

---

# PACE: Geometry-Aware Bridge Transport for Single-Cell Trajectory Inference

---

Anonymous Authors<sup>1</sup>

## Abstract

Single-cell trajectory inference from destructive time-course snapshots is fundamentally ill-posed: neither cross-time cell correspondences nor the continuous paths between snapshots are observed, so the observed snapshot distributions alone do not uniquely determine the underlying dynamics. Existing optimal transport and flow-based methods typically couple cells by Euclidean proximity at observed clock times, which can misalign trajectories when development is asynchronous and cells sampled at the same experimental time occupy different latent pseudotime stages. We propose PACE, a trajectory inference framework that selects geometry-consistent continuous transport dynamics from destructive time-course snapshots through three coupled components. First, PACE constructs a state- and time-dependent anisotropic Riemannian metric that preserves low cost along locally supported tangent directions while penalizing normal velocity components. Second, it alternates between refining cross-time couplings under the induced path-action cost and fitting endpoint-preserving neural bridges between adjacent snapshots. Third, it distills the learned bridge dynamics into a global continuous-time velocity field over cellular states. Across seven controlled and biological datasets covering nine held-out reconstruction experiments, PACE achieves the strongest overall reconstruction performance, reducing MMD,  $\mathcal{W}_1$ , and  $\mathcal{W}_2$  by 23.7% on average relative to the strongest competing baseline. PACE also improves RNA-velocity alignment by 15.4% on an embryoid body differentiation benchmark, without requiring explicit cell pairing, lineage tracing, or RNA velocity supervision during training. Code is available at <https://anonymous.4open.science/r/PACE-F444/>.

---

<sup>1</sup>Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Submitted to the 2026 Workshop on Generative and Agentic AI for Biology (ICML 2026). Do not distribute.

## 1. Introduction

Understanding how cells move from one state to another is a central problem in single-cell biology (Farrell et al., 2018; Wagner et al., 2018; Schiebinger et al., 2019; Tanay & Regev, 2017; Griffiths et al., 2018). Processes such as development, differentiation, immune activation, tumor evolution, and cellular reprogramming are not simply collections of discrete cell types, but continuous population-level transitions through high-dimensional molecular state space (Trapnell et al., 2014; Haghverdi et al., 2016; Setty et al., 2019; Cacchiarelli et al., 2018). The key questions are therefore dynamical, including which early states commit to particular fates, which intermediate states are transient but decisive, and which regulatory programs drive these transitions. Trajectory inference (Lavenant et al., 2021; Hashimoto et al., 2016) aims to answer these questions by reconstructing continuous cell-state evolution from time-course single-cell observations.

Transport- and flow-based trajectory methods provide a natural framework for modeling such dynamics (Peyré & Cuturi, 2019; Schiebinger et al., 2019; Lavenant et al., 2021). Since time-course single-cell data provide population-level snapshots rather than paired observations of the same cells (Weinreb et al., 2018; Schiebinger et al., 2019), these methods must infer how marginal distributions are coupled across time. Most existing formulations operate in the Euclidean representation space, where couplings are estimated from Euclidean endpoint costs and continuous paths or velocity fields are regularized by Euclidean kinetic energy, action, or smoothness (Neklyudov et al., 2023; Tong et al., 2023; 2020; Huguet et al., 2022). This Euclidean geometry is convenient, but it is not necessarily aligned with biological progress. Two cells can be close in expression space while lying at different developmental stages or fate branches, whereas biologically plausible motion may follow curved and locally anisotropic directions along the cell-state manifold (Moon et al., 2019; Wolf et al., 2019).

This mismatch makes trajectory inference from destructive snapshots fundamentally ambiguous (Weinreb et al., 2018; Tritschler et al., 2019): the cross-time coupling is not observed and is generally not identifiable from marginal distributions alone. A coupling may satisfy the marginal constraints while connecting cells across incompatible devel-

055 opmental programs, stalled states, or fate branches. For ex-  
 056 ample, in mouse reprogramming (Schiebinger et al., 2019),  
 057 differences in growth rates between cell types can make  
 058 marginal matching misleading, coupling apoptotic stromal  
 059 cells to rapidly expanding iPSCs. Human iPSC reprogram-  
 060 ming (Liu et al., 2020) shows the same issue at finer granu-  
 061 larity, that cells collected on the same day can contain het-  
 062 erogeneous primed-like, naive-like, and trophectoderm-like  
 063 intermediates, so clock-time adjacency alone does not de-  
 064 termine which cells should be coupled. Auxiliary measure-  
 065 ments such as RNA velocity, lineage tracing, or metabolic  
 066 labeling (La Manno et al., 2018; Bergen et al., 2020; Lange  
 067 et al., 2022) can help, but require additional experiments  
 068 and are unavailable in many datasets. Existing Euclidean  
 069 or support-based regularizers can encourage short, smooth,  
 070 or data-supported paths (Saelens et al., 2019), but they do  
 071 not directly encode which local directions of motion are  
 072 developmentally admissible.

