
Data-Conditional Diffusion Bridges

Ella Tamir
ella.tamir@aalto.fi
Dept. of Computer Science
Aalto University

Martin Trapp
martin.trapp@aalto.fi
Dept. of Computer Science
Aalto University

Arno Solin
arno.solin@aalto.fi
Dept. of Computer Science
Aalto University

Abstract

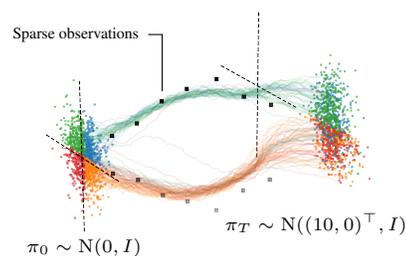
The dynamic Schrödinger bridge problem provides an appealing setting for solving constrained time-series data generation tasks posed as an iteration over optimal transport problems. Recent works have demonstrated state-of-the-art results but are limited to learning bridges with only initial and terminal constraints. Our work extends this paradigm by proposing the Iterative Smoothing Bridge (ISB). We integrate Bayesian filtering and optimal control into learning the diffusion process, enabling constrained stochastic processes governed by sparse observations at intermediate stages and terminal constraints, and assess the effectiveness of ISB on a single-cell embryo RNA data set.

1 Introduction

Generative diffusion models have gained increasing popularity and achieved impressive results in a variety of challenging application domains, such as computer vision (*e.g.*, Ho et al., 2020; Song et al., 2021; Dhariwal & Nichol, 2021), reinforcement learning (*e.g.*, Janner et al., 2022), and time series modelling (*e.g.*, Rasul et al., 2021; Vargas et al., 2021; Tashiro et al., 2021; Park et al., 2022). Recent works have explored connections between denoising diffusion models and the dynamic Schrödinger bridge problem (SBP, *e.g.*, Vargas et al., 2021; De Bortoli et al., 2021; Shi et al., 2022) to adopt iterative schemes for solving the dynamic optimal transport problem more efficiently. The solution of the SBP then acts as a denoising diffusion model in finite time and minimizes Kullback–Leibler (KL) divergence to a set reference process. Data may then be generated by time reversal of the process, *i.e.*, through the denoising process.

In many applications, the interest is not purely in modelling transport between an initial and terminal state distribution. For example, in naturally occurring generative processes, we typically observe snapshots of realizations *along intermediate stages* of individual sample trajectories (see Sec. 1). Such problems arise in medical diagnosis (*e.g.*, tissue changes and cell growth), demographic modelling, environmental dynamics, and animal movement modelling. patterns. Recently, constrained optimal control problems have been explored by adding additional fixed path constraints (Maoutsa et al., 2020; Maoutsa & Opper, 2022) or modifying the prior processes (Fernandes et al., 2021). However, defining meaningful fixed path constraints or prior processes for the optimal control problems can be challenging, while sparse observational data are accessible in many real-world applications.

In this work, we propose the *Iterative Smoothing Bridge* (ISB), an iterative method for solving control problems under constraints on both the initial and terminal distribution and sparse observational data constraints, the latter acting as a way to encourage the paths sampled from the transport process



Transport between two Gaussians at the terminal time T constrained by sparse observations (\blacksquare) at intermediate times.

to lie close to the observed data points. We perform the conditioning by leveraging the iterative pass idea from the Iterative Proportional Fitting procedure (IPFP, see Kullback, 1968; De Bortoli et al., 2021) and applying differentiable particle filtering (Reich, 2013; Corenflos et al., 2021). Integrating sequential Monte Carlo methods (*e.g.*, Doucet et al., 2001; Chopin & Papaspiliopoulos, 2020) into the IPFP framework in such a way is non-trivial and can be understood as a novel iterative version of the algorithm by Maoutsa & Opper (2022) but with more general marginal constraints and additional path constraints defined by data.

We summarize the contributions as follows. (i) We propose a novel method for solving constrained optimal transport as a bridge problem under sparse observational data constraints. (ii) Thereof, we utilize the strong connections between the constrained bridging problem and particle filtering in sequential Monte Carlo, extending those links from pure inference to learning. Additionally, (iii) we demonstrate practical efficiency by applying the iterative smoothing approach to single-cell embryo RNA modelling.

2 Background

For modelling the time dynamics, we assume a (continuous-time) state-space model consisting of a non-linear latent Itô SDE (see, *e.g.*, Øksendal, 2003; Särkkä & Solin, 2019) in $[0, T] \times \mathbb{R}^d$ with drift function $f_\theta(\cdot)$ and diffusion function $g(\cdot)$, and a Gaussian observation model, *i.e.*,

$$\mathbf{x}_0 \sim \pi_0, \quad d\mathbf{x}_t = f_\theta(\mathbf{x}_t, t) dt + g(t) d\boldsymbol{\beta}_t, \quad (1)$$

and $\mathbf{y}_k \sim N(\mathbf{y}_k | \mathbf{x}_{t_k}, \sigma^2 \mathbf{I}_d) \Big|_{t=t_k}$, where the drift function $f_\theta : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ is modelled by a neural network parameterized by $\theta \in \Theta$, diffusion $g : [0, T] \rightarrow \mathbb{R}$ and $\boldsymbol{\beta}_t$ denotes standard d -dimensional Brownian motion. \mathbf{x}_t denotes the latent process and \mathbf{y}_t denotes the observation-space process. We consider the continuous-discrete time setting, where the process is observed at discrete times t_k such that observational data can be given in terms of a collection of input–output pairs $\{(t_j, \mathbf{y}_j)\}_{j=1}^M$.

Schrödinger Bridges and Optimal Control The Schrödinger bridge problem (SBP, Schrödinger, 1932; Léonard, 2014) is an entropy-regularized optimal transport problem where the optimality is measured through the KL divergence from a reference measure \mathbb{P} to the posterior \mathbb{Q} , with fixed initial and final densities π_0 and π_T , *i.e.*, $\min_{\mathbb{Q} \in \mathcal{P}(\pi_0, \pi_T)} D_{\text{KL}}[\mathbb{Q} \| \mathbb{P}]$.

In this work, we consider only the case where the measures \mathbb{P} and \mathbb{Q} are constructed as the marginals of an SDE, *i.e.*, \mathbb{Q}_t is the probability measure of the marginal of the SDE in Eq. (1) at time t , whereas \mathbb{P}_t corresponds to the probability measure of the marginal of a reference SDE $d\mathbf{x}_t = f(\mathbf{x}_t, t) dt + g(t) d\boldsymbol{\beta}_t$, at time t , where we call f the reference drift. Under the optimal control formulation of the SBP (Caluya & Halder, 2021) the KL divergence reduces to

$$\mathbb{E} \left[\int_0^T \frac{1}{2g(t)^2} \|f_\theta(\mathbf{x}_t, t) - f(\mathbf{x}_t, t)\|^2 dt \right], \quad (2)$$

where the expectation is over paths from Eq. (1). Rüschendorf & Thomsen (1993) and Rüschendorf (1995) showed that a solution to the SBP can be obtained by iteratively solving two half-bridge problems using the Iterative Proportional Fitting procedure (IPFP) for $l = 0, 1, \dots, L$ steps, $\mathbb{Q}_{2l+1} = \arg \min_{\mathbb{Q} \in \mathcal{P}(\cdot, \pi_T)} D_{\text{KL}}[\mathbb{Q} \| \mathbb{Q}_{2l}]$ and backwards $\mathbb{Q}_{2l+2} = \arg \min_{\mathbb{Q} \in \mathcal{P}(\pi_0, \cdot)} D_{\text{KL}}[\mathbb{Q} \| \mathbb{Q}_{2l+1}]$, where \mathbb{Q}_0 is set as the reference measure, and $\mathcal{P}(\pi_0, \cdot)$ and $\mathcal{P}(\cdot, \pi_T)$ denote the sets of probability measures with only either the marginal at time 0 or time T coinciding with π_0 or π_T , respectively. Recently, the IPFP to solving Schrödinger bridges has been adapted as a machine learning problem (Bernton et al., 2019; Vargas et al., 2021; De Bortoli et al., 2021).

