Twisted Schrödinger Bridge Matching

Maxence Noble[*,1] **Marie Scheid**[*,1] **Yazid Janati**[2,3] **Eric Moulines**[3,4] **Alain Durmus**[1]
[1] CMAP, Ecole polytechnique [2] Institute of Foundation Models [3] MBZUAI [4] EPITA, LRE

Abstract

Over the past few years, diffusion-based Schrödinger bridge models have been proposed to approximate optimal transport dynamics between two prescribed boundary distributions, with successful applications to generative modeling. More precisely, these methods aim to estimate a path measure whose initial and terminal marginals match the two boundary distributions, while minimizing the Kullback–Leibler divergence with respect to a reference Markov process. In this work, we consider the generalized Schrödinger bridge problem, in which the reference process is a twisted Brownian motion, that is, a Feynman–Kac transform of a Brownian motion induced by a time-dependent differentiable potential. Building on the Iterative Markovian Fitting (IMF) paradigm, and in particular on its special case Diffusion Schrödinger Bridge Matching (DSBM), which corresponds to the zero potential case, we introduce Twisted Schrödinger Bridge Matching (TSBM), a diffusion-based method designed to handle both continuous- and discrete-time potentials. Unlike previous approaches, TSBM provides a rigorous extension of the IMF scheme to the generalized Schrödinger bridge problem. This derivation leads to a new bridge-matching loss that depends explicitly on the gradient of the potential and recovers the DSBM objective when the potential vanishes, yielding improved performance. We further introduce trajectory-based variance-reduction techniques that substantially stabilize optimization and may be useful beyond the present setting. Finally, we empirically demonstrate the benefits of TSBM for trajectory inference on a variety of experimental settings including crowd navigation, single-cell data, opinion depolarization and unpaired image translation. Code available at https://github.com/maxencenoble/twisted-sb-matching.

1 Introduction

The problem of learning a transport map between two probability distributions, a source μ and a target ν, is central in modern machine learning, with prominent applications in generative modeling. Recent *bridge matching* methods [1] address this problem dynamically, by learning a Markov process on a time interval $[0, T]$ that transports μ to ν. In continuous time, this typically amounts to learning the drift of a diffusion process. The construction starts from an interpolation process $(\mathbf{X}_t)_{t\in[0,T]}$ whose endpoints follow a coupling $\pi_{0,T}$ of (μ, ν), while its conditional law given $(\mathbf{X}_0, \mathbf{X}_T)$ is fixed and tractable, often chosen as a Brownian bridge. Under suitable conditions, the induced density path admits a Markov projection, and learning its drift reduces to a mean-square regression problem. This perspective provides a unified view of several diffusion and flow-based generative models [51, 28].

Bridge matching also recovers optimal-transport dynamics for appropriate choices of the endpoint coupling. In particular, if $\pi_{0,T}$ is the entropy-regularized optimal transport (EOT) plan between μ and ν, and if the conditional interpolation is a Brownian bridge, then the Markov projection solves the dynamical counterpart of EOT; see, e.g., [4]. Equivalently, it solves the canonical Schrödinger bridge (SB) problem[1] [47] with Brownian motion as reference path measure.

* Equal contribution. Corresponding author: maxence.noble@gmail.com
[1]Throughout this paper, we may also refer to this instance as the classic or standard SB problem.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

In general, the SB problem provides a variational framework for trajectory inference. Given a reference path measure $\mathbb{Q}^{\mathrm{ref}}$, it seeks a path measure $\mathbb{P}$ minimizing $\mathrm{KL}(\mathbb{P} \mid \mathbb{Q}^{\mathrm{ref}})$ subject to the endpoint constraints $\mathbb{P}_0 = \mu$ and $\mathbb{P}_T = \nu$. Under appropriate conditions, the solution is the unique path measure satisfying these marginal constraints, being Markov, and belonging to the reciprocal class of $\mathbb{Q}^{\mathrm{ref}}$, meaning that it has the same conditional law as $\mathbb{Q}^{\mathrm{ref}}$ given its endpoints [27]. This characterization has motivated fixed-point algorithms based on KL projections onto subsets of these constraints. In particular, the Iterative Markovian Fitting (IMF) scheme alternates reciprocal and Markovian projections, and its bridge-matching implementation in the canonical setting is Diffusion Schrödinger Bridge Matching (DSBM) [49].

In practice however, the canonical SB formulation, which takes $\mathbb{Q}^{\mathrm{ref}}$ to be Brownian motion, may be too restrictive for trajectory inference when additional information is available at intermediate times. For instance, sparse observations may provide useful information about the dynamics without being sufficient to prescribe intermediate marginal distributions [53], thereby limiting the applicability of multi-marginal SB formulations [9, 37]. This motivates considering more general reference path measures, obtained by tilting Brownian motion through a spatio-temporal potential that encodes such information (see Section 2 for a formal description). We refer to this as the *generalized* SB setting, for which the resulting problem no longer admits, in general, a tractable reduction to an EOT formulation. In this direction, [30] study a stochastic optimal control (SOC) formulation that augments the canonical kinetic-energy cost of the drift with a time-dependent state cost induced by the potential. They introduce Generalized Schrödinger Bridge Matching (GSBM), an iterative bridge-matching algorithm inspired by DSBM. Despite its strong empirical performance, the relationship between GSBM updates and the IMF scheme, which can in principle be extended to general reference path measures, remains unclear. This work addresses this gap.

Contributions. In this paper, we focus on the generalized SB problem introduced by [30] and make the following contributions:

• We reformulate this generalized SB problem through the lens of the *Twisted Schrödinger Bridge* (TSB), where the reference process is a twisted Brownian motion, i.e., a Feynman–Kac transform of Brownian motion. This formulation covers both continuous- and discrete-time state costs. We then extend the IMF scheme to each setting, obtaining update rules that differ from GSBM. In particular, we cast the reciprocal projection as a tractable variational problem with respect to the twisted Brownian bridge, and derive a Markovian projection loss that explicitly involves the state-cost gradient. Notably, unlike GSBM, this loss exactly recovers standard diffusion-model objectives in the zero-potential case: the DSBM loss in continuous time and the DDPM loss [20] in discrete time. Our Markovian objective thus yields a novel diffusion loss for learning stochastic dynamics penalized by a prescribed time-dependent potential, requiring only an input coupling of the endpoints. Importantly, this loss can also be employed independently of the IMF framework.

• Building on this formulation, we introduce *Twisted Schrödinger Bridge Matching* (TSBM), a methodology for solving the TSB problem from samples of the initial and terminal marginals. A key component is a learnable control-variate scheme that reduces variance and improves the estimation of the Markovian projection. We further adapt the efficient DSBM implementation from [5] to obtain a bidirectional version of TSBM that stabilizes the training procedure.

• We empirically validate TSBM on crowd-navigation and single-cell modeling tasks across multiple dimensions, as well as on image-to-image translation and opinion depolarization. Our results show that TSBM improves the estimation of both reciprocal and Markovian projections compared with GSBM, while achieving comparable or better performance on the SB-based objective.

Notation. For readability, we collect the notation used throughout the paper in Section A.

In the rest of the paper, we fix a scaling hyperparameter $\sigma > 0$ and denote by $\mathbb{Q}^{\sigma}$ the distribution of the process $(\sigma \mathbf{B}_t)_{t\in[0,T]}$, where $(\mathbf{B}_t)_{t\geq 0}$ is a standard Brownian motion.

2 From Standard to Twisted Schrödinger Bridge

Schrödinger Bridge problem. In its most general formulation, the Schrödinger Bridge (SB) problem consists in finding a path measure $\mathbb{P}^\star \in \mathscr{P}(\mathbf{C}_T)$ solving the constrained variational problem

$$\arg\min\{\mathrm{KL}(\mathbb{P}\|\mathbb{Q}^{\mathrm{ref}}) \;:\; \mathbb{P} \in \mathscr{P}(\mathbf{C}_T), \mathbb{P}_0 = \mu,\ \mathbb{P}_T = \nu\}\,. \tag{1}$$

Here, $\mathbb{Q}^{\text{ref}} \in \mathscr{P}(\mathbf{C}_T)$ denotes a reference path measure. From its original formulation [47] to modern developments [4], the SB problem (1) has primarily been studied with Brownian motion as the reference measure. In this setting, the conditional measure $\mathbb{Q}^{\text{ref}}_{|0,T}$ reduces to the Brownian bridge, which is tractable and can be easily sampled from.

Moreover, [27] showed in this case that the path measure formulation (1) is equivalent to the following affine Stochastic Optimal Control (SOC) problem

$$\arg\min\{\int_0^T \tfrac{1}{2}\mathbb{E}[\|u_t(\mathbf{X}^u_t)\|^2]\mathrm{d}t \,:\, u \in \mathbf{A}_T(\nu)\,, \quad (\mathbf{X}^u_t)_{t\in[0,T]} \text{ solves (3)}\}\,, \tag{2}$$

$$\mathrm{d}\mathbf{X}^u_t = u_t(\mathbf{X}^u_t)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t\,, \mathbf{X}^u_0 \sim \mu\,, \tag{3}$$

where $\mathbf{A}_T(\nu) = \{u \in \mathbf{A}_T : \mathbf{X}^u_T \sim \nu\}$ is the subset of admissible control functions $u \in \mathbf{A}_T$ (see details in Section C) enforcing the terminal marginal constraint. Then, the optimal path measure $\mathbb{P}^\star$ is induced by $\mathbf{X}^{u^\star}$, where $u^\star$ solves (2) [27, Proposition 2.10]. In the limit $\sigma \to 0$, the entropy regularization in EOT vanishes, yielding classical quadratic OT [42]. Correspondingly, (2) converges to the Benamou–Brenier formulation [2], its dynamical counterpart. We therefore refer to σ as the (entropic) regularization hyperparameter.

Twisted instance : continuous. Let $V \in \mathrm{C}^{0,1}([0,T]\times\mathbb{R}^d,\mathbb{R})$ be an arbitrary time-dependent potential. We study the instance of (1) with reference measure $\mathbb{Q}^{\text{ref}} = \mathbb{Q}^{V,\sigma}$, where $\mathbb{Q}^{V,\sigma}$ is absolutely continuous with respect to $\mathbb{Q}^\sigma$, with Radon–Nikodym derivative

$$\mathrm{d}\mathbb{Q}^{V,\sigma}(x_{[0:T]}) = \tfrac{1}{Z_V}\exp\left(-\int_0^T \tfrac{V_t(x_t)}{\sigma^2}\mathrm{d}t\right)\mathrm{d}\mathbb{Q}^\sigma(x_{[0:T]})\,, \tag{4}$$

with normalizing constant $Z_V = \int \exp\left(-\int_0^T \tfrac{V_t(x_t)}{\sigma^2}\mathrm{d}t\right)\mathrm{d}\mathbb{Q}^\sigma(x_{[0:T]})$. Such transforms of Brownian motion are referred to under various names in the probability literature, including Feynman–Kac penalization [45], general h-transform [25], exponential tilting [10], and exponential reweighting [38]. In this work, we say that $\mathbb{Q}^{V,\sigma}$ is a *twisted* version of $\mathbb{Q}^\sigma$, and we adopt the same terminology for the associated bridges, *i.e.*, $\mathbb{Q}^{V,\sigma}_{|0,T}$ is a twisted version of $\mathbb{Q}^\sigma_{|0,T}$. We refer to the resulting SB problem as the Twisted SB (TSB) problem. This form of path distribution naturally appears in filtering theory as the posterior distribution of a signal given some continuous-time observations [12, 17, 22].

Assuming $Z_V < \infty$, one can show that $\mathbb{Q}^{V,\sigma}$ is associated with a Stochastic Differential Equation (SDE) that can be written as an affine-controlled Brownian motion, although the corresponding drift is generally intractable for arbitrary V, see Section C.1. Therefore, using Girsanov's formula, solving the TSB problem is equivalent to solving the affine SOC problem introduced by [30], namely

$$\arg\min\{\int_0^T \mathbb{E}[\tfrac{1}{2}\|u_t(\mathbf{X}^u_t)\|^2 + V_t(\mathbf{X}^u_t)]\mathrm{d}t \,:\, u \in \mathbf{A}_T(\nu)\,,\, (\mathbf{X}^u_t)_{t\in[0,T]} \text{ solves (3)}\}\,. \tag{5}$$

As in the canonical setting, the optimal TSB path measure $\mathbb{P}^\star$ is induced by $\mathbf{X}^{u^\star}$, where $u^\star$ solves (5). From a physical standpoint, (5) can be interpreted as a trajectory inference problem with marginal constraints at endpoints, where V represents a penalty to be minimized along the dynamics in addition to the quadratic control cost. In particular, the canonical setting is recovered when $V \equiv 0$, and the formulation can be solved in closed form when V is quadratic [54].

Twisted instance : discrete. We may also consider discrete-time state costs. Based on K potentials $\{V_k\}_{k=1}^K$, each in $\mathrm{C}^1(\mathbb{R}^d,\mathbb{R})$, we may define $\mathbb{Q}^{V,\sigma}$ as

$$\mathrm{d}\mathbb{Q}^{V,\sigma}(x_{[0:T]}) = \tfrac{1}{Z_V}\exp\left(-\tfrac{1}{\sigma^2}\textstyle\sum_{k=1}^K V_k(x_{t_k})\right)\mathrm{d}\mathbb{Q}^\sigma(x_{[0:T]})\,, \tag{6}$$

where $Z_V = \int \exp\left(-\tfrac{1}{\sigma^2}\sum_{k=1}^K V_k(x_{t_k})\right)\mathrm{d}\mathbb{Q}^\sigma(x_{[0:T]})$, with the assumption $Z_V < \infty$. As in the continuous-time case, this class of path measures is well motivated by filtering theory [52, 13]. Analogously, the optimal TSB path measure $\mathbb{P}^\star$ solution of (1) for $\mathbb{Q}^{\text{ref}}$ given by (6) is induced by $\mathbf{X}^{u^\star}$, where $u^\star$ solves the SOC problem with the potentials $\{V_k\}_{k=1}^K$ as intermediate state cost functions, defined by

$$\arg\min\{\int_0^T \mathbb{E}[\tfrac{1}{2}\|u_t(\mathbf{X}^u_t)\|^2]\mathrm{d}t + \textstyle\sum_{k=1}^K \mathbb{E}[V_k(\mathbf{X}^u_{t_k})] \,:\, u \in \mathbf{A}_T(\nu), (\mathbf{X}^u_t)_{t\in[0,T]} \text{ solves (3)}\}. \tag{7}$$

In this setting, one may also replace the Brownian path measure by its discrete-time projection, i.e., its finite-dimensional joint distribution, leading to a fully discrete-time reformulation of the SB problem (1) under twisting penalization. This framework has received limited attention in SB-related works, but was recently studied by [18], together with a diffusion-model-based implementation, *in the canonical case only*. We extend this viewpoint to twisted reference measures in Section D, where the fully discrete-time framework is deferred due to space constraints.

Iterative Markovian Fitting & Application to the canonical SB problem. Under mild assumptions on $\mathbb{Q}^{\text{ref}}$, μ, and ν, [27] established that $\mathbb{P}^\star$ is the unique Markov path measure satisfying $\mathbb{P}^\star = \mathbb{P}^\star_{0,T}\mathbb{Q}^{\text{ref}}_{|0,T}$ together with the marginal constraints $\mathbb{P}^\star_0 = \mu$ and $\mathbb{P}^\star_T = \nu$. This characterization naturally motivates the Iterative Markovian Fitting (IMF) scheme [49], a fixed-point procedure inspired by iterative Bregman projections [6], which we now briefly recall.

Let $\mathbb{P} \in \mathscr{P}(\mathbf{C}_T)$. The *reciprocal projection* of $\mathbb{P}$ with respect to $\mathbb{Q}^{\text{ref}}$, denoted $\text{proj}_{\mathcal{R}(\mathbb{Q}^{\text{ref}})}(\mathbb{P}) \in \mathscr{P}(\mathbf{C}_T)$, is defined as

$$\arg\min\{\text{KL}(\mathbb{P}\|\Pi) \;:\; \Pi \in \mathcal{R}(\mathbb{Q}^{\text{ref}})\}\;, \tag{8}$$

where $\mathcal{R}(\mathbb{Q}^{\text{ref}})$ denotes the reciprocal class of $\mathbb{Q}^{\text{ref}}$, *i.e.*, the set of path measures of the form $\Pi = \Pi_{0,T}\mathbb{Q}^{\text{ref}}_{|0,T}$. Importantly, this projection admits a closed-form expression via the KL chain rule [26],

$$\text{proj}_{\mathcal{R}(\mathbb{Q}^{\text{ref}})}(\mathbb{P}) = \mathbb{P}_{0,T}\mathbb{Q}^{\text{ref}}_{|0,T}\;. \tag{9}$$

On the other hand, the *Markovian projection* of $\mathbb{P}$, denoted $\text{proj}_{\mathcal{M}_T}(\mathbb{P}) \in \mathscr{P}(\mathbf{C}_T)$, is defined as

$$\arg\min\{\text{KL}(\mathbb{P}\|\mathbb{M}) \;:\; \mathbb{M} \in \mathcal{M}_T\}\;, \tag{10}$$

where $\mathcal{M}_T$ is the class of Markov diffusion path measures, namely those induced by affine-control diffusion processes solving the SDE (3) for some admissible control u. In the canonical setting, it can be shown (see e.g., [50]) that, for any $\Pi \in \mathcal{R}(\mathbb{Q}^\sigma)$, the Markovian projection $\text{proj}_{\mathcal{M}_T}(\Pi)$ preserves the marginals of Π and is associated with the forward dynamics $\mathrm{d}\mathbf{X}^{\text{f}}_t = \text{b}^{\text{f}}_t(\mathbf{X}^{\text{f}}_t)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t$, with initial condition $\mathbf{X}^{\text{f}}_0 \sim \mu$, where $\text{b}^{\text{f}}_t(x_t) = \sigma^2\mathbb{E}_{\Pi_{T|t}}[\nabla_{\mathbf{X}_t}\log\mathbb{Q}^\sigma_{T|t}(\mathbf{X}_T|\mathbf{X}_t) \mid \mathbf{X}_t = x_t]$. Moreover, $\text{proj}_{\mathcal{M}_T}(\Pi)$ is the time reversal of the path measure induced by $\mathrm{d}\mathbf{X}^{\text{b}}_t = \text{b}^{\text{b}}_{T-t}(\mathbf{X}^{\text{b}}_t)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t$, with $\mathbf{X}^{\text{b}}_0 \sim \nu$, where $\text{b}^{\text{b}}_t(x_t) = \sigma^2\mathbb{E}_{\Pi_{0|t}}[\nabla_{\mathbf{X}_t}\log\mathbb{Q}^\sigma_{t|0}(\mathbf{X}_t|\mathbf{X}_0) \mid \mathbf{X}_t = x_t]$. Since the transition kernels of $\mathbb{Q}^\sigma$ are available in closed form, the drifts can be further simplified as

$$\text{b}^{\text{f}}_t(x_t) = (\mathbb{E}_{\Pi_{T|t}}[\mathbf{X}_T \mid \mathbf{X}_t = x_t] - x_t)/(T-t)\;,\;\; \text{b}^{\text{b}}_t(x_t) = (\mathbb{E}_{\Pi_{0|t}}[\mathbf{X}_0 \mid \mathbf{X}_t = x_t] - x_t)/t\;. \tag{11}$$

The ideal IMF scheme alternates between these two KL projections, thereby enforcing both the Markov and reciprocal constraints satisfied by $\mathbb{P}^\star$. Starting from an arbitrary $\mathbb{P}^{-1} \in \mathscr{P}(\mathbf{C}_T)$ with correct marginals $\mathbb{P}^{-1}_0 = \mu$ and $\mathbb{P}^{-1}_T = \nu$, one defines recursively for $n \in \mathbb{N}$

$$\mathbb{P}^{2n} = \text{proj}_{\mathcal{R}(\mathbb{Q}^{\text{ref}})}(\mathbb{P}^{2n-1})\;,\quad \mathbb{P}^{2n+1} = \text{proj}_{\mathcal{M}_T}(\mathbb{P}^{2n})\;. \tag{12}$$

By construction, the marginal constraints are preserved at each iteration. Recently, [50] provided evidence of exponential convergence of IMF sequence $(\mathbb{P}^n)_{n\in\mathbb{N}}$ toward $\mathbb{P}^\star$ under some assumptions.

The main computational bottleneck in implementing IMF in the standard case $\mathbb{Q}^{\text{ref}} = \mathbb{Q}^\sigma$, is the Markovian projection due to the fact that the controls of the associated path measures are intractable. To address this problem, [49, 40] independently cast IMF as a machine-learning procedure for approximating these controls; following [49], we refer to this approach as DSBM. Given a path measure $\Pi \in \mathcal{R}(\mathbb{Q}^\sigma)$, DSBM exploits the conditional-expectation form of the forward and backward controls b^{f} and b^{b} that define $\text{proj}_{\mathcal{M}_T}(\Pi)$, see (11), to estimate them through the regression problems

$$\arg\min\{\textstyle\int_0^T \mathbb{E}_{\Pi_{t,T}}[\|(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - v^{\theta^{\text{f}}}_t(\mathbf{X}_t)\|^2]\mathrm{d}t \;:\; \theta^{\text{f}} \in \Theta\}\;, \tag{13}$$

$$\arg\min\{\textstyle\int_0^T \mathbb{E}_{\Pi_{0,t}}[\|(\mathbf{X}_0 - \mathbf{X}_t)/t - v^{\theta^{\text{b}}}_t(\mathbf{X}_t)\|^2]\mathrm{d}t \;:\; \theta^{\text{b}} \in \Theta\}\;, \tag{14}$$

where $\{(v^{\theta^{\text{f}}}_t)_{t\in[0,T]} \;:\; \theta^{\text{f}}\}$ and $\{(v^{\theta^{\text{b}}}_t)_{t\in[0,T]} \;:\; \theta^{\text{b}}\}$ are two families of parameterized drift functions.

The DSBM objectives (13) and (14) are fully tractable. Indeed, since $\Pi \in \mathcal{R}(\mathbb{Q}^\sigma)$, samples from $\Pi_{t,T}$ or $\Pi_{0,t}$ can be generated from samples of $\Pi_{0,T}$ through Brownian bridges. After fitting parameters $\theta^{\text{f}}_\star$ and $\theta^{\text{b}}_\star$, the Markovian projection of Π is approximated either by the path measure induced by $\mathrm{d}\hat{\mathbf{X}}^{\text{f}}_t = v^{\theta^{\text{f}}_\star}_t(\hat{\mathbf{X}}^{\text{f}}_t)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t$, with $\hat{\mathbf{X}}^{\text{f}}_0 \sim \mu$, or, in the backward direction, by the time reversal of the path measure induced by $\mathrm{d}\hat{\mathbf{X}}^{\text{b}}_t = v^{\theta^{\text{b}}_\star}_{T-t}(\hat{\mathbf{X}}^{\text{b}}_t)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t$, with $\hat{\mathbf{X}}^{\text{b}}_0 \sim \nu$. Thus, in the canonical setting, IMF reduces to an iterative bridge-matching procedure with Brownian bridge as stochastic interpolant [1]. In practice, to mitigate the bias caused by imperfect marginal matching, [49] alternate between the forward loss (13) and the backward loss (14) when performing Markovian projection steps. This requires two neural networks, one for each drift direction.

3 Twisted Schrödinger Bridge Matching

A natural approach to solving the SB problem (1) with $\mathbb{Q}^{\text{ref}} = \mathbb{Q}^{V,\sigma}$—equivalently, (5) in continuous time or (7) in discrete time—is to derive the IMF scheme associated with the TSB formulation, thereby extending the algorithm originally introduced for the canonical SB problem. We therefore describe how the reciprocal and Markovian projections can be implemented in this setting. This practical realization of IMF is what we call *Twisted Schrödinger Bridge Matching* (TSBM); its updates differ from those of GSBM, as discussed in Section 4. Due to space constraints, we focus on the forward continuous-time setting and defer the backward continuous-time formulation, discrete-time results, and proofs of all propositions to Section C. Pseudo-codes are given in Section F.1.

3.1 Computation of the reciprocal projection

Let $\mathbb{P} \in \mathcal{M}_T$, and assume that we have access to samples from $\mathbb{P}_{0,T}$[2]. In principle, computing the reciprocal projection of $\mathbb{P}$ presents no difficulty, since it is explicit and given by (9). However, a key difference arises compared to the canonical setting. Although we show in Section C.1 that the twisted bridge $\mathbb{Q}^{V,\sigma}_{|0,T}$ is Markov and corresponds to a diffusion process, its drift is not tractable in general. This prevents efficient sampling from $\mathbb{Q}^{V,\sigma}_{|0,T}$, in contrast to the Brownian bridge $\mathbb{Q}^{\sigma}_{|0,T}$. To address this limitation, we approximate the twisted bridge $\mathbb{Q}^{V,\sigma}_{|0,T}$ by solving a reverse-KL problem

$$\arg\min_{\psi\in\Psi} \mathbb{E}_{\mathbb{P}_{0,T}}[\mathrm{KL}(\mathbb{Q}^{\psi}_{|0,T} \| \mathbb{Q}^{V,\sigma}_{|0,T})] \,. \tag{15}$$

Here, $\{\mathbb{Q}^{\psi}_{|0,T}\}_{\psi\in\Psi}$ denotes a parametric family of Markov bridges obtained from affine-controlled Brownian motions, with drift parameterized by ψ. This choice mirrors the structure of $\mathbb{Q}^{V,\sigma}_{|0,T}$, which itself admits an affine-controlled Brownian representation; see Section C.1 for details. Crucially, the resulting variational problem admits a tractable objective that can be implemented in practice. Below, we state the continuous-time result in the forward direction below, for $\mathbb{Q}^{V,\sigma}$ defined by (4).

Proposition 1 (Forward reciprocal projection – continuous time)**.** *Assume that the variational family $\{\mathbb{Q}^{\psi}_{|0,T}\}_{\psi\in\Psi}$ in (15) describes the pinned diffusion processes associated, for almost surely any $(x_0, x_T) \sim \mathbb{P}_{0,T}$, to forward-time parametric SDEs of the form*

$$\mathrm{d}\mathbf{X}^{0,T}_t = v^{\mathrm{f},\psi}_t(\mathbf{X}^{0,T}_t|x_0, x_T)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \,,\ \mathbf{X}^{0,T}_0 = x_0,\ \mathbf{X}^{0,T}_T = x_T \,,$$

with $v^{\mathrm{f},\psi} \in \mathbf{A}_T$. Define $\Pi^{\psi} = \mathbb{P}_{0,T}\mathbb{Q}^{\psi}_{|0,T}$. Then, the minimizer of (15), denoted by $\psi_\star$, is also the minimizer of the following loss

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{R}}(\psi) = \int_0^T \mathbb{E}_{\Pi^{\psi}_{0,t,T}}\left[\tfrac{1}{2}\|(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - v^{\mathrm{f},\psi}_t(\mathbf{X}_t|\mathbf{X}_0, \mathbf{X}_T)\|^2 + V_t(\mathbf{X}_t)\right]\mathrm{d}t \,. \tag{16}$$

Practical implementation. In practice, the variational bridge $\mathbb{Q}^{\psi}_{|0,T}$ should meet three requirements: (1) it should be easy to sample from its marginals, (2) it should allow reparameterized sampling for efficient backpropagation, and (3) it should admit tractable expressions for the forward drift $v^{\mathrm{f},\psi}$. A natural choice satisfying these properties is a stochastic interpolant [1]. In this case, $\mathbb{Q}^{\psi}_{t|0,T}$ is Gaussian, which enables efficient sampling and reparameterization, while the associated forward and backward drifts are available in closed form. To increase expressivity, one may also use mixtures of stochastic interpolants, as proposed by [15]. To reduce the cost of this variational stage, we follow the spirit of [30] and use a spline parameterization of the interpolants rather than neural networks. This allows us to reuse stored control-point trajectories when sampling from the coupling $\Pi_{0,T}$ while enjoying lightweight optimization. For more details, we refer the reader to Section F.2.

3.2 Computation of the Markovian projection

Proposition 2 characterizes the Markovian projection step of the IMF scheme for the continuous-time TSB problem in both the forward and backward directions. In particular, it shows that the corresponding drift terms explicitly depend on the state-cost gradient ∇V.

[2] In practice, this can be achieved by simulating the SDE associated with $\mathbb{P}$ in either the forward or backward direction, as in standard bridge matching approaches.

Proposition 2. *For any $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$, the path measure solution of (10) for $\mathbb{P} = \Pi$, is given by the distribution associated to the forward dynamics $\mathrm{d}\mathbf{X}_t^{\mathrm{f}} = \mathrm{b}_t^{\mathrm{f},V}(\mathbf{X}_t^{\mathrm{f}})\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t, \mathbf{X}_0^{\mathrm{f}} \sim \mu$, where*

$$\mathrm{b}_t^{\mathrm{f},V}(x_t) = \mathbb{E}_{\Pi_{|t}}[(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - (T-t)^{-1}\int_t^T (T-s)\nabla V_s(\mathbf{X}_s)|\mathbf{X}_t = x_t] \,. \tag{17}$$

Moreover, for any $t \in [0,T]$, we have $\mathrm{Law}(\mathbf{X}_t^{\mathrm{f}}) = \Pi_t$.

Informally, the forward optimal control can be viewed as a state-cost perturbation of the canonical forward control in (11), which is recovered when the potential vanishes. This perturbation is driven by $-\nabla V$ evaluated along the corresponding continuation of the trajectory: over future times $s > t$ in the forward direction, and over past times $s < t$ in the backward direction, see Section C.3. Note that the optimal control in (17) involves a state-cost integral, which would be costly to estimate directly. However, we show in Section C.3 that it admits a tractable regression characterization. This motivates the following Markovian projection loss for any $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$, defined for forward control family $\{(v^{\theta^{\mathrm{f}}})_{t\in[0,T]} : \theta^{\mathrm{f}} \in \Theta\}$:

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}) = \int_0^T (T-t)^{-1} \int_t^T \mathbb{E}_{\Pi_{t,s,T}}[\|\mathbf{b}_{t,s}^{\mathrm{f},V}(\mathbf{X}_t, \mathbf{X}_s, \mathbf{X}_T) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\|^2] \,\mathrm{d}s\mathrm{d}t \,, \tag{18}$$
$$\text{with } \mathbf{b}_{t,s}^{\mathrm{f},V}(x_t, x_s, x_T) = (x_T - x_t)/(T-t) - (T-s)\nabla V_s(x_s) \,,$$

Notably, setting $V \equiv 0$ in the TSBM bridge-matching objective (18) *exactly* recovers the DSBM objective (13). Since our objective does not depend explicitly on σ, its regression target remains unchanged in the deterministic limit $\sigma = 0$, thereby yielding in principle a twisted analogue of Rectified Flow [31].

Variance reduction via learnable control variates. In practice, the regression target $\mathbf{b}^{\mathrm{f},V}$ used in the training objective (18) can exhibit high variance, due to noisy evaluations of $\nabla V(X_s)$ along the intermediate states $(\mathbf{X}_s)_{s\in[t,T]}$. To address this issue, we introduce *control variates* into the regression target, reducing in principle its variance without introducing bias. This construction relies on the following result.

Proposition 3. *For any $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$ and $\alpha : [0,T]^2 \to \mathbb{R}$, C^1 in its second variable, denoting $\tilde{\alpha}^{\mathrm{f}}(t,s) = (T-s)\alpha(t,s)$, it holds that $\mathbb{E}_{\Pi_{|t}}[H_t^{\alpha,\mathrm{f}}|\mathbf{X}_t] = 0$ where*

$$H_t^{\alpha,\mathrm{f}} = \alpha(t,t)[\mathbf{X}_T - \mathbf{X}_t] + \int_t^T \left\{ \frac{\partial_s \tilde{\alpha}^{\mathrm{f}}(t,s)}{T-s}[\mathbf{X}_T - \mathbf{X}_s] + (T-s)\nabla V_s(\mathbf{X}_s) \int_t^s \frac{\alpha(t,u)}{T-u}\mathrm{d}u \right\} \mathrm{d}s.$$

We can therefore construct *strictly equivalent* representations of the optimal drift (17) from Proposition 2, given for any function α^{f} satisfying the assumptions of Proposition 3 by

$$\mathrm{b}_t^{\alpha^{\mathrm{f}},V}(x_t) = \mathrm{b}_t^{\mathrm{f},V}(x_t) - \mathbb{E}_{\Pi_{|t}}[H_t^{\alpha^{\mathrm{f}},\mathrm{f}} \mid \mathbf{X}_t = x_t].$$

This α-dependent drift still admits a conditional-expectation representation under $\Pi_{|t}$, again involving both the dynamical states and the gradient of the state cost. It thus induces an alternative Markovian projection loss defined as follows:

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}, \alpha^{\mathrm{f}}) = \int_0^T (T-t)^{-1} \int_t^T \mathbb{E}_{\Pi_{t,s,T}}[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t, \mathbf{X}_s, \mathbf{X}_T) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\|^2] \,\mathrm{d}s\mathrm{d}t \,,$$
$$\text{with } \mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(x_t, x_s, x_T) = (x_T - x_t)/(T-t) - \alpha^{\mathrm{f}}(t,t)[x_T - x_t] - \tfrac{\partial_s \tilde{\alpha}^{\mathrm{f}}(t,s)}{T-s}(T-t)[x_T - x_s]$$
$$- \left\{1 + (T-t)\int_t^s \tfrac{\alpha^{\mathrm{f}}(t,u)}{T-u}\mathrm{d}u\right\}(T-s)\nabla V_s(x_s) \,.$$

Although the α-based objective differs from (18), both yield the same optimal drift and are therefore equivalent in terms of optimization. In particular, setting $\alpha^{\mathrm{f}} = 0$, *i.e.*, omitting the control variate, exactly recovers the original objective (18). Moreover, $\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}$ admits a *bias–variance decomposition* into the original TSBM loss, depending only on θ, and an additional term depending only on α, which corresponds exactly to the variance of the α-based regression target. We refer to Section C.3 for details and the corresponding backward formulation involving α^{b}.

Motivated by the bias–variance decomposition, we parameterize α^{f} and α^{b} with lightweight neural networks and learn them jointly with the drifts $v^{\theta^{\mathrm{f}}}$ and $v^{\theta^{\mathrm{b}}}$, by minimizing the objectives with respect to $(\theta^{\mathrm{f}}, \alpha^{\mathrm{f}})$ or $(\theta^{\mathrm{b}}, \alpha^{\mathrm{b}})$. This preserves the optimal drift while stabilizing optimization. The exact parameterization is detailed in Section F.3. A related learnable control-variate strategy was recently proposed by [3], but only for terminal potentials V_T. Our result is more general: it defines a bi-level control variate for arbitrary pairs (t,s), and may therefore be useful beyond our SB setting.

