. Vempala. Geodesic walks in polytopes. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 927–940, Montreal Canada, June 2017. ACM. ISBN 978-1-4503-4528-6. doi: 10.1145/3055399.3055416.
- Yin Tat Lee and Santosh S Vempala. Convergence rate of riemannian hamiltonian monte carlo and faster polytope volume computation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1115–1121, 2018.
- Pierre-Louis Lions and Alain-Sol Sznitman. Stochastic differential equations with reflecting boundary conditions. *Communications on pure and applied Mathematics*, 37(4):511–537, 1984.
- Song Liu, Takafumi Kanamori, and Daniel J Williams. Estimating density models with truncation boundaries using score matching. *Journal of Machine Learning Research*, 23(186):1–38, 2022.
- Jackson Loper. Uniform ergodicity for brownian motion in a bounded convex set. *Journal of Theoretical Probability*, 33(1):22–35, 2020.
- Aaron Lou and Stefano Ermon. Reflected diffusion models. *arXiv preprint arXiv:2304.04740*, 2023.
- László Lovász and Santosh Vempala. Hit-and-run from a corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006. doi: 10.1137/S009753970544727X.
- Joseph M Lukens, Kody JH Law, Ajay Jasra, and Pavel Lougovski. A practical and efficient approach for bayesian quantum state estimation. *New Journal of Physics*, 22(6):063038, 2020.
- Sean P Meyn and Richard L Tweedie. Stability of markovian processes ii: Continuous-time processes and sampled chains. *Advances in Applied Probability*, 25(3):487–517, 1993.
- Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou, Yann Thanwerdas, Stefan Heyder, Olivier Peltre, Niklas Koep, Hadi Zaatiti, Hatem Hajri, Yann Cabanes, Thomas Gerald, Paul Chauchat, Christian Shewmake, Daniel Brooks, Bernhard Kainz, Claire Donnat, Susan Holmes, and Xavier Pennec. Geomstats: A python package for riemannian geometry in machine learning. *Journal of Machine Learning Research*, 21(223):1–9, 2020. URL <http://jmlr.org/papers/v21/19-027.html>.
- Ben J. Morris. Improved bounds for sampling contingency tables. *Random Structures & Algorithms*, 21(2):135–146, 2002. doi: <https://doi.org/10.1002/rsa.10049>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/rsa.10049>.
- Yurii Nesterov et al. *Lectures on convex optimization*, volume 137. Springer, 2018.
- Maxence Noble, Valentin De Bortoli, and Alain Durmus. Barrier Hamiltonian Monte Carlo, October 2022.
- David Nualart. *The Malliavin calculus and related topics*, volume 1995. Springer, 2006.
- Barbara Pacchiarotti, Cristina Costantini, and Flavio Sartoretto. Numerical approximation for functionals of reflecting diffusion processes. *SIAM Journal on Applied Mathematics*, 58(1):73–102, 1998.
- Frédérique Petit. Time reversal and reflected diffusions. *Stochastic Processes and their Applications*, 69(1):25–53, July 1997. ISSN 03044149. doi: 10.1016/S0304-4149(97)00035-5.
- Daniel Revuz and Marc Yor. *Continuous martingales and Brownian motion*, volume 293. Springer Science & Business Media, 2013.
- Piotr Rotkiewicz and Jeffrey Skolnick. Fast procedure for reconstruction of full-atom protein models from reduced representations. *Journal of computational chemistry*, 29(9):1460–1465, 2008.
- Laurent Saloff-Coste. Precise estimates on the rate at which certain diffusions tend to equilibrium. *Mathematische Zeitschrift*, 217(1):641–677, 1994.

- Hirohiko Shima and Katsumi Yagi. Geometry of hessian manifolds. *Differential geometry and its applications*, 7(3):277–290, 1997.
- Mykhaylo Shkolnikov and Ioannis Karatzas. Time-reversal of reflected Brownian motions in the orthant, July 2013. URL <http://arxiv.org/abs/1307.4422>.
- Anatoliy V Skorokhod. Stochastic equations for diffusion processes in a bounded region. *Theory of Probability & Its Applications*, 6(3):264–274, 1961.
- Robert L. Smith. Efficient Monte Carlo Procedures for Generating Points Uniformly Distributed over Bounded Regions. *Operations Research*, 32(6):1296–1308, December 1984. doi: 10.1287/opre.32.6.1296. URL <https://ideas.repec.org/a/inm/oropre/v32y1984i6p1296-1308.html>.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, 2019.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- Ines Thiele, Neil Swainston, Ronan M T Fleming, Andreas Hoppe, Swagatika Sahoo, Maiké K Aurich, Hulda Haraldsdóttir, Monica L Mo, Ottar Rolfsson, Miranda D Stobbe, Stefan G Thorleifsson, Rasmus Agren, Christian Bölling, Sergio Bordel, Arvind K Chavali, Paul Dobson, Warwick B Dunn, Lukas Endler, David Hala, Michael Hucka, Duncan Hull, Daniel Jameson, Neema Jamshidi, Jon J Jonsson, Nick Juty, Sarah Keating, Intawat Nookaew, Nicolas Le Novère, Naglis Malys, Alexander Mazein, Jason A Papin, Nathan D Price, Evgeni Selkov, Martin I Sigurdsson, Evangelos Simeonidis, Nikolaus Sonnenschein, Kieran Smallbone, Anatoly Sorokin, Johannes H G M van Beek, Dieter Weichart, Igor Goryanin, Jens Nielsen, Hans V Westerhoff, Douglas B Kell, Pedro Mendes, and Bernhard Ø Palsson. A community-driven global reconstruction of human metabolism. *Nature Biotechnology*, March 2013. ISSN 1087-0156. doi: 10.1038/nbt.2488.
- Brian L. Trippe, Jason Yim, Doug Tischer, Tamara Broderick, David Baker, Regina Barzilay, and Tommi Jaakkola. Diffusion probabilistic modeling of protein backbones in 3D for the motif-scaffolding problem, June 2022.
- Julen Urain, Niklas Funk, Georgia Chalvatzaki, and Jan Peters. Se (3)-diffusionfields: Learning cost functions for joint grasp and motion optimization through diffusion. *arXiv preprint arXiv:2209.03855*, 2022.
- Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E. Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, Basile I. M. Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jue Wang, Preetham Venkatesh, Isaac Sappington, Susana Vázquez Torres, Anna Lauko, Valentin De Bortoli, Emile Mathieu, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker. Broadly applicable and accurate protein design by integrating structure prediction networks and diffusion generative models, December 2022.
- R. J. Williams. Reflected Brownian motion with skew symmetric data in a polyhedral domain. *Probability Theory and Related Fields*, 75(4):459–485, August 1987. ISSN 0178-8051, 1432-2064. doi: 10.1007/BF00320328.
- Kevin E Wu, Kevin K Yang, Rianne van den Berg, James Y Zou, Alex X Lu, and Ava P Amini. Protein structure generation via folding diffusion. *arXiv preprint arXiv:2209.15611*, 2022.

Jason Yim, Brian L Trippe, Valentin De Bortoli, Emile Mathieu, Arnaud Doucet, Regina Barzilay, and Tommi Jaakkola. Se (3) diffusion model with application to protein backbone generation. *arXiv preprint arXiv:2302.02277*, 2023.

Tsuneo Yoshikawa. Manipulability of robotic mechanisms. *The international journal of Robotics Research*, 4(2):3–9, 1985.