3 The Iterative Smoothing Bridge

Given an initial and terminal distribution π_0 and π_T , we are interested in learning a data-conditional bridge between π_0 and π_T . Let $\mathcal{D} = \{(t_j, \mathbf{y}_j)\}_{j=1}^M$ be a set of M sparsely observed values (allowing for multiple observations at a single time), *i.e.*, only a few or no observations are made at each point in time and let the state-space model of interest be given by Eq. (1). Our aim is to find a parameterization of the drift function f_θ such that evolving N particles \mathbf{x}_t^i , with $\mathbf{x}_0^i \sim \pi_0$ (with $i = 1, 2, \dots, N$), according to Eq. (1) will result in samples \mathbf{x}_T^i from the terminal distribution

π_T . Inspired by the IPFP by De Bortoli et al. (2021), which decomposes the SBP into finding two half-bridges, we propose to iteratively solve the two half-bridge problems while accounting for the additional sparse observations simultaneously. For this, let

$$d\mathbf{x}_t = f_{l,\theta}(\mathbf{x}_t, t) dt + g(t) d\beta_t, \quad \mathbf{x}_0 \sim \pi_0, \quad (3)$$

$$d\mathbf{z}_t = b_{l,\phi}(\mathbf{z}_t, t) dt + g(t) d\hat{\beta}_t, \quad \mathbf{z}_0 \sim \pi_T, \quad (4)$$

denote the forward and backward SDE at iteration $l = 1, 2, \dots, L$, where $\hat{\beta}_t$ is the reverse-time Brownian motion. For simplicity, we denote $\beta_t = \hat{\beta}_t$ when the direction of the SDE is clear.

To learn the data-conditional bridge, our Iterative Smoothing Bridge (ISB) method employs the following steps: **1** evolve *forward* particle trajectories according to Eq. (3) with drift $f_{l-1,\theta}$ and filter w.r.t. the observations $\{(t_j, \mathbf{y}_j)\}_{j=1}^M$, **2** learn the drift function $b_{l,\phi}$ for the reverse-time SDE, **3** evolve *backward* particle trajectories according to Eq. (4) with the drift $b_{l,\phi}$ learned in step **2** and filter w.r.t. the observations $\{(t_j, \mathbf{y}_j)\}_{j=1}^M$, and **4** learn the drift function $f_{l,\theta}$ for the forward SDE based on the backward particles.

Step 1 (and 3): Given a fixed discretization of the time interval $[0, T]$ denoted as $\{t_k\}_{k=1}^K$ with $t_1 = 0$ and $t_K = T$, denote the time step lengths as $\Delta_k = t_{k+1} - t_k$. By truncating the Itô–Taylor series of the SDE, we can consider an Euler–Maruyama (e.g., Ch. 8 in Särkkä & Solin, 2019) type of discretization for the continuous-time problem. Following Eq. (3), we evolve the i^{th} particle at time t_k according to

$$\tilde{\mathbf{x}}_{t_k}^i = \mathbf{x}_{t_{k-1}} + f_{l-1,\theta}(\mathbf{x}_{t_{k-1}}, t_{k-1})\Delta_k + g(t_{k-1})\sqrt{\Delta_k}\boldsymbol{\xi}_k^i, \quad (5)$$

where $\boldsymbol{\xi}_k^i \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Notice that we have not yet conditioned on the observational data. In step **3**, the particles $\tilde{\mathbf{z}}_{t_k}^i$ of the backward SDE Eq. (4) are similarly obtained. The SDE dynamics sampled in steps **1** and **3** apply the learned drift functions $f_{l-1,\theta}$ and $b_{l,\phi}$ from the previous step and do not require sampling from the underlying SDE model.

For resampling, we employ a *differentiable resampling* procedure, where the particles and weights $(\tilde{\mathbf{x}}_{t_k}^i, w_{t_k}^i)$ are transported to uniformly weighted particles $(\mathbf{x}_{t_k}^i, \frac{1}{N})$ by solving an entropy-regularized optimal transport problem (Cuturi, 2013; Peyré & Cuturi, 2019; Corenflos et al., 2021) (see App. D). Through application of the ε -regularized optimal transport map $\mathbf{T}_{(\varepsilon)} \in \mathbb{R}^{N \times N}$ (see Corenflos et al., 2021) the particles are resampled via the map to $\mathbf{x}_{t_k}^i = \tilde{\mathbf{X}}_{t_k}^\top \mathbf{T}_{(\varepsilon),i}$, where $\tilde{\mathbf{X}}_{t_k} \in \mathbb{R}^{N \times d}$ denotes the stacked particles $\{\tilde{\mathbf{x}}_{t_k}^i\}_{i=1}^N$ at time t_k before resampling.

Step 2 (and 4): Given the particles $\{\mathbf{x}_{t_k}^i\}_{k=1, i=1}^{K,N}$, we now aim to learn the drift function for the respective reverse-time process. The purpose of this step is to find a mean-matching reversal of the trajectories, in other words we aim to find $f_{l,\theta}$ such that it best explains the change we observe from $\{\mathbf{x}_{t_k}^i\}$ to $\{\mathbf{x}_{t_{k+1}}^i\}$ for each trajectory $i = 1, 2, \dots, N$.

In case no observation is available at time t_k , we apply the mean-matching loss based on a Gaussian transition approximation proposed in De Bortoli et al. (2021):

$$\ell_{k+1, \text{nobs}}^i = \|b_{l,\phi}(\mathbf{x}_{t_{k+1}}^i, t_{k+1})\Delta_k - \mathbf{x}_{t_{k+1}}^i - f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}^i, t_k)\Delta_k + \mathbf{x}_{t_k}^i + f_{l-1,\theta}(\mathbf{x}_{t_k}^i, t_k)\Delta_k\|^2. \quad (6)$$

In case an observation is available at time t_k the particle values $\tilde{\mathbf{X}}_{t_k}$ will be coupled through the optimal transport map. Therefore, the transition density is a sum of Gaussian variables (see App. A for details and a derivation), and the mean-matching loss is therefore given by:

$$\begin{aligned} \ell_{k+1, \text{obs}}^i &= \|b_{l,\phi}(\mathbf{x}_{t_{k+1}}^i, t_{k+1})\Delta_k - \mathbf{x}_{t_{k+1}}^i - f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}^i, t_k)\Delta_k \\ &\quad + \sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^n + f_{l-1,\theta}(\mathbf{x}_{t_k}^n, t_k)\Delta_k)\|^2, \end{aligned}$$

In deriving the loss, we apply the property that the reverse drift should satisfy

$$b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1}) = f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k) - g(t_{k+1})^2 \nabla \ln p_{t_{k+1}}, \quad (7)$$

where $p_{t_{k+1}}$ is the particle filtering density after differential resampling at time t_{k+1} . Thus the impact of observations on the loss function is two-fold, the observations define the value of the transport matrix $\mathbf{T}_{(\varepsilon)}$ and the marginal score $\nabla \ln p_{t_{k+1}}$.