Practical implementation. Given $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$, the joint law $\Pi_{t,s,T}$ appearing in the expectation of the objective cannot be sampled directly, unlike in DSBM. This is precisely because the twisted bridge $\mathbb{Q}^{V,\sigma}_{s,t|0,T}$ is itself not directly simulable, as discussed in Section 3.1. For implementation, we therefore replace Π with its variational approximation $\Pi^{\psi^\star} = \Pi_{0,T}\mathbb{Q}^{\psi^\star}_{|0,T}$, where $\psi^\star$ is obtained at the preceding reciprocal projection step by solving (15) with coupling $\Pi_{0,T}$ using parametric stochastic interpolants. A naive implementation would sample from $\mathbb{Q}^{\psi^\star}_{s,t|0,T}$ by integrating the corresponding pinned SDEs, depending on the direction of the Markovian projection, from time t to s, hence introducing a non-negligible overhead compared with the original DSBM implementation. We avoid this bottleneck by exploiting the spline parameterization, which allows these SDE integrations to be computed almost instantaneously; see Section F.2 for details.

4 Related work

In this section, we provide a detailed comparison between TSBM and Generalized Schrödinger Bridge Matching (GSBM), introduced by [30], which is the closest existing method to ours. We focus on both the underlying framework and the practical implementation for solving the generalized SB problem in the form of (5). Other related works, which are less directly connected to TSBM than GSBM, are discussed in Section B. At a high level, GSBM and TSBM share a similar spirit: both are iterative bridge-matching methods that alternate, in the forward and backward directions, between two steps: approximating the generally intractable twisted bridge, and learning the corresponding Markovian projection. However, unlike TSBM, GSBM does not explicitly establish the connection with the IMF scheme. This leads to different update rules, which we detail below.

Reciprocal projection. Similarly to our approach, [30] consider a variational approximation of the twisted bridge based on stochastic interpolants, which is then used in their bridge-matching loss. However, the related loss derived in [30, Proposition 2] differs from our reciprocal objective (16). Instead, their objective is obtained by turning the SOC objective (2) into a conditional version

$$\mathcal{L}^f_{\mathcal{R}}(\psi) = \int_0^T \mathbb{E}_{\Pi^\psi_{0,t,T}}[\tfrac{1}{2}\|v^{\mathrm{f},\psi}_t(\mathbf{X}_t|\mathbf{X}_0,\mathbf{X}_T)\|^2 + V_t(\mathbf{X}_t)]\mathrm{d}t \;, \tag{19}$$

where $\Pi^\psi_{0,t,T} = \mathbb{Q}^\psi_{t|0,T}\mathbb{P}_{0,T}$. The key difference between (16) and (19) is the absence, in the latter, of the Brownian bridge term in the quadratic control cost used as regression target. As a result, the two objectives lead to different variational solutions under the same stochastic interpolant parameterization.

Our analysis suggests that this discrepancy may arise from an issue in the derivation of [30], which we investigate empirically in the simple zero-potential setting ($V \equiv 0$). In this case, the twisted bridge reduces to the standard Brownian bridge, whose conditional marginals $\mathbb{Q}_{t|0,T}$ are tractable and can be sampled exactly. Using a stochastic interpolant bridge parameterization, we can therefore evaluate $\mathbb{E}_{\mathbb{P}_{0,T}}[\mathrm{KL}(\mathbb{Q}^\psi_{t|0,T}\|\mathbb{Q}^{V,\sigma}_{t|0,T})]$ at any $t \in (0,T)$, providing a measure of the marginal discrepancy between the exact bridge and its variational approximation learned with either the TSBM or GSBM objective. Averaging this quantity over time yields a global measure of approximation error. Across different dimensions and values of σ, we consistently find that the TSBM objective produces more accurate bridge approximations than GSBM, favoring its use in more general settings. Full experimental details and results are provided in Section G.2.

Markovian projection. The main discrepancy between TSBM and GSBM lies in the Markovian projection loss. We detail it here in the forward direction. The GSBM bridge-matching loss stated in [30, Equation 5] is defined as

$$\arg\min\{\int_0^T \mathbb{E}_{\Pi^{\psi_\star}_{0,t,T}}[\|v^{\mathrm{f},\psi_\star}_t(\mathbf{X}_t|\mathbf{X}_0,\mathbf{X}_T) - v^{\theta^\mathrm{f}}_t(\mathbf{X}_t)\|^2]\mathrm{d}t \;:\; \theta^\mathrm{f} \in \Theta\} \;, \tag{20}$$

where $v^{\mathrm{f},\psi_\star}$ solves the variational problem (19) and is used to approximate the preceding reciprocal projection step. As in our practical Markovian objective, the expectation is taken with respect to $\Pi^{\psi_\star}$, obtained by combining the available coupling with the variational bridge rather than with the intractable twisted bridge. However, the regression target is fundamentally different. In GSBM, the loss (20) regresses onto the conditional vector field of the variational interpolant itself. If the variational approximation were exact, this objective would be valid. In the presence of a non-zero variational gap, the regression target is the approximate conditional vector field and may therefore inherit part of the variational approximation error. By contrast, our bridge-matching loss (18) avoids this source of bias, since its regression target is derived from the *exact* conditional velocity field.

Table 1: Results of crowd navigation experiments for Stunnel and Vneck, with increasing d. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

State cost	Stunnel				Vneck			
Metric	Feasibility (↓)		Optimality (↓)		Feasibility (↓)		Optimality (↓)	
Method	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM
$d = 2$	**0.006**$\pm$**0.001**	0.056$\pm$0.007	978$\pm$1.0	**956**$\pm$**1.0**	**0.008**$\pm$**0.004**	0.16$\pm$0.02	69.2$\pm$0.5	**53.1**$\pm$**0.4**
$d = 10$	**1.02**$\pm$**0.01**	1.14$\pm$0.01	762$\pm$0.9	**748**$\pm$**0.8**	**1.02**$\pm$**0.01**	1.50$\pm$0.03	254$\pm$1.0	**252**$\pm$**0.4**
$d = 50$	17.49$\pm$0.08	**17.07**$\pm$**0.06**	3592$\pm$2.9	**3570**$\pm$**2.8**	**16.4**$\pm$**0.2**	21.2$\pm$3.1	1388$\pm$2.5	**1381**$\pm$**2.8**

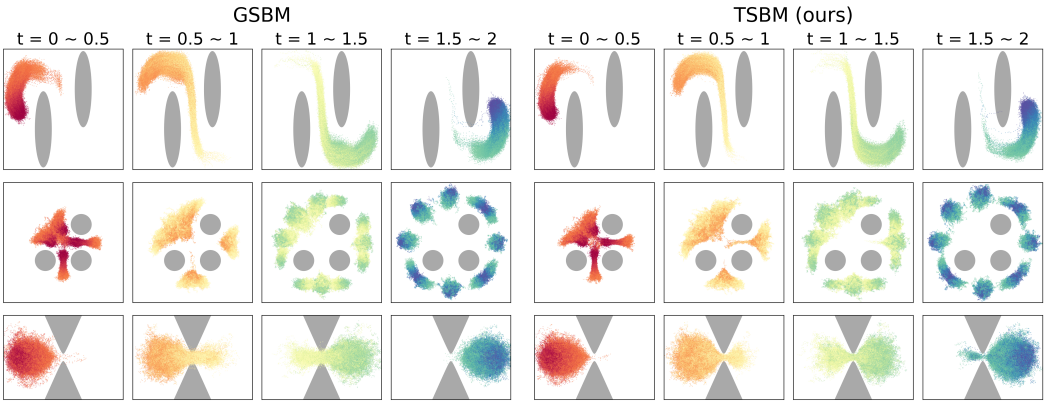

Figure 1: Learned TSB dynamics for Stunnel (top), Gmm (middle), Vneck (bottom) with $d = 2$.

5 Numerical experiments

We consider four trajectory inference tasks within the continuous-time Twisted SB framework[3]: (i) **crowd navigation**, a standard generalized SB benchmark with repulsive state costs [29, 30]; (ii) **single-cell inference** from sparse observations [55], with an attractive state cost, a setting introduced in this work; (iii) **AFHQ image-to-image unpaired translation** [49, 30], with a state cost penalizing deviations from images obtained by decoding latent-space interpolations between source and target; and (iv) **opinion depolarization** [29, 30], with a polarizing drift and a congestion potential penalizing the concentration of opinions. We compare TSBM with GSBM, its closest competitor, using the same computational backbone and, whenever possible, the same main hyperparameters. Our experiments cover several dimensions to assess the scaling behavior of both methods. We evaluate performance along two complementary axes: **optimality**, measured by the SB cost from the SOC formulation (5), and **feasibility**, measured by satisfaction of the terminal marginal constraint.

Classical unidirectional procedure and bidirectional extension. By default, we run 5 outer iterations of GSBM and TSBM, each consisting of one forward followed by one backward Markovian projection. Following [5], we also consider a *bidirectional* extension of both algorithms. Instead of alternating full forward and backward updates from the start, this variant proceeds in two phases. First, in a pretraining phase, the forward and backward Markovian projection losses, along with their spline parameters, are optimized jointly under the independent coupling $\mu \otimes \nu$. Second, in a fine-tuning phase, forward and backward reciprocal/Markovian projections alternate after fewer training epochs than in the classical procedure, set to 10% in our experiments, which is expected to yield smoother convergence. For a fair comparison, unidirectional and bidirectional variants are trained for the same total number of epochs. Full experimental details are provided in Section G.

5.1 Crowd navigation

Following [30], we validate TSBM on crowd-navigation tasks ('Stunnel', 'Gmm', 'Vneck'), modeled as Twisted SB problems with SOC formulation (5), where the state cost encodes physical constraints; see Section G.4 for details on these settings. The resulting trajectories are shown in Figure 1 for $d = 2$ and $\sigma = 2$, where TSBM appears to better capture the obstacle geometry. Complementing these qualitative results, we compare GSBM and TSBM across dimensions in Table 1, again with $\sigma = 2$. The results reveal a consistent optimality–feasibility trade-off: GSBM better matches the terminal marginal, whereas TSBM achieves lower state-cost transport. This suggests that using the state-cost gradient enforces the objective constraint more directly, but also introduces additional variance (even with the control-variate strategy), which can hinder satisfaction of the marginal constraint.

[3]Toy experiments in the discrete-time setting, absent from the experimental study of [30], are reported in Section G.3.

Table 2: Results of single-cell experiments with $\beta \in \{20, 100\}$, $\sigma = 0.1$, with increasing d. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

Hyperparameter	$\beta = 20$				$\beta = 100$			
Metric	Feasibility (↓)		Optimality (↓)		Feasibility (↓)		Optimality (↓)	
Method	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM
$d = 2$	**0.05**$\pm$0.005	**0.05**$\pm$0.003	**7.23**$\pm$0.01	7.25$\pm$0.01	**0.011**$\pm$0.002	**0.011**$\pm$0.001	**6.79**$\pm$0.01	6.82$\pm$0.01
$d = 5$	**3.05**$\pm$0.01	3.06$\pm$0.02	31.05$\pm$0.01	**29.93**$\pm$0.01	2.96$\pm$0.02	**2.81**$\pm$0.02	29.68$\pm$0.02	**29.3**$\pm$0.02
$d = 50$	**38.79**$\pm$0.04	46.31$\pm$0.06	31.75$\pm$0.01	**22.53**$\pm$0.01	**39.99**$\pm$0.03	41.92$\pm$0.04	30.76$\pm$0.02	**26.9**$\pm$0.01

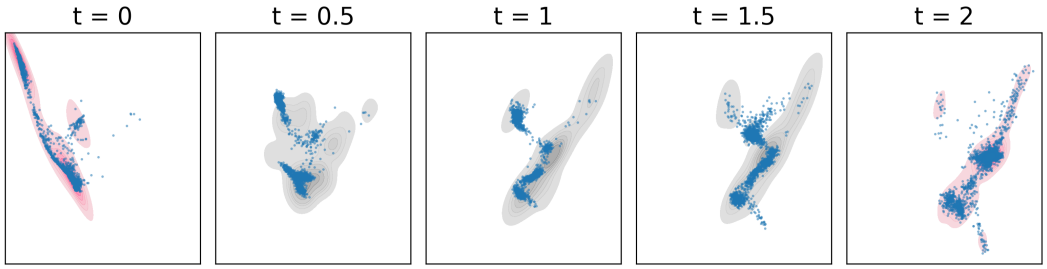

Figure 2: Snapshots of learned TSBM dynamics for 2D single-cell ($\beta = 100$, $\sigma = 0.1$) with respect to available observations (endpoint *marginal* samples in pink, intermediary *sparse* samples in grey).

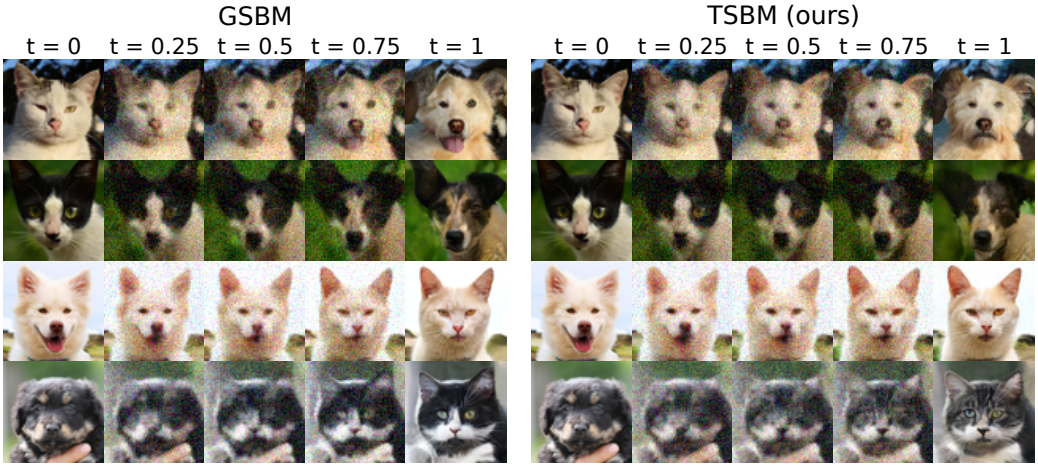

Figure 3: AFHQ translation SB trajectories: cat-to-dog (top two rows) and dog-to-cat (bottom two rows), at times $t \in \{0,\ 0.25,\ 0.5,\ 0.75,\ 1\}$ (left to right).

5.2 Singe-cell inference under sparse observations

We further evaluate the bidirectional variants of TSBM and GSBM on single-cell data [34] in dimension $d \in \{2, 5, 50\}$. We consider a *sparse-observation regime* by intentionally retaining only 25% of the intermediate observations, thereby precluding the training of a multi-marginal model. Based on this data, we design a continuous-time state cost inspired by [55, Eq. 12]. Given snapshot times $\{t_k\}_{k=1}^K$, with $t_1 = 0$ and $t_K = T$, we define V_t as a time-linear interpolation between two static state costs, which attract particles toward the available observations closest to the two snapshot times bracketing $t \in [0, T]$ (the cost being set to zero on boundary times). These static costs depend on a hyperparameter $\beta > 0$ controlling the locality of the interaction: small values of β yield nearly uniform weights over all observations, whereas large values emphasize the nearest ones. We refer to Section G.5 for the detailed expression. We emphasize that this attractive cost is not specific to single-cell data and can be applied more broadly to sparse-observation settings. As shown in Table 2, we observe the same optimality–feasibility trade-off as in crowd navigation, with a more pronounced gap in higher dimensions, while Figure 2 illustrates how sparse observations guide TSBM trajectories.

5.3 High dimensional experiments

We compare TSBM to GSBM in two high-dimensional settings. For both experiments, we follow the experimental setup used in GSBM. Further details are given in Section G.6.

	FID (↓)		Optimality (↓)	
Direction	GSBM	TSBM	GSBM	TSBM
Cat → Dog	**35.20**	41.23	2.76	**2.68**
Dog → Cat	**10.89**	11.89	2.84	**2.73**

Table 3: AFHQ unpaired translation results.

	Feasibility (↓)		Optimality (↓)	
Direction	GSBM	TSBM	GSBM	TSBM
Forward	2423.3	**2412.0**	1143.6	**1100.2**
Backward	**1722.2**	1801.1	1317.6	**867.1**

Table 4: Depolarization results.

AFHQ experiment. We study unpaired translation between cat and dog images from the AFHQ dataset [11], resized to 64×64 ($d = 4096$). Following GSBM, we encode source and target images with a pretrained VAE and obtain a reference image r_t by decoding their spherically interpolated representations at time t. The state cost penalizes the mean absolute deviation from r_t. Translation trajectories are shown in Figure 3 (additional samples in Section G.6), with TSBM producing more aligned translations that better preserve features of source images. Table 3 reports FID and optimality results. GSBM achieves lower FID in both directions, whereas TSBM yields lower optimality costs.

Depolarization experiment. We consider the opinion-depolarization experiment of [48], where polarizing interactions can split a population of 1000-dimensional opinions into groups with opposing views. Following GSBM, we parameterize the Brownian control u_t as a correction to a base polarizing drift and use a congestion potential, as in the crowd-navigation experiments, to penalize nearby opinions and discourage population concentration. Table 4 reports forward and backward feasibility and optimality results. While both methods achieve comparable feasibility, TSBM consistently yields lower optimality costs, demonstrating the benefit of our formulation in high dimensions.

5.4 Ablation studies

To further investigate the differences between GSBM and TSBM, we conduct ablation studies across all crowd-navigation tasks. Detailed results and visualizations are provided in Section G.7.

(A) **Unidirectional vs. bidirectional.** We compare the unidirectional and bidirectional variants of both methods, tracking forward and backward feasibility and optimality throughout training, following [30, Figure 2]. We find that bidirectional pretraining provides a substantial warm start, particularly in terms of feasibility, while the subsequent fine-tuning stabilizes optimality. These findings support the use of bidirectional training over the standard alternating scheme.

(B) **Lower-entropy setting.** We compare GSBM and TSBM in 2D with $\sigma = 1$, focusing on the pretraining phase of their bidirectional variants. As expected, the resulting trajectories are straighter than with $\sigma = 2$, while the relative performance trends remain consistent.

(C) **Effect of the control variate.** We assess the impact of the control variate (CV) by training TSBM with and without it. Although the resulting trajectories are visually similar, using CV generally improves the evaluation metrics, confirming its practical benefit.

(D) **About the feasibility–optimality trade-off.** To investigate the systematic performance gap between GSBM and TSBM, we conduct experiments that interchange their reciprocal and Markovian projection losses. We identify the Markovian projection as the main source of this discrepancy: the TSBM objective improves feasibility at the expense of optimality, partly due to the additional variance introduced by the state-cost gradient, which remains challenging to control even with CV. We regard this as the main numerical limitation of TSBM and leave its improvement for future work.

6 Conclusion & Perspectives

In this work, we revisit the trajectory-inference problem with time-dependent state costs and endpoint marginal constraints studied by [30], through the lens of *Twisted Schrödinger Bridge* (TSB). We extend the canonical SB setting by replacing Brownian motion with a twisted Brownian reference measure, obtained by reweighting trajectories according to the state cost. Building on this, we derive *Twisted Schrödinger Bridge Matching* (TSBM), a bridge-matching implementation of Iterative Markovian Fitting (IMF) for continuous- and discrete-time state costs. Compared with GSBM [30], TSBM introduces Markovian projection losses that leverage state-cost information to reduce bias, exactly recover DSBM [49] in the zero-potential case, support adaptive control variates and largely benefit from their bidirectional extension. We validate TSBM on crowd-navigation, single-cell, opinion depolarization and image translation tasks, where it matches or improves upon GSBM. Several limitations remain. TSBM requires differentiable state costs with informative gradients on the relevant domain; otherwise, the state-cost contribution may be underestimated in the TSBM loss. Moreover, its state-cost-informed objective is more computationally demanding than GSBM due to potential-gradient evaluations. Regarding the computational cost of iterative inference and learning, combining TSBM and GSBM with recent fast SDE inference techniques [39, 33] is a promising direction toward more lightweight SB procedures.

Acknowledgments and Disclosure of Funding

We warmly thank Stefano Peluchetti and Valentin De Bortoli for insightful discussions at the inception of this project. This work was performed using HPC resources from GENCI–IDRIS (AD011014860R1, AD011014860R2). This work received government funding administered by the National Research Agency (ANR) under the France 2030 program "Hi! PARIS", grant number ANR-23-IACL-0005. MS thanks the Chaire Deep Finance and Statistics, a joint research program between Qube R&T, Ecole polytechnique and Fondation de l'Ecole polytechnique, for financial support. AD received funding from the Fondation de l'École polytechnique as part of its "Servir la science" campaign. AD is supported by the France 2030 program with the reference ANR-25-PEIA-0001 (THEOREM project). AD is funded by the European Union (ERC-2022-SYG-OCEAN-101071601). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. AD is supported by Hi! Paris and Agence Nationale de la Recherche (Grant 11-LABX-0047).

References

[1] 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.

[2] Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

[3] Denis Blessing, Lorenz Richter, Julius Berner, Egor Malitskiy, and Gerhard Neumann. Bridge matching sampler: Scalable sampling via generalized fixed-point diffusion matching. In *International Conference on Machine Learning*, 2026.

[4] Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, pages 17695–17709, 2021.

[5] Valentin De Bortoli, Iryna Korshunova, Andriy Mnih, and Arnaud Doucet. Schrodinger bridge flow for unpaired data translation. In *Advances in Neural Information Processing Systems*, 2024.

[6] Lev M Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217, 1967.

[7] Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *International Conference on Learning Representations*, 2024.

[8] Tianrong Chen, Guan-Horng Liu, and Evangelos A Theodorou. Likelihood training of Schrödinger bridge using forward-backward SDEs theory. *International Conference on Learning Representations*, 2022.

[9] Tianrong Chen, Guan-Horng Liu, Molei Tao, and Evangelos Theodorou. Deep momentum multi-marginal schrödinger bridge. *Advances in Neural Information Processing Systems*, pages 57058–57086, 2023.

[10] Raphaël Chetrite and Hugo Touchette. Nonequilibrium markov processes conditioned on large deviations. In *Annales Henri Poincaré*, volume 16, pages 2005–2057. Springer, 2015.

[11] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 8185–8194, 2020.

[12] Pierre Del Moral and Alice Guionnet. On the stability of interacting processes with applications to filtering and genetic algorithms. In *Annales de l'Institut Henri Poincaré (B) Probability and Statistics*, volume 37, pages 155–194. Elsevier, 2001.

[13] Pierre Del Moral, Jean Jacod, and Philip Protter. The monte-carlo method for filtering with discrete-time observations. *Probability Theory and Related Fields*, 120(3):346–368, 2001.

[14] Carles Domingo i Enrich, Jiequn Han, Brandon Amos, Joan Bruna, and Ricky TQ Chen. Stochastic optimal control matching. *Advances in Neural Information Processing Systems*, pages 112459–112504, 2024.

[15] Yuanqi Du, Michael Plainer, Rob Brekelmans, Chenru Duan, Frank Noe, Carla P Gomes, Alan Aspuru-Guzik, and Kirill Neklyudov. Doob's lagrangian: A sample-efficient variational approach to transition path sampling. *Advances in Neural Information Processing Systems*, pages 65791–65822, 2024.

[16] Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trouve, and Gabriel Peyré. Interpolating between optimal transport and mmd using sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 2681–2690, 2019.

[17] M Fujisaki, G Kallianpur, and H Kunita. Stochastic differential equations for the non linear filtering problem. *Osaka J. Math*, 9:19–40, 1972.

[18] Nikita Gushchin, Daniil Selikhanovych, Sergei Kholkin, Evgeny Burnaev, and Aleksandr Korotin. Adversarial schrödinger bridge matching. *Advances in Neural Information Processing Systems*, pages 89612–89651, 2024.

[19] Jeremy Heng, Adrian N Bishop, George Deligiannidis, and Arnaud Doucet. Controlled sequential monte carlo. *The Annals of Statistics*, 48(5):2904–2929, 2020.

[20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, pages 6840–6851, 2020.

[21] Samuel Howard, Peter Potaptchik, and George Deligiannidis. Schrödinger bridge matching for tree-structured costs and entropic wasserstein barycentres. In *Advances in Neural Information Processing Systems*, 2025.

[22] G Kallianpur and C. Striebel. Estimation of stochastic systems: Arbitrary system process with additive white noise observation errors. *The Annals of Mathematical Statistics*, 39(3):785–801, 1968.

[23] Kacper Kapusniak, Peter Potaptchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael M. Bronstein, Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. In *Advances in Neural Information Processing Systems*, 2024.

[24] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in Neural Information Processing Systems*, pages 21696–21707, 2021.

[25] Christian Léonard. Stochastic derivatives and generalized h-transforms of markov processes. *arXiv preprint arXiv:1102.3172*, 2011.

[26] Christian Léonard. Some properties of path measures. *Séminaire de Probabilités XLVI*, pages 207–230, 2014.

[27] Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete & Continuous Dynamical Systems-A*, 34(4):1533–1574, 2014.

[28] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *International Conference on Learning Representations*, 2023.

[29] Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos Theodorou. Deep generalized schrödinger bridge. *Advances in Neural Information Processing Systems*, pages 9374–9388, 2022.

[30] Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, and Ricky T. Q. Chen. Generalized schrödinger bridge matching. In *International Conference on Learning Representations*, 2024.

[31] Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.

[32] Cheng Lu and Yang Song. Simplifying, stabilizing and scaling continuous-time consistency models. In *International Conference on Learning Representations*, 2025.

[33] Sam McCallum, Zander W Blasingame, Timothy Herschell, Niklas Rindtorff, Alexander Tong, and James Foster. Strong stochastic flow maps. *arXiv preprint arXiv:2606.01086*, 2026.

[34] Kevin R. Moon, David van Dijk, Zheng Wang, Scott Gigante, Daniel B. Burkhardt, William S. Chen, Kristina Yim, Antonia van den Elzen, Matthew J. Hirn, Ronald R. Coifman, Natalia B. Ivanova, Guy Wolf, and Smita Krishnaswamy. Visualizing structure and transitions for biological data exploration. *bioRxiv*, page 120378, 2019.

[35] Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *International Conference on Machine Learning*, pages 25858–25889. PMLR, 2023.

[36] Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A computational framework for solving Wasserstein lagrangian flows. In *International Conference on Machine Learning*, pages 37461–37485. PMLR, 2024.

[37] Maxence Noble, Valentin De Bortoli, Arnaud Doucet, and Alain Durmus. Tree-based diffusion Schrödinger bridge with applications to Wasserstein barycenters. *Advances in Neural Information Processing Systems*, 36:55193–55236, 2023.

[38] Nikolas Nüsken and Lorenz Richter. Solving high-dimensional hamilton–jacobi–bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial differential equations and applications*, 2(4):48, 2021.

[39] Zhengkai Pan, Peter Potaptchik, Wenxi Yao, Michael S Albergo, and Jakiw Pidstrigach. Itô maps for any-step sdes. *arXiv preprint arXiv:2606.11156*, 2026.

[40] Stefano Peluchetti. Diffusion bridge mixture transports, schrödinger bridge problems and generative modeling. *Journal of Machine Learning Research*, 24(374):1–51, 2023.

[41] Stefano Peluchetti. Non-denoising forward-time diffusions. *arXiv preprint arXiv:2312.14589*, 2023.

[42] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

[43] Aram-Alexandre Pooladian, Carles Domingo-Enrich, Ricky T. Q. Chen, and Brandon Amos. Neural optimal transport with lagrangian costs. In *Uncertainty in Artificial Intelligence*. PMLR, 2024.

[44] L Chris G Rogers and David Williams. *Diffusions, Markov processes, and martingales: Itô calculus*, volume 2. Cambridge university press, 2000.

[45] Bernard Roynette and Marc Yor. *Penalising brownian paths*, volume 1969. Springer Science & Business Media, 2009.

[46] Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your flow: Scaling continuous-time flow map distillation. In *Advances in Neural Information Processing Systems*, 2025.

[47] Erwin Schrödinger. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. *Annales de l'Institut Henri Poincaré*, 2(4):269–310, 1932.

[48] Simon Schweighofer, David Garcia, and Frank Schweitzer. An agent-based model of multi-dimensional opinion dynamics and opinion alignment. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(9), 2020.

[49] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 2023.

[50] Marta Gentiloni Silveri, Giovanni Conforti, and Alain Durmus. Exponential convergence guarantees for iterative markovian fitting. In *Advances in Neural Information Processing Systems*, 2025.

[51] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

[52] Yoshiki Takeuchi and Hajime Akashi. Nonlinear filtering formulas for discrete-time observations. *SIAM Journal on Control and Optimization*, 19(2):244–261, 1981.

[53] Ella Tamir, Martin Trapp, and Arno Solin. Transport with support: Data-conditional diffusion bridges. *Transactions on Machine Learning Research*, 2023.

[54] Alexis MH Teter, Wenqing Wang, and Abhishek Halder. Schrödinger bridge with quadratic state cost is exactly solvable. *arXiv preprint arXiv:2406.00503*, 2024.

[55] Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In *International Conference on Machine Learning*, pages 9526–9536. PMLR, 2020.

Organization of the supplementary

The appendix is organized as follows. Section A summarizes the notation used throughout the paper. Section B discusses related work on generalized SB problems and alternative formulations of the standard SB problem. Section C collects theoretical results that support the Twisted SB framework introduced in Section 2, and provides the proofs of the results stated in Section 3. Section D presents a fully discrete-time reformulation of the Twisted SB problem with discrete-time potentials, with corresponding methodology coined *Discrete-TSBM* (D-TSBM). Extending the canonical framework of [18] to twisted reference measures, we show in particular that the corresponding Markovian projection loss reduces to the widely used DDPM loss [20] in the zero-potential case, further highlighting the connection between TSBM and standard diffusion-model objectives. For clarity, the proofs of the results introduced in Section D are deferred to Section E. Section F provides implementation details for TSBM, focusing on two main aspects: the variational learning framework used to approximate the twisted bridge in the reciprocal projection step, see Section 3.1, and the control variate parameterization for the TSBM Markovian losses. Finally, Section G gives additional details on the experimental setup and reports further results.

A Notation **16**

B Related works **16**

C Theoretical results on TSBM **17**
C.1 Results on twisted measures . . . 17
C.2 Results on the reciprocal projection . . . 19
C.3 Results on the Markovian projection (continuous time) . . . 20
C.4 Results on the Markovian projection (discrete time) . . . 26

D Fully discrete-time Twisted Schrödinger Bridge Matching (D-TSBM) **28**
D.1 Setting and ideal IMF scheme . . . 28
D.2 Computation of the Reciprocal projection . . . 30
D.3 Computation of the Markovian projection . . . 31

E Proofs for D-TSBM **32**
E.1 Results on twisted measures . . . 32
E.2 Results on the discrete Twisted SB problem . . . 34
E.3 Results on the Markovian projection . . . 35

F Practical implementation of TSBM **36**
F.1 Pseudo-code for TSBM . . . 36
F.2 Variational approximation of the reciprocal projection . . . 36
F.3 Implementation of the score-based control variates . . . 38

G Experiment details and additional results **39**
G.1 General setting . . . 39
G.2 Experiments on the reciprocal projection . . . 40
G.3 Toy discrete time experiments . . . 41
G.4 Crowd navigation experiments . . . 43
G.5 Single-cell experiments . . . 45
G.6 High-dimensional experiments . . . 49
G.7 Ablation studies . . . 51

A Notation

We denote by $\mathrm{C}^{0,1}([0,T]\times\mathbb{R}^d,\mathbb{R})$ the space of time-dependent potentials on $[0,T]\times\mathbb{R}^d$ that are continuous in time and continuously differentiable in space. Similarly, $\mathrm{C}^1(\mathbb{R}^d,\mathbb{R})$ denotes the space of continuously differentiable potentials on $\mathbb{R}^d$.

For any measurable space $(\mathsf{X},\mathcal{X})$, we write $\mathscr{P}(\mathsf{X})$ for the set of probability measures on $(\mathsf{X},\mathcal{X})$. Unless stated otherwise, $\mathcal{X}$ is the Borel σ-algebra on X. We denote by $\mathbf{C}_T=\mathrm{C}([0,T],\mathbb{R}^d)$ the space of continuous paths from $[0,T]$ to $\mathbb{R}^d$; accordingly, elements of $\mathscr{P}(\mathbf{C}_T)$ are continuous-time stochastic processes, or path measures, on $[0,T]$. Given a path measure $\mathbb{P}\in\mathscr{P}(\mathbf{C}_T)$ and a time index set $\mathsf{I}\subset[0,T]$, either discrete or continuous, we denote by $\mathbb{P}_\mathsf{I}$ the law of $\mathbf{X}_\mathsf{I}$ when $\mathbf{X}\sim\mathbb{P}$. In particular, $\mathbb{P}_t$ is the marginal of $\mathbb{P}$ at time $t\in[0,T]$, while $\mathbb{P}_{0,T}$ denotes the joint law of the initial and terminal endpoints. We write $\mathbb{P}_{|0,T}(\cdot|x_0,x_T)$ for the conditional path measure of $\mathbf{X}\sim\mathbb{P}$ given $\mathbf{X}_0=x_0$ and $\mathbf{X}_T=x_T$.

We will frequently use mixtures of endpoint-conditioned bridges. Given a coupling $\pi_{0,T}\in\mathscr{P}(\mathbb{R}^d\times\mathbb{R}^d)$ and a path measure $\mathbb{Q}\in\mathscr{P}(\mathbf{C}_T)$, we denote by $\mathbb{P}=\pi_{0,T}\mathbb{Q}_{|0,T}$ the path measure defined by $\mathbb{P}(\cdot)=\int_{\mathbb{R}^d\times\mathbb{R}^d}\mathbb{Q}_{|0,T}(\cdot|x_0,x_T)\,\mathrm{d}\pi_{0,T}(x_0,x_T)$. We also denote by $\mathcal{M}_T\subset\mathscr{P}(\mathbf{C}_T)$ the set of Markov continuous-time path measures. For discrete-time processes, and for any $N\geq 0$, we use the shorthand $\mathscr{P}^{(N+2)}=\mathscr{P}((\mathbb{R}^d)^{N+2})$, which corresponds to probability measures over $(N+2)$ time steps, namely the two endpoints together with N intermediate time steps. When $\mathsf{X}=(\mathbb{R}^d)^{N+2}$, for any $x\in\mathsf{X}$ and any indices $m,n\in\{0,\dots,N+1\}$ with $m\leq n$, we write $x_{m:n}=(x_m,x_{m+1},\dots,x_n)$.

Let Leb denote the Lebesgue measure. For any nonnegative measurable function $f:\mathsf{X}\to\mathbb{R}_+$ such that $\int_\mathsf{X} f\,\mathrm{dLeb}<\infty$, we define $\mathrm{H}(f)=-\int_\mathsf{X} f\log f\,\mathrm{dLeb}\in(-\infty,+\infty]$. For any probability measure $\mu\in\mathscr{P}(\mathsf{X})$, its entropy is defined by $\mathrm{H}(\mu)=\mathrm{H}(\mathrm{d}\mu/\mathrm{dLeb})$ if $\mu\ll\mathrm{Leb}$, and $\mathrm{H}(\mu)=+\infty$ otherwise. For two arbitrary measures μ and ν on $(\mathsf{X},\mathcal{X})$, we define the Kullback–Leibler divergence by $\mathrm{KL}(\mu\|\nu)=\int_\mathsf{X}\log(\mathrm{d}\mu/\mathrm{d}\nu)\,\mathrm{d}\mu-\int_\mathsf{X}\mathrm{d}\mu+\int_\mathsf{X}\mathrm{d}\nu$ if $\mu\ll\nu$, and set $\mathrm{KL}(\mu\|\nu)=+\infty$ otherwise. Finally, $\mathrm{N}(\mathbf{m},\Sigma)$ denotes the Gaussian distribution with mean $\mathbf{m}$ and covariance matrix Σ.

