Multi-marginal temporal Schrödinger Bridge Matching from unpaired data

Thomas Gravier [*1 2] **Thomas Boyer** [*1] **Auguste Genovesio** [1]

Abstract

Many natural dynamic processes -such as in vivo cellular differentiation or disease progression- can only be observed through the lens of static sample snapshots. While challenging, reconstructing their temporal evolution to decipher underlying dynamic properties is of major interest to scientific research. Existing approaches enable data transport along a temporal axis but are poorly scalable in high dimension and require restrictive assumptions to be met. To address these issues, we propose ***Multi-marginal temporal Schrödinger Bridge Matching*** (**MMtSBM**) *from unpaired data*, extending the theoretical guarantees and empirical efficiency of Diffusion Schrödinger Bridge Matching (Shi et al., 2023) by deriving the Iterative Markovian Fitting algorithm to multiple marginals in a novel factorized fashion. Experiments show that MMtSBM retains theoretical properties on toy examples, achieves state-of-the-art performance on real-world datasets such as transcriptomic trajectory inference in 100 dimensions, and for the first time recovers couplings and dynamics in very high dimensional image settings, effectively generating temporally coherent videos from purely unpaired data. Our work establishes multi-marginal Schrödinger Bridges as a practical and theoretically principled approach for recovering hidden dynamics from static data.

code: github.com/tgravier/MMDSBM-pytorch
website: mmtsbm.notion.site

1. Introduction

The observation of many natural processes yields partial information, resulting in limited time resolution and unpaired snapshots of data. Common examples of this are single-cell sequencing and in vivo biological imaging, where existing methods are destructive and thus cannot link two observations coming from the same cell at different timestamps. The ability to recover the true underlying dynamic from time-unpaired data samples is a key motivation for developing improved methods of trajectory inference.

The modeling of this problem is inherently probabilistic, given both the variability occurring in complex natural processes and the uncertainty of the observation. We thus ask the question: *"What is the most probable evolution of an existing data point, given uncoupled samples of the same process acquired across different times?"*.
This point of view has notably been developed in the Schrödinger Bridge (SB) theory (Schrödinger, 1931). The SB is the unique stochastic process whose marginals at start and end times match given probability distributions while minimizing the Kullback–Leibler (KL) divergence w.r.t. a given reference process. The SB also happens to solve a regularized Optimal Transport (OT) problem (Léonard, 2014). Some recent works such as Chen et al. (2019); Baradat & Léonard (2020); Lavenant et al. (2024) have explored the theoretical setting of multiple marginals. Recent major advances in statistical learning of SBs have allowed using this framework between complex empirical distributions (De Bortoli et al., 2021; Wang et al., 2021), achieved important improvements in their efficiency (Shi et al., 2023; Bortoli et al., 2024), extended it to the multi-marginal setting and explored various additional constraints such as smooth trajectories (Chen et al., 2023a; Hong et al., 2025), and spline-valued trajectories (Theodoropoulos et al., 2025).
A few methods have been proposed to solve the SB problem in an applied machine learning setting. De Bortoli et al. (2021) use iterative proportional fitting (IPF) (Kullback, 1968), the general continuous analogue of the renowned Sinkhorn algorithm (Cuturi, 2013). Subsequent works have explored alternative training schemes based on likelihood bounds (Chen et al., 2023b) or on the dual algorithm of IPF: Iterative Markovian Fitting (IMF) (Shi et al., 2023).
A closely related line of work is flow matching (Lipman et al., 2023; Liu et al., 2022; Albergo & Vanden-Eijnden, 2023). These methods have explored OT variants since their inception and have been extended to the multi-marginal setting as well as connected to the Schrödinger Bridge theory (Tong et al., 2024a;b; Kapuśniak et al., 2024).
Concurrent to our work is Park & Lee (2025); we note that they do not scale to video experiments.

[*]Equal contribution [1]IBENS, ENS Ulm, PSL, Paris, France [2]ENS Paris-Saclay, Paris, France. Correspondence to: Auguste Genovesio <auguste.genovesio@bio.ens.psl.eu>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

Existing multi-marginal methods do not scale to very high dimensions such as image space. Furthermore we believe that existing multi-marginal approaches either make use of modeling assumptions that strongly restrict the class of problems they can solve, such as using spline-valued trajectories, or lack a fully theoretically sound approach.

Contributions We make the following contributions:

1. We define the multi-marginal temporal Schrödinger Bridge problem and demonstrate its fundamental properties (existence and uniqueness of the solution).

2. We introduce a novel factorized extension of the IMF algorithm presented in Shi et al. (2023) to multiple iterative marginals in a way that is efficient –because parallelized along times, and principled –because mathematically sound and with a concrete algorithm closely following theory.

3. We produce a convergence analysis of the algorithm under asymptotic hypotheses.

4. We demonstrate the soundness of the method on low-to-medium-dimensional examples, and achieve state-of-the-art results against comparable methods on 2 widely reported single-cell transcriptomic benchmarks (Moon et al., 2019; Burkhardt et al., 2022).

5. We scale up to 7 iterative marginals in a very high-dimensional *image* setting, presenting for the first time a coherent video generation algorithm from purely time-unpaired data samples.

Notations We adopt the notations from Shi et al. (2023). We denote by $\mathcal{P}(C)$ the space of *path measures*, with $\mathcal{P}(C) = \mathcal{P}(C([0,T],\mathbb{R}^d))$, where $C([0,T],\mathbb{R}^d)$ is the space of continuous functions from $[0,T]$ to $\mathbb{R}^d$. The subset of *Markov path measures* associated with the diffusion $dX_t = v_t(X_t)dt + \sigma_t dB_t$, with σ, v locally Lipschitz, is denoted $\mathcal{M}$. We denote $(B_t)_{t\geq 0}$ the d-dimensional Brownian motion. For a process $\mathbb{Q}$, the *reciprocal class* of $\mathbb{Q}$ is $\mathcal{R}(\mathbb{Q})$. For $\mathbb{P} \in \mathcal{P}(C)$, we denote by $\mathbb{P}_t$ its marginal at time t, by $\mathbb{P}_{s,t}$ the joint law at times s,t, and by $\mathbb{P}_{s|t}$ the conditional law at s given t. We write $\mathbb{P}_{|t_i,t_j} \in \mathcal{P}(C)$ for the path distribution on (t_i,t_j) given the endpoints t_i and t_j; e.g., $\mathbb{Q}_{|t_i,t_j}$ is a scaled Brownian bridge. Unless otherwise specified, ∇ refers to gradients w.r.t. x_t at time t. For a joint law $\Pi_{0,T}$ on $\mathbb{R}^d \times \mathbb{R}^d$, the *mixture of bridges measure* is $\Pi = \Pi_{0,T}\mathbb{P}_{|0,T} \in \mathcal{P}(C)$ with $\Pi(\cdot) = \int_{\mathbb{R}^d\times\mathbb{R}^d} \mathbb{P}_{|0,T}(\cdot|x_0,x_T)d\Pi_{0,T}(x_0,x_T)$. The entropy of a process w.r.t. the Brownian motion is denoted $\mathcal{H}$. Finally, for $\pi_0, \pi_T \in \mathcal{P}(X)$, the Kullback–Leibler divergence is $\mathrm{KL}(\pi_0\|\pi_T) = \int_X \log\left(\frac{d\pi_0}{d\pi_T}(x)\right) d\pi_0(x)$.

2. Background

2.1. The Schrödinger Bridge problem

The *Schrödinger Bridge problem* (Schrödinger, 1931) seeks the most likely stochastic evolution between marginals μ_0, μ_T under a reference law $\mathbb{Q}$. It admits both a *dynamic* formulation:

$$\mathbb{P}^\star = \underset{\mathbb{P}\in\mathcal{P}(C)}{\arg\min}\, \mathrm{KL}(\mathbb{P}\,\|\,\mathbb{Q}) \;\text{ s.t. } \mathbb{P}_0 = \mu_0,\; \mathbb{P}_T = \mu_T, \quad (1)$$

and a *static* formulation on couplings $\Pi \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$:

$$\begin{aligned} \Pi^\star = \underset{\Pi}{\arg\min}\, \mathrm{KL}(\Pi\,\|\,\mathbb{Q}_{0,T}) \\ \text{s.t. } \Pi_0 = \mu_0,\; \Pi_T = \mu_T. \end{aligned} \quad (2)$$

Note: Connection to Quadratic OT. If $\mathbb{Q}$ is Brownian motion, equation 2 is precisely entropy-regularized quadratic OT with cost $c(x_0,x_T) = \frac{1}{2}\|x_0 - x_T\|^2$ and regularization $\varepsilon = \sigma^2$. In the limit $\varepsilon \to 0$, this recovers classical OT, which motivates our interpolation framework.

2.2. Iterative Markovian Fitting (IMF)

The SB solution is the unique path measure that is both *Markovian* and belongs to the *reciprocal class* of $\mathbb{Q}$ while matching marginals (Léonard, 2014). This motivates the *Iterative Markovian Fitting* (IMF) algorithm (Shi et al., 2023; Peluchetti, 2023), which alternates between reciprocal and Markov projections:

$$\mathbb{P}^{2n+1} = \mathrm{proj}_{\mathcal{M}}(\mathbb{P}^{2n}),\; \mathbb{P}^{2n+2} = \mathrm{proj}_{\mathcal{R}(\mathbb{Q})}(\mathbb{P}^{2n+1}) \quad (3)$$

These projections admit KL variational characterizations (A.1) and the iterations converge to $\mathbb{P}^\star$.

In practice, IMF is implemented by learning the drift of the Markovian projection via a bridge-matching loss (see A.1). Compared to Iterative Proportional Fitting (IPF), IMF preserves both marginals simultaneously and is more efficient (details in A.1).

Shi et al. (2023) succeeds in scaling to image space but is limited to two marginals only. To the best of our knowledge, no theoretical study of their IMF algorithm in non-phase-lifted space has been made in the multi-marginal case, and the scaling capabilities of the algorithm in this setting remain unknown.

3. Multi-marginal temporal Schrödinger Bridge Matching

We start by exposing known theoretical results about multi-marginal Schrödinger Bridges (3.1).[1] Then, to the best of

[1] We originally believed Proposition 3.4 to be novel, however it appears to be already established in Baradat & Léonard (2020).

our knowledge, from 3.2 onward we present novel results. All proofs can be found in the Appendix A.6.

3.1. Multi-marginal temporal Schrödinger Bridge Problem

In the present work, we considered the time-ordered Multi-marginal Schrödinger Bridge, where the marginals are associated with an underlying temporal axis. In this setting, the goal is not simply to fit an arbitrary number of marginals, but to recover the law of a stochastic process that evolves consistently over time.

Let $0 = t_0 < t_1 < \cdots < t_K = T$ be a fixed time grid, and let $\mu_0, ..., \mu_k, ..., \mu_T \in \mathcal{P}(\mathbb{R}^d)$ denote prescribed marginals at times $(t_k)_{k=0,\dots,K}$, assuming $\mu_{t_k} \ll \mathbb{Q}_{t_k}$ for all k. Given a reference process $\mathbb{Q}$ on $C([0,T],\mathbb{R}^d)$, the multi-marginal Schrödinger Bridge problem (MMSB) is defined as

$$\begin{aligned} \mathbb{P}^\star &= \operatorname*{argmin}_{\mathbb{P}\in\mathcal{P}(C)} \mathrm{KL}(\mathbb{P}\,\|\,\mathbb{Q}) \\ \text{subject to} \quad & X_{t_k} \sim \mu_k, \quad k = 0,\dots,K \end{aligned} \tag{4}$$

Note: Connection to multi-marginal Optimal Transport If $\mathbb{Q}$ is associated with a Brownian motion, the induced reference coupling $\mathbb{Q}_{t_0,\dots,t_K}$ is characterized by independent Gaussian increments $X_{t_{i+1}} - X_{t_i} \sim \mathcal{N}(0, \sigma^2(t_{i+1} - t_i))$. By evaluating the KL term, 4 can be rewritten as:

$$\Pi^\star = \operatorname*{argmin}_{\substack{\Pi\in\mathcal{P}\left((\mathbb{R}^d)^{K+1}\right)\\ \Pi_i=\mu_{t_i}}} \left\{ \mathbb{E}_{X\sim\Pi}\left[\sum_{i=0}^{K-1} c_i(X)\right] - 2\sigma^2 T\mathcal{H}(\Pi)\right\}$$

$$\text{where} \quad c_i(X) = \frac{1}{t_{i+1}-t_i}\,\|X_{t_{i+1}} - X_{t_i}\|^2$$

This is precisely an entropy-regularized multi-marginal OT problem with a time-structured quadratic cost $c(x_0,\dots,x_K) = \sum_{i=0}^{K-1} \frac{1}{t_{i+1}-t_i}\|x_{i+1} - x_i\|^2$ and entropy-regularisation parameter $\varepsilon = 2\sigma^2$.

This formulation is particularly interesting when no better prior is available, and because of the clear interpretation it allows: when using a Brownian motion as prior, we are approaching quadratic OT. Note however that we do not rely on this assumption at all for theoretical results.

