UNDERDAMPED DIFFUSION BRIDGES WITH APPLICATIONS TO SAMPLING

Anonymous authors
Paper under double-blind review

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

$$p_{\mathrm{target}} = \frac{\rho_{\mathrm{target}}}{\mathcal{Z}}, \qquad \mathcal{Z} := \int_{\mathbb{R}^d} \rho_{\mathrm{target}}(x)\,\mathrm{d}x, \tag{1}$$

where $\rho_{\mathrm{target}} \in C(\mathbb{R}^d, \mathbb{R}_{\geq 0})$ 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), 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*

$$\mathrm{d}X_s = \nabla \log p_{\mathrm{target}}(X_s)\,\mathrm{d}s + \sqrt{2}\,\mathrm{d}W_s, \quad X_0 \sim p_{\mathrm{prior}}, \tag{2}$$

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

$$\mathrm{d}X_s = Y_s\,\mathrm{d}s, \quad X_0 \sim p_{\mathrm{prior}}, \tag{3a}$$

$$\mathrm{d}Y_s = (\nabla \log p_{\mathrm{target}}(X_s) - Y_s)\,\mathrm{d}s + \sqrt{2}\,\mathrm{d}W_s, \quad Y_0 \sim \mathcal{N}(0, \mathrm{Id}), \tag{3b}$$

where now the stationary distribution is given by $\tau(x, y) := p_{\mathrm{target}}(x)\mathcal{N}(y; 0, \mathrm{Id})$ (and $\pi(x, y) := p_{\mathrm{prior}}(x)\mathcal{N}(y; 0, \mathrm{Id})$ can be interpreted 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).

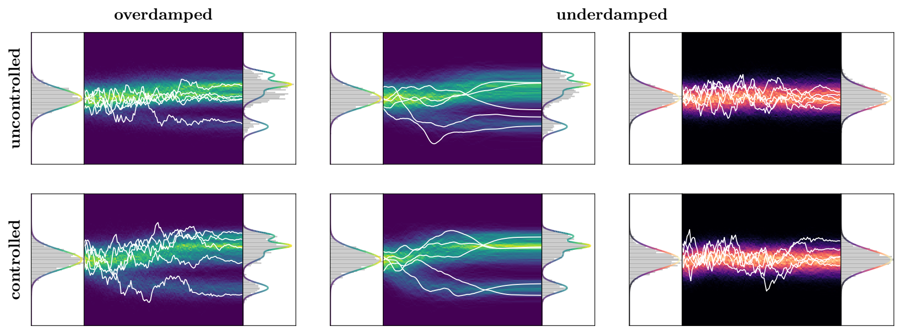

Figure 1: Illustration of uncontrolled (see (2) and (3)) and controlled (see (4) and (15)) diffusion processes in the overdamped and underdamped regime, transporting the Gaussian prior distribution to the target. For the underdamped case, we show both the positional coordinate (left/blue) as well as 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.

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 $\nabla \log p_{\mathrm{target}}$ only depends on X, but not Y). In particular, for smooth and strongly log-concave[1] targets, the number of steps to obtain KL divergence ε can be reduced from $\widetilde{\mathcal{O}}(d/\varepsilon^2)$ to $\widetilde{\mathcal{O}}(\sqrt{d}/\varepsilon)$ (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),

$$\mathrm{d}X_s = (\nabla \log p_{\mathrm{target}}(X_s) + u(X_s, s))\,\mathrm{d}s + \sqrt{2}\,\mathrm{d}W_s, \tag{4}$$

where the task is to learn $u \in C(\mathbb{R}^d \times [0,T], \mathbb{R}^d)$ as to reach $X_T \sim p_{\mathrm{target}}$ (Richter & Berner, 2024; Vargas et al., 2024); see Fig. 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. In particular, we derive the ELBO and variational gap for diffusion bridges where both forward and reverse-time processes are learned.
- **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 our 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.

1.1 Related work

Many approaches to sampling problems build an augmented target, using a sequence of densities bridging the prior and target distributions 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 Dif-*

[1] Or, more general, log-concave outside of a region.

fusion (MCD) optimized the extended target distribution to minimizing the variance of the marginal likelihood estimate (Doucet et al., 2022b). Going one step further, *Controlled Monte Carlo Diffusions* (CMCD) (Vargas et al., 2024) proposed an objective to directly optimize the transition kernels to match the annealed density.

On the other hand, there has recently also been methods prescribing the backward transition kernel, however, having an intractable sequences of densities. For instance, this includes the *Path Integral Sampler* (PIS) (Zhang & Chen, 2021; Vargas et al., 2023b), *Time-Reversed Diffusion Sampler* (DIS) (Berner et al., 2024), *Diffusion generative flow samplers* (DGFS), *Denoising Diffusion Sampler* (DDS) (Vargas et al., 2023a), as well as the *Particle Denoising Diffusion Sampler* (Phillips et al., 2024) combining the latter with SMC components. For the diffusion-based samplers, the optimal forward transition corresponds to the score of the current density, which 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 has been methods learning both kernels separately, i.e., the *(Diffusion) Bridge*[2] *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, 2022), and the works on DDS and CMCD also proposed underdamped versions.

Our proposed framework in principle encompasses all these works as special cases, (Tab. 2), noting, however, that each of the previously existing methods brings some respective additional details. 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 (App. A.9). Our unifying framework allows us to easily share integrators and training techniques for the different methods. First, we remedy tuning for all considered methods by learning hyperparameters end-to-end, also resulting in better performance (Fig. 5). Second, we improve underdamped methods with our novel integrator (Fig. 8 and Fig. 4). Third, we show how to scale DBS to more complex targets by using a suitable parametrization (Tab. 4 & Fig. 10) and divergence-free training objective (Prop. 2.3 vs. Prop. A.6). This makes our underdamped version of DBS a state-of-the-art method across a wide range of tasks (Tab. 1, Fig. 3, & Tab. 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 underdamped dynamics. We refer to Apps. 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., Fig. 2). To this end, let us define the forward and reverse-time SDEs

$$\mathrm{d}Z_s = (f + \eta\, u)\,(Z_s, s)\,\mathrm{d}s + \eta(s)\,\vec{\mathrm{d}}W_s, \qquad Z_0 \sim \pi, \tag{5}$$

$$\mathrm{d}Z_s = (f + \eta\, v)\,(Z_s, s)\,\mathrm{d}s + \eta(s)\,\overleftarrow{\mathrm{d}}W_s, \qquad Z_T \sim \tau, \tag{6}$$

on the state space $\mathbb{R}^D$, where $\vec{\mathrm{d}}W_s$ and $\overleftarrow{\mathrm{d}}W_s$ denote forward and backward Brownian motion increments (see App. 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, 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. 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 with respect to each other. This task can be approached via measures on

[2] We clarify the connection to *Schrödinger bridges* and other *diffusion bridges* in Remark A.1.

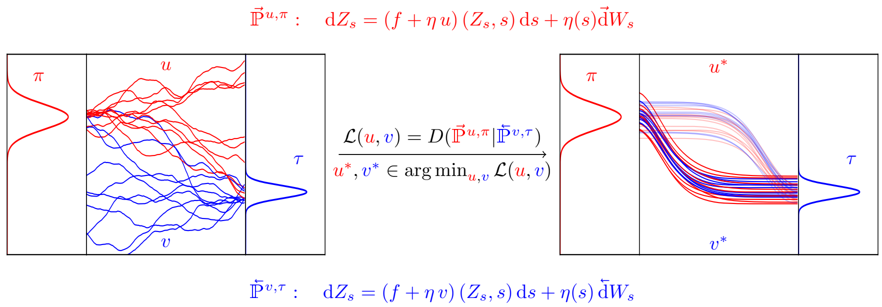

Figure 2: Illustration of our general framework for learning diffusion bridges with degenerate noise. **Left:** We consider forward and reverse-time SDEs (see (5) and (6)), starting at the (extended) prior π and target τ and controlled by u and v, respectively. **Middle:** We learn optimal controls u^* and v^* by minimizing a suitable divergence D between the associated path measures $\vec{\mathbb{P}}^{u,\pi}$ and $\overleftarrow{\mathbb{P}}^{v,\tau}$ 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 right marginals π at time $s = 0$ and τ 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 space of continuous trajectories $C([0,T], \mathbb{R}^D)$, also called *path space* (see App. A.1 for details). To this end, let us denote by $\vec{\mathbb{P}}^{u,\pi}$ the measure of the forward process (5) and by $\overleftarrow{\mathbb{P}}^{v,\tau}$ the measure of the backward process (6). We may consider the following optimization problem; see also Fig. 2.

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

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

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[4] that $\vec{\mathbb{P}}_T^{u^*,\pi} = \tau$ and $\overleftarrow{\mathbb{P}}_0^{v^*,\tau} = \pi$. 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) $\vec{\mathbb{P}}^{u,\pi} = \overleftarrow{\mathbb{P}}^{v,\tau}$.

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

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

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

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

$$\log \frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{v,\tau}}(Z) = \log \frac{\pi(Z_0)}{\tau(Z_T)} - \frac{1}{2}\int_0^T \|(\eta^+ f + u)\|^2(Z_s,s)\,\mathrm{d}s + \frac{1}{2}\int_0^T \|(\eta^+ f + v)\|^2(Z_s,s)\,\mathrm{d}s$$
$$+ \int_0^T (\eta^+ f + u)(Z_s,s)\cdot\eta^+(s)\,\vec{\mathrm{d}}Z_s - \int_0^T (\eta^+ f + v)(Z_s,s)\cdot\eta^+(s)\,\overleftarrow{\mathrm{d}}Z_s.$$

[3] We refer to App. 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 App. A.1).

[4] 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 App. A.1.

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 App. A.5. □

We refer to Prop. A.6 in the appendix for an alternative version of Prop. 2.3, which for non-degenerate noise, has been used to define previous diffusion bridge samplers (Richter & Berner, 2024). However, the latter version relies on a divergence instead of backward stochastic process which renders it prohibitive for high dimensions and does not guarantee an ELBO after discretization; see also Remark 3.2.

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* (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 App. A.9.

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 $\mathcal{L} : \mathcal{U} \times \mathcal{U} \to \mathbb{R}_{\geq 0}$ via $\mathcal{L}(u, v) := D(\vec{\mathbb{P}}^{u,\pi} | \overleftarrow{\mathbb{P}}^{v,\tau})$. A common choice is the *Kullback-Leibler* (KL) divergence, which brings the loss

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

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 (Richter et al., 2020), which we illustrate in App. A.10.3.

2.1 Implications for Generative modeling: The evidence lower bound

Contrary to the sampling setting described above, generative modeling typically assumes that one has access to samples $X \sim p_{\mathrm{target}}$, 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. (2021).