073 We propose PACE, a geometry-aware trajectory-inference  
 074 framework for selecting snapshot-consistent continuous  
 075 transport dynamics from unpaired time-course snapshots.  
 076 The intuition is that, although true cell identities and inter-  
 077 mediate paths are unobserved, local spatiotemporal neigh-  
 078 borhoods still provide a weak but useful prior on admis-  
 079 sible directions, since motion along locally supported tan-  
 080 gent directions is more plausible than motion orthogonal to  
 081 the observed manifold structure. PACE therefore replaces  
 082 Euclidean transport costs with a Riemannian path action  
 083 induced by a state- and time-dependent anisotropic metric,  
 084 which assigns lower cost to locally supported tangent motion  
 085 and higher cost to normal motion. Specifically, PACE esti-  
 086 mates local tangent *directions* from spatiotemporal neigh-  
 087 borhoods of observed cells and uses this metric to define  
 088 an anisotropic bridge transport problem, where the cost of  
 089 coupling two cells is the minimum Riemannian action of an  
 090 endpoint-conditioned path. PACE then alternates between  
 091 refining OT couplings under this path-action cost and fitting  
 092 endpoint-preserving neural bridges, allowing local geome-  
 093 try to influence both which cells are coupled and how they  
 094 move between snapshots. Finally, the learned bridge dynam-  
 095 ics are distilled into a continuous-time population velocity  
 096 field over cellular states.

097 Our contributions are threefold:

- 098 1. We introduce a time- and state-dependent Riemannian  
 099 metric based on local spatiotemporal tangent-space pro-  
 100 jections, providing a geometry-aware path-action cost  
 101 for selecting couplings between adjacent destructive  
 102 snapshots beyond Euclidean endpoint proximity.
- 103 2. We develop an iterative bridge-transport procedure that  
 104 alternates between Riemannian coupling refinement  
 105 and endpoint-preserving neural bridge fitting, allowing  
 106  
 107  
 108  
 109

cross-time correspondences and interpolant paths to be  
 selected jointly under the same geometry-aware action.

3. We distill endpoint-conditioned bridges into a global  
 continuous-time population velocity field and evaluate  
 PACE on seven datasets covering nine held-out recon-  
 struction experiments. PACE achieves the strongest  
 overall reconstruction performance, reducing MMD,  
 $\mathcal{W}_1$ , and  $\mathcal{W}_2$  by 23.7% on average relative to the  
 strongest competing baseline, improves RNA-velocity  
 alignment on an embryoid body differentiation bench-  
 mark, and shows consistent gains across component  
 ablations.

## 2. Related Work

### Optimal transport and flow-based trajectory inference.

Optimal transport provides a natural framework for cou-  
 pling unpaired snapshot distributions in single-cell analysis  
 (Schiebinger et al., 2019; Lavenant et al., 2021). Recent  
 neural extensions such as TrajectoryNet (Tong et al., 2020),  
 Wasserstein Lagrangian Flows (Neklyudov et al., 2023), and  
 Conditional Flow Matching (CFM) (Lipman et al., 2023;  
 Tong et al., 2023) learn continuous dynamics by regressing  
 vector fields along interpolants between coupled endpoints.  
 Schrodinger bridge methods (De Bortoli et al., 2021; Shi  
 et al., 2023; Chen et al., 2023) add stochasticity or entropic  
 regularization to the transport problem, while Curly Flow  
 Matching (Petrović et al., 2026) extends the framework  
 to non-gradient dynamics using approximate velocity in-  
 formation. These methods typically infer couplings from  
 Euclidean proximity or OT distances in the observed space,  
 which can misalign trajectories when cells at the same ex-  
 perimental time occupy different pseudotime stages.

