

# UNDERDAMPED DIFFUSION BRIDGES WITH APPLICATIONS TO SAMPLING

Denis Blessing<sup>\*,1</sup>, Julius Berner<sup>\*,2,†</sup>, Lorenz Richter<sup>\*,3,4</sup>, Gerhard Neumann<sup>1,5</sup>

<sup>1</sup>Karlsruhe Institute of Technology, <sup>2</sup>NVIDIA, <sup>3</sup>Zuse Institute Berlin,

<sup>4</sup>dida Datens Schmiede GmbH, <sup>5</sup>FZI Research Center for Information Technology

## ABSTRACT

We provide a general framework for learning diffusion bridges that transport prior to target distributions. It includes existing diffusion models for generative modeling, but also underdamped versions with degenerate diffusion matrices, where the noise only acts in certain dimensions. Extending previous findings, our framework allows to rigorously show that score matching in the underdamped case is indeed equivalent to maximizing a lower bound on the likelihood. Motivated by superior convergence properties and compatibility with sophisticated numerical integration schemes of underdamped stochastic processes, we propose *underdamped diffusion bridges*, where a general density evolution is learned rather than prescribed by a fixed noising process. We apply our method to the challenging task of sampling from unnormalized densities without access to samples from the target distribution. Across a diverse range of sampling problems, our approach demonstrates state-of-the-art performance, notably outperforming alternative methods, while requiring significantly fewer discretization steps and no hyperparameter tuning.

## 1 INTRODUCTION

In this paper we propose a general diffusion-based framework for sampling from a density

$$p_{\text{target}} = \frac{\rho_{\text{target}}}{\mathcal{Z}}, \quad \mathcal{Z} := \int_{\mathbb{R}^d} \rho_{\text{target}}(x) \, dx, \quad (1)$$

where  $\rho_{\text{target}} \in C(\mathbb{R}^d, \mathbb{R}_{\geq 0})$  can be evaluated pointwise, but the normalization constant  $\mathcal{Z}$  is typically intractable. This task is of great practical relevance in the natural sciences, e.g., in fields such as molecular dynamics and statistical physics (Stoltz et al., 2010; Schopmans & Friederich, 2025), but also in Bayesian statistics (Gelman et al., 2013).

Recently, multiple approaches based on diffusion processes have been proposed, where the overall idea is to learn a stochastic process in such a way that it transports an easy prior distribution to the potentially complicated target over an artificial time. Typically, the process is defined as an ordinary Itô diffusion, in particular, demanding non-degenerate noise. In this work, we aim to generalize this setting to diffusion processes with degenerate noise. This is motivated by the following model from statistical physics.

Classical sampling approaches based on stochastic processes have been extensively conducted using some version of the *overdamped Langevin dynamics*

$$dX_s = \nabla \log p_{\text{target}}(X_s) \, ds + \sqrt{2} \, dW_s, \quad X_0 \sim p_{\text{prior}}, \quad (2)$$

whose stationary distribution is given by  $p_{\text{target}}$  (under some rather mild technical assumptions on the target and on the prior  $p_{\text{prior}}$ ). Furthermore, we can define an extended dynamics by introducing an additional variable, bringing the so-called *underdamped Langevin dynamics*

$$dX_s = Y_s \, ds, \quad X_0 \sim p_{\text{prior}}, \quad (3a)$$

$$dY_s = (\nabla \log p_{\text{target}}(X_s) - Y_s) \, ds + \sqrt{2} \, dW_s, \quad Y_0 \sim \mathcal{N}(0, \text{Id}), \quad (3b)$$

<sup>\*</sup>Equal contribution. <sup>†</sup>Work partially done at Caltech.

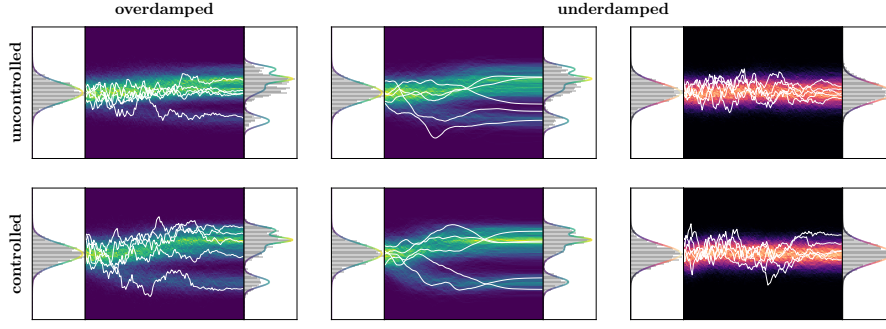


Figure 1: Illustration of uncontrolled (see (2) and (3)) and controlled (see (4) and (13)) diffusion processes in the overdamped and underdamped regime, transporting the Gaussian prior distribution to the target. For the underdamped case, we show both the positional coordinate (left/blue) as well as the velocity (right/black). While the underdamped version enjoys better convergence guarantees, both uncontrolled diffusions only converge asymptotically. Learning the control, we can achieve convergence in finite time.

where now the stationary distribution is given by  $\tau(x, y) := p_{\text{target}}(x)\mathcal{N}(y; 0, \text{Id})$  (and  $\pi(x, y) := p_{\text{prior}}(x)\mathcal{N}(y; 0, \text{Id})$  can be defined as an extended prior distribution). Intuitively, the  $y$ -variable can be interpreted as a velocity, which is coupled to the space variable  $x$  via Hamiltonian dynamics.

While both (2) and (3) converge to the desired (extended) target distribution after infinite time, their convergence speed can be exceedingly slow, in particular for multimodal targets (Eberle et al., 2019). At the same time, it has been observed numerically that the underdamped version can be significantly faster (Stoltz et al., 2010). This might be attributed to the fact that the Brownian motion is only indirectly coupled to the space variable, leading to smoother paths of  $X$  and lower discretization error in numerical integrators (since  $\nabla \log p_{\text{target}}$  only depends on  $X$ , but not on  $Y$ ). In particular, for smooth and strongly log-concave<sup>1</sup> targets, the number of steps to obtain KL divergence  $\varepsilon$  can be reduced from  $\tilde{\mathcal{O}}(d/\varepsilon^2)$  to  $\tilde{\mathcal{O}}(\sqrt{d}/\varepsilon)$  (Ma et al., 2021).

The idea of *learned* diffusion-based sampling is to reach convergence to multimodal targets after finite time. In particular, for overdamped diffusion models, the convergence rate can be shown to match the one of Langevin dynamics *without* the need for log-concavity assumptions as long as the learned model exhibits sufficiently small approximation error (Chen et al., 2022). In the overdamped setting, this can be readily formulated as adding a control function to the dynamics (2),

$$dX_s = (\nabla \log p_{\text{target}}(X_s) + u(X_s, s)) ds + \sqrt{2} dW_s, \quad (4)$$

where the task is to learn  $u \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$  as to reach  $X_T \sim p_{\text{target}}$  (Richter & Berner, 2024; Vargas et al., 2024); see Figure 1 for an illustration. It is now natural to ask the question whether we can use the same control ideas to the (typically better behaved) underdamped dynamics (3). Motivated by this guiding question this paper includes the following:

- **Controlled diffusions with degenerate noise:** Building on previous work based on path space measures, we generalize diffusion-based sampling to processes with degenerate noise, in particular including controlled underdamped Langevin equations (Section 2).
- **Underdamped methods in generative modeling:** This framework can be used to derive and analyze underdamped methods in generative modeling. We derive the ELBO and variational gap for diffusion bridges where both forward and reverse-time processes are learned (Appendix A.3).
- **Novel underdamped samplers:** Moreover, our framework culminates in underdamped versions of existing sampling methods and in particular in the novel *underdamped diffusion bridge sampler* (Section 3). In extensive numerical experiments, we can demonstrate significantly improved performance of our method.
- **Numerical integrators and ablation studies:** We provide careful ablation studies of our improvements, including the benefits of the novel integrators for controlled diffusion bridges as well as end-to-end training of hyperparameters (Section 4). We note that the latter eliminates the need for tuning and also significantly improves existing methods in the overdamped regime.

<sup>1</sup>Or, more general, log-concave outside of a region.

## 1.1 RELATED WORK

Many approaches to sampling problems build an augmented target by using a sequence of densities bridging the prior and target distribution and defining forward and backward kernels to approximately transition between the densities, often referred to as *annealed importance sampling* (AIS) (Neal, 2001). For instance, taking uncorrected overdamped Langevin kernels leads to *Unadjusted Langevin Annealing* (ULA) (Thin et al., 2021; Wu et al., 2020). Moreover, *Monte Carlo Diffusion* (MCD) optimizes the extended target distribution in order to minimize the variance of the marginal likelihood estimate (Doucet et al., 2022b). Going one step further, *Controlled Monte Carlo Diffusion* (CMCD) (Vargas et al., 2024) proposes an objective for directly optimizing the transition kernels to match the annealed density and *Sequential Controlled Langevin Diffusion* (SCLD) (Chen et al., 2025) connects this attempt with resampling strategies.

Furthermore, methods relying on diffusion processes that prescribe the backward transition kernel have recently been suggested. For instance, those include the *Path Integral Sampler* (PIS) (Zhang & Chen, 2022; Vargas et al., 2023b; Richter, 2021), *Time-Reversed Diffusion Sampler* (DIS) (Berner et al., 2024), *Diffusion generative flow samplers* (DGFS) (Zhang et al., 2024), *Denoising Diffusion Sampler* (DDS) (Vargas et al., 2023a), as well as the *Particle Denoising Diffusion Sampler* (Phillips et al., 2024). For those diffusion-based samplers, the optimal forward transition corresponds to the score of the current density, which can also be learned via its associated Fokker-Planck equation (Sun et al., 2024) or its representation via the Feynman-Kac formula (Akhound-Sadegh et al., 2024). Finally, there are methods learning both kernels separately, e.g., the *(Diffusion) Bridge<sup>2</sup> Sampler* (DBS) (Richter & Berner, 2024).

For some of the above methods improved convergence has been observed when using underdamped versions or Hamiltonian dynamics, which can be viewed as a form of momentum. In particular, ULA has been extended to *Uncorrected Hamiltonian Annealing* (UHA) (Geffner & Domke, 2021; Zhang et al., 2021), MCD has been extended to *Langevin Diffusion Variational Inference* (LDVI) (Geffner & Domke, 2023), and the works on DDS and CMCD also proposed underdamped versions.

Our framework in principle encompasses all these works as special cases (see Table 2). Moreover, we can easily derive novel algorithms using our framework, ranging from an underdamped version of DIS to an underdamped version of the Diffusion Bridge Sampler (Appendix A.10). Further, our unifying framework allows us to easily share integrators and training techniques for the different methods. First, we improve tuning for all considered methods by learning hyperparameters end-to-end, overall resulting in better performance (Figure 5). Second, we advance underdamped methods with our novel integrator (Figure 4 and Figure 9). Third, we show how to scale DBS to more complex targets by using a suitable parametrization (Table 4 & Figure 11) and divergence-free training objective (Proposition 2.3 vs. Proposition A.13). This makes our underdamped version of DBS a state-of-the-art method across a wide range of tasks (Table 1, Figure 3, & Table 3).

## 2 DIFFUSION BRIDGES WITH DEGENERATE NOISE

In this section, we lay the theoretical foundations for diffusion bridges with degenerate noise, extending the frameworks suggested in Richter & Berner (2024) and Vargas et al. (2024). Relating to the example from the introduction, we note that this includes cases where the noise only appears in certain dimensions of the stochastic process and in particular includes underdamped dynamics. We refer to Appendices A.1 and A.2 for a summary of our notation and assumptions.

The general idea of diffusion bridges is to learn a stochastic process that transports a given prior density to the prescribed target. This can be achieved via the concept of time-reversal (see, e.g., Figure 2). To this end, let us define the forward and reverse-time SDEs

$$dZ_s = (f + \eta u)(Z_s, s) ds + \eta(s) \vec{d}W_s, \quad Z_0 \sim \pi, \quad (5)$$

$$dZ_s = (f + \eta v)(Z_s, s) ds + \eta(s) \tilde{d}W_s, \quad Z_T \sim \tau, \quad (6)$$

on the state space  $\mathbb{R}^D$ , where  $\vec{d}W_s$  and  $\tilde{d}W_s$  denote forward and backward Brownian motion increments (see Appendix A.1 for details), respectively, both living in dimension  $d \leq D$ . The function  $f \in C(\mathbb{R}^D \times [0, T], \mathbb{R}^D)$  is typically fixed and maps to the full space, whereas the control functions  $u, v \in C(\mathbb{R}^D \times [0, T], \mathbb{R}^d)$  will be learned as to approach the desired bridge. In our setting,

<sup>2</sup>We clarify the connection to *Schrödinger bridges* and other *diffusion bridges* in Remark A.3.




Figure 2: Illustration of our general framework for learning diffusion bridges with degenerate noise. **Left:** We consider forward and reverse-time SDEs (see (5) and (6)), starting at the (extended) prior  $\pi$  and target  $\tau$  and being controlled by  $u$  and  $v$ , respectively. **Middle:** We learn optimal controls  $u^*$  and  $v^*$  of the corresponding SDEs (5) and (6), respectively, by minimizing a suitable divergence  $D$  between the associated path measures  $\tilde{\mathbb{P}}^{u, \pi}$  and  $\tilde{\mathbb{P}}^{v, \tau}$  on the SDE trajectories (Problem 2.1). **Right:** In general, the optimal controls are not unique and we depict an alternative solution (transparent). However, every solution leads to a perfect time-reversal and, in particular, represents a *diffusion bridge* with the correct marginals  $\pi$  at time  $s = 0$  and  $\tau$  at time  $s = T$ . We note that the trajectories can be smooth since we allow for degenerate diffusion coefficients, where the Brownian motion noise only acts in certain dimensions.

the noise coefficient  $\eta \in C([0, T], \mathbb{R}^{D \times d})$  may be degenerate in the sense that it has the shape  $\eta = (\mathbf{0}, \sigma)^\top$ , where  $\mathbf{0} \in \mathbb{R}^{D-d \times d}$  and  $\sigma \in C([0, T], \mathbb{R}^{d \times d})$  is assumed to be invertible for each  $t \in [0, T]$ . Importantly, the (scaled) control functions and the (scaled) Brownian motions operate in the same dimensions. For simplicity, we assume that  $f = (\hat{f}, \tilde{f})^\top$  with  $\tilde{f} \in C(\mathbb{R}^D \times [0, T], \mathbb{R}^d)$  and  $\hat{f} \in C(\mathbb{R}^D \times [0, T], \mathbb{R}^{D-d})$ , where  $\hat{f}$  only depends on the last  $d$  coordinates of its  $D$ -dimensional inputs. Referring to the underdamped Langevin equation (3), we may think of  $Z = (X, Y)^\top$ .

The general idea is to learn the control functions  $u$  and  $v$  such that the two processes defined in (5) and (6) are time reversals of each other. This task can be approached via measures on the space of continuous trajectories  $C([0, T], \mathbb{R}^D)$ , also called *path space* (see Appendix A.1 for details). To this end, let us denote by  $\tilde{\mathbb{P}}^{u, \pi}$  the measure of the forward process (5) and by  $\tilde{\mathbb{P}}^{v, \tau}$  the measure of the backward process (6). We may consider the following optimization problem; see also Figure 2.

**Problem 2.1** (Time-reversal). Let  $D : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$  be a divergence and let  $\mathcal{U} \subset C(\mathbb{R}^D \times [0, T], \mathbb{R}^d)$  be the set of admissible controls<sup>3</sup>. We aim to identify optimal controls  $u^*, v^*$  such that

$$u^*, v^* \in \arg \min_{u, v \in \mathcal{U} \times \mathcal{U}} D(\tilde{\mathbb{P}}^{u, \pi} | \tilde{\mathbb{P}}^{v, \tau}). \quad (7)$$

Clearly, if we can drive the divergence in (7) to zero, we have solved the time-reversal task and it readily follows for the time marginals<sup>4</sup> that  $\tilde{\mathbb{P}}_T^{u^*, \pi} = \tau$  and  $\tilde{\mathbb{P}}_0^{v^*, \tau} = \pi$ . We note that optimality in Problem 2.1 can be expressed by a local condition on the level of time marginals for any time in between prior and target.

**Lemma 2.2** (Nelson’s relation). *The following statements are equivalent:*

- (i)  $\tilde{\mathbb{P}}^{u, \pi} = \tilde{\mathbb{P}}^{v, \tau}$ .
- (ii)  $u(\cdot, t) - v(\cdot, t) = \eta^\top(t) \nabla \log \tilde{\mathbb{P}}_t^{u, \pi}$  for all  $t \in (0, T]$  and  $\tilde{\mathbb{P}}_T^{u, \pi} = \tau$ .
- (iii)  $u(\cdot, t) - v(\cdot, t) = \eta^\top(t) \nabla \log \tilde{\mathbb{P}}_t^{v, \tau}$  for all  $t \in [0, T)$  and  $\tilde{\mathbb{P}}_0^{v, \tau} = \pi$ .

*Proof.* The equivalence follows from the classical Nelson relation (Nelson, 1967; Anderson, 1982; Föllmer, 1986), which also holds for degenerate  $\eta$ ; cf. Haussmann & Pardoux (1986); Millet et al. (1989); Chen et al. (2022).  $\square$

However, we cannot directly use Lemma 2.2 to approach Problem 2.1 since the marginals  $\tilde{\mathbb{P}}_t^{v, \tau}$  and  $\tilde{\mathbb{P}}_t^{u, \pi}$  are typically intractable. To turn (7) into a feasible optimization problem, we need to explicitly compute divergences between path measures, which (in analogy to likelihood ratios) typically

<sup>3</sup>We refer to Appendix A.2 for assumptions on  $\mathcal{U}$ . We note that a *divergence*  $D$  is zero if and only if both its arguments coincide (in the space of probability measures  $\mathcal{P}$  on  $C([0, T], \mathbb{R}^D)$ ; see Appendix A.1).

<sup>4</sup>We denote the marginal of a path space measure  $\mathbb{P}$  at time  $t \in [0, T]$  by  $\mathbb{P}_t$ . Similarly, we denote by  $\mathbb{P}_{s|t}$  the conditional distribution of  $\mathbb{P}_s$  given  $\mathbb{P}_t$ ; see Appendix A.1.

involves the Radon-Nikodym derivative between those measures. A key observation is that this can indeed be achieved for forward and reverse-time processes, as stated in the following proposition.

**Proposition 2.3** (Likelihood of path measures). *Let  $\eta^+(s)$  be the pseudoinverse of  $\eta(s)$  for each  $s \in [0, T]$ . Then for  $\vec{\mathbb{P}}^{u,\pi}$ -almost every  $Z \in C([0, T], \mathbb{R}^D)$  it holds that*

$$\begin{aligned} \log \frac{d\vec{\mathbb{P}}^{u,\pi}}{d\vec{\mathbb{P}}^{v,\tau}}(Z) &= \log \frac{\pi(Z_0)}{\tau(Z_T)} - \frac{1}{2} \int_0^T \|\eta^+ f + u\|^2(Z_s, s) ds + \frac{1}{2} \int_0^T \|\eta^+ f + v\|^2(Z_s, s) ds \\ &\quad + \int_0^T (\eta^+ f + u)(Z_s, s) \cdot \eta^+(s) dZ_s - \int_0^T (\eta^+ f + v)(Z_s, s) \cdot \eta^+(s) dZ_s. \end{aligned}$$

*Proof.* Following Vargas et al. (2024, proof of Proposition 2.2), the proof applies the Girsanov theorem to the forward and reverse-time processes; see Appendix A.6.  $\square$

We refer to Proposition A.13 in the appendix for an alternative version of Proposition 2.3, which, for non-degenerate noise, has been used to define previous diffusion bridge samplers (Richter & Berner, 2024). However, this version relies on a divergence instead of backward stochastic processes, which renders it prohibitive for high dimensions and does not guarantee an ELBO after discretization; see also Remark A.7.

It is important to highlight that the optimization task in Problem 2.1 allows for infinitely many solutions. For numerical applications one may either accept this non-uniqueness (cf. Richter & Berner (2024)) or add additional constraints, such as regularizers (leading to, e.g., the so-called *Schrödinger bridge* (Vargas et al., 2021; De Bortoli et al., 2021)), a prescribed density evolution (Vargas et al., 2024) or a fixed noising process (Berner et al., 2024). Those different choices lead to different algorithms, for which we can now readily state corresponding degenerate (and thus underdamped) versions using our framework, see Appendix A.10.