B Related works

Potential-based approaches for solving the generalized SB problem. Beyond bridge-matching methods, a closely related line of work is the action-matching (AM) framework [35]. AM is a diffusion-based generative modeling approach designed to learn marginal-preserving stochastic or deterministic dynamics from observed intermediate samples. Its main idea is to parameterize the target velocity field as the gradient of a time-dependent potential $(\phi_t)_{t\in[0,T]}$, approximated by a neural network $(\phi_t^\theta)_{t\in[0,T]}$. Unlike bridge matching, which *explicitly* regresses onto a conditional drift—namely, the drift of the dynamics induced by the underlying bridge path measure—the AM objective is *implicit*: it only requires evaluating the neural time derivative $\partial_t\phi_t^\theta$, the gradient $\nabla\phi_t^\theta$, and, in the stochastic case, the Laplacian $\Delta\phi_t^\theta$.

More recently, [36] extended action matching to a range of SB-like problems, including standard, unbalanced, and generalized Schrödinger bridge formulations close to (2). Their approach casts the problem as a min–max optimization. As in TSBM, training alternates between two stages: first learning a neural interpolant, and then learning dynamics that preserve the marginals induced by this interpolant, while generating an updated coupling between endpoints. This is achieved by optimizing a joint objective based on a variant of the AM loss. In contrast to TSBM, their method does not require explicitly updating and storing the endpoint coupling; in particular, one may keep the independent coupling throughout training, reducing the memory footprint and avoiding intermediate simulations such as those used in our procedure. However, a key computational limitation is that evaluating the AM loss requires neural-network derivatives: first-order derivatives in the deterministic case, and second-order derivatives when $\sigma>0$, as in the stochastic setting considered in this paper. This restricts its practical scalability to relatively low-dimensional settings. By contrast, the TSBM loss is derivative-free with respect to the neural drift parameterization, although it may require derivatives of the prescribed state cost.

Another related contribution is Lagrangian Optimal Transport [43], which proposes a learning approach for generalized SB-type problems in the zero-noise regime, i.e., $\sigma=0$. Rather than learning the deterministic dynamics themselves, this method directly estimates the underlying optimal transport map. It also relies on a min–max formulation, alternating between the construction of a

suitable interpolant and the learning of Kantorovich potentials, parameterized by input-convex neural networks, from which the transport map is recovered. The interpolant is spline-based, similarly to the lightweight parameterization we use in our implementation. However, this approach is restricted to the deterministic setting, whereas TSBM targets stochastic dynamics. Moreover, as in the AM framework, it involves input derivatives of neural networks, which makes it most suitable for low-dimensional problems.

Diffusion-based approaches for trajectory inference with prior data. In recent years, several works have extended diffusion-based methods beyond unconditional generation between source and target distributions, by incorporating additional prior information into the learned dynamics. One direction is the multi-marginal extension of the canonical SB problem, which decomposes the trajectory into several SB sub-problems over a single time interval. This requires enough samples at the intermediate times of interest to model these intermediate marginals in the same way as the endpoint distributions. Related approaches include [9], building on the computational SB framework of [8]; [37], building on [4]; and [21], building on DSBM [49]. Intermediate samples can also be used outside the SB framework. For example, [23] propose Metric Flow Matching (MFM), an instance of Riemannian Flow Matching [7], in which the geodesic interpolant is learned from a Riemannian metric constructed using the available samples. Although this approach achieves strong empirical performance, its connection with generalized SB problems remains unclear. In the sparse-observation regime, where too few samples are available to reliably define intermediate marginals, [53] propose a method based on Bayesian filtering and optimal control, built on the diffusion backbone of [4], to generate constrained stochastic processes. Here again, the precise relation to generalized SB formulations is not explicit.

C Theoretical results on TSBM

Throughout this section, we work on the canonical measurable space $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$, where $\mathcal{B}(\mathbb{R}^d)$ denotes the Borel σ-field of $\mathbb{R}^d$. On this space, we consider a d-dimensional Brownian motion $(\mathbf{B}_t)_{t\geq 0}$. We define $\mathbf{A}_T$, the set of admissible control functions, by

$$\mathbf{A}_T = \{u \in \mathrm{C}^1([0,T]\times\mathbb{R}^d, \mathbb{R}^d) : \exists C > 0 : \forall (t,x) \in [0,T]\times\mathbb{R}^d, \|u_t(x)\| \leq C(1+\|x\|)\} \,.$$

We first recall that for any $(x_0, x_T) \in \mathbb{R}^d \times \mathbb{R}^d$, by simple application of Doob's h-transform [44], the scaled Brownian bridge $\mathbb{Q}^\sigma_{|0,T}(\cdot|x_0,x_T)$ is a pinned diffusion process equivalently associated to the forward and backward SDEs

$$\mathrm{d}\mathbf{X}^{0,T}_t = \mathbf{u}^{\mathrm{f}}_t(\mathbf{X}^{0,T}_t)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t,\ \mathbf{X}^{0,T}_0 = x_0,\ \mathbf{X}^{0,T}_T = x_T\ ,\ \mathbf{u}^{\mathrm{f}}_t(x) = \sigma^2\nabla_x \log \mathbb{Q}^\sigma_{T|t}(x_T|x)\ , \quad (21)$$

$$\mathrm{d}\mathbf{Y}^{0,T}_t = \mathbf{u}^{\mathrm{b}}_{T-t}(\mathbf{Y}^{0,T}_t)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t,\ \mathbf{Y}^{0,T}_0 = x_T, \mathbf{Y}^{0,T}_T = x_0, \mathbf{u}^{\mathrm{b}}_t(x) = \sigma^2\nabla_x \log \mathbb{Q}^\sigma_{t|0}(x|x_0). \quad (22)$$

These drifts simplify as $\mathbf{u}^{\mathrm{f}}_t(x) = (x_T - x_t)/(T-t)$ and $\mathbf{u}^{\mathrm{b}}_t(x) = (x_0 - x_t)/t$.

C.1 Results on twisted measures

We begin by recalling the following result from [25, Proposition 4.2] concerning the twisted Brownian motion $\mathbb{Q}^{V,\sigma}$ defined in (4).

Proposition 4. *Let $V \in \mathrm{C}^{0,1}([0,T]\times\mathbb{R}^d, \mathbb{R})$ be such that $Z_V < \infty$. Then it holds that (i) $\mathbb{Q}^{V,\sigma}$ is Markov, (ii) for any $0 \leq s \leq t \leq T$, we have*

$$\mathrm{d}\mathbb{Q}^{V,\sigma}_{s:t}(x_{[s:t]}) \propto \exp\left(-\textstyle\int_s^t \frac{V_u(x_u)}{\sigma^2}\mathrm{d}u\right)\mathrm{d}\mathbb{Q}^\sigma_{s:t}(x_{[s:t]})\ .$$

As a direct consequence of Proposition 4, we obtain the following unnormalized Radon–Nikodym densities:

$$\mathrm{d}\mathbb{Q}^{V,\sigma}_{t:T|t}(x_{(t:T]}|x_t) \propto \exp\left(-\textstyle\int_t^T \frac{V_u(x_u)}{\sigma^2}\mathrm{d}u\right)\mathrm{d}\mathbb{Q}^\sigma_{t:T|t}(x_{(t:T]}|x_t)\ ,$$

$$\mathrm{d}\mathbb{Q}^{V,\sigma}_{0:t|t}(x_{[0:t)}|x_t) \propto \exp\left(-\textstyle\int_0^t \frac{V_u(x_u)}{\sigma^2}\mathrm{d}u\right)\mathrm{d}\mathbb{Q}^\sigma_{0:t|t}(x_{[0:t)}|x_t)\ ,$$

$$\mathrm{d}\mathbb{Q}^{V,\sigma}_{|0,T}(x_{(0:T)}|x_0,x_T) \propto \exp\left(-\textstyle\int_0^T \frac{V_u(x_u)}{\sigma^2}\mathrm{d}u\right)\mathrm{d}\mathbb{Q}^\sigma_{|0,T}(x_{(0:T)}|x_0,x_T)\ .$$

These identities suggest that the twisted measure $\mathbb{Q}^{V,\sigma}$ and its bridge $\mathbb{Q}^{V,\sigma}_{|0,T}$ should admit stochastic differential representations. More precisely, we now show that they are associated with affine-controlled Brownian motions. To this end, we first establish a general result from which the desired statements follow as corollaries, see Corollaries 6 and 7.

Proposition 5. *Consider $\mathbb{Q}^\sigma \in \mathcal{M}_T$, associated to an SDE $\mathrm{d}\mathbf{X}_t = f_t(\mathbf{X}_t)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t$, $\mathbf{X}_0 \sim \mathbb{Q}_0^\sigma$, where $f \in \mathrm{C}^1([0,T]\times\mathbb{R}^d, \mathbb{R}^d)$ satisfies a linear growth condition, i.e., there exists $C > 0$ such that for any $(t,x) \in [0,T]\times\mathbb{R}^d$, we have $\|f_t(x)\| \leq C(1+\|x\|)$. Let $V \in \mathrm{C}^{0,1}([0,T]\times\mathbb{R}^d, \mathbb{R})$ be such that for any $t \in [0,T]$ and any $\mathbb{Q}_t^\sigma$-almost surely $x \in \mathbb{R}^d$, we have*

$$Z_V(t,x) = \mathbb{E}_{\mathbb{Q}^\sigma_{|t}}\left[\exp\left(-\int_t^T \tfrac{V_s(\mathbf{X}_s)}{\sigma^2}\mathrm{d}s\right)|\mathbf{X}_t = x\right] < \infty\,.$$

Then, the path measure $\mathbb{Q}^{V,\sigma}$, defined in (4), is Markov, associated with the affine-controlled SDE

$$\mathrm{d}\mathbf{X}_t^{\mathbf{u}} = \{f_t(\mathbf{X}_t^{\mathbf{u}}) + \sigma \mathbf{u}_t^V(\mathbf{X}_t^{\mathbf{u}})\}\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t,\ \mathbf{X}_0^{\mathbf{u}} \sim \mathbb{Q}_0^{V,\sigma} \text{ where } \mathbf{u}_t^V(x) = \sigma\nabla_x \log Z_V(t,x). \quad (23)$$

Moreover, for any arbitrary continuously differentiable function $m_t : [t,T] \to \mathbb{R}$ such that $m_t(t) = 1$, the control $\mathbf{u}_t^V$ admits a closed-form solution depending on the potential V via the path-wise reparameterization trick *[14] as*

$$\mathbf{u}_t^V(x) = \frac{\sigma}{Z_V(t,x)}\mathbb{E}_{\mathbb{Q}^\sigma_{|t}}\left[\left\{-\int_t^T m_t(s)\frac{\nabla V_s(\mathbf{X}_s)}{\sigma^2}\mathrm{d}s + \frac{1}{\sigma}\int_t^T \left(m_t(s)\nabla f_s(\mathbf{X}_s) - \partial_s m_t(s)\mathrm{I}_d\right)\mathrm{d}\mathbf{B}_s\right\}\right.$$
$$\left.\times \exp\left(-\int_t^T \frac{V_s(\mathbf{X}_s)}{\sigma^2}\mathrm{d}s\right)\middle|\mathbf{X}_t = x\right]. \quad (24)$$

Proof. Let $x_0 \sim \mathbb{Q}_0^{V,\sigma}$ and consider the following minimization problem over path measures:

$$\arg\min\{\mathrm{KL}(\mathbb{M}_{|0}(\cdot|x_0)\|\mathbb{Q}^{V,\sigma}_{|0}(\cdot|x_0))\ :\ \mathbb{M} \in \mathcal{M}_{T,c}\}\,, \quad (25)$$

where $\mathcal{M}_{T,c}$ is the set of Markov measures that are associated with controlled SDEs of the form

$$\mathrm{d}\mathbf{X}_t^{\mathbf{u}} = f_t(\mathbf{X}_t^{\mathbf{u}})\mathrm{d}t + \mathbf{u}_t(\mathbf{X}_t^{\mathbf{u}})\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t\ ,\ \mathbf{X}_0^u = x_0\ ,$$

where $\mathbf{u} \in \mathbf{A}_T$. Let $\mathbb{M}^V$ denote the solution to (25).

By Girsanov's formula, we have for any $\mathbb{M} \in \mathcal{M}_{T,c}$, the following KL decomposition

$$\mathrm{KL}(\mathbb{M}_{|0}(\cdot|x_0)\|\mathbb{Q}^{V,\sigma}_{|0}(\cdot|x_0)) = J(\mathbf{u};x_0,0) + \log Z_V\ ,$$

where

$$J(\mathbf{u};x,t) = \frac{1}{\sigma^2}\mathbb{E}\left[\int_t^T\left\{\frac{1}{2}\left\|\mathbf{u}_t(\mathbf{X}_t^{\mathbf{u}})\right\|^2 + V_t(\mathbf{X}_t^{\mathbf{u}})\right\}\mathrm{d}t \mid \mathbf{X}_t^{\mathbf{u}} = x\right]\ .$$

Hence, solving (25) is equivalent to solving the following optimization problem

$$\arg\min\{J(\mathbf{u};x_0,0)\ :\ \mathbf{u} \in \mathbf{A}_T\}\,. \quad (26)$$

In fact, this is a SOC problem associated to the cost functional $(u,x,t) \mapsto J(u;x,t)$. Combining [38, Theorem 2.2] and [38, Remark 2.6], we deduce under the stated assumptions that the unique minimizer $\mathbf{u}^V$ satisfies

$$\mathbf{u}_t^V(x) = \sigma\nabla_x \log Z_V(t,x) \quad \text{and} \quad J(\mathbf{u}^V;x_0,0) = -\log Z_V\ .$$

In particular, for any x_0, $\mathbb{Q}_0^{V,\sigma}$-almost surely,

$$\mathrm{KL}(\mathbb{M}^V_{|0}(\cdot|x_0)\|\mathbb{Q}^{V,\sigma}_{|0}(\cdot|x_0)) = 0\ ,$$

which proves the result (23). The representation (24) follows directly from [14, Proposition 1]. □

Notably, Proposition 5 extends [25, Theorem 5.4], where $\mathbb{Q}^\sigma$ is assumed to be stationary, an assumption that does not hold for Brownian motion.

Corollary 6 (Application to the Brownian motion)**.** *If $\mathbb{Q}^\sigma$ is the distribution of the scaled Brownian motion $(\sigma\mathbf{B}_t)_{t\in[0,T]}$, then $\mathbb{Q}^{V,\sigma}$ is associated with*

$$\mathrm{d}\mathbf{X}_t^{\mathbf{u}} = \mathbf{u}_t^V(\mathbf{X}_t)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t\ ,\ \mathbf{X}_0^{\mathbf{u}} \sim \mathbb{Q}_0^{V,\sigma}\ , \quad (27)$$

where

$$\mathbf{u}_t^V(x) = \mathbb{E}_{\mathbb{Q}^{V,\sigma}_{|t}}\left[-\int_t^T \nabla V_s(\mathbf{X}_s)\mathrm{d}s|\mathbf{X}_t = x\right]\ . \quad (28)$$

Proof. This is an immediate application of Proposition 5 in the case where $f \equiv 0$ and $m_t \equiv 1$ for any $t \in [0, T]$.

□

Corollary 7 (Application to the *pinned* Brownian motion)**.** *Let* $(x_0, x_T) \in \mathbb{R}^d \times \mathbb{R}^d$. *If* $\mathbb{Q}^{\sigma}_{|0,T}(\cdot|x_0, x_T)$ *is the distribution of the scaled Brownian motion* $(\sigma \mathbf{B}_t)_{t\in[0,T]}$ pinned *down at initial and terminal points* (x_0, x_T)*, then* $\mathbb{Q}^{V,\sigma}_{|0,T}(\cdot|x_0, x_T)$ *is equivalently associated with the forward and backward SDEs*

$$\mathrm{d}\mathbf{X}^{\mathbf{u}}_t = \frac{x_T - \mathbf{X}^{\mathbf{u}}_t}{T-t} + \mathbf{u}^{\mathrm{f},V}_t(\mathbf{X}^{\mathbf{u}}_t|x_T)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t \ , \mathbf{X}^{\mathbf{u}}_0 = x_0 \ , \tag{29}$$

$$\mathrm{d}\mathbf{Y}^{\mathbf{u}}_t = \frac{x_0 - \mathbf{Y}^{\mathbf{u}}_t}{T-t} + \mathbf{u}^{\mathrm{b},V}_{T-t}(\mathbf{Y}^{\mathbf{u}}_t|x_0)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t \ , \mathbf{Y}^{\mathbf{u}}_0 = x_T \ , \tag{30}$$

where

$$\mathbf{u}^{\mathrm{f},V}_t(x|x_T) = \mathbb{E}_{\mathbb{Q}^{V,\sigma}_{|t,T}}\left[-\int_t^T \frac{T-s}{T-t}\nabla V_s(\mathbf{X}_s)\mathrm{d}s|\mathbf{X}_t = x, \mathbf{X}_T = x_T\right] \ ,$$

$$\mathbf{u}^{\mathrm{b},V}_t(x|x_0) = \mathbb{E}_{\mathbb{Q}^{V,\sigma}_{|0,t}}\left[-\int_0^t \frac{s}{t}\nabla V_s(\mathbf{X}_s)\mathrm{d}s|\mathbf{X}_0 = x_0, \mathbf{X}_t = x\right] \ .$$

Proof. Obtaining the forward SDE results from the application of Proposition 5 when considering the pinned Brownian bridge (in forward time direction) as base process $\mathbb{Q}^{\sigma}$, with drift given in (21), and by choosing $m_t(s) = \frac{T-s}{T-t}$, which simplifies the stochastic integral term in (24). To obtain the backward SDE, we apply Proposition 5 to the backward pinned Brownian bridge $(\mathbb{Q}^{\sigma})^R$, whose SDE is exactly given in (22), considering as twisted path measure $\tilde{\mathbb{Q}}^{\sigma,V}(y_{[0:T]}) \propto \exp(-\int_0^T V_{T-s}(y_s)\mathrm{d}s)(\mathbb{Q}^{\sigma})^R(y_{[0:T]})$ and taking once again $m_t(s) = \frac{T-s}{T-t}$. □

C.2 Results on the reciprocal projection

Below, we restate the result of Proposition 1, additionally including the backward direction.

Proposition 8 (Reciprocal projection – continuous time)**.** *Assume that the variational family* $\{\mathbb{Q}^{\psi}_{|0,T}\}_{\psi\in\Psi}$ *in* (15) *describes the pinned diffusion processes associated, for almost surely any* $(x_0, x_T) \sim \mathbb{P}_{0,T}$*, to forward-time and backward-time parametric SDEs of the form*

$$\mathrm{d}\mathbf{X}^{0,T}_t = v^{\mathrm{f},\psi}_t(\mathbf{X}^{0,T}_t|x_0, x_T)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t \ , \ \mathbf{X}^{0,T}_0 = x_0, \ \mathbf{X}^{0,T}_T = x_T \ , \tag{31}$$

$$\mathrm{d}\mathbf{Y}^{0,T}_t = v^{\mathrm{b},\psi}_{T-t}(\mathbf{Y}^{0,T}_t|x_0, x_T)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t \ , \ \mathbf{Y}^{0,T}_0 = x_T, \ \mathbf{Y}^{0,T}_T = x_0 \ , \tag{32}$$

with $v^{\mathrm{f},\psi}, v^{\mathrm{b},\psi} \in \mathbf{A}_T$. *Then, the minimizer of* (15) *for* $\mathbb{Q}^{V,\sigma}$ *defined by* (4)*, denoted by* $\psi_\star$*, is equivalently the minimizer of the two following losses*

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{R}}(\psi) = \int_0^T \mathbb{E}_{\Pi^{\psi}_{0,t,T}}\left[\tfrac{1}{2}\|(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - v^{\mathrm{f},\psi}_t(\mathbf{X}_t|\mathbf{X}_0, \mathbf{X}_T)\|^2 + V_t(\mathbf{X}_t)\right]\mathrm{d}t \ , \tag{33}$$

$$\mathcal{L}^{\mathrm{b}}_{\mathcal{R}}(\psi) = \int_0^T \mathbb{E}_{\Pi^{\psi}_{0,t,T}}\left[\tfrac{1}{2}\|(\mathbf{X}_0 - \mathbf{X}_t)/t - v^{\mathrm{b},\psi}_t(\mathbf{X}_t|\mathbf{X}_0, \mathbf{X}_T)\|^2 + V_t(\mathbf{X}_t)\right]\mathrm{d}t \ , \tag{34}$$

where $\Pi^{\psi} = \mathbb{P}_{0,T}\mathbb{Q}^{\psi}_{|0,T}$.

Proof of Proposition 8. Consider such variational family. We first prove the result in the forward direction. For any $\psi \in \Psi$, and any $(x_0, x_T) \in \mathbb{R}^d \times \mathbb{R}^d$, we have, $\mathbb{P}_{0,T}$-almost surely,

$$\begin{aligned}
&\mathrm{KL}(\mathbb{Q}^{\psi}_{|0,T}(\cdot|x_0, x_T)\|\mathbb{Q}^{V,\sigma}_{|0,T}(\cdot|x_0, x_T)) \\
&= \mathrm{KL}(\mathbb{Q}^{\psi}_{|0,T}(\cdot|x_0, x_T)\|\mathbb{Q}^{\sigma}_{|0,T}(\cdot|x_0, x_T)) + \int_0^T \tfrac{1}{\sigma^2}\mathbb{E}_{\mathbb{Q}^{\psi}_{t|0,T}}[V_t(\mathbf{X}_t)|\mathbf{X}_0 = x_0, \mathbf{X}_T = x_T]\mathrm{d}t + \mathbf{c}(x_0, x_T) \\
&= \tfrac{1}{\sigma^2}\int_0^T \mathbb{E}_{\mathbb{Q}^{\psi}_{t|0,T}}\left[\tfrac{1}{2}\|v^{\mathrm{f},\psi}_t(\mathbf{X}_t|\mathbf{X}_0, \mathbf{X}_T) - \sigma^2\nabla \log \mathbb{Q}^{\sigma}_{T|t}(\mathbf{X}_T|\mathbf{X}_t)\|^2 + V_t(\mathbf{X}_t)|\mathbf{X}_0 = x_0, \mathbf{X}_T = x_T\right]\mathrm{d}t \\
&\quad + \mathbf{c}(x_0, x_T)
\end{aligned}$$

where $\mathbf{c}(x_0, x_T)$ is a constant independent of ψ (namely, the normalizing constant of $\mathbb{Q}^{V,\sigma}_{|0,T}(\cdot|x_0, x_T)$), and the second equality comes from Girsanov's theorem application exploiting (21). Therefore, by taking the expectation with respect to $\mathbb{P}_{0,T}$, we obtain the forward objective identity

$$\mathbb{E}_{\mathbb{P}_{0,T}}[\mathrm{KL}(\mathbb{Q}^{\psi}_{|0,T}\|\mathbb{Q}^{V,\sigma}_{|0,T})] = \frac{1}{\sigma^2}\mathcal{L}^{\mathrm{f}}_{\mathcal{R}}(\psi) + \mathbb{E}_{\mathbb{P}_{0,T}}[\mathbf{c}(\mathbf{X}_0, \mathbf{X}_T)] \; .$$

Similarly, by exploiting (22), we can obtain the backward objective identity

$$\mathbb{E}_{\mathbb{P}_{0,T}}[\mathrm{KL}(\mathbb{Q}^{\psi}_{|0,T}\|\mathbb{Q}^{V,\sigma}_{|0,T})] = \frac{1}{\sigma^2}\mathcal{L}^{\mathrm{b}}_{\mathcal{R}}(\psi) + \mathbb{E}_{\mathbb{P}_{0,T}}[\mathbf{c}(\mathbf{X}_0, \mathbf{X}_T)] \; .$$

The result of the proposition immediately follows these observations. □

We obtain a similar result for the discrete time case, which can be proved in the exact same way.

Proposition 9 (Reciprocal projection – discrete time)**.** *Assume that the variational family $\{\mathbb{Q}^{\psi}_{|0,T}\}_{\psi\in\Psi}$ in (15) describes the pinned diffusion processes associated, for almost surely any $(x_0, x_T) \sim \mathbb{P}_{0,T}$, to forward-time and backward-time parametric SDEs of the form*

$$\mathrm{d}\mathbf{X}^{0,T}_t = v^{\mathrm{f},\psi}_t(\mathbf{X}^{0,T}_t|x_0, x_T)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \; , \; \mathbf{X}^{0,T}_0 = x_0, \; \mathbf{X}^{0,T}_T = x_T \; ,$$
$$\mathrm{d}\mathbf{Y}^{0,T}_t = v^{\mathrm{b},\psi}_{T-t}(\mathbf{Y}^{0,T}_t|x_0, x_T)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \; , \; \mathbf{Y}^{0,T}_0 = x_T, \; \mathbf{Y}^{0,T}_T = x_0 \; ,$$

with $v^{\mathrm{f},\psi}, v^{\mathrm{b},\psi} \in \mathbf{A}_T$. Then, the minimizer of (15) for $\mathbb{Q}^{V,\sigma}$ defined by (6), denoted by $\psi_\star$, is equivalently the minimizer of the two following losses

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{R}}(\psi) = \int_0^T \mathbb{E}_{\Pi^{\psi}_{0,t,T}}\left[\tfrac{1}{2}\|(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - v^{\mathrm{f},\psi}_t(\mathbf{X}_t|\mathbf{X}_0, \mathbf{X}_T)\|^2\right]\mathrm{d}t + \textstyle\sum_{k=1}^K \mathbb{E}_{\Pi^{\psi}_{t_k}}[V_k(\mathbf{X}_{t_k})] \; ,$$
$$\mathcal{L}^{\mathrm{b}}_{\mathcal{R}}(\psi) = \int_0^T \mathbb{E}_{\Pi^{\psi}_{0,t,T}}\left[\tfrac{1}{2}\|(\mathbf{X}_0 - \mathbf{X}_t)/t - v^{\mathrm{b},\psi}_t(\mathbf{X}_t|\mathbf{X}_0, \mathbf{X}_T)\|^2\right]\mathrm{d}t + \textstyle\sum_{k=1}^K \mathbb{E}_{\Pi^{\psi}_{t_k}}[V_k(\mathbf{X}_{t_k})] \; ,$$

where $\Pi^{\psi} = \mathbb{P}_{0,T}\mathbb{Q}^{\psi}_{|0,T}$.

C.3 Results on the Markovian projection (continuous time)

Proposition 10 (Extended version of Proposition 2)**.** *Let $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$. Then, $\mathrm{proj}_{\mathcal{M}_T}(\Pi)$ is equivalently induced by the diffusion processes $(\mathbf{X}^{\mathrm{f}}_t)_{t\in[0,T]}$ and $(\mathbf{X}^{\mathrm{b}}_{T-t})_{t\in[0,T]}$, which are time-reversed of each other and respectively defined by the SDEs*

$$\mathrm{d}\mathbf{X}^{\mathrm{f}}_t = \mathrm{b}^{\mathrm{f},V}_t(\mathbf{X}^{\mathrm{f}}_t)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \; , \quad \mathbf{X}^{\mathrm{f}}_0 \sim \Pi_0 \; , \tag{35}$$
$$\mathrm{d}\mathbf{X}^{\mathrm{b}}_t = \mathrm{b}^{\mathrm{b},V}_{T-t}(\mathbf{X}^{\mathrm{b}}_t)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \; , \quad \mathbf{X}^{\mathrm{b}}_0 \sim \Pi_T \; , \tag{36}$$

where

$$\mathrm{b}^{\mathrm{f},V}_t(x_t) = \mathbb{E}_{\Pi_{|t}}[(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - \textstyle\int_t^T \frac{T-s}{T-t}\nabla V_s(\mathbf{X}_s)\mathrm{d}s|\mathbf{X}_t = x_t] \; , \tag{37}$$
$$\mathrm{b}^{\mathrm{b},V}_t(x_t) = \mathbb{E}_{\Pi_{|t}}[(\mathbf{X}_0 - \mathbf{X}_t)/t - \textstyle\int_0^t \frac{s}{t}\nabla V_s(\mathbf{X}_s)\mathrm{d}s|\mathbf{X}_t = x_t] \; . \tag{38}$$

In other words, the optimal drifts $\mathrm{b}^{\mathrm{f},V}$ and $\mathrm{b}^{\mathrm{b},V}$ write as conditional expectations of the forward and backward drifts driving the pinned process $\mathbb{Q}^{V,\sigma}_{|0,T}$, given in Corollary 7, as follows

$$\mathrm{b}^{\mathrm{f},V}_t(x_t) = \mathbb{E}_{\Pi_{T|t}}\left[(\mathbf{X}_T - \mathbf{X}_t)/(T-t) + \mathbf{u}^{\mathrm{f},V}_t(\mathbf{X}_t|\mathbf{X}_T)|\mathbf{X}_t = x_t\right] \; , \quad \textit{derived from (29)} \; ,$$
$$\mathrm{b}^{\mathrm{b},V}_t(x_t) = \mathbb{E}_{\Pi_{0|t}}\left[(\mathbf{X}_0 - \mathbf{X}_t)/t + \mathbf{u}^{\mathrm{b},V}_t(\mathbf{X}_t|\mathbf{X}_0)|\mathbf{X}_t = x_t\right] \; , \quad \textit{derived from (30)} \; .$$

Therefore, for any $t \in [0, T]$, we have $\mathrm{Law}(\mathbf{X}^{\mathrm{f}}_t) = \mathrm{Law}(\mathbf{X}^{\mathrm{b}}_{T-t}) = \Pi_t$.

Additionally, the forward drift $\mathrm{b}^{\mathrm{f},V}$ is approximated by $v^{\theta^{\mathrm{f}}_\star}$, where $\theta^{\mathrm{f}}_\star$ minimizes the forward objective $\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}$ given in (18), and the backward drift $\mathrm{b}^{\mathrm{b},V}$ is approximated by $v^{\theta^{\mathrm{b}}_\star}$, where $\theta^{\mathrm{b}}_\star$ minimizes

$$\mathcal{L}^{\mathrm{b}}_{\mathcal{M}_T}(\theta^{\mathrm{b}}) = \int_0^T \frac{1}{t}\int_0^t \mathbb{E}_{\Pi_{0,s,t}}\left[\left\|\mathbf{b}^{\mathrm{b},V}_{t,s}(\mathbf{X}_0, \mathbf{X}_s, \mathbf{X}_t) - v^{\theta^{\mathrm{b}}}_t(\mathbf{X}_t)\right\|^2\right]\mathrm{d}s\,\mathrm{d}t \; , \tag{39}$$

where $\Pi_{0,s,t}$ is the joint law of (X_0, X_s, X_t) under Π, and

$$\mathbf{b}^{\mathrm{b},V}_{t,s}(x_0, x_s, x_t) = \frac{x_0 - x_t}{t} - s\nabla V_s(x_s) \; .$$

Proof. The structure of the proof closely follows that of [49, Proposition 2] and relies on analogous technical assumptions, which we omit here for readability. We detail below the argument in the forward setting; the backward case follows by entirely similar arguments.

