CONSTRUCTIVE CYTOBRIDGE: SIMULATION-FREE STOCHASTIC CONTROL OF INTERACTING PARTICLES

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning continuous cell-state dynamics from unpaired single-cell snapshots is essential for understanding cellular transitions and interactions. Achieving this without trajectory simulation becomes particularly challenging in the presence of pairwise cell–cell interactions, because they couple endpoint matching with the conditional paths connecting matched cells. Here we introduce Constructive CytoBridge (CCB), a simulation-free framework for stochastic optimal control with prescribed external and pairwise potentials. By conditioning on endpoint pairs and decomposing the control cost, CCB reveals an entropy-regularized quadratic OT structure in stochastic optimal control problems with pairwise interactions. We solve the resulting coupled problem by alternating KL-proximal transport updates with optimization of neural-network-parameterized conditional paths, and then regress a velocity field and a score function to support both deterministic and stochastic inference without requiring trajectory simulation during training. Experiments on synthetic and real datasets show that CCB accurately recovers the optimal control law and reconstructs marginal distributions, while incorporating external and pairwise potentials improves interpolation at unobserved time points.

1 INTRODUCTION

Reconstructing continuous dynamics from sparse single-cell snapshots is important to studying differentiation, fate transitions, and tissue regeneration. Because sequencing is destructive, trajectory inference must recover population dynamics from unpaired observations (Zheng et al., 2017; Chen et al., 2022; Ding et al., 2022; Bunne et al., 2024; Zhang et al., 2025c). Existing approaches infer local direction or fate from molecular and lineage information (La Manno et al., 2018; Bergen et al., 2020; Lange et al., 2022; Forrow & Schiebinger, 2021), align snapshots through static OT (Schiebinger et al., 2019; Klein et al., 2025), or learn continuous population flows (Tong et al., 2020; Huguet et al., 2022; Yeo et al., 2021; Zhang et al., 2024; Peng et al., 2026b). Dynamic OT and Schrödinger Bridges provide deterministic and stochastic optimal-control formulations between fixed marginals (Benamou & Brenier, 2000; Wan et al., 2023; Chen et al., 2016; De Bortoli et al., 2021; Shi et al., 2023; Gu et al., 2025; Bunne et al., 2023a; Maddu et al., 2025). Unbalanced OT and Wasserstein–Fisher–Rao geometry further incorporate mass variation (Benamou, 2003; Chizat et al., 2018b; Liero et al., 2018; 2016; Kondratyev et al., 2016; Chizat et al., 2018a; Zhang et al., 2025a), enabling single-cell reconstruction with proliferation and apoptosis (Sha et al., 2024; Zhang et al., 2025b; Sun et al., 2025).

Training these dynamical models through neural differential equations often requires repeated numerical integration (Chen et al., 2018). Flow Matching learns vector fields by regression along conditional probability paths without trajectory simulation during training (Lipman et al., 2023; Liu et al., 2023; Albergo & Vanden-Eijnden, 2023). Building on this principle, subsequent methods select appropriate endpoint couplings to learn OT maps and Schrödinger Bridges (Tong et al., 2024a; Pooladian et al., 2023; Tong et al., 2024b; Klein et al., 2024), and extend the construction to observations at multiple times (Albergo et al., 2024; Rohbeck et al., 2025). Single-cell trajectory inference needs also account for cell proliferation and death. Several methods address the resulting unbalanced distributions (Eyring et al., 2024; Wang et al., 2025; Peng et al., 2026a). Recent work also explicitly models discrete branching events (Ying et al., 2026a). To guide the trajectories, flow matching methods incorporate data geometry (Kapuśniak et al., 2024), condition dynamics on the population (Atanackovic et al., 2025), or introduce second-order dynamics (Sun et al., 2026).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

However, cell-state transitions are shaped not only by intracellular processes but also by surrounding populations and their spatial organization. Mean-field games (Lasry & Lions, 2007), mean-field Schrödinger Bridges (Backhoff et al., 2020), and neural mean-field control solvers (Ruthotto et al., 2020; Lin et al., 2021; Liu et al., 2022) provide frameworks for such population-dependent dynamics. Zhang et al. (2025d) incorporate cell–cell interactions into trajectory inference but still rely on trajectory simulation. GSBM (Liu et al., 2024) allows state costs with mean-field interactions, but repeatedly simulates the current dynamics to update the transport plan, and its conditional bridge kinetic energy diverges. MFG-FM (Yu et al., 2026) likewise uses an outer loop of velocity regression and ODE trajectory resampling, and solves an MFG problem with a terminal control cost rather than a two-point boundary-value problem. WLF (Neklyudov et al., 2024) optimizes the action without trajectory simulation, but its sample-based formulation requires a Hamiltonian that is linearizable in the density, a condition that general pairwise interactions need not satisfy. More fundamentally, pairwise interactions change the optimization structure: unlike OT-CFM and SF^2M, whose endpoint costs and paths can be fixed in advance, pairwise interactions couple the transport plan and conditional paths because each path cost depends on the others and their weights. Both must therefore be optimized jointly to solve the optimal control problem.

Here we propose **Constructive CytoBridge (CCB)**, a simulation-free framework for stochastic optimal control of pairwise interacting particles. Under fixed marginal distributions and Fokker–Planck constraints, CCB incorporates external and pairwise potentials into the action. Endpoint conditioning then yields a formulation for jointly optimizing the transport plan and conditional paths. Regressing the resulting conditional velocity and score yields a dynamical model supporting both PF-ODE and SDE inference. Unlike methods that repeatedly update the transport plan through rollouts, CCB directly optimizes it and requires no ODE or SDE trajectory simulation during training.

Our main contributions are:

- We derive an endpoint-conditioned decomposition of the stochastic control action with pairwise interactions, revealing an entropy-regularized quadratic OT structure that couples the transport plan and conditional probability paths.
- We develop an alternating optimization scheme that couples KL-proximal transport updates with gradient-based learning of conditional probability paths.
- We demonstrate strong performance on synthetic stochastic optimal-control benchmarks, real single-cell trajectory reconstruction, and held-out-time interpolation for mean-field and spatial transcriptomics data.

2 Related Work and Preliminaries

2.1 Related Work

Flow matching for optimal control. Flow Matching connects simulation-free learning of dynamics with optimal control of distributions. Lipman et al. (2023) and Liu et al. (2023) learn continuous dynamics by regressing time-dependent vector fields along tractable conditional probability paths without trajectory simulation during training. Albergo & Vanden-Eijnden (2023) provide a related simulation-free formulation. Tong et al. (2024a) and Pooladian et al. (2023) use transport-cost-based couplings to regularize the learned flows, encouraging straighter trajectories and improving sample quality. For stochastic dynamics, Tong et al. (2024b) solve Schrödinger Bridge problems without trajectory simulation during training, while Shi et al. (2023) refine Schrödinger Bridges through iterative Markovian fitting. Subsequent work extends flow matching to data-induced geometry, general single-particle Lagrangians, and mass-varying dynamics (Kapuśniak et al., 2024; Du et al., 2026; Wang et al., 2025; Peng et al., 2026a). These methods either prescribe conditional paths for endpoint pairs or optimize each conditional path independently. Pairwise interactions break this separation because each path cost depends on other paths and their transport weights. CCB addresses this dependence by jointly optimizing the transport plan and conditional probability paths.

Interacting particles and mean-field control. Interactions couple individual particle dynamics. Mean-field games, mean-field control, and mean-field Schrödinger problems provide variational and stochastic-control formulations (Lasry & Lions, 2007; Backhoff et al., 2020). Neural solvers

address high-dimensional problems through Lagrangian trajectory optimization (Ruthotto et al., 2020), alternating training of density and value networks (Lin et al., 2021), or forward–backward SDEs (Liu et al., 2022). Neklyudov et al. (2024) optimize the action without simulation, but their sample-based formulation requires a Hamiltonian that is linearizable in the density, excluding general pairwise potentials. Liu et al. (2024) support mean-field interactions but require repeated rollouts to update the coupling, and their raw conditional-bridge kinetic cost diverges under exact endpoint conditioning. Yu et al. (2026) likewise refine the coupling through reflow-style velocity regression and trajectory resampling, but use a terminal cost rather than a two-point boundary-value formulation. In contrast, CCB minimizes the population action between fixed endpoint marginals by jointly optimizing the transport plan and conditional paths without rollouts.

Single-cell trajectory inference. The destructive nature of single-cell sequencing yields unpaired population snapshots rather than observed trajectories. Static OT methods infer cell fates, perturbation responses, and spatiotemporal correspondences across snapshots (Schiebinger et al., 2019; Bunne et al., 2023b; Klein et al., 2025). Continuous-time methods reconstruct intermediate dynamics with neural differential equations (Tong et al., 2020; Huguet et al., 2022; Koshizuka & Sato, 2023), while stochastic approaches infer dynamics by solving Schrödinger Bridge problems (Bunne et al., 2023a; Hong et al., 2025). Extensions to unbalanced dynamics model continuous mass variation or discrete birth–death events (Zhang et al., 2025b; Sun et al., 2025; Peng et al., 2026a; Ying et al., 2026a). Atanackovic et al. (2025) condition dynamics on the population but do not explicitly minimize an action with pairwise potentials, while Zhang et al. (2025d) incorporate cell–cell interactions into an unbalanced mean-field Schrödinger Bridge but rely on Neural ODE simulation. CCB instead learns interaction-aware dynamics between fixed observed marginals while remaining simulation-free throughout training.

2.2 Optimal Transport and Schrödinger Bridge

Given endpoint distributions μ_0, μ_1, consider the following optimal control problem that minimizes the Action Functional:

$$\inf_{\rho,b} \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2}\|b_t(x)\|^2 \rho_t(x)\mathrm{d}x\mathrm{d}t, \qquad \text{s.t.} \quad \partial_t \rho_t + \nabla \cdot (\rho_t b_t) = \nu \Delta \rho_t, \quad \rho_0 = \mu_0,\ \rho_1 = \mu_1. \tag{1}$$

For $\nu = 0$, the problem reduces to Dynamical Optimal Transport (Benamou & Brenier, 2000), with b_t as the transport velocity. For $\nu > 0$, it defines the Schrödinger Bridge (Chen et al., 2016) relative to zero-drift Brownian motion, with diffusion coefficient $\sqrt{2\nu}$.

Both problems admit simulation-free regression objectives constructed from static couplings. OT-CFM (Tong et al., 2024a) solves static OT with cost $c(x_0, x_1) = \frac{1}{2}\|x_1 - x_0\|^2$ and regresses the conditional velocity $x_1 - x_0$ along the linear interpolation $(1-t)x_0 + tx_1$; SF^{2}M (Tong et al., 2024b) instead adds entropy regularization with weight 2ν, samples from the conditional Brownian bridge $\mathcal{N}((1-t)x_0 + tx_1, 2\nu t(1-t)I)$, regresses the PF-ODE velocity v_θ and score s_ψ, and recovers the SDE drift through $b_{\theta,\psi} = v_\theta + \nu s_\psi$. Neither method requires trajectory simulation during training.

2.3 Wasserstein Lagrangian Flow

Adding a potential functional $\mathcal{U}[\rho]$ to the optimal control objective above yields the Wasserstein Lagrangian Flow formulation and its stochastic extension (Neklyudov et al., 2024):

$$\begin{aligned} &\inf_{\rho,b} \int_0^1 \left[\frac{1}{2}\int_{\mathbb{R}^d} \|b_t(x)\|^2 \rho_t(x)\,\mathrm{d}x - \mathcal{U}[\rho_t]\right]\mathrm{d}t, \\ &\text{s.t.} \quad \partial_t \rho_t + \nabla \cdot (\rho_t b_t) = \nu \Delta \rho_t, \qquad \rho_0 = \mu_0, \quad \rho_1 = \mu_1. \end{aligned} \tag{2}$$

When $\nu = 0$, this problem is the deterministic Wasserstein Lagrangian Flow. If the transport velocity is the gradient of a scalar field, the kinetic term is one half of the squared norm of the tangent vector under the Otto metric. The resulting functional is the classical Lagrangian action on the probability manifold. When $\nu > 0$, the problem becomes its stochastic extension, where b_t is the SDE drift with diffusion coefficient $\sqrt{2\nu}$. When the Hamiltonian is linearizable in the density, as for external potentials of the form $\mathcal{U}[\rho] = \int V(x)\rho(x)\mathrm{d}x$, WLF (Neklyudov et al., 2024) solves this problem using simulation-free dual min-max optimization.

3 METHODOLOGY OF CONSTRUCTIVE CYTOBRIDGE

In the following, we study the Interacting Stochastic Wasserstein Lagrangian Flow (IS-WLF) problem. Given endpoint distributions $\mu_0, \mu_1 \in \mathcal{P}_2(\mathbb{R}^d)$, an external potential $V(x)$, a symmetric pairwise potential $W(x, y)$, and diffusion strength $\nu > 0$, we take the SDE drift b_t as the control and solve

$$\inf_{\rho,b} \mathcal{A}(\rho, b) := \int_0^1 \left[\frac{1}{2} \int_{\mathbb{R}^d} \|b_t(x)\|^2 \rho_t(x)\, \mathrm{d}x - \int_{\mathbb{R}^d} V(x)\rho_t(x)\, \mathrm{d}x \right. \\ \left. - \frac{1}{2} \iint_{\mathbb{R}^d \times \mathbb{R}^d} W(x, y)\rho_t(x)\rho_t(y)\, \mathrm{d}x\, \mathrm{d}y \right] \mathrm{d}t. \tag{3}$$

$$\text{s.t.} \quad \partial_t \rho_t + \nabla \cdot (\rho_t b_t) = \nu \Delta \rho_t, \qquad \rho_0 = \mu_0, \quad \rho_1 = \mu_1, \qquad \rho_t(x)\mathrm{d}x \in \mathcal{P}_2(\mathbb{R}^d). \tag{4}$$

When $\nu = 0$, the system has no diffusion and reduces to a deterministic problem. When the pairwise interaction kernel $W(x, y) = 0$, it reduces to the WLF problem with only an external potential introduced above.

3.1 STATIC TRANSPORT COST VIA ENDPOINT CONDITIONING

Recovering the optimal control through simulation-free Flow Matching regression requires the optimal transport plan and the conditional velocity and score derived from the associated conditional probability paths. To jointly determine the transport plan and conditional paths, we condition on the endpoints and express the original action as a function of both.

The controlled SDE $\mathrm{d}X_t = b_t(X_t)\,\mathrm{d}t + \sqrt{2\nu}\,\mathrm{d}B_t$ induces a path measure $\mathbb{P}$ whose endpoint joint distribution is the transport plan $\pi \in \Pi(\mu_0, \mu_1)$. Conditioning $\mathbb{P}$ on an endpoint pair $z = (x_0, x_1)$ gives the conditional path measure $\mathbb{P}^z$, with time-t marginal ρ_t^z. We use the Brownian bridge measure $\mathbb{Q}^{\nu,z}$ between the same endpoints as the reference. The conditional external potential cost $C_V(z)$ and conditional interaction cost $C_W(z, z')$ are then defined by $C_V(z) = \int_0^1 \int V(x)\rho_t^z(\mathrm{d}x)\,\mathrm{d}t$ and $C_W(z, z') = \int_0^1 \iint W(x, y)\rho_t^z(\mathrm{d}x)\rho_t^{z'}(\mathrm{d}y)\,\mathrm{d}t$.

Proposition 1 (Action decomposition via endpoint conditioning)**.** With the above definitions, disintegrating $\mathbb{P}$ over the endpoint pair z yields the following decomposition of the IS-WLF action:

$$\mathcal{A}(\rho, b) \doteq 2\nu \mathrm{KL}(\pi \mid \mu_0 \otimes \mu_1) - \frac{1}{2} \iint C_W(z, z')\pi(\mathrm{d}z)\pi(\mathrm{d}z') \\ + \int \left[\frac{1}{2}\|x_1 - x_0\|^2 + 2\nu \mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) - C_V(z) \right] \pi(\mathrm{d}z). \tag{5}$$

Here $\doteq$ denotes equality up to a constant; the decomposition is proved in Appendix A.1. External and pairwise potentials can drive conditional paths away from Brownian bridges, and the conditional KL quantifies the associated control cost. From an alternative perspective, Appendix A.2 derives the conditional kinetic terms by renormalizing GSBM's conditional-bridge cost.

For empirical marginals $\widehat{\mu}_0 = \sum_{i=1}^n a_i \delta_{x_0^i}$ and $\widehat{\mu}_1 = \sum_{j=1}^m b_j \delta_{x_1^j}$ with positive normalized weights a, b, let $z_{ij} = (x_0^i, x_1^j)$ denote a pair of observed states. The transport plan is represented by $\pi_P = \sum_{i,j} P_{ij}\delta_{z_{ij}}$, where $P \in \mathcal{U}(a, b) := \{P \geq 0 : P\mathbf{1}_m = a,\ P^\top \mathbf{1}_n = b\}$. We use $\mathbb{P}^{ij}$, ρ_t^{ij}, and $\mathbb{Q}^{\nu,ij}$ as shorthand for the corresponding conditional quantities at $z = z_{ij}$. For each conditional path, the external potential cost is $C_V(z_{ij}) = \int_0^1 \int V(x)\rho_t^{ij}(\mathrm{d}x)\,\mathrm{d}t$, giving the linear cost coefficient $c_{ij} = \frac{1}{2}\|x_1^j - x_0^i\|^2 + 2\nu \mathrm{KL}(\mathbb{P}^{ij} \mid \mathbb{Q}^{\nu,ij}) - C_V(z_{ij})$. The interaction between two conditional paths is encoded by $Q_{ij,k\ell} = C_W(z_{ij}, z_{k\ell}) = \int_0^1 \iint W(x, y)\rho_t^{ij}(\mathrm{d}x)\rho_t^{k\ell}(\mathrm{d}y)\,\mathrm{d}t$.

Corollary 1 (Static quadratic OT for empirical distributions)**.** With the family of conditional probability paths fixed, empirically discretizing the objective in Proposition 1 yields the following entropy-regularized quadratic OT problem:

$$\min_{P \in \mathcal{U}(a,b)} \sum_{i,j} c_{ij} P_{ij} - \frac{1}{2} \sum_{i,j,k,\ell} Q_{ij,k\ell} P_{ij} P_{k\ell} + 2\nu \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j}. \tag{6}$$

Appendix A.3 gives the derivation. This formulation yields tractable costs for alternating updates.

The discrete objective differs from standard OT in two respects. 1) For fixed conditional paths, it is an entropy-regularized quadratic OT problem in P. 2) The optimal conditional paths are neither prescribed, as in OT-CFM and SF^{2}M, nor pre-optimizable, as in SUDO (Ying et al., 2026b). Indeed, for fixed z_{ij}, the same $\mathbb{P}^{ij}$ appears in every $Q_{ij,k\ell}$, while their weights $P_{ij}P_{k\ell}$ depend on the unknown P. Thus, the transport plan and conditional paths must be optimized jointly.

3.2 Alternating Optimization of Transport Plans and Conditional Paths

Because the transport plan and conditional probability paths are coupled, we optimize them alternately, updating one block while holding the other fixed. We next describe both updates.

Proximal update of the transport plan. For the quadratic OT subproblem with fixed conditional probability paths, we use KL-proximal linearization to reduce it to a sequence of entropy-regularized linear OT subproblems that can be solved by the Sinkhorn algorithm (Cuturi, 2013). At iteration r, we linearize the quadratic interaction term around the current transport plan $P^{(r)}$. Using the symmetry of the interaction cost $Q_{ij,k\ell} = Q_{k\ell,ij}$, we obtain the effective cost $C_{ij}^{(r)} = c_{ij} - \sum_{k,\ell} Q_{ij,k\ell} P_{k\ell}^{(r)}$ and update

$$P^{(r+1)} = \operatorname*{arg\,min}_{P \in \mathcal{U}(a,b)} \left\{ \sum_{i,j} C_{ij}^{(r)} P_{ij} + 2\nu \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j} + \frac{1}{\tau} \sum_{i,j} P_{ij} \log \frac{P_{ij}}{P_{ij}^{(r)}} \right\}. \tag{7}$$

Here $\tau > 0$ is the proximal step size. The entropy regularization coefficient 2ν in the original objective remains unchanged, while the additional KL proximal term controls the magnitude of each transport-plan update. After combining the two entropy-related terms, each update can be solved by Sinkhorn iterations.

Proposition 2 (Proximal fixed points and convergence)**.** For fixed conditional paths and any $\tau > 0$, a strictly positive $P^\star \in \mathcal{U}(a,b)$ is a fixed point of the KL-proximal update if and only if it satisfies the KKT conditions of the entropy-regularized quadratic OT objective. Moreover, with a fixed step size satisfying $\tau^{-1} > \|Q\|_{\mathrm{op}}$, the iterates converge to such a point. If the objective is convex on $\mathcal{U}(a,b)$, the resulting fixed point is also a global minimizer. Here, $\|Q\|_{\mathrm{op}}$ is the spectral norm of Q indexed over endpoint pairs. The proofs are provided in Appendices A.4–A.5.

Parameterization and update of conditional probability paths. With the endpoint transport plan fixed, for an endpoint pair $z = (x_0, x_1)$, we approximate the optimal conditional path by a Gaussian distribution $\rho_t^z = \mathcal{N}(m_t^z, \Sigma_t^z)$, where m_t^z, Σ_t^z are differentiable and $\Sigma_t^z \succ 0$. We define its path measure through the compatible affine Itô SDE

$$\mathrm{d}X_t^z = b_t^z(X_t^z)\mathrm{d}t + \sqrt{2\nu}\mathrm{d}B_t^z, \qquad b_t^z(x) = \dot{m}_t^z + \left(\frac{1}{2}\dot{\Sigma}_t^z - \nu I\right)(\Sigma_t^z)^{-1}(x - m_t^z). \tag{8}$$

The derivation is provided in Appendix A.6. To obtain the transport cost for each endpoint pair, we compute the conditional path KL introduced above.

Proposition 3 (Closed-form conditional path KL)**.** Let $\mathbb{Q}^{\nu,z}$ be the Brownian bridge with the same endpoints and diffusion coefficient. Under the conditions for Girsanov's theorem,

$$2\nu \mathrm{KL}(\mathbb{P}^z | \mathbb{Q}^{\nu,z}) = \frac{1}{2}\int_0^1 \left[\|d_t^z\|^2 + \mathrm{tr}\left(H_t^z \Sigma_t^z (H_t^z)^\top\right)\right] \mathrm{d}t, \tag{9}$$

where

$$d_t^z = \dot{m}_t^z - \frac{x_1 - m_t^z}{1-t}, \qquad H_t^z = \left(\frac{1}{2}\dot{\Sigma}_t^z - \nu I\right)(\Sigma_t^z)^{-1} + \frac{I}{1-t}. \tag{10}$$

The derivation is provided in Appendix A.7.

Not every conditional-path parameterization yields finite path KL; the behavior near the boundaries is crucial. Therefore, we choose a parameterization whose mean and covariance behave like those of a Brownian bridge near $t = 0, 1$.

Proposition 4 (Sufficient conditions for finite path KL). Let $h(t) = t(1-t)$. If

$$m_t^z = (1-t)x_0 + tx_1 + h(t)r_t^z, \qquad \Sigma_t^z = 2\nu h(t) R_t^z, \tag{11}$$

where r_t^z, R_t^z are continuously differentiable on $[0,1]$, R_t^z is symmetric positive definite, and $R_0^z = R_1^z = I$, then $\mathrm{KL}(\mathbb{P}^z|\mathbb{Q}^{\nu,z}) < \infty$. The proof is provided in Appendix A.8.

To reduce the computational cost in high dimensions, we use a diagonal covariance approximation. Let the mean and covariance networks output $r_{\phi,t}^z(x_0,x_1), g_{\phi,t}^z(x_0,x_1) \in \mathbb{R}^d$, respectively, and parameterize m_t^z, Σ_t^z as

$$m_{\phi,t}^z = (1-t)x_0 + tx_1 + h(t)r_{\phi,t}^z, \qquad \Sigma_{\phi,t}^z = \mathrm{diag}(q_{\phi,t}^z), \qquad q_{\phi,t}^z = 2\nu h(t)\exp\left(h(t)g_{\phi,t}^z\right). \tag{12}$$

The exponential is applied coordinate-wise. For network outputs continuously differentiable in time on $[0,1]$, this parameterization meets the conditions of Proposition 4 and yields finite path KL.

With P fixed, the squared endpoint-distance cost and coupling entropy are independent of ϕ. Define the instantaneous cost as $\kappa_{\phi,t}^z = \frac{1}{2}[\|d_{\phi,t}^z\|^2 + \mathrm{tr}(H_{\phi,t}^z \Sigma_{\phi,t}^z (H_{\phi,t}^z)^\top)]$. Sample $t \sim \mathcal{U}(0,1)$, $Z, Z' \overset{\text{iid}}{\sim} \pi_P$, and $\epsilon, \epsilon' \overset{\text{iid}}{\sim} \mathcal{N}(0,I)$, and set $X_{\phi,t}^Z = m_{\phi,t}^Z + \sqrt{q_{\phi,t}^Z} \odot \epsilon$ and $Y_{\phi,t}^{Z'} = m_{\phi,t}^{Z'} + \sqrt{q_{\phi,t}^{Z'}} \odot \epsilon'$. After removing terms independent of ϕ, the conditional-path update objective is

$$\min_\phi \mathcal{J}(\phi; P), \qquad \mathcal{J}(\phi; P) = \mathbb{E}_{t,Z,Z',\epsilon,\epsilon'}\left[\kappa_{\phi,t}^Z - V(X_{\phi,t}^Z) - \frac{1}{2}W(X_{\phi,t}^Z, Y_{\phi,t}^{Z'})\right]. \tag{13}$$

We optimize this objective using reparameterization gradients. The entire procedure requires only evaluating the means, variances, and their time derivatives, together with sampling from the conditional Gaussian distributions, and does not require SDE simulation.

Alternating updates of the transport plan and conditional probability paths. We alternate between updating the endpoint transport plan P and the conditional-path parameters ϕ: with ϕ fixed, we update the transport plan using KL-proximal linearization; with P fixed, we update the conditional probability paths using reparameterization gradients. The required time integrals and potential expectations are estimated by Monte Carlo sampling, while the spatial expectation in the conditional path KL is evaluated analytically using Proposition 3. For the interaction cost in the transport-plan update, we directly estimate the current transport-plan-weighted quantity $\sum_{k,\ell} Q_{ij,k\ell} P_{k\ell}^{(r)}$ rather than explicitly computing or storing the fourth-order cost tensor Q. The entire procedure requires no SDE simulation. With descent updates for both blocks, the objective values form a nonincreasing sequence and converge when the objective is bounded below.

3.3 Velocity & Score Regression and Practical Implementation

After alternating optimization, we freeze the transport plan P and conditional-path parameters ϕ, and regress the PF-ODE velocity and score following SF^{2}M (Tong et al., 2024b). We sample $t \sim U[0,1]$, an endpoint pair $Z \sim \pi_P$, and a conditional sample $X \sim \rho_{\phi,t}^Z$, and use the loss

$$\mathcal{L}(\theta,\psi) = \mathbb{E}_{t,Z,X}\left[\|v_\theta(t,X) - v_{\phi,t}^Z(X)\|^2 + \lambda(t)^2\|s_\psi(t,X) - s_{\phi,t}^Z(X)\|^2\right]. \tag{14}$$

Here, $v_{\phi,t}^z$ is the conditional PF-ODE velocity, and $s_{\phi,t}^z = \nabla_x \log \rho_{\phi,t}^z$ is the conditional score. We use the SF^{2}M weighting $\lambda(t) = \sqrt{2\nu t(1-t)}$ to reduce the effect of score singularities near the endpoints during training. Analytical expressions for the supervision targets and the derivation of conditional regression are provided in Appendix A.9. Moreover, the Markovization induced by velocity and score regression does not increase the action. The proof is provided in Appendix A.10. After training, inference uses either the PF-ODE $\mathrm{d}X_t = v_\theta(t,X_t)\mathrm{d}t$ or the SDE $\mathrm{d}X_t = [v_\theta(t,X_t) + \nu s_\psi(t,X_t)]\mathrm{d}t + \sqrt{2\nu}\mathrm{d}B_t$.