**Divergences and loss functions for sampling.** In order to solve Problem 2.1, we need to choose a divergence  $D$ , in turn leading to a loss function  $\mathcal{L} : \mathcal{U} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$  via  $\mathcal{L}(u, v) := D(\vec{\mathbb{P}}^{u,\pi} | \vec{\mathbb{P}}^{v,\tau})$ . A common choice is the *Kullback-Leibler* (KL) divergence, which brings the loss

$$\mathcal{L}_{\text{KL}}(u, v) := D_{\text{KL}}(\vec{\mathbb{P}}^{u,\pi} | \vec{\mathbb{P}}^{v,\tau}) = \mathbb{E}_{Z \sim \vec{\mathbb{P}}^{u,\pi}} \left[ \log \frac{d\vec{\mathbb{P}}^{u,\pi}}{d\vec{\mathbb{P}}^{v,\tau}}(Z) \right]. \quad (8)$$

While we will focus on the KL divergence in our experiments, we mention that our framework can be applied to arbitrary divergences. In particular, one can use divergences that allow for off-policy training and improved mode exploration, such as the log-variance divergence (Nüsken & Richter, 2021; Richter et al., 2020), see Appendix A.4. Finally, we mention that our framework can also be used to derive the ELBO and variational gap for underdamped methods in generative modeling, where one has access to samples  $X \sim p_{\text{target}}$ , and we refer to Appendix A.3 for details.

### 3 UNDERDAMPED DIFFUSION BRIDGES

In order to approach Problem 2.1 and minimize divergences (such as the KL divergence) in practice, we need to numerically approximate the Radon-Nikodym derivative in Proposition 2.3. Analogously to Vargas et al. (2024, Proposition E.1), we can discretize the appearing integrals to show that

$$\frac{d\vec{\mathbb{P}}^{u,\pi}}{d\vec{\mathbb{P}}^{v,\tau}}(Z) \approx \frac{\pi(\hat{Z}_0) \prod_{n=0}^{N-1} \vec{p}_{n+1|n}(\hat{Z}_{n+1} | \hat{Z}_n)}{\tau(\hat{Z}_N) \prod_{n=0}^{N-1} \tilde{p}_{n|n+1}(\hat{Z}_n | \hat{Z}_{n+1})}, \quad (9)$$

where the expressions for the forward and backward transition kernels  $\vec{p}$  and  $\tilde{p}$  depend on the choice of the integrator for  $Z$ . Since we have degenerate diffusion matrices, the backward kernel  $\tilde{p}$  can exhibit vanishing values, which requires careful choice of the integrators for  $Z$ . In particular, naively using an Euler-Maruyama scheme as an integrator is typically not well-suited (Leimkuhler & Reich, 2004; Neal, 2012; Doucet et al., 2022b); see also Figure 4.

We therefore consider alternative integration methods, specifically splitting schemes (Bou-Rabee & Owahdi, 2010; Melchionna, 2007), which divide the SDE into simpler parts that can be integrated individually before combining them. Such methods are particularly useful when certain parts can be solved exactly. To formalize splitting schemes, we leverage the Fokker-Planck operator framework, proposing a decomposition of the generator  $\mathcal{L}$  for diffusion processes  $Z$  of the form (5).



We can define  $\mathcal{L}$  via the (kinetic) Fokker-Planck equation<sup>5</sup>

$$\partial_t p = \mathcal{L}p \quad \text{with} \quad \mathcal{L}p = -\nabla \cdot ((f + \eta u)p) + \frac{1}{2} \text{Tr}(\eta \eta^\top \nabla^2 p), \quad (10)$$

governing the evolution of the density  $p(\cdot, t) = \mathbb{P}_t^{u, \pi}$  of the solution to the SDE in (5). In order to approximate the generator  $\mathcal{L}$ , we want to assume a suitable structure for  $f$  and  $\eta$ , such that we decompose  $\mathcal{L}$  into simpler pieces. For this, we come back to the setting of the underdamped Langevin equation stated in the introduction in equation (3). We can readily see that its controlled counterpart can be incorporated in the framework presented in Section 2 by making the choices  $D = 2d$ ,  $Z = (X, Y)^\top$ , and

$$f(x, y, s) = (y, \tilde{f}(x, s) - \frac{1}{2} \sigma \sigma^\top(s) y)^\top, \quad \eta = (\mathbf{0}, \sigma)^\top \quad (11)$$

in (5) and (6), where  $\mathbf{0} \in \mathbb{R}^{d \times d}$  and  $\tilde{f} \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$  is chosen appropriately. Following Monmarché (2021); Geffner & Domke (2023) we split the generator as  $\mathcal{L} = \mathcal{L}_A + \mathcal{L}_B + \mathcal{L}_O$  (sometimes referred to as free transport, acceleration, and damping) with

$$\mathcal{L}_A p = -y \cdot \nabla_x p, \quad \mathcal{L}_B p = -\tilde{f} \cdot \nabla_y p, \quad \mathcal{L}_O p = -\nabla_y \cdot (g p) + \frac{1}{2} \text{Tr}(\sigma \sigma^\top \nabla_y^2 p), \quad (12)$$

where  $g(x, y, s) := -\frac{1}{2} \sigma \sigma^\top(s) y + \sigma(s) u(x, y, s)$ , resulting in

$$\begin{bmatrix} dX_s \\ dY_s \end{bmatrix} = \underbrace{\begin{bmatrix} Y_s \\ 0 \end{bmatrix}}_A ds + \underbrace{\begin{bmatrix} 0 \\ \tilde{f}(X_s, s) \end{bmatrix}}_B ds + \underbrace{\begin{bmatrix} 0 \\ (-\frac{1}{2} \sigma \sigma^\top(s) Y_s + \sigma u(Z_s, s)) \end{bmatrix}}_O ds + \sigma(s) d\vec{W}_s, \quad (13)$$

where we use a standard normal for the last  $d$  components of the initial and terminal distributions following Geffner & Domke (2023), i.e.,

$$\pi(x, y) = p_{\text{prior}}(x) \mathcal{N}(y; 0, \text{Id}) \quad \text{and} \quad \tau(x, y) = p_{\text{target}}(x) \mathcal{N}(y; 0, \text{Id}). \quad (14)$$

According to the Trotter theorem (Trotter, 1959) and the Strang splitting formula (Strang, 1968), the time evolution of the system can be approximated as:

$$e^{(\mathcal{L}_A + \mathcal{L}_B + \mathcal{L}_O)t} \approx [e^{\mathcal{L}_A \Delta} e^{\mathcal{L}_B \Delta} e^{\mathcal{L}_O \Delta}]^N + \mathcal{O}(N \Delta^3), \quad (15)$$

where a finite number of time steps of length  $\Delta$  approximates the system dynamics. For a higher accuracy, symmetric splitting can be used:

$$e^{(\mathcal{L}_A + \mathcal{L}_B + \mathcal{L}_O)t} \approx [e^{\mathcal{L}_O \frac{\Delta}{2}} e^{\mathcal{L}_B \frac{\Delta}{2}} e^{\mathcal{L}_A \Delta} e^{\mathcal{L}_B \frac{\Delta}{2}} e^{\mathcal{L}_O \frac{\Delta}{2}}]^N + \mathcal{O}(N \Delta^2), \quad (16)$$

which reduces the approximation error (Yoshida, 1990). The optimal composition of terms is generally problem-dependent and has been extensively studied for uncontrolled Langevin dynamics (Monmarché, 2021). For the controlled setting, prior works often use the OBAB ordering (Geffner & Domke, 2023; Doucet et al., 2022a). In this work, we additionally consider OBABO and BAOAB, which show improved performance (cf. Section 4).

Further details on the integrators for forward and backward kernels  $\vec{p}$  and  $\bar{p}$  corresponding to these splitting schemes can be found in Appendix A.9. We refer to Algorithm 1 in the appendix for an overview of our method and to Appendix A.11 for further details. Moreover, in Appendix A.4, we comment on higher-order Langevin equations, mass matrices, discretizations of the Radon-Nikodym derivative, as well as favorable properties of the score in the underdamped setting.

## 4 NUMERICAL EXPERIMENTS

In this section, we present a comparative analysis of underdamped approaches against their overdamped counterparts. We consider five diffusion-based sampling methods, specifically, *Unadjusted Langevin Annealing* (ULA) (Thin et al., 2021; Geffner & Domke, 2021), *Monte Carlo Diffusions* (MCD) (Doucet et al., 2022b; Geffner & Domke, 2023), *Controlled Monte Carlo Diffusions* (CMCD) (Vargas et al., 2024), *Time-Reversed Diffusion Sampler* (DIS)<sup>6</sup> (Berner et al., 2024), and

<sup>5</sup>We denote by  $\text{Tr}$  the trace and by  $\nabla$  the derivative operator w.r.t. spatial variable  $z$ ; see Appendix A.1.

<sup>6</sup>It is worth noting that we do not separately consider the Denoising Diffusion Sampler (DDS) (Vargas et al., 2023a), as it can be viewed as a special case of DIS (see Appendix A.10.1 in Berner et al. (2024)).

Table 1: Results for benchmark problems of various dimensions  $d$ , averaged across four runs. Evaluation criteria include importance-weighted errors for estimating the log-normalizing constant,  $\Delta \log \mathcal{Z}$ , effective sample size ESS, Sinkhorn distance  $\mathcal{W}_2^\gamma$ , and a lower bound (LB) on  $\log \mathcal{Z}$ ; see Appendix A.11.5 for details on the metrics. The best results are highlighted in bold. Arrows ( $\uparrow$ ,  $\downarrow$ ) indicate whether higher or lower values are preferable. Blue shading indicates that the method uses the underdamped Langevin equation.

	Funnel ( $d = 10$ )			ManyWell ( $d = 50$ )		LGCP ( $d = 1600$ )	
Method	$\Delta \log \mathcal{Z} \downarrow$	ESS $\uparrow$	$\mathcal{W}_2^\gamma \downarrow$	$\Delta \log \mathcal{Z} \downarrow$	ESS $\uparrow$	$\log \mathcal{Z}$ (LB) $\uparrow$	ESS $\times 10 \uparrow$
ULA	0.310 $\pm$ 0.020	0.140 $\pm$ 0.003	169.859 $\pm$ 0.195	0.016 $\pm$ 0.003	0.179 $\pm$ 0.008	482.024 $\pm$ 0.009	0.029 $\pm$ 0.003
	0.130 $\pm$ 0.021	0.151 $\pm$ 0.016	159.212 $\pm$ 0.093	0.009 $\pm$ 0.002	0.418 $\pm$ 0.002	484.087 $\pm$ 0.063	0.030 $\pm$ 0.004
MCD	0.173 $\pm$ 0.046	0.206 $\pm$ 0.026	164.967 $\pm$ 0.334	0.005 $\pm$ 0.002	0.737 $\pm$ 0.002	483.137 $\pm$ 0.368	0.031 $\pm$ 0.004
	0.088 $\pm$ 0.008	0.375 $\pm$ 0.016	144.753 $\pm$ 0.153	0.005 $\pm$ 0.000	0.866 $\pm$ 0.012	484.933 $\pm$ 0.298	0.032 $\pm$ 0.006
CMCD	0.023 $\pm$ 0.003	0.567 $\pm$ 0.023	104.644 $\pm$ 0.710	<b>0.004<math>\pm</math>0.002</b>	0.859 $\pm$ 0.001	483.875 $\pm$ 0.275	0.032 $\pm$ 0.004
	0.268 $\pm$ 0.198	0.369 $\pm$ 0.186	148.990 $\pm$ 19.81	0.008 $\pm$ 0.003	0.585 $\pm$ 0.034	483.535 $\pm$ 0.232	0.028 $\pm$ 0.004
DIS	0.047 $\pm$ 0.003	0.498 $\pm$ 0.021	107.458 $\pm$ 0.826	0.006 $\pm$ 0.002	0.798 $\pm$ 0.002	405.686 $\pm$ 4.019	0.015 $\pm$ 0.003
	0.048 $\pm$ 0.009	0.550 $\pm$ 0.039	114.580 $\pm$ 0.457	0.005 $\pm$ 0.000	0.856 $\pm$ 0.002	diverged	diverged
DBS	0.021 $\pm$ 0.003	0.603 $\pm$ 0.014	102.653 $\pm$ 0.586	0.005 $\pm$ 0.001	0.887 $\pm$ 0.004	486.376 $\pm$ 1.020	0.032 $\pm$ 0.002
	<b>0.010<math>\pm</math>0.001</b>	<b>0.779<math>\pm</math>0.009</b>	<b>101.418<math>\pm</math>0.425</b>	0.005 $\pm$ 0.000	<b>0.898<math>\pm</math>0.002</b>	<b>497.545<math>\pm</math>0.183</b>	<b>0.174<math>\pm</math>0.017</b>

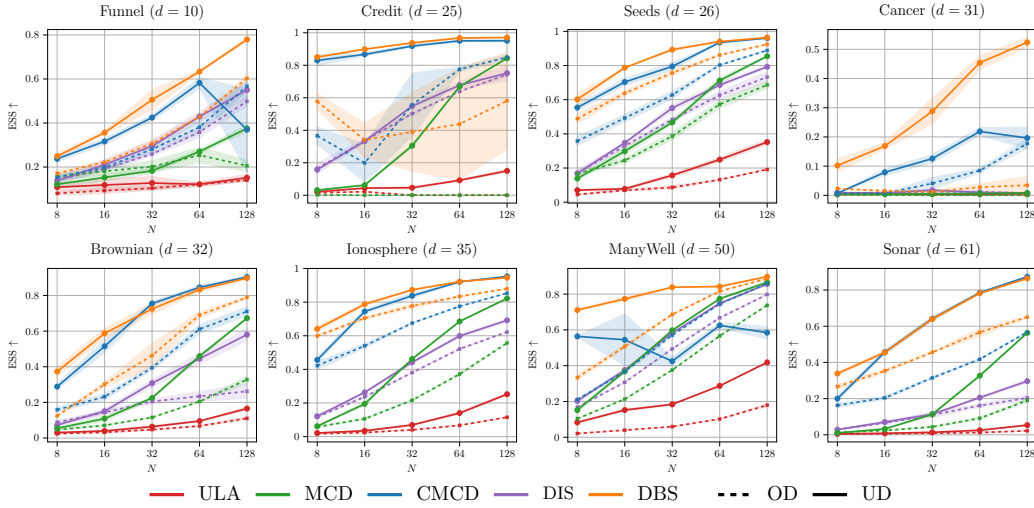


Figure 3: Effective sample size (ESS) for real-world benchmark problems of various dimensions  $d$ , averaged across four seeds. Here,  $N$  refers to the number of discretization steps. Dashed and solid lines indicate the usage of the overdamped (OD) and underdamped (UD) Langevin dynamics, respectively.

*Diffusion Bridge Sampler* (DBS) (Richter & Berner, 2024). We stress that the underdamped versions of DIS and DBS have not been considered before.

To ensure a fair comparison, all experiments are conducted under identical settings. Our code is publicly available<sup>7</sup> and our evaluation methodology adheres to the protocol suggested in Blessing et al. (2024). For a comprehensive overview of the considered benchmark problems and our experimental setup, we refer to Appendix A.11. Moreover, we provide further numerical results in Appendix A.11.6, including the comparison to competing state-of-the-art methods. In the following, we outline the main findings of our experiments.

**Underdamped vs. overdamped.** Our analysis of both real-world and synthetic benchmark problems reveals consistent improvements when using underdamped Langevin dynamics compared to its overdamped counterpart, as illustrated in Table 1 and Figure 3. The underdamped diffusion bridge sampler (DBS) demonstrates particularly impressive performance, consistently outperforming alternative methods. Remarkably, even with as few as  $N = 8$  discretization steps, it often surpasses competing methods that utilize significantly more steps.

<sup>7</sup><https://github.com/DenisBless/UnderdampedDiffusionBridges>




Figure 4: Effective sample size (ESS) and wallclock time (in seconds) of the diffusion bridge sampler (DBS) for different integration schemes, averaged across multiple benchmark problems and four seeds. Integration schemes include Euler-Maruyama (EM) for over- (OD) and underdamped (UD) Langevin dynamics and various splitting schemes (OBAB, BAOAB, OBABO).




Figure 5: Effective sample size (ESS) of the underdamped diffusion bridge sampler (DBS) for various combinations of learned parameters, averaged across multiple benchmark problems and four seeds using  $N = 64$  discretization steps. Hyperparameters include mass matrix  $M$ , diffusion matrix  $\sigma$ , terminal time  $T$ , and extended prior distribution  $\pi$ . See Figure 10 for the results with  $N = 8$  discretization steps.

**Numerical integration schemes.** We further examine various numerical schemes for the diffusion bridge sampler (DBS) introduced in Section 3. Results and a discussion for other methods can be found in Appendix A.11.6. To provide a concise overview, we present the average effective sample size (ESS) and wallclock time across all tasks, excluding LGCP, in Figure 4. Detailed results for individual benchmarks can be found in Appendix A.11.6. While it is known that classical Euler methods are not well-suited for underdamped dynamics (Leimkuhler & Reich, 2004), our findings indicate that both OBAB and BAOAB schemes offer significant improvements without incurring additional computational costs. The OBABO scheme yields the best results overall, albeit at the expense of increased computational demands due to the need for double evaluation of the control per discretization step. However, it is worth noting that in many real-world applications, target evaluations often constitute the primary computational bottleneck. In such scenarios, OBABO may be the preferred choice despite its higher computational requirements.

**End-to-end hyperparameter learning.** Finally, we examine the impact of end-to-end learning of various hyperparameters on the performance of the underdamped diffusion bridge sampler. Our investigation focuses on optimizing the (diagonal) mass matrix  $M$  (cf. Appendix A.8), diffusion matrix  $\sigma$ , terminal time  $T$ , and prior distribution  $\pi$ . Figures 5 and 10 illustrate the effective sample size, averaged across all tasks (excluding LGCP) for  $N = 64$  and  $N = 8$  diffusion steps, respectively. The results reveal that learning these parameters, particularly the terminal time and prior distribution, leads to substantial performance gains. We note that this feature improves the method’s usability and accessibility by minimizing or eliminating the need for manual hyperparameter tuning.

## 5 CONCLUSION AND OUTLOOK

In this work we have formulated a general framework for diffusion bridges including degenerate stochastic processes. In particular, we propose the novel *underdamped diffusion bridge sampler*, which achieves state-of-the-art results on multiple sampling tasks without hyperparameter tuning and only a few discretization steps. While a complete theoretical analysis is left for future work (see Remark A.9), we provide careful ablation studies showing that our improvements are due to the combination of underdamped dynamics, our novel numerical integrators, learning both the forward and backward processes as well as end-to-end learned hyperparameters. Our results also suggest to extend the method by Chen et al. (2021) and benchmark underdamped diffusion bridges for generative modeling using the ELBO derived in Lemma A.1.



## ACKNOWLEDGEMENTS

J.B. acknowledges support from the Wally Baer and Jeri Weiss Postdoctoral Fellowship. The research of L.R. was partially funded by Deutsche Forschungsgemeinschaft (DFG) through the grant CRC 1114 “Scaling Cascades in Complex Systems” (project A05, project number 235221301). D.B. acknowledges support by funding from the pilot program Core Informatics of the Helmholtz Association (HGF) and the state of Baden-Württemberg through bwHPC, as well as the HoreKa supercomputer funded by the Ministry of Science, Research and the Arts Baden-Württemberg and by the German Federal Ministry of Education and Research.

## REFERENCES

- Tara Akhond-Sadeh, Jarrod Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, Nikolay Malkin, and Alexander Tong. Iterated denoising energy matching for sampling from Boltzmann densities. In *International Conference on Machine Learning*, pp. 760–786, 2024.
- Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- Michael Arbel, Alex Matthews, and Arnaud Doucet. Annealed flow transport Monte Carlo. In *International Conference on Machine Learning*, pp. 318–330. PMLR, 2021.
- Oleg Arenz, Philipp Dahlinger, Zihan Ye, Michael Volpp, and Gerhard Neumann. A unified perspective on natural gradient variational inference with Gaussian mixture models. *arXiv preprint arXiv:2209.11533*, 2022.
- Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024.
- Denis Blessing, Xiaogang Jia, Johannes Esslinger, Francisco Vargas, and Gerhard Neumann. Beyond ELBOs: A large-scale evaluation of variational methods for sampling. In *International Conference on Machine Learning*, pp. 4205–4229. PMLR, 2024.
- Vladimir I Bogachev, Nicolai V Krylov, Michael Röckner, and Stanislav V Shaposhnikov. *Fokker–Planck–Kolmogorov Equations*, volume 207. American Mathematical Society, 2022.
- Nawaf Bou-Rabee and Houman Owhadi. Long-run accuracy of variational integrators in the stochastic context. *SIAM Journal on Numerical Analysis*, 48(1):278–297, 2010.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, et al. Jax: Autograd and xla. *Astrophysics Source Code Library*, pp. ascl–2111, 2021.
- Junhua Chen, Lorenz Richter, Julius Berner, Denis Blessing, Gerhard Neumann, and Anima Anandkumar. Sequential controlled Langevin diffusions. In *International Conference on Learning Representations*, 2025.
- Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint arXiv:2209.11215*, 2022.
- Tianrong Chen, Guan-Hong Liu, and Evangelos A Theodorou. Likelihood training of Schrödinger bridge using forward-backward SDEs theory. *arXiv preprint arXiv:2110.11291*, 2021.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- Marco Cuturi, Laetitia Meng-Papaxanthos, Yingtao Tian, Charlotte Bunne, Geoff Davis, and Olivier Teboul. Optimal transport tools (OTT): A jax toolbox for all things Wasserstein. *arXiv preprint arXiv:2201.12324*, 2022.
- Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.

- Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006.
- Bernard Delyon and Ying Hu. Simulation of conditioned diffusion and application to parameter estimation. *Stochastic Processes and their Applications*, 116(11):1660–1675, 2006.
- Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Score-based generative modeling with critically-damped Langevin diffusion. In *International Conference on Learning Representations*, 2022.
- Arnaud Doucet, Will Grathwohl, Alexander G Matthews, and Heiko Strathmann. Score-based diffusion meets annealed importance sampling. *Advances in Neural Information Processing Systems*, 35:21482–21494, 2022a.
- Arnaud Doucet, Will Grathwohl, Alexander G de G Matthews, and Heiko Strathmann. Score-based diffusion meets annealed importance sampling. In *Advances in Neural Information Processing Systems*, 2022b.
- R. Durrett. *Brownian Motion and Martingales in Analysis*. Wadsworth & Brooks/Cole Mathematics Series. Wadsworth Advanced Books & Software, 1984.
- Andreas Eberle, Arnaud Guillin, and Raphael Zimmer. Couplings and quantitative contraction rates for Langevin dynamics. *The Annals of Probability*, 2019.
- H Föllmer. Time reversal on Wiener space. *Stochastic Processes—Mathematics and Physics*, pp. 119–129, 1986.
- A Friedman. Partial differential equations of parabolic type. *Prentice-Hall*, 1964.
- Avner Friedman. Stochastic differential equations and applications. In *Stochastic differential equations*, pp. 75–148. Springer, 1975.
- Tomas Geffner and Justin Domke. MCMC variational inference via uncorrected Hamiltonian annealing. In *Advances in Neural Information Processing Systems*, 2021.
- Tomas Geffner and Justin Domke. Langevin diffusion variational inference. In *International Conference on Artificial Intelligence and Statistics*, pp. 576–593. PMLR, 2023.
- A. Gelman, J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, and D.B. Rubin. *Bayesian Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis, 2013.
- Ulrich G Haussmann and Etienne Pardoux. Time reversal of diffusions. *The Annals of Probability*, pp. 1188–1205, 1986.
- Jeremy Heng, Valentin De Bortoli, Arnaud Doucet, and James Thornton. Simulating diffusion bridges with score matching. *arXiv preprint arXiv:2111.07243*, 2021.
- 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.
- Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- Peter E Kloeden and Eckhard Platen. Stochastic differential equations. In *Numerical Solution of Stochastic Differential Equations*, pp. 103–160. Springer, 1992.
- Andrei Kolmogoroff. Über die analytischen Methoden in der Wahrscheinlichkeitsrechnung. *Mathematische Annalen*, 104:415–458, 1931.
- Takeshi Koshizuka and Issei Sato. Neural Lagrangian Schrödinger bridge: Diffusion modeling for population dynamics. In *International Conference on Learning Representations*, 2023.

- Hiroshi Kunita. *Stochastic flows and jump-diffusions*. Springer, 2019.
- Jean-François Le Gall. *Brownian motion, martingales, and stochastic calculus*. Springer, 2016.
- Ben Leimkuhler and Charles Matthews. Molecular dynamics. *Interdisciplinary applied mathematics*, 39(1), 2015.
- Benedict Leimkuhler and Sebastian Reich. *Simulating Hamiltonian dynamics*. Number 14. Cambridge university press, 2004.
- Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems*, 34(4):1533–1574, 2013.
- Christian Léonard. Some properties of path measures. *Séminaire de Probabilités XLVI*, pp. 207–230, 2014.
- Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos Theodorou. Deep generalized Schrödinger bridge. *Advances in Neural Information Processing Systems*, 35:9374–9388, 2022.
- Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos A Theodorou, and Ricky TQ Chen. Generalized Schrödinger bridge matching. *The Twelfth International Conference on Learning Representations*, 2024.
- Yi-An Ma, Niladri S Chatterji, Xiang Cheng, Nicolas Flammarion, Peter L Bartlett, and Michael I Jordan. Is there an analog of nesterov acceleration for gradient-based mcmc? *Bernoulli*, 27(3), 2021.
- Alex Matthews, Michael Arbel, Danilo Jimenez Rezende, and Arnaud Doucet. Continual repeated annealed flow transport Monte Carlo. In *International Conference on Machine Learning*, pp. 15196–15219. PMLR, 2022.
- Simone Melchionna. Design of quasisymplectic propagators for Langevin dynamics. *The Journal of chemical physics*, 127(4), 2007.
- Laurence Illing Midgley, Vincent Stimper, Gregor NC Simm, Bernhard Schölkopf, and José Miguel Hernández-Lobato. Flow annealed importance sampling bootstrap. *The Eleventh International Conference on Learning Representations*, 2023.
- Annie Millet, David Nualart, and Marta Sanz. Integration by parts and time reversal for diffusion processes. *The Annals of Probability*, pp. 208–238, 1989.
- Jesper Møller, Anne Randi Syversveen, and Rasmus Plenge Waagepetersen. Log Gaussian Cox processes. *Scandinavian journal of statistics*, 25(3):451–482, 1998.
- Pierre Monmarché. High-dimensional MCMC with a standard splitting scheme for the underdamped Langevin diffusion. *Electronic Journal of Statistics*, 15(2):4117–4166, 2021.
- Wenlong Mou, Yi-An Ma, Martin J Wainwright, Peter L Bartlett, and Michael I Jordan. High-order Langevin diffusion yields an accelerated MCMC algorithm. *Journal of Machine Learning Research*, 22(42):1–41, 2021.
- Radford M Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.
- Radford M Neal. Slice sampling. *The annals of statistics*, 31(3):705–767, 2003.
- Radford M Neal. MCMC using Hamiltonian dynamics. *arXiv preprint arXiv:1206.1901*, 2012.
- Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *International conference on machine learning*, pp. 25858–25889. PMLR, 2023.
- E Nelson. Dynamical theories of Brownian motion. *Press, Princeton, NJ*, 1967.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.

- Nikolas Nüsken and Lorenz Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial differential equations and applications*, 2(4):48, 2021.
- Grigorios A Pavliotis. Stochastic processes and applications. *Texts in applied mathematics*, 60, 2014.
- Angus Phillips, Hai-Dang Dau, Michael John Hutchinson, Valentin De Bortoli, George Deligianidis, and Arnaud Doucet. Particle denoising diffusion sampler. *International Conference on Machine Learning*, 2024.
- Lorenz Richter. *Solving high-dimensional PDEs, approximation of path space measures and importance sampling of diffusions*. PhD thesis, BTU Cottbus-Senftenberg, 2021.
- Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. In *International Conference on Learning Representations*, 2024.
- Lorenz Richter, Ayman Boustati, Nikolas Nüsken, Francisco Ruiz, and Omer Deniz Akyildiz. VarGrad: a low-variance gradient estimator for variational inference. *Advances in Neural Information Processing Systems*, 33:13481–13492, 2020.
- Simo Särkkä and Tommi Sottinen. Application of Girsanov theorem to particle filtering of discretely observed continuous-time non-linear systems. *arXiv preprint arXiv:0705.1598*, 2007.
- Moritz Schauer, Frank van der Meulen, and Harry van Zanten. Guided proposals for simulating multi-dimensional diffusion bridges. *arXiv preprint arXiv:1311.3606*, 2013.
- René L Schilling and Lothar Partzsch. *Brownian motion: an introduction to stochastic processes*. Walter de Gruyter GmbH & Co KG, 2014.
- Henrik Schopmans and Pascal Friederich. Temperature-annealed Boltzmann generators. *arXiv preprint arXiv:2501.19077*, 2025.
- Ziqiang Shi and Rujie Liu. Generative modelling with high-order Langevin dynamics. *arXiv preprint arXiv:2404.12814*, 2024.
- Gabriel Stoltz, Mathias Rousset, et al. *Free energy computations: A mathematical perspective*. World Scientific, 2010.
- Gilbert Strang. On the construction and comparison of difference schemes. *SIAM journal on numerical analysis*, 5(3):506–517, 1968.
- Jingtong Sun, Julius Berner, Lorenz Richter, Marius Zeinhofer, Johannes Müller, Kamyar Aziz-zadenesheli, and Anima Anandkumar. Dynamical measure transport and neural PDE solvers for sampling. *arXiv preprint arXiv:2407.07873*, 2024.
- Achille Thin, Nikita Kotelevskii, Alain Durmus, Eric Moulines, Maxim Panov, and Arnaud Doucet. Monte Carlo variational auto-encoders. In *International Conference on Machine Learning*, 2021.
- Hale F Trotter. On the product of semi-groups of operators. *Proceedings of the American Mathematical Society*, 10(4):545–551, 1959.
- A Süleyman Üstünel and Moshe Zakai. *Transformation of measure on Wiener space*. Springer Science & Business Media, 2013.
- Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving Schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.
- Francisco Vargas, Will Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. *The Eleventh International Conference on Learning Representations*, 2023a.
- Francisco Vargas, Andrius Ovsianas, David Fernandes, Mark Girolami, Neil D Lawrence, and Nikolas Nüsken. Bayesian learning via neural Schrödinger–Föllmer flows. *Statistics and Computing*, 33(1):3, 2023b.

- Francisco Vargas, Shreyas Padhy, Denis Blessing, and Nikolas Nüsken. Transport meets variational inference: Controlled Monte Carlo diffusions. In *International Conference on Learning Representations*, 2024.
- Hao Wu, Jonas Köhler, and Frank Noé. Stochastic normalizing flows. *Advances in Neural Information Processing Systems*, 33:5933–5944, 2020.
- Haruo Yoshida. Construction of higher order symplectic integrators. *Physics letters A*, 150(5-7): 262–268, 1990.
- Dinghuai Zhang, Ricky Tian Qi Chen, Cheng-Hao Liu, Aaron Courville, and Yoshua Bengio. Diffusion generative flow samplers: Improving learning signals through partial trajectory optimization. *The Twelfth International Conference on Learning Representations*, 2024.
- Guodong Zhang, Kyle Hsu, Jianing Li, Chelsea Finn, and Roger Grosse. Differentiable annealed importance sampling and the perils of gradient noise. In *Advances in Neural Information Processing Systems*, 2021.
- Qinsheng Zhang and Yongxin Chen. Path integral sampler: a stochastic control approach for sampling. *International Conference on Learning Representations*, 2022.



## A APPENDIX

## CONTENTS

A.1	Notation . . . . .	14
A.2	Assumptions . . . . .	15
A.3	Implications for Generative modeling: The evidence lower bound . . . . .	16
A.4	Further remarks . . . . .	17
A.5	Auxiliary results . . . . .	18
A.6	Proofs . . . . .	19
A.7	Additional statements on diffusion models . . . . .	20
A.8	Including a mass matrix . . . . .	21
A.9	Numerical discretization schemes . . . . .	22
A.9.1	Euler-Maruyama . . . . .	22
A.9.2	OBAB . . . . .	23
A.9.3	BAOAB . . . . .	24
A.9.4	OBABO . . . . .	24
A.10	Underdamped version of previous diffusion-based sampling methods . . . . .	25
A.11	Further computational details . . . . .	25
A.11.1	Algorithm . . . . .	26
A.11.2	SDE preconditioning . . . . .	26
A.11.3	Benchmark Problems . . . . .	27
A.11.4	Experimental setup . . . . .	28
A.11.5	Evaluation criteria . . . . .	29
A.11.6	Further experiments and comparisons . . . . .	30

## A.1 NOTATION

We denote by  $\text{Tr}(\Sigma)$  and  $\Sigma^+$  the trace and the (Moore-Penrose) pseudoinverse of a real-valued matrix  $\Sigma$ , by  $\|\mu\|$  the Euclidean norm of a vector  $\mu$ , and by  $\mu_1 \cdot \mu_2$  the Euclidean inner product between vectors  $\mu_1$  and  $\mu_2$ .

For a function  $p \in C(\mathbb{R}^D \times [0, T], \mathbb{R})$ , depending on the variables  $z = (x, y) \in \mathbb{R}^d \times \mathbb{R}^{D-d} \simeq \mathbb{R}^D$  and  $t \in [0, T]$ , we denote by  $\partial_t p$  its partial derivative w.r.t. the time coordinate  $t$  and by  $\nabla_x p$  and  $\nabla_y p$  its gradients w.r.t. the spatial variables  $x$  and  $y$ , respectively. Here,  $C(A, B)$  denotes the set of all continuous functions mapping from the set  $A$  to the set  $B$ . Moreover, we denote by

$$\nabla p = \begin{bmatrix} \nabla_x p \\ \nabla_y p \end{bmatrix} \quad (17)$$

the gradient w.r.t. both spatial variables  $z = (x, y)$ . We analogously denote by  $\nabla^2 p$  the Hessian of  $p$  w.r.t. the spatial variables. Similarly, we define  $\nabla \cdot f = \sum_{i=1}^D \partial_{x_i} f_i$  to be the divergence of a (time-dependent) vector field  $f = (f_i)_{i=1}^D \in C(\mathbb{R}^D \times [0, T], \mathbb{R}^D)$  w.r.t. the spatial variables.

We denote by  $\mathcal{N}(\mu, \Sigma)$  a multivariate normal distribution with mean  $\mu \in \mathbb{R}^d$  and (positive semi-definite matrix) covariance matrix  $\Sigma \in \mathbb{R}^{d \times d}$  and write  $\mathcal{N}(x; \mu, \Sigma)$  for the evaluation of its density (w.r.t. the Lebesgue measure) at  $x \in \mathbb{R}^d$ . Moreover, we denote by  $\text{Unif}([0, T])$  the uniform distribution on  $[0, T]$ . For an  $\mathbb{R}^d$ -valued random variable  $X$  with law  $\mathbb{P}$  and a function  $f \in (\mathbb{R}^d, \mathbb{R})$ , we

denote by

$$\mathbb{E}_{X \sim \mathbb{P}}[f(X)] = \int f \, d\mathbb{P} \quad (18)$$

the expected value of the random variable  $f(X)$ .

For suitable continuous stochastic processes  $Z = (Z_t)_{t \in [0, T]}$  and  $Y = (Y_t)_{t \in [0, T]}$ , we define forward and backward Itô integrals via the limits

$$\int_{\underline{t}}^{\bar{t}} X_s \cdot dY_s = \lim_{n \rightarrow \infty} \sum_{i=0}^{k_n} X_{t_i^n} \cdot (Y_{t_{i+1}^n} - Y_{t_i^n}), \quad (19)$$

$$\int_{\underline{t}}^{\bar{t}} X_s \cdot \tilde{d}Y_s = \lim_{n \rightarrow \infty} \sum_{i=0}^{k_n} X_{t_{i+1}^n} \cdot (Y_{t_{i+1}^n} - Y_{t_i^n}), \quad (20)$$

where  $\underline{t} < t_0^n < \dots < t_{k_n}^n = \bar{t}$  is an increasing sequence of subdivisions of  $[\underline{t}, \bar{t}]$  with mesh size tending to zero; see Vargas et al. (2024) for details. The relation between forward and backward integrals is given in Lemma A.12.

We denote by  $\mathcal{P}$  the set of probability measures on  $C([0, T], \mathbb{R}^D)$ , equipped with the Borel  $\sigma$ -field associated with the topology of uniform convergence on compact sets. For suitable vector fields  $u$ ,  $v$  and distributions  $\pi$ ,  $\tau$ , we denote by  $\mathbb{P}^{u, \pi} \in \mathcal{P}$  and  $\mathbb{P}^{v, \tau} \in \mathcal{P}$  the forward and reverse-time *path measures*, i.e., the laws or pushforwards on  $C([0, T], \mathbb{R}^D)$ , of the solutions  $Z = (Z_t)_{t \in [0, T]}$  to the SDEs

$$Z_t = Z_0 + \int_0^t (f + \eta u)(Z_s, s) \, ds + \int_0^t \eta(s) \, dW_s, \quad Z_0 \sim \pi, \quad (21)$$

$$Z_t = Z_T - \int_t^T (f + \eta v)(Z_s, s) \, ds - \int_t^T \eta(s) \, \tilde{d}W_s, \quad Z_T \sim \tau, \quad (22)$$

respectively. In the above,  $W$  denotes a standard  $d$ -dimensional Brownian motion satisfying the usual conditions, see, e.g., Kunita (2019). Note that we consider degenerate diffusion coefficients  $\eta$  of the form  $\eta = (0, \sigma)^\top$ . Moreover, we assume that the first components of  $f = (\hat{f}, \tilde{f})^\top$ , namely  $\hat{f}$ , satisfies  $\nabla_x \hat{f}(z, \cdot) = 0$  for every  $z = (x, y) \in \mathbb{R}^D$ , i.e.,  $\hat{f}$  only depends on  $y$  but not on  $x$ . Finally, we denote the marginal of a path space measure  $\mathbb{P}$  at time  $t \in [0, T]$  by  $\mathbb{P}_t$ , which can be interpreted as the pushforward under the evaluation  $Z \mapsto Z_t$ . Moreover, we denote by  $\mathbb{P}_{s|t}$  the conditional distribution of  $\mathbb{P}_s$  given  $\mathbb{P}_t$ .

## A.2 ASSUMPTIONS

Throughout the paper, we assume that all vector fields are smooth, i.e., for a vector field  $g$  it holds  $g \in C^\infty(\mathbb{R}^D \times [0, T], \mathbb{R}^d)$ , and satisfy a global Lipschitz condition (uniformly in time), i.e., there exists a constant  $C$  such that for all  $z_1, z_2 \in \mathbb{R}^D$  and  $t \in [0, T]$  it holds that

$$\|g(z_1, t) - g(z_2, t)\| \leq C \|z_1 - z_2\|. \quad (23)$$

These assumptions also define the set of *admissible controls*  $\mathcal{U} \subset C^\infty(\mathbb{R}^D \times [0, T], \mathbb{R}^d)$ .

Moreover, we assume that the diffusion coefficients appearing in the dimensions with the control,  $\sigma$ , are invertible for all  $t \in [0, T]$  and satisfy that  $\sigma \in C^\infty([0, T], \mathbb{R}^{d \times d})$ . Our continuity assumptions on the SDE coefficient functions and the global Lipschitz condition in (23) guarantee strong solutions with pathwise uniqueness (see, e.g., Le Gall (2016, Section 8.2)) and are sufficient for Girsanov's theorem in Theorem A.10 to hold (see, e.g., Delyon & Hu (2006)). Moreover, our conditions allow the definition of the forward and backward Itô integrals via limits of time discretizations as in (19) and (20) that are independent of the specific sequence of refinements (Vargas et al., 2024).

Finally, we assume that all SDEs admit densities of their time marginals (w.r.t. the Lebesgue measure) that are sufficiently smooth<sup>8</sup> such that we have strong solutions to the corresponding Fokker-

<sup>8</sup>Sufficient conditions for the existence of densities can be found in Millet et al. (1989, Proposition 4.1) and Haussmann & Pardoux (1986, Theorem 3.1). For time-independent SDE coefficient functions, a result by Kolmogoroff (1931) guarantees that the Fokker-Planck equation is satisfied if the density is in  $C^{2,1}(\mathbb{R}^d \times [0, T], \mathbb{R})$ ; see also Pavliotis (2014, Proposition 3.8). and Schilling & Partzsch (2014, 19.6 Proposi-

Planck equations. The existence of continuously differentiable densities and our assumptions on the SDE coefficient functions are sufficient for Nelson’s relation in Lemma 2.2 to hold; see, e.g., Millet et al. (1989). While we use the above assumptions to simplify the presentation, we note they can be significantly relaxed.

### A.3 IMPLICATIONS FOR GENERATIVE MODELING: THE EVIDENCE LOWER BOUND

Contrary to the sampling setting described in Section 1, generative modeling typically assumes that one has access to samples  $X \sim p_{\text{target}}$ , but cannot evaluate the (unnormalized) density. In this section we show how our general setup from the previous section can also be applied in this scenario. For instance, it readily brings an underdamped version of stochastic bridges (Chen et al., 2021) and serves as a theoretical foundation for underdamped diffusion models stated in Dockhorn et al. (2022).