To this end, we may approach Problem 2.1 with the forward[5] KL divergence

$$D_{\mathrm{KL}}(\vec{\mathbb{P}}^{v,\tau} | \overleftarrow{\mathbb{P}}^{u,\pi}) = \mathbb{E}_{Z \sim \vec{\mathbb{P}}^{v,\tau}}\left[\log \frac{\mathrm{d}\vec{\mathbb{P}}^{v,\tau}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{u,\pi}}(Z)\right]. \tag{9}$$

For the sake of notation, we have reversed time, which can be viewed as interchanging τ and π. Since the process corresponding to $\vec{\mathbb{P}}^{v,\tau}$ starts at the target measure τ, we indeed require samples from this measure to compute the divergence in (9). At the same time, looking at Prop. 2.3, we realize that the divergence cannot be computed directly, since τ 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 2.4 (ELBO for generative modeling)**.** *It holds that*

$$\underbrace{\mathbb{E}_{Z_0 \sim \tau}[\log \overleftarrow{\mathbb{P}}_0^{u,\pi}(Z_0)]}_{\textit{evidence / log-likelihood}} = \underbrace{D_{\mathrm{KL}}(\vec{\mathbb{P}}^{v,\tau} | \vec{\mathbb{P}}^{\widetilde{v},\tau})}_{\textit{variational gap}} + \underbrace{\mathbb{E}_{Z_0 \sim \tau}[\log \tau(Z_0)] - D_{\mathrm{KL}}(\vec{\mathbb{P}}^{v,\tau} | \overleftarrow{\mathbb{P}}^{u,\pi})}_{\textit{ELBO}}, \tag{10}$$

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

Proof. This follows from Lemma 2.2 and the chain rule for KL divergences; see App. A.5. □

Crucially, we observe that the ELBO in Lemma 2.4 does not depend on the target τ anymore as the dependency cancels between the two terms (cf. Prop. 2.3). Moreover, the variational gap is zero if and only if $v = \widetilde{v}$ 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 $\overleftarrow{\mathbb{P}}_0^{u,\pi}$

[5] While we optimize the measures in both arguments of the KL divergence, the measure $\overleftarrow{\mathbb{P}}^{u,\pi}$ corresponding to the generative process is in the second component, which is typically referred to as "forward" KL divergence.

equals the target measure τ, 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 η) using samples from the target distribution τ.

We note that for non-degenerate coefficients η, the ELBO from Lemma 2.4 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 $\vec{\mathbb{P}}_T^{0,\tau} \approx \pi$, 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 η, a corresponding *(hybrid) score matching* loss has been suggested and connected to likelihood optimization by Dockhorn et al. (2021, Appendix B.3). In the following proposition, we show that this also follows as a special case from Lemma 2.4.

Proposition 2.5 (Underdamped score matching maximizes the likelihood)**.** *For the ELBO defined in* (10) *(setting $v = 0$) it holds*

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

where the constant does not depend on u.

Proof. Following Huang et al. (2021, Appendix A), the proof combines Prop. 2.3 with Stokes' theorem; see App. A.5. Note that in our notation u learns the *negative* and *scaled* score. □

3 Underdamped Diffusion Bridges

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 Prop. 2.3. Analogously to Vargas et al. (2024, Proposition E.1), we can discretize the integrals to show that

$$\frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{v,\tau}}(Z) \approx \frac{\pi(\widehat{Z}_0)\prod_{n=0}^{N-1}\vec{p}_{n+1}(\widehat{Z}_{n+1}\big|\widehat{Z}_n)}{\tau(\widehat{Z}_N)\prod_{n=0}^{N-1}\overleftarrow{p}_n(\widehat{Z}_n\big|\widehat{Z}_{n+1})}, \tag{11}$$

where the expressions for the forward and backward transition kernels $\vec{p}_n$ and $\overleftarrow{p}_n$ depend on the choice of the integrator for Z. Since we have degenerate diffusion matrices, the backward kernel $\overleftarrow{p}$ 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 Fig. 4.

We therefore consider alternative integration methods, specifically splitting schemes (Bou-Rabee & Owhadi, 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[6]

$$\partial_t p = \mathcal{L}p \quad \text{with} \quad \mathcal{L}p = -\nabla\cdot\big((f+\eta u)p\big) + \tfrac{1}{2}\operatorname{Tr}(\eta\eta^\top\nabla^2 p) \tag{12}$$

governing the evolution of the density $p(\cdot, t) = \vec{\mathbb{P}}_t^{u,\pi}$ 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 η, 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

$$f(x, y, s) = (y, \widetilde{f}(x, s) - \tfrac{1}{2}\sigma\sigma^\top(s)y)^\top, \qquad \eta = (\mathbf{0}, \sigma)^\top \tag{13}$$

in (5) and (6), where $\mathbf{0} \in \mathbb{R}^{d\times d}$. Following Monmarché (2021); Geffner & Domke (2022) we split the generator as $\mathcal{L} = \mathcal{L}_\mathrm{A} + \mathcal{L}_\mathrm{B} + \mathcal{L}_\mathrm{O}$ (sometimes referred to as free transport, acceleration, and

[6] We denote by Tr the trace and by ∇ the del operator w.r.t. spatial variable z; see App. A.1.

Algorithm 1 Training of an underdamped diffusion sampler

Require: ▷ See App. A.10 for details

- **model:** neural networks u_θ, v_γ 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 $p_{\text{prior}} = \mathcal{N}(\zeta_\mu, \operatorname{diag}(\operatorname{softplus}(\eta_\Sigma)))$, diffusion and mass matrices $\sigma = \operatorname{diag}(\operatorname{softplus}(\eta_\sigma))$, $M = \operatorname{diag}(\operatorname{softplus}(\eta_M))$, and terminal time $T = N \operatorname{softplus}(\eta_\delta)$ 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 $k \leftarrow 0, \ldots, K-1$ **do**
  for $i \leftarrow 1, \ldots, m$ **do** ▷ Approximate cost (batched in practice)
    $\widehat{Z}_0 \sim \pi$ ▷ See (16)
    $\text{rnd}_i \leftarrow \log \pi(\widehat{Z}_0)$
    for $n \leftarrow 0, \ldots, N-1$ **do**
      $\widehat{Z}_{n+1} \leftarrow \texttt{integrate}(\widehat{Z}_n, \Theta^{(k)})$ ▷ See App. A.8
      $\text{rnd}_i \leftarrow \text{rnd}_i + \log \vec{p}_{n+1}(\widehat{Z}_{n+1} | \widehat{Z}_n) - \log \overleftarrow{p}_n(\widehat{Z}_n | \widehat{Z}_{n+1})$ ▷ See (11)
    $\text{rnd}_i \leftarrow \text{rnd}_i - \log \tilde{\tau}(\widehat{Z}_N)$
  $\widehat{\mathcal{L}} \leftarrow \frac{1}{m} \sum_{i=1}^m \text{rnd}_i$ ▷ Compute loss
  $\Theta^{(k+1)} \leftarrow \operatorname{step}\left(\Theta^{(k)}, \nabla_\Theta \widehat{\mathcal{L}}\right)$ ▷ Gradient descent
return optimized parameters $\Theta^{(K)}$

damping) with

$$\mathcal{L}_{\mathrm{A}} p = -y \cdot \nabla_x p, \quad \mathcal{L}_{\mathrm{B}} p = -\widetilde{f} \cdot \nabla_y p, \quad \mathcal{L}_{\mathrm{O}} p = -\nabla_y \cdot (gp) + \tfrac{1}{2} \operatorname{Tr}(\sigma\sigma^\top \nabla_y^2 p), \tag{14}$$

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

$$\begin{bmatrix} \mathrm{d}X_s \\ \mathrm{d}Y_s \end{bmatrix} = \underbrace{\begin{bmatrix} Y_s \\ 0 \end{bmatrix} \mathrm{d}s}_{\mathrm{A}} + \underbrace{\begin{bmatrix} 0 \\ \widetilde{f}(X_s, s) \end{bmatrix} \mathrm{d}s}_{\mathrm{B}} + \underbrace{\begin{bmatrix} 0 \\ \left(-\frac{1}{2}\sigma\sigma^\top(s) Y_s + \sigma u(Z_s, s)\right) \mathrm{d}s + \sigma(s) \mathrm{d}W_s \end{bmatrix}}_{\mathrm{O}}, \tag{15}$$

where we use a standard normal for the last d components of the initial and terminal distributions following Geffner & Domke (2022), i.e.,

$$\pi(x, y) = p_{\text{prior}}(x) \mathcal{N}(y; 0, \mathrm{Id}) \quad \text{and} \quad \tau(x, y) = p_{\text{target}}(x) \mathcal{N}(y; 0, \mathrm{Id}). \tag{16}$$

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

$$e^{(\mathcal{L}_{\mathrm{A}} + \mathcal{L}_{\mathrm{B}} + \mathcal{L}_{\mathrm{O}})t} \approx \left[e^{\mathcal{L}_{\mathrm{A}}\delta} e^{\mathcal{L}_{\mathrm{B}}\delta} e^{\mathcal{L}_{\mathrm{O}}\delta}\right]^N + \mathcal{O}(N\delta^3), \tag{17}$$

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

$$e^{(\mathcal{L}_{\mathrm{A}} + \mathcal{L}_{\mathrm{B}} + \mathcal{L}_{\mathrm{O}})t} \approx \left[e^{\mathcal{L}_{\mathrm{O}}\frac{\delta}{2}} e^{\mathcal{L}_{\mathrm{B}}\frac{\delta}{2}} e^{\mathcal{L}_{\mathrm{A}}\delta} e^{\mathcal{L}_{\mathrm{B}}\frac{\delta}{2}} e^{\mathcal{L}_{\mathrm{O}}\frac{\delta}{2}}\right]^N + \mathcal{O}(N\delta^2), \tag{18}$$

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, 2022; 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 $\overleftarrow{p}$ corresponding to these splitting schemes can be found in App. A.8. We refer to Algorithm 1 for an overview of our method and to App. A.10 for further details. A few remarks are in order (see also App. A.3).

Remark 3.1 (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 (15) 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 App. A.7. Furthermore, we conducted experiments where we learned the mass matrix, as discussed in Section 4.