For multiple observation times $t_0 < \cdots < t_K$, we optimize the transport plan and conditional probability paths separately between each pair of adjacent marginals $(\mu_{t_k}, \mu_{t_{k+1}})$. We then sample supervision signals from all segments and jointly regress a shared velocity field and score.

For large-scale datasets, to reduce mismatches caused by the restricted set of candidate matches within minibatches, we maintain a dynamic sparse coupling under the marginal constraints of all particles. At each iteration, we propose new edges between source and target particles and merge them with the current support. New edges receive positive reference mass through $\bar{P}_{ij}^{(r)} = (1 - \delta)P_{ij}^{(r)} + \delta a_i b_j, \delta \in (0, 1)$. For the uniform empirical marginals used in this work, each new edge has reference weight $\delta/(nm)$. We update the transport plan on the expanded support using the same effective cost $C_{ij}^{(r)}$ and Sinkhorn iterations. After the update, we retain the top-K highest-mass edges for each source and target and rebalance. If the marginal-error tolerance cannot be reached, we retain the expanded support. The procedure stores only edge indices and transport masses.

4 Experimental Evaluation

4.1 CCB Accurately Recovers the Optimal Control Law

We evaluate CCB on eight synthetic problems with computable ground truth, spanning two external and two interaction potentials in 10D and 20D. The marginal distributions are Gaussian, and intermediate samples are generated by the optimal controlled SDE; Appendix B.3.1 describes the datasets and ground-truth construction. We compare with GSBM (Liu et al., 2024) and adapted versions of CytoBridge (Zhang et al., 2025d) and WLF (Neklyudov et al., 2024). Adaptation is needed because CytoBridge places interactions in the Fokker–Planck dynamics rather than the action, whereas directly computable WLF requires a density-linearizable Hamiltonian; Appendix B.6 details both adaptations. We report absolute action bias and control-law error over eight rollouts. Modified WLF uses GAN-like min–max updates and can be unstable, Modified CytoBridge must balance marginal matching against action minimization, and GSBM repeatedly updates the coupling through forward–backward rollouts. In contrast, CCB avoids these difficulties by directly optimizing the transport plan while enforcing marginal constraints during path construction. Table 1 shows that CCB is best on **6/8** action and **7/8** control-law metrics, with small gaps on the others. We enable Sparse mode for the 20D experiments, where CCB leads all four control-law and three action metrics, demonstrating accurate recovery under sparse couplings. Figure 1 shows that the ground-truth distribution first expands and then contracts. CCB captures this trend, whereas GSBM does not.

Table 1: Optimal-control reconstruction on synthetic datasets. Results are mean ± standard deviation over eight rollouts. Ctrl. denotes control-law error, and Bias denotes absolute action bias.

Dataset	Mod. WLF		Mod. CytoBridge		GSBM		CCB	
	Ctrl. ($\times 10^{-1}$)↓	Bias ↓	Ctrl. ($\times 10^{-1}$)↓	Bias ↓	Ctrl. ($\times 10^{-1}$)↓	Bias ↓	Ctrl. ($\times 10^{-1}$)↓	Bias ↓
Attractive Interaction (10D)	4.720 ± 0.082	0.904 ± 0.074	1.330 ± 0.013	0.857 ± 0.048	1.740 ± 0.007	0.745 ± 0.020	**0.570 ± 0.011**	**0.110 ± 0.061**
Attractive Interaction (20D)	5.820 ± 0.091	1.930 ± 0.067	1.560 ± 0.012	2.815 ± 0.045	1.790 ± 0.015	1.703 ± 0.062	**0.380 ± 0.007**	**0.081 ± 0.003**
Repulsive Interaction (10D)	1.070 ± 0.026	0.837 ± 0.051	0.670 ± 0.011	**0.053 ± 0.058**	0.530 ± 0.006	0.418 ± 0.031	**0.320 ± 0.009**	0.064 ± 0.037
Repulsive Interaction (20D)	3.320 ± 0.096	4.111 ± 0.133	0.560 ± 0.007	0.306 ± 0.064	0.610 ± 0.006	0.597 ± 0.045	**0.130 ± 0.002**	**0.077 ± 0.006**
Linear external (10D)	**0.078 ± 0.001**	0.037 ± 0.025	0.770 ± 0.013	0.391 ± 0.048	0.170 ± 0.002	0.072 ± 0.012	0.320 ± 0.008	**0.024 ± 0.014**
Linear external (20D)	0.990 ± 0.014	**0.107 ± 0.033**	0.680 ± 0.008	0.882 ± 0.048	0.160 ± 0.001	0.239 ± 0.036	**0.120 ± 0.002**	0.147 ± 0.003
Quadratic external (10D)	7.690 ± 0.081	3.123 ± 0.093	1.400 ± 0.015	1.077 ± 0.065	1.620 ± 0.015	0.618 ± 0.039	**0.480 ± 0.012**	**0.070 ± 0.055**
Quadratic external (20D)	6.260 ± 0.091	0.225 ± 0.104	1.570 ± 0.015	2.905 ± 0.061	1.680 ± 0.013	1.416 ± 0.095	**0.370 ± 0.009**	**0.072 ± 0.007**

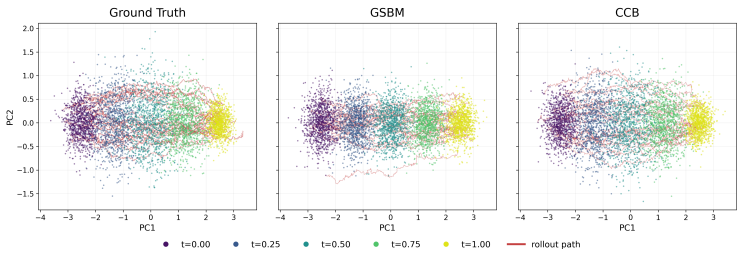

Figure 1: Optimal-control reconstruction on a 20D attractive interaction benchmark. Columns show ground truth, GSBM, and CCB in the first two PCs; colors mark five time points and red curves representative SDE rollouts.

4.2 CCB Accurately Reconstructs Dataset Marginals

As a generative model, CCB must learn the controlled dynamics while preserving the observed distributions. We evaluate CCB on six single-cell datasets for marginal reconstruction at observed time points. The two CCB settings, Attractive and Repulsive, use the same attractive and repulsive quadratic interaction potentials as in Section 4.1, respectively. Table 2 reports mean W_1 across time points. Despite optimizing interaction-aware conditional probability paths rather than the linear paths of OT-CFM, CCB remains comparable to the baselines in marginal reconstruction.

Table 2: Marginal reconstruction on real single-cell datasets, reported as mean W_1 over four rollouts.

Method / Inference	$W_1 \downarrow$					
	EMT (10D)	**EB (PCA-5D)**	**EB (PCA-50D)**	**EB (100D)**	**CITE (PCA-50D)**	**Weinreb (50D)**
OT-CFM	0.198	1.877	9.127	11.604	28.688	5.896
WFR-FM	0.223	1.946	8.889	11.311	29.132	5.853
WLF (OT)	0.237	2.784	9.892	12.107	34.344	7.104
VGFM	0.242	5.747	9.371	11.867	32.342	8.163
SF^2M	0.273 ± 0.0006	2.718 ± 0.0021	13.545 ± 0.0092	13.955 ± 0.0029	40.927 ± 0.0021	11.579 ± 0.0031
DeepRUOT	0.278 ± 0.0006	2.616 ± 0.0021	8.769 ± 0.0017	11.293 ± 0.0009	30.195 ± 0.0021	6.621 ± 0.0020
CCB (Attractive, PF-ODE)	**0.178**	1.550	7.852	11.049	**21.925**	5.171
CCB (Attractive, SDE)	0.193 ± 0.0008	1.452 ± 0.015	$\mathbf{7.723 \pm 0.017}$	11.283 ± 0.011	22.342 ± 0.032	5.409 ± 0.0075
CCB (Repulsive, PF-ODE)	0.182	1.547	7.799	10.956	21.964	**5.166**
CCB (Repulsive, SDE)	0.197 ± 0.0018	$\mathbf{1.446 \pm 0.014}$	7.758 ± 0.018	$\mathbf{10.842 \pm 0.027}$	22.031 ± 0.028	5.339 ± 0.014

4.3 Interaction and External Potentials Improve CCB Interpolation

Interpolation on data generated by Mean Field PDEs. We evaluate held-out-time interpolation on two 20D mean-field datasets: KS with pairwise attractive and repulsive interactions, and Coulomb with external and pairwise potentials. Models are trained on endpoint marginals and evaluated at three held-out times using W_1 and MMD^2; Appendix B.3.2 describes data generation. For a fair comparison, we also evaluate CytoBridge, WLF, and GSBM with the corresponding potentials. Table 3 shows that CCB attains the lowest mean W_1 with PF-ODE inference and the lowest mean MMD^2 with SDE inference on both datasets. Potential ablations on both datasets show that removing potentials markedly worsens both metrics under PF-ODE inference. Figure 2 visualizes the Coulomb ablation: removing the pairwise potential yields an overly compact cluster, whereas removing the external potential shifts and spreads it; the full model recovers both location and scale.

Table 3: Interpolation on held-out time points in mean-field datasets. Results are mean $\pm$ standard deviation over three training seeds.

Method / Inference	**KS**		**Coulomb**	
	$W_1 \downarrow$	MMD^2 $(\times 10^{-3}) \downarrow$	$W_1 \downarrow$	MMD^2 $(\times 10^{-3}) \downarrow$
OT-CFM	0.120 ± 0.0010	9.087 ± 0.045	0.222 ± 0.0014	49.596 ± 0.408
WFR-FM	0.131 ± 0.0005	9.143 ± 0.084	0.238 ± 0.0002	51.933 ± 0.257
SF^2M	0.196 ± 0.0004	9.087 ± 0.141	0.251 ± 0.0012	43.181 ± 0.475
VGFM	0.152 ± 0.0095	13.918 ± 1.996	0.220 ± 0.0054	49.371 ± 2.174
DeepRUOT	0.242 ± 0.0036	4.248 ± 0.963	0.268 ± 0.0079	32.350 ± 2.864
WLF (OT)	0.123 ± 0.0027	9.588 ± 0.494	0.235 ± 0.0035	54.389 ± 1.333
Modified WLF (with potentials)	0.257 ± 0.122	42.752 ± 35.190	0.817 ± 0.662	246.689 ± 195.446
Modified CytoBridge (with potentials)	0.119 ± 0.062	10.789 ± 9.712	0.156 ± 0.199	38.837 ± 65.635
GSBM (without potentials)	0.179 ± 0.0007	7.061 ± 0.167	0.249 ± 0.0034	42.089 ± 1.914
GSBM (with potentials)	0.163 ± 0.0004	2.103 ± 0.172	0.183 ± 0.0014	14.298 ± 0.507
CCB (without potentials, PF-ODE)	0.157 ± 0.0003	10.163 ± 0.021	0.217 ± 0.0051	47.574 ± 2.049
CCB (external potential only, PF-ODE)	–	–	0.083 ± 0.0033	5.870 ± 0.707
CCB (pairwise potential only, PF-ODE)	–	–	0.217 ± 0.0071	46.179 ± 2.765
CCB (full model, PF-ODE)	$\mathbf{0.117 \pm 0.0004}$	1.791 ± 0.031	$\mathbf{0.079 \pm 0.011}$	5.169 ± 0.522
CCB (full model, SDE)	0.168 ± 0.0016	$\mathbf{1.229 \pm 0.038}$	0.146 ± 0.0003	$\mathbf{3.934 \pm 0.035}$

Spatial Transcriptomics Interpolation Appropriate prior potentials can improve CCB interpolation even when the true potentials are unknown. To test this, we evaluate CCB on the axolotl telencephalon-regeneration spatial transcriptomics data from Wei et al (Wei et al., 2022). Models are trained on endpoint snapshots from $0.2 \to 1$, $1 \to 2$, and $2 \to 6$, with times 0.5, 1.5, and 3 held out, respectively; dataset details are in Appendix B.4.2. We use the spatial deformation potential $U_t = -\frac{\lambda}{4N(N-1)} \sum_{i \neq j} \left(\|s_i(t) - s_j(t)\|^2 - \|s_i(t_0) - s_j(t_0)\|^2 \right)^2$ as a regularization prior

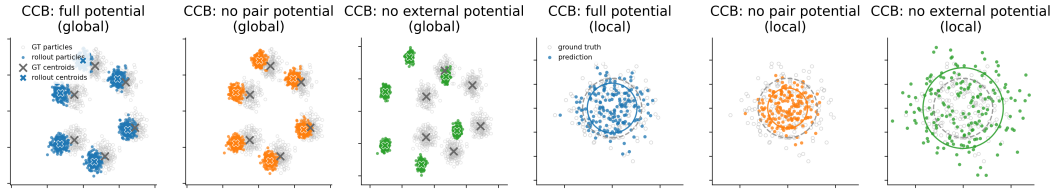

Figure 2: Potential ablation on the 20-dimensional confined repulsive Coulomb gas at $t = 0.5$. The left panels show the global view and the right panels show a local view of one cluster.

to penalize changes in pairwise squared distances from the initial state. Here, $s_i(t)$ denotes cell i's spatial coordinates at time t, t_0 the start of each interval, and N the number of cells. Because it also references the initial configuration, this potential differs slightly from the state-only potentials in the main formulation; Appendix B.5 gives the corresponding extension. Table 4 reports mean spatial W_1, expression W_1, and cell-type neighborhood TV under PF-ODE inference; metric definitions are in Appendix B.2.3. The potential improves all three metrics over the no-potential variant and achieves the lowest spatial W_1 and cell-type neighborhood TV. Figure 3 further shows that CCB better preserves tissue topology and produces clearer cell-type boundaries.

Table 4: Mean interpolation error across held-out spatial transcriptomics slices.

Method / Inference	**Spatial $W_1 \downarrow$**	**Exp. $W_1 \downarrow$**	**Neighb. TV $\downarrow$**
WFR-FM	0.258	5.064	0.485
SF^{2}M	0.263	5.057	0.497
WLF (OT)	0.271	5.064	0.572
OT-CFM	0.275	4.979	0.563
VGFM	0.302	4.849	0.536
DeepRUOT	0.325	**4.771**	0.513
GSBM	0.263	5.051	0.553
CCB (without potentials, PF-ODE)	0.302	4.873	0.565
CCB (spatial potential, PF-ODE)	**0.256**	4.826	**0.443**

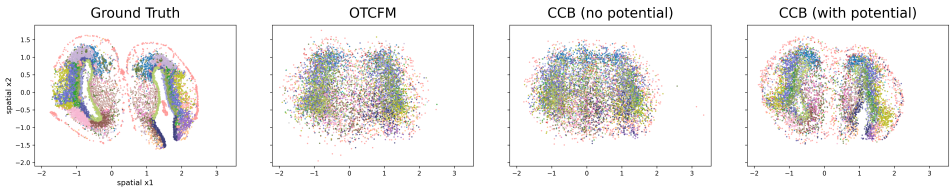

Figure 3: Axolotl interpolation at held-out $t = 0.5$: ground truth, OT-CFM, and CCB without/with the spatial potential; colors denote classifier-assigned cell types.

5 Conclusion and Limitations

CCB solves stochastic Wasserstein Lagrangian Flow with pairwise interaction potentials by jointly optimizing Conditional Probability Paths and the Transport Plan. Unlike Mean Field Game and Optimal Control solvers that update couplings through ReFlow, CCB avoids repeated SDE rollouts. Experiments show that CCB accurately recovers optimal controls, reconstructs marginal distributions, and improves interpolation when potentials are incorporated.

As with other Flow Matching methods for Optimal Control, full OT is computationally prohibitive on large datasets. Sparse mode is an effective alternative to Minibatch-OT, but its heuristic coupling needs improvement. Gaussian Conditional Probability Paths may limit expressiveness under complex external potentials. CCB also lacks a path-space interpretation of the action functionals used by Zhang et al. (2025b) or Peng et al. (2026a) and thus cannot handle unbalanced distributions. Finally, CCB minimizes action in a two-point boundary-value problem. Joint potential estimation would cause reward hacking; when needed, the method of Guan et al. (2026) can be used instead.

AI USE STATEMENT

We used AI tools to assist with code development and literature searches, including identifying references used in the proof of Proposition 2. AI tools also assisted in checking mathematical proofs. We also used AI tools for translation and grammar checking during manuscript preparation. We did not use AI tools to formulate the theory, nor write the main mathematical results. All AI-assisted content was independently checked by the authors, and all final scientific and editorial decisions were made by the authors. We take full responsibility for the final content of this work.

ETHICS STATEMENT

This work develops a computational method for reconstructing cell-population dynamics from snapshot data. Our experiments use synthetic datasets and previously published, publicly available single-cell and spatial transcriptomics datasets; no new human-subject studies or animal experiments were conducted. The data sources and preprocessing procedures are described in the appendix. Inferred trajectories depend on modeling assumptions and the coverage of the observed data, and should not be interpreted as direct evidence of individual cell lineages or causal mechanisms without independent validation. The method is intended for research use and has not been validated for clinical decision-making.

REPRODUCIBILITY STATEMENT

Our code is available in an anonymous repository at `https://anonymous.4open.science/r/CCB_under_review-7F4B`. The appendices provide detailed proofs, experimental settings, data-generation and preprocessing procedures, and algorithm pseudocode to support reproducibility.

REFERENCES

Michael Albergo, Nicholas Boffi, Michael Lindsey, and Eric Vanden-Eijnden. Multimarginal generative modeling with stochastic interpolants. In *International Conference on Learning Representations*, volume 2024, pp. 55884–55901, 2024.

Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023.

Lazar Atanackovic, Xi (Nicole) Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio, Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the wasserstein manifold. In *International Conference on Learning Representations*, volume 2025, pp. 94586–94610, 2025.

Julio Backhoff, Giovanni Conforti, Ivan Gentil, and Christian Léonard. The mean field schrödinger problem: ergodic behavior, entropy estimates and functional inequalities. *Probability Theory and Related Fields*, 178(1-2):475–530, June 2020. ISSN 1432-2064. doi: 10.1007/s00440-020-00977-8.

Jean-David Benamou. Numerical resolution of an "unbalanced" mass transport problem. *ESAIM: Mathematical Modelling and Numerical Analysis*, 37(5):851–868, Sept 2003. ISSN 1290-3841. doi: 10.1051/m2an:2003058.

Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, Jan 2000. ISSN 0945-3245. doi: 10.1007/s002110050002.

Volker Bergen, Marius Lange, Stefan Peidli, F. Alexander Wolf, and Fabian J. Theis. Generalizing rna velocity to transient cell states through dynamical modeling. *Nature Biotechnology*, 38(12):1408–1414, Aug 2020. ISSN 1546-1696. doi: 10.1038/s41587-020-0591-3.

Jérôme Bolte, Shoham Sabach, Marc Teboulle, and Yakov Vaisbourd. First order methods beyond convexity and lipschitz gradient continuity with applications to quadratic inverse problems. *SIAM Journal on Optimization*, 28(3):2131–2151, Jan 2018. ISSN 1095-7189. doi: 10.1137/17m1138558.

Charlotte Bunne, Ya-Ping Hsieh, Marco Cuturi, and Andreas Krause. The schrödinger bridge between gaussian measures has a closed form. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 5802–5833. PMLR, 25–27 Apr 2023a.

Charlotte Bunne, Stefan G. Stark, Gabriele Gut, Jacobo Sarabia del Castillo, Mitch Levesque, Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*, 20(11):1759–1768, 2023b. doi: 10.1038/s41592-023-01969-x.

Charlotte Bunne, Geoffrey Schiebinger, Andreas Krause, Aviv Regev, and Marco Cuturi. Optimal transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1):58, Aug 2024. ISSN 2662-8449. doi: 10.1038/s43586-024-00334-2.

Ao Chen, Sha Liao, Mengnan Cheng, Kailong Ma, Liang Wu, Yiwei Lai, Xiaojie Qiu, Jin Yang, Jiangshan Xu, Shijie Hao, Xin Wang, Huifang Lu, Xi Chen, Xing Liu, Xin Huang, Zhao Li, Yan Hong, Yujia Jiang, Jian Peng, Shuai Liu, Mengzhe Shen, Chuanyu Liu, Quanshui Li, Yue Yuan, Xiaoyu Wei, Huiwen Zheng, Weimin Feng, Zhifeng Wang, Yang Liu, Zhaohui Wang, Yunzhi Yang, Haitao Xiang, Lei Han, Baoming Qin, Pengcheng Guo, Guangyao Lai, Pura Muñoz-Cánoves, Patrick H. Maxwell, Jean Paul Thiery, Qing-Feng Wu, Fuxiang Zhao, Bichao Chen, Mei Li, Xi Dai, Shuai Wang, Haoyan Kuang, Junhou Hui, Liqun Wang, Ji-Feng Fei, Ou Wang, Xiaofeng Wei, Haorong Lu, Bo Wang, Shiping Liu, Ying Gu, Ming Ni, Wenwei Zhang, Feng Mu, Ye Yin, Huanming Yang, Michael Lisby, Richard J. Cornall, Jan Mulder, Mathias Uhlén, Miguel A. Esteban, Yuxiang Li, Longqi Liu, Xun Xu, and Jian Wang. Spatiotemporal transcriptomic atlas of mouse organogenesis using dna nanoball-patterned arrays. *Cell*, 185(10): 1777–1792.e21, May 2022. ISSN 0092-8674. doi: 10.1016/j.cell.2022.04.003.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. On the relation between optimal transport and schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169(2):671–691, 2016. doi: 10.1007/s10957-015-0803-z.

Lénaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. An interpolating distance between optimal transport and fisher–rao metrics. *Foundations of Computational Mathematics*, 18(1):1–44, Feb 2018a. ISSN 1615-3383. doi: 10.1007/s10208-016-9331-y.

Lénaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal transport: Dynamic and kantorovich formulations. *Journal of Functional Analysis*, 274(11): 3090–3123, June 2018b. ISSN 0022-1236. doi: 10.1016/j.jfa.2018.03.008.

David P. Cook and Barbara C. Vanderhyden. Context specificity of the emt transcriptional response. *Nature Communications*, 11(1):2142, May 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-16066-2.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems*, volume 34, pp. 17695–17709. Curran Associates, Inc., 2021.

Jun Ding, Nadav Sharon, and Ziv Bar-Joseph. Temporal modelling using single-cell transcriptomics. *Nature Reviews Genetics*, 23(6):355–368, Jan 2022. ISSN 1471-0064. doi: 10.1038/s41576-021-00444-7.

Shukai Du, Junzhe Zhang, and Yiming Li. Lagrangian flow matching: A least-action framework for principled path design. arXiv preprint arXiv:2605.15419, 2026.

Luca Eyring, Dominik Klein, Théo Uscidda, Giovanni Palla, Niki Kilbertus, Zeynep Akata, and Fabian Theis. Unbalancedness in neural monge maps improves unpaired domain translation. In *International Conference on Learning Representations*, volume 2024, pp. 383–410, 2024.

Aden Forrow and Geoffrey Schiebinger. Lineageot is a unified framework for lineage tracing and trajectory inference. *Nature Communications*, 12(1):4940, Aug 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-25133-1.

Anming Gu, Edward Chien, and Kristjan Greenewald. Partially observed trajectory inference using optimal transport and a dynamics prior. In *International Conference on Learning Representations*, volume 2025, pp. 42694–42728, 2025.

Vincent Guan, Lazar Atanackovic, and Kirill Neklyudov. A call to lagrangian action: Learning population mechanics from temporal snapshots. In *Forty-third International Conference on Machine Learning*, 2026.

Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth schrödinger bridges. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 23722–23750. PMLR, 13–19 Jul 2025.

Guillaume Huguet, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. In *Advances in Neural Information Processing Systems*, volume 35, pp. 29705–29718. Curran Associates, Inc., 2022. doi: 10.52202/068431-2154.

Kacper Kapuśniak, Peter Potaptchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Avishek Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. In *Advances in Neural Information Processing Systems*, volume 37, pp. 135011–135042. Curran Associates, Inc., 2024. doi: 10.52202/079017-4291.

Evelyn F. Keller and Lee A. Segel. Initiation of slime mold aggregation viewed as an instability. *Journal of Theoretical Biology*, 26(3):399–415, Mar 1970. ISSN 0022-5193. doi: 10.1016/0022-5193(70)90092-5.

Dominik Klein, Théo Uscidda, Fabian Theis, and Marco Cuturi. Genot: Entropic (gromov) wasserstein flow matching with applications to single-cell genomics. In *Advances in Neural Information Processing Systems*, volume 37, pp. 103897–103944. Curran Associates, Inc., 2024. doi: 10.52202/079017-3301.

Dominik Klein, Giovanni Palla, Marius Lange, Michal Klein, Zoe Piran, Manuel Gander, Laetitia Meng-Papaxanthos, Michael Sterr, Lama Saber, Changying Jing, Aimée Bastidas-Ponce, Perla Cota, Marta Tarquis-Medina, Shrey Parikh, Ilan Gold, Heiko Lickert, Mostafa Bakhti, Mor Nitzan, Marco Cuturi, and Fabian J. Theis. Mapping cells through time and space with moscot. *Nature*, 638(8052):1065–1075, Jan 2025. ISSN 1476-4687. doi: 10.1038/s41586-024-08453-2.

Stanislav Kondratyev, Léonard Monsaingeon, and Dmitry Vorotnikov. A new optimal transport distance on the space of finite radon measures. *Advances in Differential Equations*, 21(11/12):1117–1164, Nov 2016. ISSN 1079-9389. doi: 10.57262/ade/1476369298.

Takeshi Koshizuka and Issei Sato. Neural lagrangian schrödinger bridge: Diffusion modeling for population dynamics. In *The Eleventh International Conference on Learning Representations*, 2023.

Gioele La Manno, Ruslan Soldatov, Amit Zeisel, Emelie Braun, Hannah Hochgerner, Viktor Petukhov, Katja Lidschreiber, Maria E. Kastriti, Peter Lönnerberg, Alessandro Furlan, Jean Fan, Lars E. Borm, Zehua Liu, David van Bruggen, Jimin Guo, Xiaoling He, Roger Barker, Erik Sundström, Gonçalo Castelo-Branco, Patrick Cramer, Igor Adameyko, Sten Linnarsson, and Peter V. Kharchenko. Rna velocity of single cells. *Nature*, 560(7719):494–498, Aug 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0414-6.

Marius Lange, Volker Bergen, Michal Klein, Manu Setty, Bernhard Reuter, Mostafa Bakhti, Heiko Lickert, Meshal Ansari, Janine Schniering, Herbert B. Schiller, Dana Pe'er, and Fabian J. Theis. Cellrank for directed single-cell fate mapping. *Nature Methods*, 19(2):159–170, Jan 2022. ISSN 1548-7105. doi: 10.1038/s41592-021-01346-6.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, Mar 2007. ISSN 1861-3624. doi: 10.1007/s11537-007-0657-8.

Matthias Liero, Alexander Mielke, and Giuseppe Savaré. Optimal transport in competition with reaction: The hellinger–kantorovich distance and geodesic curves. *SIAM Journal on Mathematical Analysis*, 48(4):2869–2911, Jan 2016. ISSN 1095-7154. doi: 10.1137/15m1041420.

Matthias Liero, Alexander Mielke, and Giuseppe Savaré. Optimal entropy-transport problems and a new hellinger–kantorovich distance between positive measures. *Inventiones mathematicae*, 211 (3):969–1117, Mar 2018. ISSN 1432-1297. doi: 10.1007/s00222-017-0759-8.

Alex Tong Lin, Samy Wu Fung, Wuchen Li, Levon Nurbekyan, and Stanley J. Osher. Alternating the population and control neural networks to solve high-dimensional stochastic mean-field games. *Proceedings of the National Academy of Sciences*, 118(31):e2024713118, Aug 2021. doi: 10.1073/pnas.2024713118.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos Theodorou. Deep generalized schrödinger bridge. In *Advances in Neural Information Processing Systems*, volume 35, pp. 9374–9388. Curran Associates, Inc., 2022. doi: 10.52202/068431-0681.

Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, and Ricky T. Q. Chen. Generalized schrödinger bridge matching. In *International Conference on Learning Representations*, volume 2024, pp. 14527–14552, 2024.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023.

Suryanarayana Maddu, Victor Chardès, and Michael J Shelley. Learning stochastic processes with intrinsic noise from cross-sectional biological data. *Proceedings of the National Academy of Sciences*, 122(37):e2420621122, 2025. doi: 10.1073/pnas.2420621122.