To this end, we may approach Problem 2.1 with the forward<sup>9</sup> KL divergence

$$D_{\text{KL}}(\tilde{\mathbb{P}}^{v,\tau} | \tilde{\mathbb{P}}^{u,\pi}) = \mathbb{E}_{Z \sim \tilde{\mathbb{P}}^{v,\tau}} \left[ \log \frac{d\tilde{\mathbb{P}}^{v,\tau}}{d\tilde{\mathbb{P}}^{u,\pi}}(Z) \right]. \quad (24)$$

For the sake of notation, we have reversed time, which can be viewed as interchanging  $\tau$  and  $\pi$ . Since the process corresponding to  $\tilde{\mathbb{P}}^{v,\tau}$  starts at the target measure  $\tau$ , we indeed require samples from this measure to compute the divergence in (24). At the same time, looking at Proposition 2.3, we realize that the divergence cannot be computed directly, since  $\tau$  cannot be evaluated. A workaround is to instead consider an evidence lower bound (ELBO) (or, equivalently, a lower bound on the log-likelihood). In our setting, we have the following decomposition.

**Lemma A.1** (ELBO for generative modeling). *It holds that*

$$\underbrace{\mathbb{E}_{Z_0 \sim \tau} [\log \tilde{\mathbb{P}}_0^{u,\pi}(Z_0)]}_{\text{evidence / log-likelihood}} = \underbrace{D_{\text{KL}}(\tilde{\mathbb{P}}^{v,\tau} | \tilde{\mathbb{P}}^{\tilde{v},\tau})}_{\text{variational gap}} + \underbrace{\mathbb{E}_{Z_0 \sim \tau} [\log \tau(Z_0)] - D_{\text{KL}}(\tilde{\mathbb{P}}^{v,\tau} | \tilde{\mathbb{P}}^{u,\pi})}_{\text{ELBO}}, \quad (25)$$

where  $\tilde{v}(\cdot, t) - u(\cdot, t) = \eta^\top(t) \nabla \log \tilde{\mathbb{P}}_t^{u,\pi}$ .

*Proof.* This follows from Lemma 2.2 and the chain rule for KL divergences; see Appendix A.6.  $\square$

Crucially, we observe that the ELBO in Lemma A.1 does not depend on the target  $\tau$  anymore as the dependency cancels between the two terms (cf. Proposition 2.3). Moreover, the variational gap is zero if and only if  $v = \tilde{v}$  almost everywhere, i.e., the path measures are time-reversals conditioned on the same terminal condition due to Lemma 2.2. The ELBO is maximized when additionally  $\tilde{\mathbb{P}}_0^{u,\pi}$  equals the target measure  $\tau$ , i.e., if and only if we found a minimizer  $(u^*, v^*)$  of Problem 2.1. In consequence, it provides a viable objective to learn stochastic bridges in an underdamped setting (or, more generally, with degenerate noise coefficients  $\eta$ ) using samples from the target distribution  $\tau$ .

We note that for non-degenerate coefficients  $\eta$ , the ELBO from Lemma A.1 has already been derived in Chen et al. (2021); see also Richter & Berner (2024); Vargas et al. (2024). For diffusion models, i.e.,  $v = 0$  and  $f$  such that  $\tilde{\mathbb{P}}_T^{0,\tau} \approx \pi$ , this ELBO reduces to the one derived by Berner et al. (2024); Huang et al. (2021). In particular, it has been shown that maximizing the ELBO is equivalent to minimizing the *denoising score matching objective* (with a specific weighting of noise scales) typically used in practice.

For general forward and backward processes, allowing for degenerate noise, as stated in (5) and (6), the derivation of the ELBO is less explored. For (underdamped) diffusion models with degenerate  $\eta$ , a corresponding (*hybrid*) *score matching* loss has been suggested and connected to likelihood optimization by Dockhorn et al. (2022, Appendix B.3). In the following proposition, we show that this also follows as a special case from Lemma A.1.

tion). However, we note that popular results by Friedman (1964, Section 1.6) (see also Friedman (1975, Section 5) and Durrett (1984, Section 9.7)) for showing existence and uniqueness of solutions to Fokker-Planck equations require uniform ellipticity assumptions, which are not satisfied for our degenerate diffusion coefficients. We refer to Bogachev et al. (2022, Sections 6.7(ii) and 9.8(i)-(iii)) for existence and uniqueness in the degenerate case and note that we only make use of the Fokker-Planck equation for motivating our splitting schemes in Section 3.

<sup>9</sup>While we optimize the measures in both arguments of the KL divergence, the measure  $\tilde{\mathbb{P}}^{u,\pi}$ , corresponding to the generative process, is in the second component, which is typically referred to as “forward” KL divergence.

**Proposition A.2** (Underdamped score matching maximizes the likelihood). *For the ELBO defined in (25) (setting  $v = 0$ ) it holds*

$$\text{ELBO}(u) = -\frac{T}{2} \mathbb{E}_{Z \sim \tilde{\mathbb{P}}^{0,\tau}, s \sim \text{Unif}([0,T])} \left[ \|u(Z_s, s) + \eta^\top(s) \nabla \log \tilde{\mathbb{P}}_{s|0}^{0,\tau}(Z_s|Z_0)\|^2 \right] + \text{const.},$$

where the constant does not depend on  $u$ .

*Proof.* Following Huang et al. (2021, Appendix A), the proof combines Proposition 2.3 with Stokes' theorem; see Appendix A.6. Note that in our notation  $u$  learns the *negative* and *scaled* score.  $\square$

#### A.4 FURTHER REMARKS

**Remark A.3** (Stochastic bridges and bridge sampling). By *stochastic bridge* or *diffusion bridge* (also referred to as *general bridge* by Richter & Berner (2024)), we refer to an SDE that satisfies the marginals  $p_{\text{prior}}$  and  $p_{\text{target}}$  at times  $t = 0$  and  $t = T$ , respectively. For a given diffusion coefficient of the SDE, there exist infinitely many drifts satisfying these constraints. In particular, for every sufficiently regular density evolution between the prior and target, we can find a drift (given by a unique gradient field) that establishes a corresponding stochastic bridge; see, e.g., Vargas et al. (2024, Proposition 3.4) and Neklyudov et al. (2023, Appendix B.3).

However, any stochastic bridge solves our problem of sampling from  $p_{\text{target}}$  and the non-uniqueness can even lead to better performance in gradient-based optimization (Sun et al., 2024; Blessing et al., 2024). Other previous methods have obtained unique objectives by prescribing the density evolution, e.g., as diffusion process in DIS (Berner et al., 2024) or geometric annealing between prior and target in CMCD (Vargas et al., 2024).

Another popular approach for obtaining uniqueness consists of minimizing the distance<sup>10</sup> to a reference process (additionally to satisfying the marginals). In case the distance is measured via a Kullback-Leibler divergence between the path measures of the bridge and reference process, this setting is often referred to as (*dynamical*) *Schrödinger bridge problem*. In the context of samplers, reference processes have been chosen as scaled Brownian motions in PIS (Zhang & Chen, 2022) and ergodic processes in DDS (Vargas et al., 2023a); see also Richter & Berner (2024) for an overview.

A special case of such a Schrödinger bridge problem is given if the marginals  $p_{\text{prior}}$  and  $p_{\text{target}}$  are Dirac measures. Sampling from the solution to such a problem is equivalent to sampling from the reference SDE conditioned on the start and end point at the times  $t = 0$  and  $t = T$  (specified by the Dirac measures). For instance, if the reference measure is a Brownian motion, solutions are commonly referred to as *Brownian bridges*. As special cases of our considered bridges, solutions to such problems are also sometimes called *diffusion bridges* and we refer to Schauer et al. (2013); Heng et al. (2021) for further details and numerical approaches. However, our sampling problem is in some form orthogonal to such tasks: in case of a Dirac target distribution, sampling is trivial and one is interested in the conditional trajectories. For the sampling problem, the trajectories are not (directly) relevant and one is interested in samples from a general target distribution.

**Remark A.4** (Log-variance loss). As an alternative to the KL divergence in (8), we can consider the log-variance (LV) loss defined as

$$\mathcal{L}_{\text{LV}}^w(u, v) := D_{\text{LV}}^w(\tilde{\mathbb{P}}^{u,\pi}, \tilde{\mathbb{P}}^{v,\tau}) = \text{Var}_{Z \sim \tilde{\mathbb{P}}^{w,\pi}} \left[ \log \frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{v,\tau}}(Z) \right], \quad (26)$$

where the expectation is taken with respect to a path space measure corresponding to a forward process of the form (5), but with the control replaced by an arbitrary control  $w \in \mathcal{U}$ . This allows for off-policy training and avoids the need to differentiate through the simulation of the SDE. Moreover, the estimator achieves zero variance at the optimum  $(u^*, v^*)$ , see Richter et al. (2020); Nüsken & Richter (2021); Richter & Berner (2024).

**Remark A.5** (Higher order Langevin equations). We note that our general framework from Section 2 can readily be used for higher order dynamics and in particular higher order Langevin equations, where next to a position and velocity variable one considers acceleration. As argued by Shi & Liu (2024), corresponding trajectories become smoother the higher the order, which can lead

<sup>10</sup>In the context of generative modeling, also more general settings, referred to as *mean-field games* or *generalized Schrödinger bridges*, have been explored; see, e.g., Liu et al. (2022); Koshizuka & Sato (2023); Liu et al. (2024).

to improved performance of (uncontrolled) Langevin dynamics. Also, [Mou et al. \(2021\)](#) observed improved convergence of third-order Langevin dynamics for convex potentials. We leave related extensions to diffusion bridges for future work.

**Remark A.6** (Mass matrix). Previous works, such as [Geffner & Domke \(2021\)](#) and [Doucet et al. \(2022b\)](#), consider incorporating a mass matrix  $M \in C([0, T], \mathbb{R}^{d \times d})$  into the SDE formulation in (13) and terminal conditions. For simplicity, we have omitted this consideration in the current section. However, additional details on its inclusion and effects can be found in Appendix A.8. Furthermore, we conduct experiments where we learn the mass matrix, as discussed in Section 4.

**Remark A.7** (Discrete Radon-Nikodym derivative). We note that our discretization of the Radon-Nikodym derivative in (9) corresponds to a (discrete-time) Radon-Nikodym derivative between the joint distributions of the discretized forward and backward processes. In particular, we can analogously define a KL divergence which allows us to obtain a (guaranteed) lower bound for the log-normalization constant  $\log \mathcal{Z}$  in discrete time. On the other hand, this is not the case if we discretize the divergence-based Radon-Nikodym derivative in Proposition A.13 as done in previous work ([Berner et al., 2024](#); [Richter & Berner, 2024](#)). Moreover, we can still optimize the divergences between the corresponding discrete path measures as presented in (8) and Appendix A.11.6. Finally, we note that the discretized Radon-Nikodym derivative does not depend on  $\tilde{f}$  for the integrators considered in Appendix A.9. We thus choose  $\tilde{f}$  to have a good initialization for the process  $Z$ , see Appendix A.11.

**Remark A.8** (Properties of the score). Since the target density  $p_{\text{target}}$  in (14) only appears in the coordinates where  $\eta$  vanishes, Nelson’s identity in Lemma 2.2 shows that

$$u^*(x, y, T) - v^*(x, y, T) = \sigma^\top(T) \nabla_y \log \mathcal{N}(y; 0, \text{Id}), \quad (27)$$

i.e., the optimal controls  $u^*$  and  $v^*$  do not depend on the score of the target distribution,  $\nabla_x \log p_{\text{target}}$ , at terminal time  $T$ , as in the case of corresponding overdamped versions. This can lead to numerical benefits in cases where this score would attain large values, e.g., when  $p_{\text{target}}$  is essentially supported on a lower dimensional manifold ([Dockhorn et al., 2022](#); [Chen et al., 2022](#)).

**Remark A.9** (Theoretical guarantees). Our favorable findings encourage further investigation of the theoretical convergence rate of underdamped diffusion samplers. Similar to what has already been observed in generative modeling by [Dockhorn et al. \(2022\)](#), we find significant and consistent improvements over overdamped versions, in particular also for high-dimensional targets with only a few discretization steps  $N$ . However, previous results showed that (for the case  $v = 0$ ), the improved convergence rates of underdamped Langevin dynamics do not carry over to the learned setting, since (different from the score  $\nabla \log p_{\text{target}}$  in Langevin dynamics) the control  $u$  depends not only on the smooth process  $X$ , but also on  $Y$  ([Chen et al., 2022](#)). Specifically, it can be shown that a small KL divergence between the path measures generally requires the step size  $\Delta$  to scale at least linearly in  $d$  (instead of  $\sqrt{d}$ ). While the tightness of our lower bounds on  $\log \mathcal{Z}$  corresponds to such KL divergences, we believe that the results can still be reconciled with our empirical findings due to the following reasons: (1) our samplers are initialized as Langevin dynamics (see Appendix A.11) such that theoretical benefits of the underdamped case hold at least initially (2) the learning problem becomes numerically better behaved (see (27)), leading to better approximation of the optimal parameters, (3) learning both  $u$  and  $v$  as well as the prior  $\pi$ , diffusion coefficient  $\sigma$ , and terminal time  $T$  (see Figure 5) can reduce the discretization error.

## A.5 AUXILIARY RESULTS

**Theorem A.10** (Girsanov theorem). *For  $\vec{\mathbb{P}}^{u, \pi}$ -almost every  $Z \in C([0, T], \mathbb{R}^D)$  it holds that*

$$\log \frac{d\vec{\mathbb{P}}^{u, \pi}}{d\vec{\mathbb{P}}^{w, \pi}}(Z) = - \int_0^T \left( \frac{1}{2} \|u - w\|^2 + (\eta^+ f + w) \cdot (u - w) \right) (Z_s, s) ds + S \quad (28a)$$

$$= \frac{1}{2} \int_0^T (\|\eta^+ f + w\|^2 - \|\eta^+ f + u\|^2) (Z_s, s) ds + S, \quad (28b)$$

where

$$S = \int_0^T (u - w)(Z_s, s) \cdot \eta^+(s) dZ_s. \quad (29)$$



In particular, for  $Z \sim \tilde{\mathbb{P}}^{u,\pi}$  we obtain that

$$\log \frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{w,\pi}}(Z) = -\frac{1}{2} \int_0^T \|u - w\|^2(Z_s, s) ds + \int_0^T (u - w)(Z_s, s) \cdot \tilde{dB}_s. \quad (30)$$

*Proof.* See Särkkä & Sottinen (2007); Üstünel & Zakai (2013); Chen et al. (2022).  $\square$

**Theorem A.11** (Reverse-time Girsanov theorem). *For  $\tilde{\mathbb{P}}^{u,\pi}$ -almost every  $Z \in C([0, T], \mathbb{R}^D)$  it holds that*

$$\log \frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{w,\pi}}(Z) = \log \frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{w,\pi}}(Z) - \int_0^T (u - w)(Z_s, s) \cdot \eta^+(s) \tilde{dZ}_s \quad (31)$$

$$+ \int_0^T (u - w)(Z_s, s) \cdot \eta^+(s) \tilde{dZ}_s. \quad (32)$$

*Proof.* Using Theorem A.10 and the definitions in (19) and (20), we observe that  $\frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{w,\pi}}(Z)$  equals the Radon-Nikodym derivative between the path spaces measures corresponding to forward SDEs as in (5) with initial conditions  $\pi$  and all functions  $f$ ,  $u$ ,  $w$ , and  $\eta$  reversed in time, evaluated at  $t \mapsto Z_{T-t}$ . We can now substitute  $t \mapsto T - t$  to proof the claim; see also Vargas et al. (2024, Proof of Proposition 2.2).  $\square$

**Lemma A.12** (Conversion formula). *For  $Z \sim \mathbb{P}^{w,\pi}$  and suitable  $g \in C(\mathbb{R}^D \times [0, T], \mathbb{R}^D)$  it holds that*

$$\int_{\underline{t}}^{\bar{t}} g(Z_s, s) \cdot \tilde{dZ}_s = \int_{\underline{t}}^{\bar{t}} g(Z_s, s) \cdot \vec{dZ}_s + \int_{\underline{t}}^{\bar{t}} \nabla \cdot (\eta \eta^\top g)(Z_s, s) ds. \quad (33)$$

*Proof.* Similar to the conversion formula in Vargas et al. (2024, Remark 3), the result follows from combining (19) and (20). First, we rewrite the problem by observing that

$$\int_{\underline{t}}^{\bar{t}} g(Z_s, s) \cdot \tilde{dZ}_s = \int_{\underline{t}}^{\bar{t}} g(Z_s, s) \cdot \vec{dZ}_s + \int_{\underline{t}}^{\bar{t}} \tilde{g}(Z_s, s) \cdot \tilde{dW}_s - \int_{\underline{t}}^{\bar{t}} \tilde{g}(Z_s, s) \cdot \vec{dW}_s,$$

where  $\tilde{g} = \eta^\top g$ . Then we can compute

$$\begin{aligned} \int_{\underline{t}}^{\bar{t}} \tilde{g}(Z_s, s) \cdot \tilde{dW}_s &= \lim_{n \rightarrow \infty} \sum_{i=0}^{k_n} (\tilde{g}(Z_{t_{i+1}^n}, t_{i+1}^n) + \tilde{g}(Z_{t_i^n}, t_i^n)) \cdot (W_{t_{i+1}^n} - W_{t_i^n}) - \int_{\underline{t}}^{\bar{t}} \tilde{g}(Z_s, s) \cdot \vec{dW}_s \\ &= 2 \int_{\underline{t}}^{\bar{t}} \tilde{g}(Z_s, s) \circ dW_s - \int_{\underline{t}}^{\bar{t}} \tilde{g}(Z_s, s) \cdot \vec{dW}_s, \end{aligned}$$

where  $\circ$  denotes Stratonovich integration. The result now follows from the relationship between Itô and Stratonovich stochastic integrals, i.e.,

$$\int_{\underline{t}}^{\bar{t}} \tilde{g}(Z_s, s) \circ dW_s = \int_{\underline{t}}^{\bar{t}} \tilde{g}(Z_s, s) \cdot \vec{dW}_s + \frac{1}{2} \int_{\underline{t}}^{\bar{t}} \nabla \cdot (\eta \tilde{g})(Z_s, s) ds, \quad (34)$$

see, e.g., Kloeden & Platen (1992, Section 4.9).  $\square$

## A.6 PROOFS

*Proof of Proposition 2.3.* The proof follows the one by Vargas et al. (2024, proof of Proposition 2.2). Using disintegration (Léonard, 2014), we first observe that<sup>11</sup>  $\frac{d\tilde{\mathbb{P}}^{w,\tau}}{d\tilde{\mathbb{P}}^{w,\pi}}(Z) = \frac{\tau(Z_T)}{\pi(Z_0)}$  for  $w = -\eta^+ f$ . Thus, it holds that

$$\log \frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{v,\tau}}(Z) = \log \frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{w,\pi}}(Z) + \log \frac{d\tilde{\mathbb{P}}^{w,\tau}}{d\tilde{\mathbb{P}}^{v,\tau}}(Z) + \log \frac{\pi(Z_0)}{\tau(Z_T)}. \quad (35)$$

<sup>11</sup>Considering the (kinetic) Fokker-Planck equation in (10), the Lebesgue measure is an invariant measure of the SDE in (5) with control  $w = -\eta^+ f$  if and only if  $\hat{f}$  merely depends on the last  $d$  coordinates.