Classical properties of the Multi-marginal temporal Schrödinger Bridge We first demonstrate a set of classical properties that characterize MMSB (4) and guide the construction of our method.

Definition 3.1 (Static formulation)**.** *Let $\mathbb{Q}_{t_0,\dots,t_K}$ be the joint law of $\mathbb{Q}$ at $0 = t_0 < \cdots < t_K = T$. The static problem is*

$$\pi^\star = \arg\min_{\pi\in\Pi(\pi_{t_0},\dots,\pi_{t_K})} KL(\pi\,\|\,\mathbb{Q}_{t_0,\dots,t_K}),$$

where $\Pi(\pi_{t_0},\dots,\pi_{t_K})$ denotes couplings on $(\mathbb{R}^d)^{K+1}$ with marginals π_{t_i}.

The MMSB is therefore a projection of the reference law onto the set of couplings with prescribed marginals. The following results ensure that this problem is well posed and that the solution has a convenient structure.

Proposition 3.1 (Existence and uniqueness)**.** *The MMSB admits a unique solution $P^\star$.*

This guarantees that the iterative algorithms we design later target a well-defined object. Moreover, the solution can be described equivalently in both static and dynamic terms.

Proposition 3.2 (Dynamic–static equivalence)**.** *The dynamic solution $P^\star$ is determined by the static one $\pi^\star$:*

$$\pi^\star = P^\star_{t_0,\dots,t_K}, \quad P^\star = \pi^\star \otimes \mathbb{Q}(\cdot \mid X_{t_0},\dots,X_{t_K}).$$

This equivalence highlights that solving the static problem is enough to recover the full path measure. In addition, the structure of $\mathbb{Q}$ plays a key role in the nature of the solution.

Proposition 3.3 (Markovianity)**.** *If $\mathbb{Q}$ is Markov, then the MMSB solution $P^\star$ is Markov.*

These properties ensure that we can restrict our search to Markovian (and therefore reciprocal A.1) measures, which will be central to the projection algorithms introduced later. Finally, the explicit form of the solution further clarifies its structure.

Proposition 3.4 (Form of the solution)**.** *Under mild assumptions:*

$$P^\star = \pi^\star \otimes \mathbb{Q}(\cdot \mid X_{t_0},\dots,X_{t_K}),$$
$$\frac{d\pi^\star}{d\mathbb{Q}_{t_0,\dots,t_K}}(x_0,\dots,x_K) = \prod_{i=0}^{K} f_i(x_i).$$

where the f_i's are functions of the Lagrange multipliers for the marginal constraints (see A.6.6).
This factorized form motivates the use of alternating projections and parametric families of potentials in the iterative algorithm that we develop in the next section.

3.2. Iterative Markovian Fitting for Multi-marginal temporal Schrödinger Bridge

3.2.1. Multi-marginal Markov and Reciprocal projections

To construct an algorithm for MMSB, we first extend the notions of reciprocal and Markovian projections to the multi-marginal setting. The idea is to approximate the global bridge by a sequence of independent sub-bridges, and to alternate between reciprocal and Markovian structures. However, a naive implementation of this idea leads to model

forgetting, failing to converge entirely. A factorized approach is thus necessary to obtain convergence (see A.4).

Definition 3.2 (Factorized reciprocal class and projection)**.** *For each interval $[t_i, t_{i+1}]$ and endpoints (x_i, x_{i+1}), let $\mathbb{Q}^{x_i,x_{i+1}}_{[t_i,t_{i+1}]}$ denote the bridge of $\mathbb{Q}$ between x_i and x_{i+1}. Given a coupling π on $(\mathbb{R}^d)^{K+1}$, define*

$$P = \int \bigotimes_{i=0}^{K-1} \mathbb{Q}^{x_i,x_{i+1}}_{[t_i,t_{i+1}]}\, \pi(dx_0, \dots, dx_K).$$

The factorized reciprocal class*, denoted $\mathcal{R}^{\otimes}(\mathbb{Q})$, is the set of all such measures P.*

Moreover, for any $P \in \mathcal{P}(C([0,T],\mathbb{R}^d))$, the reciprocal projection *onto $\mathcal{R}^{\otimes}(\mathbb{Q})$ is defined as*

$$\Pi^\star = \mathrm{proj}_{\mathcal{R}^{\otimes}(\mathbb{Q})}(P) = P_{t_0,\dots,t_K} \bigotimes_{i=0}^{K-1} \mathbb{Q}^{x_i,x_{i+1}}_{[t_i,t_{i+1}]},$$

i.e. we keep the marginals $P_{t_0,\dots,t_K}$ at the grid points and fill the dynamics between them with independent bridges of $\mathbb{Q}$ conditioned on the endpoints (x_i, x_{i+1}).

Equivalently, $\Pi^\star$ admits the variational characterization

$$\Pi^\star = \underset{\Pi \in \mathcal{R}^{\otimes}(\mathbb{Q})}{\arg\min}\, KL(P \,\|\, \Pi).$$

Proposition 3.5 (Local reciprocal structure of the factorized class)**.** *Let $\mathbb{Q}$ be a reference Markov process and let $P \in \mathcal{R}^{\otimes}(\mathbb{Q})$ belong to the factorized reciprocal class. Then, for each subinterval $[t_{i-1}, t_i]$, the restriction of P to $C([t_{i-1}, t_i], \mathbb{R}^d)$ is in the reciprocal class of $\mathbb{Q}$ over $[t_{i-1}, t_i]$. In particular, conditionally on the endpoints $(X_{t_{i-1}}, X_{t_i})$, the law of P coincides with the bridge of $\mathbb{Q}$ between t_{i-1} and t_i.*

This class provides a tractable approximation: each subinterval is filled with the bridge of $\mathbb{Q}$, while the global coupling ensures consistency across marginals. Hence, factorized bridges inherit local reciprocity, which justifies their use as a relaxation of the true reciprocal class.

This projection enforces the prescribed marginals while completing the dynamics with local bridges. In contrast, the Markovian projection seeks a single Markov diffusion with consistent marginals.

Definition 3.3 (Markovian projection in the factorized setting)**.** *Let Π be the factorized mixture of independent Brownian bridges. For any $t \in [0,T]$, let $i(t)$ be the unique index such that $t \in [t_{i(t)}, t_{i(t)+1}]$. We employ a slight abuse of notation and subsequently write i instead of $i(t)$.*

The Markovian projection *of Π, denoted $M^\star = \mathrm{proj}_{\mathcal{M}}(\Pi)$, is the unique diffusion process*

$$dX^\star_t = \left\{ f_t(X^\star_t) + v^\star_t(X^\star_t) \right\} dt + \sigma_t\, dB_t,$$

with effective drift

$$v^\star_t(x) = \sigma_t^2\, \mathbb{E}_{\Pi_{t_{i+1}|t}} \left[\nabla \log \mathbb{Q}_t^{|t_i,t_{i+1}}(X_{t_{i+1}} \mid X_t) \,\middle|\, X_t = x \right]$$
$$\overset{\text{Brownian}}{=} \frac{\mathbb{E}_{\Pi_{t_{i+1}|t}}[X_{t_{i+1}} \mid X_t = x] - x}{t_{i+1} - t}$$

By the Markovian projection theorem of Gyöngy (1986), and as further developed in Peluchetti (2023); De Bortoli et al. (2021), the process $M^\star$ is Markov and matches the one-dimensional marginals of the original factorized law Π.

Proposition 3.6 (Variational characterization of the factorized Markovian projection)**.** *Assume that $\sigma_t > 0$. Let $M^\star = \mathrm{proj}_{\mathcal{M}}(\Pi)$ be the Markovian projection of Π as in Definition 3.3. Then:*

$$M^\star = \arg\min_{M \in \mathcal{M}} \left\{ KL(\Pi \,\|\, M) \right\},$$

and

$$KL(\Pi \,\|\, M^\star) = \frac{1}{2} \int_0^T \mathbb{E}_{\Pi_{t_i,t}} \left[\frac{1}{\sigma_t^2} \left\| V - v^\star_t(X_t) \right\|^2 \right] dt$$

with:

$$V = \sigma_t^2\, \mathbb{E}_{\Pi_{t_{i+1}|t}} \left[\nabla \log \mathbb{Q}_t^{|t_i,t_{i+1}}(X_{t_{i+1}} \mid X_t) \,\middle|\, X_t, X_{t_i} \right]$$

In addition, for any $t \in [0,T]$, the time marginal of $M^\star$ coincides with that of Π: $M^\star_t = \Pi_t$. In particular, $M^\star_{t_i} = \Pi_{t_i}$ for all grid points t_i.

Together, these results allow us to alternate between reciprocal and Markovian structures in the multi-marginal setting. Importantly, the Markovian projection admits explicit forward and backward formulations.

Proposition 3.7. *Let $\Pi \in \mathcal{R}^{\otimes}(\mathbb{Q})$. Under mild regularity conditions, the Markovian projection $M^\star = \mathrm{proj}_{\mathcal{M}}(\Pi)$ is associated with the forward SDE*

$$dX_t = \left\{ f_t(X_t) + D_f \right\} dt + \sigma_t dB_t,\ X_{t_i} \sim \mu_{t_i} \tag{5}$$

and with the backward SDE

$$dY_t = \left\{ -f_{t_{i+1}-t}(Y_t) + D_b \right\} dt + \sigma_{t_{i+1}-t} dB_t,\ Y_{t_{i+1}} \sim \mu_{t_{i+1}} \tag{6}$$

where:

$$D_f = \sigma_t^2\, \mathbb{E}_{\Pi_{t_{i+1}|t}} \left[\nabla \log \mathbb{Q}_t^{[t_i,t_{i+1}]}(X_{t_{i+1}} \mid X_t) \,\middle|\, X_t \right]$$
$$D_b = \sigma_{t_{i+1}-t}^2\, \mathbb{E}_{\Pi_{t_i}|t} \left[\nabla \log \mathbb{Q}_t^{[t_i,t_{i+1}]}(Y_{t_i} \mid Y_t) \,\middle|\, Y_t \right]$$

This key result highlights that the Markovian projection can be expressed both in the forward and in the backward direction, allowing us to design an algorithm that jointly leverages both dynamics.

Conjecture 3.1 (Analogue of Léonard (2014) Theorem 2.12)**.** *Let $\mathbb{Q}$ be a Markov reference process. Suppose that P is a Markov path measure such that*

$$P \in \mathcal{R}(\mathbb{Q}), \qquad P_{t_i} = \mu_{t_i}, \quad i = 0, \dots, K.$$

Then P coincides with the unique solution $P^\star$ of the multi-marginal Schrödinger Bridge problem (MMSB) with reference $\mathbb{Q}$.

3.2.2. Iterative Markovian Factorized Fitting

Based on Conjecture 3.1, we propose a novel algorithm called *Iterative Markovian Factorized Fitting* (IMFF) to solve multi-marginal Schrödinger Bridges. We consider a sequence $(\mathbb{P}^n)_{n\in\mathbb{N}}$ such that

$$\mathbb{P}^{2n+1} = \mathrm{proj}_{\mathcal{M}}(\mathbb{P}^{2n}), \ \mathbb{P}^{2n+2} = \mathrm{proj}_{\mathcal{R}^\otimes(\mathbb{Q})}(\mathbb{P}^{2n+1}) \quad (7)$$

with $\mathbb{P}^0$ such that $\mathbb{P}^0_{t_i} = \mu_{t_i}$ for all $i = 0, \dots, K$, and $\mathbb{P}^0 \in \mathcal{R}^\otimes(\mathbb{Q})$. These updates correspond to alternatively performing Markovian projections and factorized reciprocal projections in order to enforce all prescribed marginals.

Lemma 3.1 (Pythagorean identities in the factorized setting)**.** *Under mild assumptions, if $M \in \mathcal{M}$, $\Pi \in \mathcal{R}^\otimes(\mathbb{Q})$ and $KL(\Pi\|M) < +\infty$, we have*

$$KL(\Pi\|M) = KL(\Pi\|\,\mathrm{proj}_{\mathcal{M}}(\Pi)) + KL(\mathrm{proj}_{\mathcal{M}}(\Pi)\|M)$$

Similarly, if $KL(M\|\Pi) < +\infty$, we have

$$KL(M\|\Pi) = KL(M\|\,\mathrm{proj}_{\mathcal{R}^\otimes(\mathbb{Q})}(M)) + KL(\mathrm{proj}_{\mathcal{R}^\otimes(\mathbb{Q})}(M)\|\Pi)$$

Proposition 3.8. *Under mild assumptions, we have*

$$KL(P^{n+1}\,\|\,P^\star) \ \le \ KL(P^n\,\|\,P^\star) < \infty$$
$$\lim_{n\to\infty} KL(P^n\,\|\,P^\star) = 0$$

Hence, for the IMFF sequence $(\mathbb{P}^n)_{n\in\mathbb{N}}$, the Markov path measures $(\mathbb{P}^{2n+1})_{n\in\mathbb{N}}$ are getting closer to the factorized reciprocal class $\mathcal{R}^\otimes(\mathbb{Q})$, while the reciprocal path measures $(\mathbb{P}^{2n+2})_{n\in\mathbb{N}}$ are getting closer to the set of Markov measures. This mirrors the situation in the classical IMF setting, but now in the multi-marginal framework.