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 App. A.10.2 for details on the metrics. The best results are highlighted in bold. Arrows (↑, ↓) 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 ↑	$\mathcal{W}_2^\gamma \downarrow$	$\Delta \log \mathcal{Z} \downarrow$	ESS ↑	$\log \mathcal{Z}$ (LB) ↑	ESS × 10 ↑
ULA	0.310 ± 0.020	0.140 ± 0.003	169.859 ± 0.195	0.016 ± 0.003	0.179 ± 0.008	482.024 ± 0.009	0.029 ± 0.003
	0.130 ± 0.021	0.151 ± 0.016	159.212 ± 0.093	0.009 ± 0.002	0.418 ± 0.002	484.087 ± 0.063	0.030 ± 0.004
MCD	0.173 ± 0.046	0.206 ± 0.026	164.967 ± 0.334	0.005 ± 0.002	0.737 ± 0.002	483.137 ± 0.368	0.031 ± 0.004
	0.088 ± 0.008	0.375 ± 0.016	144.753 ± 0.153	0.005 ± 0.000	0.866 ± 0.012	484.933 ± 0.298	0.032 ± 0.006
CMCD	0.023 ± 0.003	0.567 ± 0.023	104.644 ± 0.710	$\mathbf{0.004 \pm 0.002}$	0.859 ± 0.001	483.875 ± 0.275	0.032 ± 0.004
	0.268 ± 0.198	0.369 ± 0.186	148.990 ± 19.81	0.008 ± 0.003	0.585 ± 0.034	483.535 ± 0.232	0.028 ± 0.004
DIS	0.047 ± 0.003	0.498 ± 0.021	107.458 ± 0.826	0.006 ± 0.002	0.798 ± 0.002	405.686 ± 4.019	0.015 ± 0.003
	0.048 ± 0.009	0.550 ± 0.039	114.580 ± 0.457	0.005 ± 0.000	0.856 ± 0.002	DIVERGED	DIVERGED
DBS	0.021 ± 0.003	0.603 ± 0.014	102.653 ± 0.586	0.005 ± 0.001	0.887 ± 0.004	486.376 ± 1.020	0.032 ± 0.002
	$\mathbf{0.010 \pm 0.001}$	$\mathbf{0.779 \pm 0.009}$	$\mathbf{101.418 \pm 0.425}$	0.005 ± 0.000	$\mathbf{0.898 \pm 0.002}$	$\mathbf{497.545 \pm 0.183}$	$\mathbf{0.174 \pm 0.017}$

Remark 3.2 (Discrete Radon-Nikodym derivative)**.** We note that our discretization of the Radon-Nikodym derivative in (11) 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 Prop. A.6 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 App. A.10.3. Finally, we note that the discretized Radon-Nikodym derivative does not depend on $\widetilde{f}$ for the integrators considered in App. A.8. We thus choose $\widetilde{f}$ to have a good initialization for Z, e.g., as Langevin dynamics; see App. A.10.

Remark 3.3 (Properties of the score)**.** Since the target density p_{target} in (16) only appears in the coordinates where η vanishes, Nelson's identity in Lemma 2.2 shows that

$$u^*(x, y, T) - v^*(x, y, T) = \sigma^\top(T) \nabla_y \log \mathcal{N}(y; 0, M), \tag{19}$$

i.e., the optimal controls u^* and v^* do not depend on the score of the target distribution, $\nabla_x \log p_{\mathrm{target}}$, 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_{target} is essentially supported on a lower dimensional manifold (Dockhorn et al., 2021; Chen et al., 2022).

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, 2022), *Controlled Monte Carlo Diffusions* (CMCD) (Vargas et al., 2024), *Time-Reversed Diffusion Sampler* (DIS)[7] (Berner et al., 2024), and *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 evaluation methodology adheres to the protocol suggested in Blessing et al. (2024). For a comprehensive overview of the experimental setup and additional details, we refer to App. A.10. Moreover, we provide further numerical results in App. A.10.3, including the comparison to competing state-of-the-art methods. The code is publicly available[8].

4.1 Benchmark Problems

We evaluate the different methods on various real-world and synthetic benchmark examples.

[7] 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 (Berner et al., 2024).

[8] https://anonymous.4open.science/r/UnderdampedDiffusionBridges

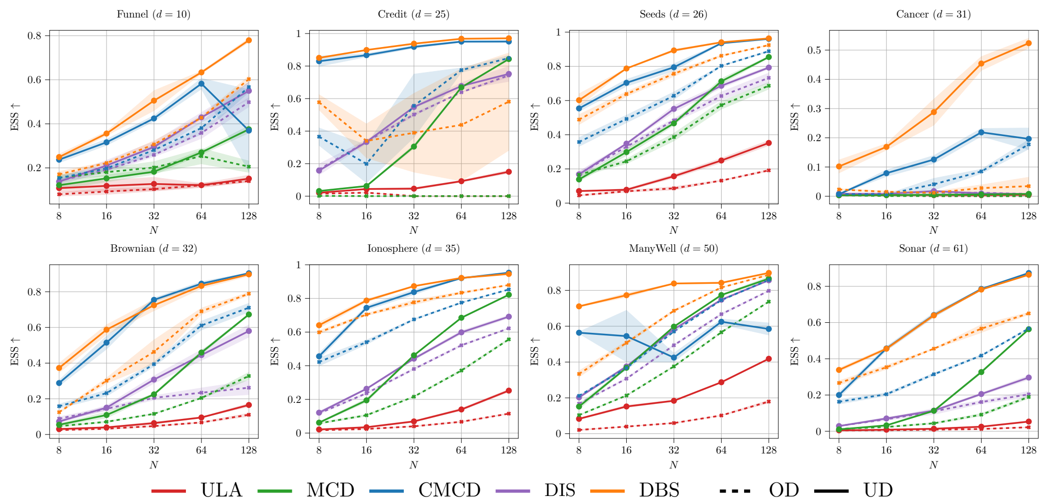

Figure 3: Effective sample size (ESS) for real-world benchmark problems of various dimensions d, averaged across four seeds. Here, N refers to the number of discretization steps. Solid/dashed lines indicate the usage of the overdamped (OD) and underdamped (UD) Langevin, respectively.

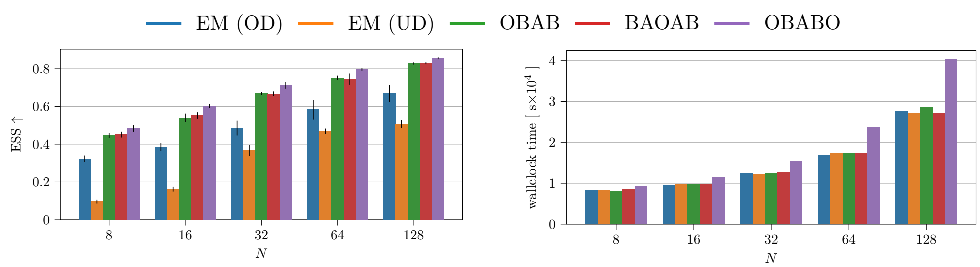

Figure 4: Effective sample size (ESS) and wallclock time of the diffusion bridge sampler (DBS) for different integration schemes, averaged across multiple benchmark problems and four seeds. Integration schemes include Euler-Maruyama (EM) for over (OD) - and underdamped (OD) Langevin and various splitting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

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, *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, *LGCP* ($d = 1600$), a high-dimensional Log Gaussian Cox process (Møller et al., 1998).

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 consider the *ManyWell* ($d = 50$) target, a highly multi-modal distribution with $2^5 = 32$ modes.

4.2 Results

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

Numerical integration schemes. Here, 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 App. A.10.3. To provide a concise overview, we present the average effective sample size (ESS) and wallclock time across all tasks, excluding LGCP, in Fig. 4. Detailed results

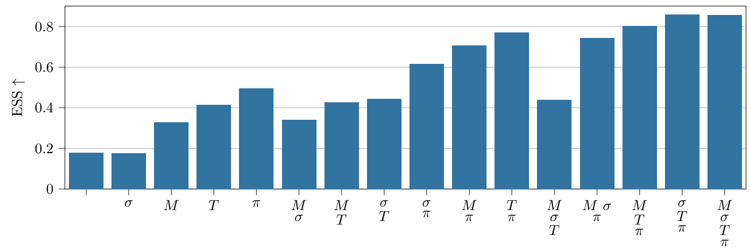

Figure 5: Effective sample size (ESS) of the underdamped diffusion bridge sampler (DBS) for various combinations of learned parameters, averaged across multiple benchmark problems and four seeds using $N = 64$ discretization steps. Haperparameters include mass matrix M, diffusion matrix σ, terminal time T, and extended prior distribution π. See Fig. 9 for the results with $N = 8$ discretization steps.

for individual benchmarks can be found in App. A.10.3. While is 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, diffusion matrix σ, terminal time T, and prior distribution π. Fig. 5 and Fig. 9 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 enhances the method's user-friendliness 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. We provide careful ablation studies showing that our improvements are due to the combination of underdamped dynamics, our novel numerical integrators, as well as end-to-end learned hyperparameters and forward and backward transitions. Our results also offer motivation to extend the method by Chen et al. (2021) and benchmark underdamped diffusion bridges for generative modeling using the ELBO derived in Lemma 2.4. Different from diffusion models, diffusion bridges require SDE simulations during training, but can also be applied to more general prior distributions.

Finally, 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. (2021), we find significant and consistent improvements over overdamped versions, in particular also for high-dimensional targets with only a few 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 $\nabla \log p_{\mathrm{target}}$ in Langevin dynamics) the control u depends not only on the smooth X but also on Y (Chen et al., 2022). Specifically, they show that a small KL divergence between the path measures generally requires the step size δ 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 the results still can be reconciled with our empirical findings due to the following reasons: (1) our samplers are initialized as Langevin dynamics (see App. A.10) such that theoretical benefits of the underdamped case hold at least initially (2) the learning problem becomes numerically better behaved (see (19)), leading to better approximation of the optimal parameters, (3) learning both u and v as well as the prior π, diffusion coefficient σ, and terminal time T (see Fig. 5) can reduce the discretization error.

References

Tara Akhound-Sadegh, Jarrid Rector-Brooks, Avishek Joey Bose, Sarthak Mittal, Pablo Lemos, Cheng-Hao Liu, Marcin Sendera, Siamak Ravanbakhsh, Gauthier Gidel, Yoshua Bengio, et al. Iterated denoising energy matching for sampling from boltzmann densities. *arXiv preprint arXiv:2402.06121*, 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. *arXiv preprint arXiv:2406.07423*, 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.

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-Horng 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. *arXiv preprint arXiv:2112.07068*, 2021.

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. 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. *arXiv preprint arXiv:2208.07743*, 2022.

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. ISBN 9781439840955. URL https://books.google.com/books?id=ZXL6AQAAQBAJ.

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. *arXiv preprint arXiv:1412.6980*, 2014.

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 *The Eleventh 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.

Benedict Leimkuhler and Sebastian Reich. *Simulating hamiltonian dynamics*. Number 14. Cambridge university press, 2004.

Christian Léonard. A survey of the schr\" odinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 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. *arXiv preprint arXiv:2310.02233*, 2023.

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? 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. *arXiv preprint arXiv:2208.01893*, 2022.

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.

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 Deligiannidis, and Arnaud Doucet. Particle denoising diffusion sampler. *arXiv preprint arXiv:2402.06320*, 2024.

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.

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.

Ziqiang Shi and Rujie Liu. Generative modelling with high-order langevin dynamics. *arXiv preprint arXiv:2404.12814*, 2024.