### Geometry-aware generative models.

The manifold hy-  
 pothesis has motivated data-dependent Riemannian metrics  
 in ambient spaces (Hauberg et al., 2012; Arvanitidis et al.,  
 2017), flow matching on known manifolds or general geome-  
 tries (Chen & Lipman, 2024), and manifold-aware OT flows  
 for single-cell trajectories (Huguet et al., 2022). Metric  
 Flow Matching (MFM) (Kapuśniak et al., 2024) is clos-  
 est to our setting: it uses task-independent *support-aware*  
 metrics such as LAND (Arvanitidis et al., 2017) and RBF  
 to pull geodesics toward the data support. PACE differs  
 in both the source and use of geometry. In asynchronous  
 reprogramming, density support alone does not determine  
 plausible cross-time transitions, since a high-density inter-  
 mediate may connect to either progressed or refractory fates  
 and Euclidean proximity can couple incompatible programs.  
 PACE instead builds a time- and state-dependent, *direction-*  
*aware* metric from local spatiotemporal tangent subspaces  
 estimated from destructive snapshots, penalizing motion  
 orthogonal to plausible developmental directions. The re-

sulting metric action is used both to learn interpolant paths and to refine cross-time couplings, whereas MFM typically assumes the endpoint pairing is fixed before interpolant learning.

### 3. Method

#### 3.1. Problem formulation

We observe single-cell point clouds at anchor times  $\{t_k\}_{k=0}^K$ , with  $t_0 = 0$ ,

$$\mathcal{X}^{(k)} = \{x_i^{(k)}\}_{i=1}^{N_k} \sim \hat{\rho}_k \subset \mathbb{R}^d, \quad k = 0, 1, \dots, K.$$

Here  $d$  denotes the representation dimension. Our goal is to infer trajectories and a velocity field transporting  $\hat{\rho}_k$  to  $\hat{\rho}_{k+1}$  between adjacent times. Once learned, the velocity field can be integrated from  $x_{t_0} \sim \hat{\rho}_0$  to generate trajectories over the time course.

**Proposition 3.1** (Ill-posedness from snapshots alone). *Given only the point clouds  $\{\mathcal{X}^{(k)}\}_{k=0}^K$ , the reconstruction of cross-time couplings and intermediate trajectories is non-unique.*

Indeed, for a single interval  $[t_k, t_{k+1}]$ , many couplings  $\pi_k \in \Pi(\hat{\rho}_k, \hat{\rho}_{k+1})$  share the same endpoint marginals. For any such coupling, each coupled endpoint pair can also be connected by infinitely many smooth paths. Thus, the observed snapshots determine neither a unique correspondence nor a unique intermediate trajectory. This observation motivates PACE to frame trajectory inference as a variational problem over couplings and paths, selected by a geometry-aware path action. The proof of Proposition 3.1 is given in Appendix D.

#### 3.2. Geometry-aware bridge transport

To address the ill-posedness in Proposition 3.1, PACE formulates trajectory reconstruction as selecting both a cross-time coupling and a family of continuous paths, chosen by minimizing a geometry-aware action. Given a time- and state-dependent metric tensor  $G_k(x, t)$  (constructed in §3.3), the cost of a path  $\gamma$ , or "path action", between endpoints  $(x, y)$  is

$$\mathcal{A}_{G_k}[\gamma] = \int_0^1 \dot{\gamma}(\tau)^\top G_k(\gamma(\tau), t_{k,\tau}) \dot{\gamma}(\tau) d\tau, \quad (1)$$

$$\gamma(0) = x, \quad \gamma(1) = y,$$

where  $\tau \in [0, 1]$  is local interpolation time and the physical time is  $t_{k,\tau} = t_k + \tau(t_{k+1} - t_k)$ . The path action first defines the geometry-aware endpoint cost, and the ideal PACE

bridge problem then transports mass using this induced cost:

$$c_{G_k}(x, y) = \inf_{\substack{\gamma(0)=x \\ \gamma(1)=y}} \mathcal{A}_{G_k}[\gamma], \quad (2)$$

$$\pi_k^* \in \arg \min_{\pi_k \in \Pi(\hat{\rho}_k, \hat{\rho}_{k+1})} \int c_{G_k}(x, y) d\pi_k(x, y).$$

Here  $c_{G_k}(x, y)$  is not a fixed Euclidean endpoint distance; it is the minimum action required to move from  $x$  to  $y$  under the metric  $G_k$ .

This formulation makes explicit how PACE differs from standard OT. If  $G_k(x, t) \equiv I$ , then Eq. (1) reduces to the Euclidean kinetic energy. For any fixed endpoint pair  $(x, y)$ , the minimum-action path is the straight constant-speed interpolant, and the induced cost becomes

$$c_{G_k}(x, y) = \|y - x\|^2.$$

In this special case, Eq. (2) reduces to standard quadratic-cost OT between  $\hat{\rho}_k$  and  $\hat{\rho}_{k+1}$ .

