Schrödinger Bridge Matching

Grigoriy Ksenofontov, Mikhail Persiianov

Moscow, 2026

Overview

Recap: Dynamic Schrödinger Bridge Problem

Bridge Matching

Diffusion Schrödinger Bridge Matching

Recap: Dynamic Schrödinger Bridge Problem

Unsupervised Domain Translation: Problem Setup

Unsupervised setting.

We observe two datasets

$$\{x^{(i)}\}_{i=1}^{M} \sim p_0, \qquad \{y^{(i)}\}_{i=1}^{M} \sim p_1,$$

with **no paired correspondences**.

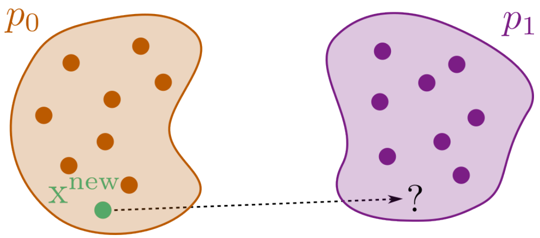

Goal: learn a mapping

$$T : \mathcal{X} \to \mathcal{Y}$$

such that

$$T_\sharp p_0 \approx p_1.$$

Key difficulty: the pushforward constraint does not uniquely determine T.

Why Static Transport Models Are Not Enough

In practice, the discussed optimal transport corresponds to learning a **one-step generator**:

$$y = T(x, z), \qquad z \sim p(z).$$

Complex distributions are difficult to generate in a single step.

Thus, modern generative models are inherently **dynamic**, i.e.:

$$x \longrightarrow x_{t_1} \longrightarrow \cdots \longrightarrow x_{t_N} \longrightarrow y$$

Examples of state-of-the-art approaches:

- Diffusion models (DDPM, DDIM, NCSN, etc.);
- Flow matching / continuous normalizing flows (FM, Rectified Flows, etc.).

QUESTION: Can we bring dynamics into Optimal Transport?

From Random Variables to Random Paths

Static OT: random variable x on state space $\mathbb{R}^d$

$$\Downarrow$$

Dynamic OT: random path[1] $x(t) \stackrel{\text{def}}{=} x_t$ on path space $C([0,1];\mathbb{R}^d)$

Evaluating the path at times $t \in [0,1]$ yields random variables $\{x_t\}_{t\in[0,1]}$, which together form a stochastic process, or formally:

$$\left(C([0,T];\mathbb{R}^d), \mathcal{B}(C([0,T];\mathbb{R}^d)), \mathbb{P}\right),$$

where $\mathbb{P}$ denotes path distribution.

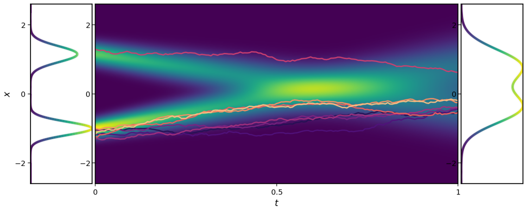

The target probabilistic object is now the path distribution.

[1]We abuse notation by using x_t both for the random path and for its realization, denoted as trajectories.

Path Distributions Induced by SDEs

Given the scaled Wiener process, we define a new path distribution $\mathbb{P}$ by combining deterministic dynamics with Wiener noise:

$$T : dx_t = \underbrace{v(x_t, t)\, dt}_{\substack{\text{drift,} \\ \text{defines deterministic motion}}} + \underbrace{\sqrt{\epsilon}\, dW_t,}_{\substack{\text{scaled Wiener noise,} \\ \text{defines stochastic motion}}} \qquad x_0 \sim p_0,$$

Over a small time step Δt, the SDE can be interpreted as

$$x_{t+\Delta t} \approx x_t + v(x_t, t)\Delta t + \sqrt{\epsilon \Delta t}\, z_t, \qquad z_t \sim \mathcal{N}(0, \mathbb{I}).$$

$$\boxed{\text{The trajectories generated by this SDE induce a path distribution } \mathbb{P} \text{ on } C([0,1]; \mathbb{R}^d).}$$

One useful notation is the <u>bridge distribution</u> $\mathbb{P}_{|x,y} \overset{\text{def}}{=} \mathbb{P}(\,\cdot | x_0 = x,\ x_1 = y)$, that is, the path distribution of the SDE conditioned on fixed endpoints.

