Multi-marginal temporal Schrödinger Bridge Matching for video generation from unpaired data

Anonymous authorsPaper under double-blind review

000

001

003

004

006

008 009 010

011 012 013

014

015

016

017

018

019

021

025

026

027 028 029

031

033

034

037

040

041

042

043

044

046

047

048

049

050

051

052

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) for video generation 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. Our work establishes multi-marginal Schrödinger bridges as a practical and principled approach for recovering hidden dynamics from static data.

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 modelization 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). 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) uses iterative proportional fitting (IPF) (Kullback, 1968), the general continuous analogue of the well known Sinkhorn algorithm (Cuturi, 2013). Subsequent works have explored alternative training schemes based on likelihood bounds

¹code: github.com/ICLRMMtDSBM/MMDSBM_ILCR | website: mmdsbm.notion.site

(Chen et al., 2023b) or on the dual algorithm of IPF: Iterative Markovian Fitting (IMF) (Shi et al., 2023).

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 This paper makes the following contributions: 1) We define the multi-marginal temporal Schrödinger Bridge problem and demonstrate its fundamental properties 2) We introduce a novel factorized extension of the IMF algorithm presented in Shi et al. (2023) to multiple iterative marginals in an efficient and principled way. 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, achieving state-of-the-art results on the widely reported Embryoid Body (EB) dataset (Moon et al., 2019). 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} . For \mathbb{Q} induced by $(\sqrt{\varepsilon}B_t)_{t\in[0,1]}, \varepsilon>0$, where $(B_t)_{t\geq0}$ is a d-dimensional Brownian motion, the reciprocal class of \mathbb{Q} is $\mathcal{R}(\mathbb{Q})$ (see Def 3.2) 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,

$$\mathbb{P}^{\star} = \underset{\mathbb{P} \in \mathcal{P}(C)}{\operatorname{arg \, min}} \, \mathrm{KL}(\mathbb{P} \, \| \, \mathbb{Q}) \, \text{ s.t. } \mathbb{P}_{0} = \mu_{0}, \, \mathbb{P}_{T} = \mu_{T}, \tag{1}$$

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

$$\Pi^* = \underset{\Pi}{\operatorname{arg\,min}} \operatorname{KL}(\Pi \parallel \mathbb{Q}_{0,T}) \text{ s.t. } \Pi_0 = \mu_0, \ \Pi_T = \mu_T.$$
(2)

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:

$$\mathbb{P}^{2n+1} = \operatorname{proj}_{\mathcal{M}}(\mathbb{P}^{2n}), \qquad \mathbb{P}^{2n+2} = \operatorname{proj}_{\mathcal{R}(\mathbb{O})}(\mathbb{P}^{2n+1}). \tag{3}$$

These projections admit KL variational characterisations (App. B), and the iteration converges to \mathbb{P}^* .

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

3 MULTI-MARGINAL TEMPORAL SCHRÖDINGER BRIDGE MATCHING

All proofs can be found in the Appendix B.5.

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

$$\mathbb{P}^{\star} = \underset{\mathbb{P} \in \mathcal{P}(C)}{\operatorname{argmin}} \operatorname{KL}(\mathbb{P} \parallel \mathbb{Q}) \quad \text{subject to} \quad X_{t_k} \sim \mu_k, \quad k = 0, \dots, N$$
 (4)

Connection to Multi-marginal Optimal Transport If \mathbb{Q} is associated with a Brownian motion, the induced reference coupling $\mathbb{Q}_{t_0,\dots,t_N}$ 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 therefore be rewritten as

$$\Pi^{\star} = \arg\min_{\Pi \in \mathcal{P}((\mathbb{R}^d)^N)} \left\{ \mathbb{E}_{X \sim \Pi} \left[\sum_{i=0}^{N-1} \frac{1}{t_{i+1} - t_i} \|X_{t_{i+1}} - X_{t_i}\|^2 \right] - 2\sigma^2 T \mathcal{H}(\Pi) : \Pi_i = \mu_{t_i}, \ \forall i \right\}$$

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

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 $\mathbb{Q}_{t_0,...,t_K}$ be the joint law of \mathbb{Q} at $0 = t_0 < \cdots < t_K = T$. The static problem is

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

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^* .

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^* *is determined by the static one* π^* :

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

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^* *is Markov.*

²This formulation is particularly interesting in practice: in many applications we only observe static snapshots μ_{t_i} without any coupling between them, and the multi-marginal Schrödinger Bridge not only provides a consistent path modeling but also induces the most natural coupling between the observed distributions.

