Article

Solving Schrödinger Bridges via Maximum Likelihood

Francisco Vargas [1,*], Pierre Thodoroff [1], Austen Lamacraft [2] and Neil Lawrence [1,*]

[1] The Computer Laboratory, Department of Computer Science and Technology, University of Cambridge, William Gates Building, 15 JJ Thomson Avenue, Cambridge CB3 0FD, UK; pt440@cam.ac.uk
[2] The Cavendish Laboratory, Deparment of Physics, The Old Schools, Trinity Ln, Cambridge CB2 1TN, UK; al200@cam.ac.uk
* Correspondence: fav25@cam.ac.uk (F.V.); ndl21@cam.ac.uk (N.L.)

Abstract: The Schrödinger bridge problem (SBP) finds the most likely stochastic evolution between two probability distributions given a prior stochastic evolution. As well as applications in the natural sciences, problems of this kind have important applications in machine learning such as dataset alignment and hypothesis testing. Whilst the theory behind this problem is relatively mature, scalable numerical recipes to estimate the Schrödinger bridge remain an active area of research. Our main contribution is the proof of equivalence between solving the SBP and an autoregressive maximum likelihood estimation objective. This formulation circumvents many of the challenges of density estimation and enables direct application of successful machine learning techniques. We propose a numerical procedure to estimate SBPs using Gaussian process and demonstrate the practical usage of our approach in numerical simulations and experiments.

Keywords: Schrödinger bridges; machine learning; stochastic control

1. Introduction

Analysis of cross-sectional data is ubiquitous in machine learning and science. Temporal data are typically sampled at discrete intervals due to technological or physical constraints. This means information between time points is lost. This motivates the need to model the stochastic evolution of a process between sampled time points. The classical Schrödinger bridge problem [1,2] finds the most likely stochastic process that evolves a distribution $\pi_0(x)$ to another distribution $\pi_1(y)$ consistently with a pre-specified Brownian motion. We consider a more general dynamical Schrödinger bridge problem for *any* pre-specified diffusion prior. Practically, this generalization allows us to exploit domain knowledge, e.g., oceanic and atmospheric flows might be interpolated from empirical measurements using previously established dynamics as priors.

In the classical set up numerical approaches to solve the Schrödinger bridge are mainly based on the Sinkhorn–Knopp algorithm [3]; however, extending those algorithms to more general diffusion priors and marginals requires complex adaptations. We introduce an iterative proportional maximum likelihood (IPML) algorithm to solve the general Schrödinger bridge problem. The IPML algorithm also obtains a good approximation for the dynamics of the underlying physical process that solves the SBP. This contrasts to previous approaches [4–6] that estimate the value that extremises the SBP objective. Practically, this means we obtain physically interpretable solutions that we can leverage for downstream tasks.

IPML's inspiration is from *probabilistic numerics* (PN) [7]. We combine a PN styled formulation with the iterative proportional fitting procedure (IPFP) [8,9]. The algorithm iteratively simulates trajectories that converge to the SBP. We prove that IPML converges in probability at each iteration. Our numerical experiments show the algorithm can be implemented with Gaussian process (GP) models of the drifts. GPs allow us to incorporate functional prior information. We demonstrate the practical use of our algorithm on real-world embryoid cells data (see Figure 4) by quantitatively and qualitatively comparing

the performance of our algorithm to state-of-the-art deep learning methods and optimal transport techniques. To summarise, the main contributions of our work are:

- We recast the iterations of the dynamic IPFP algorithm as a regression-based maximum likelihood objective. This is formalised in Theorem 1 and Observation 1. This differs to prior approaches such as [10] where their maximum likelihood formulation solves a density estimation problem. This allows the application of many regression methods from machine learning that scale well to high dimensional problems. Note that this formulation can be parametrised with any method of choice. We chose GPs; however, neural networks would also be well suited,
- We solve the aforementioned regression objectives using GPs [11] motivated by the connection between the drift of stochastic differential equations and GPs [12],
- We provide a conceptual comparison with the approach by [10] and detail why the density estimation formulation in [10] scales poorly with dimension,
- Finally we re-implement the approach by [10] and compare approaches across a series of numerical experiments. Furthermore we empirically show how our approach works well in dimensions higher than 2.

2. Technical Background

Our solution has three components. (1) We reformulate the SBP as a dynamical system giving a stochastic differential equation (SDE) with initial value (IV) and final value (FV) constraints [13]. (2) We reverse the system, reformulating the FV constraint as an IV constraint. Both IV problems are solved through a stochastic control formulation (Section 2.1.2). (3) We iterate the IV and FV constrained problems to converge to the full boundary value constrained SDEs.

2.1. Dynamic Formulation

The dynamic version of the Schrödinger bridge is written in terms of measures over the space of trajectories, which describe the stochastic dynamics defined over the unit interval.

Definition 1. *(Dynamic Schrödinger problem) The dynamic Schrödinger problem is given by*

$$\inf_{\mathbb{Q} \in \mathcal{D}(\pi_0, \pi_1)} D_{\mathrm{KL}}\big(\mathbb{Q}||\mathbb{Q}_0^\gamma\big), \tag{1}$$

where $\mathbb{Q} \in \mathcal{D}(\pi_0, \pi_1)$ is a measure with prescribed marginals of π_0, π_1 at times $0, 1$ that is $(X_0)_{\#}\mathbb{Q} = \pi_0$ and $(X_1)_{\#}\mathbb{Q} = \pi_1$. $\mathbb{Q}_0^\gamma$ is a drift augmented Brownian motion with a scalar volatility γ acting as the prior, see Figure 1. Traditionally $\mathbb{Q}_0^\gamma := \mathbb{W}^\gamma$ is a Wiener measure with volatility γ, however we consider the general setting in this work.

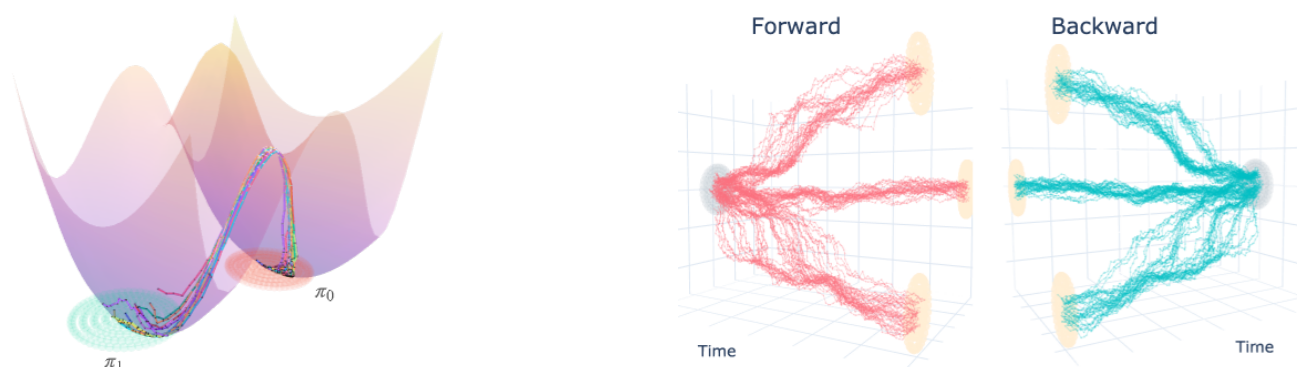**Figure 1. Left**: Learned SBP trajectories in the double well experiment of Section 5.1, with prior $\mathbb{Q}_0^\gamma$ expressed in terms of an energy landscape $U(x, y)$ as $d\mathbf{x}(t) = -\nabla_x U(\mathbf{x}(t)) + \gamma d\boldsymbol{\beta}(t)$. **Right**: Forwards and backwards diffusion of learned SBP between unimodal and multimodal boundary distributions (see Section 5 for experimental details).

The prior $\mathbb{Q}_0^\gamma$ can be written as a solution of the diffusion:

$$d\mathbf{x}(t) = \boldsymbol{b}_0^+(\mathbf{x}(t), t)dt + \sqrt{\gamma}d\boldsymbol{\beta}^+(t), \quad \mathbf{x}(0) \sim \pi_0^{\mathbb{Q}_0^\gamma}. \tag{2}$$

From here we use $\boldsymbol{b}_0$ to denote the drift of the prior $\mathbb{Q}_0^\gamma$. Note that a finite KL implies $\mathbb{Q}$ and $\mathbb{Q}_0^\gamma$ are both Itô SDEs with volatility γ. The diffusion coefficient γ is a time homogeneous constant and for the KL-divergence to be finite the process $\mathbb{Q}$ must have γ as its diffusion coefficient.

2.1.1. Time Reversal of Diffusions

When sampling solutions, IV constraints are trivially solved by initialisation of samples, but FV constraints are more problematic; however, if we can reformulate the SBP as a *time reversed diffusion* the FV constraint becomes an IV constraint (Section 3.2). Here we review how the diffusion is reversed. We reverse time (see Figure 1) in the random variable $\mathbf{x}(t)$ described by the diffusion in Equation (2) such that $\mathbf{x}^-(t) = \mathbf{x}(1-t)$. The reverse time diffusion $\mathbf{x}^-(t)$ is also an Itô process. For a modern application of time reversal in machine learning see [14].

Lemma 1 ([15])**.** *If* $\mathbf{x}^+(t)$ *obeys the SDE*

$$d\mathbf{x}^+(t) = \boldsymbol{u}^+(\mathbf{x}^+(t), t)dt + \sqrt{\gamma}d\boldsymbol{\beta}^+(t), \quad \mathbf{x}^+(0) \sim \pi_0^{\mathbb{Q}_0^\gamma},$$

then $\mathbf{x}^-(t) = \mathbf{x}^+(1-t)$ *obeys*

$$d\mathbf{x}^-(t) = \boldsymbol{u}^-(\mathbf{x}^-(t), t)dt + \sqrt{\gamma}d\boldsymbol{\beta}^-(t), \quad \mathbf{x}^-(0) = \mathbf{x}^+(1).$$

where $\boldsymbol{\beta}^-(t)$ *is a Brownian motion adapted to the reverse filtration* $(\mathcal{F}_i^-)_{i \in T}$*, that is* $\mathcal{F}_t^- \subseteq \mathcal{F}_s^-$*, $s \leq t$. Furthermore, the dual drift* $\boldsymbol{u}^-(\mathbf{x}, t)$ *satisfies Nelson's duality relation:*

$$\boldsymbol{u}^+(\mathbf{x}, t) - \boldsymbol{u}^-(\mathbf{x}, t) = \gamma \nabla_x \ln p(\mathbf{x}, t), \tag{3}$$

where $p(\mathbf{x}, t)$ *solves the associated Fokker–Planck equation.*

Proof. A variety of proofs can be found for this result [15–19]. Note the formulation of Equation (3) varies slightly across studies we use the one from [15,16] (see [15] (p. 87), definition of osmotic velocity). □

2.1.2. Stochastic Control Formulation

Now that the FV constraint has been converted to an IV constraint, we cast the problem into a stochastic control formulation to estimate the drift of each diffusion process. Following from the dynamic formulation, the control formulation casts the problem explicitly in terms of stochastic differential equations. The control formulation is used to enforce constraints as initial value problems. Furthermore, the drift based formulations of the SBP admit a reverse time formulation, which starts the chain at the end of the interval and progresses the dynamics backwards in time to the start.

Lemma 2 ([20])**.** *Let the measure* $\mathbb{Q}$ *be defined by solutions to the SDE*

$$d\mathbf{x}^+(t) = \mathbf{b}^+(t)dt + \sqrt{\gamma}d\boldsymbol{\beta}^+(t), \quad \mathbf{x}^+(0) \sim \pi_0^{\mathbb{Q}}$$

Then the KL divergence $D_{\mathrm{KL}}(\mathbb{Q}||\mathbb{Q}_0^\gamma)$ *can be decomposed in terms of either the forward or reversed diffusion as*

$$D_{\mathrm{KL}}(\mathbb{Q}||\mathbb{Q}_0^\gamma) = D_{\mathrm{KL}}(\pi_{\frac{1 \mp 1}{2}}^{\mathbb{Q}}||\pi_{\frac{1 \mp 1}{2}}^{\mathbb{Q}_0^\gamma}) + \mathbb{E}_{\mathbb{Q}}\left[\int_0^1 \frac{1}{2\gamma}||\mathbf{b}^\pm(t) - \boldsymbol{b}_0^\pm(\mathbf{x}^\pm(t), t)||^2 dt\right], \tag{4}$$

where $\mathbf{x}^-(t)$ *is the time reversal of* $\mathbf{x}^+(t)$, *and* $\mathbf{x}^+(1) \sim \pi_1^{\mathbb{Q}}$.

Proof. This theorem follows by a direct application of the disintegration theorem (Appendix B) followed by Girsanov's theorem. A detailed proof can be found in [20]. □

Using the above decompositions, we can solve the SBP by minimising either decomposition in Equation (4) over the space of random processes $\mathbf{b}^\pm(t)$ that satisfy a valid Itô SDE drift. The backwards and forwards objectives are respectively:

$$\min_{\mathbb{Q} \in \mathcal{D}(\pi_0, \pi_1)} D_{\mathrm{KL}}(\mathbb{Q} || \mathbb{Q}_0^\gamma) = \min_{\mathbf{b}^\pm \in \mathcal{B}} \mathbb{E}_{\mathbb{Q}} \left[\int_0^1 \frac{1}{2\gamma} ||\mathbf{b}^\pm(t) - b_0^\pm(\mathbf{x}^\pm(t), t)||^2 dt \right],$$

$$\text{s.t.} \quad d\mathbf{x}^\pm(t) = \mathbf{b}^\pm(t) dt + \sqrt{\gamma} \boldsymbol{\beta}^\pm(t), \quad \mathbf{x}^+(0) \sim \pi_0, \ \mathbf{x}^+(1) \sim \pi_1, \ \mathbf{x}^-(1) \sim \pi_0, \ \mathbf{x}^-(0) \sim \pi_1. \quad (5)$$

While we do not directly use the stochastic control formulations, the drift-based formulation serves as inspiration for our iterative scheme. Specifically, the existence and parametrisation of an optimal drift as shown in [10,20].

Lemma 3 ([20])**.** *The objectives in Equation (5) have optimal drifts:*

$$\mathbf{b}^{+*}(t) = \gamma \nabla_x \phi(\mathbf{x}^+(t), t), \quad \mathbf{b}^{-*}(t) = \gamma \nabla_x \hat{\phi}(\mathbf{x}^-(t), t) \quad (6)$$

where the potentials $\phi, \hat{\phi}$ *solve the Schrödinger system [20].*

This Lemma is key in formulating our ML approach to IPFP since it justifies our parametrisation of the drift in terms of a deterministic function $\boldsymbol{u}^{\pm*}$ i.e., $\mathbf{b}^{\pm*}(t) = \boldsymbol{u}^{\pm*}(\mathbf{x}^\pm(t), t)$. For a brief introduction to the Schrödinger system and potentials see Appendix A.

2.2. Iterative Proportional Fitting Procedure

We now have two boundary value constrained diffusion processes solved through a stochastic control formulation. We use the iterative proportional fitting procedure (IPFP) to alternate between the forward and backward formulation with only one initial value constraint enforced at a time (see Figure 2), such that convergence to the full boundary value problem is guaranteed.

Algorithm 1: Iterative Proportional Fitting (IPFP) [21], for the traditional SBP with Brownian prior set $\mathbb{Q}_0^\gamma := \mathbb{W}^\gamma$.