Dynamic Schrödinger Bridge Formulation

Suppose we are given marginal distributions p_0 and p_1 and a reference scaled Wiener process $\mathbb{W}^\epsilon$. Then, the dynamic Schrödinger Bridge (SB) problem[2] is defined as

$$\mathbb{P}^* = \underset{\mathbb{P}\in\Pi(p_0,p_1)}{\arg\min}\ \mathrm{KL}\left(\mathbb{P}\|\mathbb{W}^\epsilon\right), \qquad \text{(DynSB)}$$

where $\Pi(p_0, p_1)$ is the set of all path distributions with initial and terminal marginals p_0 and p_1, respectively.

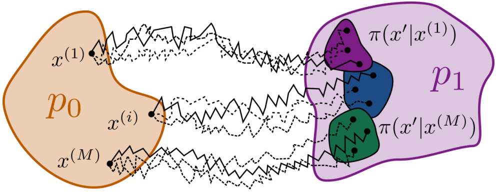

Intuitively, the problem seeks the path distribution $\mathbb{P}$ that is closest to the Wiener path distribution $\mathbb{W}^\epsilon$ in KL divergence, while matching the prescribed marginals p_0 and p_1 at times $t = 0$ and $t = 1$, respectively.

[2] Erwin Schrödinger (1931). **Über die umkehrung der naturgesetze.** Verlag der Akademie der Wissenschaften in Kommission bei Walter De Gruyter u . . .

Doob's h-transform of Path Distributions

Let $\mathbb{P}^{\mathrm{ref}}$ be a Markov reference path distribution (e.g. the Wiener path distribution $\mathbb{W}^{\epsilon}$), and let $h_t(x) \in (0, \infty)$ be a positive space-time harmonic function, such that

$$h_t(x) \stackrel{\text{def}}{=} \mathbb{E}_{\mathbb{P}^{\mathrm{ref}}}[h_1(x_1)|x_t = x] \, .$$

The Doob's h-transform[3] defines a new path distribution $\mathbb{P}^h$ by reweighting $\mathbb{P}^{\mathrm{ref}}$ with h_t:

$$\frac{d\mathbb{P}^h}{d\mathbb{P}^{\mathrm{ref}}} = \frac{h_1(x_1)}{h_0(x_0)} = \frac{h_1(x_1)}{\underbrace{\mathbb{E}_{\mathbb{P}^{\mathrm{ref}}}[h_1(x_1)|x_0 = x]}_{\text{normalization}}} . \tag{1}$$

> If the reference path distribution $\mathbb{P}^{\mathrm{ref}}$ is Markov, then its Doob-h transform $\mathbb{P}^h$ remains Markov.

[3] J. L. Doob (1984). **Classical Potential Theory and Its Probabilistic Counterpart.** Vol. 262. Springer.

Doob-h Transform as a Drift Correction

Consider a reference path distribution $\mathbb{P}^{\mathrm{ref}}$ induced by the SDE

$$dx_t = b(x_t, t)\, dt + \sigma(x_t, t)\, dW_t, \qquad a(x_t, t) \stackrel{\mathsf{def}}{=} \sigma(x, t)\sigma(x, t)^\top.$$

If the path distribution is transformed by a positive space-time harmonic function $h_t(x)$, then the Doob-h transformed diffusion has the same diffusion coefficient and modified drift

$$dx_t = \Big(b(x_t, t) + a(x_t, t)\nabla \log h_t(x_t)\Big)\, dt + \sigma(x_t, t)\, d\widetilde{W}_t.$$

