Beyond Fixed Horizons: A Theoretical Framework for Adaptive Denoising Diffusions

Anonymous authors
Paper under double-blind review

Abstract

We introduce a new class of generative diffusion models that, unlike conventional denoising diffusion models, achieve a time-homogeneous structure for both the noising and denoising processes, allowing the number of steps to adaptively adjust based on the noise level. This is accomplished by conditioning the forward process using Doob's h-transform, which terminates the process at a suitable sampling distribution at a random time. The model is particularly well suited for generating data with lower intrinsic dimensions, as the termination criterion simplifies to a first-hitting rule. A key feature of the model is its adaptability to the target data, enabling a variety of downstream tasks using a pre-trained unconditional generative model. These tasks include natural conditioning through appropriate initialisation of the denoising process and classification of noisy data.

1 Introduction

Denoising Diffusion Models (Ho et al., 2020; Song et al., 2021) have gained significant attention in recent years due to their ability to generate high-quality data samples by iteratively denoising simple distributions based on the learned dynamics of a time-reversed forward noising process initalised in the target data distribution (Niu et al., 2020; Dhariwal & Nichol, 2021; Ho et al., 2022; Watson et al., 2023; Yang et al., 2023; Evans et al., 2024). A key limitation of these models, however, is their reliance on a fixed time horizon, which introduces an artificial time dependency in the drift function of the backward process. As a result, the generative denoising process follows a predefined number of steps, regardless of the actual level of noise present along the generated path.

To overcome this limitation, we introduce a novel class of diffusion models that dynamically adapt to the state of the denoising process. By replacing the fixed deterministic time horizon with a random one and conditioning the behaviour of the forward process at a terminal time, our approach achieves greater flexibility and state awareness. The foundation of our method lies in Doob's h-transforms with respect to underlying exponential times. While the theoretical groundwork for this concept exists, its explicit application and detailed exploration – particularly in comparison to deterministic time horizons – remains underrepresented in the literature.

A key feature of our model is its inherent adaptability: the number of denoising steps dynamically adjusts based on the noise level in the data, introducing a stochastic element. This randomness not only has the potential to enhance the generation process, but also allows denoising to start from partially noisy data, naturally incorporating conditioning. Moreover, the time required for denoising serves as an intuitive measure of the distance between noisy observations and the underlying data distribution, providing the basis for tasks such as classification and anomaly detection. The model's architecture also supports natural conditioning mechanisms, allowing seamless adaptation to diverse tasks without the need for task-specific design modifications.

Thanks to its flexible, time-homogeneous structure, our model offers a fresh perspective on generative tasks that enhances the adaptability and versatility of diffusion models and establishes a robust foundation for transfer learning.

Contributions and structure Let us briefly summarize our main contributions and the structure of the paper: after discussing related work in Section 2, we first present theoretical results on the h-transform and time reversal for exponential time horizons in Section 3. Building on this foundation, we develop a universal and flexible diffusion model, accompanied by the associated learning theory, in Section 4. In this framework, we identify the *polarity* of the data distribution as a central assumption for successful learning, offering a new perspective on the manifold hypothesis. Finally, we discuss the adaptability of the model across various application domains through transfer learning in Section 5.

2 Related work

This section reviews key works relevant to extending diffusion models beyond conventional time dependencies, incorporating elements such as random horizons, Doob's h-transform, and conditioning.

Ye et al. (2022) introduce first hitting diffusion models, which use first hitting times to capture the intrinsic geometry of data manifolds. Technically, their approach is related to ours in that it uses a random time horizon (in their case via first hitting times) and makes use of Doob's h-transform. However, their method defines a backward process that ensures that the generated data lies on a predetermined manifold. In contrast, our approach inverts this perspective and uses the forward process to create a more flexible generative framework that is not constrained to predefined manifolds. Instead, the data manifold is dynamically learned.

De Bortoli et al. (2021) explore diffusion generative modeling through the lens of Schrödinger Bridges (SB) and solving transport problems. Unlike traditional methods that require running forward SDEs over long durations, SB techniques generate samples in finite time. However, their approach is limited to a deterministic time horizon, where there are well-established connections between Schrödinger Bridges and h-transforms. Further developments in this direction are also presented in Peluchetti (2023) and Shi et al. (2023). In contrast, our framework accommodates random time horizons, allowing for both finite and long durations.

The paper Zhao et al. (2024) provides a comprehensive review of approaches to conditional sampling within generative diffusion models. It discusses methods that rely on joint distributions or pre-trained marginal distributions with explicit likelihoods to generate samples conditioned on certain information, addressing challenges in areas such as Bayesian inverse problems. In these approaches, the original unconditional processes are modified in various ways to introduce conditionality. Similarly, Didi et al. (2023) unifies conditional training and sampling within a common framework based on the h-transform. However, the reliance on a fixed time horizon leads to notable differences from the approach presented in our work.

Berner et al. (2024) introduce an optimal control perspective on denoising diffusion models, demonstrating how the backward generative process naturally emerges as a solution to a suitable optimal control problem. Proposition 2 establishes a similar connection in our model.

Denker et al. (2024) use h-transform techniques to introduce fine-tuning of a pre-trained diffusion model for conditional data generation. Additionally to the natural conditioning aspect that our model provides without second stage training, we demonstrate in Section 5.2 how an analogous procedure can be implemented for structured conditioning purposes in our modeling pipeline.

Finally, Christensen et al. (2025) builds on this paper and provides a detailed high-dimensional mathematical analysis of the natural conditioning property of time-reversals of exponentially killed Brownian motions as a particular sub-class of the general class of generative models that we introduce. This highlights the immediate usefulness of our model beyond pure unconditional data generation purposes.

3 Doob's h-transform and time-inversion from random times

Doob's h-transform is a versatile mathematical technique for adjusting the dynamics of a stochastic process. It formalizes the concept of conditioning the process on specific events occurring at a random time. This section outlines the key results relevant to our approach. For a more comprehensive discussion, including relevant literature and the proofs, readers are encouraged to consult the Appendix.

3.1 Doob's h-transform for random time horizons

In our approach, we will employ h-transforms of an underlying non-degenerate and symmetric d-dimensional diffusion process

$$dZ_t = b(Z_t)\,dt + \sigma(Z_t)\,dW_t,$$

to describe both the forward and backward process. Typically, Z will be a Brownian motion or a symmetric Ornstein–Uhlenbeck process on $\mathcal{I} = \mathbb{R}^d$, but other self-dual processes are also possible. The unkilled version of Z has the r-Green kernel

$$G_r(x,y) := \int_0^\infty \mathrm{e}^{-rt} p_t(x,y)\,dt,$$

where $p_t(x,y)$ are the symmetric transition densities of Z with respect to a reference measure m. The reference measure m may, but does not necessarily, coincide with the standard Lebesgue measure. Up to scaling, G_r describes the pdf of Z at an independent exponential time. For a Brownian motion, for example, G_r can be found explicitly, see Remark 9.

To specify the idea of conditioning at a random time, we introduce two major modifications for Z:

1. The drift is modified to attract the process towards the desired states.
2. A random stopping time ζ, referred to as the *lifetime*, is introduced to terminate the process at the point where the desired distribution is attained.

Both modifications are seamlessly implemented using the h-transform. The h-transform is constructed using an r-excessive function h, which can be defined via a probability measure β on the state space $\mathcal{I}$ and a reference state $x_0 \in \mathcal{I}$. To condition the process Z to be distributed according to β at its killing time ζ when initiated from x_0, we define

$$h(x) := \int \frac{G_r(x,y)}{G_r(x_0,y)}\,\beta(dy) = \int G_r(x,y)\,\kappa(dy),$$

where $\kappa(dy) = \kappa_{x_0,\beta}(dy) = \frac{1}{G_r(x_0,y)}\beta(dy)$ is called *representing measure* of h. Based on this, the h-transformed process Z^h is now a Markov process defined by the transition kernel

$$\mathbb{P}_x(Z_t^h \in dy) = \mathbb{E}_x\left[\frac{\mathrm{e}^{-rt}h(Z_t)}{h(x)}\mathbf{1}_{\{Z_t\in dy\}}\right].$$

The essential properties of Z^h are as follows:

Proposition 1. *1. Z^h is killed at some random time ζ, the lifetime, which is an a.s. finite random variable.*

2. Outside the support of β, Z^h is an Itô diffusion with dynamics

$$dZ_t^h = b^h(Z_t^h)\,dt + \sigma(Z_t^h)\,dW_t,$$
$$b^h(y) = b(y) + \sigma(y)\sigma(y)^\top \nabla \log h(y).$$

3. The distribution of Z^h at its lifetime is supported in the support of β and given by

$$\mathbb{P}_x(Z_{\zeta-}^h \in dy) = \frac{G_r(x,y)}{h(x)}\kappa(dy).$$

In particular,

$$\mathbb{P}_{x_0}(Z_{\zeta-}^h \in dy) = \beta(dy).$$

The proof is deferred to Appendix A.

3.2 Optimality properties of the h-transform

The h-transform is characterised by inherent optimality properties which underline its role as a canonical approach. We now demonstrate that it emerges as the solution to a stochastic stopping and control problem. In this context, the running cost component can be interpreted as minimizing the expected lifetime of the process while penalizing large drifts. By framing the h-transform within this stochastic control problem, we establish a connection to the KL divergence.

Proposition 2. *We write $k(u) = r + \frac{1}{2}\|u\|^2$, $g(y) = -\log h(y)$. For an (admissible) control u consider the controlled process*

$$dZ_t^u = (b(Z_t^u) + \sigma(Z_t^u)u_t)\,dt + \sigma(Z_t^u)\,dW_t,\ t < \tau_\beta,$$

where $\tau_\beta = \tau_\beta^u$ denotes the first entrance time into the support of β. The stochastic stopping and control problem of minimizing

$$J(u,\tau,x) := \mathbb{E}_x\left[\int_0^\tau k(u_t)\,dt + g(Z_{\tau-}^u)\right], \qquad v(x) = \inf_{u,\ \tau\le\tau_\beta^u} J(u,\tau,x) \tag{1}$$

in both u and $\tau \le \tau_\beta$ is solved as follows:

1. *$v(x) = g(x) = -\log h(x)$.*
2. *All pairs (u^*, τ) for the Markov feedback control u^* given by*

$$u^*(x) \;=\; -\sigma(x)^\top \nabla g(x),$$

   and arbitrary stopping times $\tau \le \tau_\beta$ are optimal.

3. *For any admissible u,*

$$J(u,\tau,x) - v(x) = \mathrm{KL}\left(\mathbb{P}_x^{Z^u}\big|_{\mathcal{F}_{\tau_\beta}} \,\middle\|\, \mathbb{P}_x^{Z^h}\big|_{\mathcal{F}_{\tau_\beta}}\right),$$

   so the Kullback–Leibler divergence between the controlled law and the h-transformed law (up to τ_β) equals the variational gap.

The proof can be found in Appendix B. We note here that for the cases we are interested in later, the first entrance time into the support just happens to coincide with the lifetime ζ.

3.3 Connection to time-inversion

Compared to the deterministic case, the time reversal from a random time is much less elaborated in the literature. In terms of the h-transform, however, the results are very clear: for the h-transformed process Z^h with (finite) lifetime ζ, we consider the time-reversed process

$$\overleftarrow{Z}_s^h := Z_{\zeta-s}^h,\ 0 < s < \zeta,$$

killed at $s = \zeta$. Let also α be a fixed initial distribution of Z^h, and define

$$\overleftarrow{h}(x) = \int \frac{G_r(x,y)}{h(y)}\,\alpha(dy). \tag{2}$$

We then have the following key result, whose proof is given in Appendix A.

