
Learning non-equilibrium diffusions with Schrödinger bridges: from exactly solvable to simulation-free

Stephen Y. Zhang*
University of Melbourne

Michael P. H. Stumpf
University of Melbourne

Abstract

We consider the Schrödinger bridge problem which, given ensemble measurements of the initial and final configurations of a stochastic dynamical system and some prior knowledge on the dynamics, aims to reconstruct the “most likely” evolution of the system compatible with the data. Most existing literature assume Brownian reference dynamics, and are implicitly limited to modelling systems driven by the gradient of a potential energy. We depart from this regime and consider reference processes described by a multivariate Ornstein-Uhlenbeck process with generic drift matrix $A \in \mathbb{R}^{d \times d}$. When A is asymmetric, this corresponds to a non-equilibrium system in which non-gradient forces are at play: this is important for applications to biological systems, which naturally exist out-of-equilibrium. In the case of Gaussian marginals, we derive explicit expressions that characterise exactly the solution of both the static and dynamic Schrödinger bridge. For general marginals, we propose MVOU-OTFM, a simulation-free algorithm based on flow and score matching for learning an approximation to the Schrödinger bridge. In application to a range of problems based on synthetic and real single cell data, we demonstrate that MVOU-OTFM achieves higher accuracy compared to competing methods, whilst being significantly faster to train.

1 Introduction

We are interested in reconstruction of stochastic dynamics of individuals from static population snapshots. This is a central problem with applications arising across the natural and social sciences, whenever longitudinal tracking of individuals over time is either impossible or impractical [56, 35, 48, 43, 34, 1]. In simple terms, consider a system of indistinguishable particles $x_t \in \mathbb{R}^d$ undergoing some unobserved temporal dynamics. The practitioner observes the system to have distribution $x_0 \sim \rho_0$ at an initial time $t = 0$ and later to be $x_1 \sim \rho_1$ at a final time $t = 1$. The question is then: can we, under some suitable assumptions, reconstruct the *continuous-time* behaviour of the system for the unobserved time interval $0 < t < 1$?

The Schrödinger bridge problem (SBP), by now a centrepiece of the theoretical literature on this topic, places this task on a theoretical footing in terms of a mathematical formulation of this problem as a large deviations principle on the path space [28] and a (stochastic) least action principle intimately linked to optimal transportation theory [12]. Given a *reference process* that encodes prior knowledge on the dynamics, the SBP can be understood as identifying the “most likely” evolution of the system that is compatible with the snapshot observations.

The SBP and related topics have enjoyed a great deal of recent interest from both applied and theoretical perspectives. Many applications arise from biological modelling of cell dynamics [7, 42]. The majority of existing work assumes, either explicitly or implicitly, that the system of interest is *potential driven*, that is, driven by the gradient of a potential energy [26, 13, 52, 5, 51, 57, 6]. In fact,

*Correspondence to syz@syz.id.au




Figure 1: The Schrödinger bridge problem with a multivariate Ornstein-Uhlenbeck reference process (1) can be solved via a generalised entropic transport problem and characterisation of the $\mathbb{Q}$ -bridges. For non-Gaussian endpoints, score and flow matching provide a route to building neural approximations without simulation.

most studies do not consider *any* prior drift, corresponding to the setting where the reference process is purely an isotropic Brownian motion. A few studies [5, 45, 15] allow for scalar Ornstein-Uhlenbeck (OU) processes, motivated by applications in the generative modelling domain. Since scalar OU dynamics have unidimensional drift which can always be written as the gradient of a scalar function, all of these are also potential driven systems.

Changing the choice for the reference process provides a route to dealing with systems which are not potential driven, i.e. driven by a drift that is a non-conservative vector field. Abundant motivations for modelling these dynamics arises from biological systems and other kinds of active matter; these exist naturally far from equilibrium [60, 37, 19, 17], exhibiting irreversible dynamics at a non-equilibrium steady state [27]. While a few studies [53, 44, 25] allow for reference processes described by a SDE with a *generic* drift term, they rely on computationally expensive simulation procedures such as numerical integration and suffer from accuracy issues in high dimensions. Here we explore a complementary approach and concentrate our attention on a family of *linear* reference dynamics as a middle ground between physical relevance and analytical tractability. Specifically, we consider reference processes arising in a class of linear drift-diffusion dynamics described by the SDE

$$d\mathbf{X}_t = \mathbf{A}(\mathbf{X}_t - \mathbf{m}) dt + \boldsymbol{\sigma} d\mathbf{B}_t. \quad (1)$$

These processes are also known as *multivariate Ornstein-Uhlenbeck* (mvOU) processes, often used as a model of non-equilibrium systems in the physics literature [27, 20, 18]. Indeed, for asymmetric drift matrix $\mathbf{A}$ the system (1) has a drift that is no longer the gradient of a potential. As an example in the case of isotropic diffusion $\boldsymbol{\sigma} = \sigma \mathbf{I}$, when $\mathbf{A}$ is asymmetric and has all eigenvalues with negative real part, a non-equilibrium steady state exists where the dynamics are irreversible and exhibit nonzero probability currents while the population density is unchanged [18, 20]. While allowing a broader range of reference dynamics, the statistics of (1) remain analytically tractable – in particular it is a Markovian Gaussian process with explicit formulae for its mean, covariance, and transition kernel. This provides avenues to efficient solution of the SBP that sidesteps the need for numerical integration, and as we will see, even analytical expressions in the Gaussian case. In particular, the family of SBP problems that we model is strictly larger than the standard (Brownian) SBP, which we recover as a special case when $\mathbf{A}, \mathbf{m} = \mathbf{0}$.

The contributions of our paper are twofold. Leveraging results on mvOU processes, we characterise for Gaussian endpoints the Gaussian Schrödinger Bridge (GSB) with reference dynamics described by (1). To handle general marginals, we develop a *simulation free* training procedure based on score and flow matching [52] to solve the SBP. We conclude by demonstrating our results in a range of synthetic and real data examples. We find that our approach solves the generalised SBP faster and more accurately than comparable simulation-based algorithms, at the cost of assuming the form (1) of the reference process. Our findings highlight the tradeoff between (i) analytically tractable but more restrictive models and (ii) more expressive non-parametric models, which are much less tractable to learn, both in terms of computational cost and statistical complexity.

2 Background and related work

2.1 Schrödinger bridges

In what follows we provide a concise summary of the SBP. For details we refer readers to [28, 12] for in-depth discussion. We work in $\mathcal{X} \subseteq \mathbb{R}^d$ and denote by $\mathcal{P}(\mathcal{X})$ the space of probability distributions on $\mathcal{X}$. Let $C([0, 1], \mathcal{X})$ denote the space of continuous paths $\omega_t : [0, 1] \mapsto \mathcal{X}$ valued in $\mathcal{X}$, and informally we will refer to *path measures* as probability measures supported on $C([0, 1], \mathcal{X})$. Viewing stochastic processes as random variables valued in $C([0, 1], \mathcal{X})$, path measures prescribe their law. Let $\mathbb{Q}$ be a general path measure describing a *reference* Markov stochastic process. For prescribed initial and final marginals $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$, the Schrödinger bridge problem can be written

$$\min_{\mathbb{P} \in \mathcal{P}(C([0, 1], \mathcal{X})) : \mathbb{P}_0 = \rho_0, \mathbb{P}_1 = \rho_1} \text{KL}(\mathbb{P} | \mathbb{Q}), \quad (\text{SBP-dyn})$$

where the minimum is taken over all candidate processes $\mathbb{P}$ absolutely continuous with respect to $\mathbb{Q}$ that are compatible with the observed data at $t = 0, 1$. This *dynamic* form of the SBP, while elegant, is unwieldy for practical purposes owing to the formulation on path space. A well known result [28] connects the dynamical SBP with its static counterpart:

$$\min_{\mathbb{P}_{01} \in \Pi(\rho_0, \rho_1)} \text{KL}(\mathbb{P}_{01} | \mathbb{Q}_{01}). \quad (\text{SBP-static})$$

Furthermore, the static and dynamic SBP are interchangeable via the disintegration identity regarding the optimal law $\mathbb{P}^*$:

$$\mathbb{P}^*(\cdot) = \int d\mathbb{P}_{01}^*(x_0, x_1) \mathbb{Q}^{(x_0, x_1)}(\cdot), \quad (2)$$

where $\mathbb{Q}^{xy}$ denotes the law of the $\mathbb{Q}$ -bridge conditioned at $(0, x_0)$ and $(1, x_1)$. In other words, solution of (SBP-dyn) amounts to solving (SBP-static) for $\mathbb{P}_{01}^*$ followed by construction of $\mathbb{P}^*$ as per (2) by taking mixtures of $\mathbb{Q}$ -bridges.

The problem (SBP-static) can be reformulated as an entropy-regularised optimal transportation problem [12] where the effective cost matrix is the log-transition kernel under $\mathbb{Q}$. This admits efficient solution via the Sinkhorn-Knopp algorithm [14] in the discrete case. This provides a practical roadmap to constructing dynamical Schrödinger bridges by first solving (SBP-static), then using construction of $\mathbb{Q}$ -bridges with (2) to build a solution of (SBP-dyn).

2.2 Probability flows and (score, flow)-matching

Probability flows Consider a generic drift-diffusion process in d dimensions with drift $\mathbf{v}_t(\mathbf{x})$ and diffusivity $\boldsymbol{\sigma}_t$, described by an Itô diffusion whose marginal densities evolve following the corresponding Fokker-Planck equation (FPE):

$$d\mathbf{X}_t = \mathbf{v}_t(\mathbf{X}_t) dt + \boldsymbol{\sigma}_t d\mathbf{B}_t, \quad \partial_t p_t(\mathbf{x}) = -\nabla \cdot (p_t(\mathbf{x}) \mathbf{v}_t(\mathbf{x})) + \nabla \cdot (\mathbf{D}_t \nabla p_t(\mathbf{x})), \quad (3)$$

where $\mathbf{D}_t = \frac{1}{2} \boldsymbol{\sigma}_t \boldsymbol{\sigma}_t^\top$ is the diffusivity matrix. Equivalently, the FPE can be rewritten in the form of a continuity equation involving a *probability flow* field $\mathbf{u}_t$ [30, 3, 2, 52, 31]:

$$\partial_t p_t(\mathbf{x}) = -\nabla \cdot [p_t(\mathbf{x}) \mathbf{u}_t(\mathbf{x})], \quad \mathbf{u}_t(\mathbf{x}) = \mathbf{v}_t(\mathbf{x}) - \mathbf{D}_t \nabla_{\mathbf{x}} \log p_t(\mathbf{x}), \quad (4)$$

By recognising the form of the continuity equation in (4), it is apparent that the family of *marginal distributions* $\{p_t(\mathbf{x})\}_{t \geq 0}$ generated by the dynamics specified in (3) are also generated by the *probability flow (PF)-ODE*:

$$\dot{\mathbf{X}}_t = \mathbf{u}_t(\mathbf{X}_t), \quad \mathbf{X}_0 \sim p_0 \quad (5)$$

The central quantity that allows us to convert between the SDE (3) and the PF-ODE (4) is the gradient of the log-density $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) =: \mathbf{s}_t(\mathbf{x})$, also known as the *score function*. Under mild regularity conditions, knowledge of the score $\mathbf{s}(\mathbf{x})$ allows for sampling from $p(\mathbf{x})$ via Langevin dynamics: the SDE $d\mathbf{X}_t = \frac{1}{2} \mathbf{s}(\mathbf{x}) dt + d\mathbf{W}_t$ has stationary distribution $p(\mathbf{x})$ [46]. While typically one needs to resort to approximations to learn $\mathbf{s}$ [55, 46, 47], in setting of the Brownian (and as we will see, the mvOU) bridge, one has closed form expressions for the score and the objective.

Conditional flow matching Stated in its original form, let $t \mapsto p_t(\mathbf{x})$ be a family of marginals on $t \in [0, 1]$ satisfying the continuity equation $\partial_t p_t(\mathbf{x}) = -\nabla \cdot (p_t(\mathbf{x}) \mathbf{u}_t(\mathbf{x}))$, where $\mathbf{u}_t(\mathbf{x})$ is a time-dependent vector field. Suppose further that $p_t(\mathbf{x})$ admits a representation as a *mixture*

$$p_t(\mathbf{x}) = \int dq(\mathbf{z}) p_{t|z}(\mathbf{x}|\mathbf{z})$$

where $\mathbf{z} \sim q(\mathbf{z})$ is some latent variable and $t \mapsto p_{t|z}(\mathbf{x})$ are called the *conditional probability paths*. We introduce *conditional flow fields* $\mathbf{u}_{t|z}$ that generate the conditional probability paths, i.e. $\partial_t p_{t|z}(\mathbf{x}) = -\nabla \cdot (p_{t|z}(\mathbf{x}) \mathbf{u}_{t|z}(\mathbf{x}|\mathbf{z}))$. Then, the insight presented in [30, Theorem 1] is that, in fact,

$$\mathbf{u}_t(\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z})} \left[\frac{p_{t|z}(\mathbf{x}|\mathbf{z})}{p_t(\mathbf{x})} \mathbf{u}_{t|z}(\mathbf{x}|\mathbf{z}) \right],$$

as can be easily verified by checking the continuity equation. When $p_t(\mathbf{x})$ is not tractable but $q(\mathbf{z})$, $p_{t|z}(\mathbf{x})$ and $\mathbf{u}_{t|z}(\mathbf{x})$ are, as is the case for the Schrödinger bridge (2), conditional flow matching is useful: Lipman et al. [30, Theorem 2] prove that minimising

$$L_{\text{CFM}}(\theta) = \mathbb{E}_{t \sim U[0,1], \mathbf{z} \sim q} \mathbb{E}_{\mathbf{x} \sim p_{t|z}} \|\mathbf{v}_{\theta,t}(\mathbf{x}) - \mathbf{u}_{t|z}(\mathbf{x})\|^2 \quad (6)$$

is equivalent to regression on the true (marginal) vector field $\mathbf{u}_t$.

Simulation-free Schrödinger bridges Previous work [52, 39] has exploited the connection (2) between the dynamic and static SBP problems to create solutions to (SBP-dyn) without the need to numerically simulate SDEs. In particular, [52] proposed to utilise score matching and flow matching simultaneously to learn the probability flow and score of the dynamical SBP with a Brownian reference process. Crucially, they exploit availability of closed form expressions for the Brownian bridge and minimise the objective

$$L_{[\text{SF}]^2 M}(\theta, \varphi) = \mathbb{E}_{\substack{t, (x_0, x_1) \sim U[0,1] \otimes \pi \\ \mathbf{z} \sim p_{t|(x_0, x_1)}}} \left[\|\mathbf{v}_{\theta,t}(\mathbf{z}) - \mathbf{u}_{t|(x_0, x_1)}(\mathbf{z})\|^2 + \lambda_t \|\mathbf{s}_{\varphi,t}(\mathbf{z}) - \mathbf{s}_{t|(x_0, x_1)}(\mathbf{z})\|^2 \right], \quad (7)$$

where $\pi = \mathbb{P}_{01}^*$ is the optimal coupling solving (SBP-static). In the above $\mathbf{u}_{t|(x_0, x_1)}$, $\mathbf{s}_{t|(x_0, x_1)}$ are, respectively, the flow and score of the Brownian bridge conditioned on (x_0, x_1) , for which there are readily accessible closed form expressions [52, Eq. 8]: $\mathbf{u}_t(\mathbf{x}) = \frac{1-2t}{2t(1-t)}(\mathbf{x} - \bar{\mathbf{x}}_t) + (\mathbf{x}_1 - \mathbf{x}_0)$ and $\mathbf{s}_t(\mathbf{x}) = (\sigma^2 t(1-t))^{-1}(\bar{\mathbf{x}}_t - \mathbf{x})$ with $\bar{\mathbf{x}}_t = tx_1 + (1-t)x_0$. Motivating the use of the conditional objective (7) in practice, the authors prove [52, Theorem 3.2] that minimising (7) is equivalent to regressing against the SB flow and score. Since $\mathbb{P}_{01}^*$ can be obtained by solving an static entropic optimal transport problem.

2.3 Related work

The generalisation of dynamical optimal transport and the Schrödinger bridge to linear reference dynamics has been studied [11, 9, 10] from the viewpoint of stochastic optimal control. However, as was pointed out by [5], these studies have primarily focused on theoretical aspects of the problem, such as existence and uniqueness. In particular, the result for the Gaussian case [9] is in terms of a system of coupled matrix differential equations and does not lend itself to straightforward computation. More generally, forward-backward SDEs corresponding to a continuous-time iterative proportional fitting scheme [53, 44] have been proposed for general reference processes.

For the Gaussian case, the availability of closed form solutions is by now classical [36, 49], and more recent work provides analytical characterisations of Gaussian *entropy-regularised* transport [5, 22, 32]. [5] provides explicit formulae that characterise the bridge marginals as well as the force (bridge control) for a Brownian reference as well as a class of *scalar* OU references. To the authors' knowledge, however, all explicit formulae for Gaussian Schrödinger bridges assume either Brownian or scalar OU dynamics. The application of flow matching [30, 3] as a simulation-free technique to learn Schrödinger bridges was proposed in [52], but the authors restrict consideration to a Brownian reference process. Recently [33] apply flow matching for stochastic linear control systems, however they are concerned with interpolating distributions and do not study the SBP.

Finally, concurrent work [38] proposes a related flow matching scheme for approximately solving the $\mathbb{Q}$ -SBP where the reference process $\mathbb{Q}$ is described by a general, potentially non-linear diffusion. This comes however at the cost of using learned neural approximations for bridges and log-transition densities of $\mathbb{Q}$. Our work treats the complementary case of linear reference dynamics, in which we may avail ourselves of analytical formulae for these quantities.

3 Schrödinger bridges for non-equilibrium systems

3.1 Multivariate Ornstein-Uhlenbeck bridges

As is evident from (2), the dynamical formulation of the $\mathbb{Q}$ -SBP amounts to solution of the *static* $\mathbb{Q}$ -SBP (SBP-static) together with characterisation of the $\mathbb{Q}$ -bridges, i.e. the reference process $\mathbb{Q}$ conditioned on initial and terminal endpoints. For $\mathbb{Q}$ described by a linear SDE of the form (1), explicit formulae for the $\mathbb{Q}$ -bridges are available and we state these in the form of the following theorem.

Theorem 1 (SDE characterisation of mvOU bridge, adapted from results in [8]). *Consider the d -dimensional mvOU process (1). Conditioning on $(0, \mathbf{x}_0)$ and $(T, \mathbf{x}_T)$, the bridges of this process $\mathbf{Y}_t = \mathbf{X}_t | \{\mathbf{X}_0 = \mathbf{x}_0, \mathbf{X}_T = \mathbf{x}_T\}, 0 \leq t \leq T$ are generated by the SDE*

$$d\mathbf{Y}_t = (\mathbf{A}(\mathbf{Y}_t - \mathbf{m}) + \mathbf{c}_{t|(\mathbf{x}_0, \mathbf{x}_T)}) dt + \boldsymbol{\sigma} d\mathbf{B}_t, \quad \mathbf{c}_{t|(\mathbf{x}_0, \mathbf{x}_T)} = -\boldsymbol{\Lambda}_t^{-1}(\mathbf{Y}_t - \mathbf{k}_t). \quad (8)$$

where $\boldsymbol{\Lambda}_t = \int_0^{T-t} e^{-s\mathbf{A}} e^{-s\mathbf{A}^\top} ds$ and $\mathbf{k}_t = e^{-(T-t)\mathbf{A}}(\mathbf{x}_T - \mathbf{m}) + \mathbf{m}$.