The result now follows by applying the Girsanov theorem; see Theorem A.10 and Theorem A.11.  $\square$

*Proof of Lemma A.1.* Using Lemma 2.2 and the chain rule for the KL divergence, we observe that

$$D_{\text{KL}}(\tilde{\mathbb{P}}^{v,\tau} | \tilde{\mathbb{P}}^{u,\pi}) = D_{\text{KL}}(\tilde{\mathbb{P}}^{v,\tau} | \tilde{\mathbb{P}}^{\tilde{v},\tilde{\tau}}) = D_{\text{KL}}(\tilde{\mathbb{P}}^{v,\tau} | \tilde{\mathbb{P}}^{\tilde{v},\tau}) + D_{\text{KL}}(\tau | \tilde{\mathbb{P}}_0^{u,\pi}), \quad (36)$$

where  $\tilde{\tau} = \tilde{\mathbb{P}}_0^{u,\pi}$ . We note that the Girsanov theorem (see Theorem A.10) implies that the variational gap can equivalently be written as

$$D_{\text{KL}}(\tilde{\mathbb{P}}^{v,\tau} | \tilde{\mathbb{P}}^{\tilde{v},\tau}) = \mathbb{E}_{Z \sim \tilde{\mathbb{P}}^{v,\tau}} \left[ \frac{1}{2} \int_0^T \|v(Z_s, s) - u(Z_s, s) + \eta^\top(s) \nabla \log \tilde{\mathbb{P}}_s^{u,\pi}(Z_s)\|^2 ds \right],$$

see also Vargas et al. (2024, Appendix C).  $\square$

*Proof of Proposition A.2.* The proof extends the ones by Huang et al. (2021, Appendix A), Berner et al. (2024, Lemma A.11), and (Vargas et al., 2024, Appendix C.2) to the case of degenerate diffusion coefficients  $\eta$ . Using Proposition A.13 and a Monte Carlo approximation, we first observe that, for the case  $v = 0$ , the ELBO can be represented as

$$\text{ELBO} = \mathbb{E}_{Z \sim \tilde{\mathbb{P}}^{0,\tau}} \left[ \log \pi(Z_T) - \int_0^T \left( \frac{1}{2} \|u\|^2 - \nabla \cdot (\eta u + f) \right) (Z_s, s) ds \right] \quad (37)$$

$$= -T \mathbb{E}_{Z \sim \tilde{\mathbb{P}}^{0,\tau}, s \sim \text{Unif}([0, T])} \left[ \left( \frac{1}{2} \|u\|^2 - \nabla \cdot (\eta u) \right) (Z_s, s) \right] + \text{const.}, \quad (38)$$

where the last expression can be viewed as an extension of *implicit score matching* (Hyvärinen & Dayan, 2005) to degenerate  $\eta$ .

Completing the square and using the tower property in (37), it remains to show that

$$\mathbb{E}[r(Z_s) | Z_0] = -\mathbb{E}[\nabla \cdot (\eta u)(Z_s, s) | Z_0] \quad (39)$$

for fixed  $s \in [0, T]$ , where we used the abbreviations

$$p(z) := \mathbb{P}_{s|0}^{0,\tau}(z | Z_0) \quad \text{and} \quad r(z) = u(z, s) \cdot (\eta^\top(s) \nabla \log p(z)) = (\eta(s)u(z, s)) \cdot \frac{\nabla p(z)}{p(z)}. \quad (40)$$

Under suitable assumptions, the statement in (39) follows from the computation

$$\mathbb{E}[r(Z_s) | Z_0] = \int_{\mathbb{R}^d} r(z) p(z) dz = \underbrace{\int_{\mathbb{R}^d} \nabla \cdot (\eta u p)(z, s) dz}_{=0} - \int_{\mathbb{R}^d} \nabla \cdot (\eta u)(z, s) p(z) dz \quad (41)$$

$$= -\mathbb{E}[\nabla \cdot (\eta u)(Z_s, s) | Z_0], \quad (42)$$

where we used identities for divergences and Stokes' theorem.  $\square$

## A.7 ADDITIONAL STATEMENTS ON DIFFUSION MODELS

The following proposition is an alternative version of Proposition 2.3, which, instead of backward integrations, depends on the divergence operation and does not rely on computing the pseudoinverse of  $\eta$ .

**Proposition A.13** (Radon-Nikodym derivative). *For a process  $Z \sim \tilde{\mathbb{P}}^{w,\pi}$  as defined in (5) it holds*

$$\begin{aligned} \log \frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{v,\tau}}(Z) &= \log \frac{\pi(Z_0)}{\tau(Z_T)} + \int_0^T \left( (u - v) \cdot \left( w - \frac{u + v}{2} \right) - \nabla \cdot (f + \eta v) \right) (Z_s^w, s) ds \\ &\quad + \int_0^T (u - v)(Z_s, s) \cdot dW_s. \end{aligned}$$

*Proof.* This follows from combining Proposition 2.3 with Lemma A.12. In particular, note that for  $Z \sim \vec{\mathbb{P}}^{w,\pi}$  it holds that

$$\begin{aligned} \log \frac{d\vec{\mathbb{P}}^{u,\pi}}{d\vec{\mathbb{P}}^{v,\tau}}(Z) &= \log \frac{\pi(Z_0)}{\tau(Z_T)} - \frac{1}{2} \int_0^T \|(\eta^+ f + u)\|^2(Z_s, s) ds + \frac{1}{2} \int_0^T \|(\eta^+ f + v)\|^2(Z_s, s) ds \\ &\quad + \int_0^T (\eta^+ f + u)(Z_s, s) \cdot \eta^+(s) dZ_s - \int_0^T (\eta^+ f + v)(Z_s, s) \cdot \eta^+(s) dZ_s \\ &= \log \frac{\pi(Z_0)}{\tau(Z_T)} - \frac{1}{2} \int_0^T \|(\eta^+ f + u)\|^2(Z_s, s) ds + \frac{1}{2} \int_0^T \|(\eta^+ f + v)\|^2(Z_s, s) ds \\ &\quad + \int_0^T (u - v)(Z_s, s) \cdot \eta^+(s) dZ_s - \int_0^T \nabla \cdot (\eta \eta^+ f + \eta v)(Z_s, s) ds \\ &= \log \frac{\pi(Z_0)}{\tau(Z_T)} + \int_0^T \left( (u - v) \cdot \left( w - \frac{u + v}{2} \right) - \nabla \cdot (\eta \eta^+ f + \eta v) \right) (Z_s^w, s) ds \\ &\quad + \int_0^T (u - v)(Z_s, s) \cdot dW_s. \end{aligned}$$

We note that  $\eta \eta^+ = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \text{Id}_{d \times d} \end{pmatrix}$ . Together with our assumption that the first component  $\hat{f}$  of  $f = (\hat{f}, \tilde{f})^\top$  only depends on the last  $d$  coordinates, we obtain that  $\nabla \cdot (\eta \eta^+ f) = \nabla \cdot f$ , which proves the claim.  $\square$

**Remark A.14** (PDE perspective). Similar to Berner et al. (2024); Sun et al. (2024), we can also derive the expression in Proposition A.13 using the underlying PDEs. To this end, we recall that the density  $p$  of the solution  $Z$  to the SDE in stated (5) is governed by the (kinetic) Fokker-Planck equation in (10). Using the Hopf-Cole transformation  $V := \log p$ , we get the Hamilton-Jacobi-Bellman equation

$$\partial_t V = -\text{div}(f + \eta u) - \nabla V \cdot (f + \eta u) + \frac{1}{2} \|\eta^\top \nabla V\|^2 + \frac{1}{2} \text{Tr}(\eta \eta^\top \nabla^2 V).$$

Moreover, Itô's formula implies that

$$\begin{aligned} V(Z_T, T) - V(Z_0, 0) &= \int_0^T \left( \partial_s V + \frac{1}{2} \text{Tr}(\eta \eta^\top \nabla^2 V) + (f + \eta w) \cdot \nabla V \right) (Z_s, s) ds \\ &\quad + \int_0^T \nabla V(Z_s, s) \cdot \eta(s) dW_s, \end{aligned}$$

where  $Z \sim \vec{\mathbb{P}}^{w,\pi}$ . Following the same computations as in Proposition 3.1 in Sun et al. (2024) and minimizing the squared residual of the above Itô formula, we obtain the loss

$$\mathcal{L}^w(u, v) = \mathbb{E}_{Z \sim \vec{\mathbb{P}}^{w,\pi}} \left[ \left( \log \frac{d\vec{\mathbb{P}}^{u,\pi}}{d\vec{\mathbb{P}}^{v,\tau}}(Z) \right)^2 \right],$$

where the Radon-Nikodym derivative is equivalent to the one given in Proposition A.13.

#### A.8 INCLUDING A MASS MATRIX

In Section 3, we omitted the mass matrix  $M$  for simplicity. Here, we give further details on the SDEs when the mass matrix is incorporated. It can be incorporated in the framework presented in Section 2 by making the choices  $D = 2d$ ,  $Z = (X, Y)^\top$  and

$$f(x, y, s) = (y, \tilde{f}(x, y, s) - \frac{1}{2} \sigma \sigma^\top(s) y)^\top, \quad \eta = (\mathbf{0}_d, \sigma M^{1/2})^\top \quad (43)$$

in (5) and (6), where  $\mathbf{0}_d \in \mathbb{R}^{d \times d}$ ,  $\tilde{f} \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$  is chosen appropriately and  $\sigma, M \in C([0, T], \mathbb{R}^{d \times d})$ . For the terminal conditions, the standard normal for the last  $d$  components of the initial and terminal distributions is replaced by a Gaussian whose covariance matrix is given by the mass, i.e.,

$$\pi(x, y) = p_{\text{prior}}(x) \mathcal{N}(y; 0, M) \quad \text{and} \quad \tau(x, y) = p_{\text{target}}(x) \mathcal{N}(y; 0, M). \quad (44)$$

We, therefore, get the forward and reverse-time processes

$$dX_s = M^{-1}Y_s ds, \quad X_0 \sim p_{\text{prior}}, \quad (45a)$$

$$dY_s = \left( \tilde{f}(Z_s, s) - \frac{1}{2}\sigma\sigma^\top(s)Y_s + \sigma M^{1/2}u(Z_s, s) \right) ds + \sigma(s)M^{1/2}\tilde{d}W_s, \quad Y_0 \sim \mathcal{N}(0, M), \quad (45b)$$

and

$$dX_s = M^{-1}Y_s ds, \quad X_T \sim p_{\text{target}}, \quad (46a)$$

$$dY_s = \left( \tilde{f}(Z_s, s) - \frac{1}{2}\sigma\sigma^\top(s)Y_s + \sigma M^{1/2}v(Z_s, s) \right) ds + \sigma(s)M^{1/2}\tilde{d}W_s, \quad Y_T \sim \mathcal{N}(0, M). \quad (46b)$$

In a similar spirit to the diffusion matrix  $\sigma$ , one can also learn the mass matrix. However, our experiments (Section 4) showed little improvements when doing so.

#### A.9 NUMERICAL DISCRETIZATION SCHEMES

In this section we provide details on the numerical integration schemes discussed in this work, called OBAB, BAOAB, and OBABO. In particular, we derive the transition kernels  $\vec{p}$  and  $\tilde{p}$  for computing the discrete-time approximation of the Radon-Nikodym derivative as

$$\frac{d\tilde{\mathbb{P}}^{u,\pi}}{d\tilde{\mathbb{P}}^{v,\tau}}(Z) \approx \frac{\pi(\hat{Z}_0) \prod_{n=0}^{N-1} \vec{p}_{n+1|n}(\hat{Z}_{n+1}|\hat{Z}_n)}{\tau(\hat{Z}_N) \prod_{n=0}^{N-1} \tilde{p}_{n|n+1}(\hat{Z}_n|\hat{Z}_{n+1})}. \quad (47)$$

We note that such splitting schemes are well-studied in the uncontrolled setting, see, e.g., Section 7 in [Leimkuhler & Matthews \(2015\)](#) or Section 2.2.3.2 in [Stoltz et al. \(2010\)](#). The controlled setting, and in particular the approximation of the Radon-Nikodym derivative between path space measures, has to the best of our knowledge only been considered for OBAB yet [Geffner & Domke \(2023\)](#); [Doucet et al. \(2022b\)](#).

For convenience, let us recall the following splitting for the forward SDE that is used throughout this section, i.e.,

$$\begin{bmatrix} dX_s \\ dY_s \end{bmatrix} = \underbrace{\begin{bmatrix} Y_s \\ 0 \end{bmatrix}}_{\vec{A}} ds + \underbrace{\begin{bmatrix} 0 \\ \tilde{f}(X_s, s) \end{bmatrix}}_{\vec{B}} ds + \underbrace{\begin{bmatrix} 0 \\ (-\frac{1}{2}\sigma\sigma^\top(s)Y_s + \sigma u(Z_s, s)) ds + \sigma(s)\tilde{d}W_s \end{bmatrix}}_{\vec{O}}, \quad (48)$$

and let us use the following splitting for the reverse SDE

$$\begin{bmatrix} dX_s \\ dY_s \end{bmatrix} = \underbrace{\begin{bmatrix} Y_s \\ 0 \end{bmatrix}}_{\vec{A}} ds + \underbrace{\begin{bmatrix} 0 \\ \tilde{f}(X_s, s) \end{bmatrix}}_{\vec{B}} ds + \underbrace{\begin{bmatrix} 0 \\ (-\frac{1}{2}\sigma\sigma^\top(s)Y_s + \sigma v(Z_s, s)) ds + \sigma(s)\tilde{d}W_s \end{bmatrix}}_{\vec{O}}. \quad (49)$$

Here, we use arrows to indicate whether the corresponding splitting belongs to the generative or inference SDE. To simplify the notation, we define  $\sigma_n := \sigma(n\Delta)$ ,  $f_n := \tilde{f}(X_{n\Delta}, n\Delta)$ ,  $\vec{p}_{n+1|n} := \vec{p}_{n+1|n}(\hat{Z}_{n+1}|\hat{Z}_n)$  and analogously for the backward transition  $\tilde{p}_{n|n+1}$ .

##### A.9.1 EULER-MARUYAMA

We follow [Geffner & Domke \(2023\)](#) and leverage a semi-implicit Euler-Maruyama (EM) scheme, where the velocity update is computed first and then used to move the position, i.e.,

$$\hat{Y}_{n+1} = \hat{Y}_n(1 - \frac{1}{2}\sigma_n\sigma_n^\top\Delta) + \sigma_n u(\hat{Z}_n, n\Delta)\Delta + \tilde{f}_n\Delta + \sigma_n\sqrt{\Delta}\xi_n, \quad (50)$$

$$\hat{X}_{n+1} = \hat{X}_n + \hat{Y}_{n+1}\Delta, \quad (51)$$

with  $\xi_n \sim \mathcal{N}(0, I)$  and step size  $\Delta > 0$ . We further define  $\hat{Z}_{n+1} = \Phi(\hat{X}_n, \hat{Y}_{n+1}) := (\hat{X}_n + \hat{Y}_{n+1}\Delta, \hat{Y}_{n+1})^\top$  which helps to simplify the discrete-time approximation of the Radon-Nikodym derivative as shown in the following. We obtain the forward transition density

$$\vec{p}_{n+1|n} = \mathcal{N}\left(\hat{Y}_{n+1} \middle| \hat{Y}_n(1 - \frac{1}{2}\sigma_n\sigma_n^\top\Delta) + \sigma_n u(\hat{Z}_n, n\Delta)\Delta + \tilde{f}_n\Delta, \sigma_n\sigma_n^\top\Delta\right) \times \delta_{\Phi(\hat{X}_n, \hat{Y}_{n+1})}(\hat{Z}_{n+1}),$$

where  $\delta$  is the Dirac delta distribution. Integrating the reverse process (49) using the semi-implicit EM scheme, we obtain

$$\hat{X}_n = \hat{X}_{n+1} - \hat{Y}_{n+1}\Delta \quad (52)$$

$$\hat{Y}_n = \hat{Y}_{n+1}(1 + \frac{1}{2}\sigma_{n+1}\sigma_{n+1}^\top\Delta) - \sigma_{n+1}v(\hat{Z}_{n+1}, (n+1)\Delta)\Delta - \tilde{f}_{n+1}\Delta + \sigma_{n+1}\sqrt{\Delta}\xi_{n+1}, \quad (53)$$

with corresponding transition density

$$\begin{aligned} \tilde{p}_{n|n+1} = & \mathcal{N}\left(\hat{Y}_n \left| \hat{Y}_{n+1}(1 + \frac{1}{2}\sigma_{n+1}\sigma_{n+1}^\top\Delta) - \sigma_{n+1}v(\hat{Z}_{n+1}, (n+1)\Delta)\Delta - \tilde{f}_{n+1}\Delta, \sigma_{n+1}\sigma_{n+1}^\top\Delta\right.\right) \\ & \times \delta_{\Phi^{-1}(\hat{Z}_{n+1})}(\hat{X}_n, \hat{Y}_{n+1}), \end{aligned} \quad (54)$$

resulting in the following ratio between forward and backward transitions

$$\frac{\vec{p}_{n+1|n}}{\vec{p}_{n|n+1}} = \frac{\mathcal{N}\left(\hat{Y}_{n+1} \left| \hat{Y}_n(1 - \frac{1}{2}\sigma_n\sigma_n^\top\Delta) + \sigma_n u(\hat{Z}_n, n\Delta)\Delta + \tilde{f}_n\Delta, \sigma_n\sigma_n^\top\Delta\right.\right)}{\mathcal{N}\left(\hat{Y}_n \left| \hat{Y}_{n+1}(1 + \frac{1}{2}\sigma_{n+1}\sigma_{n+1}^\top\Delta) - \sigma_{n+1}v(\hat{Z}_{n+1}, (n+1)\Delta)\Delta - \tilde{f}_{n+1}\Delta, \sigma_{n+1}\sigma_{n+1}^\top\Delta\right.\right)}, \quad (55)$$

as the ratio between the two Dirac delta distribution cancel.

### A.9.2 OBAB

Composing the splitting terms as  $\vec{\text{OBAB}}$  yields the integrator

$$\hat{Y}'_n = \hat{Y}_n(1 - \frac{1}{2}\sigma_n\sigma_n^\top\Delta) + \sigma_n u(\hat{Z}_n, n\Delta)\Delta + \sigma_n\sqrt{\Delta}\xi_n, \quad \xi_n \sim \mathcal{N}(0, I) \quad (56a)$$

$$\left. \begin{aligned} \hat{Y}''_n &= \hat{Y}'_n + \tilde{f}_n \frac{\Delta}{2} \\ \hat{X}_{n+1} &= \hat{X}_n + \hat{Y}''_n \Delta \\ \hat{Y}_{n+1} &= \hat{Y}''_n + \tilde{f}_{n+1} \frac{\Delta}{2} \end{aligned} \right\} \Phi \quad (56b)$$

with  $\hat{Z}_{n+1} = \Phi(\hat{X}_n, \hat{Y}'_n)$ . The resulting forward transition is given by

$$\vec{p}_{n+1|n} = \delta_{\Phi(\hat{X}_n, \hat{Y}'_n)}(\hat{Z}_{n+1}) \mathcal{N}\left(\hat{Y}'_n \left| \hat{Y}_n(1 - \frac{1}{2}\sigma_n\sigma_n^\top\Delta) + \sigma_n u(\hat{Z}_n, n\Delta)\Delta, \sigma_n\sigma_n^\top\Delta\right.\right).$$

The inference SDE, i.e.,  $\vec{\text{OBAB}}$ , is integrated as

$$\left. \begin{aligned} \hat{Y}''_n &= \hat{Y}_{n+1} - \tilde{f}_{n+1} \frac{\Delta}{2} \\ \hat{X}_n &= \hat{X}_{n+1} - \hat{Y}''_n \Delta \\ \hat{Y}'_n &= \hat{Y}''_n - \tilde{f}_n \frac{\Delta}{2} \end{aligned} \right\} \Phi^{-1} \quad (57)$$

$$\hat{Y}_n = \hat{Y}'_n(1 + \frac{1}{2}\sigma_n\sigma_n^\top\Delta) - \sigma_n v(\hat{Z}'_n, n\Delta)\Delta + \sigma_n\sqrt{\Delta}\xi_n, \quad \xi_n \sim \mathcal{N}(0, I), \quad (58)$$

with  $(\hat{X}_n, \hat{Y}'_n) = \Phi^{-1}(\hat{Z}_{n+1})$ , giving the backward transition

$$\tilde{p}_{n|n+1} = \delta_{\Phi^{-1}(\hat{Z}_{n+1})}(\hat{X}_n, \hat{Y}'_n) \mathcal{N}\left(\hat{Y}_n \left| \hat{Y}'_n(1 + \frac{1}{2}\sigma_n\sigma_n^\top\Delta) - \sigma_n v(\hat{Z}'_n, n\Delta)\Delta, \sigma_n\sigma_n^\top\Delta\right.\right).$$

This results in the following ratio between forward and backward transitions

$$\frac{\vec{p}_{n+1|n}}{\vec{p}_{n|n+1}} = \frac{\mathcal{N}\left(\hat{Y}'_n \left| \hat{Y}_n(1 - \frac{1}{2}\sigma_n\sigma_n^\top\Delta) + \sigma_n u(\hat{Z}_n, n\Delta)\Delta, \sigma_n\sigma_n^\top\Delta\right.\right)}{\mathcal{N}\left(\hat{Y}_n \left| \hat{Y}'_n(1 + \frac{1}{2}\sigma_n\sigma_n^\top\Delta) - \sigma_n v(\hat{Z}'_n, n\Delta)\Delta, \sigma_n\sigma_n^\top\Delta\right.\right)}. \quad (59)$$