PACE departs from this Euclidean case by using a state- and time-dependent anisotropic metric  $G_k(x, t)$ . The induced cost  $c_{G_k}(x, y)$  is no longer determined only by endpoint distance; it depends on the entire path and on how the *path velocity* aligns with local developmental geometry. As a result, the minimum-action cost generally has no closed-form solution and cannot be reduced to a fixed Euclidean endpoint cost. The next section defines the spatiotemporal metric  $G_k(x, t)$ , and § 3.4 describes how PACE approximates this problem with neural bridges and alternating coupling updates.

**Finite-dimensional approximation.** Problem (2) is an ideal infinite-dimensional formulation. PACE makes it tractable through two approximations (§3.4): (i) the path family is parameterized by neural bridges  $\gamma_\theta$  that satisfy endpoint constraints by construction; (ii) the action integral is approximated on a finite time grid. These yield a finite-dimensional alternating optimization over bridge parameters and correspondence variables.

#### 3.3. Time-dependent spatiotemporal tangent metric

The metric should encode a simple prior: admissible motion should follow directions locally supported by the observed geometry, rather than cut across the cell-state manifold. Although local neighborhoods do not reveal the true direction of time, their dominant variation directions provide an undirected tangent approximation to the set of plausible state changes. PACE therefore treats tangent motion as low cost and penalizes velocity components in the locally estimated normal subspace. Because the local composition and geometry of snapshots can change across experimental time, the admissible subspaces are indexed by both state and time, yielding a time- and state-dependent metric  $G_k(x, t)$ .

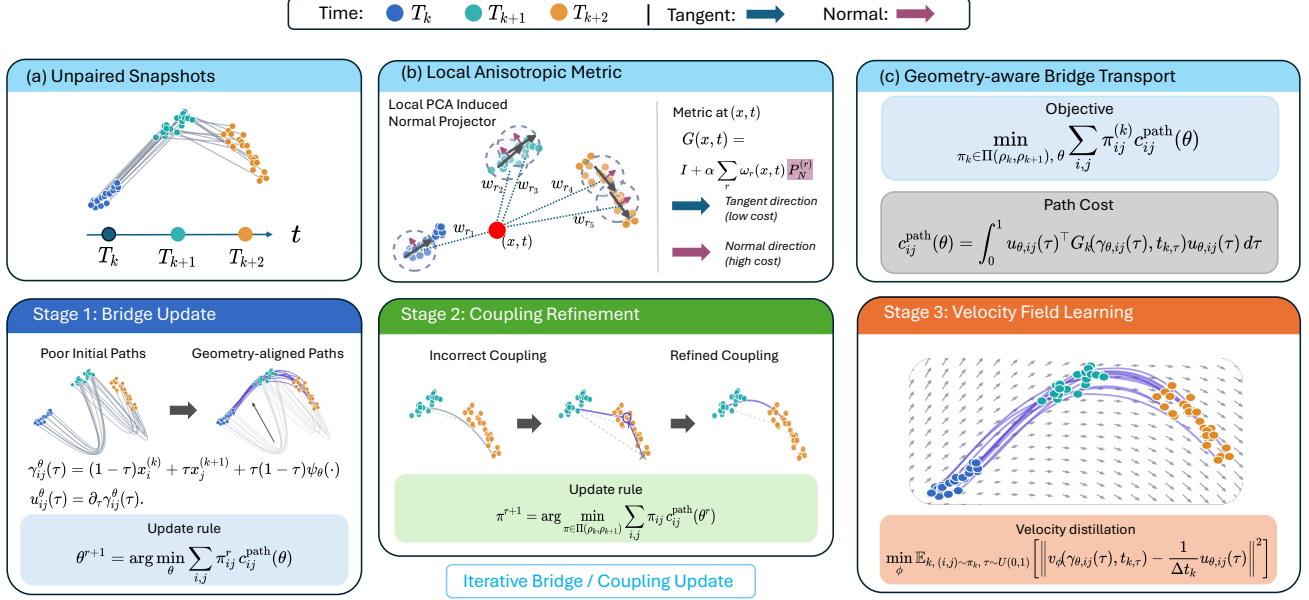


Figure 1. **Overview of PACE.** PACE uses local PCA to construct an anisotropic metric  $G_k(x, t) = I + \alpha C_N^{(k)}(x, t)$ , trains endpoint-preserving neural bridges under the corresponding path-action cost, iteratively refines cross-time couplings, and distills the learned bridge dynamics into a global velocity field for trajectory inference from unpaired snapshots.