Tommi Sottinen and Simo Särkkä. Application of girsanov theorem to particle filtering of discretely observed continuous-time non-linear systems. 2008.

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 Azizzadenesheli, 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, Will Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. *arXiv preprint arXiv:2302.13834*, 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 *The Twelfth 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.

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. *arXiv preprint arXiv:2111.15141*, 2021.

A Appendix

Contents

A.1 Notation

We denote by $\operatorname{Tr}(\Sigma)$ and Σ^+ the trace and the (Moore-Penrose) pseudoinverse of a real-valued matrix Σ, by $\|\mu\|$ the Euclidean norm of a vector μ, and by $\mu_1 \cdot \mu_2$ the Euclidean inner product between vectors μ_1 and μ_2.

For a function $p\colon \mathbb{R}^D \times [0,T] \to \mathbb{R}$, 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 $\partial_t p$ it 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. Moreover, we denote by

$$\nabla p = \begin{bmatrix} \nabla_x p \\ \nabla_y p \end{bmatrix} \tag{20}$$

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 $f = (f_i)_{i=1}^D \colon \mathbb{R}^D \times [0,T] \to \mathbb{R}^D$ 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 $\operatorname{Unif}([0,T])$ the uniform distribution on $[0,T]$. For an $\mathbb{R}^d$-valued random variable X with law $\mathbb{P}$ and a function $f\colon \mathbb{R}^d \to \mathbb{R}$, we denote by

$$\mathbb{E}_{X\sim\mathbb{P}}[f(X)] = \int f \,\mathrm{d}\mathbb{P} \tag{21}$$

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

For suitable 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

$$\int_{\underline{t}}^{\overline{t}} X_s \cdot \vec{\mathrm{d}} Y_s = \lim_{n\to\infty} \sum_{i=0}^{k_n} X_{t_i^n} \cdot (Y_{t_{i+1}^n} - Y_{t_i^n}), \tag{22}$$

$$\int_{\underline{t}}^{\overline{t}} X_s \cdot \overleftarrow{\mathrm{d}} Y_s = \lim_{n\to\infty} \sum_{i=0}^{k_n} X_{t_{i+1}^n} \cdot (Y_{t_{i+1}^n} - Y_{t_i^n}), \tag{23}$$

where $\underline{t} < t_0^n < \cdots < t_{k_n}^n = \overline{t}$ is an increasing sequence of subdivisions of $[\underline{t}, \overline{t}]$ with mesh tending to zero; see Vargas et al. (2024) for details. The relation between forward and backward integrals is given in Lemma A.5.

We denote by $\mathcal{P}$ the set of probability measures on $C([0,T],\mathbb{R}^D)$, equipped with the Borel σ-field associated with the topology of uniform convergence on compact sets. For suitable vector fields u, v and distributions π, τ, 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

$$Z_t = Z_0 + \int_0^t (f + \eta\, u)\,(Z_s, s)\,\mathrm{d}s + \int_0^t \eta(s)\,\vec{\mathrm{d}} W_s, \qquad Z_0 \sim \pi, \tag{24}$$

$$Z_t = Z_T - \int_t^T (f + \eta\, v)\,(Z_s, s)\,\mathrm{d}s - \int_t^T \eta(s)\,\overleftarrow{\mathrm{d}} W_s, \qquad Z_T \sim \tau, \tag{25}$$

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 η of the form $\eta = (\mathbf{0}, \sigma)^\top$. 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

$$\|g(z_1,t) - g(z_2,t)\| \le C\|z_1 - z_2\|. \tag{26}$$

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, σ, 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 (26) guarantee strong solutions with pathwise uniqueness (see, e.g., Le Gall (2016, Section 8.2)) and are sufficient for Girsanov's theorem in Thm. A.3 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 (22) and (23) 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[9] such that we have strong solutions to the corresponding Fokker-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

[9] 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 Proposition). 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.

et al. (1989). While we use the above assumptions to simplify the presentation, we note they can be significantly relaxed.

A.3 Further remarks

Remark A.1 (Stochastic bridges and bridge sampling)**.** By *stochastic bridge* or *diffusion bridge* (also referred to as *general bridge* by Richter & Berner (2024)), we refer to a SDE that satisfies the marginals p_{prior} and p_{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_{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 of obtaining uniqueness consists of minimizing the distance[10] 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 DIS (Zhang & Chen, 2021) 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_{prior} and p_{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.2 (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 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.

A.4 Auxiliary results

Theorem A.3 (Girsanov theorem)**.** *For $\vec{\mathbb{P}}^{u,\pi}$-almost every $Z \in C([0,T],\mathbb{R}^D)$ it holds that*

$$\log \frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}}{\mathrm{d}\vec{\mathbb{P}}^{w,\pi}}(Z) = -\int_0^T \left(\frac{1}{2}\|u-w\|^2 + (\eta^+ f + w)\cdot(u-w)\right)(Z_s,s)\,\mathrm{d}s + S \tag{27}$$

$$= \frac{1}{2}\int_0^T \left(\|\eta^+ f + w\|^2 - \|\eta^+ f + u\|^2\right)(Z_s,s)\,\mathrm{d}s + S, \tag{28}$$

where

$$S = \int_0^T (u-w)(Z_s,s)\cdot\eta^+(s)\,\vec{\mathrm{d}}Z_s. \tag{29}$$

[10] 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. (2023).

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

$$\log \frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}}{\mathrm{d}\vec{\mathbb{P}}^{w,\pi}}(Z) = -\frac{1}{2}\int_0^T \|u-w\|^2(Z_s,s)\,\mathrm{d}s + \int_0^T (u-w)(Z_s,s)\cdot \vec{\mathrm{d}}B_s. \tag{30}$$

Proof. See Sottinen & Särkkä (2008); Chen et al. (2022); Üstünel & Zakai (2013). □

Theorem A.4 (Reverse-time Girsanov theorem). *For $\vec{\mathbb{P}}^{u,\pi}$-almost every $Z \in C([0,T],\mathbb{R}^D)$ holds that*

$$\log \frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{u,\pi}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{w,\pi}}(Z) = \log \frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}}{\mathrm{d}\vec{\mathbb{P}}^{w,\pi}}(Z) - \int_0^T (u-w)(Z_s,s)\cdot \eta^+(s)\,\vec{\mathrm{d}}Z_s \tag{31}$$

$$+ \int_0^T (u-w)(Z_s,s)\cdot \eta^+(s)\,\overleftarrow{\mathrm{d}}Z_s. \tag{32}$$

Proof. Using Thm. A.3 and the definitions in (22) and (23), we observe that $\frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{u,\pi}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{w,\pi}}(Z)$ equals the Radon-Nikodym derivative between the path spaces measures corresponding to forward SDEs as in (5) with initial conditions π and all functions f, u, w, and η 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). □

Lemma A.5 (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*

$$\int_{\underline{t}}^{\bar{t}} g(Z_s,s)\cdot \overleftarrow{\mathrm{d}}Z_s = \int_{\underline{t}}^{\bar{t}} g(Z_s,s)\cdot \vec{\mathrm{d}}Z_s + \int_{\underline{t}}^{\bar{t}} \nabla\cdot(\eta\eta^\top g)(Z_s,s)\,\mathrm{d}s. \tag{33}$$

Proof. Similar to the conversion formula in Vargas et al. (2024, Remark 3), the result follows from combining (22) and (23). First, we rewrite the problem by observing that

$$\int_{\underline{t}}^{\bar{t}} g(Z_s,s)\cdot \overleftarrow{\mathrm{d}}Z_s = \int_{\underline{t}}^{\bar{t}} g(Z_s,s)\cdot \vec{\mathrm{d}}Z_s + \int_{\underline{t}}^{\bar{t}} \widetilde{g}(Z_s,s)\cdot \overleftarrow{\mathrm{d}}W_s - \int_{\underline{t}}^{\bar{t}} \widetilde{g}(Z_s,s)\cdot \vec{\mathrm{d}}W_s,$$

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

$$\begin{aligned}\int_{\underline{t}}^{\bar{t}} \widetilde{g}(Z_s,s)\cdot \overleftarrow{\mathrm{d}}W_s &= \lim_{n\to\infty}\sum_{i=0}^{k_n}\big(\widetilde{g}(Z_{t_{i+1}^n},t_{i+1}^n) + \widetilde{g}(Z_{t_i^n},t_i^n)\big)\cdot\big(W_{t_{i+1}^n} - W_{t_i^n}\big) - \int_{\underline{t}}^{\bar{t}} \widetilde{g}(Z_s,s)\cdot \vec{\mathrm{d}}W_s \\ &= 2\int_{\underline{t}}^{\bar{t}} \widetilde{g}(Z_s,s)\circ \mathrm{d}W_s - \int_{\underline{t}}^{\bar{t}} \widetilde{g}(Z_s,s)\cdot \vec{\mathrm{d}}W_s,\end{aligned}$$

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

$$\int_{\underline{t}}^{\bar{t}} \widetilde{g}(Z_s,s)\circ \mathrm{d}W_s = \int_{\underline{t}}^{\bar{t}} \widetilde{g}(Z_s,s)\cdot \vec{\mathrm{d}}W_s + \frac{1}{2}\int_{\underline{t}}^{\bar{t}} \nabla\cdot(\eta\widetilde{g})(Z_s,s)\,\mathrm{d}s, \tag{34}$$

see, e.g., Kloeden & Platen (1992, Section 4.9). □

A.5 PROOFS

Proof of Prop. 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 $\frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{w,\tau}}{\mathrm{d}\vec{\mathbb{P}}^{w,\pi}}(Z) = \frac{\tau(Z_T)}{\pi(Z_0)}$ for $w = -\eta^+ f$. Thus, it holds that

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

The result now follows by applying the Girsanov theorem; see Thm. A.3 and Thm. A.4. □

Proof of Lemma 2.4. Using Lemma 2.2 and the chain rule for the KL divergence, we observe that

$$D_{\mathrm{KL}}(\vec{\mathbb{P}}^{v,\tau}|\overleftarrow{\mathbb{P}}^{u,\pi}) = D_{\mathrm{KL}}(\vec{\mathbb{P}}^{v,\tau}|\vec{\mathbb{P}}^{\widetilde{v},\widetilde{\tau}}) = D_{\mathrm{KL}}(\vec{\mathbb{P}}^{v,\tau}|\vec{\mathbb{P}}^{\widetilde{v},\tau}) + D_{\mathrm{KL}}(\tau|\overleftarrow{\mathbb{P}}_0^{u,\pi}), \tag{36}$$

where $\widetilde{\tau} = \overleftarrow{\mathbb{P}}_0^{u,\pi}$. We note that the Girsanov theorem (see Thm. A.3) implies that the variational gap can equivalently be written as

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

see also Vargas et al. (2024, Appendix C). □

Proof of Prop. 2.5. 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 η. Using Prop. A.6 and a Monte Carlo approximation, we first observe that, for the case $v=0$, the ELBO can be represented as

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

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

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

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

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

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

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

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

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

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

where we used identities for divergences and Stokes' theorem. □

A.6 Additional statements on diffusion models

Proposition A.6 (Radon-Nikodym derivative)**.** *For a process $Z\sim\vec{\mathbb{P}}^{w,\pi}$ as defined in (5) it holds*

$$\log\frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{v,\tau}}(Z) = \log\frac{\pi(Z_0)}{\tau(Z_T)} + \int_0^T\left((u-v)\cdot\left(w - \frac{u+v}{2}\right) - \nabla\cdot(\eta\eta^+ f + \eta v)\right)(Z_s^w,s)\,\mathrm{d}s$$
$$+ \int_0^T (u-v)(Z_s,s)\cdot\vec{\mathrm{d}}W_s,$$

where we note that $\eta\eta^+ = \begin{pmatrix}\mathbf{0}_d & \mathbf{0}_d\\ \mathbf{0}_d & \mathrm{Id}_{d\times d}\end{pmatrix}$.