### A.9.3 BAOAB

Composing the splitting terms as  $\vec{\text{B}}\vec{\text{A}}\vec{\text{O}}\vec{\text{A}}\vec{\text{B}}$  yields the integrator

$$\left. \begin{aligned} \hat{Y}'_n &= \hat{Y}_n + \tilde{f}_n \frac{\Delta}{2} \\ \hat{X}'_n &= \hat{X}_n + \hat{Y}'_n \frac{\Delta}{2} \end{aligned} \right\} \Phi_1 \quad (60)$$

$$\hat{Y}''_n = \hat{Y}'_n (1 - \frac{1}{2} \sigma_n \sigma_n^\top \Delta) + \sigma_n u(\hat{X}'_n, \hat{Y}'_n, n\Delta) \Delta + \sigma_n \sqrt{\Delta} \xi_n \quad (61)$$

$$\left. \begin{aligned} \hat{X}_{n+1} &= \hat{X}'_n + \hat{Y}''_n \frac{\Delta}{2} \\ \hat{Y}_{n+1} &= \hat{Y}''_n + \tilde{f}_{n+1} \frac{\Delta}{2} \end{aligned} \right\} \Phi_2 \quad (62)$$

with  $\xi_n \sim \mathcal{N}(0, I)$ ,  $(\hat{X}'_n, \hat{Y}'_n) = \Phi_1(\hat{Z}_n)$ , and  $\hat{Z}_{n+1} = \Phi_2(\hat{X}'_n, \hat{Y}'_n)$ . Hence, we obtain the forward transition density

$$\begin{aligned} \tilde{p}_{n+1|n} &= \delta_{\Phi_2(\hat{X}'_n, \hat{Y}'_n)}(\hat{Z}_{n+1}) \times \delta_{\Phi_1(\hat{Z}_n)}(\hat{X}'_n, \hat{Y}'_n) \\ &\times \mathcal{N}\left(\hat{Y}''_n \left| \hat{Y}'_n (1 - \frac{1}{2} \sigma_n \sigma_n^\top \Delta) + \sigma_n u(\hat{X}'_n, \hat{Y}'_n, n\Delta) \Delta, \sigma_n \sigma_n^\top \Delta\right.\right). \end{aligned} \quad (63)$$

For  $\vec{\text{B}}\vec{\text{A}}\vec{\text{O}}\vec{\text{A}}\vec{\text{B}}$  we obtain

$$\left. \begin{aligned} \hat{Y}''_n &= \hat{Y}_{n+1} - \tilde{f}_{n+1} \frac{\Delta}{2} \\ \hat{X}'_n &= \hat{X}_{n+1} - \hat{Y}''_n \frac{\Delta}{2} \end{aligned} \right\} \Phi_2^{-1} \quad (64)$$

$$\hat{Y}'_n = \hat{Y}''_n (1 + \frac{1}{2} \sigma_n \sigma_n^\top \Delta) - \sigma_n v(\hat{X}'_n, \hat{Y}''_n, n\Delta) \Delta + \sigma_n \sqrt{\Delta} \xi_n \quad (65)$$

$$\left. \begin{aligned} \hat{X}_n &= \hat{X}'_n - \hat{Y}'_n \frac{\Delta}{2} \\ \hat{Y}_n &= \hat{Y}'_n - \tilde{f}_n \frac{\Delta}{2} \end{aligned} \right\} \Phi_1^{-1} \quad (66)$$

with  $(\hat{X}'_n, \hat{Y}'_n) = \Phi_2^{-1}(\hat{Z}_{n+1})$  and  $\hat{Z}_n = \Phi_1^{-1}(\hat{X}'_n, \hat{Y}'_n)$ . Moreover, we have

$$\begin{aligned} \tilde{p}_{n|n+1} &= \delta_{\Phi_1^{-1}(\hat{X}'_n, \hat{Y}'_n)}(\hat{Z}_n) \times \delta_{\Phi_2^{-1}(\hat{Z}_{n+1})}(\hat{X}'_n, \hat{Y}'_n) \\ &\times \mathcal{N}\left(\hat{Y}'_n \left| \hat{Y}''_n (1 + \frac{1}{2} \sigma_n \sigma_n^\top \Delta) - \sigma_n v(\hat{X}'_n, \hat{Y}''_n, n\Delta) \Delta, \sigma_n \sigma_n^\top \Delta\right.\right). \end{aligned} \quad (67)$$

We therefore obtain the following ratio between forward and backward transitions as

$$\frac{\tilde{p}_{n+1|n}}{\tilde{p}_{n|n+1}} = \frac{\mathcal{N}\left(\hat{Y}''_n \left| \hat{Y}'_n (1 - \frac{1}{2} \sigma_n \sigma_n^\top \Delta) + \sigma_n u(\hat{X}'_n, \hat{Y}'_n, n\Delta) \Delta, \sigma_n \sigma_n^\top \Delta\right.\right)}{\mathcal{N}\left(\hat{Y}'_n \left| \hat{Y}''_n (1 + \frac{1}{2} \sigma_n \sigma_n^\top \Delta) - \sigma_n v(\hat{X}'_n, \hat{Y}''_n, n\Delta) \Delta, \sigma_n \sigma_n^\top \Delta\right.\right)}. \quad (68)$$

### A.9.4 OBABO

Composing the splitting terms as  $\vec{\text{O}}\vec{\text{B}}\vec{\text{A}}\vec{\text{B}}\vec{\text{O}}$  yields the integrator

$$\hat{Y}'_n = \hat{Y}_n (1 - \frac{1}{4} \sigma_n \sigma_n^\top \Delta) + \sigma_n u(\hat{Z}_n, n\Delta) \frac{\Delta}{2} + \sigma_n \sqrt{\frac{\Delta}{2}} \xi_n^{(1)} \quad (69)$$

$$\left. \begin{aligned} \hat{Y}''_n &= \hat{Y}'_n + \tilde{f}_n \frac{\Delta}{2} \\ \hat{X}_{n+1} &= \hat{X}_n + \hat{Y}''_n \Delta \\ \hat{Y}'''_n &= \hat{Y}''_n + \tilde{f}_{n+1} \frac{\Delta}{2} \end{aligned} \right\} \Phi \quad (70)$$

$$\hat{Y}_{n+1} = \hat{Y}'''_n (1 - \frac{1}{4} \sigma_n \sigma_n^\top \Delta) + \sigma_n u(\hat{X}_{n+1}, \hat{Y}'''_n, (n + \frac{1}{2})\Delta) \frac{\Delta}{2} + \sigma_n \sqrt{\frac{\Delta}{2}} \xi_n^{(2)} \quad (71)$$

with  $\xi_n^{(1)}, \xi_n^{(2)} \sim \mathcal{N}(0, I)$  and  $(\hat{X}_{n+1}, \hat{Y}'''_n) = \Phi(\hat{X}_n, \hat{Y}'_n)$ . The resulting forward transition density is given by

$$\begin{aligned} \tilde{p}_{n+1|n} &= \mathcal{N}\left(\hat{Y}_{n+1} \left| \hat{Y}'''_n (1 - \frac{1}{4} \sigma_n \sigma_n^\top \Delta) + \sigma_n u(\hat{X}_{n+1}, \hat{Y}'''_n, (n + \frac{1}{2})\Delta) \frac{\Delta}{2}, \frac{1}{2} \sigma_n \sigma_n^\top \Delta\right.\right) \\ &\times \delta_{\Phi(\hat{X}_n, \hat{Y}'_n)}(\hat{X}_{n+1}, \hat{Y}'''_n) \\ &\times \mathcal{N}\left(\hat{Y}'_n \left| \hat{Y}_n (1 - \frac{1}{4} \sigma_n \sigma_n^\top \Delta) + \sigma_n u(\hat{Z}_n, n\Delta) \frac{\Delta}{2}, \frac{1}{2} \sigma_n \sigma_n^\top \Delta\right.\right). \end{aligned} \quad (72)$$

The inference SDE, i.e.,  $\tilde{\mathbf{O}}\tilde{\mathbf{B}}\tilde{\mathbf{A}}\tilde{\mathbf{B}}\tilde{\mathbf{O}}$ , is integrated as

$$\hat{Y}_n''' = \hat{Y}_{n+1}(1 + \frac{1}{4}\sigma_n\sigma_n^\top\Delta) - \sigma_nv(\hat{Z}_{n+1}, (n+1)\Delta)\frac{\Delta}{2} + \sigma_n\sqrt{\frac{\Delta}{2}}\xi_n^{(2)} \quad (73)$$

$$\left. \begin{aligned} \hat{Y}_n'' &= \hat{Y}_n''' - \tilde{f}_{n+1}\frac{\Delta}{2} \\ \hat{X}_n &= \hat{X}_{n+1} - \hat{Y}_n''\Delta \\ \hat{Y}_n' &= \hat{Y}_n'' - \tilde{f}_n\frac{\Delta}{2} \end{aligned} \right\} \Phi^{-1} \quad (74)$$

$$\hat{Y}_n = \hat{Y}_n'(1 + \frac{1}{4}\sigma_n\sigma_n^\top\Delta) - \sigma_nv(\hat{X}_n, \hat{Y}_n', (n + \frac{1}{2})\Delta)\frac{\Delta}{2} + \sigma_n\sqrt{\frac{\Delta}{2}}\xi_n^{(1)}, \quad (75)$$

with  $(\hat{X}_n, \hat{Y}_n') = \Phi^{-1}(\hat{X}_{n+1}, \hat{Y}_n''')$ , and gives the following backward transition densities

$$\begin{aligned} \tilde{p}_{n|n+1} &= \mathcal{N}\left(\hat{Y}_n' \middle| \hat{Y}_n'(1 + \frac{1}{4}\sigma_n\sigma_n^\top\Delta) - \sigma_nv(\hat{X}_n, \hat{Y}_n', (n + \frac{1}{2})\Delta)\frac{\Delta}{2}, \frac{1}{2}\sigma_n\sigma_n^\top\Delta\right) \\ &\quad \times \delta_{\Phi^{-1}(\hat{X}_{n+1}, \hat{Y}_n''')}(\hat{X}_n, \hat{Y}_n') \\ &\quad \times \mathcal{N}\left(\hat{Y}_n''' \middle| \hat{Y}_{n+1}(1 + \frac{1}{4}\sigma_n\sigma_n^\top\Delta) - \sigma_nv(\hat{Z}_{n+1}, (n+1)\Delta)\frac{\Delta}{2}, \frac{1}{2}\sigma_n\sigma_n^\top\Delta\right), \end{aligned} \quad (76)$$

resulting in the following ratio between forward and backward transitions

$$\begin{aligned} \frac{\tilde{p}_{n+1|n}}{\tilde{p}_{n|n+1}} &= \frac{\mathcal{N}\left(\hat{Y}_{n+1} \middle| \hat{Y}_n'''(1 - \frac{1}{4}\sigma_n\sigma_n^\top\Delta) + \sigma_nu(\hat{X}_{n+1}, \hat{Y}_n''', (n + \frac{1}{2})\Delta)\frac{\Delta}{2}, \frac{1}{2}\sigma_n\sigma_n^\top\Delta\right)}{\mathcal{N}\left(\hat{Y}_n''' \middle| \hat{Y}_{n+1}(1 + \frac{1}{4}\sigma_n\sigma_n^\top\Delta) - \sigma_nv(\hat{Z}_{n+1}, (n+1)\Delta)\frac{\Delta}{2}, \frac{1}{2}\sigma_n\sigma_n^\top\Delta\right)} \\ &\quad \times \frac{\mathcal{N}\left(\hat{Y}_n' \middle| \hat{Y}_n(1 - \frac{1}{4}\sigma_n\sigma_n^\top\Delta) + \sigma_nu(\hat{Z}_n, n\Delta)\frac{\Delta}{2}, \frac{1}{2}\sigma_n\sigma_n^\top\Delta\right)}{\mathcal{N}\left(\hat{Y}_n' \middle| \hat{Y}_n'(1 + \frac{1}{4}\sigma_n\sigma_n^\top\Delta) - \sigma_nv(\hat{X}_n, \hat{Y}_n', (n + \frac{1}{2})\Delta)\frac{\Delta}{2}, \frac{1}{2}\sigma_n\sigma_n^\top\Delta\right)}. \end{aligned}$$

#### A.10 UNDERDAMPED VERSION OF PREVIOUS DIFFUSION-BASED SAMPLING METHODS

In this section we outline how our framework in Section 2 includes previous diffusion-based sampling methods. First, we note that setting the drift  $\tilde{f}$  and controls  $u$  and  $v$  in (13) to specific values recovers underdamped methods of ULA, MCD, and CMCD, see Table 2. Moreover, we can also introduce reference processes with controls  $\tilde{u}$  and  $\tilde{v}$  that satisfy

$$\frac{d\tilde{\mathbb{P}}^{\tilde{u}, \tilde{\pi}}}{d\tilde{\mathbb{P}}^{\tilde{v}, \tilde{\tau}}} \equiv 1, \quad (77)$$

where  $\tilde{\pi}$  and  $\tilde{\tau}$  are known reference distributions. In other words, this assumes having knowledge of a perfect time-reversal for specific controls  $\tilde{u}, \tilde{v}$  and marginals  $\tilde{\pi}, \tilde{\tau}$  (cf. Section 3.3 in Richter & Berner (2024)). We remark that these processes take a role similar to the Brownian motion used in the proof of Proposition 2.3. In particular, by applying Proposition 2.3 twice, we obtain that

$$\begin{aligned} \log \frac{d\tilde{\mathbb{P}}^{u, \pi}}{d\tilde{\mathbb{P}}^{v, \tau}}(Z) &= \log \frac{d\tilde{\mathbb{P}}^{u, \pi}}{d\tilde{\mathbb{P}}^{v, \tau}}(Z) - \log \frac{d\tilde{\mathbb{P}}^{\tilde{u}, \tilde{\pi}}}{d\tilde{\mathbb{P}}^{\tilde{v}, \tilde{\tau}}}(Z) \\ &= \log \frac{\pi(Z_0)}{\tilde{\pi}(Z_0)} - \log \frac{\tau(Z_T)}{\tilde{\tau}(Z_T)} \\ &\quad + \frac{1}{2} \int_0^T ((v - \tilde{v}) \cdot (2\eta^+ f + v + \tilde{v}) - (u - \tilde{u}) \cdot (2\eta^+ f + u + \tilde{u}))(Z_s, s) ds \\ &\quad + \int_0^T (u - \tilde{u})(Z_s, s) \cdot \eta^+(s) dZ_s - \int_0^T (v - \tilde{v})(Z_s, s) \cdot \eta^+(s) dZ_s. \end{aligned}$$

Several previous methods, such as versions of PIS and DDS, can be recovered by fixing  $v$  and using the choices  $\tilde{v} = v$  as well as  $\tilde{\pi} = \pi$ , which significantly simplifies the above expression (Vargas et al., 2024; Richter & Berner, 2024).

#### A.11 FURTHER COMPUTATIONAL DETAILS

In this section we provide additional computational details.

### A.11.1 ALGORITHM

In Algorithm 1 we display the training process of our underdamped diffusion bridge sampler.

### A.11.2 SDE PRECONDITIONING

Here, we introduce a preconditioning technique that modifies the SDEs to improve numerical behavior while preserving the optimal control for both overdamped and underdamped diffusions.

**Overdamped diffusions.** Consider the controlled forward and backward pair of SDEs given by

$$dX_s = (f + \sigma u)(X_s, s) ds + \sigma(s) d\tilde{W}_s, \quad X_0 \sim p_{\text{prior}}, \quad (78)$$

$$dX_s = (f + \sigma v)(X_s, s) ds + \sigma(s) d\tilde{W}_s, \quad X_T \sim p_{\text{target}}, \quad (79)$$

with forward and backward Brownian motion  $d\tilde{W}_s$  and  $d\tilde{W}_s$ ,  $f \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$ , diffusion coefficient  $\sigma \in C([0, T], \mathbb{R}^{d \times d})$ , control functions  $u, v \in C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$  and path measures  $\tilde{\mathbb{P}}^{u, p_0}$  and  $\tilde{\mathbb{P}}^{v, p_T}$  corresponding to (78) and (79), respectively, with  $p_0 := p_{\text{prior}}$  and  $p_T := p_{\text{target}}$ . Note that for denoising diffusion sampler such as DIS, we have<sup>12</sup>  $f(x, \cdot) = 2\sigma\sigma^\top x$  and  $v = 0$ , i.e., an Ornstein-Uhlenbeck (OU) process. While such a choice for  $f$  ensures for a suitable  $\sigma$  that  $\tilde{\mathbb{P}}_0^{v, p_T} \approx \mathcal{N}(0, I)$ , it may lead to a poor initialization for the path measure  $\tilde{\mathbb{P}}^{u, p_0}$ . Similarly, for DBS it helps significantly to choose  $f = \nabla \log p_{\text{target}}$  or  $f = \nabla \log \nu$  where  $\nu(x, s) \propto p_{\text{prior}}^{1-\beta(s)}(x) p_{\text{target}}^{\beta(s)}(x)$ ; see Appendix A.11.6 to obtain a good initialization for  $\tilde{\mathbb{P}}^{u, p_0}$ . However, this may again lead to a poor initialization for  $\tilde{\mathbb{P}}^{v, p_T}$ . We can alleviate these problems by preconditioning of the SDEs: For DIS, we set  $\sigma u = \sigma \tilde{u} - 2f$  with  $f(x, \cdot) = 2\sigma\sigma^\top x$ , resulting in the forward SDE

$$dX_s = (-f + \sigma \tilde{u})(X_s, s) ds + \sigma(s) d\tilde{W}_s, \quad X_0 \sim p_{\text{prior}}. \quad (80)$$

Note that when initializing  $\tilde{u} = 0$ , we have  $\tilde{\mathbb{P}}_t^{\tilde{u}, p_0} = \mathcal{N}(0, I)$  for all  $t \in [0, T]$  which may lead to more stable training. Similarly, for ULA, MCD, and DBS, we set  $\sigma v = \sigma \tilde{v} - 2f$  to obtain the backward SDE

$$dX_s = (-f + \sigma \tilde{v})(X_s, s) ds + \sigma(s) d\tilde{W}_s, \quad X_T \sim p_{\text{target}}, \quad (81)$$

where  $f = \nabla \log \nu$  for ULA and MCD. Moreover, DBS uses preconditioning for all choices of  $f$  discussed in Appendix A.11.6. The impact of preconditioning is illustrated in Figure 6. A few remarks are in order.

**Remark A.15** (Preconditioning for CMCD). The controlled Monte Carlo diffusion sampler (CMCD, Vargas et al. (2024)) prescribes a path of densities  $(\nu(\cdot, t))_{t \in [0, T]}$  and uses  $f = \frac{1}{2}\sigma\sigma^\top \nabla \log \nu$ . Therefore, Nelsons relation (see Lemma 2.2) yields  $u - v = \sigma^\top \nabla \log \nu$  resulting in the pair of SDEs

$$dX_s = (\frac{1}{2}\sigma\sigma^\top \nabla \log \nu + \sigma u)(X_s, s) ds + \sigma(s) d\tilde{W}_s, \quad X_0 \sim p_{\text{prior}}, \quad (82)$$

$$dX_s = (-\frac{1}{2}\sigma\sigma^\top \nabla \log \nu + \sigma u)(X_s, s) ds + \sigma(s) d\tilde{W}_s, \quad X_T \sim p_{\text{target}}, \quad (83)$$

by replacing  $v = u - \sigma^\top \nabla \log \nu$ . As a consequence, CMCD does not require preconditioning as the sign flip of  $f$  naturally occurs due to Nelsons relation.

**Remark A.16** (Preconditioning for DIS). The time-reversed diffusion sampler (DIS, Berner et al. (2024)) and related methods (Zhang & Chen, 2022; Vargas et al., 2023a) incorporate  $\nabla \log p_{\text{target}}$  (or  $\nabla \log \nu$ ) into the SDE by treating it as part of the control, i.e.,

$$u(x, s) = u_1(x, s) + u_2(s) \nabla \log \rho_{\text{target}}(x). \quad (84)$$

However, we found that this approach can lead to worse numerical behavior. As a result, we opted not to use this form of preconditioning.

<sup>12</sup>Please note that the sign differs from most existing literature as the process initialized at the target starts at  $t = T$  instead of  $t = 0$ .