The overall objective function is a combination of both loss functions, with the respective mean-matching loss depending on whether t_k is an observation time. The final loss function is written as:

$$\ell(\phi) = \sum_{i=1}^N \left[\sum_{k=1}^K \ell_{k,\text{obs}}^i(\phi) \mathbb{I}_{y_{t_k} \neq \emptyset} + \ell_{k,\text{nobs}}^i(\phi) \mathbb{I}_{y_{t_k} = \emptyset} \right], \quad (8)$$

where \mathbb{I}_{cond} denotes an indicator function that returns ‘1’ iff the condition is true, and ‘0’ otherwise. Consequently, the parameters ϕ of $b_{l,\phi}$ are learned by minimizing Eq. (8) through gradient descent. In practice, a cache of trajectories $\{\mathbf{x}_{t_k}^i\}_{k=1, i=1}^{K,N}$ is maintained through training of the drift functions, and refreshed at a fixed number of inner loop iterations, as in De Bortoli et al. (2021), avoiding differentiation over the SDE generation computational graph. In the single-cell experiment, the cache size is set to 1000 and it is refreshed every 500 iterations. The calculations for step 4 follow similarly. We present a high-level description of the ISB steps in Alg. 1.

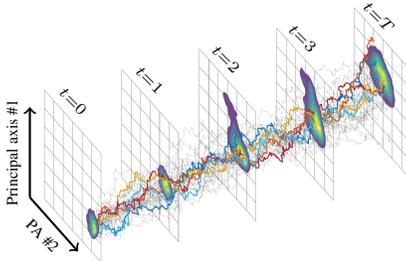
4 Experiments

To assess the properties and performance of the ISB, we present a single-cell embryo RNA modelling experiment to demonstrate how the iterative learning procedure can incorporate both observational data and terminal constraints.

We parametrize the forward and backward drift functions f_θ and b_ϕ as neural networks, and use a MLP block design as in De Bortoli et al. (2021). The latent state SDE was simulated by Euler–Maruyama with a fixed time-step of 0.01 over 100 steps and 1000 particles, full details are included in App. B.

Table 1: Results for the single-cell embryo RNA experiment. We compare ISB to TrajectoryNet, IPML, and our implementation of IPFP. Unlike the other methods, our model is able to utilize the intermediate data distributions while training.

METHOD	Earth mover’s distance				
	$t=0$	$t=1$	$t=2$	$t=3$	$t=T$
TrajectoryNet	0.62	1.15	1.49	1.26	0.99
IPML	0.34	1.13	1.35	1.01	0.49
IPFP (no obs)	0.57	1.53	1.86	1.32	0.85
ISB (single-cell obs)	0.57	1.04	1.24	0.94	0.83



(b) Iterative Smoothing Bridge

Single-cell embryo RNA-seq We evaluated our approach on an Embryoid body scRNA-seq time course (Tong et al., 2020). The data consists of RNA measurements collected over five time ranges from a developing human embryo system. No trajectory information is available, instead we only have access to snapshots of RNA data. This leads to a data set over 5 time ranges, the first from days 0–3 and the last from days 15–18. In the experiment, we followed the protocol by Tong et al. (2020), reduced the data dimensionality to $d = 5$ using PCA, and used the first and last time ranges as the initial and terminal constraints. All other time ranges are considered observational data. Contrary to the other experiments, intermediate data are imprecise (only a time range of multiple days is known) but abundant.

We set the diffusion $g(t)$ to a constant level, and learned the ISB using a zero drift and compared it against an unconditional bridge obtained through the IPFP (De Bortoli et al., 2021)—see Sec. 4. The ISB learns to generate trajectories with marginals closer to the observed data while performing comparably to the IPFP at the initial and terminal stages. This improvement is also verified numerically in Table 1, showing that the ISB obtains a lower Earth mover’s distance between the generated marginals and the observational data than IPFP. Additionally, Table 1 lists the performance of previous works that do not use the intermediate data during training (Tong et al., 2020) or only use it to construct an informative reference drift (Vargas et al., 2021), see App. B.1 for details. In both cases, ISB outperforms the other approaches w.r.t. the intermediate marginal distributions ($t = 1, 2, 3$), while IPML (Vargas et al., 2021) outperforms ISB at the initial and terminal stages due to its data-driven reference drift. Notice that while we reduced the dimensionality via PCA to 5 for fair comparisons to Vargas et al. (2021), the ISB model would also allow modelling the full state-space model, with observations in the high-dimensional gene space and a latent SDE.

References

- Espen Bernton, Jeremy Heng, Arnaud Doucet, and Pierre E. Jacob. Schrödinger bridge samplers. *arXiv preprint arXiv:1912.13170*, 2019.
- Kenneth F. Caluya and Abhishek Halder. Reflected Schrödinger bridge: Density control with path constraints. In *American Control Conference, ACC 2021*, pp. 1137–1142. IEEE, 2021.
- Nicolas Chopin and Omiros Papaspiliopoulos. *An Introduction to Sequential Monte Carlo*. Springer, 2020.
- Adrien Corenflos, James Thornton, George Deligiannidis, and Arnaud Doucet. Differentiable particle filtering via entropy-regularized optimal transport. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021*, volume 139 of *Proceedings of Machine Learning Research*, pp. 2100–2111. PMLR, 2021.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems 26 (NIPS)*, pp. 2292–2300. Curran Associates, Inc., 2013.
- Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, pp. 17695–17709. Curran Associates, Inc., 2021.
- Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pp. 8780–8794. Curran Associates, Inc., 2021.
- Arnaud Doucet, Nando De Freitas, and Neil James Gordon. *Sequential Monte Carlo methods in practice*. Statistics for Engineering and Information Science. Springer, 2001.
- David Lopes Fernandes, Francisco Vargas, Carl Henrik Ek, and Neill DF Campbell. Shooting Schrödinger’s cat. In *Proceedings of the Fourth Symposium on Advances in Approximate Bayesian Inference (AABI)*, 2021.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pp. 6840–6851. Curran Associates, Inc., 2020.
- Michael Janner, Yilun Du, Joshua B. Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International Conference on Machine Learning (ICML)*, volume 162, pp. 9902–9915. PMLR, 2022.
- Solomon Kullback. Probability densities with given marginals. *The Annals of Mathematical Statistics*, 39(4):1236–1243, 1968.
- Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete & Continuous Dynamical Systems*, 34(4):1533–1574, 2014.
- Dimitra Maoutsa and Manfred Opper. Deterministic particle flows for constraining stochastic nonlinear systems. *Physical Review Research*, 4(4), 2022.
- Dimitra Maoutsa, Sebastian Reich, and Manfred Opper. Interacting particle solutions of Fokker–Planck equations through gradient–log–density estimation. *Entropy*, 22:802, 07 2020.
- Bernt Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer, New York, NY, sixth edition, 2003.
- Sung Woo Park, Kyungjae Lee, and Junseok Kwon. Neural markov controlled SDE: Stochastic optimization for continuous-time data. In *International Conference on Learning Representations (ICLR)*, 2022.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.

- Kashif Rasul, Abdul-Saboor Sheikh, Ingmar Schuster, Urs M. Bergmann, and Roland Vollgraf. Multivariate probabilistic time series forecasting via conditioned normalizing flows. In *International Conference on Learning Representations (ICLR)*, 2021.
- Sebastian Reich. A nonparametric ensemble transform method for Bayesian inference. *SIAM Journal on Scientific Computing*, 35, 2013.
- Ludger Ruschendorf. Convergence of the iterative proportional fitting procedure. *The Annals of Statistics*, 23(4):1160–1174, 1995.
- L. Rüschemdorf and W. Thomsen. Note on the Schrödinger equation and I-projections. *Statistics & Probability Letters*, 17(5):369–375, 1993.
- Simo Särkkä and Arno Solin. *Applied Stochastic Differential Equations*. Cambridge University Press, Cambridge, UK, 2019.
- E. Schrödinger. Sur la théorie relativiste de l’électron et l’interprétation de la mécanique quantique. *Annales de l’institut Henri Poincaré*, 2(4):269–310, 1932.
- Yuyang Shi, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. Conditional simulation using diffusion Schrödinger bridges. In *38th Conference on Uncertainty in Artificial Intelligence*. UAI, 2022.
- Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pp. 1415–1428. Curran Associates, Inc., 2021.
- Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pp. 24804–24816. Curran Associates, Inc., 2021.
- Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9526–9536. PMLR, 2020.
- Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving Schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.