Since $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$, we have $\Pi_{|0} = \phi_{T|0}\mathbb{Q}^{V,\sigma}_{|0}$, where $\phi_{T|0} = \frac{\mathrm{d}\Pi_{T|0}}{\mathrm{d}\mathbb{Q}^{V,\sigma}_{T|0}}$. In particular, for any almost surely $x_0 \sim \Pi_0$, $\Pi_{|0}(\cdot|x_0)$ is obtained from $\mathbb{Q}^{V,\sigma}_{|0}(\cdot|x_0)$ through a Doob's h-transform, and is therefore associated with the SDE

$$\mathrm{d}\mathbf{X}^{\mathbf{u}}_t = \mathbf{u}^V_t(\mathbf{X}^{\mathbf{u}}_t)\mathrm{d}t + \sigma^2\nabla\log\phi_{t|0}(\mathbf{X}^{\mathbf{u}}_t|\mathbf{X}^{\mathbf{u}}_0)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \ , \ \mathbf{X}^{\mathbf{u}}_0 = x_0 \ ,$$

where $\mathbf{u}^V$ denotes the drift of the SDE associated with $\mathbb{Q}^{V,\sigma}$, see (28), and the potential $\phi_{t|0}$ is given by

$$\phi_{t|0}(x_t|x_0) = \mathbb{E}_{\mathbb{Q}^{V,\sigma}_{T|t}}[\phi_{T|0}(\mathbf{X}_T|x_0) \mid \mathbf{X}_t = x_t] \ .$$

Observing that $\phi_{t|0}(x_t|x_0) = \frac{\mathrm{d}\Pi_{t|0}}{\mathrm{d}\mathbb{Q}^{V,\sigma}_{t|0}}(x_t|x_0)$, we may apply Tweedie's formula to obtain

$$\nabla\log\phi_{t|0}(x_t|x_0) = \mathbb{E}_{\Pi_{T|0,t}}[\nabla_{\mathbf{X}_t}\log\mathbb{Q}^{V,\sigma}_{T|t}(\mathbf{X}_T|\mathbf{X}_t) \mid \mathbf{X}_t = x_t, \mathbf{X}_0 = x_0] \ .$$

We now compute the score $\nabla\log\mathbb{Q}^{V,\sigma}_{T|t}$. By definition,

$$\mathrm{d}\mathbb{Q}^{V,\sigma}_{T|t}(x_T|x_t) = \frac{1}{Z_t(x_t)}\overbrace{\int\exp\left(-\int_t^T\frac{V_s(x_s)}{\sigma^2}\mathrm{d}s\right)\mathrm{d}\mathbb{Q}^{\sigma}(x_{(t,T)}|x_t,x_T)}^{F(x_t,x_T)}\mathrm{d}\mathbb{Q}^{\sigma}_{T|t}(x_T|x_t)$$
$$\text{where } Z_t(x_t) = \int\exp\left(-\int_t^T\frac{V_s(x'_s)}{\sigma^2}\mathrm{d}s\right)\mathrm{d}\mathbb{Q}^{\sigma}(x'_{(t,T]}|x_t) \ .$$

Therefore,

$$\nabla_{x_t}\log\mathbb{Q}^{V,\sigma}_{T|t}(x_T|x_t) = -\nabla\log Z_t(x_t) + \nabla_{x_t}\log\mathbb{Q}^{\sigma}_{T|t}(x_T|x_t) + \nabla_{x_t}\log F(x_t,x_T) \ .$$

To compute $\nabla_{x_t}\log F(x_t,x_T)$, we use the Brownian bridge reparameterization inside $F(x_t,x_T)$. This yields

$$F(x_t,x_T) = \int_0^{T-t}\exp\left(-\int_t^T\frac{1}{\sigma^2}V_s\left(\frac{s-t}{T-t}x_T + \frac{T-s}{T-t}x_t + \sigma\{\mathbf{B}_{s-t} - \frac{s-t}{T-t}\mathbf{B}_{T-t}\}\right)\right)\mathrm{d}\mathbf{B}_u \ .$$

Differentiating under the integral sign, we obtain

$$\nabla_{x_t}\log F(x_t,x_T) = \frac{1}{\sigma^2}\mathbb{E}_{\mathbb{Q}^{V,\sigma}_{|t,T}}\left[-\int_t^T\frac{T-s}{T-t}\nabla V_s(\mathbf{X}_s)\mathrm{d}s \mid \mathbf{X}_t = x_t, \mathbf{X}_T = x_T\right] \ .$$

Using the same reparameterization argument for $Z_t(x_t)$, we further obtain

$$\nabla_{x_t}\log Z_t(x_t) = \frac{1}{\sigma^2}\mathbb{E}_{\mathbb{Q}^{V,\sigma}_{|t}}\left[-\int_t^T\nabla V_s(\mathbf{X}_s)\mathrm{d}s \mid \mathbf{X}_t = x_t\right] = \frac{\mathbf{u}^V_t(x_t)}{\sigma^2} \ .$$

Combining the above identities, we conclude that $\Pi_{|0}(\cdot|x_0)$ is associated with the SDE

$$\mathrm{d}\mathbf{X}_t = \underbrace{\mathbb{E}_{\Pi_{|0,t}}\left[\sigma^2\nabla_{x_t}\log\mathbb{Q}^{\sigma}_{T|t}(\mathbf{X}_T|\mathbf{X}_t) - \int_t^T\frac{T-s}{T-t}\nabla V_s(\mathbf{X}_s)\mathrm{d}s \mid \mathbf{X}_t, \mathbf{X}_0\right]}_{\tilde{\mathbf{b}}^{\mathrm{f},V}_t(\mathbf{X}_t,\mathbf{X}_0)}\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \ , \ \mathbf{X}_0 = x_0 \ .$$

Let $b \in \mathbf{A}_T$, and denote by $\mathbb{M}^b$ the path measure associated with

$$\mathrm{d}\mathbf{X}_t = b_t(\mathbf{X}_t)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \ , \quad \mathbf{X}_0 \sim \Pi_0 \ .$$

Since $\mathbb{M}_0^b = \Pi_0$, the Markovian projection objective writes by Girsanov's theorem as

$$\mathrm{KL}(\Pi \| \mathbb{M}) = \frac{1}{2\sigma^2} \int_0^T \mathbb{E}_{\Pi_{0,t}} \left[\left\| \tilde{\mathrm{b}}_t^{\mathrm{f},V}(\mathbf{X}_t, \mathbf{X}_0) - b_t(\mathbf{X}_t) \right\|^2 \right] \mathrm{d}t \ .$$

This expression is minimized when $b_t = \mathrm{b}_t^{\mathrm{f},V}$, defined by (37), proving that $\mathrm{proj}_{\mathcal{M}_T}(\Pi)$ is associated with (35). Then, the marginal-preserving property of the Markovian projection of Π follows directly from two facts: the optimal drift is given by conditional expectations of the conditional drifts of $\mathbb{Q}_{|0,T}^{V,\sigma}$ under $\Pi_{0,T|t}$, and Π satisfies the reciprocal property. Combining these observations, one can show, by applying [41, Theorem 2], that Π and $\mathrm{proj}_{\mathcal{M}_T}(\Pi)$ solve the same Fokker–Planck equation and therefore share the same time marginals.

To obtain a computable formulation of the Markovian projection step, we replace the original Markovian projection step (10) by the variational problem

$$\arg\min\{\mathrm{KL}(\mathrm{proj}_{\mathcal{M}_T}(\Pi) \| \mathbb{M}_{\theta^{\mathrm{f}}}) \ : \ \theta^{\mathrm{f}} \in \Theta\} \ ,$$

where $\mathbb{M}_{\theta^{\mathrm{f}}}$ is associated with

$$\mathrm{d}\mathbf{X}_t = v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t \ , \quad \mathbf{X}_0 \sim \Pi_0 \ .$$

For any $\theta^{\mathrm{f}} \in \Theta$, by Girsanov's theorem, we have

$$\begin{aligned} \mathrm{KL}(\mathrm{proj}_{\mathcal{M}_T}(\Pi) \| \mathbb{M}_{\theta^{\mathrm{f}}}) &= \frac{1}{2\sigma^2} \int_0^T \mathbb{E}_{\mathrm{proj}_{\mathcal{M}_T}(\Pi)_t} \left[\left\| \mathrm{b}_t^{\mathrm{f},V}(\mathbf{X}_t) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t) \right\|^2 \right] \mathrm{d}t \\ &= \frac{1}{2\sigma^2} \int_0^T \mathbb{E}_{\Pi_t} \left[\left\| \mathrm{b}_t^{\mathrm{f},V}(\mathbf{X}_t) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t) \right\|^2 \right] \mathrm{d}t \\ &= \frac{1}{2\sigma^2} \{ \mathcal{L}_{\mathcal{M}_T}^{\mathrm{f}}(\theta^{\mathrm{f}}) - \mathrm{c}^{\mathrm{f}} \} \ , \end{aligned} \tag{40}$$

where

$$\begin{aligned} \mathrm{c}^{\mathrm{f}} = \int_0^T (T-t)^{-1} \int_t^T \mathbb{E}_{\Pi_t} \Big[\mathbb{E}_{\Pi_{s,T|t}} \Big[\Big\| \mathbf{b}_{t,s}^{\mathrm{f},V}(\mathbf{X}_t, \mathbf{X}_s, \mathbf{X}_T) \Big\|^2 \Big] \\ - \Big\| (T-t)^{-1} \int_t^T \mathbb{E}_{\Pi_{u,T|t}} [\mathbf{b}_{t,u}^{\mathrm{f},V}(\mathbf{X}_t, \mathbf{X}_u, \mathbf{X}_T)] \mathrm{d}u \Big\|^2 \Big] \, \mathrm{d}s\mathrm{d}t \ , \end{aligned}$$

with $\mathbf{b}_{t,s}^{\mathrm{f},V}(x_t, x_s, x_T) = (x_T - x_t)/(T-t) - (T-s)\nabla V_s(x_s)$. Since c^{f} does not depend on θ^{f}, minimizing $\mathcal{L}_{\mathcal{M}_T}^{\mathrm{f}}(\theta^{\mathrm{f}})$ with respect to θ^{f} is equivalent to minimizing $\mathrm{KL}(\mathrm{proj}_{\mathcal{M}_T}(\Pi) \| \mathbb{M}_{\theta^{\mathrm{f}}})$. This yields a variational approximation of the true Markovian projection. □

Derivations with learnable control variates. We first provide the proof of Proposition 3.

Proof of Proposition 3. Let $\alpha : [0,T]^2 \to \mathbb{R}$ be continuously differentiable with respect to its second variable. We assume throughout that the quantities involved are sufficiently integrable to justify the use of Itô's formula, conditional expectations, Fubini's theorem, and the martingale property of the stochastic integral. In particular, for every fixed $t \in [0,T)$, we assume

$$\begin{aligned} &\int_t^T |\alpha(t,s)|^2 \mathrm{d}s < \infty \ , \\ &\mathbb{E}_{\Pi_{|t}} \left[\int_t^T |\alpha(t,s)| \, \|\mathbf{u}_s^{\mathrm{f},V}(\mathbf{X}_s | \mathbf{X}_T)\| \mathrm{d}s + \int_t^T \int_s^T |\alpha(t,s)| \tfrac{T-r}{T-s} \|\nabla V_r(\mathbf{X}_r)\| \mathrm{d}r\mathrm{d}s \, \Big| \, \mathbf{X}_t \right] < \infty \ . \end{aligned}$$

Fix $t \in [0,T)$, and define the following notation

$$\alpha_s = \alpha(t,s), \qquad \tilde{\alpha}_s = (T-s)\alpha_s, \qquad s \in [t,T] \ .$$

Since $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$, for almost surely any $(x_t, x_T) \sim \Pi_{t,T}$, we know from Corollary 7 that the pinned path measure $\Pi_{(t,T)|t,T}$ is Markov, induced by the pinned SDE defined on $[t,T]$ by

$$\mathrm{d}\mathbf{X}_s = \left\{ \frac{x_T - \mathbf{X}_s}{T-s} + \mathbf{u}_s^{\mathrm{f},V}(\mathbf{X}_s | x_T) \right\} \mathrm{d}s + \sigma \mathrm{d}\mathbf{B}_s \ , \ \mathbf{X}_t = x_t \ , \ \mathbf{X}_s = x_T \ .$$

We introduce the auxiliary stochastic process

$$G_s = \alpha_s(x_T - \mathbf{X}_s), \qquad s \in [t, T].$$

By Itô's formula, we obtain

$$\mathrm{d}G_s = \partial_s\alpha_s(x_T - \mathbf{X}_s)\mathrm{d}s - \alpha_s\mathrm{d}\mathbf{X}_s,$$

hence

$$\mathrm{d}G_s = \left(\partial_s\alpha_s - \frac{\alpha_s}{T-s}\right)(x_T - \mathbf{X}_s)\mathrm{d}s - \alpha_s\mathbf{u}_s^{\mathrm{f},V}(\mathbf{X}_s|x_T)\mathrm{d}s - \sigma\alpha_s\mathrm{d}\mathbf{B}_s\ .$$

Since

$$\frac{\partial_s\tilde{\alpha}_s}{T-s} = \partial_s\alpha_s - \frac{\alpha_s}{T-s}\ ,$$

we have

$$\mathrm{d}G_s = \frac{\partial_s\tilde{\alpha}_s}{T-s}(x_T - \mathbf{X}_s)\mathrm{d}s - \alpha_s\mathbf{u}_s^{\mathrm{f},V}(\mathbf{X}_s|x_T)\mathrm{d}s - \sigma\alpha_s\mathrm{d}\mathbf{B}_s\ .$$

Integrating from t to T, and using $G_T = \alpha_T(x_T - x_T) = 0$, we obtain

$$\alpha_t(x_T - x_t) + \int_t^T \frac{\partial_s\tilde{\alpha}_s}{T-s}(x_T - \mathbf{X}_s)\mathrm{d}s = \int_t^T \alpha_s\mathbf{u}_s^{\mathrm{f},V}(\mathbf{X}_s|x_T)\mathrm{d}s + \sigma\int_t^T \alpha_s\mathrm{d}\mathbf{B}_s\ .$$

We next take the conditional expectation of the expressions above with respect to $\Pi_{(t,T]|t}$. In particular, the stochastic term vanishes and we obtain

$$\mathbb{E}_{\Pi_{|t}}\left[\alpha_t(\mathbf{X}_T - \mathbf{X}_t) + \int_t^T \tfrac{\partial_s\tilde{\alpha}_s}{T-s}(\mathbf{X}_T - \mathbf{X}_s)\mathrm{d}s \;\middle|\; \mathbf{X}_t = x_t\right] = \mathbb{E}_{\Pi_{|t}}\left[\int_t^T \alpha_s\mathbf{u}_s^{\mathrm{f},V}(\mathbf{X}_s|\mathbf{X}_T)\mathrm{d}s \;\middle|\; \mathbf{X}_t = x_t\right].$$

Using the definition of $\mathbf{u}_s^{\mathrm{f},V}$, the tower property, and Fubini's theorem, the right-hand side becomes

$$\begin{aligned}
&\mathbb{E}_{\Pi_{|t}}\left[\int_t^T \alpha_s\mathbf{u}_s^{\mathrm{f},V}(\mathbf{X}_s|\mathbf{X}_T)\mathrm{d}s \;\middle|\; \mathbf{X}_t = x_t\right] \\
&= \int_t^T \alpha_s\,\mathbb{E}_{\Pi_{|t}}\left[\mathbb{E}_{\Pi_{|s,T}}\left[-\int_s^T \frac{T-r}{T-s}\nabla V_r(\mathbf{X}_r)\mathrm{d}r \;\middle|\; \mathbf{X}_s, \mathbf{X}_T\right] \;\middle|\; \mathbf{X}_t = x_t\right]\mathrm{d}s \\
&= -\mathbb{E}_{\Pi_{|t}}\left[\int_t^T \alpha_s\int_s^T \frac{T-r}{T-s}\nabla V_r(\mathbf{X}_r)\mathrm{d}r\,\mathrm{d}s \;\middle|\; \mathbf{X}_t = x_t\right] \\
&= -\mathbb{E}_{\Pi_{|t}}\left[\int_t^T \left(\int_t^s \frac{\alpha_u}{T-u}\mathrm{d}u\right)(T-s)\nabla V_s(\mathbf{X}_s)\mathrm{d}s \;\middle|\; \mathbf{X}_t = x_t\right].
\end{aligned}$$

Thus,

$$\begin{aligned}
\mathbb{E}_{\Pi_{|t}}\Bigg[&\alpha(t,t)(\mathbf{X}_T - \mathbf{X}_t) + \int_t^T \frac{\partial_s\tilde{\alpha}(t,s)}{T-s}(\mathbf{X}_T - \mathbf{X}_s)\mathrm{d}s \\
&+ \int_t^T \left(\int_t^s \frac{\alpha(t,u)}{T-u}\mathrm{d}u\right)(T-s)\nabla V_s(\mathbf{X}_s)\mathrm{d}s \mid \mathbf{X}_t\Bigg] = \mathbb{E}_{\Pi_{|t}}[H_t^{\alpha,\mathrm{f}}|\mathbf{X}_t] = 0.
\end{aligned}$$

□

Backward control-variate identity. For α^{b} satisfying the same regularity and integrability assumptions as in the forward case, let $\tilde{\alpha}^{\mathrm{b}}(t,s) = s\alpha^{\mathrm{b}}(t,s)$ and define

$$H_t^{\alpha,\mathrm{b}} = \alpha(t,t)[\mathbf{X}_0 - \mathbf{X}_t] - \int_0^t \left\{\frac{\partial_s\tilde{\alpha}^{\mathrm{b}}(t,s)}{s}[\mathbf{X}_0 - \mathbf{X}_s] + s\nabla V_s(\mathbf{X}_s)\int_s^t \frac{\alpha(t,u)}{u}\mathrm{d}u\right\}\mathrm{d}s\ ,$$

Applying the same Itô argument to the time-reversed pinned process gives $\mathbb{E}_{\Pi_{|t}}\left[H_t^{\alpha^{\mathrm{b}},\mathrm{b}} \mid \mathbf{X}_t\right] = 0$, for almost every $t \in (0,T)$.

Informal explanation of Proposition 3. At first sight, the quantity $H_t^{\alpha,\mathrm{f}}$ introduced in Proposition 3 may appear somewhat ad hoc. We explain here, in the forward case, why they arise naturally as *score-based* quantities.

Recall from (37) that the drift of the forward Markovian projection, $\mathrm{b}_t^{\mathrm{f},V}$, is written as a conditional expectation over the future segment $(X_s)_{s\in(t,T]} \sim \Pi_{(t,T]|t}$. Consider a time discretization $(s_k)_{k=0}^{K+1}$ of $[t,T]$, with $s_0 = t$, $s_{K+1} = T$, and constant step size $\delta > 0$. A discrete-time approximation of $\Pi_{(t,T]|t}$ is obtained through its finite-dimensional projection:

$$\tilde{\Pi}_{s_1,\dots,s_{K+1}|t}(x_{s_1:s_{K+1}} \mid x_t) = \int_{\mathbb{R}^d} \mathbb{Q}^{V,\sigma}_{s_1,\dots,s_K|t,T}(x_{s_1:s_K} \mid x_t, x_T)\mathrm{d}\Pi_{T|t}(x_T \mid x_t) \ .$$

By the Markov property of $\mathbb{Q}^{V,\sigma}_{|0,T}$, this conditional path measure factorizes as

$$\mathbb{Q}^{V,\sigma}_{s_1,\dots,s_K|t,T}(x_{s_1:s_K} \mid x_t, x_T) = \prod_{k=1}^K \mathbb{Q}^{V,\sigma}_{s_k|t,s_{k+1}}(x_{s_k} \mid x_t, x_{s_{k+1}}) \ .$$

Consequently, for each $k \in \{1,\dots,K\}$, the usual zero-score identity gives

$$\int_{\mathbb{R}^d} \nabla \log \mathbb{Q}^{V,\sigma}_{s_k|t,s_{k+1}}(x_{s_k} \mid x_t, x_{s_{k+1}})\, \mathrm{d}\mathbb{Q}^{V,\sigma}_{s_k|t,s_{k+1}}(x_{s_k} \mid x_t, x_{s_{k+1}}) = 0 \ .$$

Combining this identity with the definition of $\tilde{\Pi}$ yields

$$\mathbb{E}_{\tilde{\Pi}_{|t}}\left[\sum_{k=1}^K \delta\sigma^2\alpha(t,s_k)\nabla \log \mathbb{Q}^{V,\sigma}_{s_k|t,s_{k+1}}(\mathbf{X}_{s_k} \mid \mathbf{X}_t, \mathbf{X}_{s_{k+1}}) \,\middle|\, \mathbf{X}_t\right] = 0 \ .$$

In the continuous-time limit $\delta \to 0$, this identity becomes precisely the forward identity in Proposition 3, thereby providing a score-based interpretation of the proposed quantities.

Relying on Proposition 3 and its backward counterpart, we provide in Proposition 11 strictly equivalent expressions for the optimal control drifts derived in Proposition 10 based on forward and backward functions α^{f} and α^{b}, formally acting as control variates.

Proposition 11 (Markovian projection with control variates – continuous time)**.** *Let $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$, α^{f}, $\alpha^{\mathrm{b}} : [0,T]^2 \to \mathbb{R}$, continuously differentiable in their second variable. Define the velocity fields*

$$\begin{aligned}\mathrm{b}_t^{\alpha^{\mathrm{f}},V}(x_t) = \mathbb{E}_{\Pi_{|t}}\Big[&\tfrac{\mathbf{X}_T-\mathbf{X}_t}{T-t} - \alpha^{\mathrm{f}}(t,t)[\mathbf{X}_T-\mathbf{X}_t] - \int_t^T \tfrac{\partial_s\tilde{\alpha}^{\mathrm{f}}(t,s)}{T-s}[\mathbf{X}_T-\mathbf{X}_s]\mathrm{d}s \\ &- \int_t^T \left\{\tfrac{1}{T-t} + \int_t^s \tfrac{\alpha^{\mathrm{f}}(t,u)}{T-u}\mathrm{d}u\right\}(T-s)\nabla V_s(\mathbf{X}_s)\mathrm{d}s \,\middle|\, \mathbf{X}_t = x_t\Big] \ , \end{aligned} \tag{41}$$

$$\begin{aligned}\mathrm{b}_t^{\alpha^{\mathrm{b}},V}(x_t) = \mathbb{E}_{\Pi_{|t}}\Big[&\tfrac{\mathbf{X}_0-\mathbf{X}_t}{t} - \alpha^{\mathrm{b}}(t,t)[\mathbf{X}_0-\mathbf{X}_t] + \int_0^t \tfrac{\partial_s\tilde{\alpha}^{\mathrm{b}}(t,s)}{s}[\mathbf{X}_0-\mathbf{X}_s]\mathrm{d}s \\ &- \int_0^t \left\{\tfrac{1}{t} - \int_s^t \tfrac{\alpha^{\mathrm{b}}(t,u)}{u}\mathrm{d}u\right\} s\nabla V_s(\mathbf{X}_s)\mathrm{d}s \,\middle|\, \mathbf{X}_t = x_t\Big] \ , \end{aligned} \tag{42}$$

where $\tilde{\alpha}^{\mathrm{b}}(t,s) = s\alpha^{\mathrm{b}}(t,s)$ and $\tilde{\alpha}^{\mathrm{f}}(t,s) = (T-s)\alpha^{\mathrm{f}}(t,s)$. Then, we have the equality $\mathrm{b}^{\alpha^{\mathrm{f}},V} = \mathrm{b}^{\mathrm{f},V}$ and $\mathrm{b}^{\alpha^{\mathrm{b}},V} = \mathrm{b}^{\mathrm{b},V}$, where the drifts $\mathrm{b}^{\mathrm{f},V}$ and $\mathrm{b}^{\mathrm{b},V}$ are respectively defined in (37) and (38).

Proof. The forward result derives from the linearity of the expectation, observing that

$$\mathrm{b}_t^{\alpha^{\mathrm{f}},V}(x_t) = \mathrm{b}_t^{\mathrm{f},V}(x_t) - \mathbb{E}_{\Pi_{|t}}[H_t^{\alpha^{\mathrm{f}},\mathrm{f}} \mid \mathbf{X}_t = x_t] \ , \ \mathrm{b}_t^{\alpha^{\mathrm{b}},V}(x_t) = \mathrm{b}_t^{\mathrm{b},V}(x_t) - \mathbb{E}_{\Pi_{|t}}[H_t^{\alpha^{\mathrm{b}},\mathrm{b}} \mid \mathbf{X}_t = x_t].$$

□

Since the obtained α-based optimal drifts write as conditional expectations, we propose to learn them via parameterized families via a joint minimization loss that also includes the scalar coefficients α^{f} and α^{b}. Interestingly, we show these joint losses follow a bias-variance decomposition, where the bias term is the original TSBM loss, while the variance term is minimized with respect to α^{f} or α^{b}. This explains the interest of using learnable control variates in our setting.

Proposition 12 (TSBM bridge matching losses with control variates – continuous time)**.** *Let $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$. Consider the respective forward and backward objectives*

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}},\alpha^{\mathrm{f}}) = \int_0^T (T-t)^{-1}\int_t^T \mathbb{E}_{\Pi_{t,s,T}}\left[\left\|\mathbf{b}^{\alpha^{\mathrm{f}},V}_{t,s}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\right\|^2\right]\mathrm{d}s\mathrm{d}t \ , \tag{43}$$

$$\mathcal{L}^{\mathrm{b}}_{\mathcal{M}_T}(\theta^{\mathrm{b}},\alpha^{\mathrm{b}}) = \int_0^T t^{-1}\int_0^t \mathbb{E}_{\Pi_{0,s,t}}\left[\left\|\mathbf{b}^{\alpha^{\mathrm{b}},V}_{t,s}(\mathbf{X}_0,\mathbf{X}_s,\mathbf{X}_t) - v_t^{\theta^{\mathrm{b}}}(\mathbf{X}_t)\right\|^2\right]\mathrm{d}s\mathrm{d}t \ , \tag{44}$$

where

$$\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(x_t,x_s,x_T) = (x_T-x_t)/(T-t) - \alpha^{\mathrm{f}}(t,t)[x_T-x_t] - \frac{\partial_s\tilde{\alpha}^{\mathrm{f}}(t,s)}{T-s}(T-t)[x_T-x_s] \quad (45)$$
$$- \left\{1+(T-t)\int_t^s \frac{\alpha^{\mathrm{f}}(t,u)}{T-u}\mathrm{d}u\right\}(T-s)\nabla V_s(x_s)\,,$$

$$\mathbf{b}_{t,s}^{\alpha^{\mathrm{b}},V}(x_0,x_s,x_t) = (x_0-x_t)/t - \alpha^{\mathrm{b}}(t,t)[x_0-x_t] + \frac{\partial_s\tilde{\alpha}^{\mathrm{b}}(t,s)}{s}t[x_0-x_s] \quad (46)$$
$$- \left\{1-t\int_s^t \frac{\alpha^{\mathrm{b}}(t,u)}{u}\mathrm{d}u\right\}s\nabla V_s(x_s)\,,$$

with $\tilde{\alpha}^{\mathrm{f}}(t,s) = (T-s)\alpha^{\mathrm{f}}(t,s)$ *and* $\tilde{\alpha}^{\mathrm{b}}(t,s) = s\alpha^{\mathrm{b}}(t,s)$*, where* α^{f} *and* $\alpha^{\mathrm{b}}\,[0,T]^2 \to \mathbb{R}$ *are continuously differentiable in their second variable.*

Then, up to additive numerical constants independent of θ^{f}*,* θ^{b}*,* α^{f}*,* α^{b}*, the following bias-variance loss decomposition holds*

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}},\alpha^{\mathrm{f}}) = \mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}) + \mathcal{L}^{\mathrm{f}}_{CV}(\alpha^{\mathrm{f}})\,,$$
$$\mathcal{L}^{\mathrm{b}}_{\mathcal{M}_T}(\theta^{\mathrm{b}},\alpha^{\mathrm{b}}) = \mathcal{L}^{\mathrm{b}}_{\mathcal{M}_T}(\theta^{\mathrm{b}}) + \mathcal{L}^{\mathrm{b}}_{CV}(\alpha^{\mathrm{b}})\,,$$

with

$$\mathcal{L}^{\mathrm{f}}_{CV}(\alpha^{\mathrm{f}}) = \int_0^T (T-t)^{-1}\int_t^T \mathbb{E}_{\Pi_t}\left[\mathbb{E}_{\Pi_{s,T|t}}\left[\left\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\right\|^2\right]\right.$$
$$\left. - \left\|(T-t)^{-1}\int_t^T \mathbb{E}_{\Pi_{u,T|t}}[\mathbf{b}_{t,u}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_u,\mathbf{X}_T)]\mathrm{d}u\right\|^2\right]\mathrm{d}s\mathrm{d}t\,,$$

$$\mathcal{L}^{\mathrm{b}}_{CV}(\alpha^{\mathrm{b}}) = \int_0^T t^{-1}\int_0^t \mathbb{E}_{\Pi_t}\left[\mathbb{E}_{\Pi_{0,s|t}}\left[\left\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{b}},V}(\mathbf{X}_0,\mathbf{X}_s,\mathbf{X}_t)\right\|^2\right]\right.$$
$$\left. - \left\|t^{-1}\int_0^t \mathbb{E}_{\Pi_{0,u|t}}[\mathbf{b}_{t,u}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_0,\mathbf{X}_u,\mathbf{X}_t)]\mathrm{d}u\right\|^2\right]\mathrm{d}s\mathrm{d}t\,.$$

Therefore, minimizing $\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}},\alpha^{\mathrm{f}})$ *with respect to* $(\theta^{\mathrm{f}},\alpha^{\mathrm{f}})$ *gives* $\theta^{\mathrm{f}}_\star$*, the original minimizer from* (18)*, and* $\alpha^{\mathrm{f}}_\star$ *that minimizes* $\mathcal{L}^{\mathrm{f}}_{CV}$*, i.e., the scalar variance of the regression target. The same result applies for the backward direction.*

Proof of Proposition 12. We only prove the forward identity, the backward one being identical after time reversal. By Proposition 11, for almost surely any $x_t \sim \Pi_t$,

$$\mathrm{b}_t^{\mathrm{f},V}(x_t) = \mathrm{b}_t^{\alpha^{\mathrm{f}},V}(x_t) = \frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\mid \mathbf{X}_t = x_t\right]\mathrm{d}s.$$

Let us write

$$m_t = \mathrm{b}_t^{\mathrm{f},V}(\mathbf{X}_t), \qquad y_t = v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t).$$

We begin with the following computation

$$\frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T) - y_t\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s$$
$$= \frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s + \frac{1}{T-t}\int_t^T \|y_t\|^2\mathrm{d}s$$
$$- \frac{2}{T-t}\int_t^T \left\langle \mathbb{E}_{\Pi_{s,T|t}}\left[\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\mid \mathbf{X}_t\right], y_t\right\rangle \mathrm{d}s$$
$$= \frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s + \|y_t\|^2 - 2\langle m_t, y_t\rangle$$
$$= \frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s + \|m_t - y_t\|^2 - \|m_t\|^2\,.$$

Rearranging the obtain terms, we have

$$\|m_t - y_t\|^2 = \frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T) - y_t\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s \\ - \left\{\frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s - \|m_t\|^2\right\}.$$

We now take the expectation with respect to $\mathbf{X}_t \sim \Pi_t$, and integrate over $t \in [0,T]$, to get

$$\begin{aligned}
&\int_0^T \mathbb{E}_{\Pi_t}\left[\|\mathrm{b}_{t,s}^{\mathrm{f},V}(\mathbf{X}_t) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\|^2\right]\mathrm{d}t \\
&\quad= \int_0^T \mathbb{E}_{\Pi_t}\left[\frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s\right]\mathrm{d}t \\
&\qquad- \int_0^T \mathbb{E}_{\Pi_t}\left[\frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s - \|\mathrm{b}_t^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t)\|^2\right]\mathrm{d}t \\
&\quad= \int_0^T \frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{t,s,T}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s\mathrm{d}t \\
&\qquad- \int_0^T \mathbb{E}_{\Pi_t}\left[\frac{1}{T-t}\int_t^T \mathbb{E}_{\Pi_{s,T|t}}\left[\|\mathbf{b}_{t,s}^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t,\mathbf{X}_s,\mathbf{X}_T)\|^2 \mid \mathbf{X}_t\right]\mathrm{d}s - \|\mathrm{b}_t^{\alpha^{\mathrm{f}},V}(\mathbf{X}_t)\|^2\right]\mathrm{d}t\,.
\end{aligned}$$

This amounts exactly to the identity

$$\int_0^T \mathbb{E}_{\Pi_t}\left[\|\mathrm{b}_t^{\mathrm{f},V}(\mathbf{X}_t) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\|^2\right]\mathrm{d}t = \mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}},\alpha^{\mathrm{f}}) - \mathcal{L}^{\mathrm{f}}_{\mathrm{CV}}(\alpha^{\mathrm{f}})\,.$$

Recall that, by (40), there exists a constant c^f, independent of θ^{f} and α^{f}, such that

$$\int_0^T \mathbb{E}_{\Pi_t}\left[\|\mathrm{b}_t^{\mathrm{f},V}(\mathbf{X}_t) - v_t^{\theta^{\mathrm{f}}}(\mathbf{X}_t)\|^2\right]\mathrm{d}t = \mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}) + \mathrm{c}^{\mathrm{f}}\,.$$

Combining the two identities yields

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}},\alpha^{\mathrm{f}}) = \mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}) + \mathcal{L}^{\mathrm{f}}_{\mathrm{CV}}(\alpha^{\mathrm{f}}) + \mathrm{c}^{\mathrm{f}}\,.$$

□

C.4 Results on the Markovian projection (discrete time)

In this section, we assume that $\mathbb{Q}^{V,\sigma}$ is given by (6). As in the continuous-time setting, we derive the optimal control drifts of the Markovian projection, along with their equivalent α-based reformulation and the corresponding loss functions. We notably show that the potential-based control term admits an analogous interpretation: it can be written as an average contribution over future trajectory states at which the state cost is activated. We omit the proofs as they closely parallel those of the continuous-time case.

Proposition 13 (Markovian projection – discrete time)**.** *Let $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$. Then, $\mathrm{proj}_{\mathcal{M}_T}(\Pi)$ is equivalently induced by the diffusion processes $(\mathbf{X}_t^{\mathrm{f}})_{t\in[0,T]}$ and $(\mathbf{X}_{T-t}^{\mathrm{b}})_{t\in[0,T]}$, which are time-reversed of each other and respectively defined by the SDEs*

$$\mathrm{d}\mathbf{X}_t^{\mathrm{f}} = \mathrm{b}_t^{\mathrm{f},V}(\mathbf{X}_t^{\mathrm{f}})\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t\,, \quad \mathbf{X}_0^{\mathrm{f}} \sim \Pi_0\,,$$
$$\mathrm{d}\mathbf{X}_t^{\mathrm{b}} = \mathrm{b}_{T-t}^{\mathrm{b},V}(\mathbf{X}_t^{\mathrm{b}})\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t\,, \quad \mathbf{X}_0^{\mathrm{b}} \sim \Pi_T\,,$$

where

$$\mathrm{b}_t^{\mathrm{f},V}(x_t) = \mathbb{E}_{\Pi_{|t}}\left[(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - \tfrac{1}{K_f(t)}\textstyle\sum_{t_k>t}(T-t_k)\nabla V_k(\mathbf{X}_{t_k})|\mathbf{X}_t = x_t\right]\,, \tag{47}$$

$$\mathrm{b}_t^{\mathrm{b},V}(x_t) = \mathbb{E}_{\Pi_{|t}}\left[(\mathbf{X}_0 - \mathbf{X}_t)/t - \tfrac{1}{K_b(t)}\textstyle\sum_{t_k<t} t_k\nabla V_k(\mathbf{X}_{t_k})|\mathbf{X}_t = x_t\right]\,, \tag{48}$$

where $K_f(t) = |\{k \in \{1,\dots,K\} : t_k > t\}|$ and $K_b(t) = |\{k \in \{1,\dots,K\} : t_k < t\}|$.