---

    input: $\pi_0(x), \pi_1(y), \mathbb{Q}_0^\gamma$

1 $\mathbb{Q}_0^* = \mathbb{Q}_0^\gamma$

2 $i = 0$

3 **repeat**

4 $i := i + 1$

5 $\mathbb{P}_i^* = \arg\inf_{\mathbb{P} \in \mathcal{D}(\cdot, \pi_1)} D_{\mathrm{KL}}(\mathbb{P} || \mathbb{Q}_{i-1}^*)$

6 $\mathbb{Q}_i^* = \arg\inf_{\mathbb{Q} \in \mathcal{D}(\pi_0, \cdot)} D_{\mathrm{KL}}(\mathbb{Q} || \mathbb{P}_i^*)$

7 **until** *convergence*;

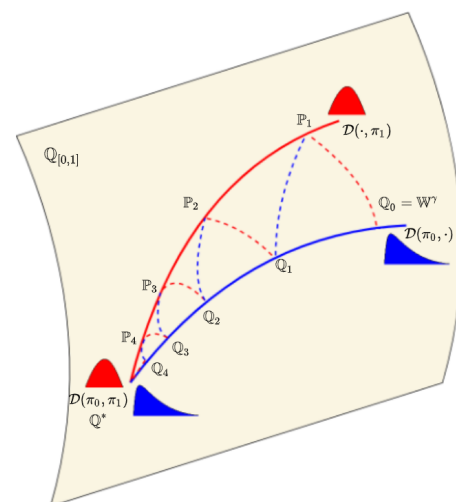**8** **return** $\mathbb{Q}_i^*, \mathbb{P}_i^*$

Figure 2. Left: IPFP algorithm. **Right**: illustration of the convergence of the IPFP Algorithm.

The measure theoretic version of IPFP we introduce is an extension of the continuous IPFP, initially proposed by [8] to a more general setting over general probability measures (see [21,22]). The convergence of IPFP has been shown in [9] and further extensions and

results have been presented in [21,22]. The idea behind this family of approaches is to alternate minimising KL between the two marginal distribution constraints,

$$\mathbb{P}_i^* = \underset{\mathbb{P}\in\mathcal{D}(\cdot,\pi_1)}{\arg\inf}\ D_{\mathrm{KL}}(\mathbb{P}||\mathbb{Q}_{i-1}^*) \quad \mathbb{Q}_i^* = \underset{\mathbb{Q}\in\mathcal{D}(\pi_0,\cdot)}{\arg\inf}\ D_{\mathrm{KL}}(\mathbb{Q}||\mathbb{P}_i^*),$$

until convergence. The quantities $\mathbb{P}_i^*$ and $\mathbb{Q}_i^*$ in the algorithm of Figure 2 are known as half bridges [10] and can be expressed in closed form in terms of known quantities,

$$\mathbb{P}_i^*\left(A_{[0,1)} \times A_1\right) = \int_{A_{[0,1)}\times A_1} \frac{d\pi_1}{dp_1^{\mathbb{Q}_{i-1}^*}}d\mathbb{Q}_{i-1}^*, \quad \mathbb{Q}_i^*\left(A_0 \times A_{(0,1]}\right) = \int_{A_0\times A_{(0,1]}} \frac{d\pi_0}{dp_0^{\mathbb{P}_i^*}}d\mathbb{P}_i^*, \quad (7)$$

where $p_1^{\mathbb{Q}_{i-1}^*}$ and $p_0^{\mathbb{P}_i^*}$ are the marginals of $\mathbb{Q}_{i-1}^*$ and $\mathbb{P}_i^*$ at times 1 and 0, respectively. While a closed form expression to estimate half bridges is known, its components are usually not available in closed form and require approximations. Note that the time reversed formulations (Section 2.1.2) and a simpler variant of Lemma 3 also apply to the half bridge problems. Applying the disintegration theorem (see Appendix B) we can reduce the IPFP introduced to the more popular instance of IPFP proposed by [8]. Furthermore for the case of discrete measures this algorithm reduces to the Sinkhorn–Knopp algorithm [3].

3. Methodology

In this section we introduce IPML by providing the theoretical foundations of our algorithm and convergence guarantees. Then, we propose a practical implementation of IPML based on a Bayesian non-parametric model (GP). These components allow us to solve the general SBP problem.

3.1. Approximate Half Bridge Solving as Optimal Drift Estimation

We present a novel approach to approximately solve the empirical Schrödinger bridge problem by exploiting the closed form expressions of the half bridge problem. Rather than parametrising the measures in the half bridge and solving the optimisation numerically, we seek to directly approximate the measure that extremises the half bridge objective. We achieve this by using Gaussian processes [11] to estimate the drift of the trajectories sampled from the optimal half bridge measure.

We start from the following observation, which tells us how to sample from the optimal half bridge distribution (see Appendix C for proof).

Observation 1. *We can parametrise a measure $\mathbb{Q}$ with its drift as either of the SDEs,*

$$d\mathbf{x}^{\pm}(t) = \mathbf{b}^{\pm}(t) + \sqrt{\gamma}d\boldsymbol{\beta}^{\pm}(t), \ \ \mathbf{x}^{\pm}(0) \sim \pi_{0,1}^{\mathbb{Q}}.$$

Then, we can sample from the solution to the half bridges,

$$\mathbb{P}^{*-} = \underset{\mathbb{P}\in\mathcal{D}(\cdot,\pi_1)}{\arg\inf} D_{\mathrm{KL}}(\mathbb{P}||\mathbb{Q}), \ \ \mathbb{P}^{*+} = \underset{\mathbb{P}\in\mathcal{D}(\pi_0,\cdot)}{\arg\inf} D_{\mathrm{KL}}(\mathbb{P}||\mathbb{Q}),$$

via simulating trajectories (e.g., using the Euler–Maruyama (EM) method) following the SDEs

$$d\mathbf{x}^{\pm}(t) = \mathbf{b}^{\pm}(t) + \sqrt{\gamma}d\boldsymbol{\beta}^{\pm}(t), \ \ \mathbf{x}^{\pm}(0) \sim \pi_{0,1}.$$

Solutions to the above SDEs are distributed according to $\mathbb{P}^{*-}$ and $\mathbb{P}^{*+}$, respectively.

Intuitively, we are performing a cut-and-paste-styled operation by cutting the dynamics of the shortened (unconstrained) time interval and pasting the constraint to it at the corresponding boundary.

Theorem 1. *(Consistency of Reverse-MLE Formulations) Let $\left\{ \left(x_{t_k}^{(n)+} \right)_{k=0}^{T} \right\}_{n=0}^{N}$ be sampled/ discretised trajectories from the SDE that represents the half bridge measure $\mathbb{P}^+$:*

$$d\mathbf{x}^+(t) = \boldsymbol{u}_0^+(\mathbf{x}^+(t), t) + \sqrt{\gamma} d\boldsymbol{\beta}^+(t), \quad \mathbf{x}(0) \sim \pi_0.$$

Then carrying out maximum likelihood estimation of the time reversed drift $\boldsymbol{u}_0^-$ on time reversed samples:

$$\prod_n p\left(\left(x_{t_k}^{(n)+} \right)_{k=0}^{T} \Big| \boldsymbol{u}^- \right) \propto \prod_n p(x_{t_T}^{(n)+}) \prod_{k=0}^{T} \mathcal{N}\left(x_{t_k-\Delta t}^{(n)+} \Big| x_{t_k}^{(n)+} - \Delta t \boldsymbol{u}^- \left(x_{t_k}^{(n)+}, 1 - t_k \right), \gamma \Delta t \right), \quad (8)$$

where $\boldsymbol{u}^-$ is the drift of our estimator SDE:

$$d\mathbf{x}^-(t) = \boldsymbol{u}^-(\mathbf{x}^-(t), t) + \sqrt{\gamma} d\boldsymbol{\beta}^-(t), \quad \mathbf{x}^-(t)(0) \sim p, \quad (9)$$

converges in probability to the to the true dual drift $\boldsymbol{u}_0^-(\boldsymbol{x}, t) = \boldsymbol{u}_0^+(\boldsymbol{x}, t) - \gamma \nabla_{\boldsymbol{x}} \ln p(\boldsymbol{x}, t)$ where $\boldsymbol{u}_0^+(\boldsymbol{x}, t)$ is the optimal half bridge drift for $\mathbb{P}^+$ as $N \to \infty, \Delta t \to 0$ (see Appendix D for proof).

Note that w.l.o.g. the above result also holds for estimating the forward drift from backward samples. The combination of Observation 1 and Theorem 1 constitute one the main contributions of this work as they allow us to solve half bridges with a simple regression objective. Finally it is important to highlight that in practice for most interesting SDEs we can only sample approximately using schemes such as the Euler Mayurama (EM) method, for these settings our proof of Theorem 1 does not hold, however we believe it may be possible to extend/adapt the result for the approximate EM sample case.

3.2. On the Need for Time Reversal

We have already provided the intuition as to how time reversal enables the exchange of initial value constraints for final value constraints. In what follows we provide a more detailed technical review of this important simplification. Consider the first half bridge problem for a given IPFP iteration, $\mathbb{P}^{*-} = \arg\inf_{\mathbb{P} \in \mathcal{D}(\cdot, \pi_1)} D_{\text{KL}}(\mathbb{P}||\mathbb{Q})$. When formulated without time reversal in terms of the forward drift and the forward diffusion it reduces to a stochastic control problem which is subject to dynamics,

$$d\mathbf{x}^+(t) = \mathbf{b}^+(t) + \sqrt{\gamma} d\boldsymbol{\beta}^+(t), \quad \mathbf{x}(1) \sim \pi_1,$$

which is a forward SDE with a terminal hitting condition. Unlike a forward SDE with an initial value problem the above SDE is not trivial to sample from, one approach would be to solve the backwards Kolmogorov equation subject to the terminal condition; however, this approach requires (1) solving a parabolic PDE that typically involves mesh-based methods that do not scale well in high dimensions and (2) carrying out density estimation on the samples from π_1, also problematic in high dimensions; however, if we consider the time reversed process, the terminal condition becomes an initial value problem, $d\mathbf{x}^-(t) = \mathbf{b}^-(t) + \sqrt{\gamma} d\boldsymbol{\beta}^-(t), \mathbf{x}(0) \sim \pi_1$, for which it is easy to sample consistent trajectories via the Euler–Maruyama (EM) discretisation without mesh-based methods or additional density estimations.

3.3. Iterative Proportional Maximum Likelihood (IPML)

Combining the KL minimisation routines in the original IPFP algorithm with our MLE-based drift estimation provides Figure 3. The routine DriftFit fits a dual drift on the EM sampled trajectories, and can be parametrised by any function estimation procedure with consistency guarantees. The SDESolve routine generates M trajectories using the EM method.

Algorithm 2: IPML.

  input: $\pi_0(x), \pi_1(y), \mathbb{Q}_0^\gamma$

1 Initialise: $i := 0 \; \mathbb{Q}_0^* := \mathbb{Q}_0^\gamma$

2 $\bar{u}_{\mathbb{Q}_0}^-(x,t) := \text{ObtainBackwardDrift}(\mathbb{Q}_0^*)$

3 repeat

4 $i := i + 1$

5 $\left\{ x_{[1:T]}^{(m)-} \right\}_m := \text{SDESolve}\left(-\bar{u}_{\mathbb{Q}_{i-1}}^-, \gamma, \pi_1, \Delta t \right)$

6 $\bar{u}_{\mathbb{P}_i}^+ := \text{DriftFit}\left(\text{Reverse}\left(\left\{ x_{[1:T]}^{(m)-} \right\}_m \right), \frac{\gamma}{\Delta t} \right)$

7 $\left\{ x_{[1:T]}^{(n)+} \right\}_n := \text{SDESolve}\left(\bar{u}_{\mathbb{P}_i}^+, \gamma, \pi_0, \Delta t \right)$

8 $\bar{u}_{\mathbb{Q}_i}^- := \text{DriftFit}\left(\text{Reverse}\left(\left\{ x_{[1:T]}^{(n)+} \right\}_n \right), \frac{\gamma}{\Delta t} \right)$

9 until *convergence*;

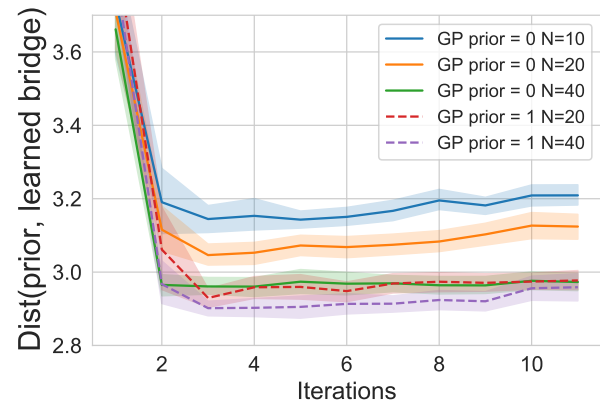**10 return** $\bar{u}_{\mathbb{Q}_i}^-, \bar{u}_{\mathbb{P}_i}^+$

Figure 3. Distance between the prior and the learned bridge as a function of the iterations, number of samples used (N) and whether the prior is used for the GP ($T = \frac{1}{\Delta t} = 100$).

Drift Estimation with Gaussian Processes: One can choose any parametric or non parametric model to carry out the DriftFit routine. In this section, we briefly introduce our implementation. Our DriftFit routine is based on the work in [12,23] that uses GPs to estimate the drift of SDEs from observations. We refer to this routine as GPDriftFit. We can restate the regression problems given by Equation (8) and its reversed counterpart in the following form:

$$\frac{x_{t_k}^{(n)-} - x_{t_k - \Delta t}^{(n)-}}{\Delta t} = u^+\left(x_{(t_k - \Delta t)}^{(n)-}, 1 - (t_k - \Delta t) \right) + \frac{\sqrt{\gamma}}{\Delta t}\epsilon,$$

$$\frac{x_{t_l - \Delta t}^{(m)+} - x_{t_l}^{(m)+}}{\Delta t} = -u^-\left(x_{t_l}^{(m)+}, 1 - t_l \right) + \frac{\sqrt{\gamma}}{\Delta t}\epsilon,$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbb{I}_d)$. Placing GP priors on the drift functions $u^- \sim \mathcal{GP}$ and $u^+ \sim \mathcal{GP}$, we arrive at standard multi-output GP formulation. Following [12,24], we assume the dimensions of the drift function are independent (assuming that the drift dimensions are decoupled is the least restrictive assumption as coupling them would impose a form of regularization). (equivalent to imposing a block-diagonal kernel matrix on a multi-output GP) and thus we fit a separate GP for each dimension. See Appendix E.1 for the specification of the predictive means.

Similar to [12,23], we are not making use of the predictive variances (using the predictive mean as an estimate for the drift can also be interpreted as a form of kernel ridge regression under the empirical risk minimisation framework). Instead, we simply use the predictive mean as an estimate of the drift and subsequently use that estimate to perform the EM method, thus effectively we could interpret this approach as a form of kernel ridge regression under the empirical risk minimisation framework.

Note that the advantage of the Bayesian GP interpretation of this procedure is that the GP posterior under certain conditions is naturally conjugate to the posterior of the SDE drift when modelled with a GP prior [23,25]. Furthermore, the Bayesian non-parametric formulation allows for the encoding of the prior drift function as the mean function of the GP prior (i.e., $u \sim \mathcal{GP}(b_0^{\mathbb{Q}_0^\gamma}, k)$). This becomes helpful when trying to sample unlikely paths and evaluating the fitted drift at samples that were not observed by the GP during training, here we want the fitted drift to fall back to the prior rather than to a Brownian motion given by a GP prior with 0 mean.

We now have the relevant ingredients to carry out IPML as specified in Figure 3. The computational cost of IPML with a Gaussian process is detailed in Appendix H.5. The

majority of the computation is spent fitting the GP in DriftFit at each iteration and scales as a function of the time discretization used as well as the size of π_0 and π_1.