We remark that this characterisation can be found implicitly in the results of [8]; however, their results were conceived for a more general setting of time-varying coefficients and hence is stated in terms of the state transition matrix, which does not in general admit an explicit formula. In our case, computation of the control term $\boldsymbol{\Lambda}_t$ relies upon a unidimensional integral as a key quantity and crucially this is independent of the endpoints, meaning it incurs a one-off computational cost. Although the derivation of Theorem 1 follows the same procedure as [8], we provide details in the Appendix since the work of Chen and Georgiou [8] is written for a control audience. Using the SDE characterisation (8) of the $\mathbb{Q}$ -bridge and the fact that it remains a Gaussian process, we obtain explicit formulae for the *score* and *flow* fields as well as mean and covariance functions of the $\mathbb{Q}$ -bridge. One can furthermore check that the Brownian bridge formulae are recovered as a special case.

Theorem 2 (Score and flow for multivariate Ornstein-Uhlenbeck bridge). *For the mvOU bridge conditioned on $(0, \mathbf{x}_0)$, $(T, \mathbf{x}_T)$, denote by $p_{t|(\mathbf{x}_0, \mathbf{x}_T)}$ the density at time $0 < t < T$. Then, score function and probability flow of the bridge are respectively*

$$\mathbf{s}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) = \boldsymbol{\Sigma}_{t|(\mathbf{x}_0, \mathbf{x}_T)}^{-1}(\boldsymbol{\mu}_{t|(\mathbf{x}_0, \mathbf{x}_T)} - \mathbf{x}) \quad (9)$$

$$\mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) = \mathbf{A}(\mathbf{x} - \mathbf{m}) + \mathbf{c}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) - \frac{1}{2} \boldsymbol{\sigma} \boldsymbol{\sigma}^\top \mathbf{s}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}). \quad (10)$$

where $\mathbf{c}_{t|(\mathbf{x}_0, \mathbf{x}_T)}$ is the bridge control from Theorem 1 and $(\boldsymbol{\mu}_{t|(\mathbf{x}_0, \mathbf{x}_T)}, \boldsymbol{\Sigma}_{t|(\mathbf{x}_0, \mathbf{x}_T)})$ are the mean and covariance of $p_{t|(\mathbf{x}_0, \mathbf{x}_T)}$:

$$\boldsymbol{\Sigma}_{t|(\mathbf{x}_0, \mathbf{x}_T)} = \boldsymbol{\Phi}_t - \boldsymbol{\Phi}_t e^{(T-t)\mathbf{A}^\top} \boldsymbol{\Phi}_T^{-1} e^{(T-t)\mathbf{A}} \boldsymbol{\Phi}_t =: \boldsymbol{\Omega}_t, \quad (11)$$

$$\boldsymbol{\mu}_{t|(\mathbf{x}_0, \mathbf{x}_T)} = \boldsymbol{\mu}_t^{\mathbf{x}_0} + \boldsymbol{\Phi}_t e^{(T-t)\mathbf{A}^\top} \boldsymbol{\Phi}_T^{-1} (\mathbf{x}_T - \boldsymbol{\mu}_T^{\mathbf{x}_0}). \quad (12)$$

In the above, $\boldsymbol{\mu}_t^{\mathbf{x}}$ denotes the mean at time t of an unconditioned mvOU process started from $(0, \mathbf{x})$, and the function $\boldsymbol{\Phi}_t = \int_0^t e^{(t-s)\mathbf{A}} \boldsymbol{\sigma} \boldsymbol{\sigma}^\top e^{(t-s)\mathbf{A}^\top} ds$ is independent of the endpoints $(\mathbf{x}_0, \mathbf{x}_T)$. Consequently, $\boldsymbol{\Sigma}_{t|(\mathbf{x}_0, \mathbf{x}_T)}$ depends only on t , so we write $\boldsymbol{\Omega}_t = \boldsymbol{\Sigma}_{t|(\cdot, \cdot)}$ to make this explicit.

While all our theoretical results are stated in terms of a generic diffusion matrix $\boldsymbol{\sigma}$, we remark that in practice one can always assume $\boldsymbol{\sigma}$ to be diagonal. This is formalised in the following informal lemma, for which we provide a formal statement and proof in Appendix A.1.

Lemma 1. *Up to a orthogonal change of coordinates, any linear SDE with generic drift and diffusion matrix, is equal in its law to another linear SDE with a transformed drift and diagonal diffusion matrix.*

3.2 The Gaussian case

We derive explicit formulae for Gaussian Schrödinger bridges with *general* linear reference dynamics. Our results generalise those in [5] to mvOU reference processes, and we verify in the appendix that we recover several results of [5, Theorem 3, Table 1] as special cases. Both derivations use the disintegration property (2) together with analytical expressions for the optimal coupling and bridges.

Theorem 3 (Characterisation of mvOU-GSB). *Consider a mvOU reference process $\mathbb{Q}$ described by (1) and Gaussian initial and terminal marginals at times $t = 0$ and $t = T$ respectively:*

$$\rho_0 = \mathcal{N}(\mathbf{a}, \mathbf{A}), \rho_T = \mathcal{N}(\mathbf{b}, \mathbf{B}).$$

In what follows, we write Σ_t for the covariance of the unconditioned process (1) started at a point mass. Define the transformed means and covariances

$$\begin{aligned}\bar{\mathbf{a}} &= \Sigma_T^{-1/2} (e^{TA}(\mathbf{a} - \mathbf{m}) + \mathbf{m}), & \bar{\mathbf{A}} &= \Sigma_T^{-1/2} e^{TA} \mathbf{A} e^{TA^\top} \Sigma_T^{-1/2} \\ \bar{\mathbf{b}} &= \Sigma_T^{-1/2} \mathbf{b}, & \bar{\mathbf{B}} &= \Sigma_T^{-1/2} \mathbf{B} \Sigma_T^{-1/2},\end{aligned}$$

and let $\bar{\mathbf{C}}$ be the cross-covariance term of the Gaussian entropic optimal transport plan [21, Theorem 1] between $\bar{\rho}_0 = \mathcal{N}(\bar{\mathbf{a}}, \bar{\mathbf{A}})$ and $\bar{\rho}_T = \mathcal{N}(\bar{\mathbf{b}}, \bar{\mathbf{B}})$ with unit diffusivity, i.e.

$$\bar{\mathbf{C}} = \bar{\mathbf{A}}^{1/2} \left(\bar{\mathbf{A}}^{1/2} \bar{\mathbf{B}} \bar{\mathbf{A}}^{1/2} + \frac{1}{4} \mathbf{I} \right)^{1/2} \bar{\mathbf{A}}^{-1/2} - \frac{1}{2} \mathbf{I}. \quad (13)$$

Setting $\Gamma_t = \Phi_t e^{(T-t)A} \Phi_T^{-1}$ for brevity, define

$$\mathfrak{A}_t = \left(e^{-(T-t)A} - \Gamma_t \right) \Sigma_T^{1/2}, \quad \mathfrak{B}_t = \Gamma_t \Sigma_T^{1/2}, \quad \mathfrak{c}_t = \left(\mathbf{I} - e^{-(T-t)A} \right) \mathbf{m} \quad (14)$$

Then the $\mathbb{Q}$ -GSB is a Markovian Gaussian process for $0 \leq t \leq T$ with mean and covariance

$$\nu_t = \mathbb{E}[Y_t] = \mathfrak{A}_t \bar{\mathbf{a}} + \mathfrak{B}_t \bar{\mathbf{b}} + \mathfrak{c}_t, \quad (15)$$

$$\Xi_{st} = \mathbb{V}[Y_s, Y_t] = \Omega_{st} + \mathfrak{A}_s \bar{\mathbf{A}} \mathfrak{A}_t^\top + \mathfrak{A}_s \bar{\mathbf{C}} \mathfrak{B}_t^\top + \mathfrak{B}_s \bar{\mathbf{C}}^\top \mathfrak{A}_t^\top + \mathfrak{B}_s \bar{\mathbf{B}} \mathfrak{B}_t^\top. \quad (16)$$

Furthermore, the $\mathbb{Q}$ -GSB (SBP-dyn) is described by a SDE

$$dY_t = (\dot{\nu}_t + S_t^\top \Xi_t^{-1} (Y_t - \nu_t)) dt + \sigma dB_t, \quad Y_t \sim \mathcal{N}(\mathbf{a}, \mathbf{A}). \quad (17)$$

where $S_t = (\partial_t \Omega_{t,t'})(t) + \mathfrak{A}_t \dot{\bar{\mathbf{A}}} \mathfrak{A}_t^\top + \mathfrak{A}_t \dot{\bar{\mathbf{C}}} \mathfrak{B}_t^\top + \mathfrak{B}_t \dot{\bar{\mathbf{C}}}^\top \mathfrak{A}_t^\top + \mathfrak{B}_t \dot{\bar{\mathbf{B}}} \mathfrak{B}_t^\top$.

Importantly, we show in the Appendix that we recover from the results of Theorem 3 two prior results of [5] – specifically, we obtain results for a Brownian and scalar OU reference process as special cases. Furthermore, in [5] the authors remark that the quantity $S_t^\top \Xi_t^{-1}$, determining the GSB drift, is *symmetric* although $S_t^\top$ is asymmetric in general. The interpretation of this result is that the GSB is driven by a time-varying potential and is thus of gradient type. By contrast, we observe empirically that the mvOU-GSB drift is *asymmetric* when the underlying reference process $\mathbb{Q}$ is not of gradient type.

3.3 The Non-Gaussian case – score and flow matching

When the endpoints (ρ_0, ρ_T) are non-Gaussian, direct analytical characterisation of the solution to (SBP-dyn) is out of reach. Here we show that combining the exact characterisation of the $\mathbb{Q}$ -bridges (Section 3.1) with score and flow matching (Section 2.2) yields a simulation-free estimator of the Schrödinger bridge when $\mathbb{Q}$ is a mvOU process (1). Specifically, we propose to first exploit the static problem (SBP-static) which can be efficiently solved via Sinkhorn-Knopp iterations [14] with an analytically tractable cost function. Next, we use the property (2) together with the characterisation of $\mathbb{Q}$ -bridges presented in Theorem 2. The approach of [52], which considered a Brownian reference, thus arises as a special case for $\mathbf{A} = 0, \mathbf{m} = 0, \sigma = \sigma \mathbf{I}$.

Proposition 1. When $\mathbb{Q}$ is a mvOU process (1), (SBP-static) shares the same minimiser as

$$\min_{\pi \in \Pi(\rho_0, \rho_T)} \frac{1}{2} \int (\mathbf{x}_T - \boldsymbol{\mu}_T^{\mathbf{x}_0})^\top \Sigma_T^{-1} (\mathbf{x}_T - \boldsymbol{\mu}_T^{\mathbf{x}_0}) d\pi(\mathbf{x}_0, \mathbf{x}_T) + H(\pi | \rho_0 \otimes \rho_T). \quad (18)$$

We remind that $\mathbf{x} \mapsto \boldsymbol{\mu}_T^{\mathbf{x}_0}$ is affine and computation of Σ_T relies on a one-off time integration that does not depend on the endpoints (Section A.1). Equipped with π the solution to (18), we parameterise the unknown probability flow and score function of (SBP-dyn) by $\mathbf{u}_t^\theta(\mathbf{x})$ and $\mathbf{s}_t^\varphi(\mathbf{x})$ where (θ, φ) are trainable parameters. We seek to minimise the loss

$$L(\theta, \varphi) = \mathbb{E}_{\substack{t, (x_0, x_T) \sim U[0, T] \otimes \pi \\ \mathbf{z} \sim p_{t|(x_0, x_T)}}} \left[\|\mathbf{u}_t^\theta(\mathbf{z}) - \mathbf{u}_{t|(x_0, x_T)}(\mathbf{z})\|^2 + \lambda_t \|\mathbf{s}_t^\varphi(\mathbf{z}) - \mathbf{s}_{t|(x_0, x_T)}(\mathbf{z})\|^2 \right], \quad (19)$$

In the practical case when (ρ_0, ρ_T) are discrete, sampling from $t, (x_0, x_T) \sim U[0, T] \otimes \pi$ is trivial once (18) is solved. Sampling from $p_{t|(x_0, x_T)}$ and evaluating $\mathbf{u}_{t|(x_0, x_T)}, \mathbf{s}_{t|(x_0, x_T)}$ amounts to invoking the results of Theorem 2. Again, computations of these quantities involve one-off solution of a 1D matrix-valued integral: this can be solved to a desired accuracy and then cached and queried so it poses negligible computational cost, and this is what we do in practice. As in the Brownian case, the following result establishes the connection of the conditional score and flow matching loss (19) to (SBP-dyn), and is a generalisation and restatement of results from [52, Theorem 3.2, Proposition 3.4].

Algorithm 1 mvOU-OTFM: score and flow matching for mvOU-Schrödinger Bridges

Input: Samples $\{\mathbf{x}_i\}_{i=1}^N, \{\mathbf{x}'_j\}_{j=1}^{N'}$ from source and target distribution at times $t = 0, T$, mvOU reference parameters $(\mathbf{A}, \mathbf{m}, \mathbf{D} = \frac{1}{2}\boldsymbol{\sigma}\boldsymbol{\sigma}^\top)$, batch size B .
Initialise: Probability flow field $\mathbf{u}_t^\theta(\mathbf{x})$, score field $\mathbf{s}_t^\varphi(\mathbf{x})$.
 $\hat{\rho}_0 \leftarrow N^{-1} \sum_{i=1}^N \delta_{\mathbf{x}_i}, \hat{\rho}_1 \leftarrow N'^{-1} \sum_{i=1}^{N'} \delta_{\mathbf{x}'_i}$ *Form empirical marginals*
 $\pi \leftarrow \text{sinkhorn}(\mathbf{C}, \hat{\rho}_0, \hat{\rho}_1, \text{reg} = 1.0)$ *Solve (18) with log-kernel cost*
while not converged **do**
 $\{(\mathbf{x}_i, \mathbf{x}'_i)\}_{i=1}^B \leftarrow \text{sample}(\pi), \{t_i\}_{i=1}^B \leftarrow \text{sample}(U[0, T])$
 for $1 \leq j \leq B$ **do**
 $\mathbf{z} \leftarrow \mathcal{N}(\boldsymbol{\mu}_{t_j | (\mathbf{x}_j, \mathbf{x}'_j)}, \boldsymbol{\Sigma}_{t_j | (\mathbf{x}_j, \mathbf{x}'_j)})$ *Sample from mvOU-bridge (Thm. 2)*
 $\ell_j \leftarrow \|\mathbf{u}_{t_j}^\theta(\mathbf{z}) - \mathbf{u}_{t_j}^\theta(\mathbf{x}_j, \mathbf{x}'_j)\|^2 + \lambda_{t_j} \|\mathbf{s}_{t_j}^\varphi(\mathbf{z}) - \mathbf{s}_{t_j}^\varphi(\mathbf{x}_j, \mathbf{x}'_j)\|^2$ *(Flow, score)-matching loss*
 end for
 $L = B^{-1} \sum_{j=1}^B \ell_j$
 $(\theta, \varphi) \leftarrow \text{Step}(\nabla_\theta L, \nabla_\varphi L)$
end while

Theorem 4 (mvOU-OTFM solves (SBP-dyn)). *If $p_t(\mathbf{x}) > 0$ for all $(t, \mathbf{x})$, (19) shares the same gradients as the unconditional loss: $L(\theta, \varphi) = \mathbb{E}_{t \in U[0, T]} \mathbb{E}_{\mathbf{z} \sim p_t} [\|\mathbf{u}_t^\theta(\mathbf{z}) - \mathbf{u}_t^\theta(\mathbf{z})\|^2 + \lambda_t \|\mathbf{s}_t^\varphi(\mathbf{z}) - \mathbf{s}_t^\varphi(\mathbf{z})\|^2]$. Furthermore, if the coupling π solves (SBP-static) as given in Proposition 1, then for $(\mathbf{u}_t^\theta, \mathbf{s}_t^\varphi)$ achieving the global minimum of (19), the solution to (SBP-dyn) is given by the SDE $d\mathbf{X}_t = (\mathbf{u}_t^\theta(\mathbf{X}_t) + \mathbf{D}\mathbf{s}_t^\varphi(\mathbf{X}_t)) dt + \boldsymbol{\sigma} d\mathbf{B}_t$.*

We stress here that the linearity assumption is only imposed on the reference process $\mathbb{Q}$ in (1), and not on the SBP solution $\mathbb{P}$. In the general non-Gaussian case the mvOU-SBP dynamics will still be nonlinear, much in the same way as for the Brownian SBP.

Further, a non-linear reference dynamics can be connected to linear dynamics of the form (1) by linearisation about a fixed point of the drift. That is, for a smooth reference drift $\mathbf{f}(\mathbf{x})$ with $\mathbf{x}_0$ a fixed point, $\mathbf{f}(\mathbf{x}_0) = \mathbf{0}$, one has $\mathbf{f}(\mathbf{x}) = (\partial_x \mathbf{f})(\mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0) + O(\|\mathbf{x} - \mathbf{x}_0\|^2)$. This naturally motivates the study of mvOU processes with $\mathbf{A} = (\partial_x \mathbf{f})(\mathbf{x}_0)$ and $\mathbf{m} = \mathbf{x}_0$.

4 Results

Gaussian marginals: benchmarking accuracy As a first visual illustration of Theorem 3, we show in Fig. 2 the family of marginals solving (SBP-dyn) for the same pair of Gaussian marginals in dimension $d = 10$, where the reference process $\mathbb{Q}$ is taken to be (a)(i) a mvOU process with high-dimensional rotational drift and (b)(i) a standard Brownian motion. While both provide valid interpolations of ρ_0 to ρ_1 , it is immediately clear that the dynamics are completely different. To provide some further insight, we show in (a)(ii) and (b)(ii) the time-dependent vector field generating each bridge. According to Theorem 3, both processes are time-dependent linear SDEs of the form (17). For the mvOU reference the asymmetric nature of the drift matrix is evident, while for the Brownian reference the symmetry of the drift matrix corresponds to potential driven dynamics [5].

In the Gaussian case, our explicit formulae for the OU-GSB allow us to empirically quantify the accuracy of our neural solver. Since existing computational methods for solving the $\mathbb{Q}$ -Schrödinger bridge for non-gradient $\mathbb{Q}$ are limited, we compare mvOU-OTFM against Iterative Proportional Maximum Likelihood (IPML) [53] and Neural Lagrangian Schrödinger Bridge (NLSB) [25]. The former is based on a forward-backward SDE characterisation of the $\mathbb{Q}$ -SB and alternates between simulation and drift estimation using Gaussian processes, while the latter uses a Neural SDE framework. For d ranging from 2 to 100, we solve (SBP-dyn) with each method and report in Table 1 the average marginal error, measured in the Bures-Wasserstein metric [49] and the average vector field error in L^2 . In all cases, we find that mvOU-OTFM achieves the highest accuracy among all solvers considered. Among the others, NLSB is generally the most accurate followed by IPML, while BM-OTFM performs the poorest since the reference process is misspecified.

In addition to being more accurate, mvOU-OTFM is extremely fast to train owing to being simulation free, taking approximately 1-2 minutes to train on CPU for $d = 50$, regressing the score and flow networks against (19). NLSB took 15+ minutes to train on GPU requiring backpropagation through

SDE solves, while IPFP requires iterative SDE simulation. On the other hand, our approach directly exploits the exact solution to $\mathbb{Q}$ -bridges rather than relying on numerical approximations.