**Local normal subspaces.** At each anchor point  $x_r$  observed at time  $t_r$ , PACE estimates an anchor-wise local normal direction by finding its  $m_{\text{nn}}$  nearest spatial neighbors within the same snapshot (including  $x_r$  itself), computing a Gaussian-kernel weighted covariance of the neighbor cloud, and extracting the minimum-variance principal direction  $n_r$ . In two dimensions the normal projector is simply  $P_N^{(r)} = n_r n_r^\top$ ; in higher dimensions PACE adaptively selects the tangent-subspace dimension and builds the normal projector as the orthogonal complement (Appendix K).

**Spatiotemporal metric construction.** For a query point  $(x, t)$  in segment  $k$ , PACE interpolates these anchor-wise local normal projectors across state and time using space-time Gaussian weights over all anchor points  $r$ :

$$\omega_r^{(k)}(x, t) \propto \exp\left(-\frac{\|x - x_r\|^2}{(h_x^{(k)})^2} - \frac{|t - t_r|^2}{(h_t^{(k)})^2}\right), \quad (3)$$

with normalization chosen so that  $\sum_r \omega_r^{(k)}(x, t) = 1$ , where the bandwidths  $h_x^{(k)}$  and  $h_t^{(k)}$  are estimated adaptively for each segment (see Appendix L). These weights define an averaged normal projector field and the corresponding metric tensor

$$C_N^{(k)}(x, t) = \sum_r \omega_r^{(k)}(x, t) P_N^{(r)},$$

$$G_k(x, t) = I + \alpha C_N^{(k)}(x, t), \quad \alpha > 0.$$

The corresponding velocity cost is  $v^\top G_k(x, t) v = \|v\|^2 + \alpha v^\top C_N^{(k)}(x, t) v$ . Because  $C_N^{(k)}$  is a positive-semidefinite average of local normal projectors, the second term increases the cost most strongly for velocities aligned with nearby estimated normal directions.

**Proposition 3.2** (Normal-subspace penalizing property). *For every segment  $k$ , query point  $(x, t)$ , and velocity  $v \in \mathbb{R}^d$ ,*

$$v^\top G_k(x, t) v = \|v\|^2 + \alpha v^\top C_N^{(k)}(x, t) v \geq \|v\|^2. \quad (4)$$

*Moreover, if  $C_N^{(k)}(x, t) = P_N(x, t)$  is an exact orthogonal projector onto the normal subspace  $T_{x,t}^\perp$ , and  $v = v_T + v_N$  with  $v_T \in T_{x,t}$  and  $v_N \in T_{x,t}^\perp$ , then*

$$v^\top G_k(x, t) v = \|v_T\|^2 + (1 + \alpha) \|v_N\|^2. \quad (5)$$

Proposition 3.2 shows that the interpolated metric always preserves at least the Euclidean velocity cost, and that tangent motion retains its Euclidean cost while normal motion is penalized by a factor  $1 + \alpha$  in the ideal projector case. Thus, the metric turns the qualitative principle of geometry-aligned cellular motion into an explicit action functional.

### 3.4. Finite-dimensional optimization by alternating bridge and OT coupling updates

PACE approximates the anisotropic bridge problem by alternating between an endpoint-conditioned bridge and a cross-time coupling. The coupling determines which endpoint pairs are used to train the bridge, while the bridge

induces a path-action cost for updating the coupling. This creates a bootstrap mechanism in which more plausible couplings provide better endpoint supervision, and better bridges provide a geometry-aware approximation to the endpoint cost used for OT refinement. Since solving a separate minimum-action path problem for every candidate endpoint pair is infeasible under the state- and time-dependent metric  $G_k(x, t)$ , PACE amortizes path optimization with a shared endpoint-preserving neural bridge  $\gamma_\theta(x, y, \tau)$ . See Appendix J for a detailed discussion of this alternating procedure.