Moreover, for any $t \in [0,T]$, we have $\mathrm{Law}(\mathbf{X}^{\mathrm{f}}_t) = \mathrm{Law}(\mathbf{X}^{\mathrm{b}}_{T-t}) = \Pi_t$. Additionally, $\mathrm{b}^{\mathrm{f},V}$ and $\mathrm{b}^{\mathrm{b},V}$ can be approximated by $v^{\theta^{\mathrm{f}}_\star}$ and $v^{\theta^{\mathrm{b}}_\star}$, where $\theta^{\mathrm{f}}_\star$ and $\theta^{\mathrm{b}}_\star$ are respectively the minimizers of forward and backward objectives $\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}$ and $\mathcal{L}^{\mathrm{b}}_{\mathcal{M}_T}$, defined by

$$\begin{aligned}
\mathcal{L}^{f}_{\mathcal{M}_T}(\theta^{\mathrm{f}}) &= \int_0^{t_K} \frac{1}{K_f(t)} \sum_{t_k > t} \mathbb{E}_{\Pi_{t,t_k,T}} \left[\left\| \mathbf{b}^{\mathrm{f},V}_{t,k}(\mathbf{X}_t, \mathbf{X}_{t_k}, \mathbf{X}_T) - v^{\theta^{\mathrm{f}}}_t(\mathbf{X}_t) \right\|^2 \right] \mathrm{d}t \\
&\quad + \int_{t_K}^{T} \mathbb{E}_{\Pi_{t,T}} \left[\left\| (\mathbf{X}_T - \mathbf{X}_t)/(T-t) - v^{\theta^{\mathrm{f}}}_t(\mathbf{X}_t) \right\|^2 \right] \mathrm{d}t \,, \\
\mathcal{L}^{b}_{\mathcal{M}_T}(\theta^{\mathrm{b}}) &= \int_{t_1}^{T} \frac{1}{K_b(t)} \sum_{t_k < t} \mathbb{E}_{\Pi_{0,t_k,t}} \left[\left\| \mathbf{b}^{\mathrm{b},V}_{t,k}(\mathbf{X}_0, \mathbf{X}_{t_k}, \mathbf{X}_t) - v^{\theta^{\mathrm{b}}}_t(\mathbf{X}_t) \right\|^2 \right] \mathrm{d}t \\
&\quad + \int_{0}^{t_1} \mathbb{E}_{\Pi_{0,t}} \left[\left\| (\mathbf{X}_0 - \mathbf{X}_t)/t - v^{\theta^{\mathrm{b}}}_t(\mathbf{X}_t) \right\|^2 \right] \mathrm{d}t \,,
\end{aligned} \tag{49}$$

with the following regression targets

$$\begin{aligned}
\mathbf{b}^{\mathrm{f},V}_{t,k}(x_t, x_{t_k}, x_T) &= (x_T - x_t)/(T-t) - (T-t_k)\nabla V_k(x_{t_k}) \,, \\
\mathbf{b}^{\mathrm{b},V}_{t,k}(x_0, x_{t_k}, x_t) &= (x_0 - x_t)/t - t_k \nabla V_k(x_{t_k}) \,.
\end{aligned}$$

Proposition 14 (Markovian projection with control variates – discrete time)**.** *Let $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$, and α^{f}, $\alpha^{\mathrm{b}} : [0,T]^2 \to \mathbb{R}$, continuously differentiable in their second variable.. Define the velocity fields*

$$\begin{aligned}
\mathrm{b}^{\alpha^{\mathrm{f}},V}_t(x_t) &= \mathbb{E}_{\Pi_{|t}} \Big[\tfrac{\mathbf{X}_T - \mathbf{X}_t}{T-t} - \alpha^{\mathrm{f}}(t,t)[\mathbf{X}_T - \mathbf{X}_t] - \textstyle\int_t^T \frac{\partial_s \tilde{\alpha}^{\mathrm{f}}(t,s)}{T-s}[\mathbf{X}_T - \mathbf{X}_s]\mathrm{d}s \\
&\qquad - \textstyle\sum_{t_k > t} \left\{ \frac{1}{K_f(t)} + \int_t^{t_k} \frac{\alpha^{\mathrm{f}}(t,u)}{K_f(u)} \mathrm{d}u \right\} (T-t_k)\nabla V_k(\mathbf{X}_{t_k}) | \mathbf{X}_t = x_t \Big] \,, \\
\mathrm{b}^{\alpha^{\mathrm{b}},V}_t(x_t) &= \mathbb{E}_{\Pi_{|t}} \Big[\tfrac{\mathbf{X}_0 - \mathbf{X}_t}{t} - \alpha^{\mathrm{b}}(t,t)[\mathbf{X}_0 - \mathbf{X}_t] + \textstyle\int_0^t \frac{\partial_s \tilde{\alpha}^{\mathrm{b}}(t,s)}{s}[\mathbf{X}_0 - \mathbf{X}_s]\mathrm{d}s \\
&\qquad - \textstyle\sum_{t_k < t} \left\{ \frac{1}{K_b(t)} - \int_{t_k}^t \frac{\alpha^{\mathrm{b}}(t,u)}{K_b(u)} \mathrm{d}u \right\} t_k \nabla V_k(\mathbf{X}_{t_k}) | \mathbf{X}_t = x_t \Big] \,,
\end{aligned}$$

where $\tilde{\alpha}^{\mathrm{b}}(t,s) = s\alpha^{\mathrm{b}}(t,s)$ and $\tilde{\alpha}^{\mathrm{f}}(t,s) = (T-s)\alpha^{\mathrm{f}}(t,s)$. Then, we have the equality $\mathrm{b}^{\alpha^{\mathrm{f}},V} = \mathrm{b}^{\mathrm{f},V}$ and $\mathrm{b}^{\alpha^{\mathrm{b}},V} = \mathrm{b}^{\mathrm{b},V}$, where the drifts $\mathrm{b}^{\mathrm{f},V}$ and $\mathrm{b}^{\mathrm{b},V}$ are respectively defined in (47) and (48)[4].

Proposition 15 (TSBM bridge matching losses with control variates – discrete time)**.** *Let $\Pi \in \mathcal{R}(\mathbb{Q}^{V,\sigma})$. Consider the respective forward and objectives*

$$\begin{aligned}
\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}, \alpha^{\mathrm{f}}) &= \textstyle\int_0^{t_K} \mathrm{d}t \sum_{t_k > t} \frac{1}{K_f(t)} \int_{t_{k-1}(t)}^{t_k} \frac{\mathrm{d}s}{t_k - t_{k-1}} \int_{t_K}^T \frac{\mathrm{d}r}{T - t_K} \mathbb{E}_{\Pi_{t,s,t_k,r,T}} \left[\left\| \mathbf{b}^{\alpha^{\mathrm{f}},V}_{t,s,t_k,r} - v^{\theta^{\mathrm{f}}}_t(\mathbf{X}_t) \right\|^2 \right] \\
&\quad + \textstyle\int_{t_K}^T \mathrm{d}t \int_t^T \frac{\mathrm{d}s}{T-t} \mathbb{E}_{\Pi_{t,s,T}} \left[\left\| \mathbf{b}^{\alpha^{\mathrm{f}}}_{t,s} - v^{\theta^{\mathrm{f}}}_t(\mathbf{X}_t) \right\|^2 \right] \,,
\end{aligned} \tag{50}$$

$$\begin{aligned}
\mathcal{L}^{\mathrm{b}}_{\mathcal{M}_T}(\theta^{\mathrm{b}}, \alpha^{\mathrm{b}}) &= \textstyle\int_{t_1}^T \mathrm{d}t \sum_{t_k < t} \frac{1}{K_b(t)} \int_{t_k}^{t_{k+1}(t)} \frac{\mathrm{d}s}{t_{k+1} - t_k} \int_0^{t_1} \frac{\mathrm{d}r}{t_1} \mathbb{E}_{\Pi_{0,r,s,t_k,t}} \left[\left\| \mathbf{b}^{\alpha^{\mathrm{b}},V}_{t,s,t_k,r} - v^{\theta^{\mathrm{b}}}_t(\mathbf{X}_t) \right\|^2 \right] \\
&\quad + \textstyle\int_0^{t_1} \mathrm{d}t \int_0^t \frac{\mathrm{d}s}{t} \mathbb{E}_{\Pi_{0,s,t}} \left[\left\| \mathbf{b}^{\alpha^{\mathrm{b}}}_{t,s} - v^{\theta^{\mathrm{b}}}_t(\mathbf{X}_t) \right\|^2 \right] \,,
\end{aligned} \tag{51}$$

with the following notation: in the forward direction, we set

$$t_{k-1}(t) = \begin{cases} t, & \text{if } k = \arg\min\{k' : t_{k'} > t\}, \\ t_{k-1}, & \text{otherwise}, \end{cases}$$

while in the backward direction, we set

$$t_{k+1}(t) = \begin{cases} t, & \text{if } k = \arg\max\{k' : t_{k'} < t\}, \\ t_{k+1}, & \text{otherwise}. \end{cases}$$

[4]Note that the α-based expressions of the drifts are still valid in out-of-potential settings ($t > t_K$ in forward case, $t < t_0$ in backward case) : the sum over t_k is simply 0.

Moreover, we write

$$\mathbf{b}^{\alpha^{\mathrm{f}}}_{t,s} = \mathbf{b}^{\alpha^{\mathrm{f}},0}_{t,s}(\mathbf{X}_t, \mathbf{X}_s, \mathbf{X}_T), \qquad \mathbf{b}^{\alpha^{\mathrm{b}}}_{t,s} = \mathbf{b}^{\alpha^{\mathrm{b}},0}_{t,s}(\mathbf{X}_0, \mathbf{X}_s, \mathbf{X}_t),$$

where these quantities are inherited from the continuous-time case with $V \equiv 0$, see Proposition 12. The discrete-time α-based regression targets are then given by

$$\begin{aligned}
\mathbf{b}^{\alpha^{\mathrm{f}},V}_{t,s,t_k,r} &= \tfrac{\mathbf{X}_T - \mathbf{X}_t}{T-t} - \alpha^{\mathrm{f}}(t,t)[\mathbf{X}_T - \mathbf{X}_t] - \tfrac{\partial_s \tilde{\alpha}^{\mathrm{f}}(t,s)}{T-s} K_f(t)(t_k - t_{k-1})[\mathbf{X}_T - \mathbf{X}_s] \\
&\quad - \tfrac{\partial_s \tilde{\alpha}^{\mathrm{f}}(t,r)}{T-r}(T - t_K)[\mathbf{X}_T - \mathbf{X}_r] - \left\{1 + K_f(t) \textstyle\int_t^{t_k} \frac{\alpha^{\mathrm{f}}(t,u)}{K_f(u)} \mathrm{d}u\right\}(T - t_k)\nabla V_k(\mathbf{X}_{t_k}) \,, \\
\mathbf{b}^{\alpha^{\mathrm{b}},V}_{t,s,t_k,r} &= \tfrac{\mathbf{X}_0 - \mathbf{X}_t}{t} - \alpha^{\mathrm{b}}(t,t)[\mathbf{X}_0 - \mathbf{X}_t] + \tfrac{\partial_s \tilde{\alpha}^{\mathrm{b}}(t,s)}{s} K_b(t)(t_{k+1} - t_k)[\mathbf{X}_0 - \mathbf{X}_s] \\
&\quad + \tfrac{\partial_s \tilde{\alpha}^{\mathrm{f}}(t,r)}{r} t_1[\mathbf{X}_0 - \mathbf{X}_r] - \left\{1 - K_b(t) \textstyle\int_{t_k}^t \frac{\alpha^{\mathrm{b}}(t,u)}{K_b(u)} \mathrm{d}u\right\} t_k \nabla V_k(\mathbf{X}_{t_k}) \,,
\end{aligned}$$

with $\tilde{\alpha}^{\mathrm{f}}(t,s) = (T-s)\alpha^{\mathrm{f}}(t,s)$ and $\tilde{\alpha}^{\mathrm{b}}(t,s) = s\alpha^{\mathrm{b}}(t,s)$, where α^{f} and $\alpha^{\mathrm{b}}\,[0,T]^2 \to \mathbb{R}$ are continuously differentiable in their second variable.

Then, up to additive numerical constants independent of θ^{f} and α^{f}, the following bias-variance loss decomposition holds

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}, \alpha^{\mathrm{f}}) = \mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}) + \mathcal{L}^{\mathrm{f}}_{CV}(\alpha^{\mathrm{f}}) \,,$$

with

$$\begin{aligned}
\mathcal{L}^{\mathrm{f}}_{CV}(\alpha^{\mathrm{f}}) = \textstyle\int_0^{t_K} \mathrm{d}t \sum_{t_k > t} \frac{1}{K_f(t)} \int_{t_{k-1}}^{t_k} \frac{\mathrm{d}s}{t_k - t_{k-1}} \int_{t_K}^T \frac{\mathrm{d}r}{T - t_K} \mathbb{E}_{\Pi_t} \Bigg[\mathbb{E}_{\Pi_{s,t_k,r,T|t}} \left[\left\| \mathbf{b}^{\alpha^{\mathrm{f}},V}_{t,s,t_k,r} \right\|^2 \right] \\
- \left\| \textstyle\sum_{t_k > t} \frac{1}{K_f(t)} \int_{t_{k-1}}^{t_k} \frac{\mathrm{d}s}{t_k - t_{k-1}} \int_{t_K}^T \frac{\mathrm{d}r}{T - t_K} \mathbb{E}_{\Pi_{s,t_k,r,T|t}}[\mathbf{b}^{\alpha^{\mathrm{f}},V}_{t,s,t_k,r}] \right\|^2 \Bigg] \\
+ \textstyle\int_{t_K}^T (T-t)^{-1} \int_t^T \mathbb{E}_{\Pi_t} \left[\mathbb{E}_{\Pi_{s,T|t}} \left[\left\| \mathbf{b}^{\alpha^{\mathrm{f}}}_{t,s} \right\|^2 \right] - \left\| (T-t)^{-1} \int_t^T \mathbb{E}_{\Pi_{u,T|t}}[\mathbf{b}^{\alpha^{\mathrm{f}}}_{t,u}] \mathrm{d}u \right\|^2 \right] \mathrm{d}s\mathrm{d}t \,.
\end{aligned}$$

Therefore, minimizing $\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}, \alpha^{\mathrm{f}})$ with respect to $(\theta^{\mathrm{f}}, \alpha^{\mathrm{f}})$ gives $\theta^{\mathrm{f}}_\star$, the original minimizer from (49), and $\alpha^{\mathrm{f}}_\star$ that minimizes $\mathcal{L}^{\mathrm{f}}_{CV}$, i.e., the scalar variance of the regression target. The same result applies for the backward direction.

D Fully discrete-time Twisted Schrödinger Bridge Matching (D-TSBM)

In practice, numerical methods for solving SB problems rely on translating continuous-time objectives to a discrete-time setting, which can introduce unavoidable bias. To mitigate this issue, we reformulate the SB problem directly in discrete time. This strategy was recently proposed by [18] in the canonical SB case; here, we extend it to the twisted setting, which is the focus of this paper. In Section D.1, we introduce the resulting discrete-time formulation and derive an ideal IMF scheme aligned with the continuous-time procedure reviewed in Section 2. We then translate this ideal scheme into a practical algorithm, namely the discrete-time version of TSBM (D-TSBM), with the only main change being the Markovian projection, described in Section D.3. Proofs of the propositions stated in this section are deferred to Section E.

Throughout this section, we consider a time-discretization grid $\{t_n\}_{n=0}^{N+1}$ of $[0,T]$ with $(N+2)$ increasing time points, such that $t_0 = 0 < t_1 < \ldots < t_{N+1} = T$, and we define the step sizes $\delta_n = t_n - t_{n-1}$ for $n \in \{1, \ldots, N+1\}$.

D.1 Setting and ideal IMF scheme

Discrete-time twisted reference measure. Let $q^\sigma \in \mathscr{P}^{(N+2)}$ denote the joint distribution given by the finite-dimensional projection of $\mathbb{Q}^\sigma$ at times $\{t_n\}_{n=0}^{N+1}$. In particular, q^σ is Markov and can be written as

$$q^\sigma(x_{0:N+1}) = \delta_0(x_0) \prod_{n=1}^{N+1} q^\sigma_{n|n-1}(x_n|x_{n-1}) \,, \text{ where } q^\sigma_{n|n-1}(\cdot|x_{n-1}) = \mathrm{N}(x_{n-1}, \sigma^2(t_n - t_{n-1})\mathrm{I}_d) \,.$$

To define a discrete-time analogue of the SB problem (1), we introduce a twisted version of q^σ that will serve as reference. To this end, we replace the continuous-time potential V by a collection of N potentials $\{g_n\}_{n=1}^N$ satisfying $g_n \in \mathrm{C}^1(\mathbb{R}^d, (0,\infty))$, and we define the *reference* distribution $q^{g,\sigma} \in \mathscr{P}^{(N+2)}$, absolutely continuous with respect to q^σ, by

$$q^{g,\sigma}(x_{0:N+1}) = \frac{1}{Z_g}\left(\prod_{n=1}^N g_n(x_n)\right) q^\sigma(x_{0:N+1}) \,, \tag{52}$$

with normalizing constant $Z_g = \int_{(\mathbb{R}^d)^{N+2}} \prod_{n=1}^N g_n(x_n) \mathrm{d}q^\sigma(x_{0:N+1})$. Thus, $q^{g,\sigma}$ is a Feynman–Kac model with prior q^σ; in line with the continuous-time setting, we refer to $q^{g,\sigma}$ as a twisted version of q^σ. Such twisted measures are ubiquitous in hidden Markov modeling and often simulated by Sequential Monte Carlo (SMC), see e.g., [19]. Intuitively, each potential g_n provides guidance on top of the prior transition kernel $q^\sigma_{n|n-1}$, with the uninformative limit case given by $g_n \equiv 1$. In Section E, we show that, as in the continuous-time setting, both $q^{g,\sigma}$ and its associated bridge $q^{g,\sigma}_{|0,N+1}$ are Markov. Moreover, their transition kernels are explicit but typically intractable, and can be written as Doob's h-transforms of the corresponding kernels of q^σ and $q^\sigma_{|0,N+1}$, respectively, induced by the potentials $\{g_n\}_{n=1}^N$.

Discrete-time Twisted Schrödinger Bridge problem. Based on (52), we define the discrete-time TSB problem as

$$\arg\min\{\mathrm{KL}(p\|q^{g,\sigma}) \,:\, p \in \mathscr{P}^{(N+2)}, p_0 = \mu,\ p_{N+1} = \nu\} \,. \tag{53}$$

By the KL chain rule, the solution of (53), denoted $p^\star$, is the finite-dimensional projection of the continuous-time solution $\mathbb{P}^N$ at times $\{t_n\}_{n=0}^{N+1}$, where $\mathbb{P}^N$ solves (1) with reference measure $\mathbb{Q}^{\mathrm{ref}} = \mathbb{Q}^{g,N}$, where

$$\mathrm{d}\mathbb{Q}^{g,N}(x_{[0:T]}) \propto \prod_{n=1}^N g_n(x_{t_n}) \mathrm{d}\mathbb{Q}^\sigma(x_{[0:T]}) \,. \tag{54}$$

This observation allows us to make explicit the correspondence between the continuous-time formulation (1) and the discrete-time formulation (53).

Proposition 16 (Informal version)**.** *Let $V \in \mathrm{C}^{0,1}([0,T]\times\mathbb{R}^d, \mathbb{R})$ be such that $Z_V < \infty$. Define the potentials $g_n = \exp(-\delta_n V_{t_n}/\sigma^2)$ for any $n \in \{1,\dots,N\}$. Then, in the asymptotic regime $\delta_n \to 0$ (equivalently $N \to \infty$), $\mathbb{P}^N$ converges to $\mathbb{P}^\star$.*

We now restate [18, Theorem 3.1], which was originally formulated for the canonical setting but remains valid here under the same assumptions. In particular, this result provides a discrete-time characterization of $p^\star$, consistent with the continuous-time formulation established in [27].

Proposition 17 (From [18])**.** *Let $p \in \mathscr{P}^{(N+2)}$ such that $p \in \mathcal{M}_N$, $p = p_{0,N+1} q^{g,\sigma}_{1:N|0,N+1}$, $p_0 = \mu$, and $p_{N+1} = \nu$. Then, p is a solution to (53).*

This result motivates the design of a discrete-time IMF scheme, which we detail now.

Discrete-time ideal IMF scheme. Let $p \in \mathscr{P}^{(N+2)}$. We define the reciprocal projection of p with respect to $q^{g,\sigma}$, denoted $\mathrm{proj}_{\mathcal{R}(q^{\sigma,g})}(p) \in \mathscr{P}^{(N+2)}$, as the solution of

$$\arg\min\{\mathrm{KL}(p\|r) \,:\, r \in \mathcal{R}(q^{g,\sigma})\} \,, \tag{55}$$

where $\mathcal{R}(q^{g,\sigma})$ denotes the class of distributions $r \in \mathscr{P}^{(N+2)}$ that are mixtures of $q^{g,\sigma}$-bridges, *i.e.*, $r = r_{0,N+1} q^{g,\sigma}_{1:N|0,N+1}$. In analogy with the continuous-time setting, we refer to this class as the set of reciprocal distributions. Following [18, Proposition 3.3], this projection admits the explicit form

$$\mathrm{proj}_{\mathcal{R}(q^{\sigma,g})}(p) = p_{0,N+1} q^{g,\sigma}_{1:N|0,N+1} \,. \tag{56}$$

We also define the Markovian projection of p, denoted $\mathrm{proj}_{\mathcal{M}_N}(p) \in \mathscr{P}^{(N+2)}$, as the solution of

$$\arg\min\{\mathrm{KL}(p\|m) \,:\, m \in \mathcal{M}_N\} \,, \tag{57}$$

where $\mathcal{M}_N$ denotes the set of Markov distributions $m \in \mathscr{P}^{(N+2)}$ admitting the same factorization as q^σ, *i.e.*,

$$m(x_0, x_1, \dots, x_{N+1}) = m_0(x_0) \prod_{n=1}^{N+1} m_{n|n-1}(x_n|x_{n-1}) \ .$$

In general, this minimization does not yield a closed-form solution. However, when $p \in \mathcal{R}(q^{g,\sigma})$, we can expand the KL objective explicitly, as stated in the following lemma.

Lemma 18. *Let $p \in \mathcal{R}(q^{g,\sigma})$ and $m \in \mathcal{M}_N$. Then, the following decomposition holds*

$$\begin{aligned}\mathrm{KL}(p\|m) = \mathrm{KL}(p_0\|m_0) + \mathbb{E}_{p_{0,N}}[\mathrm{KL}(p_{N+1|0}(\cdot|X_0)\|m_{N+1|N}(\cdot|X_N))] \\ + \sum_{n=1}^{N} \mathbb{E}_{p_{n-1,N+1}}[\mathrm{KL}(q^{g,\sigma}_{n|n-1,N+1}(\cdot|X_{n-1}, X_{N+1})\|m_{n|n-1}(\cdot|X_{n-1}))] \ .\end{aligned}$$

Using the equivalent backward factorization

$$m(x_0, x_1, \dots, x_{N+1}) = m_{N+1}(x_{N+1}) \prod_{n=1}^{N+1} m_{n-1|n}(x_{n-1}|x_n) \ ,$$

we derive a symmetric result.

Lemma 19. *Let $p \in \mathcal{R}(q^{g,\sigma})$ and $m \in \mathcal{M}_N$. Then, the following decomposition holds*

$$\begin{aligned}\mathrm{KL}(p\|m) = \mathrm{KL}(p_{N+1}\|m_{N+1}) + \mathbb{E}_{p_{1,N+1}}[\mathrm{KL}(p_{0|N+1}(\cdot|X_{N+1})\|m_{0|1}(\cdot|X_1))] \\ + \sum_{n=2}^{N+1} \mathbb{E}_{p_{0,n}}[\mathrm{KL}(q^{g,\sigma}_{n-1|0,n}(\cdot|X_0, X_n)\|m_{n-1|n}(\cdot|X_n))] \ .\end{aligned}$$

In particular, the KL objective in Lemma 19 is equivalent to the discrete-time DDPM variational objective derived in [20, Equation 5] in the case of uninformative potentials, highlighting a tight connection between bridge matching procedures and diffusion-model training objectives beyond the continuous-time setting.

We can now define the ideal discrete-time IMF procedure for solving (53) by alternating between the Markovian and Reciprocal projections defined above. Starting from an arbitrary $p^{-1} \in \mathscr{P}^{(N+2)}$ satisfying $p_0^{-1} = \mu$ and $p_{N+1}^{-1} = \nu$, we define recursively for $n \in \mathbb{N}$

$$p^{2n} = \mathrm{proj}_{\mathcal{R}(q^{\sigma,g})}(p^{2n-1}) \ , \quad p^{2n+1} = \mathrm{proj}_{\mathcal{M}_N}(p^{2n}) \ . \tag{58}$$

D.2 Computation of the Reciprocal projection

Let $p \in \mathcal{M}_N$, and assume that samples from $p_{0,N+1}$ are available. Observe that $\mathrm{proj}_{\mathcal{R}(q^{\sigma,g})}(p)$ coincides with the finite-dimensional projection of the reciprocal path measure $p_{0,N+1}\mathbb{Q}^{g,N}_{|0,T}$, where $\mathbb{Q}^{g,N}$ is defined in (54). Since the conditional bridge $\mathbb{Q}^{g,N}_{|0,T}$ is intractable, we replace the discrete reciprocal projection step with the following *continuous-time* variational problem:

$$\arg\min_{\psi\in\Psi} \mathbb{E}_{p_{0,N+1}}\left[\mathrm{KL}(\mathbb{Q}^{\psi}_{|0,T}\|\mathbb{Q}^{g,N}_{|0,T})\right] \ , \tag{59}$$

using the same variational bridge family $\{\mathbb{Q}^{\psi}_{|0,T}\}_{\psi\in\Psi}$ as in the continuous-time setting. The following proposition, which provides tractable objectives, follows directly from Proposition 9.

Proposition 20. *Assume that the variational family $\{\mathbb{Q}^{\psi}_{|0,T}\}_{\psi\in\Psi}$ in (15) describes the pinned diffusion processes associated, for almost surely any $(x_0, x_T) \sim \mathbb{P}_{0,T}$, to forward-time and backward-time parametric SDEs of the form*

$$\begin{aligned}\mathrm{d}\mathbf{X}^{0,T}_t &= v^{\mathrm{f},\psi}_t(\mathbf{X}^{0,T}_t|x_0, x_T)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \ , \ \mathbf{X}^{0,T}_0 = x_0, \ \mathbf{X}^{0,T}_T = x_T \ , \\ \mathrm{d}\mathbf{Y}^{0,T}_t &= v^{\mathrm{b},\psi}_{T-t}(\mathbf{Y}^{0,T}_t|x_0, x_T)\mathrm{d}t + \sigma\mathrm{d}\mathbf{B}_t \ , \ \mathbf{Y}^{0,T}_0 = x_T, \ \mathbf{Y}^{0,T}_T = x_0 \ ,\end{aligned}$$

with $v^{\mathrm{f},\psi}, v^{\mathrm{b},\psi} \in \mathbf{A}_T$. *Then, the minimizer of* (59) *for* $\mathbb{Q}^{g,N}$ *defined by* (54)*, denoted by* $\psi_\star$*, is equivalently the minimizer of the two following losses*

$$\begin{aligned}
\mathcal{L}^{\mathrm{f}}_{\mathcal{R}}(\psi) &= \int_0^T \mathbb{E}_{\Pi^{\psi}_{0,t,T}}\left[\tfrac{1}{2}\|(\mathbf{X}_T - \mathbf{X}_t)/(T-t) - v^{\mathrm{f},\psi}_t(\mathbf{X}_t|\mathbf{X}_0, \mathbf{X}_T)\|^2\right] \mathrm{d}t \\
&\quad - \sigma^2 \textstyle\sum_{n=1}^N \mathbb{E}_{\Pi^{\psi}_{t_n}}\left[\log g_n(\mathbf{X}_{t_n})\right] , \\
\mathcal{L}^{\mathrm{b}}_{\mathcal{R}}(\psi) &= \int_0^T \mathbb{E}_{\Pi^{\psi}_{0,t,T}}\left[\tfrac{1}{2}\|(\mathbf{X}_0 - \mathbf{X}_t)/t - v^{\mathrm{b},\psi}_t(\mathbf{X}_t|\mathbf{X}_0, \mathbf{X}_T)\|^2\right] \mathrm{d}t \\
&\quad - \sigma^2 \textstyle\sum_{n=1}^N \mathbb{E}_{\Pi^{\psi}_{t_n}}\left[\log g_n(\mathbf{X}_{t_n})\right] ,
\end{aligned}$$

where $\Pi^{\psi} = p_{0,N+1}\mathbb{Q}^{\psi}_{|0,T}$. *Moreover,* $\Pi^{\psi_\star}_{t_0,t_1,\ldots,t_{N+1}}$ *is a valid variational approximation of the discrete reciprocal measure* $\mathrm{proj}_{\mathcal{R}(q^{\sigma,g})}(p)$.

D.3 Computation of the Markovian projection

In this subsection, we focus on the practical implementation of the Markovian projection (57), restricted to reciprocal distributions $p \in \mathcal{R}(q^{g,\sigma})$, which is precisely the setting encountered in the IMF scheme (58).

As in the continuous-time setting, we restrict the Markovian projection to a prescribed parametric family indexed by $\theta \in \Theta$, and reformulate the KL objectives of Lemma 18 and Lemma 19 as tractable training losses. More specifically, we consider Markov distributions with Gaussian transition kernels, equipped with a denoiser-based parameterization consistent with standard diffusion-model practice and applicable for arbitrary N.

Before introducing this parameterization, we recall the general expression of the discrete-time Brownian bridge marginals: for any triplet $(i,j,k) \in \{0,\ldots,N+1\}$ such that $j < i < k$,

$$q^{\sigma}_{i|j,k}(x_i|x_j,x_k) = \mathrm{N}\left(\alpha^{j,k}_i x_j + \beta^{j,k}_i x_k, \sigma^2(\gamma^{j,k}_i)^2 \mathrm{I}_d\right) , \tag{60}$$

$$\text{where } \alpha^{j,k}_i = \frac{t_k - t_i}{t_k - t_j} , \ \beta^{j,k}_i = \frac{t_i - t_j}{t_k - t_j} , \ \gamma^{j,k}_i = \sqrt{\frac{(t_i - t_j)(t_k - t_i)}{t_k - t_j}} . \tag{61}$$

Denoiser-like Markovian parameterization. Inspecting the KL objectives in Lemma 18 and Lemma 19, a natural strategy would be to parameterize Markov kernels using the twisted bridge $q^{g,\sigma}_{|0,N+1}$, mirroring the DDPM parameterization [20] in the canonical case where $g_n \equiv 1$. However, since $q^{g,\sigma}_{|0,N+1}$ is intractable for general potentials, this approach is not directly applicable. We therefore retain the original DDPM parameterization based on the Brownian bridge $q^{\sigma}_{|0,N+1}$. For the forward Markovian projection, we consider

$$\begin{aligned}
m^{\theta^{\mathrm{f}}}_{n|n-1}(x_n|x_{n-1}) &= q^{\sigma}_{n|n-1,N+1}(x_n|x_{n-1}, x^{\theta^{\mathrm{f}}}_{N+1}(t_{n-1},x_{n-1})) \quad \text{if } n \in \{1,\ldots,N\}, \\
m^{\theta^{\mathrm{f}}}_{N+1|N}(x_{N+1}|x_N) &= \mathrm{N}(x_{N+1}; x^{\theta^{\mathrm{f}}}_{N+1}(t_N,x_N), \sigma^2_{N+1}\mathrm{I}_d) ,
\end{aligned} \tag{62}$$

where $x^{\theta^{\mathrm{f}}}_{N+1}$ is intended to predict x_{N+1}, and $\sigma_{N+1} > 0$ is left as a hyperparameter.

Symmetrically, for the backward projection, we set

$$\begin{aligned}
m^{\theta^{\mathrm{b}}}_{n-1|n}(x_{n-1}|x_n) &= q^{\sigma}_{n-1|0,n}(x_{n-1}|x^{\theta^{\mathrm{b}}}_0(t_n,x_n), x_n) \qquad \text{if } n \in \{2,\ldots,N+1\} , \\
m^{\theta^{\mathrm{b}}}_{0|1}(x_0|x_1) &= \mathrm{N}(x_0; x^{\theta^{\mathrm{b}}}_0(t_1,x_1), \sigma^2_0 \mathrm{I}_d) ,
\end{aligned} \tag{63}$$

where $x^{\theta^{\mathrm{b}}}_0$ is intended to predict x_0, and $\sigma_0 > 0$ is left as a hyperparameter. This variational family yields computable Markovian projection objectives for any $N \geq 1$.

Proposition 21. *Let $p \in \mathcal{R}(q^{g,\sigma})$. Define $\theta^{\mathrm{f}}_\star$ and $\theta^{\mathrm{b}}_\star$ as the minimizers over Θ of the following forward and backward objectives*

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_N}(\theta^{\mathrm{f}}) = \sum_{n=1}^N \tfrac{1}{N-n+1} \sum_{k=n}^N \mathbb{E}_{p_{n-1,k,N+1}} \Big[\| A^{\mathrm{f}}_n \sigma^2 (T-t_k) \nabla \log g_k(X_k) + B^{\mathrm{f}}_n \left(X_{N+1} - x^{\theta^{\mathrm{f}}}_{N+1}(t_{n-1}, X_{n-1}) \right) \|^2 \Big] + \tfrac{\sigma^2}{\sigma^2_{N+1}} \mathbb{E}_{p_{N,N+1}} \left[\left\| X_{N+1} - x^{\theta^{\mathrm{f}}}_{N+1}(t_N, X_N) \right\|^2 \right] , \tag{64}$$

$$\mathcal{L}^{\mathrm{b}}_{\mathcal{M}_N}(\theta^{\mathrm{b}}) = \sum_{n=2}^{N+1} \tfrac{1}{n-1} \sum_{k=1}^{n-1} \mathbb{E}_{p_{0,k,n}} \left[\left\| A^{\mathrm{b}}_n \sigma^2 t_k \nabla \log g_k(X_k) + B^{\mathrm{b}}_n \left(X_0 - x^{\theta^{\mathrm{b}}}_0(t_n, X_n) \right) \right\|^2 \right] + \tfrac{\sigma^2}{\sigma^2_0} \mathbb{E}_{p_{0,1}} \left[\left\| X_0 - x^{\theta^{\mathrm{b}}}_0(t_1, X_1) \right\|^2 \right] , \tag{65}$$

where $A^{\mathrm{f}}_n = \gamma^{n-1,N+1}_n \times \frac{N-n+1}{T-t_n}$, $B^{\mathrm{f}}_n = \frac{\beta^{n-1,N+1}_n}{\gamma^{n-1,N+1}_n}$, $A^{\mathrm{b}}_n = \gamma^{0,n}_{n-1} \times \frac{n-1}{t_{n-1}}$ and $B^{\mathrm{b}}_n = \frac{\alpha^{0,n}_{n-1}}{\gamma^{0,n}_{n-1}}$.

Then, the Markovian projection of p restricted to the denoiser-based family, denoted $m^\star$, is equivalently defined by (a) $m^\star_0 = \mu$ and forward kernels (62) parameterized by $x^{\theta^{\mathrm{f}}_\star}_{N+1}$ or (b) $m^\star_{N+1} = \nu$ and backward kernels (63) parameterized by $x^{\theta^{\mathrm{b}}_\star}_0$.

In practice, $x^{\theta^{\mathrm{f}}}_{N+1}$ and $x^{\theta^{\mathrm{b}}}_0$ are implemented as neural networks. In the case $g_n \equiv 1$, we recover exactly the classical DDPM objective for the backward loss (65), as expected.

Link with continuous-time setting. Consider the denoiser-paramaterization inherited from the canonical diffusion setting, defined by

$$v^{\theta^{\mathrm{f}}}_t(x) = \frac{x^{\theta^{\mathrm{f}}}_T(t,x) - x}{T-t} , \quad v^{\theta^{\mathrm{b}}}_t(x) = \frac{x^{\theta^{\mathrm{b}}}_0(t,x) - x}{t} ,$$

where $x^{\theta^{\mathrm{f}}}_T$ and $x^{\theta^{\mathrm{b}}}_0$ are predictors of the terminal state x_T and the initial state x_0, respectively, given (t,x). Under this parameterization, the TSBM Markovian objectives (18) and (39) admit the simplified forms

$$\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_T}(\theta^{\mathrm{f}}) = \int_0^T (T-t)^{-1} \int_t^T \mathbb{E}_{\Pi_{t,s,T}} \left[\left\| -(T-s)\nabla V_s(\mathbf{X}_s) + \tfrac{1}{T-t} \left(\mathbf{X}_T - x^{\theta^{\mathrm{f}}}_T(t, \mathbf{X}_t) \right) \right\|^2 \right] \mathrm{d}s\mathrm{d}t , \tag{66}$$

$$\mathcal{L}^{\mathrm{b}}_{\mathcal{M}_T}(\theta^{\mathrm{b}}) = \int_0^T t^{-1} \int_0^t \mathbb{E}_{\Pi_{0,s,t}} \left[\left\| -s\nabla V_s(\mathbf{X}_s) + \tfrac{1}{t} \left(\mathbf{X}_0 - x^{\theta^{\mathrm{b}}}_0(t, \mathbf{X}_t) \right) \right\|^2 \right] \mathrm{d}s\mathrm{d}t . \tag{67}$$

When setting $g_n = \exp(-\delta_n V_{t_n}/\sigma^2)$ to match the continuous-time formulation (see Proposition 16) and taking $\delta_n = o(1)$, the following approximations hold

$$A^{\mathrm{f}}_n \sigma^2 (T-t_k) \nabla \log g_k(X_k) \approx -\sqrt{\delta_n}(T-t_k) \nabla V_{t_k}(X_k) , \quad B^{\mathrm{f}}_n \approx \frac{\sqrt{\delta_n}}{T-t_n} ,$$

$$A^{\mathrm{b}}_n \sigma^2 t_k \nabla \log g_k(X_k) \approx -\sqrt{\delta_n} t_k \nabla V_{t_k}(X_k) , \quad B^{\mathrm{b}}_n \approx \frac{\sqrt{\delta_n}}{t_n} .$$

Hence, it is clear that the denoiser-like TSBM objectives (66) and (67) can respectively be seen as the time-continuous limit of objectives (64) and (65). Note also, that in this case, the regression targets become independent of σ (with suitable choices of σ_0 and σ_{N+1}), thereby connecting to the deterministic regime.