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, Dec 2019. ISSN 1546-1696. doi: 10.1038/s41587-019-0336-3.

Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A computational framework for solving Wasserstein lagrangian flows. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 37461–37485. PMLR, 21–27 Jul 2024.

Qiangwei Peng, Zihan Wang, Junda Ying, Yuhao Sun, Qing Nie, Lei Zhang, Tiejun Li, and Peijie Zhou. Wfr-fm: Simulation-free dynamic unbalanced optimal transport. In *International Conference on Learning Representations*, volume 2026, pp. 18943–18976, 2026a.

Qiangwei Peng, Peijie Zhou, and Tiejun Li. stvcr: spatiotemporal dynamics of single cells. *Nature Methods*, 23(3):542–553, Mar 2026b. ISSN 1548-7105. doi: 10.1038/s41592-026-03010-3.

Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky T. Q. Chen. Multisample flow matching: Straightening flows with minibatch couplings. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 28100–28127. PMLR, 23–29 Jul 2023.

Martin Rohbeck, Edward De Brouwer, Charlotte Bunne, Jan-Christian Huetter, Anne Biton, Kelvin Chen, Aviv Regev, and Romain Lopez. Modeling complex system dynamics with flow matching across time and conditions. In *International Conference on Learning Representations*, volume 2025, pp. 12164–12188, 2025.

Lars Ruthotto, Stanley J. Osher, Wuchen Li, Levon Nurbekyan, and Samy Wu Fung. A machine learning framework for solving high-dimensional mean field game and mean field control problems. *Proceedings of the National Academy of Sciences*, 117(17):9183–9193, Apr 2020. doi: 10.1073/pnas.1922204117.

Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, Lia Lee, Jenny Chen, Justin Brumbaugh, Philippe Rigollet, Konrad Hochedlinger, Rudolf Jaenisch, Aviv Regev, and Eric S. Lander. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943.e22, Feb 2019. ISSN 0092-8674. doi: 10.1016/j.cell.2019.01.006.

Yutong Sha, Yuchi Qiu, Peijie Zhou, and Qing Nie. Reconstructing growth and dynamic trajectories from single-cell transcriptomics data. *Nature Machine Intelligence*, 6(1):25–39, Jan 2024. ISSN 2522-5839. doi: 10.1038/s42256-023-00763-w.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. In *Advances in Neural Information Processing Systems*, volume 36, pp. 62183–62223. Curran Associates, Inc., 2023. doi: 10.52202/075280-2717.

Marlon Stoeckius, Christoph Hafemeister, William Stephenson, Brian Houck-Loomis, Pratip K Chattopadhyay, Harold Swerdlow, Rahul Satija, and Peter Smibert. Simultaneous epitope and transcriptome measurement in single cells. *Nature Methods*, 14(9):865–868, July 2017. ISSN 1548-7105. doi: 10.1038/nmeth.4380.

Yuhao Sun, Zhenyi Zhang, Zihan Wang, Tiejun Li, and Peijie Zhou. Variational regularized unbalanced optimal transport: Single network, least action. In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 16963–17005. Curran Associates, Inc., 2025. doi: 10.52202/085713-0572.

Yuhao Sun, Zekun Wu, Zixun Huang, and Peijie Zhou. Tracingflow: A simulation-free trajectory inference framework based on second-order dynamics. arXiv preprint arXiv:2608.21070, 2026.

Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. TrajectoryNet: A dynamic optimal transport network for modeling cellular dynamics. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9526–9536. PMLR, 13–18 Jul 2020.

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. *Transactions on Machine Learning Research*, 2024a. ISSN 2835-8856. Expert Certification.

Alexander Y. Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and flow matching. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 1279–1287. PMLR, 02–04 May 2024b.

Wei Wan, Yuejin Zhang, Chenglong Bao, Bin Dong, and Zuoqiang Shi. A scalable deep learning approach for solving high-dimensional dynamic optimal transport. *SIAM Journal on Scientific Computing*, 45(4):B544–B563, Aug 2023. ISSN 1095-7197. doi: 10.1137/22m1490831.

Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-growth flow matching for single-cell dynamics modeling. In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 160552–160588. Curran Associates, Inc., 2025. doi: 10.52202/085713-5363.

Xiaoyu Wei, Sulei Fu, Hanbo Li, Yang Liu, Shuai Wang, Weimin Feng, Yunzhi Yang, Xiawei Liu, Yan-Yun Zeng, Mengnan Cheng, Yiwei Lai, Xiaojie Qiu, Liang Wu, Nannan Zhang, Yujia Jiang, Jiangshan Xu, Xiaoshan Su, Cheng Peng, Lei Han, Wilson Pak-Kin Lou, Chuanyu Liu, Yue Yuan, Kailong Ma, Tao Yang, Xiangyu Pan, Shang Gao, Ao Chen, Miguel A. Esteban, Huanming Yang,

Jian Wang, Guangyi Fan, Longqi Liu, Liang Chen, Xun Xu, Ji-Feng Fei, and Ying Gu. Single-cell stereo-seq reveals induced progenitor cells involved in axolotl brain regeneration. *Science*, 377 (6610):eabp9444, Sept 2022. ISSN 1095-9203. doi: 10.1126/science.abp9444.

Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D. Camargo, and Allon M. Klein. Lineage tracing on transcriptional landscapes links state to fate during differentiation. *Science*, 367(6479): eaaw3381, Feb 2020. ISSN 1095-9203. doi: 10.1126/science.aaw3381.

Grace Hui Ting Yeo, Sachit D. Saksena, and David K. Gifford. Generative modeling of single-cell time series with prescient enables prediction of cell trajectories with interventions. *Nature Communications*, 12(1):3222, May 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-23518-w.

Junda Ying, Yuxuan Wang, Bowen Yang, Peijie Zhou, and Lei Zhang. Beyond continuity: Simulation-free reconstruction of discrete branching dynamics from single-cell snapshots. In *Forty-third International Conference on Machine Learning*, 2026a.

Junda Ying, Yuxuan Wang, Bowen Yang, Peijie Zhou, and Lei Zhang. Simulation-free unbalanced dynamic optimal transport with general growth penalty. arXiv preprint arXiv:2609.04710, 2026b.

Jiajia Yu, Junghwan Lee, Yao Xie, and Xiuyuan Cheng. High-dimensional mean-field games by particle-based flow matching. In *International Conference on Learning Representations*, volume 2026, pp. 119161–119185, 2026.

Jiaqi Zhang, Erica Larschan, Jeremy Bigness, and Ritambhara Singh. scnode: generative model for temporal single cell transcriptomic data prediction. *Bioinformatics*, 40(Supplement 2):ii146–ii154, September 2024. ISSN 1367-4811. doi: 10.1093/bioinformatics/btae393.

Stephen Zhang, Suryanarayana Maddu, Xiaojie Qiu, and Victor Chardès. Inferring stochastic dynamics with growth from cross-sectional data. In *Advances in Neural Information Processing Systems*, volume 38, pp. 126512–126545. Curran Associates, Inc., 2025a. doi: 10.52202/085713-4219.

Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through regularized unbalanced optimal transport. In *International Conference on Learning Representations*, volume 2025, pp. 19888–19919, 2025b.

Zhenyi Zhang, Yuhao Sun, Qiangwei Peng, Tiejun Li, and Peijie Zhou. Integrating dynamical systems modeling with spatiotemporal scrna-seq data analysis. *Entropy*, 27(5):453, Apr 2025c. ISSN 1099-4300. doi: 10.3390/e27050453.

Zhenyi Zhang, Zihan Wang, Yuhao Sun, Tiejun Li, and Peijie Zhou. Modeling cell dynamics and interactions with unbalanced mean field schrödinger bridge. In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 139518–139558. Curran Associates, Inc., 2025d. doi: 10.52202/085713-4659.

Grace X. Y. Zheng, Jessica M. Terry, Phillip Belgrader, Paul Ryvkin, Zachary W. Bent, Ryan Wilson, Solongo B. Ziraldo, Tobias D. Wheeler, Geoff P. McDermott, Junjie Zhu, Mark T. Gregory, Joe Shuga, Luz Montesclaros, Jason G. Underwood, Donald A. Masquelier, Stefanie Y. Nishimura, Michael Schnall-Levin, Paul W. Wyatt, Christopher M. Hindson, Rajiv Bharadwaj, Alexander Wong, Kevin D. Ness, Lan W. Beppu, H. Joachim Deeg, Christopher McFarland, Keith R. Loeb, William J. Valente, Nolan G. Ericson, Emily A. Stevens, Jerald P. Radich, Tarjei S. Mikkelsen, Benjamin J. Hindson, and Jason H. Bielas. Massively parallel digital transcriptional profiling of single cells. *Nature Communications*, 8(1):14049, Jan 2017. ISSN 2041-1723. doi: 10.1038/ncomms14049.

A Derivations and Technical Details

A.1 Derivation of Proposition 1

Proposition 1 (Action decomposition via endpoint conditioning)**.** Let $\mathbb{P}$ be the path measure of the controlled SDE, $\pi \in \Pi(\mu_0, \mu_1)$ its endpoint transport plan, $\mathbb{P}^z$ the conditional path measure given $z = (x_0, x_1)$, and $\mathbb{Q}^{\nu,z}$ the Brownian bridge between the same endpoints. With the conditional potential costs C_V and C_W defined in the main text, the IS-WLF action decomposes as

$$\begin{aligned}\mathcal{A}(\rho, b) \doteq\ & 2\nu \mathrm{KL}(\pi \mid \mu_0 \otimes \mu_1) \\ & + \int \left[\frac{1}{2}\|x_1 - x_0\|^2 + 2\nu \mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) - C_V(z)\right] \pi(\mathrm{d}z) \\ & - \frac{1}{2} \iint C_W(z, z')\pi(\mathrm{d}z)\pi(\mathrm{d}z').\end{aligned}$$

Derivation. Let $\mathbb{P}$ be the path measure of the controlled process $\mathrm{d}X_t = b_t(X_t)\,\mathrm{d}t + \sqrt{2\nu}\,\mathrm{d}B_t$, and let $\mathbb{Q}^\nu$ be the zero-drift Brownian reference measure with the same initial distribution μ_0 and diffusion coefficient. All spatial integrals below are taken over $\mathbb{R}^d$. By Girsanov's formula, the kinetic energy in the action satisfies

$$K(\rho, b) := \frac{1}{2}\int_0^1 \int \|b_t(x)\|^2 \rho_t(x)\,\mathrm{d}x\,\mathrm{d}t = 2\nu \mathrm{KL}(\mathbb{P} \mid \mathbb{Q}^\nu).$$

Let $z = (x_0, x_1)$, let π be the endpoint transport plan of $\mathbb{P}$, and let $\mathbb{P}^z$ and $\mathbb{Q}^{\nu,z}$ denote the conditional path measure and Brownian bridge given endpoints z, respectively. Conditioning on the endpoints and applying the chain rule for KL divergence gives

$$\mathrm{KL}(\mathbb{P} \mid \mathbb{Q}^\nu) = \mathrm{KL}(\pi \mid \mathbb{Q}^\nu_{01}) + \int \mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z})\pi(\mathrm{d}z),$$

where $\mathbb{Q}^\nu_{01}$ is the endpoint joint distribution of the reference process, whose conditional transition density is

$$q_\nu(x_1 \mid x_0) = (4\pi\nu)^{-d/2} \exp\left(-\frac{\|x_1 - x_0\|^2}{4\nu}\right).$$

Substituting this heat kernel into the endpoint KL and using the fact that the two marginals of π are fixed as μ_0 and μ_1 yields

$$2\nu \mathrm{KL}(\pi \mid \mathbb{Q}^\nu_{01}) \doteq 2\nu \mathrm{KL}(\pi \mid \mu_0 \otimes \mu_1) + \frac{1}{2}\int \|x_1 - x_0\|^2 \pi(\mathrm{d}z).$$

The omitted terms arise from the normalization factor of the heat kernel and the entropy of the fixed target distribution, neither of which depends on the endpoint transport plan or the conditional paths.

We next conditionally decompose the potential energies. Let $\rho^z_t(x)$ denote the marginal density at time t of the conditional path $\mathbb{P}^z$. By the law of total probability,

$$\rho_t(x) = \int \rho^z_t(x)\pi(\mathrm{d}z).$$

Substituting this into the external potential energy and exchanging the order of integration gives

$$\begin{aligned}\int_0^1 \int V(x)\rho_t(x)\,\mathrm{d}x\,\mathrm{d}t &= \int_0^1 \int V(x)\left[\int \rho^z_t(x)\pi(\mathrm{d}z)\right]\mathrm{d}x\,\mathrm{d}t \\ &= \int \underbrace{\left[\int_0^1 \int V(x)\rho^z_t(x)\,\mathrm{d}x\,\mathrm{d}t\right]}_{C_V(z)} \pi(\mathrm{d}z).\end{aligned}$$

Thus, $C_V(z)$ is the conditional external potential cost for the endpoint pair z. For the pairwise potential energy, let $z' = (y_0, y_1)$ and condition the two density factors separately, giving

$$\rho_t(x)\rho_t(y) = \iint \rho^z_t(x)\rho^{z'}_t(y)\pi(\mathrm{d}z)\pi(\mathrm{d}z').$$

Substituting this into the pairwise potential energy and exchanging the order of integration gives

$$\frac{1}{2}\int_0^1 \iint W(x,y)\rho_t(x)\rho_t(y)\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}t = \frac{1}{2}\iint \underbrace{\left[\int_0^1 \iint W(x,y)\rho_t^z(x)\rho_t^{z'}(y)\,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}t\right]}_{C_W(z,z')} \pi(\mathrm{d}z)\pi(\mathrm{d}z').$$

Here $C_W(z,z')$ is the conditional interaction cost between two endpoint pairs. Substituting the decompositions of the kinetic energy, external potential energy, and pairwise potential energy back into the original action yields

$$\begin{aligned}\mathcal{A}(\rho,b) \doteq\ & 2\nu\mathrm{KL}(\pi \mid \mu_0\otimes\mu_1)\\ &+\int\left[\frac{1}{2}\|x_1-x_0\|^2 + 2\nu\mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) - C_V(z)\right]\pi(\mathrm{d}z)\\ &-\frac{1}{2}\iint C_W(z,z')\pi(\mathrm{d}z)\pi(\mathrm{d}z').\end{aligned}$$

The notation $\doteq$ indicates equality up to a constant that depends only on the fixed endpoint distributions and diffusion coefficient and therefore does not affect optimization.

□

A.2 Alternative Derivation of the Conditional-Bridge Kinetic Cost

The conditional-bridge kinetic cost used in GSBM (Liu et al., 2024), $\frac{1}{2}\int_0^1 \mathbb{E}_{\mathbb{P}^z}\|b_t^z(X_t)\|^2\,\mathrm{d}t$, diverges under exact endpoint conditioning with nonzero diffusion. We recover a finite cost by subtracting its logarithmic divergence. The result below shows that this renormalized cost agrees with the conditional kinetic cost used by CCB, up to a fixed constant. For fixed endpoints $z=(x_0,x_1)$, let $\mathbb{P}^z$ denote a conditional bridge with diffusion coefficient $\sqrt{2\nu}$, $\nu>0$, and let $\mathbb{Q}^{\nu,z}$ denote the Brownian bridge with the same endpoints and diffusion coefficient.

Proposition (Renormalized conditional-bridge kinetic cost)**.** Consider a conditional bridge $\mathbb{P}^z$ with finite relative entropy $\mathrm{KL}(\mathbb{P}^z\|\mathbb{Q}^{\nu,z}) < \infty$ and the same leading-order mean-square behavior as the Brownian bridge near the final endpoint:

$$\lim_{\epsilon\downarrow 0}\frac{\mathbb{E}_{\mathbb{P}^z}\|X_{1-\epsilon}-x_1\|^2}{\epsilon} = 2\nu d.$$

The truncated cost $K_\epsilon^z = \frac{1}{2}\int_0^{1-\epsilon}\mathbb{E}_{\mathbb{P}^z}\|b_t^z(X_t)\|^2\,\mathrm{d}t$ diverges as $\epsilon\downarrow 0$, with logarithmic term

$$\mathrm{CT}_\epsilon = \nu d\log\frac{1}{\epsilon}.$$

Subtracting this term gives a finite renormalized cost:

$$K_{\mathrm{ren}}^z := \lim_{\epsilon\downarrow 0}\left(K_\epsilon^z - \mathrm{CT}_\epsilon\right) = \frac{1}{2}\|x_1-x_0\|^2 + 2\nu\mathrm{KL}(\mathbb{P}^z\|\mathbb{Q}^{\nu,z}) - \nu d.$$

Thus, after discarding the fixed constant $-\nu d$, the renormalized result agrees with the squared endpoint-distance and conditional path-space KL terms used by CCB.

Proof. Write the reference Brownian-bridge drift as $\beta_t^z(x) = (x_1-x)/(1-t)$, and let $R_t^z = \mathbb{E}_{\mathbb{P}^z}\|X_t-x_1\|^2$. Unless stated otherwise, all expectations below are taken under $\mathbb{P}^z$.

For $f(x)=\|x-x_1\|^2$, we have $\nabla f(x) = 2(x-x_1)$ and $\nabla^2 f(x) = 2I_d$. Itô's formula gives

$$\mathrm{d}\|X_t-x_1\|^2 = \left[2(X_t-x_1)^\top b_t^z(X_t) + 2\nu d\right]\mathrm{d}t + 2\sqrt{2\nu}(X_t-x_1)^\top \mathrm{d}B_t.$$

The second-order correction is $\frac{1}{2}\mathrm{tr}[(2\nu I_d)(2I_d)] = 2\nu d$. Taking expectations eliminates the stochastic-integral term and yields

$$\dot{R}_t^z = 2\mathbb{E}\left[(X_t-x_1)^\top b_t^z(X_t)\right] + 2\nu d,$$

or equivalently,

$$\mathbb{E}\left[(X_t - x_1)^\top b_t^z(X_t)\right] = \frac{1}{2}\dot{R}_t^z - \nu d.$$

Because $b_t^z(x) - \beta_t^z(x) = b_t^z(x) + (x - x_1)/(1-t)$, expanding the square and taking expectations gives

$$\frac{1}{2}\mathbb{E}\|b_t^z(X_t) - \beta_t^z(X_t)\|^2 = \frac{1}{2}\mathbb{E}\|b_t^z(X_t)\|^2 + \frac{\mathbb{E}[(X_t - x_1)^\top b_t^z(X_t)]}{1-t} + \frac{R_t^z}{2(1-t)^2}.$$

Substitution of the cross term gives

$$\frac{1}{2}\mathbb{E}\|b_t^z(X_t) - \beta_t^z(X_t)\|^2 = \frac{1}{2}\mathbb{E}\|b_t^z(X_t)\|^2 + \frac{\dot{R}_t^z}{2(1-t)} - \frac{\nu d}{1-t} + \frac{R_t^z}{2(1-t)^2}.$$

Using

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{R_t^z}{1-t}\right) = \frac{\dot{R}_t^z}{1-t} + \frac{R_t^z}{(1-t)^2},$$

we rearrange the preceding identity as

$$\frac{1}{2}\mathbb{E}\|b_t^z(X_t)\|^2 = \frac{1}{2}\mathbb{E}\|b_t^z(X_t) - \beta_t^z(X_t)\|^2 + \frac{\nu d}{1-t} - \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{R_t^z}{1-t}\right).$$

Define the relative control energy $D_\epsilon^z = \frac{1}{2}\int_0^{1-\epsilon}\mathbb{E}\|b_t^z(X_t) - \beta_t^z(X_t)\|^2\,\mathrm{d}t$. Integrating the preceding equality and using

$$\int_0^{1-\epsilon}\frac{\nu d}{1-t}\,\mathrm{d}t = \nu d\log\frac{1}{\epsilon}$$

and

$$-\frac{1}{2}\int_0^{1-\epsilon}\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{R_t^z}{1-t}\right)\mathrm{d}t = -\frac{R_{1-\epsilon}^z}{2\epsilon} + \frac{R_0^z}{2},$$

we obtain, since $R_0^z = \|x_1 - x_0\|^2$, the exact identity

$$K_\epsilon^z = D_\epsilon^z + \nu d\log\frac{1}{\epsilon} + \frac{1}{2}\|x_1 - x_0\|^2 - \frac{R_{1-\epsilon}^z}{2\epsilon}.$$

It remains to relate the relative control energy to the path-space KL divergence. Let $T_\epsilon = 1 - \epsilon$, let $\mathbb{P}_\epsilon^z$ and $\mathbb{Q}_\epsilon^{\nu,z}$ be the restrictions of the two conditional bridges to $[0, T_\epsilon]$, and define

$$a_t^z = \frac{b_t^z(X_t) - \beta_t^z(X_t)}{\sqrt{2\nu}}.$$

Writing $B_t^{\mathbb{Q},z}$ for the driving Brownian motion of the reference bridge, Girsanov's theorem gives

$$\log\frac{\mathrm{d}\mathbb{P}_\epsilon^z}{\mathrm{d}\mathbb{Q}_\epsilon^{\nu,z}} = \int_0^{T_\epsilon}(a_t^z)^\top\mathrm{d}B_t^{\mathbb{Q},z} - \frac{1}{2}\int_0^{T_\epsilon}\|a_t^z\|^2\,\mathrm{d}t.$$

Under $\mathbb{P}^z$, the process $B_t^{\mathbb{P},z} = B_t^{\mathbb{Q},z} - \int_0^t a_s^z\,\mathrm{d}s$ is Brownian motion. Substituting $\mathrm{d}B_t^{\mathbb{Q},z} = \mathrm{d}B_t^{\mathbb{P},z} + a_t^z\,\mathrm{d}t$ yields

$$\log\frac{\mathrm{d}\mathbb{P}_\epsilon^z}{\mathrm{d}\mathbb{Q}_\epsilon^{\nu,z}} = \int_0^{T_\epsilon}(a_t^z)^\top\mathrm{d}B_t^{\mathbb{P},z} + \frac{1}{2}\int_0^{T_\epsilon}\|a_t^z\|^2\,\mathrm{d}t.$$

Taking expectations under $\mathbb{P}^z$ eliminates the stochastic-integral term, and hence

$$\mathrm{KL}(\mathbb{P}_\epsilon^z\|\mathbb{Q}_\epsilon^{\nu,z}) = \frac{1}{4\nu}\int_0^{1-\epsilon}\mathbb{E}\|b_t^z(X_t) - \beta_t^z(X_t)\|^2\,\mathrm{d}t = \frac{D_\epsilon^z}{2\nu}.$$

Letting $\epsilon \downarrow 0$ recovers the full time interval and gives

$$D_\epsilon^z \longrightarrow 2\nu\mathrm{KL}(\mathbb{P}^z\|\mathbb{Q}^{\nu,z}).$$

Together with the endpoint condition $R_{1-\epsilon}^z/(2\epsilon) \to \nu d$, this yields

$$K_\epsilon^z = \nu d\log\frac{1}{\epsilon} + \frac{1}{2}\|x_1 - x_0\|^2 + 2\nu\mathrm{KL}(\mathbb{P}^z\|\mathbb{Q}^{\nu,z}) - \nu d + o(1).$$

All terms except $\nu d\log(1/\epsilon)$ have finite limits, whereas the latter diverges to $+\infty$. Minimal subtraction removes only this logarithmically divergent term, leaving the finite renormalized cost stated in the proposition.

□

A.3 Derivation of Corollary 1

Corollary 1 (Static quadratic OT for empirical distributions)**.** For empirical marginals $\widehat{\mu}_0 = \sum_i a_i \delta_{x_0^i}$ and $\widehat{\mu}_1 = \sum_j b_j \delta_{x_1^j}$, and fixed conditional paths, removing the endpoint-dependent constant from Proposition 1 yields

$$\min_{P \in \mathcal{U}(a,b)} \sum_{i,j} c_{ij} P_{ij} - \frac{1}{2} \sum_{i,j,k,\ell} Q_{ij,k\ell} P_{ij} P_{k\ell} + 2\nu \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j}.$$

Derivation. Restrict the endpoint transport plan in Proposition 1 to the empirical support $z_{ij} = (x_0^i, x_1^j)$. Let $z = (x_0, x_1)$, and represent the transport plan using Dirac functions:

$$\pi_P(\mathrm{d}z) = \sum_{i,j} P_{ij} \delta(z - z_{ij})\, \mathrm{d}z, \qquad \delta(z - z_{ij}) = \delta(x_0 - x_0^i)\delta(x_1 - x_1^j).$$

Integrating over x_1 and x_0, respectively, gives the source and target marginals $\sum_i (\sum_j P_{ij}) \delta_{x_0^i}$ and $\sum_j (\sum_i P_{ij}) \delta_{x_1^j}$. Therefore, the endpoint constraints are equivalent to $\sum_j P_{ij} = a_i$ and $\sum_i P_{ij} = b_j$. Together with the nonnegativity constraint, this gives $P \in \mathcal{U}(a,b) := \{P \geq 0 : \ P\mathbf{1}_m = a,\ P^\top \mathbf{1}_n = b\}$.

With the family of conditional paths fixed, denote the single-path cost in Proposition 1 by $c(z) = \frac{1}{2}\|x_1 - x_0\|^2 + 2\nu \mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) - C_V(z)$, and let $c_{ij} = c(z_{ij})$. Substituting the Dirac representation of the transport plan and using its sifting property gives

$$\int c(z) \pi_P(\mathrm{d}z) = \sum_{i,j} P_{ij} \int c(z) \delta(z - z_{ij})\, \mathrm{d}z = \sum_{i,j} P_{ij} c(z_{ij}) = \sum_{i,j} c_{ij} P_{ij}.$$

For the pairwise interaction term, substitute the Dirac representation for both copies of the endpoint transport plan and integrate successively over z and z', yielding

$$\begin{aligned}
&\iint C_W(z, z') \pi_P(\mathrm{d}z) \pi_P(\mathrm{d}z') \\
&= \sum_{i,j,k,\ell} P_{ij} P_{k\ell} \iint C_W(z, z') \delta(z - z_{ij}) \delta(z' - z_{k\ell})\, \mathrm{d}z\, \mathrm{d}z' \\
&= \sum_{i,j,k,\ell} P_{ij} P_{k\ell} \int C_W(z_{ij}, z') \delta(z' - z_{k\ell})\, \mathrm{d}z' \\
&= \sum_{i,j,k,\ell} P_{ij} P_{k\ell} C_W(z_{ij}, z_{k\ell}).
\end{aligned}$$

Let $Q_{ij,k\ell} = C_W(z_{ij}, z_{k\ell})$. The interaction term above then becomes $\sum_{i,j,k,\ell} Q_{ij,k\ell} P_{ij} P_{k\ell}$.

Finally, the product empirical distribution has the same discrete support, with Dirac representation

$$(\widehat{\mu}_0 \otimes \widehat{\mu}_1)(\mathrm{d}z) = \sum_{i,j} a_i b_j \delta(z - z_{ij})\, \mathrm{d}z.$$

The probability masses of the transport plan π_P and the product empirical distribution at the support point z_{ij} are P_{ij} and $a_i b_j$, respectively. Hence, their KL divergence is

$$\mathrm{KL}(\pi_P \mid \widehat{\mu}_0 \otimes \widehat{\mu}_1) = \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j},$$

where the sample weights a_i, b_j are positive, and the corresponding term is defined to be zero when $P_{ij} = 0$.

Substituting the three results above into the empirical discretization of the objective in Proposition 1 and minimizing over the feasible transport plans $\mathcal{U}(a,b)$ gives

$$\min_{P \in \mathcal{U}(a,b)} \sum_{i,j} c_{ij} P_{ij} - \frac{1}{2} \sum_{i,j,k,\ell} Q_{ij,k\ell} P_{ij} P_{k\ell} + 2\nu \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j}.$$

This is precisely the entropy-regularized quadratic OT subproblem in Corollary 1 with the conditional paths fixed.