**Endpoint-preserving neural bridge.** For an endpoint pair  $(x, y) \in \mathbb{R}^d \times \mathbb{R}^d$  and local time  $\tau \in [0, 1]$ , we define

$$\gamma_\theta(x, y, \tau) = (1 - \tau)x + \tau y + \tau(1 - \tau)\psi_\theta(x, y, \tau), \quad (6)$$

where  $\psi_\theta : \mathbb{R}^d \times \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$  is a neural network. The factor  $\tau(1 - \tau)$  enforces  $\gamma_\theta(x, y, 0) = x$  and  $\gamma_\theta(x, y, 1) = y$ , so the network only controls the interior deformation of the path (Petrović et al., 2026). The bridge velocity is  $u_\theta(x, y, \tau) = \partial_\tau \gamma_\theta(x, y, \tau)$ , computed by automatic differentiation. When  $\psi_\theta = 0$ , the bridge reduces to the straight Euclidean interpolant.

**Bridge learning under fixed coupling.** Assume that, for each adjacent snapshot pair, a coupling  $\pi_k \in \Pi(\hat{\rho}_k, \hat{\rho}_{k+1})$  is given. During the bridge update, endpoint pairs  $(x_i^{(k)}, x_j^{(k+1)}) \sim \pi_k$  are sampled from the current coupling. We write  $t_{k,\tau} = t_k + \tau(t_{k+1} - t_k)$ .

PACE trains the endpoint-preserving bridge by minimizing

$$\mathcal{L}_{\text{bridge}}(\theta) = \lambda_{\text{metric}} \mathcal{L}_{\text{metric}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}. \quad (7)$$

The main term  $\mathcal{L}_{\text{metric}}$  is the anisotropic bridge action induced by the spatiotemporal metric  $G_k$ . An optional regularization term  $\mathcal{L}_{\text{reg}}$  penalizes large normal components and incoherent cross-segment velocities (Appendix G).

The metric-action loss is

$$\begin{aligned} \mathcal{L}_{\text{metric}} = \mathbb{E}_{k,(x,y) \sim \pi_k} \left[ \frac{1}{T} \sum_{\ell=1}^T u_\theta^\top G_k(\gamma_\theta(x, y, \tau_\ell), t_{k,\tau_\ell}) \right. \\ \left. \times \nu_\theta(x, y, \tau_\ell) \right], \quad \nu_\theta = \nu_\theta(x, y, \tau_\ell). \end{aligned} \quad (8)$$

**Coupling update under fixed bridge.** After the bridge model has been updated under the current coupling, PACE recomputes the coupling using the full path action rather than Euclidean endpoint distance. For a candidate source-target pair  $(x_i^{(k)}, x_j^{(k+1)})$ , the path-action cost is

$$c_{ij}^{\text{path}}(\theta) = \frac{1}{M} \sum_{m=1}^M u_{ij}^\top G_k(\gamma_{ij}, t_{k,\tau_m}) \nu_{ij}, \quad (9)$$

where  $\nu_{ij} = \nu_\theta(x_i^{(k)}, x_j^{(k+1)}, \tau_m)$  and  $\gamma_{ij} = \gamma_\theta(x_i^{(k)}, x_j^{(k+1)}, \tau_m)$ . The coupling is then updated by solving

$$\pi_k \leftarrow \arg \min_{\pi_k \in \Pi(\hat{\rho}_k, \hat{\rho}_{k+1})} \sum_{i,j} c_{ij}^{\text{path}}(\theta) \pi_{ij}^{(k)}. \quad (10)$$

**Alternating optimization.** Given an initial coupling, PACE alternates between two blocks:

- Bridge update.** Fix the current couplings  $\{\pi_k\}$ , sample endpoint pairs  $(x_i^{(k)}, x_j^{(k+1)}) \sim \pi_k$ , and update  $\theta$  by taking gradient steps on Eq. (7).
- Coupling update.** Fix the current bridge  $\gamma_\theta$ , compute the path-action cost matrix in Eq. (9), and update each  $\pi_k$  by solving the OT problem in Eq. (10).

The bridge update learns low-action endpoint-conditioned paths under the current coupling, while the coupling update selects source-target transport using the learned path action. In the idealized finite-dimensional setting where the bridge block is solved exactly for the metric-action objective, these two updates form a monotone block-coordinate descent scheme; Appendix F states this property and relates it to the stochastic neural implementation. In practice, the coupling update is triggered periodically, for example every  $R_0$  epochs.

### 3.5. Distilling endpoint-conditioned bridges into a global velocity field

The alternating optimization in §3.4 yields endpoint-conditioned bridges  $\gamma_\theta(x, y, \tau)$  and  $\tau$ -velocities  $u_\theta(x, y, \tau) = \partial_\tau \gamma_\theta(x, y, \tau)$ . These bridges are pair-specific: to query a velocity at an arbitrary state  $z$  and time  $t$ , one would need to know which endpoint pair  $(x, y)$  and interpolation parameter  $\tau$  generated  $z$ . PACE therefore distills the bridge dynamics into a global velocity field  $v_\phi(x, t)$  that can be evaluated without reference to a particular endpoint pair.

