

---

# SCALABLE SAMPLING VIA GENERALIZED FIXED-POINT DIFFUSION MATCHING

Denis Blessing<sup>1</sup>, Lorenz Richter<sup>2,3</sup>, Julius Berner<sup>4</sup>, Egor Malitskiy<sup>1</sup>, Gerhard Neumann<sup>1</sup>

<sup>1</sup>Karlsruhe Institute of Technology, <sup>2</sup>Zuse Institute Berlin, <sup>3</sup>dida, <sup>4</sup>NVIDIA

## ABSTRACT

Sampling from unnormalized densities using diffusion models has emerged as a powerful paradigm. However, while recent approaches that use least-squares ‘matching’ objectives have improved scalability, they often necessitate significant trade-offs, such as restricting prior distributions or relying on unstable optimization schemes. By generalizing these methods as special forms of fixed-point iterations rooted in Nelson’s relation, we develop a new method that addresses these limitations. Our approach enables learning a stochastic transport map between arbitrary prior and target distributions with a single, scalable, and stable objective. Furthermore, we introduce a damped variant of this iteration that incorporates a regularization term to mitigate mode collapse. Empirically, we demonstrate that our method enables sampling at unprecedented scales while preserving mode diversity, achieving state-of-the-art results on complex synthetic densities and high-dimensional molecular benchmarks.

## 1 INTRODUCTION

We consider the problem of sampling from a target distribution with density

$$p_{\text{target}}(x) = \frac{\rho_{\text{target}}(x)}{\mathcal{Z}},$$

where  $\rho_{\text{target}} \in C(\mathbb{R}^d, \mathbb{R}_{\geq 0})$  can be evaluated pointwise and  $\mathcal{Z} = \int_{\mathbb{R}^d} \rho_{\text{target}}(x) dx$  is the associated normalizing constant, which is generally intractable. This problem is fundamental in computational science and underlies a wide range of applications, including molecular dynamics Noé et al. (2019); Nam et al. (2025), statistical physics Hénin et al. (2022); Faulkner & Livingstone (2024), and Bayesian inference Neal (1993); Gelman et al. (2013). Our goal is to learn a stochastic transport map from a tractable prior distribution  $p_{\text{prior}}$ , e.g., a Gaussian, to the target  $p_{\text{target}}$ . Formally, we seek a vector field  $u$  such that the stochastic differential equation (SDE)

$$dX_t = \sigma(t)u(X_t, t)dt + \sigma(t)dB_t, \quad X_0 \sim p_{\text{prior}}, \quad (1)$$

yields a terminal marginal  $\mathbb{P}_T^u = p_{\text{target}}$ , where  $\mathbb{P}^u$  denotes the induced path measure.

Recent approaches based on *stochastic optimal control* (SOC) address this problem by introducing a corresponding time-reversed process

$$dY_t = -\sigma(t)v(Y_t, t)dt + \sigma(t)\bar{d}B_t, \quad Y_T \sim p_{\text{target}}, \quad (2)$$

with path measure  $\bar{\mathbb{P}}^v$  (Richter & Berner, 2024; Vargas et al., 2024; Blessing et al., 2025b). These methods attempt to align the forward and backward path measures by minimizing a divergence  $D(\mathbb{P}^u, \bar{\mathbb{P}}^v)$  over the controls  $u$  and  $v$ . This optimization yields (generally non-unique) optimal controls  $u^*$  and  $v^*$  that transport the prior to the target and vice versa. However, practical implementation of such objectives is often computationally prohibitive. The optimization requires the repeated simulation and storage of full trajectories of the SDE (1) to compute gradients – a bottleneck that becomes particularly challenging in high-dimensional settings. In contrast, approaches based on bridge processes and their associated path measures admit substantially more efficient simulation procedures. Specifically, one can consider path measures of the form<sup>1</sup>

$$\Pi^*(X) = p_{\text{prior}}(X_0)p_{\text{target}}(X_T)\mathbb{P}_{|0,T}(X|X_0, X_T), \quad (3)$$

---

<sup>1</sup>This can be viewed as the *reciprocal projection* of a process with independent coupling onto the reference process  $\mathbb{P}$ ; see Theorem 2.1.

Table 1: Overview: Our framework considers target path measures of the form  $\Pi^* = \Pi_{0,T}^* \mathbb{P}_{|0,T}$  with different couplings  $\Pi_{0,T}^*$ . We provide expressions for  $\xi$  that satisfy  $u^*(x, t) = \mathbb{E}_{\Pi^*} [\xi(X, t) | X_t = x]$  and are used for the proposed fixed-point iteration scheme in Algorithm 1. For an overview of the different expressions for  $\xi$  see Table 3 in Appendix D.

Coupling		$\xi$	Markovian	Algorithm
General	$\Pi_{0,T}^*$	Prop. 2.6	✗	
Schrödinger (half)	$\Pi_{0,T}^* = \mathbb{P}_0 \otimes \Pi_T^*$	Prop. 2.7	✓	AS Havens et al. (2025)
Schrödinger	$\Pi_{0,T}^* = \Pi_{0,T}^{\text{SB}}$	Prop. 2.9	✓	ASBS Liu et al. (2025)
Independent	$\Pi_{0,T}^* = \Pi_0^* \otimes \Pi_T^*$	Prop. 2.10	✗	BMS (ours)

where  $\mathbb{P}_{|0,T}$  denotes a tractable reference process conditioned on its endpoints, e.g., a Brownian bridge. The measure  $\Pi^*$  can be represented as the distribution of a non-Markovian stochastic process of the form

$$dX_t = \sigma(t)\xi(X, t)dt + \sigma(t)dB_t, \quad X_0 \sim p_{\text{prior}}, \quad (4)$$

where the functional  $\xi$  may depend on the entire path, and in particular on the terminal value  $X_T$ , e.g.,

$$\xi(X, t) = \sigma(t) \left( \int_t^T \sigma^2(s) ds \right)^{-1} (X_T - X_t) \quad (5)$$

in the case of a Brownian bridge  $\mathbb{P}_{|0,T}$ ; see, Albergo et al. (2025); Shi et al. (2023); Liu et al. (2023). By construction, the dependence in (3) enforces the terminal marginal  $X_T \sim p_{\text{target}}$ ; however, the process  $X$  in (4) can only be simulated when samples  $X_T \sim p_{\text{target}}$  are available.

The key idea pursued here is to *Markovianize* (see Theorem 2.2) the non-Markovian dynamics in (4) by minimizing suitable loss functionals, yielding a Markovian control  $u$  with associated path measure  $\mathbb{P}^u$  as in (1). The objective is to match the time marginals, namely  $\mathbb{P}_t^u = \Pi_t^*$  for all  $t \in [0, T]$ .

In this work, we combine the two approaches described above and introduce a scalable algorithm for Markovianization, which we term *Bridge Matching Sampler* (BMS). This method does not require samples from the target distribution, but only access to its unnormalized density. In this sense, our method can be viewed as a generalization of recently proposed sampling algorithms based on least-squares fixed-point iterations (often referred to as *matching* algorithms) arising from stochastic optimal control theory Havens et al. (2025); Liu et al. (2025); Nam et al. (2025). In contrast to these approaches, however, our framework accommodates general prior distributions and reference measures and avoids alternating optimization schemes, which are frequently unstable in practice. To further enhance stability and robustness, we introduce a damped variant of the fixed-point iteration that regularizes the update steps. As a result, BMS enables diffusion-based sampling methods to scale to higher-dimensional and more challenging problems, while outperforming existing state-of-the-art samplers in both accuracy and computational efficiency.

## 2 GENERALIZED FIXED-POINT DIFFUSION MATCHING

**Notation.** We denote by  $\mathcal{U} \subset C(\mathbb{R}^d \times [0, T], \mathbb{R}^d)$  the set of admissible controls and by  $\mathcal{P}$  the set of all probability measures on  $C([0, T], \mathbb{R}^d)$ . Moreover, we define  $\mathcal{M} \subset \mathcal{P}$  to be the subset of Markov measures, i.e., the laws of all Markov processes. We define the path space measure  $\mathbb{P} \in \mathcal{P}$  as the law of a  $\mathbb{R}^d$ -valued stochastic process  $X = (X)_{t \in [0, T]}$  and we denote by  $\mathbb{P}_s$  its marginal distribution at time  $s$ . Moreover, we denote by  $\mathbb{P}_{|0,T}$  the path measure  $\mathbb{P}$  conditioned on the random variables  $X_0$  and  $X_T$ . We refer to Appendix A for further details on our notation and assumptions.

**Overview.** To provide a roadmap for the theoretical developments that follow, let us first present a high-level overview of our framework. We propose a fixed-point iteration that alternates between the construction of valid bridge processes and their Markovianization to update the control fields. Conceptually, each iteration  $i$  consists of three steps (see also Algorithm 1 for a summary):

1. *Simulation and coupling:* We sample  $X_0 \sim p_{\text{prior}}$  from the prior and simulate the SDE in (1) using the current control  $u_i$  to obtain  $X_T$ . Together with a sample from the prior, this defines the current coupling.

2. *Reciprocal projection and target drift*: Conditioned on the coupling, we use a reference process  $\mathbb{P}_{|0,T}$  to define a bridge process  $\Pi^t$ . Our generalized target score identity in Theorem 2.5 can now be used to derive the drift  $\xi(X, t)$  for the process in (4) that corresponds to the target measure  $\Pi^*$ ; see Table 1.
3. *Markovianization*: Finally, we obtain  $u_{i+1}$  by minimizing a least-squares regression objective that fits the Markovian control  $u_i$  to the target drift  $\xi$  evaluated at samples drawn from the bridge process.

To present our general framework, which alternates between the construction of bridge processes – also known as *reciprocal projections* – and the computation of Markovian projections, we first introduce the relevant definitions.

**Definition 2.1** (Reciprocal class and reciprocal projection). A path measure  $\Pi \in \mathcal{P}$  is in the *reciprocal class*  $\mathcal{R}(\mathbb{P})$  of the reference process  $\mathbb{P} \in \mathcal{M}$  if  $\Pi = \Pi_{0,T}\mathbb{P}_{|0,T}$ . We define the *reciprocal projection* of  $\Pi$  onto  $\mathcal{R}(\mathbb{P})$  as

$$\text{proj}_{\mathcal{R}(\mathbb{P})}(\Pi) := \Pi_{0,T}\mathbb{P}_{|0,T}. \quad (6)$$

It is important to note that, for suitable reference measures  $\mathbb{P}$ , the conditioned process  $\mathbb{P}_{|0,T}$  can be simulated efficiently. For instance, when  $\mathbb{P}$  corresponds to (scaled) Brownian motion, the conditioned measure  $\mathbb{P}_{|0,T}$  is the classical *Brownian bridge*, for which the non-Markovian drift  $\xi$  admits the closed-form expression stated in (5). Moreover, for any  $t \in [0, T]$ , the marginal distribution  $X_t \sim \mathbb{P}_{t|0,T}$  can be sampled directly, without integrating full trajectories. This eliminates the need to store entire paths in memory and leads to substantial computational savings, especially in high-dimensional settings.

In our general framework, we consider a (possibly non-Markovian) target path measure  $\Pi^*$  subject to the boundary constraints

$$p_{\text{prior}}(x_0) = \int_{\mathbb{R}^d} \Pi_{0,T}^*(x_0, x_T) dx_T, \quad \text{and} \quad p_{\text{target}}(x_T) = \int_{\mathbb{R}^d} \Pi_{0,T}^*(x_0, x_T) dx_0. \quad (7)$$

Together with the bridge process introduced above, these constraints define a target path measure

$$\Pi^* = \Pi_{0,T}^*\mathbb{P}_{|0,T} \in \mathcal{R}(\mathbb{P}), \quad (8)$$

which can be represented as the law of the stochastic process

$$dX_t = \sigma(t)\xi(X, t)dt + \sigma(t)dB_t, \quad X_0 \sim p_{\text{prior}}. \quad (9)$$

Here, the drift  $\xi$  is a *path-dependent functional*, potentially depending on the entire trajectory  $(X_t)_{t \in [0, T]}$ . In the absence of samples from  $p_{\text{target}}$ , constructing a valid  $\xi$  is non-trivial. We will provide a feasible representation later. Our objective is to approximate  $\Pi^*$  by a Markov diffusion process of the form

$$dX_t = \sigma(t)u^*(X_t, t)dt + \sigma(t)dB_t, \quad (10)$$

whose time marginals coincide with those of the non-Markovian process (9). To formalize this approximation, we next introduce the notion of *Markovian projections*; see, e.g., Gyöngy (1986); Shi et al. (2023); Liu et al. (2022); Peluchetti (2023).

**Definition 2.2** (Markovian projection). Let  $\Pi^*$  be the path measure induced by the solution to the path-dependent SDE (9). The *Markovian projection* of  $\Pi^*$  is a path measure  $\mathbb{P}^{u^*}$  corresponding to the solution of the Markovian SDE (10) such that for every  $t \in [0, T]$ , the marginal distributions of both measures coincide, i.e.

$$\Pi_t^* = \mathbb{P}_t^{u^*}. \quad (11)$$

**Lemma 2.3** (Formula for Markov control). *Suppose  $X$  solves the path-dependent SDE (9) where  $\xi(X, t)$  is a functional of the path  $X$  and  $\Pi^*$  is the associated non-Markovian path measure. The drift  $u^*$  of the corresponding Markovian process is given by the conditional expectation*

$$u^*(x, t) = \mathbb{E}_{\Pi^*} [\xi(X, t) | X_t = x]. \quad (12)$$

Furthermore, it can be shown that

$$u^* = \arg \min_{u \in \mathcal{U}} D_{\text{KL}}(\Pi^* | \mathbb{P}^u). \quad (13)$$

We refer to Brunick & Shreve (2013) for a proof. While we do not have access to samples from  $\Pi^*$ , the Markovian projection formula (12) naturally suggests the fixed-point iteration

$$u_{i+1} = \Phi(u_i), \quad (14)$$

where

$$\Phi(u_i) := \mathbb{E}_{\Pi^i} [\xi(X, t) | X_t] = \arg \min_{u \in \mathcal{U}} \mathbb{E}_{\Pi^i} \left[ \int_0^T \frac{1}{2} \|\xi(X, t) - u(X_t, t)\|^2 dt \right]$$

and  $\Pi^i := \Pi_{0,T}^i \mathbb{P}_{|0,T} \in \mathcal{R}(\mathbb{P})$ . In other words, the method alternates between reciprocal projections and Markovianization, see also Algorithm 1. By construction, if  $u_i = u^*$ , then  $u_{i+1} = u^*$ , and hence  $u^*$  is a fixed point of the iteration. To make the above fixed-point iteration tractable, however, we need an expression for the non-Markovian drift  $\xi$ .

## 2.1 GENERAL COUPLINGS

To derive expressions for the path-dependent drift  $\xi$  in the non-Markovian SDE (9), we first recall that the optimal Markovian drifts  $u^*$  and  $v^*$  satisfy Nelson's identity

$$u^*(\cdot, t) + v^*(\cdot, t) = \sigma(t) \nabla \log \mathbb{P}_t^{u^*} = \sigma(t) \nabla \log \Pi_t^*, \quad (15)$$

where  $\Pi_t^*$  denotes the marginal of the target path measure at time  $t$ . Since the time marginals are fixed and by uniqueness of the Markovian projection, the drifts  $u^*$  and  $v^*$  are uniquely determined for a given coupling  $\Pi_{0,T}^*$ . They can be characterized explicitly as follows. We refer to Appendix B.1 in the appendix for the proof.

**Proposition 2.4** (Markovian projections). *Let  $\Pi^* \in \mathcal{R}(\mathbb{P})$ . The unique optimal drifts  $u^*$  and  $v^*$  corresponding to the forward and backward Markovian projections are*

$$\begin{aligned} u^*(x, t) &= \mathbb{E}_{\Pi_{T|t}^*} \left[ \sigma(t) \nabla_{X_t} \log \mathbb{P}_{T|t}(X_T | X_t) \Big| X_t = x \right], \quad \text{and} \\ v^*(x, t) &= \mathbb{E}_{\Pi_{0|t}^*} \left[ \sigma(t) \nabla_{X_t} \log \mathbb{P}_{t|0}(X_t | X_0) \Big| X_t = x \right]. \end{aligned}$$

Following the approach of bridge matching Liu et al. (2022); Peluchetti (2023), one might be tempted to set

$$\xi(X, t) = \sigma(t) \nabla_{X_t} \log \mathbb{P}_{T|t}(X_T | X_t), \quad (16)$$

to perform the fixed-point iteration in (14). However,  $\Phi(u_i)$  evaluates the expectation with respect to  $\Pi^i$  rather than the true target measure  $\Pi^*$ . Since the Markovian projection is marginal preserving, the coupling would remain static across all iterations, resulting in  $\Pi_{0,T}^i = \Pi_{0,T}^0$  for all  $i$ . Intuitively, this behavior is to be expected: when the expectation is taken over  $\Pi^i$ , the expressions in Theorem 2.4 do not incorporate any information about the prior or target boundary densities.

Instead of relying on Theorem 2.4 alone, we leverage Nelson's identity (15) and treat  $v^*$  and  $\nabla \log \Pi_t^*$  separately. While  $v^*$  has already been characterized in Theorem 2.4, we derive an identity for  $\nabla \log \Pi_t^*$  in the following proposition. We term this identity *generalized target score identity* (TSI) since it allows us to recover and generalize many of the *target score identities* in the literature De Bortoli et al. (2024).

**Proposition 2.5** (Generalized target score identity). *Let the reference measure  $\mathbb{P}$  be induced by scaled Brownian motion,  $dX_t = \sigma(t)dB_t$ , and define  $\kappa(t) := \int_0^t \sigma^2(s) ds$ ,  $\gamma(t) := \frac{\kappa(t)}{\kappa(T)}$ . Let  $c \in C([0, T], (0, 1])$  be an arbitrary function<sup>2</sup>. Then, for the target path measure  $\Pi^* = \Pi_{0,T}^* \mathbb{P}_{|0,T}$ , it holds*

$$\nabla_x \log \Pi_t^*(x) = \mathbb{E}_{\Pi_{0,T|t}^*} \left[ \frac{1 - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) \Big| X_t = x \right].$$

Combining Theorem 2.4 and Theorem 2.5, we can now get an expression for the non-Markovian drift  $\xi$ .

<sup>2</sup>The function  $c \in C([0, T], (0, 1])$  can be seen as a control variate; see Appendix D.2 for further details.

**Proposition 2.6** (Path-dependent drift of general target measure). *Under the assumptions of Theorem 2.5, the path-dependent drift  $\xi$  corresponding to the non-Markovian SDE (9) that induces the target measure  $\Pi^*$  is given by  $\sigma(t)^{-1}\xi(X, t) =$*

$$\frac{1 - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) - \nabla_{X_t} \log \mathbb{P}_{t|0}(X_t|X_0)$$

While one could, in principle, use Theorem 2.6 to perform the fixed-point iteration suggested in Algorithm 1, in practice we typically do not have access to the coupling scores  $\nabla \log \Pi_{0,T}^*$ . However, for certain special forms of  $\Pi_{0,T}^*$ , these scores can be computed explicitly, leading to tractable algorithms for learning the stochastic transport map  $u^*$ . The remainder of this section focuses on these tractable cases.

## 2.2 SCHRÖDINGER HALF BRIDGES

*Schrödinger half bridges* (SHBs) solve problems of the form

$$\min_{u \in \mathcal{U}} \mathbb{E}_{\mathbb{P}^u} \left[ \int_0^T \frac{1}{2} \|u(X_t, t)\|^2 dt \right] \text{ s.t. } X_T \sim \Pi_T^* = p_{\text{target}},$$

i.e., they are characterized by a kinetic optimal drift given a marginal constraint<sup>3</sup>. Using that

$$D_{\text{KL}}(\mathbb{P}^u | \mathbb{P}) = \mathbb{E}_{\mathbb{P}^u} \left[ \int_0^T \frac{1}{2} \|u(X_t, t)\|^2 dt \right], \quad (17)$$

and by the chain rule of the KL divergence it is straightforward to see that the minimum is attained when the conditional measures match,  $\mathbb{P}_{|T}^u = \mathbb{P}_{|T}$ . Consequently, the optimal joint coupling is given by  $\Pi_{0,T}^* = \Pi_T^* \mathbb{P}_{0|T}$ . However, this implies that the initial marginal is  $\Pi_0^* = \mathbb{E}_{\Pi_T^*}[\mathbb{P}_{0|T}]$ , which generally depends on the target  $\Pi_T^*$  and differs from the reference prior  $\mathbb{P}_0$ . To ensure that the optimal model actually starts at the prior (i.e.,  $\Pi_0^* = p_{\text{prior}}$ ), the reference must satisfy the independence condition  $\mathbb{P}_{0|T} = \mathbb{P}_0$ , a property commonly known as *memorylessness* Domingo-Enrich et al. (2025). For instance, this can be enforced by using a deterministic initial condition  $X_0 = x_0$ , i.e.  $p_{\text{prior}} = \delta_{x_0}$ . Using the Brownian motion reference process and  $x_0 = 0$  – a setting that is considered frequently in the literature – the generalized TSI in Theorem 2.5 then becomes

$$\nabla \log \Pi_t^*(x) = \mathbb{E}_{\Pi_{T|t}^*} \left[ \frac{1}{\gamma(t)} \nabla_{X_T} \log \Pi_T^*(X_T) | X_t = x \right].$$

Combining this with the general results for  $v^*$  in Theorem 2.4 yields the following result, see Appendix B.3 for the proof.