Proof. This follows from combining Prop. 2.3 with Lemma A.5. Note that for $Z \sim \vec{\mathbb{P}}^{w,\pi}$ it holds that

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

□

A.7 Including a mass matrix

In Section 3, we omitted the mass matrix M for simplicity. Here, 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

$$
f(x,y,s) = (y, \widetilde{f}(x,y,s) - \tfrac{1}{2}\sigma\sigma^\top(s)y)^\top, \qquad \eta = (\mathbf{0}_d, \sigma M^{1/2})^\top \tag{43}
$$

in (5) and (6), where $\mathbf{0}_d \in \mathbb{R}^{d\times d}$, 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.,

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

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

$$
\mathrm{d}X_s = M^{-1}Y_s\,\mathrm{d}s, \qquad X_0 \sim p_{\mathrm{prior}}, \tag{45a}
$$

$$
\mathrm{d}Y_s = \left(\widetilde{f}(Z_s,s) - \tfrac{1}{2}\sigma\sigma^\top(s)Y_s + \sigma M^{1/2}u(Z_s,s)\right)\mathrm{d}s + \sigma(s)M^{1/2}\,\vec{\mathrm{d}}W_s, \quad Y_0 \sim \mathcal{N}(0,M), \tag{45b}
$$

and

$$
\mathrm{d}X_s = M^{-1}Y_s\,\mathrm{d}s, \qquad X_T \sim p_{\mathrm{target}}, \tag{46a}
$$

$$
\mathrm{d}Y_s = \left(\widetilde{f}(Z_s,s) - \tfrac{1}{2}\sigma\sigma^\top(s)Y_s - \sigma M^{1/2}v(Z_s,s)\right)\mathrm{d}s + \sigma(s)M^{1/2}\,\overleftarrow{\mathrm{d}}W_s, \quad Y_T \sim \mathcal{N}(0,M). \tag{46b}
$$

In a similar spirit to the diffusion matrix σ, one can also learn the mass matrix. However, our experiments (Section 4) showed little improvements.

A.8 Numerical discretization schemes

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

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

For convenience, we recall the following split for the generative SDE that is used throughout this section, i.e.,

$$\begin{bmatrix} \mathrm{d}X_s \\ \mathrm{d}Y_s \end{bmatrix} = \underbrace{\begin{bmatrix} Y_s \\ 0 \end{bmatrix} \mathrm{d}s}_{\vec{\mathrm{A}}} + \underbrace{\begin{bmatrix} 0 \\ \widetilde{f}(X_s, s) \end{bmatrix} \mathrm{d}s}_{\vec{\mathrm{B}}} + \underbrace{\begin{bmatrix} 0 \\ \left(-\frac{1}{2}\sigma\sigma^\top(s)Y_s + \sigma u(Z_s, s)\right)\, \mathrm{d}s + \sigma(s)\vec{\mathrm{d}}W_s \end{bmatrix}}_{\vec{\mathrm{O}}}, \tag{48}$$

and use the following split for the inference SDE

$$\begin{bmatrix} \mathrm{d}X_s \\ \mathrm{d}Y_s \end{bmatrix} = \underbrace{\begin{bmatrix} Y_s \\ 0 \end{bmatrix} \mathrm{d}s}_{\overleftarrow{\mathrm{A}}} + \underbrace{\begin{bmatrix} 0 \\ \widetilde{f}(X_s, s) \end{bmatrix} \mathrm{d}s}_{\overleftarrow{\mathrm{B}}} + \underbrace{\begin{bmatrix} 0 \\ \left(-\frac{1}{2}\sigma\sigma^\top(s)Y_s - \sigma v(Z_s, s)\right)\, \mathrm{d}s + \sigma(s)\overleftarrow{\mathrm{d}}W_s \end{bmatrix}}_{\overleftarrow{\mathrm{O}}}. \tag{49}$$

Here, we use arrows to indicate whether the corresponding split belongs to the generative or inference SDE. To simplify the notation, we define $\sigma_s \coloneqq \sigma(s)$ and $\widetilde{f}_s \coloneqq \widetilde{f}(X_s, s)$.

A.8.1 OBAB

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

$$\widehat{Y}_n' = \widehat{Y}_n(1 + \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta} u(\widehat{Z}_n, n\delta)\delta + \sigma_{n\delta}\sqrt{\delta}\xi_n, \quad \xi_n \sim \mathcal{N}(0, I) \tag{50a}$$

$$\left.\begin{aligned} \widehat{Y}_n'' &= \widehat{Y}_n' + \widetilde{f}_n \tfrac{\delta}{2} \\ \widehat{X}_{n+1} &= \widehat{X}_n + \widehat{Y}_n''\delta \\ \widehat{Y}_{n+1} &= \widehat{Y}_n'' + \widetilde{f}_{n+1}\tfrac{\delta}{2} \end{aligned}\right\} \Phi \tag{50b}$$

$$\tag{50c}$$

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

$$\vec{p}_{n+1}(\widehat{Z}_{n+1}|\widehat{Z}_n) = \delta_{\Phi(\widehat{X}_n, \widehat{Y}_n')}(\widehat{Z}_{n+1})\mathcal{N}\left(\widehat{Y}_n'|\widehat{Y}_n(1 + \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}u(\widehat{Z}_n, n\delta)\delta, \sigma_{n\delta}\sigma_{n\delta}^\top\delta\right).$$

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

$$\left.\begin{aligned} \widehat{Y}_n'' &= \widehat{Y}_{n+1} - \widetilde{f}_{n+1}\tfrac{\delta}{2} \\ \widehat{X}_n &= \widehat{X}_{n+1} - \widehat{Y}_n''\delta \\ \widehat{Y}_n' &= \widehat{Y}_n'' - \widetilde{f}_n\tfrac{\delta}{2} \end{aligned}\right\} \Phi^{-1} \tag{51}$$

$$\widehat{Y}_n = \widehat{Y}_n'(1 - \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}v(\widehat{Z}_n', n\delta)\delta + \sigma_{n\delta}\sqrt{\delta}\xi_n, \quad \xi_n \sim \mathcal{N}(0, I), \tag{52}$$

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

$$\overleftarrow{p}_n(\widehat{Z}_n|\widehat{Z}_{n+1}) = \delta_{\Phi^{-1}(\widehat{Z}_{n+1})}(\widehat{X}_n, \widehat{Y}_n')\mathcal{N}\left(\widehat{Y}_n|\widehat{Y}_n'\left(1 - \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right) + \sigma_{n\delta}v(\widehat{Z}_n', n\delta)\delta, \sigma_{n\delta}\sigma_{n\delta}^\top\delta\right),$$

resulting in the following ratio between forward and backward transitions

$$\frac{\vec{p}_{n+1}(\widehat{Z}_{n+1}|\widehat{Z}_n)}{\overleftarrow{p}_n(\widehat{Z}_n|\widehat{Z}_{n+1})} = \frac{\mathcal{N}\left(\widehat{Y}_n'|\widehat{Y}_n(1 + \frac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}u(\widehat{Z}_n, n\delta)\delta, \sigma_{n\delta}\sigma_{n\delta}^\top\delta\right)}{\mathcal{N}\left(\widehat{Y}_n|\widehat{Y}_n'\left(1 - \frac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right) + \sigma_{n\delta}v(\widehat{Z}_n', n\delta)\delta, \sigma_{n\delta}\sigma_{n\delta}^\top\delta\right)}. \tag{53}$$

A.8.2 BAOAB

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

$$\left.\begin{aligned} \widehat{Y}_n' &= \widehat{Y}_n + \widetilde{f}_n \tfrac{\delta}{2} \\ \widehat{X}_n' &= \widehat{X}_n + \widehat{Y}_n' \tfrac{\delta}{2} \end{aligned}\right\} \Phi_1 \tag{54}$$

$$\widehat{Y}_n'' = \widehat{Y}_n'(1 + \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top \delta) + \sigma_{n\delta} u(\widehat{X}_n', \widehat{Y}_n', n\delta)\delta + \sigma_{n\delta}\sqrt{\delta}\xi_n \tag{55}$$

$$\left.\begin{aligned} \widehat{X}_{n+1} &= \widehat{X}_n' + \widehat{Y}_n'' \tfrac{\delta}{2} \\ \widehat{Y}_{n+1} &= \widehat{Y}_n'' + \widetilde{f}_{n+1} \tfrac{\delta}{2} \end{aligned}\right\} \Phi_2 \tag{56}$$

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

$$\vec{p}_{n+1}(\widehat{Z}_{n+1}|\widehat{Z}_n) = \delta_{\Phi_2(\widehat{X}_n', \widehat{Y}_n'')}(\widehat{Z}_{n+1}) \tag{57}$$

$$\times \mathcal{N}\left(\widehat{Y}_n'' | \widehat{Y}_n'(1 + \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top \delta) + \sigma_{n\delta} u(\widehat{X}_n', \widehat{Y}_n', n\delta)\delta, \sigma_{n\delta}\sigma_{n\delta}^\top \delta\right) \tag{58}$$

$$\times \delta_{\Phi_1(\widehat{Z}_n)}(\widehat{X}_n', \widehat{Y}_n'). \tag{59}$$