4. Related Methodology

In this section we carry out a conceptual comparison with two pre-existing numerical approaches for solving the static SBP. While our goal is to solve the dynamic SBP, the solution of the static SBP can be used to construct that of the dynamic SBP [10], and this connection is central to our discussion. We would also like to highlight that an algorithm akin to IPML has been proposed concurrently and independently by [26], the main difference with our algorithm is that they estimate the drifts of the SDEs using neural networks score matching while we use using Gaussian processes and maximum-likelihood-based ideas.

4.1. Sinkhorn–Knop Algorithms

Within the machine learning community the static SBP with a Brownian motion prior is popularly known as entropic optimal transport [4]. In this formulation there are no trajectories and the empirical distributions are treated as discrete measures $[\boldsymbol{\pi_0}]_j = \frac{1}{N}$ and $[\boldsymbol{\pi_1}]_i = \frac{1}{M}$ thus the objective is given by $\min_{Q \in \mathcal{D}(\boldsymbol{\pi_0}, \boldsymbol{\pi_1})} \langle Q, C^{\mathbb{Q}_0^\gamma} \rangle + \gamma h(Q)$, where h is the discrete entropy, $\langle \cdot, \cdot \rangle$ computes the dot product between two matrices. The cost matrix C corresponds to the log transition density induced by the prior SDE: $C_{ij}^{\mathbb{Q}_0^\gamma} = \ln p^{\mathbb{Q}_0^\gamma}(\boldsymbol{y}_i | \boldsymbol{x}_j)$, which in the case of a Brownian motion prior reduces to a scaled Euclidean distance.

Once the problem has been discretized as described, the Sinkhorn–Knopp algorithm [3] can be applied directly to fit an optimal discrete transport map (and discrete SBP potentials $\phi, \hat{\phi}$) between the two distributions. Recent work [6] has showed empirical success in forming a continuous approximation of the SBP potentials using the log-sumexp formula [6]; however, it still remains to formally analyse the accuracy of the logsumexp potentials.

Estimating the cost matrix $C_{ij}^{\mathbb{Q}_0^\gamma}$ required by the Sinkhorn algorithm requires a mixture of both density estimation and further simulation of the prior. Furthermore, once we have the logsumexp potentials, additional high dimensional integrals must be estimated every time we wish to evaluate the optimal drift. Details are discussed in Appendix F. In addition, the Sinkhorn–Knopp algorithm still faces challenges in high dimensional spaces, specially for small values of γ [27]. Many proposed enhancements and literature [4,5] have focused on cost functions that implicitly require a Brownian motion prior and thus do not apply to our general setting.

4.2. Data-Driven Schrödinger Bridge (DDSB)

The method proposed by [10] is perhaps the most similar approach to our approach and consists of iterating two coupled density estimation objectives (see Appendix G) fitted at the marginals until convergence. While conceptually similar to our approach, there are three key differences. As with Sinkhorn–Knopp-based methods, their approach aims to solve the static SBP. Once converged, it requires further approximation to estimate the optimal drift. The coupled maximum likelihood formulation of the static half bridges in [10] is based on un-normalised density estimation with respect to the SBP potentials $\phi, \hat{\phi}$. The coupling of $\phi, \hat{\phi}$ in these objectives does not directly admit the application of modern methods in density estimation since it does not allow us to freely parametrise MLE estimators for the boundaries, and thus, neural density estimators such as [28–30] cannot be taken advantage of to circumvent the computation of the partition function. Regression problems are ubiquitous in machine learning (ML) and thus why we believe that this formulation can be very impactful as it allows us to leverage all these methods from ML. Experimentally, regression methods have been observed to scale better to high dimensional problems than density estimation methods, we observed similar evidence as we were unable to scale up the DDSB method beyond two dimensions.

Additionally, to compute the normalising term in the DDSB objective, we have to estimate a multidimensional integral that is not taken with respect to a probability distribution. This poses a difficult challenge in high dimensions. For a more detailed commentary, see Appendix G. Note that the method by [20] uses importance sampling to estimate these quantities which performs poorly beyond two dimensions.

It worth noting that we are interested in a method that can obtain a dynamic interpolation between two distributions. A downside of solving the static bridge either by Sinkhorn or [10] is that it does not directly provide us with an estimate of the optimal dynamics. In order to obtain an estimate of the optimal drift we require a series of approximations to estimate the integrals in Equations (A70) and (A71), thus every time we evaluate the optimal drift using these approaches we have to simulate the prior SDE $\mathcal{O}(N + M)$ times. This makes the runtime of obtaining the drift and dynamical interpolation expensive, since for each Euler step we take we have to simulate another SDE and backpropagate through it to evaluate the drift.

Finally we would like to highlight that a series of modern approaches to generative modelling [26,31–33] motivate both empirically and theoretically the gain in accuracy obtained in generative modelling tasks when using a dynamical approach rather than a static one.

5. Numerical Experiments

In this section, we demonstrate the capability of IPML to solve the Schrödinger bridge while efficiently incorporating priors on a range of different tasks from synthetic experiments to embryo cells (code supporting experiments can be found at https://github.com/franciscovargas/GP_Sinkhorn, accessed on the 22 August 2021).

5.1. Simple 1D and 2D Distribution Alignment

The first experiment considered is a simple alignment experiment where π_0 and π_1 are either unimodal or bimodal Gaussian distributions (see Appendix H.1 for exact details on the distributions). In Table 1 we compare the accuracy of the fitted marginals with our implementation of the Data-Driven Schrödinger Bridge (DDSB) by [10]. The scoring metrics used are the Earth mover's distance (EMD) as well as a Kolmogorov–Smirnov (KS) statistic on both sample sets.

Table 1. Performance comparison of the fitted marginals between DDSB and IPML on unimodal and bimodal experiment.

| | Unimodal | | | | Bimodal | | | |
| | π_0 | | π_1 | | π_0 | | π_1 | |
	KS	EMD	KS	EMD	KS	EMD	KS	EMD
DDSB	0.17	0.34	0.19	0.13	0.18	0.12	0.07	0.04
IPML	0.06	0.10	0.13	0.04	0.05	0.04	0.07	0.15

We were unable to get DDSB to work well when π_0 and π_1 were distant from each other. In these distant settings, DDSB collapses the mass of the marginals to a single data point (see Appendix H.1.1), as a result for the unimodal experiment we had to set $\gamma = 100$ for the DDSB approach to yield sensible results. We can observe IPML obtains better marginals overall and at a lower value of $\gamma = 1$.

We carried out 2D experiments with our approach to show the ability of our method to diffuse from a simple uni-modal distribution to a multi-modal distribution as in Figure 1 where our learned bridge successfully splits. Furthermore, we can visually observe how the forward and backwards trajectories are mirror images of each other as expected.

5.2. 2D Double Well Experiments

In the double well experiment, we illustrate how to incorporate an arbitrary functional prior and learn the distribution over paths connecting π_0 and π_1. In order to encode

prior information, we experiment with the potential well illustrated in Figures 1 and A3. The boundary distributions π_1, π_0 are taken to be Gaussian distributions centred at the centre of each well, respectively (See Appendix H.2 for the experiment specification and IPML parameters).

The motivation behind this experiment is to show that the SBP with this prior follows low energy (according to the well's potential function) trajectories for configurations of particles sampled at the wells. Intuitively, we can expect the learned trajectories to avoid the high energy peak located at $x = (0, 0)$ and go via the "passes" on either side. Note that if we estimated the optimal transport (OT) geodesics between π_0 and π_1 or similarly ran IPML with a Brownian motion prior, the learned optimal trajectories would go right through the middle, which is the highest energy path between both wells.

The prior is incorporated in the algorithm in two different ways, (1) by having the first drift $\mathbb{Q}_0^*$ to follow the negative derivative of the potential function $d\mathbf{x}(t) = -\nabla_{\mathbf{x}} U(\mathbf{x}) + \gamma d\boldsymbol{\beta}(t)$ and (2) by setting the mean function of the GP used to fit the drift as the negative derivative of the potential function. As illustrated in Figure 1, the learned trajectories between the two wells avoid the main high energy region and go through via lower energy passes, thus respecting the potential prior; however, the behaviour of the trajectory will differ depending on the choice of kernel, as illustrated in Figure A3, underlying the need for its careful consideration. We compared our approach to DDSB using the mean squared error distance to the prior in our evaluation. The results were **DDSB: 22.1, IPML (no prior): 19.9, IPML: (prior) 18.7**. As we can see, IPML considerably outperforms DDSB.

5.3. Finite Sample/Iteration Convergence

Theorem 1 provides us with asymptotic guarantees; however, it does not extend to the finite sample and discretisation case. To highlight the importance of finite effects on IPML, we carried out this analysis empirically. In Figure 3 we plot an empirical estimate of the error term in the control formulation of the SB (Equation (5)). This term is effectively the mean squared error between the learned drift and the prior drift (gradient field of the well). We analyse this metric for different values of N (number of samples) and Δ_t (discretization factor). We observe that IPML quickly reaches a low error valley, then, the cumulative error from the successive finite sample MLE can be observed and the drift starts to slowly deviate from the prior, this motivates early stopping. As N is increased, IPML achieves lower error faster and deviates less from the prior in later iterations. Finally, we observe placing the drift prior via the GP has a significant effect in improving the error and its convergence. Additionally in Appendix E we detail how this question could be approached from a theoretical perspective while underlining its significance and difficulty as illustrated by the lack of such analyses in related algorithms [10,22].

5.4. Single Cell—Embryo Body (EB) Data Set

We perform an experiment on an embryoid body scRNA-seq time course [34]. Single-cell RNA sequencing enables accurate identification of cells at specific time points; however, all cells are destroyed by measurement. This prevents modelling single-cell trajectory and instead we rely on modelling the data-manifold at discrete time points. The datasets consist of 5 time points illustrated in Figure A4. To evaluate the performance of the algorithms, we fit the models at the endpoints ($T = 1, 5$) and predict the intermediate frames. The metric used is the Earth mover's distance between the data at intermediate frames and the predicted distribution. We evaluate the performance at the endpoints by considering the prediction of the forward model at $T = 5$ and the backward model at $T = 1$.

We compare the performance with two methods. The first one, TrajectoryNet [34], uses continuous normalising flows with a soft constraint based on optimal transport. The second leverages the McCann interpolant [35,36] to interpolate the discrete OT solution; this corresponds to the linear interpolation induced by the transport map.

The results are summarised in Table 2. In most frames, IPML outperforms TrajectoryNet and performs similarly to OT. As the noise (volatility) in the trajectories goes

to 0, IPML theoretically converges to the OT solution with linear geodesics. So without any additional prior information, we would not expect IPML to outperform OT; however, when dealing with finite data, the DriftFit procedure allows for some non-linearity in the trajectories if it improves the fit. As a result, we do see some differences in the convex hull displayed in Figure 4 where we observe a better coverage of the single cell observations. We hypothesize this explains the improvement in performance vs. OT for frame 4. The performance of IPML may be further improved by incorporating domain-specific knowledge as a prior. It would outperform OT in those settings.

Table 2. Earth Moving Distance (EMD) on the EB data. EXP stands for the exponential kernel and EQ for exponentiated quadratic. The column "full" represents EMD averaged over all frames whereas "path" is averaged over the intermediate frames ($T = 2, 3, 4$).

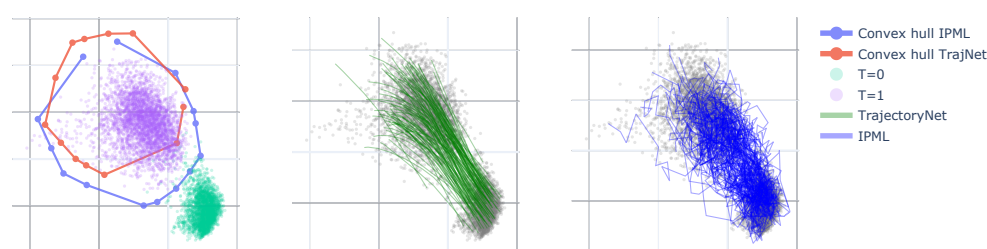| | T = 1 | T = 2 | T = 3 | T = 4 | T = 5 | Mean | |
						Path	Full
TrajectoryNet	0.62	1.15	1.49	1.26	0.99	1.30	1.18
IPML EQ	0.38	1.19	1.44	1.04	0.48	1.22	1.02
IPML EXP	0.34	1.13	1.35	1.01	0.49	1.16	0.97
OT	Na	1.13	1.10	1.11	Na	1.11	Na

Figure 4. The left plot represent the first and last frame of the cell's data and the convex hull of the forward model for TrajectoryNet and IPML. The right plots represent sample trajectories for Trajectorynet and IPML.

5.5. Motion Capture

In this experiment, we demonstrate how IPML can be used to model human motion from sensor data (data from The CMU Graphics Lab Motion Capture Database funded by the NSF (http://mocap.cs.cmu.edu, accessed on the 22 August 2021)). The motion corresponds to a basketball movement where the subject raises both arms simultaneously as illustrated in Figure 5 and each sensor corresponds to an oriented angle. We focus on modelling the right shoulder and elbow where the starting distribution corresponds to the leftmost image and the ending one the rightmost. This results in a 4-dimensional space as we model both position and velocity for each sensor. We compare the fit of IPML using a Brownian and a 2nd order linear ODE (Langevin) prior. The experimental details can be found in Appendix H.3. As illustrated in Figure 5, IPML is able to approximately model the dynamics using both priors. The Brownian prior displays noisy trajectories as expected in contrast to the Langevin prior that, by construction, smooths out the predicted positions (due to the 2nd order term). We observe that the Langevin prior approximates the step function nature of the true trajectory more closely than the Brownian prior, additionally, we can see that it also produces slightly better alignments.
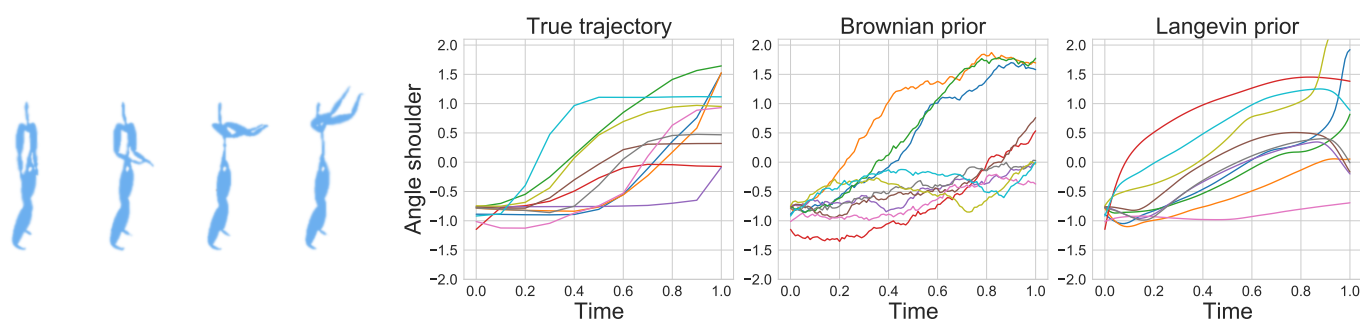
Figure 5. Left: 3D animation of the basketball signal motion modeled in Section 5.4 (motion from left to right). **Right**: Trajectories of the shoulder's oriented angle sensor through the motion. The two plots on the right demonstrate IPML's fit using a Brownian and Langevin prior.

6. Limitations and Opportunities

In this work, we propose to use GPs to estimate the drift; however, IPML could also be used with a parametric function estimator (e.g., a neural network). This could be useful with high-dimensional data where GPs may underperform. The main advantage of using a GP is its capacity to incorporate functional priors via the mean function. This can be useful in applications such as molecular dynamics where a potential function may be available. In contrast, implementing functional regularization in a neural network would require approximating a non-trivial high-dimensional integral to estimate the mean squared error between the parametric estimator and the functional prior.