**Proposition 2.7** (Path-dependent drift for Schrödinger half bridges). *Under the assumptions of Theorem 2.5, the path-dependent drift  $\xi$  corresponding to the non-Markovian SDE (9) that induces the target measure  $\Pi^* = \Pi_T^* \mathbb{P}_{|T}$  is given by*

$$\sigma(t)^{-1}\xi(X, t) = \nabla_{X_T} \log \frac{\Pi_T^*(X_T)}{\mathbb{P}_T(X_T)}. \quad (18)$$

We note that the result in Theorem 2.7, together with the associated fixed-point iteration, coincides with the *adjoint sampling* (AS) method introduced by Havens et al. (2025). While their derivation is based on a stochastic optimal control (SOC) formulation combined with the adjoint method, the same scheme arises in our setting as a special case of the general framework developed here. However, as mentioned above, the memorylessness assumptions severely restricts the choice of the reference process and the prior distribution and, in practice, these choices lead to instabilities due to large diffusion coefficients  $\sigma$ . In particular, for the Dirac delta prior, we need a sufficiently large coefficient to guarantee initial exploration of the sampler. For a Gaussian prior, one requires sufficiently large diffusion coefficients such that the distribution of the time-reversed process approximates the prior at finite time  $T$ . In the following, we present a method that allows for the use of arbitrary prior distributions and reference processes, thereby generalizing the previously introduced Schrödinger half-bridges.

<sup>3</sup>One can also consider constraining the marginal at time  $t = 0$ .

### 2.3 SCHRÖDINGER BRIDGES

Similar to half-bridges, *Schrödinger bridges* (SBs) seek a kinetically optimal drift, however, containing both boundary constraints, i.e.,

$$\min_{u \in \mathcal{U}} \mathbb{E} \left[ \int_0^T \frac{1}{2} \|u(X_t, t)\|^2 dt \right] \text{ s.t. } \begin{cases} X_0 \sim \Pi_0^* = p_{\text{prior}}, \\ X_T \sim \Pi_T^* = p_{\text{target}}. \end{cases}$$

The solution to the SB problem can be characterized using the Schrödinger system in the following proposition; see, e.g., Léonard (2013); Chen et al. (2016), and Appendix B.2 in the appendix.

**Proposition 2.8** (Schrödinger system). *Let  $\varphi, \hat{\varphi} \in C(\mathbb{R}^d \times [0, T]; \mathbb{R})$  satisfy the Schrödinger system, i.e.,*

$$\varphi_t(x) = \int \mathbb{P}_{T|t}(x_T|x) \varphi_T(x_T) dx_T \quad \text{and} \quad \hat{\varphi}_t(x) = \int \mathbb{P}_{t|0}(x|x_0) \hat{\varphi}_0(x_0) dx_0, \quad (19)$$

with  $\varphi_0 \hat{\varphi}_0 = \Pi_0^*$ ,  $\varphi_T \hat{\varphi}_T = \Pi_T^*$ ,  $u^* = \sigma \nabla \log \varphi$  and  $v^* = \sigma \nabla \log \hat{\varphi}$ . Moreover, the Schrödinger coupling satisfies

$$\Pi_{0,T}^* = \Pi_{0,T}^{\text{SB}} := \mathbb{P}_{T|0} \hat{\varphi}_0 \varphi_T. \quad (20)$$

Combining the generalized TSI in Theorem 2.5 and (20), it is straightforward to define a TSI for SBs and derive the path-dependent drift, see Appendix B.2 in the appendix for further details, where we also state a more general version in Theorem B.5.

**Proposition 2.9** (Path-dependent drift for Schrödinger bridges). *Under the assumptions of Theorem 2.5, the path-dependent drift  $\xi$  corresponding to the non-Markovian SDE (9) that induces the target measure  $\Pi^* = \Pi_{0,T}^{\text{SB}} \mathbb{P}_{|0,T}$  is given by*

$$\sigma(t)^{-1} \xi(X, t) = \nabla_{X_T} \log \frac{\Pi_T^*(X_T)}{\hat{\varphi}_T(X_T)}. \quad (21)$$

The result in Theorem 2.9, together with the associated fixed-point iteration, recovers the *Adjoint Schrödinger Bridge Sampler* (ASBS) introduced by Liu et al. (2025), which was originally derived from a stochastic optimal control perspective. While Schrödinger bridges (SBs) allow for arbitrary prior distributions  $p_{\text{prior}}$ , their associated fixed-point iteration requires access to the terminal potential  $\hat{\varphi}_T$ , which is typically unavailable. To address this, Liu et al. (2025) proposed a variant based on *iterative proportional fitting* (IPF) Kullback (1968), which alternates between updating the drift  $u$  and the terminal potential  $\hat{\varphi}_T$ . Because these updates are interdependent, the resulting algorithm can be unstable. In the following, we propose a method that enables the use of arbitrary prior and target distributions without requiring such alternating optimization schemes.

### 2.4 INDEPENDENT COUPLINGS

Recalling Theorem 2.6, which states a path-dependent drift for general target measures, let us now consider the special case when the coupling is assumed to be independent,

$$\Pi_{0,T}^* = \Pi_0^* \otimes \Pi_T^* = p_{\text{prior}} \otimes p_{\text{target}}.$$

In this case, the coupling scores  $\nabla \log \Pi_{0,T}^*$  become tractable, which directly yields the following proposition, see also Appendix B.4.

**Proposition 2.10** (Target score and drift for independent couplings). *Under the assumptions of Theorem 2.5, and additionally assuming that the coupling  $\Pi_{0,T}^*$  factorizes as  $\Pi_{0,T}^* = \Pi_0^* \otimes \Pi_T^*$ , the following identity holds:*

$$\nabla_x \log \Pi_t^*(x) = \mathbb{E}_{\Pi_{0,T|t}^*} \left[ \frac{1-c(t)}{1-\gamma(t)} \nabla_{X_0} \log \Pi_0^*(X_0) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_T^*(X_T) \Big| X_t = x \right].$$

Consequently, the associated drift is given by

$$\sigma(t)^{-1} \xi(X, t) = \frac{1-c(t)}{1-\gamma(t)} \nabla_{X_0} \log \Pi_0^*(X_0) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_T^*(X_T) - \nabla_{X_t} \log \mathbb{P}_{t|0}(X_t|X_0). \quad (22)$$




Figure 1: Ramachandran plots with  $10^6$  samples for Alanine dipeptide for Molecular Dynamics (MD) data, ASBS Liu et al. (2025), and our method, BMS, for damping values  $\eta \in \{0, 10\}$ . ASBS and BMS are trained on all-atom coordinates solely using energy evaluations.

With Theorem 2.10, we obtain a tractable path-dependent drift  $\xi$  corresponding to the non-Markovian target measure  $\Pi^*$ , which can be directly employed in Algorithm 1. This allows us to pick arbitrary prior distributions and reference processes, including arbitrary diffusion coefficients  $\sigma$ , while still retaining a single, scalable matching objective. Together with the damping outlined in the next section, this defines our proposed *Bridge Matching Sampler* (BMS). A detailed description of the algorithm is provided in Algorithm 1 in Appendix E, which additionally highlights how the integration of replay buffers enables data reuse, thereby significantly increasing the method’s sample efficiency.

## 2.5 DAMPED FIXED-POINT DIFFUSION MATCHING

In order to update the control gradually and thereby improve the robustness of the fixed-point iteration, we extend (14) by introducing the *damped fixed-point iteration*

$$u_{i+1} = \alpha \Phi(u_i) + (1 - \alpha)u_i, \quad (23)$$

where  $\alpha \in (0, 1]$  is the step size. This damped scheme admits a variational characterization, as formalized in the following proposition, see Appendix B.5 for the proof.

**Proposition 2.11** (Variational characterization of damped fixed-point iteration). *Consider the fixed-point operator  $\Phi$  defined in (14). Then, the damped fixed-point iteration (23) satisfies that*

$$u_{i+1} = \arg \min_{u \in \mathcal{U}} \left\{ \mathbb{E}_{\Pi^i} \left[ \int_0^T \frac{1}{2} \|\xi(X, t) - u(X_t, t)\|^2 dt \right] + \eta \mathbb{E}_{\Pi^i} \left[ \int_0^T \frac{1}{2} \|u_i(X_t, t) - u(X_t, t)\|^2 dt \right] \right\},$$

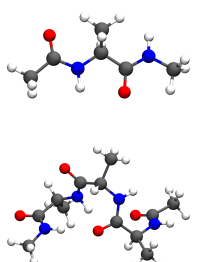
with  $\Pi^i = \Pi_{0,T}^i \mathbb{P}_{|0,T}$  and damping parameter  $\eta = \frac{1-\alpha}{\alpha}$ .

This variational perspective shows that the damped iteration regularizes the current drift to remain close to the previous iterate  $u_i$  in the  $L^2$  sense. Similar regularization techniques are common in proximal point methods Parikh et al. (2014); Hertrich & Gruhlke (2024) and trust-region methods Conn et al. (2000); Schulman et al. (2015), which have recently gained traction in the diffusion model literature Guo et al. (2025); Blessing et al. (2025a). In practice, damping is integrated into Algorithm 1 by modifying the Markovianization step and storing the parameters of the previous iteration to evaluate  $u_i$ . As shown in our experiments (Section 3), this modification substantially improves the performance of both, baseline methods and our proposed approach, particularly in high-dimensional settings.

## 3 NUMERICAL EXPERIMENTS

We evaluate the performance of our proposed BMS by comparing it to key instances of our framework: the *adjoint sampler* (AS) Havens et al. (2025) and the *adjoint Schrödinger bridge sampler* (ASBS) Liu et al. (2025). Our evaluation spans three distinct testbeds: high-dimensional synthetic Gaussian mixture models,  $n$ -body identical particle systems, and two molecular peptide systems; see Appendix F. Detailed descriptions of the experimental setup, benchmark problems, and evaluation metrics are provided in Appendix E, with supplementary results available in Appendix F.

Table 2: Comparison of performance metrics on Alanine Dipeptide (ALA2) and Tetrapeptide (ALA4). We report the Jensen–Shannon divergence of the joint backbone dihedral angle distribution ( $\phi, \psi$ ) ( $D_{JS}$ ), the root-mean-squared error between the corresponding free-energy surfaces ( $\text{RMSE}_{\Delta F}$ ), and the Wasserstein-2 distance between the first two TIC components ( $\text{TICA-}W_2$ ). Runs marked with † diverged during training; for these, results from the final checkpoint are shown. All metrics are averaged over three random seeds. The best-performing method is highlighted in bold.



System	Method	$\eta$	$(\phi, \psi)\text{-}D_{JS} \downarrow$	$\text{RMSE}_{\Delta F} \downarrow$	$\text{TICA-}W_2 \downarrow$
ALA2 ( $d = 66$ )	ASBS <sup>Liu et al. (2025)</sup>	0	0.232 $\pm$ 0.018	0.204 $\pm$ 0.003	1.263 $\pm$ 0.140
	ASBS-damped (ours)	10	0.111 $\pm$ 0.017	0.128 $\pm$ 0.008	1.526 $\pm$ 0.150
	BMS (ours)	0	0.311 $\pm$ 0.050	0.742 $\pm$ 0.514	2.707 $\pm$ 0.268
	BMS-damped (ours)	10	<b>0.068</b> $\pm$ 0.009	<b>0.079</b> $\pm$ 0.000	<b>0.900</b> $\pm$ 0.066
ALA4 ( $d = 126$ )	ASBS <sup>†</sup> <sup>Liu et al. (2025)</sup>	0	0.489 $\pm$ 0.016	1.368 $\pm$ 0.836	474.8 $\pm$ 41.91
	ASBS-damped (ours)	50	0.230 $\pm$ 0.003	0.457 $\pm$ 0.073	<b>1.606</b> $\pm$ 0.028
	BMS <sup>†</sup> (ours)	0	0.502 $\pm$ 0.027	0.811 $\pm$ 0.158	280.3 $\pm$ 1.526
	BMS-damped (ours)	50	<b>0.228</b> $\pm$ 0.011	<b>0.396</b> $\pm$ 0.012	1.697 $\pm$ 0.023

**Molecular benchmarks: Peptide systems.** We evaluate our method on two increasingly complex molecular benchmarks: Alanine Dipeptide (ALA2;  $d = 66$ ) and Alanine Tetrapeptide (ALA4;  $d = 126$ ). Notably, we train our model directly on Cartesian coordinates, a departure from existing literature which typically relies on prior domain knowledge – such as internal coordinates Midgley et al. (2022); Liu et al. (2025), collective variables – to enhance exploration Nam et al. (2025), or pretraining on molecular dynamics (MD) data with high temperatures Akhound-Sadegh et al. (2025). To assess the quality of the generated conformations, we analyze Ramachandran plots, which represent the joint distribution of the backbone dihedral angles  $\phi$  and  $\psi$ . These angles characterize the local conformation of the molecular backbone and are essential for identifying stable states. We compute the Jensen-Shannon divergence of pairs of backbone dihedral angles between MD data and samples generated by the model ( $D_{JS}$ ), averaged over all residues. Similarly, we compute the root-mean-squared error between the corresponding free-energy surfaces ( $\text{RMSE}_{\Delta F}$ ). Additionally, we compute the Wasserstein-2 metric between the first two TIC components ( $\text{TICA-}W_2$ ) to ensure the capture of global conformational dependencies. Qualitative results (Ramachandran and TICA visualizations) are presented in Figures 1 and 7 for ALA2 and Figure 8 for ALA4. Quantitative results are shown Table 2 and Figure 10.

While our method achieves state-of-the-art performance across both systems, the results underscore the critical role of the proposed damped-fixed point iteration scheme. Specifically, omitting the damping component leads to poor performance and significant training instability in both ASBS and BMS, as evidenced in Figure 10. Note that we were not able to obtain results for ALA4 without damping due to numerical instabilities. A notable characteristic across both systems is their mode-covering behavior. Unlike traditional sampling techniques that often suffer from mode-collapse – typically driven by the mode-seeking nature of the reverse KL divergence – our approach successfully maintains probability mass across all target regions. We conjecture that this mass-covering property is a direct result of our fixed-point iteration; as shown in (13), the fixed point corresponds to the Markovian projection, which is inherently linked to a forward KL objective.

## 4 CONCLUSION

In this work, we formulate diffusion-based sampling from unnormalized target densities as a generalized fixed-point iteration grounded in Nelson’s relation and a generalized target score identity (see Theorem 2.5). This perspective provides a unified theoretical framework encompassing several recently proposed matching-based algorithms, including Adjoint Sampling Havens et al. (2025) and Adjoint Schrödinger Bridge Samplers Liu et al. (2025); Nam et al. (2025). Building on this foundation, we introduce a new scalable and stable algorithm, termed the *Bridge Matching Sampler* (BMS), which addresses key limitations of existing approaches – most notably alternating objectives and restrictive choices of reference processes or priors. To further improve robustness and avoid mode collapse in high dimensions, we propose a damped fixed-point iteration. Empirically, BMS achieves state-of-the-art performance on challenging synthetic targets and high-dimensional molecular benchmarks, and scales successfully to highly multimodal distributions in dimensions up to  $d = 2500$ .

---

## 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.
- Tara Akhound-Sadegh, Jungyoon Lee, Avishek Joey Bose, Valentin De Bortoli, Arnaud Doucet, Michael M Bronstein, Dominique Beaini, Siamak Ravanbakhsh, Kirill Neklyudov, and Alexander Tong. Progressive inference-time annealing of diffusion models for sampling from Boltzmann densities. *arXiv preprint arXiv:2506.16471*, 2025.
- Michael Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025.
- Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024.
- Julius Berner, Lorenz Richter, Marcin Sendera, Jarrid Rector-Brooks, and Nikolay Malkin. From discrete-time policies to continuous-time diffusion samplers: Asymptotic equivalences and faster training. *arXiv preprint arXiv:2501.06148*, 2025.
- 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.
- Denis Blessing, Julius Berner, Lorenz Richter, Carles Domingo-Enrich, Yuanqi Du, Arash Vahdat, and Gerhard Neumann. Trust region constrained measure transport in path space for stochastic optimal control and inference. *arXiv preprint arXiv:2508.12511*, 2025a.
- Denis Blessing, Julius Berner, Lorenz Richter, and Gerhard Neumann. Underdamped diffusion bridges with applications to sampling. In *The Thirteenth International Conference on Learning Representations*, 2025b.
- Gerard Brunick and Steven Shreve. Mimicking an Itô process by a solution of a stochastic differential equation. *The Annals of Applied Probability*, 23(4), August 2013. ISSN 1050-5164. doi: 10.1214/12-aap881. URL <http://dx.doi.org/10.1214/12-AAP881>.
- Junhua Chen, Lorenz Richter, Julius Berner, Denis Blessing, Gerhard Neumann, and Anima Anandkumar. Sequential controlled Langevin diffusions. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- 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.
- Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169(2):671–691, 2016.
- Andrew R Conn, Nicholas IM Gould, and Philippe L Toint. *Trust region methods*. SIAM, 2000.
- Valentin De Bortoli, Michael Hutchinson, Peter Wirnsberger, and Arnaud Doucet. Target score matching. *arXiv preprint arXiv:2402.08667*, 2024.
- Carles Domingo-Enrich, Jiequn Han, Brandon Amos, Joan Bruna, and Ricky T. Q. Chen. Stochastic optimal control matching. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Carles Domingo-Enrich, Michal Drozdal, Brian Karrer, and Ricky T. Q. Chen. Adjoint matching: Fine-tuning flow and diffusion generative models with memoryless stochastic optimal control. In *The Thirteenth International Conference on Learning Representations*, 2025.

- 
- 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*, 2022.
- Peter Eastman, Raimondas Galvelis, Raúl P Peláez, Charles RA Abreu, Stephen E Farr, Emilio Gallicchio, Anton Gorenko, Michael M Henry, Frank Hu, Jing Huang, et al. Openmm 8: molecular dynamics simulation with machine learning potentials. *The Journal of Physical Chemistry B*, 128(1):109–116, 2023.
- Ezra Erives, Bowen Jing, Peter Holderrieth, and Tommi Jaakkola. Continuously tempered diffusion samplers. In *Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025.
- Michael F Faulkner and Samuel Livingstone. Sampling algorithms in statistical physics: a guide for statistics and machine learning. *Statistical Science*, 39(1):137–164, 2024.
- Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, et al. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
- 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.
- Louis Grenioux, Maxence Noble, and Marylou Gabrié. Improving the evaluation of samplers on multi-modal targets. *arXiv preprint arXiv:2504.08916*, 2025.
- Timofei Gritsaev, Nikita Morozov, Kirill Tamogashev, Daniil Tiapkin, Sergey Samsonov, Alexey Naumov, Dmitry Vetrov, and Nikolay Malkin. Adaptive destruction processes for diffusion samplers. *arXiv preprint arXiv:2506.01541*, 2025.
- Wei Guo, Jaemoo Choi, Yuchen Zhu, Molei Tao, and Yongxin Chen. Proximal diffusion neural sampler. *arXiv preprint arXiv:2510.03824*, 2025.
- István Gyöngy. Mimicking the one-dimensional marginal distributions of processes having an Itô differential. *Probability theory and related fields*, 71(4):501–516, 1986.
- Aaron Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Brandon Wood, Daniel Levine, Bin Hu, Brandon Amos, Brian Karrer, et al. Adjoint sampling: Highly scalable diffusion samplers via adjoint matching. *arXiv preprint arXiv:2504.11713*, 2025.
- Jiajun He, Wenlin Chen, Mingtian Zhang, David Barber, and José Miguel Hernández-Lobato. Training neural samplers with reverse diffusive KL divergence. *arXiv preprint arXiv:2410.12456*, 2024.
- Jérôme Hénin, Tony Lelièvre, Michael R Shirts, Omar Valsson, and Lucie Delemotte. Enhanced sampling methods for molecular dynamics simulations. *arXiv preprint arXiv:2202.04164*, 2022.
- Johannes Hertrich and Robert Gruhlke. Importance corrected neural JKO sampling. *arXiv preprint arXiv:2407.20444*, 2024.
- Xunpeng Huang, Hanze Dong, Yifan Hao, Yian Ma, and Tong Zhang. Monte Carlo sampling without isoperimetry: A reverse diffusion approach. *arXiv preprint arXiv:2307.02037*, 2023.
- Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.
- Khaled Kahouli, Romuald Elie, Klaus-Robert Müller, Quentin Berthet, Oliver T Unke, and Arnaud Doucet. Control variate score matching for diffusion models. *arXiv preprint arXiv:2512.20003*, 2025.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022.