Figure 2: **Gaussian Schrödinger Bridges.** (a) (i) Marginals of the Gaussian Schrödinger bridge ($d = 10$) with mvOU reference (mvOU-GSB) (ii) Time dependent vector field generating the bridge. (b) Same as (a) but for Brownian reference.

d	Marginal error (BW_2^2)					Force error (L^2)			
	mvOU-OTFM	BM-OTFM	IPML ($\leftarrow$)	IPML ($\rightarrow$)	NLSB	mvOU-OTFM	BM-OTFM	IPML	NLSB
2	0.19 ± 0.17	8.40 ± 0.77	5.55 ± 1.53	5.65 ± 1.41	1.21 ± 0.18	3.56 ± 0.25	12.23 ± 0.27	10.31 ± 0.45	7.59 ± 0.33
5	0.23 ± 0.16	9.06 ± 0.66	6.34 ± 5.64	3.24 ± 0.98	1.16 ± 0.26	3.66 ± 0.19	12.27 ± 0.12	10.08 ± 0.75	7.72 ± 0.15
10	0.59 ± 0.36	8.93 ± 0.55	3.00 ± 0.73	3.00 ± 0.63	1.36 ± 0.13	3.82 ± 0.13	12.33 ± 0.13	10.96 ± 1.11	8.15 ± 0.25
25	0.84 ± 0.22	8.81 ± 1.19	6.97 ± 2.01	5.03 ± 0.60	2.97 ± 1.24	4.67 ± 0.16	12.54 ± 0.16	12.42 ± 0.80	11.02 ± 2.41
50	2.21 ± 0.36	11.74 ± 0.37	9.03 ± 0.21	8.32 ± 0.63	6.39 ± 0.13	6.25 ± 0.17	13.26 ± 0.15	14.14 ± 1.00	14.40 ± 0.02
100	6.84 ± 0.78	15.14 ± 0.95	16.19 ± 1.87	14.38 ± 0.38	17.40 ± 0.13	10.45 ± 0.18	15.53 ± 0.14	15.30 ± 0.39	16.49 ± 0.07

Table 1: Marginal error measured in Bures-Wasserstein metric and force error measured in L^2 norm for the Schrödinger bridge learned between two Gaussian measures in varying dimension d .




Figure 3: **Non-Gaussian marginals.** (a)(i) Sampled trajectories from the mvOU-SB learned between non-Gaussian marginals in $d = 10$, shown in (t, x_0) coordinates with score field shown in background. (a)(ii, iii) Sampled SDE (stochastic) and PF-ODE (deterministic) trajectories shown in (x_0, x_1) coordinates, and reference drift shown in background. (b) Same as (a) but for Brownian reference.

Non-Gaussian marginals We demonstrate an application of mvOU-OTFM for general non-Gaussian marginals in Fig. 3, where the reference dynamics are the same as those in Fig. 2. Again, the contrast between the mvOU and Brownian reference dynamics is clear, and we show the distinction between sampling of trajectories via the SDE (3) and PF-ODE formulations (4).

Repressilator dynamics with iterated reference As an application to a system of biophysical importance, we consider the repressilator [16], a model of a synthetic gene circuit exhibiting oscillatory, and hence non-equilibrium, dynamics in cells. This system is composed of three genes in a loop, in which each gene represses the activity of the next (Fig. 4(a)). This model, implemented as a SDE (162), was also studied by [44], which introduced a method, “Schrödinger Bridge with Iterative Reference Refinement” (SBIRR) building upon IPFP [53] to solve a series of Schrödinger bridge problems whilst simultaneously learning an improved reference $\mathbb{Q}$ within a parametric family. In particular, they propose to alternate between solving the $\mathbb{Q}$ -SBP, using IPFP as a subroutine, and fitting of a parametric global drift describing $\mathbb{Q}$. The same approach can be used with mvOU-OTFM, which we present in Algorithm 2 in the Appendix: we propose to apply Algorithm 1 as a subroutine, and solve a regularised linear regression problem to update the mvOU reference parameters (A, m) at each outer step.




Figure 4: **Repressilator dynamics.** (a) Repressilator population snapshots and trajectories. (b) (i) Ground truth vector field. (ii, iii) Inferred multi-marginal SB vector field with (fitted mvOU, Brownian) references. (c) Ground truth linearisation of system and drift A learned by mvOU-OTFM. (d) Leave-one-out error by iteration. (e) Illustration of leave-one-out interpolation between two example timepoints p_{i-1} (blue), p_{i+1} (green) with learned mvOU reference vs. Brownian reference.

Error metric	Iterate 0	Iterate 1	Iterate 2	Iterate 3	Iterate 4	SBIRR (mvOU)	SBIRR (MLP)
EMD	3.38 ± 1.52	2.22 ± 1.12	1.59 ± 0.66	1.49 ± 0.64	1.40 ± 0.57	2.10 ± 0.74	1.67 ± 0.95
Energy distance	1.86 ± 1.06	1.29 ± 0.86	1.03 ± 0.65	0.95 ± 0.58	0.89 ± 0.55	1.39 ± 0.82	1.10 ± 0.86

Table 2: Repressilator leave-one-out interpolation error for mvOU-OTFM and SBIRR.

We use a similar setup to [44] and sample snapshots for $T = 10$ instants from the repressilator system. We show in Fig. 4 the sampled data along with example trajectories. We run Algorithm 2 for 5 iterations and we compare in Fig. 4(b) the ground truth vector field to the SB vector field, both the fitted mvOU reference (at the final iterate of Alg. 2) and for a Brownian reference (as the first iterate of Alg. 2 or equivalently the output of [52]). As was found by [44, 58], iterated fitting of a global autonomous vector field for the reference process leverages multi-timepoint information and allows reconstruction of dynamics better adapted to the underlying system.

We reason that the “best” mvOU process to describe the dynamics should resemble the linearisation of the repressilator system about its fixed point. Comparing the Jacobian of the repressilator system to the fitted drift matrix A in Fig. 4(c), shows a clear resemblance – in particular the cyclic pattern of activation and inhibition is recovered. For a quantitative assessment of performance, we show in Fig. 4(d) and Table 2 the averaged leave-one-out error for a marginal interpolation task: for each $1 < i < T$, Alg. 2 is run on the $T - 1$ remaining timepoints with timepoint t_i held out. The trained model is evaluated on predicting t_i from t_{i-1} . For both the earth mover’s distance (EMD) and energy distance [41] we find that prediction accuracy improves with each additional iteration of Alg. 2.

We compare to SBIRR [44] with two choices of reference process: (i) mvOU and (ii) a general reference family parameterised by a feedforward neural network. The former scenario is thus comparable to mvOU-OTFM, whilst the latter considers a wider class of reference processes. We find that both methods offer an improvement on the null reference (iteration 0 of Alg. 2), and as expected SBIRR with a neural reference family outperforms the mvOU reference family. However, mvOU-OTFM achieves a higher accuracy than either of these methods which suggests that, while SBIRR is able to handle general reference processes and is thus more flexible, mvOU-OTFM trains more accurately. Furthermore, mvOU-OTFM is significantly faster, training in 1-2 minutes on CPU. By comparison, training of SBIRR requires at least several minutes with GPU acceleration.

Dynamics-resolved single cell data We next apply our framework to a single-cell RNA sequencing (scRNA-seq) dataset [4] of the cell cycle, where partial experimental quantification of the vector field is possible via metabolic labelling, i.e. for each observed cell state x_i we also have a velocity estimate $\hat{v}_i$ [40]. We consider $N = 2793$ cells in $d = 30$ PCA dimensions (Fig. 5(a)(i)). For each cell




Figure 5: **Cell cycle scRNA-seq.** (a) Streamlines of (i) transcriptomic vector field calculated from metabolic labelling data (ii) Fitted mvOU reference process, and Schrödinger bridge drift v_{SB} learned by OTFM with (iii) mvOU reference and (iv) Brownian reference. (b) Marginal interpolation error between first and last snapshot as a function of the reference velocity scale parameter, γ .

an experimental reading of its progression along the cell cycle is available via fluorescence, and cells are binned into $T = 5$ snapshots using this coordinate. From the paired state-velocity data $\{x_i, \hat{v}_i\}$ we fit parameters $(\mathbf{A}, \mathbf{m})$ of a mvOU process (1) by ridge regression (Fig. 5(a)(ii)). A limitation of “RNA velocity” is that estimation of vector field direction is often significantly more reliable than magnitude [59]. We introduce therefore an additional “scale” parameter $\gamma > 0$ by which to scale the mvOU drift $\gamma \mathbf{A}$. We reason that this amounts to a choice of matching the disparate and a priori unknown timescales of the snapshots and velocity data.

In order to select an appropriate scale, for $0 \leq \gamma \leq 100$ we use the mvOU-GSB to interpolate between the first and last snapshots (see Fig. 7 for an illustration) and calculate the Bures-Wasserstein distance between the interpolant and each intermediate marginal (Fig. 5(b)). From this it is clear that choosing γ too small or too large leads to a mvOU-GSB that does not fit well the data, while a broad range of approximately $30 \leq \gamma \leq 70$ yields good behaviour. We fix $\gamma = 50$ and use mvOU-OTFM to solve a multi-marginal SBP with reference parameters $(\gamma \mathbf{A}, \mathbf{m})$. From Fig. 5(a)(iii) we see that the expected cyclic behaviour is recovered. On the other hand, using a Brownian reference (equivalently, setting $\gamma = 0$) fails to do so.

Finally, to further demonstrate robustness of our approach to the dimension d , we carry out the interpolation analysis for $d = 50, 100$ and provide numerical results in Table 4. These settings of moderately high dimensionality are typical for single cell studies, where PCA projection to 20-50 dimensions is routinely used prior to any downstream analysis. In all cases, we are able to run mvOU-OTFM within minutes on CPU.

5 Discussion

Motivated by applications to non-equilibrium dynamics, we study the Schrödinger bridge problem where the reference process is taken to be in a family of multivariate Ornstein-Uhlenbeck (mvOU) processes. These processes are able to capture some key behaviours of non-equilibrium systems whilst still being analytically tractable. Leveraging the tractability, we derive an exact solution of the bridge between Gaussian measures for mvOU reference. We extend our approach to non-Gaussian measures using a score and flow matching approach that avoids the need for costly numerical simulation. We showcase the improvement in both accuracy and speed of our approach in a variety of settings, both low and high-dimensional scenarios.

Limitations A key limitation of our method is that our analytical expressions involve matrix inversions and powers of positive definite matrices which scale as $O(d^3)$ with the dimension d . The same limitation was identified in [5] as also applying to their approach. In practice, we find that application to problems with $d \leq 100$ work well. Scaling our methods to even higher dimensional settings is left to future work, and we expect that tools from the Gaussian process literature will be helpful for this. Regarding flow matching, the link between Alg. 1 and (SBP-dyn) requires the solution to (SBP-static) to be computed exactly. In all our experiments, we achieve this by computing couplings using all available samples. In large-scale settings where the number of data points N is very large or possibly streaming, a common approach is to use *minibatch* couplings [50]. In such settings, the use of minibatches results in learned flows being biased. An in-depth discussion of this is provided in [24], which advocates for the infrequent computation of OT couplings on large batch sizes rather than the commonly used per-iteration minibatch OT.

References

- [1] Guy J Abel and Nikola Sander. Quantifying global international migration flows. *Science*, 343(6178):1520–1522, 2014.
- [2] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- [3] Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.
- [4] Nico Battich, Joep Beumer, Buys De Barbanson, Lenno Krenning, Chloé S Baron, Marvin E Tanenbaum, Hans Clevers, and Alexander Van Oudenaarden. Sequencing metabolically labeled transcripts in single cells reveals mrna turnover strategies. *Science*, 367(6482):1151–1156, 2020.
- [5] Charlotte Bunne, Ya-Ping Hsieh, Marco Cuturi, and Andreas Krause. The schrödinger bridge between gaussian measures has a closed form. In *International Conference on Artificial Intelligence and Statistics*, pages 5802–5833. PMLR, 2023.
- [6] Charlotte Bunne, Laetitia Papaxanthos, Andreas Krause, and Marco Cuturi. Proximal optimal transport modeling of population dynamics. In *International Conference on Artificial Intelligence and Statistics*, pages 6511–6528. PMLR, 2022.
- [7] Charlotte Bunne, Geoffrey Schiebinger, Andreas Krause, Aviv Regev, and Marco Cuturi. Optimal transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1):58, 2024.
- [8] Yongxin Chen and Tryphon Georgiou. Stochastic bridges of linear systems. *IEEE Transactions on Automatic Control*, 61(2):526–531, 2015.
- [9] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal steering of a linear stochastic system to a final probability distribution, part i. *IEEE Transactions on Automatic Control*, 61(5):1158–1169, 2015.
- [10] Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. Optimal steering of a linear stochastic system to a final probability distribution, part ii. *IEEE Transactions on Automatic Control*, 61(5):1170–1180, 2016.
- [11] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal transport over a linear dynamical system. *IEEE Transactions on Automatic Control*, 62(5):2137–2152, 2016.
- [12] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrodinger bridge. *Siam Review*, 63(2):249–313, 2021.
- [13] Lénaïc Chizat, Stephen Zhang, Matthieu Heitz, and Geoffrey Schiebinger. Trajectory inference via mean-field langevin in path space. *Advances in Neural Information Processing Systems*, 35:16731–16742, 2022.
- [14] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- [15] Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- [16] Michael B Elowitz and Stanislas Leibler. A synthetic oscillatory network of transcriptional regulators. *Nature*, 403(6767):335–338, 2000.
- [17] Xiaona Fang, Karsten Kruse, Ting Lu, and Jin Wang. Nonequilibrium physics in biology. *Reviews of Modern Physics*, 91(4):045004, 2019.
- [18] Matthieu Gilson, Enzo Tagliazucchi, and Rodrigo Cofré. Entropy production of multivariate ornstein-uhlenbeck processes correlates with consciousness levels in the human brain. *Physical Review E*, 107(2):024121, 2023.

- [19] Federico S Gnesotto, Federica Mura, Jannes Gladrow, and Chase P Broedersz. Broken detailed balance and non-equilibrium dynamics in living systems: a review. *Reports on Progress in Physics*, 81(6):066601, 2018.
- [20] Claude Godrèche and Jean-Marc Luck. Characterising the nonequilibrium stationary states of ornstein–uhlenbeck processes. *Journal of Physics A: Mathematical and Theoretical*, 52(3):035002, 2018.
- [21] Hicham Janati. *Advances in Optimal transport and applications to neuroscience*. PhD thesis, Institut Polytechnique de Paris, 2021.
- [22] Hicham Janati, Boris Muzellec, Gabriel Peyré, and Marco Cuturi. Entropic optimal transport between unbalanced gaussian measures has a closed form. *Advances in neural information processing systems*, 33:10468–10479, 2020.
- [23] Donald E Kirk. *Optimal control theory: an introduction*. Courier Corporation, 2004.
- [24] Michal Klein, Alireza Mousavi-Hosseini, Stephen Zhang, and Marco Cuturi. On fitting flow models with large sinkhorn couplings. *arXiv preprint arXiv:2506.05526*, 2025.
- [25] Takeshi Koshizuka and Issei Sato. Neural lagrangian schrödinger bridge: Diffusion modeling for population dynamics. *arXiv preprint arXiv:2204.04853*, 2022.
- [26] Hugo Lavenant, Stephen Zhang, Young-Heon Kim, and Geoffrey Schiebinger. Towards a mathematical theory of trajectory inference. *arXiv preprint arXiv:2102.09204*, 2021.
- [27] Melvin Lax. Fluctuations from the nonequilibrium steady state. *Reviews of modern physics*, 32(1):25, 1960.
- [28] Christian Léonard. A survey of the schrodinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.
- [29] Anders Lindquist. On feedback control of linear stochastic systems. *SIAM Journal on Control*, 11(2):323–343, 1973.
- [30] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [31] Suryanarayana Maddu, Victor Chardès, Michael Shelley, et al. Inferring biological processes with intrinsic noise from cross-sectional data. *arXiv preprint arXiv:2410.07501*, 2024.
- [32] Anton Mallasto, Augusto Gerolin, and Hà Quang Minh. Entropy-regularized 2-wasserstein distance between gaussian measures. *Information Geometry*, 5(1):289–323, 2022.
- [33] Yuhang Mei, Mohammad Al-Jarrah, Amirhossein Taghvaei, and Yongxin Chen. Flow matching for stochastic linear control systems. *arXiv preprint arXiv:2412.00617*, 2024.
- [34] Christopher E Miles, Scott A McKinley, Fangyuan Ding, and Richard B Lehoucq. Inferring stochastic rates from heterogeneous snapshots of particle positions. *Bulletin of mathematical biology*, 86(6):74, 2024.
- [35] Boris Muzellec, Richard Nock, Giorgio Patrini, and Frank Nielsen. Tsallis regularized optimal transport and ecological inference. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- [36] Ingram Olkin and Friedrich Pukelsheim. The distance between two random vectors with given dispersion matrices. *Linear Algebra and its Applications*, 48:257–263, 1982.
- [37] Juan MR Parrondo, Christian Van den Broeck, and Ryoichi Kawai. Entropy production and the arrow of time. *New Journal of Physics*, 11(7):073008, 2009.
- [38] Katarina Petrović, Lazar Atanackovic, Kacper Kapusniak, Michael M Bronstein, Joey Bose, and Alexander Tong. Curly flow matching for learning non-gradient field dynamics. In *ICLR 2025 Workshop on Machine Learning for Genomics Explorations*, 2025.

- [39] Aram-Alexandre Pooladian and Jonathan Niles-Weed. Plug-in estimation of schrödinger bridges. *arXiv preprint arXiv:2408.11686*, 2024.
- [40] Xiaojie Qiu, Yan Zhang, Jorge D Martin-Rufino, Chen Weng, Shayan Hosseinzadeh, Dian Yang, Angela N Pogson, Marco Y Hein, Kyung Hoi Joseph Min, Li Wang, et al. Mapping transcriptomic vector fields of single cells. *Cell*, 185(4):690–711, 2022.
- [41] Maria L Rizzo and Gábor J Székely. Energy distance. *wiley interdisciplinary reviews: Computational statistics*, 8(1):27–38, 2016.
- [42] Geoffrey Schiebinger. Reconstructing developmental landscapes and trajectories from single-cell data. *Current Opinion in Systems Biology*, 27:100351, 2021.
- [43] Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- [44] Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal schrödinger bridges with iterative reference refinement. *arXiv preprint arXiv:2408.06277*, 2024.
- [45] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *arXiv preprint arXiv:2303.16852*, 2023.
- [46] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [47] Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. In *Uncertainty in artificial intelligence*, pages 574–584. PMLR, 2020.
- [48] Andrew M Stuart and Marie-Therese Wolfram. Inverse optimal transport. *SIAM Journal on Applied Mathematics*, 80(1):599–619, 2020.
- [49] Asuka Takatsu. Wasserstein geometry of gaussian measures. 2011.
- [50] Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.
- [51] Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectory-net: A dynamic optimal transport network for modeling cellular dynamics. In *International conference on machine learning*, pages 9526–9536. PMLR, 2020.
- [52] Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. *arXiv preprint arXiv:2307.03672*, 2023.
- [53] Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.
- [54] Pat Vatiwutipong and Nattakorn Phewchean. Alternative way to derive the distribution of the multivariate ornstein–uhlenbeck process. *Advances in Difference Equations*, 2019:1–7, 2019.
- [55] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [56] Caleb Weinreb, Samuel Wolock, Betsabeh K Tusi, Merav Socolovsky, and Allon M Klein. Fundamental limits on dynamic inference from single-cell snapshots. *Proceedings of the National Academy of Sciences*, 115(10):E2467–E2476, 2018.
- [57] Grace Hui Ting Yeo, Sachit D Saksena, and David K Gifford. Generative modeling of single-cell time series with prescient enables prediction of cell trajectories with interventions. *Nature communications*, 12(1):3222, 2021.