A particular useful extension would be to adapt the SBP to work with more general forms of volatility functions that are not constant. This can be used, for example, in enforcing positivity constraints on a stochastic process via a geometric Brownian motion prior; this has applications in modelling biological signals such as transcription factors [37]. Work in this direction would require extending the theory of SBPs to non constant volatility functions.

Another promising setting is when multiple frames of data are available rather than just two boundary conditions. The IPFP algorithm trivially adapts to multiple constraints [21] rather than just initial and terminal distributions. Future work could explore experiments similar to the one presented in [34] where multiple frames are considered during training and the performance is measured using a leave-one-out procedure.

7. Conclusions

We have presented IPML, a method to solve the Schrödinger bridge for arbitrary diffusion priors. We presented theoretical results guaranteeing convergence in the limit of infinite data. We devised a practical application of the algorithm using Gaussian processes and presented several experiments on a variety of problems from synthetic to biological data. The approach opens up opportunities in science, where oftentimes, prior knowledge about the temporal evolution of a process has been developed but needs to be combined with data-driven methods to scale up to modern problems.

Assumptions Across Proofs

For abstraction purposes we will list the set of assumptions assumed across all results in this appendix:

- All SDEs considered have $L-$Lipchitz diffusion coefficients as well as satisfying linear growth
- The optimal drifts are elements of a compact space and thus satisfy the HJB equations. Note for the proposes of Theorem 1 this can be relaxed using notions of $\Gamma-$convergence,
- All SDEs satisfy a Fokker–Plank equation and thus $p(\boldsymbol{x}, t) = \mathrm{Law}(X_t)$ is differentiable with respect to $\boldsymbol{x}$,
- The boundary distributions are bounded in $\mathscr{L}_2(\pi_i)$ that is $\exists C < \infty$, *s.t.* $\mathbb{E}_{\pi_i}[|X|^2] \leq C$.

Appendix A. Brief Introduction to the Schrödinger System and Potentials

The Schrödinger system and its potentials are mentioned when introducing some of the results and connections of the full Schrödinger bridge problem. In this section, we will provide a brief introduction to how the system arises from the original static Schrödinger bridge. For brevity we will use $\boldsymbol{x} = \boldsymbol{x}(0)$ and $\boldsymbol{y} = \boldsymbol{x}(1)$ to denote the boundaries.

The static Schrödinger bridge problem can be derived from the full dynamic bridge by marginalising out the dynamics via the Disintegration Theorem and focusing on the problem only concerning what is happening at the boundaries. The static SBP is formulated as:

$$\underset{q(\boldsymbol{x},\boldsymbol{y})\in\mathcal{D}(\pi_0,\pi_1)}{\arg\min} \quad D_{\mathrm{KL}}(q||p^{\mathbb{Q}^\gamma}) \tag{A1}$$

where $p^{\mathbb{Q}^\gamma}(\boldsymbol{x}, \boldsymbol{y})$ is the prior joint distribution for the boundary and can be obtained by solving the FPK equation corresponding to the prior SDE. Now if we write down the Lagrangian for the above problem we arrive at:

$$D_{\mathrm{KL}}(q||p^{\mathbb{Q}^\gamma}) + \int \lambda(\boldsymbol{x})\left[\int q(\boldsymbol{x},\boldsymbol{y})d\boldsymbol{y} - \pi_0(\boldsymbol{x})\right]d\boldsymbol{x} + \int \mu(\boldsymbol{y})\left[\int q(\boldsymbol{x},\boldsymbol{y})d\boldsymbol{x} - \pi_1(\boldsymbol{y})\right]d\boldsymbol{y} \tag{A2}$$

which via performing the appropriate variations wrt to q leads to the optimal solution:

$$q^*(\boldsymbol{x},\boldsymbol{y}) = \exp\left(\ln p^{\mathbb{Q}^\gamma}(\boldsymbol{x}) - 1 - \lambda(\boldsymbol{x})\right)p^{\mathbb{Q}^\gamma}(\boldsymbol{y}|\boldsymbol{x})\exp(-\mu(\boldsymbol{y})), \tag{A3}$$

which when relabeling the terms containing the Lagrange multipliers we obtain:

$$q^*(\boldsymbol{x},\boldsymbol{y}) = \hat{\phi}(\boldsymbol{x})p^{\mathbb{Q}^\gamma}(\boldsymbol{y}|\boldsymbol{x})\phi(\boldsymbol{y}), \tag{A4}$$

such that:

$$\hat{\phi}(\boldsymbol{x})\int p^{\mathbb{Q}^\gamma}(\boldsymbol{y}|\boldsymbol{x})\phi(\boldsymbol{y})d\boldsymbol{y} = \pi_0(\boldsymbol{x}) \tag{A5}$$

$$\phi(\boldsymbol{y})\int p^{\mathbb{Q}^\gamma}(\boldsymbol{y}|\boldsymbol{x})\hat{\phi}(\boldsymbol{x})d\boldsymbol{x} = \pi_1(\boldsymbol{y}), \tag{A6}$$

Furthermore, if we relabel the potentials to indicate the times they correspond to we arrive at the Schrödinger system:

$$\phi(x,0)\hat{\phi}(x,0) = \pi_0(x), \quad \phi(y,1)\hat{\phi}(y,1) = \pi_1(y), \tag{A7}$$

where

$$\hat{\phi}(x,0) = \hat{\phi}(x), \quad \phi(y,1) = \phi(y), \tag{A8}$$

$$\phi(x,0) = \int p^{\mathbb{Q}^\gamma}(y|x)\phi(y)dy, \quad \hat{\phi}(y,1) = \int p^{\mathbb{Q}^\gamma}(y|x)\hat{\phi}(x)dx. \tag{A9}$$

The above functional system is refered to as the Schrödinger system and the Schrödinger potentials are given by $\phi, \hat{\phi}$. Furthermore the time interpolates for the potentials can be obtained by the propagations:

$$\phi(x,t) = \int \phi(z(1))p^{\mathbb{Q}_0^\gamma}(z(1)|x(t))dz(1) \tag{A10}$$

$$\hat{\phi}(y,t) = \int \hat{\phi}(z(0))p^{\mathbb{Q}_0^\gamma}(y(t)|z(0))dz(0) \tag{A11}$$

For a more rigorous and extensive introduction please see [20].

Appendix B. Disintegration Theorem—Product Rule for Measures

In this section, we present the Disintegration Theorem in the context of probability measures, which serves as the extension of the product rule to measures that do not admit the traditional product rule. Furthermore we will provide a proof for a direct lemma of the Disintegration Theorem that is more analogous to the standard product rule. Similar to the product rule these theorems are essential for decomposing and manipulating path measures and thus is needed for most results pertaining to the dynamic SBP.

Theorem A1. *(Disintegration Theorem for continuous probability measures):*
For a probability space $(Z, \mathcal{B}(Z), \mathbb{P})$, where Z is a product space: $Z = Z_x \times Z_y$ and

- $Z_x \subseteq \mathbb{R}^d, Z_y \subseteq \mathbb{R}^{d'}$,
- $\pi_i : Z \to Z_i$ *is a measurable function known as the canonical projection operator (i.e., $\pi_x(z_x, z_y) = z_x$ and $\pi_x^{-1}(z_x) = \{y | \pi_x(z_x) = z\}$),*

there exists a measure $\mathbb{P}_{y|x}(\cdot|x)$, such that

$$\int_{Z_x \times Z_y} f(x,y)d\mathbb{P}(y) = \int_{Z_x} \int_{Z_y} f(x,y)d\mathbb{P}_{y|x}(y|x)d\mathbb{P}(\pi^{-1}(x)), \tag{A12}$$

where $P_x(\cdot) = \mathbb{P}(\pi^{-1}(\cdot))$ is a probability measure, typically referred to as a push-forward measure, and corresponds to the marginal distribution.

A direct consequence of the above instance of the disintegration theorem is, with $f(x,y) = \mathbf{1}_{A_x \times A_y}(x,y)$,

$$\mathbb{P}(A_x \times A_y) = \int_{A_x} \mathbb{P}(A_y|x)d\mathbb{P}_x(x). \tag{A13}$$

We can see that, in the context of probability measures, the above is effectively analogous to the product rule.

RN Derivative Disintegration

We now have the required ingredients to show the following:

Lemma A1. *(RN-derivative product rule) Given two probability measures defined on the same product space,* $(Z_x \times Z_y, \mathcal{B}(Z_x \times Z_y), \mathbb{P})$ *and* $(Z_x \times Z_y, \mathcal{B}(Z_x \times Z_y), \mathbb{Q})$, *the Radon–Nikodym derivative* $\frac{d\mathbb{P}}{d\mathbb{Q}}(x, y)$ *can be decomposed as*

$$\frac{d\mathbb{P}}{d\mathbb{Q}}(x, y) = \frac{d\mathbb{P}_{y|x}}{d\mathbb{Q}_{y|x}}(y)\frac{d\mathbb{P}_x}{d\mathbb{Q}_x}(x). \tag{A14}$$

Proof. Note that this is an instructive sketch for a well known result [10,20,38,39]. We assume that $d\mathbb{Q}_{y|x} >> d\mathbb{P}_{y|x}$, which is not a trivial fact, yet it follows from the disintegration theorem in this particular setting (see Theorem 1.6 b in [39]) . Starting from

$$\mathbb{P}(A_x \times A_y) = \int_{A_x} \mathbb{P}(A_y|x)d\mathbb{P}_x(x),$$

we apply the Radon–Nikodym theorem to $\mathbb{P}(A_y|x)$ and then to P_x:

$$\mathbb{P}(A_x \times A_y) = \int_{A_x} \int_{A_y} \frac{d\mathbb{P}_{y|x}}{d\mathbb{Q}_{y|x}}(y)d\mathbb{Q}_{y|x}(y)d\mathbb{P}_x(x)$$

$$= \int_{A_x} \left(\int_{A_y} \frac{d\mathbb{P}_{y|x}}{d\mathbb{Q}_{y|x}}(y)d\mathbb{Q}_{y|x}(y) \right) \frac{d\mathbb{P}_x}{d\mathbb{Q}_x}(x)d\mathbb{Q}_x(x)$$

$$= \int_{A_x} \int_{A_y} \frac{d\mathbb{P}_x}{d\mathbb{Q}_x}(x)\frac{d\mathbb{P}_{y|x}}{d\mathbb{Q}_{y|x}}(y)d\mathbb{Q}_{y|x}(y)d\mathbb{Q}_x(x).$$

Now, via the disintegration we have that

$$\int_{A_x \times A_y} \frac{d\mathbb{P}_x}{d\mathbb{Q}_x}(x)\frac{d\mathbb{P}_{y|x}}{d\mathbb{Q}_{y|x}}(y)d\mathbb{Q}(x, y) = \int_{A_x} \int_{A_y} \frac{d\mathbb{P}_x}{d\mathbb{Q}_x}(x)\frac{d\mathbb{P}_{y|x}}{d\mathbb{Q}_{y|x}}(y)d\mathbb{Q}_{y|x}(y)d\mathbb{Q}_x(x).$$

Thus, we conclude that

$$\mathbb{P}(A_x \times A_y) = \int_{A_x \times A_y} \frac{d\mathbb{P}_x}{d\mathbb{Q}_x}(x)\frac{d\mathbb{P}_{y|x}}{d\mathbb{Q}_{y|x}}(y)d\mathbb{Q}(x, y),$$

which, via the Radon–Nikodym theorem, implies

$$\frac{d\mathbb{P}}{d\mathbb{Q}}(x, y) = \frac{d\mathbb{P}_{y|x}}{d\mathbb{Q}_{y|x}}(y)\frac{d\mathbb{P}_x}{d\mathbb{Q}_x}(x).$$

Note that swapping $\mathbb{Q}$ for the Lebesgue measure would result in the standard product rule for probability density functions. $\square$

Appendix C. Proof Sketches for Half Bridges

In this section we provide a proof sketch for the closed form solution of the half bridges as well as a proof for Observation 1.

Theorem A2. *The forward half bridge admits the following solution:*

$$\mathbb{P}^*\left(A_0 \times A_{(0,1]} \right) = \int_{A_0 \times A_{(0,1]}} \frac{d\pi_0}{d\pi_0^\mathbb{Q}}d\mathbb{Q}. \tag{A15}$$

Proof. Via the disintegration theorem, we have the following decomposition of KL:

$$D_{\mathrm{KL}}(\mathbb{P}||\mathbb{Q}) = D_{\mathrm{KL}}(\pi_0^\mathbb{P}||\pi_0^\mathbb{Q}) + \mathbb{E}_{\pi_0^\mathbb{P}}[D_{\mathrm{KL}}(\mathbb{P}(\cdot|x)||\mathbb{Q}(\cdot|x))].$$

Thus, via matching terms accordingly, we can construct $\mathbb{P}^*$ by setting $\mathbb{P}(\cdot|x) = \mathbb{Q}(\cdot|x)$ and matching the constraints:

$$\mathbb{P}^*\left(A_0 \times A_{(0,1]}\right) = \int_{A_0} \mathbb{Q}(A_{(0,1]}|x) d\pi_0(x), \tag{A16}$$

$$\mathbb{P}^* = \int_{A_0} \frac{d\pi_0}{d\pi_0^{\mathbb{Q}}}(x) \mathbb{Q}(\cdot|x) d\pi_0^{\mathbb{Q}}(x)$$

$$= \int_{A_0 \times A_{(0,1]}} \frac{d\pi_0}{d\pi_0^{\mathbb{Q}}}(x) d\mathbb{Q}. \tag{A17}$$

□

Observation A1. *We can parametrise a measure $\mathbb{Q}$ with its drift as the solution to either of the following SDEs:*

$$d\mathbf{x}^{\pm}(t) = \mathbf{b}^{\pm}(t) + \sqrt{\gamma} d\boldsymbol{\beta}^{\pm}(t),$$
$$\mathbf{x}^+(0) \sim \pi_0^{\mathbb{Q}}, \, \mathbf{x}(1) \sim \pi_1^{\mathbb{Q}}.$$

Then, we can sample from the solution to the following half bridges:

$$\mathbb{P}^{*-} = \arg\inf_{\mathbb{P} \in \mathcal{D}(\cdot, \pi_1)} D_{\mathrm{KL}}(\mathbb{P}||\mathbb{Q}),$$
$$\mathbb{P}^{*+} = \arg\inf_{\mathbb{P} \in \mathcal{D}(\pi_0, \cdot)} D_{\mathrm{KL}}(\mathbb{P}||\mathbb{Q}),$$

via simulating trajectories following the SDEs

$$d\mathbf{x}^-(t) = \mathbf{b}^-(t) + \sqrt{\gamma} d\boldsymbol{\beta}^-(t), \quad \mathbf{x}(1) \sim \pi_1,$$
$$d\mathbf{x}^+(t) = \mathbf{b}^+(t) + \sqrt{\gamma} d\boldsymbol{\beta}^+(t), \quad \mathbf{x}(0) \sim \pi_0.$$

Paths sampled from the above SDEs will be distributed according to $\mathbb{P}^{*-}$ and $\mathbb{P}^{*+}$, respectively.