- 
- Minkyu Kim, Kiyoungh Seong, Dongyeop Woo, Sungsoo Ahn, and Minsu Kim. On scalable and efficient training of diffusion samplers. *arXiv preprint arXiv:2505.19552*, 2025.
- Minsu Kim, Sanghyeok Choi, Taeyoung Yun, Emmanuel Bengio, Leo Feng, Jarrid Rector-Brooks, Sungsoo Ahn, Jinkyoo Park, Nikolay Malkin, and Yoshua Bengio. Adaptive teachers for amortized samplers. *arXiv preprint arXiv:2410.01432*, 2024.
- Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- JooHwan Ko and Tomas Geffner. Latent target score matching, with an application to simulation-based inference. In *Machine Learning and the Physical Sciences Workshop, NeurIPS*, 2025.
- Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative learning for symmetric densities. In *International conference on machine learning*, pp. 5361–5370. PMLR, 2020.
- Peter Kollman, Richard Dixon, Wendy Cornell, Thomas Fox, Chris Chipot, and Andrew Pohorille. The development/application of a ‘minimalist’ organic/biochemical molecular mechanic force field using a combination of ab initio calculations and experimental data. In *Computer simulation of biomolecular systems: Theoretical and experimental applications*, pp. 83–96. Springer, 1997.
- Solomon Kullback. Probability densities with given marginals. *The Annals of Mathematical Statistics*, 39(4):1236–1243, 1968.
- Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.
- Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. Equiformerv2: Improved equivariant transformer for scaling to higher-degree representations. *arXiv preprint arXiv:2306.12059*, 2023.
- Guan-Hong Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos A Theodorou, and Ricky TQ Chen. Generalized Schrödinger bridge matching. *arXiv preprint arXiv:2310.02233*, 2023.
- Guan-Hong Liu, Jaemoon Choi, Yongxin Chen, Benjamin Kurt Miller, and Ricky TQ Chen. Adjoint Schrödinger Bridge Sampler. *arXiv preprint arXiv:2506.22565*, 2025.
- Xingchao Liu, Lemeng Wu, Mao Ye, and Qiang Liu. Let us build bridges: Understanding and extending diffusion generative models. *arXiv preprint arXiv:2208.14699*, 2022.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Dimitra Maoutsa, Sebastian Reich, and Manfred Opper. Interacting particle solutions of Fokker-Planck equations through gradient-log-density estimation. *Entropy*, 22(8):802, 2020.
- 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.
- Juno Nam, Bálint Máté, Artur P Toshev, Manasa Kaniselvan, Rafael Gómez-Bombarelli, Ricky TQ Chen, Brandon Wood, Guan-Hong Liu, and Benjamin Kurt Miller. Enhancing diffusion-based sampling with molecular collective variables. *arXiv preprint arXiv:2510.11923*, 2025.
- Radford M Neal. Probabilistic inference using Markov chain Monte Carlo methods. 1993.
- Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019.
- Nikolas Nüsken and Lorenz Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial differential equations and applications*, 2:1–48, 2021.

- 
- Alexey Onufriev, Donald Bashford, and David A Case. Exploring protein native states and large-scale conformational changes with a modified generalized born model. *Proteins: Structure, Function, and Bioinformatics*, 55(2):383–394, 2004.
- RuiKang OuYang, Bo Qiang, and José Miguel Hernández-Lobato. BNEM: A Boltzmann sampler based on bootstrapped noised energy matching. *arXiv preprint arXiv:2409.09787*, 2024.
- Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and trends® in Optimization*, 1(3):127–239, 2014.
- Stefano Peluchetti. Non-denoising forward-time diffusions. *arXiv preprint arXiv:2312.14589*, 2023.
- Lorenz Richter. *Solving high-dimensional PDEs, approximation of path space measures and importance sampling of diffusions*. PhD thesis, BTU Cottbus-Senftenberg, 2021.
- Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. In *International Conference on Learning Representations*, 2024.
- L. C. G. Rogers and David Williams. *Diffusions, Markov Processes and Martingales*. Cambridge Mathematical Library. Cambridge University Press, 2 edition, 2000.
- Sebastian Sanokowski, Lukas Gruber, Christoph Bartmann, Sepp Hochreiter, and Sebastian Lehner. Rethinking losses for diffusion bridge samplers. *arXiv preprint arXiv:2506.10982*, 2025.
- Henrik Schopmans and Pascal Friederich. Temperature-annealed Boltzmann generators. *arXiv preprint arXiv:2501.19077*, 2025.
- Mathias Schreiner, Ole Winther, and Simon Olsson. Implicit transfer operator learning: Multiple time-resolution models for molecular dynamics. *Advances in Neural Information Processing Systems*, 36:36449–36462, 2023.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897. PMLR, 2015.
- Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *International conference on machine learning*, pp. 9377–9388. PMLR, 2021.
- Marcin Sendera, Minsu Kim, Sarthak Mittal, Pablo Lemos, Luca Scimeca, Jarrid Rector-Brooks, Alexandre Adam, Yoshua Bengio, and Nikolay Malkin. Improved off-policy training of diffusion samplers. *Advances in Neural Information Processing Systems*, 37:81016–81045, 2024.
- Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion Schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36:62183–62223, 2023.
- Zhekun Shi, Longlin Yu, Tianyu Xie, and Cheng Zhang. Diffusion-PINN sampler. *arXiv preprint arXiv:2410.15336*, 2024.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- Charlie B Tan, Majdi Hassan, Leon Klein, Saifuddin Syed, Dominique Beaini, Michael M Bronstein, Alexander Tong, and Kirill Neklyudov. Amortized sampling with transferable normalizing flows. *arXiv preprint arXiv:2508.18175*, 2025.
- Surya T Tokdar and Robert E Kass. Importance sampling: a review. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(1):54–60, 2010.
- 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.
- Chenguang Wang, Xiaoyu Zhang, Kaiyuan Cui, Weichen Zhao, Yongtao Guan, and Tianshu Yu. Importance weighted score matching for diffusion samplers with enhanced mode coverage. *arXiv preprint arXiv:2505.19431*, 2025.
- Sangwoong Yoon, Himchan Hwang, Hyeokju Jeong, Dong Kyu Shin, Che-Sang Park, Sehee Kweon, and Frank Chongwoo Park. Value gradient sampler: Sampling as sequential decision making. *arXiv preprint arXiv:2502.13280*, 2025.
- James Matthew Young, Paula Cordero-Encinar, Sebastian Reich, Andrew Duncan, and O Deniz Akyildiz. Diffusion path samplers via sequential Monte Carlo. *arXiv preprint arXiv:2601.21951*, 2026.
- Dinghuai Zhang, Ricky TQ Chen, Cheng-Hao Liu, Aaron Courville, and Yoshua Bengio. Diffusion generative flow samplers: Improving learning signals through partial trajectory optimization. In *The Twelfth International Conference on Learning Representations*, 2024.
- Leo Zhang, Peter Potapchik, George Deligiannidis, Arnaud Doucet, Hai-Dang Dau, and Saifuddin Syed. Generalised parallel tempering: Flexible replica exchange via flows and diffusions. In *Frontiers in Probabilistic Inference: Learning meets Sampling*, 2025a.
- Leo Zhang, Peter Potapchik, Jiajun He, Yuanqi Du, Arnaud Doucet, Francisco Vargas, Hai-Dang Dau, and Saifuddin Syed. Accelerated parallel tempering via neural transports. *arXiv preprint arXiv:2502.10328*, 2025b.
- Qinsheng Zhang and Yongxin Chen. Path Integral Sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022.

---

APPENDIX

<b>A Assumptions and auxiliary results</b>	<b>15</b>
A.1 Additional notation . . . . .	15
A.2 Technical assumptions . . . . .	15
A.3 Useful identities . . . . .	15
<b>B Proofs</b>	<b>17</b>
B.1 Proofs for general couplings . . . . .	17
B.2 Proofs for Schrödinger bridges . . . . .	19
B.3 Proofs for Schrödinger half bridges . . . . .	24
B.4 Proofs for independent couplings . . . . .	26
B.5 Proofs for damped fixed-point iterates . . . . .	26
B.6 Proofs for control variates . . . . .	27
<b>C Related work</b>	<b>28</b>
<b>D Further algorithmic details</b>	<b>28</b>
D.1 Importance sampling and probability flow ordinary differential equation . . . . .	28
D.2 Variance reduction and control variates . . . . .	30
D.3 Loss weighting and numerical stability . . . . .	32
D.4 Reference process . . . . .	33
D.5 Detailed algorithmic description . . . . .	33
<b>E Experimental setup</b>	<b>33</b>
E.1 Gaussian mixture models . . . . .	33
E.2 $n$ -body identical particle systems . . . . .	35
E.3 Molecular benchmarks: Peptide systems . . . . .	36
<b>F Further numerical results</b>	<b>40</b>
F.1 Gaussian mixture models . . . . .	40
F.2 $n$ -body identical particle systems . . . . .	40

## A ASSUMPTIONS AND AUXILIARY RESULTS

### A.1 ADDITIONAL NOTATION

For vectors  $v_1, v_2 \in \mathbb{R}^d$ , we denote by  $\|v\|$  the Euclidean norm and by  $v_1 \cdot v_2$  the Euclidean inner product. For a real-valued matrix  $A$ , we denote by  $\text{Tr}(A)$  and  $A^\top$  its trace and transpose.

For a sufficiently smooth function  $f: \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}$ , we denote by  $\nabla f = \nabla_x f$  its gradient w.r.t. the spatial variables  $x$  and by  $\partial_t f$  and  $\partial_{x_i} f$  its partial derivatives w.r.t. the time coordinate  $t$  and the spatial coordinate  $x_i$ , respectively.

We denote by  $\mathcal{N}(\mu, \Sigma)$  a multivariate normal distribution with mean  $\mu \in \mathbb{R}^d$  and covariance matrix  $\Sigma \in \mathbb{R}^{d \times d}$ . For random variables  $X_1, X_2$ , we denote by  $\mathbb{E}[X_1]$  and  $\text{Var}[X_1]$  the expectation and variance of  $X_1$  and by  $\mathbb{E}[X_1|X_2]$  the conditional expectation of  $X_1$  given  $X_2$ .

### A.2 TECHNICAL ASSUMPTIONS

Throughout our work, we make the same assumptions as in Nüsken & Richter (2021); Domingo-Enrich et al. (2024); Shi et al. (2023), which are needed for all the objects considered to be well-defined and for the Propositions to hold. Namely, we assume that:

- (i) The set  $\mathcal{U}$  of *admissible controls* is given by

$$\mathcal{U} = \{u \in C^1(\mathbb{R}^d \times [0, T]; \mathbb{R}^d) \mid \exists C > 0, \forall (x, s) \in \mathbb{R}^d \times [0, T], u(x, s) \leq C(1 + \|x\|)\}. \quad (24)$$

- (ii) The diffusion coefficient  $\sigma(t)$  is a deterministic, continuously differentiable, and scalar-valued function of time, implying isotropic noise. To ensure the process is non-degenerate, we assume  $\sigma(t) > 0$  for all  $t \in [0, T]$ .
- (iii) For the target path measure  $\Pi^* = \Pi_{0,T}^* \mathbb{P}_{0,T} \in \mathcal{R}(\mathbb{P})$ , we assume

- (a) For all  $x_0 \in \mathbb{R}^d$ , we have  $\Pi_{T|0}^* \ll \mathbb{P}_{T|0}$ .

- (b) The Radon-Nikodym derivative  $r_T(x_T|x_0) = \frac{d\Pi_{T|0}^*}{d\mathbb{P}_{T|0}}(x_T)$  is bounded for all  $x_T \in \mathbb{R}^d$ .

- (c) The mapping  $(t, x) \mapsto \mathbb{E}_{\Pi^*}[\xi(X, t)|X_t = x]$  is locally Lipschitz and there exists  $C > 0$  and  $\psi \in C([0, T], \mathbb{R}_+)$  such that for any  $t \in [0, T]$  and  $x \in \mathbb{R}^d$ , we have

$$\|\mathbb{E}_{\Pi_{T|t}^*}[\nabla \log \mathbb{P}_{T|t}(X_T|X_t)|X_t = x]\| \leq C\psi(t)(1 + \|x\|). \quad (25)$$

- (iv) The target boundary coupling  $\Pi_{0,T}^*(x_0, x_T)$  is such that for any  $t \in (0, T)$ , the product  $\Pi_{0,T}^*(x_0, x_T)\mathbb{P}_{t|0,T}(x|x_0, x_T)$  and its first-order spatial derivatives vanish at the limit as  $\|x\| \rightarrow \infty$  or  $\|x_T\| \rightarrow \infty$ .

### A.3 USEFUL IDENTITIES

Throughout this section, we consider the following stochastic processes:

- (i) **Reference Process** ( $X \sim \mathbb{P}$ ): We fix it to a scaled Brownian motion:  $dX_t = \sigma(t)dB_t$  with  $X_0 \sim p_{\text{prior}}$ . We define the cumulative variance  $\kappa(t) = \int_0^t \sigma^2(s)ds$  and the relative variance  $\gamma(t) = \kappa(t)/\kappa(T)$ .
- (ii) **Controlled Process** ( $X \sim \mathbb{P}^u$ ): The Markovian diffusion defined by  $dX_t = \sigma(t)u(X_t, t)dt + \sigma(t)dB_t$ .
- (iii) **Path-Dependent Process** ( $X \sim \Pi^*$  or  $\Pi^i$ ): The non-Markovian process defined by  $dX_t = \sigma(t)\xi(X, t)dt + \sigma(t)dB_t$ .

**Theorem A.1** (Girsanov's theorem for path measures). *Let  $u, v \in \mathcal{U}$ . Then the Radon-Nikodym derivative between  $\mathbb{P}^u$  and  $\mathbb{P}^v$ , evaluated along  $X$ , is given by*

$$\log \frac{d\mathbb{P}^u}{d\mathbb{P}^v}(X) = \int_0^T \sigma^{-1}(u - v)(X_t, t) \cdot dX_t - \frac{1}{2} \int_0^T (\|u\|^2 - \|v\|^2)(X_t, t)dt. \quad (26)$$

*Proof.* See, e.g., Lemma A.1 in Nüsken & Richter (2021) or Appendix E in Vargas et al. (2024).  $\square$

**Proposition A.2** (Nelson's relation). *Let  $\Pi^* \in \mathcal{P}$  be the target path measure with marginal densities  $\Pi_t^*$ . Let  $u^*$  and  $v^*$  be the unique optimal forward and backward controls such that the process  $X$  driven by  $dX_t = \sigma(t)u^*(X_t, t)dt + \sigma(t)dB_t$  and the time-reversed process  $Y$  driven by  $dY_t = -\sigma(t)v^*(Y_t, t)dt + \sigma(t)d\bar{B}_t$  satisfy  $\text{Law}(X) = \text{Law}(Y) = \Pi^*$ . Then, the controls satisfy the following duality relation:*

$$u^*(x, t) + v^*(x, t) = \sigma(t)\nabla_x \log \Pi_t^*(x). \quad (27)$$

*Proof.* See e.g. Vargas et al. (2024) Proposition 2.1.  $\square$

**Lemma A.3** (Markovian Projection for path measures (Brunick & Shreve, 2013)). *Let  $\Pi^* \in \mathcal{P}$  be a (possibly non-Markovian) path measure induced by the path-dependent SDE in (9) with drift functional  $\xi(X, t)$ . Then, there exists a unique Markov measure  $\mathbb{P}^{u^*} \in \mathcal{M}$  induced by the Markovian SDE*

$$dX_t = \sigma(t)u^*(X_t, t)dt + \sigma(t)dB_t \quad (28)$$

with  $u^*(x, t) = \mathbb{E}_{\Pi^*}[\xi(X, t)|X_t = x]$ . The measure  $\mathbb{P}^{u^*}$  is the Markovian projection of  $\Pi^*$  in the sense that their time-marginals coincide for all  $t \in [0, T]$ :  $\mathbb{P}_t^{u^*} = \Pi_t^*$ .

*Proof.* See e.g., Corollary 3.7 in Brunick & Shreve (2013).  $\square$

**Lemma A.4** (Fokker-Planck Equation). *Let  $X = (X_t)_{t \in [0, T]}$  be a Markov process satisfying the SDE*

$$dX_t = \sigma(t)u(X_t, t)dt + \sigma(t)dB_t, \quad X_0 \sim \Pi_0^*, \quad (29)$$

where  $u \in \mathcal{U}$  is an admissible control and  $\sigma(t) \in C^1([0, T], \mathbb{R}_{>0})$  is a deterministic diffusion coefficient. Then, under Assumption (iv) the evolution of the marginal density  $\Pi_t^*$  is governed by the Fokker-Planck equation

$$\partial_t \Pi_t^*(x) = -\nabla \cdot (\Pi_t^*(x)\sigma(t)u(x, t)) + \frac{1}{2}\sigma(t)^2 \Delta \Pi_t^*(x). \quad (30)$$

*Proof.* See e.g., Berner et al. (2024) section 2.1.

**Lemma A.5** (Orthogonal Projection in  $L^2$ ). *Let  $\xi$  be a path-dependent random variable such that  $\mathbb{E}_{\Pi^i}[\xi^2] < \infty$ , and let  $\Phi(X_t) := \mathbb{E}_{\Pi^i}[\xi|X_t]$  be its Markovian projection. For any measurable Markovian function  $u(X_t, t)$ , the following decomposition holds:*

$$\mathbb{E}_{\Pi^i}[\|\xi(X, t) - u(X_t, t)\|^2] = \mathbb{E}_{\Pi^i}[\|\xi(X, t) - \Phi(X_t)\|^2] + \mathbb{E}_{\Pi^i}[\|\Phi(X_t) - u(X_t, t)\|^2]. \quad (31)$$

*Proof.* We begin by adding and subtracting the projection  $\Phi(X_t)$  inside the norm, i.e.

$$\mathbb{E}_{\Pi^i}[\|\xi(X, t) - u(X_t, t)\|^2] = \mathbb{E}_{\Pi^i}[\|(\xi(X, t) - \Phi(X_t)) + (\Phi(X_t) - u(X_t, t))\|^2]. \quad (32)$$

Expanding the square, we obtain

$$\begin{aligned} \mathbb{E}_{\Pi^i}[\|\xi(X, t) - u(X_t, t)\|^2] &= \mathbb{E}_{\Pi^i}[\|\xi(X, t) - \Phi(X_t)\|^2] + \mathbb{E}_{\Pi^i}[\|\Phi(X_t) - u(X_t, t)\|^2] \\ &\quad + 2\mathbb{E}_{\Pi^i}[\langle \xi(X, t) - \Phi(X_t), \Phi(X_t) - u(X_t, t) \rangle]. \end{aligned} \quad (33)$$

We now show that the cross-term vanishes using the tower property. Let  $g(X_t) := \Phi(X_t) - u(X_t, t)$ , which is a function solely of the state at time  $t$ :

$$\begin{aligned} \mathbb{E}_{\Pi^i}[\langle \xi(X, t) - \Phi(X_t), g(X_t) \rangle] &= \mathbb{E}_{\Pi^i}[\mathbb{E}_{\Pi^i}[\langle \xi(X, t) - \Phi(X_t), g(X_t) \rangle | X_t]] \\ &= \mathbb{E}_{\Pi^i}[\langle \mathbb{E}_{\Pi^i}[\xi(X, t) | X_t] - \Phi(X_t), g(X_t) \rangle]. \end{aligned} \quad (34)$$

By the definition of the projection,  $\Phi(X_t) = \mathbb{E}_{\Pi^i}[\xi(X, t) | X_t]$ . Therefore, the inner term is

$$\mathbb{E}_{\Pi^i}[\langle \Phi(X_t) - \Phi(X_t), g(X_t) \rangle] = 0. \quad (35)$$

The cross-term vanishes, leaving only the sum of the squared distances, which concludes the proof.  $\square$

**Lemma A.6.** Let  $\Pi^*, \mathbb{P} \in \mathcal{P}$  such that  $\Pi^* \ll \mathbb{P}$ . Let  $X$  denote the path and  $X_t$  the process value at time  $t$ . Let  $\Pi_t^*$  and  $\mathbb{P}_t$  denote the marginal distributions at time  $t$ . Then,  $\mathbb{P}_t$ -almost surely

$$\frac{d\Pi_t^*}{d\mathbb{P}_t}(x) = \mathbb{E}_{\mathbb{P}} \left[ \frac{d\Pi^*}{d\mathbb{P}}(X) \middle| X_t = x \right]. \quad (36)$$

*Proof.* Let  $h : \mathbb{R}^d \rightarrow \mathbb{R}$  be an arbitrary bounded measurable test function. We compute the expectation  $\mathbb{E}_{\Pi^*}[h(X_t)]$  in two distinct ways.

1. Via the marginal measure  $\Pi_t^*$ : Using the definition of the Radon-Nikodym derivative on the marginal space  $\mathbb{R}^d$ , we have

$$\mathbb{E}_{\Pi^*}[h(X_t)] = \int_{\mathbb{R}^d} h(x) d\Pi_t^*(x) = \int_{\mathbb{R}^d} h(x) \frac{d\Pi_t^*}{d\mathbb{P}_t}(x) d\mathbb{P}_t(x). \quad (37)$$

2. Via the path measure  $\Pi^*$ : Using the Radon-Nikodym derivative  $L = \frac{d\Pi^*}{d\mathbb{P}}$  on the full path space and the tower property of conditional expectation, we get

$$\mathbb{E}_{\Pi^*}[h(X_t)] = \mathbb{E}_{\mathbb{P}} \left[ h(X_t) \frac{d\Pi^*}{d\mathbb{P}}(X) \right] \quad (38a)$$

$$= \mathbb{E}_{\mathbb{P}} \left[ \mathbb{E}_{\mathbb{P}} \left[ h(X_t) \frac{d\Pi^*}{d\mathbb{P}}(X) \middle| X_t \right] \right] \quad (38b)$$

$$= \mathbb{E}_{\mathbb{P}} \left[ h(X_t) \mathbb{E}_{\mathbb{P}} \left[ \frac{d\Pi^*}{d\mathbb{P}}(X) \middle| X_t \right] \right]. \quad (38c)$$

Note that we pulled  $h(X_t)$  out of the inner expectation because it is  $\sigma(X_t)$ -measurable. We can rewrite (38) as an integral over the marginal measure  $\mathbb{P}_t$ , i.e.

$$\mathbb{E}_{\Pi^*}[h(X_t)] = \int_{\mathbb{R}^d} h(x) \mathbb{E}_{\mathbb{P}} \left[ \frac{d\Pi^*}{d\mathbb{P}}(X) \middle| X_t = x \right] d\mathbb{P}_t(x). \quad (39)$$

Comparing (37) and (39), we have

$$\int_{\mathbb{R}^d} h(x) \frac{d\Pi_t^*}{d\mathbb{P}_t}(x) d\mathbb{P}_t(x) = \int_{\mathbb{R}^d} h(x) \mathbb{E}_{\mathbb{P}} \left[ \frac{d\Pi^*}{d\mathbb{P}}(X) \middle| X_t = x \right] d\mathbb{P}_t(x). \quad (40)$$

Since this equality holds for any bounded measurable  $h$ , the integrands must be equal  $\mathbb{P}_t$ -almost everywhere.  $\square$

## B PROOFS

In this section we state the proofs of the statements from the main part of the paper and add some additional statements as well.