- [58] Stephen Y Zhang. Joint trajectory and network inference via reference fitting. In *Machine Learning in Computational Biology*, pages 72–85. PMLR, 2024.
- [59] Shijie C Zheng, Genevieve Stein-O’Brien, Leandros Boukas, Loyal A Goff, and Kasper D Hansen. Pumping the brakes on rna velocity by understanding and interpreting rna velocity estimates. *Genome biology*, 24(1):246, 2023.
- [60] Ligang Zhu, Songlin Yang, Kun Zhang, Hong Wang, Xiaona Fang, and Jin Wang. Uncovering underlying physical principles and driving forces of cell differentiation and reprogramming from single-cell transcriptomics. *Proceedings of the National Academy of Sciences*, 121(34):e2401540121, 2024.

A Theoretical results

Quantity	Description	Introduced in
$\mathbf{A} \in \mathbb{R}^{d \times d}, \mathbf{m} \in \mathbb{R}^d$	Multivariate Ornstein-Uhlenbeck drift parameters	Eq. (1)
$\mathbf{D} = \frac{1}{2} \boldsymbol{\sigma} \boldsymbol{\sigma}^\top$	Multivariate Ornstein-Uhlenbeck diffusivity	Eq. (1)
$\mathbb{P}$	Schrödinger bridge path measure	Section 2.1
$\mathbb{Q}$	Reference process path measure	Section 2.1
$\boldsymbol{\mu}_t^{x_0}, \boldsymbol{\Sigma}_t$	Unconditional mean & cov. of mvOU process started at x_0	Section 3.1
$p_t (x_0, x_T)$	Conditional density of mvOU bridge	Section 3.1
$c_t (x_0, x_T)$	mvOU bridge control between (x_0, x_T)	Section 3.1
$\boldsymbol{\mu}_t (x_0, x_T), \boldsymbol{\Sigma}_t (x_0, x_T) = \boldsymbol{\Omega}_t$	Conditional mean and covariance of mvOU bridge	Section 3.1
$\boldsymbol{\Omega}_{st}$	Conditional covariance process of mvOU bridge	Section 3.1
$\mathbf{u}_t (x_0, x_T), \mathbf{s}_t (x_0, x_T)$	Conditional probability flow and score field of mvOU bridge	Section 3.1
$\mathcal{N}(\mathbf{a}, \mathbf{A}), \mathcal{N}(\mathbf{b}, \mathbf{B})$	Initial and terminal mvOU-GSB marginals	Section 3.2
$\mathcal{N}(\bar{\mathbf{a}}, \bar{\mathbf{A}}), \mathcal{N}(\bar{\mathbf{b}}, \bar{\mathbf{B}})$	Transformed mvOU-GSB marginals	Section 3.2
$\bar{\mathbf{C}}$	Cross-covariance of entropic transport plan	Section 3.2
$\mathfrak{A}_t, \mathfrak{B}_t, \mathbf{c}_t$	Key quantities for the mvOU-GSB	Section 3.2
ν_t	mvOU-GSB mean process	Section 3.2
Ξ_t, Ξ_{st}	mvOU-GSB variance and covariance process	Section 3.2
$\mathbf{S}_t^\top \Xi_t^{-1}$	SDE drift matrix of mvOU-GSB	Section 3.2.

Table 3: Glossary of some key notations and quantities used in the statements of our theoretical results.

A.1 Some calculations on multivariate Ornstein-Uhlenbeck processes

For convenience, we first collect some results about multivariate Ornstein-Uhlenbeck processes of the form (1), a detailed discussion of mvOU processes can be found in e.g. [54]. For a *time-invariant* process with drift matrix $\mathbf{A}$ and diffusion $\boldsymbol{\sigma}$, we have that

$$(X_t | X_0 = x_0) \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t),$$

where

$$\boldsymbol{\mu}_t = e^{t\mathbf{A}}x_0, \quad \boldsymbol{\Sigma}_t = \int_0^t e^{(t-\tau)\mathbf{A}}\boldsymbol{\sigma}\boldsymbol{\sigma}^\top e^{(t-\tau)\mathbf{A}^\top} d\tau. \quad (20)$$

If $X_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and writing $X_{0t} = X_0 \oplus X_t$, some tedious but straightforward applications of conditional expectations and the tower property reveal that

$$X_{0t} \sim \mathcal{N}(\boldsymbol{\mu}_{0t}, \boldsymbol{\Sigma}_{0t}), \quad \boldsymbol{\mu}_{0t} = \boldsymbol{\mu}_0 \oplus e^{t\mathbf{A}}\boldsymbol{\mu}_0, \quad \boldsymbol{\Sigma}_{0t} = \begin{bmatrix} \boldsymbol{\Sigma}_0 & \boldsymbol{\Sigma}_0 e^{t\mathbf{A}^\top} \\ e^{t\mathbf{A}}\boldsymbol{\Sigma}_0 & \boldsymbol{\Sigma}_t + e^{t\mathbf{A}}\boldsymbol{\Sigma}_0 e^{t\mathbf{A}^\top} \end{bmatrix}. \quad (21)$$

More generally, for a process $dX_t = (\mathbf{A}X_t + \mathbf{b})dt + \boldsymbol{\sigma}dB_t$, one has:

$$\boldsymbol{\mu}_t = e^{t\mathbf{A}}x_0 + (e^{t\mathbf{A}} - \mathbf{I})(\mathbf{A}^{-1}\mathbf{b}), \quad \boldsymbol{\Sigma}_t = \int_0^t e^{(t-s)\mathbf{A}}\boldsymbol{\sigma}\boldsymbol{\sigma}^\top e^{(t-s)\mathbf{A}^\top} ds. \quad (22)$$

The expressions for the covariance of the mvOU process can be obtained from the fact that the covariance evolves following a *Lyapunov equation*. For the case of constant coefficients, we state the following result.

Lyapunov equation solution It is easy to verify that $\dot{\mathbf{G}}_t = \mathbf{A}\mathbf{G}_t + \mathbf{G}_t\mathbf{A}^\top + \mathbf{Q}_t$ has solution

$$\mathbf{G}_t = \int_0^t e^{(t-s)\mathbf{A}}\mathbf{Q}_s e^{(t-s)\mathbf{A}^\top} ds.$$

A.2 Bridges of multivariate Ornstein-Uhlenbeck processes

This problem has been studied by Chen et al. in [8], however the material therein is geared towards a control audience and considers a more general case where all coefficients are time-dependent. We will re-derive the results that we will need for processes of the form (1). Additionally, while in practice we typically use $\boldsymbol{\sigma} = \sigma\mathbf{I}$, we state some results in the general setting of non-isotropic noise.

Derivation of the dynamical OU-bridge (Theorem 1) Consider a mvOU process

$$d\mathbf{X}_t = (\mathbf{A}\mathbf{X}_t + \mathbf{b}) dt + \boldsymbol{\sigma} d\mathbf{B}_t. \quad (23)$$

Now form the *controlled* version pinned at $(0, \mathbf{x}_0)$ and $(T, \mathbf{x}_T)$, where $\mathbf{u}_t(\mathbf{X}_t)$ is an additional force arising from the conditioning of the process at an endpoint:

$$d\mathbf{X}_t = (\mathbf{A}\mathbf{X}_t + \mathbf{b} + \mathbf{u}_t(\mathbf{X}_t)) dt + \boldsymbol{\sigma} d\mathbf{B}_t. \quad (24)$$

The main result that we need is that $\mathbf{u}_t(\mathbf{x})$ is itself a linear time-dependent field whose coefficients are independent of $(\mathbf{x}_0, \mathbf{x}_T)$. In the above we consider the simplest case of classical optimal control, where the control is *directly* added to the system drift and there is no unobserved system.

Consider the system started at $(0, \mathbf{x}_0)$ and pinned to $(T, \mathbf{x}_T)$. For now, we make all statements for time-dependent coefficients and the diffusion is non-isotropic. We then form the Lagrangian:

$$\min_{\mathbf{u}} \mathbb{E} \left\{ \int_0^T \|\mathbf{u}_t\|^2 dt + \sup_{\lambda} \lambda \|\mathbf{X}_T - \mathbf{x}_T\|^2 \right\} = \sup_{\lambda} \min_{\mathbf{u}} \mathbb{E} \left\{ \int_0^T \|\mathbf{u}_t\|^2 dt + \lambda \|\mathbf{X}_T - \mathbf{x}_T\|^2 \right\}$$

Fixing some $\lambda > 0$, the resulting problem is known as a *linear-quadratic-Gaussian* problem [29] in the control literature:

$$\min_{\mathbf{u}} \mathbb{E} \left\{ \frac{1}{2} \int_0^T \|\mathbf{u}_t\|^2 dt + \frac{\lambda}{2} \|\mathbf{X}_T - \mathbf{x}_T\|^2 \right\}, \quad (25)$$

for which we can write a HJB equation:

$$0 = \partial_t V_t(\mathbf{x}) + \min_{\mathbf{u}} \left\{ \tilde{\mathcal{A}}_t[V_t](\mathbf{x}, \mathbf{u}_t(\mathbf{x})) + \frac{1}{2} \|\mathbf{u}_t(\mathbf{x})\|^2 \right\}, \quad (26)$$

subject to the terminal boundary condition $V_T(\mathbf{x}) = \frac{\lambda}{2} \|\mathbf{x} - \mathbf{x}_T\|^2$, and the operators $\tilde{\mathcal{A}}, \mathcal{A}$ are defined as

$$\tilde{\mathcal{A}}_t[f] = \langle \mathbf{A}_t \mathbf{x} + \mathbf{b}_t + \mathbf{u}_t, \nabla_x f \rangle + \frac{1}{2} \nabla_x \cdot (\boldsymbol{\sigma}_t \boldsymbol{\sigma}_t^\top \nabla_x f), \quad (27)$$

$$\mathcal{A}_t[f] = \langle \mathbf{A}_t \mathbf{x} + \mathbf{b}_t, \nabla_x f \rangle + \frac{1}{2} \nabla_x \cdot (\boldsymbol{\sigma}_t \boldsymbol{\sigma}_t^\top \nabla_x f), \quad (28)$$

these are the generators of the controlled (24) and uncontrolled (23) SDEs respectively. Substituting all these in, we find that the HJBE is

$$0 = \partial_t V_t(\mathbf{x}) + \mathcal{A}_t[V_t](\mathbf{x}) + \min_{\mathbf{u}} \left[\langle \mathbf{u}_t, \nabla_x V_t \rangle + \frac{1}{2} \|\mathbf{u}_t\|^2 \right]. \quad (29)$$

It follows from the first order condition on the “inner” problem that $\mathbf{u}_t = -\nabla_x V_t$, so

$$0 = \partial_t V_t(\mathbf{x}) + \mathcal{A}_t[V_t](\mathbf{x}) - \frac{1}{2} \|\nabla_x V_t(\mathbf{x})\|^2. \quad (30)$$

Use an ansatz that the value function is quadratic:

$$V_t(\mathbf{x}) = \frac{1}{2} \mathbf{x}^\top \mathbf{M}_t \mathbf{x} + \mathbf{c}_t^\top \mathbf{x} + d_t \implies \nabla_x V_t(\mathbf{x}) = \mathbf{M}_t \mathbf{x} + \mathbf{c}_t. \quad (31)$$

Then:

$$\mathcal{A}_t[V_t] = \mathbf{x}^\top (\mathbf{A}_t^\top \mathbf{M}_t) \mathbf{x} + \langle \mathbf{A}_t^\top \mathbf{c}_t + \mathbf{M}_t \mathbf{b}_t, \mathbf{x} \rangle + \langle \mathbf{b}_t, \mathbf{c}_t \rangle + \frac{1}{2} \text{tr}(\boldsymbol{\sigma}_t \boldsymbol{\sigma}_t^\top \mathbf{M}_t). \quad (32)$$

Now note that the quadratic form only depends on the *symmetric* part of the matrix, i.e.

$$\mathbf{x}^\top \mathbf{A}_t^\top \mathbf{M}_t \mathbf{x} = \frac{1}{2} \mathbf{x}^\top (\mathbf{A}_t^\top \mathbf{M}_t + \mathbf{M}_t \mathbf{A}_t) \mathbf{x}.$$

The HJBE is therefore

$$\begin{aligned} 0 = & \frac{1}{2} \mathbf{x}^\top \dot{\mathbf{M}}_t \mathbf{x} + \dot{\mathbf{c}}_t^\top \mathbf{x} + \dot{d}_t \\ & + \frac{1}{2} \mathbf{x}^\top (\mathbf{A}_t^\top \mathbf{M}_t + \mathbf{M}_t \mathbf{A}_t) \mathbf{x} + (\mathbf{A}_t^\top \mathbf{c}_t + \mathbf{M}_t^\top \mathbf{b}_t)^\top \mathbf{x} + \mathbf{b}_t^\top \mathbf{c}_t + \frac{1}{2} \text{tr}(\boldsymbol{\sigma}_t \boldsymbol{\sigma}_t^\top \mathbf{M}_t) \\ & - \frac{1}{2} (\mathbf{M}_t \mathbf{x} + \mathbf{c}_t)^\top (\mathbf{M}_t \mathbf{x} + \mathbf{c}_t) \end{aligned} \quad (33)$$

So

$$0 = \dot{M}_t + A_t^\top M_t + M_t A_t - M_t^2, \quad (34)$$

$$0 = \dot{c}_t + A_t^\top c_t + M_t b_t - M_t c_t. \quad (35)$$

with the boundary condition $M_T = \lambda I$ and $c_t = -\lambda x_T$.

Let $\Lambda_t = M_t^{-1}$, so that $\dot{\Lambda}_t = -M_t^{-1} \dot{M}_t M_t^{-1} = -\Lambda_t \dot{M}_t \Lambda_t$. Substituting, we find that

$$\dot{\Lambda}_t = \Lambda_t A_t^\top + A_t \Lambda_t - I. \quad (36)$$

Actually we can rewrite the value function in a different way, which is more convenient for us:

$$V_t(x) = \frac{1}{2}(x - k_t)^\top M_t(x - k_t) + \text{const}$$

in which case $c_t = -M_t k_t$, and so we have $k_T = x_T$. The corresponding ODE for k_t is

$$\dot{k}_t = A_t k_t + b_t. \quad (37)$$

Let us now specialise to the case of time-invariant coefficients, which allows us to write down explicit expressions for the solutions:

$$\dot{k}_t = A k_t + b \implies k_t = e^{-(T-t)A}(x_T + A^{-1}b) - A^{-1}b. \quad (38)$$

Rewriting the SDE drift to have the form $Ax + b = A(x - m)$ we have that $b = -Am \implies A^{-1}b = -m$, when A is nonsingular. Then:

$$k_t = e^{-(T-t)A}(x_T - m) + m. \quad (39)$$

Now we deal with the quadratic term Λ_t . Let $G_\tau = \Lambda_{T-\tau}$, so that $\partial_\tau G_\tau = -\dot{\Lambda}_{T-\tau}$. Then G_τ satisfies

$$\partial_\tau G_\tau = -G_\tau A^\top - A G_\tau + I, \quad G_0 = 0. \quad (40)$$

So

$$G_\tau = \int_0^\tau e^{-(\tau-s)A} e^{-(\tau-s)A^\top} ds. \quad (41)$$

Substituting back, we find that

$$\Lambda_t = \int_0^{T-t} e^{-(T-t-s)A} e^{-(T-t-s)A^\top} ds = \int_0^{T-t} e^{-sA} e^{-sA^\top} ds. \quad (42)$$

The bridge control is therefore

$$u_{t|(x_0, x_T)} = -\Lambda_t^{-1}(x - k_t) \quad (43)$$

For the general case of non-constant coefficients, we express them as solutions of ODEs with terminal boundary conditions.

$$\dot{k}_t = A_t k_t + b_t, \quad k_T = x_T, \quad (44)$$

$$\dot{\Lambda}_t = \Lambda_t A_t^\top + A_t \Lambda_t - I, \quad \Lambda_T = 0. \quad (45)$$

In practice, both of these equations can be approximately solved by numerical integration in time.

Finally we remark here that the diffusion component does not play a role in the control, which is classical.