Proof. (Sketch) W.l.o.g., consider the decomposition of the KL divergence that follows from the disintegration Theorem (Appendix B):

$$D_{\mathrm{KL}}(\mathbb{P}||\mathbb{Q}) = D_{\mathrm{KL}}(\pi_0^{\mathbb{P}}||\pi_0^{\mathbb{Q}}) + \mathbb{E}_{\pi_0^{\mathbb{P}}}[D_{\mathrm{KL}}(\mathbb{P}(\cdot|x)||\mathbb{Q}(\cdot|x))].$$

Furthermore, the disintegration $\mathbb{Q}(\cdot|\mathbf{x}(0))$ is a solution to the dynamics $d\mathbf{x}^+(t) = \mathbf{b}^+(t) + \sqrt{\gamma} d\boldsymbol{\beta}^+(t)$. We can make the term $\mathbb{E}_{\pi_0^{\mathbb{P}}}[D_{\mathrm{KL}}(\mathbb{P}(\cdot|x)||\mathbb{Q}(\cdot|x))]$ go to 0 by setting $\mathbb{P}(\cdot|\mathbf{x}(0)) = \mathbb{Q}(\cdot|\mathbf{x}(0))$, it is clear the dynamics of $\mathbb{P}(\cdot|\mathbf{x}(0))$ follows $d\mathbf{x}^+(t) = \mathbf{b}^+(t) + \sqrt{\gamma} d\boldsymbol{\beta}^+(t)$. What is left is to attach the constraint via $\mathbf{x}(0) \sim \pi_0(\mathbf{x}(0))$ which brings $D_{\mathrm{KL}}(\pi_0^{\mathbb{P}}||\pi_0^{\mathbb{Q}})$ to $D_{\mathrm{KL}}(\pi_0||\pi_0^{\mathbb{Q}}) = D_{\mathrm{KL}}(\mathbb{P}^{*+}||\mathbb{Q})$ coinciding with the half bridge minima as per [20,22]. This is simply enforcing the constraints via the product rule (Disintegration Theorem) and then matching the remainder of the unconstrained interval with the disintegration for the reference distribution $\mathbb{Q}$, following Equation (7). □

Appendix D. Proof Sketch for Reverse-MLE Consistency

The proof sketch for Theorem 1 will show how the likelihood converges in the large data and small time step limit to an optimisation of the KL divergence between two reverse time diffusions, from here one can use the standard arguments to show this quantity is minimised when the two measures describe the same stochastic process or equivalently when the drifts are equal.

Lemma A2. *Normalizing the time reversed likelihood with the true discretised backwards SDE density does not affect the maximum likelihood estimate. That is:*

$$\underset{\boldsymbol{u}}{\arg\max}\sum_n \sum_{k=0}^T \ln \mathcal{N}\left(\boldsymbol{x}_{t_k-\Delta t}^{(n)+}\middle|\boldsymbol{\mu}_{t_k-\Delta t}^{(n)+},\gamma\Delta t\right)$$

$$= \underset{\boldsymbol{u}}{\arg\max}\frac{1}{N}\sum_n \sum_{k=0}^T \left(\ln \mathcal{N}\left(\boldsymbol{x}_{t_k-\Delta t}^{(n)+}\middle|\boldsymbol{\mu}_{t_k-\Delta t}^{(n)+},\gamma\Delta t\right) - \ln \mathcal{N}\left(\boldsymbol{x}_{t_k-\Delta t}^{(n)+}\middle|\boldsymbol{x}_{t_k}^{(n)+} - \Delta t\boldsymbol{u}_0^-\left(\boldsymbol{x}_{t_k}^{(n)+},1-t_k\right),\gamma\Delta t\right)\right) \tag{A18}$$

Proof. The term $-\ln \mathcal{N}\left(\boldsymbol{x}_{t_k-\Delta t}^{(n)+}\middle|\boldsymbol{x}_{t_k}^{(n)+} - \Delta t\boldsymbol{u}_0^-\left(\boldsymbol{x}_{t_l}^{(n)+},1-t_k\right),\gamma\Delta t\right)$ does not depend on $\boldsymbol{u}$ and thus is an additive constant. $\square$

Theorem A3. *(Consistency of Reverse-MLE Formulations) Let $\left\{\left(\boldsymbol{x}_{t_k}^{(n)+}\right)_{k=0}^T\right\}_{n=0}^N$ be sampled/ discretised trajectories from the SDE that represents the half bridge measure $\mathbb{P}^+$:*

$$d\mathbf{x}^+(t) = \boldsymbol{u}_0^+(\mathbf{x}^+(t),t) + \sqrt{\gamma}d\boldsymbol{\beta}^+(t), \quad \mathbf{x}(0) \sim \pi_0.$$

Then carrying out maximum likelihood estimation of the time reversed drift $\boldsymbol{u}_0^-$ on time reversed samples:

$$\prod_n p\left(\left(\boldsymbol{x}_{t_k}^{(n)+}\right)_{k=0}^T\middle|\boldsymbol{u}^-\right) \propto \tag{A19}$$

$$\prod_n p(\boldsymbol{x}_{t_T}^{(n)+})\prod_{k=0}^T \mathcal{N}\left(\boldsymbol{x}_{t_k-\Delta t}^{(n)+}\middle|\boldsymbol{\mu}_{t_k-\Delta t}^{(n)+},\gamma\Delta t\right),$$

where,

$$\boldsymbol{\mu}_{t_k-\Delta t}^{(n)+} = \boldsymbol{x}_{t_k}^{(n)+} - \Delta t\boldsymbol{u}^-\left(\boldsymbol{x}_{t_k}^{(n)+},1-t_k\right) \tag{A20}$$

and $\boldsymbol{u}^-$ is the drift of our estimator SDE:

$$d\mathbf{x}^-(t) = \boldsymbol{u}^-(\mathbf{x}^-(t),t) + \sqrt{\gamma}d\boldsymbol{\beta}^-(t), \quad \mathbf{x}^-(t)(0) \sim p, \tag{A21}$$

converges in probability to the to the true dual drift $\boldsymbol{u}_0^-(\boldsymbol{x},t) = \boldsymbol{u}_0^+(\boldsymbol{x},t) - \gamma\nabla_{\boldsymbol{x}}\ln p(\boldsymbol{x},t)$ where $\boldsymbol{u}_0^+(x,t)$ is the optimal half bridge drift for $\mathbb{P}^+$ as $N \to \infty, \Delta t \to 0$. Under the assumptions:

- *The density $p(\boldsymbol{x},t)$ is differentiable with respect to $\boldsymbol{x}$,*
- *The optimal drift lies in a compact space (this can be relaxed using notions of $\Gamma-$convergence),*
- *The prior drift coefficient is $L-$Lipchitz and satisfies linear growth.*

Proof. For the interest of brevity let $\boldsymbol{u}_t^{(n)-} = \boldsymbol{u}^-\left(\boldsymbol{x}_{t_k}^{(n)+},1-t_k\right)$. Taking logs and applying Lemma A2 to Equation (A19) yields:

$$\frac{1}{N}\sum_n \sum_{k=0}^T \left(\ln \mathcal{N}\left(\boldsymbol{x}_{t_k-\Delta t}^{(n)+}\middle|\boldsymbol{\mu}_{t_k-\Delta t}^{(n)+},\gamma\Delta t\right) - \ln \mathcal{N}\left(\boldsymbol{x}_{t_k-\Delta t}^{(n)+}\middle|\boldsymbol{x}_{t_k}^{(n)+} - \Delta t\boldsymbol{u}_0^-\left(\boldsymbol{x}_{t_k}^{(n)+},1-t_k\right),\gamma\Delta t\right) + \ln\frac{p(\mathbf{x}_{t_T})}{q(\mathbf{x}_{t_T})}\right) \tag{A22}$$

$$\frac{1}{N}\sum_n \frac{1}{2\gamma\Delta t}\sum_{k=0}^T \left(-\left\|\boldsymbol{x}_{t_k-\Delta t}^{(n)+} - (\boldsymbol{x}_{t_k}^{(n)+} - \Delta t\boldsymbol{u}_t^{(n)-})\right\|^2 + \left\|\boldsymbol{x}_{t_k-\Delta t}^{(n)+} - (\boldsymbol{x}_{t_k}^{(n)+} - \Delta t\boldsymbol{u}_{0t}^{(n)-})\right\|^2 + \ln\frac{p(\mathbf{x}_{t_T})}{q(\mathbf{x}_{t_T})}\right) \tag{A23}$$

where $q(\mathbf{x}_{t_T})$ represents the terminal distribution of the forward SDE (i.e., $\mathbf{x}^+(1) \sim \pi$). Now we can equivalently write the above expression in terms of the time reversed samples (i.e., $\boldsymbol{x}_{t_i}^- = \boldsymbol{x}_{t_{n-i}}^+$):

$$\frac{1}{N}\sum_n \frac{1}{2\gamma\Delta t}\sum_{k=0}^T \left(-\left\|\boldsymbol{x}_{t_k}^{(n)-} - (\boldsymbol{x}_{t_k-\Delta t}^{(n)-} - \Delta t\boldsymbol{u}_t^{(n)-})\right\|^2 + \left\|\boldsymbol{x}_{t_k}^{(n)-} - (\boldsymbol{x}_{t_k-\Delta t}^{(n)} - \Delta t\boldsymbol{u}_{0t}^{(n)-})\right\|^2 + \ln\frac{p(\mathbf{x}_{t_0}^{(n)-})}{q(\mathbf{x}_{t_0}^{(n)-})}\right) \tag{A24}$$

Expanding the squares and re-arranging:

$$\frac{1}{N}\sum_n \frac{1}{\gamma}\sum_{k=0}^{T}\left((\Delta x^{(n)-})^{\top}(u_{1-t}^{(n)-} - u_{0(1-t)}^{(n)-}) - \frac{1}{2}(||u_{1-t}^{(n)-}||^2 - ||u_{0(1-t)}^{(n)-}||^2)\Delta t + \ln\frac{p(x_{t_0}^{(n)-})}{q(x_{t_0}^{(n)-})}\right), \qquad (A25)$$

where $\Delta x^{(n)-} = x_{t_k}^{(n)-} - x_{t_k-\Delta t}^{(n)-}$. Now we can consider the limit of Equation (A25):

$$\lim_{N\to\infty}\lim_{\Delta t\to 0}\frac{1}{N}\sum_n \frac{1}{\gamma}\sum_{k=0}^{T}\left((\Delta x^{(n)-})^{\top}(u_{1-t}^{(n)-} - u_{0(1-t)}^{(n)-}) - \frac{1}{2}(||u_{1-t}^{(n)-}||^2 - ||u_{0(1-t)}^{(n)-}||^2)\Delta t + \ln\frac{p(x_{t_0}^{(n)-})}{q(x_{t_0}^{(n)-})}\right), \qquad (A26)$$

$$\lim_{N\to\infty}\frac{1}{N}\sum_n \frac{1}{\gamma}\lim_{\Delta t\to 0}\sum_{k=0}^{T}\left((\Delta x^{(n)-})^{\top}(u_{1-t}^{(n)-} - u_{0(1-t)}^{(n)-}) - \frac{1}{2}(||u_{1-t}^{(n)-}||^2 - ||u_{0(1-t)}^{(n)-}||^2)\Delta t + \ln\frac{p(x_{t_0}^{(n)-})}{q(x_{t_0}^{(n)-})}\right), \qquad (A27)$$

We can write the sum over the time grid as a stochastic integral if we express the inner terms using the continuous time approximation from Lemma A3:

$$\lim_{N\to\infty}\frac{1}{N}\sum_n \frac{1}{\gamma}\lim_{\Delta t\to 0}\int_0^1 (u_{1-t}^{(n)-} - u_{0(1-t)}^{(n)-})^{\top}d\hat{x}^{(n)-}(t) - \frac{1}{2}\int_0^1(||u_{1-t}^{(n)-}||^2 - ||u_{0(1-t)}^{(n)-}||^2)dt + \ln\frac{p(x_{t_0}^{(n)-})}{q(x_{t_0}^{(n)-})}, \qquad (A28)$$

From [15,16,18] it follows that $\hat{x}^{(n)-}(t)$ is a semi-martingale (w.r.t. to the backwards filtration, see [17,40]) then in the limit the stochastic integrals are taken with respect to the true time reversed stochastic process adapted to the backwards filtration $(\mathcal{F}_i^-)_{i\in T}$ (see Theorem 2.13 in [41]):

$$\lim_{\Delta t\to 0}\int_0^1 (u_{1-t}^{(n)-} - u_{0(1-t)}^{(n)-})^{\top}d\hat{x}^{(n)-}(t) \xrightarrow{P} \int_0^1 (u^-(x^{(n)-}(t),t) - u_0^-(x^{(n)-}(t),t))^{\top}dx^{(n)-}(t) \qquad (A29)$$

$$\lim_{\Delta t\to 0}\int_0^1 (||u_{1-t}^{(n)-}||^2 - ||u_{0(1-t)}^{(n)-}||^2)dt \xrightarrow{P} \int_0^1 (||u^-(x^{(n)-}(t),t)||^2 - ||u_0^-(x^{(n)-}(t),t)||^2)dt, \qquad (A30)$$

now we have:

$$\lim_{N\to\infty}\frac{1}{N}\sum_n \frac{1}{\gamma}\int_0^1 (u^-(x^{(n)-}(t),t) - u_0^-(x^{(n)-}(t),t))^{\top}dx^{(n)-}(t)$$

$$-\frac{1}{2\gamma}\int_0^1(||u^-(x^{(n)-}(t),t)||^2 - ||u_0^-(x^{(n)-}(t),t)||^2)dt + \ln\frac{p(x^{(n)-}(0))}{q(x^{(n)-}(0))}, \quad (A31)$$

where by Lemma A3 each random function $x^{(n)-}(t)$ is sampled i.i.d from the SDE:

$$d\mathbf{x}^-(t) = u_0^-(\mathbf{x}^-(t),t) + \sqrt{\gamma}d\beta^+(t), \quad \mathbf{x}^-(0)\sim q, \qquad (A32)$$

we can now apply the weak law of large numbers (WLLN):

$$\mathbb{E}_{\mathbb{Q}}\left[\frac{1}{\gamma}\int_0^1 (u^-(x^{(n)-}(t),t) - u_0^-(x^{(n)-}(t),t))^{\top}dx^{(n)-}(t)\right.$$

$$\left.-\frac{1}{2\gamma}\int_0^1(||u^-(x^{(n)-}(t),t)||^2 - ||u_0^-(x^{(n)-}(t),t)||^2)dt + \ln\frac{p(x^{(n)-}(0))}{q(x^{(n)-}(0))}\right], \qquad (A33)$$

Now using that the log RN derivative between the estimator SDE and the true SDE is given by Girsanov's theorem [20,42,43]:

$$
\ln \frac{d\mathbb{Q}}{d\mathbb{P}_{\boldsymbol{u}}} = \frac{1}{\gamma} \int_0^1 (\boldsymbol{u}_0^-(\boldsymbol{x}^{(n)-}(t), t) - \boldsymbol{u}^-(\boldsymbol{x}^{(n)-}(t), t))^\top d\boldsymbol{x}^{(n)-}(t)
$$
$$
- \frac{1}{2\gamma} \int_0^1 (||\boldsymbol{u}_0^-(\boldsymbol{x}^{(n)-}(t), t)||^2 - ||\boldsymbol{u}^-(\boldsymbol{x}^{(n)-}(t), t)||^2) dt - \ln \frac{p(\mathbf{x}^{(n)-}(0))}{q(\mathbf{x}^{(n)-}(0))}, \quad (A34)
$$

we arrive at:

$$
\mathbb{E}_{\mathbb{Q}}\left[-\ln \frac{d\mathbb{Q}}{d\mathbb{P}_{\boldsymbol{u}}} \right] = -D(\mathbb{Q}||\mathbb{P}_{\boldsymbol{u}}). \quad (A35)
$$

Thus using $\frac{1}{N}\mathcal{L}(\boldsymbol{u})$ to denote the negative normalized log-likelihood we have shown the following pointwise (i.e., for $\boldsymbol{u} \in \mathcal{B}$) convergence in probability:

$$
\left| \frac{1}{N}\mathcal{L}(\boldsymbol{u}) - D(\mathbb{Q}||\mathbb{P}_{\boldsymbol{u}}) \right| \xrightarrow{P} 0. \quad (A36)
$$