Theorem 3.2. *Under mild assumptions, the IMFF sequence $(\mathbb{P}^n)_{n\in\mathbb{N}}$ admits at least one fixed point $\mathbb{P}^\star$, and we have:*

$$\lim_{n\to+\infty} KL(\mathbb{P}^n\,\|\,\mathbb{P}^\star) = 0$$

Moreover, denoting by $\mathbb{P}^{\mathrm{MMSB}}$ the solution of (MMSB) and by $\mathbb{P}^{\mathrm{pair}}$ the gluing of pairwise Schrödinger Bridges, the limit of the IMFF sequence satisfies the inequality:

$$KL(\mathbb{P}^{\mathrm{MMSB}}\,\|\,\mathbb{Q}) \ = \ KL(\mathbb{P}^\star\,\|\,\mathbb{Q}) \ \le \ KL(\mathbb{P}^{\mathrm{pair}}\,\|\,\mathbb{Q})$$

where $\mathbb{Q}$ is the chosen reference process. Thus, $\mathbb{P}^\star$ is the multi-marginal Schrödinger Bridge.

3.2.3. Theoretical algorithm

The Markovian projection necessitates learning one neural drift per direction. Concretely, we solve

$$\underset{\theta}{\mathrm{argmin}}\, \mathbb{E}_b\Big[\big\|v_\theta(X_t,t) - \sigma_t^2\,\mathbb{E}\big[G_t(X_{t_{i+1}})\,\big|\,X_t\big]\big\|^2\Big] \quad (8)$$

for the *forward* drift v_θ, and

$$\underset{\phi}{\mathrm{argmin}}\, \mathbb{E}_b\Big[\big\|v_\phi(Y_t,t) - \sigma_{t_{i+1}-t}^2\,\mathbb{E}\big[G_t(Y_{t_i})\,\big|\,Y_t\big]\big\|^2\Big] \quad (9)$$

for the *backward* drift v_ϕ, where $\mathbb{E}_b$ is over the batch and:

$$G_t(Z) = \nabla \log \mathbb{Q}_t^{[t_{i(t)},t_{i(t)+1}]}(Z \mid Z_t)$$

We summarize in Algorithm 1 our method and provide a practical implementation of IMFF in A.3.

Proposition 3.9. *Suppose the families of functions $\{v_\theta : \theta \in \Theta\}$ and $\{v_\phi : \phi \in \Phi\}$ are rich enough to represent the optimal forward and backward drifts. Let $(P^n, M^n)_{n\in\mathbb{N}}$ be the sequence produced by Algorithm 1. Then, as $n \to \infty$, we have convergence towards an approximate multi-marginal Schrödinger Bridge. Moreover, the Markov law M^n coincides in the limit with the intermediate approximate MMSB solution lying between the true multi-marginal Schrödinger Bridge and the pairwise construction.*

4. Experiments

For all experiments, we employ Brownian motion $(\sigma_t B_t)_{0\le t\le T}$ for the reference measure $\mathbb{Q}$ and $T = N - 1$ where N is the number of marginals. All trainings start after a warmup phase like in Shi et al. (2023), detailed in A.3. Videos for most experiments can be found at mmtsbm.notion.site.

4.1. MMtSBM recovers the exact OT between Gaussian mixtures

In this 2D experiment akin to Liu et al. (2022), we used $N = 3$ mixtures of two standard Gaussians as marginals.

In this configuration the optimal transport *between each pair of marginals* can be computed exactly: it is a pure translation of each Gaussian component inside the mixtures, as we verified with POT (see Figure 7). After only the warm-up phase (akin to flow matching (Lipman et al., 2023)), we can see that the learned transport maps *mix* the Gaussian components of the mixtures, resulting in intersecting trajectories as can be seen in the top row of Figure 1. However, *after* the SB learning phase of MMtSBM, we can see in the bottom row that the learned trajectories do *not* intersect each other anymore and that MMtSBM yields the expected exact

Algorithm 1 Iterative Markovian Factorized Fitting (IMFF)

Input: time grid $0 = t_0 < \cdots < t_K = T$, marginals $(\mu_{t_i})_{i=0}^K$, reference process $\mathbb{Q}$, number of iterations N
Init: choose $\mathbb{P}^0 \in \mathcal{R}^{\otimes}(Q)$ with $\mathbb{P}^0_{t_i} = \mu_{t_i}$ for all i
for $n = 0, \ldots, N-1$ **do**
  Backward Markovian step: learn drift v_ϕ via SDE equation 6, yielding $\mathbb{P}^{2n+1}$ with t_i updated and t_{i+1} from $\mu_{t_{i+1}}$.
  Forward reciprocal projection: $\mathbb{P}^{2n+1} \leftarrow \mathrm{proj}_{\mathcal{R}^{\otimes}(Q)}(\mathbb{P}^{2n+1})$ (cf. Def. 3.2), filling bridges with $\mathbb{Q}$ using t_i from $\mathbb{P}^{2n+1}$ and t_{i+1} from the dataset.
  Forward Markovian step: learn drift v_θ via SDE equation 5, yielding $\mathbb{P}^{2n+2}$ with t_{i+1} updated and t_i from μ_{t_i}.
  Backward reciprocal projection: $\mathbb{P}^{2n+2} \leftarrow \mathrm{proj}_{\mathcal{R}^{\otimes}(Q)}(\mathbb{P}^{2n+2})$ (cf. Def. 3.2), filling bridges with $\mathbb{Q}$ using t_{i+1} from $\mathbb{P}^{2n+2}$ and t_i from the dataset.
end for
Output: learned drifts (v_ϕ, v_θ)

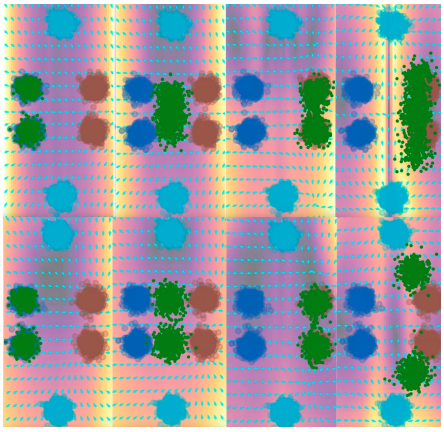

Figure 1. Top row: epoch 0 (only noisy flow matching). Bottom row: epoch 5 (after MMtSBM training). From left to right: snapshots at times $(0, 0.5, 1, 1.3)$. True marginal times $(t_0, t_1, t_2) = (0, 1, 2)$. The order of the 3 true marginals is: $t_0 =$ dark blue; $t_1 =$ red; $t_2 =$ light blue. Generated samples are in green. In the background is the quiver plot of the learned score network.

optimal transport map: pure translations between Gaussian components.

This observation is consistent with the theory: the warm-up phase preserves only the Markov property, while the final learned coupling additionally also preserves the reciprocal property, thus corresponding to the true SB. We empirically observe that the optimality emerges gradually along MMtSBM training epochs: trajectories get *rectified* from epoch 1, become optimal around epoch 5, and consistently remain so after. We will now confirm these visual findings with quantitative metrics in 4.2.

4.2. MMtSBM achieves good usual SB metrics

To quantitatively verify that MMtSBM recovers the correct multi-marginal SB in terms of both 1) static coupling and 2) energy minimization, we extended the now classical "Moons" and "8Gaussians" experiments found in Tong et al. (2024a) and Shi et al. (2023) to our temporal multi-marginal setting in Table 1 (see Figure 8). Choosing $N = 4$, we considered ($\mathcal{N} \to$ Moons $\to \mathcal{N} \to$ Moons), and ($\mathcal{N} \to$ 8Gaussians $\to \mathcal{N} \to$ 8Gaussians). To assess 1) we report the $\mathcal{W}_2$ distance of generations vs test set data at target marginal time(s), averaging along the $N - 1 = 3$ target times for MMtSBM and comparing this to the single bridge setting. To assess 2) we report the full path energy $\mathbb{E}\left[\int_0^T \|v(t, \mathbf{Z}_t)\|^2 \, dt\right]$ where Z_t is the process simulated along the ODE drift 10.

Table 1. Comparison in terms of static coupling ("$\mathcal{W}_2$") and energy minimization ("Path Energy"). The rows marked "×3" correspond to the hypothetical case where the energy of a single bridge is simply tripled, and are included as an ideal baseline for comparison with our actual multi-bridge setting. All metrics apart from ours are from Shi et al. (2023).

	Model	$\mathcal{W}_2$	Path Energy
Moons	Single bridge	$0.144_{\pm 0.024}$	*$1.580_{\pm 0.036}$*
	Single bridge ×3	–	4.740
	MMtSBM (ours)	$0.148_{\pm 0.041}$	$5.350_{\pm 0.085}$
8 $\mathcal{N}$	Single bridge	$0.338_{\pm 0.091}$	*$14.810_{\pm 0.255}$*
	Single bridge ×3	–	44.430
	MMtSBM (ours)	$0.352_{\pm 0.084}$	$46.920_{\pm 0.285}$

We observe that despite a much more complex *time-varying* true transport map to be learned, MMtSBM achieves almost as low $\mathcal{W}_2$ distances as the simple single-bridge setting (3% to 4%), and that our full path energy is within 13% to 6% of the ideal extrapolation of the single bridge result. This validates that MMtSBM manages to approach the true SB in practice.

4.3. MMtSBM scales to $50d$ Gaussian transport

We next proceed to scaling our method to dimension $d = 50$. We follow the setting of Shi et al. (2023) and consider a Gaussian-to-Gaussian transport experiment, ex-

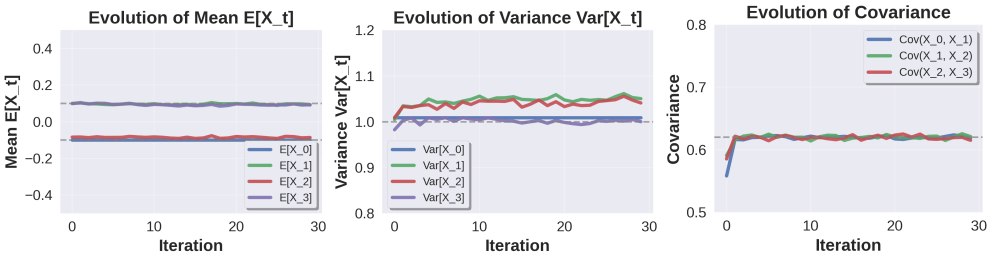

Figure 2. Evolution of mean, variance, and covariance in the multi-marginal $50d$ Gaussian case. Dashed lines are the theoretical values.

tended to our multi-marginal case. Specifically, we prescribe four Gaussian marginals at times $t = 0, 1, 2, 3$: $\mu_0 = \mathcal{N}(-0.1 \cdot \mathbf{1}_d, I_d), \mu_1 = \mathcal{N}(0.1 \cdot \mathbf{1}_d, I_d), \mu_2 = \mathcal{N}(-0.1 \cdot \mathbf{1}_d, I_d), \mu_3 = \mathcal{N}(0.1 \cdot \mathbf{1}_d, I_d)$ where $\mathbf{1}_d \in \mathbb{R}^d$ denotes the vector of all ones, and I_d is the $d \times d$ identity matrix. Since no closed-form solution is available for the static multi-marginal SB, we compare our method to the sequence of theoretical results for each *pairwise* SB (Bunne et al., 2023). As shown in Figure 2, the mean converges rapidly to the prescribed values (0.1 or -0.1) across all four marginals. The variance is slightly more difficult to match: for interior marginals the process tends to overestimate the standard deviation. In contrast, the covariance is consistently well reproduced by our method and remains stable across all three transitions. Interestingly, the covariance converges only after the warmup stage, confirming the added value of the subsequent OT phases. Overall, these results show that MMtSBM scales effectively to the multi-marginal Gaussian setting in $d = 50$.

4.4. MMtSBM achieves SOTA results on $100d$ transcriptomic benchmarks

Table 2. MMD and SWD of generations vs test set for the $d = 100$ EB benchmark. Our generations start from $\mu^{\text{test}}_{t=0}$. Top table: per-marginal metrics. Bottom table: average over all marginals. Others' results are from Chen et al. (2023a). Our error margins are over 10 evaluations while DMSB's are over 3. Best value in **bold**.