In the Wiener reference ($\mathbb{P}^{\mathrm{ref}} = \mathbb{W}^\epsilon$) case we obtain $a(x, t) = \epsilon\mathbb{I}$ leading to

$$dx_t = \sqrt{\epsilon}\, dW_t, \qquad \text{under } \mathbb{W}^\epsilon,$$

this becomes

$$dx_t = \epsilon\nabla \log h_t(x_t)\, dt + \sqrt{\epsilon}\, d\widetilde{W}_t. \qquad \text{under } \mathbb{P}^h$$

Doob-h transform keeps the noise unchanged and only **adds a drift correction**.

Bridge Matching

Diffusion Bridges

Consider a reference path distribution $\mathbb{P}^{\mathrm{ref}}$ induced by the SDE

$$dx_t = b(x_t, t)\,dt + \sigma(x_t, t)\,dW_t,$$
$$a(x_t, t) \stackrel{\text{def}}{=} \sigma(x_t, t)\sigma(x_t, t)^\top.$$

Its bridge distribution $\mathbb{P}^{\mathrm{ref}}_{|x,y}$ admits a Doob-h transform with known SDE:

$$\frac{d\mathbb{P}^{\mathrm{ref}}_{|x,y}}{d\mathbb{P}^{\mathrm{ref}}_{|x}} = \frac{h_1(x_1)}{\mathbb{E}_{\mathbb{P}^{\mathrm{ref}}}\left[h_1(x_1)|x_0 = x\right]} = \frac{\delta_y(x_1)}{q^{\mathrm{ref}}(y|x)}.$$

Or in the SDE form:

$$\mathbb{P}^{\mathrm{ref}}_{|x,y} : dx_t = \left[b(x_t, t) + a(x_t, t)\nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t)\right] dt + \sigma(x_t, t)dW_t,$$

where $q^{\mathrm{ref}}(y|x_t) = \mathcal{N}(y; x_t, (1 - t)a(x_t, t))$

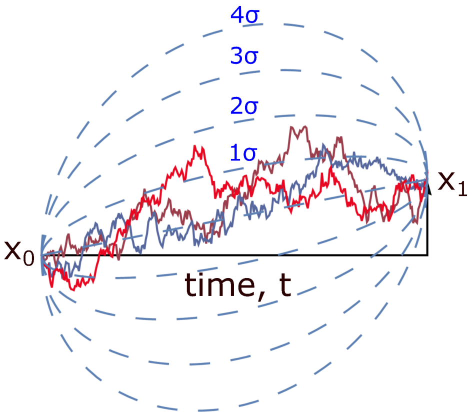

Reciprocal Process (Mixture of Bridges)

Consider a joint distribution $\pi(x_0, x_1) \in \Pi(p_0, p_1)$. Given bridge processes $\mathbb{P}^{\mathrm{ref}}_{|x_0,x_1}$ we can define the following mixture of bridges that is called reciprocal process:

$$\mathbb{Q} \stackrel{\text{def}}{=} \int \mathbb{P}^{\mathrm{ref}}_{|x_0,x_1} \pi(x_0, x_1)\, dx_0\, dx_1.$$

Sampling trajectory $x_0, \ldots, x_t, \ldots, x_1$ from $\mathbb{Q}$:

1. Sample $(x_0, x_1) \sim \pi(x_0, x_1)$;
2. Sample trajectory from $\mathbb{P}^{\mathrm{ref}}_{|x_0,x_1}$.

Key limitation: $\mathbb{Q}$ is **not representable as an SDE**

$\Downarrow$

no direct sampling of $q^{\mathbb{Q}}(x_1|x_0)$ or $q^{\mathbb{Q}}(x_0|x_1)$.