For $\overleftarrow{\mathrm{B}}\overleftarrow{\mathrm{A}}\overleftarrow{\mathrm{O}}\overleftarrow{\mathrm{A}}\overleftarrow{\mathrm{B}}$ we obtain

$$\left.\begin{aligned} \widehat{Y}_n'' &= \widehat{Y}_{n+1} - \widetilde{f}_{n+1} \tfrac{\delta}{2} \\ \widehat{X}_n' &= \widehat{X}_{n+1} - \widehat{Y}_n'' \tfrac{\delta}{2} \end{aligned}\right\} \Phi_2^{-1} \tag{60}$$

$$\widehat{Y}_n' = \widehat{Y}_n''(1 - \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top \delta) + \sigma_{n\delta} v(\widehat{X}_n', \widehat{Y}_n'', n\delta)\delta + \sigma_{n\delta}\sqrt{\delta}\xi_n \tag{61}$$

$$\left.\begin{aligned} \widehat{X}_n &= \widehat{X}_n' - \widehat{Y}_n' \tfrac{\delta}{2} \\ \widehat{Y}_n &= \widehat{Y}_n' - \widetilde{f}_n \tfrac{\delta}{2} \end{aligned}\right\} \Phi_1^{-1} \tag{62}$$

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

$$\overleftarrow{p}_n(\widehat{Z}_n|\widehat{Z}_{n+1}) = \delta_{\Phi_1^{-1}(\widehat{X}_n', \widehat{Y}_n')}(\widehat{Z}_n) \tag{63}$$

$$\times \mathcal{N}\left(\widehat{Y}_n' | \widehat{Y}_n''(1 - \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top \delta) + \sigma_{n\delta} v(\widehat{X}_n', \widehat{Y}_n'', n\delta)\delta, \sigma_{n\delta}\sigma_{n\delta}^\top \delta\right) \tag{64}$$

$$\times \delta_{\Phi_2^{-1}(\widehat{Z}_{n+1})}(\widehat{X}_n', \widehat{Y}_n''). \tag{65}$$

We, therefore, obtain the following ratio between forward and backward transitions:

$$\frac{\vec{p}_{n+1}(\widehat{Z}_{n+1}|\widehat{Z}_n)}{\overleftarrow{p}_n(\widehat{Z}_n|\widehat{Z}_{n+1})} = \frac{\mathcal{N}\left(\widehat{Y}_n'' | \widehat{Y}_n'(1 + \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top \delta) + \sigma_{n\delta} u(\widehat{X}_n', \widehat{Y}_n', n\delta)\delta, \sigma_{n\delta}\sigma_{n\delta}^\top \delta\right)}{\mathcal{N}\left(\widehat{Y}_n' | \widehat{Y}_n''(1 - \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top \delta) + \sigma_{n\delta} v(\widehat{X}_n', \widehat{Y}_n'', n\delta)\delta, \sigma_{n\delta}\sigma_{n\delta}^\top \delta\right)}. \tag{66}$$

A.8.3 OBABO

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

$$\widehat{Y}_n' = \widehat{Y}_n(1 + \tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top \delta) + \sigma_{n\delta} u(\widehat{Z}_n, n\delta)\tfrac{\delta}{2} + \sigma_{n\delta}\sqrt{\tfrac{\delta}{2}}\xi_n^{(1)} \tag{67}$$

$$\left.\begin{aligned} \widehat{Y}_n'' &= \widehat{Y}_n' + \widetilde{f}_n \tfrac{\delta}{2} \\ \widehat{X}_{n+1} &= \widehat{X}_n + \widehat{Y}_n'' \delta \\ \widehat{Y}_n''' &= \widehat{Y}_n'' + \widetilde{f}_{n+1} \tfrac{\delta}{2} \end{aligned}\right\} \Phi \tag{68}$$

$$\widehat{Y}_{n+1} = \widehat{Y}_n'''(1 + \tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top \delta) + \sigma_{n\delta} u(\widehat{X}_{n+1}, \widehat{Y}_n''', (n + \tfrac{1}{2})\delta)\tfrac{\delta}{2} + \sigma_{n\delta}\sqrt{\tfrac{\delta}{2}}\xi_n^{(2)} \tag{69}$$

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

$$\begin{aligned}\vec{p}_{n+1}(\widehat{Z}_{n+1}|\widehat{Z}_n) =&\mathcal{N}\left(\widehat{Y}_{n+1}|\widehat{Y}_n^{'''}(1+\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}u(\widehat{X}_{n+1},\widehat{Y}_n^{'''},(n+\tfrac{1}{2})\delta)\tfrac{\delta}{2}, \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right)\\ &\times \delta_{\Phi(\widehat{X}_n,\widehat{Y}_n^{'})}(\widehat{X}_{n+1},\widehat{Y}_n^{'''})\\ &\times \mathcal{N}\left(\widehat{Y}_n^{'}|\widehat{Y}_n(1+\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}u(\widehat{Z}_n, n\delta)\tfrac{\delta}{2}, \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right).\end{aligned}$$

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

$$\widehat{Y}_n^{'''} = \widehat{Y}_{n+1}(1-\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}v(\widehat{Z}_{n+1},(n+1)\delta)\tfrac{\delta}{2} + \sigma_{n\delta}\sqrt{\tfrac{\delta}{2}}\xi_n^{(2)} \tag{70}$$

$$\left.\begin{aligned}\widehat{Y}_n^{''} &= \widehat{Y}_n^{'''} - \widetilde{f}_{n+1}\tfrac{\delta}{2}\\ \widehat{X}_n &= \widehat{X}_{n+1} - \widehat{Y}_n^{''}\delta\\ \widehat{Y}_n^{'} &= \widehat{Y}_n^{''} - \widetilde{f}_n\tfrac{\delta}{2}\end{aligned}\right\}\Phi^{-1} \tag{71}$$

$$\widehat{Y}_n = \widehat{Y}_n^{'}(1-\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}v(\widehat{X}_n,\widehat{Y}_n^{'},(n+\tfrac{1}{2})\delta)\tfrac{\delta}{2} + \sigma_{n\delta}\sqrt{\tfrac{\delta}{2}}\xi_n^{(1)}, \tag{72}$$

with $(\widehat{X}_n, \widehat{Y}_n^{'}) = \Phi^{-1}(\widehat{X}_{n+1}, \widehat{Y}_n^{'''})$, giving the following backward transitions

$$\overleftarrow{p}_n(\widehat{Z}_n|\widehat{Z}_{n+1}) =\mathcal{N}\left(\widehat{Y}_n|\widehat{Y}_n^{'}(1-\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}v(\widehat{X}_n,\widehat{Y}_n^{'},(n+\tfrac{1}{2})\delta)\tfrac{\delta}{2}, \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right) \tag{73}$$

$$\times\, \delta_{\Phi^{-1}(\widehat{X}_{n+1},\widehat{Y}_n^{'''})}(\widehat{X}_n,\widehat{Y}_n^{'}) \tag{74}$$

$$\times\, \mathcal{N}\left(\widehat{Y}_n^{'''}|\widehat{Y}_{n+1}(1-\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}v(\widehat{Z}_{n+1},(n+1)\delta)\tfrac{\delta}{2}, \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right), \tag{75}$$

resulting in the following ratio between forward and backward transitions

$$\begin{aligned}\frac{\vec{p}_{n+1}(\widehat{Z}_{n+1}|\widehat{Z}_n)}{\overleftarrow{p}_n(\widehat{Z}_n|\widehat{Z}_{n+1})} =& \frac{\mathcal{N}\left(\widehat{Y}_{n+1}|\widehat{Y}_n^{'''}(1+\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}u(\widehat{X}_{n+1},\widehat{Y}_n^{'''},(n+\tfrac{1}{2})\delta)\tfrac{\delta}{2}, \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right)}{\mathcal{N}\left(\widehat{Y}_n^{'''}|\widehat{Y}_{n+1}(1-\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}v(\widehat{Z}_{n+1},(n+1)\delta)\tfrac{\delta}{2}, \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right)}\\ &\times \frac{\mathcal{N}\left(\widehat{Y}_n^{'}|\widehat{Y}_n(1+\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}u(\widehat{Z}_n,n\delta)\tfrac{\delta}{2}, \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right)}{\mathcal{N}\left(\widehat{Y}_n|\widehat{Y}_n^{'}(1-\tfrac{1}{4}\sigma_{n\delta}\sigma_{n\delta}^\top\delta) + \sigma_{n\delta}v(\widehat{X}_n,\widehat{Y}_n^{'},(n+\tfrac{1}{2})\delta)\tfrac{\delta}{2}, \tfrac{1}{2}\sigma_{n\delta}\sigma_{n\delta}^\top\delta\right)}.\end{aligned}$$

A.9 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 $\widetilde{f}$ and controls u and v in (15) to specific values recovers underdamped methods of ULA, MCD, and CMCD, see Tab. 2. Moreover, we can also introduce reference processes with controls $\widetilde{u}$ and $\widetilde{v}$ satisfying that

$$\frac{\mathrm{d}\vec{\mathbb{P}}^{\widetilde{u},\widetilde{\pi}}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{\widetilde{v},\widetilde{\tau}}} \equiv 1, \tag{76}$$

where $\widetilde{\pi}$ and $\widetilde{\tau}$ are known reference distributions. In other words, we have knowledge of a perfect time-reversal for specific controls $\widetilde{u}, \widetilde{v}$ and marginals $\widetilde{\pi}, \widetilde{\tau}$. We remark that these processes take a role similar to the Brownian motion used in the proof of Prop. 2.3. In particular, by applying Prop. 2.3

twice, we obtain that

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

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

A.10 Further computational details

A.10.1 Experimental setup

Here, we provide further details on our experimental setup. Moreover, we provide an algorithmic description of for 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, 2014), gradient clipping with a value of 1, and a learning rate scheduler that starts at 5×10^{-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^θ and v^γ with parameters θ and γ, respectively, were parameterized as two-layer neural networks with 128 neurons. Unlike Zhang & Chen (2021), we did not include the score of the target density as part of the parameterized control functions u^θ and v^γ. Inspired by Nichol & Dhariwal (2021), we applied a cosine-square scheduler for the discretization step size: $\delta = a\cos^2\left(\frac{\pi}{2}\frac{n}{N}\right)$, where $a : [0,\infty) \to (0,\infty)$ is learned. The diffusion matrix σ and the mass matrix M were parameterized as diagonal matrices, and we learned the parameters μ and Σ for the prior distribution $p_{\text{prior}} = \mathcal{N}(\mu,\Sigma)$, with Σ also set as a diagonal matrix. We enforced non-negativity of a and made σ, M, and Σ positive semidefinite via an element-wise softplus transformation.

For the methods that use geometric annealing (see Tab. 2), that is, $\nu(x,s) \propto p_{\text{prior}}^{1-\beta(s)}(x) p_{\text{target}}^{\beta(s)}(x)$, where $\beta\colon [0,T] \to [0,1]$ is a monotonically increasing function satisfying $\beta(0) = 0$ and $\beta(T) = 1$, we additionally learn the annealing schedule β. 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

$$\beta(n\delta) = \frac{\sum_z \text{softplus}(b(n'\delta))}{\sum_{n'\delta \le n\delta} \text{softplus}(b(n'\delta))}, \tag{77}$$

where $\mathrm{softplus}$ ensures non-negativity. Further, we fix $\beta(0) = 0$ and $\beta(T) = 1$. which ensures $\beta(n'\delta) \le \beta(n\delta)$ when $n' \le n$. We initialized b such that β is a linear interpolation between 0 and 1. Note that if not otherwise specified, we use $\nabla_x \log \nu$ as x-component for the drift $\widetilde{f}$.