**Distillation objective.** For each segment  $k$ , we sample  $(x, y) \sim \pi_k$  and  $\tau \sim \text{Uniform}[0, 1]$ , set  $x_{t_{k,\tau}} = \gamma_\theta(x, y, \tau)$  with  $t_{k,\tau} = t_k + \tau \Delta t_k$ , and match the global velocity field:

$$\begin{aligned} \mathcal{L}_{\text{distill}}(\phi) = \mathbb{E}_{k,(x,y) \sim \pi_k} \mathbb{E}_\tau \left[ \|v_\phi(x_t, t_{k,\tau}) - \frac{u_\theta}{\Delta t_k}\|^2 \right], \\ x_t = \gamma_\theta(x, y, \tau), \quad u_\theta = u_\theta(x, y, \tau). \end{aligned} \quad (11)$$

where  $\Delta t_k = t_{k+1} - t_k$  converts  $\tau$ -velocity to physical-time velocity. Once  $v_\phi$  is learned, continuous trajectories are generated by integrating

$$\frac{dx_t}{dt} = v_\phi(x_t, t), \quad x_0 \sim \hat{\rho}_0. \quad (12)$$




Figure 2. Overview of the 2D benchmark datasets. Points are colored by observed time for Ocean (Petrović et al., 2026), Two-Branch, EB PHATE (Moon et al., 2019), Schiebinger2019 (Schiebinger et al., 2019), and iPSC-Liu (Liu et al., 2020).

Table 1. Per-timepoint results on Ocean (2D) (Petrović et al., 2026), holdout  $t \in \{1, 3, 5, 7\}$ .

Method	MMD ↓				$\mathcal{W}_1$ ↓				$\mathcal{W}_2$ ↓			
	$t=1$	$t=3$	$t=5$	$t=7$	$t=1$	$t=3$	$t=5$	$t=7$	$t=1$	$t=3$	$t=5$	$t=7$
Action Matching	0.7429	1.0662	0.9465	1.1289	0.3549	0.4823	0.2383	0.5073	0.3662	0.4846	0.2425	0.5143
Aligned CFM	0.7989	0.8754	1.0073	0.8272	0.0976	0.1644	0.3024	0.2245	0.1017	0.1673	0.3046	0.2262
CURLY	1.0545	0.8373	0.5869	0.8663	0.3766	0.1521	0.1077	0.2638	0.3771	0.1537	0.1102	0.2646
DMSB	1.1260	<u>0.6089</u>	<u>0.2542</u>	<u>0.4962</u>	0.3060	<u>0.1097</u>	<u>0.0435</u>	<u>0.1097</u>	0.3065	<u>0.1150</u>	<u>0.0473</u>	<u>0.1142</u>
MFN	0.8482	0.7859	0.6928	0.5823	<u>0.0960</u>	0.1433	0.1359	0.1279	<u>0.0991</u>	0.1452	0.1385	0.1302
OT-CFM	0.9410	0.8169	0.6827	0.6305	0.1227	0.1428	0.1289	0.1321	0.1255	0.1459	0.1308	0.1334
<b>PACE(ours)</b>	<b>0.4504</b>	<b>0.2588</b>	<b>0.0735</b>	<b>0.2779</b>	<b>0.0399</b>	<b>0.0398</b>	<b>0.0268</b>	<b>0.0505</b>	<b>0.0440</b>	<b>0.0434</b>	<b>0.0362</b>	<b>0.0534</b>

## 4. Experiment

The experiments evaluate PACE through five questions. First, on controlled 2D trajectories, can PACE recover smooth and branching geometry without paired identities or velocity supervision (Figure 2; Tables 1 and A.2)? Second, when reference velocities are available only for evaluation, do the learned dynamics align with held-out velocity fields on Ocean and EB PHATE (Figure 3; Appendix Tables A.3 and A.4)? Third, on biological time courses, does PACE improve held-out reconstruction for single-cell differentiation and reprogramming from destructive snapshots (Table 2)? Fourth, does PACE remain effective as the representation dimension increases (Tables 3)? Finally, which components of PACE account for the observed gains (Figure 5)?