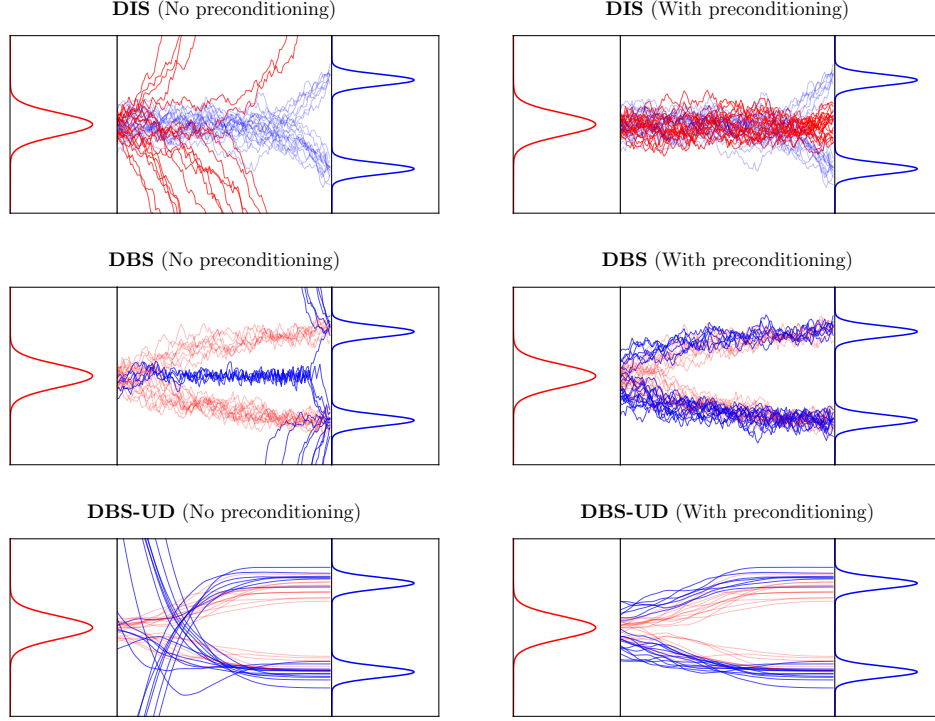


Figure 6: Illustration of the effect of preconditioning as explained in Appendix A.11.2 at initialization, i.e. with  $u, v = 0$ . The red trajectories start from the prior distribution (uni-modal Gaussian distribution) and are integrated forward in time (from left to right), whereas the blue trajectories start at the target (bi-modal distribution) and move backward in time (from right to left). The non-transparent trajectories correspond to the SDE that is affected by preconditioning.

**Underdamped diffusions.** We consider the pair of SDEs given by (5) and (6) where the drift is chosen as in (11), i.e.,

$$f(x, y, s) = \left( y, \tilde{f}(x, s) - \frac{1}{2} \sigma \sigma^\top(s) y \right)^\top. \quad (85)$$

In a similar fashion to the overdamped case, we can precondition the underdamped SDEs: Assume for now that  $\tilde{f}, u = 0$ . Then, (85) results in an OU process for  $Y$  which is desirable as both initial and terminal density of the velocity are Gaussian, see (14). However, the sign flip when performing the time-reversal may again lead to a bad initialization for the path measure  $\tilde{\mathbb{P}}^{v, \tau}$ . We therefore transform the control in the backward process (6) as  $\eta v(\cdot, y, \cdot) = \eta \tilde{v}(\cdot, y, \cdot) + \sigma \sigma^\top y$ , resulting in the backward SDE

$$dY_s = \left( \tilde{f} + \eta \tilde{v} \right) (Y_s, s) ds + \frac{1}{2} \sigma(s) \sigma^\top(s) Y_s ds + \eta(s) d\tilde{W}_s, \quad Y_T \sim \mathcal{N}(0, \text{Id}). \quad (86)$$

The impact of preconditioning for the underdamped setting is illustrated in Figure 6.

### A.11.3 BENCHMARK PROBLEMS

In this section we describe the real-world and synthetic benchmark examples on which we evaluate the different methods.

**Real-world benchmark problems.** We consider seven real-world benchmark problems: Four Bayesian inference tasks, namely *Credit* ( $d = 25$ ), *Cancer* ( $d = 31$ ), *Ionosphere* ( $d = 35$ ), and *Sonar* ( $d = 61$ ). Additionally, we choose *Seeds* ( $d = 26$ ) and *Brownian* ( $d = 32$ ), where the goal is to perform inference over the parameters of a random effect regression model, and the time discretization of a Brownian motion, respectively. Lastly, we consider *LGCP* ( $d = 1600$ ), a high-dimensional Log Gaussian Cox process (Møller et al., 1998).

**Algorithm 1** Training of underdamped diffusion bridge sampler

---

**Require:** ▷ See Appendix A.11 for details

- **model:** neural networks  $u_\theta, v_\gamma$  with initial parameters  $\theta^{(0)}, \gamma^{(0)}$
- **fixed hyperparameters:** number of gradient steps  $K$ , number of discretization steps  $N$ , batch size  $m$ , optimizer method `step`, integrator method `integrate`
- **learned hyperparameters:** prior distribution  $p_{\text{prior}} = \mathcal{N}(\zeta_\mu, \text{diag}(\text{softplus}(\eta_\Sigma)))$ , diffusion and mass matrices  $\sigma = \text{diag}(\text{softplus}(\eta_\sigma))$  and  $M = \text{diag}(\text{softplus}(\eta_M))$ , and terminal time  $T = N \text{softplus}(\eta_\Delta)$  with initial parameters  $\eta_\mu^{(0)}, \eta_\Sigma^{(0)}, \eta_\sigma^{(0)}, \eta_M^{(0)}, \eta_\Delta^{(0)}$

$\Theta^{(0)} = \{\theta^{(0)}, \gamma^{(0)}, \eta_\mu^{(0)}, \eta_\Sigma^{(0)}, \eta_\sigma^{(0)}, \eta_M^{(0)}, \eta_\Delta^{(0)}\}$ ,  $\pi = p_{\text{prior}} \otimes \mathcal{N}(0, M)$ ,  $\tilde{\tau} = \rho_{\text{target}} \otimes \mathcal{N}(0, M)$

**for**  $k \leftarrow 0, \dots, K-1$  **do**

**for**  $i \leftarrow 1, \dots, m$  **do** ▷ Approximate cost (batched in practice)

$\hat{Z}_0 \sim \pi$  ▷ See (14)

$\text{rnd}_{i,0} \leftarrow \log \pi(\hat{Z}_0)$

**for**  $n \leftarrow 0, \dots, N-1$  **do**

$\hat{Z}_{n+1} \leftarrow \text{integrate}(\hat{Z}_n, \Theta^{(k)})$  ▷ See Appendix A.9

$\text{rnd}_{i,n+1} \leftarrow \text{rnd}_{i,n} + \log \tilde{p}_{n+1|n}(\hat{Z}_{n+1} | \hat{Z}_n) - \log \tilde{p}_{n+1}(\hat{Z}_n | \hat{Z}_{n+1})$  ▷ See (9)

$\text{rnd}_{i,N} \leftarrow \text{rnd}_{i,N} - \log \tilde{\tau}(\hat{Z}_N)$

$\hat{\mathcal{L}} \leftarrow \frac{1}{m} \sum_{i=1}^m \text{rnd}_{i,N}$  ▷ Compute loss

$\Theta^{(k+1)} \leftarrow \text{step}(\Theta^{(k)}, \nabla_{\Theta} \hat{\mathcal{L}})$  ▷ Gradient descent

**return** optimized parameters  $\Theta^{(K)}$

---

**Synthetic benchmark problems.** We consider two synthetic benchmark problems in this work: The challenging *Funnel* distribution ( $d = 10$ ) introduced by Neal (2003), whose shape resembles a funnel, where one part is tight and highly concentrated, while the other is spread out over a wide region. Moreover, we choose the *ManyWell* ( $d = 50$ ) target, a highly multi-modal distribution with  $2^5 = 32$  modes.

## A.11.4 EXPERIMENTAL SETUP

Here, we provide further details on our experimental setup. Moreover, we provide an algorithmic description of the training of an underdamped diffusion sampler in Algorithm 1.

**General setting.** All experiments are conducted using the Jax library (Bradbury et al., 2021). Our default experimental setup, unless specified otherwise, is as follows: We use a batch size of 2000 (halved if memory-constrained) and train for 140k gradient steps to ensure approximate convergence. We use the Adam optimizer (Kingma & Ba, 2015), gradient clipping with a value of 1, and a learning rate scheduler that starts at  $5 \times 10^{-3}$  and uses a cosine decay starting at 60k gradient steps. We utilized 128 discretization steps and the EM and OBABO schemes to integrate the overdamped and underdamped Langevin equations, respectively. The control functions  $u_\theta$  and  $v_\gamma$  with parameters  $\theta$  and  $\gamma$ , respectively, were parameterized as two-layer neural networks with 128 neurons. Unlike Zhang & Chen (2022), we did not include the score of the target density as part of the parameterized control functions  $u_\theta$  and  $v_\gamma$ . Inspired by Nichol & Dhariwal (2021), we applied a cosine-square scheduler for the discretization step size, i.e.,  $\Delta = a \cos^2\left(\frac{\pi n}{2N}\right)$  at step  $n$ , where  $a : [0, \infty) \rightarrow (0, \infty)$  is learned. Furthermore, we use preconditioning as explained in Appendix A.11.2. The diffusion matrix  $\sigma$  and the mass matrix  $M$  were parameterized as diagonal matrices, and we learned the parameters  $\mu$  and  $\Sigma$  for the prior distribution  $p_{\text{prior}} = \mathcal{N}(\mu, \Sigma)$ , with  $\Sigma$  also set as a diagonal matrix. We enforced non-negativity of  $a$  and made  $\sigma$ ,  $M$ , and  $\Sigma$  positive semidefinite via an element-wise softplus transformation.

For the methods that use geometric annealing (see Table 2), that is,  $\nu(x, s) \propto p_{\text{prior}}^{1-\beta(s)}(x) p_{\text{target}}^{\beta(s)}(x)$ , where  $\beta : [0, T] \rightarrow [0, 1]$  is a monotonically increasing function satisfying  $\beta(0) = 0$  and  $\beta(T) = 1$ , we additionally learn the annealing schedule  $\beta$ . Similar to prior works (Doucet et al., 2022b), we parameterize an increasing sequence of  $N$  steps using unconstrained parameters  $b(s)$ . We map these to our annealing schedule with

$$\beta(n\Delta) = \frac{\sum_{n' \leq n} \text{softplus}(b(n'\Delta))}{\sum_{n=1}^N \text{softplus}(b(n\Delta))}, \quad (87)$$



Table 2: Comparison of different diffusion-based sampling methods based on  $\tilde{f}$ ,  $u$ ,  $v$ ,  $\nu$  as defined in the text.

Method	$\tilde{f}$	$u$	$v$
ULA	$\sigma\sigma^\top \nabla_x \log \nu$	0	0
MCD	$\sigma\sigma^\top \nabla_x \log \nu$	0	learned
CMCD	$\frac{1}{2}\sigma\sigma^\top \nabla_x \log \nu$	learned	$\sigma^\top \nabla_x \log \nu - u$
DIS	$-\sigma\sigma^\top \nabla_x \log p_{\text{prior}}$	learned	0
DBS	arbitrary	learned	learned

where softplus ensures non-negativity. Further, we fix  $\beta(0) = 0$  and  $\beta(T) = 1$ , ensuring  $\beta(n'\Delta) \leq \beta(n\Delta)$  when  $n' \leq n$ . We initialized  $b$  such that  $\beta$  is a linear interpolation between 0 and 1.

Moreover, we initialized  $\sigma = M = \Sigma = \text{Id}$  and  $\mu = 0$  for all experiments. In the case of the *Brownian*, *LGCP*, and *ManyWell* tasks, we set  $a = 0.1$ , while for the remaining benchmark problems, we chose  $a = 0.01$  to avoid numerical instabilities encountered with  $a = 0.1$ .

**Evaluation protocol and model selection.** We follow the evaluation protocol of prior work (Blessing et al., 2024) and evaluate all performance criteria 100 times during training, using 2000 samples for each evaluation. To smooth out short-term fluctuations and to obtain more robust results within a single run, we apply a running average with a window of 5 evaluations. We conduct each experiment using four different random seeds and average the best results of each run.

**Benchmark problem details.** All benchmark problems, with the exception of *ManyWell*, were taken from the benchmark suite of Blessing et al. (2024). In their work, the authors used an uninformative prior for the parameters in the Bayesian logistic regression models for the *Credit* and *Cancer* tasks, which frequently caused numerical instabilities. To maintain the challenge of the tasks while ensuring stability, we opted for a Gaussian prior with zero mean and variance of 100. For more detailed descriptions of the tasks, we refer readers to Blessing et al. (2024).

The *ManyWell* target involves a  $d$ -dimensional *double well* potential, corresponding to the (unnormalized) density

$$\rho_{\text{target}}(x) = \exp \left( -\sum_{i=1}^m (x_i^2 - \delta)^2 - \frac{1}{2} \sum_{i=m+1}^d x_i^2 \right),$$

with  $m \in \mathbb{N}$  representing the number of combined double wells (resulting in  $2^m$  modes), and a separation parameter  $\delta \in (0, \infty)$  (see also Wu et al. (2020)). In our experiments, we set  $d = 50$ ,  $m = 5$  and  $\delta = 2$ . Since  $\rho_{\text{target}}$  factorizes across dimensions, we can compute a reference solution for  $\log \mathcal{Z}$  via numerical integration, as described in Midgley et al. (2023).

#### A.11.5 EVALUATION CRITERIA

Here, we provide further information on how our evaluation criteria are computed. To evaluate our metrics, we consider  $n = 2 \times 10^3$  samples  $(x^{(i)})_{i=1}^n$ .

**Effective sample size (ESS).** We compute the (normalized) ESS as

$$\text{ESS} := \frac{(\sum_{i=1}^n w^{(i)})^2}{n \sum_{i=1}^n (w^{(i)})^2}, \quad (88)$$

where  $(w^{(i)})_{i=1}^n$  are the unnormalized importance weights of the samples  $(Z^{(i)})_{i=1}^n$  in path space given as

$$w^{(i)} := \mathcal{Z} \frac{d\tilde{\mathbb{P}}^{v,\tau}}{d\tilde{\mathbb{P}}^{u,\pi}}(Z^{(i)}). \quad (89)$$

Table 3: Results for lower bounds on  $\log \mathcal{Z}$  for various real-world benchmark problems. Higher values indicate better performance. The best results are highlighted in bold. Blue shading indicates that the method uses underdamped Langevin dynamics. Red shading indicate competing state-of-the-art methods.

Method	Credit	Seeds	Cancer	Brownian	Ionosphere	Sonar
DBS	$-585.524 \pm 0.414$ $-585.112 \pm 0.001$	$-73.437 \pm 0.001$ $-73.423 \pm 0.001$	$-83.395 \pm 4.184$ <b><math>-77.881 \pm 0.014</math></b>	$1.081 \pm 0.004$ <b><math>1.136 \pm 0.001</math></b>	$-111.673 \pm 0.002$ <b><math>-111.636 \pm 0.001</math></b>	$-108.595 \pm 0.006$ <b><math>-108.458 \pm 0.004</math></b>
GMMVI	<b><math>-585.098 \pm 0.000</math></b>	<b><math>-73.415 \pm 0.002</math></b>	$-77.988 \pm 0.054$	$1.092 \pm 0.006$	$-111.832 \pm 0.007$	$-108.726 \pm 0.007$
SMC	$-698.403 \pm 4.146$	$-74.699 \pm 0.100$	$-194.059 \pm 0.613$	$-1.874 \pm 0.622$	$-114.751 \pm 0.238$	$-111.355 \pm 1.177$
CRAFT	$-594.795 \pm 0.411$	$-73.793 \pm 0.015$	$-95.737 \pm 1.067$	$0.886 \pm 0.053$	$-112.386 \pm 0.182$	$-115.618 \pm 1.316$
FAB	$-585.102 \pm 0.001$	$-73.418 \pm 0.002$	$-78.287 \pm 0.835$	$1.031 \pm 0.010$	$-111.678 \pm 0.003$	$-108.593 \pm 0.008$

Note that the dependence on  $\mathcal{Z}$  vanishes when replacing  $\frac{d\tilde{\mathbb{P}}^{v,\tau}}{d\tilde{\mathbb{P}}^{u,\pi}}$  with the expression in Proposition 2.3.

**Sinkhorn distance.** We estimate the Sinkhorn distance  $\mathcal{W}_\gamma^2$  (Cuturi, 2013), i.e., an entropy regularized optimal transport distance between a set of samples from the model and target using the Jax ott library (Cuturi et al., 2022).

**Log-normalizing constant.** For the computation of the log-normalizing constant  $\log \mathcal{Z}$  in the general diffusion bridge setting, we note that for any  $u, v \in \mathcal{U}$  it holds that

$$\mathbb{E}_{Z \sim \tilde{\mathbb{P}}^{u,\pi}} \left[ \frac{d\tilde{\mathbb{P}}^{v,\tau}}{d\tilde{\mathbb{P}}^{u,\pi}}(Z) \right] = 1. \quad (90)$$

Together with Proposition 2.3, this shows that

$$\log \mathcal{Z} = \mathbb{E}_{Z \sim \tilde{\mathbb{P}}^{u,\pi}} \left[ \log \frac{d\tilde{\mathbb{P}}^{v,\tau}}{d\tilde{\mathbb{P}}^{u,\pi}}(Z) + \frac{\tilde{\tau}(Z_T)}{\pi(Z_0)} \right], \quad (91)$$

where  $\tilde{\tau}(Z_T) = \rho_{\text{target}}(X_T) \mathcal{N}(0, \text{Id})$  and  $\tilde{\mathbb{P}}^{u,\pi}_{\cdot|0}$  denotes the path space measure of the process  $Z$  with initial condition  $Z_0 = \hat{Z}_0 \in \mathbb{R}^{2d}$  (analogously for  $\tilde{\mathbb{P}}^{v,\tau}_{\cdot|T}$ ), see e.g. Léonard (2013).

If  $u = u^*$  and  $v = v^*$ , the expression in the expectation is almost surely constant, which implies

$$\log \mathcal{Z} = \log \frac{d\tilde{\mathbb{P}}^{v^*,\tau}}{d\tilde{\mathbb{P}}^{u^*,\pi}}(Z) + \frac{\tilde{\tau}(Z_T)}{\pi(Z_0)} \quad (92)$$

If we only have approximations of  $u^*$  and  $v^*$ , Jensen’s inequality shows that the right-hand side in (92) yields a lower bound to  $\log \mathcal{Z}$ . For other methods, the log-normalizing constants can be computed analogously, by replacing  $u, v$  accordingly, see e.g. Berner et al. (2024) for DIS. Our experiments use the lower bound as an estimator for  $\log \mathcal{Z}$  when labeled with “LB”.

#### A.11.6 FURTHER EXPERIMENTS AND COMPARISONS

**Comparison with competing methods.** We extend our evaluation by comparing DBS against several state-of-the-art techniques, including *Gaussian Mixture Model Variational Inference* (GMMVI) (Arenz et al., 2022), *Sequential Monte Carlo* (SMC) (Del Moral et al., 2006), *Continual Repeated Annealed Flow Transport* (CRAFT) (Arbel et al., 2021; Matthews et al., 2022), and *Flow Annealed Importance Sampling Bootstrap* (FAB) (Midgley et al., 2023). The results, presented in Table 3, are primarily drawn from Blessing et al. (2024), where hyperparameters were carefully optimized. Since our experimental setup differs for the *Credit* and *Cancer* tasks (detailed in Section A.11), we adhered to the tuning recommendations provided by Blessing et al. (2024). Across most tasks, we observe that the underdamped variants of DBS and CMCD consistently yield similar or tighter bounds on  $\log \mathcal{Z}$  compared to the competing methods, without the necessity for hyperparameter tuning. Notably, the underdamped version of DBS consistently performs well across *all* tasks and demonstrates robustness, as evidenced by the low variance between different random seeds.




Figure 7: Effective sample size (ESS) for various methods (ULA, MCD, CMCD, DIS) and different integration schemes, averaged across multiple benchmark problems and four seeds. Integration schemes include Euler-Maruyama (EM) for over- (OD) and underdamped (OD) Langevin dynamics and various splitting schemes (OBAB, BAOAB, OBABO).

**Choice of integrator.** To complement the results from Section 4, we conducted an ablation study evaluating the performance and runtime of different integrators for ULA, MCD, CMCD, and DIS. The results are presented in Figures 7 and 8. Consistent with previous findings, the OBABO integrator delivers the best overall performance, with the exception of ULA. We hypothesize that in the case of ULA, simulating the controlled part (O) twice offers little advantage, as both the forward and backward processes are uncontrolled for this method.

**Additional results for DBS.** We present further details regarding the results discussed in Section 4. Specifically, we provide a breakdown of the performance of different integration schemes across all tasks in Figure 9 (ESS values) and Table 5 ( $\log \mathcal{Z}$  (LB) values). Overall, we observe a notable improvement in performance with (symmetric) splitting schemes compared to Euler-Maruyama discretization. However, as the number of discretization steps increases, the performance differences between OBAB, BAOAB, and OBABO become less pronounced. Interestingly, OBABO tends to yield substantial performance gains when the number of discretization steps is small. Furthermore, we examine the impact of parameter learning for  $N = 8$  discretization steps, with the results shown in Figure 10. Surprisingly, while learning either the terminal time  $T$  or the parameters of the prior distribution yields modest improvements, learning both leads to a remarkable  $5\times$  performance increase.