### B.1 PROOFS FOR GENERAL COUPLINGS

**Proposition B.1** (Denoising score identity). *For any path measure  $\Pi^* \in \mathcal{R}(\mathbb{P})$  with coupling  $\Pi_{0,T}^*$ , the score of the marginal density  $\Pi_t^*$  satisfies the Denoising Score Identity (DSI)*

$$\nabla_x \log \Pi_t^*(x) = \mathbb{E}_{\Pi_{0,T}^*} \left[ \nabla_{X_t} \log \mathbb{P}_{t|0,T}(X_t | X_0, X_T) \middle| X_t = x \right]. \quad (41)$$

*Proof.* We consider the marginal density  $\Pi_t^*(x)$ , which for a general coupling  $\Pi_{0,T}^*(x_0, x_T)$ , is given by

$$\Pi_t^*(x) = \int_{\mathbb{R}^{2d}} \mathbb{P}_{t|0,T}(x | x_0, x_T) \Pi_{0,T}^*(x_0, x_T) dx_0 dx_T. \quad (42)$$

We differentiate  $\Pi_t^*(x)$  with respect to  $x$ . With Assumption (iii)b, we can interchange differentiation and integration, such that the gradient operates directly on the transition kernel  $\mathbb{P}_{t|0,T}$ :

$$\nabla_x \Pi_t^*(x) = \int_{\mathbb{R}^{2d}} \nabla_x \mathbb{P}_{t|0,T}(x | x_0, x_T) \Pi_{0,T}^*(x_0, x_T) dx_0 dx_T. \quad (43)$$

Using the log-derivative identity  $\nabla_x \mathbb{P}_{t|0,T} = \mathbb{P}_{t|0,T} \nabla_x \log \mathbb{P}_{t|0,T}$ , we rewrite the gradient as

$$\nabla_x \Pi_t^*(x) = \int_{\mathbb{R}^{2d}} (\nabla_x \log \mathbb{P}_{t|0,T}(x|x_0, x_T)) \mathbb{P}_{t|0,T}(x|x_0, x_T) \Pi_{0,T}^*(x_0, x_T) dx_0 dx_T. \quad (44)$$

To recover the score  $\nabla_x \log \Pi_t^*(x)$ , we divide both sides by the marginal density  $\Pi_t^*(x)$ . By Bayes' theorem, the posterior density of the boundaries given the current state is given as

$$\Pi_{0,T|t}^*(x_0, x_T|x) = \frac{\mathbb{P}_{t|0,T}(x|x_0, x_T) \Pi_{0,T}^*(x_0, x_T)}{\Pi_t^*(x)}. \quad (45)$$

Substituting this identity into the integral expression (44), we get:

$$\nabla_x \log \Pi_t^*(x) = \int_{\mathbb{R}^{2d}} \nabla_x \log \mathbb{P}_{t|0,T}(x|x_0, x_T) \Pi_{0,T|t}^*(x_0, x_T|x) dx_0 dx_T \quad (46a)$$

$$= \mathbb{E}_{\Pi_{0,T|t}^*} [\nabla_{X_t} \log \mathbb{P}_{t|0,T}(X_t|X_0, X_T)|X_t = x]. \quad (46b)$$

This concludes the proof.  $\square$

*Proof of Theorem 2.4.* Applying Doob's  $h$ -transforms Rogers & Williams (2000), conditioning the base process on a terminal state  $X_T$  yields the (path-dependent) drift  $\sigma(t) \nabla_{X_t} \log \mathbb{P}_{T|t}(X_T|X_t)$ . Using the Markovian projection, we identify the optimal forward control

$$u^*(x, t) = \sigma(t) \mathbb{E}_{\Pi_{T|t}^*} [\nabla_{X_t} \log \mathbb{P}_{T|t}(X_T|X_t)|X_t = x]. \quad (47)$$

Further, by Nelson's identity, it holds that

$$u^*(x, t) + v^*(x, t) = \sigma(t) \nabla_x \log \Pi_t^*(x). \quad (48)$$

We apply Theorem B.1 for  $\nabla_x \log \Pi_t^*$ :

$$\nabla_x \log \Pi_t^*(x) = \mathbb{E}_{\Pi_{0,T|t}^*} [\nabla_{X_t} \log \mathbb{P}_{t|0,T}(X_t|X_0, X_T)|X_t = x]. \quad (49)$$

For the Markovian reference process  $\mathbb{P}$ , the bridge transition density factorizes as  $\mathbb{P}_{t|0,T}(x|x_0, x_T) \propto \mathbb{P}_{t|0}(x|x_0) \mathbb{P}_{T|t}(x_T|x)$ . Taking the gradient of the log-density yields

$$\nabla_x \log \mathbb{P}_{t|0,T}(x|X_0, X_T) = \nabla_x \log \mathbb{P}_{t|0}(x|X_0) + \nabla_x \log \mathbb{P}_{T|t}(X_T|x), \quad (50)$$

and substituting into (49), we obtain

$$\begin{aligned} \sigma(t) \nabla_x \log \Pi_t^*(x) &= \sigma(t) \mathbb{E}_{\Pi_{0|t}^*} [\nabla_{X_t} \log P_{t|0}(X_t|X_0)|X_t = x] \\ &\quad + \sigma(t) \mathbb{E}_{\Pi_{T|t}^*} [\nabla_{X_t} \log P_{T|t}(X_T|X_t)|X_t = x] \end{aligned} \quad (51a)$$

$$= \sigma(t) \mathbb{E}_{\Pi_{0|t}^*} [\nabla_{X_t} \log P_{t|0}(X_t|X_0)|X_t = x] + u^*(x, t) \quad (51b)$$

By Nelson relation, we get the expression for the backward drift  $v^*(x, t)$ .  $\square$

*Proof of Theorem 2.5.* By Theorem B.1, the marginal density  $\Pi_t^*(x) = \int_{\mathbb{R}^{2d}} \Pi_{0,T}^*(x_0, x_T) \mathbb{P}_{t|0,T}(x|x_0, x_T) dx_0 dx_T$  is given by

$$\nabla_x \log \Pi_t^*(x) = \mathbb{E}_{\Pi_{0,T|t}^*} [\nabla_x \log \mathbb{P}_{t|0,T}(X_t|X_0, X_T)|X_t = x]. \quad (52)$$

For the reference process  $dX_t = \sigma(t) dB_t$ , the transition density is Gaussian, namely

$$\mathbb{P}_{t|0,T}(X_t|X_0, X_T) = \mathcal{N}(X_t; (1 - \gamma(t))X_0 + \gamma(t)X_T, \kappa(T)\gamma(t)(1 - \gamma(t))I), \quad (53)$$

see also Appendix D.4. We can therefore compute

$$\nabla_{X_T} \log P_{t|0,T}(X_t|X_0, X_T) = \frac{1}{1 - \gamma(t)} \nabla_{X_0} \log P_{t|0,T}(X_t|X_0, X_T) \quad (54)$$

$$= \frac{1}{\gamma(t)} \nabla_{X_T} \log P_{t|0,T}(X_t|X_0, X_T) \quad (55)$$

We integrate by parts and apply Assumption (iv) to eliminate the boundary terms:

$$\begin{aligned} & \int_{\mathbb{R}^{2d}} \nabla_{x_0} P_{t|0,T}(x_t|x_0, x_T) \Pi_{0,T}^*(x_0, x_T) dx_0 dx_T \\ &= - \int_{\mathbb{R}^{2d}} P_{t|0,T}(x_t|x_0, x_T) \nabla_{x_0} \Pi_{0,T}^*(x_0, x_T) dx_0 dx_T \end{aligned} \quad (56)$$

Using the log-derivative identity  $\nabla_{x_0} \Pi_{0,T}^*(x_0, x_T) = \Pi_{0,T}^*(x_0, x_T) \nabla_{x_0} \log \Pi_{0,T}^*(x_0, x_T)$ , we get

$$\mathbb{E}_{\Pi_{0,T|t}^*} [\nabla_{X_t} \log \mathbb{P}_{t|0,T}(X_t|X_0, X_T) | X_t = x] \quad (57)$$

$$= \mathbb{E}_{\Pi_{0,T|t}^*} \left[ \frac{1}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) | X_t = x \right]. \quad (58)$$

By symmetry, applying this to  $X_T$  yields

$$\mathbb{E}_{\Pi_{0,T|t}^*} [\nabla_{X_t} \log \mathbb{P}_{t|0,T}(X_t|X_0, X_T) | X_t = x] = \mathbb{E}_{\Pi_{0,T|t}^*} \left[ \frac{1}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) | X_t = x \right]. \quad (59)$$

Since both (57) and (59) equal the same marginal score  $\nabla_x \log \Pi_t^*$ , any convex combination is also an unbiased estimator, which concludes the proof.  $\square$

*Proof of Theorem 2.6.* We start with Nelson's identity (Theorem A.2)

$$u^*(x, t) = -v^*(x, t) + \sigma(t) \nabla_x \log \Pi_t^*(x). \quad (60)$$

By Theorem 2.4 the backward drift is given by

$$v^*(x, t) = \sigma(t) \mathbb{E}_{\Pi_{0,t}^*} [\nabla_{X_t} \log \mathbb{P}_{t|0}(X_t|X_0) | X_t = x]. \quad (61)$$

Next, we substitute the marginal score  $\nabla_x \log \Pi_t^*(x)$  using the generalized TSI (Theorem 2.5), namely

$$\nabla_x \log \Pi_t^*(x) = \mathbb{E}_{\Pi_{0,T|t}^*} \left[ \frac{1 - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) \middle| X_t = x \right]. \quad (62)$$

Substituting the expressions for  $v^*$  and  $\nabla \log \Pi_t^*(x)$  back into Nelson's identity, we obtain

$$\begin{aligned} u^*(x, t) &= -\sigma(t) \mathbb{E}_{\Pi^*} [\nabla \log P_{t|0}(X_t|X_0) | X_t = x] \\ &+ \sigma(t) \mathbb{E}_{\Pi^*} \left[ \frac{1 - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) \right. \\ &\quad \left. + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) \middle| X_t = x \right] \end{aligned} \quad (63a)$$

$$\begin{aligned} &= \sigma(t) \mathbb{E}_{\Pi^*} \left[ -\nabla \log P_{t|0}(X_t|X_0) + \frac{1 - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) \right. \\ &\quad \left. + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) \middle| X_t = x \right]. \end{aligned} \quad (63b)$$

Identifying the term inside the expectation as  $\xi(X, t)$ , we arrive at the result.  $\square$

## B.2 PROOFS FOR SCHRÖDINGER BRIDGES

We consider the Schrödinger Bridge (SB) problem with reference measure  $\mathbb{P}$  induced by the SDE  $dX_t = \sigma(t) dB_t$ . This problem is defined as the entropy-regularized optimal transport problem

$$\min_{\mathbb{P}^u \in \mathcal{P}} D_{\text{KL}}(\mathbb{P}^u | \mathbb{P}) \quad \text{subject to} \quad \mathbb{P}_0^u = \Pi_0^*, \mathbb{P}_T^u = \Pi_T^*. \quad (64)$$

This is equivalent to the stochastic optimal control problem

$$\min_{u \in \mathcal{U}} \mathbb{E} \left[ \int_0^T \frac{1}{2} \|u(X_t, t)\|^2 dt \right] \text{ s.t. } \begin{cases} X_0 \sim \Pi_0^* = p_{\text{prior}}, \\ X_T \sim \Pi_T^* = p_{\text{target}}, \end{cases} \quad (65)$$

under the controlled dynamics

$$dX_t = \sigma(t)u(X_t, t)dt + \sigma(t)dB_t. \quad (66)$$

The following derivations of the Schrödinger system and its associated optimal controls are standard in the literature, and we provide them here primarily for the reader's convenience.

Using Girsanov's theorem and the Fokker planck equation, the evolution of the marginal density  $\Pi_t^*(x)$  can be enforced via the Lagrange multiplier  $\psi(x, t)$ :

$$\begin{aligned} \mathcal{L}(\Pi^*, u, \psi) = & \int_0^T \int_{\mathbb{R}^d} \left( \frac{1}{2} \|u(x, t)\|^2 \Pi_t^*(x) \right. \\ & + \psi(x, t) [\partial_t \Pi_t^*(x) + \sigma(t) \nabla \cdot (\Pi_t^*(x) u(x, t)) \\ & \left. - \frac{1}{2} \sigma^2(t) \text{Tr}(\nabla^2 \Pi_t^*(x)) \right] dx dt \end{aligned} \quad (67)$$

**Lemma B.2.** *The optimal marginal density  $\Pi_t^*$  and the optimal control  $u^*$  satisfy:*

$$\begin{cases} \partial_t \psi + \frac{1}{2} \sigma^2 \text{Tr}(\nabla^2 \psi) + \frac{1}{2} \sigma^2 \|\nabla \psi\|^2 = 0, \\ \partial_t \Pi_t^* + \sigma \nabla \cdot (\Pi_t^* u^*) = \frac{1}{2} \sigma^2 \text{Tr}(\nabla^2 \Pi_t^*), \end{cases} \quad (68)$$

where the optimal control is given by  $u^* = \sigma \nabla \psi_t$ .

*Proof.* Using integration by parts (where boundary terms vanish by Assumption (iv)), we can rewrite the term containing the divergence inside the Lagrangian:

$$\int_{\mathbb{R}^d} \psi(x, t) \sigma(t) \nabla \cdot (\Pi_t^*(x) u(x, t)) dx = - \int_{\mathbb{R}^d} \sigma(t) \Pi_t^*(x) u(x, t) \cdot \nabla \psi(x, t) dx. \quad (69)$$

The terms in the Lagrangian involving the control  $u$  are

$$\int_0^T \int_{\mathbb{R}^d} \Pi_t^*(x) \left( \frac{1}{2} \|u(x, t)\|^2 - \sigma(t) u(x, t) \cdot \nabla \psi(x, t) \right) dx dt. \quad (70)$$

Taking the functional derivative with respect to  $u$  and setting it to zero, we get

$$\frac{\delta \mathcal{L}}{\delta u} = \Pi_t^*(u - \sigma \nabla \psi) = 0. \quad (71)$$

Hence, we obtain the optimal control

$$u^* = \sigma \nabla \psi. \quad (72)$$

Substituting  $u^*$  into the Lagrangian yields

$$\begin{aligned} \mathcal{L} = & \int_0^T \int_{\mathbb{R}^d} \left( \frac{1}{2} \sigma^2(t) \|\nabla \psi(x, t)\|^2 \Pi_t^*(x) + \psi(x, t) \partial_t \Pi_t^*(x) \right. \\ & \left. + \sigma^2(t) \psi(x, t) \nabla \cdot (\Pi_t^*(x) \nabla \psi(x, t)) - \frac{1}{2} \sigma^2(t) \psi(x, t) \text{Tr}(\nabla^2 \Pi_t^*(x)) \right) dx dt. \end{aligned} \quad (73)$$

We isolate  $\Pi_t^*$  in every term. For that, we use integration by parts on the last three terms:

$$\int_0^T \int_{\mathbb{R}^d} \psi(x, t) \partial_t \Pi_t^*(x) dx = - \int_0^T \int_{\mathbb{R}^d} \Pi_t^*(x) \partial_t \psi(x, t) dx, \quad (74)$$

$$\int_{\mathbb{R}^d} \sigma^2(t) \psi(x, t) \nabla \cdot (\Pi_t^*(x) \nabla \psi(x, t)) dx = - \int_{\mathbb{R}^d} \sigma^2(t) \Pi_t^*(x) (\nabla \psi(x, t)) \cdot \nabla \psi(x, t) dx \quad (75)$$

$$= - \int_{\mathbb{R}^d} \sigma^2(t) \Pi_t^*(x) \|\nabla \psi(x, t)\|^2 dx, \quad (76)$$

$$\int_{\mathbb{R}^d} \sigma^2(t) \psi(x, t) \text{Tr}(\nabla^2 \Pi_t^*(x)) dx = \int_{\mathbb{R}^d} \sigma^2(t) \Pi_t^*(x) \text{Tr}(\nabla^2 \psi(x, t)) dx, \quad (77)$$

where the boundary terms vanish by Assumption (iv). Substituting these back into the Lagrangian gives

$$\begin{aligned} \mathcal{L} = & \int_0^T \int_{\mathbb{R}^d} \Pi_t^*(x) \left( \frac{1}{2} \sigma^2(t) \|\nabla \psi(x, t)\|^2 - \partial_t \psi(x, t) \right. \\ & \left. - \sigma^2(t) \|\nabla \psi(x, t)\|^2 - \frac{1}{2} \sigma^2(t) \operatorname{Tr}(\nabla^2 \psi(x, t)) \right) dx dt \end{aligned} \quad (78a)$$

$$\begin{aligned} = & \int_0^T \int_{\mathbb{R}^d} \Pi_t^*(x) \left( -\partial_t \psi(x, t) - \frac{1}{2} \sigma^2(t) \|\nabla \psi(x, t)\|^2 \right. \\ & \left. - \frac{1}{2} \sigma^2(t) \operatorname{Tr}(\nabla^2 \psi(x, t)) \right) dx dt. \end{aligned} \quad (78b)$$

Taking the functional derivative with respect to  $\Pi_t^*$  and setting it to zero gives the first PDE in (68) (namely the Hamilton-Jacobi-Bellman equation). If we substitute the optimal control  $u^* = \sigma \nabla \psi$  into the Fokker-Planck equation of the general controlled process, we obtain the second PDE:

$$\partial_t \Pi_t^* + \sigma \nabla \cdot (\Pi_t^* u^*) = \frac{1}{2} \sigma^2 \operatorname{Tr}(\nabla^2 \Pi_t^*). \quad (79)$$

□

**Lemma B.3.** *The nonlinear system of PDEs in (68) can be linearized via the Hopf-Cole transformation to satisfy the linear system*

$$\begin{cases} \partial_t \varphi = -\frac{1}{2} \sigma^2 \Delta \varphi, \\ \partial_t \widehat{\varphi} = \frac{1}{2} \sigma^2 \Delta \widehat{\varphi}, \end{cases} \quad (80)$$

subject to the coupled boundary conditions

$$\begin{cases} \varphi_0 \widehat{\varphi}_0 = \Pi_0^*, \\ \varphi_T \widehat{\varphi}_T = \Pi_T^*. \end{cases} \quad (81)$$

*Proof.* Define the mapping  $(\Pi_t^*, \psi) \mapsto (\varphi, \widehat{\varphi})$  with

$$\varphi_t = \exp(\psi_t), \quad (82)$$

$$\widehat{\varphi}_t = \Pi_t^* \exp(-\psi_t) = \frac{\Pi_t^*}{\varphi_t}. \quad (83)$$

From  $\psi = \log \varphi$ , we calculate the derivatives

$$\nabla \psi = \frac{\nabla \varphi}{\varphi}, \quad \Delta \psi = \frac{\Delta \varphi}{\varphi} - \frac{\|\nabla \varphi\|^2}{\varphi^2}, \quad \partial_t \psi = \frac{\partial_t \varphi}{\varphi}. \quad (84)$$

Substituting these into the second PDE from (68) yields

$$\partial_t \psi + \frac{1}{2} \sigma^2 \operatorname{Tr}(\nabla^2 \psi) + \frac{1}{2} \sigma^2 \|\nabla \psi\|^2 = 0 \quad (85a)$$

$$\Leftrightarrow \frac{\partial_t \varphi}{\varphi} + \frac{1}{2} \sigma^2 \left( \frac{\Delta \varphi}{\varphi} - \frac{\|\nabla \varphi\|^2}{\varphi^2} \right) + \frac{1}{2} \sigma^2 \frac{\|\nabla \varphi\|^2}{\varphi^2} = 0. \quad (85b)$$

Multiplying by  $\varphi$ , we obtain

$$\partial_t \varphi + \frac{1}{2} \sigma^2 \Delta \varphi = 0, \quad (86)$$

which is the first PDE in (81). For  $\widehat{\varphi}$ , we use the product relation  $\Pi_t^* = \varphi \widehat{\varphi}$  and substitute it into the Fokker-Planck equation of the optimal density:

$$\partial_t (\varphi \widehat{\varphi}) + \sigma^2 \nabla \cdot (\varphi \widehat{\varphi} \nabla \log \varphi) = \frac{1}{2} \sigma^2 \Delta (\varphi \widehat{\varphi}). \quad (87)$$