Proposition 3. *1. $\overleftarrow{Z}^h$ has the same distribution as $Z^{\overleftarrow{h}}$ with initialisation $Z_0^{\overleftarrow{h}} \sim Z_{\zeta-}^h$; in particular,*

$$d\overleftarrow{Z}_s^h = \overleftarrow{b}^h(\overleftarrow{Z}_s^h)\,ds + \sigma(\overleftarrow{Z}_s^h)\,d\overleftarrow{W}_s,$$
$$\overleftarrow{b}^h(y) = b(y) + \sigma(y)\sigma(y)^\top \nabla \log \overleftarrow{h}(y),$$

$\overleftarrow{W}$ a Brownian motion, outside the support of α.

2. *$\overleftarrow{Z^h}$ is killed on the support of α.*

3. *When started in $x \notin \operatorname{supp}\alpha$, the distribution of $Z^{\overleftarrow{h}}$ at its lifetime ζ is given by*

$$\mathbb{P}_x(Z^{\overleftarrow{h}}_{\zeta-} \in dy) = \frac{G_r(x,y)}{\overleftarrow{h}(x)h(y)}\,\alpha(dy).$$

In particular, for $\mathbb{P}_x(Z^h_{\zeta-} \in \cdot)$-a.e. $x \notin \operatorname{supp}\alpha$ it holds for the time reversed process $\overleftarrow{Z^h}$ that

$$\mathbb{P}_x(\overleftarrow{Z^h}_{\zeta-} \in dy) = \frac{G_r(x,y)}{\overleftarrow{h}(x)h(y)}\,\alpha(dy).$$

The interpretation of the time reversal in terms of the h-transform has interesting consequences. Let us recall the concept of a polar set: for Z^h, $A \subseteq \mathcal{I}$ is called *polar* if $\mathbb{P}_x(Z^h_t \in A \text{ for some } t > 0) = 0$ for all $x \in \mathcal{I}$. Typical examples of polar sets for diffusions are sets on low-dimensional submanifolds. For instance, if Z^h is a d-dimensional Brownian motion, any set with Hausdorff dimension less than $d-2$ is polar, see e.g. Mörters & Peres (2010); a related result for general diffusions is also provided in Ramasubramanian (1988).

Proposition 4. *Assume that the support of the initial distribution α is polar for Z^h. Then, $\overleftarrow{Z^h}$ is killed at the first entry into the support of α. In particular, the distribution of $\overleftarrow{Z^h}$ is entirely characterised by the drift $\overleftarrow{b^h}$ and the support of α.*

Proof. The assumption implies that the starting distribution of Z^h lies in a polar set, which is not visited during the rest of the lifetime. This means that the time-reversed process $\overleftarrow{Z^h}$ only visits this set at the time of killing. □

3.4 Key examples

1. In the case $\beta(dy) = rG_r(x_0,y)\,m(dy)$ (the distribution of Z at the exponential time when started in x_0), we have

$$h(x) = \int \frac{G_r(x,y)}{G_r(x_0,y)}\,\beta(dy) = \int rG_r(x,y)\,m(dy) = 1.$$

Then, Z^h has the same dynamics as Z and ζ is an independent $\operatorname{Exp}(r)$-time. The time-reversed dynamics are determined by $\overleftarrow{h}(x) = \int G_r(x,y)\,\alpha(dy)$.

2. In the case $\beta(dy) = \delta_{x_1}$ for some $x_1 \in \mathcal{I}$, we have

$$h(x) = G_r(x,x_1)/c_{x_1},\ c_{x_1} = G_r(x_0,x_1).$$

h is thus independent of the starting point, except for a multiplicative constant that is irrelevant for the h-transform, so we may choose $c_{x_1} = 1$ w.l.o.g. Z^h is thus an exponential bridge killed in x_1, no matter where the starting point was, with drift

$$b_h(x) = b(x) + \sigma(x)\sigma(x)^\top \nabla \log G_r(x,x_1).$$

In this case,

$$\overleftarrow{h}(x) = \int \frac{G_r(x,y)}{G_r(x_1,y)}\,\alpha(dy).$$

3. In case $Z = W$ is a Brownian motion, we consider radially symmetric functions $h(x) = f(|x|)$ of the form

$$h(x) = \int_{\partial B_R} G_r(x,y)\,\sigma_R(dy),$$

where σ_R denotes the surface measure on the sphere with radius R. We obtain that (except for scaling)

$$f(y) = y^{-\nu} I_\nu(y\sqrt{2r}), \quad \nu = \frac{d-2}{2},$$

I_ν is a modified Bessel function of the first kind, and this results in the forward process Z^h with drift

$$b^h(x) = \nabla \log h(x) = \sqrt{2r} \frac{I_{\nu+1}(\|x\|\sqrt{2r})}{I_\nu(\|x\|\sqrt{2r})} \frac{x}{\|x\|},$$

killed when exiting B_R (in the sense of a last exit). For large dimensions d and moderate values of $\|x\|$, it holds that $b^h(x) \approx \frac{2r}{d} x$ (see 10.41 in DLMF), so that the forward process Z^h can be approximated using the (non-stationary) Ornstein–Uhlenbeck process dynamics

$$d\widetilde{Z}_t = \frac{2r}{d} \widetilde{Z}_t \, dt + dW_t. \tag{3}$$

4 The proposed model

Standard diffusion generative models rely on a deterministic time horizon, leading to time-dependent backward processes and inefficiencies in the forward process due to excessive noise application. Two promising directions for improving these models can be identified: (1) replacing the deterministic time horizon with a randomised one to achieve time-homogeneous backward dynamics, and (2) introducing conditioning in the forward process to reduce the need for extensive noise application.

We achieve the simultaneous implementation of both modifications in a unified framework employing the h-transform described above. To this end, we use an appropriate choice of the process Z^h as our forward process and $\overleftarrow{Z^h}$ as backward process, as discussed next.

4.1 Three possible implementations for unconditional sampling

The examples from Section 3.4 immediately suggest possible realisations of our framework. Specifically, we use Z^h as a forward process initialised in the data distribution α to learn the drift of the backward process $\overleftarrow{Z^h}$ (see the following sections for details). To generate (unconditional) samples from α, we require suitable initial distributions for $\overleftarrow{Z^h}$, which are (approximately) given in the following examples:

1. **Long exponential time horizons:** If we choose an ergodic diffusion Z (e.g., an Ornstein–Uhlenbeck process) and set $h = 1$, we can pick r small and let ζ follow an $\mathrm{Exp}(r)$ distribution. Under this setup, $\mathbb{P}_\alpha(Z^h_{\zeta-} \in dy)$ is close to the stationary distribution of Z. Thus, we may initialize the backward process $\overleftarrow{Z^h}$ in the stationary distribution of Z (if it exists).
2. **Exponential bridge:** In the second example, the forward process Z^h is an exponential bridge targeting a specific state x_1. Here, we can sample from the data distribution by starting the backward process $\overleftarrow{Z^h}$ at $Z^h_{\zeta-} = x_1$.
3. **Hitting a large sphere:** In the third example, if α is compactly supported and the radius R is chosen sufficiently large, $\mathbb{P}_\alpha(Z^h_{\zeta-} \in dy)$ is approximately a uniform distribution over the sphere, which can then serve as an initial distribution for $\overleftarrow{Z^h}$. If helpful, the forward process can be approximated by (3), which allows a more direct sampling.

In the idealised setting where α is known analytically, generation via the backward process for unconditional sampling, i.e., sampling from α, proceeds as described in Algorithm 1. This algorithm is derived from the results presented in Section 3.3. The simulation of the backward process up to the lifetime ζ can be implemented using standard numerical methods, such as the Euler–Maruyama scheme. In the most general setup, the backward process terminates at a randomised Markovian stopping time of the form

$$\zeta = \inf\{t : A^h_t \geq E\},$$

Algorithm 1 Idealised generation when no learning is necessary

Input: $r > 0$, Green kernel G_r of diffusion Z, forward transform h, target distribution α, backward initial distribution $\beta' \approx \mathbb{P}_\alpha(Z^h_{\zeta-})$
Set $\overleftarrow{h}(x) = \int \frac{G_r(x,y)}{h(y)}\,\alpha(dy)$
Generation:
draw $x \sim \beta'$
Simulate path $(y_t)_{t\in[0,\zeta)}$ of h-transform $Y = \overleftarrow{Z^h}$ initialised in x until lifetime ζ, with dynamics given in Proposition 3;
Output: $y_{\zeta-}$

where E is an independent exponential random variable and A^h_t is an additive functional supported within the data distribution, which describes the measure of killing. For a detailed mathematical discussion and numerical approximations of A^h, see Bally (1989); Stoica (1992); Christensen & Schultz (2024). This framework directly enables simulation but poses challenges for estimation. However, the situation becomes significantly simpler when the data distribution α resides within a polar set, as will be discussed in the following section.

4.2 Polarity hypothesis

The target distribution α is typically unknown in practice. Consequently, we lack direct access to $\overleftarrow{h}$ as defined in (2), which is required to employ the backward generating process $Y = \overleftarrow{Z^h}$. Thus, $\overleftarrow{h}$ must be inferred from the data. As discussed in Section 3, both the drift of the backward process and the mechanism governing killing must be learned. While the killing mechanism can be described by a measure on the state space $\mathcal{I}$ (the killing measure), this adds significant complexity to the learning process.

A natural approach is to terminate the backward process as soon as a meaningful element of the data distribution is encountered. Proposition 4 offers a criterion for when this is feasible: specifically, when the data distribution is concentrated in a polar set for the forward process, such as a lower-dimensional manifold. This aligns naturally with the manifold hypothesis: this well-explored concept in the literature assumes that high-dimensional data typically lie on or near a low-dimensional manifold (or a union of such) (Loaiza-Ganem et al., 2024) and has been empirically verified for image data (Pope et al., 2021; Brown et al., 2023).

In our framework, we adopt a slightly more general assumption, referred to as the *polarity hypothesis*, which posits that the data resides in a polar set. This generalisation allows us to disentangle learning of the drift and the killing mechanism, as detailed in the following sections.