Time	DMSB (Chen et al., 2023a) MMD ↓	DMSB (Chen et al., 2023a) SWD ↓	**MMtSBM** (ours) MMD ↓	**MMtSBM** (ours) SWD ↓
t_1	0.021	0.114	**0.016**	**0.104**
t_2	0.029	0.155	**0.020**	**0.139**
t_3	0.038	0.190	**0.020**	**0.127**
t_4	0.034	0.155	**0.020**	**0.143**
Average	0.032 $\pm 3e{-3}$	0.160 $\pm 2e{-2}$	**0.019** $\pm 4e{-4}$	**0.130** $\pm 2e{-3}$

Algorithm	MMD ↓	SWD ↓
NLSB (Koshizuka & Sato, 2023)	0.66	0.54
MIOFlow (Huguet et al., 2022)	0.23	0.35
DMSB (Chen et al., 2023a)	0.032 $\pm 3e{-3}$	0.16 $\pm 2e{-2}$
MMtSBM (ours)	**0.019** $\pm 4e{-4}$	**0.130** $\pm 2e{-3}$

We next evaluate our method on the Embryoid Body (EB) (Moon et al., 2019) and MULTI (Lance et al., 2022) benchmarks, two trajectory inference tasks on real single-cell RNA-seq data. We project RNA counts to their first $d = 100$ principal components for each of the $N = 5$ and $N = 4$ marginals, respectively.[2] We report the Maximum Mean Discrepancy (MMD) and Sliced Wasserstein Distance (SWD) for EB in Table 2 and Table 4, and the Wasserstein-1 distance for MULTI in Table 3. For the EB benchmark we use two evaluation settings: either we train MMtSBM on all marginals (Table 2), or we leave out odd-indexed timesteps (Table 4). For the MULTI benchmark we leave-out one of either intermediate times ($t = 1$ or $t = 2$) during training.

Table 4. MMD and SWD of generations vs held-out times for the $d = 100$ *unnormalized* EB benchmark. Our generations start from $\mu^{\text{test}}_{t=0}$. Figures for SBIRR (Shen et al., 2025), MMFM (Rohbeck et al., 2025), DMSB (Chen et al., 2023a) and 3MSBM (Theodoropoulos et al., 2025) are from Theodoropoulos et al. (2025). Our standard deviations are over 3 runs.

	MMD t_1 ↓	SWD t_1 ↓	MMD t_3 ↓	SWD t_3 ↓
SBIRR	0.71$\pm$0.08	0.80$\pm$0.06	0.73$\pm$0.06	0.91$\pm$0.05
MMFM	0.37$\pm$0.02	0.59$\pm$0.04	0.35$\pm$0.04	0.76$\pm$0.04
DMSB	0.38$\pm$0.04	0.58$\pm$0.06	0.36$\pm$0.07	0.54$\pm$0.06
3MSBM	0.18$\pm$0.01	0.48$\pm$0.04	0.14$\pm$0.04	0.38$\pm$0.03
MMtSBM (ours)	**0.17**$\pm$0.00	**0.45**$\pm$0.03	**0.06**$\pm$0.00	**0.31**$\pm$0.01

On the EB benchmark, our method consistently outperforms baselines on all settings and marginals. It significantly reduces the average MMD by **-41%** and the SWD by **-19%** compared to DMSB when training on all marginals (Table 2), demonstrating superior distribution fitting. It also achieves the best performance on fully held-out times (Table 4), where the interpolative prior also matters.

On the MULTI benchmark, we again reach significantly better average $\mathcal{W}_1$ distances than the directly comparable literature[3], beating the previous directly comparable state-of-the-art ([SF]^{2}M-Geo) by **-15%** with a high statistical significance, and even outperforming methods leveraging

[2] We actually reuse preprocessed data from Tong et al. (2020) and Tong et al. (2024b).

[3] Comparable literature: mainly methods computing the Schrödinger Bridge. Also: methods performing trajectory *inference*, instead of *interpolation*, and those not needing a precomputed OT plan –our goal is to learn a generalizable entropic one!

Table 3. $\mathcal{W}_1$ of generations vs left-out set for the $d = 100$ MULTI benchmark. Reported figures are the average between left-out days 3 and 4. Our error margin is over 3 training runs. Best value in **bold**, second best underlined. See A.5.5 for details & comments.

Method		$\mathcal{W}_1$ (↓)
Schrödinger Bridge		
WLF-SB	(Neklyudov et al., 2024)	$55.065_{\pm 5.499}$
$[SF]^2$M-Exact	(Tong et al., 2024b)	$52.888_{\pm 1.986}$
$[SF]^2$M-Geo	(Tong et al., 2024b)	$52.203_{\pm 1.957}$
MMtSBM	(ours)	$44.542_{\pm 0.637}$
No precomputed OT conditioning		
I-CFM	(Tong et al., 2024a)	$57.262_{\pm 3.855}$
I-MFM$_{\text{RBF}}$	(Kapuśniak et al., 2024)	$54.197_{\pm 1.408}$
WLF-UOT	(Neklyudov et al., 2024)	$54.222_{\pm 5.827}$
I-CDC-FM	(Bamberger et al., 2026)	$54.419_{\pm 0.629}$
MMtSBM	(ours)	$44.542_{\pm 0.637}$

Method		$\mathcal{W}_1$ (↓)
Wasserstein Gradient Flows (WGF)		
WLF-SB	(Neklyudov et al., 2024)	$55.065_{\pm 5.499}$
WLF-OT	(Neklyudov et al., 2024)	$55.416_{\pm 6.097}$
WLF-UOT	(Neklyudov et al., 2024)	$54.222_{\pm 5.827}$
WGF with knowledge of the left-out marginal		
WLF-(OT+potential)	(Neklyudov et al., 2024)	$47.365_{\pm 0.051}$
WLF-(UOT+potential)	(Neklyudov et al., 2024)	$45.231_{\pm 0.010}$
Flow Matching with exact OT conditioning		
OT-CFM	(Tong et al., 2024a)	$54.814_{\pm 5.858}$
OT-MFM$_{\text{RBF}}$	(Kapuśniak et al., 2024)	$50.906_{\pm 4.627}$
OT-CDC-FM	(Bamberger et al., 2026)	$52.043_{\pm 1.948}$
Metric-aware *interpolation* with exact OT conditioning		
GAGA	(Sun et al., 2025)	$\mathbf{27.04}_{\pm \mathbf{2.95}}$

precomputed OT plans and specialized to the transcriptomic setting. This demonstrates the applicability of MMtSBM on pure cellular trajectory inference, despite the absence of restrictive modeling such as spline-valued trajectories, explicitly precomputed OT plan, ad-hoc particularizations, or start *and* end true points trajectory pinning.

4.5. MMtSBM recovers continuous video dynamics from unpaired data

We now evaluate our method on image-space datasets, where the goal is to recover continuous trajectories (*ie videos*) from completely unpaired temporal snapshots.

4.5.1. MNIST DIGIT MORPHING

We conducted experiments on the MNIST dataset of hand-written digits, transporting digits in decreasing order: $4 \rightarrow 3 \rightarrow 2 \rightarrow 1 \rightarrow 0$. The algorithm was trained directly in image space, in dimension $28 \times 28 = 784$. As shown in Figure 3, MMtSBM exhibits clear digit morphing, sometimes reusing pixel structures (e.g., the top of the 3 to form the top of the 2), which is what is expected from OT in pixel space. This experiment thus demonstrates that MMtSBM manages to learn a complex temporal OT map in image space directly.

4.5.2. "BIOTINE" CELL CULTURE

The (in-house) "biotine" dataset consists of 3-channel 128×128 fluorescence images (GFP, membrane, nucleus) of A549 lung epithelial cultured cells, treated with biotin, and imaged over 90 minutes at 7 discrete time steps.

Figure 4 shows the ground truth unpaired dynamic. We can clearly observe fluorescence loss in the cytoplasmic area, corresponding to the green channel. Interestingly, in the generated video (Figure 5), contrary to the above MNIST experiment, a mostly static *positional* evolution is observed.

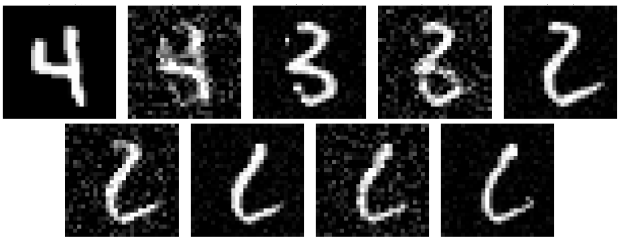

Figure 3. Video generated by MMtSBM on MNIST, backward direction. Starting image is from the test set. Left-to-right, top-to-bottom order: generation at time $t = 4, 3.5, 3, 2.5, 2, 1.5, 1, 0.5, 0$. Integer times are marginal times.

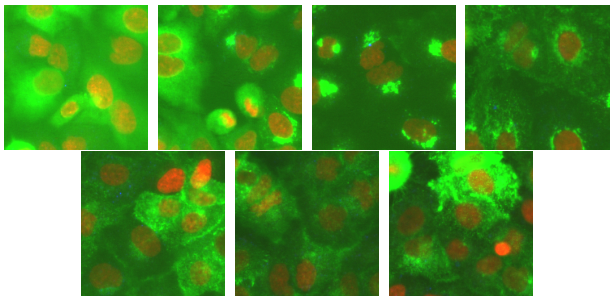

Figure 4. Ground truth biotine examples at training marginal times $t = 0, 1, 2, 3, 4, 5, 6$. Left-to-right, top-to-bottom order.

This stems from the fact that cell position is not a statistically varying information on the biotine dataset, and this (nonexistent) signal is thus simply not seen by our purely unpaired method, resulting in non-moving cells. MMtSBM rather reconstructs the OT trajectory in pixel space, yielding very close cellular position while still accurately matching the time-varying phenotype (mainly: the fluorescence loss in the cytoplasm). As a baseline reference for future works,

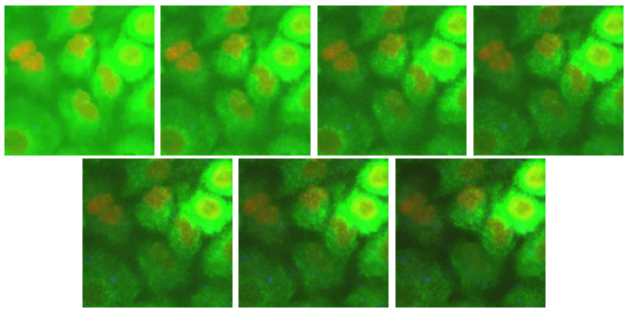

Figure 5. Video generated by MMtSBM on biotine, forward direction. Top-left true starting image from the test set, then left-to-right, top-to-bottom order: generations every $\Delta t = 1$.

Table 5. dinov2-vit-b-14-KIDs for each time, and all times at once.

Time	$t=1$	$t=2$	$t=3$	$t=4$	$t=5$	$t=6$	all times
KID mean (↓)	11.1	13.0	20.1	23.5	26.0	27.7	17.1
KID std	±0.23	±0.20	±0.23	±0.25	±0.29	±0.31	±0.32

we report in Table 5 the KID (Bińkowski et al., 2021) values of true vs generated samples, using DINOv2 (Oquab et al., 2024) as the feature extractor.

4.5.3. KTH Actions

KTH Actions (Schuldt et al., 2004) is a 160×120 small video dataset of human actions with varying subjects and conditions. We selected the "running", "jogging" and "walking" categories, and reprocessed the dataset for globally coherent right-to-left motion by flipping the originally left-to-right videos, using $N = 5$ marginals.

Figure 6. Video generated by MMtSBM on KTH Actions, forward direction. Top-left true starting image from the test set, then left-to-right, top-to-bottom order: generations every $\Delta t = 1$.

As can be seen in Figure 6, MMtSBM clearly learns the global right-to-left motion, demonstrating the capability of our method to generate coherent non-static subject movements from purely unpaired data. Although the generation quality appears limited, we hypothesize that it is due to limited data size (1200 videos total) and compute budget.

To the best of our knowledge, this is the first demonstration of any method performing video generation from purely unpaired data. Together, this provides evidence for both the scalability to very high-dimensional data and for the fidelity to the underlying biological process of MMtSBM.

5. Discussion

In this work we introduce MMtSBM, a novel method that solves the multi-marginal temporal Schrödinger Bridge problem, adapting Bridge Matching (Shi et al., 2023) to our setting. We demonstrate the theoretical soundness of both our modeling and algorithm. We show that MMtSBM indeed produces transport maps that are close to the true OT plan in toy experiments and verify its correct behavior in low-dim experiments. We achieve state-of-the-art results in 2 widely reported single-cell transcriptomic benchmarks, and for the first time demonstrate a method producing temporally coherent videos from purely unpaired data, hoping to lead to many future scientific applications.

Our method has a number of limitations. Compared to competing methods (Tong et al., 2024b; Gushchin et al., 2024), MMtSBM requires simulating full trajectories during training, which is computationally costly. IMFF is also an iterative optimization algorithm that operates on two distinct networks, resulting in potentially unstable training dynamics. It also importantly necessitates fine-tuning the noise scale hyperparameter, for which we have not designed a principled law. Finally, MMtSBM learns a rather naive dynamic induced by the Brownian interpolant. While it appears to yield excellent distribution matching, it also generates noisy trajectories directly in data space.