Applying the product rule and using  $u^* = \sigma^2 \frac{\nabla \varphi}{\varphi}$ , we get

$$\widehat{\varphi} \partial_t \varphi + \varphi \partial_t \widehat{\varphi} + \sigma^2 \nabla \cdot (\widehat{\varphi} \nabla \varphi) = \frac{1}{2} \sigma^2 (\widehat{\varphi} \Delta \varphi + 2 \nabla \varphi \cdot \nabla \widehat{\varphi} + \varphi \Delta \widehat{\varphi}). \quad (88)$$

Expanding the divergence term  $\nabla \cdot (\widehat{\varphi} \nabla \varphi) = \nabla \widehat{\varphi} \cdot \nabla \varphi + \widehat{\varphi} \Delta \varphi$  and substituting  $\partial_t \varphi = -\frac{1}{2} \sigma^2 \Delta \varphi$ , the terms involving  $\Delta \varphi$  and  $\nabla \varphi \cdot \nabla \widehat{\varphi}$  cancel out, leaving

$$\varphi \partial_t \widehat{\varphi} = \frac{1}{2} \sigma^2 \varphi \Delta \widehat{\varphi}. \quad (89)$$

Dividing by  $\varphi$  yields the second PDE. The boundary constraints follow since  $\Pi_t^* = \varphi \widehat{\varphi}$ . □

*Proof of Theorem 2.8.* By Theorem B.2, the optimal control is  $u^*(x, t) = \sigma(t)\nabla\psi_t(x)$ . Using the Hopf-Cole transformation  $\psi = \log \varphi$ , taking the gradient yields  $u^* = \sigma\nabla \log \varphi$  and by symmetry,  $v^* = \sigma\nabla \log \widehat{\varphi}$ . We further have the following equations for  $\varphi$  and  $\widehat{\varphi}$ :

$$\partial_t \varphi + \frac{1}{2}\sigma \Delta \varphi = 0, \quad (90)$$

$$\partial_t \widehat{\varphi} - \frac{1}{2}\sigma \Delta \widehat{\varphi} = 0. \quad (91)$$

By the Feynman-Kac formula, the solution to the PDE can be represented as the conditional expectation of the terminal value,

$$\varphi_t(x) = \mathbb{E}_{\mathbb{P}}[\varphi_T(X_T)|X_t = x] = \int_{\mathbb{R}^d} \varphi_T(x_T) \mathbb{P}_{T|t}(x_T|x) dx_T, \quad (92)$$

and similarly for the forward equation,

$$\widehat{\varphi}_t(x) = \int_{\mathbb{R}^d} \mathbb{P}_{t|0}(x|x_0) \widehat{\varphi}_0(x_0) dx_0. \quad (93)$$

It remains to show the factorization of the optimal coupling  $\Pi_{0,T}^*$ . Conditioned on the initial state  $X_0$ , the optimal conditional path measure  $\Pi_{|0}^*$  is related to the reference path measure  $\mathbb{P}_{|0}$  via the Radon-Nikodym derivative Léonard (2013):

$$\frac{d\Pi_{|0}^*}{d\mathbb{P}_{|0}}(X) = \frac{\varphi_T(X_T)}{\varphi_0(X_0)}. \quad (94)$$

By the Markov property of the reference process, the conditional transition  $\Pi_{T|0}^*$  is the reference transition  $\mathbb{P}_{T|0}$  weighted by the potential ratio:

$$\Pi_{T|0}^*(x_T|x_0) = \mathbb{P}_{T|0}(x_T|x_0) \frac{\varphi_T(x_T)}{\varphi_0(x_0)} \quad (95)$$

The joint distribution  $\Pi_{0,T}^*$  is the product of this transition density and the initial marginal  $\Pi_0^*$ . Using the boundary condition  $\Pi_0^* = \varphi_0 \widehat{\varphi}_0$  from the Schrödinger system, we can compute

$$\Pi_{0,T}^*(x_0, x_T) = \Pi_0^*(x_0) \Pi_{T|0}^*(x_T|x_0) \quad (96a)$$

$$= \varphi_0(x_0) \widehat{\varphi}_0(x_0) \mathbb{P}_{T|0}(x_T|x_0) \frac{\varphi_T(x_T)}{\varphi_0(x_0)} \quad (96b)$$

$$= \widehat{\varphi}_0(x_0) \mathbb{P}_{T|0}(x_T|x_0) \varphi_T(x_T), \quad (96c)$$

which concludes the statement.  $\square$

**Proposition B.4.** *Consider the Schrödinger Bridge problem associated with the reference process  $dX_t = \sigma(t)dB_t$ . Let  $\Pi^* \in \mathcal{R}(\mathbb{P})$  be the unique bridge measure with marginals  $\Pi_0^*$  and  $\Pi_T^*$ . The optimal forward drift  $u^*$  and backward drift  $v^*$  are characterized by the Schrödinger potentials  $(\varphi, \widehat{\varphi})$  and satisfy the identities*

$$u^*(x, t) = \sigma(t) \mathbb{E}_{\Pi_{T|t}^*} \left[ \nabla_{X_T} \log \frac{\Pi_T^*(X_T)}{\widehat{\varphi}_T(X_T)} \middle| X_t = x \right], \quad (97)$$

$$v^*(x, t) = \sigma(t) \mathbb{E}_{\Pi_{0|t}^*} \left[ \nabla_{X_0} \log \frac{\Pi_0^*(X_0)}{\varphi_0(X_0)} \middle| X_t = x \right]. \quad (98)$$

*Proof.* For the optimal forward control  $u^*$  holds that

$$u^*(x, t) = \sigma(t) \nabla_x \log \varphi_t(x) = \sigma(t) \frac{\nabla_x \varphi_t(x)}{\varphi_t(x)}. \quad (99)$$

The potential  $\varphi_t(x)$  evolves backward in time, yielding

$$\varphi_t(x) = \int_{\mathbb{R}^d} \mathbb{P}_{T|t}(x_T|x) \varphi_T(x_T) dx_T. \quad (100)$$

For the reference process  $dX_t = \sigma(t)dB_t$ , the transition density  $\mathbb{P}_{T|t}(x_T|x)$  depends only on the difference  $(x_T - x)$ . Explicitly,  $\mathbb{P}_{T|t}(x_T|x) = \mathcal{N}(x_T; x, \Sigma_{T|t})$ . This implies the identity

$$\nabla_x \mathbb{P}_{T|t}(x_T|x) = -\nabla_{x_T} \mathbb{P}_{T|t}(x_T|x). \quad (101)$$

We differentiate the integral form of  $\varphi$  with respect to  $x$ :

$$\nabla_x \varphi_t(x) = \int_{\mathbb{R}^d} \nabla_x \mathbb{P}_{T|t}(x_T|x) \varphi_T(x_T) dx_T \quad (102a)$$

$$= \int_{\mathbb{R}^d} (-\nabla_{x_T} \mathbb{P}_{T|t}(x_T|x)) \varphi_T(x_T) dx_T \quad (102b)$$

$$= \int_{\mathbb{R}^d} \mathbb{P}_{T|t}(x_T|x) \nabla_{x_T} \varphi_T(x_T) dx_T. \quad (102c)$$

Use the log-derivative trick  $\nabla \varphi = \varphi \nabla \log \varphi$ , we get

$$\nabla_x \varphi_t(x) = \int_{\mathbb{R}^d} \mathbb{P}_{T|t}(x_T|x) \varphi_T(x_T) \nabla_{x_T} \log \varphi_T(x_T) dx_T. \quad (103)$$

We substitute this back into the expression for  $u^*$  and divide by  $\varphi_t(x)$ . We identify the bridge density  $\Pi_{T|t}^*$  using the reciprocal relation for Schrödinger Bridges:

$$\Pi_{T|t}^*(x_T|x) = \frac{\mathbb{P}_{T|t}(x_T|x) \varphi_T(x_T)}{\varphi_t(x)}. \quad (104)$$

Thus we get

$$u^*(x, t) = \sigma(t) \int_{\mathbb{R}^d} \underbrace{\frac{\mathbb{P}_{T|t}(x_T|x) \varphi_T(x_T)}{\varphi_t(x)}}_{=\Pi_{T|t}^*(x_T|x)} \nabla_{x_T} \log \varphi_T(x_T) dx_T \quad (105a)$$

$$= \sigma(t) \mathbb{E}_{\Pi_{T|t}^*} [\nabla_{X_T} \log \varphi_T(X_T) | X_t = x]. \quad (105b)$$

At time  $T$ , the Schrödinger system boundary condition is  $\Pi_T^*(x) = \varphi_T(x) \widehat{\varphi}_T(x)$ . Solving for  $\varphi_T(x)$ , we get

$$\varphi_T(x) = \frac{\Pi_T^*(x)}{\widehat{\varphi}_T(x)}. \quad (106)$$

Substituting this into the expectation yields the result

$$u^*(x, t) = \sigma(t) \mathbb{E}_{\Pi_{T|t}^*} \left[ \nabla \log \frac{\Pi_T(X_T)}{\widehat{\varphi}_T(X_T)} \Big| X_t = x \right]. \quad (107)$$

The proof for  $v^*$  is analogous but proceeds forward in time.  $\square$

**Proposition B.5** (Path-dependent drift for Schrödinger bridges). *Under the assumptions of Theorem 2.5, the path-dependent drift  $\xi$  corresponding to the non-Markovian SDE (9) that induces the target measure  $\Pi^* = \Pi_{0,T}^{\text{SB}} \mathbb{P}_{|0,T}$  is given by*

$$\begin{aligned} \sigma(t)^{-1} \xi(X, t) &= \frac{\gamma(t) - c(t)}{\gamma(t)(1 - \gamma(t))} \nabla_{X_0} \log P_{T|0}(X_T|X_0) \\ &\quad + \frac{\gamma(t) - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \frac{\Pi_0^*(X_0)}{p_0(X_0)} \\ &\quad + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \frac{\Pi_T^*(X_T)}{p_T(X_T)} \end{aligned} \quad (108)$$

*Proof.* Let us start from the general identity for  $\xi$ , stated in Theorem 2.6,

$$\begin{aligned} \sigma^{-1}(t) \xi(X, t) &= -\nabla_{X_t} \log P_{t|0}(X_t|X_0) \\ &\quad + \frac{1 - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) \\ &\quad + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) \end{aligned} \quad (109)$$

As stated in Theorem 2.8, for a Schrödinger Bridge, the optimal coupling factorizes as

$$\Pi_{0,T}^*(X_0, X_T) = \mathbb{P}_{T|0}(X_T|X_0)\widehat{\varphi}_0(X_0)\varphi_T(X_T). \quad (110)$$

Hence, we can compute

$$\nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) = \nabla_{X_0} \log \mathbb{P}_{T|0}(X_T|X_0) + \nabla_{X_0} \log \widehat{\varphi}_0(X_0), \quad (111)$$

$$\nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) = \nabla_{X_T} \log \mathbb{P}_{T|0}(X_T|X_0) + \nabla_{X_T} \log \varphi_T(X_T). \quad (112)$$

Plugging this into (109) yields

$$\begin{aligned} \sigma^{-1}(t)\xi(X, t) &= -\nabla_{X_t} \log P_{t|0}(X_t|X_0) \\ &\quad + \frac{1-c(t)}{1-\gamma(t)} \nabla_{X_0} \log P_{T|0}(X_T|X_0) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log P_{T|0}(X_T|X_0) \\ &\quad + \frac{1-c(t)}{1-\gamma(t)} \nabla_{X_0} \log \psi_0(X_0) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \varphi_T(X_T) \end{aligned} \quad (113)$$

For the reference process  $dX_t = \sigma(t)dB_t$ , the transition density is  $\mathbb{P}_{T|0}(X_T|X_0) = \mathcal{N}(X_T; X_0, \kappa(T)I)$ . Hence,

$$\nabla_{X_T} \log \mathbb{P}_{T|0}(X_T|X_0) = -\nabla_{X_0} \log \mathbb{P}_{T|0}(X_T|X_0), \quad (114)$$

and therefore

$$\frac{1-c(t)}{1-\gamma(t)} \nabla_{X_0} \log \mathbb{P}_{T|0}(X_T|X_0) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \mathbb{P}_{T|0}(X_T|X_0) \quad (115)$$

$$= \frac{\gamma(t)-c(t)}{\gamma(t)(1-\gamma(t))} \nabla_{X_0} \log \mathbb{P}_{T|0}(X_T|X_0). \quad (116)$$

Further, using Theorem 2.4, we have

$$v^*(x, t) = \sigma(t)\mathbb{E}_{\Pi_{0,t}^*}[\nabla_{X_t} \log P_{t|0}(X_t|X_0)|X_t = x], \quad (117)$$

and from Theorem B.4

$$v^*(x, t) = \sigma(t)\mathbb{E}_{\Pi_{0,t}^*} \left[ \nabla_{X_0} \log \frac{\Pi_0^*(X_0)}{\varphi_0(X_0)} \Big| X_t = x \right] = \sigma(t)\mathbb{E}_{\Pi_{0,t}^*} \left[ \nabla_{X_0} \log \widehat{\varphi}_0(X_0) \Big| X_t = x \right]. \quad (118)$$

We can therefore compute

$$\mathbb{E}_{\Pi_{0,t}^*} \left[ -\nabla_{X_t} \log P_{t|0}(X_t|X_0) + \frac{1-c(t)}{1-\gamma(t)} \nabla_{X_0} \log \widehat{\varphi}_0(X_0) \Big| X_t = x \right] \quad (119)$$

$$= \mathbb{E}_{\Pi_{0,t}^*} \left[ \frac{\gamma(t)-c(t)}{1-\gamma(t)} \nabla_{X_0} \log \widehat{\varphi}_0(X_0) \Big| X_t = x \right]. \quad (120)$$

Finally, combining (113), (116) and (120), and noting again the Markovian projection formula from Theorem 2.3, we obtain

$$\begin{aligned} \sigma^{-1}(t)\xi(X, t) &= \frac{\gamma(t)-c(t)}{\gamma(t)(1-\gamma(t))} \nabla_{X_0} \log P_{T|0}(X_T|X_0) \\ &\quad + \frac{\gamma(t)-c(t)}{1-\gamma(t)} \nabla_{X_0} \log \psi_0(X_0) + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \varphi_T(X_T). \end{aligned} \quad (121)$$

□

### B.3 PROOFS FOR SCHRÖDINGER HALF BRIDGES

Under the memoryless condition, a Schrödinger Half Bridge can be formulated as a special case of a Schrödinger Bridge. Recall that we have the following minimization problem for SHBs:

$$\min_{u \in \mathcal{U}} \mathbb{E}_{\mathbb{P}^u} \left[ \int_0^T \frac{1}{2} \|u(X_t, t)\|^2 dt \right] \quad \text{s.t.} \quad X_T \sim \Pi_T^* = p_{\text{target}}. \quad (122)$$

This optimization yields the optimal joint coupling  $\Pi_{0,T}^* = \Pi_T^* \mathbb{P}_{0|T}$ . By the memoryless condition, the backward transition of the reference process becomes independent of the terminal state and thus  $\mathbb{P}_{0|T} = \mathbb{P}_0$ . Consequently, we also have  $\Pi_0^* = \mathbb{P}_0$ .

Since both the initial and terminal marginals are determined, we can identify the corresponding potentials for the Schrödinger system  $\Pi_t^* = \varphi_t \widehat{\varphi}_t$ .

- **Initial Potentials:** To satisfy  $\Pi_0^* = \varphi_0 \hat{\varphi}_0 = \mathbb{P}_0$ , we have the initial potentials as  $\hat{\varphi}_0 = \mathbb{P}_0$  and  $\varphi_0 = 1$ .
- **Terminal Potentials:** Under the reference dynamics, the forward potential evolves exactly as the reference marginal, yielding  $\hat{\varphi}_T = \mathbb{P}_T$ . To satisfy the terminal constraint  $\Pi_T^* = \varphi_T \hat{\varphi}_T$ , the terminal backward potential must therefore be  $\varphi_T = \Pi_T^*/\mathbb{P}_T$ .

**Lemma B.6** (Path-dependent drift for Schrödinger Half Bridges). *Let  $dX_t = \sigma(t)dB_t$  be the reference process with path measure  $\mathbb{P}$ . Consider the target path measure  $\Pi^* \in \mathcal{R}(\mathbb{P})$  defined by the Schrödinger Half Bridge coupling  $\Pi_{0,T}^* = \Pi_T^*(X_T)\mathbb{P}_{0|T}(X_0|X_T)$ . The unique optimal control  $u^*$  is given by:*

$$u^*(x, t) = \sigma(t)\mathbb{E}_{\Pi^*}[\xi(X, t)|X_t = x], \quad (123)$$

where

$$\begin{aligned} \sigma(t)^{-1}\xi(X, t) &= \frac{c(t) - \gamma(t)}{\gamma(t)} \nabla_{X_0} \log P_0(X_0) \\ &+ \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \frac{\Pi_T^*(X_T)}{P_T(X_T)} \\ &+ \frac{\gamma(t) - c(t)}{\gamma(t)(1 - \gamma(t))} \nabla_{X_0} \log P_{0|T}(X_0|X_T) \end{aligned} \quad (124)$$

for any  $c = c(t) \in [0, 1]$ .

*Proof.* For a general Schrödinger Bridge, the path-dependent drift  $\xi$  is given by Theorem B.5:

$$\begin{aligned} \sigma(t)^{-1}\xi(X, t) &= \frac{\gamma(t) - c(t)}{\gamma(t)(1 - \gamma(t))} \nabla_{X_0} \log P_{T|0}(X_T|X_0) \\ &+ \frac{\gamma(t) - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \frac{\Pi_0^*(X_0)}{p_0(X_0)} \\ &+ \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \frac{\Pi_T^*(X_T)}{p_T(X_T)} \end{aligned} \quad (125)$$

Substituting the known couplings for Half Bridges yields

$$\begin{aligned} \sigma(t)^{-1}\xi(X, t) &= \frac{\gamma(t) - c(t)}{\gamma(t)(1 - \gamma(t))} \nabla_{X_0} \log P_{T|0}(X_T|X_0) \\ &+ \frac{\gamma(t) - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log P_0(X_0) \\ &+ \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \frac{\Pi_T^*(X_T)}{P_T(X_T)} \end{aligned} \quad (126)$$

We use Bayes' theorem to expand the score of the forward transition:  $\nabla_{X_0} \log \mathbb{P}_{T|0}(X_T|X_0) = \nabla_{X_0} \log \mathbb{P}_{0|T}(X_0|X_T) - \nabla_{X_0} \log \mathbb{P}_0(X_0)$ . After substituting and grouping the terms of  $\nabla_{X_0} \log \mathbb{P}_0(X_0)$  together, we have

$$\begin{aligned} \sigma(t)^{-1}\xi(X, t) &= \frac{c(t) - \gamma(t)}{\gamma(t)} \nabla_{X_0} \log P_0(X_0) \\ &+ \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \frac{\Pi_T^*(X_T)}{P_T(X_T)} \\ &+ \frac{\gamma(t) - c(t)}{\gamma(t)(1 - \gamma(t))} \nabla_{X_0} \log P_{0|T}(X_0|X_T) \end{aligned} \quad (127)$$

□

*Proof of Theorem 2.7.* This directly follows from Theorem B.6 setting  $c(t) = \gamma(t) \forall t \in [0, T]$ . □

#### B.4 PROOFS FOR INDEPENDENT COUPLINGS

*Proof of Theorem 2.10.* This follows directly from Theorem 2.6, which for the general case states

$$\begin{aligned} \xi(X, t) = & -\sigma(t) \left( \nabla_x \log P_{t|0}(X_t|X_0) \right. \\ & + \frac{1-c(t)}{1-\gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) \\ & \left. + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) \right) \end{aligned} \quad (128)$$

For independent couplings, we have  $\Pi_{0,T}^*(x_0, x_T) = \Pi_0^*(x_0) \otimes \Pi_T^*(x_T)$  and can calculate

$$\nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) = \nabla_{X_0} \log \Pi_0^*(X_0), \quad (129)$$

$$\nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) = \nabla_{X_T} \log \Pi_T^*(X_T). \quad (130)$$

Substituting these gradients back into the expression for  $\xi$  concludes the proof.  $\square$

#### B.5 PROOFS FOR DAMPED FIXED-POINT ITERATES

*Proof of Theorem 2.11.* The operator  $\Phi(u_i)$  is defined as the Markovian projection of the drift  $\xi$  under the measure  $\Pi^i$ :

$$\Phi(u_i)(x, t) = \mathbb{E}_{\Pi^i}[\xi(X, t)|X_t = x]. \quad (131)$$

Let  $\mathcal{J}$  be the objective functional

$$\mathcal{J}(u) = \mathbb{E}_{\Pi^i} \left[ \int_0^T \left( \frac{1}{2} \|\xi(X, t) - u(X_t, t)\|^2 + \frac{\eta}{2} \|u_i(X_t, t) - u(X_t, t)\|^2 \right) dt \right]. \quad (132)$$

Using the bias-variance decomposition (see Theorem A.5), the first term of the objective can be decomposed as:

$$\begin{aligned} \mathbb{E}_{\Pi^i} [\|\xi(X, t) - u(X_t, t)\|^2 | X_t = x] \\ = \mathbb{E}_{\Pi^i} [\|\xi(X, t) - \Phi(u_i)(X_t, t)\|^2 | X_t = x] \\ + \|\Phi(u_i)(x, t) - u(x, t)\|^2. \end{aligned} \quad (133)$$