Then if $\mathcal{B}$, is compact plus additional continuity and boundedness assumptions on $\mathcal{L}(\boldsymbol{u})$ a stronger form of uniform convergence in probability can be attained:

$$
\sup_{b \in \mathcal{B}} \left| \frac{1}{N}\mathcal{L}(\boldsymbol{u}) - D(\mathbb{Q}||\mathbb{P}_{\boldsymbol{u}}) \right| \xrightarrow{P} 0, \quad (A37)
$$

from the above it directly follows that:

$$
\left| \frac{1}{N}\mathcal{L}(\boldsymbol{u}_N^*) - D(\mathbb{Q}||\mathbb{P}_{\boldsymbol{u}_{N'}^*}) \right| \xrightarrow{P} 0, \quad (A38)
$$

where $\boldsymbol{u}_N^*$ is the MLE estimate of the dual drift at N samples. Which implies:

$$
\lim_{N \xrightarrow{P} \infty} D(\mathbb{Q}||\mathbb{P}_{\boldsymbol{u}_N^*}) = 0, \quad (A39)
$$

$$
\lim_{N \xrightarrow{P} \infty} \boldsymbol{u}_N^* = \boldsymbol{u}_0^- \quad (A40)
$$

See Chapter 4.5 of [44] for a more detailed discussion on the required assumptions of $\mathcal{L}$. □

Extending Theorem 1 to EM Samples

Note that the above proof holds for discretised samples from the original SDE, however we have been unable to extend it to approximate sampling schemes such as EM. We believe that the following result motivates the possibility that Theorem 1 holds when the samples are obtained from a consistent scheme such as EM:

Lemma A3. *(Convergence of discrete time reversal) The time reversal of discrete Euler-Mayurama samples:*

$$
\boldsymbol{x}_{t_i}^- = \boldsymbol{x}_{t_{n-i}}^+ \quad (A41)
$$

converges in probability to the solutions of the time reversed diffusion:

$$
d\mathbf{x}^-(t) = \mathbf{b}^-(t)dt + \sqrt{\gamma}d\boldsymbol{\beta}^-(t) \quad (A42)
$$

where $\mathbf{x}^-(t) = \mathbf{x}^+(1-t)$

Proof. First let's consider the continuous time step-wise approximation induced by the the EM samples:

$$\hat{\mathbf{x}}(t) = \sum_{i=1}^{T} x_{t_{i-1}} \delta_{t \in [t_{i-1}, t_i)},$$

From [45] it is a well known result that under the standard regularity assumptions on the drift $u(x, t)$ (i.e., Lipchitz continuity in t and and x) that the above approximation converges in probability to the SDE solution $\mathbf{x}^+(t)$ as $\Delta t \to 0$.

Observing that that the time reversal of the above corresponds to the approximation induced by the reverse samples we now consider the continuous time reversed approximation:

$$\hat{\mathbf{x}}^-(t) = \hat{\mathbf{x}}(1 - t) = \sum_{i=1}^{T} x_{t_{i-1}} \delta_{1 - t \in [t_{i-1}, t_i)},$$

Using $\hat{\mathbf{x}}(t)$ converges in probability to $\mathbf{x}^+(t)$ we have that $\forall t \in [0, 1]$, $\epsilon > 0$:

$$\lim_{\Delta t \to 0} P \left(\sup_{t \in [0,1]} \left| \hat{\mathbf{x}}(t) - \mathbf{x}^+(t) \right| > \epsilon \right) = 0, \tag{A43}$$

we carry out the following substitution $s = 1 - t$ then $\forall s \in [0, 1]$, $\epsilon > 0$:

$$\lim_{\Delta t \to 0} P \left(\sup_{t \in [0,1]} \left| \hat{\mathbf{x}}(1 - s) - \mathbf{x}^+(1 - s) \right| > \epsilon \right) = 0, \tag{A44}$$

$$\lim_{\Delta t \to 0} P \left(\sup_{t \in [0,1]} \left| \hat{\mathbf{x}}^-(s) - \mathbf{x}^-(s) \right| > \epsilon \right) = 0, \tag{A45}$$

□

Now it requires to be proved that $\hat{\mathbf{x}}^{(n)-}(t)$ is a semi-martingale which may follow from Lemma A3 as $\Delta t \to 0$. Another potential proof strategy would be to try and exploit the strong convergence properties of the Euler scheme to show that the stochastic integral from Theorem 1 also converges in the case of Euler samples. It may be the case that a correction term needs to be introduced in order for the integrals built from the reversed EM sample to converge to the backwards integrals. This remains an interesting question for future work.

Appendix E. Towards a Finite Sample Analysis of Approximate IPFP Schemes

We use the term approximate IPFP schemes for methodologies such as the one we present in this paper (i.e., IPML) where steps 5, 6 of IPFP algorithm (Figure 2) are replaced with inexact approximations. In this section we will present a rough sketch that takes the first step towards formally analysing approximate IPFP schemes in the finite sample/discretisation regime. This will serve to illustrate many of the challenges that still remain in this analysis as well as provide some initial results.

Conjecture A1. *(Heuristic Finite Sample Bound)*
Given that:

* *The exact IPFP at the ith iteration can be bounded from above as:*

$$\|u_*^+ - u_i^+\| \leq \frac{S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma}}{i} \tag{A46}$$

* *The approximate IPFP projection operators $\hat{\mathcal{P}}_\pi^\pm : \mathcal{X} \to \mathcal{X}$ are $K-$Lipchitz*

- *and the approximate finite sample projection error can be bounded by a constant:*

$$||u_i^+ - \hat{\mathcal{P}}_{\pi_{0,1}}[u_{i-1}^+]|| \le \epsilon_{\Delta T, N}, \tag{A47}$$

Then it follows that the approximate IPFP iteration error can be bounded from above by:

$$||u_*^+ - \hat{u}_i^+|| \le \epsilon_{\Delta T, N} i + \frac{S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma}}{i} \tag{A48}$$

when $\mathcal{P}_\pi^\pm$ are non-expansive operators ($K = 1$), and:

$$||u_*^+ - \hat{u}_i^+|| \le \epsilon_{\Delta T, N} \frac{1 - K^i}{1 - K} + \frac{S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma}}{i} \tag{A49}$$

when $K \ne 1$. Where $S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma} = \inf_{\mathbb{Q} \in \mathcal{D}(\pi_0, \pi_1)} D_{\mathrm{KL}}(\mathbb{Q}||\mathbb{Q}_0^\gamma)$, and $||.||$ is the $\mathscr{L}_2(\mathbb{Q}^*)$ norm error as per Equation (4), that is it is just the half bridge in Line 6 of the algorithm in Figure 2 written in terms of the drifts.

Proof. (Sketch) For notational simplicity we will relabel the half bridges in the algorithm of Figure 2 as projection operators $\mathcal{P}_\pi^\pm : \mathcal{X} \to \mathcal{X}$ in a function space $\mathcal{X}$. This leads to the following iterates:

$$u_i^- = \mathcal{P}_{\pi_1}^-[u_{i-1}^+] \tag{A50}$$
$$u_i^+ = \mathcal{P}_{\pi_0}^+[u_i^-] \tag{A51}$$

For the approximate IPFP we have the iterates:

$$\hat{u}_i^- = \hat{\mathcal{P}}_{\pi_1}^-[\hat{u}_{i-1}^+] \tag{A52}$$
$$\hat{u}_i^+ = \hat{\mathcal{P}}_{\pi_0}^+[\hat{u}_i^-]. \tag{A53}$$

Furthermore, to simplify the analysis further we will consider the composition of the two projection operators (i.e., $\mathcal{P}_{\pi_{0,1}} = \mathcal{P}_{\pi_0}^+ \circ \mathcal{P}_{\pi_1}^-$) combined into a single operator (for both exact and approximate projections):

$$u_i^+ = \mathcal{P}_{\pi_{0,1}}[u_{i-1}^+] = \mathcal{P}_{\pi_0}^+[\mathcal{P}_{\pi_1}^-[u_{i-1}^+]] \tag{A54}$$

We will now proceed to bound the error:

$$||u_*^+ - \hat{u}_i^+|| \le ||u_i^+ - \hat{u}_i^+|| + ||u_*^+ - u_i^+|| \tag{A55}$$

From the assumptions it follows that:

$$||u_*^+ - u_i^+|| \le \frac{S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma}}{i} \tag{A56}$$

yielding:

$$||u_*^+ - \hat{u}_i^+|| \le ||u_i^+ - \hat{u}_i^+|| + \frac{S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma}}{i}, \tag{A57}$$

Now we will proceed to analyse the first term:

$$||u_i^+ - \hat{u}_i^+|| \le ||u_i^+ - \hat{\mathcal{P}}_{\pi_{0,1}}[u_{i-1}^+]|| + ||\hat{\mathcal{P}}_{\pi_{0,1}}[u_{i-1}^+] - \hat{u}_i^+|| \tag{A58}$$
$$\le ||u_i^+ - \hat{\mathcal{P}}_{\pi_{0,1}}[u_{i-1}^+]|| + ||\hat{\mathcal{P}}_{\pi_{0,1}}[u_{i-1}^+] - \hat{\mathcal{P}}_{\pi_{0,1}}[\hat{u}_{i-1}^+]|| \tag{A59}$$
$$\le \epsilon_{\Delta T, N} + K||u_{i-1}^+ - \hat{u}_{i-1}^+|| \tag{A60}$$

Now we can expand the recurrence until the first iteration, yielding:

$$||\boldsymbol{u}_i^+ - \hat{\boldsymbol{u}}_i^+|| \le \sum_{j=0}^{i} \epsilon_{\Delta T,N} K^j \tag{A61}$$

When $K \ne 1$ we attain the bound:

$$||\boldsymbol{u}_i^+ - \hat{\boldsymbol{u}}_i^+|| \le \epsilon_{\Delta T,N} \frac{(1 - K^i)}{1 - K} \tag{A62}$$

When the projection operator $\hat{\mathcal{P}}_{\pi_{0,1}}$ is non expansive this gives the bound:

$$||\boldsymbol{u}_i^+ - \hat{\boldsymbol{u}}_i^+|| \le \sum_{j=0}^{i} \epsilon_{\Delta T,N} K^j \tag{A63}$$

$$\le \epsilon_{\Delta T,N} i \tag{A64}$$

$\square$

Note that the above sketch is more of a strategy to outline the challenges and steps required for this analysis rather than a proof itself. We believe that the formal finite sample analysis on approximate IPFP schemes merits its own separate work and is thus not the focus of this work; however, in order to highlight its importance we have provided a template towards this formal analysis following a strategy similar to [46]. This analysis allows us to understand what conditions are needed to be shown and how they will affect the finite sample convergence rates.

We will now proceed to discuss each of the assumptions that are required as elements of the above analysis strategy.

Assumption A1. *The exact IPFP at the ith iteration can be bounded from above as:*

$$||\boldsymbol{u}_*^+ - \boldsymbol{u}_i^+|| \le \frac{S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma}}{i} \tag{A65}$$

This assumption/conjecture is somewhat reasonable since Proposition 2.1 in [22] proves formally that the sum of the KL errors of the *i*th marginals are bounded by $\frac{S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma}}{i}$, furthermore via the disintegration theorem one can show that the half bridge error terms are in fact bounded by the sum of the error of the *i*th marginals [20]. Thus it seems feasible to extend the result in [22] to the path error we use in our conjecture.

Assumption A2. *The approximate IPFP projection operators $\hat{\mathcal{P}}_\pi^\pm : \mathcal{X} \to \mathcal{X}$ are $K-$Lipchitz. Specifically we are interested in the cases where they may be non-expansive and contractive.*

This assumption is specific to the approximation method used to estimate the half bridges. It requires going into rigorous details on the nature of the approximation and imposing suitable regularity assumptions. Not that ideally we want the operator to be contractive or at the very least not expansive since for Lipchitz constants greater than one the cumulative error would potentially accumulate exponentially.

It is interesting to note that in the contractive case we can attain the bound:

$$||\boldsymbol{u}_*^+ - \hat{\boldsymbol{u}}_i^+|| \le \frac{\epsilon_{\Delta T,N}}{1 - K} + \frac{S_{\pi_{0,1}}^{\mathbb{Q}_0^\gamma}}{i} \tag{A66}$$

where the error term is constant with respect to the iteration number. Additionally the bound in Equation (A48) also trivially applies.

Assumption A3. *The approximate finite sample projection error can be bounded by a constant:*

$$||\boldsymbol{u}_i^+ - \hat{\mathcal{P}}_{\pi_{0,1}}[\boldsymbol{u}_{i-1}^+]|| \le \epsilon_{\Delta T, N}, \tag{A67}$$

This should typically be something that one shows; however, as seen in [46], it is something that is often assumed in these proof strategies. The reason for assuming the result is that the projection varies depending on how we approximate the KL minimisation, in our case we combine our proposed IPML with drift estimation via Gaussian processes. The error of the drift estimation (i.e., $\epsilon_{\Delta T, N}$) via this approach is not straightforward to obtain and is still an area of active research see [23,25] for more details. As mentioned earlier this is not the focus of our work thus it is reasonable for us to assume that the drift estimation machinery proposed by prior authors is sound and should be possible to bound with a reasonable error.

Appendix E.1. GPDriftFit Implementation Details

In this section we provide the closed form predictive formulas used by our method to estimate the drift. Our GP formulation for the half bridge approximations yield the following mean per dimension d estimates of the drift:

$$[\bar{\boldsymbol{u}}^+(\boldsymbol{x},t)]_d = \text{vec}\big(\boldsymbol{k}_d^+(\boldsymbol{x}\oplus t)\big)^\top \Big(\tilde{\boldsymbol{K}}_d^+ + \frac{\gamma}{\Delta t}\mathbb{I}_{MT}\Big)^{-1}\text{vec}(\boldsymbol{Y}_d^+), \tag{A68}$$

$$[\boldsymbol{u}^-(\boldsymbol{x},t)]_d = \text{vec}\big(\boldsymbol{k}_d^-(\boldsymbol{x}\oplus t)\big)^\top \Big(\tilde{\boldsymbol{K}}_d^- + \frac{\gamma}{\Delta t}\mathbb{I}_{NT}\Big)^{-1}\text{vec}(\boldsymbol{Y}_d^-), \tag{A69}$$

where

$$[\tilde{\boldsymbol{K}}_d^\pm]_{i \cdot T+l, i' \cdot T+l'} = [\boldsymbol{K}_d^\pm]_{ii'll'},$$

$$[\boldsymbol{K}_d^+]_{mm'll'} = K_d\big(\boldsymbol{x}_{t_l-\Delta t}^{(m)-}\oplus(t_l-\Delta t), \boldsymbol{x}_{t_{l'}-\Delta t}^{(m')-}\oplus(t_{l'}-\Delta t)\big),$$

$$[\boldsymbol{K}_d^-]_{nn'kk'} = K_d\big(\boldsymbol{x}_{t_k}^{(n)+}\oplus t_k, \boldsymbol{x}_{t_{k'}}^{(n')+}\oplus t_{k'}\big),$$

$$[\boldsymbol{k}_d^+(\boldsymbol{x}\oplus t)]_{lm} = K_d\big(\boldsymbol{x}\oplus t, \boldsymbol{x}_{t_l-\Delta t}^{(m)-}\oplus(t_l-\Delta t)\big),$$

$$[\boldsymbol{k}_d^-(\boldsymbol{x}\oplus t)]_{kn} = K_d\big(\boldsymbol{x}\oplus t, \boldsymbol{x}_{t_k}^{(n)-}\oplus(t_k)\big),$$

$$[\boldsymbol{Y}_d^\pm]_{li} = \left[\frac{\pm\boldsymbol{x}_{t_l}^{(i)\mp} \mp \boldsymbol{x}_{t_l-\Delta t}^{(i)\mp}}{\Delta t}\right]_d,$$

where $\oplus$ is the concatenation operator, vec is the standard vectorisation operator, and $K_d : \mathbb{R}^{d+1} \times \mathbb{R}^{d+1} \to \mathbb{R}$ is a valid kernel function.