In future works we would like to investigate other regularizations, such as lifting the process to acceleration space to obtain smoother interpolation trajectories, or exploring other empirical reference processes than the Brownian motion. We also intend to investigate learning the transport map in a latent space. We would also like to explore using the single network theory developed in Bortoli et al. (2024) for efficiency gains, as well as simulation-free methods.

Impact Statement

This paper presents work whose goal is to advance the field of machine learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

Albergo, M. S. and Vanden-Eijnden, E. Building normalizing flows with stochastic interpolants, 2023. URL https://arxiv.org/abs/2209.15571.

Bamberger, J., Jones, I., Duncan, D., Bronstein, M. M., Vandergheynst, P., and Gosztolai, A. Carré du champ flow matching: better quality-generalisation tradeoff in generative models, 2026. URL https://arxiv.org/abs/2510.05930.

Baradat, A. and Léonard, C. Minimizing relative entropy of path measures under marginal constraints, 2020. URL https://arxiv.org/abs/2001.10920.

Bińkowski, M., Sutherland, D. J., Arbel, M., and Gretton, A. Demystifying mmd gans, 2021. URL https://arxiv.org/abs/1801.01401.

Bortoli, V. D., Korshunova, I., Mnih, A., and Doucet, A. Schrödinger bridge flow for unpaired data translation, 2024. URL https://arxiv.org/abs/2409.09347.

Bunne, C., Stark, S. G., Gut, G., Stelzer, E. H., Rätsch, G., and Cuturi, M. Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*, 20(12):1820–1829, 2023.

Burkhardt, D., Luecken, M., Benz, A., Holderrieth, P., Bloom, J., Lance, C., Chow, A., and Holbrook, R. Open problems - multimodal single-cell integration. https://kaggle.com/competitions/open-problems-multimodal, 2022. Kaggle.

Chen, T., Liu, G.-H., Tao, M., and Theodorou, E. A. Deep momentum multi-marginal schrödinger bridge, 2023a. URL https://arxiv.org/abs/2303.01751.

Chen, T., Liu, G.-H., and Theodorou, E. A. Likelihood training of schrödinger bridge using forward-backward sdes theory, 2023b. URL https://arxiv.org/abs/2110.11291.

Chen, Y., Conforti, G., Georgiou, T. T., and Ripani, L. Multi-marginal schrödinger bridges. In Nielsen, F. and Barbaresco, F. (eds.), *Geometric Science of Information*, pp. 725–732, Cham, 2019. Springer International Publishing. ISBN 978-3-030-26980-7.

Csiszár, I. I-divergence geometry of probability distributions and minimization problems. *The Annals of Probability*, 3 (1):146–158, 1975. doi: 10.1214/aop/1176996454. URL https://doi.org/10.1214/aop/1176996454.

Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2013.

De Bortoli, V., Thornton, J., Heng, J., and Doucet, A. Diffusion schrödinger bridge with applications to score-based generative modeling, 2021. URL https://arxiv.org/abs/2106.01357.

Gushchin, N., Kholkin, S., Burnaev, E., and Korotin, A. Light and optimal schrödinger bridge matching. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 17100–17122. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/gushchin24a.html.

Gyöngy, I. Mimicking the one-dimensional marginal distributions of processes having an itô differential. *Probability Theory and Related Fields*, 71:501–516, 1986.

Hong, W., Shi, Y., and Niles-Weed, J. Trajectory inference with smooth schrödinger bridges, 2025. URL https://arxiv.org/abs/2503.00530.

Huguet, G., Magruder, D. S., Tong, A., Fasina, O., Kuchroo, M., Wolf, G., and Krishnaswamy, S. Manifold interpolating optimal-transport flows for trajectory inference, 2022. URL https://arxiv.org/abs/2206.14928.

Kapuśniak, K., Potaptchik, P., Reu, T., Zhang, L., Tong, A., Bronstein, M., Bose, A. J., and Giovanni, F. D. Metric flow matching for smooth interpolations on the data manifold, 2024. URL https://arxiv.org/abs/2405.14780.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization, 2017. URL https://arxiv.org/abs/1412.6980.

Koshizuka, T. and Sato, I. Neural lagrangian schrödinger bridge: Diffusion modeling for population dynamics, 2023. URL https://arxiv.org/abs/2204.04853.

Kullback, S. Probability Densities with Given Marginals. *The Annals of Mathematical Statistics*, 39(4):1236 – 1243, 1968. doi: 10.1214/aoms/1177698249. URL https://doi.org/10.1214/aoms/1177698249.

Lance, C., Luecken, M. D., Burkhardt, D. B., Cannoodt, R., Rautenstrauch, P., Laddach, A., Ubingazhibov, A., Cao, Z.-J., Deng, K., Khan, S., Liu, Q., Russkikh, N., Ryazantsev, G., Ohler, U., data integration competition participants, N. . M., Pisco, A. O., Bloom, J., Krishnaswamy, S., and Theis, F. J. Multimodal single cell data integration challenge: Results and lessons learned. In Kiela, D., Ciccone, M., and Caputo, B. (eds.), *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*, volume 176 of *Proceedings of Machine Learning Research*, pp. 162–176. PMLR, 06–14 Dec 2022. URL https://proceedings.mlr.press/v176/lance22a.html.

Lavenant, H., Zhang, S., Kim, Y.-H., and Schiebinger, G. Toward a mathematical theory of trajectory inference. *The Annals of Applied Probability*, 34(1A):428 – 500, 2024. doi: 10.1214/23-AAP1969. URL https://doi.org/10.1214/23-AAP1969.

Léonard, C. From the schrödinger problem to the monge–kantorovich problem. *Journal of Functional Analysis*, 262(4):1879–1920, 2012.

Léonard, C. A survey of the schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems - A*, 34(4):1533–1574, 2014.

Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling, 2023. URL https://arxiv.org/abs/2210.02747.

Liu, X., Gong, C., and Liu, Q. Flow straight and fast: Learning to generate and transfer data with rectified flow, 2022. URL https://arxiv.org/abs/2209.03003.

Loshchilov, I. and Hutter, F. Decoupled weight decay regularization, 2019. URL https://arxiv.org/abs/1711.05101.

Moon, K. R., van Dijk, D., Wang, Z., Gigante, S., Burkhardt, D. B., Chen, W. S., Yim, K., van den Elzen, A., Hirn, M. J., Coifman, R. R., Ivanova, N. B., Wolf, G., and Krishnaswamy, S. Visualizing structure and transitions in high-dimensional biological data. *Nature Biotechnology*, 37(12):1482–1492, 2019. ISSN 1546-1696. doi: 10.1038/s41587-019-0336-3. URL https://doi.org/10.1038/s41587-019-0336-3.

Neklyudov, K., Brekelmans, R., Tong, A., Atanackovic, L., Liu, Q., and Makhzani, A. A computational framework for solving wasserstein lagrangian flows, 2024. URL https://arxiv.org/abs/2310.10649.

Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.-Y., Li, S.-W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., and Bojanowski, P. Dinov2: Learning robust visual features without supervision, 2024. URL https://arxiv.org/abs/2304.07193.

Park, B. and Lee, J. Multi-marginal schrödinger bridge matching, 2025. URL https://arxiv.org/abs/2510.16587.

Peluchetti, S. Diffusion bridge mixture transports, schrödinger bridge problems and generative modeling. *J. Mach. Learn. Res.*, 24:374:1–374:51, 2023. URL https://api.semanticscholar.org/CorpusID:257912618.

Rohbeck, M., Bunne, C., Brouwer, E. D., Huetter, J.-C., Biton, A., Chen, K. Y., Regev, A., and Lopez, R. Modeling complex system dynamics with flow matching across time and conditions. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=hwnObmOTrV.

Schrödinger, E. Über die umkehrung der naturgesetze. *Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse*, pp. 144–153, 1931.

Schuldt, C., Laptev, I., and Caputo, B. Recognizing human actions: a local svm approach. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 3, pp. 32–36 Vol.3, 2004. doi: 10.1109/ICPR.2004.1334462.

Shen, Y., Berlinghieri, R., and Broderick, T. Multi-marginal schrödinger bridges with iterative reference refinement. In Li, Y., Mandt, S., Agrawal, S., and Khan, E. (eds.), *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 3817–3825. PMLR, 03–05 May 2025. URL https://proceedings.mlr.press/v258/shen25b.html.

Shi, Y., De Bortoli, V., Campbell, A., and Doucet, A. Diffusion schrödinger bridge matching, 2023. URL https://arxiv.org/abs/2303.16852.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021.

Sun, X., Liao, D., MacDonald, K., Zhang, Y., Liu, C., Huguet, G., Wolf, G., Adelstein, I., Rudner, T. G. J., and Krishnaswamy, S. Geometry-aware generative autoencoders for warped riemannian metric learning and generative modeling on data manifolds, 2025. URL https://arxiv.org/abs/2410.12779.

Theodoropoulos, P., Saravanos, A. D., Theodorou, E. A., and Liu, G.-H. Momentum multi-marginal schrödinger bridge matching, 2025. URL https://arxiv.org/abs/2506.10168.

Tong, A., Huang, J., Wolf, G., van Dijk, D., and Krishnaswamy, S. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics, 2020. URL https://arxiv.org/abs/2002.04461.

Tong, A., Fatras, K., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks, J., Wolf, G., and Bengio, Y. Improving and generalizing flow-based generative models with minibatch optimal transport, 2024a. URL https://arxiv.org/abs/2302.00482.

Tong, A., Malkin, N., Fatras, K., Atanackovic, L., Zhang, Y., Huguet, G., Wolf, G., and Bengio, Y. Simulation-free schrödinger bridges via score and flow matching, 2024b. URL https://arxiv.org/abs/2307.03672.

Wang, G., Jiao, Y., Xu, Q., Wang, Y., and Yang, C. Deep generative learning via schrödinger bridge, 2021. URL https://arxiv.org/abs/2106.10410.

A. Appendix

A.1. Additional Background

Reciprocal projection. The reciprocal class $\mathcal{R}(\mathbb{Q})$ consists of mixtures of $\mathbb{Q}$-bridges. For $\mathbb{P} \in \mathcal{P}(C)$,

$$\mathrm{proj}_{\mathcal{R}(\mathbb{Q})}(\mathbb{P}) = \mathbb{P}_{0,T}\, \mathbb{Q}_{|0,T}.$$

Markovian projection. The Markov class $\mathcal{M}$ consists of diffusions $\mathrm{d}X_t = v(t, X_t)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t$. The projection $\mathrm{proj}_{\mathcal{M}}(\Pi)$ has drift

$$\mathrm{d}X_t = \left[\frac{\mathbb{E}_{\Pi}[X_T \mid X_t] - X_t}{T - t}\right]\mathrm{d}t + \sigma\,\mathrm{d}B_t.$$

Variational formulations. Both projections solve KL problems:

$$\mathrm{proj}_{\mathcal{R}(\mathbb{Q})}(\mathbb{P}) = \underset{\Pi \in \mathcal{R}(\mathbb{Q})}{\arg\min}\, \mathrm{KL}(\mathbb{P} \,\|\, \Pi), \qquad \mathrm{proj}_{\mathcal{M}}(\Pi) = \underset{M \in \mathcal{M}}{\arg\min}\, \mathrm{KL}(\Pi \,\|\, M).$$

Bridge matching. In practice, the Markov drift is learned by minimising

$$\mathcal{L}(\theta) = \int_0^T \mathbb{E}_{(X_0,X_T)\sim\Pi_{0,T},\, X_t\sim\mathbb{Q}(\cdot|X_0,X_T)}\left[\|v_\theta(X_t,t) - \tfrac{X_T - X_t}{T-t}\|^2\right] dt.$$

Iterative Proportional Fitting (IPF). IPF alternately enforces marginals by KL minimisation:

$$\mathbb{P}^{2n+1} = \underset{\mathbb{P}:\mathbb{P}_T=\mu_T}{\arg\min}\, \mathrm{KL}(\mathbb{P} \,\|\, \mathbb{P}^{2n}), \quad \mathbb{P}^{2n+2} = \underset{\mathbb{P}:\mathbb{P}_0=\mu_0}{\arg\min}\, \mathrm{KL}(\mathbb{P} \,\|\, \mathbb{P}^{2n+1}).$$

Unlike IMF, this requires caching full trajectories.

A.2. Other properties on IMFF or MMSB

Proposition A.1 (Markov implies reciprocal)**.** *Any Markov measure on $C([0,T],\mathbb{R}^d)$ is reciprocal. Hence $P^\star \in \mathcal{R}(Q)$. See Proposition 2.3 in Léonard (2012).*

Proposition A.2 (Sampling with ODE probability flow)**.** *Given the forward and backward drifts of the multi-marginal Schrödinger bridge, one can simulate trajectories using the probability flow ODE ((Song et al., 2021)):*

$$\frac{dX_t}{dt} = f_t(X_t) - \tfrac{1}{2}\sigma_t^2 \nabla \log p_t(X_t).$$

Although the score function $\nabla \log p_t$ is not directly available, (De Bortoli et al., 2021) show that it can be equivalently recovered by averaging the forward and backward drifts:

$$v_t(x) \;=\; \tfrac{1}{2}\Big(v_t^{\mathrm{fwd}}(x) + v_t^{\mathrm{bwd}}(x)\Big) \tag{10}$$

Simulating the ODE with drift v_t thus yields a deterministic sampling procedure that preserves the marginals of the stochastic bridge, providing an efficient and numerically stable alternative to direct SDE simulation.