Since the first term does not depend on  $u$ , minimizing the original objective functional Equation (132) is equivalent to minimizing

$$\int_0^T \mathbb{E}_{\Pi^i} \left[ \frac{1}{2} \|\Phi(u_i)(X_t, t) - u(X_t, t)\|^2 + \frac{\eta}{2} \|u_i(X_t, t) - u(X_t, t)\|^2 \right] dt. \quad (134)$$

We minimize the integrand pointwise for almost every  $(x, t)$  by taking the variational derivative with respect to  $u$  and setting it to zero, i.e.

$$\nabla_u \left( \frac{1}{2} \|\Phi(u_i) - u\|^2 + \frac{\eta}{2} \|u_i - u\|^2 \right) = -(\Phi(u_i) - u) - \eta(u_i - u) \stackrel{!}{=} 0. \quad (135)$$

Rearranging terms to isolate  $u$  gives

$$u = \frac{1}{1+\eta} \Phi(u_i) + \frac{\eta}{1+\eta} u_i. \quad (136)$$

By substituting the definition of the damping parameter  $\eta = \frac{1-\alpha}{\alpha}$ , we find that  $\frac{1}{1+\eta} = \alpha$  and  $\frac{\eta}{1+\eta} = 1 - \alpha$ . Thus, the minimizer exactly recovers the damped fixed-point update rule:

$$u_{i+1}(x, t) = \alpha \Phi(u_i)(x, t) + (1 - \alpha) u_i(x, t). \quad (137)$$

$\square$

## B.6 PROOFS FOR CONTROL VARIATES

*Proof of Theorem D.1.* We seek to minimize the conditional variance term from the decomposition with respect to  $c(t)$ . Let us isolate the integrand:

$$\begin{aligned} J(c) &= \mathbb{E}_{\Pi_t^*} \left[ \mathbb{V}_{\Pi_t^*} (\sigma^{-1}(t)\xi^c(X, t|X_t)) \right] \\ &= \mathbb{E}_{\Pi^*} \left[ \left\| \sigma^{-1}(t)\xi^c(X, t) - \mathbb{E}_{\Pi^*} [\sigma^{-1}(t)\xi^c(X, t)|X_t] \right\|^2 \right]. \end{aligned} \quad (138)$$

Substituting our definition  $\sigma^{-1}(t)\xi^c = G_0(X) - G_v(X) - c(t)(G_0(X) - G_T(X))$  into the expected conditional variance yields

$$\begin{aligned} J(c) &= \mathbb{E}_{\Pi^*} \left[ \left\| (G_0(X) - G_v(X) - c(t)(G_0(X) - G_T(X))) \right. \right. \\ &\quad \left. \left. - \mathbb{E}_{\Pi^*} [G_0(X) - G_v(X) - c(t)(G_0(X) - G_T(X))|X_t] \right\|^2 \right] \end{aligned} \quad (139)$$

Crucially, because  $G_0$  and  $G_T$  are both unbiased estimators of the marginal score, their conditional expectations are equal. Therefore, the conditional expectation of their difference evaluates to zero:

$$\mathbb{E}_{\Pi^*} [G_0(X) - G_T(X)|X_t] = \mathbb{E}_{\Pi^*} [G_0(X)|X_t] - \mathbb{E}_{\Pi^*} [G_T(X)|X_t] = 0. \quad (140)$$

This simplifies our objective to

$$J(c) = \mathbb{E}_{\Pi^*} \left[ \left\| (G_0(X) - G_v(X) - \mathbb{E}_{\Pi^*} [G_0(X) - G_v(X)|X_t]) - c(t)(G_0(X) - G_T(X)) \right\|^2 \right]. \quad (141)$$

Expanding the squared norm using the inner product, and applying our definitions for scalar variance and covariance, we obtain a quadratic function in  $c(t)$ , namely

$$\begin{aligned} J(c) &= \mathbb{E}_{\Pi^*} \left[ \left\| G_0(X) - G_v(X) - \mathbb{E}_{\Pi^*} [G_0(X) - G_v(X)|X_t] \right\|^2 \right] \\ &\quad - 2c(t)\mathbb{E}_{\Pi^*} \left[ \langle G_0(X) - G_v(X) - \mathbb{E}_{\Pi^*} [G_0(X) - G_v(X)|X_t], \right. \\ &\quad \left. G_0(X) - G_T(X) \rangle \right] \\ &\quad + c(t)^2 \mathbb{V}(G_0 - G_T). \end{aligned} \quad (142)$$

By the law of total expectation, and recalling that  $\mathbb{E}_{\Pi^*} [G_0(X) - G_T(X)|X_t] = 0$ , the cross-term simplifies to exactly the covariance between  $(G_0 - G_v)$  and  $(G_0 - G_T)$ :

$$\mathbb{E}_{\Pi^*} \left[ \langle G_0(X) - G_v(X) - \mathbb{E}_{\Pi^*} [G_0(X) - G_v(X)|X_t], G_0(X) - G_T(X) \rangle \right] \quad (143a)$$

$$= \text{Cov}(G_0 - G_v, G_0 - G_T) - \mathbb{E}_{\Pi^*} \left[ \langle \mathbb{E}_{\Pi^*} [G_0(X) - G_v(X)|X_t], \mathbb{E}_{\Pi^*} [G_0(X) - G_T(X)|X_t] \rangle \right] \quad (143b)$$

$$= \text{Cov}(G_0 - G_v, G_0 - G_T). \quad (143c)$$

To find the minimum, we differentiate  $J(c)$  with respect to the scalar  $c(t)$  and set the derivative to zero:

$$\frac{\partial J(c)}{\partial c(t)} = -2 \text{Cov}(G_0 - G_v, G_0 - G_T) + 2c(t) \text{Var}(G_0 - G_T) = 0. \quad (144)$$

Solving for  $c(t)$  yields  $c^*(t) = \frac{\text{Cov}(G_0 - G_v, G_0 - G_T)}{\text{Var}(G_0 - G_T)}$ . Finally, using the bilinearity of covariance to expand the numerator and denominator produces the full expression in the proposition.  $\square$

*Proof of Theorem D.2.* The left term on the right-hand side of the bias-variance decomposition in (154) is independent of  $c$  due to  $\nabla_c \mathbb{E}_{\Pi^*} [\xi^c(X, t)|X_t] = 0$ . Minimizing the total objective  $\mathbb{E}_{\Pi^*} \left[ \int_0^T \frac{1}{2} \|u(X_t, t) - \xi^c(X, t)\|^2 dt \right]$  with respect to  $c$  is therefore equivalent to minimizing the variance penalty  $\mathbb{E}_{\Pi^*} \left[ \int_0^T \frac{1}{2} \|\xi^c(X, t) - \mathbb{E}[\xi^c(X, t)|X_t]\|^2 dt \right]$ . Using the equality in (155) concludes the proof.  $\square$

---

## C RELATED WORK

**Diffusion-based sampling.** Sampling from unnormalized densities has traditionally relied on Markov Chain Monte Carlo (MCMC) methods. Early learnable diffusion-based approaches utilized Schrödinger-Föllmer diffusions Zhang & Chen (2022); Richter (2021); Vargas et al. (2023b) or adaptations of denoising diffusion models to the energy-based setting Berner et al. (2024); Vargas et al. (2023a); Doucet et al. (2022); Huang et al. (2023). From the perspective of stochastic optimal control, many of these approaches can be viewed as finding controls that minimize divergences between forward and backward processes Richter & Berner (2024); Vargas et al. (2024); Blessing et al. (2025b). A series of recent works improved the objectives, training schemes, and exploration of diffusion-based samplers Zhang et al. (2024); He et al. (2024); OuYang et al. (2024); Erives et al. (2025); Wang et al. (2025); Yoon et al. (2025); Sanokowski et al. (2025); Gritsaev et al. (2025); Shi et al. (2024); Berner et al. (2025); Kim et al. (2024; 2025); Chen et al. (2025); Zhang et al. (2025a;b); Sendera et al. (2024). However, while effective, these methods typically require to keep SDE simulations in memory for the gradient computation, which is computationally expensive.

To address this, Havens et al. (2025) introduced *adjoint sampling* (AS), a framework that leverages the adjoint method to enable memory-efficient gradient estimation. This allowed them to scale diffusion-based sampling to amortized conformer generation across many molecular system. However, AS relies on a memoryless condition, restricting the reference process and prior distributions to specific forms (e.g., Dirac priors). This condition necessitates large diffusion coefficients to achieve good performance, which can leads to instability.

**Diffusion bridges.** To enable sampling between arbitrary boundary distributions, the problem can be framed as learning a Schrödinger Bridge (SB). Shi et al. (2023) introduced *Diffusion Schrödinger Bridge Matching* (DSBM) for the data-driven setting, utilizing Iterative Markovian Fitting that alternates between projecting onto Markovian and reciprocal classes. Recently, Liu et al. (2025) extended this to energy-based sampling with the *Adjoint Schrödinger Bridge Sampler* (ASBS), which generalizes AS by relaxing the memoryless condition, allowing for arbitrary priors. However, ASBS employs an alternating optimization scheme between the control and a corrector potential, which can be empirically unstable due to the interdependence of the objectives. While our method retains the flexibility of ASBS, it only relies on a single, scalable objective. To address mode collapse in complex landscapes, Nam et al. (2025) proposed *Well-Tempered ASBS* (WT-ASBS), which augments diffusion samplers with metadynamics-inspired bias potentials along collective variables. While our work focuses on the stability and scalability of the core sampling algorithm, our method is compatible with such enhanced sampling techniques.

**Target score estimation.** A core component of matching-based samplers is the estimation of the regression target. De Bortoli et al. (2024) formally introduced the *Target Score Identity* (TSI) and *Target Score Matching* (TSM), demonstrating that score estimation can be significantly improved by leveraging the known structure of the target score  $\nabla \log p_{\text{target}}$ . This identity underpins the regression targets used in AS, ASBS, and in our framework.

## D FURTHER ALGORITHMIC DETAILS

### D.1 IMPORTANCE SAMPLING AND PROBABILITY FLOW ORDINARY DIFFERENTIAL EQUATION

**Importance sampling in path space.** Our primary goal is to estimate the expectation of an observable  $\Omega \in C(\mathbb{R}^d, \mathbb{R})$  under a target distribution, denoted as  $\mathbb{E}_{X_T \sim p_{\text{target}}} [\Omega(X_T)]$ . Since the optimized process  $\mathbb{P}^u$  serves as an approximation to the true target dynamics (i.e.,  $\mathbb{P}_T^u \approx p_{\text{target}}$ ), we can correct for the resulting bias by employing importance sampling. While importance sampling is classically applied to probability densities, it extends naturally to path space Berner et al. (2024), allowing for the reweighting of entire continuous trajectories of a stochastic process. We can express the target expectation as an integral over the proposal measure  $\mathbb{P}^u$ :

$$\mathbb{E}_{X_T \sim p_{\text{target}}} [\Omega(X_T)] = \mathbb{E}_{\mathbb{P}^u} \left[ \frac{d\mathbb{P}^v}{d\mathbb{P}^u}(X) \Omega(X_T) \right]. \quad (145)$$

Here, the Radon-Nikodym derivative (RND),  $w(X) := \frac{d\mathbb{P}^v}{d\mathbb{P}^u}(X)$ , is given by an extension of Girsanov's theorem Nüsken & Richter (2021); Vargas et al. (2024); Richter & Berner (2024):

$$w(X) = \frac{p_{\text{target}}(X_T)}{p_{\text{prior}}(X_0)} \exp \left( \int_0^T \sigma^{-1}(v - u)(X_t, t) \cdot dX_t - \frac{1}{2} \int_0^T (\|v\|^2 - \|u\|^2)(X_t, t) dt \right). \quad (146)$$

Practical implementation faces two distinct challenges. First, we typically only have access to  $p_{\text{target}}$  up to an intractable normalization constant  $\mathcal{Z}$ . We therefore resort to self-normalized importance sampling Tokdar & Kass (2010):

$$\frac{\mathbb{E}_{\mathbb{P}^u} [\tilde{w}(X) \Omega(X_T)]}{\mathbb{E}_{\mathbb{P}^u} [\tilde{w}(X)]}, \quad (147)$$

where  $\tilde{w} = w\mathcal{Z}$ . The second challenge is that evaluating the RND requires access to both the forward and the backward drift  $u, v$ , respectively. However, our proposed fixed-point iteration yields only  $u$ , making the direct computation of (146) impossible. To resolve this, we leverage Nelson's identity (15) to formulate a regression problem for learning  $v \approx v^*$  and the score  $s \approx \sigma \nabla \log \Pi^*$  separately, instead of  $u$ :

$$\mathcal{L}(v, s) := \mathbb{E}_{\Pi^i} \left[ \int_0^T \frac{1}{2} \|\xi^v(X, t) - v(X_t, t)\|^2 dt \right] + \mathbb{E}_{\Pi^i} \left[ \int_0^T \frac{1}{2} \|\xi^s(X, t) - s(X_t, t)\|^2 dt \right], \quad (148)$$

with

$$v^*(x, t) = \mathbb{E}_{\Pi^*} [\xi^v(X, t) | X_t = x] \quad \text{and} \quad \sigma(t) \nabla \log \Pi_t^*(x) = \mathbb{E}_{\Pi^*} [\xi^s(X, t) | X_t = x], \quad (149)$$

where  $\xi^v = \sigma \nabla \log \mathbb{P}_{t|0}$  and  $\xi^s$  is inferred via the generalized target score identity. Implementation-wise, it is possible to employ a shared neural backbone to approximate  $v$  and  $s$  simultaneously, potentially improving parameter efficiency. However, we leave the exploration of such architectural optimizations for future work. Finally, the forward process (1) can be simulated as

$$dX_t = \sigma(t) (s - v)(X_t, t) dt + \sigma(t) dB_t, \quad X_0 \sim p_{\text{prior}}. \quad (150)$$

**Probability flow ODE.** An alternative approach to importance sampling is to leverage the Probability Flow ODE (PF-ODE) Maoutsa et al. (2020); Song et al. (2020); Chen et al. (2021). This deterministic process shares the same marginal densities  $(\Pi_t^*)_{t \in [0, T]}$  as the optimally controlled stochastic process but follows a smooth trajectory defined by:

$$dX_t = \underbrace{\left( \sigma u^* - \frac{1}{2} \sigma^2 \nabla \log \Pi_t^* \right)}_{:= f(X_t, t)} (X_t, t) dt. \quad (151)$$

By leveraging Nelson's identity (15), we can express the drift term  $f$  in equivalent forms involving the backward control  $v$ , i.e.,

$$\sigma u^* - \frac{1}{2} \sigma^2 \nabla \log \Pi_t^* = -\sigma v^* + \frac{1}{2} \sigma^2 \nabla \log \Pi_t^* = \frac{1}{2} \sigma (u^* - v^*), \quad (152)$$

suggesting that one needs two quantities among  $u, v, \nabla \log \Pi^*$  to simulate (151). Consequently, one can use the objective (148) to learn  $v$  and the score  $\nabla \log \Pi^*$  separately. A key advantage of the ODE formulation is the ability to compute exact log-likelihoods via the *instantaneous change of variables* formula Chen et al. (2018). The evolution of the log-density is governed by the divergence of the drift field:

$$\log \mathbb{P}_T^u(X_T) = \log p_{\text{prior}}(X_0) - \int_0^T \nabla \cdot f(X_t, t) dt, \quad (153)$$

where  $\nabla \cdot f = \text{Tr}(\nabla f)$  denotes the divergence. While this allows for exact density estimation, computing the trace of the Jacobian  $\nabla f$  is computationally expensive for high-dimensional problems. To mitigate this cost, one can trade exactness for efficiency by employing unbiased estimators, such as Hutchinson's trace estimator Hutchinson (1989). Having access to  $\mathbb{P}_T^u$  allows for importance sampling via (147) with  $w = p_{\text{target}} / \mathbb{P}_T^u$ .

## D.2 VARIANCE REDUCTION AND CONTROL VARIATES

The efficiency of learning the optimal drift  $u^*$  is limited by the stochasticity of the target drift  $\xi$ . To formalize this, we observe that the matching objective decomposes into two distinct components via Theorem A.5:

$$\mathbb{E}_{\Pi^*} \left[ \int_0^T \frac{1}{2} \|u(X_t, t) - \xi(X, t)\|^2 dt \right] = \underbrace{\mathbb{E}_{\Pi^*} \left[ \int_0^T \frac{1}{2} \|u(X_t, t) - \mathbb{E}_{\Pi^*}[\xi(X, t)|X_t]\|^2 dt \right]}_{\text{bias}} + \underbrace{\mathbb{E}_{\Pi^*} \left[ \int_0^T \frac{1}{2} \|\xi(X, t) - \mathbb{E}_{\Pi^*}[\xi(X, t)|X_t]\|^2 dt \right]}_{\text{variance}}. \quad (154)$$

While the variance term

$$\mathbb{E}_{\Pi^*} \left[ \int_0^T \frac{1}{2} \|\xi(X, t) - \mathbb{E}_{\Pi^*}[\xi(X, t)|X_t]\|^2 dt \right] = \int_0^T \frac{1}{2} \mathbb{E}_{\Pi^*} \left[ \text{Var}_{\Pi_t^*}(\xi(X, t|X_t)) \right] dt, \quad (155)$$

does not change the optimal drift  $u^*$ , it adds noise to the empirical gradient. To minimize its contribution, we first recall the definition of  $\xi$ :

$$\begin{aligned} \sigma^{-1}(t)\xi^c(X, t) &= \frac{1 - c(t)}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) \\ &\quad + \frac{c(t)}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) \\ &\quad - \nabla_{X_t} \log P_{t|0}(X_t|X_0) \end{aligned} \quad (156)$$

where we make the dependence of  $\xi$  on  $c$  explicit. In what follows, we show that  $c$  can be used to reduce the empirical variance.

**Control variate interpretation of  $c$ .** We start by making the following definitions:

$$G_0(X) := \frac{1}{1 - \gamma(t)} \nabla_{X_0} \log \Pi_{0,T}^*, \quad G_T(X) := \frac{1}{\gamma(t)} \nabla_{X_T} \log \Pi_{0,T}^*, \quad G_v(X) := \nabla_{X_t} \log \mathbb{P}_{t|0}. \quad (157)$$

Using these definitions, we can express the scaled target field as  $\sigma^{-1}\xi^c = (1 - c)G_0 + cG_T - G_v$ . Crucially, both  $G_0$  and  $G_T$  are unbiased estimators of the true score  $\nabla_x \log \Pi_t^*(x)$  when conditioned on  $X_t$ . This implies:

$$\begin{aligned} \mathbb{E}_{\Pi^*} [\sigma^{-1}(t)\xi^c(X, t) | X_t = x] &= \mathbb{E}_{\Pi^*} [(1 - c(t))G_0 + c(t)G_T - G_v | X_t = x] \\ &= \underbrace{\mathbb{E}_{\Pi^*} [G_0 - G_v | X_t = x]}_{=\sigma^{-1}(t)u^*(x,t)} - c(t) \underbrace{\mathbb{E}_{\Pi^*} [G_0 - G_T | X_t = x]}_{=0}. \end{aligned} \quad (158)$$

The term  $\Delta G := G_0 - G_T$  acts as a zero-mean control variate. We can therefore optimize  $c$  to minimize the variance contribution in (154) by solving:

$$c^* = \arg \min_c \int_0^T \frac{1}{2} \mathbb{E}_{\Pi_t^*} \left[ \text{Var}_{\Pi_t^*}(\xi^c(X, t|X_t)) \right] dt. \quad (160)$$

**Proposition D.1** (Optimal scalar-valued control variate). *For any time  $t \in [0, T]$ , noting that the estimators are zero-mean, we define the scalar variance (equivalent to the trace of the covariance matrix) as  $\text{Var}(A) := \mathbb{E}_{\Pi^*} [|A|^2]$  and the scalar cross-covariance as  $\text{Cov}(A, B) := \mathbb{E}_{\Pi^*} [A, B]$ . The optimal control variate  $c^*(t)$  that minimizes the conditional variance of  $\xi^c$  is given by:*

$$c^*(t) = \frac{\text{Var}(G_0) - \text{Cov}(G_0, G_T) - \text{Cov}(G_0, G_v) + \text{Cov}(G_T, G_v)}{\text{Var}(G_0) + \text{Var}(G_T) - 2\text{Cov}(G_0, G_T)}. \quad (161)$$

**Joint optimization of the drift and control variate.** In practice, evaluating the exact variances and covariances in the optimal schedule  $c^*(t)$  is intractable. However, we can circumvent this limitation and still achieve variance reduction by jointly learning the drift  $u$  and the schedule  $c$ . This approach is formally justified by the following proposition, which demonstrates that minimizing the fixed-point diffusion matching objective with respect to  $c$  exactly minimizes the conditional variance of the target.