E Proofs for D-TSBM

E.1 Results on twisted measures

In this section, we present general results concerning the Markovian decomposition of the twisted bridge $q^{g,\sigma}_{|0,N+1}$ relative to the discrete-time Brownian bridge $q^{\sigma}_{|0,N+1}$. In particular, we derive closed-form expressions for the corresponding quantities. Although explicit, these expressions remain computationally intractable in practice.

Lemma 22. *The bridge $q^{g,\sigma}_{1:N|0,N+1} \in \mathcal{M}_{N-2}$ admits the following forward Markov decomposition*

$$q^{g,\sigma}_{1:N|0,N+1}(x_{1:N}|x_0, x_{N+1}) = \prod_{n=1}^{N} q^{g,\sigma}_{n|n-1,N+1}(x_n|x_{n-1}, x_{N+1})$$

where for any $n \in \{1, \dots, N\}$, we have

$$q^{g,\sigma}_{n|n-1,N+1}(x_n|x_{n-1}, x_{N+1}) = \frac{h_n(x_n|x_{N+1})g_n(x_n)}{h_{n-1}(x_{n-1}|x_{N+1})} q^{\sigma}_{n|n-1,N+1}(x_n|x_{n-1}, x_{N+1})$$

with

$$h_n(x_n|x_{N+1}) = \int_{(\mathbb{R}^d)^{N-n}} \prod_{k=n+1}^{N} g_k(x_k) \mathrm{d}q^{\sigma}_{k|k-1,N+1}(x_k|x_{k-1}, x_{N+1}) \, .$$

Lemma 23. *For any $n \in \{1, \dots, N\}$, the score of the forward bridge kernel $q^{g,\sigma}_{n|n-1,N+1}$ is given by*

$$\begin{aligned}\nabla \log q^{g,\sigma}_{n|n-1,N+1}(x_n|x_{n-1}, x_{N+1}) &= \nabla \log q^{\sigma}_{n|n-1,N+1}(x_n|x_{n-1}, x_{N+1}) \\ &+ \sum_{k=n}^{N} \mathbb{E}_{q^{g,\sigma}_{k|n,N+1}} \left[\frac{T - t_k}{T - t_n} \nabla \log g_k(X_k) | X_n = x_n, X_{N+1} = x_{N+1} \right]\end{aligned}$$

Proof. Since we have for any $n \in \{1, \dots, N\}$,

$$\begin{aligned}\nabla \log q^{g,\sigma}_{n|n-1,N+1}(x_n|x_{n-1}, x_{N+1}) &= \nabla \log q^{\sigma}_{n|n-1,N+1}(x_n|x_{n-1}, x_{N+1}) \\ &+ \nabla \log h_n(x_n|x_{N+1}) + \nabla \log g_n(x_n) \, ,\end{aligned}$$

proving this result amounts to prove the identity

$$\nabla \log h_n(x_n|x_{N+1}) = \sum_{k=n+1}^{N} \mathbb{E}_{q^{g,\sigma}_{k|n,N+1}} \left[\frac{T - t_k}{T - t_n} \nabla \log g_k(X_k) | X_n = x_n, X_{N+1} = x_{N+1} \right] \, . \tag{68}$$

We prove this equivalent result via backward recursion on n. First note that this result stands trivially for $n = N$, since $h(x_N|x_{N+1}) = 1$, and therefore, $\nabla \log h_N(x_N|x_{N+1}) = 0$, validating the result. Assume (68) is verified at rank $(n+1)$ for some $n \in \{1, \dots, N-1\}$. At rank n, we obtain by Lemma 22 the following recursion relation

$$h_n(x_n|x_{N+1}) = \int h_{n+1}(x_{n+1}|x_{N+1}) g_{n+1}(x_{n+1}) \mathrm{d}q^{\sigma}_{n+1|n,N+1}(x_{n+1}|x_n, x_{N+1}) \, .$$

Hence, using the reparameterization trick in the discrete Brownian bridge $q^{\sigma}_{n+1|n,N+1}$ along with Tweedie formula, we obtain

$$\begin{aligned}\nabla \log h_n(x_n|x_{N+1}) &= \frac{T - t_{n+1}}{T - t_n} \mathbb{E}_{q^{g,\sigma}_{n+1|n,N}} \Bigg[\nabla \log h_{n+1}(X_{n+1}|x_{N+1}) \\ &\qquad + \nabla \log g_{n+1}(X_{n+1}) \Big| X_n = x_n, X_{N+1} = x_{N+1} \Bigg] \, , \\ &= \frac{T - t_{n+1}}{T - t_n} \sum_{k=n+2}^{N} \mathbb{E}_{q^{g,\sigma}_{k|n,N+1}} \left[\frac{T - t_k}{T - t_{n+1}} \nabla \log g_k(X_k) | X_n = x_n, X_{N+1} = x_{N+1} \right] \\ &\quad + \mathbb{E}_{q^{g,\sigma}_{n+1|n,N+1}} \left[\frac{T - t_{n+1}}{T - t_n} \nabla \log g_{n+1}(X_{n+1}) | X_n = x_n, X_{N+1} = x_{N+1} \right] \, ,\end{aligned}$$

where the simplification comes from the recursion assumption from rank $(n+1)$. Hence, we obtain the result at rank n, which proves the result for any $n \in \{1, \dots, N\}$. □

Lemma 24. *Symmetrically, $q^{g,\sigma}_{1:N|0,N+1} \in \mathcal{M}_{N-2}$ admits the following backward Markov decomposition*

$$q^{g,\sigma}_{1:n|0,N+1}(x_{1:N}|x_0, x_{N+1}) = \prod_{n=2}^{N+1} q^{g,\sigma}_{n-1|0,n}(x_{n-1}|x_0, x_n)$$

where for any $n \in \{2, \dots, N+1\}$

$$q^{g,\sigma}_{n-1|0,n}(x_{n-1}|x_0, x_n) = \frac{h_{n-1}(x_{n-1}|x_0) g_{n-1}(x_{n-1})}{h_n(x_n|x_0)} q^{\sigma}_{n-1|0,n}(x_{n-1}|x_0, x_n) \ ,$$

where

$$h_n(x_n|x_0) = \int_{(\mathbb{R}^d)^{n-1}} \prod_{k=2}^{n} g_{k-1}(x_{k-1}) \mathrm{d}q^{\sigma}_{k-1|0,k}(x_{k-1}|x_0, x_k) \ .$$

Lemma 25. *For any $n \in \{2, \dots, N+1\}$, the score of the forward bridge kernel $q^{g,\sigma}_{n-1|0,n}$ is given by*

$$\nabla \log q^{g,\sigma}_{n-1|0,n}(x_{n-1}|x_0, x_n) = \nabla \log q^{\sigma}_{n-1|0,n}(x_{n-1}|x_0, x_n) \tag{69}$$

$$+ \sum_{k=1}^{n-1} \mathbb{E}_{q^{g,\sigma}_{k|0,n-1}} \left[\frac{t_k}{t_{n-1}} \nabla \log g_k(X_k) | X_0 = x_0, X_{n-1} = x_{n-1} \right] \tag{70}$$

Proof. The proof structure follows the same recursion principle as for Lemma 23, but with increasing n, and exploiting the formula from Lemma 24. □

E.2 Results on the discrete Twisted SB problem

Proposition 26 (Formal version of Proposition 16)**.** *Let $V \in \mathrm{C}^{0,1}([0,T] \times \mathbb{R}^d, \mathbb{R})$ be such that $Z_V < \infty$. Define the potentials $g_n = \exp(-\delta_n V_{t_n}/\sigma^2)$ for any $n \in \{1, \dots, N\}$. Denote by $\mathbb{P}^N$ the solution to the continuous-time SB problem (1) with reference measure given by*

$$\mathrm{d}\mathbb{Q}^{g,N}(x_{[0:T]}) = \frac{1}{Z_N} \prod_{n=1}^{N} g_n(x_{t_n}) \mathrm{d}\mathbb{Q}^{\sigma}(x_{[0:T]}) \ , \quad \textit{where} \quad Z_N = \int \prod_{n=1}^{N} g_n(x_{t_n}) \mathrm{d}\mathbb{Q}^{\sigma}(x_{[0:T]}) \ .$$

Assume that $\mathbb{P}^N$ converges to a certain $\mathbb{P}^\infty \in \mathscr{P}(\mathbf{C}_T)$ as $N \to \infty$ in the weak topology. Then, $\mathbb{P}^\infty$ is solution to (1) with reference measure $\mathbb{Q}^{V,\sigma}$.

Proof. Denote by $\mathbb{P}^\star$ the solution to (1) with reference measure $\mathbb{Q}^{V,\sigma}$. First note that for any $\mathbb{P} \in \mathscr{P}(\mathbf{C}_T)$,

$$\mathbb{E}_{\mathbb{P}}[\log \frac{\mathrm{d}\mathbb{Q}^{V,\sigma}}{\mathrm{d}\mathbb{Q}^{g,N}}] = \mathbb{E}_{\mathbb{P}} \left[\exp \left(\sum_{n=0}^{N} \int_{t_n}^{t_{n+1}} \frac{V_{t_{n+1}}(\mathbf{X}_{t_{n+1}}) - V_t(\mathbf{X}_t)}{\sigma^2} \mathrm{d}t \right) \right] + \log \frac{Z_N}{Z_V} \to 0 \tag{71}$$

using the dominated convergence theorem. Then, we have

$$\begin{aligned}
\mathrm{KL}(\mathbb{P}^\infty \| \mathbb{Q}^{V,\sigma}) &\leq \liminf \mathrm{KL}(\mathbb{P}^N \| \mathbb{Q}^{g,N}) && \text{(KL property)} \\
&\leq \liminf \mathrm{KL}(\mathbb{P}^\star \| \mathbb{Q}^{g,N}) && \text{(definition of } \mathbb{Q}^{g,N}\text{)} \\
&= \mathrm{KL}(\mathbb{P}^\star \| \mathbb{Q}^{V,\sigma}) + \liminf \mathbb{E}_{\mathbb{P}^\star}[\log \frac{\mathrm{d}\mathbb{Q}^{V,\sigma}}{\mathbb{Q}^{g,N}}] \\
&= \mathrm{KL}(\mathbb{P}^\star \| \mathbb{Q}^{V,\sigma}) \ .
\end{aligned}$$

Then, we conclude by uniqueness of $\mathbb{P}^\star$ that $\mathbb{P}^\infty = \mathbb{P}^\star$. □

E.3 Results on the Markovian projection

Proofs of Lemma 18 and Lemma 19. These results follow from the observation that any $p \in \mathcal{R}(q^{g,\sigma})$ writes

$$\begin{aligned} p_{0:N+1} &= p_{0,N+1} q^{g,\sigma}_{1:N|0,N+1} \\ &= p_0 p_{N+1|0} \prod_{n=1}^{N} q^{g,\sigma}_{n|n-1,N+1} && \text{(forward decomposition)} \\ &= p_{N+1} p_{0|N+1} \prod_{n=2}^{N+1} q^{g,\sigma}_{n-1|0,n} && \text{(backward decomposition)} , \end{aligned}$$

using the Markov structure of the twisted bridge, see Lemmas 22 and 24. Then, we obtain the KL decompositions by applying the KL chain rule. □

Proof of Proposition 21. In this proof, we solely consider the forward setting, since the backward result can be proved with very similar tools. Let $\theta^{\mathrm{f}} \in \Theta$, and denote by $m^{\theta^{\mathrm{f}}} \in \mathcal{M}_N$ the Markovian measure obtained via forward denoiser-based parameterization (62) with θ^{f}. In this case, we have for any $n \in \{1, \dots, N\}$, almost surely any $(x_{n-1}, x_{N+1}) \sim p_{n-1,N+1}$, up to additive constants independent from θ^{f},

$$\begin{aligned} &\mathrm{KL}(q^{g,\sigma}_{n|n-1,N+1}(\cdot|x_{n-1}, x_{N+1}) \| m^{\theta^{\mathrm{f}}}_{n|n-1}(\cdot|x_{n-1})) \\ &\quad \overset{c}{=} \frac{1}{2\sigma^2 (\gamma^f_n)^2} \mathbb{E}\left[\left\| X_n - \alpha^{\mathrm{f}}_n x_{n-1} - \beta^{\mathrm{f}}_n \tilde{x}_{N+1} \right\|^2 \right] \\ &\quad \overset{c}{=} \frac{1}{2} \mathbb{E}\left[\left\| -\sigma \gamma^{\mathrm{f}}_n \nabla \log q^{\sigma}_{n|n-1,N+1}(X_n|x_{n-1}, x_{N+1}) + \frac{\beta^{\mathrm{f}}_n}{\sigma \gamma^{\mathrm{f}}_n} (x_{N+1} - \tilde{x}_{N+1}) \right\|^2 \right] \end{aligned}$$

where $\tilde{x}_{N+1} = x^{\theta^{\mathrm{f}}}_{N+1}(t_{n-1}, x_{n-1})$, and the expectation is taken with respect to $q^{g,\sigma}_{n|n-1,N+1}(\cdot|x_{n-1}, x_{N+1})$. Using the result from Lemma 23, combined with the fact that

$$\mathbb{E}_{q^{g,\sigma}_{n|n-1,N+1}}[\nabla \log q^{g,\sigma}_{n|n-1,N+1}(X_n|x_{n-1}, x_{N+1})] = 0 ,$$

we obtain, up to additive constants independent from θ^f,

$$\begin{aligned} &\mathrm{KL}(q^{g,\sigma}_{n|n-1,N+1}(\cdot|x_{n-1}, x_{N+1}) \| m^{\theta^{\mathrm{f}}}_{n|n-1}(\cdot|x_{n-1})) \\ &\quad \overset{c}{=} \tfrac{1}{2} \mathbb{E}\left[\left\| \sigma \gamma^{\mathrm{f}}_n \textstyle\sum_{k=n}^{N} \mathbb{E}\left[\frac{T - t_k}{T - t_n} \nabla \log g_k(X_k) | X_n = x_n, X_{N+1} = x_{N+1} \right] + \frac{\beta^{\mathrm{f}}_n}{\sigma \gamma^{\mathrm{f}}_n} (x_{N+1} - \tilde{x}_{N+1}) \right\|^2 \right] \\ &\quad \overset{c}{=} \tfrac{1}{2(N-n+1)\sigma^2} \textstyle\sum_{k=n}^{N} \mathbb{E}\left[\left\| \sigma^2 \gamma^{\mathrm{f}}_n (N - n + 1) \frac{T - t_k}{T - t_n} \nabla \log g_k(X_k) + \frac{\beta^{\mathrm{f}}_n}{\gamma^{\mathrm{f}}_n} (x_{N+1} - \tilde{x}_{N+1}) \right\|^2 \right] . \end{aligned}$$

By deriving this identity for all $n \in \{1, \dots, N\}$, and combining the results, we conclude that

$$\mathrm{KL}(p \| m^{\theta^{\mathrm{f}}}) = \frac{1}{2\sigma^2} \mathcal{L}^{\mathrm{f}}_{\mathcal{M}_N}(\theta^{\mathrm{f}}) + \mathrm{C}^{\mathrm{f}} ,$$

where C^{f} is a constant independent from θ^{f}. This equivalence shows that minimizing $\mathcal{L}^{\mathrm{f}}_{\mathcal{M}_N}(\theta^{\mathrm{f}})$ ensures that $\theta^{\mathrm{f}}_\star$ yields the forward Markovian projection of p within the Markov measures with denoiser-like parameterization. □

Algorithm 1 Twisted Schrödinger Bridge Matching (TSBM) - Continuous time state cost

Input: source dataset μ, target dataset ν, regularization $\sigma > 0$, number of outer iterations $N \in \mathbb{N}$
Start from the independent coupling $\hat{\mathbb{P}}_{0,T} = \mu \otimes \nu$
for $n = 0$ to $N - 1$ **do**
Set dir $=$ f if $n \bmod 2 = 0$ else dir $=$ b
Learn $\mathbb{Q}^{\psi_\star}_{|0,T}$ by running `ReciprocalProj`(dir), see Algorithm 2 ▷ Reciprocal projection
Learn $v^{\theta^{\mathrm{dir}}_\star}_\star$ by running `MarkovProj`(dir, $\psi_\star$), see Algorithm 3 ▷ Markovian projection
if dir $=$ f
then Solve $\mathrm{d}\hat{\mathbf{X}}^{\mathrm{f}}_t = v^{\theta^{\mathrm{f}}_\star}_t(\hat{\mathbf{X}}^{\mathrm{f}}_t)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t, \hat{\mathbf{X}}^{\mathrm{f}}_0 \sim \hat{\mathbb{P}}_0$ and set $(\hat{\mathbf{X}}_0, \hat{\mathbf{X}}_T) = (\hat{\mathbf{X}}^{\mathrm{f}}_0, \hat{\mathbf{X}}^{\mathrm{f}}_T)$
else Solve $\mathrm{d}\hat{\mathbf{X}}^{\mathrm{b}}_t = v^{\theta^{\mathrm{b}}_\star}_{T-t}(\hat{\mathbf{X}}^{\mathrm{b}}_t)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t, \hat{\mathbf{X}}^{\mathrm{b}}_0 \sim \hat{\mathbb{P}}_T$ and set $(\hat{\mathbf{X}}_0, \hat{\mathbf{X}}_T) = (\hat{\mathbf{X}}^{\mathrm{b}}_T, \hat{\mathbf{X}}^{\mathrm{b}}_0)$
Update the coupling as $\hat{\mathbb{P}}_{0,T} = \mathrm{Law}(\hat{\mathbf{X}}_0, \hat{\mathbf{X}}_T)$
end for
Output: Approximate solutions $(\hat{\mathbf{X}}^{\mathrm{f}}_t)_{t\in[0,T]}$ and $(\hat{\mathbf{X}}^{\mathrm{b}}_t)_{t\in[0,T]}$ to (5), respectively in the forward and backward time directions

Algorithm 2 `ReciprocalProj`(dir)

Input: dir $\in \{\mathrm{f}, \mathrm{b}\}$, coupling $\hat{\mathbb{P}}_{0,T}$
repeat
Sample $(\mathbf{X}_0, \mathbf{X}_T) \sim \hat{\mathbb{P}}_{0,T}$, $t \sim \mathrm{U}(0, T)$
Sample $\mathbf{X}_t \sim \mathbb{Q}^{\psi}_{t|0,T}(\cdot|\mathbf{X}_0, \mathbf{X}_T)$
Compute $v^{\mathrm{dir},\psi}$, see Section F.2
Minimize $\mathcal{L}^{\mathrm{dir}}_{\mathcal{R}}(\psi)$, see (16)–(34)
until it converges
Output: learned bridge $\mathbb{Q}^{\psi_\star}_{|0,T} \approx \mathbb{Q}^{V,\sigma}_{|0,T}$

Algorithm 3 `MarkovProj`(dir, $\psi_\star$)

Input: dir $\in \{\mathrm{f}, \mathrm{b}\}$, coupling $\hat{\mathbb{P}}_{0,T}$, bridge $\mathbb{Q}^{\psi_\star}_{|0,T}$
repeat
Sample $(\mathbf{X}_0, \mathbf{X}_T) \sim \hat{\mathbb{P}}_{0,T}$, $t \sim \mathrm{U}(0, T)$
Sample $s \sim \mathrm{U}(t, T)$ if dir $=$ f else $s \sim \mathrm{U}(0, t)$
Sample $(\mathbf{X}_t, \mathbf{X}_s) \sim \mathbb{Q}^{\psi_\star}_{t,s|0,T}(\cdot|\mathbf{X}_0, \mathbf{X}_T)$
Minimize $\mathcal{L}^{\mathrm{dir}}_{\mathcal{M}_T}(\theta^{\mathrm{dir}})$, see (18)–(39)
until it converges
Output: learned $v^{\theta^{\mathrm{dir}}_\star}_\star$ for $\mathrm{proj}_{\mathcal{M}_T}(\hat{\mathbb{P}}_{0,T}\mathbb{Q}^{\psi_\star}_{|0,T})$

F Practical implementation of TSBM

F.1 Pseudo-code for TSBM

We provide the pseudo-code of the TSBM methodology in Algorithm 1 (for the continuous time version), further detailing the reciprocal projection step in Algorithm 2 and the Markovian projection step in Algorithm 3.

F.2 Variational approximation of the reciprocal projection

In this section, we detail the construction of the variational family used in Section 3.1, relying on the stochastic interpolants introduced by [1]. We assume access to conditioning pairs $(x_0, x_T) \sim \mathbb{P}_{0,T}$ and emphasize that the families described below can be straightforwardly parallelized across multiple conditioning pairs in large data batches.

Gaussian variational family. We say that the stochastic process $\mathbb{P}_{|0,T}(\cdot|x_0, x_T)$ is a Gaussian stochastic interpolant between x_0 and x_T [1] if its marginals satisfy

$$\begin{aligned} \mathbb{P}_{t|0,T}(\cdot|x_0, x_T) &= \mathrm{N}(I_t, \gamma^2_t \mathrm{I}_d) \ , \\ \text{where } I \in \mathcal{I}(x_0, x_T) &= \{I \in \mathrm{C}^1([0,T], \mathbb{R}^d) \ : \ I_0 = x_0 \ , I_T = x_T\} \ , \\ \gamma \in \Gamma &= \{\gamma \in \mathrm{C}^1([0,T], (0,\infty)) \ : \ \gamma_0 = \gamma_T = 0\} \ . \end{aligned}$$

We denote by $\mathrm{SI}(x_0, x_T)$ the resulting class of stochastic interpolants.

For any $I \in \mathcal{I}(x_0, x_T)$ and $\gamma \in \Gamma$, we define the time-dependent velocity fields $(t, x) \mapsto v_t^{\mathrm{f}}(x|x_0, x_T)$ (forward direction) and $(t, x) \mapsto v_t^{\mathrm{b}}(x|x_0, x_T)$ (backward direction) by

$$v_t^{\mathrm{f}}(x|x_0, x_T) = \partial_t I_t + \left(\frac{\partial_t \gamma_t}{\gamma_t} - \frac{\sigma^2}{2\gamma_t^2}\right)(x - I_t) ,$$
$$v_t^{\mathrm{b}}(x|x_0, x_T) = -\partial_t I_t + \left(-\frac{\partial_t \gamma_t}{\gamma_t} - \frac{\sigma^2}{2\gamma_t^2}\right)(x - I_t) .$$

As shown in [1], the SDEs

$$\mathrm{d}\mathbf{X}_t^{0,T} = v_t^{\mathrm{f}}(\mathbf{X}_t^{0,T}|x_0, x_T)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t \, , \mathbf{X}_0^{0,T} = x_0 \, , \tag{72}$$
$$\mathrm{d}\mathbf{Y}_t^{0,T} = v_{T-t}^{\mathrm{b}}(\mathbf{Y}_t^{0,T}|x_0, x_T)\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t \, , \mathbf{Y}_0^{0,T} = x_T \, , \tag{73}$$

induce identical path measures up to time reversal, and share the same marginals as $\mathbb{P}_{|0,T}(\cdot|x_0, x_T)$. In particular, the boundary conditions are satisfied almost surely, namely $\mathbf{X}_T^{0,T} = x_T$ and $\mathbf{Y}_T^{0,T} = x_0$.

Motivated by this characterization, we define our variational family as a parametric subset of $\mathrm{SI}(x_0, x_T)$ of the form

$$\mathbb{Q}_{|0,T}^{\psi}(\cdot|x_0, x_T) = \mathrm{N}(I_t^{\psi}, (\gamma_t^{\psi})^2 \mathrm{I}_d) \, ,$$

where the mean path $I^{\psi} \in \mathcal{I}(x_0, x_T)$ and the standard deviation $\gamma^{\psi} \in \Gamma$ are parameterized by $\psi \in \Psi$. In this case, $\mathbb{Q}_{|0,T}^{\psi}(\cdot|x_0, x_T)$ is associated with forward SDE (72) and backward SDE (73) where

$$v_t^{\mathrm{f},\psi}(x|x_0, x_T) = \partial_t I_t^{\psi} + \left(\frac{\partial_t \gamma_t^{\psi}}{\gamma_t^{\psi}} - \frac{\sigma^2}{2(\gamma_t^{\psi})^2}\right)(x - I_t^{\psi}) ,$$
$$v_t^{\mathrm{b},\psi}(x|x_0, x_T) = -\partial_t I_t^{\psi} + \left(-\frac{\partial_t \gamma_t^{\psi}}{\gamma_t^{\psi}} - \frac{\sigma^2}{2(\gamma_t^{\psi})^2}\right)(x - I_t^{\psi}) .$$

Expression of the transition kernels. The SDEs (72) and (73) can be integrated exactly as follows. The forward and backward transition kernels are Gaussian distributions, respectively given by

$$\mathbb{Q}_{s|t,T}^{\psi}(\cdot|x_t, x_T) = \mathrm{N}\left(e^{-\frac{\sigma^2}{2}J^{\psi}(t,s)}\frac{\gamma_s^{\psi}}{\gamma_t^{\psi}}(x_t - I_t^{\psi}) + I_s^{\psi}, (\gamma_s^{\psi})^2\left[1 - e^{-\frac{\sigma^2}{2}J^{\psi}(t,s)}\right]\right) \, , \; s \in [t, T] \, , \tag{74}$$

$$\mathbb{Q}_{s|0,t}^{\psi}(\cdot|x_0, x_t) = \mathrm{N}\left(e^{-\frac{\sigma^2}{2}J^{\psi}(s,t)}\frac{\gamma_s^{\psi}}{\gamma_t^{\psi}}(x_t - I_t^{\psi}) + I_s^{\psi}, (\gamma_s^{\psi})^2\left[1 - e^{-\frac{\sigma^2}{2}J^{\psi}(s,t)}\right]\right) \, , \; s \in [0, t] \, , \tag{75}$$

where $J^{\psi}(r_1, r_2) = \int_{r_1}^{r_2} \frac{1}{(\gamma_u^{\psi})^2}\mathrm{d}u$. For arbitrary γ^{ψ}, J^{ψ} is however not tractable.

Combination with spline parameterization. Following [30], we parameterize the mean path I^{ψ} as a linear d-dimensional spline and the standard deviation γ^{ψ} as a linear one-dimensional spline defined through control points. More precisely, we fix $C \geq 1$ intermediate control times $\{t_i\}_{i=1}^C$ in $(0, T)$ such that

$$0 < t_1 < t_2 < \ldots < t_C < T \, .$$

For a given conditioning pair (x_0, x_T), we define

$$I_t^{\psi} = \text{Linear Spline}(t; x_0, \{X_{t_i}\}_{i=1}^C, x_T), \qquad \gamma_t^{\psi} = \text{Linear Spline}(t; 0, \{\gamma_{t_i}\}_{i=1}^C, 0) \, ,$$

where the control points $\{X_{t_i}\}_{i=1}^C \subset \mathbb{R}^d$ and $\{\gamma_{t_i}\}_{i=1}^C \subset (0, \infty)$ are the parameters to optimize, which is much more lightweight than a neural network parameterization. In practice, the number of control points C can be chosen significantly smaller than the number of discretization steps used to simulate the SDEs. This yields a favorable trade-off between expressive power of the interpolant and computational efficiency. The second advantage of this method is that we can leverage samples generated at intermediary timesteps between endpoints when updating the coupling $\mathbb{P}_{0,T}$ via SDE simulation, allowing for well-informed initialization of the spline at every reciprocal projection step (except the first where we initialize with the linear interpolant). Similarly, we initialize the standard deviation with the one from the Brownian bridge $\mathbb{Q}_{|0,T}^{\sigma}$ at the very first step, and save the learned control standard deviations for initialization of future reciprocal steps.

Optimizing the sampling computational cost. In practice, we first perform the variational optimization over the full training dataset of endpoint pairs $(x_0, x_T) \sim \mathbb{P}_{0,T}$, before running the subsequent Markovian projection step. For each pair, we store the learned control points, namely the mean and standard deviation, both of which depend on (x_0, x_T), and treat them as additional data features. Crucially, we also precompute and store the values $\{1/(\gamma^{\psi}_{u_\ell})^2\}_{\ell=1}^{L}$ on a fine grid $\{u_\ell\}_{\ell=1}^{L}$ of $[0, T]$. These operations are performed on large batches and benefit from efficient parallelization.

Then, during Markovian optimization, for any training pair (x_0, x_T), we can immediately sample $x_t \sim \mathbb{Q}^{\psi}_{t|0,T}(\cdot \mid x_0, x_T)$. Moreover, using the precomputed grid values to approximate J^{ψ} by a Riemann sum, we can efficiently sample either $x_s \sim \mathbb{Q}^{\psi}_{s|t,T}(\cdot \mid x_t, x_T)$ in the forward direction, for $s \in [t, T]$, via (74), or $x_s \sim \mathbb{Q}^{\psi}_{s|0,t}(\cdot \mid x_0, x_t)$ in the backward direction, for $s \in [0, t]$, via (75). Consequently, this additional sampling step, which is absent from GSBM and DSBM, introduces only a small computational overhead.

Extension to Gaussian mixture variational family. Assume that we are provided with M Gaussian stochastic interpolants described by $\{I^{\psi_m}, \gamma^{\psi_m}\}_{m=1}^{M}$ with associated drift functions $\{v^{\mathrm{f},\psi_m}, v^{\mathrm{b},\psi_m}\}_{m=1}^{M}$. Following [15, Proposition 4], we may build Gaussian mixture stochastic interpolant with the following setting

$$
\begin{aligned}
\mathbb{Q}^{\psi}_{t|0,T} &= \sum_{m=1}^{M} w_m \mathbb{Q}^{\psi_m}_{t|0,T} \,, \\
v^{\mathrm{f},\psi}_t(x|x_0, x_T) &= \sum_{m=1}^{M} \tilde{w}_m(x|x_0, x_T) v^{\mathrm{f},\psi_m}_t(x|x_0, x_T) \,, \\
v^{\mathrm{b},\psi}_t(x|x_0, x_T) &= \sum_{m=1}^{M} \tilde{w}_m(x|x_0, x_T) v^{\mathrm{b},\psi_m}_t(x|x_0, x_T) \,, \\
\text{with } \tilde{w}_m(x|x_0, x_T) &= \frac{w_m \mathbb{Q}^{\psi_m}_{t|0,T}(x|x_0, x_T)}{\sum_{j=1}^{M} w_j \mathbb{Q}^{\psi_j}_{t|0,T}(x|x_0, x_T)} \,,
\end{aligned}
$$

where $\{w_m\}_{m=1}^{M}$ are mixture weights, that can be parameterized as well. We leave the practical implementation of this setting for future work.

F.3 Implementation of the score-based control variates

Continuous time setting. In practice, we consider a light polynomial parameterization for the control variates given by

$$
\alpha^{\mathrm{f}}(t, s) = (T - s) \sum_{k=0}^{K} \phi^{\mathrm{f}}_k(t) \left(\frac{s - t}{T - t} \right)^k \,, \; \alpha^{\mathrm{b}}(t, s) = s \sum_{k=0}^{K} \phi^{\mathrm{b}}_k(t) \left(\frac{s - t}{t} \right)^k \,, \tag{76}
$$

where $\phi^{\mathrm{f}}, \phi^{\mathrm{b}} : [0, T] \to \mathbb{R}^{K+1}$ are Fourier time-embedding neural networks with shared hyperparameter $K \geq 1$. The intuition is as follows, described for the forward direction, with the backward case being analogous. First, the control variate should have a limited effect for large values of t, when most of the trajectory has already been simulated; this motivates the rescaling by $(T - s)$, since s is sampled in $[t, T]$. Second, for a fixed t, the control variate should leverage information from sufficiently far in the future to influence the estimate at time t; this motivates the dependence on the lag $(s - t)$, which increases as s moves further away from t. With this parameterization, the three coefficient terms involving α^{f} in (45) and α^{b} in (46) remain fully tractable while relying on a lightweight parametrization.

Discrete time setting. We use the same parametrization as in the continuous-time setting for the schedules α^{f} and α^{b}. The only difference lies in the computation of the integral terms, $\int_t^{t_k} \frac{\alpha^{\mathrm{f}}(t,u)}{K_f(u)} \mathrm{d}u$ in the forward case and $\int_{t_k}^{t} \frac{\alpha^{\mathrm{b}}(t,u)}{K_b(u)} \mathrm{d}u$ in the backward case. Since K_f and K_b are discontinuous on $[0, T]$, these integrals are not directly handled as in the continuous-time setting. We therefore decompose them over the intermediate intervals $[t_i, t_{i+1}]$, on which K_f and K_b are constant, and compute the corresponding contributions in parallel.

G Experiment details and additional results

G.1 General setting

We train the crowd-navigation and single-cell models on a single NVIDIA RTX PRO 6000 GPU, the opinion-depolarization models on 8 AMD Instinct MI210 GPUs, and the image-translation models on 32 AMD Instinct MI210 GPUs. Training takes at most a few hours across all settings, with comparable computational costs for TSBM and GSBM.

Data. For the endpoint marginals, both training and test datasets contain 2048 independent samples.

Markovian projection training. Across all experiments except image translation, we use the same neural architecture for the Markovian projection, inspired by [30]. The drift is parameterized by an MLP with four hidden layers of width 128, Fourier time embeddings, skip connections, and SiLU activations, resulting in approximately 100,000 trainable parameters. To balance expressivity and computational cost, we use a single architecture for both forward and backward directions: the two networks share the first layer as a common backbone, and then branch into two direction-specific heads. This choice is motivated by the fact that the forward and backward dynamics explore the same state space and can therefore benefit from a shared representation. For image translation, we adopt a more suitable architecture, also proposed by [30], consisting of two direction-specific UNets with base width 128 and four residual blocks per resolution level. All networks are trained from scratch.