Appendix E.2. Coupled vs. Decoupled Drift Estimators

A careful reader may enquire if the decoupled drift parametrisation we have used imposes any limitations on the generality of the drifts that our method can estimate in contrast to an approach that couples/correlates drift dimensions. While it may seem counter-intuitive, correlating the GP outputs (i.e., coupling the drift dimensions) is more restrictive than the decoupled approach we have taken. This point is discussed and motivated in [47,48], where it is illustrated that the coupling of the drift outputs correspond to regularising the RKHS hypothesis space. So in fact the decoupled approach we take in this work is more flexible (less limited) than coupling the drift outputs as discussed in [47,48]. It is important to note that the regularisation effects of coupling may be desirable in certain physical systems where we wish to impose a constraint; however, this would come at a large computational cost. Finally note that our method parametrises the drift $[\boldsymbol{u}]_i = f_i(\boldsymbol{x},t)$ with a function that depends on all dimensions of the dynamical process $\mathbf{x}(t)$,

that is each drift coordinate depends on the entire input space and thus no simplifying assumptions have been made in this setting.

Appendix F. Estimates Required by the Sinkhorn–Knop Algorithm

In order to adapt the Sinkhorn–Knop algorithm to a general prior we require the following considerations:

- The Sinkhorn–Knop algorithm requires computing the cost $C_{ij}^{\mathbb{Q}_0^\gamma} = \ln p^{\mathbb{Q}_0^\gamma}(\boldsymbol{y}_i|\boldsymbol{x}_j)$. For more general SDE priors whose transition densities are not available in closed form can be challenging:
  - First we have to estimate the empirical distribution for $\ln p^{\mathbb{Q}_0^\gamma}(\boldsymbol{y}_i|\boldsymbol{x}_j)$ by generating multiple $\boldsymbol{x}(1)$ samples from the priors SDE for each $\boldsymbol{x}_j$ in the dataset:

$$d\mathbf{x}(t) = \boldsymbol{u}(\mathbf{x}(t),t)dt + \gamma d\boldsymbol{\beta}(t), \quad \mathbf{x}(0) = \boldsymbol{x}_j$$

  - Once we have samples $\{\mathbf{x}_k\}_j$ we must carry out density estimation for each j in order to evaluate $\ln p^{\mathbb{Q}_0^\gamma}(\boldsymbol{y}_i|\boldsymbol{x}_j)$.
- The Sinkhorn–Knop algorithm produces discrete potentials, which can be interpolated using the logsumexp formula; however, how do we go from these static potentials to time dependant trajectories?
  - From [10] we can obtain expressions for the optimal drift as a function of the time extended potentials:

$$\phi(\boldsymbol{x},t) = \int \phi(\boldsymbol{z}(1)) p^{\mathbb{Q}_0^\gamma}(\boldsymbol{z}(1)|\boldsymbol{x}(t)) d\boldsymbol{z}(1) \tag{A70}$$

$$\hat{\phi}(\boldsymbol{y},t) = \int \hat{\phi}(\boldsymbol{z}(0)) p^{\mathbb{Q}_0^\gamma}(\boldsymbol{y}(t)|\boldsymbol{z}(0)) d\boldsymbol{z}(0) \tag{A71}$$

The first integral can be approximated using the Montecarlo approximation via sampling from the prior SDE to draw samples from $p^{\mathbb{Q}_0^\gamma}(\boldsymbol{y}(t)|\boldsymbol{z}(0))$; however, computing the $\hat{\phi}$ potential integral is less clear and requires careful thought.
  - Note that in order to estimate the drift we would have to simulate the SDE prior every time we want to evaluate the drift, which itself will be run in another SDE simulation to generate optimal trajectories.

Appendix G. Approximations Required by DDSB

The approach in [10] iterates the following two objectives.

$$\hat{\beta}_i^* = \arg\max_{\hat{\beta}} \frac{1}{M} \sum_s \ln \hat{\phi}_0(\boldsymbol{x}_s;\hat{\beta}) \phi_0^{(i)}(\boldsymbol{x}_s) \tag{A72}$$

$$- \int \hat{\phi}_0(\boldsymbol{x};\hat{\beta}) \phi_0^{(i)}(\boldsymbol{x}) d\boldsymbol{x}, \quad \boldsymbol{x}_s \sim \pi_0(\boldsymbol{x}),$$

and

$$\beta_i^* = \arg\max_{\beta} \frac{1}{N} \sum_s \ln \phi_1(\boldsymbol{y}_s;\beta) \hat{\phi}_1^{(i)}(\boldsymbol{y}_s) \tag{A73}$$

$$- \int \phi_1(\boldsymbol{y};\beta) \hat{\phi}_1^{(i)}(\boldsymbol{y}) d\boldsymbol{y}, \quad \boldsymbol{y}_s \sim \pi_1(\boldsymbol{y}).$$

where $\hat{\phi}_0, \phi_1$ are parametric functions aimed at estimating the potentials, and:

$$\int \hat{\phi}_0(\boldsymbol{x};\hat{\beta}) \phi_0^{(i)}(\boldsymbol{x}) d\boldsymbol{x} = \int \phi_1(\boldsymbol{y};\beta) \hat{\phi}_1^{(i)}(\boldsymbol{y}) d\boldsymbol{y}$$

$$= \int \hat{\phi}_0(\boldsymbol{x};\hat{\beta}) \int p^{\mathbb{Q}_0^\gamma}(\boldsymbol{y}|\boldsymbol{x}) \phi_1^{(i)}(\boldsymbol{x};\beta) d\boldsymbol{y} d\boldsymbol{x}, \tag{A74}$$

Note that in Equation (A74) the outer most integral is not taken with respect to a probability distribution and thus does not admit standard approximations. The authors in [10] propose the method of importance sampling [49]; however, this does not scale well to higher dimensions. This is contrasting to our approach where all integrals are expectations with respect to the empirical distribution and the SDEs being fitted.
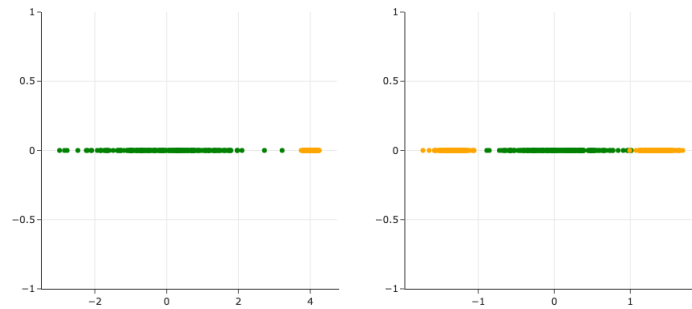

Figure A1. Unimodal experiment on the **left** and bimodal on the **right**.

Appendix H. Experiments

In this section we provide in depth details regarding the parameter configurations of our approach plus our experimental setup. For all the experiments below, unless stated otherwise, we used the exponential kernel with a lengthscale set at the default value of 1. We ran IPML for 5 iterations with a discretization factor $\Delta_t = 0.01$ and $\gamma = 1$. The iteration number (5) was selected by observing that the algorithm converged after that number in all experiments considered.

Appendix H.1. 1D Experiment

The prior $\mathbb{Q}_0^\gamma$ in IPM (the practical implementation of $\mathbb{Q}_0^\gamma$ will be added in the main paper during the revision period) is specified by a prior drift and an initial value distribution $\pi_0^{\mathbb{Q}_0^\gamma}$. We set $\pi_0^{\mathbb{Q}_0^\gamma} = \pi_0$ and the prior to a Brownian motion $\mathbb{W}^\gamma$ unless stated otherwise. The initial value distribution does not affect the argmax of the SBP and is thus a free parameter. Note that $\pi_0^{\mathbb{Q}_0^\gamma}$ is used in practice to estimate the backwards drift of the prior before entering the IPFP loop by first sampling from $\mathbb{Q}_0^\gamma$ using the SDESolve subroutine. It is possible to encourage more exploration of the space by increasing the variance of $\pi_0^{\mathbb{Q}_0^\gamma}$.

For the unimodal boundary distributions we used:

$$\pi_0 \sim \mathcal{N}(0,1), \ \ \pi_1 \sim \mathcal{N}\left(4, 0.1^2\right) \tag{A75}$$

and for the bimodal experiment we used:

$$\pi_0 \sim \mathcal{N}(0,1) \ \ \ \pi_1 \sim \frac{1}{2}\mathcal{N}(1.8, 0.6^2) + \frac{1}{2}\mathcal{N}(-1.9, 0.6^2). \tag{A76}$$

Appendix H.1.1. Delta Collapse in DDSB

Here we will illustrate a common failure case of the approach in [10]. That is when the distributions π_0 and π_1 are distant from each other the methodology proposed in [10] breaks for suitable values of γ (i.e., $\gamma = 1, 2, 3...$). Using the same marginals π_0, π_1 as in Experiment 1 we re-train the method by [10] with $\gamma = 1$. In Figure A2 we can see the results of this experiment. We can see the marginals learned by DDSB collapse all the mass in a very small low density region for the true marginals. We found this phenomena to occur often and harder to overcome in 2-dimensional experiments and thus did not pursue further comparisons with this method. Note we normalised the delta spike of the DDSB marginals to have a range in $[0,1]$ so we could compare to the true density, its actual value

is on the order of 10^{11}, which confirms that the model is collapsing to a point mass (dirac delta function).
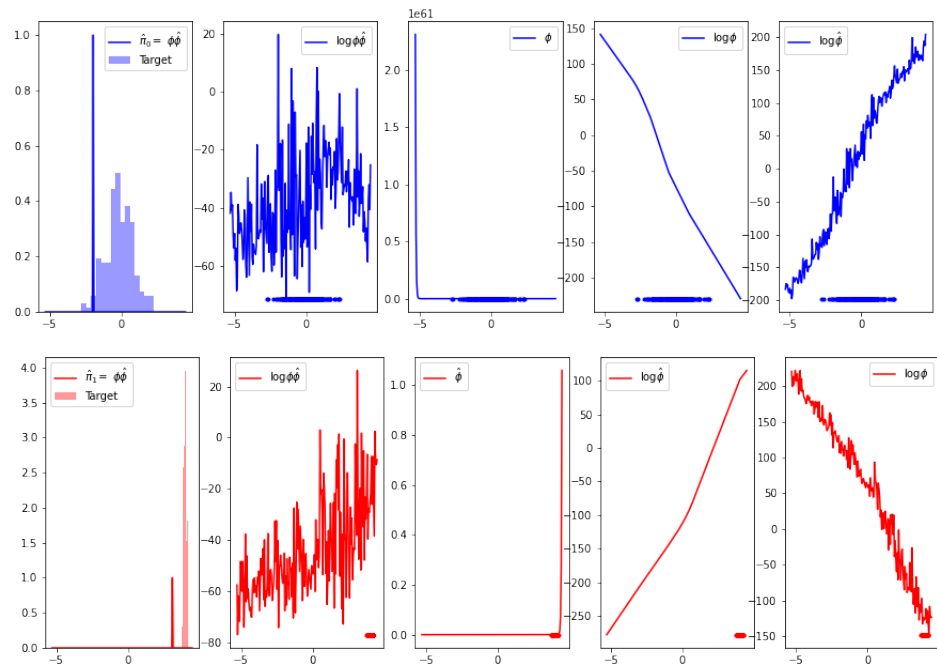

Figure A2. Schrödinger Bridge results using the method by [10] on unimodal to bimodal Gaussian 1D data. This example illustrates the Dirac delta collapse of the marginals.

Appendix H.2. Well Experiment

We used the following potential to model the double wells:

$$U(x,y) = \frac{5}{2}(x^2 - 1)^2 + y^2 + \frac{1}{\delta}\exp\left(-\frac{x^2 + y^2}{\delta}\right), \tag{A77}$$

furthermore, we used the boundary distributions:

$$\pi_0 \sim \mathcal{N}\left(\begin{pmatrix} -1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0.0125 & 0 \\ 0 & 0.15 \end{pmatrix}\right), \quad \pi_1 \sim \mathcal{N}\left(\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0.0125 & 0 \\ 0 & 0.15 \end{pmatrix}\right)$$

where we selected the boundary distributions to be located at the wells and visually follow a similar spread/curvature as the contour of the well. For these experiments we use $\gamma = 3$. As highlighted previously $\pi_0^{\mathbb{Q}_0^\gamma}$ is a free parameter, which in this section we set as:

$$\pi_0^{\mathbb{Q}_0^\gamma} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0.5 & 0 \\ 0 & 0.5 \end{pmatrix}\right), \tag{A78}$$

we select this as the marginal prior since its samples cover the space over which the wells is defined well, and helps estimate the backwards drift of the prior more accurately. As mentioned in the paper, using different kernels as well as lengthscales resulted in slightly different behaviour on the well experiment where the trajectory would not split as illustrated in Figure A3. The splitting pattern was obtained using an exponential kernel with a lower lengthscale (0.25) underlying the need for careful consideration of the kernel and its parameters.
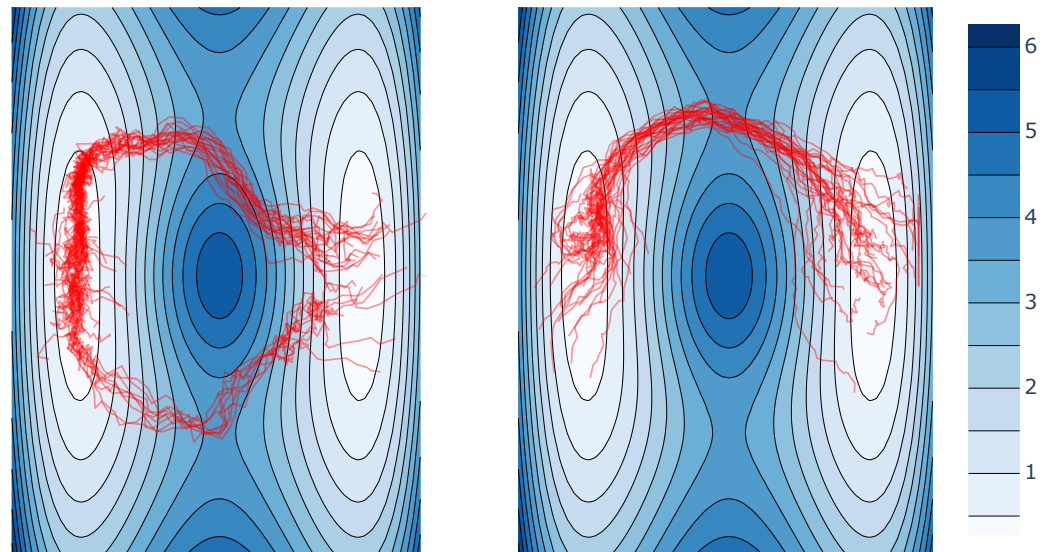
Figure A3. IPML forward trajectories on potential well. Exponentiated kernel on the **left** and exponential quadratic on the **right**.

Discussion on Kernel Choice

We can observe that when using EQ kernel, the trajectories do not cross both passes but instead choose one. We conjecture that this is due to the fitted predictive means having a preference for simpler functions that send nearby particles in the same direction. We empirically verify this conjecture by experimenting with alternative kernels to EQ (i.e., exponential) that model a wider class of functions. We believe that these alternative kernels are less prone to discouraging splitting since, when viewed within the kernel ridge regression framework, they span an RKHS of functions that do not have the smoothness constraint of the EQ kernel.

The data used for this experiment were taken directly from the GitHub repository published by [34]. PCA is first applied to the data and the first five components are selected. We used the code available at https://github.com/KrishnaswamyLab/TrajectoryNet/tree/master/TrajectoryNet (accessed on the 22 August 2021) to reproduce the performance of TrajectoryNet as well as the optimal transport baseline.