4.3 Learning the drift

Here we focus on the case where the polarity hypothesis is satisfied, such that by Proposition 4 the lifetime of $\overleftarrow{Z^h}$ is given by the first hitting time of $\Omega := \operatorname{supp}\alpha$, which we assume to be known for the moment. Based on Proposition 3, we aim to fit the data to a class of time-homogeneous diffusion processes

$$
\begin{aligned}
dY^\theta_s &= \overleftarrow{b^\theta}(Y^\theta_s)\,ds + \sigma(Y^\theta_s)\,dW_s, \quad Y^\theta_0 \sim \mathbb{P}_\alpha(Z^h_{\zeta-} \in \cdot),\\
\overleftarrow{b^\theta}(y) &= b(y) + \sigma(y)\sigma(y)^\top s^\theta(y),
\end{aligned}
$$

induced by a suitable class of candidate functions $\mathcal{S} = \{s^\theta : \theta \in \Theta\}$, and for an optimization objective and corresponding optimiser θ^* to be determined below, run the backward process Y^{θ^*} until an appropriate stopping time. Ideally, we would like to stop in the first hitting time of Ω by Y^{θ^*}, i.e., in $\zeta^{\theta^*} = \tau_\Omega(Y^{\theta^*})$, denoting $\tau_A(X) := \inf\{t : X_t \in A\}$ for a process X and a set $A \subset \mathbb{R}^d$. However, in the typical situation where Ω is a lower dimensional manifold with no simple structure, designing a candidate class $\mathcal{S}$ that guarantees $\mathbb{P}(\zeta^{\theta^*} < \infty) = 1$ is a difficult and highly problem-specific task (e.g., a C^2 submanifold Ω is polar for non-degenerate diffusions if its Hausdorff dimension is no larger than $d-2$ (Friedman, 1975, Chapter 11), (Ramasubramanian, 1988)). Instead, for some small $\varepsilon > 0$, we consider the closed ε-environment

$\Omega_\varepsilon \coloneqq \{x \in \mathbb{R}^d : d(x, \Omega) \leq \varepsilon\}$ and run Y^{θ^*} until

$$\zeta_\varepsilon^{\theta^*} = \tau_{\Omega_\varepsilon}(Y^{\theta^*}).$$

Thus, we are not targeting α directly, but the distribution of $\overleftarrow{Z}^h_{\zeta_\varepsilon}$ for

$$\zeta_\varepsilon \coloneqq \tau_{\Omega_\varepsilon}(\overleftarrow{Z}^h).$$

This is justified for small $\varepsilon > 0$, since by continuity of the sample paths, we have that $\tau_{\Omega_\varepsilon}(\overleftarrow{Z}^h)$ almost surely increases to ζ as $\varepsilon \downarrow 0$ if Ω is closed, and therefore obtain the a.s. convergence $\lim_{\varepsilon \downarrow 0} \overleftarrow{Z}^h_{\zeta_\varepsilon} = \overleftarrow{Z}^h_{\zeta-} \sim \alpha$. This is comparable to early stopping of the generating process in standard diffusion models with deterministic time horizon, which implies that the generative model targets the original data set blurred by some small Gaussian noise. On a forward time scale, the first entrance time in Ω^ε of the backward process $\overleftarrow{Z}^h$ corresponds to the last exit time σ_ε of Ω_ε by the forward process Z^h, that is,

$$\sigma_\varepsilon \coloneqq \sup\{t < \zeta : Z^h_t \in \Omega_\varepsilon\} = \zeta - \zeta_\varepsilon.$$

As a natural optimisation objective, we therefore target the Kullback–Leibler divergence between the law $\overleftarrow{\mathbb{P}}^\varepsilon$ of the theoretical generating process $\overleftarrow{Z}^h$ killed in ζ_ε and the law $\mathbb{P}^{\theta,\varepsilon}$ of the parametrised generating process Y^θ killed at first entrance into Ω_ε, i.e., at $\zeta_\varepsilon^\theta \coloneqq \tau_{\Omega_\varepsilon}(Y^\theta)$. To this end, we set

$$\begin{aligned}
\theta^* &\in \operatorname*{arg\,min}_{\theta \in \mathcal{S}} \mathcal{L}_{\mathrm{ex}}(\theta), \\
\mathcal{L}_{\mathrm{ex}}(\theta) &\coloneqq \mathbb{E}\Big[\int_0^{\zeta_\varepsilon} \|s^\theta(\overleftarrow{Z}^h_t) - \nabla \log \overleftarrow{h}(\overleftarrow{Z}^h_t)\|^2 \, dt\Big] \\
&= \mathbb{E}\Big[\int_{\sigma_\varepsilon}^{\zeta} \|s^\theta(Z^h_t) - \nabla \log \overleftarrow{h}(Z^h_t)\|^2 \, dt\Big]
\end{aligned}$$

which is the natural analog to the *explicit score matching* objective in standard diffusion models. Then, given sufficient integrability conditions, Girsanov's theorem indeed yields that

$$\begin{aligned}
\mathrm{KL}\big(\overleftarrow{\mathbb{P}}^\varepsilon \,\|\, \mathbb{P}^{\theta,\varepsilon}\big) &= \frac{1}{2}\mathbb{E}\Big[\int_0^{\zeta_\varepsilon} \|\sigma^{-1}(\overleftarrow{Z}^h_t)(\overleftarrow{b}^h(\overleftarrow{Z}^h_t) - \overleftarrow{b}^\theta(\overleftarrow{Z}^h_t))\|^2 \, dt\Big] \\
&\leq \frac{C_\sigma}{2}\mathcal{L}_{\mathrm{ex}}(\theta),
\end{aligned}$$

for $C_\sigma \coloneqq \sup_{x \in \mathbb{R}^d} \|\sigma^{-1}(x)\|^2$. Consequently,

$$\mathrm{KL}\big(\mathbb{P}(\overleftarrow{Z}^h_{\zeta_\varepsilon} \in \cdot) \,\|\, \mathbb{P}(Y^{\theta^*}_{\zeta_\varepsilon^{\theta^*}} \in \cdot)\big) \lesssim \min_{\theta \in \mathcal{S}} \mathbb{E}\Big[\int_0^{\zeta_\varepsilon} \|s^\theta(\overleftarrow{Z}^h_t) - \nabla \log \overleftarrow{h}(\overleftarrow{Z}^h_t)\|^2 \, dt\Big],$$

showing that if $\mathcal{S}$ is rich enough to guarantee a high approximation quality of $\nabla \log \overleftarrow{h}$ on $(\Omega_\varepsilon)^{\mathrm{c}}$ (e.g., a suitable class of neural networks), our generative process produces high quality samples for slightly blurred target data. However, the general inaccessibility of $\overleftarrow{h}$, and thus of $\mathcal{L}_{\mathrm{ex}}$, requires us to determine a tractable training objective that is comparable to $\mathcal{L}_{\mathrm{ex}}$. This is provided by our next result, which in our proposed model can be understood as an analog of the *denoising score matching loss* (Vincent, 2011) employed in standard diffusion models.

Proposition 5. *Given sufficient integrability conditions, there exists a constant C independent of θ, such that*

$$\mathcal{L}_{\mathrm{ex}}(\theta) = \mathcal{L}(\theta) - C_\varepsilon(\theta) + C,$$

where

$$\mathcal{L}(\theta) \coloneqq \mathbb{E}\Big[\int_{\sigma_\varepsilon}^{\zeta} \|s^\theta(Z^h_t) - \nabla_{Z^h_t} \log G_r(Z^h_t, Z^h_0)\|^2 \, dt\Big]$$

and

$$C_\varepsilon(\theta) \coloneqq 2\mathbb{E}\Big[\int_0^{\sigma_\varepsilon} \big\langle s_\theta(Z^h_t), \nabla_{Z^h_t} \log \tfrac{G_r(Z^h_t, Z^h_0)}{\overleftarrow{h}(Z^h_t)} \big\rangle \, dt\Big].$$

Algorithm 2 Generation for unknown $\overleftarrow{h}$ and known polar data support Ω

Input: data $\{y_i\}_{i=1}^n \overset{iid}{\sim} \alpha$, $N \in \mathbb{N}$, $\varepsilon > 0, r > 0$, Green kernel G_r of diffusion Z, enlarged data support Ω_ε, forward transform h, backward initialisation $\beta' \approx \mathbb{P}_\alpha(Z^h_{\zeta-} \in \cdot)$ with $\operatorname{supp} \beta' \subset \Omega_\varepsilon^c$, function class $\mathcal{S}$.
for $j = 1$ **to** N **do**
draw y_{i_j} with replacement from $\{y_i\}_{i=1}^n$ and simulate path $(z_t^{h,i_j})_{t\in[0,\zeta^{i_j})}$ of h-transform Z^h started in y_{i_j};
end for
Training: Learn $\widehat{\theta} = \arg\min_{\theta\in\mathcal{S}} \widehat{\mathcal{L}}(\theta)$ for

$$\widehat{\mathcal{L}}(\theta) = \frac{1}{N}\sum_{j=1}^N \int_{\sigma_\varepsilon^{i_j}}^{\zeta^{i_j}} \|s^\theta(z_t^{h,i_j}) - \nabla_1 \log G_r(z_t^{h,i_j}, y_{i_j})\|^2 \, dt$$

Generation:
Draw x from β'
Simulate path $(y_t^{\widehat{\theta}})_{t\in[0,\tau_{\Omega_\varepsilon}(y^{\widehat{\theta}})]}$ of $Y^{\widehat{\theta}}$ initialised in x
Output: $\widehat{\theta}, y^{\widehat{\theta}}_{\tau_{\Omega_\varepsilon}(y^{\widehat{\theta}})}$

Since we assumed Ω to be polar, $C_\varepsilon(\theta)$ vanishes as $\varepsilon \to 0$, so that for small $\varepsilon > 0$ we expect minimisers of $\mathcal{L}$ – which is accessible for a given Green kernel G_r and can be efficiently approximated by a Monte-Carlo estimator $\widehat{\mathcal{L}}$ based on simulated forward trajectories of Z^h – to be also approximate minimisers of $\mathcal{L}_{\text{ex}}$. This motivates the generative Algorithm 2, which outputs an estimated drift parameter $\widehat{\theta}$ and a corresponding sample from the estimated backward generative process $Y^{\widehat{\theta}}$ initialised in a distribution β' approximating the true forward terminal distribution $\mathbb{P}_\alpha(Z^h_{\zeta-} \in \cdot)$. We note once again that a significant difference from most other methods is that there is no time dependence to be learned. This is also a feature of the first hitting model from Ye et al. (2022), which, however, relies on non-polarity of Ω and access to the Ω-dependent Poisson kernel $\mathbb{P}_x(Z_{\tau_\Omega} \in dz)$ for training and generation purposes. Usually, this is analytically tractable only for simple sets Ω. In contrast, our learning and generation step only requires access to the Green kernel G_r of the *unconditioned* forward process Z, which is independent of Ω and therefore allows for enhanced flexibility.

4.4 (Implicit) manifold learning

We now turn to the most general setting, where the polar data support $\Omega = \operatorname{supp} \alpha$ is unknown. Since the generative mechanism proposed in Algorithm 2 fundamentally requires Ω_ε as input in order to know where to stop the generative process, we must infer Ω_ε from the data. To this end, we discuss two different strategies: first, a traditional statistical plug-in approach that simply replaces Ω_ε in Algorithm 1 with a separately obtained estimate $\widehat{\Omega_\varepsilon}$. The second approach is based on the observation that under the polarity hypothesis, the unknown drift component $\nabla \log \overleftarrow{h}$ implicitly encodes the geometry of the data manifold Ω, cf. Proposition 8. This motivates an adaptive stopping criterion that only requires an estimate of $\nabla \log \overleftarrow{h}$ as input and may be interpreted as a first hitting time of an implicit manifold estimate.

Approach based on separation of manifold and drift estimation A first—but perhaps naïve—strategy is to build an estimator $\widehat{\Omega_\varepsilon}$ based on the given data in a pre-processing step, and then use a plug-in approach to draw samples based on Algorithm 2. This is formalised in Algorithm 3.