Markovian Projection: From Reciprocal to Markov Process

Idea: find the closest Markov process[4] $\mathbb{P}$ by KL projection:

$$\mathbb{P} = \arg\min_{\widetilde{\mathbb{P}}} \mathrm{KL}\left(\mathbb{Q} \| \widetilde{\mathbb{P}}\right) \quad \text{such that}$$

$$\widetilde{\mathbb{P}} : dx_t = \Big[\underbrace{b(x_t, t)}_{\substack{\text{reference drift} \\ \text{from } \mathbb{P}^{\mathrm{ref}}}} + a(x_t, t) \quad \underbrace{u(x_t, t)}_{\substack{\text{Markov correction} \\ \text{(to be optimized)}}} \Big] dt + \sigma(x_t, t)\, dW_t .$$

Markovian projection keeps the reference dynamics of $\mathbb{P}^{\mathrm{ref}}$ and chooses the Markov correction $u(x_t, t)$ that makes $\mathbb{P}$ closest to the reciprocal process $\mathbb{Q}$ in KL.

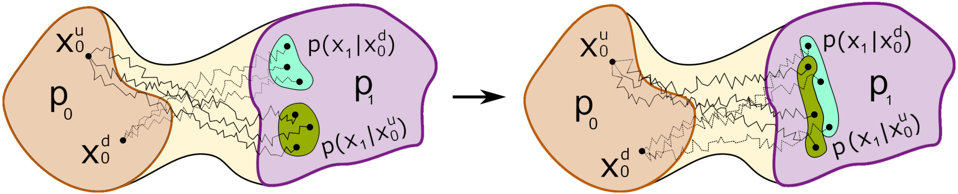

[4] There are many processes satisfying this property. .

Optimal Drift of the KL Projection

Each reference bridge $\mathbb{P}^{\mathrm{ref}}_{|x,y}$ has drift

$$b(x_t, t) + a(x_t, t)\nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t).$$

However, in the reciprocal process $\mathbb{Q}$ the endpoint x_1 remains random, so $\mathbb{Q}$ is **non-Markov**. To recover a Markov process, we project $\mathbb{Q}$ onto Markov dynamics by minimizing $\mathrm{KL}\left(\mathbb{Q}\|\widetilde{\mathbb{P}}\right)$. By Girsanov's theorem, this is equivalent to the following drift-matching objective:

$$\mathbb{E}_{\mathbb{Q}} \int_0^1 \left\| u(x_t, t) - \nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t) \right\|^2_{a(x_t,t)} dt.$$

Since $u(x_t, t)$ must depend only on x_t, the MSE minimizer is the conditional expectation[5]:

$$u^*(x_t, t) = \mathbb{E}_{y\sim q^{\mathbb{Q}}(x_1|x_t)} \nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t).$$

Hence, the projected Markov process $\mathbb{P}$ satisfies

$$dx_t = \Big[b(x_t, t) + a(x_t, t) \underbrace{\mathbb{E}_{y\sim q^{\mathbb{Q}}(x_1|x_t)} \nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t)}_{=:u(x_t,t)} \Big] dt + \sigma(x_t, t)\, dW_t.$$

[5] General fact; see the theory of Bregman divergences.

Why This Projection Is Exactly the Right Markov Process?

Recall that the reciprocal process $\mathbb{Q}$ is a mixture of bridges:

$$\mathbb{Q} = \mathbb{E}_{\pi(x,y)}\left[\mathbb{P}^{\mathrm{ref}}_{|x,y}\right] \qquad \Rightarrow \qquad q_t^{\mathbb{Q}}(x_t) = \mathbb{E}_{\pi(x,y)}\left[q^{\mathrm{ref}}(x_t|x,y)\right].$$

Each bridge satisfies the Kolmogorov-Fokker-Planck (KFP) equation:

$$\partial_t q^{\mathrm{ref}}(x_t|x,y) = -\nabla\cdot\Big(\big[b(x_t,t)+a(x_t,t)\nabla_{x_t}\log q^{\mathrm{ref}}(y|x_t)\big]\,q^{\mathrm{ref}}(x_t|x,y)\Big)+\tfrac{1}{2}\partial_{ij}\Big(a_{ij}\,q^{\mathrm{ref}}(x_t|x,y)\Big).$$