A.3. Concrete Algorithms

We always start trainings with a warmup phase, akin to Shi et al. (2023). It allows MMtSBM to start rectifying the trajectories from a non-random state, which could be complicated because the IMFF phase uses the forward/backward network to train the backward/forward one.

Algorithm 2 Warmup (our algorithm)

Input: Subdivision $\{0 = t_0 < t_1 < \cdots < t_n = T\}$, datasets $\{\pi_{t_i}\}$, networks v_θ, v_ϕ, initial params θ, ϕ, batch size B, warmup steps N_{warmup}
Define bridges $\mathcal{B} = \{(t_i, t_{i+1})\}$, $b \leftarrow B/|\mathcal{B}|$
for direction $\in$ {forward, backward} **do**
  for $n \in [\![0, N_{\text{warmup}}]\!]$ **do**
    for all $(t_i, t_{i+1}) \in \mathcal{B}$ **in parallel do**
      Sample $(X_{t_i}, X_{t_{i+1}}) \sim (\pi_{t_i} \otimes \pi_{t_{i+1}})^{\otimes b}$, $t_{(i)} \sim \text{Unif}[t_i, t_{i+1}]^{\otimes b}$
    end for
    Aggregate $X_{\text{init}}, X_{\text{final}}, t$; Sample $Z \sim \mathcal{N}(0, I)^{\otimes B}$
    $X_t \leftarrow \text{Interp}_t(X_{\text{init}}, X_{\text{final}}, Z)$ ▷ cf. equation 11
    Update θ if forward with ℓ^{fwd} equation 13, else ϕ with ℓ^{bwd} equation 14
  end for
end for
Output: Warmup parameters θ, ϕ

Algorithm 3 Iterative Markovian Factorized Fitting (IMFF) (our algorithm)

Input: Subdivision $\{0 = t_0 < t_1 < \cdots < t_n = T\}$, datasets $\{\pi_{t_i}\}$, networks v_θ, v_ϕ, warmup params θ, ϕ, batch size B, finetune steps N_{finetune}, inner steps N_{inner}
Define bridges $\mathcal{B} = \{(t_i, t_{i+1})\}$, $b \leftarrow B/|\mathcal{B}|$
for $N \in [\![0, N_{\text{finetune}}]\!]$ **do**
  for all $(t_i, t_{i+1}) \in \mathcal{B}$ **in parallel do**
    Sample $(X_{t_i}, X_{t_{i+1}})$ from $(\pi_{t_i} \otimes \pi_{t_{i+1}})^{\otimes b}$
    Sample $t_{(i)} \sim \text{Unif}[t_i, t_{i+1}]^{\otimes b}$
  end for
  Aggregate $X_{\text{init}}, X_{\text{final}}, t$
  for direction $\in$ {backward, forward} **do**
    for $n \in [\![0, N_{\text{inner}}]\!]$ **do**
      Sample $Z \sim \mathcal{N}(0, I)^{\otimes B}$
      if direction = forward **then**
        $\hat{X}_{\text{init}} \leftarrow \text{SDE}(X_{\text{final}}, v_\phi)$ ▷ cf. equation 12
        $X_t \leftarrow \text{Interp}_t(\hat{X}_{\text{init}}, X_{\text{final}}, Z)$ ▷ cf. equation 11
        Update θ with ℓ^{fwd} equation 13
      else
        $\hat{X}_{\text{final}} \leftarrow \text{SDE}(X_{\text{init}}, v_\theta)$ ▷ cf. equation 12
        $X_t \leftarrow \text{Interp}_t(X_{\text{init}}, \hat{X}_{\text{final}}, Z)$ ▷ cf. equation 11
        Update ϕ with ℓ^{bwd} equation 14
      end if
    end for
  end for
end for
Output: Finetuned parameters θ, ϕ

A.4. Critical Implementation Considerations

A naive implementation of the algorithm quickly led to the *forgetting* of paths between marginals as training progressed. To overcome this, we developed a fully vectorized implementation that ensures stable learning across all intervals. This design is essential for the quality of our solution. Key components are detailed below.

A.4.1. Scalability with High Dimensions and Many Marginals

Both Markovian and reciprocal projections are implemented in a fully vectorized manner. Instead of looping over intervals, all pairs are aggregated into global vectors and processed simultaneously on GPU.

At iteration n, for interval $[t_i, t_{i+1}]$, pairs are sampled as

$$z_i \sim (M^n)_{t_i}, \quad z_{i+1} \sim \mu_{i+1} \quad \text{(forward)}, \qquad z_{i+1} \sim (M^n)_{t_{i+1}}, \quad z_i \sim \mu_i \quad \text{(backward)}.$$

Pairs from all intervals form two batched vectors $(Z_{\text{init}}, Z_{\text{final}})$. Each bridge is then simulated in parallel as

$$X_t^{(b)} = (1-s)\, z_{\text{init}}^{(b)} + s\, z_{\text{final}}^{(b)} + \sigma_t \sqrt{s(1-s)}\, \xi^{(b)}, \qquad \xi^{(b)} \sim \mathcal{N}(0, I).$$

This parallelization makes multi-marginal training feasible at scale.

A.4.2. Masking and Time Discretization

The horizon $[0, T]$ is discretized into N_{total} steps, allocated proportionally to interval length:

$$N_i = \left\lfloor N_{\text{total}} \tfrac{t_{i+1}-t_i}{T} \right\rfloor, \quad dt_i = \pm \tfrac{\Delta\tau}{t_{i+1}-t_i}, \; \Delta\tau = \tfrac{T_{\max}-T_{\min}}{N_{\text{total}}}.$$

This ensures consistent integration with bounded cost.

Since N_i varies across intervals, all trajectories are embedded into a common tensor of shape `(num_bridges,max_N)` with binary masks:

$$z_{k+1}^{(b)} = z_k^{(b)} + v(z_k^{(b)}, t_k^{(b)})\, dt^{(b)} + \sigma_{t_k^{(b)}} \sqrt{dt^{(b)}}\, \xi^{(b)},$$

updated only where `mask`=1. This allows heterogeneous bridges to evolve in a single GPU loop.

A.4.3. Interpolation Operator and Losses

For each bridge (t_i, t_{i+1}) and batch B, define

$$\mathbf{s} = \frac{\mathbf{t} - t_{\text{init}}}{t_{\text{final}} - t_{\text{init}}} \in [0,1]^B.$$

Then the interpolation is

$$\text{Interp}_{\mathbf{t}}(X_{\text{init}}, X_{\text{final}}, Z) = (1-\mathbf{s}) \odot X_{\text{init}} + \mathbf{s} \odot X_{\text{final}} + \sqrt{\varepsilon(1-\mathbf{s}) \odot \mathbf{s}} \odot Z, \tag{11}$$

with $\odot$ the elementwise product.

We also define a generic simulation operator for SDEs. Given an initial condition X_{init} and a drift $v_{\text{direction}}$ (either forward or backward), we denote

$$\text{SDE}(X_{\text{init}}, v_{\text{direction}}): \quad dX_t = v_{\text{direction}}(t, X_t)\, dt + \sigma_t\, dB_t, \qquad X_{t_{\text{init}}} = X_{\text{init}}. \tag{12}$$

This operator returns a trajectory $(X_t)_{t \in [t_{\text{init}}, t_{\text{final}}]}$.

Forward/backward losses enforce vectorized drift consistency:

$$\ell^{\text{fwd}}(\theta; \mathbf{t}, X_{\text{final}}, X_t) = \tfrac{1}{B} \left\| v_\theta(\mathbf{t}, X_t) - \tfrac{X_{\text{final}} - X_t}{t_{\text{final}} - \mathbf{t}} \right\|^2 \tag{13}$$

$$\ell^{\text{bwd}}(\phi; \mathbf{t}, X_{\text{init}}, X_t) = \tfrac{1}{B} \left\| v_\phi(\mathbf{t}, X_t) - \tfrac{X_{\text{init}} - X_t}{\mathbf{t} - t_{\text{init}}} \right\|^2 \tag{14}$$

A.4.4. Time-Dependent Drift Networks

The drifts v_θ, v_ϕ are parameterized by networks with explicit time encodings (sinusoidal, Gaussian Fourier, FiLM). This enables (i) generalization across intervals through parallel training, and (ii) sensitivity to local temporal position, ensuring bridge consistency and global coherence.

A.5. Experiments details

The Adam (Kingma & Ba, 2017) or AdamW (Loshchilov & Hutter, 2019) optimizer is used throughout experiments. Unless stated otherwise, we employ a learning rate of $2 * 10^{-4}$, and SiLU activations are applied on each layer.

A.5.1. Exact OT between Gaussian mixtures ↑↑

In Figure 7 we can see the (*exact*) "glued" OT plan empirically computed with POT. Observe how the global trajectory transports each Gaussian component of the mixture to a single other Gaussian component of the next marginal, yielding paths without any crossing. Note that the *true multi-marginal* transport plan remains unknown even in this simple Gaussian mixture setting.

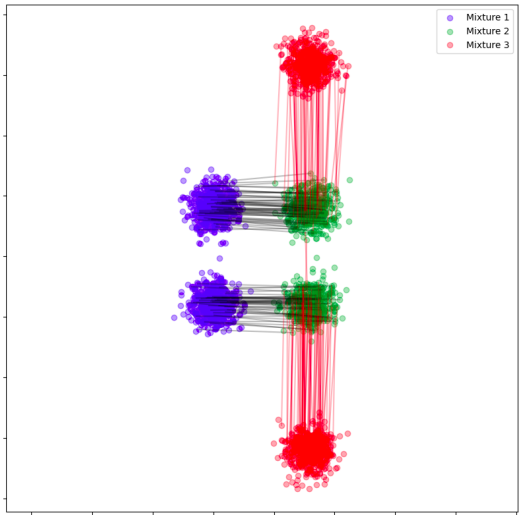

Figure 7. Here we computed the OT plan between each pair of adjacent marginals empirically, in red and black lines. This plan can serve as a good proxy for the true multi-marginal plan.

A.5.2. 8Gaussians and Moons experiment ↑↑

We used the same experimental setting as (Shi et al., 2023), except that we increase the batch size proportionally to the number of intermediate bridges. The 2-Wasserstein distance are computed with `pot` and the integrated path energy are computed with $\mathbb{E}\left[\int_0^T \|v(t, \mathbf{Z}_t)\|^2 \, dt\right]$ where Z_t is the process simulated along the ODE drift 10.

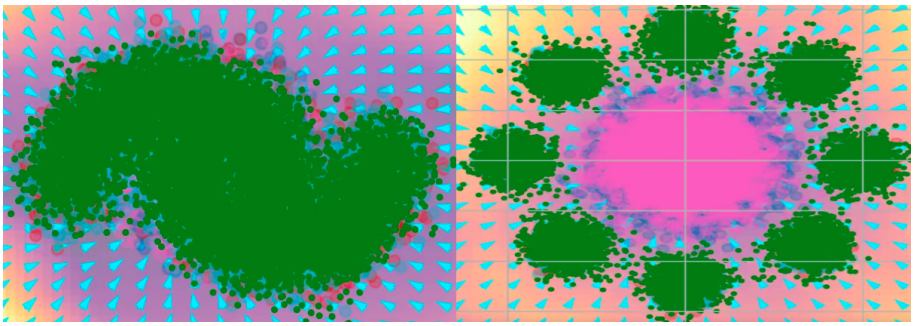

Figure 8. Third marginal fitting for the moons and 8-Gaussian trajectories. Blue vectors indicate the drift direction, with gradient intensity showing vector field strength; green points denote moving samples, and pink highlights the Gaussian fitted along the trajectory.

A.5.3. $50d$ GAUSSIAN EXPERIMENTS ↑↑

On an NVIDIA A100 GPU, the full training took approximately 300 minutes for 30 outer iterations, each with 10,000 training steps and 20 diffusion steps per bridge.

A.5.4. $100d$ TRANSCRIPTOMIC EXPERIMENTS: EMBRYOID BODY ↑↑

The dataset comprises 5 timepoints (each of them being the aggregation of 2 days, from day 0 to day 24) covering the progression from a homogeneous stem-cell population toward mesoderm, endoderm, and ectoderm precursors. The Embryoid Body dataset thus constitutes a realistic and challenging testbed for Schrödinger Bridge methods, combining high dimensionality, non-Gaussian distributions, and branching lineages. We preprocessed the data following Tong et al. (2020).

During MMtSBM training all timepoints are standardized (zero mean and unit variance). We then used two distinct settings to fairly evaluate our method: in the first one we compute metrics in the normalized space, while in the second one we denormalize the data before computing metrics in order to create a comparable setting to Theodoropoulos et al. (2025). In both cases 1000 samples were withheld from each timepoint to form a test set used for evaluating the Maximum Mean Discrepancy (MMD) and the Sliced Wasserstein Distance (SWD) between test set and generated samples. Additionally, in Table 4, odd-indexed timepoints are fully removed from the training data.