Stochastic gradient descent (SGD) is performed with Adam optimizer, with initial learning rate 3×10^{-4}. To improve training stability in high dimensions (opinion depolarizatoin and image translation), we follow GSBM and use an exponential moving average (EMA) with decay 0.999. In practice, the TSBM and GSBM objectives are optimized using Monte Carlo estimates of the corresponding losses. At each optimization step, we sample tuples $(\mathbf{X}_0, \mathbf{X}_T, t)$ with batch size $B = 512$. Following the recommendation of [24], we use a low-discrepancy sampler for t to reduce the variance of the loss estimator. Time samples are drawn on $[\varepsilon, T - \varepsilon]$, with $\varepsilon = 10^{-4}$, to avoid numerical issues at the boundaries. Each Markovian projection is trained for 2500 epochs.

TSBM-specific hyperparameters and implementation details. For TSBM, we further reduce the variance of the Markovian projection loss by sampling several pairs $(s, \mathbf{X}_s)$, in practice 16, from the variational bridge conditionally on $(t, \mathbf{X}_t)$. This provides a richer exploration of the landscape induced by ∇V. We also use a low-discrepancy sampler for s, again with boundary margin $\varepsilon = 10^{-4}$.

Noisy evaluations of ∇V can increase the variance of the loss, particularly at the beginning of training. To alleviate this issue, we draw inspiration from recent warmup strategies for consistency models [32, 46], whose objectives also involve potential-gradient terms, and optimize an *annealed* version of the TSBM regression target. More precisely, we rescale the ∇V contribution in the regression target by a factor that increases linearly from 0 to 1 until a prescribed fraction $\epsilon \in (0, 1]$ of the optimization is reached. Concretely, at the k-th SGD epoch, the forward regression target in (18) is replaced by

$$\mathbf{b}^{\mathrm{f},V}_{t,s}(x_t, x_s, x_T, k) = \frac{x_T - x_t}{T - t} - (T - s)\epsilon_k \nabla V_s(x_s) \ , \ \text{with } \epsilon_k = \min\left(\epsilon, \frac{k}{k_{\max}}\right) \ ,$$

where $k_{\max}$ denotes the warmup horizon, chosen smaller than the total number of SGD epochs. This schedule progressively introduces the state-cost contribution into the TSBM objective. Equivalently, this can be interpreted as replacing the reference measure $\mathbb{Q}^{V,\sigma}$ with the annealed reference

$$\mathrm{d}\mathbb{Q}^{V,\sigma,k}(x_{[0:T]}) \propto \exp\left(-\int_0^T \tfrac{V_t(x_t)}{\sigma^2}\mathrm{d}t\right)^{\epsilon_k} \mathrm{d}\mathbb{Q}^{\sigma}(x_{[0:T]}) \ ,$$

where the tilting is progressively introduced: it starts from $\epsilon_k = 0$, corresponding to the Brownian reference $\mathbb{Q}^{\sigma}$, and increases toward $\epsilon_k = \epsilon$, which is closer to the desired state-cost-based reference $\mathbb{Q}^{V,\sigma}$. We emphasize that this strategy *only applies during the first Markovian projection in each direction*; afterward, the original regression target is kept as such. Importantly, we do not rescale the state cost during the spline optimization step, as this led to worse empirical performance. In our experiments, we use this strategy for every task except single-cell data, with $\epsilon = 0.9$ and a warmup lasting 50% of the training epochs, corresponding to $k_{\max} = 1250$.

Although the contributions of the ∇V terms in the TSBM loss are theoretically damped by the factors s in the backward direction and $T - s$ in the forward direction, we sometimes observe undesirable boundary effects induced by ∇V, which can degrade satisfaction of the marginal constraints. To mitigate this issue, we replace the TSBM regression target with the GSBM target whenever $t/T \leq 0.1$

or $t/T \geq 0.9$. This leverages the strong marginal-matching properties of GSBM while retaining the influence of ∇V over the central portion of the trajectory. In practice, this simple substitution substantially improves numerical stability, and we therefore use it by default.

Learnable control variates. We use two separate networks, one for the forward direction and one for the backward direction. Following the parameterization in (76), each network uses an MLP with Fourier embeddings and a single hidden layer of width 128. We set the polynomial degree to 4 in all experiments; larger degrees did not improve performance. Each control-variate network contains approximately 15,000 trainable parameters. We optimize it jointly with the drift network using Adam with learning rate 10^{-4}. We use these control variates by default in all TSBM experiments. For all of them, the joint optimization was stable and did not require any additional stabilization procedure.

Reciprocal projection. For the spline parameterization, we use the same number of control points as in [30], independently of the dimension: 15 control points for the mean spline and 30 for the standard-deviation spline. The spline parameters are optimized with Adam. At the first reciprocal projection step, we run 1500 gradient steps; for subsequent reciprocal projections, we use 200 gradient steps. Spline optimization is performed using batches of size 1024 sampled from $\mathbb{P}_{0,T}$. For each batch, we sample 100 time points $t \in [0, T]$ to compute a Monte Carlo estimate of the reciprocal projection loss.

Bidirectional setting. For a fair comparison with the classical unidirectional variant, the pretraining stage is run for 2500 epochs with learning rate 3×10^{-4}, optimizing jointly the forward and backward objectives. The fine-tuning stage then uses a total of 10,000 epochs in each time direction, matching the training budget of four Markovian projection steps in the unidirectional setting. Thus, both variants use the same total number of forward and backward training epochs.

The main difference between the two settings lies in the frequency at which the training direction is alternated. In the bidirectional variant, we alternate more frequently to ensure regular coupling updates. In practice, every 250 epochs, we simulate samples in the current direction and use them as the training dataset for the opposite direction. This corresponds to 10% of the refresh period used in the classical procedure.

During fine-tuning, we reduce the learning rate of the drift network to 3×10^{-5}, following the recommandations from [5]. For TSBM, we also reduce the learning rate of the control variate to 10^{-5} during this phase.

Test setup. For evaluation and sample generation, we discretize $[0, T]$ using 200 uniformly spaced time steps. All test metrics are computed with batch size 2048. Following the terminology of [30], we report two complementary metrics to assess how accurately the Twisted Schrödinger Bridge problem is solved, both for continuous-time and discrete-time versions:

- **Feasibility**: in the forward direction, this measures the discrepancy between the generated terminal samples $\hat{\mathbf{X}}_T$ and ground-truth samples from the target marginal ν. Conversely, in the backward direction, it measures the discrepancy between samples $\hat{\mathbf{X}}_0$ generated by the backward SDE and ground-truth samples from μ. For all experiments except image translation, we compute this discrepancy using the Sinkhorn divergence implemented in the `SamplesLoss` Python package [16], with default hyperparameters; for image translation, we rather use the FID score.

- **Optimality**: this measures how well the stochastic optimal control cost (2) or (5) associated with the Twisted Schrödinger Bridge problem is minimized. We estimate a Riemann-discretized version of this cost in expectation along the generated dynamics, either in the forward or backward direction.

G.2 Experiments on the reciprocal projection

In this section, we aim to assess how the discrepancy between the TSBM and GSBM reciprocal objectives, given in (16) and (19), affects the approximation of the twisted bridge. Since it is generally intractable, its approximation error cannot usually be measured exactly. As explained in Section 4, we therefore consider the zero-potential case $V \equiv 0$, where the twisted bridge reduces to the tractable Brownian bridge. We set $T = 1$, and evaluate both methods across dimension $d \in \{2, 10, 50, 100\}$ and regularization $\sigma \in \{1, 2, 5\}$. In each setting, the reciprocal projection is performed with respect to $\pi_{0,1} = \mu \otimes \nu$, where $\mu = \mathrm{N}(-10 \cdot \mathbf{1}_d, 0.2\mathrm{I}_d)$ and $\nu = \mathrm{N}(10 \cdot \mathbf{1}_d, 0.2\mathrm{I}_d)$, with $\mathbf{1}_d$ being the d-dimensional vector with all entries equal to one.

Table 5: Accuracy of Brownian bridge estimation using the TSBM and GSBM reciprocal losses for increasing dimension d and regularization σ. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

Regularization	$\sigma = 1$		$\sigma = 2$		$\sigma = 5$	
Method	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM
$d = 2$	0.09	**0.02**	0.05	**0.02**	0.01	**0.0008**
$d = 10$	0.17	**0.10**	0.11	**0.08**	0.03	**0.004**
$d = 50$	0.58	**0.50**	0.41	**0.38**	0.14	**0.02**
$d = 100$	1.09	**1.00**	**0.77**	**0.77**	0.27	**0.04**

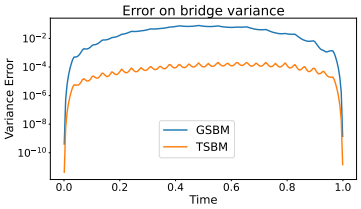

Figure 4: GSBM and TSBM learning error of the Brownian bridge variance over time ($d = 2, \sigma = 2$).

For both the TSBM and GSBM objectives, we run 1500 optimization steps on a test dataset constructed from samples of the independent coupling, and verify convergence in each setting. Following the standard spline routine used in GSBM and TSBM, we initialize the spline as the Brownian bridge, which is precisely the bridge to be recovered. Despite this favorable initialization, the GSBM loss underperforms relative to TSBM, highlighting a limitation of its reciprocal projection step. The forward KL-based metric $\mathbb{E}_{\pi_{0,1}}[\mathrm{KL}(\mathbb{Q}^{\sigma}_{t|0,1} \| \mathbb{Q}^{\psi^\star}_{t|0,1})]$ averaged over 200 uniformly sampled values of $t \in [0, 1]$ is reported in Table 5.

TSBM consistently yields a more accurate approximation than GSBM, especially with high σ, supporting our variational formulation. Figure 4 further illustrates this performance gap by showing the evolution over time of the error on the learned bridge variance in this simple setting.

G.3 Toy discrete time experiments

In this section, we present 2D toy experiments illustrating the performance of TSBM with discrete-time state costs, whose Markovian and reciprocal projection steps are described in Section C.4 and Section C.2, respectively. In all settings, the source and target distributions are independent Gaussian distributions. We mimic a sparse-observation scenario by defining, at each time t_k, a quadratic state cost of the form

$$V_k(x) = \frac{1}{M_k} \sum_{m=1}^{M_k} \frac{\|x - y_m^k\|^2}{2\gamma^2} \,,$$

where $\mathcal{Y}^k = \{y_m^k\}_{m=1}^{M_k} \subset \mathbb{R}^d$ denotes the set of observed points and $\gamma > 0$ is a bandwidth parameter. This cost attracts trajectories toward the observations in $\mathcal{Y}^k$ with uniform weights, while γ controls the interaction strength.

In the experiments below, we set $M_k = 50$ at all observation times and use $\gamma = 0.05$. Since the state cost is simple, we do not apply the annealing strategy for TSBM in this setting; empirically, doing so degraded convergence and led to worse numerical results. Moreover, we evaluate only the unidirectional variants of TSBM and GSBM, and use by default the control variate TSBM loss detailed in Proposition 15.

Single-observation setting. We first consider a single intermediate observation set at time $t_1 = 0.5$, with horizon $T = 1$, regularization $\sigma = 2$, source distribution $\mu = \mathrm{N}((-10, 0), 0.2\mathrm{I}_d)$, and target distribution $\nu = \mathrm{N}((10, 0), 0.5\mathrm{I}_d)$. The observations are sampled from $\mathrm{N}((0, 10), 0.5\mathrm{I}_d)$. Although the source and target distributions are aligned along the x-axis, the observations are deliberately placed far away along the y-axis, forcing the Brownian-bridge linear interpolation to be twisted in order to satisfy the additional constraint.

We observe that TSBM and GSBM generate very similar trajectories in both forward and backward directions, as shown in Figure 5. This visual similarity is reflected in the metrics reported in the left column of Table 6: the feasibility scores are identical, while optimality is comparable in the forward direction and only slightly worse for TSBM in the backward direction.

Multi-observation setting. We next consider a more challenging setting with two intermediate observation sets at times $t_1 = 0.4$ and $t_2 = 1.6$, with horizon $T = 2$, regularization $\sigma = 2$, source distribution $\mu = \mathrm{N}((-10, -5), 0.2\mathrm{I}_d)$, and target distribution $\nu = \mathrm{N}((10, 2.5), 0.2\mathrm{I}_d)$. The two observation sets are sampled independently from $\mathrm{N}((-3, 2.5), 0.5\mathrm{I}_d)$ and $\mathrm{N}((5, 7.5), 0.5\mathrm{I}_d)$, respectively. As in the single-observation setting, the observations are placed so as to induce a clear deviation from the linear interpolation between source and target distributions.

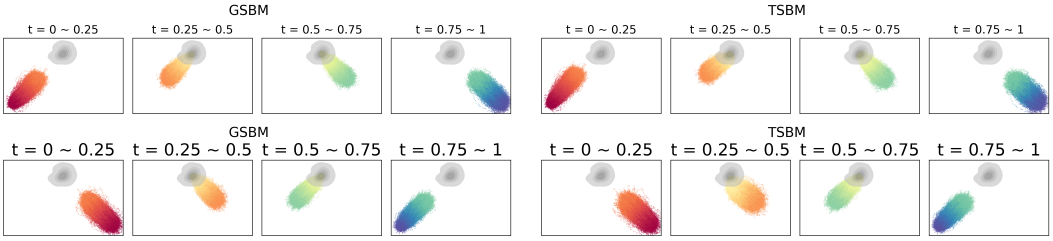

Figure 5: Learned TSB dynamics in the **single-observation** discrete-time setting. The intermediate observation distribution at $t_1 = 0.5$ is shown in grey. **Top**: forward time direction. **Bottom**: backward time direction.

Table 6: Results of 2D toy discrete-time experiments (single- and multi-observation settings). Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

Setting	Single-observation (full training)				Multi-observation (full training)			
Metric	Feasibility (↓)		Optimality (↓)		Feasibility (↓)		Optimality (↓)	
Method	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM
Forward	**0.01**$\pm$**0.003**	0.02$\pm$0.003	**651**$\pm$**1.0**	**649**$\pm$**1.0**	**0.005**$\pm$**0.001**	0.016$\pm$0.002	**636**$\pm$**0.8**	696$\pm$1.7
Backward	**0.003**$\pm$**0.001**	**0.004**$\pm$**0.001**	**651**$\pm$**1.3**	686$\pm$1.4	**0.006**$\pm$**0.001**	0.010$\pm$0.001	**636**$\pm$**1.0**	666$\pm$1.0

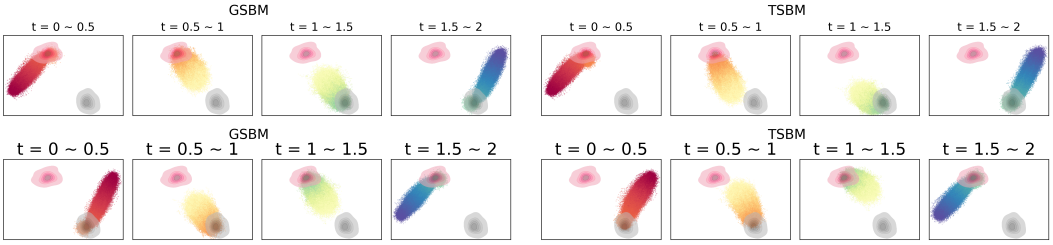

Figure 6: Learned TSB dynamics in the **multi-observation** discrete-time setting. The intermediate observation distributions at $t_1 = 0.4$ and $t_2 = 1.6$ are respectively shown in pink and grey. **Top**: forward time direction. **Bottom**: backward time direction.

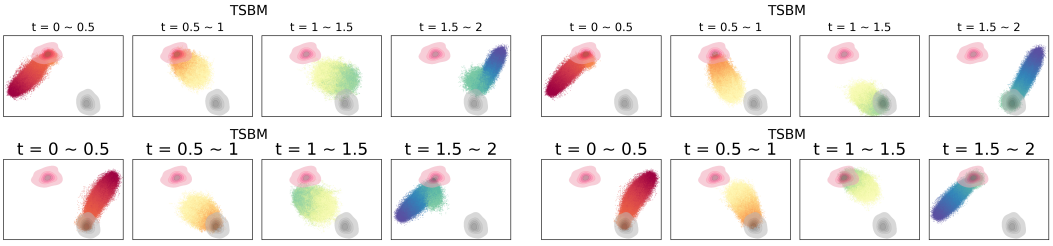

Figure 7: **Effect of control variate** on the learned TSBM dynamics in the **multi-observation** discrete-time setting (only 1 outer iteration is performed). The intermediate observation distributions at $t_1 = 0.4$ and $t_2 = 1.6$ are respectively shown in pink and grey. **Top**: forward time direction. **Bottom**: backward time direction. **(Left):** Without control variate. **(Right):** With control variate.

Table 7: Results of 2D toy multi-observation ablation study. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

Setting	Multi-observation (only 1 outer iteration)			
Metric	Feasibility (↓)		Optimality (↓)	
Method	TSBM (no CV)	TSBM	TSBM (no CV)	TSBM
Forward	0.07$\pm$0.005	**0.04**$\pm$**0.004**	4315$\pm$19.2	**795**$\pm$**2.7**
Backward	0.05$\pm$0.005	**0.02**$\pm$**0.002**	1706$\pm$9.5	**720**$\pm$**1.8**

We observe that TSBM successfully steers trajectories toward the first observation set, but deviates during the second part of the dynamics and slightly underperforms GSBM, which more accurately targets the second observation set. This behavior is confirmed by the optimality and feasibility metrics reported in the right column of Table 6, where TSBM achieves worse scores.

In this multi-observation setting, we also perform an ablation study highlighting the crucial role of the control variate (CV) in reducing the variance of the Markovian loss. We run TSBM for a single outer iteration, comparing the default version with CV to a variant without CV. Forward and backward trajectories are shown in Figure 7, and the corresponding metrics are reported in Table 7. While both variants satisfy the marginal constraints, removing the CV substantially degrades TSBM, as reflected by the large gap in optimality.

On the underperformance of TSBM in the discrete-time setting. Although TSBM and GSBM yield visually comparable trajectory estimates in Figures 5 and 6, we observe that TSBM systematically underperforms in this setting, despite the simplicity of the state cost. In contrast to the continuous-time case, TSBM does not reliably minimize the Twisted SB objective and exhibits large fluctuations during optimization, even when using the control variate strategy. Experiments conducted in the bidirectional setting led to the same conclusion.

We conjecture that this behavior stems from the ill-conditioning of the discrete-time Markovian projection loss, whose variance is high because samples from $\mathbf{X}_s \mid \mathbf{X}_t$ explore only a restricted region of the state space. This contrasts with the continuous-time formulation, where the state cost provides more distributed information along the trajectory. While introducing time dependence over large portions of the interval may seem to increase complexity, it appears in practice to regularize the problem by providing informative knowledge about future trajectory states. In the discrete-time formulation, this effect is absent: state-cost information is available only at the observation times. Improving the discrete-time formulation so as to make TSBM competitive with GSBM is left for future work.

G.4 Crowd navigation experiments

We consider the crowd navigation benchmarks of [29, 30], namely 'Stunnel', 'Gmm', and 'Vneck'. These tasks model the transport of a population of particles from a source distribution to a target distribution while avoiding obstacles, encoded through a structured state cost. They therefore naturally fit within the SOC formulation in (5).

For each benchmark, the state cost is decomposed into two components: an *obstacle* cost L_{obstacle}, which encodes the geometry of the environment, and an *interaction* cost $L_{\text{interaction}}$, which depends on the current population distribution. Specifically,

$$V_t(x_t) = \lambda_{\text{obs}} L_{\text{obstacle}}(x_t) + \lambda_{\text{int}} L_{\text{interaction}}(x_t; \hat{\mathbb{P}}_t) \ ,$$

where $\lambda_{\text{obs}} > 0$ and $\lambda_{\text{int}} > 0$ control the relative strength of the two terms. The two components are defined as follows.

• *Obstacle cost.* The obstacle term is task-specific and time-independent. In [29, 30], it is defined through sharp softplus-based penalties, which vanish outside the obstacle and increase approximately linearly inside it. In our experiments, we use a Gaussianized version of these penalties in order to provide more informative trajectory-level feedback through the ∇V term in the TSBM loss. This modification preserves the qualitative geometry of the obstacles while extending the influence of the cost beyond the immediate obstacle boundary, making the resulting gradients informative both inside and near the obstacles.

• *Interaction cost.* The interaction term, when activated, is time-dependent and is not part of the static obstacle geometry. It is computed from the current estimated population distribution $\hat{\mathbb{P}}_t$. Following [29, 30], we first consider the negative-entropy-like cost

$$L_{\text{interaction}}(x_t; \hat{\mathbb{P}}_t) = \log \hat{\mathbb{P}}_t(x_t).$$

Increasing λ_{int} encourages higher marginal entropy along the path and therefore broadens the support of the dynamics. Since $\hat{\mathbb{P}}_t$ is intractable, we follow [30] and approximate it with a variational particle-based estimate,

$$\hat{\mathbb{P}}_t(x_t) \approx \sum_{(x_0, x_T) \in \text{batch}} \mathbb{Q}_t^{\psi}(x_t \mid x_0, x_T) \ , \tag{77}$$

where the endpoint pairs (x_0, x_T) are taken from the current data batch, either during variational optimization of the reciprocal projection or in the Markovian bridge-matching loss.

Here, $\mathbb{Q}^{\psi}(\cdot \mid x_0, x_T)$ denotes the spline bridge parametrization associated with the pair (x_0, x_T), either being optimized in the reciprocal step or already learned in the Markovian step. This empirical approximation is accurate in practice due to the large batch sizes used, at least 512 samples.

We also consider the congestion cost

$$L_{\text{interaction}}(x_t; \hat{\mathbb{P}}_t) = \mathbb{E}_{y_t \sim \hat{\mathbb{P}}_t}\left[\frac{2}{\|x_t - y_t\|^2 + 1}\right]. \tag{78}$$

In this case, increasing λ_{int} encourages stronger dispersion of the generated particles. In practice, as in [30], we approximate this expectation by Monte Carlo using a single particle $y_t \sim \hat{\mathbb{P}}_t$ for each x_t.

High-dimensional extension. The crowd navigation costs are originally defined in 2D through softplus penalties based on the distance between the evaluated point and the obstacle centers, with ellipses for Stunnel, circles for Gmm, and triangles for Vneck. To assess how the methods scale with the ambient dimension, we extend each geometry to dimension $d \in \{10, 50\}$ as follows.

For an obstacle center with 2D coordinates (c_1, c_2), we define its lifted counterpart as

$$\tilde{c} = (c_1, c_1, \ldots, c_1, c_2) \in \mathbb{R}^d,$$

where the first coordinate is repeated over the first $d-1$ longitudinal dimensions, while the last coordinate corresponds to the transverse dimension. This construction preserves the qualitative 2D geometry while allowing the ambient dimension to vary.

Distances to obstacle centers are then computed using the RMS-normalized squared distance

$$\Delta_d(x, \tilde{c}; s_{\text{long}}, s_{\text{trans}}) = \frac{2}{d}\left[\sum_{i=1}^{d-1}\left(\frac{x_i - \tilde{c}_i}{s_{\text{long}}}\right)^2 + \left(\frac{x_d - \tilde{c}_d}{s_{\text{trans}}}\right)^2\right],$$

where $s_{\text{long}} > 0$ and $s_{\text{trans}} > 0$ control the longitudinal and transverse scalings. These parameters are chosen so that the effective obstacle width remains comparable as d increases, while recovering the standard 2D geometry when $d = 2$.

Experimental setting. Following [30], we use a single interaction cost for each crowd navigation benchmark: entropy for Vneck, and congestion for Stunnel and Gmm. In 2D, we keep the same values of λ_{obs} and λ_{int} as in [30], where these weights have a natural interpretation in terms of physical modeling. In higher dimensions, we focus only on the obstacle cost, since the interaction cost becomes prohibitively memory-intensive. By default, we use $T = 2$ and $\sigma = 2$. All weight values are reported in Table 8.

Table 8: State-cost weights used in the crowd navigation experiments. Here λ_{ent} indicates the weight of entropy interaction cost while λ_{cgst} indicates the weight of the congestion interaction cost.

Setting	Dimension	λ_{obs}	λ_{ent}	λ_{cgst}
Stunnel	$d = 2$	1500	0	50
	$d = 10$	1000	0	0
	$d = 50$	800	0	0
Gmm	$d = 2$	1500	0	5
	$d = 10$	30	0	0
	$d = 50$	35	0	0
Vneck	$d = 2$	3000	8	0
	$d = 10$	10000	0	0
	$d = 50$	10000	0	0

Additional results (main experiment). In complement of Table 1 giving only results of forward dynamics, we provide in Table 9 the metrics computed with *backward* generated dynamics. In the same manner, we give in Figure 8 the backward counterpart of the 2D forward dynamics of Figure 1. We draw the same conclusion exposed in the main : TSBM better integrates the state cost than GSBM when optimizing the SB cost, but slightly underperforms at source marginal constraint.

Table 9: **Backward** results of crowd navigation experiments for Stunnel and Vneck, with increasing d. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

State cost	Stunnel				Vneck			
Metric	Feasibility (↓)		Optimality (↓)		Feasibility (↓)		Optimality (↓)	
Method	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM
$d = 2$	$0.040_{\pm 0.005}$	$\mathbf{0.030}_{\pm \mathbf{0.005}}$	$977_{\pm 0.8}$	$\mathbf{953}_{\pm \mathbf{0.9}}$	$\mathbf{0.02}_{\pm \mathbf{0.005}}$	$0.13_{\pm 0.03}$	$70.1_{\pm 0.4}$	$\mathbf{56.9}_{\pm \mathbf{0.4}}$
$d = 10$	$\mathbf{1.12}_{\pm \mathbf{0.01}}$	$\mathbf{1.13}_{\pm \mathbf{0.01}}$	$759_{\pm 0.8}$	$\mathbf{746}_{\pm \mathbf{0.7}}$	$\mathbf{1.05}_{\pm \mathbf{0.01}}$	$1.37_{\pm 0.02}$	$\mathbf{257}_{\pm \mathbf{1.3}}$	$259_{\pm 0.5}$
$d = 50$	$\mathbf{17.33}_{\pm \mathbf{0.07}}$	$18.06_{\pm 0.09}$	$3589_{\pm 2.7}$	$\mathbf{3577}_{\pm \mathbf{2.8}}$	$\mathbf{18.1}_{\pm \mathbf{0.1}}$	$20.3_{\pm 2.6}$	$1386_{\pm 2.3}$	$\mathbf{1380}_{\pm \mathbf{2.4}}$

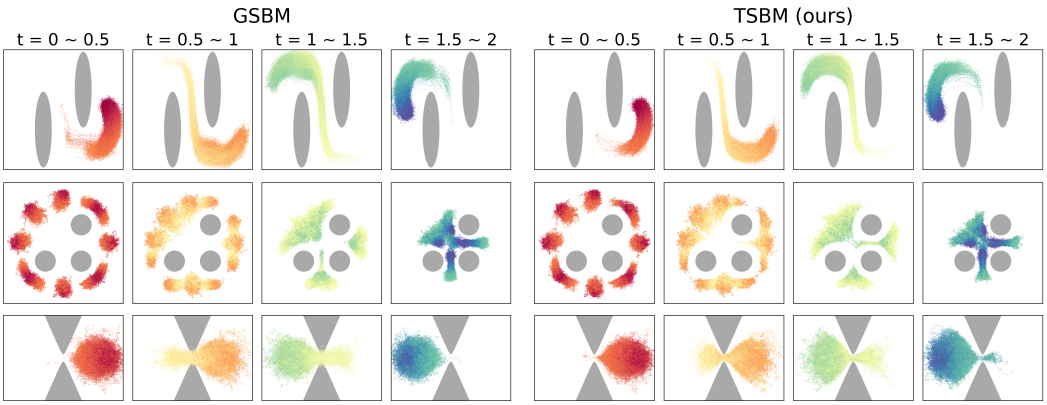

Figure 8: **Backward** learned TSB dynamics for Stunnel (top), Gmm (middle), Vneck (bottom) with $d = 2$. This is complementary to Figure 1.

G.5 Single-cell experiments

Single-cell trajectory inference aims to reconstruct the temporal evolution of a cell population from snapshot measurements. In typical single-cell RNA sequencing experiments, cells are profiled at several developmental or experimental time points, but the measurement process is destructive: a cell observed at one time cannot be observed again later. As a result, data collected at different times does not form individual trajectories. Consequently, the dataset provides marginal distributions over cellular states at selected time points, but no pairing between cells observed at different times. The task is to infer a continuous-time population dynamics whose time marginals are compatible with these snapshots.

Embryoid body dataset. We use the embryoid body (EB) single-cell RNA-seq dataset analyzed in the PHATE study of [34] and later used by [55]. This dataset consists of a developmental single-cell RNA-seq time course, in which each cell is represented by a high-dimensional gene-expression vector and annotated with its collection time. More generally, single-cell datasets can be viewed as matrices whose rows correspond to cells and whose columns correspond to molecular features, such as genes in scRNA-seq.

For the EB dataset, the biological goal is to infer how a population of cells evolves over developmental time, as cells differentiate and branch toward distinct lineages. From a trajectory-inference perspective, the observations collected at each day define empirical population distributions over cellular states. The dataset contains 18,203 cells, with snapshots at days 0, 6, 12, 18, and 24. We perform the experiments in 2-, 5- and 50-dimensional PCA spaces.

In our formulation, days 0 and 24 define the initial and terminal marginals, respectively, while the intermediate snapshots at days 6, 12, and 18 are treated as intermediate observations. We rescale the physical time interval $[0, 24]$ to a computational horizon $[0, T]$ with $T = 2$, so that these intermediate observations correspond to $t = 0.5$, $t = 1$, and $t = 1.5$.

Sparse-observation regime. Rather than using the full set of available cells at intermediate observation times, as is common in the generative modeling literature, e.g., [55, 36, 23], *we deliberately retain only a subset of these observations*, fixed to 25%, as illustrated in Figure 9, and incorporate them through a Twisted SB formulation.

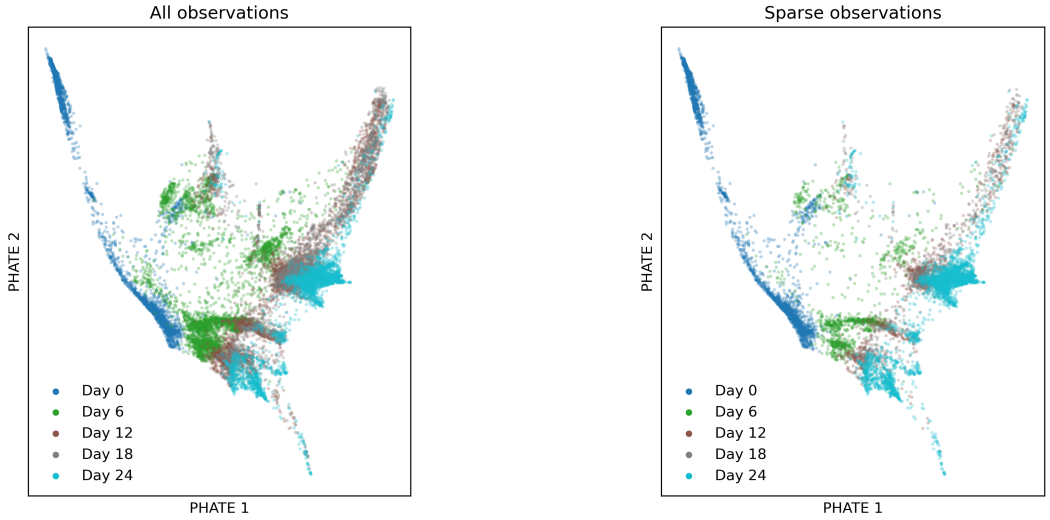

Figure 9: 2D Single-cell data splitted by snapshot times : endpoint samples (day 0, day 24) and intermediary samples (day 6, day 12, day 18). **(Left):** Multi-marginal setting : all intermediary observations are available. **(Right):** Sparse-observation setting considered in our experiments : only 25% of the intermediary observations are available, but the endpoint marginals are complete.

This choice is motivated by practical considerations rather than being purely artificial. In real single-cell studies, measurements can be costly, invasive, or technically limited, so observations may only be available at a few time points or for a restricted number of cells. The problem is therefore better viewed not as a fully specified multi-marginal interpolation task, but as a partially observed trajectory-inference problem, a regime for which TSBM is particularly well suited.

In this setting, the Brownian prior provides a baseline dynamics that bridges the source and target distributions without using additional biological information. Twisting this prior with a data-informed state cost then steers the dynamics toward regions compatible with the observed snapshots, while avoiding the hard enforcement of intermediate marginals that would arise in standard bridge-matching approaches. By contrast, when all intermediate marginal constraints are imposed, multi-marginal methods are more appropriate, as they are explicitly designed to interpolate between temporally ordered marginals.

Our experiments therefore focus on the sparse single-cell regime, where the goal is not to match a dense sequence of marginals, but to infer plausible continuous-time population dynamics from limited snapshot information. We rely on the SB formulation through the SOC objective in (5), with a state cost specifically designed for this setting and detailed below. This choice precludes a direct comparison with methods that exploit all intermediate data, but enables a fair comparison with GSBM, which is trained under the same formulation as TSBM.

Design of the single-cell state cost. We now describe in detail the construction of the state cost used in our experiments, inspired by the seminal work of [55]. We present its definition from an abstract perspective, as it applies more broadly to sparse-observation settings beyond the single-cell experiments considered here.

Let K denote the number of snapshot times, including the boundary times, and write $\{t_k\}_{k=1}^K \subset [0,T]$, with $t_1 = 0$ and $t_K = T$, for the corresponding increasing sequence of observation times ($K = 5$ in our experiment). For each $k \in \{1,\dots,K\}$, let

$$\mathcal{Z}^k = \{z_1^k, \dots, z_{M_k}^k\} \subset \mathbb{R}^d$$

denote the collection of samples observed at time t_k, and set $\mathcal{Z} = \bigcup_{k=1}^K \mathcal{Z}^k$. The endpoint samples are assumed to be fully observed, in contrast with the intermediate snapshots, which are only sparsely available. For $x \in \mathbb{R}^d$ and $z \in \mathbb{R}^d$, define

$$r(x,z) = (\|x - z\| - h)_+ \,, \tag{79}$$

where $h > 0$ is a distance threshold.

We recall the single-cell density potential introduced by [55], inspired by biological priors,

$$U(x) = \sum_{j=1}^{k_{NN}} r(x, z_{j(x)}) \ , \tag{80}$$

where $z_{j(x)}$ denotes the j-th nearest neighbor in $\mathcal{Z}$ of x in Euclidean distance, and $k_{NN} \geq 1$ is a nearest-neighbor hyperparameter. We propose adapting this potential into an informative and differentiable state cost for our sparse-observation setting.