**Proposition D.2.** *Let  $\xi^c$  be defined as in (156). It holds that*

$$c^* = \arg \min_c \mathbb{E}_{\Pi^*} \left[ \int_0^T \frac{1}{2} \|u(X_t, t) - \xi^c(X, t)\|^2 dt \right] \quad (162)$$

$$= \arg \min_c \int_0^T \frac{1}{2} \mathbb{E}_{\Pi_t^*} \left[ \text{Var}_{\Pi_t^*} (\xi^c(X, t|X_t)) \right] dt. \quad (163)$$

Because  $c(t)$  strictly weights a zero-mean control variate, the conditional expectation of our target – and therefore the optimal drift  $u^*$  – remains perfectly invariant to its value. Looking at the decomposition of our main matching loss, this implies that any gradient taken with respect to  $c$  exclusively targets the irreducible conditional variance of  $\xi^c$ . Consequently, following Ko & Geffner (2025); Kahouli et al. (2025); Young et al. (2026) we can parameterize the control variate as a learnable function  $c(t)$  (e.g., a lightweight neural network) and jointly train it alongside the drift network  $u$ , i.e.,

$$u^*, c^* = \arg \min_{u, c} \mathbb{E}_{\Pi^*} \left[ \int_0^T \frac{1}{2} \|u(X_t, t) - \xi^c(X, t)\|^2 dt \right], \quad (164)$$

using the exact same objective. By treating the variance reduction as an end-to-end learning problem, the model can dynamically discover the optimal interpolation schedule during optimization.

Note, that Equation (158) only holds if the expectation is taken with respect to  $\Pi^*$ . However, when performing our proposed fixed-point iteration (14), we only have access to samples from  $\Pi^i$ . It is therefore unclear if a joint optimization (164) with samples from  $\Pi^i$  leads to variance reduction. We leave the numerical investigation for future work. For completeness, we provide a parameterization of the control variate  $c$  that avoids numerical instabilities at the boundaries  $t = 0$  and  $t = T$  below.

While Equation (164) assumes access to the true coupling  $\Pi^*$ , our fixed-point iteration (14) relies on samples from the intermediate distribution  $\Pi^i$ . The extent to which joint optimization preserves its variance-reduction properties under this distributional shift remains an open question, which we leave for future numerical study. For completeness, we detail below a parameterization of  $c$  designed to prevent numerical instabilities at the boundaries  $t = 0$  and  $t = T$ .

**Parameterization of the learned control variate.** A critical consideration when learning  $c$  is the numerical stability of the target drift  $\xi^c$  at the temporal boundaries. As defined in (156), the terms weighting the boundary score estimators are scaled by  $(1 - \gamma(t))^{-1}$  and  $\gamma(t)^{-1}$ . Under standard interpolation schemes where  $\gamma(0) = 0$  and  $\gamma(T) = 1$ , these fractions introduce division-by-zero singularities at  $t = 0$  and  $t = T$ , respectively. To prevent the empirical gradient from exploding, the control variate must strictly satisfy the boundary conditions  $c(0) = 0$  and  $c(T) = 1$ . This ensures the numerators vanish at the boundaries, allowing the limits to remain finite. To enforce this constraint by design while still retaining full expressivity, we parameterize the learnable schedule as

$$c^\phi(t) = \gamma(t) + \gamma(t)(1 - \gamma(t))\text{NN}^\phi(t), \quad (165)$$

where  $\text{NN}^\phi(t)$  is a neural network with learnable parameters  $\phi$ . In this formulation, the base term  $\gamma(t)$  provides a stable linear-equivalent baseline, while the multiplicative envelope  $\gamma(t)(1 - \gamma(t))$  acts as a gating mechanism. It zeroes out the network’s contribution exactly at the boundaries, satisfying  $c^\phi(0) = 0$  and  $c^\phi(T) = 1$ . Substituting  $c^\phi(t)$  back into the coefficients of (156) yields

$$\frac{c^\phi(t)}{\gamma(t)} = 1 + (1 - \gamma(t))\text{NN}^\phi(t) \quad \text{and} \quad \frac{1 - c^\phi(t)}{1 - \gamma(t)} = 1 - \gamma(t)\text{NN}^\phi(t), \quad (166)$$

and therefore (156) becomes

$$\sigma^{-1}(t)\xi^\phi(X, t) = \left(1 - \gamma(t)\text{NN}^\phi(t)\right) \nabla_{X_0} \log \Pi_{0,T}^*(X_0, X_T) \quad (167)$$

$$+ \left(1 + (1 - \gamma(t))\text{NN}^\phi(t)\right) \nabla_{X_T} \log \Pi_{0,T}^*(X_0, X_T) - \nabla_{X_t} \log \mathbb{P}_{t|0}(X_t|X_0). \quad (168)$$

Because both simplified coefficients are well-defined and bounded for all  $t \in [0, T]$ , this parameterization completely circumvents the boundary explosion, ensuring stable joint optimization of  $u_\theta$  and  $c_\phi$ .

### D.3 LOSS WEIGHTING AND NUMERICAL STABILITY

In this section, we analyze the variance properties of the matching objective and derive a weighting schedule to mitigate numerical instabilities near  $t = 0$ . We introduce a time-dependent weighting function  $\omega(t) > 0$  to the primary matching objective. This weighting balances the contribution of the loss over time without altering the optimal solution  $u^* \in \mathcal{U}$  at the critical points of the objective. The weighted loss is given by

$$\mathcal{L}(u) = \mathbb{E}_{\Pi} \left[ \int_0^T \omega(t) \|u(X_t, t) - \xi(X, t)\|^2 dt \right]. \quad (169)$$

The target vector field  $\xi(X, t)$  typically includes the score of the forward transition kernel  $\nabla \log \mathbb{P}_{t|0}(X_t|X_0)$ , which dominates the variance as  $t \rightarrow 0$ . For a general Gaussian reference process with variance schedule  $\kappa(t) = \int_0^t \sigma^2(s) ds$ , the transition kernel is given by  $\mathbb{P}_{t|0}(X_t|X_0) = \mathcal{N}(X_t; X_0, \kappa(t)I)$ . Consequently, the scaled score term is

$$\sigma(t) \nabla_{X_t} \log \mathbb{P}_{t|0}(X_t|X_0) = \sigma(t) \frac{X_0 - X_t}{\kappa(t)}. \quad (170)$$

Substituting the interpolant  $X_t = \frac{\kappa(T) - \kappa(t)}{\kappa(T)} X_0 + \frac{\kappa(t)}{\kappa(T)} X_T + \left( B_{\kappa(t)} - \frac{\kappa(t)}{\kappa(T)} B_{\kappa(T)} \right)$  into the score expression allows us to isolate the stochastic component. Note that the stochastic part corresponds to a Brownian bridge with variance  $\frac{\kappa(t)(\kappa(T) - \kappa(t))}{\kappa(T)}$  which can be represented as  $\sqrt{\frac{\kappa(t)(\kappa(T) - \kappa(t))}{\kappa(T)}} \varepsilon$ , where  $\varepsilon \sim \mathcal{N}(0, I)$ . The expansion of the score term yields

$$\sigma(t) \nabla_{X_t} \log \mathbb{P}_{t|0}(X_t|X_0) = \underbrace{\frac{\sigma(t)}{\kappa(T)} (X_0 - X_T)}_{\text{drift}} - \underbrace{\sigma(t) \sqrt{\frac{\kappa(T) - \kappa(t)}{\kappa(T)\kappa(t)}}}_{\text{noise}} \varepsilon. \quad (171)$$

We now analyze the variance of the noise term, denoted by  $\Sigma_{\text{noise}}^2(t)$ , as  $t \rightarrow 0$ . Since  $\kappa(0) = 0$ , the term  $\kappa(t)$  in the denominator causes a singularity. Approximating  $\kappa(T) - \kappa(t) \approx \kappa(T)$  for small  $t$ , the variance behaves as

$$\Sigma_{\text{noise}}^2(t) = \sigma^2(t) \frac{\kappa(T) - \kappa(t)}{\kappa(T)\kappa(t)} \approx \frac{\sigma^2(t)}{\kappa(t)}. \quad (172)$$

Assuming  $\sigma(t)$  is roughly constant near  $t = 0$ ,  $\kappa(t) \approx \sigma^2 t$ , implying that the variance explodes as  $\mathcal{O}(1/t)$ . To enforce numerical stability, we select  $\omega(t)$  to counteract this variance, ensuring the effective loss contribution remains bounded. The optimal choice, often referred to as the *noise-prediction* weighting, is the inverse of the asymptotic variance factor:

$$\omega(t) = \frac{\kappa(t)}{\sigma^2(t)}. \quad (173)$$

Applying this weighting to the loss function neutralizes the singularity. By absorbing the weight into the norm, we obtain a numerically stable parameterization, namely

$$\mathcal{L}(u) = \mathbb{E}_{\Pi} \left[ \int_0^T \left\| \frac{\sqrt{\kappa(t)}}{\sigma(t)} u(X_t, t) - \frac{\sqrt{\kappa(t)}}{\sigma(t)} \xi(X, t) \right\|^2 dt \right] \quad (174a)$$

$$= \mathbb{E}_{\Pi} \left[ \int_0^T \left\| \hat{u}(X_t, t) - \left( \frac{\sqrt{\kappa(t)}}{\kappa(T)} (X_0 - X_T) - \sqrt{1 - \frac{\kappa(t)}{\kappa(T)}} \varepsilon + \sqrt{\kappa(t)} \nabla \log \Pi_t^*(X_t) \right) \right\|^2 dt \right], \quad (174b)$$

with  $\hat{u}(x, t) = \frac{\sqrt{\kappa(t)}}{\sigma(t)} u(x, t)$ . Here, the coefficient of the noise  $\varepsilon$  becomes  $\sqrt{1 - \kappa(t)/\kappa(T)}$ , which approaches 1 as  $t \rightarrow 0$ , resulting in a stable, unit-variance objective throughout training. However, while the weighting  $\omega(t)$  stabilizes training, it necessitates rescaling the drift of the forward SDE in (1) to recover the drift  $u$ . Substituting  $u(x, t) = \frac{\sigma(t)}{\sqrt{\kappa(t)}} \hat{u}(x, t)$  back into the forward process gives

$$dX_t = \frac{\sigma^2(t)}{\sqrt{\kappa(t)}} \hat{u}(X_t, t) dt + \sigma(t) dB_t. \quad (175)$$

Table 3: Overview: our framework considers target path measures of the form  $\Pi^* = \Pi_{0,T}^* \mathbb{P}_{|0,T}$  with different couplings  $\Pi_{0,T}^*$ . We provide expressions for  $\xi$  that satisfy  $u^*(x, t) = \sigma(t) \mathbb{E}_{\Pi^*} [\xi(X, t) | X_t = x]$  and are used for the proposed fixed-point iteration scheme in Algorithm 1.

Coupling		Formula for $\sigma(t)^{-1} \xi(X)$
Bridge dep.	$\Pi_{0,T}^*$	$\frac{1-c(t)}{1-\gamma(t)} \nabla_{x_0} \log \Pi_{0,T}^*(X_0, X_T) + \frac{c(t)}{\gamma(t)} \nabla_{x_T} \log \Pi_{0,T}^*(X_0, X_T) - \nabla_{x_t} \log \mathbb{P}_{t 0}(X_t   X_0)$
SHB	$\Pi_{0,T}^* = \mathbb{P}_0 \otimes \Pi_T^*$	$\frac{\gamma(t)-c(t)}{1-\gamma(t)} \nabla_{x_0} \log \mathbb{P}_0(X_0) + \frac{c(t)}{\gamma(t)} \nabla_{x_T} \log \frac{\Pi_T^*(X_T)}{\mathbb{P}_T(X_T)} + \frac{\gamma(t)-c(t)}{\gamma(t)(1-\gamma(t))} \nabla_{x_0} \log \mathbb{P}_{0 T}(X_0   X_T)$
SB	$\Pi_{0,T}^* = \Pi_{0,T}^{\text{SB}}$	$\frac{\gamma(t)-c(t)}{1-\gamma(t)} \nabla_{x_0} \log \frac{\Pi_0^*(X_0)}{\varphi_0(X_0)} + \frac{c(t)}{\gamma(t)} \nabla_{x_T} \log \frac{\Pi_T^*(X_T)}{\varphi_T(X_T)} + \frac{\gamma(t)-c(t)}{\gamma(t)(1-\gamma(t))} \nabla_{x_0} \log \mathbb{P}_{T 0}(X_T   X_0)$
Bridge indep.	$\Pi_{0,T}^* = \Pi_0^* \otimes \Pi_T^*$	$\frac{1-c(t)}{1-\gamma(t)} \nabla_{x_0} \log \Pi_0^*(X_0) + \frac{c(t)}{\gamma(t)} \nabla_{x_T} \log \Pi_T^*(X_T) - \nabla_{x_t} \log \mathbb{P}_{t 0}(X_t   X_0)$

This formulation reveals a numerical challenge near the boundary  $t = 0$ . Assuming  $\sigma(t) \approx \sigma$  is constant near the origin, we have  $\kappa(t) \approx \sigma^2 t$ . Consequently, the effective drift coefficient scales as

$$\frac{\sigma^2(t)}{\sqrt{\kappa(t)}} \approx \frac{\sigma^2}{\sigma \sqrt{t}} \propto \frac{1}{\sqrt{t}}. \quad (176)$$

This  $t^{-1/2}$  singularity implies that the drift velocity diverges as  $t \rightarrow 0$ . Standard fixed-step solvers (e.g., Euler-Maruyama) may exhibit instability or large discretization errors if evaluated strictly at  $t = 0$ . In practice, this is mitigated by starting the integration at a small cutoff time, e.g.,  $t = 10^{-3}$ .

#### D.4 REFERENCE PROCESS

In this work, we use a reference measure  $\mathbb{P}$  induced by scaled Brownian motion. The dynamics are governed by the driftless SDE  $dX_t = \sigma(t) dB_t$ , where  $\sigma(t)$  is a time-dependent noise schedule. To characterize the path properties, we define the cumulative variance schedule as  $\kappa(t) := \int_0^t \sigma^2(s) ds$ . When conditioned on fixed boundary values  $X_0 = x_0$  and  $X_T = x_T$ , the process admits a closed-form representation known as a stochastic interpolant Albergo et al. (2025):

$$X_t = \frac{\kappa(T) - \kappa(t)}{\kappa(T)} x_0 + \frac{\kappa(t)}{\kappa(T)} x_T + \left( B_{\kappa(t)} - \frac{\kappa(t)}{\kappa(T)} B_{\kappa(T)} \right). \quad (177)$$

Using  $\gamma(t) := \frac{\kappa(t)}{\kappa(T)}$ , we recover the simplified notation

$$X_t = (1 - \gamma(t)) x_0 + \gamma(t) x_T + \mathcal{B}_t^\gamma, \quad (178)$$

where  $\mathcal{B}_t^\gamma$  is a Gaussian process with mean zero and covariance  $\kappa(T) \gamma(t) (1 - \gamma(t)) I$ . This analytic form allows for the direct simulation of bridge states  $X_t \sim \mathbb{P}_{t|0,T}(\cdot | x_0, x_T)$  without requiring numerical integration of the SDE. The relevant transition densities are given by

$$\mathbb{P}_{t|0,T}(x_t | x_0, x_T) = \mathcal{N}(x_t | (1 - \gamma(t)) x_0 + \gamma(t) x_T, \kappa(T) \gamma(t) (1 - \gamma(t)) I), \quad (179)$$

$$\mathbb{P}_{t|0}(x_t | x_0) = \mathcal{N}(x_t | x_0, \kappa(t) I), \quad (180)$$

$$\mathbb{P}_{T|t}(x_T | x_t) = \mathcal{N}(x_T | x_t, \kappa(T) (1 - \gamma(t)) I). \quad (181)$$

#### D.5 DETAILED ALGORITHMIC DESCRIPTION

We provide a detailed description of our framework in Algorithm 1.

## E EXPERIMENTAL SETUP

### E.1 GAUSSIAN MIXTURE MODELS

**Target density specifications.** We consider a *Gaussian mixture model (GMM)* target of the form

$$p_{\text{GMM}}(x) = \sum_{k=1}^K \pi_k \mathcal{N}(x | \mu_k, \Sigma_k), \quad (182)$$

where  $\mu_k \in \mathbb{R}^d$ ,  $\Sigma_k \in \mathbb{R}^{d \times d}$ ,  $\pi_k \geq 0$ , and  $\sum_{k=1}^K \pi_k = 1$ . For our experiments, we set  $\pi_k = 1/K$ ,  $\Sigma_k = I$ . The means  $\mu_k$  are uniformly sampled from  $[-K, K]^d$ .

---

**Algorithm 1** Damped Fixed-Point Diffusion Matching (detailed version)

---

**Require:** Neural network  $u_\theta$  with parameters  $\theta$ , diffusion schedule  $\sigma$ , target path measure  $\Pi^*$ , buffer size  $K$ , damping  $\eta$ , number of outer steps  $I$ , number of gradient steps  $M$  per outer step

**for**  $i \leftarrow 0, \dots, I - 1$  **do**

$u_i \leftarrow u_\theta$  (detached)

*Simulation:* Simulate  $K$  trajectories  $\{X^{(k)} \sim \mathbb{P}^{u_i}\}_{k=1}^K$

*Coupling:* Let  $\Pi_{0,T}^i = \begin{cases} \mathbb{P}_{0,T}^{u_i}, & \text{ASBS Liu et al. (2025)} \\ \mathbb{P}_0 \otimes \mathbb{P}_T^{u_i}, & \text{AS Havens et al. (2025)} \\ \mathbb{P}_0^{u_i} \otimes \mathbb{P}_T^{u_i}, & \text{BMS (ours)} \end{cases}$

*Buffer:* Initialize  $\mathcal{B} = \{(X_0^{(k)}, X_T^{(k)}) \sim \Pi_{0,T}^i\}_{k=1}^K$

**for**  $m \leftarrow 1, \dots, M$  **do**

Estimate  $\mathcal{L}(\theta) = \mathbb{E}_{(X_0, X_T) \sim \mathcal{B}, X_t \sim \mathbb{P}_{t|0,T}} \left[ \frac{1}{2} \|\xi(X, t) - u_\theta(X_t, t)\|^2 + \frac{\eta}{2} \|u_i(X_t, t) - u_\theta(X_t, t)\|^2 \right]$

$\theta \leftarrow \text{step}(\theta, \nabla_\theta \mathcal{L}(\theta))$

**end for**

**end for**

**return** optimized control  $u_I$

---




Figure 2: Comparison of sliced TVD values on Gaussian mixture models across varying dimensions  $d$  and damping values  $\eta$ . Figure 2a shows that ASBS exhibits instability at low damping values (N/A) and performance degradation as dimensionality increases. Figure 2b demonstrates that our method maintains stable training and consistently avoids mode collapse up to  $d = 2500$ , significantly scaling beyond the  $d = 50$  range common in recent literature. All values are averaged over three random seeds.

**Evaluation criteria.** The *mode TVD* metric is proposed by Blessing et al. (2025a) and inspired by recent work Blessing et al. (2024); Grenioux et al. (2025). It uses the ground truth mixing weights  $\pi_i$ ,  $k \in \{1, \dots, K\}$ , along with a partition  $\{S_1, \dots, S_i\}$  of  $\mathbb{R}^d$ , where each region  $S_i \subset \mathbb{R}^d$  corresponds to the  $k$ -th mode of the target distribution given as

$$S_i = \{x \in \mathbb{R}^d \mid \arg \max_j \pi_j \mathcal{N}(x \mid \mu_j, \Sigma_j) = k\}. \quad (183)$$

We estimate the empirical mixing weights using

$$\hat{\pi}_i = \frac{\mathbb{E}_{\mathbb{P}^{u_i}} [\mathbb{1}_{S_i}(X_T)]}{\sum_{k'=1}^K \mathbb{E}_{\mathbb{P}^{u_i}} [\mathbb{1}_{S_{k'}}(X_T)]}. \quad (184)$$

Using these estimates, we compute the total variation distance (TVD) between the empirical and true mode weights as

$$\text{mode TVD} = \sum_{k=1}^K |\pi_i - \hat{\pi}_i|, \quad (185)$$

where we approximate the expectation in (184) using  $2k$  samples. Note that mode TVD  $\in [0, 1]$  where a value of 0 indicates that the model learned the same mixture weights as the target and therefore did not suffer from any mode-collapse.

Additionally, we use *sliced TVD* to assess the accuracy of the density fitting by comparing empirical




Figure 3: Comparison of *mode TVD*, *sliced TVD* and Wasserstein-2  $W_2$  on a Gaussian mixture model target in  $d = 100$  with different number of modes. All values are averaged over three random seeds.

densities against the analytically available marginals. For a set of model samples  $\{x_n\}_{n=1}^N$  and target samples  $\{y_m\}_{m=1}^M$ , we project the data onto  $P$  random unit vectors  $\{\theta_p\}_{p=1}^P$  sampled uniformly from the unit sphere  $\mathbb{S}^{d-1}$ . For each projection, we compute the 1D total variation distance between the resulting histograms:

$$\text{sliced TVD} = \frac{1}{P} \sum_{p=1}^P \text{TVD}(\text{hist}(\{x_n \cdot \theta_p\}), \text{hist}(\{y_m \cdot \theta_p\})). \quad (186)$$

In our experiments, we use  $P = 100$  projections and 50 bins to estimate these 1D distributions. Lastly, we compare the Wasserstein-2 ( $W_2$ ) optimal transport distance between  $2k$  model and target samples using the POT library Flamary et al. (2021).