We trained a network of about $300k$ parameters for 20 outer iterations with 20,000 inner iterations.

We show in Table 6 the performance advantage of our method compared to an iterative algorithm such as Chen et al. (2023a).

Table 6. Training and sampling times for Chen et al. (2023a) and MMtSBM (ours) in dimension 100.

	DSBM (Chen et al., 2023a)		MMtSBM (ours)	
number of marginals	4	5	4	5
Train time	33min	44 min	20 min	32 min
Sampling time	2.00s	2.02s	2.00 s	2.00 s

A.5.5. $100d$ TRANSCRIPTOMIC EXPERIMENTS: MULTI ↑↑

We reused the preprocessed data from (Tong et al., 2024b). We do not whiten it. We conducted a minimal sweep to select the best σ (0.3). The network is a simple 3-layers MLP with around $500k$ parameters and we employ 150 discretization time steps in total. Metrics are computed over $1k$ true test samples vs $1k$ generated samples, where these generations themselves come from the previous test marginal (μ_{i-1}^{test} if i is the left-out time). We trained 3 models with different seeds for each left-out time (either $t = 1$ or $t = 2$, corresponding to days 3 and 4). Our reported standard deviation is the pooled variance of the best same-hyperparameters $\sigma = 0.3$ models over 2 groups, each group corresponding to a left-out time. Other papers seem to have reported the overall variance, which we think makes less sense given the structure of the problem.

Table 7. Per-group statistics with pooled standard deviation $s_{\text{pooled}} = \sqrt{\sum(n_i - 1)s_i^2 \;/\; \sum(n_i - 1)}$, where n_i and s_i are the sample size and standard deviation of each group.

Group	Number of runs	Mean	Std
Leave-out & test $t = 1$	3	37.026	0.822
Leave-out & test $t = 2$	3	52.059	0.367
Global	6	44.542	0.637

About other methods reported in Table 3: only I-CFM, I-MFM$_{\text{RBF}}$, and MMtSBM (ours) do not rely on a precomputed OT plan, be it exact or approximate. GAGA (Sun et al., 2025) performs *interpolation between 2 true pinned endpoints* in the latent space of a metric-aware autoencoder trained with the true exact OT plan; we thus still claim SOTA, either within methods solving the SB, or within methods doing "pure" trajectory *inference* (without a pinned true endpoint).

A.5.6. BIOTINE CELL CULTURE ↑↑

We perform learning directly in image space at $3 \times 128 \times 128$ definition with 3M parameters UNets. We also experimented with learning in a VAE latent space but produced images were more blurry.

The model trains in only 5 hours and subsequently generates an entire 602-frame trajectory in just 32 seconds, demonstrating

Table 8. Training and sampling statistics for video generation on the Biotine dataset.

Dataset	Biotine
Dimension	$128 \times 128 \times 3 = 49,152$
Number of marginals	7
Training time	5 h
Number of epochs	5
Sampling time	32 s
Generated frames	602

both low training cost and highly efficient sampling.

A.5.7. KTH Actions ↑↑

We perform learning directly in image space at $1 \times 160 \times 120$ definition with 14M parameters purely convolutional UNets. For training we use the AdamW optimizer, $100,000$ warmup steps, a noise scale $\sigma = 0.5$, a learning rate of $2e-4$, $10,000$ inner steps and 30 outer iterations. Videos are generated with 400 steps.

A.6. Proofs

A.6.1. Definition 3.2

Proof of variational proposition in Definition 3.2 (variational characterization). By the additive property of the KL divergence (Léonard, 2014), for any $P \in \mathcal{P}(C([0,T],\mathbb{R}^d))$ and $\Pi \in \mathcal{R}^{\otimes}(\mathbb{Q})$, we can write

$$KL(P \,\|\, \Pi) = KL(P_{t_0,\dots,t_K} \,\|\, \Pi_{t_0,\dots,t_K}) + \mathbb{E}_{P_{t_0,\dots,t_K}}\Big[KL\Big(P^{x_0,\dots,x_K}_{[0,T]} \,\|\, \bigotimes_{i=0}^{K-1} \mathbb{Q}^{x_i,x_{i+1}}_{[t_i,t_{i+1}]}\Big)\Big],$$

where $P^{x_0,\dots,x_K}_{[0,T]}$ denotes the conditional law of P given its values at the grid points $(t_0,\dots,t_K)$.

Restricting to Π such that $\Pi_{t_0,\dots,t_K} = P_{t_0,\dots,t_K}$ cancels the first KL term, and then the minimizer is uniquely obtained by replacing the conditional path law of P with the tensor product of Q-bridges between each (x_i, x_{i+1}).

Hence the optimal projection is

$$\Pi^\star = P_{t_0,\dots,t_K} \bigotimes_{i=0}^{K-1} \mathbb{Q}^{x_i,x_{i+1}}_{[t_i,t_{i+1}]},$$

which is exactly the definition of the factorized reciprocal projection. □

A.6.2. Definition 3.3

Proof of proposition in the Definition 3.3 in the Brownian case. By Definition 3.3, the effective drift is

$$v_t^\star(x) = \sigma_t^2\, \mathbb{E}_{\Pi_{t_{i+1}|t}}\Big[\nabla \log Q_t^{|t_i,t_{i+1}}(X_{t_{i+1}} \mid X_t) \,\Big|\, X_t = x\Big].$$

For a Brownian reference process, the transition kernel is Gaussian,

$$Q_t^{|t_i,t_{i+1}}(y \mid x) = \frac{1}{(2\pi\sigma^2(t_{i+1}-t))^{d/2}} \exp\Big(-\tfrac{\|y-x\|^2}{2\sigma^2(t_{i+1}-t)}\Big),$$

so that

$$\nabla_x \log Q_t^{|t_i,t_{i+1}}(y \mid x) = \frac{y-x}{\sigma^2(t_{i+1}-t)}.$$

Plugging this into the definition yields

$$v_t^\star(x) = \sigma_t^2\, \mathbb{E}\Big[\tfrac{X_{t_{i+1}}-x}{\sigma^2(t_{i+1}-t)} \,\Big|\, X_t = x\Big].$$

In the Brownian case $\sigma_t^2 = \sigma^2$, which simplifies to

$$v_t^\star(x) = \frac{\mathbb{E}[X_{t_{i+1}} \mid X_t = x] - x}{t_{i+1} - t},$$

as claimed. □

A.6.3. Proposition 3.1

Proof of Proposition 3.1. The feasible set

$$\mathcal{A} = \{P : P \ll Q,\ P_{t_i} = \mu_{t_i},\ i = 0, \dots, K\}$$

is convex. Since the functional $P \mapsto D_{\mathrm{KL}}(P\|Q)$ is strictly convex, there is at most one minimizer.

To show existence, observe that $\mathcal{A}$ is non-empty. Indeed, consider any coupling γ of $(\mu_{t_0}, \dots, \mu_{t_K})$. For each pair (x_i, x_{i+1}), let $Q^{x_i,x_{i+1}}_{[t_i,t_{i+1}]}$ denote the Brownian bridge of Q conditioned on $X_{t_i} = x_i$ and $X_{t_{i+1}} = x_{i+1}$. Then the measure

$$P = \int \bigotimes_{i=0}^{K-1} Q^{x_i,x_{i+1}}_{[t_i,t_{i+1}]}\, d\gamma(x_0, \dots, x_K)$$

belongs to $\mathcal{A}$. Hence the admissible set is non-empty.

Therefore, (MMSB) admits a unique solution $P^\star$. □

A.6.4. Proposition 3.2

Proof of Proposition 3.2. The argument is identical to Proposition 2.10 in Léonard (2014), extended to the multi-marginal setting. For any admissible path measure $P \ll Q$, the additivity property of the relative entropy gives

$$KL(P \,\|\, Q) = KL(P_{t_0,\dots,t_K} \,\|\, Q_{t_0,\dots,t_K}) + \mathbb{E}_{P_{t_0,\dots,t_K}}[KL(P(\cdot \mid X_{t_0}, \dots, X_{t_K}) \,\|\, Q(\cdot \mid X_{t_0}, \dots, X_{t_K}))]\,.$$

Since the second term is always nonnegative, minimizing the dynamic problem is equivalent to minimizing the static one. Moreover, the inequality becomes an equality if and only if

$$P(\cdot \mid X_{t_0}, \dots, X_{t_K}) = Q(\cdot \mid X_{t_0}, \dots, X_{t_K}), \quad P_{t_0,\dots,t_K}\text{-a.s.}$$

Hence the optimal dynamic solution $P^\star$ is uniquely obtained from the optimal static solution $\pi^\star$ by gluing the conditional bridges of Q, which establishes the equivalence. □

A.6.5. Proposition 3.3

Proof of Proposition 3.3. We follow the argument of (Léonard, 2014, Prop. 2.10). Fix an intermediate time t_k with $0 < k < n$. For any $Q \in \mathcal{P}(\Omega)$ and $z \in X$, set

$$Q^{t_k,z}_{[0,t_k]} := Q(X_{[0,t_k]} \in \cdot \mid X_{t_k} = z), \qquad Q^{t_k,z}_{[t_k,1]} := Q(X_{[t_k,1]} \in \cdot \mid X_{t_k} = z).$$

Let $\mu \in \mathcal{P}(X)$ and for each $z \in X$ prescribe $Q^<_z \in \mathcal{P}(\Omega_{[0,t_k]} \cap \{X_{t_k} = z\})$, $Q^>_z \in \mathcal{P}(\Omega_{[t_k,1]} \cap \{X_{t_k} = z\})$. By the entropy additivity property (see formula (A.8) in Léonard (2014)), the measure

$$P^* = \int_X Q^<_z \otimes Q^>_z\, \mu(dz)$$

is the unique minimizer of $H(\cdot \mid R)$ under these constraints, and it satisfies

$$P^*_{[t_k,1]}(\cdot \mid X_{[0,t_k]}) = P^*_{[t_k,1]}(\cdot \mid X_{t_k}).$$

This is exactly the Markov property at time t_k.

Now apply this to $Q = \widehat{P}$, the solution of the multi-marginal Schrödinger problem. If $\widehat{P}$ were not Markov, one could construct a measure P^* with the same time-marginal constraints but strictly smaller entropy, a contradiction with the definition of a minimizer. Since t_k was arbitrary, $\widehat{P}$ must be Markov at all grid times $t_0, \dots, t_n$, hence Markov on $[0, 1]$. □

A.6.6. PROPOSITION 3.4

Proof of Proposition 3.4. The argument is a direct extension of Theorem 2.8 and Proposition 2.10 in Léonard (2014).

Assume that the reference law $Q_{t_0,\dots,t_K}$ satisfies the usual regularity conditions: (i) each one-time marginal coincides with a reference measure m; (ii) there exists a nonnegative function A such that

$$Q_{t_0,\dots,t_K}(dx_0,\dots,dx_K) \;\geq\; \exp\Big(-\sum_{i=0}^{K} A(x_i)\Big)\, m(dx_0)\cdots m(dx_K);$$

(iii) there exists B such that

$$\int_{\mathcal{X}^{K+1}} \exp\Big(-\sum_{i=0}^{K} B(x_i)\Big)\, Q_{t_0,\dots,t_K}(dx_0,\dots,dx_K) < \infty;$$

(iv) either $m^{\otimes(K+1)} \ll Q_{t_0,\dots,t_K}$ or the converse holds. Suppose further that the prescribed marginals $(\pi_{t_0},\dots,\pi_{t_K})$ satisfy $H(\pi_{t_i}\,|\,m) < \infty$,

$$\sum_{i=0}^{K} \int (A+B)(x)\, d\pi_{t_i}(x) < \infty,$$

and that they are internal in the sense of Proposition 2.6 of (Léonard, 2014).

Under these assumptions, the dual problem is well posed. Introducing Lagrange multipliers $(\varphi_i)_{i=0}^{K}$ for the marginal constraints, convex duality shows that the minimizer $\pi^\star$ of the static problem is absolutely continuous with respect to $Q_{t_0,\dots,t_K}$ with density

$$\frac{d\pi^\star}{dQ_{t_0,\dots,t_K}}(x_0,\dots,x_K) = \exp\Big(\sum_{i=0}^{K} \varphi_i(x_i)\Big).$$

Defining $f_i(x_i) := e^{\varphi_i(x_i)}$ yields the factorized form

$$\frac{d\pi^\star}{dQ_{t_0,\dots,t_K}}(x_0,\dots,x_K) = \prod_{i=0}^{K} f_i(x_i).$$

□

A.6.7. PROPOSITION 3.6

Proof of Proposition 3.6. The argument is the same as in the two-marginal case (Shi et al., 2023, Prop. 2), except that all computations must now be performed interval by interval along the grid $t_0 < \cdots < t_K$. Under Assumptions A1–A3, the Doob–h transform is well-defined on each interval $[t_i, t_{i+1}]$ and Lemma 11 of (Shi et al., 2023) applies verbatim. The only change is that the terminal conditioning in the backward equation is at t_{i+1} instead of T. This yields the drift

$$v_t^{\Pi}(x) = \sigma_t^2\, \mathbb{E}_{\Pi}\Big[\nabla \log Q^{t_i,t_{i+1}}_{t_{i+1}|t}(X_{t_{i+1}} \mid X_t)\,\big|\, X_{t_i}, X_t\Big], \qquad t \in [t_i, t_{i+1}].$$

Hence the dynamics of Π is piecewise independent: its increment on $[t_i, t_{i+1}]$ depends only on the local bridge $Q^{t_i,t_{i+1}}$.