A Method Details

We present the details of the objective function derivation in App. A.1 and explain the connection of the backward drift function to Hamilton–Jacobi equations in App. A.2. In App. A.3, we discuss the behaviour of our model at the limit $M \rightarrow \infty$, that is, when the observations fully represent the marginal densities of the stochastic process.

A.1 Deriving the Mean-matching Loss at Observation Times

Recall that the forward loss is written as

$$\ell(\phi) = \sum_{i=1}^N \left[\sum_{k=1}^K \ell_{k,\text{obs}}^i(\phi) \mathbb{I}_{y_{t_k} \neq \emptyset} + \ell_{k,\text{nobs}}^i(\phi) \mathbb{I}_{y_{t_k} = \emptyset} \right], \quad (9)$$

where the loss at observations $\ell_{k,\text{obs}}^i(\phi)$ and loss elsewhere $\ell_{k,\text{nobs}}^i(\phi)$ are

$$\ell_{k+1,\text{nobs}}^i = \|b_{l,\phi}(\mathbf{x}_{t_{k+1}}^i, t_{k+1})\Delta_k - \mathbf{x}_{t_{k+1}}^i - f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}^i, t_k)\Delta_k + \mathbf{x}_{t_k}^i + f_{l-1,\theta}(\mathbf{x}_{t_k}^i, t_k)\Delta_k\|^2, \quad (10)$$

$$\begin{aligned} \ell_{k+1,\text{obs}}^i &= \|b_{l,\phi}(\mathbf{x}_{t_{k+1}}^i, t_{k+1})\Delta_k - \mathbf{x}_{t_{k+1}}^i - f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}^i, t_k)\Delta_k \\ &\quad + \frac{1}{C_{\varepsilon,i}} \sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^n + f_{l-1,\theta}(\mathbf{x}_{t_k}^n, t_k)\Delta_k)\|^2, \end{aligned} \quad (11)$$

For convenience, we state the backward loss functions which follow similarly to their forward versions. The backward loss is defined as

$$\overleftarrow{\ell}(\theta) = \sum_{i=1}^N \left[\sum_{k=1}^K \overleftarrow{\ell}_{k,\text{obs}}^i(\theta) \mathbb{I}_{y_{t_k} \neq \emptyset} + \overleftarrow{\ell}_{k,\text{nobs}}^i(\theta) \mathbb{I}_{y_{t_k} = \emptyset} \right], \quad (12)$$

where the loss at observations $\overleftarrow{\ell}_{k,\text{obs}}^i(\theta)$ and loss elsewhere $\overleftarrow{\ell}_{k,\text{nobs}}^i(\theta)$ are

$$\overleftarrow{\ell}_{k+1,\text{nobs}}^i = \|f_{l,\theta}(\mathbf{x}_{t_{k+1}}^i, t_{k+1})\Delta_k - \mathbf{x}_{t_{k+1}}^i - b_{l,\phi}(\mathbf{x}_{t_{k+1}}^i, t_k)\Delta_k + \mathbf{x}_{t_k}^i + b_{l,\theta}(\mathbf{x}_{t_k}^i, t_k)\Delta_k\|^2, \quad (13)$$

$$\begin{aligned} \overleftarrow{\ell}_{k+1,\text{obs}}^i &= \|f_{l,\theta}(\mathbf{x}_{t_{k+1}}^i, t_{k+1})\Delta_k - \mathbf{x}_{t_{k+1}}^i - b_{l,\phi}(\mathbf{x}_{t_{k+1}}^i, t_k)\Delta_k \\ &\quad + \frac{1}{C_{\varepsilon,i}} \sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^n + b_{l,\phi}(\mathbf{x}_{t_k}^n, t_k)\Delta_k)\|^2. \end{aligned} \quad (14)$$

Proposition 1. *Define the forward SDE as*

$$d\mathbf{x}_t = f_{l,\theta}(\mathbf{x}_t, t) dt + g(t) d\beta_t, \quad \mathbf{x}_0 \sim \pi_0, \quad (15)$$

and a backward SDE drift as

$$b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1}) = f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k) - g(t_{k+1})^2 \nabla \ln p_{t_{k+1}}, \quad (16)$$

where $p_{t_{k+1}}$ is the particle filtering density after differential resampling at time t_{k+1} . Then $b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1})$ minimizes the loss function

$$\begin{aligned} \ell_{k+1,\text{obs}}^i &= \|b_{l,\phi}(\mathbf{x}_{t_{k+1}}^i, t_{k+1})\Delta_k - \mathbf{x}_{t_{k+1}}^i - f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}^i, t_k)\Delta_k \\ &\quad + \frac{1}{C_{\varepsilon,i}} \sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^n + f_{l-1,\theta}(\mathbf{x}_{t_k}^n, t_k)\Delta_k)\|^2, \end{aligned} \quad (17)$$

where we denote $C_{\varepsilon,i} = \frac{1}{g(t_{k+1})^2 \Delta_k} \text{Var} \left(\sum_{n=1}^N T_{(\varepsilon),i,n} \tilde{\mathbf{x}}_{t_{k+1}}^n \right)$, and $\{\tilde{\mathbf{x}}_{t_{k+1}}^i\}_{i=1}^N$ are the particles before resampling.

Proof sketch. Our objective is to find a backward drift function $b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1})$ as in Eq. (16). Notice that at observation times t_k , this is not equivalent to finding the reverse drift of the SDE forward transition and differential resampling combined, since the drift function $f_{l-1,\theta}$ alone does not map the particles $\{\mathbf{x}_{t_k}^i\}_{i=1}^N$ to the particles $\{\mathbf{x}_{t_{k+1}}^i\}_{i=1}^N$. We will derive a loss function for learning

the backward drift as in Eq. (16) below, leaving the discussion on why it is a meaningful choice of a backward drift to App. A.2. Our derivation closely follows the proof of Proposition 3 in De Bortoli et al. (2021), but we provide the details here for the sake of completeness.

First, we give the transition density $p_{\mathbf{x}_{t_k} | \mathbf{x}_{t_{k-1}}^i}(\mathbf{x}_k)$ and apply it to derive the observation time loss $\ell_{k,\text{obs}}^i$. The derivation for the loss $\ell_{k,\text{no obs}}^i$ is skipped since it is as in the proof of Proposition 3 in De Bortoli et al. (2021). Suppose that at t_k , there are observations. By definition, the particles before resampling $\{\tilde{\mathbf{x}}_{t_{k+1}}^i\}_{i=1}^N$ are generated by the Gaussian transition density

$$p(\tilde{\mathbf{x}}_{t_{k+1}} | \mathbf{x}_{t_k}^i) = \mathbf{N}(\tilde{\mathbf{x}}_{t_{k+1}} | \mathbf{x}_{t_k}^i + \delta_k f_l(\mathbf{x}_{t_k}^i, t_k), g(t_{k+1})^2 \Delta_k \mathbf{I}). \quad (18)$$