These properties ensure that we can restrict our search to Markovian (and therefore reciprocal B.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*,

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

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.

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

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

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

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

i.e. we keep the marginals $P_{t_0,...,t_K}$ 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, Π^* admits the variational characterization

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

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 sub-interval 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 $M^* = \operatorname{proj}_{\mathcal{M}}(\Pi)$, is the unique diffusion process

$$dX_t^{\star} = \left\{ f_t(X_t^{\star}) + v_t^{\star}(X_t^{\star}) \right\} dt + \sigma_t dB_t,$$

with effective drift

$$v_t^{\star}(x) = \sigma_t^2 \mathbb{E}_{\Pi_{t_{i+1}\mid t}} \Big[\nabla \log \mathbb{Q}_t^{\mid t_i, t_{i+1}\mid} (X_{t_{i+1}}\mid X_t) \mid X_t = x \Big] \stackrel{\textit{Brownian}}{=} \frac{\mathbb{E}_{\Pi_{t_{i+1}\mid t}} [X_{t_{i+1}}\mid X_t = x] - x}{t_{i+1} - t}$$

By the Markovian projection theorem of Gyöngy (1986), and as further developed in Peluchetti (2023); De Bortoli et al. (2021), the process M^* 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 $M^* = \operatorname{proj}_{\mathcal{M}}(\Pi)$ be the Markovian projection of Π as in Definition 3.3. Then we work under the assumption that the variational characterization still holds in the factorized setting ((Shi et al., 2023; Gyöngy, 1986; Peluchetti, 2023)), i.e.

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

and

$$KL(\Pi \parallel M^{\star}) = \frac{1}{2} \int_{0}^{T} \mathbb{E}_{\Pi_{t_{i},t}} \left[\frac{1}{\sigma_{t}^{2}} \left\| \sigma_{t}^{2} \mathbb{E}_{\Pi_{t_{i+1}\mid t}} \left[\nabla \log \mathbb{Q}_{t}^{\mid t_{i},t_{i+1}} \left(X_{t_{i+1}} \mid X_{t} \right) \mid X_{t} \right] - v_{t}^{\star}(X_{t}) \right\|^{2} \right] dt$$

In addition, for any $t \in [0, T]$, the time marginal of M^* coincides with that of Π : $M_t^* = \Pi_t$. In particular, $M_{t_i}^* = \Pi_{t_i}$ 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 $M^* = \operatorname{proj}_{\mathcal{M}}(\Pi)$ is associated with the forward SDE

$$dX_{t} = \left\{ f_{t}(X_{t}) + \sigma_{t}^{2} \mathbb{E}_{\Pi_{t_{i+1}}|t} \left[\nabla \log \mathbb{Q}_{t}^{[t_{i}, t_{i+1}]} (X_{t_{i+1}} \mid X_{t}) \mid X_{t} \right] \right\} dt + \sigma_{t} dB_{t}, \quad X_{t_{i}} \sim \mu_{t_{i}} \quad (5)$$

and with the backward SDE

$$dY_{t} = \left\{ -f_{t_{i+1}-t}(Y_{t}) + \sigma_{t_{i+1}-t}^{2} \mathbb{E}_{\Pi_{t_{i}}|t} \left[\nabla \log \mathbb{Q}_{t}^{[t_{i},t_{i+1}]}(Y_{t_{i}} \mid Y_{t}) \mid Y_{t} \right] \right\} dt + \sigma_{t_{i+1}-t} dB_{t}, \quad Y_{t_{i+1}} \sim \mu_{t_{i+1}}$$

$$\tag{6}$$

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 leverage 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^* 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

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

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

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

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

$$KL(M|\Pi) = KL(M|\operatorname{proj}_{\mathcal{R}^{\otimes}(\mathbb{O})}(M)) + KL(\operatorname{proj}_{\mathcal{R}^{\otimes}(\mathbb{O})}(M)|\Pi)$$

Proposition 3.8. *Under mild assumptions, we have*

270

271

272 273 274

275

276

277

278

279

280

281 282

283

284 285

286

287

288

289 290

291 292

293

295 296

297 298

299 300 301

302 303

304

305

306 307

308

309

310

311

312

313

314

315

316 317

318 319

320

321

322

323

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

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}^* , and we have

$$\lim_{n \to +\infty} KL(\mathbb{P}^n \parallel \mathbb{P}^*) = 0$$

Moreover, denoting by \mathbb{P}^{MMSB} the true multi-marginal Schrödinger bridge and by \mathbb{P}^{pair} the classical Schrödinger bridge constructed pairwise, the limiting solution satisfies the inequality

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

where \mathbb{Q} is the chosen reference process. Thus, \mathbb{P}^* can be interpreted as an approximate multimarginal Schrödinger Bridge, lying between the optimal multi-marginal solution and the pairwise construction.