This discrete objective also admits a path-space relative entropy interpretation. We choose as the reference process a *mixture of Brownian bridges connecting the empirical data endpoints*: a discrete endpoint pair is first sampled according to fixed weights, and a continuous-space stochastic trajectory is then generated along the corresponding Brownian bridge. Let $D_{ij} = \frac{1}{2}\|x_1^j - x_0^i\|^2$, and define the reference endpoint weights and normalization constant by

$$r_{ij} = \frac{a_i b_j \exp[-D_{ij}/(2\nu)]}{\mathcal{Z}_\nu}, \qquad \mathcal{Z}_\nu = \sum_{i,j} a_i b_j \exp[-D_{ij}/(2\nu)].$$

The corresponding reference path measure is $\widehat{\mathbb{Q}}^\nu = \sum_{i,j} r_{ij}\mathbb{Q}^{\nu,ij}$, where $\mathbb{Q}^{\nu,ij}$ is the Brownian bridge from x_0^i to x_1^j with diffusion coefficient $\sqrt{2\nu}$. The reference weights r_{ij} are fixed, whereas the transport weights P_{ij} and conditional paths $\mathbb{P}^{ij}$ of the candidate process are optimization variables.

For the candidate path measure $\mathbb{P}_P = \sum_{i,j} P_{ij}\mathbb{P}^{ij}$, assume that each conditional path has finite KL divergence relative to the corresponding Brownian bridge. By the endpoint-conditioned decomposition of relative entropy,

$$\mathrm{KL}(\mathbb{P}_P \mid \widehat{\mathbb{Q}}^\nu) = \sum_{i,j} P_{ij} \log \frac{P_{ij}}{r_{ij}} + \sum_{i,j} P_{ij}\mathrm{KL}(\mathbb{P}^{ij} \mid \mathbb{Q}^{\nu,ij}).$$

Substituting r_{ij} and using $\sum_{i,j} P_{ij} = 1$ gives

$$2\nu\mathrm{KL}(\mathbb{P}_P \mid \widehat{\mathbb{Q}}^\nu) = \sum_{i,j} \left[D_{ij} + 2\nu\mathrm{KL}(\mathbb{P}^{ij} \mid \mathbb{Q}^{\nu,ij})\right] P_{ij} + 2\nu \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j} + 2\nu \log \mathcal{Z}_\nu.$$

Let $\rho_{P,t} = \sum_{i,j} P_{ij}\rho_t^{ij}$ denote the time marginal of the candidate path measure. Adding the external and pairwise potential energies, define

$$\widehat{\mathcal{A}}(\mathbb{P}_P) = 2\nu\mathrm{KL}(\mathbb{P}_P \mid \widehat{\mathbb{Q}}^\nu) - \int_0^1 \int V(x)\rho_{P,t}(\mathrm{d}x)\,\mathrm{d}t - \frac{1}{2}\int_0^1 \iint W(x,y)\rho_{P,t}(\mathrm{d}x)\rho_{P,t}(\mathrm{d}y)\,\mathrm{d}t.$$

Using the potential decompositions above together with the definitions of c_{ij} and $Q_{ij,k\ell}$, we obtain

$$\widehat{\mathcal{A}}(\mathbb{P}_P) = \sum_{i,j} c_{ij}P_{ij} - \frac{1}{2}\sum_{i,j,k,\ell} Q_{ij,k\ell}P_{ij}P_{k\ell} + 2\nu \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j} + 2\nu \log \mathcal{Z}_\nu.$$

Therefore, the discrete objective in Corollary 1 and the path-space objective relative to the reference bridge mixture $\widehat{\mathbb{Q}}^\nu$ differ only by the finite constant $2\nu \log \mathcal{Z}_\nu$. This constant depends only on the fixed empirical endpoints, sample weights, and diffusion coefficient, and therefore does not affect optimization of the endpoint transport plan or conditional paths.

□

A.4 Proof of Proposition 2: Fixed Points

Proposition 2 (Proximal fixed points and convergence)**.** For fixed conditional paths and any $\tau > 0$, a strictly positive $P^\star \in \mathcal{U}(a,b)$ is a fixed point of the KL-proximal update if and only if it satisfies the KKT conditions of the entropy-regularized quadratic OT objective. Moreover, with a fixed step size satisfying $\tau^{-1} > \|Q\|_{\mathrm{op}}$, the iterates converge to such a point. If the objective is convex on $\mathcal{U}(a,b)$, the resulting fixed point is also a global minimizer. Here, $\|Q\|_{\mathrm{op}}$ is the spectral norm of Q indexed over endpoint pairs.

Proof. We compare the KKT conditions of the original quadratic OT problem with those of its KL-proximal subproblem.

KKT Conditions of the Original Problem The original transport-plan problem is $\min_{P\in\mathcal{U}(a,b)} F(P)$, where

$$F(P) = \sum_{i,j} c_{ij}P_{ij} - \frac{1}{2}\sum_{i,j,k,\ell} Q_{ij,k\ell}P_{ij}P_{k\ell} + 2\nu\sum_{i,j} P_{ij}\log\frac{P_{ij}}{a_ib_j}.$$

Introduce multipliers α_i for the row marginal constraints $\sum_j P_{ij} = a_i$, multipliers β_j for the column marginal constraints $\sum_i P_{ij} = b_j$, and multipliers $\lambda_{ij} \geq 0$ for the nonnegativity constraints $-P_{ij} \leq 0$. The corresponding Lagrangian is

$$\mathcal{L}(P,\alpha,\beta,\lambda) = F(P) + \sum_i \alpha_i\left(\sum_j P_{ij} - a_i\right) + \sum_j \beta_j\left(\sum_i P_{ij} - b_j\right) - \sum_{i,j}\lambda_{ij}P_{ij}.$$

Fix an index pair (p,q). The derivative of the linear term is c_{pq}. For the quadratic term, P_{pq} can appear in either the first or the second factor of the product, and therefore

$$\frac{\partial}{\partial P_{pq}}\left[-\frac{1}{2}\sum_{i,j,k,\ell} Q_{ij,k\ell}P_{ij}P_{k\ell}\right] = -\frac{1}{2}\sum_{k,\ell} Q_{pq,k\ell}P_{k\ell} - \frac{1}{2}\sum_{i,j} Q_{ij,pq}P_{ij} = -\sum_{k,\ell} Q_{pq,k\ell}P_{k\ell},$$

where the final equality uses $Q_{ij,pq} = Q_{pq,ij}$ and relabels the summation indices.

For the entropy term, when $P_{pq} > 0$,

$$\frac{\partial}{\partial P_{pq}}\left[P_{pq}\log\frac{P_{pq}}{a_pb_q}\right] = \log\frac{P_{pq}}{a_pb_q} + P_{pq}\frac{1}{P_{pq}} = \log\frac{P_{pq}}{a_pb_q} + 1.$$

The derivatives of the row-constraint, column-constraint, and nonnegativity terms are α_p, β_q, and $-\lambda_{pq}$, respectively. Combining these results and replacing (p,q) by (i,j) gives

$$\frac{\partial\mathcal{L}}{\partial P_{ij}} = c_{ij} - \sum_{k,\ell} Q_{ij,k\ell}P_{k\ell} + 2\nu\left(\log\frac{P_{ij}}{a_ib_j} + 1\right) + \alpha_i + \beta_j - \lambda_{ij}.$$

Therefore, at a strictly positive transport plan, the KKT conditions of the original problem consist of

$$\begin{aligned}
&c_{ij} - \sum_{k,\ell} Q_{ij,k\ell}P_{k\ell} + 2\nu\left(\log\frac{P_{ij}}{a_ib_j} + 1\right) + \alpha_i + \beta_j - \lambda_{ij} = 0,\\
&\sum_j P_{ij} = a_i, \qquad \sum_i P_{ij} = b_j, \qquad P_{ij} \geq 0,\\
&\lambda_{ij} \geq 0, \qquad \lambda_{ij}P_{ij} = 0.
\end{aligned}$$

Since we consider $P_{ij} > 0$, complementary slackness gives $\lambda_{ij} = 0$. The stationarity condition therefore simplifies to

$$c_{ij} - \sum_{k,\ell} Q_{ij,k\ell}P_{k\ell} + 2\nu\left(\log\frac{P_{ij}}{a_ib_j} + 1\right) + \alpha_i + \beta_j = 0.$$

KKT Conditions of the Proximal Subproblem Let $R > 0$ denote the current transport plan, and define $C_{ij}(R) = c_{ij} - \sum_{k,\ell} Q_{ij,k\ell}R_{k\ell}$. The update minimizes

$$G_R(P) = \sum_{i,j} C_{ij}(R)P_{ij} + 2\nu\sum_{i,j} P_{ij}\log\frac{P_{ij}}{a_ib_j} + \frac{1}{\tau}\sum_{i,j} P_{ij}\log\frac{P_{ij}}{R_{ij}}, \qquad P\in\mathcal{U}(a,b).$$

Here R is fixed, so $C_{ij}(R)$ is also constant when differentiating with respect to the subproblem variable P. Introducing multipliers $\widetilde{\alpha}_i, \widetilde{\beta}_j, \widetilde{\lambda}_{ij}$ for the same constraints gives

$$\widetilde{\mathcal{L}}_R = G_R(P) + \sum_i \widetilde{\alpha}_i\left(\sum_j P_{ij} - a_i\right) + \sum_j \widetilde{\beta}_j\left(\sum_i P_{ij} - b_j\right) - \sum_{i,j}\widetilde{\lambda}_{ij}P_{ij}.$$

The derivative of the proximal term is

$$\frac{\partial}{\partial P_{ij}}\left[\frac{1}{\tau}P_{ij}\log\frac{P_{ij}}{R_{ij}}\right]=\frac{1}{\tau}\left(\log\frac{P_{ij}}{R_{ij}}+1\right).$$

Therefore,

$$\frac{\partial\widetilde{\mathcal{L}}_R}{\partial P_{ij}}=C_{ij}(R)+2\nu\left(\log\frac{P_{ij}}{a_ib_j}+1\right)+\frac{1}{\tau}\left(\log\frac{P_{ij}}{R_{ij}}+1\right)+\widetilde{\alpha}_i+\widetilde{\beta}_j-\widetilde{\lambda}_{ij}.$$

The KKT conditions of the subproblem are

$$\begin{aligned}
&C_{ij}(R)+2\nu\left(\log\frac{P_{ij}}{a_ib_j}+1\right)+\frac{1}{\tau}\left(\log\frac{P_{ij}}{R_{ij}}+1\right)+\widetilde{\alpha}_i+\widetilde{\beta}_j-\widetilde{\lambda}_{ij}=0,\\
&\sum_j P_{ij}=a_i,\qquad \sum_i P_{ij}=b_j,\qquad P_{ij}\geq 0,\\
&\widetilde{\lambda}_{ij}\geq 0,\qquad \widetilde{\lambda}_{ij}P_{ij}=0.
\end{aligned}$$

Equivalence of the KKT Conditions. If $P^\star>0$ is a fixed point of the proximal update, then setting $R=P^\star$ makes the solution of the subproblem $P=P^\star$ as well. In this case $\widetilde{\lambda}_{ij}=0$ and $\log(P^\star_{ij}/R_{ij})=0$. Substituting these into the stationarity condition of the subproblem gives

$$c_{ij}-\sum_{k,\ell}Q_{ij,k\ell}P^\star_{k\ell}+2\nu\left(\log\frac{P^\star_{ij}}{a_ib_j}+1\right)+\frac{1}{\tau}+\widetilde{\alpha}_i+\widetilde{\beta}_j=0.$$

Define $\alpha_i=\widetilde{\alpha}_i+\tau^{-1}$, $\beta_j=\widetilde{\beta}_j$, and $\lambda_{ij}=0$. The equation above then becomes

$$c_{ij}-\sum_{k,\ell}Q_{ij,k\ell}P^\star_{k\ell}+2\nu\left(\log\frac{P^\star_{ij}}{a_ib_j}+1\right)+\alpha_i+\beta_j=0.$$

The fixed point itself satisfies the marginal constraints and positivity, so all KKT conditions of the original problem hold.

Conversely, suppose $P^\star>0$ satisfies the KKT conditions of the original problem, with corresponding multipliers α_i,β_j and $\lambda_{ij}=0$. Set the current transport plan of the proximal subproblem to $R=P^\star$ and choose

$$\widetilde{\alpha}_i=\alpha_i-\frac{1}{\tau},\qquad \widetilde{\beta}_j=\beta_j,\qquad \widetilde{\lambda}_{ij}=0.$$

At the candidate solution $P=P^\star$, the derivative of the subproblem Lagrangian is

$$\left.\frac{\partial\widetilde{\mathcal{L}}_{P^\star}}{\partial P_{ij}}\right|_{P=P^\star}=c_{ij}-\sum_{k,\ell}Q_{ij,k\ell}P^\star_{k\ell}+2\nu\left(\log\frac{P^\star_{ij}}{a_ib_j}+1\right)+\frac{1}{\tau}+\left(\alpha_i-\frac{1}{\tau}\right)+\beta_j=0,$$

where the final equality is exactly the stationarity condition of the original problem. The remaining primal feasibility, dual feasibility, and complementary slackness conditions also hold directly, so $P^\star$ satisfies all KKT conditions of the proximal subproblem.

Finally, we verify that these conditions indeed determine the unique solution of the subproblem. Since $C(R)$ is fixed, the subproblem contains no quadratic interaction term, and its Hessian is diagonal:

$$\frac{\partial^2 G_R}{\partial P_{ij}\partial P_{k\ell}}=\frac{2\nu+\tau^{-1}}{P_{ij}}\,\delta_{ik}\delta_{j\ell}.$$

Since $2\nu+\tau^{-1}>0$, and $p\log p$ extended by $0\log 0=0$ is strictly convex on the nonnegative domain, G_R is strictly convex over the convex feasible set $\mathcal{U}(a,b)$. The KKT conditions above are therefore sufficient to guarantee that $P^\star$ is the unique minimizer. Hence, starting the update from $R=P^\star$ returns $P^\star$ itself, proving that it is a fixed point of the proximal update.

□

A.5 Proof of Proposition 2: Convergence

Proof. We invoke Theorem 6.2 of Bolte et al. (Bolte et al., 2018). For a bounded sequence of iterates, if the objective $\widetilde{F}$ incorporating the feasibility constraints satisfies the Kurdyka–Łojasiewicz property, then verifying the following three conditions guarantees that the entire sequence has finite length and converges to a stationary point.

1. **Sufficient decrease.** There exists a constant $c_1 > 0$ such that the decrease in the objective at each iteration satisfies
$$\widetilde{F}(P^{(r)}) - \widetilde{F}(P^{(r+1)}) \geq c_1 \|P^{(r+1)} - P^{(r)}\|_{\mathrm{F}}^2.$$
2. **Relative error bound.** There exist a constant $c_2 > 0$ and $w^{(r+1)} \in \partial \widetilde{F}(P^{(r+1)})$ such that the first-order optimality residual satisfies
$$\|w^{(r+1)}\|_{\mathrm{F}} \leq c_2 \|P^{(r+1)} - P^{(r)}\|_{\mathrm{F}}.$$
3. **Continuity of the objective along convergent subsequences.** For any convergent subsequence $P^{(r_k)} \to \overline{P}$,
$$\widetilde{F}(P^{(r_k)}) \longrightarrow \widetilde{F}(\overline{P}).$$

We first prove that the transport plans admit a uniform positive lower bound, which avoids the singularity of the entropy term at the zero-mass boundary. We then verify the three conditions above and the Kurdyka–Łojasiewicz property of the objective. Note that the Kurdyka–Łojasiewicz property is distinct from the Kullback–Leibler divergence.

Proof. With the family of conditional probability paths fixed, let $\gamma = \tau^{-1}$ and $L = \|Q\|_{\mathrm{op}}$, and write the original objective as $F(P) = f(P) + 2\nu h(P)$, where

$$f(P) = \sum_{i,j} c_{ij} P_{ij} - \frac{1}{2} \sum_{i,j,k,\ell} Q_{ij,k\ell} P_{ij} P_{k\ell}, \qquad h(P) = \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j}.$$

We work throughout with the full finite-support setting: $\nu > 0$, all marginal weights are strictly positive, all costs are finite, the initial transport plan is strictly positive, and each proximal subproblem is solved exactly. Since $Q_{ij,k\ell} = Q_{k\ell,ij}$,

$$\frac{\partial f}{\partial P_{ij}} = c_{ij} - \frac{1}{2} \sum_{k,\ell} Q_{ij,k\ell} P_{k\ell} - \frac{1}{2} \sum_{k,\ell} Q_{k\ell,ij} P_{k\ell} = c_{ij} - \sum_{k,\ell} Q_{ij,k\ell} P_{k\ell}.$$

Therefore, the effective cost satisfies $C^{(r)} = \nabla f(P^{(r)})$.

We first show that all iterates remain uniformly away from the zero-mass boundary. Combining the entropy terms in the proximal subproblem gives the Sinkhorn kernel (Cuturi, 2013)

$$K_{ij}^{(r)} = (a_i b_j)^{\frac{2\nu}{2\nu+\gamma}} (P_{ij}^{(r)})^{\frac{\gamma}{2\nu+\gamma}} \exp\left(-\frac{C_{ij}^{(r)}}{2\nu+\gamma}\right).$$

If $P^{(r)} > 0$, then this kernel is strictly positive, and its row and column scaling produces a strictly positive $P^{(r+1)}$. Hence, starting from $P^{(0)} > 0$, every iterate remains strictly positive.

Let the objective of the proximal subproblem be

$$G_r(P) = \sum_{i,j} C_{ij}^{(r)} P_{ij} + 2\nu \sum_{i,j} P_{ij} \log \frac{P_{ij}}{a_i b_j} + \gamma \sum_{i,j} P_{ij} \log \frac{P_{ij}}{P_{ij}^{(r)}}.$$

Its Lagrangian is

$$\mathcal{L}_r(P, \widetilde{\alpha}, \beta) = G_r(P) + \sum_i \widetilde{\alpha}_i \left(\sum_j P_{ij} - a_i\right) + \sum_j \beta_j \left(\sum_i P_{ij} - b_j\right).$$

Since the minimizer is strictly positive, the multipliers for the nonnegativity constraints vanish. Using $\frac{\mathrm{d}}{\mathrm{d}p}[p\log(p/q)] = \log(p/q) + 1$, differentiating with respect to P_{ij} and setting the derivative to zero gives

$$C_{ij}^{(r)} + 2\nu\left(\log\frac{P_{ij}^{(r+1)}}{a_i b_j} + 1\right) + \gamma\left(\log\frac{P_{ij}^{(r+1)}}{P_{ij}^{(r)}} + 1\right) + \widetilde{\alpha}_i^{(r+1)} + \beta_j^{(r+1)} = 0.$$

Define a new row multiplier $\alpha_i^{(r+1)} := \widetilde{\alpha}_i^{(r+1)} + \gamma$. The condition above can then be written as

$$C_{ij}^{(r)} + 2\nu\left(\log\frac{P_{ij}^{(r+1)}}{a_i b_j} + 1\right) + \gamma\log\frac{P_{ij}^{(r+1)}}{P_{ij}^{(r)}} + \alpha_i^{(r+1)} + \beta_j^{(r+1)} = 0.$$

The row constraints are equality constraints, so their multipliers may take arbitrary real values; thus, this redefinition does not change the meaning of the first-order conditions.

Add the first-order conditions at positions (i, j) and (k, ℓ) and subtract those at (i, ℓ) and (k, j). The row and column multipliers cancel:

$$\alpha_i + \alpha_k - \alpha_i - \alpha_k = 0, \qquad \beta_j + \beta_\ell - \beta_\ell - \beta_j = 0.$$

The marginal-weight terms also cancel because

$$\log\frac{(a_i b_j)(a_k b_\ell)}{(a_i b_\ell)(a_k b_j)} = 0.$$

Define the logarithmic cross-ratio by

$$\chi_{ij,k\ell}^{(r)} = \log\frac{P_{ij}^{(r)} P_{k\ell}^{(r)}}{P_{i\ell}^{(r)} P_{kj}^{(r)}}.$$

Combining the four first-order conditions yields

$$(2\nu + \gamma)\chi_{ij,k\ell}^{(r+1)} - \gamma\chi_{ij,k\ell}^{(r)} + C_{ij}^{(r)} + C_{k\ell}^{(r)} - C_{i\ell}^{(r)} - C_{kj}^{(r)} = 0.$$

Since $\sum_{k,\ell} P_{k\ell}^{(r)} = 1$, the effective costs admit an iteration-independent bound:

$$|C_{ij}^{(r)}| \le |c_{ij}| + \sum_{k,\ell} |Q_{ij,k\ell}| P_{k\ell}^{(r)} \le M,$$

where $M := \max_{i,j} |c_{ij}| + \max_{i,j,k,\ell} |Q_{ij,k\ell}|$. Let $\omega = \gamma/(2\nu + \gamma) < 1$. Then

$$|\chi_{ij,k\ell}^{(r+1)}| \le \omega|\chi_{ij,k\ell}^{(r)}| + \frac{4M}{2\nu + \gamma}.$$

Recursively,

$$|\chi_{ij,k\ell}^{(r)}| \le \omega^r|\chi_{ij,k\ell}^{(0)}| + \frac{4M}{2\nu + \gamma}\sum_{s=0}^{r-1}\omega^s = \omega^r|\chi_{ij,k\ell}^{(0)}| + \frac{4M}{2\nu}(1 - \omega^r).$$

Therefore, letting $B := \max\{\max_{i,j,k,\ell} |\chi_{ij,k\ell}^{(0)}|, 4M/(2\nu)\}$, all iterates satisfy $|\chi_{ij,k\ell}^{(r)}| \le B$, and hence

$$P_{ij}^{(r)} P_{k\ell}^{(r)} \ge e^{-B} P_{i\ell}^{(r)} P_{kj}^{(r)}.$$

Summing over k, ℓ and using the total-mass and marginal constraints gives

$$P_{ij}^{(r)} \ge e^{-B}\left(\sum_{\ell} P_{i\ell}^{(r)}\right)\left(\sum_{k} P_{kj}^{(r)}\right) = e^{-B} a_i b_j.$$

Thus, there exists a uniform positive lower bound

$$\delta := e^{-B}\min_{i,j} a_i b_j > 0, \qquad P_{ij}^{(r)} \ge \delta.$$

Sufficient Decrease. We next verify sufficient decrease. Let $\Delta^{(r)} = P^{(r+1)} - P^{(r)}$ denote the update difference, and write the divergence as $D(P, R) = \sum_{i,j} P_{ij} \log(P_{ij}/R_{ij})$. Using the current iterate $P^{(r)}$ as a comparison point in the proximal subproblem, the inequality $G_r(P^{(r+1)}) \leq G_r(P^{(r)})$ gives

$$\sum_{i,j} C_{ij}^{(r)} \Delta_{ij}^{(r)} + 2\nu\big[h(P^{(r+1)}) - h(P^{(r)})\big] + \gamma D(P^{(r+1)}, P^{(r)}) \leq 0.$$

Since f is quadratic, its expansion is exact:

$$f(P^{(r+1)}) - f(P^{(r)}) = \sum_{i,j} C_{ij}^{(r)} \Delta_{ij}^{(r)} - \frac{1}{2} \sum_{i,j,k,\ell} Q_{ij,k\ell} \Delta_{ij}^{(r)} \Delta_{k\ell}^{(r)}.$$

By the definition of the spectral norm,

$$-\frac{1}{2} \sum_{i,j,k,\ell} Q_{ij,k\ell} \Delta_{ij}^{(r)} \Delta_{k\ell}^{(r)} \leq \frac{L}{2} \|\Delta^{(r)}\|_{\mathrm{F}}^2.$$

Combining the results above yields

$$F(P^{(r+1)}) - F(P^{(r)}) \leq -\gamma D(P^{(r+1)}, P^{(r)}) + \frac{L}{2} \|\Delta^{(r)}\|_{\mathrm{F}}^2.$$

Since P and R have the same total mass, $D(P, R) = h(P) - h(R) - \langle \nabla h(R), P - R \rangle$. Using the second derivative of the entropy function, this can be written as

$$D(P, R) = \int_0^1 (1 - s) \sum_{i,j} \frac{(P_{ij} - R_{ij})^2}{R_{ij} + s(P_{ij} - R_{ij})} \, \mathrm{d}s.$$

Every component of a feasible transport plan is at most 1, so the denominator is at most 1. Therefore,

$$D(P, R) \geq \int_0^1 (1 - s) \|P - R\|_{\mathrm{F}}^2 \, \mathrm{d}s = \frac{1}{2} \|P - R\|_{\mathrm{F}}^2.$$

Hence, when $\gamma > L$,

$$F(P^{(r)}) - F(P^{(r+1)}) \geq \frac{\gamma - L}{2} \|P^{(r+1)} - P^{(r)}\|_{\mathrm{F}}^2.$$

This verifies the sufficient-decrease condition with $c_1 = (\gamma - L)/2 > 0$. With the convention $0 \log 0 = 0$, F is continuous on the compact feasible set and is therefore bounded below. Consequently, the objective values converge, and

$$\sum_{r=0}^{\infty} \|P^{(r+1)} - P^{(r)}\|_{\mathrm{F}}^2 < \infty.$$

Relative Error Bound. We next verify the relative error bound. To apply the convergence theorem cited above, define the auxiliary set $\mathcal{K} = \{P \in \mathcal{U}(a, b) : P_{ij} \geq \delta/2\}$ and let $\widetilde{F} = F + \iota_{\mathcal{K}}$, where the indicator function $\iota_{\mathcal{K}}$ is zero inside the set and $+\infty$ outside. Since all iterates satisfy $P_{ij}^{(r)} \geq \delta$, the additional lower-bound constraints are inactive at the iterates and therefore do not change the updates or their first-order conditions, while $\widetilde{F}(P^{(r)}) = F(P^{(r)})$.

The gradient of the original objective is

$$\frac{\partial F}{\partial P_{ij}}(P) = c_{ij} - \sum_{k,\ell} Q_{ij,k\ell} P_{k\ell} + 2\nu \left(\log \frac{P_{ij}}{a_i b_j} + 1 \right).$$

At $P^{(r+1)}$, adding the row and column multipliers above, define

$$w_{ij}^{(r+1)} := \frac{\partial F}{\partial P_{ij}}(P^{(r+1)}) + \alpha_i^{(r+1)} + \beta_j^{(r+1)}.$$

The row and column multipliers form a normal vector to the marginal constraints, so $w^{(r+1)} \in \partial\widetilde{F}(P^{(r+1)})$. Substituting the proximal first-order condition gives

$$w_{ij}^{(r+1)} = -\sum_{k,\ell} Q_{ij,k\ell}\Delta_{k\ell}^{(r)} - \gamma \log \frac{P_{ij}^{(r+1)}}{P_{ij}^{(r)}}.$$

Since every component is at least δ, the mean value theorem gives

$$|\log P_{ij}^{(r+1)} - \log P_{ij}^{(r)}| \le \delta^{-1}|P_{ij}^{(r+1)} - P_{ij}^{(r)}|.$$

Therefore,

$$\|w^{(r+1)}\|_{\mathrm{F}} \le \left(L + \frac{\gamma}{\delta}\right)\|P^{(r+1)} - P^{(r)}\|_{\mathrm{F}}.$$

This verifies the relative error condition with $c_2 = L + \gamma/\delta$.

Continuity of the Objective. Finally, $\mathcal{K}$ is a compact polytope contained entirely in the strictly positive region of transport plans, and F consists of polynomial and logarithmic functions in a neighborhood of $\mathcal{K}$, hence is real analytic there. Therefore, $\widetilde{F} = F + \iota_{\mathcal{K}}$ satisfies the Kurdyka–Łojasiewicz property. All iterates lie in $\mathcal{K}$, which guarantees boundedness. For any convergent subsequence $P^{(r_k)} \to \overline{P}$, continuity of F and $\overline{P} \in \mathcal{K}$ give

$$\widetilde{F}(P^{(r_k)}) = F(P^{(r_k)}) \longrightarrow F(\overline{P}) = \widetilde{F}(\overline{P}).$$

Thus, the continuity condition for the objective is also verified.