Taking expectation over $\pi(x,y)$ yields

$$\partial_t q_t^{\mathbb{Q}}(x_t) = -\nabla\cdot\Big(b(x_t,t)\,q_t^{\mathbb{Q}}(x_t)+a(x_t,t)\underbrace{\int \nabla_{x_t}\log q^{\mathrm{ref}}(y|x_t)\,q^{\mathrm{ref}}(x_t|x,y)\,\pi(x,y)\,dx\,dy}_{=q_t^{\mathbb{Q}}(x_t)\,u^*(x_t,t)}\Big)+\tfrac{1}{2}\partial_{ij}\Big(a_{ij}\,q_t^{\mathbb{Q}}(x_t)\Big).$$

By uniqueness of the KFP equation, the Markovian projection $\mathbb{P}$ has the same marginals[6] as $\mathbb{Q}$, so now we can map p_0 to p_1.

[6] D. Gyöngy (1986). **"Mimicking the one-dimensional marginal distributions of processes having an Itô differential"**. In: *Probability Theory and Related Fields* 71.4, pp. 501–516.

Bridge Matching: What Remains to Be Learned?

The Markovian projection $\mathbb{P}$ gives an SDE that maps

$$p_0 \to p_1.$$

However, simulating this SDE requires the drift

$$u(x_t, t).$$

The drift $u(x_t, t)$ is generally **unknown**, but it can be **learned** from bridge samples[7].

In particular, Bridge Matching (BM) shows that

$$u^* = \arg\min_u \mathbb{E}_{y,t,x_t}\left[\|u(x_t,t) - \nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t)\|^2\right],$$

$$\text{where } (x,y) \sim \pi(x_0, x_1), \qquad t \sim U[0,1], \qquad x_t \sim q^{\mathrm{ref}}(x_t|x,y).$$

[7] An analogous statement holds for the backward drift and the reverse-time process.

Bridge Dynamics under Scaled Wiener Reference Process

The scaled Wiener process is the simplest reference process:

$$\mathbb{P}^{\mathrm{ref}}: \qquad dx_t = \sqrt{\epsilon}\, dW_t.$$

Conditioning on fixed endpoints $x_0 = x$ and $x_1 = y$ yields the corresponding bridge process

$$\mathbb{P}^{\mathrm{ref}}_{|x,y}: \qquad dx_t = \Big[\epsilon\, u^*(x_t, t; y)\Big] dt + \sqrt{\epsilon}\, dW_t,$$

where the exact bridge correction is

$$u^*(x_t, t; y) = \frac{y - x_t}{\epsilon(1-t)}.$$

The corresponding bridge distribution is Gaussian:

$$q^{\mathrm{ref}}(x_t|x, y) = \mathcal{N}((1-t)x + ty,\, \epsilon t(1-t)\mathbb{I})\,.$$

Thus, under the scaled Wiener reference, the bridge remains Gaussian and differs from the reference process only by the drift correction $\epsilon\, u^*(x_t, t; y)$.

Schrödinger Bridge: Markov and Reciprocal

Notably, the Schrödinger Bridge $\mathbb{P}^*$ is both Markov and reciprocal.

Markov property

- The reference process $\mathbb{P}^{\mathrm{ref}}$ is Markov.
- The SB is obtained as a Doob-h transform of $\mathbb{P}^{\mathrm{ref}}$.
- Therefore, $\mathbb{P}^*$ remains Markov.

$$\frac{d\mathbb{P}^*}{d\mathbb{P}^{\mathrm{ref}}} = f(x_0)g(x_1).$$

Reciprocal property

- By optimality, $\mathbb{P}^*$ solves both dynamic and static SB problems.
- Hence, its bridges coincide with those of $\mathbb{P}^{\mathrm{ref}}$.
- Therefore, $\mathbb{P}^*$ is reciprocal.

$$\mathbb{P}^*_{|x,y} = \mathbb{P}^{\mathrm{ref}}_{|x,y}.$$

QUESTION: Can Bridge Matching be used to recover the Schrödinger Bridge?