Moreover, we initialized $\sigma = M = \Sigma = \mathrm{Id}$ 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 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.

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

METHOD	$\widetilde{f}$	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	$\sigma\sigma^\top \nabla_x \log p_{\mathrm{prior}}$	LEARNED	0
DBS	ARBITRARY	LEARNED	LEARNED

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 ρ_{target} factorizes across dimensions, we can compute a reference solution for $\log \mathcal{Z}$ via numerical integration, as described in Midgley et al. (2022).

A.10.2 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 further compute the (normalized) ESS as

$$\text{ESS} \coloneqq \frac{\left(\sum_{i=1}^n w^{(i)}\right)^2}{n \sum_{i=1}^n \left(w^{(i)}\right)^2}, \tag{78}$$

where $(w^{(i)})_{i=1}^n$ are the importance weights of the samples $(x^{(i)})_{i=1}^n$ in path space.

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

$$\mathbb{E}_{Z \sim \vec{\mathbb{P}}^{u,\pi}}\left[\log \frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{v,\tau}}(Z)\right] = 1. \tag{79}$$

Together with Prop. 2.3, this shows that

$$\log \mathcal{Z} = \mathbb{E}_{Z \sim \vec{\mathbb{P}}^{u,\pi}}\left[\log \frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}_{\cdot|0}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{v,\tau}_{\cdot|T}}(Z) + \frac{\pi(Z_0)}{\widetilde{\tau}(Z_T)}\right], \tag{80}$$

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

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	$-585.524_{\pm 0.414}$	$-73.437_{\pm 0.001}$	$-83.395_{\pm 4.184}$	$1.081_{\pm 0.004}$	$-111.673_{\pm 0.002}$	$-108.595_{\pm 0.006}$
	$-585.112_{\pm 0.001}$	$-73.423_{\pm 0.001}$	$\mathbf{-77.881_{\pm 0.014}}$	$\mathbf{1.136_{\pm 0.001}}$	$\mathbf{-111.636_{\pm 0.001}}$	$\mathbf{-108.458_{\pm 0.004}}$
GMMVI	$\mathbf{-585.098_{\pm 0.000}}$	$\mathbf{-73.415_{\pm 0.002}}$	$-77.988_{\pm 0.054}$	$1.092_{\pm 0.006}$	$-111.832_{\pm 0.007}$	$-108.726_{\pm 0.007}$
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	$-585.102_{\pm 0.001}$	$-73.418_{\pm 0.002}$	$-78.287_{\pm 0.835}$	$1.031_{\pm 0.010}$	$-111.678_{\pm 0.003}$	$-108.593_{\pm 0.008}$

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

$$\log \mathcal{Z} = \log \frac{\mathrm{d}\vec{\mathbb{P}}^{u^*,\pi}_{\cdot|0}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{v^*,\tau}_{\cdot|T}}(Z) + \frac{\pi(Z_0)}{\widetilde{\tau}(Z_T)} \tag{81}$$

If we only have approximations of u^* and v^*, Jensen's inequality shows that the right-hand side in (81) 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.10.3 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., 2022). 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.10), 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.

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 Fig. 6 and Fig. 7. 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 Fig. 8 (ESS values) and Tab. 6 ($\log \mathcal{Z}$ (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 Fig. 9. 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 $\widetilde{f}$ 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 β.

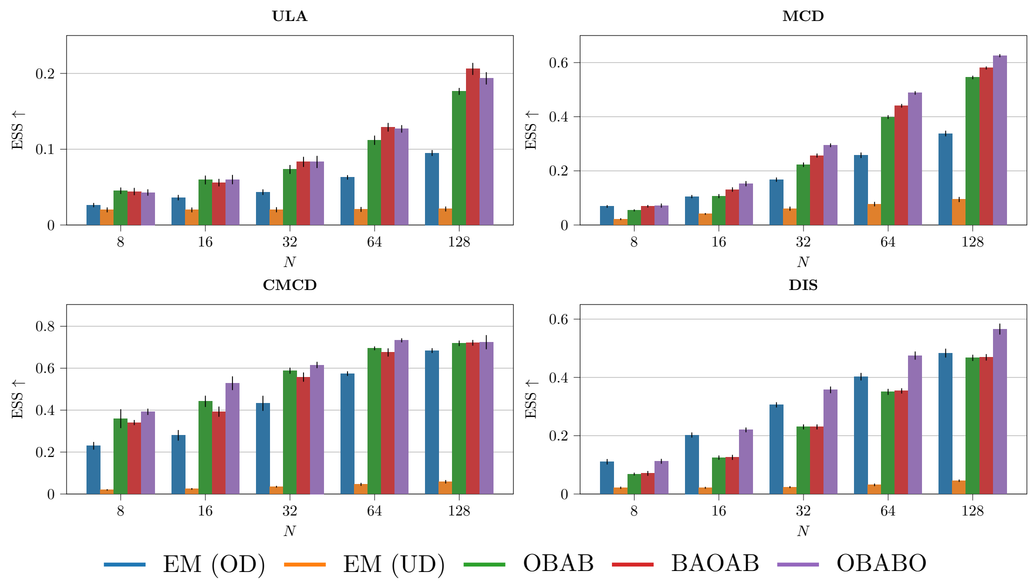

Figure 6: Effective sample size (ESS) for various methods (ULA, MCD, CMCD, DIS) and different integration schemes, averaged across multiple benchmark problems and four seeds. Integration schemes include Euler-Maruyama (EM) for over (OD) - and underdamped (OD) Langevin and various splitting schemes (OBAB, BAOAB, OBABO).

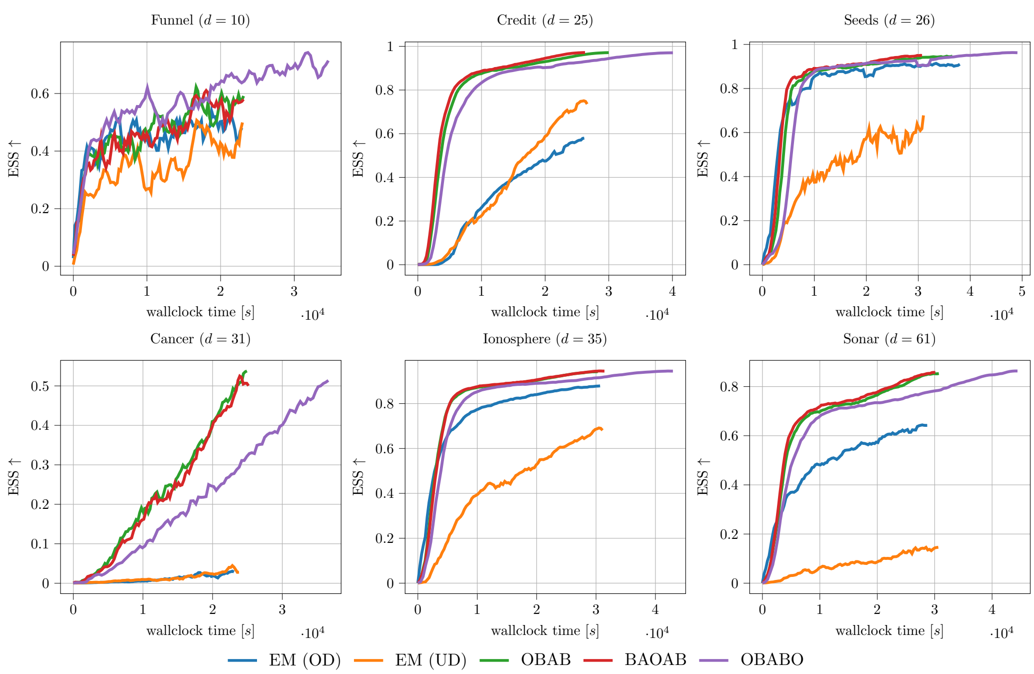

Figure 7: Effective sample size (ESS) over wallclock time of the diffusion bridge sampler (DBS) with 128 diffusion steps for different integration schemes, multiple benchmark problems, and four seeds. Integration schemes include Euler-Maruyama (EM) for over (OD) – and underdamped (UD) Langevin and various splitting schemes (OBAB, BAOAB, OBABO).

The results of these experiments are presented in Tab. 4 and Figure Fig. 10. The findings suggest that the most consistent performance is achieved when using the learned geometric annealing path as the drift $\widetilde{f}$. Interestingly, using the score of the target distribution ($\nabla_x \log p_{\text{target}}$) resulted in

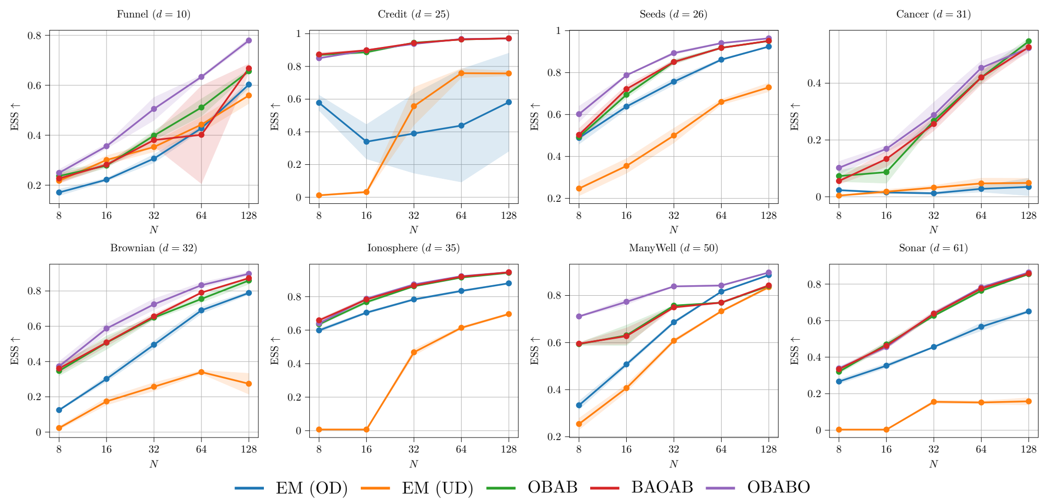

Figure 8: Effective sample size (ESS) of the diffusion bridge sampler (DBS) for different integration schemes, multiple benchmark problems, and four seeds. Integration schemes include Euler-Maruyama (EM) for over (OD) - and underdamped (UD) Langevin and various splitting schemes (OBAB, BAOAB, OBABO).