To construct $\widehat{\Omega_\varepsilon}$ given a data sample $(Y_i)_{i=1}^n \overset{iid}{\sim} \alpha$, we can consider the following two natural approaches:

First, we may target Ω_ε directly by blurring the data via $Y_i^\varepsilon = Y_i + \eta_i^\varepsilon$, where $(\eta_i^\varepsilon)_{i=1}^n$ is some iid noise that is independent of the data sample and absolutely continuous with support $B(0,\varepsilon)$, e.g., $\eta_i^\varepsilon \sim \mathcal{U}(B(0,\varepsilon))$. Then, Y_i^ε is absolutely continuous with support Ω_ε and density $\pi^\varepsilon(x) = \int_\Omega \pi_{\eta^\varepsilon}(x-y)\,\alpha(dy)$. Recovering the

Algorithm 3 Full "naïve" generative algorithm

Input: data $\{y_i\}_{i=1}^n \overset{iid}{\sim} \alpha$, $\varepsilon > 0, r > 0$, Green kernel G_r of diffusion Z, estimator $\widehat{\Omega_\varepsilon}$ of enlarged data support, forward transform h, backward initial distribution $\beta' \approx \mathbb{P}_\alpha(Z^h_{\zeta_-})$ with $\operatorname{supp}\beta' \subset \widehat{\Omega_\varepsilon}^{\mathrm{c}}$, function class $\mathcal{S}$.
Set $\Omega_\varepsilon \leftarrow \widehat{\Omega_\varepsilon}$
Generation: Run Algorithm 2
Output: $\widehat{\theta}, y^{\widehat{\theta}}_{\tau_{\widehat{\Omega_\varepsilon}}(y^{\widehat{\theta}})}$

compact support Ω_ε of such a data sample is a well-studied statistical problem. The perhaps most common approaches are plug-in estimators Cuevas & Fraiman (1997) that set

$$\widehat{\Omega_\varepsilon} = \{\widehat{\pi^\varepsilon} > \delta_n\}$$

for some nonparametric estimator $\widehat{\pi^\varepsilon}$ of π^ε and a tuning parameter $\delta_n \to 0$, or the simple and intuitive *Devroye–Wise* estimator Devroye & Wise (1980) given by

$$\widehat{\Omega_\varepsilon} = \bigcup_{i=1}^n B(Y_i^\varepsilon, \delta_n).$$

Minimax optimality of these estimators with respect to metrics such as the Hausdorff metric or the symmetric difference volume has been established under different assumptions on the data density π^ε and the geometry of Ω_ε (Korostelëv & Tsybakov, 1993; Mammen & Tsybakov, 1995; Cuevas & Fraiman, 1997; Cuevas & Rodríguez-Casal, 2004; Biau et al., 2008). The rates, however, generally slow down exponentially in terms of the ambient data dimension d, which motivates the second approach that allows to exploit directly the lower-dimensional structure of Ω.

For this second approach, instead of estimating Ω_ε directly, we may first construct an estimator $\widehat{\Omega}$ of the true data support Ω based on the unmodified data sample $(Y_i)_{i=1}^n$ and then set

$$\widehat{\Omega_\varepsilon} = (\widehat{\Omega})_\varepsilon.$$

This approach might be more promising for high-dimensional data since a lot of progress has been made in recent years on statistical theory and algorithmic implementation for support estimation of distributions concentrated on lower-dimensional manifolds. Important contributions that provide estimators $\widehat{\Omega}$ with provable convergence rates only depending on the smoothness s and the dimension $d' < d$ of the data manifold include Genovese et al. (2012); Ma & Fu (2012); Kim & Zhou (2015); Aamari & Levrard (2018; 2019); Aamari et al. (2023). While computationally more involved than the simple estimators based on noised data discussed above, the estimators from Aamari & Levrard (2018; 2019); Aamari et al. (2023) are constructive and implementable.

Fully integrated drift and support estimation The plug-in approach is statistically sound, but it is rather unclear if it will lead to convincing results in actual implementations. Indeed, part of the empirical success of standard diffusion models is explained by an implicit adaptation to the data support via the learned score networks Stanczuk et al. (2024). In this spirit, we propose a third strategy that only relies on information obtainable from the approximation of the log-gradient of the reverse h-transform. This is based on the intuition that $\nabla \log \overleftarrow{h}(x)$ must explode when x approaches the low-dimensional data support Ω to ensure that the reverse diffusion process is killed in finite time. To formalise this intuition in the following, let us first give a useful characterisation of $\nabla \log \overleftarrow{h}$. The proof is given in Appendix C.

Lemma 6. *For all $x \notin \operatorname{supp}\alpha$, it holds that*

$$\nabla \log \overleftarrow{h}(x) = \mathbb{E}\left[\nabla_x \log G_r(x, Z^{\overleftarrow{h}}_{\zeta_-}) \mid Z^{\overleftarrow{h}}_0 = x\right]. \tag{4}$$

In particular, for $\mathbb{P}_\alpha(Z^h_{\zeta_-} \in \cdot)$-a.e. x, it holds that

$$\nabla \log \overleftarrow{h}(x) = \mathbb{E}_\alpha\left[\nabla_x \log G_r(Z^h_0, x) \mid Z^h_{\zeta_-} = x\right]. \tag{5}$$

Remark 7. (i) If the forward process is simply an exponentially killed diffusion (i.e., $h = 1$), the representation of $\nabla \log \overleftarrow{h}$ as a conditional expectation can be used to motivate a simplified denoising score matching objective that only requires simulation of the forward process at its exponential lifetime, cf. (Christensen et al., 2025, Theorem 2.7) for the case of killed Brownian motion.

(i) Formula (5) is analogue to the well-known score representation

$$\nabla \log p_t(x) = \mathbb{E}_\alpha \left[\nabla_x \log p_t(\vec{X}_0, x) \mid \vec{X}_t = x \right]$$

in conventional diffusion models for a homogeneous forward process $\vec{X}$ with transition densities $p_t(x, y)$, which is linked to Tweedie's formula Robbins (1956).

Let us for the moment consider the Green kernel G_r of a Brownian motion. Then, we have

$$\nabla_x \log G_r(x, y) = -\sqrt{2r}\, \mathrm{sign}(x - y) \frac{K_{\nu+1}(\sqrt{2r}\|x - y\|)}{K_\nu(\sqrt{2r}\|x - y\|)}, \quad \nu = \frac{d-2}{2}, \ \mathrm{sign}(x) := \frac{x}{\|x\|},$$

and therefore from above, for any $x \notin \operatorname{supp} \alpha$,

$$\nabla \log \overleftarrow{h}(x) = \sqrt{2r}\, \mathbb{E}^x \left[\mathrm{sign}\left(Z^{\overleftarrow{h}}_{\zeta-} - x \right) \frac{K_{\nu+1}(\sqrt{2r}\|Z^{\overleftarrow{h}}_{\zeta-} - x\|)}{K_\nu(\sqrt{2r}\|Z^{\overleftarrow{h}}_{\zeta-} - x\|)} \right] \approx \frac{d}{\mathrm{e}}\, \mathbb{E}_x \left[\frac{\mathrm{sign}(Z^{\overleftarrow{h}}_{\zeta-} - x)}{\left\| Z^{\overleftarrow{h}}_{\zeta-} - x \right\|} \right], \tag{6}$$

for large d, cf. 10.41 in DLMF. A more careful analysis furthermore reveals that the scaling factor d in (6) can be replaced by $d - \dim \operatorname{supp} \alpha$ under mild regularity assumptions on the data manifold. This demonstrates quite clearly how the reverse process is dragged onto the data manifold. To exploit this connection for our purposes, we can use the explosive behavior of $\nabla \log \overleftarrow{h}$ near $\operatorname{supp} \alpha$ to obtain an insightful characterisation of the backward killing time, as the following result shows.

Proposition 8. *Suppose that*

(i) the drift b of Z is locally bounded and its diffusion coefficient $\sigma\sigma^\top$ is uniformly elliptic;

(ii) $\operatorname{supp} \alpha$ is polar for both Z^h and Y solving $\sigma(Y_t)\, \mathrm{d}B_t$.

Then, the killing time ζ of $\overleftarrow{Z^h}$ is a.s. given by

$$\zeta = \inf \left\{ t \geq 0 : \sup_{s \leq t} \left\| \nabla \log \overleftarrow{h}(Z^{\overleftarrow{h}}_s) \right\| = \infty \right\} = \inf \left\{ t \geq 0 : \left\| \nabla \log \overleftarrow{h}(\overleftarrow{Z^h}) \right\|_{L^2([0,t])} = \infty \right\}.$$

Proof. The proof proceeds along the same lines as the one of Theorem 2.8 in Christensen et al. (2025), where the result is proved for the particular case of the forward process Z^h being given by an exponentially killed Brownian motion. □

Most importantly in our context, the above considerations lead to simple and interpretable stopping criteria for the reverse process based on an estimate $s^{\widehat{\theta}}$ of $\nabla \log \overleftarrow{h}(x)$. For instance, when Z is a Brownian motion and we target killing on the slightly enlarged data support Ω_ε, then $Z^{\overleftarrow{h}}_{\zeta-} \in \Omega$ implies for any $x \notin \Omega_\varepsilon$ that

$$\|\nabla \log \overleftarrow{h}(x)\| \approx \frac{d}{\mathrm{e}} \left\| \mathbb{E}_x \left[\frac{\mathrm{sign}(Z^{\overleftarrow{h}}_{\zeta-} - x)}{\left\| Z^{\overleftarrow{h}}_{\zeta-} - x \right\|} \right] \right\| \leq \frac{d}{\mathrm{e}}\, \mathbb{E}_x \left[\frac{1}{\|Z^{\overleftarrow{h}}_{\zeta-} - x\|} \right] \leq \frac{d}{\mathrm{e}\varepsilon}.$$

Combining this with Proposition 8, we may therefore reasonly stop the data-generating mechanism as soon as $\|s^{\widehat{\theta}}(Y^{\widehat{\theta}}_t)\| \gtrsim \frac{d}{\varepsilon}$. This translates into a score-based reverse first killing time

$$\widehat{\zeta_\varepsilon} := \inf \left\{ t \geq 0 : \|s^{\widehat{\theta}}(Y^{\widehat{\theta}}_t)\| \geq \gamma d \varepsilon^{-1} \right\},$$

Algorithm 4 Full generative algorithm without separate support estimation

Input: data $\{y_i\}_{i=1}^n \overset{iid}{\sim} \alpha$, $N \in \mathbb{N}$, $\varepsilon > 0, \underline{T} \geq 0, \gamma > 0, r > 0$, (approximation of) Green kernel G_r of diffusion Z, enlarged data support Ω_ε, forward transform h, backward initialization $\beta' \approx \mathbb{P}_\alpha(Z^h_{\zeta-} \in \cdot)$, function class $\mathcal{S}$.

for $j = 1$ **to** N **do**

Draw y_{i_j} uniformly with replacement from $\{y_i\}_{i=1}^n$ and simulate path $(z_t^{h,i_j})_{t\in[0,\zeta^{i_j})}$ of h-transform Z^h started in y_{i_j};

end for

Training: Learn $\widehat{\theta} = \arg\min_{\theta\in\mathcal{S}} \widehat{\mathcal{L}}(\theta)$ for

$$\widehat{\mathcal{L}}(\theta) = \frac{1}{N}\sum_{j=1}^{N} \int_{\underline{T}\wedge\zeta^{i_j}}^{\zeta^{i_j}} \left\| s^\theta(z_t^{h,i_j}) - \nabla_1 \log G_r(z_t^{h,i_j}, y_{i_j}) \right\|^2 dt$$

Generation:

Draw x from β'

Simulate path $(y_t^{\widehat{\theta}})_{t\in[0,\widehat{\zeta}_\varepsilon]}$ of $Y^{\widehat{\theta}}$ initialised in x until $\widehat{\zeta}_\varepsilon = \inf\{t \geq 0 : \|s^{\widehat{\theta}}(y_t^{\widehat{\theta}})\| \geq \gamma d\varepsilon^{-1}\}$

Output: $\widehat{\theta}, y^{\widehat{\theta}}_{\widehat{\zeta}}$

for some hyperparameter γ or, equivalently, the first hitting time of the implicit data support estimator

$$\widehat{\Omega_\varepsilon} := \left\{x \in \mathbb{R}^d : \|s^{\widehat{\theta}}(x)\| \geq \gamma d\varepsilon^{-1}\right\}.$$

Such a stopping rule encodes mathematically, what a human observer would do intuitively when tracking the sequentially generated data:

> *stop the generation process as soon the observed ordered set of pixels visually resembles a high quality image.*

To obtain an estimator $s^{\widehat{\theta}}$ based on Proposition 5 without prior estimation of Ω_ε, we simply replace the lower time index $\sigma_\varepsilon := \sup\{t < \zeta : Z_t^h \in \Omega_\varepsilon\}$ in the training objective with $\zeta \wedge \underline{T}$ for some hard-coded small value $\underline{T} \geq 0$. For specific forward processes such as exponentially killed diffusions, more elaborate alternatives are theoretically justified, see (Christensen et al., 2025, Theorem 2.7) and Remark 7. The full generative algorithm is given in Algorithm 4.