Recall that the resampled particles are defined as a weighted average of all the particles, $\mathbf{x}_{t_k}^i = \sum_{n=1}^N \tilde{\mathbf{x}}_{t_k}^n T_{(\varepsilon),i,n}$. Thus, the transition density from $\{\mathbf{x}_{t_k}^i\}_{i=1}^N$ to the particles $\{\mathbf{x}_{t_{k+1}}^i\}_{i=1}^N$ is also a Gaussian,

$$p(\mathbf{x}_{t_{k+1}}^i | \mathbf{x}_{t_k}^i) = \mathbf{N}(\tilde{\mathbf{x}}_{t_{k+1}} | \sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^n + \Delta_k f_{l-1,\theta}(\mathbf{x}_{t_k}^n, t_k)), g(t_{k+1})^2 \Delta_k C_{\varepsilon,i} \mathbf{I}_d). \quad (19)$$

We will derive the loss function Eq. (17) by modifying the mean matching proof in De Bortoli et al. (2021) by the transition mean Eq. (19) and the backward drift definition Eq. (16). Using the particle filtering approximation, the marginal density can be decomposed as $p_{t_{k+1}}(\mathbf{x}_{k+1}) = \sum_{i=1}^N p_{t_k}(\mathbf{x}_k^i) p_{\mathbf{x}_{k+1} | \mathbf{x}_k^i}(\mathbf{x}_{k+1})$. By substituting the transition density Eq. (19) it follows that

$$p_{t_{k+1}}(\mathbf{x}_{t_{k+1}}) = \frac{1}{Z} \sum_{i=1}^N p_{t_k}(\mathbf{x}_{t_k}^i) \exp\left(-\frac{\left\| \left(\sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^i + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) \right) - \mathbf{x}_{t_{k+1}} \right\|^2}{2g(t_{k+1})^2 C_{\varepsilon,i} \Delta_k}\right), \quad (20)$$

where Z is the normalization constant of Eq. (19). As in the proof of Proposition 3 of De Bortoli et al. (2021), we derive an expression for the score function. Since $\nabla \ln p_{t_{k+1}}(\mathbf{x}_{t_{k+1}}) = \frac{\nabla_{\mathbf{x}_{t_{k+1}}} p_{t_{k+1}}(\mathbf{x}_{t_{k+1}})}{p_{t_{k+1}}(\mathbf{x}_{t_{k+1}})}$, we first manipulate $\nabla_{\mathbf{x}_{t_{k+1}}} p_{t_{k+1}}(\mathbf{x}_{t_{k+1}})$,

$$\begin{aligned} \nabla_{\mathbf{x}_{t_{k+1}}} p_{t_{k+1}}(\mathbf{x}_{t_{k+1}}) &= \frac{1}{Z} \sum_{i=1}^N \nabla_{\mathbf{x}_{t_{k+1}}} p(\mathbf{x}_{t_k}^i) \exp\left(-\frac{\left\| \left(\sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^i + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) \right) - \mathbf{x}_{t_{k+1}} \right\|^2}{2g(t_{k+1})^2 C_{\varepsilon,i} \Delta_k}\right) \\ &= \frac{1}{Z} \left(\sum_{i=1}^N p(\mathbf{x}_{t_k}^i) \left(\sum_{n=1}^N \frac{1}{g(t_{k+1})^2 \Delta_k C_{\varepsilon,i}} (T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^i + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) - \mathbf{x}_{t_{k+1}}) \right) \right. \\ &\quad \left. \exp\left(-\frac{\left\| \left(\sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^i + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) \right) - \mathbf{x}_{t_{k+1}} \right\|^2}{2g(t_{k+1})^2 C_{\varepsilon,i} \Delta_k}\right) \right). \end{aligned} \quad (21)$$

Substituting $p_{t_k}(\mathbf{x}_k^i) = \frac{p_{t_{k+1}}(\mathbf{x}_{t_{k+1}}) p_{\mathbf{x}_{k+1} | \mathbf{x}_k^i}(\mathbf{x}_{k+1})}{p_{\mathbf{x}_k^i | \mathbf{x}_{k+1}}(\mathbf{x}_k^i)}$ to the equation above gives

$$\nabla_{\mathbf{x}_{t_{k+1}}} p_{t_{k+1}}(\mathbf{x}_{t_{k+1}}) = p_{t_{k+1}}(\mathbf{x}_{t_{k+1}}) \sum_{i=1}^N p_{\mathbf{x}_{k+1} | \mathbf{x}_k^i}(\mathbf{x}_k^i) \left(\sum_{n=1}^N \frac{(T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^i + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) - \mathbf{x}_{t_{k+1}})}{g(t_{k+1})^2 \Delta_k C_{\varepsilon,i}} \right), \quad (22)$$

and dividing by $p_{t_{k+1}}(\mathbf{x}_{t_{k+1}})$ yields

$$\nabla \ln p_{t_{k+1}}(\mathbf{x}_{t_{k+1}}) = \sum_{i=1}^N p_{\mathbf{x}_k^i | \mathbf{x}_{k+1}}(\mathbf{x}_k^i) \left(\sum_{n=1}^N \frac{(T_{(\varepsilon),i,n}(\mathbf{x}_{t_k}^i + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) - \mathbf{x}_{t_{k+1}})}{g(t_{k+1})^2 \Delta_k C_{\varepsilon,i}} \right). \quad (23)$$

Substituting Eq. (23) to the definition of the optimal backward drift Eq. (16) gives

$$\begin{aligned}
b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1}) &= f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k) - g(t_{k+1})^2 \nabla \ln p_{t_{k+1}}(\mathbf{x}_{k+1}) \\
&= f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k) \\
&\quad - g(t_{k+1})^2 \sum_{i=1}^N p_{\mathbf{x}_{t_k^i} | \mathbf{x}_{t_{k+1}}}(\mathbf{x}_{t_{k+1}}) \left(\sum_{n=1}^N \frac{(T_{(\varepsilon),i,n}(\mathbf{x}_{t_k^i} + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) - \mathbf{x}_{t_{k+1}})}{g(t_{k+1})^2 \Delta_k C_{\varepsilon,i}} \right), \tag{24}
\end{aligned}$$

where taking $f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k)$ inside the sum yields

$$\begin{aligned}
b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1}) &= \sum_{i=1}^N p_{\mathbf{x}_{t_k^i} | \mathbf{x}_{t_{k+1}}}(\mathbf{x}_{t_{k+1}}) \\
&\quad \left(\frac{1}{C_{\varepsilon,i}} \left(\sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k^i} + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) \right) - \frac{\mathbf{x}_{t_{k+1}}}{C_{\varepsilon,i}} - \Delta_k f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k) \right) / \Delta_k. \tag{25}
\end{aligned}$$

Multiplying the equation above by Δ_k gives

$$\Delta_k b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1}) = \left(\sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k^n} + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) \right) - \frac{\mathbf{x}_{t_{k+1}}}{C_{\varepsilon,i}} - \Delta_k f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k). \tag{26}$$

Thus we may set the objective for finding the optimal backward drift $b_{l,\phi}$ as

$$\begin{aligned}
\ell_{k+1,\text{no obs}}^i &= \| b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1}) \Delta_k - \frac{\mathbf{x}_{t_{k+1}}}{C_{\varepsilon,i}} - f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k) \Delta_k \\
&\quad + \frac{1}{C_{\varepsilon,i}} \sum_{n=1}^N T_{(\varepsilon),i,n}(\mathbf{x}_{t_k^n} + f_{l-1,\theta}(\mathbf{x}_{t_k}, t_k)) \Delta_k \|^2. \tag{27}
\end{aligned}$$

□

Notice that if the weights before resampling are uniform, then $T_{(\varepsilon)} = \mathbf{I}_N$, and for all $i \in 1, 2, \dots, N$ it holds that $C_{\varepsilon,i} = 1$, since all but one of the terms in the sum $\frac{1}{g(t_{k+1})^2} \text{Var} \left(\sum_{n=1}^N T_{(\varepsilon),i,n} \tilde{\mathbf{x}}_{t_{k+1}}^n \right)$ vanish. Similarly, for one-hot weights $C_{\varepsilon,i} = 1$. In practice, we set the constant $C_{\varepsilon,i} = 1$ as in ?? and observe good empirical performance with the simplified loss function.