3.2.3 Theoretical algorithm

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

$$\theta^{\star} = \underset{\theta}{\operatorname{arg\,min}} \ \mathbb{E}_{\operatorname{batch}} \Big[\Big\| v_{\theta}(X_t, t) - \sigma_t^2 \, \mathbb{E} \big[\nabla \log \mathbb{Q}_t^{[t_{i(t)}, t_{i(t)+1}]} (X_{t_{i(t)+1}} \mid X_t) \, \big| \, X_t \big] \Big\|^2 \Big]$$
(8)

for the *forward* drift v_{θ} , and

$$\phi^{\star} = \underset{\phi}{\operatorname{arg\,min}} \ \mathbb{E}_{\operatorname{batch}} \Big[\big\| v_{\phi}(Y_t, t) - \sigma_{t_{i(t)+1}-t}^2 \, \mathbb{E} \big[\nabla \log \mathbb{Q}_t^{[t_{i(t)}, t_{i(t)+1}]} (Y_{t_{i(t)}} \mid Y_t) \, \big| \, Y_t \big] \big\|^2 \Big] \quad (9)$$

for the *backward* drift v_{ϕ} . We summarize in Algorithm 1 our method.

Algorithm 1 Iterative Markovian Factorized Fitting (IMFF)

- 1: **Input:** time grid $0 = t_0 < \cdots < t_K = T$, marginals $(\mu_{t_i})_{i=0}^K$, reference process \mathbb{Q} , number of iterations
- 2: Init: choose $\mathbb{P}^0 \in \mathcal{R}^{\otimes}(Q)$ with $\mathbb{P}^0_{t_i} = \mu_{t_i}$ for all i 3: 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} fixed from $(\mu_{t_{i+1}})$.
- Forward reciprocal projection: $\mathbb{P}^{2n+1} \leftarrow \operatorname{proj}_{\mathcal{R}^{\otimes}(Q)}(\mathbb{P}^{2n+1})$ (cf. Def. 3.2), filling bridges with Qusing 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 fixed from (μ_{t_i}) .
- **Backward reciprocal projection:** $\mathbb{P}^{2n+2} \leftarrow \operatorname{proj}_{\mathcal{R} \otimes (Q)}(\mathbb{P}^{2n+2})$ (cf. Def. 3.2), filling bridges with Q using t_{i+1} from \mathbb{P}^{2n+2} and t_i from the dataset.
- 8: end for
- 9: **Output:** learned drifts (v_{ϕ}, v_{θ})

A practical implementation of IMFF is provided in Appendix B.2.

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 $(\sigma_t B_t)_{0 \le t \le T}$ as the reference measure \mathbb{Q} and T = N - 1 where N is the number of marginals.³

4.1 MMTSBM RECOVERS THE KNOWN EXACT OT BETWEEN GAUSSIAN MIXTURES

In this 2D experiment, we used N=3 mixtures of two standard Gaussian as marginals. In this configuration the optimal transport between each pair of marginals is known exactly: it is a pure translation of each Gaussian components inside the mixtures [CITE]. After only the warm-up phase (akin to flow matching [CITE], as said before), 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 of Figure 1 that the learned trajectories do not intersect each other anymore and that MMtSBM yields the exact static optimal transport map: pure translations between Gaussians.

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.

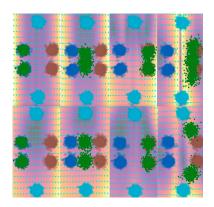


Figure 1: Top row: epoch 0, (only noisy flow matching). Bottom row: epoch 5, (after MMtSBM training). From left to right: snapshots at times (0,0.5,1,1.3). True marginal times $(t_0,t_1,t_2)=(0,1,2)$. The order of the 3 true marginals is: $t_0=\operatorname{dark}$ blue; $t_1=\operatorname{red}$; $t_2=\operatorname{light}$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

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. (2024) and Shi et al. (2023) to our temporal multi-marginal setting in Table 1. Choosing N=4, we considered ($\mathcal{N}\to \text{Moons}\to \mathcal{N}\to \text{Moons}$), and ($\mathcal{N}\to \text{8Gaussians}\to \mathcal{N}\to \text{8Gaussians}$). 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.