4.5 Noise-Schedules as a special case of radial h-transforms

In classical diffusion models, the forward noise schedule α_t is of key importance for sampling. In our time-homogeneous setting this effect can be reproduced in high dimension by choosing a suitable h-transform for the forward process, exemplified here for a Brownian motion Z.

Consider the standard forward noising

$$x_t = \sqrt{\alpha_t}\, x_0 + \sqrt{1-\alpha_t}\, z, \qquad z \sim \mathcal{N}(0, I_d),\ \|x_0\| = o(\sqrt{d}),$$

for which the law of large numbers gives $\|x_t\|/\sqrt{d} \longrightarrow \rho^*(t) := \sqrt{1-\alpha_t}$ as $d \longrightarrow \infty$. The limiting radial trajectory $\rho^*(t)$ satisfies

$$\dot{\rho}^*(t) = -\frac{\dot{\alpha}_t}{2\sqrt{1-\alpha_t}}. \tag{7}$$

To reproduce this within our framework, define Ψ by

$$\Psi'(\rho) = \tfrac{1}{2\rho} - \dot{\rho}^*((\rho^*)^{-1}(\rho)), \qquad \Psi(\rho_0) = 0,$$

on a radial range corresponding to $t \in [\tau, T] \subset (0, \infty)$, where ρ^* is smooth and strictly monotone. If we can choose a positive radial r-potential h whose high-dimensional log-profile satisfies

$$-\frac{1}{d} \log h(x) \approx \Psi(\rho), \qquad \rho = \|x\|/\sqrt{d},$$

and assuming these log-asymptotics is differentiable on the considered radial range, then

$$\nabla \log h(x) \approx -\sqrt{d}\, \Psi'(\rho)\, \frac{x}{\|x\|}.$$

For $\rho_t = \|Z_t^h\|/\sqrt{d}$ we then have the approximation

$$d\rho_t \approx \Big(\frac{d-1}{2d\,\rho_t} - \Psi'(\rho_t)\Big)\, dt + \tfrac{1}{\sqrt{d}}\, dB_t$$

where B is a one-dimensional Brownian motion. As $d \longrightarrow \infty$, we obtain the autonomous ODE

$$\dot{\rho} = \frac{1}{2\rho} - \Psi'(\rho) = \dot{\rho}^*((\rho^*)^{-1}(\rho)). \tag{8}$$

Hence, if we set $\rho(t) = \rho^*(t)$, then $(\rho^*)^{-1}(\rho(t)) = t$ (since ρ^* is strictly monotone) and the autonomous ODE (8) reduces to $\dot{\rho}(t) = \dot{\rho}^*(t)$, i.e. (8) and (7) describe the same radial trajectory (possibly up to a time shift determined by the initial condition).

Thus the two descriptions differ in form: the classical schedule induces an explicit time-dependent drift for ρ^*, while the h-transform encodes the same effect through a space-dependent drift $-\Psi'(\rho)$. By the choice of Ψ, both ODEs (8) and (7) coincide along the trajectory $\rho^*(t)$ in the high-dimensional limit; consequently, in high dimensions, a single radial h reproduces the schedule-driven radial behavior in a time-homogeneous formulation.

5 Features of the model

5.1 Natural conditioning

In the procedure described above, we have explained how to sample from the data distribution α using the time-homogeneous backward processes $Y = \overleftarrow{Z^h}$ or a learned version thereof when initialised from the terminal distribution of the forward process. The flexibility of our proposed model compared to existing diffusion models, however, arises from the observation that due to its time-homogeneous nature, the backward process Y can be started in an interpretable way directly from *any* initial state x, regardless of the specific time, along the following lines:

1. Initialize the backward process Y at a chosen state x, which could be:
   - A noisy version of a sample from α, or
   - Any point of interest in the state space.
2. Simulate the backward process Y until the lifetime ζ, generating the trajectory $\{Y_t\}_{t \le \zeta}$.
3. Extract the value $Y_{\zeta-}$ as the final sample. Note that this value lies within the support of α, but is not necessarily distributed exactly according to α.
4. Optionally, repeat the process for multiple initial states x to analyze how the sampling distribution varies with the initialisation.

In Bayesian terminology, unconditional sampling corresponds to drawing from the prior distribution α, whereas the current task involves sampling from $\mathbb{P}_x(Y_{\zeta-} \in dy)$, the posterior distribution conditioned on the input data x.

It is reasonable to expect that $Y_{\zeta-}$ will typically be close to x. Indeed, Proposition 3 states that

$$\mathbb{P}_x(Y_{\zeta-} \in dy) = \frac{1}{\overleftarrow{h}(x)} \frac{G_r(x,y)}{h(y)} \, \alpha(dy).$$

If $h(y)$ is approximately constant over the support of α, the sampling procedure approximates drawing from the measure proportional to $G_r(x,y)\,\alpha(dy)$ (up to a normalizing constant). Typically, $G_r(x,y)$ decreases with the distance between x and y, as illustrated in Remark 9. Consequently, points y closer to x appear with higher frequency than those farther away.

In other words, if the backward process Y is started near the support of α, the resulting sample from α will tend to be close to x. Here, the choice of the diffusion process Z implicitly defines a corresponding notion of distance. So by choosing the starting point of Y, we get natural conditioning, and in exactly the same (learned) model as for unconditional sampling.

5.2 Fine tuning for conditional data generation

In addition to the natural conditioning properties discussed above, we may also use our model to explicitly sample from a conditional distribution of interest, which is only represented through a smaller subset of samples having certain properties. More specifically, suppose that instead of sampling from α we would like to sample from the posterior $\alpha(dy \mid u) \propto \pi(u \mid y)\,\alpha(dy)$, where $\pi(u)$ is the density of some latent random variable U of interest and assume that the likelihood satisfies $\pi(u \mid \cdot) \in L^2(\alpha)$. To do so, we aim at lifting a generative model from our unconditional pre-trained model (i.e., using the learned score $\nabla \log \overleftarrow{h}$), for which more data is available. Ideally, we would like to simulate a backward process with drift determined by $\nabla \log \overleftarrow{h}(\cdot \mid u)$, where $\overleftarrow{h}(x \mid u) := \int \frac{G_r(x,y)}{h(y)} \, \alpha(dy \mid u)$. This can be calculated similarly to Lemma 6 as follows

$$\begin{aligned}
\nabla \log \overleftarrow{h}(x \mid u) &= \frac{\int \nabla_x G_r(x,y) \frac{1}{h(y)} \pi(u \mid y)\,\alpha(dy)}{\int G_r(x,y) \frac{1}{h(y)} \pi(u \mid y)\,\alpha(dy)} \\
&= \frac{\int \nabla_x \log G_r(x,y) \pi(u \mid y) \frac{G_r(x,y)}{\overleftarrow{h}(x) h(y)} \, \alpha(dy)}{\int \pi(u \mid y) \frac{G_r(x,y)}{\overleftarrow{h}(x) h(y)} \, \alpha(dy)} \\
&= \frac{\mathbb{E}_x \left[\nabla_x \log G_r(x, Z^{\overleftarrow{h}}_{\zeta-}) \pi(u \mid Z^{\overleftarrow{h}}_{\zeta-}) \right]}{\mathbb{E}_x \left[\pi(u \mid Z^{\overleftarrow{h}}_{\zeta-}) \right]} \\
&= \mathbb{E}^u_x \left[\nabla_x \log G_r(x, Z^{\overleftarrow{h}}_{\zeta-}) \right],
\end{aligned} \tag{9}$$

where $\mathbb{P}_x = \mathbb{P}(\,\cdot \mid Z^{\overleftarrow{h}}_0 = x)$ and $\mathbb{P}^u_x$ is a probability measure defined by the Radon–Nikodym derivative

$$\frac{d\,\mathbb{P}^u_x}{d\,\mathbb{P}_x} := \frac{\pi\big(u \mid Z^{\overleftarrow{h}}_{\zeta-}\big)}{\mathbb{E}_x \left[\pi\big(u \mid Z^{\overleftarrow{h}}_{\zeta-}\big) \right]},$$

relative to $\mathbb{P}_x$. Recalling that Lemma 6 shows $\nabla \log h(x) = \mathbb{E}_x \left[\nabla_x \log G_r(x, Z^{\overleftarrow{h}}_{\zeta-}) \right]$, we see that the conditional log gradient of the reverse h-transform is obtained via a change of measure induced by the likelihood $\pi(u \mid y)$ and we obtain the decomposition

$$\nabla \log \overleftarrow{h}(x \mid u) = \nabla \log \overleftarrow{h}(x) + \mathbb{E}_x \left[\nabla_x \log G_r(x, Z^{\overleftarrow{h}}_{\zeta-}) \frac{\pi(u \mid Z^{\overleftarrow{h}}_{\zeta-}) - \mathbb{E}_x[\pi(u \mid Z^{\overleftarrow{h}}_{\zeta-})]}{\mathbb{E}_x[\pi(u \mid Z^{\overleftarrow{h}}_{\zeta-})]} \right].$$

Furthermore, (9) yields

$$\begin{aligned}
\nabla \log \overleftarrow{h}(x \mid u) &= \frac{\mathbb{E}_{Z^h_0 \sim \alpha} \left[\nabla_x \log G_r(x, Z^h_0) \pi(u \mid Z^h_0) \,\middle|\, Z^h_{\zeta-} = x \right]}{\mathbb{E}_{Z^h_0 \sim \alpha} \left[\pi(u \mid Z^h_0) \,\middle|\, Z^h_{\zeta-} = x \right]} \\
&= \int \nabla_x \log G_r(x,y) \; \mathbb{P}_{Z^h_0 \sim \alpha(dy \mid u)} \left(Z^h_0 \in dy \,\middle|\, Z^h_{\zeta-} = x \right) \\
&= \mathbb{E}_{Z^h_0 \sim \alpha(dy \mid u)} \left[\nabla_{Z^h_{\zeta-}} \log G_r(Z^h_{\zeta-}, Z^h_0) \,\middle|\, Z^h_{\zeta-} = x \right],
\end{aligned}$$

Algorithm 5 Conditional generation via fine tuning of unconditional algorithm

Input: data $\{y_i\}_{i=1}^n \overset{iid}{\sim} \alpha$, $M \in \mathbb{N}$, $\varepsilon, \delta, \gamma > 0, r > 0$, Green kernel G_r of diffusion Z, forward transform h, backward initialisation $\beta'_u \approx \mathbb{P}_{\alpha(dy|u)}(Z^h_{\zeta-} \in \cdot)$, function class $\mathcal{D}$, unconditional log-gradient h-transform estimator $\widehat{s} = s^{\widehat{\theta}} \approx \nabla \log \overleftarrow{h}$.
for $j = 1$ **to** M **do**
Draw y_{i_j} uniformly from $\{y_1, \dots, y_n\}$ and simulate terminal value $z^{h,i_j}_{\zeta^{i_j}_-}$ of h-transform Z^h started in $y_{i_j} \sim \alpha(dy \mid u)$;
end for
Training: Learn $\widehat{D}_u \in \arg\min_{D \in \mathcal{D}} \widehat{\mathcal{L}}_u(D)$ for

$$\widehat{\mathcal{L}}_u(D) = \frac{1}{M} \sum_{j=1}^{M} \| D(z^{h,i_j}_{\zeta^{i_j}_-}) + \widehat{s}(z^{h,i_j}_{\zeta^{i_j}_-}) - \nabla_{z^{h,i_j}_{\zeta-}} \log G_r(z^{h,i_j}_{\zeta^{i_j}_-}, y_{i_j}) \|^2 \mathbf{1}_{(\delta,\infty)}(\| z^{h,i_j}_{\zeta^{i_j}_-} - y_{i_j} \|)$$

set $\widehat{s}(\cdot \mid u) = \widehat{D}_u + s^{\widehat{\theta}}$
Generation:
Draw x from β'_u
Simulate path $(\widehat{y}^u_t)_{t \in [0, \widehat{\zeta}_\varepsilon]}$ of

$$dY^u_t = (b(Y^u_t) + \sigma\sigma^\top(Y^u_t)\widehat{s}(Y^u_t \mid u))\, dt + \sigma(Y^u_t)\, dW_t, \quad Y^u_0 = x,$$

until $\widehat{\zeta}_\varepsilon = \inf\{t \geq 0 : \|\widehat{s}(\widehat{Y}^u_t \mid u) \geq \gamma d \varepsilon^{-1}\}$
Output: $\widehat{y}^u_{\widehat{\zeta}_\varepsilon}$

where we used that by Bayes' theorem,

$$\pi(u \mid y)\, \mathbb{P}_{Z^h_0 \sim \alpha}(Z^h_0 \in dy \mid Z^h_{\zeta-} = x) \propto \mathbb{P}_{Z^h_0 \sim \alpha(dy|u)}(Z^h_0 \in dy \mid Z^h_{\zeta-} = x).$$