“Static” Ornstein-Uhlenbeck bridge statistics (Theorem 2) As was studied by [8] for a more general scenario of time-varying processes, the mvOU process and its conditioned versions are Gaussian processes. Under these assumptions, (X_s, X_t) has joint mean

$$\mu_s \oplus \mu_t, \quad (46)$$

and covariance

$$\begin{bmatrix} \Phi_s & \Phi_s e^{(t-s)A^\top} \\ e^{(t-s)A} \Phi_s & \Phi_t \end{bmatrix}. \quad (47)$$

where Φ_t is the covariance of the process started from a point mass at time t :

$$\Phi_t = \int_0^t e^{(t-s)A} \sigma \sigma^\top e^{(t-s)A^\top} ds, \quad \Phi_t = \Phi_{t-s} + e^{(t-s)A} \Phi_s e^{(t-s)A^\top}. \quad (48)$$

By taking the Schur complement, we have that $X_s | X_t$ is distributed with variance

$$\mathbb{V}[X_s | X_t] = \Phi_s - \Phi_s e^{(t-s)A^\top} \Phi_t^{-1} e^{(t-s)A} \Phi_s \quad (49)$$

and mean

$$\mathbb{E}[X_s | X_t] = \mu_s + \Phi_s e^{(t-s)A^\top} \Phi_t^{-1} (X_t - \mu_t). \quad (50)$$

More generally, we can consider the three-point correlation (X_r, X_s, X_t) . Then the mean is $\mu_r \oplus \mu_s \oplus \mu_t$ and the covariance is

$$\begin{bmatrix} \Phi_r & \Phi_r e^{(s-r)A^\top} & \Phi_r e^{(t-r)A^\top} \\ \cdot & \Phi_s & \Phi_s e^{(t-s)A^\top} \\ \cdot & \cdot & \Phi_t \end{bmatrix} \quad (51)$$

Using the Schur complement again, we have that $(X_r, X_s) | X_t$ has covariance

$$\mathbb{V}[X_r, X_s | X_t] = \Phi_r e^{(s-r)A^\top} - \Phi_r e^{(t-r)A^\top} \Phi_t^{-1} e^{(t-s)A} \Phi_s. \quad (52)$$

For the case of time-varying coefficients, we need to introduce the state transition matrix Ψ_{ts} [29] which describes the deterministic aspects of evolution between $s < t$ under $(A_t)_t$. That is, for a dynamics $\dot{x}_t = A_t x_t$ one has $x_t = \Psi_{t,s} x_s$ and $\Psi(t, s) = \Psi(t, r) \Psi(r, s)$ for $t \geq r \geq s$. In this case there is not an explicit expression for Φ_t . Instead, let $(\Phi_t)_t$ be the unconditional variance evolutions for such a process started at $t = 0$, obtained by solving

$$\dot{\Phi}_t = A_t \Phi_t + \Phi_t A_t^\top + \sigma_t \sigma_t^\top, \quad \Phi_0 = 0. \quad (53)$$

Then the more general result from [8] for the covariance of (X_r, X_s, X_r) is

$$\begin{bmatrix} \Phi_r & \Phi_r \Psi_{sr}^\top & \Phi_r \Psi_{tr}^\top \\ \Psi_{sr} \Phi_r & \Phi_s & \Phi_s \Psi_{ts}^\top \\ \Psi_{tr} \Phi_r & \Psi_{ts} \Phi_s & \Phi_t \end{bmatrix} \quad (54)$$

From this result, the same Schur complement computation yields expressions for the conditional covariances and mean:

$$\mathbb{V}[X_r, X_s | X_t] = \Phi_r \Psi_{sr}^\top - \Phi_r \Psi_{tr}^\top \Phi_t^{-1} \Psi_{ts} \Phi_s, \quad (55)$$

$$\mathbb{V}[X_s | X_t] = \Phi_s - \Phi_s \Psi_{ts}^\top \Phi_t^{-1} \Psi_{ts} \Phi_s, \quad (56)$$

$$\mathbb{E}[X_s | X_t] = \mu_s + \Phi_s \Psi_{ts}^\top \Phi_t^{-1} (X_t - \mu_t) \quad (57)$$

Formal statement and proof of Lemma 1

Lemma (Formal statement of Lemma 1). *Let $dX_t = AX_t dt + \sigma dB_t$ be a linear SDE with generic drift $A \in \mathbb{R}^{d \times d}$ and diffusion $\sigma \in \mathbb{R}^{d \times d}$. There exists an orthogonal matrix $U \in \mathbb{R}^{d \times d}$ such that $Y_t = U^\top X_t$ obeys a linear SDE with transformed drift matrix $U^\top A U$ and a diagonal diffusion matrix $\Lambda = \text{diag}(\lambda)$, where $\lambda = (\lambda_1, \dots, \lambda_d) \geq 0$.*

Proof. Let $D = \frac{1}{2} \sigma \sigma^\top$. This is positive semidefinite and therefore can its spectral decomposition can be written as $D = \frac{1}{2} U \Lambda^2 U^\top$ where U is orthogonal and $\Lambda = \text{diag}(\lambda) \geq 0$. Consider the SDEs

$$dX_t = AX_t dt + \sigma dB_t \quad (a.)$$

and

$$dX_t = AX_t dt + U \Lambda U^\top dB_t. \quad (b.)$$

Let $f(x)$ be a twice differentiable test function. SDE (a.) has generator $(\mathcal{L}f)(x) = \langle \partial_x f(x), Ax \rangle + \frac{1}{2} \text{tr}(\sigma^\top \partial_{xx} f(x) \sigma)$. Computing the generator for SDE (b.), $(\tilde{\mathcal{L}}f)(x) = \langle \partial_x f(x), Ax \rangle + \frac{1}{2} \text{tr}(U \Lambda U^\top \partial_{xx} f(x) U \Lambda U^\top)$. Note that $\text{tr}(U \Lambda U^\top \partial_{xx} f(x) U \Lambda U^\top) = \text{tr}(\partial_{xx} f(x) U \Lambda^2 U^\top) = \text{tr}(\sigma^\top \partial_{xx} f(x) \sigma)$. Hence both SDEs share the same generator, and so are equal in law.

Since orthogonal transformations of Brownian increments are also Brownian increments, we have that $U^\top dB_t = d\tilde{B}_t$. Substituting into (b.), we have $dX_t = AX_t dt + U \Lambda d\tilde{B}_t$. Then,

$$d(U^\top X_t) = U^\top A X_t dt + \Lambda d\tilde{B}_t \implies dY_t = U^\top A U Y_t dt + \Lambda d\tilde{B}_t.$$

□

A.3 Simulation-free Schrödinger bridges with linear reference dynamics

Having characterised the solution of the SBP for general reference processes and now that we have derived the score and flow for the Ornstein-Uhlenbeck bridge, the path is clear towards a simulation-free scheme for learning Schrödinger bridges where the reference dynamics are given by a *linear* SDE.

For solution of the static SBP problem (18) (equivalently, (SBP-static)), the transition kernel of the reference dynamics is given by

$$\begin{aligned} \frac{dQ_{0t}}{d\mathbf{x} \otimes d\mathbf{x}}(\mathbf{x}_0, \mathbf{x}_t) &= \frac{1}{Z_t} \exp \left(-\frac{1}{2}(\mathbf{x}_t - \boldsymbol{\mu}_t^{\mathbf{x}_0})^\top \boldsymbol{\Sigma}_t^{-1}(\mathbf{x}_t - \boldsymbol{\mu}_t^{\mathbf{x}_0}) \right) \frac{dQ_0}{d\mathbf{x}}(\mathbf{x}_0) \\ \implies -\log \left(\frac{dQ_{0t}}{d\mathbf{x} \otimes d\mathbf{x}} \right) &= \frac{1}{2}(\mathbf{x}_t - \boldsymbol{\mu}_t^{\mathbf{x}_0})^\top \boldsymbol{\Sigma}_t^{-1}(\mathbf{x}_t - \boldsymbol{\mu}_t^{\mathbf{x}_0}) + \log Z_t - \log \left(\frac{dQ_0}{d\mathbf{x}}(\mathbf{x}_0) \right), \end{aligned} \quad (58)$$

where $\boldsymbol{\mu}_t^{\mathbf{x}_0}$ denotes the mean at time t conditional on $(0, \mathbf{x}_0)$. Notably, the last two terms depend only on $\mathbf{x}_0$. It is a classical result (and very easy to show) that these kinds of terms do not affect the minimiser of (18) and so they are immaterial. The cost function to use is thus effectively

$$C(\mathbf{x}_0, \mathbf{x}_t) = \frac{1}{2}(\mathbf{x}_t - \boldsymbol{\mu}_t^{\mathbf{x}_0})^\top \boldsymbol{\Sigma}_t^{-1}(\mathbf{x}_t - \boldsymbol{\mu}_t^{\mathbf{x}_0}). \quad (59)$$

Once a solution π of the static SBP is on hand, we want to utilise the stochastic regression objective on the conditional score and flow. Sampling $(\mathbf{x}_0, \mathbf{x}_t)$ from π , recall that $\mathbf{P}^{(\mathbf{x}_0, \mathbf{x}_1)} = \mathbf{Q}^{(\mathbf{x}_0, \mathbf{x}_1)}$ so we want to sample from the $\mathbf{Q}$ -bridge using (11), (12):

$$t \sim U[0, 1], \quad \mathbf{x}_t^{(\mathbf{x}_0, \mathbf{x}_1)} \sim \mathcal{N}(\boldsymbol{\mu}_{t|(\mathbf{x}_0, \mathbf{x}_1)}, \boldsymbol{\Sigma}_{t|(\mathbf{x}_0, \mathbf{x}_1)}). \quad (60)$$

Equations (9) and (12) give us the score and flow at $\mathbf{x}_t$.

A.4 Characterisation of the $\mathbb{Q}$ -GSB (Theorem 3)

We will construct the $\mathbb{Q}$ -GSB utilising its characterisation (2). Our approach is similar to that of [5] – we first obtain explicit formulae for the static $\mathbb{Q}$ -GSB (SBP-static) and then build towards the dynamical $\mathbb{Q}$ -GSB (SBP-dyn) by using the characterisation of $\mathbb{Q}$ -bridges.

Static $\mathbb{Q}$ -GSB The standard Gaussian EOT problem has a well known solution [32, 21, 5, 22]. Here we use the notations of [21] and define

$$\text{OT}_{\sigma^2}^{\otimes}(\alpha, \beta) := \min_{\pi \in \Pi(\alpha, \beta)} \frac{1}{2} \int \|\mathbf{x} - \mathbf{y}\|_2^2 d\pi(\mathbf{x}, \mathbf{y}) + \sigma^2 H(\pi | \alpha \otimes \beta), \quad (61)$$

where $\sigma > 0$ is the regularisation level and $\alpha = \mathcal{N}(\mathbf{a}, \mathbf{A})$ and $\beta = \mathcal{N}(\mathbf{b}, \mathbf{B})$. The solution to (61) in the Gaussian case is given by [21, Theorem 1]

$$\text{OT}_{\sigma^2}^{\otimes}(\alpha, \beta) = \frac{1}{2} \|\mathbf{a} - \mathbf{b}\|_2^2 + \frac{1}{2} \text{B}_{\sigma^2}^{\otimes}(\mathbf{A}, \mathbf{B}) \quad (62)$$

$$\text{B}_{\sigma^2}^{\otimes}(\mathbf{A}, \mathbf{B}) = \text{tr}(\mathbf{A}) + \text{tr}(\mathbf{B}) - 2\text{tr}(\mathbf{C}) + \sigma^2 \log \det \left(\frac{1}{\sigma^2} \mathbf{C} + \mathbf{I} \right), \quad (63)$$

$$\mathbf{C} = \mathbf{A}^{1/2} \left(\mathbf{A}^{1/2} \mathbf{B} \mathbf{A}^{1/2} + \frac{\sigma^4}{4} \mathbf{I} \right)^{1/2} \mathbf{A}^{-1/2} - \frac{\sigma^2}{2} \mathbf{I}. \quad (64)$$

We are, of course, interested in a slightly different problem, namely, for $\rho_0 = \mathcal{N}(\mathbf{a}, \mathbf{A})$ and $\rho_T = \mathcal{N}(\mathbf{b}, \mathbf{B})$, we seek to solve $\min_{\pi \in \Pi(\rho_0, \rho_T)} \text{KL}(\pi | \mathbb{Q}_{0,T})$. Expanding the definition of KL and

simplifying, we get

$$(\text{SBP-static}) = \min_{\pi \in \Pi(\rho_0, \rho_T)} \int d\pi \log \left(\frac{d\pi}{d\mathbb{Q}_{0,T}} \right) \quad (65)$$

$$= \min_{\pi \in \Pi(\rho_0, \rho_T)} - \int d\pi \log \left(\frac{d\mathbb{Q}_{0,T}}{dx_0 \otimes dx_T} \right) + H(\pi|\rho_0 \otimes \rho_T) + H(\rho_0|dx_0) + H(\rho_T|dx_T) \quad (66)$$

$$\simeq \min_{\pi \in \Pi(\rho_0, \rho_T)} - \int d\pi \log \left(\frac{d\mathbb{Q}_{0,T}}{dx_0 \otimes dx_T} \right) + H(\pi|\rho_0 \otimes \rho_T) \quad (67)$$

$$\simeq \min_{\pi \in \Pi(\rho_0, \rho_T)} - \int d\pi(x_0, x_T) \log \left(\frac{d\mathbb{Q}_{T|0}(x_T|x_0)}{dx_T} \right) + H(\pi|\rho_0 \otimes \rho_T) \quad (68)$$

where by $\simeq$ we denote equality of the objective up to additive constants which do not affect the minimiser π . We note that the last line is exactly (61) with $\alpha, \beta = \rho_0, \rho_T$ and $\sigma = 1$ and a modified cost. Further, recognise $d\mathbb{Q}_{T|0}(x_T|x_0)/dx_T$ as the $\mathbb{Q}$ -transition density, and $\mathbb{Q}_{T|0} = \mathcal{N}(\boldsymbol{\mu}_T^{\mathbf{x}_0}, \boldsymbol{\Sigma}_T)$. Thus,

$$(\text{SBP-static}) \simeq \min_{\pi \in \Pi(\rho_0, \rho_T)} \frac{1}{2} \int (\mathbf{x}_T - \boldsymbol{\mu}_T^{\mathbf{x}_0})^\top \boldsymbol{\Sigma}_T^{-1} (\mathbf{x}_T - \boldsymbol{\mu}_T^{\mathbf{x}_0}) d\pi(\mathbf{x}_0, \mathbf{x}_T) + H(\pi|\rho_0 \otimes \rho_T) \quad (69)$$

where we remind that $\boldsymbol{\mu}_t^{\mathbf{x}_0}$ is the mean at time t conditional on starting at $(0, \mathbf{x}_0)$, and $\boldsymbol{\Sigma}_t$ is the covariance at time t , started at a point mass, i.e. $\boldsymbol{\Sigma}_0 = 0$.

In particular, we have $\boldsymbol{\Sigma}_t = \boldsymbol{\Phi}_t$, using the notation of Theorem 2. Recall that in the case of constant coefficients, we have the following explicit expressions:

$$\boldsymbol{\mu}_t^{\mathbf{x}_0} = e^{t\mathbf{A}}(\mathbf{x}_0 - \mathbf{m}) + \mathbf{m}, \quad (70)$$

$$\boldsymbol{\Phi}_t = \int_0^t e^{(t-s)\mathbf{A}} \boldsymbol{\sigma} \boldsymbol{\sigma}^\top e^{(t-s)\mathbf{A}^\top} ds. \quad (71)$$

Define the following changes of variable $(\mathbf{x}_0, \mathbf{x}_T) \mapsto (\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_T)$, injective whenever $\boldsymbol{\Sigma}_T$ is positive definite:

$$\bar{\mathbf{x}}_0 = T_0(\mathbf{x}_0) = \boldsymbol{\Sigma}_T^{-1/2} \boldsymbol{\mu}_T^{\mathbf{x}_0},$$

$$\bar{\mathbf{x}}_T = T_T(\mathbf{x}_T) = \boldsymbol{\Sigma}_T^{-1/2} \mathbf{x}_T$$

Let $\bar{\rho}_0 = (T_0)_\# \rho_0$ and $\bar{\rho}_T = (T_T)_\# \rho_T$, and $\bar{\pi} = (T_0 \times T_T)_\# \pi$ and note that the relative entropy is invariant to this change of coordinates. So the problem (69) is equivalent to

$$\min_{\bar{\pi} \in \Pi(\bar{\rho}_0, \bar{\rho}_T)} \frac{1}{2} \int \|\bar{\mathbf{x}}_T - \bar{\mathbf{x}}_0\|_2^2 d\bar{\pi}(\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_T) + H(\bar{\pi}|\bar{\rho}_0 \otimes \bar{\rho}_T). \quad (72)$$

This is exactly $\text{OT}_{\sigma^2=1}^\otimes(\bar{\rho}_0, \bar{\rho}_T)$ as per (61) with $\bar{\rho}_0 = \mathcal{N}(\bar{\mathbf{a}}, \bar{\mathbf{A}})$ and $\bar{\rho}_T = \mathcal{N}(\bar{\mathbf{b}}, \bar{\mathbf{B}})$. The transformed means and covariances are

$$\bar{\mathbf{a}} = \boldsymbol{\Sigma}_T^{-1/2} (e^{T\mathbf{A}}(\mathbf{a} - \mathbf{m}) + \mathbf{m})$$

$$\bar{\mathbf{A}} = \boldsymbol{\Sigma}_T^{-1/2} e^{T\mathbf{A}} \mathbf{A} e^{T\mathbf{A}^\top} \boldsymbol{\Sigma}_T^{-1/2}$$

$$\bar{\mathbf{b}} = \boldsymbol{\Sigma}_T^{-1/2} \mathbf{b}$$

$$\bar{\mathbf{B}} = \boldsymbol{\Sigma}_T^{-1/2} \mathbf{B} \boldsymbol{\Sigma}_T^{-1/2}$$

We remark that for time-dependent coefficients our approach still applies, however computation of the transformation coefficients will be in terms of the general state transition matrix instead of matrix exponentials. In what follows, we focus on results for constant coefficients for simplicity and their practical relevance.

The solution to (72) is therefore

$$(\bar{\mathbf{x}}_0, \bar{\mathbf{x}}_T) \sim \bar{\pi} = \mathcal{N} \left(\begin{bmatrix} \bar{\mathbf{a}} \\ \bar{\mathbf{b}} \end{bmatrix}, \begin{bmatrix} \bar{\mathbf{A}} & \bar{\mathbf{C}} \\ \bar{\mathbf{C}}^\top & \bar{\mathbf{B}} \end{bmatrix} \right),$$

where $\bar{\mathbf{C}}$ is the cross-covariance the transport plan, given by (64).

Dynamic $\mathbb{Q}$ -GSB : marginals and covariance Let $X_t|(x_0, x_T)$ denote the $\mathbb{Q}$ -bridge pinned at (x_0, x_T) . Then by Theorem 2,

$$X_t|(x_0, x_T) = \mathcal{N}(\mu_{t|(x_0, x_T)}, \Sigma_{t|(x_0, x_T)}) = \mathcal{N}(\mu_{t|(x_0, x_T)}, \Omega_t).$$

Invoking the inverse mappings T_0^{-1}, T_T^{-1} and substituting the expression for $\mu_{t|(x_0, x_T)}$ from 2, we get

$$\mu_{t|(T_0^{-1}(\bar{x}_0), T_T^{-1}(\bar{x}_T))} = \underbrace{(e^{-(T-t)A} - \Gamma_t) \Sigma_T^{1/2} \bar{x}_0}_{\mathfrak{A}_t} + \underbrace{\Gamma_t \Sigma_T^{1/2} \bar{x}_T}_{\mathfrak{B}_t} + \underbrace{(I - e^{-(T-t)A}) m}_{\mathfrak{C}_t} \quad (73)$$

$$= \mathfrak{A}_t \bar{x}_0 + \mathfrak{B}_t \bar{x}_T + \mathfrak{C}_t. \quad (74)$$

where $\Gamma_t = \Phi_t e^{(T-t)A^\top} \Phi_T^{-1}$. For clarity, we state first a general formula:

Lemma 2. Let $(X_0, X_1) \sim \mathcal{N}\left(\begin{bmatrix} \mu_0 \\ \mu_1 \end{bmatrix}, \begin{bmatrix} \Sigma_{00} & \Sigma_{01} \\ \Sigma_{10} & \Sigma_{11} \end{bmatrix}\right)$. Let

$$Y_{X_0, X_1} \sim \mathcal{N}(AX_0 + BX_1 + c, \Omega).$$

Then, $\mathbb{E}[Y] = A\mu_0 + B\mu_1 + c$ and $\mathbb{V}[Y] = \Omega + A\Sigma_{00}A^\top + B\Sigma_{11}B^\top + A\Sigma_{01}B^\top + B\Sigma_{01}A^\top$.