A.2 Connection to Hamilton–Jacobi Equations

We connect the backward drift function $b_{l,\phi}(\mathbf{x}_{t_{k+1}}, t_{k+1}) = f_{l-1,\theta}(\mathbf{x}_{t_{k+1}}, t_k) - g(t_{k+1})^2 \nabla \ln p_{t_{k+1}}(\mathbf{x}_{t_{k+1}})$ to the Hamilton–Jacobi equations for stochastic control through following the setting of Maoutsa & Opper (2022), which applies the drift $f_{l-1,\theta}(\mathbf{x}_t, t) - g(t)^2 \nabla \ln p_t(\mathbf{x}_t)$ for a backwards SDE initialized at π_T .

Consider a stochastic control problem with a path constraint $U(\mathbf{x}_t, t)$, optimizing the following loss function,

$$\mathcal{J} = \frac{1}{N} \sum_{i=1}^N \int_{t=0}^T \frac{1}{2g(t)^2} \| f_{\theta}(\mathbf{x}_t^i, t) - f(\mathbf{x}_t^i, t) \|^2 + U(\mathbf{x}_t^i, t) dt - \ln \chi(\mathbf{x}_T^i), \tag{28}$$

with the paths, \mathbf{x}_t^i sampled as trajectories from the SDE

$$\mathbf{x}_0 \sim \pi_0, \quad d\mathbf{x}_t = f_{l-1,\theta}(\mathbf{x}_t, t) dt + g(t) d\beta_t, \tag{29}$$

and the loss $\ln \chi(\mathbf{x}_T^i)$ measures distance from the distribution π_T . Since we set the path constraint via observational data, our method resembles setting $U(\mathbf{x}_t^i, t) = 0$ when t is not an observation time, and $U(\mathbf{x}_t^i) = -\log \mathbf{p}(\mathbf{y} | \mathbf{x}_t^i)$, where $\mathbf{p}(\mathbf{y} | \mathbf{x}_t^i)$ is the observation model.

Let $q_t(\mathbf{x})$ denote the marginal density of the controlled (drift f_{θ}) SDE at time t . In Maoutsa & Opper (2022), the marginal density is decomposed as

$$q_t(\mathbf{x}) = \varphi_t(\mathbf{x}) p_t(\mathbf{x}), \tag{30}$$

Algorithm 1 The Iterative Smoothing Bridge

Input: Marginal constraints (π_0, π_T) , observations $\mathcal{D} = \{(t_j, \mathbf{y}_j)\}_{j=1}^M$, initial drift function $f_{0,\theta}$, iterations L , discretization steps K , number of particles N , observation noise schedule $\kappa(l)$

Output: Learned forward and backward drift (f_θ, b_ϕ)

	for $l = 1$ to L do	
	Initialize forward particles $\{\mathbf{x}_0^i\}_{i=1}^N \sim \pi_0$	
Forward process	for $k = 1$ to K do	
	Generate $\{\mathbf{x}_k^i\}_{i=1}^N$ using $\{\mathbf{x}_{k-1}^i\}_{i=1}^N$	▷ Eq. (3)
	if Observations at t_k then	
	$\{\mathbf{x}_k^i\}_{i=1}^N \leftarrow \mathbf{DiffResample}(\{\mathbf{x}_k^i\}_{i=1}^N, \kappa(l))$	
	end if	
	end for	
	Optimize the forward loss function w.r.t. ϕ	▷ Eq. (8)
	Initialize backward particles $\{\mathbf{z}_K^i\}_{i=1}^N \sim \pi_T$	
Backward process	for $k = K$ to 1 do	
	Generate $\{\mathbf{z}_{k-1}^i\}_{i=1}^N$ using $\{\mathbf{z}_k^i\}_{i=1}^N$	▷ Eq. (4)
	if Observations at t_k then	
	$\{\mathbf{z}_{k-1}^i\}_{i=1}^N \leftarrow \mathbf{DiffResample}(\{\mathbf{z}_{k-1}^i\}_{i=1}^N, \kappa(l))$	
	end if	
	end for	
	Optimize the backwards loss function w.r.t. θ	▷ Eq. (12)
	end for	

where $\varphi_t(\mathbf{x})$ is a solution to a backwards Fokker-Planck-Kolmogorov (FPK) partial differential equation starting from $\varphi_T(\mathbf{x}) = \pi_T$, and the density evolves as in

$$\frac{d\varphi_t(\mathbf{x})}{dt} = -\mathcal{L}_f^\dagger \varphi_t(\mathbf{x}) + U(\mathbf{x}, t)\varphi_t(\mathbf{x}), \quad (31)$$

where \mathcal{L}_f^\dagger is the adjoint FPK operator to the uncontrolled system. The density $p_t(\mathbf{x})$ corresponds to the forward filtering problem, initialized with π_0 ,

$$\frac{dp_t(\mathbf{x})}{dt} = \mathcal{L}_f(p_t(\mathbf{x})) - U(\mathbf{x}, t)p_t(\mathbf{x}), \quad (32)$$

where \mathcal{L}_f is the FPK operator of the uncontrolled SDE (with drift f). The particle filtering trajectories $\{\mathbf{x}_{t_k}^i\}^i$ generated in our method are samples from the density defined by Eq. (32). In the context of our method, the path constraint matches the log-weights of particle filtering at observation times and is zero elsewhere.

In Maoutsa & Opper (2022), a backward evolution for q_t is applied, using the backwards time $\tilde{q}_{T-\tau}(\mathbf{x}) = q_\tau(\mathbf{x})$, yielding a backwards SDE starting from $\tilde{q}_0(\mathbf{x}) = \{\mathbf{x}_T^i\}_{i=1}^N$, reweighted according to π_T . The backward samples from \tilde{q} are generated following the SDE dynamics

$$d\mathbf{x}_\tau^i = (f(\mathbf{x}_\tau^i, T - \tau) + g(t)^2 \nabla \ln p_{T-\tau}(\mathbf{x}_\tau^i)) dt + g(t) d\beta_\tau. \quad (33)$$

We have thus selected the backward drift $b_{l,\phi}$ to match the drift of $\tilde{q}_t(x)$, the backward controlled density. Intuitively, our choice of $b_{l,\phi}$ is a drift which generates the smoothed particles when initialized at $\{\mathbf{x}_T^i\}_{i=1}^N$, the terminal state of the forward SDE. The discrepancy between π_T and the distribution induced by $\{\mathbf{x}_T^i\}_{i=1}^N$ then motivates the use of an iterative scheme after learning to simulate from $q_t(x)$.

A.3 Observing the Full Marginal Density

Suppose that at time t_k , we let the number of observations grow unbounded. We analyse the behaviour of our model at the resampling step, at the limit $M \rightarrow \infty$ for the number of observations and $\sigma \rightarrow 0$ for the observation noise. When applying the bootstrap proposal, recall that we combined the multiple observations to compute the log-weights as

$$\log w_{t_k}^i = -\frac{1}{2\sigma^2} \sum_{\mathbf{y}_j \in \mathcal{D}_{i,t_k}^H} \|\mathbf{x}_{t_k}^i - \mathbf{y}_j\|^2, \quad (34)$$

which works well in practice for the sparse-data settings we have considered. Below we analyse the behaviour of an alternative way to combine the weights and show that given an infinite number of observations, it creates samples from the true underlying distribution.