This may be interpreted as an analogue of the score matching identity for the posterior score in diffusion models. Similarly to the fine-tuning approach from Denker et al. (2024), given a pre-processed estimator $s^{\widehat{\theta}}$ of $\nabla \log \overleftarrow{h}$ and a smaller subset of training data for the conditional distribution $\alpha(dy \mid u)$, we may now aim at approximating only the difference

$$D_u(x) := \nabla \log \overleftarrow{h}(x \mid u) - \nabla \log \overleftarrow{h}(x) = \mathbb{E}\left[\nabla_x \log G_r(x, Z^{\overleftarrow{h}}_{\zeta-}) \frac{\pi(u|Z^{\overleftarrow{h}}_{\zeta-}) - \mathbb{E}[\pi(u|Z^{\overleftarrow{h}}_{\zeta-})|Z^{\overleftarrow{h}}_0 = x]}{\mathbb{E}[\pi(u|Z^{\overleftarrow{h}}_{\zeta-})|Z^{\overleftarrow{h}}_0 = x]} \,\middle|\, Z^{\overleftarrow{h}}_0 = x \right],$$

of the log-gradient of the conditional and unconditional h-transform. From (Christensen et al., 2025, Lemma A.1), it follows that for any $\delta > 0$, D_u coincides on $(\Omega_\delta)^c$ with the minimiser of

$$D \mapsto \mathbb{E}_{Z^h_0 \sim \alpha(dy|u)} \left[\left\| D(Z^h_{\zeta-}) + \nabla \log \overleftarrow{h}(Z^h_{\zeta-}) - \nabla_{Z^h_{\zeta-}} \log G_r(Z^h_0, Z^h_{\zeta-}) \right\|^2 \mathbf{1}_{(\delta,\infty)}(\| Z^h_{\zeta-} - Z^h_0 \|) \right],$$

in the space of measurable functions, where conditioning on $\{\|Z^h_{\zeta-} - Z^h_0\| > \delta\}$ ensures that the expectation is finite. Given $s^{\widehat{\theta}} \approx \nabla \log \overleftarrow{h}$, this may be implemented as an empirical version of

$$\arg\min_{D \in \mathcal{D}} \mathbb{E}_{Z^h_0 \sim \alpha(dy|u)} \left[\left\| D(Z^h_{\zeta-}) + s^{\widehat{\theta}}(Z^h_{\zeta-}) - \nabla_{Z^h_{\zeta-}} \log G_r(Z^h_0, Z^h_{\zeta-}) \right\|^2 \mathbf{1}_{(\delta,\infty)}(\| Z^h_{\zeta-} - Z^h_0 \|) \right],$$

for some class of approximating functions $\mathcal{D}$. Given such minimiser $\widehat{D}_u$, we then define a conditional score estimator as

$$\widehat{s}(y \mid u) := \widehat{D}_u(y) + s^{\widehat{\theta}}(y).$$

The corresponding conditional generation procedure is summarised in Algorithm 5.

As an alternative to this fine-tuning approach, we may use the probabilistic characterisation (9) directly, provided that we have access to the likelihood $\pi(u \mid y)$ or a good approximation thereof (e.g., a classifier). To this end, we may first approximate the expected likelihood function

$$x \mapsto \mathbb{E}_x[\pi(u \mid Z^{\overleftarrow{h}}_{\zeta_-})] = \mathbb{E}_{Z^h_0 \sim \alpha}[\pi(u \mid Z^h_0) \mid Z^h_{\zeta_-} = x]$$

which minimises the L^2-loss

$$\ell_u \mapsto \mathbb{E}_{Z^h_0 \sim \alpha}\left[|\ell_u(Z^h_{\zeta_-}) - \pi(u \mid Z^h_0)|^2\right],$$

in $L^2(\alpha)$. Denote the solution of the associated nonparametric regression problem (based on a subset $\{y_1, \dots, y_{n_1}\}$ of the unconditional data set $\{y_1, \dots, y_n\}$ and some class of candidate likelihood functions $\mathcal{P}_u$) by $\widehat{\ell}_u$, and similarly, using the same reasoning as above based on (Christensen et al., 2025, Lemma A.1), estimate $\nabla \log \overleftarrow{h}(\cdot \mid u)$ by a minimiser of

$$s_u \mapsto \frac{1}{M} \sum_{j=1}^{M} \left\| s_u\big(z^{h,i_j}_{\zeta^{i_j}_-}\big) - \nabla_2 G_r(y_{i_j}, z^{h,i_j}_{\zeta^{i_j}_-}) \frac{\pi(u \mid y_{i_j})}{\widehat{\ell}_u(z^{h,i_j}_{\zeta^{i_j}_-}) \vee \varepsilon} \right\|^2 \mathbf{1}_{\left\{\|z^{h,i_j}_{\zeta^{i_j}_-} - y_{i_j}\| > \delta\right\}}$$

within some approximating function class $\mathcal{S}_u$. Here, δ, ε are regularisation parameters, the indices i_j are uniformly drawn with replacement from the remaining data indices $\{n_1 + 1, \dots, n\}$ and $z^{h,i_j}_{\zeta^{i_j}_-}$ is a simulated value of the h-transform Z^h at its lifetime when initialised in the data point y_{i_j}. Note here that the estimation procedure neither requires an estimator of $\nabla \log \overleftarrow{h}$ nor explicit access to samples of $\alpha(dy \mid u)$. The separate estimation of the denominator ℓ_u may however cause numerical challenges. These are more pronounced if $\pi(u \mid Z^h_0)$ is concentrated around 0 under $\mathbb{P}_{Z^h_0 \sim \alpha}$ corresponding to scarce data availability for the condition $U = u$.

5.3 Distances to the distribution and anomaly detection

We now address the question of how to measure the distance of an input x from the data distribution α. The previous discussion provides a natural measure for this, namely the time it takes the forward process Y to transform the input x into a sample from α.

We can identify x as an anomaly if the average lifetime exceeds a certain threshold $\underline{T}$ using the Monte-Carlo procedure given in Algorithm 6.

Algorithm 6 Anomaly detection

Input: state x to be evaluated, learned score $s^{\widehat{\theta}}$ based on Algorithm 4, threshold $\underline{T} > 0$, # of Monte-Carlo runs N, precision parameters $\gamma, \varepsilon > 0$
Initialize $\overline{\zeta} = 0$
if $\|s^{\widehat{\theta}}(x)\| < \gamma d \varepsilon^{-1}$ **then**
  for $i = 1$ **to** N **do**
    Initialize $Y^{\widehat{\theta}}_0 = x$
    repeat
      Simulate next step of $Y^{\widehat{\theta}}$
    until $\|s^{\widehat{\theta}}(Y^{\widehat{\theta}}_t)\| \geq \gamma d \varepsilon^{-1}$
    Update $\overline{\zeta} \leftarrow \overline{\zeta} + t/N$
  end for
end
if $\overline{\zeta} > \underline{T}$ **then**
  Classify as anomaly
end

This can also serve as a criterion for determining whether new training data necessitates a modification of our trained model or not.

5.4 Class sampling and classification

The model can be naturally extended when the data can be decomposed into subclasses, i.e., when α is a mixture of distributions $\alpha_1, \ldots, \alpha_n$ with disjoint supports Ω^i. In this scenario, there is a corresponding decomposition $\overleftarrow{h} = \sum_{i=1}^n \overleftarrow{h}_i$. The lifetime ζ is then determined as the minimum of the first entry times ζ_i into the supports of the α_i. With unconditional sampling, this provides information about the specific class from which the sampled image originates.

The truly interesting aspect of class decomposition becomes evident when we consider the conditional setting as the previous observations induce the natural classification Algorithm 7 based on the (learned) unconditional model that records the class i responsible for killing the backward process initialised in x.

Algorithm 7 Classification

Input: state x to be classified, classes $i = 1, \ldots, K$, estimates $\widehat{\Omega^i}$ of class supports Ω^i, learned backward model $Y^{\widehat{\theta}}$ for $\widehat{\Omega_\varepsilon} = \bigcup_{i=1}^K (\widehat{\Omega^i})_\varepsilon$ based on Algorithm 3
if $x \notin \widehat{\Omega_\varepsilon}$ **then**
  Initialize $Y_0^{\widehat{\theta}} = x$
  repeat
    Simulate next step of $Y^{\widehat{\theta}}$
  until $Y_t^{\widehat{\theta}} \in (\widehat{\Omega^i})_\varepsilon$ for some $i \in \{1, \ldots, K\}$
else
  Find $i \in \{1, \ldots, K\}$ s.t. $i \in \arg\min_{i=1,\ldots,K} d(x, \widehat{\Omega^i})$
end
Output: class i

The presented conditional diffusion model thus naturally enables classification through transfer learning. Refinements are also conceivable. For instance, by starting several runs from x we can estimate the conditional class distribution, which allows the construction of statistical tests. Analogously to the case of anomaly detection, the average lifetime $\overline{\zeta}$ can then be regarded as a measure of reliability.

5.5 General transfer learning

The idea of transfer learning presented in the previous section can be applied more generally. For instance, if the classes are not predefined, but clusters are to be learned from the data, the typical approach is to perform clustering with known algorithms only on the support of α, which is a significantly lower-dimensional task. Similar approaches may be possible for tasks in Reinforcement Learning.

6 Conclusion

We have introduced a novel class of diffusion models grounded in Doob's h-transform. This theoretical framework provides a general and adaptable foundation for time-homogeneous diffusion models that eliminates the necessity to model artificial time dependencies and thereby enhances interpretability. A key innovation in our model is the introduction of the polarity hypothesis, which closely parallels the manifold hypothesis in machine learning. This property enables an intuitive and efficient mechanism for determining the termination of the denoising process, addressing a crucial challenge in implementing our generative model.

Another notable feature of our framework is a methodological simplification for learning the dynamics of the denoising process. Unlike conventional diffusion models that require estimating temporally inhomogeneous dynamics, we must learn a time-*independent* backward drift, which is achieved by a denoising score matching procedure. While this comes at a cost of having to learn the now random simulation time for the generative process as well, we argue that no separate estimation strategy is needed for this purpose, but that a simple explosion criterion for the estimated backward drift along the generated path yields an adaptive termination rule.

The time-homogeneous nature of the model opens up numerous opportunities for practical applications. While its utility will depend on the specific requirements of each use case, the theoretical framework established here provides a robust foundation for experimental exploration, particularly in the context of transfer learning.