**Choice of drift for DBS.** The drift term  $\tilde{f}$  in the diffusion bridge sampler (DBS) can be freely chosen. To explore the impact of different drift choices, we conducted an ablation study. We tested several options: no drift,  $\nabla_x \log p_{\text{prior}}$ ,  $\nabla_x \log p_{\text{target}}$ , and a geometric annealing path, represented by  $\nabla_x \log \nu$ , where  $\nu(x, s) \propto p_{\text{prior}}^{1-\beta(s)}(x) p_{\text{target}}^{\beta(s)}(x)$ . We also tested using a learned function for  $\beta$ . The results of these experiments are presented in Table 4 and Figure 11. The findings suggest that the most consistent performance is achieved when using the learned geometric annealing path as the drift  $\tilde{f}$ . Interestingly, using the score of the target distribution,  $\nabla_x \log p_{\text{target}}$ , resulted in worse performance compared to no drift for overdamped DBS and only marginal improvements for underdamped DBS.

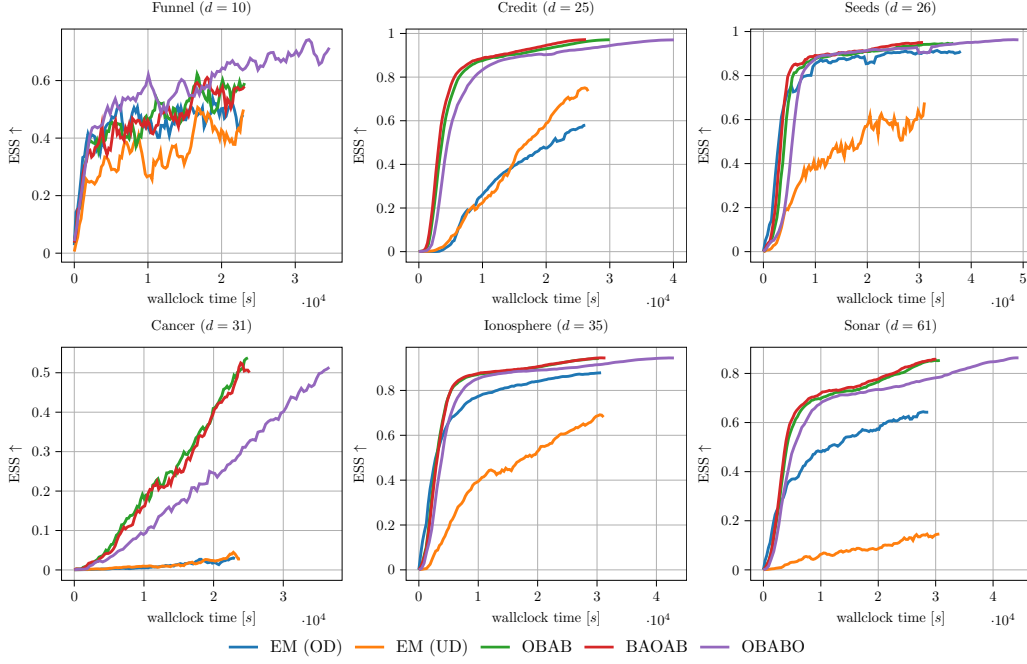


Figure 8: Effective sample size (ESS) over wallclock time of the diffusion bridge sampler (DBS) with 128 diffusion steps for different integration schemes, multiple benchmark problems, and four seeds. Integration schemes include Euler-Maruyama (EM) for over- (OD) and underdamped (UD) Langevin dynamics and various splitting schemes (OBAB, BAOAB, OBABO).

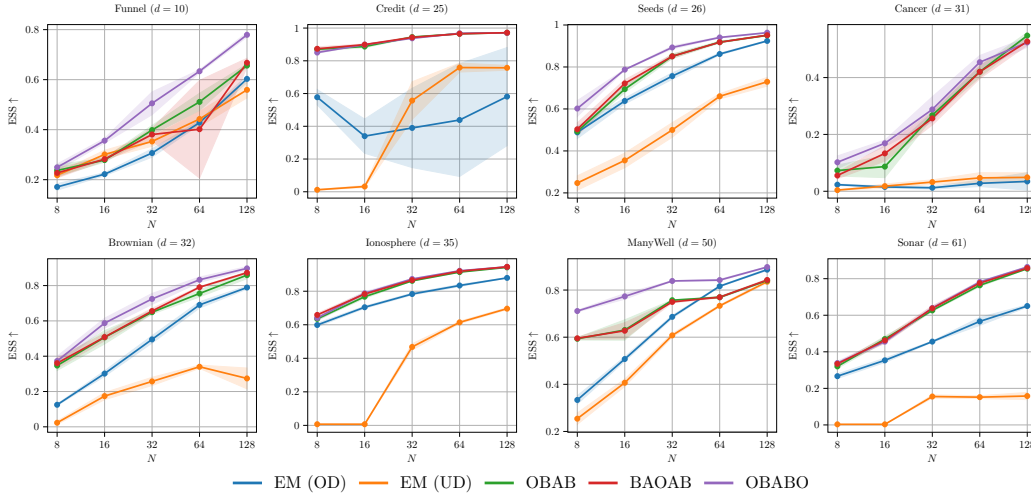


Figure 9: Effective sample size (ESS) of the diffusion bridge sampler (DBS) for different integration schemes, multiple benchmark problems, and four seeds. Integration schemes include Euler-Maruyama (EM) for over- (OD) and underdamped (UD) Langevin dynamics and various splitting schemes (OBAB, BAOAB, OBABO).




Figure 10: Effective sample size (ESS) of the underdamped diffusion bridge sampler (DBS) for various combinations of learned parameters, averaged across multiple benchmark problems and four seeds with  $N = 8$  discretization steps. Parameters include mass matrix  $M$ , diffusion matrix  $\sigma$ , terminal time  $T$ , and extended prior distribution  $\pi$ .

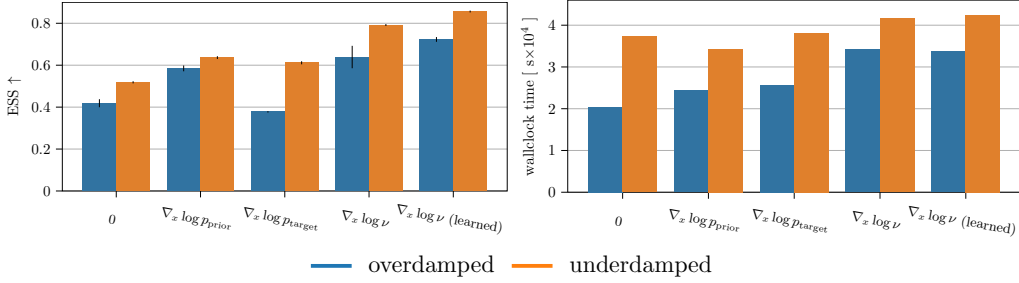


Figure 11: Effective sample size (ESS) and wallclock time for various drifts  $\tilde{f}$  of the underdamped and overdamped diffusion bridge sampler, averaged across multiple benchmark problems and four seeds. Here,  $\nu(x, s) \propto p_{\text{prior}}^{1-\beta(s)}(x)p_{\text{target}}^{\beta(s)}(x)$ , where “learned” indicates that  $\beta$  is learned (end-to-end).

Table 4: Lower bounds on  $\log \mathcal{Z}$  for different drift function  $\tilde{f}$  for DBS on various benchmark problems. Higher values indicate better performance. The best results are highlighted in bold. Here,  $\nu(x, s) \propto p_{\text{prior}}^{1-\beta(s)}(x)p_{\text{target}}^{\beta(s)}(x)$ , where “learned” indicates that  $\beta$  is learned (end-to-end). Blue shading indicates that the method uses underdamped Langevin dynamics.

$\tilde{f}$	Funnel	Credit	Seeds	Cancer	Brownian	Ionosphere	ManyWell	Sonar
0	-0.212 $\pm$ 0.001 -0.155 $\pm$ 0.004	-585.208 $\pm$ 0.008 -585.155 $\pm$ 0.007	-73.501 $\pm$ 0.001 -73.505 $\pm$ 0.009	-81.712 $\pm$ 0.151 -81.307 $\pm$ 0.114	0.466 $\pm$ 0.096 0.449 $\pm$ 0.042	-111.778 $\pm$ 0.005 -111.845 $\pm$ 0.007	38.609 $\pm$ 0.829 42.771 $\pm$ 0.002	-108.936 $\pm$ 0.014 -109.718 $\pm$ 0.013
$\nabla_x \log p_{\text{prior}}$	-0.216 $\pm$ 0.001 -0.145 $\pm$ 0.005	-585.173 $\pm$ 0.005 -585.146 $\pm$ 0.004	-73.483 $\pm$ 0.001 -73.460 $\pm$ 0.000	-81.792 $\pm$ 0.142 -81.080 $\pm$ 0.520	0.787 $\pm$ 0.011 0.972 $\pm$ 0.004	-111.741 $\pm$ 0.002 -111.760 $\pm$ 0.003	42.772 $\pm$ 0.000 <b>42.787<math>\pm</math>0.001</b>	-108.893 $\pm$ 0.050 -109.035 $\pm$ 0.025
$\nabla_x \log p_{\text{target}}$	-0.186 $\pm$ 0.001 <b>-0.096<math>\pm</math>0.004</b>	-685.852 $\pm$ 2.400 -585.271 $\pm$ 0.022	-73.467 $\pm$ 0.000 -73.445 $\pm$ 0.006	-126.194 $\pm$ 15.528 -81.250 $\pm$ 0.219	0.901 $\pm$ 0.004 1.061 $\pm$ 0.004	-111.979 $\pm$ 0.018 -111.826 $\pm$ 0.005	N/A 42.782 $\pm$ 0.000	-109.463 $\pm$ 0.010 -109.264 $\pm$ 0.025
$\nabla_x \log \nu$	-0.183 $\pm$ 0.002 -0.110 $\pm$ 0.000	-4990.364 $\pm$ 4405.152 -585.127 $\pm$ 0.000	-73.442 $\pm$ 0.000 -73.432 $\pm$ 0.000	-83.981 $\pm$ 2.105 -78.086 $\pm$ 0.015	1.055 $\pm$ 0.010 1.106 $\pm$ 0.001	-111.678 $\pm$ 0.000 -111.661 $\pm$ 0.001	42.772 $\pm$ 0.003 42.756 $\pm$ 0.013	-108.616 $\pm$ 0.005 -108.530 $\pm$ 0.002
$\nabla_x \log \nu$ (learned)	-0.175 $\pm$ 0.003 -0.102 $\pm$ 0.003	-585.166 $\pm$ 0.017 <b>-585.112<math>\pm</math>0.000</b>	-73.438 $\pm$ 0.000 <b>-73.422<math>\pm</math>0.001</b>	-78.853 $\pm$ 0.168 <b>-77.866<math>\pm</math>0.007</b>	1.074 $\pm$ 0.005 <b>1.137<math>\pm</math>0.001</b>	-111.673 $\pm$ 0.001 <b>-111.636<math>\pm</math>0.000</b>	42.769 $\pm$ 0.002 42.765 $\pm$ 0.005	-108.593 $\pm$ 0.008 <b>-108.454<math>\pm</math>0.003</b>

Table 5: Results for lower bounds on  $\log \mathcal{Z}$  for various benchmark problems, integration methods, and discretization steps  $N$  for DBS. Higher values indicate better performance. The best results are highlighted in bold. Blue shading indicates that the method uses underdamped Langevin dynamics.

Integrator	Funnel ( $d = 10$ )	Credit ( $d = 25$ )	Seeds ( $d = 26$ )	Cancer ( $d = 31$ )	Brownian ( $d = 32$ )	Ionosphere ( $d = 35$ )	ManyWell ( $d = 50$ )	Sonar ( $d = 61$ )
$N = 8$								
EM (OD)	-0.860 $\pm$ 0.010	-585.400 $\pm$ 0.054	-73.643 $\pm$ 0.003	-80.960 $\pm$ 0.169	0.198 $\pm$ 0.074	-111.858 $\pm$ 0.003	42.162 $\pm$ 0.002	-109.046 $\pm$ 0.017
EM (UD)	-0.725 $\pm$ 0.001	-590.417 $\pm$ 0.859	-73.852 $\pm$ 0.036	-96.286 $\pm$ 2.349	-3.185 $\pm$ 1.159	-123.426 $\pm$ 0.022	42.002 $\pm$ 0.018	-137.601 $\pm$ 0.025
OBAB	-0.670 $\pm$ 0.003	-585.168 $\pm$ 0.004	-73.607 $\pm$ 0.005	<b>-79.167<math>\pm</math>0.176</b>	0.801 $\pm$ 0.004	-111.822 $\pm$ 0.005	42.502 $\pm$ 0.008	-108.937 $\pm$ 0.015
BAOAB	-0.674 $\pm$ 0.008	<b>-585.164<math>\pm</math>0.010</b>	-73.603 $\pm$ 0.011	-79.252 $\pm$ 0.183	0.807 $\pm$ 0.003	<b>-111.811<math>\pm</math>0.007</b>	42.497 $\pm$ 0.004	<b>-108.906<math>\pm</math>0.019</b>
OBABO	<b>-0.557<math>\pm</math>0.004</b>	-585.179 $\pm$ 0.005	<b>-73.560<math>\pm</math>0.008</b>	-78.951 $\pm$ 0.050	<b>0.835<math>\pm</math>0.017</b>	-111.818 $\pm$ 0.009	<b>42.625<math>\pm</math>0.002</b>	-108.933 $\pm$ 0.007
$N = 16$								
EM (OD)	-0.645 $\pm$ 0.002	-585.792 $\pm$ 0.243	-73.561 $\pm$ 0.003	-81.628 $\pm$ 0.345	0.683 $\pm$ 0.010	-111.780 $\pm$ 0.005	42.460 $\pm$ 0.004	-108.902 $\pm$ 0.009
EM (UD)	-0.568 $\pm$ 0.007	-587.429 $\pm$ 0.323	-73.752 $\pm$ 0.018	-82.696 $\pm$ 2.850	0.421 $\pm$ 0.036	-123.426 $\pm$ 0.022	42.354 $\pm$ 0.013	-137.601 $\pm$ 0.025
OBAB	-0.491 $\pm$ 0.004	-585.153 $\pm$ 0.004	-73.520 $\pm$ 0.004	-79.118 $\pm$ 0.723	0.943 $\pm$ 0.004	-111.735 $\pm$ 0.006	42.546 $\pm$ 0.048	-108.766 $\pm$ 0.010
BAOAB	-0.490 $\pm$ 0.003	-585.149 $\pm$ 0.003	-73.516 $\pm$ 0.003	-78.685 $\pm$ 0.261	0.944 $\pm$ 0.004	-111.726 $\pm$ 0.007	42.548 $\pm$ 0.037	<b>-108.754<math>\pm</math>0.009</b>
OBABO	<b>-0.381<math>\pm</math>0.007</b>	<b>-585.149<math>\pm</math>0.002</b>	<b>-73.491<math>\pm</math>0.003</b>	<b>-78.454<math>\pm</math>0.027</b>	<b>0.977<math>\pm</math>0.003</b>	<b>-111.725<math>\pm</math>0.002</b>	<b>42.685<math>\pm</math>0.009</b>	-108.760 $\pm$ 0.010
$N = 32$								
EM (OD)	-0.452 $\pm$ 0.002	-585.941 $\pm$ 0.778	-73.503 $\pm$ 0.001	-84.032 $\pm$ 2.197	0.898 $\pm$ 0.008	-111.730 $\pm$ 0.005	42.626 $\pm$ 0.004	-108.758 $\pm$ 0.005
EM (UD)	-0.425 $\pm$ 0.006	-585.388 $\pm$ 0.120	-73.627 $\pm$ 0.003	-80.207 $\pm$ 0.338	0.595 $\pm$ 0.014	-111.973 $\pm$ 0.009	42.552 $\pm$ 0.009	-109.378 $\pm$ 0.026
OBAB	-0.346 $\pm$ 0.003	<b>-585.126<math>\pm</math>0.002</b>	-73.465 $\pm$ 0.005	-78.224 $\pm$ 0.014	1.024 $\pm$ 0.004	-111.680 $\pm$ 0.002	42.665 $\pm$ 0.005	-108.612 $\pm$ 0.006
BAOAB	-0.347 $\pm$ 0.003	-585.127 $\pm$ 0.002	-73.463 $\pm$ 0.003	-78.206 $\pm$ 0.008	1.035 $\pm$ 0.004	-111.677 $\pm$ 0.003	42.661 $\pm$ 0.006	-108.602 $\pm$ 0.006
OBABO	<b>-0.249<math>\pm</math>0.003</b>	-585.129 $\pm$ 0.004	<b>-73.448<math>\pm</math>0.002</b>	<b>-78.189<math>\pm</math>0.069</b>	<b>1.048<math>\pm</math>0.005</b>	<b>-111.673<math>\pm</math>0.004</b>	<b>42.729<math>\pm</math>0.002</b>	<b>-108.601<math>\pm</math>0.008</b>
$N = 64$								
EM (OD)	-0.295 $\pm$ 0.002	-586.567 $\pm$ 1.871	-73.463 $\pm$ 0.002	-80.890 $\pm$ 1.226	1.027 $\pm$ 0.001	-111.692 $\pm$ 0.004	42.718 $\pm$ 0.004	-108.661 $\pm$ 0.005
EM (UD)	-0.328 $\pm$ 0.009	-585.231 $\pm$ 0.012	-73.554 $\pm$ 0.003	-79.747 $\pm$ 0.382	0.702 $\pm$ 0.017	-111.837 $\pm$ 0.009	42.661 $\pm$ 0.006	-109.410 $\pm$ 0.019
OBAB	-0.228 $\pm$ 0.002	-585.116 $\pm$ 0.001	-73.441 $\pm$ 0.002	-77.968 $\pm$ 0.005	1.082 $\pm$ 0.002	-111.652 $\pm$ 0.002	42.683 $\pm$ 0.003	-108.517 $\pm$ 0.005
BAOAB	-0.606 $\pm$ 0.643	-585.116 $\pm$ 0.001	-73.441 $\pm$ 0.001	-77.979 $\pm$ 0.011	1.091 $\pm$ 0.002	-111.650 $\pm$ 0.002	42.684 $\pm$ 0.004	-108.509 $\pm$ 0.003
OBABO	<b>-0.164<math>\pm</math>0.005</b>	<b>-585.113<math>\pm</math>0.002</b>	<b>-73.431<math>\pm</math>0.003</b>	<b>-77.945<math>\pm</math>0.010</b>	<b>1.104<math>\pm</math>0.003</b>	<b>-111.648<math>\pm</math>0.002</b>	<b>42.730<math>\pm</math>0.004</b>	<b>-108.501<math>\pm</math>0.003</b>
$N = 128$								
EM (OD)	-0.187 $\pm$ 0.003	-585.524 $\pm$ 0.414	-73.437 $\pm$ 0.001	-83.395 $\pm$ 4.184	1.081 $\pm$ 0.004	-111.673 $\pm$ 0.002	42.760 $\pm$ 0.003	-108.595 $\pm$ 0.006
EM (UD)	-0.249 $\pm$ 0.003	-585.235 $\pm$ 0.009	-73.508 $\pm$ 0.005	-79.704 $\pm$ 0.177	0.684 $\pm$ 0.038	-111.786 $\pm$ 0.006	42.731 $\pm$ 0.002	-109.351 $\pm$ 0.075
OBAB	-0.151 $\pm$ 0.003	<b>-585.112<math>\pm</math>0.001</b>	-73.428 $\pm$ 0.001	<b>-77.856<math>\pm</math>0.007</b>	1.121 $\pm$ 0.004	-111.637 $\pm$ 0.001	42.731 $\pm$ 0.002	-108.459 $\pm$ 0.001
BAOAB	-0.159 $\pm$ 0.005	<b>-585.112<math>\pm</math>0.001</b>	-73.428 $\pm$ 0.001	-77.874 $\pm$ 0.010	1.131 $\pm$ 0.002	-111.637 $\pm$ 0.002	42.733 $\pm$ 0.006	<b>-108.457<math>\pm</math>0.004</b>
OBABO	<b>-0.103<math>\pm</math>0.003</b>	<b>-585.112<math>\pm</math>0.001</b>	<b>-73.423<math>\pm</math>0.001</b>	-77.881 $\pm$ 0.014	<b>1.136<math>\pm</math>0.001</b>	<b>-111.636<math>\pm</math>0.001</b>	<b>42.763<math>\pm</math>0.002</b>	-108.458 $\pm$ 0.004