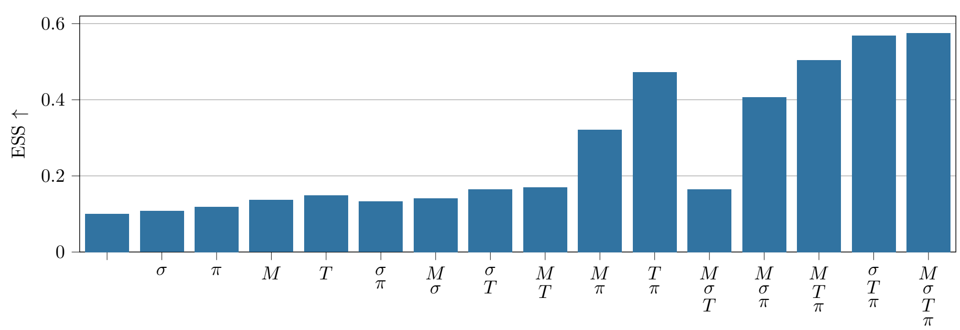

Figure 9: Effective sample size (ESS) of the underdamped diffusion bridge sampler (DBS) for various combinations of learned parameters, averaged across multiple benchmark problems and four seeds and $N = 8$ discretization steps. Parameters include mass matrix M, diffusion matrix σ, terminal time T, and extended prior distribution π.

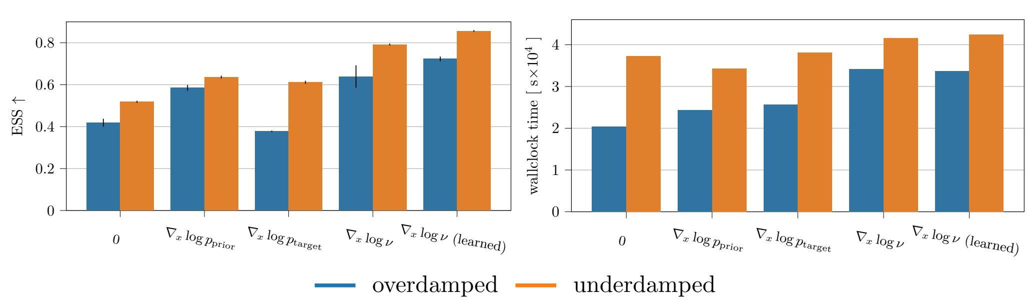

Figure 10: Effective sample size (ESS) and wallclock time for various drifts $\widetilde{f}$ of the underdamped and overdamped diffusion bridge sampler, averaged across multiple benchmark problems and four seeds. Here, $\nu(x, s) \propto p_{\text{prior}}^{1-\beta(s)}(x) p_{\text{target}}^{\beta(s)}(x)$, where '(learned)' indicates the β is learned (end-to-end).

worse performance compared to no drift for overdamped DBS and only marginal improvements for underdamped DBS.

Table 4: Lower bounds on $\log \mathcal{Z}$ for different drift function $\widetilde{f}$ 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 the β is learned (end-to-end). Blue shading indicates that the method uses underdamped Langevin dynamics.

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

Table 5: Results for *Funnel* and *ManyWell* using the Kullback-Leibler $\mathcal{L}_{\text{KL}}$ and log-variance loss $\mathcal{L}_{\text{LV}}$, averaged across four seeds. Evaluation criteria include importance-weighted errors for estimating the log-normalizing constant $\Delta \log Z$, effective sample size ESS, and Sinkhorn distance $\mathcal{W}_2^\gamma$. The best results are highlighted in bold. Arrows (↑, ↓) indicate whether higher or lower values are preferable, respectively. Blue shading indicates that the method uses the underdamped Langevin equation.

	Funnel ($d = 10$)			ManyWell ($d = 50$)	
Loss	$\Delta \log Z \downarrow$	ESS $\uparrow$	$\mathcal{W}_2^\gamma \downarrow$	$\Delta \log Z \downarrow$	ESS $\uparrow$
$\mathcal{L}_{\text{KL}}$	0.021 ± 0.003	0.603 ± 0.014	102.653 ± 0.586	0.005 ± 0.001	0.887 ± 0.004
	$\mathbf{0.010} \pm \mathbf{0.001}$	$\mathbf{0.779} \pm \mathbf{0.009}$	$\mathbf{101.418} \pm \mathbf{0.425}$	$\mathbf{0.005} \pm \mathbf{0.000}$	0.898 ± 0.002
$\mathcal{L}_{\text{LV}}$	0.504 ± 0.003	0.618 ± 0.025	117.679 ± 0.156	0.006 ± 0.001	0.866 ± 0.003
	0.593 ± 0.003	0.565 ± 0.393	123.587 ± 0.183	$\mathbf{0.005} \pm \mathbf{0.000}$	$\mathbf{0.942} \pm \mathbf{0.002}$

Log-Variance Loss. As an alternative to the KL divergence in (8), we can consider the log-variance (LV) loss:

$$\mathcal{L}_{\text{LV}}(u,v) := D_{\text{LV}}\left(\vec{\mathbb{P}}^{u,\pi}, \overleftarrow{\mathbb{P}}^{v,\tau}\right) = \operatorname{Var}_{Z \sim \vec{\mathbb{P}}^{w,\pi}} \left[\log \frac{\mathrm{d}\vec{\mathbb{P}}^{u,\pi}}{\mathrm{d}\overleftarrow{\mathbb{P}}^{v,\tau}}(Z)\right], \tag{82}$$

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^*) (Richter & Berner, 2024).

We conducted a preliminary comparison between the KL and LV losses, with results shown in Tab. 5. The findings are mixed: while the LV loss achieves superior performance on the multimodal *ManyWell* target, it falls behind the KL loss on the *Funnel* target. We plan to explore this further in future work, including investigating the impact of end-to-end learned parameters and degenerate diffusion matrices.

Table 6: 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	$-0.670_{\pm 0.003}$	$-585.168_{\pm 0.004}$	$-73.607_{\pm 0.005}$	$\mathbf{-79.167}_{\pm \mathbf{0.176}}$	$0.801_{\pm 0.004}$	$-111.822_{\pm 0.005}$	$42.502_{\pm 0.008}$	$-108.937_{\pm 0.015}$
BAOAB	$-0.674_{\pm 0.008}$	$\mathbf{-585.164}_{\pm \mathbf{0.010}}$	$-73.603_{\pm 0.011}$	$-79.252_{\pm 0.183}$	$0.807_{\pm 0.003}$	$\mathbf{-111.811}_{\pm \mathbf{0.007}}$	$42.497_{\pm 0.004}$	$\mathbf{-108.906}_{\pm \mathbf{0.019}}$
OBABO	$\mathbf{-0.557}_{\pm \mathbf{0.004}}$	$-585.179_{\pm 0.005}$	$\mathbf{-73.560}_{\pm \mathbf{0.008}}$	$-78.951_{\pm 0.050}$	$\mathbf{0.835}_{\pm \mathbf{0.017}}$	$-111.818_{\pm 0.009}$	$\mathbf{42.625}_{\pm \mathbf{0.002}}$	$-108.933_{\pm 0.007}$
	$N = 16$							
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	$-0.490_{\pm 0.003}$	$-585.149_{\pm 0.003}$	$-73.516_{\pm 0.003}$	$-78.685_{\pm 0.261}$	$0.944_{\pm 0.004}$	$-111.726_{\pm 0.007}$	$42.548_{\pm 0.037}$	$\mathbf{-108.754}_{\pm \mathbf{0.009}}$
OBABO	$\mathbf{-0.381}_{\pm \mathbf{0.007}}$	$\mathbf{-585.149}_{\pm \mathbf{0.002}}$	$\mathbf{-73.491}_{\pm \mathbf{0.003}}$	$\mathbf{-78.454}_{\pm \mathbf{0.027}}$	$\mathbf{0.977}_{\pm \mathbf{0.003}}$	$\mathbf{-111.725}_{\pm \mathbf{0.002}}$	$\mathbf{42.685}_{\pm \mathbf{0.009}}$	$-108.760_{\pm 0.010}$
	$N = 32$							
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	$-0.346_{\pm 0.003}$	$\mathbf{-585.126}_{\pm \mathbf{0.002}}$	$-73.465_{\pm 0.005}$	$-78.224_{\pm 0.014}$	$1.024_{\pm 0.004}$	$-111.680_{\pm 0.002}$	$42.665_{\pm 0.005}$	$-108.612_{\pm 0.006}$
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	$\mathbf{-0.249}_{\pm \mathbf{0.003}}$	$-585.129_{\pm 0.004}$	$\mathbf{-73.448}_{\pm \mathbf{0.002}}$	$\mathbf{-78.189}_{\pm \mathbf{0.069}}$	$\mathbf{1.048}_{\pm \mathbf{0.005}}$	$\mathbf{-111.673}_{\pm \mathbf{0.004}}$	$\mathbf{42.729}_{\pm \mathbf{0.002}}$	$\mathbf{-108.601}_{\pm \mathbf{0.008}}$
	$N = 64$							
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	$\mathbf{-0.164}_{\pm \mathbf{0.005}}$	$\mathbf{-585.113}_{\pm \mathbf{0.002}}$	$\mathbf{-73.431}_{\pm \mathbf{0.003}}$	$\mathbf{-77.945}_{\pm \mathbf{0.010}}$	$\mathbf{1.104}_{\pm \mathbf{0.003}}$	$\mathbf{-111.648}_{\pm \mathbf{0.002}}$	$\mathbf{42.730}_{\pm \mathbf{0.004}}$	$\mathbf{-108.501}_{\pm \mathbf{0.003}}$
	$N = 128$							
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	$-0.151_{\pm 0.003}$	$\mathbf{-585.112}_{\pm \mathbf{0.001}}$	$-73.428_{\pm 0.001}$	$\mathbf{-77.856}_{\pm \mathbf{0.007}}$	$1.121_{\pm 0.004}$	$-111.637_{\pm 0.001}$	$42.731_{\pm 0.002}$	$-108.459_{\pm 0.001}$
BAOAB	$-0.159_{\pm 0.005}$	$\mathbf{-585.112}_{\pm \mathbf{0.001}}$	$-73.428_{\pm 0.001}$	$-77.874_{\pm 0.010}$	$1.131_{\pm 0.002}$	$-111.637_{\pm 0.002}$	$42.733_{\pm 0.006}$	$\mathbf{-108.457}_{\pm \mathbf{0.004}}$
OBABO	$\mathbf{-0.103}_{\pm \mathbf{0.003}}$	$\mathbf{-585.112}_{\pm \mathbf{0.001}}$	$\mathbf{-73.423}_{\pm \mathbf{0.001}}$	$-77.881_{\pm 0.014}$	$\mathbf{1.136}_{\pm \mathbf{0.001}}$	$\mathbf{-111.636}_{\pm \mathbf{0.001}}$	$\mathbf{42.763}_{\pm \mathbf{0.002}}$	$-108.458_{\pm 0.004}$