Proposition 2. Let $\{\mathbf{x}_{t_k}^i\}_{i=1}^N$ be a set of particles and $\{\mathbf{y}_j\}_{j=1}^M$ the observations at time t_k . Assume that the observations have been sampled from a density ρ_{t_k} and that for all i it holds that $\mathbf{x}_{t_k}^i \in \text{supp}(\rho_{t_k})$. Define the particle weights as

$$\log w_{t_k, \sigma, M}^i = \log \left(\frac{1}{Z |\mathcal{D}_{i, t_k}^{H(M)}|} \sum_{\mathbf{y}_j \in \mathcal{D}_{i, t_k}^{H(M)}} \exp(-\|\mathbf{x}_{t_k}^i - \mathbf{y}_j\|^2 / 2\sigma^2) \right), \quad (35)$$

where Z is the normalization constant of the observation model Gaussian $p(\mathbf{y} | \mathbf{x}_{t_k}^i)$. Then for each particle $\mathbf{x}_{t_k}^i$, its weight satisfies

$$\lim_{\sigma \rightarrow 0} \lim_{M \rightarrow \infty} w_{t_k, \sigma, M}^i = \rho_{t_k}(\mathbf{x}_{t_k}^i). \quad (36)$$

Proof sketch. We drop the σ and $H(M)$ from the weight notation for simplicity of notation, but remark that the particle filtering weights are dependent on both quantities. Consider the number of particles N fixed, and denote the d -dimensional sphere centered at $\mathbf{x}_{t_k}^i$ as $B(\mathbf{x}_{t_k}^i, r)$. Since each particle $\mathbf{x}_{t_k}^i$ lies in the support of the true underlying marginal density ρ_{t_k} , then for any radius $r > 0$ such that $B(\mathbf{x}_{t_k}^i, r) \in \text{supp}(\rho_{t_k})$, and $H > 0$, we may choose M high enough so that the points $\mathbf{y}_j \in \mathcal{D}_{i, t_k}^H$ satisfy $\mathbf{y}_j \in B(\mathbf{x}_{t_k}^i, r)$. It follows from Eq. (35) that

$$w_{t_k}^i = \frac{1}{Z |\mathcal{D}_{i, t_k}^{H(M)}|} \sum_{\mathbf{y}_j \in \mathcal{D}_{i, t_k}^{H(M)}} \exp(-\|\mathbf{x}_{t_k}^i - \mathbf{y}_j\|^2 / 2\sigma^2). \quad (37)$$

For any $r > 0$ and with observation noise $\sigma = cr$, we may set $c, H(M)$ so that the sum above approximates the integral

$$w_{r, t_k}^i \approx \frac{1}{|B(\mathbf{x}_{t_k}^i, r)|} \int_{B(\mathbf{x}_{t_k}^i, r)} p(\mathbf{y} | \mathbf{x}_{t_k}^i) \rho_{t_k}(\mathbf{y}) \, d\mathbf{y}. \quad (38)$$

By applying the Lebesgue differentiation theorem, we obtain that for almost every $\mathbf{x}_{t_k}^i$, we have $\lim_{r \rightarrow 0} w_{r, t_k}^i = \rho_{t_k}(\mathbf{x}_{t_k}^i)$, since as $\sigma \rightarrow 0$, the density $p(\mathbf{y} | \mathbf{x}_{t_k}^i)$ collapses to the Dirac delta of $\mathbf{x}_{t_k}^i$. \square

Prop. 2 can be interpreted as the infinite limit of a kernel density estimate of the true underlying distribution. Resampling accurately reweights the particles so that the probability of resampling particle $\mathbf{x}_{t_k}^i$ is proportional to the density ρ_{t_k} compared to the other particles. Notice that the result does not guarantee that the particles will cover the support of ρ_{t_k} , since we did not assume that the drift initialization generates a marginal density at time t_k covering its support.

B Experimental Details

The observational data consists of 10 points selected from the Schrödinger bridge trajectories, all observed at $t \in [0.25, 0.5, 0.75]$ with an exponential observation noise schedule $\kappa(l) = 1.25^{l-1}$. The ISB was run for 6 epochs and initialized with a drift from the pre-trained Schrödinger bridge model from the unconstrained problem.

The observational data consists of 10 points which lie evenly distributed on a circle, observed at $t = 0.5$ with an exponential observational noise schedule $\kappa(l) = 0.5 \cdot 1.25^{l-1}$. The ISB was run for 6 epochs and initialized with a drift from the pre-trained Schrödinger bridge model from the unconstrained problem.

The observational data consists of 6 points, with pairs being observed at times $t \in [0.4, 0.5, 0.6]$. We used a bilinear observational noise schedule with a linear decay for the first half of the iterations from $\kappa(0)^2 = 4$ to $\kappa(L/2)^2 = 1$ and a linear ascend for the second half of the iterations from $\kappa(L/2)^2 = 1$ to $\kappa(L)^2 = 4$. The ISB ran for 6 epochs, with a zero drift initialization.

B.1 Single-Cell Data Set

We directly use the preprocessed data from the TrajectoryNet (Tong et al., 2020) repository and run the ISB model on a MacBook Pro CPU. A major difference between our implementation and Vargas et al. (2021) is the reference drift. We set the reference drift to zero, which means that we utilize the intermediate data only as observations in the state-space model. On the contrary, Vargas et al. (2021) fits a mixture model of 15 Gaussians on the combined data set (across all measurement times) and sets the reference drift to the gradient of the log-likelihood of the mixture model. Effectively, such a reference drift aids in keeping the SDE trajectories within the support of the combined data set. We remark that if the intermediate observed marginals had clearly disjoint support, combining all the data would cause the mixture model to have ‘gaps’ and could cause an unstable reference model drift. Thus we consider our approach of setting the reference drift to zero as more generally applicable.

As in Vargas et al. (2021), we set the process noise to $g(t) = 1$ and model the SDE between time $t \in [0, 4]$. The learning rate is set to 0.001 with a batch size of 256 and the number of neural network training iterations equal to 5000. We apply the ISB for 6 iterations. We perform filtering using 1000 points from the intermediate data sets, but compute the Earth mover’s distance by comparing it to all available data. As the observational data at $T = 1, 2, 3$ consists of a high number of data points, the parameters H (number of nearest neighbours) and σ (observation noise) need to be carefully set. We set $H = 10$ to only include the close neighbourhood of each particle and set the observation noise schedule as constant 0.7.

C Computational Considerations

Below we discuss some important computational considerations in detail, analyzing the limit $L \rightarrow \infty$ from the perspective of setting the observation noise schedule in App. C.1.

C.1 Discussion on Observation Noise

When letting $L \rightarrow \infty$, the choice of observation noise should be carefully planned in order for the ISB procedure to have a stationary point. Here we explain why an unbounded observation noise schedule $\kappa(l)$ implies convergence to the IPF method for uncontrolled Schrödinger bridges (De Bortoli et al., 2021), when using a nearest neighbour bootstrap filter as the proposal density.

Proposition 3. *Let $\Omega \in \mathbb{R}^d$ be a bounded domain where both the observations and SDE trajectories lie, and let the particle filtering weights $\{w_{l,t_k}^i\}_{i=1}^N$ at ISB iteration l be*

$$\log w_{l,t_k}^i = -\frac{1}{2\kappa(l)^2} \sum_{\mathbf{y}_j \in \mathcal{D}_{t_k}^H} \|\mathbf{x}_{t_k}^i - \mathbf{y}_j\|^2. \quad (39)$$

If the schedule $\kappa(l)$ is unbounded with respect to l , then for any δ there exists l' such that for the normalized weights it holds

$$|\hat{w}_{l',t_k}^i - \frac{1}{N}| \leq \delta. \quad (40)$$

Proof sketch. Since $\kappa(l)$ is unbounded, for any $S > 0 \exists l'$ such that $\kappa(l') \geq S$. We choose the value of S so that the following derivation yields Eq. (40).