_ a	341	D 1 E
Setting	\mathcal{W}_2	Path Energy
Moons (single bridge)	0.144 ± 0.024	1.580 ± 0.036
Moons $\times 3$	_	4.740
Moons (MMtSBM)	$0.148{\scriptstyle\pm0.041}$	$5.350{\scriptstyle\pm0.085}$
8 Gaussians (single bridge)	0.338 ± 0.091	14.810 ± 0.255
8 Gaussians ×3	_	44.430
8 Gaussians (MMtSBM)	$0.352{\scriptstyle\pm0.084}$	$46.920{\scriptstyle\pm0.285}$

Table 1: Comparison of static coupling (W₂ column) and energy minimization (Path Energy column). 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.

We observe that despite a much more complex *time-varying* true transport map to be learned, MMtSBM achieves almost as low W_2 distances than 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.

³More examples for each experiment can be found at mmdsbm.notion.site.

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, extended 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.

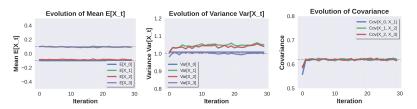


Figure 2: Evolution of mean, variance, and covariance in the multi-marginal 50d Gaussian transport. Dash lines are the theoretical true values.

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 BENCHMARK

We next evaluate our method on the the TrajectoryNet benchmark (Tong et al., 2020) which uses real single-cell RNA-seq embryoid body differentiation data from Moon et al. (2019). We project RNA counts to their first 100 principal components for each of the N=5 marginals. We report in Table 2 the Maximum Mean Discrepency (MMD) and Sliced Wasserstein Distance (SWD).

	DMSB		MMtSBM (ours)	
Time	MMD ↓	SWD \downarrow	$MMD \downarrow$	SWD \downarrow
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	0.030	0.160	0.019	0.130

Algorithm	MMD↓	$\text{SWD}\downarrow$
NLSB	0.66	0.54
MIOFlow	0.23	0.35
DMSB	0.03	0.20
MMtSBM (ours)	0.02	0.13

Table 2: Results on test set of embryoid body RNA-seq data (d = 100). Left table: per-marginal metrics. Right table: average over all target marginals; results from Chen et al. (2023a).

Our method consistently outperforms baselines on all marginals, reducing the average MMD from 0.03 to **0.02** and the SWD from 0.20 to **0.13**. This demonstrates that enforcing all marginal constraints simultaneously as we do yields sharper and more consistent interpolations across developmental stages, setting a new state of the art on this benchmark.

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 EXPERIMENT

We conducted experiments on the MNIST dataset, 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.



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

4.5.2 BIOTINE EXPERIMENT

The (in-house) "biotine" dataset consists of 3-channel fluorescence images (GFP, membrane, nucleus) of A549 lung epithelial cells cultured in 384-well plates, treated with biotin, and imaged at 19 discrete time steps. Figure 4 shows the unpaired dynamic we have at hand.

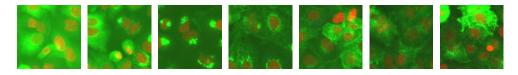


Figure 4: Ground truth biotine examples at training marginal times t = 0, 1, 2, 3, 4, 5, 6, from left to right.

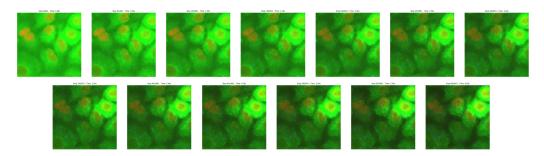


Figure 5: Video generated by MMtSBM on biotine, forward direction. To be read in reading order: top left \rightarrow top right, then bottom left \rightarrow bottom right. Generations at times t=0,0.5,1,1.5,...,5.5,6. Top-left starting image is from the test set.

In Figure 5 we can see that MMtSBM successfully recovers a plausible generative coupling between exsting timestamps, preserving both spatial and temporal coherence and yielding smooth, low-noise dynamics that extends the starting test image into a continuous video that respects the global phenotype seen in ground truth data (see Figure 4). To the best of our knowledge, this is the first demonstration of a multi-marginal Schrödinger Bridge performing video generation from purely unpaired data. 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 its theoretical soundness as well as it efficiency. We achieve state-of-the-art results in a common literature benchmark and present for the first time an algorithm that produces temporarily coherent videos from purely unpaired data, hoping to lead to many future applications in the scientific domain.