Application to (74) gives us the following formula for the variance of the bridge at time t :

$$\mathbb{V}[X_t] = \Omega_t + \mathfrak{A}_t \bar{\mathcal{A}} \mathfrak{A}_t^\top + \mathfrak{B}_t \bar{\mathcal{B}} \mathfrak{B}_t^\top + \mathfrak{A}_t \bar{\mathcal{C}} \mathfrak{B}_t^\top + \mathfrak{B}_t \bar{\mathcal{C}}^\top \mathfrak{A}_t^\top =: \Xi_t. \quad (75)$$

Substituting and simplifying, we get the following expressions for each of the terms:

$$\mathfrak{A}_t \bar{\mathcal{A}} \mathfrak{A}_t^\top = (I - \Gamma_t e^{(T-t)A}) e^{tA} \mathcal{A} e^{tA^\top} (I - \Gamma_t e^{(T-t)A})^\top, \quad (76)$$

$$\mathfrak{B}_t \bar{\mathcal{B}} \mathfrak{B}_t^\top = \Gamma_t \mathcal{B} \Gamma_t^\top, \quad (77)$$

$$\mathfrak{A}_t \bar{\mathcal{C}} \mathfrak{B}_t^\top = (e^{-(T-t)A} - \Gamma_t) \Sigma_T^{1/2} \bar{\mathcal{C}} \Sigma_T^{1/2} \Gamma_t^\top. \quad (78)$$

Similarly, the mean can be computed as

$$\mathbb{E}[X_t] = \mathfrak{A}_t \bar{a} + \mathfrak{B}_t \bar{b} + \mathfrak{C}_t \quad (79)$$

$$= \left(e^{-(T-t)A} - \Gamma_t\right) e^{tA} (a - m) + \Gamma_t (b - m) + m =: \nu_t. \quad (80)$$

To calculate the covariance of the SB process, we use again the disintegration property (2). We have from (52) for the $\mathbb{Q}$ -bridges and $0 < s < t < T$:

$$\mathbb{V}[X_s, X_t | X_0, X_T] = \Phi_s e^{(t-s)A^\top} - \Phi_s e^{(T-s)A^\top} \Phi_T^{-1} e^{(T-t)A} \Phi_t =: \Omega_{s,t}. \quad (81)$$

Now,

$$\mathbb{V}[X_s, X_t] = \mathbb{E}[(X_s - \nu_s)(X_t - \nu_t)^\top] \quad (82)$$

$$= \mathbb{E}[X_s X_t^\top] - \nu_s \nu_t^\top \quad (83)$$

$$= \mathbb{E}[\mathbb{E}[X_s X_t^\top | X_0, X_T]] - \nu_s \nu_t^\top. \quad (84)$$

and

$$\mathbb{E}[X_s X_t^\top | X_0, X_T] = \mathbb{V}[X_s, X_t | X_0, X_T] + \mu_{s|(X_0, X_T)} \mu_{t|(X_0, X_T)}^\top, \quad (85)$$

in which the first (variance) term doesn't actually depend on (X_0, X_T) . So,

$$\mathbb{E}[X_s X_t^\top] = \Omega_{st} + \mathbb{E}_{(X_0, X_T)}[\mu_{s|(X_0, X_T)} \mu_{t|(X_0, X_T)}^\top]. \quad (86)$$

In fact let's switch to the "mapped" coordinates $\bar{X}_0, \bar{X}_T$ for the endpoints. We abuse notation and omit the inverse map in what follows. Then, from previously,

$$\mu_{t|(\bar{X}_0, \bar{X}_T)} = \mathfrak{A}_t \bar{X}_0 + \mathfrak{B}_t \bar{X}_T + \mathfrak{C}_t. \quad (87)$$

Expanding, collecting and cancelling terms, we have that

$$\mathbb{E}[\mu_{s|(\bar{X}_0, \bar{X}_T)} \mu_{t|(\bar{X}_0, \bar{X}_T)}^\top] - \nu_s \nu_t^\top \quad (88)$$

$$= \mathfrak{A}_s \mathbb{V}[\bar{X}_0] \mathfrak{A}_t^\top + \mathfrak{A}_s \mathbb{V}[\bar{X}_0, \bar{X}_T] \mathfrak{B}_t^\top + \mathfrak{B}_s \mathbb{V}[\bar{X}_T, \bar{X}_0] \mathfrak{A}_t^\top + \mathfrak{B}_s \mathbb{V}[\bar{X}_T] \mathfrak{B}_t^\top \quad (89)$$

$$= \mathfrak{A}_s \bar{\mathcal{A}} \mathfrak{A}_t^\top + \mathfrak{A}_s \bar{\mathcal{C}} \mathfrak{B}_t^\top + \mathfrak{B}_s \bar{\mathcal{C}}^\top \mathfrak{A}_t^\top + \mathfrak{B}_s \bar{\mathcal{B}} \mathfrak{B}_t^\top. \quad (90)$$

Putting everything together, we get

$$\mathbb{V}[X_s, X_t] = \Omega_{st} + \mathfrak{A}_s \bar{\mathcal{A}} \mathfrak{A}_t^\top + \mathfrak{A}_s \bar{\mathcal{C}} \mathfrak{B}_t^\top + \mathfrak{B}_s \bar{\mathcal{C}}^\top \mathfrak{A}_t^\top + \mathfrak{B}_s \bar{\mathcal{B}} \mathfrak{B}_t^\top =: \Xi_{s,t}. \quad (91)$$

Dynamic Q-GSB : SDE representation We proceed via the generator route also used by [5]. Expanding the covariance (91), we find that

$$\mathbb{V}[\mathbf{X}_t, \mathbf{X}_{t+h}] = \Xi_{t,t+h} = \Omega_{t,t+h} + \mathfrak{A}_t \bar{\mathbf{A}} \mathfrak{A}_{t+h}^\top + \mathfrak{A}_t \bar{\mathbf{C}} \mathfrak{B}_{t+h}^\top + \mathfrak{B}_t \bar{\mathbf{C}}^\top \mathfrak{A}_{t+h}^\top + \mathfrak{B}_t \bar{\mathbf{B}} \mathfrak{B}_{t+h}^\top \quad (92)$$

$$= \Xi_t + h \left\{ (\partial_{t'} \Omega_{t,t'})(t) + \mathfrak{A}_t \bar{\mathbf{A}} \dot{\mathfrak{A}}_t^\top + \mathfrak{A}_t \bar{\mathbf{C}} \dot{\mathfrak{B}}_t^\top + \mathfrak{B}_t \bar{\mathbf{C}}^\top \dot{\mathfrak{A}}_t^\top + \mathfrak{B}_t \bar{\mathbf{B}} \dot{\mathfrak{B}}_t^\top \right\} + \dots \quad (93)$$

$$= \Xi_t + h \mathbf{S}_t + \dots \quad (94)$$

where we have set the key quantity

$$\mathbf{S}_t = (\partial_{t'} \Omega_{t,t'})(t) + \mathfrak{A}_t \bar{\mathbf{A}} \dot{\mathfrak{A}}_t^\top + \mathfrak{A}_t \bar{\mathbf{C}} \dot{\mathfrak{B}}_t^\top + \mathfrak{B}_t \bar{\mathbf{C}}^\top \dot{\mathfrak{A}}_t^\top + \mathfrak{B}_t \bar{\mathbf{B}} \dot{\mathfrak{B}}_t^\top \quad (95)$$

Now, $(\mathbf{X}_{t+h}, \mathbf{X}_t)$ are jointly Gaussian with mean and covariance

$$\begin{bmatrix} \boldsymbol{\nu}_{t+h} \\ \boldsymbol{\nu}_t \end{bmatrix}, \begin{bmatrix} \Xi_{t+h} & \Xi_{t+h,t} \\ \Xi_{t,t+h} & \Xi_t \end{bmatrix}. \quad (96)$$

So

$$\mathbb{E}[\mathbf{X}_{t+h} | \mathbf{X}_t] = \boldsymbol{\nu}_{t+h} + \Xi_{t,t+h}^\top \Xi_t^{-1} (\mathbf{X}_t - \boldsymbol{\nu}_t) \quad (97)$$

$$= \mathbf{X}_t + h [\dot{\boldsymbol{\nu}}_t + \mathbf{S}_t^\top \Xi_t^{-1} (\mathbf{X}_t - \boldsymbol{\nu}_t)] \quad (98)$$

$$\mathbb{V}[\mathbf{X}_{t+h} | \mathbf{X}_t] = \Xi_{t+h} - \Xi_{t,t+h}^\top \Xi_t^{-1} \Xi_{t,t+h} \quad (99)$$

$$= h(\dot{\Sigma}_t - (\mathbf{S}_t + \mathbf{S}_t^\top)). \quad (100)$$

Let $\mathbf{x} \mapsto u(\mathbf{x})$ be a twice differentiable test function. Then:

$$\mathbb{E}[u(\mathbf{X}_{t+h}) | \mathbf{X}_t = \mathbf{x}] \quad (101)$$

$$= \mathbb{E}[u(\mathbf{X}_t) + (\partial_x u(\mathbf{X}_t))^\top (\mathbf{X}_{t+h} - \mathbf{X}_t) +$$

$$\frac{1}{2} (\mathbf{X}_{t+h} - \mathbf{X}_t)^\top (\partial_{xx}^2 u(\mathbf{X}_t)) (\mathbf{X}_{t+h} - \mathbf{X}_t) \Big| \mathbf{X}_t = \mathbf{x}] \quad (102)$$

$$= u(\mathbf{x}) + h(\partial_x u(\mathbf{x}))^\top (\dot{\boldsymbol{\nu}}_t + \mathbf{S}_t^\top \Xi_t^{-1} (\mathbf{x} - \boldsymbol{\nu}_t)) +$$

$$\frac{1}{2} \mathbb{E} \left[(\mathbf{X}_{t+h} - \mathbf{X}_t)^\top (\partial_{xx}^2 u(\mathbf{X}_t)) (\mathbf{X}_{t+h} - \mathbf{X}_t) \Big| \mathbf{X}_t = \mathbf{x} \right] \quad (103)$$

$$= u(\mathbf{x}) + h(\partial_x u(\mathbf{x}))^\top (\dot{\boldsymbol{\nu}}_t + \mathbf{S}_t^\top \Xi_t^{-1} (\mathbf{x} - \boldsymbol{\nu}_t)) + \frac{h}{2} \text{tr}[(\partial_{xx}^2 u(\mathbf{x}))(\dot{\Sigma}_t - (\mathbf{S}_t + \mathbf{S}_t^\top))] + \dots \quad (104)$$

Subtracting $u(\mathbf{x})$ and taking the limit $h \rightarrow 0$, it is clear from the definition of the generator that the SDE drift is

$$\dot{\boldsymbol{\nu}}_t + \mathbf{S}_t^\top \Xi_t^{-1} (\mathbf{x} - \boldsymbol{\nu}_t). \quad (105)$$

This takes the same form as the equation found in [5], however our formulae for $\mathbf{S}_t$ allow us to apply it to any linear reference SDE, not necessarily ones with scalar drift. Additionally, we empirically verify that for asymmetric $\mathbf{A}$ the matrix $\mathbf{S}_t$ is asymmetric. This contrasts with the symmetric nature of the drift for the gradient-type setting.

Now we want to work out what each of these terms are in practice. Effectively we need to compute $\dot{\mathfrak{A}}_t$, $\dot{\mathfrak{B}}_t$ and $(\partial_{t'} \Omega_{t,t'})(t)$:

$$\dot{\mathfrak{A}}_t = (\mathbf{A} e^{-(T-t)\mathbf{A}} - \dot{\Gamma}_t) \Sigma_T^{1/2}, \quad (106)$$

$$\dot{\mathfrak{B}}_t = \dot{\Gamma}_t \Sigma_T^{1/2}, \quad (107)$$

$$\dot{\Gamma}_t = \left(e^{t\mathbf{A}} \boldsymbol{\sigma} \boldsymbol{\sigma}^\top e^{T\mathbf{A}^\top} - \Phi_t \mathbf{A}^\top e^{(T-t)\mathbf{A}^\top} \right) \Phi_T^{-1}, \quad (108)$$

$$(\partial_{t'} \Omega_{t,t'})(t) = \Phi_t \left[\mathbf{A}^\top - e^{(T-t)\mathbf{A}^\top} \Phi_T^{-1} e^{(T-t)\mathbf{A}} \left(-\mathbf{A} \Phi_t + e^{t\mathbf{A}} \boldsymbol{\sigma} \boldsymbol{\sigma}^\top e^{t\mathbf{A}^\top} \right) \right]. \quad (109)$$

Remark on the case of time-varying coefficients In this case we can still assume without loss of generality that we still work on the time interval $[0, T]$, by shifting the time coordinate if necessary. Let $\Psi(t, s)$ be the state transition matrix associated with $\mathbf{A}_t$ and Φ_t be the solution at time t to (53). Then we still have $\Sigma_T = \Phi_T$ for the transition kernel covariance in the cost, i.e. the covariance started from a point mass. For the mean of the reference process started from $\mathbf{x}_0$, we have the generalised formula

$$\mu_t(\mathbf{x}_0) = \Psi(t, 0)\mathbf{x}_0 - \int_0^t \Psi(t, s)\mathbf{A}_s\mathbf{m}_s \, ds, \quad (110)$$

and we remark that its inverse μ_t^{-1} always exists since Ψ is never singular:

$$\mu_t^{-1}(\mathbf{y}) = \Psi(t, 0)^{-1} \left(\mathbf{y} + \int_0^t \Psi(t, s)\mathbf{A}_s\mathbf{m}_s \, ds \right). \quad (111)$$

Then

$$\bar{\mathbf{a}} = \Sigma_T^{-1/2} \mu_T^{\mathbf{a}}, \quad (112)$$

$$\bar{\mathcal{A}} = \Sigma_T^{-1/2} \Psi(t, 0) \mathcal{A} \Psi(t, 0)^\top \Sigma_T^{-1/2}, \quad (113)$$

$$\bar{\mathbf{b}} = \Sigma_T^{-1/2} \mathbf{b}, \quad (114)$$

$$\bar{\mathcal{B}} = \Sigma_T^{-1/2} \mathcal{B} \Sigma_T^{-1/2}. \quad (115)$$

Recall that the general expression for the $\mathbb{Q}$ -bridge conditioned on $(0, \mathbf{x}_0), (T, \mathbf{x}_T)$ is

$$\mu_{t|(x_0, x_1)} = \mu_t(\mathbf{x}_0) + \Phi_t \Psi(T, t)^\top \Phi_T^{-1} (\mathbf{x}_T - \mu_T(\mathbf{x}_0)). \quad (116)$$

Letting $\mathbf{x}_0 = \mu_T^{-1}(\Sigma_T^{1/2} \bar{\mathbf{x}}_0)$ and $\mathbf{x}_T = \Sigma_T^{1/2} \bar{\mathbf{x}}_T$ and substituting the expression for μ_t^{-1} we have

$$\begin{aligned} \mu_{t|(T_0^{-1}(\bar{\mathbf{x}}_0), T_T^{-1}(\bar{\mathbf{x}}_T))} &= (\Psi(t, T) - \Phi_t \Psi(T, t)^\top \Phi_T^{-1}) \Sigma_T^{1/2} \bar{\mathbf{x}}_0 \\ &\quad + \Phi_t \Psi(T, t)^\top \Phi_T^{-1} \Sigma_T^{1/2} \bar{\mathbf{x}}_T \\ &\quad + \int_t^T \Psi(t, s) \mathbf{A}_s \mathbf{m}_s \, ds. \end{aligned} \quad (117)$$

Set $\Gamma_t = \Phi_t \Psi(T, t)^\top \Phi_T^{-1}$, then

$$\mathfrak{A}_t = (\Psi(t, T) - \Gamma_t) \Sigma_T^{1/2}, \quad (118)$$

$$\mathfrak{B}_t = \Gamma_t \Sigma_T^{1/2}, \quad (119)$$

$$\mathfrak{c}_t = \int_t^T \Psi(t, s) \mathbf{A}_s \mathbf{m}_s \, ds. \quad (120)$$

Case: Brownian motion Let's check this against the results of Bunne and Hsieh [5], who consider a class of reference processes

$$d\mathbf{Y}_t = (c_t \mathbf{Y}_t + \alpha_t) dt + g_t d\mathbf{B}_t.$$

This corresponds to a special case of ours, where the drift is scalar-valued. Setting $g_t = \omega, \alpha_t = c_t = 0$, their results give us the marginal parameters of the GSB connecting $\mathcal{N}(\mu_0, \Sigma_0)$ and $\mathcal{N}(\mu_1, \Sigma_1)$. In what follows, using their results and notations of [5, Table 1], we have that $r_t = t, \bar{r}_t = 1 - t, \zeta = 0, \kappa(t, t') = \omega^2 t, \rho_t = t$. Then:

$$\mu_t = \bar{r}_t \mu_0 + r_t \mu_1 + \zeta(t) - r_t \zeta(1), \quad (121)$$

$$= (1 - t) \mu_0 + t \mu_1, \quad (122)$$

$$\Sigma_t = \bar{r}_t^2 \Sigma_0 + r_t^2 \Sigma_1 + r_t \bar{r}_t (C_\sigma + C_\sigma^\top) + \kappa(t, t') (1 - \rho_t) \mathbf{I} \quad (123)$$

$$= (1 - t)^2 \Sigma_0 + t^2 \Sigma_1 + t(1 - t)(C_\sigma + C_\sigma^\top) + \omega^2 t(1 - t) \mathbf{I}. \quad (124)$$

Let us compute the same quantities using our formulae. We set $\mathbf{A} = 0, \mathbf{m} = 0$ and $T = 1$. Then, $\Phi_t = t\omega^2 \mathbf{I}, \Gamma_t = t \mathbf{I}, \Sigma_T^{1/2} = \omega \mathbf{I}$. Then,

$$\bar{\mathcal{A}} = \omega^{-2} \mathcal{A}, \quad \bar{\mathcal{B}} = \omega^{-2} \mathcal{B}.$$

Also from the underbraced expressions, we have

$$\mathfrak{A}_t = \omega(1-t)\mathbf{I}, \quad \mathfrak{B}_t = \omega t\mathbf{I}.$$

Assembling our results, we get

$$\begin{aligned} \mathbb{E}[\mathbf{X}_t] &= (\mathbf{I} - \Gamma_t)\mathbf{a} + \Gamma_t\mathbf{b} \\ &= (1-t)\mathbf{a} + t\mathbf{b} \\ \mathbb{V}[\mathbf{X}_t] &= \Omega_t + \mathfrak{A}_t\overline{\mathbf{A}}\mathfrak{A}_t^\top + \mathfrak{B}_t\overline{\mathbf{B}}\mathfrak{B}_t^\top + \mathfrak{A}_t\overline{\mathbf{C}}\mathfrak{B}_t^\top + \mathfrak{B}_t\overline{\mathbf{C}}\mathfrak{A}_t^\top \\ &= \Omega_t + (1-t)^2\mathbf{A} + t^2\mathbf{B} + \omega^2t(1-t)(\overline{\mathbf{C}} + \overline{\mathbf{C}}^\top). \end{aligned}$$

Now $\overline{\mathbf{C}}$ is the transport plan obtained from the *scaled* input Gaussians, of variance $\overline{\mathbf{A}} = \mathbf{A}/\omega^2$, $\overline{\mathbf{B}} = \mathbf{B}/\omega^2$ (in the case of Brownian reference process, there is no linear mapping via the flow map). For the scaled measures, EOT is calculated with $\varepsilon = \sigma^2 = 1$. It stands to reason that once we scale everything back, we should get $\overline{\mathbf{C}} = \mathbf{C}/\omega^2$. To verify this, recall that for the untransformed problem

$$\arg \min_{\pi \in \Pi(\alpha, \beta)} \frac{1}{2} \int \frac{1}{\omega^2} \|x - y\|^2 d\pi + H(\pi | \alpha \otimes \beta) = \arg \min_{\pi \in \Pi(\alpha, \beta)} \frac{1}{2} \int \|x - y\|^2 d\pi + \omega^2 H(\pi | \alpha \otimes \beta).$$

Then for the RHS problem, we have that

$$\begin{aligned} \mathbf{C} &= \mathbf{A}^{1/2}(\mathbf{A}^{1/2}\mathbf{B}\mathbf{A}^{1/2} - \frac{\omega^4}{4}\mathbf{I})^{1/2}\mathbf{A}^{-1/2} - \frac{\omega^2}{2}\mathbf{I} \\ &= \omega^2(\overline{\mathbf{A}}^{1/2}(\overline{\mathbf{A}}^{1/2}\overline{\mathbf{B}}\overline{\mathbf{A}}^{1/2} + \frac{1}{4}\mathbf{I})^{1/2}\overline{\mathbf{A}}^{-1/2} - \frac{1}{2}\mathbf{I}) \\ &= \omega^2\overline{\mathbf{C}}. \end{aligned}$$

Also it is easy to verify that $\Omega_t = \Sigma_{t|(x_0, x_1)} = \omega^2t(1-t)\mathbf{I}$. So,

$$\mathbb{V}[\mathbf{X}_t] = \omega^2t(1-t)\mathbf{I} + (1-t)^2\mathbf{A} + t^2\mathbf{B} + t(1-t)(\mathbf{C} + \mathbf{C}^\top), \quad (125)$$

where $\mathbf{C}$ is the EOT plan between the unscaled measures $\mathcal{N}(\mathbf{a}, \mathbf{A})$, $\mathcal{N}(\mathbf{b}, \mathbf{B})$ with $\varepsilon = \omega^2$. This is exactly the same result as the one derived in [5].