Applying Theorem 6.2 above gives

$$\sum_{r=0}^{\infty}\|P^{(r+1)} - P^{(r)}\|_{\mathrm{F}} < \infty, \qquad P^{(r)} \longrightarrow P^{\star}, \qquad 0 \in \partial\widetilde{F}(P^{\star}).$$

Since $P_{ij}^{\star} \ge \delta$, the auxiliary lower-bound constraints remain inactive at the limit. Therefore, the stationary condition above is exactly the KKT condition of the original problem: there exist row and column multipliers $\alpha_i^{\star}, \beta_j^{\star}$ such that

$$c_{ij} - \sum_{k,\ell} Q_{ij,k\ell}P_{k\ell}^{\star} + 2\nu\left(\log\frac{P_{ij}^{\star}}{a_i b_j} + 1\right) + \alpha_i^{\star} + \beta_j^{\star} = 0, \qquad P^{\star} \in \mathcal{U}(a,b).$$

Therefore, the entire sequence of transport plans converges to a strictly positive KKT stationary point of the original entropy-regularized quadratic OT problem. By the characterization in Appendix A.4, this point is also a fixed point of the KL-proximal update. If the objective is convex on $\mathcal{U}(a,b)$, the KKT conditions are sufficient for global optimality, so the resulting fixed point is also a global minimizer.

□

A.6 Construction of the Conditional Diffusion

Suppose m_t^z, Σ_t^z are differentiable and $\Sigma_t^z \succ 0$. Then the affine Itô SDE

$$\mathrm{d}X_t^z = b_t^z(X_t^z)\mathrm{d}t + \sqrt{2\nu}\mathrm{d}B_t^z, \qquad b_t^z(x) = \dot{m}_t^z + \left(\frac{1}{2}\dot{\Sigma}_t^z - \nu I\right)(\Sigma_t^z)^{-1}(x - m_t^z)$$

is compatible with $\rho_t^z = \mathcal{N}(m_t^z, \Sigma_t^z)$.

Derivation. Fix an endpoint pair $z = (x_0, x_1)$ and define

$$A_t^z = \left(\frac{1}{2}\dot{\Sigma}_t^z - \nu I\right)(\Sigma_t^z)^{-1}.$$

Consider the affine SDE

$$\mathrm{d}X_t^z = [\dot{m}_t^z + A_t^z(X_t^z - m_t^z)]\,\mathrm{d}t + \sqrt{2\nu}\,\mathrm{d}B_t^z.$$

Let $Y_t^z = X_t^z - m_t^z$. Then

$$\mathrm{d}Y_t^z = A_t^z Y_t^z \,\mathrm{d}t + \sqrt{2\nu}\,\mathrm{d}B_t^z.$$

At any interior time $s \in (0,1)$, initialize $X_s^z \sim \mathcal{N}(m_s^z, \Sigma_s^z)$. Taking expectations gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[Y_t^z] = A_t^z \mathbb{E}[Y_t^z].$$

Since $\mathbb{E}[Y_s^z] = 0$, this linear equation yields $\mathbb{E}[Y_t^z] = 0$, and hence $\mathbb{E}[X_t^z] = m_t^z$.

Furthermore, by Itô's product rule,

$$\mathrm{d}(Y_t^z (Y_t^z)^\top) = (\mathrm{d}Y_t^z)(Y_t^z)^\top + Y_t^z (\mathrm{d}Y_t^z)^\top + (\mathrm{d}Y_t^z)(\mathrm{d}Y_t^z)^\top.$$

Using $(\mathrm{d}B_t^z)(\mathrm{d}B_t^z)^\top = I\,\mathrm{d}t$ and defining $\Gamma_t^z = \mathbb{E}[Y_t^z (Y_t^z)^\top]$, taking expectations gives

$$\dot{\Gamma}_t^z = A_t^z \Gamma_t^z + \Gamma_t^z (A_t^z)^\top + 2\nu I.$$

Since both Σ_t^z and $\dot{\Sigma}_t^z$ are symmetric,

$$A_t^z \Sigma_t^z = \frac{1}{2}\dot{\Sigma}_t^z - \nu I, \qquad \Sigma_t^z (A_t^z)^\top = \frac{1}{2}\dot{\Sigma}_t^z - \nu I.$$

Therefore,

$$A_t^z \Sigma_t^z + \Sigma_t^z (A_t^z)^\top + 2\nu I = \dot{\Sigma}_t^z.$$

Thus, Σ_t^z satisfies the same linear matrix differential equation and initial condition as Γ_t^z, so $\Gamma_t^z = \Sigma_t^z$. Since an affine SDE initialized from a Gaussian distribution remains Gaussian, we obtain $X_t^z \sim \mathcal{N}(m_t^z, \Sigma_t^z)$.

□

A.7 Derivation of Proposition 3

Proposition 3 (Closed-form conditional path KL)**.** Let $\mathbb{Q}^{\nu,z}$ be the Brownian bridge with the same endpoints and diffusion coefficient as $\mathbb{P}^z$. Under the conditions for Girsanov's theorem,

$$2\nu \mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) = \frac{1}{2}\int_0^1 \left[\|d_t^z\|^2 + \mathrm{tr}\left(H_t^z \Sigma_t^z (H_t^z)^\top\right)\right] \mathrm{d}t,$$

where

$$d_t^z = \dot{m}_t^z - \frac{x_1 - m_t^z}{1-t}, \qquad H_t^z = \left(\frac{1}{2}\dot{\Sigma}_t^z - \nu I\right)(\Sigma_t^z)^{-1} + \frac{I}{1-t}.$$

Derivation. The drift of the reference Brownian bridge $\mathbb{Q}^{\nu,z}$ is $b_t^{\mathrm{BB},z}(x) = (x_1 - x)/(1-t)$. Define

$$\Delta b_t^z(x) = b_t^z(x) - b_t^{\mathrm{BB},z}(x), \qquad u_t^z(x) = \frac{\Delta b_t^z(x)}{\sqrt{2\nu}}.$$

To avoid the singularity at the terminal time, first consider any $T < 1$. Under the conditions for Girsanov's theorem, the path measures of the two processes on $[0,T]$ satisfy

$$\log \frac{\mathrm{d}\mathbb{P}^z_{[0,T]}}{\mathrm{d}\mathbb{Q}^{\nu,z}_{[0,T]}} = \int_0^T u_t^z(X_t^z) \cdot \mathrm{d}B_t^{\mathbb{Q},z} - \frac{1}{2}\int_0^T \|u_t^z(X_t^z)\|^2 \,\mathrm{d}t.$$

Under $\mathbb{P}^z$, $\mathrm{d}B_t^{\mathbb{Q},z} = \mathrm{d}B_t^{\mathbb{P},z} + u_t^z(X_t^z)\,\mathrm{d}t$. Substituting this into the expression above gives

$$\log \frac{\mathrm{d}\mathbb{P}^z_{[0,T]}}{\mathrm{d}\mathbb{Q}^{\nu,z}_{[0,T]}} = \int_0^T u_t^z(X_t^z) \cdot \mathrm{d}B_t^{\mathbb{P},z} + \frac{1}{2}\int_0^T \|u_t^z(X_t^z)\|^2 \,\mathrm{d}t.$$

Taking expectation under $\mathbb{P}^z$, the square-integrable stochastic integral has zero expectation, and therefore

$$\mathrm{KL}(\mathbb{P}^z_{[0,T]} \mid \mathbb{Q}^{\nu,z}_{[0,T]}) = \frac{1}{4\nu}\int_0^T \mathbb{E}_{\rho_t^z}\left[\|\Delta b_t^z(X_t^z)\|^2\right] \mathrm{d}t.$$

The two processes have the same initial point, so there is no additional initial relative-entropy term. Letting $T \uparrow 1$ and using the limiting property of relative entropy with respect to increasing path information yields

$$\mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) = \frac{1}{4\nu} \int_0^1 \mathbb{E}_{\rho_t^z} \left[\|\Delta b_t^z(X_t^z)\|^2 \right] \mathrm{d}t.$$

This equality allows both sides to take the value $+\infty$; the time integral is understood as an improper integral at the endpoints.

We next evaluate the spatial expectation analytically. Expanding the reference drift as

$$\frac{x_1 - x}{1-t} = \frac{x_1 - m_t^z}{1-t} - \frac{x - m_t^z}{1-t},$$

we obtain

$$\Delta b_t^z(x) = d_t^z + H_t^z(x - m_t^z),$$

where

$$d_t^z = \dot{m}_t^z - \frac{x_1 - m_t^z}{1-t}, \qquad H_t^z = \left(\frac{1}{2}\dot{\Sigma}_t^z - \nu I \right) (\Sigma_t^z)^{-1} + \frac{I}{1-t}.$$

Let $\xi_t^z = X_t^z - m_t^z$. From the Gaussian marginal distribution, $\mathbb{E}[\xi_t^z] = 0$ and $\mathbb{E}[\xi_t^z (\xi_t^z)^\top] = \Sigma_t^z$. Therefore,

$$\begin{aligned} \mathbb{E}\left[\|d_t^z + H_t^z \xi_t^z\|^2 \right] &= \|d_t^z\|^2 + 2(d_t^z)^\top H_t^z \mathbb{E}[\xi_t^z] + \mathbb{E}[(\xi_t^z)^\top (H_t^z)^\top H_t^z \xi_t^z] \\ &= \|d_t^z\|^2 + \mathrm{tr}\left((H_t^z)^\top H_t^z \Sigma_t^z \right) \\ &= \|d_t^z\|^2 + \mathrm{tr}\left(H_t^z \Sigma_t^z (H_t^z)^\top \right). \end{aligned}$$

Substituting this into the path-space KL expression and multiplying by 2ν gives

$$2\nu \mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) = \frac{1}{2} \int_0^1 \left[\|d_t^z\|^2 + \mathrm{tr}\left(H_t^z \Sigma_t^z (H_t^z)^\top \right) \right] \mathrm{d}t.$$

□

A.8 Proof of Proposition 4

Proposition 4 (Sufficient conditions for finite path KL)**.** Let $h(t) = t(1-t)$. If

$$m_t^z = (1-t)x_0 + t x_1 + h(t) r_t^z, \qquad \Sigma_t^z = 2\nu h(t) R_t^z,$$

where r_t^z, R_t^z are continuously differentiable on $[0,1]$, R_t^z is symmetric positive definite, and $R_0^z = R_1^z = I$, then $\mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) < \infty$.

Proof. Let $h(t) = t(1-t)$. For the mean parameterization in the proposition,

$$m_t^z = (1-t)x_0 + t x_1 + h(t) r_t^z,$$

direct differentiation gives

$$\dot{m}_t^z = x_1 - x_0 + (1-2t) r_t^z + h(t) \dot{r}_t^z, \qquad \frac{x_1 - m_t^z}{1-t} = x_1 - x_0 - t r_t^z.$$

Therefore,

$$d_t^z = (1-t) r_t^z + h(t) \dot{r}_t^z = (1-t)(r_t^z + t \dot{r}_t^z).$$

Since r_t^z is continuously differentiable, d_t^z is bounded on $[0,1]$.

For the covariance parameterization $\Sigma_t^z = 2\nu h(t) R_t^z$,

$$\dot{\Sigma}_t^z = 2\nu \left[h'(t) R_t^z + h(t) \dot{R}_t^z \right], \qquad (\Sigma_t^z)^{-1} = \frac{(R_t^z)^{-1}}{2\nu h(t)}.$$

Substituting these into the definition of H_t^z gives

$$H_t^z = \frac{h'(t)}{2h(t)} I + \frac{1}{2} \dot{R}_t^z (R_t^z)^{-1} - \frac{(R_t^z)^{-1}}{2h(t)} + \frac{I}{1-t} = \frac{1}{2} \dot{R}_t^z (R_t^z)^{-1} + \frac{I - (R_t^z)^{-1}}{2h(t)},$$

where we used $h'(t) + 2t = 1$.

Since R_t^z is continuous and positive definite on a compact interval, there exists $\ell_z > 0$ such that $R_t^z \succeq \ell_z I$, and hence $\|(R_t^z)^{-1}\|_2 \le \ell_z^{-1}$. Let $K_z = \sup_t \|\dot{R}_t^z\|_2 < \infty$. Since $R_0^z = R_1^z = I$,

$$R_t^z - I = \int_0^t \dot{R}_s^z \, \mathrm{d}s = -\int_t^1 \dot{R}_s^z \, \mathrm{d}s,$$

so $\|R_t^z - I\|_2 \le K_z \min(t, 1-t)$. Using $I - (R_t^z)^{-1} = (R_t^z)^{-1}(R_t^z - I)$, we obtain

$$\left\| \frac{I - (R_t^z)^{-1}}{2h(t)} \right\|_2 \le \frac{K_z}{2\ell_z} \frac{\min(t, 1-t)}{t(1-t)} \le \frac{K_z}{\ell_z}.$$

Therefore, H_t^z is uniformly bounded over the entire interior time interval. This also shows that the drift difference $\Delta b_t^z(x) = d_t^z + H_t^z(x - m_t^z)$ grows at most linearly in x, with uniformly bounded coefficients.

Suppose $\|d_t^z\| \le D_z$, $\|H_t^z\|_2 \le E_z$, and $\mathrm{tr}(R_t^z) \le S_z$. Then

$$\mathrm{tr}\left(H_t^z \Sigma_t^z (H_t^z)^\top\right) \le \|H_t^z\|_2^2 \, \mathrm{tr}(\Sigma_t^z) \le 2\nu E_z^2 S_z h(t).$$

Hence, the integrand in the KL expression satisfies

$$\int_0^1 \left[\|d_t^z\|^2 + \mathrm{tr}\left(H_t^z \Sigma_t^z (H_t^z)^\top\right)\right] \mathrm{d}t \le D_z^2 + \frac{\nu E_z^2 S_z}{3} < \infty.$$

It remains to justify the change of measure at the endpoints. On any $[0, T]$ with $T < 1$, the coefficients of the Brownian-bridge drift and the affine drift above are bounded, so Girsanov's theorem applies by standard localization. The estimates above further provide a uniform upper bound, with respect to T, on the relative entropy of these restricted path measures. Consequently, their likelihood ratios are uniformly integrable, and the change of measure extends to the full interval $[0, 1]$. The resulting path measure is absolutely continuous with respect to the reference Brownian bridge and inherits its fixed endpoints.

Applying Proposition 3, we finally obtain

$$\mathrm{KL}(\mathbb{P}^z \mid \mathbb{Q}^{\nu,z}) \le \frac{D_z^2}{4\nu} + \frac{E_z^2 S_z}{12} < \infty.$$

□

A.9 Derivation of the Conditional Velocity and Score

Conditional velocity and score expressions. For $\rho_{\phi,t}^z = \mathcal{N}(m_{\phi,t}^z, \Sigma_{\phi,t}^z)$, the conditional score and PF-ODE velocity are

$$s_{\phi,t}^z(x) = -(\Sigma_{\phi,t}^z)^{-1}(x - m_{\phi,t}^z), \qquad v_{\phi,t}^z(x) = \dot{m}_{\phi,t}^z + \frac{1}{2}\dot{\Sigma}_{\phi,t}^z(\Sigma_{\phi,t}^z)^{-1}(x - m_{\phi,t}^z).$$

Derivation. Fix an endpoint pair z. The log-density of the conditional Gaussian distribution is

$$\log \rho_{\phi,t}^z(x) = -\frac{d}{2}\log(2\pi) - \frac{1}{2}\log\det \Sigma_{\phi,t}^z - \frac{1}{2}(x - m_{\phi,t}^z)^\top (\Sigma_{\phi,t}^z)^{-1}(x - m_{\phi,t}^z).$$

Differentiating with respect to x gives the conditional score:

$$s_{\phi,t}^z(x) = \nabla_x \log \rho_{\phi,t}^z(x) = -(\Sigma_{\phi,t}^z)^{-1}(x - m_{\phi,t}^z).$$

By the construction in Appendix A.6, the conditional SDE drift is

$$b_{\phi,t}^z(x) = \dot{m}_{\phi,t}^z + \left(\frac{1}{2}\dot{\Sigma}_{\phi,t}^z - \nu I\right)(\Sigma_{\phi,t}^z)^{-1}(x - m_{\phi,t}^z).$$

Using $\nabla \rho_{\phi,t}^z = \rho_{\phi,t}^z s_{\phi,t}^z$, its Fokker–Planck equation can be rewritten as

$$\partial_t \rho_{\phi,t}^z = -\nabla \cdot (\rho_{\phi,t}^z b_{\phi,t}^z) + \nu \Delta \rho_{\phi,t}^z = -\nabla \cdot \left[\rho_{\phi,t}^z (b_{\phi,t}^z - \nu s_{\phi,t}^z)\right].$$

Therefore, the conditional PF-ODE velocity is

$$v_{\phi,t}^z(x) = b_{\phi,t}^z(x) - \nu s_{\phi,t}^z(x) = \dot{m}_{\phi,t}^z + \frac{1}{2}\dot{\Sigma}_{\phi,t}^z(\Sigma_{\phi,t}^z)^{-1}(x - m_{\phi,t}^z).$$

□

A.10 Regression Does Not Increase the Action

The Markovization induced by velocity and score regression does not increase the action.

Proof. Fix the transport plan P and conditional-path parameters ϕ, and let the mixture path measure be $\mathbb{P}_\phi = \int \mathbb{P}^z_\phi \pi_P(\mathrm{d}z)$, with time marginal $\rho_{\phi,t}$. According to Theorems 3.1 and 3.2 of SF^{2}M (Tong et al., 2024b), exact regression recovers the posterior means of the conditional velocity and score. Therefore, the recovered SDE drift satisfies

$$b^\star(t,x) = v^\star(t,x) + \nu s^\star(t,x) = \mathbb{E}[b^Z_{\phi,t}(x) \mid X_t = x].$$

The process with this drift and diffusion coefficient $\sqrt{2\nu}$ is the Markovization of $\mathbb{P}_\phi$. Let its marginals be $\widehat{\rho}_t$. Then $\widehat{\rho}_t = \rho_{\phi,t}$.

Let $\mathcal{F}^X_t$ denote the natural filtration of the mixture process, and let $\beta_{\phi,t}$ denote its drift with respect to this natural filtration. By conditional projection of the drift and the tower property of conditional expectation,

$$\beta_{\phi,t} = \mathbb{E}[b^Z_{\phi,t}(X_t) \mid \mathcal{F}^X_t], \qquad b^\star(t,X_t) = \mathbb{E}[\beta_{\phi,t} \mid X_t].$$

The orthogonal decomposition of conditional expectation gives

$$\mathbb{E}_{\mathbb{P}_\phi}\|\beta_{\phi,t}\|^2 - \mathbb{E}_{\rho_{\phi,t}}\|b^\star(t,X_t)\|^2 = \mathbb{E}_{\mathbb{P}_\phi}\|\beta_{\phi,t} - b^\star(t,X_t)\|^2 \geq 0.$$

Thus, Markovization does not increase the kinetic energy. On the other hand, since all time marginals remain unchanged after regression, the external potential energy satisfies

$$\int_0^1 \int V(x)\widehat{\rho}_t(\mathrm{d}x)\,\mathrm{d}t = \int_0^1 \int V(x)\rho_{\phi,t}(\mathrm{d}x)\,\mathrm{d}t,$$

and the pairwise potential energy also satisfies

$$\frac{1}{2}\int_0^1 \iint W(x,y)\widehat{\rho}_t(\mathrm{d}x)\widehat{\rho}_t(\mathrm{d}y)\,\mathrm{d}t = \frac{1}{2}\int_0^1 \iint W(x,y)\rho_{\phi,t}(\mathrm{d}x)\rho_{\phi,t}(\mathrm{d}y)\,\mathrm{d}t.$$

Therefore, neither potential-energy term in the action changes. Defining the action of the pre-regression mixture process using the kinetic energy associated with drift $\beta_{\phi,t}$ minus the potential terms above, we obtain

$$\mathcal{A}_{\text{before}} - \mathcal{A}_{\text{after}} = \frac{1}{2}\int_0^1 \mathbb{E}_{\mathbb{P}_\phi}\|\beta_{\phi,t} - b^\star(t,X_t)\|^2\,\mathrm{d}t \geq 0.$$

Therefore, the action does not increase after regression.

□

B Experimental Details

B.1 Experimental Settings and Parameters

All CCB and baseline experiments are conducted on a single NVIDIA A100 GPU. CCB uses four standard MLPs to parameterize the mean and diagonal covariance of the conditional probability paths, the PF-ODE velocity field $v_\theta(t,x)$, and the score $s_\psi(t,x)$. The mean and covariance MLPs take the time t and the pair (x_0,x_1) as inputs and output $r^z_{\phi,t}$ and $g^z_{\phi,t}$, respectively. These outputs determine the conditional mean $m^z_{\phi,t}$ and variance vector $q^z_{\phi,t}$ through the boundary-compatible parameterization in the main text, with $\Sigma^z_{\phi,t} = \mathrm{diag}(q^z_{\phi,t})$. The velocity and score MLPs instead take (t,x) as inputs. We construct the SDE drift as $b(t,x) = v_\theta(t,x) + \nu s_\psi(t,x)$. The diffusion coefficient is uniformly set to $\sigma = \sqrt{2\nu}$.

For the 10D and 20D synthetic datasets with known ground truth, we set $\nu = 0.05$ to match the ground-truth diffusion. The Two-signal KS and Coulomb–harmonic datasets are generated by deterministic mean-field dynamics, so we use a small noise level of $\nu = 0.001$. For real single-cell and spatial transcriptomics datasets, we select the noise level according to the scale of the preprocessed data. EMT and the STVCR axolotl dataset have substantially smaller scales, for which we

use $\nu = 0.005$; the remaining datasets have larger scales, for which we use $\nu = 0.4$. Importantly, to ensure a fair comparison, all stochastic methods are evaluated with the same noise intensity on each dataset.

The noise parameters and coupling configurations for each experiment are summarized in Table 5. Here, τ denotes the proximal-update parameter, set to 20 in all experiments. In Sparse mode, q denotes the number of new edges proposed for each source and target particle at each iteration, and K denotes the maximum number of highest-transport-mass edges retained for each particle during bidirectional pruning. The final support is the union of the edges retained from both sides. "—" indicates Dense coupling, for which sparse edge proposal and pruning are not used. All configurations in the table correspond to CCB.

Table 5: Noise and coupling configurations used by CCB.

Setting	Noise ν	Proposed q	Retained K	Proximal τ
Attr. quadratic interaction (10D)	0.05	—	—	20
Linear external potential (10D)	0.05	—	—	20
Quadratic external potential (10D)	0.05	—	—	20
Repulsive quadratic interaction (10D)	0.05	—	—	20
Attr. quadratic interaction (20D)	0.05	128	32	20
Linear external potential (20D)	0.05	128	32	20
Quadratic external potential (20D)	0.05	128	32	20
Repulsive quadratic interaction (20D)	0.05	128	32	20
EMT (10D)	0.005	—	—	20
EB (5D)	0.4	64	16	20
EB (50D)	0.4	64	16	20
EB (100D)	0.4	64	16	20
CITE-seq (50D)	0.4	64	16	20
Weinreb (50D)	0.4	32	16	20
Two-signal KS (20D)	0.001	64	16	20
Coulomb–harmonic (20D)	0.001	—	—	20
STVCR Axolotli (32D)	0.005	64	16	20

B.2 Evaluation Metrics

B.2.1 Evaluation Metrics for Experiment 1

In Experiment 1, we use the absolute action bias and the dimension-normalized mean squared error of the SDE drift to evaluate reconstruction accuracy of the control objective and control law, respectively. For each trained model, we perform $R = 8$ independent SDE rollouts:

$$\mathrm{d}X_t = b_\theta(t, X_t)\mathrm{d}t + \sqrt{2\nu}\mathrm{d}B_t, \qquad X_0 \sim \mu_0.$$

Let $\{X_{t,i}^{(r)}\}_{i=1}^N$ denote the particles in the r-th rollout, and let $b^\star$ and $\mathcal{A}^\star$ denote the analytical optimal drift and action, respectively.

Absolute Action Bias. The action estimate for the r-th rollout is

$$\widehat{\mathcal{A}}^{(r)} = \int_0^T \left[\frac{1}{2N} \sum_{i=1}^N \|b_\theta(t, X_{t,i}^{(r)})\|^2 - \frac{1}{N} \sum_{i=1}^N V(X_{t,i}^{(r)}) - \frac{1}{2N^2} \sum_{i,j=1}^N W(X_{t,i}^{(r)}, X_{t,j}^{(r)}) \right] \mathrm{d}t.$$

The corresponding absolute action bias is

$$\mathcal{E}_A^{(r)} = \left| \widehat{\mathcal{A}}^{(r)} - \mathcal{A}^\star \right|.$$

Control-law Error. Along the model's own SDE rollout, we compare the learned drift with the analytical optimal drift at the same times and particle locations:

$$\mathcal{E}_b^{(r)} = \frac{1}{dTN} \int_0^T \sum_{i=1}^N \left\| b_\theta(t, X_{t,i}^{(r)}) - b^\star(t, X_{t,i}^{(r)}) \right\|^2 \mathrm{d}t.$$

For either metric, we report the mean and sample standard deviation:

$$\overline{\mathcal{E}} = \frac{1}{R} \sum_{r=1}^R \mathcal{E}^{(r)}, \qquad s_{\mathcal{E}} = \sqrt{\frac{1}{R-1} \sum_{r=1}^R \left(\mathcal{E}^{(r)} - \overline{\mathcal{E}}\right)^2}.$$

B.2.2 Evaluation Metrics for Experiment 2

We use the Wasserstein-1 distance to evaluate the discrepancy between generated and observed distributions. At observation time t_k, let $\{x_i^{(k)}\}_{i=1}^{N_k}$ denote the true cell states and $\{\widehat{x}_j^{(k)}\}_{j=1}^{M_k}$ the generated particles. The true distribution is represented as $\mu_k = \sum_i a_i^{(k)} \delta_{x_i^{(k)}}$, where $a_i^{(k)} = 1/N_k$, and the generated distribution as $\widehat{\mu}_k = \sum_j w_j^{(k)} \delta_{\widehat{x}_j^{(k)}}$. For unbalanced methods that output particle masses $\widetilde{w}_j^{(k)}$, we use normalized weights $w_j^{(k)} = \widetilde{w}_j^{(k)} / \sum_\ell \widetilde{w}_\ell^{(k)}$; all other methods use uniform weights $w_j^{(k)} = 1/M_k$.

We compute the discrete optimal transport distance

$$W_1(\mu_k, \widehat{\mu}_k) = \min_{P \in \mathcal{U}(a^{(k)}, w^{(k)})} \sum_{i=1}^{N_k} \sum_{j=1}^{M_k} P_{ij} \|x_i^{(k)} - \widehat{x}_j^{(k)}\|_2,$$

where $\mathcal{U}(a^{(k)}, w^{(k)}) = \{P \geq 0 : \sum_j P_{ij} = a_i^{(k)},\ \sum_i P_{ij} = w_j^{(k)}\}$. The Euclidean distance in the feature space of the corresponding experiment is used as the transport cost.

B.2.3 Evaluation Metrics for Experiment 3

At the three held-out time points, we separately evaluate the spatial distribution, expression distribution, and cell-type neighborhood structure of the generated slices. At a given time, let the true cells be $x_i = (s_i, g_i)$ and the generated cells be $\widehat{x}_j = (\widehat{s}_j, \widehat{g}_j)$, where $s_i, \widehat{s}_j \in \mathbb{R}^2$ are spatial coordinates and $g_i, \widehat{g}_j \in \mathbb{R}^{d-2}$ are expression features, with d denoting the total dimension. True cells use uniform weights $a_i = 1/N$. Let b_j denote the normalized weights of generated cells; $b_j = 1/M$ for equally weighted generation, while methods that predict masses use their normalized output masses. The corresponding empirical spatial and expression distributions are

$$\mu^{\mathrm{sp}} = \sum_{i=1}^N a_i \delta_{s_i}, \qquad \widehat{\mu}^{\mathrm{sp}} = \sum_{j=1}^M b_j \delta_{\widehat{s}_j}, \qquad \mu^{\mathrm{expr}} = \sum_{i=1}^N a_i \delta_{g_i}, \qquad \widehat{\mu}^{\mathrm{expr}} = \sum_{j=1}^M b_j \delta_{\widehat{g}_j}.$$

Spatial W1 and Expression W1. We compute Wasserstein-1 distances separately over spatial coordinates and expression features:

$$W_1^{\mathrm{spatial}}(\mu^{\mathrm{sp}}, \widehat{\mu}^{\mathrm{sp}}) = \min_{P \in \mathcal{U}(a,b)} \sum_{i,j} P_{ij} \|s_i - \widehat{s}_j\|_2,$$

$$W_1^{\mathrm{expression}}(\mu^{\mathrm{expr}}, \widehat{\mu}^{\mathrm{expr}}) = \min_{P \in \mathcal{U}(a,b)} \sum_{i,j} P_{ij} \|g_i - \widehat{g}_j\|_2.$$

Here $\mathcal{U}(a,b) = \{P \geq 0 : \sum_j P_{ij} = a_i,\ \sum_i P_{ij} = b_j\}$. Spatial W_1 uses Euclidean distance between two-dimensional spatial coordinates and measures reconstruction of the spatial cell distribution. Expression W_1 uses Euclidean distance between the $(d-2)$-dimensional expression features and measures reconstruction of the expression-state distribution. Their optimal couplings are solved independently.