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 reference processes than the Brownian motion. We also intend to investigate learning the transport map in a VAE latent space, as well as definitively proving our theoretical conjectures. We would also like to explore using the single network theory developed in Bortoli et al. (2024) for efficiency gains.

REFERENCES

486

487

488

489 490

491

492

493

494

495 496

497

498 499

500

501

502

504 505

506

507

508 509

510 511

512

513

514

515516

517

518

519

520

521 522

523 524

525

526

527

528

529

530

531

532

534 535

536

537

538

- Valentin De Bortoli, Iryna Korshunova, Andriy Mnih, and Arnaud Doucet. Schrödinger bridge flow for unpaired data translation, 2024. URL https://arxiv.org/abs/2409.09347.
- Charlotte Bunne, Stefan G. Stark, Gabriele Gut, Eran H.K. Stelzer, Gunnar Rätsch, and Marco Cuturi. Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*, 20(12):1820–1829, 2023.
- Tianrong Chen, Guan-Horng Liu, Molei Tao, and Evangelos A. Theodorou. Deep momentum multi-marginal schrödinger bridge, 2023a. URL https://arxiv.org/abs/2303.01751.
- Tianrong Chen, Guan-Horng Liu, and Evangelos A. Theodorou. Likelihood training of schrödinger bridge using forward-backward sdes theory, 2023b. URL https://arxiv.org/abs/2110.11291.
- Imre Csiszár. I-divergence geometry of probability distributions and imization problems. TheAnnals Probability, 3(1):146-158, 1975. of 10.1214/aop/1176996454. URL https://projecteuclid. org/journals/annals-of-probability/volume-3/issue-1/ I-Divergence-Geometry-of-Probability-Distributions-and-Minimization-Problems/ 10.1214/aop/1176996454.full.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2013.
- Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling, 2021. URL https://arxiv.org/abs/2106.01357.
- István Gyöngy. Mimicking the one-dimensional marginal distributions of processes having an itô differential. *Probability Theory and Related Fields*, 71:501–516, 1986.
- Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth schrödinger bridges, 2025. URL https://arxiv.org/abs/2503.00530.
- S. Kullback. 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.
- Christian Léonard. From the schrödinger problem to the monge–kantorovich problem. *Journal of Functional Analysis*, 262(4):1879–1920, 2012.
- Christian Léonard. 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.
- Kevin R. Moon, David van Dijk, Zheng Wang, Scott Gigante, Daniel B. Burkhardt, William S. Chen, Kristina Yim, Antonia van den Elzen, Matthew J. Hirn, Ronald R. Coifman, Natalia B. Ivanova, Guy Wolf, and Smita Krishnaswamy. Visualizing structure and transitions 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.
- Stefano Peluchetti. 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.
- Erwin Schrödinger. Über die umkehrung der naturgesetze. Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse, pp. 144–153, 1931.
- Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching, 2023. URL https://arxiv.org/abs/2303.16852.
 - Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021.

Panagiotis Theodoropoulos, Augustinos D. Saravanos, Evangelos A. Theodorou, and Guan-Horng Liu. Momentum multi-marginal schrödinger bridge matching, 2025. URL https://arxiv. org/abs/2506.10168. Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics, 2020. URL https: //arxiv.org/abs/2002.04461. Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport, 2024. URL https://arxiv.org/abs/2302.00482. Gefei Wang, Yuling Jiao, Qian Xu, Yang Wang, and Can Yang. Deep generative learning via schrödinger bridge, 2021. URL https://arxiv.org/abs/2106.10410.

A APPENDIX

B ADDITIONAL BACKGROUND

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

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

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

$$dX_t = \left[\frac{\mathbb{E}_{\Pi}[X_T \mid X_t] - X_t}{T - t} \right] dt + \sigma dB_t.$$

Variational formulations. Both projections solve KL problems:

$$\operatorname{proj}_{\mathcal{R}(\mathbb{Q})}(\mathbb{P}) = \operatorname*{arg\,min}_{\Pi \in \mathcal{R}(\mathbb{Q})} \operatorname{KL}(\mathbb{P} \parallel \Pi), \qquad \operatorname{proj}_{\mathcal{M}}(\Pi) = \operatorname*{arg\,min}_{M \in \mathcal{M}} \operatorname{KL}(\Pi \parallel M).$$

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

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

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

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

Unlike IMF, this requires caching full trajectories.