**Algorithm specifications.** For the all GMM experiments we used the same setting for ASBS and BMS: A residual network with 6 layers and 512 hidden units and Fourier time embedding is used. The source distribution is modeled as a normal distribution which is adapted to the support of the target by setting  $\mathcal{N}(0, K^2)$ . For optimization, the Adam optimizer Kingma (2014) with a learning rate of  $10^{-4}$  and (value) gradient clipping at 1 is used. Furthermore, we leverage a batchsize of 1024 and a buffer size of  $30k$ . We perform  $1k$  outer steps where each outer step entails  $1k$  gradient steps on the current buffer. The diffusion process is integrated using the Euler-Maruyama method with 100 discretization steps. For the GMM experiments, we used a constant noise scheduler, i.e.  $\sigma(t) = \sigma$ . For both ASBS and BMS, we tested values in  $\{1.5, 2.5, 4\}$  and found that  $\sigma = 2.5$  works best for both methods. In our attempts to compare to adjoint sampling Havens et al. (2025) we required larger values for  $\sigma$  in order to roughly cover the support of the target. Particularly, we set  $\sigma = K$ , as the means are sampled from  $[-K, K]^d$ . We conjecture that these large  $\sigma$  values caused the instabilities in training. For all algorithms, we used  $c = \gamma$  when computing  $\xi$ .

## E.2 $n$ -BODY IDENTICAL PARTICLE SYSTEMS

**Target density specifications.** Following the configurations in Köhler et al. (2020) and Akhound-Sadeh et al. (2024), we evaluate our method on two primary particle systems defined by their potential energy functions  $E$  as

$$p_{\text{target}}(x) = \frac{e^{-E(x)}}{\mathcal{Z}}. \quad (187)$$

The *Double Well Potential (DW-4)* models a system of  $n = 4$  particles in  $\mathbb{R}^2$  ( $d = 8$ ). The energy is defined by a pairwise distance potential

$$E_{\text{DW}}(x) = \frac{1}{\tau} \sum_{i < j} [a(d_{ij} - d_0) + b(d_{ij} - d_0)^2 + c(d_{ij} - d_0)^4], \quad (188)$$

where  $d_{ij} = \|x_i - x_j\|_2$ . We set parameters  $a = 0, b = -4, c = 0.9$ , and temperature  $\tau = 1$ .

The *Lennard-Jones (LJ)* potential describes the intermolecular interactions of  $n$  particles in  $\mathbb{R}^3$ , where we consider systems of  $n = 13$  particles ( $d = 39$ ) and  $n = 55$  particles ( $d = 165$ ). The total energy




Figure 4: Samples obtained with BMS, projected on the first two dimensions of a GMM with  $d = 100$  and 100 modes.

Table 4: Comparison on  $n$ -body particle systems across varying dimensions  $d$ . We report the Wasserstein-2 distance between model and long run MCMC samples ( $W_2$ ) and energy values ( $E(\cdot) W_2$ ). Results are averaged over three random seeds, with the best performing method highlighted in bold. Baseline results are quoted from Liu et al. (2025).

Method	DW-4 ( $d = 8$ )		LJ-13 ( $d = 39$ )		LJ-55 ( $d = 165$ )	
	$W_2 \downarrow$	$E(\cdot) W_2 \downarrow$	$W_2 \downarrow$	$E(\cdot) W_2 \downarrow$	$W_2 \downarrow$	$E(\cdot) W_2 \downarrow$
PIS Zhang & Chen (2022)	$0.68 \pm 0.28$	$0.65 \pm 0.25$	$1.93 \pm 0.07$	$18.02 \pm 1.12$	$4.79 \pm 0.45$	$228.70 \pm 131.27$
DDS Vargas et al. (2023a)	$0.92 \pm 0.11$	$0.90 \pm 0.37$	$1.99 \pm 0.13$	$24.61 \pm 8.99$	$4.60 \pm 0.09$	$173.09 \pm 18.01$
iDEM Akhound-Sadegh et al. (2024)	$0.70 \pm 0.06$	$0.55 \pm 0.14$	$1.61 \pm 0.01$	$30.78 \pm 24.46$	$4.69 \pm 1.52$	$93.53 \pm 16.31$
AS Havens et al. (2025)	$0.62 \pm 0.06$	$0.55 \pm 0.12$	$1.67 \pm 0.01$	$2.40 \pm 1.25$	$4.50 \pm 0.05$	$58.04 \pm 20.98$
ASBS Liu et al. (2025)	$0.43 \pm 0.05$	<b><math>0.20 \pm 0.11</math></b>	$1.59 \pm 0.03$	$1.99 \pm 1.01$	$4.00 \pm 0.03$	$28.10 \pm 8.15$
<b>BMS (ours)</b>	<b><math>0.38 \pm 0.01</math></b>	$0.21 \pm 0.12$	<b><math>1.54 \pm 0.00</math></b>	<b><math>0.53 \pm 0.04</math></b>	<b><math>3.80 \pm 0.01</math></b>	<b><math>2.36 \pm 0.45</math></b>




Figure 5: Visualization of interatomic distances and energy histograms of different  $n$ -body identical particle systems. The blue and red curves are generated from our BMS method, and compared to the orange and green curves obtained from a long run MCMC.

$E_{\text{Tot}}(x)$  is the sum of the LJ pairwise potential and a harmonic oscillator potential  $E_{\text{osc}}$  used to prevent system drift

$$E_{\text{Tot}}(x) = \frac{\varepsilon}{\tau} \sum_{i < j} \left[ \left( \frac{r_m}{d_{ij}} \right)^{12} - 2 \left( \frac{r_m}{d_{ij}} \right)^6 \right] + c_{\text{osc}} \sum_i \|x_i - x_{\text{COM}}\|^2, \quad (189)$$

where  $x_{\text{COM}}$  denotes the center of mass. Consistent with Köhler et al. (2020) and Akhound-Sadegh et al. (2024), we use  $r_m = 1$ ,  $\tau = 1$ ,  $\varepsilon = 1$ , and an oscillator scale  $c_{\text{osc}} = 1.0$ .

**Evaluation criteria.** For the evaluation criteria, we refer to Havens et al. (2025) (Appendix E.4 Reported Metrics) for detailed information on the evaluation criteria.

**Algorithm specifications.** We integrated our method into the code base of Liu et al. (2025) (see [github.com/facebookresearch/adjoint\\_samplers](https://github.com/facebookresearch/adjoint_samplers)) and re-used their hyperparameter settings and model architecture. As such, we refer to Liu et al. (2025) for a detailed specification of architectures and parameters. For tuning our method, we varied the scale of the initial distribution  $\mathcal{N}(0, \sigma_{\text{init}}^2 I)$  and  $\sigma_{\text{max}}$  of the geometric noise scheduler, i.e.,

$$\sigma(t) = \sigma_{\text{min}} \left( \frac{\sigma_{\text{max}}}{\sigma_{\text{min}}} \right)^{1-t} \sqrt{2 \log \frac{\sigma_{\text{max}}}{\sigma_{\text{min}}}}. \quad (190)$$

Specifically, we tested  $\sigma_{\text{init}} \in \{0.5, 1, 2\}$  and  $\sigma_{\text{max}} \in \{0.5, 1, 1.5, 2\}$ . We obtained the following parameter, which led to the best performance: For DW-4 we got  $\sigma_{\text{init}} = 2$ ;  $\sigma_{\text{max}} = 1.5$ , for LJ-13  $\sigma_{\text{init}} = 1$ ;  $\sigma_{\text{max}} = 1$  and LJ-55  $\sigma_{\text{init}} = 0.5$ ;  $\sigma_{\text{max}} = 0.5$ . Lastly, for all algorithms, we used  $c = \gamma$  when computing  $\xi$ . The results reported are without using damping, i.e.,  $\eta$  is set to zero.

### E.3 MOLECULAR BENCHMARKS: PEPTIDE SYSTEMS

**Target density specifications.** The target density for the Alanine dipeptide and tetrapeptide systems follow the configuration in Nam et al. (2025) and are defined over the configuration space of their 22 and 42 constituent atoms, respectively. In this work, the model is trained directly on the raw Cartesian coordinates of the molecule purely from energy evaluations. The distribution of molecular states is

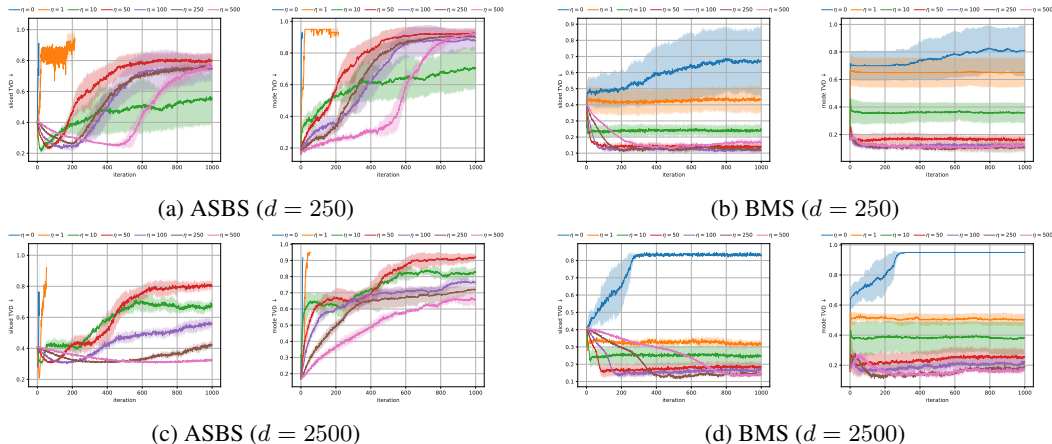


Figure 6: Comparison of *mode TVD* and *sliced TVD* on Gaussian mixture models across varying damping values  $\eta$  for  $d = 250$  and  $d = 2500$  visualized over training iterations. All values are averaged over three random seeds.

governed by the Boltzmann distribution, which defines the unnormalized target density as

$$p_{\text{target}}(x) \propto \exp\left(-\frac{E(x)}{k_B T}\right), \quad (191)$$

where  $E(x)$  represents the potential energy of the molecular configuration,  $k_B$  is the Boltzmann constant, and  $\tau$  is the thermodynamic temperature which is set to  $\tau = 300K$ . The potential energy of the system uses the Amber ff96 force field Kollman et al. (1997) with GBSA-OBC implicit solvation Onufriev et al. (2004) and uses OpenMM Eastman et al. (2023). Moreover, we follow Nam et al. (2025) to constrain the chirality of the generated molecules.

**Evaluation criteria.** To quantify the performance of our sampler on the Alanine dipeptide and tetrapeptide systems, we evaluate the learned density against ground-truth Molecular Dynamics (MD) data using Ramachandran plots, which serve as a standard visualization in structural biology to represent the joint distribution of the backbone dihedral angles  $\phi$  and  $\psi$  Noé et al. (2019). To measure the local backbone structural fidelity, we compared the probability distributions of the backbone dihedral angles,  $\phi$  and  $\psi$ . For each residue  $i$ , we computed the 2D empirical probability density functions,  $P_i(\phi, \psi)$  for the reference MD data and  $P_i^{\text{MD}}(\phi, \psi)$  for the generated samples. These distributions were discretized onto a  $50 \times 50$  grid spanning  $[-\pi, \pi] \times [\pi, \pi]$ . The similarity between the two distributions for the  $i$ th residue was quantified using the Jensen-Shannon divergence, defined as

$$D_{\text{JS}}^i(P_i|P_i^{\text{MD}}) = \frac{1}{2}D_{\text{KL}}(P_i|M_i) + \frac{1}{2}D_{\text{KL}}(P_i^{\text{MD}}|M_i), \quad (192)$$

where  $M_i = \frac{1}{2}(P_i + P_i^{\text{MD}})$  is the mixture distribution, and  $D_{\text{KL}}$  denotes the Kullback-Leibler divergence. The final reported metric is the average Jensen-Shannon divergence across all  $N_{\text{res}}$  residues:

$$(\phi, \psi) - D_{\text{JS}} := \frac{1}{N_{\text{res}}} \sum_{i=1}^{N_{\text{res}}} D_{\text{JS}}(P_i|Q_i). \quad (193)$$

While  $D_{\text{JS}}$  measures the divergence of probability distributions, we also evaluated the thermodynamic alignment by comparing the implied Free-Energy Surfaces (FES) of the backbone dihedrals. Using the same  $50 \times 50$  discrete probability distributions  $P_i(\phi, \psi)$  obtained above, the free energy  $F_i$  for a given bin  $b$  is estimated via the Boltzmann inversion

$$F_i(b) := -k_B T \log(P_i(b) + \varepsilon), \quad (194)$$

where  $k_B T \approx 0.596$  kcal/mol (assuming  $T = 300$  K) and  $\varepsilon = 10^{-10}$  is a small constant added for numerical stability. To account for relative free-energy differences rather than absolute values, each surface was zero-centered by subtracting its global minimum, i.e.,  $F'_i(b) = F_i(b) - \min_b F_i(b)$ . To




Figure 7: TICA plots with  $10^6$  samples for Alanine dipeptide for Molecular Dynamics (MD) data, ASBS Liu et al. (2025), and our proposed method, BMS, for damping values  $\eta \in \{0, 10\}$ .

avoid noise from highly under-sampled, high-energy regions, we restricted our comparison to a valid subspace  $\mathcal{M}_i$ . A bin was included in  $\mathcal{M}_i$  if its probability in *either* the reference or the generated ensemble exceeded  $10^{-4}$ . The Root-Mean-Squared Error for residue  $i$  is then computed as

$$\text{RMSE}_{\Delta F, i} = \sqrt{\frac{1}{|\mathcal{M}_i|} \sum_{b \in \mathcal{M}_i} \left( F'_{\text{ref}, i}(b) - F'_{\text{gen}, i}(b) \right)^2}. \quad (195)$$

The overall FES discrepancy is reported as the mean  $\text{RMSE}_{\Delta F, i}$  across all residues. To capture global conformational dynamics and long-range structural correlations, we evaluated the generated ensembles in a reduced, kinetically relevant latent space. We employed Time-lagged Independent Component Analysis (TICA), which identifies the slowest collective motions in a molecular system by finding projections that maximize the autocorrelation of the time-series MD data at a given lag time. We refer to Tan et al. (2025) for a more detailed description.

All results for  $D_{\text{JS}}$  and  $\text{RMSE}_{\Delta F}$  except those presented in Figure 10 are computed using  $10^6$  model and MD data samples. The latter are taken from Schopmans & Friederich (2025). Due to computational constraints, we used  $10^4$  samples for computing the TICA- $W_2$  metric. Lastly, the values in Figure 10 are computed during training on the replay buffer and therefore use  $\approx 16k$  samples.

**Algorithm specifications.** Our implementation builds on the codebase by Nam et al. (2025) (see [github.com/facebookresearch/wt-asbs](https://github.com/facebookresearch/wt-asbs)), where we removed the collective variables as we are interested in the performance without using prior knowledge. We use the same hyperparameter setting for ASBS and our method, as detailed below.

For the model architecture, the PaiNN architecture Schütt et al. (2021) was used, an E(3)-equivariant graph neural network, modified to include time-conditioning inputs and vector product layers Schreiner et al. (2023) to break parity symmetry, effectively restricting the model to SE(3)-equivariance Liao et al. (2023). To improve training stability, equivariant layer normalization is applied. For ALA2 with  $\eta = 0$ , we follow the setting of Nam et al. (2025) and employ 4 layers, a hidden dimension of 128 and 250 diffusion steps, as we were facing numerical instabilities when using larger architectures. For  $\eta = 10$ , we used 6 layers, a hidden dimension of 256, and 50 diffusion steps, training stably. For ALA4, we also use 6 layers, a hidden dimension of 256, and 50 diffusion steps. For the damped version we set  $\eta = 50$ . A radial cutoff of 8.0 Å is used for graph construction for all methods and systems. For all algorithms, we used  $c = \gamma$  when computing  $\xi$ . The noise schedule is defined by the EDM variance exploding schedule Karras et al. (2022), i.e.,

$$\sigma(t) = \left[ (1-t)\sigma_{\text{max}}^{1/\rho} + t\sigma_{\text{min}}^{1/\rho} \right]^\rho, \quad (196)$$

with a minimum noise level  $\sigma_{\text{min}} = 0.001$  Å, a maximum noise level  $\sigma_{\text{max}} = 6.0$  Å, and an exponent  $\rho = 3.0$ . The source distribution is modeled as a normal distribution  $\mathcal{N}(0, 1 \text{ Å}^2)$ . For optimization, the AdamW optimizer Loshchilov & Hutter (2017) with a weight decay of  $10^{-3}$ . We employ a cosine annealing learning rate schedule, decaying from an initial learning rate of  $1 \times 10^{-4}$  to a final learning rate of  $3 \times 10^{-5}$ . To stabilize training, energy gradients are clipped at a maximum value of 100.0. The training utilizes a replay buffer with a capacity of 16,384 samples. At each outer step, we draw 2,048 samples from the buffer and perform  $L = 100$  gradient updates. The diffusion process is




Figure 8: Ramachandran (Figure 8a) and TICA (Figure 8b) plots with  $10^6$  samples for Alanine Tetrapeptide for molecular dynamics (MD) data, ASBS Liu et al. (2025), and our proposed method, BMS, for damping values  $\eta = 50$ .




Figure 9: Comparison of *mode TVD* and *Wasserstein-2* ( $W_2$ ) values on Gaussian mixture models across varying dimensions  $d$  and damping values  $\eta$ . Figure 9a shows that ASBS exhibits instability at low damping values (N/A) and performance degradation as dimensionality increases. Figure 9b demonstrates that our method maintains stable training and consistently avoids mode collapse up to  $d = 2500$ , significantly scaling beyond the  $d = 50$  range common in the recent literature. All values are averaged over three random seeds.

integrated using the Euler-Maruyama method with 250 discretization steps. For ALA2 and ALA4, we trained for 5k and 8k outer steps, respectively. For ASBS, we follow Nam et al. (2025) and perform 500 corrector matching steps after every 1000 outer steps. Note that we do not count these corrector matching steps towards the outer iterations, meaning that we provide more computational budget to ASBS. For model selection, we always picked the last checkpoint.

## F FURTHER NUMERICAL RESULTS

Here, we provide additional results that complement the results in Section 3.

### F.1 GAUSSIAN MIXTURE MODELS

In this experiment, we consider Gaussian mixture targets consisting of 20 components with unit variance and means sampled uniformly from  $[-20, 20]^d$ . We assess the models based on their ability to capture all modes and the accuracy of their density fitting. Specifically, we utilize *mode TVD* – the total variation distance between empirical and ground-truth mixture weights – to quantify mode collapse. Additionally, we use the *Wasserstein* ( $W_2$ ) distance to compare the model’s empirical density with ground truth samples from the target.

As shown in Figure 9, our analysis focuses on a comparison with ASBS, as the standard adjoint sampler did not produce stable or competitive results in this setting. While many contemporary studies focus on dimensions ranging from  $d = 2$  to 50, we push the complexity to  $d = 2500$ . We observe that while ASBS is prone to instability at low damping values  $\eta$ , our method remains consistently stable, although it may exhibit some mode collapse. Notably, as the dimensionality increases, ASBS suffers from a clear degradation in performance, whereas our method maintains consistent and robust results across the entire range of tested dimensions. We additionally report the sliced total variation distance (TVD); see Appendix E.1 for further details. The results are shown in Figure 2 and are consistent with the findings in Section 3. Additionally, we test ASBS and our method on an increasing number of modes  $K$  sampled from  $[-K, K]^d$ . The results are shown in Figure 3 and indicate that our method’s performance is less degrading when increasing the number of modes. As sampling multimodal densities in such high dimensions was previously unattainable, we conjecture that damping will be a key ingredient in the future of generative sampling methods.

### F.2 $n$ -BODY IDENTICAL PARTICLE SYSTEMS

We extend our evaluation to  $n$ -body systems of increasing complexity, spanning from DW-4 ( $d = 8$ ) to the high-dimensional LJ-55 ( $d = 165$ ). The results in Table 4 demonstrate that our approach consistently outperforms established baselines, including AS and ASBS, across both geometric *Wasserstein-2* ( $W_2$ ) and energy value distributions ( $E(\cdot) W_2$ ). Critically, the most substantial benefits of our method emerge in the high-dimensional  $d = 165$  regime, where it maintains superior accuracy in energy modeling while other methods begin to degrade. Furthermore, our approach exhibits

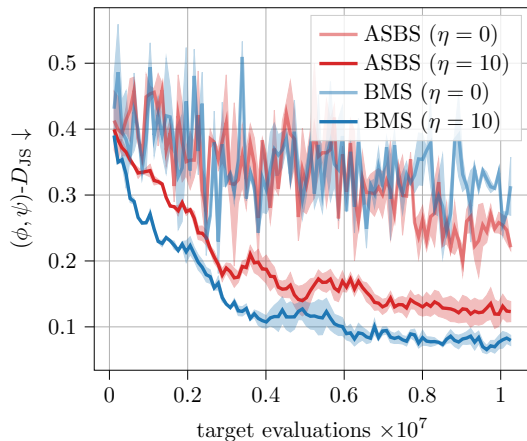


Figure 10: Jensen-Shannon divergence for the joint distribution of backbone dihedral angles  $(\phi, \psi)$ - $D_{JS}$  for Alanine Dipeptide over the course of training for ASBS Liu et al. (2025) and our proposed method, BMS, for damping values  $\eta \in \{0, 10\}$ .

significantly reduced variance across random seeds, underscoring its stability compared to existing approaches. For reference, we provide visualizations of the energy histograms and interatomic distances in Figure 5. Moreover, we visualize the performance over training iterations in Figure 6

**Molecular benchmarks: Peptide systems.** Here, we provide additional visualization for peptide systems. Specifically, Figure 7 shows the TICA plot for ALA2 and Figure 8 shows the TICA and Ramachandran plots for ALA4.