Cell-type Neighborhood TV. We train a cell-type classifier using the expression features and type labels of real cells from all time points, and use it to predict cell types for generated cells. We then identify the 15 spatially nearest neighbors of each cell in both the real and generated slices and compute the neighbor-type composition conditioned on the type of the center cell. Let $p_t(c' \mid c)$ denote the proportion of neighbors of type c' around center cells of type c in the real slice, and let $\widehat{p}_t(c' \mid c)$ denote the corresponding proportion in the generated slice. We compute the total variation distance between these conditional distributions, weighted by the abundance $\alpha_t(c)$ of each center-cell type in the real slice:

$$\mathrm{TV}_{\mathrm{type}}(t) = \frac{1}{2} \sum_{c} \alpha_t(c) \sum_{c'} \left| p_t(c' \mid c) - \widehat{p}_t(c' \mid c) \right|.$$

This metric compares local spatial neighborhood relationships among different cell types rather than only the global type proportions of the entire slice. Lower values indicate that the cell-type neighborhood structure of the generated slice is closer to that of the real slice. The classifier is used only for post-generation evaluation; cell-type labels, including those at held-out time points, are not used in potential construction, transport-plan optimization, or flow-network training.

B.3 Construction of Synthetic Datasets

B.3.1 Construction of Synthetic Datasets in Experiment 1

We consider stochastic optimal control problems with either external or pairwise interaction potentials on $t \in [0, 1]$, governed by the SDE $\mathrm{d}X_t = b_t(X_t)\mathrm{d}t + \sqrt{2\nu}\mathrm{d}B_t$, with $\nu = 0.05$. We use dimensions $d = 10$ and $d = 20$, with endpoint marginals $\mu_0 = \mathcal{N}(-2.5e_1, 0.1I_d)$ and $\mu_1 = \mathcal{N}(2.5e_1, 0.05I_d)$, where e_1 is the unit vector along the first coordinate. We refer to the positive and negative quadratic kernels as attractive and repulsive, respectively, according to the induced variational force $-\nabla_x W$. We consider the following four potential settings, with $\kappa = 4$ and $f = e_1$ throughout:

Table 6: External and pairwise potentials for the synthetic optimal-control problems.

Setting	External Potential $V(x)$	Pairwise Potential $W(x, y)$
Attractive Quadratic Interaction	0	$\frac{\kappa}{2}\|x - y\|^2$
Repulsive Quadratic Interaction	0	$-\frac{\kappa}{2}\|x - y\|^2$
Linear External Potential	$-f^\top x$	0
Quadratic External Potential	$\frac{\kappa}{2}\|x\|^2$	0

We first show that, for the four potentials above and Gaussian endpoints, the optimal distribution has the Gaussian form $\rho_t^\star = \mathcal{N}(m_t, s_t^2 I_d)$. For any feasible process, let $m_t = \mathbb{E}[X_t]$ denote its mean and $s_t^2 = d^{-1}\mathbb{E}\|X_t - m_t\|^2$ its average coordinate variance. By Itô's formula,

$$\dot{m}_t = \mathbb{E}[b_t(X_t)], \qquad d(s_t \dot{s}_t) = \mathbb{E}[(X_t - m_t)^\top b_t(X_t)] + \nu d.$$

Decomposing the drift into its mean and centered components gives $\mathbb{E}\|b_t(X_t)\|^2 = \|\dot{m}_t\|^2 + \mathbb{E}\|b_t(X_t) - \dot{m}_t\|^2$. By the Cauchy–Schwarz inequality,

$$\mathbb{E}\|b_t(X_t) - \dot{m}_t\|^2 \geq \frac{\left(\mathbb{E}[(X_t - m_t)^\top (b_t(X_t) - \dot{m}_t)]\right)^2}{d s_t^2} = d\left(\dot{s}_t - \frac{\nu}{s_t}\right)^2.$$

Equality is attained by the affine drift $b_t(x) = \dot{m}_t + (\dot{s}_t/s_t - \nu/s_t^2)(x - m_t)$. Starting from a Gaussian initial distribution, this drift generates Gaussian marginals with mean m_t and covariance $s_t^2 I_d$.

Therefore, computing the GT reduces to a two-point boundary-value problem for the mean m_t and standard deviation s_t. The expected values of the four potentials also depend only on m_t and s_t:

$$U(m,s) = \begin{cases} \frac{\kappa d}{2} s^2, & \text{attractive quadratic interaction,} \\ -\frac{\kappa d}{2} s^2, & \text{repulsive quadratic interaction,} \\ -f^\top m, & \text{linear external potential,} \\ \frac{\kappa}{2}(\|m\|^2 + d s^2), & \text{quadratic external potential.} \end{cases}$$

The original problem therefore reduces to

$$\min_{m,s>0} \int_0^1 \left[\frac{1}{2}\|\dot{m}_t\|^2 + \frac{d}{2}\left(\dot{s}_t - \frac{\nu}{s_t}\right)^2 - U(m_t, s_t) \right] \mathrm{d}t,$$

with $m_0 = -2.5e_1$, $m_1 = 2.5e_1$, $s_0 = \sqrt{0.1}$, and $s_1 = \sqrt{0.05}$.

Expanding the kinetic term for the standard deviation gives

$$\frac{d}{2}\left(\dot{s}_t - \frac{\nu}{s_t}\right)^2 = \frac{d}{2}\dot{s}_t^2 + \frac{d\nu^2}{2s_t^2} - d\nu \frac{\mathrm{d}}{\mathrm{d}t}\log s_t.$$

The integral of the last term is $-d\nu \log(s_1/s_0)$, which depends only on the fixed endpoints and therefore does not affect the optimal path. Omitting this term gives the Lagrangian

$$L(m, s, \dot{m}, \dot{s}) = \frac{1}{2}\|\dot{m}\|^2 + \frac{d}{2}\dot{s}^2 + \frac{d\nu^2}{2s^2} - U(m, s).$$

For the mean variable, $\partial_{\dot{m}} L = \dot{m}$ and $\nabla_m L = -\nabla_m U$. Substituting these into the Euler–Lagrange equation $\frac{\mathrm{d}}{\mathrm{d}t}\partial_{\dot{m}} L - \nabla_m L = 0$ gives

$$\ddot{m}_t + \nabla_m U(m_t, s_t) = 0.$$

For the standard deviation variable, $\partial_{\dot{s}} L = d\dot{s}$, while

$$\partial_s L = \frac{d\nu^2}{2}(-2s^{-3}) - \partial_s U = -\frac{d\nu^2}{s^3} - \partial_s U.$$

Therefore, $\frac{\mathrm{d}}{\mathrm{d}t}\partial_{\dot{s}} L - \partial_s L = 0$ gives

$$d\ddot{s}_t + \frac{d\nu^2}{s_t^3} + \partial_s U(m_t, s_t) = 0.$$

Equivalently,

$$\ddot{m}_t = -\nabla_m U(m_t, s_t), \qquad \ddot{s}_t = -\frac{\nu^2}{s_t^3} - \frac{1}{d}\partial_s U(m_t, s_t).$$

For both interaction potentials, the mean satisfies $\ddot{m}_t = 0$, so $m_t = (1-t)m_0 + t m_1$. Under the linear external potential, $\ddot{m}_t = f$, giving $m_t = (1-t)m_0 + t m_1 - \frac{1}{2}t(1-t)f$. Under the quadratic external potential, $\ddot{m}_t = -\kappa m_t$, yielding

$$m_t = \frac{\sin(\sqrt{\kappa}(1-t))}{\sin(\sqrt{\kappa})} m_0 + \frac{\sin(\sqrt{\kappa}t)}{\sin(\sqrt{\kappa})} m_1.$$

The standard deviation s_t satisfies the two-point boundary-value problem

$$\ddot{s}_t = \gamma s_t - \frac{\nu^2}{s_t^3}, \qquad s_0 = \sqrt{0.1}, \quad s_1 = \sqrt{0.05},$$

where $\gamma = -\kappa$ for the attractive interaction and quadratic external potential, $\gamma = \kappa$ for the repulsive interaction, and $\gamma = 0$ for the linear external potential. After obtaining the positive solution s_t, the

covariance is $\Sigma_t^\star = s_t^2 I_d$. We directly solve this second-order boundary-value problem for s_t and then compute the optimal control law

$$b^\star(t,x) = \dot{m}_t + \left(\frac{\dot{s}_t}{s_t} - \frac{\nu}{s_t^2}\right)(x - m_t),$$

which is substituted back into the full action to obtain

$$\mathcal{A}^\star = \int_0^1 \left[\frac{1}{2}\|\dot{m}_t\|^2 + \frac{d}{2}\left(\dot{s}_t - \frac{\nu}{s_t}\right)^2 - U(m_t, s_t)\right] \mathrm{d}t.$$

Table 7 reports the ground-truth actions used to compute the absolute action biases in Table 1.

Table 7: Ground-truth actions $\mathcal{A}^\star$ for the synthetic optimal-control benchmarks.

Potential setting	**10D**	**20D**
Attractive Interaction	10.511	8.522
Repulsive Interaction	14.129	15.758
Linear external	12.845	13.231
Quadratic external	6.037	4.049

B.3.2 Construction of Synthetic Datasets in Experiment 3

We first describe the initial-data generation shared by both datasets. In each case, the initial distribution is a six-cluster Gaussian mixture in $\mathbb{R}^d$. We sample a $d \times d$ matrix from a standard Gaussian distribution and obtain an orthogonal matrix through QR decomposition. Its first two rows are taken as orthonormal unit vectors q_1, q_2. The six cluster centers are evenly placed on a circle of radius $L = 1$:

$$m_k = L\left[\cos\left(\frac{2\pi k}{6}\right) q_1 + \sin\left(\frac{2\pi k}{6}\right) q_2\right], \qquad k = 0, \ldots, 5.$$

Particles are assigned sequentially to the six clusters, and isotropic Gaussian noise is added around the corresponding centers:

$$Y_i = m_{i \bmod 6} + \sigma\epsilon_i, \qquad \sigma = \frac{L}{c\sqrt{d}}, \qquad \epsilon_i \overset{\text{iid}}{\sim} \mathcal{N}(0, I_d), \qquad i = 0, \ldots, N-1.$$

Here N is the number of particles and $c > 0$ is a concentration parameter. Larger c produces more compact clusters, and the root-mean-square radius of the within-cluster noise is L/c. Finally, we shift the empirical centroid of all particles to the origin:

$$X_i(0) = Y_i - \frac{1}{N}\sum_{j=0}^{N-1} Y_j.$$

Two-signal Keller–Segel Equation. Inspired by the Keller–Segel model of cellular chemotaxis (Keller & Segel, 1970), this model uses two signals with different interaction ranges to capture both local repulsion and long-range attraction. We use the two-scale Gaussian interaction potential

$$W_{\mathrm{KS}}(r) = B\exp\left(-\frac{r^2}{2\ell_r^2}\right) - A\exp\left(-\frac{r^2}{2\ell_a^2}\right),$$

where r is the inter-particle distance, $B = 0.8$ and $\ell_r = 0.22$ control short-range repulsion, and $A = 1.2$ and $\ell_a = 0.9$ control longer-range attraction. We additionally introduce a constant drift $u = 0.55e_1$, where e_1 is the unit vector along the first coordinate. The corresponding diffusion-free mean-field PDE is

$$\partial_s \rho_s = \nabla \cdot [\rho_s \nabla (W_{\mathrm{KS}} * \rho_s)] - u \cdot \nabla \rho_s.$$

Here ρ_s is the particle distribution at physical time s, and the convolution is defined as $(W_{\mathrm{KS}} * \rho_s)(x) = \int W_{\mathrm{KS}}(\|x-y\|)\rho_s(\mathrm{d}y)$. Let $r_{ij} = \|X_i - X_j\|$. Substituting the equally weighted empirical distribution and excluding self-interactions gives the particle dynamics

$$\frac{\mathrm{d}X_i}{\mathrm{d}s} = u + \frac{1}{N}\sum_{j\neq i}\left[\frac{B}{\ell_r^2}e^{-r_{ij}^2/(2\ell_r^2)} - \frac{A}{\ell_a^2}e^{-r_{ij}^2/(2\ell_a^2)}\right](X_i - X_j).$$

For this experiment, we use $d = 20$, $N = 4096$, and concentration $c = 5$. Samples are initialized according to the six-cluster construction above, producing dynamics that combine local expansion, inter-cluster attraction, and global translation. Physical time is $s \in [0, 3]$ and is normalized as $t = s/3$. Only the endpoint marginals at $t = 0, 1$ are used for training, while $t = 0.25, 0.5, 0.75$ are held out for interpolation evaluation.

Coulomb Particles with an Off-center Harmonic External Potential. This model describes the overdamped motion of charged particles in an off-center harmonic trap: the external potential drives particles toward the trap center, while pairwise repulsion promotes local expansion. We define the external and regularized pairwise potentials as

$$V(x) = \frac{a}{2}\|x - x_{\mathrm{c}}\|^2, \qquad W_{\mathrm{C}}(r) = \frac{\beta}{d-2}(r^2+\varepsilon^2)^{-(d-2)/2},$$

where $a = 1.5$ is the external-potential strength, $x_{\mathrm{c}} = 1.5e_1$ is the trap center, and $\varepsilon = 0.35$ is the regularization scale. To avoid directly specifying an interaction coefficient that is difficult to interpret, we choose a reference distance $r_\star = 0.3$ and set

$$\beta = a(r_\star^2 + \varepsilon^2)^{d/2}.$$

With this choice, at distance $r_\star$, the radial coefficient $\beta/(r_\star^2+\varepsilon^2)^{d/2}$ of the unit-mass interaction kernel equals the restoring-force coefficient a of the external potential. The corresponding diffusion-free mean-field PDE is

$$\partial_s \rho_s = \nabla \cdot \left[\rho_s \nabla \left(V + W_{\mathrm{C}} * \rho_s\right)\right],$$

where $(W_{\mathrm{C}} * \rho_s)(x) = \int W_{\mathrm{C}}(\|x-y\|)\rho_s(\mathrm{d}y)$. Substituting the equally weighted empirical distribution and excluding self-interactions gives the particle dynamics

$$\frac{\mathrm{d}X_i}{\mathrm{d}s} = -a(X_i - x_{\mathrm{c}}) + \frac{\beta}{N}\sum_{j\neq i}\frac{X_i - X_j}{(r_{ij}^2+\varepsilon^2)^{d/2}}.$$

For this experiment, we use $d = 20$, $N = 1024$, and concentration $c = 16$, producing more compact initial clusters using the same six-cluster initialization. Physical time is $s \in [0, 1]$ and we set $t = s$. Again, only the endpoint marginals at $t = 0, 1$ are used for training, while $t = 0.25, 0.5, 0.75$ are held out for interpolation evaluation.

Figure 4 visualizes the ground-truth particle trajectories for both datasets.

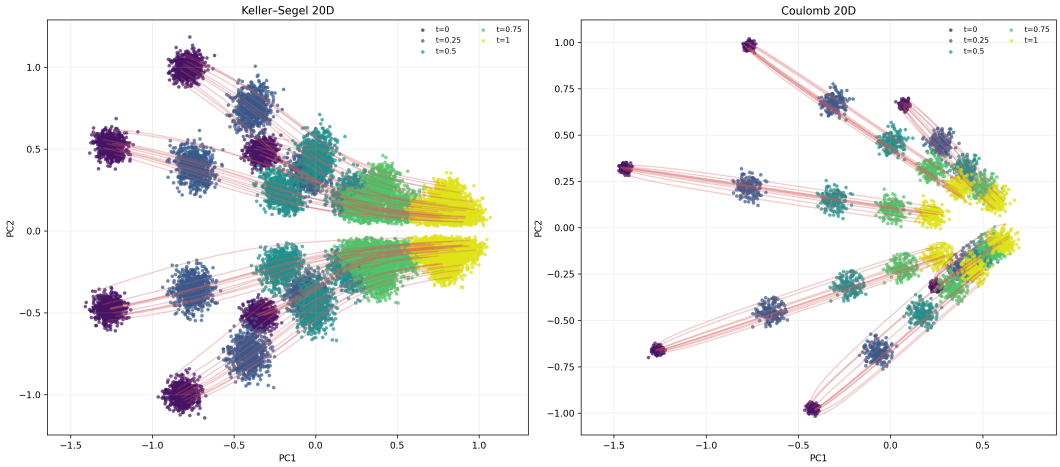

Figure 4: Ground-truth particle trajectories for the 20D mean-field datasets: Keller–Segel (left) and Coulomb (right), projected onto the first two principal components. Colors indicate normalized times $t = 0, 0.25, 0.5, 0.75, 1$, and red curves show representative particle trajectories. Only the endpoint marginals are used for training; the three intermediate time points are held out for interpolation evaluation.

B.4 Real Datasets

B.4.1 Real Datasets in Experiment 2

EB (Embryoid Bodies). The EB dataset comes from the human embryonic stem-cell embryoid-body differentiation experiments of Moon et al. (Moon et al., 2019). It records transcriptional-state changes as cells differentiate into different germ layers and their derived lineages over 27 days of culture, providing an in vitro model for studying early embryonic development and multilineage fate decisions. The dataset used here follows the WFR-FM preprocessing (Peng et al., 2026a) and contains 16,819 cells across 5 time groups, with 2,381, 4,163, 3,278, 3,665, and 3,332 cells, respectively. We conduct experiments using 5D, 50D, and 100D representations to evaluate marginal reconstruction across different representation dimensions.

CITE. The CITE dataset records the in vitro differentiation of human $CD34^+$ hematopoietic stem and progenitor cells used by SF^2M (Tong et al., 2024b) and WFR-FM (Peng et al., 2026a), with four observations on culture days 2, 3, 4, and 7. CITE-seq (Stoeckius et al., 2017) jointly measures the transcriptome and cell-surface proteins; in this experiment, we use only the gene-expression modality with a 50D PCA representation. The dataset used here contains 31,240 cells, with 7,476, 6,999, 9,511, and 7,254 cells at the four time points, respectively.

EMT (Epithelial–Mesenchymal Transition). The EMT dataset comes from the cancer cell-state transition experiments of Cook and Vanderhyden (Cook & Vanderhyden, 2020). We use their TGFB1-induced A549 lung cancer cell data, which describe transcriptional changes as cells transition from epithelial states through intermediate states to mesenchymal states. The dataset contains 3,133 cells across 4 observation-time groups, with 577, 885, 788, and 883 cells, respectively. We use the same 10D autoencoder representation as WFR-FM (Peng et al., 2026a).

Weinreb. The Weinreb dataset comes from an in vitro differentiation experiment on mouse hematopoietic stem and progenitor cells (Weinreb et al., 2020). The original study used LARRY, which combines heritable expressed DNA barcodes with single-cell transcriptomic measurements, to trace clonal descendants and study the relationship between early transcriptional states and subsequent cell fates. The dataset used here follows the WFR-FM setup (Peng et al., 2026a) and contains 49,302 cells with lineage-tracing information across 3 time points, with 4,638, 14,985, and 29,679 cells, respectively. We use 10D, 50D, and 100D PCA representations for distribution reconstruction experiments.

B.4.2 Real Dataset in Experiment 3

Axolotl (Spatial Transcriptomics of Brain Regeneration). This dataset comes from the axolotl telencephalon injury-regeneration experiments of Wei et al. (Wei et al., 2022) and is also the spatial transcriptomics dataset used by stVCR (Peng et al., 2026b). After partial resection of the dorsal tissue of the left telencephalon, the original study jointly measured spatial cell locations and gene expression using Stereo-seq to investigate post-injury cell-state changes, neuronal replenishment, and spatial tissue reconstruction. The public data record is STDS0000056. The dataset used here contains 66,425 cells across 7 regeneration time groups, with 7,668, 8,106, 9,440, 9,676, 11,319, 9,252, and 10,964 cells, respectively. We concatenate two-dimensional spatial coordinates with 30-dimensional expression features to obtain a 32D representation. Using the time labels provided by the dataset, we train independently on the intervals $0.2 \to 1$, $1 \to 2$, and $2 \to 6$, while holding out times 0.5, 1.5, and 3 for interpolation evaluation.

B.5 Adapting the Method to the Spatial Transcriptomics Dataset

The external potential $V(x)$ and pairwise potential $W(x, y)$ in the main text depend only on the current particle states, and their expectations are determined by the current marginal distribution ρ_t. The deformation penalty used in the spatial transcriptomics experiment additionally compares the spatial distance of the same particle pair at the current and initial times, and therefore depends on the correspondence between each particle and its initial state. To accommodate this dependence, we retain the original action structure and augment the state with the initial spatial coordinates, which remain fixed over time. We describe the formulation below on the interval $[0, 1]$.

Let $x \in \mathbb{R}^d$ denote the current cell state, including its two-dimensional spatial coordinates and expression features, and let $\mathsf{S} : \mathbb{R}^d \to \mathbb{R}^2$ be the projection that extracts the spatial coordinates. Let $\xi = \mathsf{S}X_0 \in \mathbb{R}^2$ denote the particle's initial spatial coordinate, and let $\eta_t(\mathrm{d}\xi, \mathrm{d}x)$ denote the joint distribution of ξ and the current state X_t. Let μ_0 and μ_1 denote the prescribed initial and target state distributions, respectively, and let μ_0^{sp} denote the distribution of spatial coordinates extracted from μ_0. The joint distribution can be written as

$$\eta_t(\mathrm{d}\xi, \mathrm{d}x) = \mu_0^{\mathrm{sp}}(\mathrm{d}\xi)\, p_t(\mathrm{d}x \mid \xi),$$

where $p_t(\mathrm{d}x \mid \xi)$ denotes the distribution of the particle state at time t conditioned on the initial spatial coordinate ξ. Integrating out ξ gives the current state distribution, while integrating out x gives the initial spatial distribution, which remains fixed over time:

$$\rho_t(\mathrm{d}x) = \int_{\xi \in \mathbb{R}^2} \eta_t(\mathrm{d}\xi, \mathrm{d}x), \qquad \int_{x \in \mathbb{R}^d} \eta_t(\mathrm{d}\xi, \mathrm{d}x) = \mu_0^{\mathrm{sp}}(\mathrm{d}\xi).$$

Allowing the control to depend on the initial label, the controlled process becomes

$$\mathrm{d}X_t = b_t(\xi, X_t)\mathrm{d}t + \sqrt{2\nu}\mathrm{d}B_t, \qquad \xi = \mathsf{S}X_0,$$

where $b_t(\xi, x) \in \mathbb{R}^d$ is the controlled drift, $\nu > 0$ is the diffusion parameter, and B_t is a d-dimensional standard Brownian motion. The initial label ξ remains fixed throughout the process, while only the current state x is transported and diffused.

We continue to use W to denote the pairwise kernel, but extend its inputs to two augmented states (ξ, x) and (ξ', y), where ξ' and y denote the initial spatial coordinate and current state of another particle, respectively. The augmented optimal control problem is

$$\inf_{\eta, b} \mathcal{A}(\eta, b) := \int_0^1 \left[\frac{1}{2} \int \|b_t(\xi, x)\|^2 \eta_t(\mathrm{d}\xi, \mathrm{d}x) - \frac{1}{2} \iint W((\xi, x), (\xi', y)) \eta_t(\mathrm{d}\xi, \mathrm{d}x) \eta_t(\mathrm{d}\xi', \mathrm{d}y) \right] \mathrm{d}t.$$

The integration domain of a single augmented state is $\mathbb{R}^2 \times \mathbb{R}^d$, and the pairwise integral is taken over two copies of this space. The joint distribution satisfies the Fokker–Planck equation

$$\partial_t \eta_t + \nabla_x \cdot (\eta_t b_t) = \nu \Delta_x \eta_t,$$

where $\nabla_x \cdot$ and Δ_x denote the divergence and Laplacian with respect to the current state x only.

The endpoint constraints are

$$\eta_0(\mathrm{d}\xi, \mathrm{d}x) = \delta_{\mathsf{S}x}(\mathrm{d}\xi)\mu_0(\mathrm{d}x), \qquad \int_{\xi \in \mathbb{R}^2} \eta_1(\mathrm{d}\xi, \mathrm{d}x) = \mu_1(\mathrm{d}x).$$

Here $\delta_{\mathsf{S}x}$ is the Dirac probability measure concentrated at $\mathsf{S}x$. Thus, initially x follows μ_0 and its label equals $\mathsf{S}x$, while the terminal current-state marginal equals μ_1.

These constraints preserve the prescribed endpoint marginals.

For the spatial transcriptomics experiment, we define

$$W((\xi, x), (\xi', y)) = -\frac{\lambda}{2}\left[\|\mathsf{S}x - \mathsf{S}y\|^2 - \|\xi - \xi'\|^2\right]^2,$$

where $\lambda > 0$ controls the strength of the spatial deformation penalty; we use $\lambda = 7$ in this experiment. Under the kinetic-minus-potential convention, the pairwise term in the action becomes the nonnegative penalty

$$\mathcal{R}(\eta_t) = \frac{\lambda}{4} \iint \left[\|\mathsf{S}x - \mathsf{S}y\|^2 - \|\xi - \xi'\|^2\right]^2 \eta_t(\mathrm{d}\xi, \mathrm{d}x)\eta_t(\mathrm{d}\xi', \mathrm{d}y).$$

For N equally weighted particles, let $X_i(t)$ denote the current state of particle i, and let $\xi_i = \mathsf{S}X_i(0)$ denote its initial spatial coordinate. Then

$$\mathcal{R}(\eta_t) = \frac{\lambda}{4N(N-1)} \sum_{i \neq j} \left[\|\mathsf{S}X_i(t) - \mathsf{S}X_j(t)\|^2 - \|\xi_i - \xi_j\|^2\right]^2.$$

This term penalizes changes in pairwise squared distances relative to the initial configuration, while remaining invariant to global translations or rotations that preserve these distances. It depends on the joint distribution η_t of the initial and current states rather than only on the current marginal ρ_t.

This extension preserves the basic alternating optimization structure between the transport plan and conditional probability paths. For an endpoint pair (x_0, x_1), the initial spatial label is directly determined by $\xi = \mathsf{S}x_0$. We therefore only need to replace the original interaction kernel with the augmented-state kernel, while the construction of the conditional-bridge drift and the KL computation relative to the Brownian bridge remain unchanged. Finally, we use the initial spatial position ξ as an additional condition for the neural networks and learn the conditional PF-ODE velocity field and score using $v_\theta(t, x, \xi)$ and $s_\psi(t, x, \xi)$, respectively, where θ and ψ are network parameters. During training, we use the same conditional regression loss as in the main text. During inference, each particle's initial spatial position ξ is fixed, and its subsequent states are generated using either the PF-ODE or SDE.

B.6 Modified Baselines

B.6.1 Modified CytoBridge

To adapt CytoBridge (Zhang et al., 2025d) to the IS-WLF problem studied in this work, we retain its original velocity network, score network, and three-stage training procedure: drift pretraining, score-network training, and joint fine-tuning.

We disable the Growth branch and no longer model mass variation. Therefore, all particles use uniform weights, and the OT constraints during training use balanced OT rather than the unbalanced OT formulation (Chizat et al., 2018b) and its Sinkhorn-type solver. We also move the interaction from an explicit interaction force in the dynamical equation into the action. The interaction is no longer added directly to the particle drift, but instead influences network training through the potential-energy term.

The modified CytoBridge solves

$$\min_{\rho,b} \int_0^T \left[\frac{1}{2} \int_{\mathbb{R}^d} \|b_t(x)\|^2 \rho_t(x)\,\mathrm{d}x - \int_{\mathbb{R}^d} V(x)\rho_t(x)\,\mathrm{d}x \right. \\ \left. - \frac{1}{2} \iint_{\mathbb{R}^d \times \mathbb{R}^d} W(x,y)\rho_t(x)\rho_t(y)\,\mathrm{d}x\,\mathrm{d}y \right] \mathrm{d}t.$$

This optimization problem is subject to the endpoint constraints and Fokker–Planck equation

$$\rho_0 = \mu_0, \qquad \rho_T = \mu_T, \qquad \partial_t \rho_t = -\nabla \cdot (\rho_t b_t) + \nu \Delta \rho_t.$$

This is exactly the IS-WLF problem considered in this work.