Case: Centered scalar OU process $\mathbf{A} = -\lambda\mathbf{I}$, $\mathbf{m} = \mathbf{0}$ From [5, Table 1], we have that $\bar{r}_t = \sinh(\lambda t)/\sinh(\lambda)$, $\bar{r}_t = \sinh(\lambda t) \coth(\lambda t) - \sinh(\lambda t) \coth(\lambda)$, $\zeta = 0$. The mean is then

$$\boldsymbol{\mu}_t = \bar{r}_t\boldsymbol{\mu}_0 + r_t\boldsymbol{\mu}_1. \quad (126)$$

With $\kappa(t, t') = \omega^2 e^{-\lambda t} \sinh(\lambda t)/\lambda$ and $\rho_t = e^{-\lambda(1-t)} \sinh(\lambda t)/\sinh(\lambda)$, the variance is

$$\Sigma_t = \bar{r}_t^2 \Sigma_0 + r_t \Sigma_1 + r_t \bar{r}_t (\mathbf{C}_{\sigma^*} + \mathbf{C}_{\sigma^*}^\top) + \kappa(t, t)(1 - \rho_t)\mathbf{I}, \quad (127)$$

where $\sigma^{*2} = \omega^2 \sinh(\lambda)/\lambda$. For convenience, we will check things term by term and using Mathematica.

Check first the mean. We have:

$$\Phi_t = \omega^2 \left(\frac{1 - e^{-2\lambda t}}{2\lambda} \right) \mathbf{I}, \quad \Gamma_t = e^{-(1-t)\lambda} \left(\frac{1 - e^{-2\lambda t}}{1 - e^{-2\lambda}} \right) \mathbf{I} \quad (128)$$

Then, plugging into our formula,

$$\mathbb{E}[\mathbf{X}_t] = \underbrace{e^{-\lambda}(e^{(1-t)\lambda}\mathbf{I} - \Gamma_t)}_{=\bar{r}_t} \mathbf{a} + \underbrace{\Gamma_t}_{=r_t} \mathbf{b}. \quad (129)$$

Now the variance. We have that

$$\Sigma_1 = \omega^2 \left(\frac{e^{-\lambda} \sinh(\lambda)}{\lambda} \right) \mathbf{I}, \quad (130)$$

$$\overline{\mathbf{A}} = \frac{\lambda e^{-\lambda}}{\omega^2 \sinh(\lambda)} \mathbf{A} = \alpha \mathbf{A}, \quad (131)$$

$$\overline{\mathbf{B}} = \frac{\lambda}{\omega^2 e^{-\lambda} \sinh(\lambda)} \mathbf{B} = \beta \mathbf{B}, \quad (132)$$

$$\mathfrak{A}_t = \left(e^{(1-t)\lambda} \mathbf{I} - \Gamma_t \right) \Sigma_1^{1/2}, \quad (133)$$

$$\mathfrak{B}_t = \Gamma_t \Sigma_1^{1/2}. \quad (134)$$

From here it is straightforward to verify that

$$\mathfrak{A}_t \overline{\mathcal{A}}_t^\top = \bar{r}_t^2 \mathcal{A}, \quad \mathfrak{B}_t \overline{\mathcal{B}}_t^\top = r_t^2 \mathcal{B}. \quad (135)$$

Now note that

$$\bar{\mathcal{C}} = \overline{\mathcal{A}}^{1/2} \left(\overline{\mathcal{A}}^{1/2} \overline{\mathcal{B}} \overline{\mathcal{A}}^{1/2} + \frac{1}{4} \mathbf{I} \right)^{1/2} \overline{\mathcal{A}}^{-1/2} - \frac{1}{2} \mathbf{I} \quad (136)$$

$$= \sqrt{\alpha\beta} \left[\mathcal{A}^{1/2} \left(\mathcal{A}^{1/2} \mathcal{B} \mathcal{A}^{1/2} + \frac{(\alpha\beta)^{-1}}{4} \mathbf{I} \right)^{1/2} \mathcal{A}^{-1/2} - \frac{(\alpha\beta)^{-1/2}}{2} \mathbf{I} \right] \quad (137)$$

$$= \sqrt{\alpha\beta} \mathcal{C}_{1/\sqrt{\alpha\beta}}, \quad (138)$$

where $\mathcal{C}_{1/\sqrt{\alpha\beta}}$ denotes the EOT covariance for entropic regularisation level $1/\sqrt{\alpha\beta}$. Then:

$$\mathfrak{A}_t \bar{\mathcal{C}} \mathfrak{B}_t^\top + \mathfrak{B}_t \bar{\mathcal{C}}^\top \mathfrak{A}_t^\top = \underbrace{(e^{\lambda(1-t)} \mathbf{I} - \Gamma_t) \Gamma_t \Sigma_1 \sqrt{\alpha\beta} (\mathcal{C}_{1/\sqrt{\alpha\beta}} + \mathcal{C}_{1/\sqrt{\alpha\beta}}^\top)}_{r_t \bar{r}_t} \quad (139)$$

It is also easy to check that

$$\frac{1}{\sqrt{\alpha\beta}} = \frac{\omega^2 \sinh(\lambda)}{\lambda}.$$

Finally,

$$\Omega_t = \Sigma_{t|(x_0, x_1)} = \Phi_t (\mathbf{I} - e^{-2\lambda(1-t)} \Phi_t \Phi_1^{-1}) = \kappa(t, t) (1 - \rho_t) \mathbf{I}. \quad (140)$$

We have verified that all the terms in the expression for the variance agree.

A.5 Proof of Theorem 4

Proof. We follow the arguments of [52] with application to the generalised Schrödinger bridge – not much changes for $\mathbb{Q}$ as the reference and we reproduce in detail the arguments as follows. For the equality of gradients, it suffices to show this for the flow matching component of the loss, since the score matching component can be handled using the exact same arguments [52]. Let $\mathbf{u}_{t|(x_0, x_T)}$ be the conditional flow between (x_0, x_T) and $\mathbf{u}_t$ be the SB flow defined by

$$\mathbf{u}_t(\mathbf{x}) = \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_T) \sim \pi} \left[\frac{p_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x})}{p_t(\mathbf{x})} \mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) \right], \quad p_t(\mathbf{x}) = \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_T) \sim \pi} p_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}).$$

Let $\mathbf{u}_t^\theta$ be the neural flow approximation with parameter θ . Assuming that $p_t(\mathbf{x}) > 0$ for all $(t, \mathbf{x})$ we then have:

$$\nabla_\theta \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_T) \sim \pi} \mathbb{E}_{\mathbf{x} \sim p_{t|(\mathbf{x}_0, \mathbf{x}_T)}} \|\mathbf{u}_t^\theta(\mathbf{x}) - \mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x})\|^2 - \nabla_\theta \mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} \|\mathbf{u}_t^\theta(\mathbf{x}) - \mathbf{u}_t(\mathbf{x})\|^2 \quad (141)$$

$$= \nabla_\theta \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_T) \sim \pi} \mathbb{E}_{\mathbf{x} \sim p_{t|(\mathbf{x}_0, \mathbf{x}_T)}} [\|\mathbf{u}_t^\theta(\mathbf{x})\|^2 + \|\mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x})\|^2 - 2\langle \mathbf{u}_t^\theta(\mathbf{x}), \mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) \rangle] \quad (142)$$

$$- \nabla_\theta \mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} [\|\mathbf{u}_t^\theta(\mathbf{x})\|^2 + \|\mathbf{u}_t(\mathbf{x})\|^2 - 2\langle \mathbf{u}_t^\theta(\mathbf{x}), \mathbf{u}_t(\mathbf{x}) \rangle] \quad (143)$$

$$= 2\nabla_\theta \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_T) \sim \pi} \mathbb{E}_{\mathbf{x} \sim p_{t|(\mathbf{x}_0, \mathbf{x}_T)}} \langle \mathbf{u}_t^\theta(\mathbf{x}), \mathbf{u}_t(\mathbf{x}) \rangle - \langle \mathbf{u}_t^\theta(\mathbf{x}), \mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) \rangle, \quad (144)$$

where in the last line terms not depending on θ are zero and we use the fact that $\mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_T)} \mathbb{E}_{\mathbf{x} | (\mathbf{x}_0, \mathbf{x}_T)} f(\mathbf{x}) = \mathbb{E}_{\mathbf{x}} f(\mathbf{x})$. Now, using the relation between $\mathbf{u}_t$ and $\mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}$,

$$\mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} \langle \mathbf{u}_t^\theta(\mathbf{x}), \mathbf{u}_t(\mathbf{x}) \rangle = \int d p_t(\mathbf{x}) \langle \mathbf{u}_t^\theta(\mathbf{x}), \mathbf{u}_t(\mathbf{x}) \rangle \quad (145)$$

$$= \int d p_t(\mathbf{x}) \left\langle \mathbf{u}_t^\theta(\mathbf{x}), \int d\pi(\mathbf{x}_0, \mathbf{x}_T) \frac{p_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x})}{p_t(\mathbf{x})} \mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) \right\rangle \quad (146)$$

$$= \int d\pi(\mathbf{x}_0, \mathbf{x}_T) \int d\mathbf{x} p_t(\mathbf{x}) \left\langle \mathbf{u}_t^\theta(\mathbf{x}), \frac{p_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x})}{p_t(\mathbf{x})} \mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) \right\rangle \quad (147)$$

$$= \int d\pi(\mathbf{x}_0, \mathbf{x}_T) \int d p_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) \langle \mathbf{u}_t^\theta(\mathbf{x}), \mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) \rangle \quad (148)$$

$$= \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_T) \sim \pi} \mathbb{E}_{\mathbf{x} \sim p_{t|(\mathbf{x}_0, \mathbf{x}_T)}} \langle \mathbf{u}_t^\theta(\mathbf{x}), \mathbf{u}_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x}) \rangle. \quad (149)$$

So we conclude that the two gradients are equal. Clearly, the *unconditional* loss can be rewritten over candidate flow and score fields $(\hat{\mathbf{u}}, \hat{\mathbf{s}})$ as

$$(\hat{\mathbf{u}}, \hat{\mathbf{s}}) \mapsto \|\hat{\mathbf{u}} - \mathbf{u}\|_{L^2(\mathrm{d}p_t(\mathbf{x})\mathrm{d}t)}^2 + \|\lambda_t(\hat{\mathbf{s}} - \mathbf{s})\|_{L^2(\mathrm{d}p_t(\mathbf{x})\mathrm{d}t)}^2,$$

where we denote by $\|h_t(\mathbf{x})\|_{L^2(\mathrm{d}p_t(\mathbf{x})\mathrm{d}t)}^2 = \int_0^1 \mathrm{d}t \int \mathrm{d}p_t(\mathbf{x}) \|h_t(\mathbf{x})\|_{L^2(\mathbb{R}^d)}^2$ for a test function $h_t : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$. This loss is zero iff $\hat{\mathbf{u}} = \mathbf{u}$ and $\hat{\mathbf{s}} = \mathbf{s}$ ($\mathrm{d}p_t \times \mathrm{d}t$)-almost everywhere.

Let $\mathbb{P}$ denote the law of the Schrödinger bridge as per (SBP-dyn) and write $p_t(\mathbf{x}) = \mathbb{P}_t$ to mean its marginal at time t . Then in (2) and for $\mathbb{Q}$ a mvOU process (1), identifying:

- $\mathrm{d}\mathbb{P}_{0T}^* = \pi$ where π is prescribed in Proposition 1, and
- $\mathbb{Q}_t^{x_0, x_T} = p_{t|(\mathbf{x}_0, \mathbf{x}_T)}(\mathbf{x})$ where $p_{t|(\mathbf{x}_0, \mathbf{x}_T)}$ is defined as in Theorem 2,

substituting all these into (2) one has

$$p_t(\mathbf{x}) = \int \mathrm{d}\pi(\mathbf{x}_0, \mathbf{x}_T) p_{t|(\mathbf{x}_0, \mathbf{x}_T)}.$$

Consequently, there exists $\mathbf{u}_t(\mathbf{x})$ such that $\partial_t p_t(\mathbf{x}) = -\nabla \cdot (p_t(\mathbf{x}) \mathbf{u}_t(\mathbf{x}))$. Writing $\mathbf{s}_t(\mathbf{x}) = \nabla_x \log p_t(\mathbf{x})$ to be the score, recognising terms from the probability flow ODE, it follows that that the SDE

$$\mathrm{d}\mathbf{X}_t = (\mathbf{u}_t(\mathbf{X}_t) + D\mathbf{s}_t(\mathbf{X}_t))\mathrm{d}t + \sigma \mathrm{d}\mathbf{B}_t \quad (150)$$

generates the marginals $p_t(\mathbf{x})$ of the Schrödinger bridge. Since $\mathbb{P}$ is characterised as a mixture of $\mathbb{Q}$ -bridges, it follows that $\mathbf{X}_t$ defined by the SDE (150) generates the Markovianisation of $\mathbb{P}$ (see Appendix B of [52]). Moreover, $\mathbb{P}$ is the *unique* process that is both Markov and a mixture of $\mathbb{Q}$ -bridges [28, 45]. $\square$

A.6 Additional background

Hamilton-Jacobi-Bellman equation Consider the *deterministic* control problem on $[0, T]$:

$$V_0(x_0) = \min_u \int_0^T C(x_t, u_t) \mathrm{d}t + D(x_T). \quad (151)$$

subject to a dynamics $\dot{x}_t = F(x_t, u_t)$. $V_t(x)$ is the *value* function for the problem starting at (t, x) up to the final time T . The function $x_T \mapsto D(x_T)$ specifies a cost on the final value. The corresponding Hamilton-Jacobi-Bellman (HJB) equation [23, Section 3.11] is

$$\partial_t V_t(x) + \min_{u_t} \{\langle \nabla_x V_t(x), F(x, u_t(x)) \rangle + C(x, u_t(x))\} = 0, \quad (152)$$

subject to the final boundary condition $V_T(x) = D(x)$. A heuristic derivation is as follows. Considering a time interval $(t, t + \delta t)$ and a path x_t , it's clear that

$$V_t(x_t) = \min_u \left\{ \int_t^{t+\delta t} C(x_s, u_s) \mathrm{d}s + V_{t+\delta t}(x_{t+\delta t}) \right\}. \quad (153)$$

Using Taylor expansion and the constraint $\dot{x}_t = F(x_t, u_t)$ we have to leading order that

$$V_{t+\delta t}(x_{t+\delta t}) = V_t(x_t) + (\partial_t V_t(x_t) + \langle \nabla_x V_t(x_t), F(x_t, u_t) \rangle) \delta t + O(\delta t^2). \quad (154)$$

Substituting back and also approximating the integral, we find that

$$V_t(x_t) = \min_u C(x_t, u_t) \delta t + V_t(x_t) + (\partial_t V_t(x_t) + \langle \nabla_x V_t(x_t), F(x_t, u_t) \rangle) \delta t \quad (155)$$

Cancelling terms, rearranging and taking a limit $\delta t \downarrow 0$, we get the desired result.

Stochastic Hamilton-Jacobi-Bellman Now we consider the *stochastic* variant of the HJB, in which case X_t is driven by an SDE of the form

$$dX_t = F(X_t, u_t) dt + \sigma_t dB_t. \quad (156)$$

The value function is thus in expectation:

$$V_0(x_0) = \min_u \mathbb{E} \left\{ \int_0^T C(X_t, u_t) dt + D(X_T) \right\}. \quad (157)$$

Carrying out the same expansion as before, we have

$$V_t(x_t) = \min_u \mathbb{E} \left\{ \int_t^{t+\delta t} C(X_s, u_s) ds + V_{t+\delta t}(X_{t+\delta t}) \right\}. \quad (158)$$

Note that in the above, $(X_{t+\delta t} | X_t = x_t)$ is a random variable and hence so is $V_{t+\delta t}(X_{t+\delta t})$ which is why it appears in the expectation. Using Itô's formula to expand this, we get

$$\begin{aligned} V_{t+\delta t}(X_{t+\delta t}) &= V_t(X_t) + \left[\partial_t V_t(X_t) + \frac{1}{2} \nabla \cdot (\sigma_t \sigma_t^\top \nabla_x V_t(X_t)) \right] dt + \langle \nabla_x V_t(X_t), dX_t \rangle \\ &= V_t(X_t) + \left[\partial_t V_t(X_t) + \frac{1}{2} \nabla \cdot (\sigma_t \sigma_t^\top \nabla_x V_t(X_t)) + \langle \nabla_x V_t(X_t), F(X_t, u_t) \rangle \right] dt \\ &\quad + \langle \nabla_x V_t(X_t), \sigma_t dB_t \rangle. \end{aligned}$$

Plugging this in, cancelling terms, and noting that the final term has zero expectation, we find that

$$0 = \partial_t V_t(X_t) + \min_u \{ \mathcal{A}[V_t](X_t, u_t) + C(X_t, u_t) \} \quad (159)$$

where $\mathcal{A}$ is the generator of the SDE governing X_t , i.e.

$$\mathcal{A}[f](x, u) = \langle F(x, u), \nabla_x f(x) \rangle + \frac{1}{2} \nabla \cdot (\sigma_t \sigma_t^\top \nabla_x f). \quad (160)$$

The HJB equation for stochastic control problem is therefore

$$0 = \partial_t V_t(x) + \min_u \{ \mathcal{A}[V_t](x, u_t(x)) + C(x, u_t(x)) \}, \quad (161)$$

i.e. this is the same as the HJB for the deterministic case, except with a diffusive term arising from the stochasticity.

B Experiment details

For (mvOU, BM)-OTFM and IPFP, all computations were carried out by CPU (8x Intel Xeon Gold 6254). NLSB and SBIRR computations were accelerated using a single NVIDIA L40S GPU. Code is available at https://github.com/zsteve/mvOU_SBP.