Diffusion Schrödinger Bridge Matching

Iterative Markovian Fitting (IMF)[8,9]: Pretraining Phase

Key idea: Iteratively perform BM, while updating the coupling after each step.

Step 1: Apply BM with some initial coupling, for example $\pi^0(x_0, x_1) = p_0(x_0)p_1(x_1)$. Then the first drift correction is learned by

$$u^1 = \arg\min_u \mathbb{E}_{(x,y)\sim\pi^0(x_0,x_1)} \mathbb{E}_{t\sim U[0,1]} \mathbb{E}_{x_t\sim q^{\mathrm{ref}}(x_t|x,y)} \left|\left|u(x_t,t) - \nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t)\right|\right|^2.$$

Result: The learned drift defines a Markov path distribution.

$$\mathbb{P}^1 : dx_t = \Big[\underbrace{b(x_t,t) + a(x_t,t)u^1(x_t,t)}_{=:v^1(x_t,t)}\Big]dt + \sigma(x_t,t)dW_t$$

IMF now repeats the same BM step, but with the updated coupling $\pi^1(x_0, x_1)$ instead of the initial one.

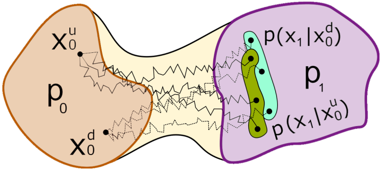

[8] Stefano Peluchetti (2023). **"Diffusion bridge mixture transports, Schrödinger bridge problems and generative modeling".** In: *Journal of Machine Learning Research* 24.374, pp. 1–51.
[9] Yuyang Shi et al. (2023). **"Diffusion schrödinger bridge matching".** In: *Advances in Neural Information Processing Systems* 36, pp. 62183–62223.

Iterative Markovian Fitting (IMF): Fine-tuning Phase

Step $l+1$: Apply BM with $(x,y)\sim\pi^l(x_0,x_1)$. Then the $l+1$-th drift correction is learned by

$$u^{l+1} = \arg\min_u \mathop{\mathbb{E}}_{(x,y)\sim\pi^l(x_0,x_1)} \mathop{\mathbb{E}}_{t\sim U[0,1]} \mathop{\mathbb{E}}_{x_t\sim q^{\mathrm{ref}}(x_t|x,y)} \left|\left|u(x_t,t) - \nabla_{x_t}\log q^{\mathrm{ref}}(y|x_t)\right|\right|^2.$$

Result: The learned drift defines a Markov path distribution.

$$\mathbb{P}^{l+1} : dx_t = \Big[\underbrace{b(x_t,t) + a(x_t,t)u^{l+1}(x_t,t)}_{=:v^{l+1}(x_t,t)}\Big]dt + \sigma(x_t,t)dW_t$$

Theorem (Convergence of IMF to the Schrödinger Bridge[10])

Let $\{\mathbb{P}^l\}_{l\geq 0}$ be the sequence of Markov path distributions produced by IMF, and let $\{v^l\}_{l\geq 0}$ be the corresponding learned drifts. Then, as $l\to\infty$,

$$\mathbb{P}^l \longrightarrow \mathbb{P}^*,$$

where $\mathbb{P}^*$ is the Schrödinger Bridge between p_0 and p_1 with respect to the reference process $\mathbb{P}^{\mathrm{ref}}$. Equivalently, $\mathbb{P}^*$ is induced by the limiting drift v^*.

[10] Yuyang Shi et al. (2023). **"Diffusion schrödinger bridge matching".** In: *Advances in Neural Information Processing Systems* 36, pp. 62183–62223.

Diffusion Schödinger Bridge Matching (DSBM)[11]

Parameterization.

$$u_\theta(x_t, t) \approx u^*(x_t, t) = \mathbb{E}_{y\sim\pi^*(x_1|x_t)} \nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t), \text{ where } q^{\mathrm{ref}} \text{ given by } \mathbb{W}^\epsilon$$

Training (iteration l).