References

Eddie Aamari and Clément Levrard. Stability and minimax optimality of tangential Delaunay complexes for manifold reconstruction. *Discrete Comput. Geom.*, 59(4):923–971, 2018. ISSN 0179-5376,1432-0444.

Eddie Aamari and Clément Levrard. Nonasymptotic rates for manifold, tangent space and curvature estimation. *Ann. Statist.*, 47(1):177–204, 2019.

Eddie Aamari, Catherine Aaron, and Clément Levrard. Minimax boundary estimation and estimation with boundary. *Bernoulli*, 29(4):3334–3368, 2023.

Vlad Bally. Approximation models for the continuous additive functionals of multidimensional brownian motion. *Stochastic processes and their applications*, 32(2):331–345, 1989.

Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL `https://openreview.net/forum?id=oYIjw37pTP`.

Gérard Biau, Benoît Cadre, and Bruno Pelletier. Exact rates in density support estimation. *J. Multivariate Anal.*, 99(10):2185–2207, 2008.

Andrei N. Borodin and Paavo Salminen. *Handbook of Brownian Motion – Facts and Formulae*. Probability and its Applications. Birkhäuser Verlag, Basel, second edition, 2002. ISBN 3-7643-6705-9.

Bradley CA Brown, Anthony L. Caterini, Brendan Leigh Ross, Jesse C Cresswell, and Gabriel Loaiza-Ganem. Verifying the union of manifolds hypothesis for image data. In *International Conference on Learning Representations*, 2023.

Sören Christensen and Boy Schultz. On the existence of Markovian randomized equilibria in Dynkin games of war-of-attrition-type. *arXiv preprint arXiv:2406.09820*, 2024.

Sören Christensen, Jan Kallsen, Claudia Strauch, and Lukas Trottner. Model-free filtering in high dimensions via projection and score-based diffusions, 2025. URL `https://arxiv.org/abs/2510.23197`.

Kai Lai Chung and John B Walsh. *Markov processes, Brownian motion, and time symmetry*, volume 249. Springer Science & Business Media, 2005.

Antonio Cuevas and Ricardo Fraiman. A plug-in approach to support estimation. *Ann. Statist.*, 25(6): 2300–2312, 1997. ISSN 0090-5364,2168-8966.

Antonio Cuevas and Alberto Rodríguez-Casal. On boundary estimation. *Adv. in Appl. Probab.*, 36(2): 340–354, 2004. ISSN 0001-8678,1475-6064.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Alexander Denker, Francisco Vargas, Shreyas Padhy, Kieran Didi, Simon V Mathis, Riccardo Barbano, Vincent Dutordoir, Emile Mathieu, Urszula Julia Komorowska, and Pietro Lio. DEFT: Efficient fine-tuning of diffusion models by learning the generalised h-transform. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL `https://openreview.net/forum?id=AKBTFQhCjm`.

Luc Devroye and Gary L. Wise. Detection of abnormal behavior via nonparametric estimation of the support. *SIAM J. Appl. Math.*, 38(3):480–488, 1980.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Kieran Didi, Francisco Vargas, Simon Mathis, Vincent Dutordoir, Emile Mathieu, Urszula Julia Komorowska, and Pietro Lio. A framework for conditional diffusion modelling with applications in motif scaffolding for protein design. In *NeurIPS 2023 Workshop on New Frontiers of AI for Drug Discovery and Development*, 2023.

DLMF. NIST Digital Library of Mathematical Functions. `https://dlmf.nist.gov/`, Release 1.2.2 of 2024-09-15. URL `https://dlmf.nist.gov/`. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds.

A. Erdélyi, W. Magnus, F. Oberhettinger, and F.G. Tricomi. *Tables of Integral Transforms*. McGraw-Hill, New York, 1954.

Zach Evans, CJ Carr, Josiah Taylor, Scott H Hawley, and Jordi Pons. Fast timing-conditioned latent audio diffusion. In *International Conference on Machine Learning*, 2024.

Avner Friedman. *Stochastic differential equations and applications. Vol. 1*, volume Vol. 28 of *Probability and Mathematical Statistics*. Academic Press [Harcourt Brace Jovanovich, Publishers], New York-London, 1975.

Christopher R. Genovese, Marco Perone-Pacifico, Isabella Verdinelli, and Larry Wasserman. Minimax manifold estimation. *J. Mach. Learn. Res.*, 13:1263–1291, 2012.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.

Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 8633–8646, 2022.

Arlene K. H. Kim and Harrison H. Zhou. Tight minimax rates for manifold estimation under Hausdorff loss. *Electron. J. Stat.*, 9(1):1562–1582, 2015.

A. P. Korostelëv and A. B. Tsybakov. *Minimax theory of image reconstruction*, volume 82 of *Lecture Notes in Statistics*. Springer-Verlag, New York, 1993.

N. N. Lebedev. *Special functions and their applications*. Dover publications, New York, 1972.

Gabriel Loaiza-Ganem, Brendan Leigh Ross, Rasa Hosseinzadeh, Anthony L. Caterini, and Jesse C. Cresswell. Deep generative models through the lens of the manifold hypothesis: A survey and new connections. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.

Yunqian Ma and Yun Fu (eds.). *Manifold learning theory and applications*. CRC Press, Boca Raton, FL, 2012. ISBN 978-1-4398-7109-6.

E. Mammen and A. B. Tsybakov. Asymptotical minimax recovery of sets with smooth boundaries. *Ann. Statist.*, 23(2):502–524, 1995.

Peter Mörters and Yuval Peres. *Brownian motion*, volume 30. Cambridge University Press, 2010.

Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *International Conference on Artificial Intelligence and Statistics*, pp. 4474–4484. PMLR, 2020.

Bernt Øksendal and Agnes Sulem. Stochastic control of jump diffusions. In *Applied Stochastic Control of Jump Diffusions*, pp. 93–155. Springer, 2019.

Stefano Peluchetti. Diffusion bridge mixture transports, Schrödinger bridge problems and generative modeling. *J. Mach. Learn. Res.*, 24:Paper No. [374], 51, 2023. ISSN 1532-4435,1533-7928.

Phil Pope, Chen Zhu, Ahmed Abdelkader, Micah Goldblum, and Tom Goldstein. The intrinsic dimension of images and its impact on learning. In *International Conference on Learning Representations*, 2021.

S. Ramasubramanian. Hitting of submanifolds by diffusions. *Probab. Theory Related Fields*, 78(1):149–163, 1988. ISSN 0178-8051,1432-2064.

Herbert Robbins. An empirical Bayes approach to statistics. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1954–1955, vol. I*, pp. 157–163. Univ. California Press, Berkeley-Los Angeles, Calif., 1956.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. In *Advances in Neural Information Processing Systems*, volume 36, pp. 62183–62223, 2023.

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.

Jan Pawel Stanczuk, Georgios Batzolis, Teo Deveney, and Carola-Bibiane Schönlieb. Diffusion models encode the intrinsic dimension of data manifolds. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 46412–46440. PMLR, 21–27 Jul 2024. URL `https://proceedings.mlr.press/v235/stanczuk24a.html`.

L. Stoica. The approximation of additive functionals of diffusion processes. *Stochastics Stochastics Rep.*, 38 (3):185–200, 1992.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Comput.*, 23(7): 1661–1674, 2011. ISSN 0899-7667,1530-888X.

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, Sergey Ovchinnikov, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.

Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.

Mao Ye, Lemeng Wu, and Qiang Liu. First hitting diffusion models for generating manifold, graph and categorical data. *Advances in Neural Information Processing Systems*, 35:27280–27292, 2022.

Zheng Zhao, Ziwei Luo, Jens Sjölund, and Thomas B Schön. Conditional sampling within generative diffusion models. *arXiv preprint arXiv:2409.09650*, 2024.

A Doob's h-transform and lifetime of a diffusion process

The following is based on Chung & Walsh (2005), which provides a level of generality necessary for our purposes. For a more explicit discussion in the univariate case, see also Borodin & Salminen (2002).

We first note that the Green kernel can be found explicitly for the Brownian motion:

Remark 9. Let $Z = W$ be a standard d-dimensional Brownian motion, m Lebesgue measure, then

$$p_t(x, y) = (2\pi t)^{-d/2} \exp\left(-\frac{|x-y|^2}{2t}\right).$$

In this case, (Erdélyi et al., 1954, p. 146) yields

$$G_r(x, y) = G_r(|x-y|) = (2\pi)^{-d/2} 2 \left(\frac{|x-y|^2}{2r}\right)^{(2-d)/4} K_{(d-2)/2}\left(|x-y|\sqrt{2r}\right),$$

where K_ν is the modified Bessel function of the second kind as defined in Lebedev (1972), p. 109. The function $G_r(\cdot)$ has a pole in 0 and is decreasing.

Now we give the remaining proofs:

Proof of Proposition 1. The finiteness of ζ follows from (Chung & Walsh, 2005, Theorem 13.50):

$$\mathbb{P}_x(\zeta < \infty) = \int_0^\infty \int \mathrm{e}^{-rt} \frac{1}{h(x)} p_t(x, y)\kappa(dy)dt = \frac{1}{h(x)} \int G_r(x, y)\kappa(dy) = 1.$$

The second statement follows by using the generator identity

$$\mathbb{A}^h f(x) = \frac{1}{h(x)} \left(\mathbb{A}(h(x)f(x)) - rh(x)f(x)\right),$$

taking into account that h is r-harmonic (i.e., $(\mathbb{A} - r)h(x) = 0$) outside the support of β.

The third statement follows from (Chung & Walsh, 2005, Theorem 13.39) (after correction of an obvious typo), where the special case $x = x_0$ yields

$$\mathbb{P}_{x_0}(Z^h_{\zeta-} \in dy) = \frac{G_r(x_0, y)}{h(x_0)}\kappa(dy) = G_r(x_0, y)\kappa(dy) = \beta(dy).$$

□

Proof of Proposition 3. The first claim holds by combining (Chung & Walsh, 2005, Theorem 13.34) with Proposition 1. The other two are direct consequences of Proposition 1. □

B Stochastic control

For the proof of Proposition 2, we start with the following

Lemma 10. *Let h be as above and $g(x) := -\log h(x)$. Then, outside of the support of β, it holds that*

$$\mathbb{A}g(x) + r - \frac{1}{2}\nabla g(x)^\top \sigma(x)\sigma^\top(x)\nabla g(x) = 0.$$

Proof. Outside of $\operatorname{supp}\kappa$, it holds that $\mathbb{A}h(x) - rh(x) = 0$, and a straightforward calculation shows that

$$\mathbb{A}g(x) = -r + \frac{1}{2}\nabla g(x)^\top \sigma(x)\sigma^\top(x)\nabla g(x).$$

□

Now, we follow the standard verification approach to stochastic stopping and control problems, see Øksendal & Sulem (2019), Chapter 5, where we assume the appropriate standard regularity assumptions for the objects without explicitly mentioning them. Outside of $\operatorname{supp}\kappa$, using the notation

$$\mathbb{A}^u v(x) = \langle \sigma(x)u + b(x), \nabla v(x)\rangle + \frac{1}{2}\operatorname{Tr}\left[\Sigma(x)\nabla^2 v(x)\right],$$

where $\Sigma(x) = \sigma(x)\sigma(x)^\top$, it holds

$$\begin{aligned}\inf_u \left(\mathbb{A}^u v(x) + k(u)\right) =& \langle b(x), \nabla v(x)\rangle + \frac{1}{2}\operatorname{Tr}\left[\Sigma(x)\nabla^2 v(x)\right] + r + \inf_u \left(\langle \sigma(x)u, \nabla v(x)\rangle + \frac{1}{2}\|u\|^2\right)\\ =& \langle b(x), \nabla v(x)\rangle + \frac{1}{2}\operatorname{Tr}\left[\Sigma(x)\nabla^2 v(x)\right] + r - \frac{1}{2}\|\sigma(x)^\top \nabla v(x)\|^2\end{aligned}$$

with minimiser $u = -\sigma(x)^\top \nabla v(x)$. Using $v(x) = g(x) = -\log h(x)$ and Lemma 10, we see that $g(x)$ fulfills the HJB equation for problem (1). Following the steps in the verification procedure in Øksendal & Sulem (2019), p. 92ff., it is easily seen that this solution to the HJB equation is indeed a solution to (1), proving the first part of Proposition 2.