Let $S = \sqrt{0.5R^{-1}|\mathcal{D}_{t_k}^H| \text{diam}(\Omega)^2}$, and apply the property that $\|\mathbf{x}_{t_k}^i - \mathbf{y}_j\|^2 \leq \text{diam}(\Omega)^2$ to Eq. (39),

$$\begin{aligned} \log w_{l',t_k}^i &\geq -\frac{1}{2S^2} \sum_{\mathbf{y}_j \in \mathcal{D}_{t_k}^H} \|\mathbf{x}_{t_k}^i - \mathbf{y}_j\|^2 \\ &\geq -\frac{\sum_{\mathbf{y}_j \in \mathcal{D}_{t_k}^H} \|\mathbf{x}_{t_k}^i - \mathbf{y}_j\|^2}{R^{-1}|\mathcal{D}_{t_k}^H| \text{diam}(\Omega)^2} \geq -\frac{\sum_{\mathbf{y}_j \in \mathcal{D}_{t_k}^H} \text{diam}(\Omega)^2}{R^{-1}|\mathcal{D}_{t_k}^H| \text{diam}(\Omega)^2} \geq -R. \end{aligned} \quad (41)$$

The bound above is for the unnormalized weights, and the normalized log-weights are defined as

$$\log \hat{w}_{l',t_k}^i = \log w_{l',t_k}^i - \log \left(\sum_{j=1}^N \exp(\log w_{l',t_k}^j) \right), \quad (42)$$

where for the normalizing constant it holds that

$$\log \left(\sum_{j=1}^N \exp(\log w_{l',t_k}^j) \right) \leq \log \left(\sum_{j=1}^N 1 \right) = \log(N), \quad (43)$$

since w_{l',t_k}^j is the value of a probability density and thus always $w_{l',t_k}^j \leq 1$. Combining Eq. (42), Eq. (41) and Eq. (43), it follows that

$$\log \hat{w}_{l',t_k}^i - (-\log(N)) \geq -R, \quad (44)$$

where taking exponentials on both sides gives

$$\hat{w}_{l',t_k}^i - \frac{1}{N} \geq -(1 - \exp(-R)) \frac{1}{N}. \quad (45)$$

Since the weights are normalized, even the largest particle weight \hat{w}_{l',t_k}^j can differ from $\frac{1}{N}$ as much as every smaller weight in total lies under $\frac{1}{N}$,

$$\hat{w}_{l',t_k}^j \leq \frac{1}{N} + (N-1) \left((1 - \exp(-R)) \frac{1}{N} \right), \quad (46)$$

implying that for any weight \hat{w}_{l',t_k}^j , it holds that

$$|\hat{w}_{l',t_k}^j - \frac{1}{N}| \leq (N-1) \left((1 - \exp(-R)) \frac{1}{N} \right) \leq 1 - \exp(-R), \quad (47)$$

and selecting $R = -\log(1 - \delta)$ is sufficient for $\delta < 1$. \square

Effectively, the above derivation implies that for an unbounded observation noise schedule $\kappa(l)$, the particle weights will converge to uniform weights. Since performing differentiable resampling on uniform weights implies that $\mathbf{T}_{(\varepsilon)} = \mathbf{I}_N$, the ISB method trajectory generation step and the objective in training the backward drift converge to those of the IPF method for solving unconstrained Schrödinger bridges. Intuitively, this means that at the limit $L \rightarrow \infty$, our method will focus on reversing the trajectories and matching the terminal distribution while not further utilizing information from the observations.

D Differentiable Resampling

In the ISB model steps 1 and 3 presented in Sec. 3, we applied differentiable resampling (see Corenflos et al., 2021). Resampling itself is a basic block of particle filtering. A differentiable resampling step transports the particles and weights $(\tilde{\mathbf{x}}_{t_k}^i, w_{t_k}^i)$ to a uniform distribution over a set of particles through applying the *differentiable* ensemble transport map $\mathbf{T}_{(\varepsilon)}$, that is

$$(\tilde{\mathbf{x}}_{t_k}^i, w_{t_k}^i) \rightarrow (\tilde{\mathbf{X}}_{t_k}^\top \mathbf{T}_{(\varepsilon),i}, 1/N) = (\mathbf{x}_{t_k}^i, 1/N), \quad (48)$$

where $\tilde{\mathbf{X}}_{t_k} \in \mathbb{R}^{N \times d}$ denotes the stacked particles $\{\tilde{\mathbf{x}}_{t_k}^i\}_{i=1}^N$ at time t_k before resampling and $\mathbf{x}_{t_k}^i$ denotes the particles post resampling. Here we give the definition of the map $\mathbf{T}_{(\varepsilon)}$ and review the regularized optimal transport problem which has to be solved to compute it. We partly follow the presentation in Sections 2 and 3 of Corenflos et al. (2021), but directly apply the notation we use for particles and weights and focus on explaining the transport problem rather than the algorithm used to solve it.

The standard particle filtering resampling step consists of sampling N particles from the categorical distribution defined by the weights $\{w_{t_k}^i\}_{i=1}^N$, resulting in the particles with large weights being most likely to be repeated multiple times. A result from Reich (2013) gives the property that the random resampling step can be approximated by a deterministic ensemble transform \mathbf{T} . In heuristic terms, the ensemble transform map will be selected so that the particles $\{\mathbf{x}_{t_k}^i\}_{i=1}^N$ will be transported with minimal cost, while allowing all the weights to be uniform.

Let μ and ν be atomic measures, $\mu = \sum_{i=1}^N w_{t_k}^i \delta_{\tilde{\mathbf{x}}_{t_k}^i}$ and $\nu = \sum_{i=1}^N N^{-1} \delta_{\mathbf{x}_{t_k}^i}$, where δ_x is the Dirac delta at x . Then μ is the particle filtering distribution before resampling. Define the elements of a

cost matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$ as $C_{i,j} = \|\tilde{\mathbf{x}}_{t_k}^i - \tilde{\mathbf{x}}_{t_k}^j\|^2$, and the 2-Wasserstein distance between two atomic measures as

$$\mathcal{W}_2^2(\mu, \nu) = \min_{P \in S(\mu, \nu)} \sum_{i=1}^N \sum_{j=1}^N C_{i,j} P_{i,j}. \quad (49)$$

Above the optimal matrix \mathbf{P} is to be found within $S(\mu, \nu)$, which is a space consisting of mixtures of N particles to N particles such that the marginals coincide with the weights of μ and ν , formally

$$S(\mu, \nu) = \left\{ \mathbf{P} \in [0, 1]^{N \times N} \mid \sum_{i=1}^N P_{i,j} = w_{t_k}^i, \sum_{j=1}^N P_{i,j} = \frac{1}{N} \right\}. \quad (50)$$

The entropy-regularized Wasserstein distance with regularization parameter ε is then

$$\mathcal{W}_{2,\varepsilon}^2 = \min_{\mathbf{P} \in S(\mu, \nu)} \sum_{i=1}^N \sum_{j=1}^N P_{i,j} \left(C_{i,j} + \varepsilon \log \frac{P_{i,j}}{w_{t_k}^i \cdot \frac{1}{N}} \right). \quad (51)$$

The unique minimizing transport map of the above Wasserstein distance is denoted by $\mathbf{P}_\varepsilon^{\text{OPT}}$, and the ensemble transport map is then set as $\mathbf{T}_{(\varepsilon)} = N\mathbf{P}_\varepsilon^{\text{OPT}}$. This means that we can find the matrix $\mathbf{T}_{(\varepsilon)}$ via minimizing the regularized Wasserstein distance, which is done by applying the iterative Sinkhorn algorithm for entropy-regularized optimal transport (Cuturi, 2013).