1. Sample endpoints, bridge time, and state:

$$(x, y) \sim \pi_\theta(x_0, x_1), \quad t \sim U[0, 1], \qquad x_t \sim q^{\mathrm{ref}}(x_t|x, y).$$

2. Update network:

$$\min_\theta \mathbb{E}_{(x,y),t,x_t} \left|\left| u_\theta(x_t, t) - \nabla_{x_t} \log q^{\mathrm{ref}}(y|x_t) \right|\right|^2.$$

- **Error accumulation.** Each step uses generated samples from the previous iterate, which leads to **error accumulation**. This can be **mitigated** by jointly training the forward and Markov representations of the path distribution $\mathbb{P}^l$.
- **Not simulation-free.** Each step requires **simulating the SDE**.

[11] Yuyang Shi et al. (2023). **"Diffusion schrödinger bridge matching".** In: *Advances in Neural Information Processing Systems* 36, pp. 62183–62223.

Unpaired Image-to-Image Translation

(a) `cat` (p_0) $\rightarrow$ `wild` (p_1)

(b) `wild` (p_1) $\rightarrow$ `cat` (p_0)

Unpaired Image-to-Image Translation

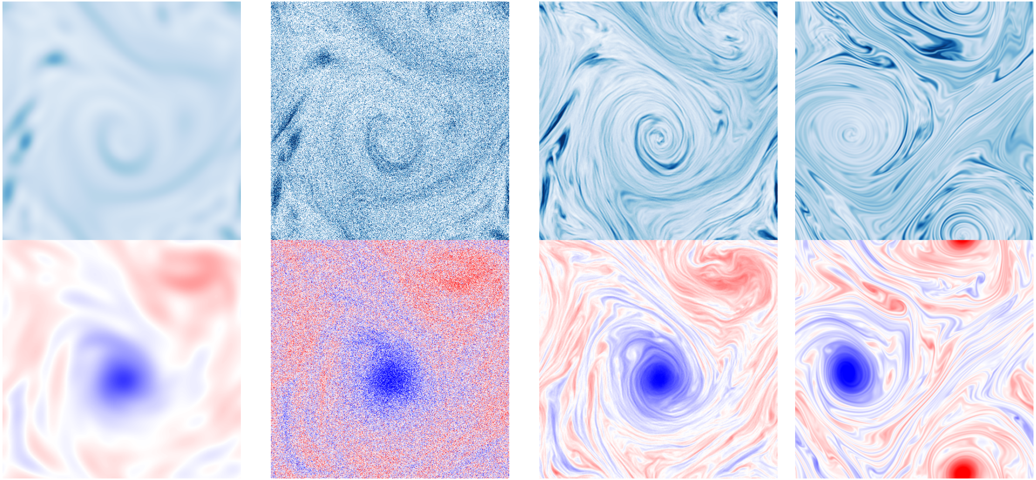

(a) p_0 (b) $q_\theta(x_t)$ (c) $\pi_\theta(x_1|x_0)$ (d) p_1

Conclusions & Takeaways

- The SB solution $\mathbb{P}^*$ is both **Markov** and **reciprocal**.
- A reciprocal process can be viewed as a **mixture of bridges**, but it is generally **not representable as an SDE**.
- The **Markovian projection** recovers the closest Markov path distribution $\mathbb{P}$ by averaging bridge drifts, which yields an SDE with the same marginals as the reciprocal process.
- **Bridge Matching** turns the unknown Markov correction $u(x_t, t)$ into a regression problem over bridge samples.
- **IMF** repeatedly updates the coupling and, in the limit, converges to the Schrödinger Bridge $\mathbb{P}^*$.
- **DSBM** provides a practical neural implementation of this idea.

Main message: Schrödinger Bridge Matching learns the Schrödinger Bridge by combining reciprocal bridge structure, Markovian projection, and bridge-matching regression.