Now, we study the connection to the KL divergence and write $Y = Z^h$. Note that

$$\left.\frac{d\mathbb{P}^{Z^u}}{d\mathbb{P}^Y}\right|_{\mathcal{F}_\zeta} = \exp\left(\frac{1}{2}\int_0^\zeta \|u_t\|^2 dt - M_\zeta\right),$$

for some (local) martingale M with $M_0 = 0$. For the killed process Y, by the definition of the h-transform,

$$\left.\frac{d\mathbb{P}^{Z}}{d\mathbb{P}^Y}\right|_{\mathcal{F}_\zeta} = \frac{\mathrm{e}^{r\zeta} h(Z_0)}{h(Z_\zeta)}.$$

This yields

$$\log \left.\frac{d\mathbb{P}^{Z^u}}{d\mathbb{P}^Y}\right|_{\mathcal{F}_\zeta} = \log \left.\frac{d\mathbb{P}^{Z^u}}{d\mathbb{P}^Z}\frac{d\mathbb{P}^{Z}}{d\mathbb{P}^Y}\right|_{\mathcal{F}_\zeta} = \frac{1}{2}\int_0^\zeta \|u_t\|^2 dt + r\zeta + \log h(Z_0^u) - \log h(Z_\zeta^u) + M_\zeta.$$

We obtain that

$$\mathrm{KL}(\mathbb{P}_x^{Z^u}\big|_{\mathcal{F}_\zeta} \,\|\, \mathbb{P}_x^{Y}\big|_{\mathcal{F}_\zeta}) = \mathbb{E}_x \log \left.\frac{d\mathbb{P}^{Z^u}}{d\mathbb{P}^Z}\right|_{\mathcal{F}_\zeta} = J(u,x) + \log h(x) \geq v(x) + \log h(x) = 0$$

with equality for $u = u^*$, so that the KL divergence is the variational gap in problem (1).

C Score matching and manifold learning

Proof of Proposition 5. It holds that

$$\begin{aligned}\mathcal{L}_{\mathrm{ex}}(\theta) &= \mathbb{E}\Big[\int_0^{\zeta_\varepsilon} \|\nabla_y \log \overleftarrow{h}(\overleftarrow{Z_t^h}) - s_\theta(\overleftarrow{Z_t^h})\|^2\, dt\Big]\\ &= \mathbb{E}\Big[\int_0^{\zeta_\varepsilon} |s_\theta(\overleftarrow{Z_t^h})|^2\, dt\Big] - 2\mathbb{E}\Big[\int_0^{\zeta_\varepsilon} \langle \nabla_y \log \overleftarrow{h}(\overleftarrow{Z_t^h}), s_\theta(\overleftarrow{Z_t^h})\rangle\, dt\Big] + C,\end{aligned} \tag{10}$$

where C is independent of θ. Let $\widetilde{\beta}(dy) = \mathbb{P}_\alpha(Z^h_{\zeta-} \in dy)$. Since the time reversal of the h-transform $Z^{\overleftarrow{h}}$ started in $\widetilde{\beta}$ is equal in law to the h-transform Z^h started in α, we obtain from Proposition 3 that

$$\frac{G_r(x,y)}{\overleftarrow{h}(x)}\,\widetilde{\beta}(dx) = h(y)\frac{G_r(y,x)}{h(y)\overleftarrow{h}(x)}\,\widetilde{\beta}(dx) = h(y)\,\mathbb{P}_y(Z^h_{\zeta-} \in dx).$$

Using this and assuming sufficient integrability properties that allow the application of Fubini's theorem and pulling the derivative inside the integral, we can calculate for the second term as follows:

$$\begin{aligned}
\mathbb{E}\Big[\int_0^{\zeta} \big\langle \nabla_y \log \overleftarrow{h}(\overleftarrow{Z_t^h}), s_\theta(\overleftarrow{Z_t^h})\big\rangle\, dt\Big] &= \int\int \langle \nabla_y \log \overleftarrow{h}(y), s_\theta(y)\rangle \frac{\overleftarrow{h}(y)}{\overleftarrow{h}(x)} G_r(x,y)\, m(dy)\, \widetilde{\beta}(dx) \\
&= \int \langle \nabla_y \log \overleftarrow{h}(y), s_\theta(y)\rangle \overleftarrow{h}(y) h(y)\, m(dy) \\
&= \int \langle \nabla_y \overleftarrow{h}(y), s_\theta(y)\rangle h(y)\, m(dy) \\
&= \int\int \langle \nabla_y \frac{G_r(y,z)}{h(z)}, s_\theta(y)\rangle h(y)\, m(dy)\, \alpha(dz) \\
&= \int\int \langle \nabla_y \log G_r(y,z), s_\theta(y)\rangle \frac{h(y)}{h(z)} G_r(y,z)\, m(dy)\, \alpha(dz) \\
&= \int\int \langle \nabla_y \log G_r(y,z), s_\theta(y)\rangle \frac{h(y)}{h(z)} G_r(z,y)\, m(dy)\, \alpha(dz) \\
&= \int \mathbb{E}_z\Big[\int_0^{\zeta} \langle \nabla_{Z_t^h} \log G_r(Z_t^h, Z_0^h), s_\theta(Z_t^h)\rangle dt\Big]\, \alpha(dz) \\
&= \mathbb{E}_\alpha\Big[\int_0^{\zeta} \langle \nabla_{Z_t^h} \log G_r(Z_t^h, Z_0^h), s_\theta(Z_t^h)\rangle\, dt\Big].
\end{aligned}$$

This gives

$$\begin{aligned}
&\mathbb{E}\Big[\int_0^{\zeta_\varepsilon} \big\langle \nabla_y \log \overleftarrow{h}(\overleftarrow{Z_t^h}), s_\theta(\overleftarrow{Z_t^h})\big\rangle\, dt\Big] \\
&= \mathbb{E}\Big[\int_0^{\zeta} \big\langle \nabla_y \log \overleftarrow{h}(\overleftarrow{Z_t^h}), s_\theta(\overleftarrow{Z_t^h})\big\rangle\, dt\Big] - \mathbb{E}\Big[\int_{\zeta_\varepsilon}^{\zeta} \big\langle \nabla_y \log \overleftarrow{h}(\overleftarrow{Z_t^h}), s_\theta(\overleftarrow{Z_t^h})\big\rangle\, dt\Big] \\
&= \mathbb{E}\Big[\int_0^{\zeta} \langle \nabla_{Z_t^h} \log G_r(Z_t^h, Z_0^h), s_\theta(Z_t^h)\rangle\, dt\Big] - \mathbb{E}\Big[\int_0^{\sigma_\varepsilon} \big\langle \nabla_y \log \overleftarrow{h}(Z_t^h), s_\theta(Z_t^h)\big\rangle\, dt\Big] \\
&= \mathbb{E}\Big[\int_{\sigma_\varepsilon}^{\zeta} \langle \nabla_{Z_t^h} \log G_r(Z_t^h, Z_0^h), s_\theta(Z_t^h)\rangle\, dt\Big] + \mathbb{E}\Big[\int_0^{\sigma_\varepsilon} \big\langle \nabla_{Z_t^h} \log \tfrac{G_r(Z_t^h, Z_0^h)}{\overleftarrow{h}(Z_t^h)}, s_\theta(Z_t^h)\big\rangle\, dt\Big].
\end{aligned}$$

Moreover, the first term satisfies

$$\mathbb{E}\Big[\int_0^{\zeta_\varepsilon} \|s_\theta(\overleftarrow{Z_t^h})\|^2\, dt\Big] = \mathbb{E}\Big[\int_{\sigma_\varepsilon}^{\zeta} \|s_\theta(Z_t^h)\|^2\, dt\Big].$$

Substituting these results into (10), we obtain

$$\mathcal{L}_{\text{ex}}(\theta) = \mathbb{E}\Big[\int_{\sigma_\varepsilon}^{\zeta} \|\nabla_{Z_t^h} \log G_r(Z_t^h, Z_0^h) - s_\theta(Z_t^h)\|^2\, dt\Big] - 2\mathbb{E}\Big[\int_0^{\sigma_\varepsilon} \big\langle \nabla_{Z_t^h} \log \tfrac{G_r(Z_t^h, Z_0^h)}{\overleftarrow{h}(Z_t^h)}, s_\theta(Z_t^h)\big\rangle\, dt\Big] + C',$$

where C' is a constant independent of θ. □

Proof of Lemma 6. Using (2) and Proposition 3, we find for any $x \notin \operatorname{supp} \alpha$,

$$\begin{aligned}
\nabla \log \overleftarrow{h}(x) = \frac{1}{\overleftarrow{h}(x)} \int \nabla_x G_r(x,y) \frac{1}{h(y)}\, \alpha(dy) &= \int \nabla_x \log G_r(x,y) \frac{G_r(x,y)}{\overleftarrow{h}(x) h(y)}\, \alpha(dy) \\
&= \mathbb{E}_x\left[\nabla_x \log G_r(x, Z_{\zeta-}^{\overleftarrow{h}})\right].
\end{aligned}$$

By Proposition 3, the time reversed process $\overleftarrow{Z^h}$ has the same distribution as the diffusion process $Z^{\overleftarrow{h}}$ started in the forward terminal distribution $\beta_h = \mathbb{P}_\alpha(Z^h_{\zeta_-} \in \cdot)$. For any Borel sets A, B it follows

$$\begin{aligned}
\int_B \mathbb{P}_\alpha(Z^h_0 \in A \mid Z^h_{\zeta_-} = x)\, \mathbb{P}_\alpha(Z^h_{\zeta_-} \in dx) &= \mathbb{P}_\alpha(Z^h_0 \in A, Z^h_{\zeta_-} \in B) \\
&= \mathbb{P}_{Z^{\overleftarrow{h}}_0 \sim \beta_h}(Z^{\overleftarrow{h}}_{\zeta_-} \in A, Z^{\overleftarrow{h}}_0 \in B) \\
&= \int_B \mathbb{P}(Z^{\overleftarrow{h}}_{\zeta_-} \in A \mid Z^{\overleftarrow{h}}_0 = x)\, \beta_h(dx) \\
&= \int_B \mathbb{P}(Z^{\overleftarrow{h}}_{\zeta_-} \in A \mid Z^{\overleftarrow{h}}_0 = x)\, \mathbb{P}_\alpha(Z^h_{\zeta_-} \in dx),
\end{aligned}$$

and thus, $\mathbb{P}(Z^{\overleftarrow{h}}_{\zeta_-} \in \mathrm{d}y \mid Z^{\overleftarrow{h}}_0 = x)$ is a regular conditional probability for Z^h_0 given $Z^h_{\zeta_-}$ w.r.t. $\mathbb{P}_\alpha$. Consequently, uniqueness of regular conditional probabilities gives

$$\mathbb{P}_\alpha(Z^h_0 \in dy \mid Z^h_{\zeta_-} = x) = \mathbb{P}(Z^{\overleftarrow{h}}_{\zeta_-} \in dy \mid Z^{\overleftarrow{h}}_0 = x), \quad \text{for } \mathbb{P}_\alpha(Z^h_{\zeta_-} \in \cdot)\text{-a.e. } x.$$

This together with (4) and symmetry of G_r yields (5). □