B.1 Gaussian benchmarking

We consider settings of dimension $d \in \{2, 5, 10, 25, 50\}$ with data consisting of an initial and terminal Gaussian distribution and a mvOU reference process constructed as follows. We sample a $d \times 2$ submatrix U from a random $d \times d$ orthogonal matrix. Then for the reference process, A is constructed as the $d \times d$ matrix

$$A = U \begin{bmatrix} 0 & 1 \\ -2.5 & 0 \end{bmatrix} U^\top.$$

Similarly, we let $m = U \begin{bmatrix} 1 & -1 \end{bmatrix}^\top$. For the marginals, we take $\mathcal{N}(\mu_0, \Sigma_0)$ and $\mathcal{N}(\mu_1, \Sigma_1)$ where

$$\mu_0 = U \begin{bmatrix} -2.5 & -0.5 \end{bmatrix}^\top, \quad \Sigma_0 = U \begin{bmatrix} 0.1 & 0.005 \\ 0.005 & 0.1 \end{bmatrix} U^\top + 0.1I.$$

$$\mu_1 = U \begin{bmatrix} 0.5 & 2.5 \end{bmatrix}^\top, \quad \Sigma_1 = U \begin{bmatrix} 1.1 & -2 \\ -2 & 1.1 \end{bmatrix} U^\top + 0.1I.$$

For each d , the initial and final marginals are approximated by $N = 128$ samples $\{x_0^i\}_{i=1}^N \sim \mathcal{N}(\mu_0, \Sigma_0)$, $\{x_1^i\}_{i=1}^N \sim \mathcal{N}(\mu_1, \Sigma_1)$. We fix $\sigma = I$ and compute the exact marginal parameters $(\mu_t, \Sigma_t)_{t \in [0,1]}$ using the formulas of Theorem 3 and reference parameters (A, m) . Additionally, we compute the SDE drift $v_{SB}(t, x)$ of the mvOU-GSB using (17).

(mvOU, BM)-OTFM We apply mvOU-OTFM as per Algorithm 1 to learn a conditional flow matching approximation to the Schrödinger bridge. We choose to parameterise the probability flow and score fields using two feed-forward neural networks $\mathbf{u}_\theta(t, \mathbf{x}) = \text{NN}_\theta(d+1, d)(t, \mathbf{x})$ and $\mathbf{s}_\varphi(t, \mathbf{x}) = \text{NN}_\varphi(d+1, d)(t, \mathbf{x})$, each with $[64, 64, 64]$ hidden dimensions and ReLU activations. We use a batch size of 64 and learning rate 10^{-2} for 2,500 iterations using the AdamW optimiser. For BM-OTFM, the same training procedure was used except the reference process was taken to be Brownian motion with unit diffusivity.

IPFP We use the IPFP implementation provided by the authors of [53], specifically using the variant of their algorithm based on Gaussian process approximations to the Schrödinger bridge drift [53, Algorithm 2]. We provide the mvOU reference process as the prior drift function, i.e. $\mathbf{x} \mapsto \mathbf{A}(\mathbf{x} - \mathbf{m})$ and employ an exponential kernel for the Gaussian process vector field approximation. We run the IPFP algorithm for 10 iterations, which is double that used in the original publication. Since IPFP explicitly constructs forward and reverse processes, at each time $0 \leq t \leq 1$, IPFP outputs *two* estimates of the Schrödinger bridge marginal, one from each process. We consider both outputs and distinguish between them using the $(\rightarrow, \leftarrow)$ symbols in Table 1. For the vector field estimate we employ only the forward drift.

NLSB We use the NLSB implementation provided by the authors of [25]. Following [12, Section 4.6], for a mvOU reference process $d\mathbf{X}_t = \mathbf{f}_t dt + \sigma d\mathbf{B}_t$, the generalised SB problem (SBP-dyn) can be rewritten as a stochastic control problem

$$\min_{\mathbf{u}_t} \mathbb{E} \left[\int_0^1 \frac{1}{2} \|\mathbf{u}_t\|_2^2 dt \right], \text{ subject to } d\mathbf{X}_t = (\mathbf{f}_t + \mathbf{u}_t) dt + \sigma d\mathbf{B}_t.$$

On the other hand, the Lagrangian SB problem [25, Definition 3.1] is

$$\min_{\mathbf{v}_t} \mathbb{E} \left[\int_0^1 L(t, \mathbf{x}, \mathbf{v}_t(\mathbf{x})) dt \right], \text{ subject to } d\mathbf{X}_t = \mathbf{v}_t dt + \sigma d\mathbf{B}_t.$$

This shows that the appropriate Lagrangian to use is $L(t, \mathbf{x}, \mathbf{v}_t(\mathbf{x})) = \frac{1}{2} \|\mathbf{v}_t(\mathbf{x}) - \mathbf{f}_t(\mathbf{x})\|_2^2$. We apply this in the mvOU setting by taking $\mathbf{f}_t(\mathbf{x}) = \mathbf{A}(\mathbf{x} - \mathbf{m})$. The NLSB approach backpropagate through solution of the SDE to directly learn a neural approximation of the SB drift $\mathbf{v}_t$. We train NLSB the same hyperparameters as used in the original publication, using the Adam optimiser with learning rate 10^{-3} for a total of 2,500 epochs.

Metrics We measure both the approximation error of Gaussian Schrödinger bridge marginals as well as the error in vector field estimation. For the marginal reconstruction, for each method and for each $0 \leq t \leq 1$ we sample points integrated forward in time and compute the mean $\hat{\boldsymbol{\mu}}_t$ and variance $\hat{\boldsymbol{\Sigma}}_t$ from samples. We then compare to the ground truth marginals $(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ computed from exact formulas using the Bures-Wasserstein metric (62). For the vector field, for each method and for each time $0 \leq t \leq 1$ we sample $M = 1024$ points from the ground truth SB marginal $\mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ and empirically estimate $\|\hat{\mathbf{v}}_{\text{SB}} - \mathbf{v}_{\text{SB}}\|_{\mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)}$ from samples. All experiments were repeated over 5 independent runs and summary statistics for metrics are shown in Table 1.

B.2 Gaussian mixture example

We consider initial and terminal marginals sampled from Gaussian mixtures. At time $t = 0$, we sample from

$$\mathcal{N}(\mathbf{U}[-0.5 \quad -0.5]^\top, 0.01\mathbf{U}\mathbf{U}^\top), \quad \mathcal{N}(\mathbf{U}[0.5 \quad 0.5]^\top, 0.0625\mathbf{U}\mathbf{U}^\top)$$

in a 1:1 ratio, and at time $t = 1$ we sample in the same fashion from

$$\mathcal{N}(\mathbf{U}[-2.5 \quad -2.5]^\top, 0.01\mathbf{U}\mathbf{U}^\top), \quad \mathcal{N}(\mathbf{U}[2.5 \quad 2.5]^\top, 0.25\mathbf{U}\mathbf{U}^\top).$$

Here, the matrix $\mathbf{U}$ and reference process parameters $(\mathbf{A}, \mathbf{m})$ are the same as used for the previous Gaussian example for $d = 10$. All other training details are the same as for the Gaussian example.

B.3 Repressilator example

Based on the system studied in [44], we simulate stochastic trajectories from the system

$$\begin{aligned} \frac{dx_1}{dt} &= \left(\frac{\beta}{1 + (x_3/k)^n} - \gamma x_1 \right) dt + \sigma dB_t^{(1)} \\ \frac{dx_2}{dt} &= \left(\frac{\beta}{1 + (x_1/k)^n} - \gamma x_2 \right) dt + \sigma dB_t^{(2)}, \\ \frac{dx_3}{dt} &= \left(\frac{\beta}{1 + (x_2/k)^n} - \gamma x_3 \right) dt + \sigma dB_t^{(3)}, \end{aligned} \quad \begin{bmatrix} \beta \\ n \\ k \\ \gamma \\ \sigma \end{bmatrix} = \begin{bmatrix} 10 \\ 3 \\ 1 \\ 1 \\ 0.1 \end{bmatrix}, \quad (162)$$

with initial condition $\mathcal{N}([1, 1, 2], 0.01\mathbf{I})$ and simulated for the time interval $t \in [0, 10]$ using the Euler-Maruyama discretisation. Snapshots were sampled at $T = 10$ time points evenly spaced on $[0, 10]$, each comprising of 100 samples.

We employ Algorithm 2 to this data, running for 5 iterations starting from an initial Brownian reference process $\mathbf{A} = \mathbf{0}, \mathbf{m} = \mathbf{0}$. For the inner loop running mvOU-OTFM (Algorithm 1), we parameterise the probability flow and score as in the Gaussian example, with feed-forward networks of hidden dimension [64, 64, 64] and ReLU activations. We run mvOU-OTFM for 1,000 iterations using the AdamW optimiser, a batch size of 64 and a learning rate of 10^{-2} . At each step of the outer loop, we employ ridge regression to fit the updated mvOU reference parameters $(\mathbf{A}, \mathbf{m})$. We do this using the standard RidgeCV method implemented in the `scikit-learn` package, which automatically selects the regularisation parameter.

We also carry out hold-one-out runs where for $2 \leq i \leq 9$ (i.e. all time-points except for the very first and last), Algorithm 2 is applied to $T - 1$ snapshots with the snapshot at t_i held out. Once the mvOU reference parameters are learned, forward integration of the learned mvOU-SB is used to predict the marginal at t_i . We report the reconstruction error in terms of the earth-mover distance (EMD) and energy distance [41]. Table 2 shows results averaged over held-out timepoints, and full results (split by timepoint) are shown in Table 5.

Since the system marginals are unimodal we reason that they can be reasonably well approximated by Gaussians. For each $2 \leq i \leq 9$ we fit multivariate Gaussians to the snapshots at t_{i-1}, t_i, t_{i+1} and use the results of Theorem 3 with the fitted mvOU reference output by Algorithm 2 to solve the mvOU-GSB between $p_{t_{i-1}}, p_{t_{i+1}}$. This is illustrated in Figure 4 for $i = 4$, and we show the full results for all timepoints in Figure 6 in comparison with the standard Brownian GSB.

SBIRR We apply SBIRR [44] using the implementation provided with the original publication, which notably includes an improved implementation of IPFP [53] that utilises GPU acceleration. We provide the reference vector field $\mathbf{x} \mapsto \mathbf{A}(\mathbf{x} - \mathbf{m})$ and seek to learn a reference process in one of two families: (1) mvOU processes, i.e. we consider the family of reference drifts $\hat{\mathbf{A}}(\mathbf{x} - \hat{\mathbf{m}})$ where $\hat{\mathbf{A}}, \hat{\mathbf{m}}$ are to be fit, or (2) general drifts, i.e. we parameterise the drift using a feed-forward neural network with hidden dimensions [64, 64, 64]. For each choice of reference family, we run Algorithm 1 of [44] for 5 outer iterations and 10 inner IPFP iterations, as was also done in the original paper.

B.4 Cell cycle scRNA-seq

The metabolic labelled cell cycle dataset of [4] is obtained and preprocessed following the tutorial available with the Dynamo [40] package. This gives a dataset of $N = 2,793$ cells, embedded in 30 PCA dimensions. In addition to transcriptional state $\{\mathbf{x}_i\}_{i=1}^N$, Dynamo uses metabolic labelling data to predict the transcriptional velocity $\{\hat{\mathbf{v}}_i\}_{i=1}^N$ for each cell.

To fit the reference process parameters $(\mathbf{A}, \mathbf{m})$, we use again ridge regression via the RidgeCV method in `scikit-learn`. We train mvOU-OTFM using Algorithm 1 with $\sigma = 0.3$, parameterising the probability flow and score as previously using feed-forward networks of hidden dimensions [64, 64, 64] and train with a batch size of 64, learning rate of 10^{-2} for a total of 1,000 iterations.




Figure 6: **Repressilator mvOU-GSB interpolation.** Using the learned mvOU reference process, we interpolate between p_{i-1} (blue) and p_{i+1} (green). Middle timepoint p_i is shown in red.




Figure 7: **Cell cycle mvOU-GSB interpolation.** Using the learned mvOU reference process and scale factor $\gamma = 50$, we interpolate between the first snapshot p_1 (blue) and last snapshot p_T (green). All computations are done in $d = 30$ and shown in leading 2 PCs.

		Scale factor γ										
	t	0.0	10.0	20.0	30.0	40.0	50.0	60.0	70.0	80.0	90.0	100.0
$d = 50$	0.25	40.18	30.8	22.55	16.34	12.77	11.07	10.26	10.11	10.69	11.93	13.65
	0.5	28.59	18.47	11.53	9.62	12.36	15.32	15	12.52	10.64	11.02	13.96
	0.75	15.31	10.31	7.47	7.06	8.18	8.9	8.5	7.9	8.09	9.42	11.92
$d = 100$	0.25	44.49	36.82	29.97	24.19	19.99	17.48	16.13	15.29	14.78	14.7	15.18
	0.5	32.48	23.75	17.12	13.55	13.79	16.84	19.78	20.21	18.25	15.7	14.22
	0.75	19.33	14.95	12	10.74	11.19	12.68	13.98	14.36	14.03	13.69	13.82

Table 4: Single cell interpolation results for $d = 50, 100$.

Algorithm 2 Iterated reference fitting with mvOU-OTFM

Input: Samples $\{\mathbf{x}_i^{t_j}\}_{i=1}^{N_j}$ from multiple snapshots at times $\{t_j\}_{j=1}^T$, initial mvOU reference parameters $(\mathbf{A}, \mathbf{m})$, diffusivity $\mathbf{D} = \frac{1}{2}\boldsymbol{\sigma}\boldsymbol{\sigma}^\top$
Initialise: Probability flow field $\mathbf{u}_t^\theta(\mathbf{x})$, score field $\mathbf{s}_t^\varphi(\mathbf{x})$.
Define: $\text{FitReference}(\mathbf{X}, \mathbf{V}) := \arg \min_{\mathbf{A}, \mathbf{m}} \|\mathbf{V} - \mathbf{A}(\mathbf{X} - \mathbf{m})\|_2^2 + \lambda \|\mathbf{A}\|_F^2 + \gamma \|\mathbf{m}\|^2$
 $\hat{\rho}_j \leftarrow N^{-1} \sum_{i=1}^N \delta_{\mathbf{x}_i^{t_j}}, 1 \leq j \leq T$ *Form empirical marginals*
while not converged **do**
 $(\mathbf{u}_t^\theta, \mathbf{s}_t^\theta) \leftarrow \text{fitOTFM}_{(\mathbf{A}, \mathbf{m}, \mathbf{D})}(\hat{\rho}_1, \dots, \hat{\rho}_T)$ *Fit flow and score with reference parameters*
 $\mathbf{v}_i^{t_j} \leftarrow (\mathbf{u}_{t_j}^\theta + \mathbf{D}\mathbf{s}_{t_j}^\varphi)(\mathbf{x}_i^{t_j}), \quad 1 \leq i \leq N_j, 1 \leq j \leq T$ *Get SDE drift*
 $\mathbf{A}, \mathbf{m} \leftarrow \text{FitReference}(\{(\mathbf{v}_i^{t_j})_{i=1}^{N_j}\}_{j=1}^T, \{(\mathbf{x}_i^{t_j})_{i=1}^{N_j}\}_{j=1}^T)$ *Update reference parameters*
end while

Error metric	t	Leave-one-out marginal interpolation error						SBIRR (mvOU)	SBIRR (MLP)
		Iterate 0	Iterate 1	Iterate 2	Iterate 3	Iterate 4			
EMD	1	3.59 ± 0.17	3.14 ± 0.12	2.28 ± 0.15	2.08 ± 0.11	2.02 ± 0.11	2.24 ± 0.28	2.71 ± 0.41	
	2	5.20 ± 0.47	2.59 ± 0.29	1.62 ± 0.28	1.27 ± 0.25	1.13 ± 0.16	3.13 ± 0.62	2.29 ± 0.92	
	3	3.23 ± 0.24	1.42 ± 0.18	1.10 ± 0.14	0.86 ± 0.08	0.83 ± 0.10	2.67 ± 0.85	1.39 ± 0.55	
	4	1.48 ± 0.20	0.52 ± 0.05	0.47 ± 0.05	0.47 ± 0.06	0.48 ± 0.06	1.38 ± 0.46	0.94 ± 0.28	
	5	2.50 ± 0.40	1.43 ± 0.65	1.12 ± 0.32	1.29 ± 0.11	1.21 ± 0.13	1.63 ± 0.17	1.40 ± 0.71	
	6	6.18 ± 0.41	3.42 ± 0.50	2.18 ± 0.28	1.91 ± 0.38	1.75 ± 0.17	2.40 ± 0.32	1.96 ± 1.41	
	7	2.56 ± 0.25	3.09 ± 1.53	2.13 ± 0.46	2.23 ± 0.55	1.93 ± 0.50	1.45 ± 0.20	0.51 ± 0.13	
	8	2.29 ± 0.26	2.12 ± 0.09	1.82 ± 0.33	1.81 ± 0.31	1.81 ± 0.16	1.93 ± 0.61	2.14 ± 0.40	
Energy	1	3.00 ± 0.09	2.75 ± 0.07	2.24 ± 0.10	2.10 ± 0.07	2.05 ± 0.08	2.20 ± 0.18	2.56 ± 0.26	
	2	3.53 ± 0.21	2.27 ± 0.19	1.66 ± 0.20	1.38 ± 0.21	1.26 ± 0.13	2.61 ± 0.33	2.05 ± 0.61	
	3	2.29 ± 0.17	1.31 ± 0.16	1.00 ± 0.10	0.80 ± 0.06	0.78 ± 0.07	2.10 ± 0.52	1.25 ± 0.43	
	4	0.93 ± 0.14	0.25 ± 0.06	0.17 ± 0.04	0.15 ± 0.02	0.14 ± 0.01	0.99 ± 0.38	0.84 ± 0.25	
	5	1.10 ± 0.18	0.53 ± 0.27	0.45 ± 0.13	0.55 ± 0.04	0.51 ± 0.07	0.64 ± 0.10	0.66 ± 0.37	
	6	2.49 ± 0.12	1.22 ± 0.10	1.13 ± 0.15	1.01 ± 0.22	0.91 ± 0.11	1.43 ± 0.17	0.90 ± 0.64	
	7	0.97 ± 0.13	1.39 ± 0.65	1.02 ± 0.21	1.05 ± 0.25	0.91 ± 0.25	0.65 ± 0.09	0.17 ± 0.08	
	8	0.54 ± 0.11	0.57 ± 0.03	0.56 ± 0.11	0.55 ± 0.10	0.55 ± 0.04	0.48 ± 0.08	0.36 ± 0.06	

Table 5: Full results for repressilator example

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 this paper's contributions and scope.

Guidelines:

- The answer NA 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 NA 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: We discuss the limitations of our work, specifically on the linearity of the reference process and the computational cost of our algorithm.

Guidelines:

- The answer NA 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: We provide all assumptions and complete proofs for all theoretical results stated.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulae, 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: In our paper we provide all information needed to reproduce our main experimental results.

Guidelines:

- The answer NA 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 data and code with instructions to reproduce our experimental results in our supplemental material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.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://nips.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, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: These details are all available in our supplemental material.

Guidelines:

- The answer NA 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: We provide error bars for the main experimental results of our paper.

Guidelines:

- The answer NA 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: These details are all available in our supplemental material.

Guidelines:

- The answer NA 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: All research in this paper conforms with the NeurIPS code of ethics.

Guidelines:

- The answer NA 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: [NA]

Justification: Our work is foundational and we do not foresee any direct path to any societal implications arising from our work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA 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., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

- The answer NA 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 assets used in our research are appropriately credited.

Guidelines:

- The answer NA 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: [\[Yes\]](#)

Justification: We provide the full code and data in our supplementary materials.

Guidelines:

- The answer NA 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: [\[NA\]](#)

Justification: Not applicable.

Guidelines:

- The answer NA 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: [\[NA\]](#)

Justification: Not applicable.

Guidelines:

- The answer NA 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 rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: No LLM usage towards this paper.

Guidelines:

- The answer NA 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 (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.