B.1 OTHER PROPERTIES ON IMFF OR MMSB

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

Proposition B.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)):

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

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:

$$v_t(x) = \frac{1}{2} \left(v_t^{\text{fwd}}(x) + v_t^{\text{bwd}}(x) \right)$$
 (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.

B.2 Concrete Algorithms

648

649 650

651

652

653

654

655

656

657

658

659

660

661

662

663

664 665 666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690 691

692 693

694

696

697 698

699

700

701

13: **end for**

Algorithm 2 Warmup Phase (our algorithm) 1: 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} 2: Output: Warmup parameters θ, ϕ 3: Define bridges $\mathcal{B} = \{(t_i, t_{i+1})\}, b \leftarrow B/|\mathcal{B}|$ 4: for direction \in {forward, backward} do 5: for $n \in [0, N_{\text{warmup}}]$ do 6: for all $(t_i, t_{i+1}) \in \mathcal{B}$ in parallel do

```
7: Sample (X_{t_i}, X_{t_{i+1}}) \sim (\pi_{t_i} \otimes \pi_{t_{i+1}})^{\otimes b}, t_{(i)} \sim \operatorname{Unif}[t_i, t_{i+1}]^{\otimes b}
8: end for
9: Aggregate X_{\operatorname{init}}, X_{\operatorname{final}}, t; Sample Z \sim \mathcal{N}(0, I)^{\otimes B}
10: X_t \leftarrow \operatorname{Interp}_t(X_{\operatorname{init}}, X_{\operatorname{final}}, Z) \triangleright cf. equation 11
11: Update \theta if forward with \ell^{\operatorname{fwd}} equation 13, else \phi with \ell^{\operatorname{bwd}} equation 14
12: end for
```

Algorithm 3 MM-IMF Phase (our algorithm)

```
1: Input: Subdivision \{0 = t_0 < t_1 < \dots < t_n = T\}, datasets \{\pi_{t_i}\}, networks v_{\theta}, v_{\phi}, warmup
      params \theta, \phi, batch size B, finetune steps N_{\text{finetune}}, inner steps N_{\text{inner}}
 2: Output: Finetuned parameters \theta, \phi
 3: Define bridges \mathcal{B} = \{(t_i, t_{i+1})\}, b \leftarrow B/|\mathcal{B}|
 4: for N \in [0, N_{\text{finetune}}] do
 5:
            for all (t_i, t_{i+1}) \in \mathcal{B} in parallel do
 6:
                   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}
 7:
 8:
            end for
 9:
            Aggregate X_{\text{init}}, X_{\text{final}}, t
            \textbf{for } direction \in \{backward, forward\} \ \textbf{do}
10:
                   for n \in [0, N_{\text{inner}}] do
11:
                         if direction = forward then
12:
13:
                               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
14:
                               Update \theta with \ell^{\text{fwd}} equation 13
15:
16:
                               X_{\text{final}} \leftarrow \text{SDE}(X_{\text{init}}, v_{\theta})
                                                                                                                                       ⊳ cf. equation 12
17:
                               X_t \leftarrow \operatorname{Interp}_t(X_{\operatorname{init}}, \hat{X}_{\operatorname{final}}, Z)
Update \phi with \ell^{\operatorname{bwd}} equation 14
18:
                                                                                                                                       ⊳ cf. equation 11
19:
                         end if
20:
21:
                   end for
             end for
22:
23: end for
```

B.3 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.

B.3.1 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

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

B.3.2 MASKING AND TIME DISCRETIZATION

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

$$N_i = \left| N_{\mathrm{total}} rac{t_{i+1} - t_i}{T}
ight|, \quad dt_i = \pm rac{\Delta au}{t_{i+1} - t_i}, \; \Delta au = rac{T_{\mathrm{max}} - T_{\mathrm{min}}}{N_{\mathrm{total}}}.$$

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.

B.3.3 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

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

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

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

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

Forward/backward losses enforce vectorized drift consistency:

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

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

B.3.4 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.

B.4 EXPERIMENTS DETAILS

The Adam optimizer is used with a learning rate of $2*10^{-4}$, and SiLU activations are applied on each layers unless stated otherwise. All experiments are executed on computing clusters that utilize GPU resources.