The same interval-wise independence holds for any Markov $M \in \mathcal{M}$, whose SDE also factorizes on the grid. Thus both Π and M have product decompositions over the intervals, and their Radon–Nikodym derivative factorizes multiplicatively,

$$\frac{d\Pi}{dM} = \prod_{i=0}^{K-1} \frac{d\Pi^{(i)}}{dM^{(i)}}.$$

Taking logarithms and integrating with respect to Π gives the additivity of the relative entropy,

$$KL(\Pi\|M) = \sum_{i=0}^{K-1} KL(\Pi^{(i)}\|M^{(i)}).$$

For each interval $[t_i, t_{i+1}]$, using the conditional expectation identity as in the proof of (Shi et al., 2023), we have for every $t \in [t_i, t_{i+1}]$,

$$\mathbb{E}_{\Pi_{t_i,t}}\left[\left\|\sigma_t^2\, \mathbb{E}_{\Pi_{t_{i+1}|t_i,t}}[\nabla \log Q_{t_{i+1}|t}^{t_i,t_{i+1}}(X_{t_{i+1}} \mid X_t) \mid X_t, X_{t_i}] - v_t(X_t)\right\|^2\right]$$
$$\geq \mathbb{E}_{\Pi_{t_i,t}}\left[\left\|\sigma_t^2 \mathbb{E}_{\Pi_{t_{i+1}|t}}[\nabla \log Q_{t_{i+1}|t}^{t_i,t_{i+1}}(X_{t_{i+1}} \mid X_t) \mid X_t, X_{t_i}] - v_t^\star(X_t)\right\|^2\right],$$

where the optimal drift is defined by the orthogonal projection

$$v_t^\star(x) = \sigma_t^2\, \mathbb{E}_{\Pi_{t_{i+1}|t}}\left[\nabla \log Q_{t_{i+1}|t}^{t_i,t_{i+1}}(X_{t_{i+1}} \mid X_t) \mid X_t = x_t\right].$$

Using (Léonard, 2012), Theorem 2.3 on each interval and summing the contributions gives

$$KL(\Pi \| M^\star) = \frac{1}{2}\sum_{i=0}^{K-1}\int_{t_i}^{t_{i+1}} \mathbb{E}_{\Pi_t}\left[\|v_t^\Pi(X_t) - v_t^\star(X_t)\|^2/\sigma_t^2\right] dt.$$

Finally, the same Fokker–Planck uniqueness argument as in (Shi et al., 2023) ensures that $M_t^\star = \Pi_t$ for all $t \in [t_i, t_{i+1}]$ and all i. Since the grid points are included, this implies $M^\star = \Pi$, which concludes the proof. □

A.6.8. Lemma 3.1

Proof of Lemma 3.1. For the Markovian part, the equality follows analogously to the proof of (Shi et al., 2023).

For each interval $[t_i, t_{i+1}]$, the same quadratic expansion gives

$$2\,KL\Big(\Pi^{(i)} \,\Big\|\, M^{(i)}\Big) = 2\,KL\Big(\Pi^{(i)} \,\Big\|\, \mathrm{proj}_{\mathcal{M}}(\Pi)^{(i)}\Big) + 2\,KL\Big(\mathrm{proj}_{\mathcal{M}}(\Pi)^{(i)} \,\Big\|\, M^{(i)}\Big).$$

Summing this identity over $i = 0, \ldots, K-1$, using the interval-wise independence, yields

$$2\,KL(\Pi \| M) = 2\,KL(\Pi \,\|\, \mathrm{proj}_{\mathcal{M}}(\Pi)) + 2\,KL(\mathrm{proj}_{\mathcal{M}}(\Pi) \,\|\, M),$$

which is the desired result.

For the factorized reciprocal part :

Let $\Pi \in \mathcal{R}^{\otimes}(Q)$ and denote by

$$\Pi^\star = \mathrm{proj}_{\mathcal{R}^{\otimes}(Q)}(\mathbb{P}) = \mathbb{P}_{t_0,\ldots,t_K} \otimes_{i=0}^{K-1} Q_{[t_i,t_{i+1}]}^{x_i,x_{i+1}}.$$

We have the Radon–Nikodym factorization

$$\frac{d\mathbb{P}}{d\Pi} = \frac{d\mathbb{P}}{d\Pi^\star} \cdot \frac{d\Pi^\star}{d\Pi}(X_{t_0}, \ldots, X_{t_K}).$$

By integrating w.r.t. $\mathbb{P}$ and applying Csiszár's Pythagorean identity (Csiszár, 1975, Eq. 2.6), we obtain

$$KL(\mathbb{P}\|\Pi) = KL(\mathbb{P}\|\Pi^\star) + \int \log \frac{d\Pi^\star}{d\Pi}(x_0, \ldots, x_K)\, d\mathbb{P}_{t_0,\ldots,t_K}.$$

Since $\mathbb{P}_{t_0,\ldots,t_K} = \Pi^\star_{t_0,\ldots,t_K}$, the second term equals

$$\int \log \frac{d\Pi^\star}{d\Pi}(x_0, \ldots, x_K)\, d\Pi^\star_{t_0,\ldots,t_K} = KL(\Pi^\star\|\Pi).$$

Thus

$$KL(\mathbb{P}\|\Pi) = KL(\mathbb{P}\|\Pi^\star) + KL(\Pi^\star\|\Pi),$$

which concludes the proof. □

A.6.9. PROPOSITION 3.7

Proof of Proposition 3.7. It follows from the fact that the time-reversal map $\mathcal{T} : \Omega \to \Omega$ is a bijection, and by reversibility of the reference process $\mathbb{Q}$ we have, for any probability measure $\mathbb{P} \in \mathcal{P}(C)$,

$$KL(\mathbb{P} \,\|\, \mathbb{Q}) = KL(\mathbb{P} \circ \mathcal{T} \,\|\, \mathbb{Q} \circ \mathcal{T}) = KL(\mathbb{P} \circ \mathcal{T} \,\|\, \mathbb{Q}).$$

To prove the direction "$\Longrightarrow$", assume $\mathbb{P} \in \mathcal{R}^{\otimes}(\mathbb{Q})$ is the minimizer of the forward problem. Then, for any $\Pi \in \mathcal{R}^{\otimes}(\mathbb{Q})$ we have $\Pi \circ \mathcal{T} \in \mathcal{R}^{\otimes}(\mathbb{Q})$, and

$$KL(\Pi \,\|\, \mathbb{Q}) = KL(\Pi \circ \mathcal{T} \,\|\, \mathbb{Q} \circ \mathcal{T}) \;\; \geq \;\; KL(\mathbb{P} \circ \mathcal{T} \,\|\, \mathbb{Q} \circ \mathcal{T}) = KL(\mathbb{P} \,\|\, \mathbb{Q}).$$

The reverse direction follows by symmetry, replacing $\mathbb{P}$ with $\mathbb{P} \circ \mathcal{T}$. Thus, working with forward or backward processes is equivalent up to the bijection $\mathcal{T}$, and the KL minimization problem is unchanged. In particular, this justifies that alternating forward and backward steps in the IMFF algorithm is well-defined and analogous to IPF. □

A.6.10. PROPOSITION 3.8

Proof of Proposition 3.8, first claim. As a reminder, we follow the same argument as in (Shi et al., 2023) and (De Bortoli et al., 2021). Applying Lemma 3.1, for any $N \in \mathbb{N}$ we obtain

$$KL(\mathbb{P}^0 \,\|\, \mathbb{P}^\star) = KL(\mathbb{P}^0 \,\|\, \mathbb{P}^1) + KL(\mathbb{P}^1 \,\|\, \mathbb{P}^2) + \cdots + KL(\mathbb{P}^N \,\|\, \mathbb{P}^\star).$$

Since each term is nonnegative, we deduce the monotonicity

$$KL(\mathbb{P}^{n+1} \,\|\, \mathbb{P}^\star) \;\leq\; KL(\mathbb{P}^n \,\|\, \mathbb{P}^\star),$$

and boundedness $KL(\mathbb{P}^n \,\|\, \mathbb{P}^\star) \leq KL(\mathbb{P}^0 \,\|\, \mathbb{P}^\star) < \infty$. This proves the first claim. □

Proof of Proposition 3.8, second claim. We proceed by induction, adapting the argument of (De Bortoli et al., 2021, Appendix C.8).

At initialization, we choose $\mathbb{P}^0 \in \mathcal{R}^{\otimes}(\mathbb{Q})$ with $\mathbb{P}^0_{t_i} = \mu_{t_i}$ for all i. We also define $M^0 = \mathrm{proj}_{\mathcal{M}}(\mathbb{P}^0)$.

By construction (Algorithm 1), the IMFF sequence alternates:

$$\mathbb{P}^{2n+1} = \mathrm{proj}_{\mathcal{M}}(\mathbb{P}^{2n}), \qquad \mathbb{P}^{2n+2} = \mathrm{proj}_{\mathcal{R}^{\otimes}(\mathbb{Q})}(\mathbb{P}^{2n+1}).$$

Suppose now that $\mathbb{P}^{2n}$ satisfies the claim. By definition, $\mathbb{P}^{2n+1} \in \mathcal{M}$ and $\mathbb{P}^{2n+2} \in \mathcal{R}^{\otimes}(\mathbb{Q})$. From Lemma 3.1, we then have

$$KL(\mathbb{P}^{2n+1} \,\|\, P^\star) \leq KL(\mathbb{P}^{2n} \,\|\, P^\star), \qquad KL(\mathbb{P}^{2n+2} \,\|\, P^\star) \leq KL(\mathbb{P}^{2n+1} \,\|\, P^\star).$$

Hence, $(KL(\mathbb{P}^n \,\|\, P^\star))_{n\in\mathbb{N}}$ is a nonincreasing sequence bounded below by 0, and is therefore convergent. Moreover, by induction we have $\mathbb{P}^n \in \mathcal{M} \cap \mathcal{R}^{\otimes}(\mathbb{Q})$ for all n, so the limit must coincide with $P^\star$, the unique measure in this intersection with prescribed marginals.

Finally, note that in Algorithm 1 the forward and backward Markovian steps are time-reversals of each other (they follow the same law under the change of variable $t \mapsto T - t$). Therefore, alternating a backward step with a forward reciprocal projection, or a forward step with a backward reciprocal projection, is equivalent from the viewpoint of convergence analysis. All the arguments above apply symmetrically in both directions, and the resulting sequence $(\mathbb{P}^n)_{n\in\mathbb{N}}$ still converges.

We conclude that

$$\lim_{n\to\infty} KL(\mathbb{P}^n \,\|\, P^\star) = 0,$$

and $P^\star$ is indeed the weak solution produced by the IMFF algorithm, proving the second claim. □

A.6.11. THEOREM 3.2

Proof of Theorem 3.2. As a reminder, the argument is the same as in (Shi et al., 2023) and (De Bortoli et al., 2021), but adapted to the multi-marginal setting.

By Proposition 3.8, the sequence $(\mathbb{P}^n)_{n\in\mathbb{N}}$ is bounded in KL divergence with respect to $\mathbb{P}^\star$, hence relatively compact under weak convergence. Thus, it admits a subsequence $(\mathbb{P}^{n_j})_j$ converging weakly to some limit $\mathbb{P}^\infty$. By construction, $\mathbb{P}^\infty \in \mathcal{M} \cap \mathcal{R}^\otimes(Q)$ and matches the marginals $(\mu_{t_i})_{i=0}^K$, so by uniqueness of the weak MMSB solution we must have $\mathbb{P}^\infty = \mathbb{P}^\star$.

By lower semicontinuity of KL, this implies

$$\lim_{n\to\infty} KL(\mathbb{P}^n \,\|\, \mathbb{P}^\star) = 0.$$

Finally, the inequality

$$KL(\mathbb{P}^{\mathrm{MMSB}} \,\|\, Q) \;\leq\; KL(\mathbb{P}^\star \,\|\, Q) \;\leq\; KL(\mathbb{P}^{\mathrm{pair}} \,\|\, Q)$$

is justified because $\mathbb{P}^{\mathrm{MMSB}}$ is the global minimizer (hence gives the smallest KL), while $\mathbb{P}^\star$ is the best Markovian candidate in $\mathcal{M} \cap \mathcal{R}^\otimes(Q)$, and therefore lies below the pairwise construction obtained by gluing local bridges. □