B.6.2 Modified WLF

The algorithm in the main text of WLF (Neklyudov et al., 2024) writes the Hamiltonian in a form that is linear or linearizable in the density ρ_t, and therefore cannot be directly applied to pairwise interactions. Instead, we start from Eq. (31) in its appendix. Let $\langle f, \mu \rangle = \int f(x)\,\mathrm{d}\mu(x)$, and define the endpoint term $\mathcal{B}(s) = \langle s_1, \mu_1 \rangle - \langle s_0, \mu_0 \rangle$ and the time-integral term $\mathcal{I}(\rho, s) = \int_0^1 [\langle \partial_t s_t, \rho_t \rangle + \mathcal{H}[\rho_t, s_t, t]]\,\mathrm{d}t$. Equation (31) is

$$\inf_{\rho \in \Gamma(\mu_0, \mu_1)} \sup_s \{\mathcal{B}(s) - \mathcal{I}(\rho, s)\}.$$

This formulation is valid for arbitrary potential energies. The relevant Hamiltonian has the form kinetic energy plus potential energy. After applying the Fokker–Planck reparameterization used in the original paper for the Schrödinger Bridge, the time-integral terms become $\mathbb{E}_{\rho_t}[\partial_t \Phi + \frac{1}{2}\|\nabla_x \Phi\|^2 + \nu \Delta_x \Phi + V]$, together with the pairwise energy $\mathcal{U}(\rho_t) = \frac{1}{2} \iint W(x,y)\rho_t(x)\rho_t(y)\,\mathrm{d}x\,\mathrm{d}y$. Here Φ is the drift potential of the forward SDE, and V is an optional external potential.

All these terms can be estimated by Monte Carlo sampling. The endpoint terms are estimated using sample averages from μ_0 and μ_1, respectively. For the time-integrated terms, let $g = 1, \ldots, G$ index time samples t_g, and let $i = 1, \ldots, m$ index particles x_{gi} generated by the network at that time. The terms $\partial_t \Phi$, $\frac{1}{2}\|\nabla_x \Phi\|^2$, and V are averaged over these Gm points, while $\nu \Delta_x \Phi$ is estimated using $\nu \varepsilon^\top \nabla_x^2 \Phi\, \varepsilon$ with a random vector ε. The pairwise potential term uses distinct particle pairs from the same time point, with $i, j = 1, \ldots, m$ and $i \neq j$:

$$\widehat{\mathcal{U}} = \frac{1}{2Gm(m-1)} \sum_{g=1}^{G} \sum_{i \neq j} W(x_{gi}, x_{gj}).$$

Equation (31) can then be optimized using these sample estimates; the SDE drift after training is $b = \nabla_x \Phi$.

B.7 Ablation Studies and Sensitivity Analysis

B.7.1 Effectiveness of Jointly Updating the Transport Plan and Conditional Probability Paths

To evaluate the effect of joint updates, we compare the full CCB model with an ablation that fixes the transport plan and updates only the conditional probability paths on the four 20D synthetic datasets from Experiment 1. The full model recomputes the transport costs and updates the transport plan as the conditional paths change, allowing endpoint matching and path optimization to adapt to each other. We use the same absolute action bias and control-law error as in Experiment 1. Table 8 reports the results.

Table 8: Effect of jointly updating the transport plan and conditional probability paths.

Setting	Fixed: Act.	CCB: Act.	Fixed: Ctrl.	CCB: Ctrl.
Attractive Interaction	0.211 ± 0.0028	$\mathbf{0.081 \pm 0.0033}$	0.058 ± 0.0007	$\mathbf{0.038 \pm 0.0007}$
Repulsive Interaction	0.086 ± 0.0079	$\mathbf{0.077 \pm 0.0063}$	0.014 ± 0.0002	$\mathbf{0.013 \pm 0.0002}$
Linear external	0.148 ± 0.0053	$\mathbf{0.147 \pm 0.0033}$	0.018 ± 0.0003	$\mathbf{0.012 \pm 0.0002}$
Quadratic external	0.209 ± 0.0036	$\mathbf{0.072 \pm 0.0067}$	0.056 ± 0.0009	$\mathbf{0.037 \pm 0.0009}$

Full CCB achieves lower mean errors on both metrics across all four datasets, with particularly large improvements in action bias for the attractive quadratic interaction and quadratic external potential settings. This indicates that optimizing conditional paths under a fixed transport plan is insufficient to match the performance of the full model. Adjusting the transport costs and transport plan according to the updated paths further reduces action bias and improves reconstruction of the optimal control law.

B.7.2 Sensitivity to the Number of Edges in Sparse Mode

To study how sparse-support size affects solution accuracy and computational cost, we construct a 50D synthetic dataset with an attractive quadratic pairwise interaction potential and a total of 20,000 particles. We jointly vary the number of proposed edges q and retained edges K per source and target particle at each iteration, and compare training time, absolute action error, and control-law error across different configurations.

Table 9: Sensitivity to the sparse-support size on the 50D attractive-interaction dataset.

q	K	Time (s)	Action	Act. error	Ctrl. MSE
32	8	324.9	2.884 ± 0.0121	0.328 ± 0.0121	0.060 ± 0.0004
43	11	386.2	2.856 ± 0.0143	0.300 ± 0.0143	0.057 ± 0.0004
64	**16**	**499.1**	2.801 ± 0.0113	0.245 ± 0.0113	0.055 ± 0.0002
128	32	834.4	2.746 ± 0.0154	0.190 ± 0.0154	0.051 ± 0.0003
256	64	1509.8	2.657 ± 0.0143	0.101 ± 0.0143	0.047 ± 0.0003
384	96	2174.9	2.634 ± 0.0160	0.078 ± 0.0160	0.046 ± 0.0002
640	160	3461.3	2.587 ± 0.0128	0.031 ± 0.0128	0.042 ± 0.0003

As the numbers of proposed and retained edges increase, the mean values of both errors decrease consistently, but training time also increases substantially. Increasing (q, K) from $(32, 8)$ to $(640, 160)$ reduces the absolute action error from 0.328 to 0.031 and the control-law MSE from 0.060 to 0.042, while increasing training time from 324.9 s to 3461.3 s. These results show that enlarging the candidate-edge set and retained support improves solution accuracy at additional computational cost. Balancing accuracy and efficiency, we recommend $(q, K) = (64, 16)$ as the standard configuration; tasks that prioritize solution accuracy can further increase the number of edges.

B.7.3 Runtime Across Data Dimensions

We compare the training times of CCB, GSBM (Liu et al., 2024), and Modified WLF (Neklyudov et al., 2024) with interaction potentials on the 5D, 50D, and 100D representations of the EB dataset. All experiments are conducted on a single NVIDIA A100 GPU.

Table 10: Training time in seconds on the EB dataset.

Data	CCB	GSBM	Modified WLF
EB (5D)	**755.752**	974.093	1117.395
EB (50D)	**780.563**	881.587	1112.743
EB (100D)	**838.174**	895.461	1117.881

The results show that CCB has the shortest training time at all three dimensions. From 5D to 100D, its runtime increases by only approximately 10.9%. The runtime of Modified WLF remains nearly unchanged with dimensionality: its network width and depth are fixed, so increasing the dimension mainly affects input- and output-related computations; moreover, each iteration samples only a fixed number of particles and estimates the energy by Monte Carlo, so the per-iteration computational cost does not increase substantially with the full dataset size. These results show that CCB has relatively low training overhead in these settings and can accommodate higher-dimensional data representations.

C Additional Discussion

C.1 Connections to Related Algorithms

C.1.1 Relation to GSBM

Different coupling-update mechanisms. GSBM (Liu et al., 2024) updates its coupling from the endpoint joint distribution of trajectories simulated under the current dynamics, then reoptimizes the conditional paths and regresses the drift. Its implementation alternates forward and backward directions: an SDE rollout generates the coupling, the paths are optimized, and the opposite drift is trained. Thus, although Gaussian conditional-path fitting can be simulation-free, the overall procedure still requires repeated rollouts. CCB instead explicitly optimizes the transport plan by alternating a static OT subproblem with conditional-path optimization, and regresses the velocity and score only after both have been optimized. No dynamical rollout is used to update its coupling.

Different conditional-bridge costs. Proposition 2 of GSBM defines the conditional-bridge kinetic energy by directly integrating the squared conditional SDE drift:

$$K^z = \frac{1}{2}\int_0^1 \mathbb{E}_{\mathbb{P}^z}\|b_t^z(X_t)\|^2\,\mathrm{d}t,$$

where $z=(x_0,x_1)$, $\mathrm{d}X_t = b_t^z(X_t)\mathrm{d}t + \sqrt{2\nu}\mathrm{d}B_t$, and $X_0=x_0$, $X_1=x_1$; see Proposition 2 and Eq. (6) of GSBM (Liu et al., 2024). For $\nu>0$, exact endpoint conditioning makes the control singular and this energy divergent. Indeed, the reference Brownian bridge has drift $\beta_t^z(x)=(x_1-x)/(1-t)$ and instantaneous kinetic energy

$$\frac{1}{2}\mathbb{E}_{\mathbb{Q}^{\nu,z}}\|\beta_t^z(X_t)\|^2 = \frac{1}{2}\|x_1-x_0\|^2 + d\nu\frac{t}{1-t}.$$

Its time integral is $+\infty$. CCB instead combines the squared endpoint distance with a conditional path-space KL divergence. Appendix A.2 shows that, under an endpoint regularity condition, minimal-subtraction renormalization of the original kinetic energy yields exactly this part of the CCB cost, up to a fixed constant.

C.1.2 Relation to CytoBridge

CCB and CytoBridge (Zhang et al., 2025d) both model stochastic interacting-particle dynamics, but solve different problems.

Different locations of the interaction term. To use consistent notation, let the diffusion coefficient be $\sqrt{2\nu}$ and denote the additional control drift of CytoBridge by u_t. Its objective is

$$\inf_{\rho,u,g,\Phi}\int_0^1\int_{\mathbb{R}^d}\left[\frac{1}{2}\|u_t(x)\|^2 + \alpha\Psi(g_t(x))\right]\rho_t(x)\,\mathrm{d}x\,\mathrm{d}t,$$

where Ψ is the growth cost and α is its weight. Define the interaction drift

$$F_\Phi(\rho_t,x) = -\int_{\mathbb{R}^d} k(x,y)\nabla_x\Phi(x-y)\rho_t(y)\,\mathrm{d}y,$$

where Φ is the interaction potential and k is the interaction weight. The dynamical constraint is then

$$\partial_t\rho_t + \nabla\cdot\left[\rho_t\left(u_t + F_\Phi(\rho_t,\cdot)\right)\right] = \nu\Delta\rho_t + g_t\rho_t.$$

Thus, the interaction enters the Fokker–Planck equation directly, while the kinetic-energy term penalizes the additional control u_t rather than the total drift $u_t + F_\Phi(\rho_t, \cdot)$, as specified in Definition 4.1 of CytoBridge (Zhang et al., 2025d).

By contrast, CCB places the interaction in the action:

$$\mathcal{A}(\rho, b) = \int_0^1 \left[\frac{1}{2} \int \|b_t(x)\|^2 \rho_t(x)\,\mathrm{d}x - \int V(x)\rho_t(x)\,\mathrm{d}x \right. \\ \left. - \frac{1}{2} \iint W(x,y)\rho_t(x)\rho_t(y)\,\mathrm{d}x\,\mathrm{d}y \right] \mathrm{d}t.$$

It minimizes this action subject to $\partial_t \rho_t + \nabla \cdot (\rho_t b_t) = \nu \Delta \rho_t$, $\rho_0 = \mu_0$, and $\rho_1 = \mu_1$, where b_t is the total drift. The two methods therefore represent interactions through the dynamical constraint and optimization objective, respectively; they are not simply different solvers for the same control problem.

Different mass-conservation assumptions. CytoBridge (Zhang et al., 2025d) supports an unbalanced formulation, using the source term $g_t \rho_t$ to model cell proliferation and death and adding a growth cost to its objective. The current CCB formulation transports probability distributions with conserved mass and does not yet include birth or death processes.

Different research objectives. CytoBridge (Zhang et al., 2025d) learns cell-state transitions, growth rates, and an interaction potential from snapshots to infer unknown population dynamics. CCB instead treats interactions as given mechanism information or a regularization prior and solves the least-action problem under fixed endpoint marginals; it constructs the corresponding optimal control rather than identifying unknown interactions.

C.1.3 Relation to Wasserstein Lagrangian Flows

CCB and Wasserstein Lagrangian Flows (WLF) (Neklyudov et al., 2024) both minimize an action and avoid simulating dynamical trajectories during training, but they differ in their applicability and optimization strategies.

Different classes of admissible potential energies. The general WLF dual framework permits broad potential energies, but its sample-based solver requires a Hamiltonian that is linear in the density or can be linearized by a change of variables:

$$\mathcal{H}(\rho_t, s_t, t) = \int \mathbb{H}(x, s_t, t)\rho_t(x)\,\mathrm{d}x,$$

where $\mathbb{H}$ is independent of ρ_t. Since $\mathcal{H} = \mathcal{K}^* + \mathcal{U}$, the potential must obey the same requirement. External potentials $\mathcal{U}(\rho) = \int V(x)\rho(x)\,\mathrm{d}x$ do so, whereas general pairwise potentials $\mathcal{U}(\rho) = \frac{1}{2}\iint W(x,y)\rho(x)\rho(y)\,\mathrm{d}x\,\mathrm{d}y$ are quadratic in the density. CCB handles them through pairwise expectations between conditional paths, without linearizing the potential.

Different optimization structures and uses of couplings. WLF (Neklyudov et al., 2024) uses a GAN-like min–max structure: a generator learns $\rho_{\eta,t}$ while another network learns the dynamical dual potential $s_{\theta,t}$, without an explicit optimal endpoint coupling. Alternating descent and ascent can inherit GAN-like instability because insufficient inner dual optimization may distort the path update. CCB instead maintains a transport plan and alternates coupling and conditional-path updates under one minimization objective. It then samples endpoint pairs from the optimized coupling and learns velocity and score by conditional regression, avoiding adversarial min–max optimization.

C.1.4 Relation to A Call to Lagrangian Action

Different dynamical descriptions. The Wasserstein Lagrangian Mechanics (WLM) framework proposed in *A Call to Lagrangian Action* explicitly evolves particle positions and velocities through damped second-order dynamics:

$$\frac{\mathrm{d}X_t}{\mathrm{d}t} = v_t, \qquad \frac{\mathrm{d}v_t}{\mathrm{d}t} = -\nabla_x \frac{\delta \mathcal{U}}{\delta \rho}(\rho_t)(X_t) - \gamma v_t,$$

where $\mathcal{U}$ is the population potential energy and $\gamma \geq 0$ is the damping coefficient. This formulation retains inertia and includes both undamped systems and the overdamped limit (Guan et al., 2026). CCB instead uses a first-order controlled diffusion without an explicit inertial velocity variable,

$$\mathrm{d}X_t = b_t(X_t)\,\mathrm{d}t + \sqrt{2\nu}\,\mathrm{d}B_t,$$

and treats the drift b_t as the control when minimizing the action under a given potential energy.

Different unknowns and constraints. From observed marginals and a prescribed initial velocity, WLM (Guan et al., 2026) learns an unknown population potential and interaction mechanism by simulating second-order dynamics and matching later marginals, then uses the learned system for interpolation and extrapolation. CCB instead takes the potentials as known, imposes no initial-velocity condition, and solves for an optimal control under $\rho_0 = \mu_0$ and $\rho_1 = \mu_1$. Thus, WLM identifies a mechanical mechanism from data, whereas CCB solves least-action dynamics under a given mechanism and endpoint constraints.

C.2 Pseudocode of Constructive CytoBridge

Algorithm 1 Joint optimization of the transport plan and conditional probability paths

Require: Samples $\{x_0^i, a_i\}_{i=1}^n$, $\{x_1^j, b_j\}_{j=1}^m$; potentials V, W; $\nu > 0$, $\tau > 0$.
Require: Mode Dense or Sparse; iteration counts $N_{\mathrm{alt}}, N_P, N_\phi$; batch size B; path learning rate η_ϕ.
Require: Sinkhorn tolerance $\varepsilon_{\mathrm{marg}}$; sparse settings in Algorithm 3.
Ensure: Optimized transport plan P and conditional-path parameters ϕ.

Initialize a feasible transport plan $P \in \mathcal{U}(a, b)$ and conditional-path parameters ϕ.
Set $z_{ij} = (x_0^i, x_1^j)$, $\pi_P = \sum_{i,j} P_{ij}\delta_{z_{ij}}$, and $\gamma = \tau^{-1}$.
Let $\mathcal{E} = \mathrm{supp}(P)$ and $\rho_{\phi,t}^z = \mathcal{N}(m_{\phi,t}^z, \mathrm{diag}(q_{\phi,t}^z))$.
Set $\kappa_{\phi,t}^z = \frac{1}{2}[\|d_{\phi,t}^z\|^2 + \mathrm{tr}(H_{\phi,t}^z \Sigma_{\phi,t}^z (H_{\phi,t}^z)^\top)]$ as in Proposition 3.
for $r = 1, \ldots, N_{\mathrm{alt}}$ **do**
    Hold ϕ fixed while updating the transport plan.
    if mode is Sparse **then**
        Update $(P, \mathcal{E})$ once using Algorithm 3.
    end if
    for $u = 1, \ldots, N_P$ **do** ▷ Refine the plan on the fixed support.
        Estimate $\widehat{C}_{ij} = \frac{1}{2}\|x_1^j - x_0^i\|^2 + \widehat{\mathbb{E}}[\kappa_{\phi,t}^{z_{ij}} - V(X) - W(X, Y)]$ on $\mathcal{E}$,
        with $t \sim \mathcal{U}(0, 1)$, $X \sim \rho_{\phi,t}^{z_{ij}}$, $Z' \sim \pi_P$, and $Y \sim \rho_{\phi,t}^{Z'}$.
        $P \leftarrow \mathrm{Prox}_{\mathcal{E}}(\widehat{C}, P)$, using the KL-proximal solver in Algorithm 3.
    end for
    Hold P fixed while updating the conditional paths.
    for $u = 1, \ldots, N_\phi$ **do**
        Draw B samples of $t_b \sim \mathcal{U}(0, 1)$, $Z_b, Z_b' \overset{\mathrm{iid}}{\sim} \pi_P$, and $\epsilon_b, \epsilon_b' \overset{\mathrm{iid}}{\sim} \mathcal{N}(0, I_d)$.
        $X_b \leftarrow m_{\phi,t_b}^{Z_b} + \sqrt{q_{\phi,t_b}^{Z_b}} \odot \epsilon_b, \quad Y_b \leftarrow m_{\phi,t_b}^{Z_b'} + \sqrt{q_{\phi,t_b}^{Z_b'}} \odot \epsilon_b'$.
        $\widehat{\mathcal{J}} \leftarrow B^{-1} \sum_{b=1}^B [\kappa_{\phi,t_b}^{Z_b} - V(X_b) - \frac{1}{2} W(X_b, Y_b)]$.
        $\phi \leftarrow \phi - \eta_\phi \nabla_\phi \widehat{\mathcal{J}}$, differentiating through both X_b and Y_b.
    end for
end for
return P, ϕ.

Algorithm 2 Velocity and score regression with frozen conditional probability paths

Require: Optimized P, ϕ; $\nu > 0$; initial velocity/score parameters θ, ψ; weighting $\lambda(t) = \sqrt{2\nu t(1-t)}$.
Require: Steps N_{reg}; supervision batch size B; learning rates η_v, η_s; score-loss weight $\omega_s > 0$; variance floor $\varepsilon_q \geq 0$.
Ensure: PF-ODE velocity v_θ and score s_ψ.
Freeze P, ϕ; detach regression targets; interpret operations on q element-wise.
for $r = 1, \ldots, N_{\mathrm{reg}}$ **do**
  Sample pairs $Z_b \sim \pi_P$, stratified times $t_b \in (0,1)$, and fresh noise $\epsilon_b \sim \mathcal{N}(0, I_d)$.
  for $b = 1, \ldots, B$ **do**
    Evaluate or retrieve cached $m_b = m^{Z_b}_{\phi,t_b}$, $q_b = q^{Z_b}_{\phi,t_b}$, and derivatives $\dot{m}_b, \dot{q}_b$.
    $X_b \leftarrow m_b + \sqrt{q_b} \odot \epsilon_b$.
    $v_b^{\mathrm{tar}} \leftarrow \dot{m}_b + \frac{1}{2}[\dot{q}_b/(q_b + \varepsilon_q)] \odot (X_b - m_b)$.
    $s_b^{\mathrm{tar}} \leftarrow -(X_b - m_b)/(q_b + \varepsilon_q)$.
  end for
  $\widehat{\mathcal{L}}_v \leftarrow (Bd)^{-1} \sum_{b=1}^{B} \|v_\theta(t_b, X_b) - v_b^{\mathrm{tar}}\|^2$.
  $\widehat{\mathcal{L}}_s \leftarrow (Bd)^{-1} \sum_{b=1}^{B} \lambda(t_b)^2 \|s_\psi(t_b, X_b) - s_b^{\mathrm{tar}}\|^2$.
  Backpropagate $\widehat{\mathcal{L}}_v + \omega_s \widehat{\mathcal{L}}_s$; clip gradients separately for θ and ψ.
  Update θ, ψ with AdamW using learning rates η_v, η_s.
end for
return v_θ, s_ψ; the SDE drift is $b_{\theta,\psi}(t,x) = v_\theta(t,x) + \nu s_\psi(t,x)$.

Algorithm 3 Sparse OT with support expansion and bidirectional Top-K pruning

Require: Feasible sparse plan P with support $\mathcal{E}$; all samples and marginals a, b; fixed path parameters ϕ and potentials V, W.
Require: $\nu > 0$, $\gamma = \tau^{-1} > 0$, exploration weight $\delta \in (0,1)$; proposal count q_{prop}; retention count K; marginal tolerance $\varepsilon_{\mathrm{marg}}$.
Ensure: Updated sparse plan and support, satisfying the marginal tolerance.
Propose q_{prop} uniformly random neighbors per source and target, obtaining new edges $\mathcal{E}_{\mathrm{new}}$.
Expand the support: $\mathcal{E}^+ \leftarrow \mathcal{E} \cup \mathcal{E}_{\mathrm{new}}$.
Treat $P_{ij} = 0$ on newly added edges.
for each $(i,j) \in \mathcal{E}^+$ **do**
  $\bar{P}_{ij} \leftarrow (1-\delta)P_{ij} + \delta a_i b_j$.
  Estimate $\widehat{C}_{ij} = \frac{1}{2}\|x_1^j - x_0^i\|^2 + \widehat{\mathbb{E}}[\kappa^{z_{ij}}_{\phi,t} - V(X) - W(X,Y)]$,
  with $t \sim \mathcal{U}(0,1)$, $X \sim \rho^{z_{ij}}_{\phi,t}$, $Z' \sim \pi_P$, and $Y \sim \rho^{Z'}_{\phi,t}$.
end for
$P^+ \leftarrow \mathrm{Prox}_{\mathcal{E}^+}(\widehat{C}, \bar{P})$.
Retain bidirectional Top-K edges by mass: $\mathcal{E}_{\mathrm{keep}} \leftarrow \mathrm{TopK}_{\mathrm{row}}(P^+) \cup \mathrm{TopK}_{\mathrm{col}}(P^+)$.
Attempt $P_{\mathrm{keep}} \leftarrow \mathrm{Prox}_{\mathcal{E}_{\mathrm{keep}}}(\widehat{C}, P^+)$ with the same costs restricted to $\mathcal{E}_{\mathrm{keep}}$.
if the solve reaches the marginal tolerance **then**
  return $P_{\mathrm{keep}}, \mathcal{E}_{\mathrm{keep}}$.
else
  return $P^+, \mathcal{E}^+$ ▷ Reject pruning and retain the expanded plan.
end if
function $\mathrm{Prox}_{\mathcal{E}}(C, R)$
  $A_{ij} \leftarrow (a_i b_j)^{\frac{2\nu}{2\nu+\gamma}} R_{ij}^{\frac{\gamma}{2\nu+\gamma}} \exp(-C_{ij}/(2\nu+\gamma))$ for $(i,j) \in \mathcal{E}$.
  Set weights outside $\mathcal{E}$ to zero and apply log-domain Sinkhorn scaling to marginals a, b.
  return $\widehat{P}$ if $\frac{1}{2}\max\{\|\widehat{P}\mathbf{1}_m - a\|_1, \|\widehat{P}^\top \mathbf{1}_n - b\|_1\} \leq \varepsilon_{\mathrm{marg}}$; otherwise report failure.
end function

D Detailed Results by Time Point

D.1 Marginal Reconstruction on Real Single-cell Datasets

Tables 11–16 report W_1 and W_2 at individual observation times, complementing the time-averaged W_1 results in Table 2. The time indices label the evaluated observations after the initial snapshot. Lower values indicate better reconstruction. Deterministic methods are reported as single values; stochastic methods are reported as mean $\pm$ standard deviation over four rollouts. Method names and ordering follow the main text. For each metric, the lowest mean at each time point is highlighted in bold.

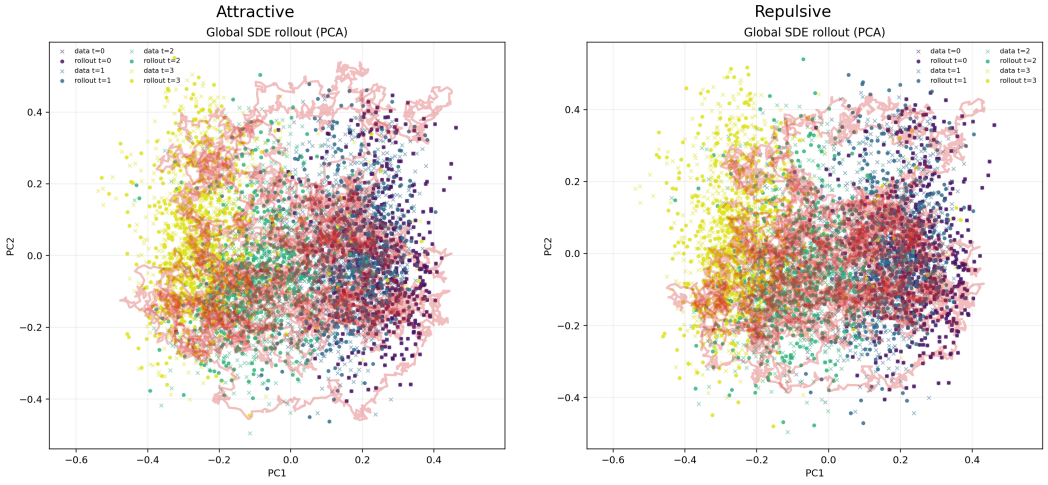

Figure 5: SDE rollouts of CCB on EMT (10D) with attractive (left) and repulsive (right) interactions, starting from the initial snapshot and visualized in a global PCA projection. Crosses denote observed data, circles denote rollout samples, and colors indicate time points. Red curves show representative sample paths.

Table 11: Marginal reconstruction at individual time points on EMT (10D). The left and right blocks report W_1 and W_2, respectively. Stochastic results are mean $\pm$ standard deviation over four rollouts; lower is better for both metrics.