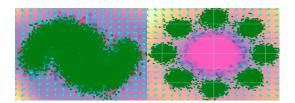


Figure 6: Third marginal fitting for the moons and 8-Gaussian trajectories. Blue vectors indicate the drift direction, with gradient intensity showing vector field strength; green points denote moving samples, and pink highlights the Gaussian fitted along the trajectory.

B.4.1 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.

B.4.2 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.

B.4.3 100D TRANSCRIPTOMIC EXPERIMENTS

The dataset comprises 5 timepoints (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).

All datasets were standardized (zero mean and unit variance), and from each dataset 1000 samples were withheld 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.

We trained a network of about 3×10^5 parameters for 20 outer iterations with 20,000 inner iterations.

B.4.4 MNIST EXPERIMENTS

We experimented 2 approaches for the MNIST dataset: a MLP with flattened image vectors of dimension ($28 \times 28 = 784$, and a UNet with image-shape data of shape (28, 28).

B.4.5 BIOTINE EXPERIMENTS

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

B.5 PROOFS

Proof of variationnal 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

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

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

Restricting to Π such that $\Pi_{t_0,...,t_K} = P_{t_0,...,t_K}$ 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

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

which is exactly the definition of the factorized reciprocal projection.

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

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

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

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

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

$$v_t^{\star}(x) = \sigma_t^2 \mathbb{E}\left[\frac{X_{t_{i+1}} - x}{\sigma^2(t_{i+1} - t)} \mid X_t = x\right].$$

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.

Proof of Proposition 3.1. The feasible set

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

is convex and closed under the weak topology. 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 $Q_{[t_i,t_{i+1}]}^{x_i,x_{i+1}}$ denote the Brownian bridge of Q conditioned on $X_{t_i}=x_i$ and $X_{t_{i+1}}=x_{i+1}$. Then the measure

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

belongs to A. Hence the admissible set is non-empty.

Therefore, (MMSB) admits a unique solution P^* .

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

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

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^* is uniquely obtained from the optimal static solution π^* by gluing the conditional bridges of Q, which establishes the equivalence.

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

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

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,\ldots,t_n , hence Markov on [0,1].

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 $Q_{t_0,...,t_K}$ 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

$$Q_{t_0,...,t_K}(dx_0,...,dx_K) \ge \exp\left(-\sum_{i=0}^K A(x_i)\right) m(dx_0) \cdots m(dx_K);$$

(iii) there exists B such that

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

(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 $H(\pi_{t_i}\mid m)<\infty$,

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

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 π^* of the static problem is absolutely continuous with respect to Q_{t_0,\dots,t_K} with density

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

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

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

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

For the factorized reciprocal part:

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

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

We have the Radon-Nikodym factorization

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

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

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

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

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

Thus

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

П

which concludes the proof.

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

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

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.

Proof of Proposition 3.8. 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

$$KL(\mathbb{P}^0 \parallel \mathbb{P}^*) = KL(\mathbb{P}^0 \parallel \mathbb{P}^1) + KL(\mathbb{P}^1 \parallel \mathbb{P}^2) + \dots + KL(\mathbb{P}^N \parallel \mathbb{P}^*).$$

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 $KL(\mathbb{P}^n \parallel \mathbb{P}^*) \leq KL(\mathbb{P}^0 \parallel \mathbb{P}^*) < \infty$. This proves the claim.

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}^* , hence relatively compact under weak convergence. Thus, it admits a subsequence $(\mathbb{P}^{n_j})_j$ converging weakly to some limit \mathbb{P}^{∞} . 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 $\mathbb{P}^{\infty} = \mathbb{P}^*$.

By lower semicontinuity of KL, this implies

$$\lim_{n\to\infty} KL(\mathbb{P}^n \parallel \mathbb{P}^*) = 0.$$

Finally, the inequality

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

is justified because $\mathbb{P}^{\mathrm{MMSB}}$ is the global minimizer (hence gives the smallest KL), while \mathbb{P}^* 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.

Proof of Proposition 3.8. 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 $M^0 = \operatorname{proj}_{\mathcal{M}}(\mathbb{P}^0)$.

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

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

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

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

Hence, $(KL(\mathbb{P}^n \parallel P^*))_{n \in \mathbb{N}}$ 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^* , 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

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

and P^* is indeed the weak solution produced by the IMFF algorithm.