We first replace the non-smooth thresholded distance (79) by its smooth counterpart

$$r(x, z) = \mathrm{Softplus}(\|x - z\| - h) \ .$$

Then, for each snapshot set $\mathcal{Z}^k$, we recursively define on $j = 1, \dots, k_{NN}$ the unnormalized weights

$$\tilde{w}_m^{k,j}(x) = \Big(1 - \textstyle\sum_{j'<j} w_m^{k,j'}(x)\Big)_+ \exp\big(-\beta r(x, z_m^k)\big) \ , \qquad \forall z_m^k \in \mathcal{Z}^k \ ,$$

with the convention that the empty sum is zero, along with their normalized version

$$w_m^{k,j}(x) = \frac{\tilde{w}_m^{k,j}(x)}{\sum_{m'=1}^{M_k} \tilde{w}_{m'}^{k,j}(x)}, \qquad \forall z_m^k \in \mathcal{Z}^k,$$

where $\beta > 0$ is an inverse-temperature parameter. We then define the snapshot-specific state cost

$$V_k(x) = \frac{1}{k_{NN}} \sum_{j=1}^{k_{NN}} \sum_{m=1}^{M_k} w_m^{k,j}(x)\, r(x, z_m^k) \ .$$

With this construction, observations closer to x receive larger weights. In the limit $\beta \to +\infty$, the weights concentrate on the minimizers of $z \mapsto r(x, z)$ over $\mathcal{Z}^k$. Combined with the recursive masking in the definition of the weights, this recovers the contribution of the k_{NN} nearest neighbors of x within $\mathcal{Z}^k$, and *therefore matches the potential* (80) *restricted to that snapshot* (up to renormalization by k_{NN}). In contrast, when $\beta \to 0$, the weights become nearly uniform over $\mathcal{Z}^k$, making the state cost weakly informative. Thus, β controls the locality of the interaction with sparse observations, and mediates a trade-off between fidelity to the observed samples and generalization. In Figure 10, we illustrate the effect of β on the simulated trajectories. We observe that lower values of β induce unlikely trajectories with high transport cost, namely trajectories connecting distant data clusters. This behavior is expected, since uniform weighting over the observed data (obtained with low β) no longer exploits data locality. This motivates rather using large β, in line with the approach of [55].

To exploit sparse observations across the full simulation time interval in our Twisted Schrödinger Bridge setting, we introduce a time-dependent state cost obtained by linearly interpolating between neighboring snapshot costs. For $t \in [0, T]$, let $k^-(t)$ and $k^+(t)$ denote the indices of the closest observation times immediately before and after t, respectively. We define

$$V_t(x_t) = \begin{cases} \dfrac{t_{k^+(t)} - t}{t_{k^+(t)} - t_{k^-(t)}} V_{k^-(t)}(x_t), & \text{if } k^+(t) = K, \\ \dfrac{t - t_{k^-(t)}}{t_{k^+(t)} - t_{k^-(t)}} V_{k^+(t)}(x_t), & \text{if } k^-(t) = 1, \\ \dfrac{t - t_{k^-(t)}}{t_{k^+(t)} - t_{k^-(t)}} V_{k^+(t)}(x_t) + \dfrac{t_{k^+(t)} - t}{t_{k^+(t)} - t_{k^-(t)}} V_{k^-(t)}(x_t), & \text{otherwise.} \end{cases}$$

This construction yields a state cost that is continuous in time and differentiable with respect to the state variable. Intuitively, at any time $t \in [0, T]$, V_t combines information from nearby past and future observations, while reducing to the corresponding snapshot-specific cost at observation times. We note that more flexible, learnable interpolation weights could also be considered; however, in our experiments, the linear interpolation already provided satisfactory performance. To avoid introducing an artificial influence from the endpoint marginals, which have a different status from the sparse intermediate observations, we set their contribution to zero.

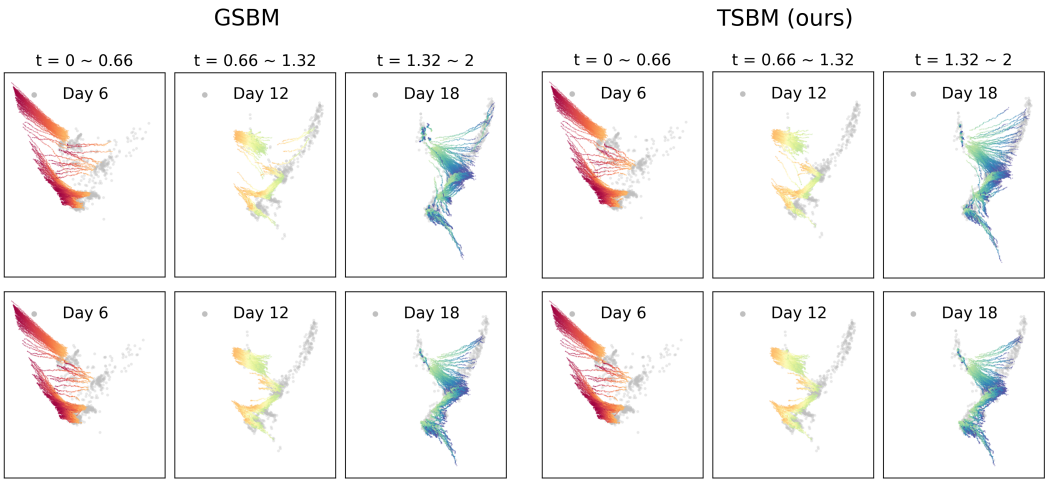

Figure 10: Learned TSB dynamics for 2D single-cell data for $\beta = 20$ (top) and $\beta = 100$ (bottom). Available intermediary observations that influence the trajectories are marked in grey.

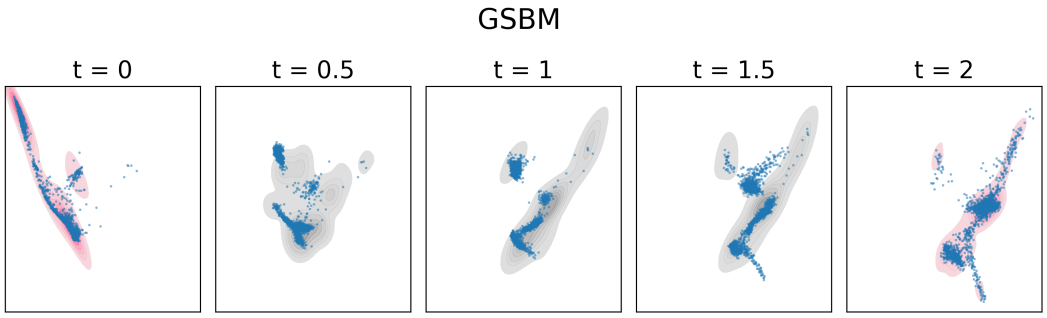

Figure 11: Snapshots of learned GSBM dynamics for 2D single-cell ($\beta = 100$, $\sigma = 0.1$) with respect to available observations (endpoint *marginal* samples in pink, intermediary *sparse* samples in grey). This is complementary to Figure 2.

Experimental setting. For $d = 2$, we use $h = 0.1$ and $k_{NN} = 5$, following the recommendation of [55]. In higher-dimensional settings, we keep $k_{NN} = 5$, which provides a favorable trade-off between computational cost and expressivity, and choose h manually according to the characteristic nearest-neighbor scale observed in the data. Specifically, we use $h = 1$ for $d = 5$ and $h = 7$ for $d = 50$. For numerical stability, we implement the state cost using logit weights rather than the normalized weights directly. To obtain smooth single-cell trajectories, we choose $\sigma = 0.1$. In this setting, *we only consider the bidirectional variant*, which was shown in the crowd navigation experiments to stabilize the training of TSBM, making it the most suitable choice for real-world applications.

Additional results. In Figure 11, we report the GSBM counterpart of Figure 2 from Section 5. Interestingly, the estimated intermediate marginals appear to connect observation clusters more effectively under TSBM than under GSBM, in agreement with the 2D optimality metric reported in Table 2. For illustration, we also show the corresponding 2D snapshots for $\beta = 20$ in Figure 12. In this case, no clear visual difference is observed between GSBM and TSBM, consistently with the metrics in Table 2.

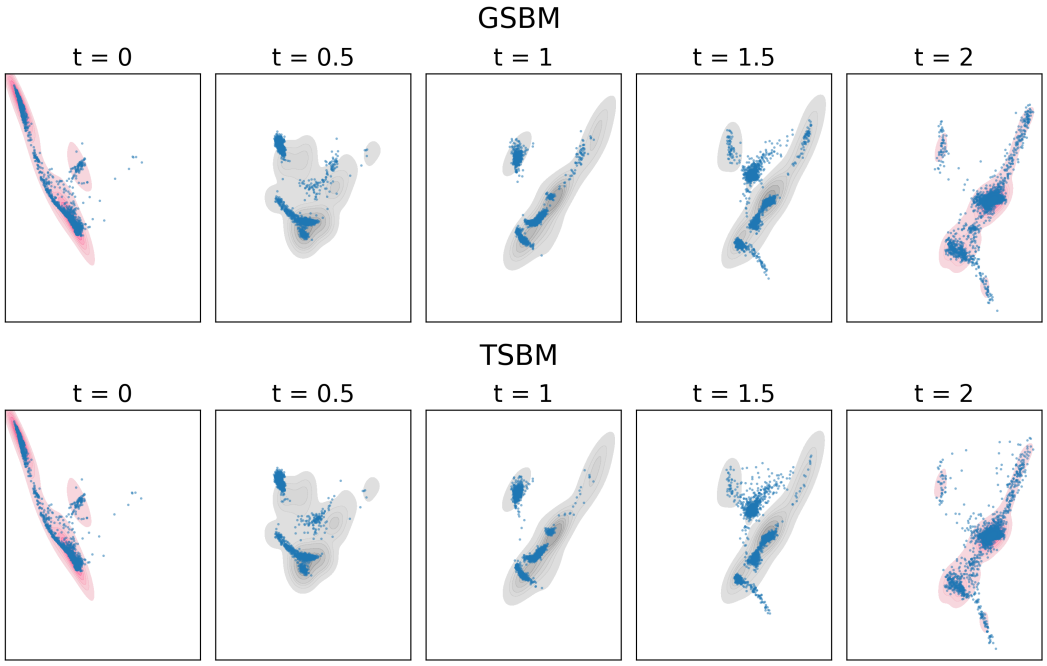

Figure 12: Snapshots of learned GSBM dynamics for 2D single-cell ($\beta = 20$, $\sigma = 0.1$) with respect to available observations (endpoint *marginal* samples in pink, intermediary *sparse* samples in grey). **(Top):** GSBM, **(Bottom):** TSBM.

G.6 High-dimensional experiments

For both experiments described below, we set $T = 1$ and $\sigma = 0.5$.

AFHQ image translation. We consider bidirectional translation between the cat and dog classes of the AFHQ dataset, containing 5153 and 4739 training images, respectively. Images are resized to 64×64 and normalized to $[-1, 1]$.

Following GSBM, we construct the state cost using a pretrained variational autoencoder (VAE), with encoder E and decoder D. Given an endpoint pair (x_0, x_1), we interpolate between their latent representations $z_0 = E(x_0)$ and $z_1 = E(x_1)$ using spherical linear interpolation (Slerp):

$$z_t = \mathrm{Slerp}(t, z_0, z_1) = \frac{\sin((1-t)\Omega)}{\sin \Omega} z_0 + \frac{\sin(t\Omega)}{\sin \Omega} z_1, \qquad \Omega = \arccos\langle z_0, z_1 \rangle,$$

where we follow the latent normalization convention of GSBM. Decoding this interpolation yields a reference trajectory $r_t(x_0, x_1) = D(z_t)$, which we use to define the state cost

$$V_t(x) = \frac{5}{d} \|x - r_t(x_0, x_1)\|_1, \qquad d = 3 \times 64 \times 64.$$

Here, the dependence on (x_0, x_1) is implicit in our notation. Indeed, the potential is coupled to the stochastic interpolant used to parameterize the bridge: each evaluation of $V_t(x_t)$ uses an intermediate state sampled from $\mathbb{Q}^{\psi}_{t|0,1}(\cdot \mid x_0, x_1)$ and the corresponding reference trajectory, with endpoint pairs drawn either from the original data or from the generated samples.

Quantitative results are reported in Table 3. Additional qualitative examples, with five outputs per method in each translation direction, are provided in Figures 13 and 14.

Opinion depolarization. This experiment investigates whether TSBM can counteract polarizing interactions in a high-dimensional population of opinions. Each agent's opinion is represented by a vector in $\mathbb{R}^d$ where $d = 1000$, with initial and terminal distributions $\mu = \mathcal{N}(0, \mathrm{diag}(4, 0.25, \ldots, 0.25))$ and $\nu = \mathcal{N}(0, 4I_d)$. The objective is to steer the initially anisotropic population toward an isotropic distribution, while the congestion state cost (78), associated to weight $\lambda_{\text{cgst}} = 10$, discourages the concentration of similar opinions.

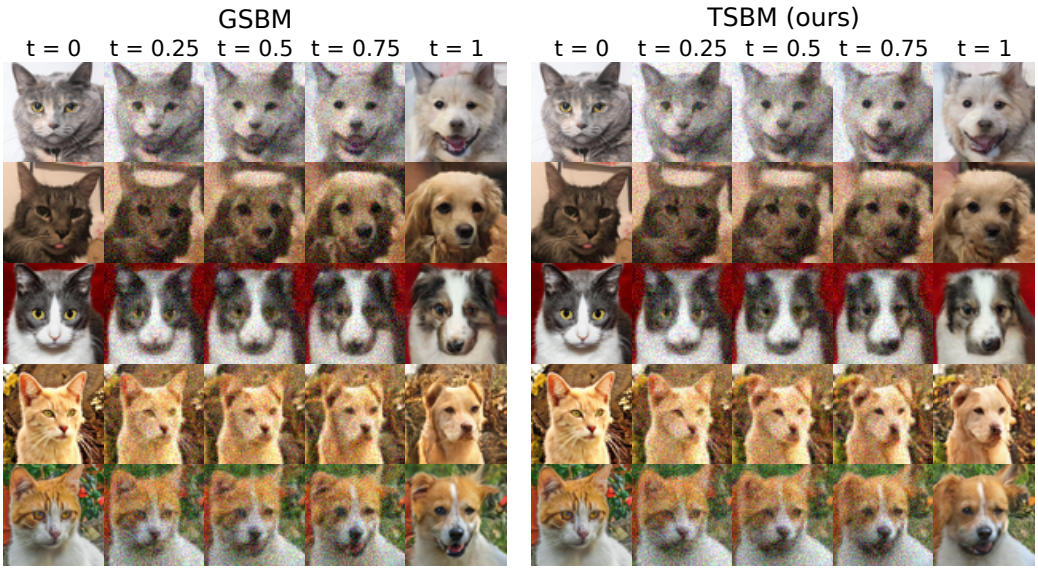

Figure 13: Additional **AFHQ cat-to-dog translations** for GSBM and TSBM.

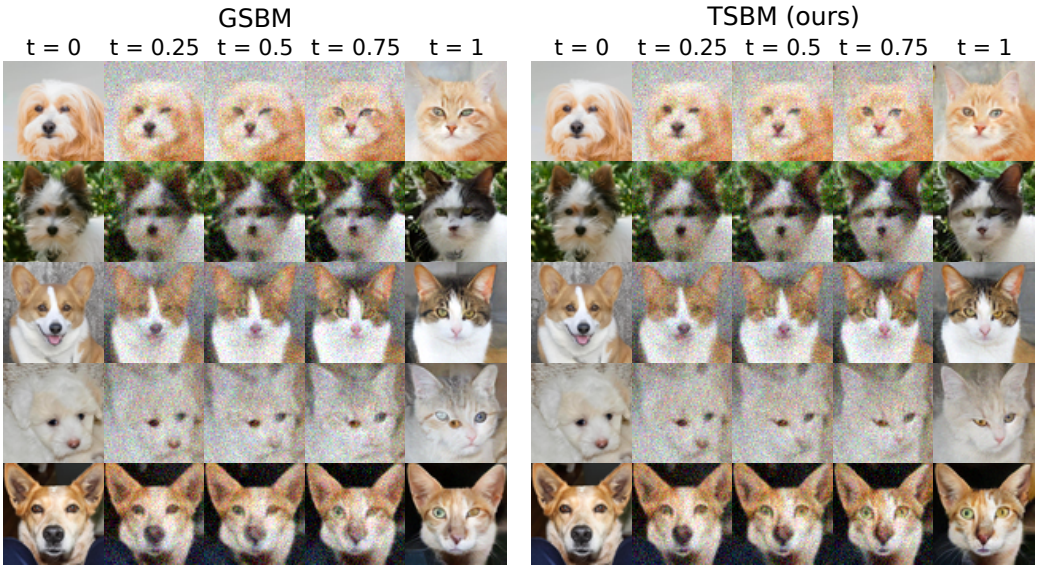

Figure 14: Additional **AFHQ dog-to-cat translations** for GSBM and TSBM.

Unlike the previous experiments, where the uncontrolled dynamics are purely Brownian, we follow GSBM and consider a polarizing mean-field drift $f_{\text{polarize}}(x_t, \hat{\mathbb{P}}_t)$, defined in [30, Eq. (37)]. Driven by a common random signal, this drift attracts opinions responding similarly and repels those responding oppositely, promoting the formation of polarized groups. To incorporate these dynamics into the Twisted SB framework, we parameterize the Brownian control as $u_t^\theta(x_t) = f_{\text{polarize}}(x_t, \hat{\mathbb{P}}_t) + v_t^\theta(x_t)$, where only the residual control v^θ is learned. For the Markovian projection, we retain the respective GSBM and TSBM objectives, fitting u^θ to the conditional variational regression target in GSBM, see (20), or the twisted regression target in TSBM, see (18). In both cases, the distribution $\hat{\mathbb{P}}_t$ entering f_{polarize} is estimated from mini-batches sampled from the variational bridge, following (77). The reciprocal projection, however, differs more substantially: TSBM retains our theoretically motivated variational objective (16), whereas GSBM replaces its original objective (19) with

$$\mathcal{L}_{\mathcal{R}}^{f}(\psi) = \int_0^T \mathbb{E}_{\Pi_{0,t,T}^{\psi}} \left[\left\| v_t^{\mathrm{f},\psi}(\mathbf{X}_t | \mathbf{X}_0, \mathbf{X}_T) - f_{\text{polarize}}(\mathbf{X}_t, \hat{\mathbb{P}}_t) \right\|^2 + V_t(\mathbf{X}_t) \right] \mathrm{d}t \ ,$$

without providing a theoretical justification for this modification. Consequently, we hypothesize that TSBM's improved performance in this experiment may arise from differences in both the reciprocal and Markovian projection objectives.

G.7 Ablation studies

(A) **Unidirectional vs. bidirectional.** We report the evolution of feasibility and optimality throughout training for the 2D-Stunnel, 2D-Gmm, and 2D-Vneck tasks in Figures 15 to 17, respectively. To complement these convergence plots, we report the final metric values across all crowd-navigation tasks for the forward and backward directions in Tables 10 and 11, respectively. Overall, bidirectional training substantially improves the stability and smoothness of TSBM convergence, without necessarily improving its final performance. In contrast, GSBM benefits less consistently from bidirectional training, which can even degrade performance in some cases, notably on 2D-Gmm.

(B) **Lower-entropy setting.** We report forward and backward feasibility and optimality results for the lower-entropy setting ($\sigma = 1$) in Tables 12 and 13, respectively. Consistent with the higher-entropy setting, TSBM achieves lower SB costs than GSBM. The corresponding forward trajectories, shown in Figure 18, are also straighter than those obtained with $\sigma = 2$.

(C) **Effect of the control variate.** We compare TSBM with and without the control variate (CV) in the 2D setting, restricting the ablation to the first outer iteration of unidirectional training. Forward and backward feasibility and optimality results are reported in the rightmost columns of Tables 12 and 13, respectively. The results generally favor the use of CV, despite the visually similar forward trajectories shown in Figure 19.

(D) **About the feasibility–optimality trade-off.** We investigate whether TSBM's lower marginal accuracy relative to GSBM stems from the reciprocal or Markovian projection. We denote their reciprocal projections by [RT] (Reciprocal Twisted) and [RG] (Reciprocal Generalized), and their Markovian projections by [MT] (Markovian Twisted), [MTcv] (with control variates), and [MG] (Markovian Generalized). We consider four combinations: [RT/MT] (TSBM without control variates), [RT/MTcv] (TSBM with control variates), [RT/MG] (hybrid algorithm), [RG/MG] (GSBM). We evaluate these variants on the Gmm, Stunnel, and Vneck benchmarks over eight random seeds, with results reported in Table 14.

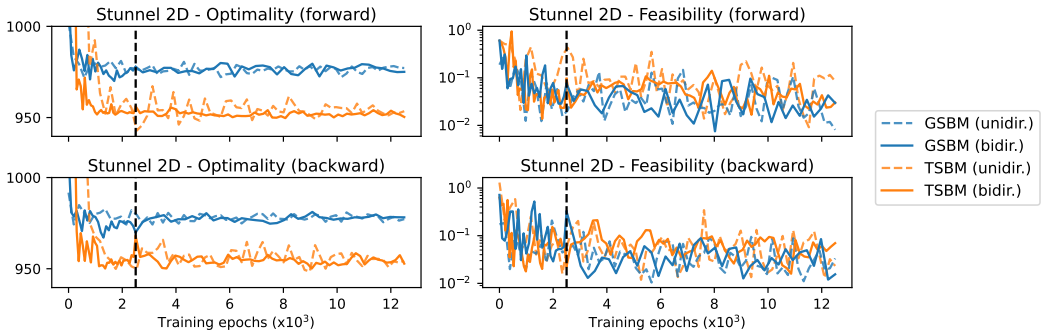

Figure 15: Evolution of **forward and backward metrics** for TSBM and GSBM on **2D-Stunnel** with $\sigma = 2$. The vertical dotted line marks the end of bidirectional pretraining or, equivalently, the first unidirectional outer iteration.

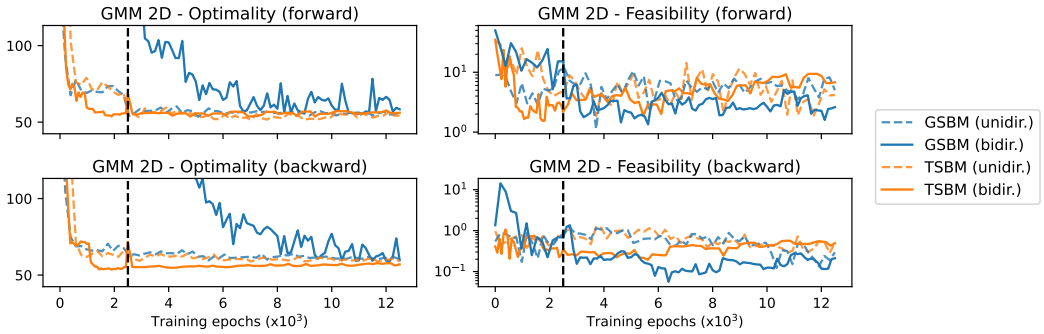

Figure 16: Evolution of **forward and backward metrics** for TSBM and GSBM on **2D-Gmm** with $\sigma = 2$. The vertical dotted line marks the end of bidirectional pretraining or, equivalently, the first unidirectional outer iteration.

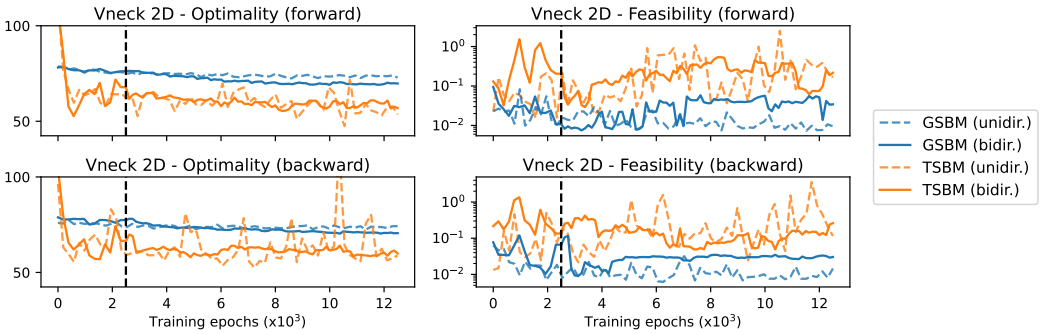

Figure 17: Evolution of **forward and backward** metrics for TSBM and GSBM on **2D-Vneck** with $\sigma = 2$. The vertical dotted line marks the end of bidirectional pretraining or, equivalently, the first unidirectional outer iteration.

Table 10: **Unidirectional vs. bidirectional (forward)**: results of crowd navigation experiments for 2D-Stunnel, 2D-GMM and 2D-Vneck. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

Setting	Unidirectional ($\sigma = 2$, full training)				Bidirectional ($\sigma = 2$, full training)			
Metric	Feasibility (↓)		Optimality (↓)		Feasibility (↓)		Optimality (↓)	
Method	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM
2D-Stunnel	**0.006**±0.001	0.056±0.007	978±1.0	**956**±1.0	**0.02**±0.002	**0.02**±0.003	977±1.2	**952**±0.9
2D-Gmm	3.66±0.45	**2.83**±0.30	55.5±0.4	**54.4**±0.3	**2.77**±0.50	7.10±0.50	58.2±1.9	**55.0**±0.3
2D-Vneck	**0.008**±0.004	0.16±0.02	69.2±0.5	**53.1**±0.4	**0.04**±0.02	0.20±0.02	67.5±0.6	**55.2**±0.3

Table 11: **Unidirectional vs. bidirectional (backward)**: results of crowd navigation experiments for 2D-Stunnel, 2D-GMM and 2D-Vneck. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

Setting	Unidirectional ($\sigma = 2$, full training)				Bidirectional ($\sigma = 2$, full training)			
Metric	Feasibility (↓)		Optimality (↓)		Feasibility (↓)		Optimality (↓)	
Method	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM	GSBM	TSBM
2D-Stunnel	0.040±0.005	**0.030**±0.005	977±0.8	**953**±0.9	**0.01**±0.002	0.07±0.007	978±0.8	**952**±1.0
2D-Gmm	**0.31**±0.03	**0.30**±0.03	58.9±0.3	**58.2**±0.3	**0.31**±0.05	0.61±0.05	53.1±3.8	**55.1**±0.2
2D-Vneck	**0.02**±0.005	0.13±0.03	70.1±0.4	**56.9**±0.4	**0.03**±0.02	0.32±0.04	67.0±0.3	**57.2**±0.2

Table 12: **Forward** results of crowd navigation **ablation studies** for Stunnel, Gmm and Vneck with $d = 2$. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

Setting	Bidirectional ($\sigma = 1$, only pretraining)				Unidirectional ($\sigma = 2$, only 1 outer iteration)			
Metric	Feasibility (↓)		Optimality (↓)		Feasibility (↓)		Optimality (↓)	
Method	GSBM	TSBM	GSBM	TSBM	TSBM (no CV)	TSBM	TSBM (no CV)	TSBM
2D-Stunnel	0.09±0.02	**0.03**±0.003	958±0.9	**949**±0.8	**0.18**±0.02	0.25±0.02	964±1.0	**957**±1.0
2D-Gmm	**4.92**±0.83	7.68±0.63	**61.8**±3.2	**62.3**±0.3	11.5±1.3	**4.6**±0.8	68.8±0.3	**65.5**±0.5
2D-Vneck	**0.03**±0.004	0.09±0.01	64.9±0.2	**48.0**±0.2	0.03±0.002	**0.02**±0.005	**60.2**±0.3	**60.5**±0.3

Table 13: **Backward** results of crowd navigation **ablation studies** for Stunnel, Gmm and Vneck with $d = 2$. Results are averaged over 8 metric evaluations. Bold indicates the best result in each setting.

Setting	Bidirectional ($\sigma = 1$, only pretraining)				Unidirectional ($\sigma = 2$, only 1 outer iteration)			
Metric	Feasibility (↓)		Optimality (↓)		Feasibility (↓)		Optimality (↓)	
Method	GSBM	TSBM	GSBM	TSBM	TSBM (no CV)	TSBM	TSBM (no CV)	TSBM
2D-Stunnel	0.51±0.02	**0.02**±0.002	952±0.6	**947**±0.5	0.05±0.005	**0.02**±0.001	**948**±0.7	951±0.9
2D-Gmm	**0.13**±0.02	0.28±0.02	67.9±12.1	**54.2**±0.3	**0.40**±0.04	0.46±0.07	**59.0**±0.3	62.4±0.2
2D-Vneck	**0.08**±0.006	0.46±0.06	66.0±0.3	**56.6**±0.4	**0.06**±0.002	**0.06**±0.006	58.4±0.2	**57.9**±0.4

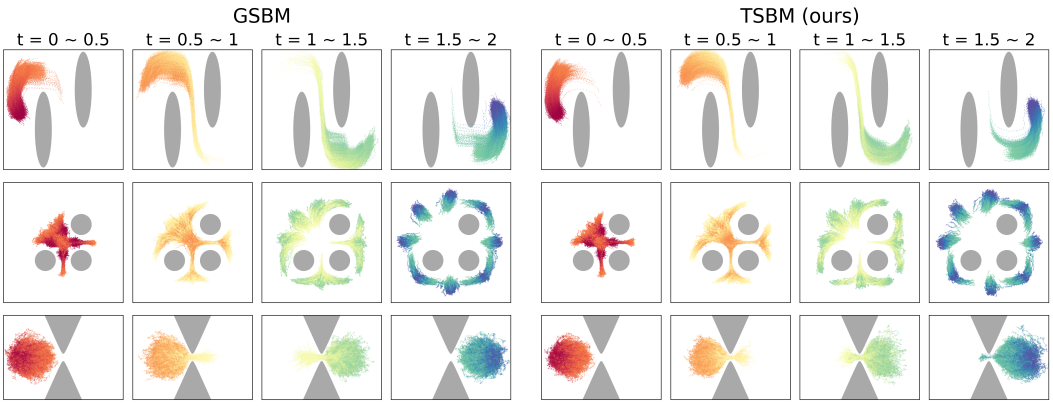

Figure 18: Learned TSB dynamics for Stunnel (top), Gmm (middle), Vneck (bottom) with $d = 2$ and $\sigma = 1$. These results are obtained after the stage of pretraining within the bidirectional variant.

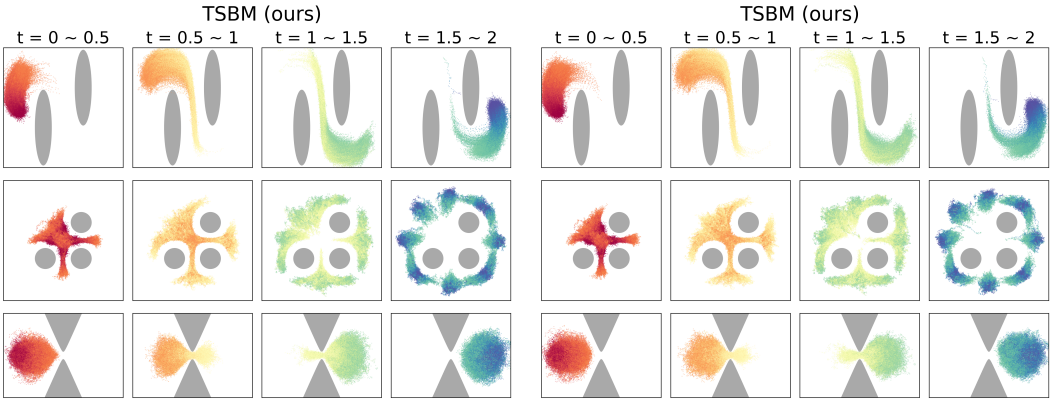

Figure 19: Learned TSBM dynamics for Stunnel (top), Gmm (middle), Vneck (bottom) with $d = 2$ and $\sigma = 1$. **(Left):** Without control variate. **(Right):** With control variate. These results are obtained after the first Markovian projection of the unidirectional variant.

Table 14: **Ablation study of the Reciprocal and Markovian projection objectives.** Results are averaged over 8 random seeds and reported in the forward (f) and backward (b) directions. Bold indicates the best result for each setting, direction, and metric.

Metric		Feasibility (↓)				Optimality (↓)			
Setting	Direction	[RT/MTcv]	[RT/MT]	[RT/MG]	[RG/MG]	[RT/MTcv]	[RT/MT]	[RT/MG]	[RG/MG]
GMM	Forward	2.1±0.5	3.1±0.6	**0.8**±0.3	**0.8**±0.2	54.5±0.6	**53.2**±0.5	56.6±0.6	56.3±0.6
GMM	Backward	0.11±0.02	0.16±0.02	**0.059**±0.006	0.079±0.006	55.5±0.4	**54.7**±0.4	58.4±0.4	58.3±0.4
Stunnel	Forward	0.061±0.008	0.027±0.005	**0.023**±0.005	0.024±0.003	**953.0**±1.5	955.6±1.6	976.4±1.5	976.1±1.7
Stunnel	Backward	0.056±0.009	0.055±0.008	0.030±0.005	**0.026**±0.005	**951.3**±1.1	954.8±0.9	975.6±0.8	975.6±0.9
Vneck	Forward	0.07±0.02	0.15±0.02	**0.0346**±0.0348	0.037±0.030	58.5±0.4	**55.1**±0.4	67.6±0.5	67.3±0.7
Vneck	Backward	0.52±0.06	0.15±0.03	0.021±0.016	**0.017**±0.011	**55.4**±0.5	58.1±0.4	68.0±1.3	67.7±1.3

Comparing [RT/MG] with [RG/MG] reveals similar feasibility and optimality, suggesting that the reciprocal projection is not responsible for the observed gap. Instead, the main difference arises from the Markovian projection: [RT/MG] achieves better feasibility, whereas [RT/MT] yields lower optimality costs. Indeed, [MG] regresses on the drift of the learned variational bridge, introducing a bias relative to the exact twisted projection unless this bridge is exact. In contrast, [MT] avoids this bias but requires estimating the time integral of ∇V along trajectories, increasing optimization variance. Under our training budget, this appears to favor optimality at the expense of marginal accuracy. Finally, comparing [RT/MT] with [RT/MTcv] shows that control variates affect this trade-off without consistently improving both metrics, making their benefit dependent on the metric of interest.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims made in the Abstract and Introduction accurately reflect the theoretical and experimental results established in the paper. Our theory is developed in Section 3 and our method is illustrated trough some numerical experiments in Section 5. We provide a detailed comparison with related work and a thorough discussion of our contributions in Section 4.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses the limitations of the work performed in the conclusion section.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: The paper provides the full set of assumptions for its main results and and their proof in the Appendix.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: All the experimental setup as well as details are provided in the supplement.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example

(a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.

(b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

(c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: We provide the code to reproduce our experiments and the used real datasets are public.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Details are in the Appendix.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.

- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Error bars and standard deviations are provided when needed.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Details are in the Appendix.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We confirm that the research conducted our paper conform, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The conclusion discusses the broader impact of our work.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: We do not release pre-trained models or scraped datasets.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All the datasets and publicly available models used are properly used and credited.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not introduce or release a new dataset, codebase, checkpoint, or benchmark asset as part of the submission.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.

- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing, annotation work commissioned by the authors, or research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The work does not involve human subjects, so IRB approval or an equivalent review is not applicable.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs are not an important, original, or non-standard component of the proposed methodology.

    Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.