Method / Inference	$W_1 \downarrow$			$W_2 \downarrow$		
	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$
OT-CFM	0.155	0.204	0.236	**0.170**	0.217	0.248
WFR-FM	0.209	0.228	0.233	0.216	0.235	0.242
WLF	0.234	0.239	0.237	0.241	0.245	0.245
VGFM	0.238	0.245	0.244	0.253	0.254	0.258
SF^{2}M / SDE	0.258 $\pm$ 0.0010	0.278 $\pm$ 0.0008	0.283 $\pm$ 0.0009	0.266 $\pm$ 0.0014	0.287 $\pm$ 0.0010	0.294 $\pm$ 0.0012
DeepRUOT / SDE	0.259 $\pm$ 0.0003	0.281 $\pm$ 0.0011	0.295 $\pm$ 0.0019	0.266 $\pm$ 0.0002	0.290 $\pm$ 0.0009	0.306 $\pm$ 0.0024
CCB Attractive / PF-ODE	**0.154**	**0.182**	**0.197**	0.183	**0.211**	0.230
CCB Attractive / SDE	0.189 $\pm$ 0.0004	0.193 $\pm$ 0.0012	0.198 $\pm$ 0.0010	0.206 $\pm$ 0.0010	0.215 $\pm$ 0.0017	**0.220 $\pm$ 0.0010**
CCB Repulsive / PF-ODE	0.164	0.183	0.198	0.192	0.213	0.229
CCB Repulsive / SDE	0.193 $\pm$ 0.0008	0.196 $\pm$ 0.0028	0.201 $\pm$ 0.0018	0.211 $\pm$ 0.0014	0.218 $\pm$ 0.0038	0.222 $\pm$ 0.0020

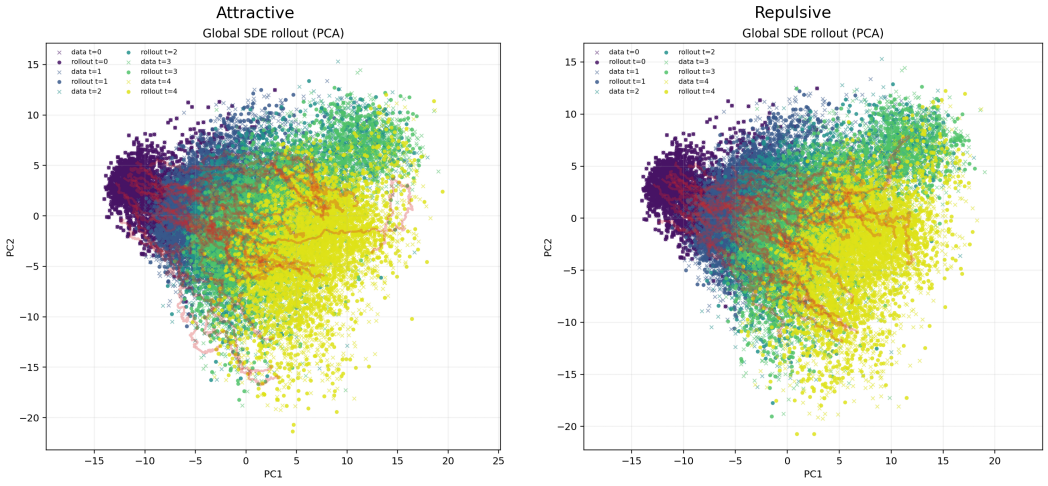

Figure 6: SDE rollouts of CCB on EB (PCA-50D) with attractive (left) and repulsive (right) interactions, starting from the initial snapshot and visualized in a global PCA projection. Crosses denote observed data, circles denote rollout samples, and colors indicate time points. Red curves show representative sample paths.

Table 12: Marginal reconstruction at individual time points on EB (PCA-5D). The left and right blocks report W_1 and W_2, respectively. Stochastic results are mean $\pm$ standard deviation over four rollouts; lower is better for both metrics.

Method / Inference	$W_1 \downarrow$				$W_2 \downarrow$			
	$t=1$	$t=2$	$t=3$	$t=4$	$t=1$	$t=2$	$t=3$	$t=4$
OT-CFM	1.617	1.903	1.988	2.001	1.769	2.078	2.156	2.175
WFR-FM	1.639	1.960	2.089	2.096	1.811	2.161	2.306	2.342
WLF	2.258	2.962	3.241	2.674	2.509	3.491	3.698	2.934
VGFM	3.135	6.060	7.064	6.729	3.672	6.985	7.651	7.143
SF^2M / SDE	2.382 $\pm$ 0.0016	2.742 $\pm$ 0.0021	2.912 $\pm$ 0.0043	2.834 $\pm$ 0.0018	2.669 $\pm$ 0.0018	3.080 $\pm$ 0.0036	3.300 $\pm$ 0.0059	3.333 $\pm$ 0.0035
DeepRUOT / SDE	2.143 $\pm$ 0.0020	2.868 $\pm$ 0.0020	2.700 $\pm$ 0.0048	2.754 $\pm$ 0.0075	2.421 $\pm$ 0.0029	3.184 $\pm$ 0.0006	2.962 $\pm$ 0.0064	3.028 $\pm$ 0.0077
CCB Attractive / PF-ODE	1.313	1.531	1.703	1.654	1.468	1.740	1.905	1.904
CCB Attractive / SDE	1.301 $\pm$ 0.0061	1.435 $\pm$ 0.0132	1.556 $\pm$ 0.0322	**1.514 $\pm$ 0.0083**	1.516 $\pm$ 0.0112	1.757 $\pm$ 0.0268	1.876 $\pm$ 0.0534	1.850 $\pm$ 0.0160
CCB Repulsive / PF-ODE	1.312	1.534	1.697	1.646	**1.466**	**1.730**	1.898	1.896
CCB Repulsive / SDE	**1.299 $\pm$ 0.0077**	**1.417 $\pm$ 0.0189**	**1.549 $\pm$ 0.0238**	1.518 $\pm$ 0.0051	1.527 $\pm$ 0.0064	1.753 $\pm$ 0.0427	**1.871 $\pm$ 0.0400**	**1.838 $\pm$ 0.0175**

Table 13: Marginal reconstruction at individual time points on EB (PCA-50D). The left and right blocks report W_1 and W_2, respectively. Stochastic results are mean $\pm$ standard deviation over four rollouts; lower is better for both metrics.

Method / Inference	$W_1 \downarrow$				$W_2 \downarrow$			
	$t=1$	$t=2$	$t=3$	$t=4$	$t=1$	$t=2$	$t=3$	$t=4$
OT-CFM	7.578	8.776	9.561	10.596	7.701	8.943	9.786	10.941
WFR-FM	7.612	8.674	9.252	10.018	7.726	8.833	9.443	10.237
WLF	8.646	9.855	10.409	10.659	8.829	10.115	10.675	10.897
VGFM	8.270	9.270	9.755	10.191	8.434	9.492	9.979	10.402
SF^2M / SDE	9.109 $\pm$ 0.0010	11.297 $\pm$ 0.0038	14.261 $\pm$ 0.0117	19.513 $\pm$ 0.0228	9.266 $\pm$ 0.0010	11.605 $\pm$ 0.0054	15.669 $\pm$ 0.0246	24.596 $\pm$ 0.0622
DeepRUOT / SDE	7.955 $\pm$ 0.0025	8.683 $\pm$ 0.0024	8.992 $\pm$ 0.0030	9.445 $\pm$ 0.0031	8.142 $\pm$ 0.0024	8.907 $\pm$ 0.0022	9.199 $\pm$ 0.0042	9.618 $\pm$ 0.0032
CCB Attractive / PF-ODE	**5.447**	**7.229**	8.748	9.986	**5.900**	**7.572**	8.969	10.174
CCB Attractive / SDE	6.619 $\pm$ 0.0141	7.665 $\pm$ 0.0081	**7.966 $\pm$ 0.0182**	8.644 $\pm$ 0.0273	6.802 $\pm$ 0.0120	7.888 $\pm$ 0.0148	**8.198 $\pm$ 0.0222**	8.874 $\pm$ 0.0263
CCB Repulsive / PF-ODE	5.597	7.283	8.592	9.723	5.962	7.596	8.801	9.925
CCB Repulsive / SDE	6.802 $\pm$ 0.0064	7.656 $\pm$ 0.0139	8.104 $\pm$ 0.0162	**8.471 $\pm$ 0.0353**	6.963 $\pm$ 0.0085	7.868 $\pm$ 0.0249	8.314 $\pm$ 0.0329	**8.708 $\pm$ 0.0487**

Table 14: Marginal reconstruction at individual time points on EB (100D). The left and right blocks report W_1 and W_2, respectively. Stochastic results are mean $\pm$ standard deviation over four rollouts; lower is better for both metrics.

Method / Inference	$W_1 \downarrow$				$W_2 \downarrow$			
	$t=1$	$t=2$	$t=3$	$t=4$	$t=1$	$t=2$	$t=3$	$t=4$
OT-CFM	9.954	11.142	11.882	13.438	10.056	11.270	12.100	14.032
WFR-FM	9.927	11.036	11.541	12.741	10.016	11.153	11.685	12.990
WLF	10.697	11.957	12.485	13.289	10.829	12.142	12.674	13.446
VGFM	10.592	11.730	12.137	13.007	10.719	11.901	12.299	13.138
SF^{2}M / SDE	11.279 $\pm$ 0.0004	13.212 $\pm$ 0.0010	14.708 $\pm$ 0.0037	16.620 $\pm$ 0.0085	11.412 $\pm$ 0.0003	13.399 $\pm$ 0.0012	14.960 $\pm$ 0.0040	17.233 $\pm$ 0.0141
DeepRUOT / SDE	10.269 $\pm$ 0.0017	11.219 $\pm$ 0.0025	11.321 $\pm$ 0.0028	12.362 $\pm$ 0.0030	10.415 $\pm$ 0.0017	11.402 $\pm$ 0.0022	11.490 $\pm$ 0.0031	12.499 $\pm$ 0.0032
CCB Attractive / PF-ODE	**7.935**	10.246	11.999	14.015	**8.208**	10.445	12.144	14.146
CCB Attractive / SDE	9.587 $\pm$ 0.0050	11.074 $\pm$ 0.0158	11.675 $\pm$ 0.0126	12.797 $\pm$ 0.0117	9.685 $\pm$ 0.0047	11.199 $\pm$ 0.0193	11.802 $\pm$ 0.0167	12.919 $\pm$ 0.0172
CCB Repulsive / PF-ODE	8.101	**10.238**	11.781	13.704	8.312	**10.399**	11.922	13.858
CCB Repulsive / SDE	9.521 $\pm$ 0.0176	10.529 $\pm$ 0.0287	**11.018** $\pm$ **0.0323**	**12.302** $\pm$ **0.0280**	9.609 $\pm$ 0.0205	10.657 $\pm$ 0.0321	**11.145** $\pm$ **0.0352**	**12.435** $\pm$ **0.0307**

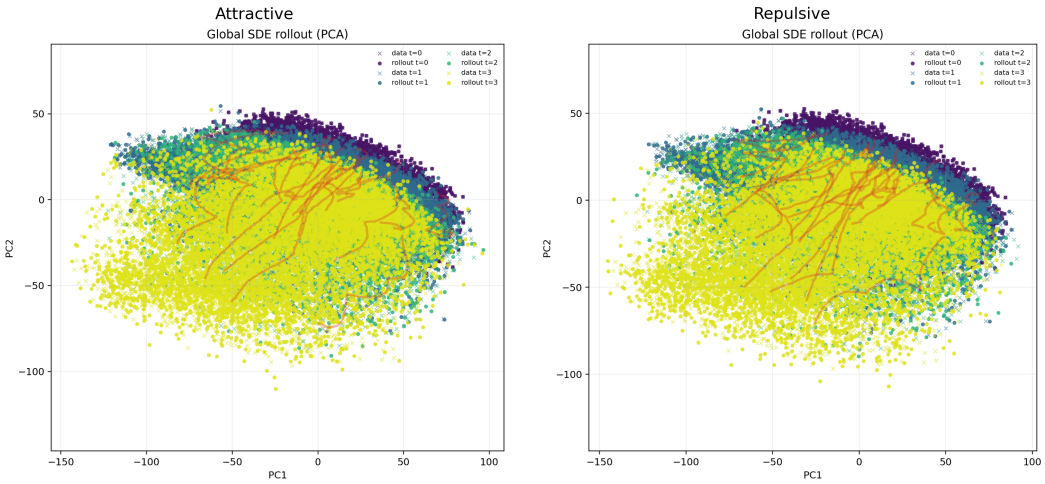

Figure 7: SDE rollouts of CCB on CITE (PCA-50D) with attractive (left) and repulsive (right) interactions, starting from the initial snapshot and visualized in a global PCA projection. Crosses denote observed data, circles denote rollout samples, and colors indicate time points. Red curves show representative sample paths.

Table 15: Marginal reconstruction at individual time points on CITE (PCA-50D). The left and right blocks report W_1 and W_2, respectively. Stochastic results are mean $\pm$ standard deviation over four rollouts; lower is better for both metrics.

Method / Inference	$W_1 \downarrow$			$W_2 \downarrow$		
	$t=1$	$t=2$	$t=3$	$t=1$	$t=2$	$t=3$
OT-CFM	27.449	27.992	30.623	28.017	28.688	31.385
WFR-FM	27.638	27.753	32.005	28.254	28.603	33.130
WLF	32.295	32.205	38.531	33.324	33.222	40.211
VGFM	29.935	30.599	36.492	30.759	31.416	37.740
SF^{2}M / SDE	32.164 $\pm$ 0.0003	41.358 $\pm$ 0.0038	49.257 $\pm$ 0.0046	33.238 $\pm$ 0.0005	43.851 $\pm$ 0.0041	52.497 $\pm$ 0.0085
DeepRUOT / SDE	27.897 $\pm$ 0.0015	28.360 $\pm$ 0.0013	34.328 $\pm$ 0.0045	28.655 $\pm$ 0.0017	29.235 $\pm$ 0.0012	36.272 $\pm$ 0.0055
CCB Attractive / PF-ODE	**12.185**	24.095	29.495	**13.811**	24.719	30.157
CCB Attractive / SDE	14.145 $\pm$ 0.0127	24.304 $\pm$ 0.0150	28.576 $\pm$ 0.0685	15.478 $\pm$ 0.0229	24.876 $\pm$ 0.0134	29.221 $\pm$ 0.0884
CCB Repulsive / PF-ODE	13.178	24.126	28.588	14.588	24.701	29.159
CCB Repulsive / SDE	14.853 $\pm$ 0.0199	**23.898** $\pm$ **0.0241**	**27.343** $\pm$ **0.0390**	16.052 $\pm$ 0.0245	**24.449** $\pm$ **0.0301**	**27.947** $\pm$ **0.0531**

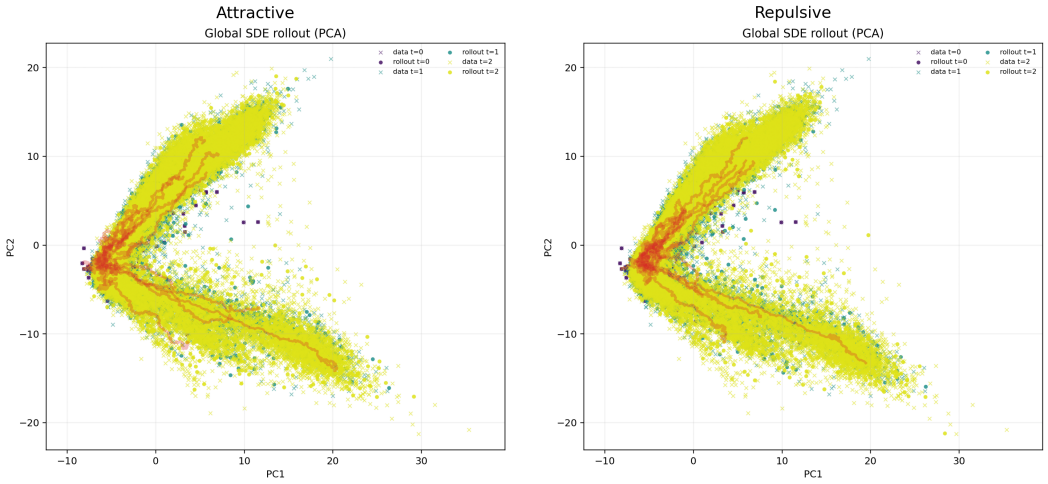

Figure 8: SDE rollouts of CCB on Weinreb (50D) with attractive (left) and repulsive (right) interactions, starting from the initial snapshot and visualized in a global PCA projection. Crosses denote observed data, circles denote rollout samples, and colors indicate time points. Red curves show representative sample paths.

Table 16: Marginal reconstruction at individual time points on Weinreb (50D). The left and right blocks report W_1 and W_2, respectively. Stochastic results are mean $\pm$ standard deviation over four rollouts; lower is better for both metrics.

Method / Inference	$W_1 \downarrow$		$W_2 \downarrow$	
	$t = 1$	$t = 2$	$t = 1$	$t = 2$
OT-CFM	5.659	6.134	6.903	7.251
WFR-FM	5.489	6.217	6.933	7.805
WLF	6.824	7.385	8.353	8.853
VGFM	7.382	8.944	9.341	11.215
SF^{2}M / SDE	8.568 $\pm$ 0.0015	14.590 $\pm$ 0.0050	10.020 $\pm$ 0.0013	15.598 $\pm$ 0.0052
DeepRUOT / SDE	6.311 $\pm$ 0.0018	6.932 $\pm$ 0.0035	7.988 $\pm$ 0.0015	8.541 $\pm$ 0.0037
CCB Attractive / PF-ODE	**4.762**	5.581	6.370	6.954
CCB Attractive / SDE	5.035 $\pm$ 0.0044	5.783 $\pm$ 0.0106	6.533 $\pm$ 0.0333	6.975 $\pm$ 0.0573
CCB Repulsive / PF-ODE	4.789	**5.542**	**6.209**	6.804
CCB Repulsive / SDE	4.979 $\pm$ 0.0154	5.698 $\pm$ 0.0116	6.330 $\pm$ 0.1183	**6.798** $\pm$ **0.0530**

D.2 Held-out-time Interpolation on Mean-field Datasets

Figure 9 shows continuous SDE rollouts of CCB on the two 20D mean-field datasets. Models are trained on the endpoint marginals at $t = 0$ and $t = 1$; Tables 17 and 18 report interpolation errors at the three held-out time points $t = 0.25, 0.5, 0.75$. The left and right blocks report W_1 and MMD^2, respectively. Stochastic results are reported as mean $\pm$ standard deviation, while deterministic results are single values. Lower is better for both metrics, and the lowest mean at each time point is highlighted in bold.

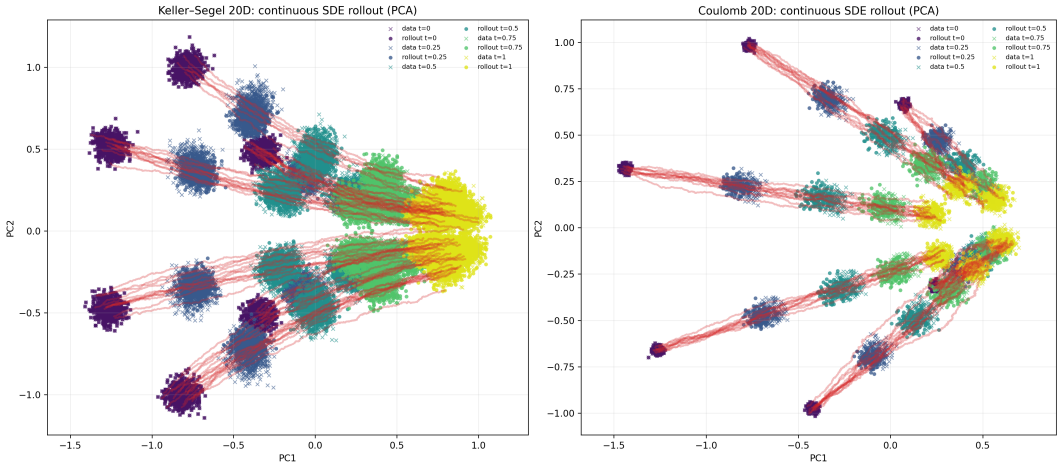

Figure 9: Continuous SDE rollouts of CCB on the 20D Keller–Segel (left) and Coulomb (right) datasets, starting from the initial marginal and visualized in a global PCA projection. Crosses denote ground-truth samples, circles denote rollout samples, and colors indicate $t = 0, 0.25, 0.5, 0.75, 1$. Red curves show representative sample paths. Only the endpoint marginals are used for training.

Table 17: Interpolation at the three held-out time points on Keller–Segel (20D). The left and right blocks report W_1 and MMD^2, respectively. Stochastic results are mean $\pm$ standard deviation; lower is better. MMD^2 means and standard deviations are reported in units of 10^{-3}.

Method / Inference	$W_1 \downarrow$			MMD^2 $(\times 10^{-3}) \downarrow$		
	$t = 0.25$	$t = 0.5$	$t = 0.75$	$t = 0.25$	$t = 0.5$	$t = 0.75$
OT-CFM	**0.0655**	0.1503	0.1426	1.423	14.602	11.120
WFR-FM	0.0816	0.1602	0.1502	1.871	14.892	10.591
SF^{2}M / SDE	0.1330 $\pm$ 0.0002	0.2192 $\pm$ 0.0002	0.2360 $\pm$ 0.0003	**1.151** $\pm$ **0.006**	14.848 $\pm$ 0.026	11.490 $\pm$ 0.099
VGFM	0.0990	0.1798	0.1591	3.639	19.608	14.724
DeepRUOT / SDE	0.2223 $\pm$ 0.0003	0.2509 $\pm$ 0.0004	0.2628 $\pm$ 0.0006	1.428 $\pm$ 0.059	7.132 $\pm$ 0.106	7.103 $\pm$ 0.154
WLF	0.0678	0.1547	0.1518	1.676	15.662	12.355
CCB (without potentials, PF-ODE)	0.1109	0.1931	0.1668	2.216	15.718	12.563
CCB (full model, PF-ODE)	0.1142	**0.1288**	**0.1079**	3.234	0.688	1.540
CCB (full model, SDE)	0.1395 $\pm$ 0.0001	0.1733 $\pm$ 0.0004	0.1940 $\pm$ 0.0003	2.036 $\pm$ 0.017	**0.320** $\pm$ **0.008**	**1.481** $\pm$ **0.023**

Table 18: Interpolation at the three held-out time points on Coulomb (20D). The left and right blocks report W_1 and MMD^2, respectively. Stochastic results are mean $\pm$ standard deviation; lower is better. MMD^2 means and standard deviations are reported in units of 10^{-3}.

Method / Inference	$W_1 \downarrow$			MMD^2 $(\times 10^{-3}) \downarrow$		
	$t = 0.25$	$t = 0.5$	$t = 0.75$	$t = 0.25$	$t = 0.5$	$t = 0.75$
OT-CFM	0.2283	0.2606	0.1802	46.548	64.132	38.748
WFR-FM	0.2440	0.2744	0.1962	49.040	66.927	40.709
SF^{2}M / SDE	0.2494 $\pm$ 0.0003	0.2842 $\pm$ 0.0006	0.2192 $\pm$ 0.0007	41.397 $\pm$ 0.244	56.472 $\pm$ 0.303	31.860 $\pm$ 0.387
VGFM	0.2183	0.2475	0.1856	41.698	60.593	40.906
DeepRUOT / SDE	0.2607 $\pm$ 0.0016	0.3074 $\pm$ 0.0017	0.2565 $\pm$ 0.0022	28.268 $\pm$ 0.461	46.713 $\pm$ 0.621	28.980 $\pm$ 1.046
WLF	0.2338	0.2716	0.1988	48.332	69.219	45.128
CCB (without potentials, PF-ODE)	0.2230	0.2525	0.1676	44.580	59.552	34.852
CCB (external only, PF-ODE)	0.1033	0.0955	0.0542	7.409	8.159	3.208
CCB (pairwise only, PF-ODE)	0.2172	0.2505	0.1673	41.167	56.627	34.509
CCB (full model, PF-ODE)	**0.0864**	**0.0816**	**0.0514**	5.468	6.266	**2.861**
CCB (full model, SDE)	0.1343 $\pm$ 0.0008	0.1556 $\pm$ 0.0008	0.1462 $\pm$ 0.0009	**3.401** $\pm$ **0.162**	**5.131** $\pm$ **0.247**	3.149 $\pm$ 0.216

D.3 Held-out-time Interpolation on Spatial Transcriptomics Data

Figure 10 compares the held-out ground-truth slices with CCB interpolation using the spatial deformation potential as a regularization prior. The three held-out times $t = 0.5, 1.5, 3$ correspond to the endpoint intervals $[0.2, 1]$, $[1, 2]$, and $[2, 6]$, respectively. Table 19 reports spatial W_1, expression W_1, and cell-type neighborhood TV at each time point; metric definitions are given in Appendix B.2.3. Lower is better for all metrics, and the lowest value at each time point is highlighted in bold.

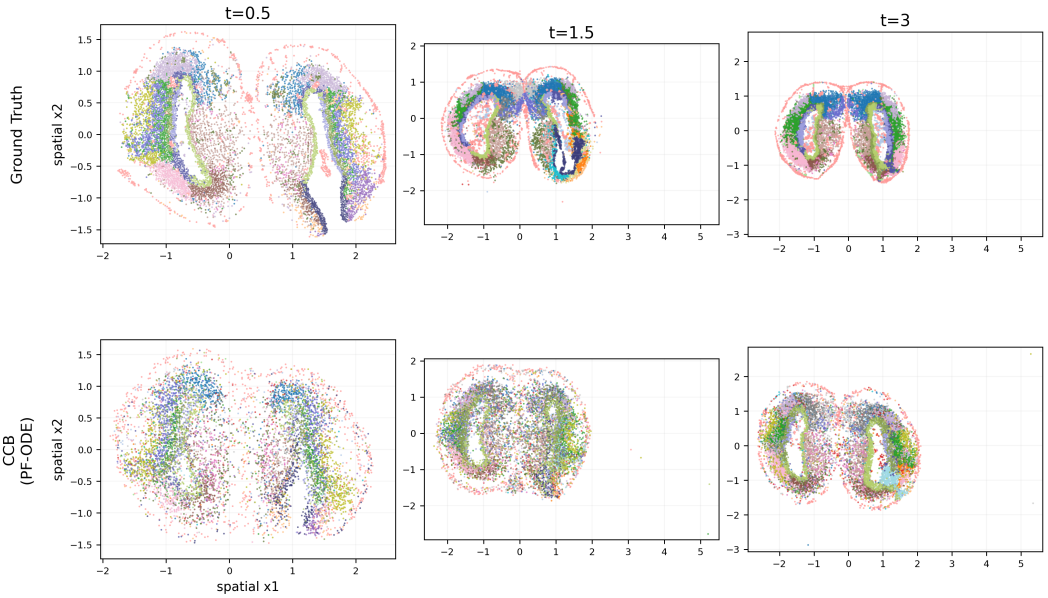

Figure 10: Spatial transcriptomics interpolation at the three held-out time points $t = 0.5, 1.5, 3$ (columns from left to right). The top row shows ground-truth slices, and the bottom row shows CCB PF-ODE interpolation with the spatial deformation potential as a regularization prior. Axes represent spatial coordinates. Colors denote cell types with a shared color scheme across panels; generated cells are labeled by the cell-type classifier.

Table 19: Spatial transcriptomics interpolation at individual held-out time points. The three blocks report spatial W_1, expression W_1, and cell-type neighborhood TV. Lower is better for all metrics.

Method / Inference	**Spatial $W_1 \downarrow$**			**Expression $W_1 \downarrow$**			**Neighborhood TV $\downarrow$**		
	$t = 0.5$	$t = 1.5$	$t = 3$	$t = 0.5$	$t = 1.5$	$t = 3$	$t = 0.5$	$t = 1.5$	$t = 3$
WFR-FM	0.2258	0.2531	0.2936	5.3123	4.9348	4.9450	0.4669	0.5518	0.4370
SF^{2}M	0.2429	0.2668	0.2787	5.2586	5.0466	4.8666	0.4464	0.5692	0.4739
WL-OT	0.2731	0.2730	0.2659	5.2466	5.0274	4.9177	0.5606	0.6252	0.5314
OT-CFM	0.2884	0.2727	0.2634	5.1279	4.9309	4.8768	0.5595	0.6147	0.5140
VGFM	0.2996	0.3322	0.2737	5.0255	4.8342	4.6883	0.5332	0.5682	0.5078
DeepRUOT	0.3830	0.3419	0.2511	**4.8595**	**4.7725**	**4.6799**	0.5046	0.5707	0.4634
CCB (without potentials, PF-ODE)	0.3561	0.3015	**0.2472**	5.0052	4.8479	4.7657	0.5531	0.6065	0.5342
CCB (spatial potential, PF-ODE)	**0.2005**	**0.2395**	0.3283	4.9732	4.8211	4.6825	**0.4004**	**0.5420**	**0.3876**