Appendix H.3. Motion Experiment

For this experiment we use a second order SDE prior (i.e., Langevin dynamics). This can be encoded in first order SDE using the companion/controllable form of a dynamical system:

$$d \begin{pmatrix} \mathbf{x}(t) \\ \mathbf{v}(t) \end{pmatrix} = \begin{pmatrix} \mathbf{v}(t) \\ \mathbf{K}\mathbf{x}(t) - \beta\mathbf{v}(t) \end{pmatrix} dt + \gamma \begin{pmatrix} \mathbf{0}_d \\ \mathbf{1}_d \end{pmatrix} d\boldsymbol{\beta}(t) \tag{A79}$$

For simplicity we set the dampening factor β to 0 and $\gamma = 0.5$. As for the spring matrix we explored the following decoupled and coupled forms:

$$\mathbf{K} = k\mathbb{I} \quad \mathbf{K} = k\frac{3}{2}\mathbb{I} - \frac{k}{2}\mathbf{1}\mathbf{1}^\top \tag{A80}$$

where k defines the frequency of the solution via $\omega = \sqrt{k/m}$ in our setting we ignore the mass (imagine is at absorbed in the matrix $\mathbf{K}$ and work directly with setting ω. Since the movement we select moves from rest to a high position and does not return to its starting point/or repeat itself we set $\omega \approx \pi$ that corresponds to 0.5 Hz which describes

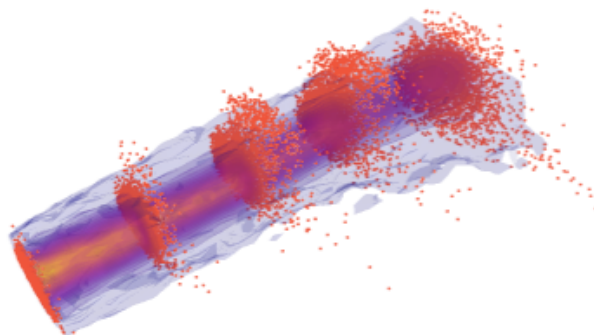half a revolution. We visually confirmed with the true trajectories that this seems like a reasonable prior.

Figure A4. Fitted SBP distribution on the start and end data slices with Brownian prior on single-cell human Embryo data from [34]. Observations depicted as red point clouds. See Section 5.3 for experimental details.

Appendix H.4. Cell Experiment

For the Brownian motion prior that we use as a baseline comparison in Appendix H.3 we explored several values of γ and selected the best performing one (visually) of $\gamma = 0.3$. Note we do pick a smaller value of γ for the Brownian motion and this is partly to compensate that the Langevin prior trajectories will always look smooth by construction thus we pick a not so noisy Brownian motion that still gives good results.

Note that the volatility term here is singular and in our mocap application setting it is exactly the vector:

$$\gamma \begin{pmatrix} \mathbf{0}_d \\ \mathbf{1}_d \end{pmatrix} \tag{A81}$$

where the dimension d is given by the number of sensors we use to fit the motion. A careful reader may wonder weather the SBP machinery still applies to such a sparsely structured diffusion and in fact it does. Following Theorem 4 of [43] we can see how the Radon–Nikodym derivative is finite and can be expressed in an almost identical fashion to the original controlled SBP formulation and thus justifying the existence of the SBP in this scenario as well as the application of the IPML algorithm and its derivatives.
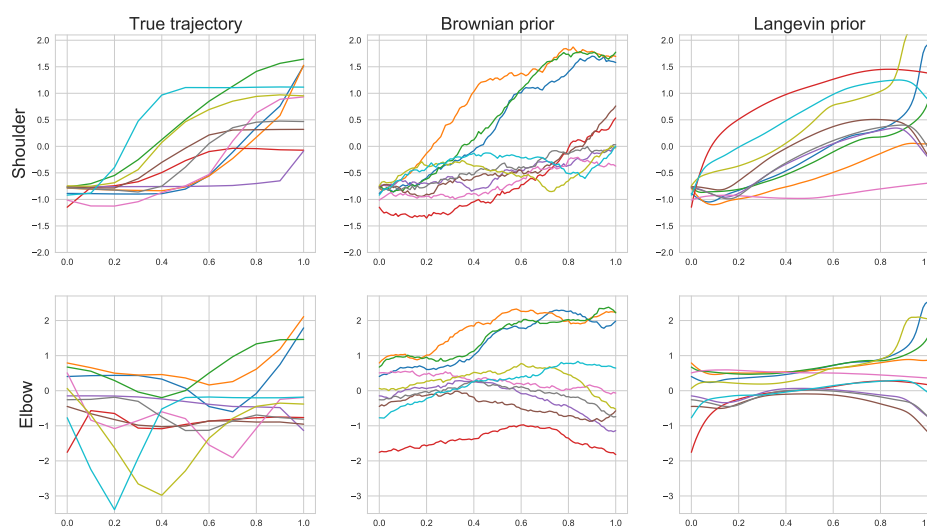

Figure A5. Trajectories of the shoulder and arm oriented angle sensor through the motion. The two plots on the right demonstrate IPML's fit using a Brownian and Langevin prior.

Appendix H.5. Computational Resources

Appendix H.5.1. Computational Costs

The main computational cost of IPML comes from fitting a Gaussian process to model the drift (DriftFit). Assuming the cardinality of $|\pi_0| = |\pi_1| = N$ and a discretization factor Δ_t, at each iteration, IPML requires fitting a Gaussian process on $\frac{N}{\Delta_t}$ samples. It is a well-known fact that GPs have a cubic time complexity making the costs of the DriftFit subroutine $O((\frac{N}{\Delta_t})^3)$. This could be scaled down by using GP approximations (Nystrom) and was successful in our preliminary results. The computational costs of the SDESolve sub-routine in comparison is $O((\frac{N}{\Delta_t})^2)$ due to the GP predictions costs. Finally, the memory complexity of the algorithm is $O((\frac{N}{\Delta_t})^2)$ due to the GP fitting in DriftFit. The running time of IPML on the machine described in the section below for the well experiment is around 5 min. For the cell experiment, the running time is between 1 and 2 h depending on the discretization factor chosen.

Appendix H.5.2. Infrastructure Used

The experiments were performed on Compute Canada clusters. Specifically, they ran on a CPU cluster composed of Intel CPUs. Each node contains 20 Intel Skylake cores (2.4 GHz, AVX512), for a total of 40 cores per node and 202 GB of RAM. The computational costs for the toy experiments (Well and 1D) was quite small; however, the RAM consumption of the cell experiment was significant due to the use of a Gaussian process.

References

1. Schrödinger, E. *Uber die Umkehrung der Naturgesetze*; Akademie der Wissenschaften: Berlin, Germany, 1931; pp. 144–153.
2. Schrödinger, E. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. *Annales de l'Institut Henri Poincaré* **1932**, *2*, 269–310.
3. Sinkhorn, R.; Knopp, P. Concerning nonnegative matrices and doubly stochastic matrices. *Pac. J. Math.* **1967**, *21*, 343–348. [CrossRef]
4. Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. In Proceedings of the 26th International Conference on Neural Information Processing Systems, Lake Tahoe, NV, USA, 5–10 December 2013.
5. Feydy, J.; Séjourné, T.; Vialard, F.X.; Amari, S.I.; Trouvé, A.; Peyré, G. Interpolating between optimal transport and MMD using Sinkhorn divergences. In Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics, Okinawa, Japan, 16 April 2019; pp. 2681–2690.
6. Chizat, L.; Roussillon, P.; Léger, F.; Vialard, F.X.; Peyré, G. Faster Wasserstein Distance Estimation with the Sinkhorn Divergence. In Proceedings of the 2020 Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 6–12 December 2020; Volume 33.
7. Hennig, P.; Osborne, M.A.; Girolami, M. Probabilistic numerics and uncertainty in computations. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **2015**, *471*, 20150142. [CrossRef]
8. Kullback, S. Probability densities with given marginals. *Ann. Math. Stat.* **1968**, *39*, 1236–1243. [CrossRef]
9. Ruschendorf, L. Convergence of the iterative proportional fitting procedure. *Ann. Stat.* **1995**, *23*, 1160–1174. [CrossRef]
10. Pavon, M.; Tabak, E.G.; Trigila, G. The Data Driven Schrödinger Bridge. *arXiv* **2018**, arXiv:1610.02588v4.
11. Williams, C.K.; Rasmussen, C.E. *Gaussian Processes for Machine Learning*; MIT Press: Cambridge, MA, USA, 2006; Volume 2.
12. Ruttor, A.; Batz, P.; Opper, M. Approximate Gaussian process inference for the drift function in stochastic differential equations. In Proceedings of the 26th International Conference on Neural Information Processing Systems, Lake Tahoe, NV, USA, 5–10 December 2013; pp. 2040–2048.
13. Øksendal, B. *Stochastic Differential Equations*; Springer: Berlin/Heidelberg, Germany, 2003; pp. 65–84.
14. Song, Y.; Sohl-Dickstein, J.; Kingma, D.P.; Kumar, A.; Ermon, S.; Poole, B. Score-Based Generative Modeling through Stochastic Differential Equations. *arXiv* **2020**, arXiv:2011.13456.
15. Nelson, E. *Dynamical Theories of Brownian Motion*; Princeton University Press: Princeton, NJ, USA, 1967; Volume 3.
16. Anderson, B.D. Reverse-time diffusion equation models. *Stoch. Process. Appl.* **1982**, *12*, 313–326. [CrossRef]
17. Elliott, R.J.; Anderson, B.D. Reverse time diffusions. *Stoch. Process. Appl.* **1985**, *19*, 327–339. [CrossRef]
18. Follmer, H. An entropy approach to the time reversal of diffusion processes. *Lect. Notes Control Inf. Sci.* **1984**, *69*, 156–163.
19. Haussmann, U.; Pardoux, E. Time reversal of diffusion processes. In *Stochastic Differential Systems Filtering and Control*; Springer: Berlin/Heidelberg, Germany, 1985; pp. 176–182.
20. Pavon, M.; Wakolbinger, A. On free energy, stochastic control, and Schrödinger processes. In *Modeling, Estimation and Control of Systems with Uncertainty*; Springer: New York, NY, USA, 1991; pp. 334–348.

21. Cramer, E. Probability measures with given marginals and conditionals: I-projections and conditional iterative proportional fitting. *Stat. Decis.—Int. J. Stoch. Methods Models* **2000**, *18*, 311–330.
22. Bernton, E.; Heng, J.; Doucet, A.; Jacob, P.E. Schrödinger Bridge Samplers. *arXiv* **2019**, arXiv:1912.13170.
23. Papaspiliopoulos, O.; Pokern, Y.; Roberts, G.O.; Stuart, A.M. Nonparametric estimation of diffusions: A differential equations approach. *Biometrika* **2012**, *99*, 511–531. [CrossRef]
24. Batz, P.; Ruttor, A.; Opper, M. Approximate Bayes learning of stochastic differential equations. *Phys. Rev. E* **2018**, *98*, 022109. [CrossRef]
25. Pokern, Y.; Stuart, A.M.; van Zanten, J.H. Posterior consistency via precision operators for Bayesian nonparametric drift estimation in SDEs. *Stoch. Process. Appl.* **2013**, *123*, 603–628. [CrossRef]
26. De Bortoli, V.; Thornton, J.; Heng, J.; Doucet, A. Diffusion Schrödinger Bridge with Applications to Score-Based Generative Modeling. *arXiv* **2021**, arXiv:2106.01357.
27. Feydy, J. Geometric Data Analysis, beyond Convolutions. Ph.D. Thesis, Université Paris-Saclay, Gif-sur-Yvette, France, 2020. Available online https://www.math.ens.fr/$\sim$feydy (accessed on 22 August 2021).
28. Kingma, D.P.; Welling, M. Auto-encoding variational bayes. *arXiv* **2013**, arXiv:1312.6114.
29. Papamakarios, G.; Pavlakou, T.; Murray, I. Masked autoregressive flow for density estimation. *arXiv* **2017**, arXiv:1705.07057.
30. Papamakarios, G. Neural density estimation and likelihood-free inference. *arXiv* **2019**, arXiv:1910.13233.
31. Wang, G.; Jiao, Y.; Xu, Q.; Wang, Y.; Yang, C. Deep Generative Learning via Schrödinger Bridge. *arXiv* **2021**, arXiv:2106.10410.
32. Huang, J.; Jiao, Y.; Kang, L.; Liao, X.; Liu, J.; Liu, Y. Schrödinger-Föllmer Sampler: Sampling without Ergodicity. *arXiv* **2021**, arXiv:2106.10880.
33. Kingma, D.P.; Salimans, T.; Poole, B.; Ho, J. Variational Diffusion Models. *arXiv* **2021**, arXiv:2107.00630.
34. Tong, A.; Huang, J.; Wolf, G.; Van Dijk, D.; Krishnaswamy, S. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In Proceedings of the International Conference on Machine Learning, online, 12–18 July 2020; pp. 9526–9536.
35. McCann, R.J.; Guillen, N. Five lectures on optimal transportation: Geometry, regularity and applications. In *Analysis and Geometry of Metric Measure Spaces: Lecture Notes of the Séminaire de Mathématiques Supérieure (SMS) Montréal*; American Mathematical Society: Providence, RI, USA, 2011; pp. 145–180.
36. Schiebinger, G.; Shu, J.; Tabaka, M.; Cleary, B.; Subramanian, V.; Solomon, A.; Gould, J.; Liu, S.; Lin, S.; Berube, P.; et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell* **2019**, *176*, 928–943. [CrossRef]
37. Sanguinetti, G.; Lawrence, N.D.; Rattray, M. Probabilistic inference of transcription factor concentrations and gene-specific regulatory activities. *Bioinformatics* **2006**, *22*, 2775–2781. [CrossRef]
38. Léonard, C. A survey of the Schrödinger problem and some of its connections with optimal transport. *arXiv* **2013**, arXiv:1308.0215.
39. Léonard, C. Some properties of path measures. In *Séminaire de Probabilités XLVI*; Springer: Cham, Switzerland, 2014; pp. 207–230.
40. Kunitha, H. On backward stochastic differential equations. *Stochastics* **1982**, *6*, 293–313. [CrossRef]
41. Revuz, D.; Yor, M. *Continuous Martingales and Brownian Motion*; Springer: Berlin/Heidelberg, Germany, 2013; Volume 293.
42. Kailath, T. The structure of Radon-Nikodym derivatives with respect to Wiener and related measures. *Ann. Math. Stat.* **1971**, *42*, 1054–1067. [CrossRef]
43. Sottinen, T.; Särkkä, S. Application of Girsanov theorem to particle filtering of discretely observed continuous-time non-linear systems. *Bayesian Anal.* **2008**, *3*, 555–584. [CrossRef]
44. Levy, B.C. *Principles of Signal Detection and Parameter Estimation*; Springer: Boston, MA, USA, 2008.
45. Gyöngy, I.; Krylov, N. Existence of strong solutions for Itô's stochastic equations via approximations. *Probab. Theory Relat. Fields* **1996**, *105*, 143–158. [CrossRef]
46. Sra, S. Scalable nonconvex inexact proximal splitting. In Proceedings of the Advances in Neural Information Processing Systems, Lake Tahoe, NV, USA, 3–8 December 2012; Volume 25, pp. 530–538.
47. Álvarez, M.A.; Rosasco, L.; Lawrence, N.D. Kernels for Vector-Valued Functions: A Review. *Found. Trends Mach. Learn.* **2012**, *4*, 195–266. [CrossRef]
48. Evgeniou, T.; Micchelli, C.A.; Pontil, M.; Shawe-Taylor, J. Learning multiple tasks with kernel methods. *J. Mach. Learn. Res.* **2005**, *6*, 615–637.
49. Martino, L.; Elvira, V.; Louzada, F. Effective sample size for importance sampling based on discrepancy measures. *Signal Process.* **2017**, *131*, 386–401. [CrossRef]