### 4.1. Experimental setup

**Evaluation protocol.** All experiments use the same held-out snapshot protocol: methods observe only the training time points and are evaluated on the held-out time points listed in each table caption. We report distributional reconstruction quality using maximum mean discrepancy (MMD), 1-Wasserstein distance ( $\mathcal{W}_1$ ), and 2-Wasserstein distance ( $\mathcal{W}_2$ ); lower is better, bold denotes the best result, and underline denotes the second best.

**Datasets.** Our benchmark suite includes controlled 2D temporal point clouds and biological single-cell time-course datasets. The controlled datasets are designed to isolate geometric behavior in low-dimensional temporal data, including Ocean (Shen et al., 2025; Petrović et al., 2026) and a branching toy dataset, without relying on cell identities or velocity supervision. The biological datasets test destructive snapshot settings in which individual cell identities are not shared across time, covering embryoid

body differentiation (Moon et al., 2019), mouse reprogramming (Schiebinger et al., 2019), induced trophoblast stem-cell reprogramming (Liu et al., 2020), and multimodal CITE-seq/Multiome time courses (Lance et al., 2022). Figure 2 gives a visual overview of the 2D datasets used in the main low-dimensional experiments. Detailed dataset sources and preprocessing choices are provided in Appendix A.

**Baselines.** We compare against representative transport- and flow-based trajectory models, covering action-based continuous dynamics (Neklyudov et al., 2023), minibatch optimal-transport flow matching (Tong et al., 2023), metric/geodesic flow matching (Kapuśniak et al., 2024), pseudo-velocity-corrected flows (Petrović et al., 2026), adversarial multi-marginal interpolant learning (Kviman et al., 2026), and stochastic Schrödinger bridge dynamics (Chen et al., 2023). Implementation details for baselines and PACE are provided in Appendices B and C. Qualitative prediction-versus-held-out overlays for the same 2D runs are shown in Appendix Figures A.1–A.5.

### 4.2. Controlled 2D trajectories




Figure 3. Velocity-alignment diagnostics on Ocean (Petrović et al., 2026) and EB PHATE (Moon et al., 2019). Ocean compares learned velocities with held-out simulator velocities, while EB PHATE compares learned 2D velocities with the RNA-velocity reference. Bars report cosine distance and normalized  $L_2$  error; lower is better.

We evaluate PACE on two controlled 2D settings, including Ocean, a rotational non-gradient benchmark from Curly Flow Matching (Petrović et al., 2026), and a Two-Branch toy benchmark. Methods observe only unpaired point positions; simulator velocities and latent identities are not used for training. For Ocean, we also hold out evaluation snapshots and compare the learned velocity field at those times with the simulator velocity field, testing whether the inferred dynamics recover the correct directionality rather than only matching held-out marginals.

Tables 1 and A.2 show that PACE gives the strongest overall held-out reconstruction. It is best on all Ocean metrics and remains strongest on most Two-Branch entries, with CURLY slightly better only for late-branch  $\mathcal{W}_2$ . The Ocean panel in Figure 3 further checks the learned velocity direction against held-out simulator velocities; the per-timepoint diagnostic values are reported in Appendix Table A.3. Together, these results indicate that the geometry-aware bias

helps recover both smooth rotational motion and simple branching structure from snapshots alone.

### 4.3. Low-dimensional single-cell trajectories

We next ask whether the same behavior carries over to biological time courses represented in low-dimensional embeddings, where the true cell identities are unobserved and geometric structure must be inferred from destructive snapshots. Table 2 summarizes the time-averaged low-dimensional biological results across EB PHATE, iPSC-Liu, and Schiebinger2019; the corresponding per-timepoint results are reported in Appendix Tables A.5, A.6, and A.7. On EB PHATE, PACE obtains the best MMD,  $\mathcal{W}_1$ , and  $\mathcal{W}_2$ , showing that the method improves held-out marginal reconstruction in a manifold-aware PHATE embedding. The EB PHATE panel in Figure 3 adds an independent directionality check by comparing the learned 2D velocity field at the held-out snapshot with RNA velocity, which is used only for evaluation. Appendix Table A.4 shows that PACE also best aligns with this external velocity proxy, suggesting that the gain reflects both endpoint reconstruction and plausible local flow direction. On iPSC-Liu 2D, PACE is best on every reported metric. On Schiebinger2019, PACE is best on seven of nine detailed metric-timepoint entries and remains close on the remaining final-time scores. Overall, the low-dimensional biological results show that PACE transfers the controlled-geometry gains to both differentiation and reprogramming snapshots. Compared with the strongest baseline for each dataset–metric pair, PACE reduces MMD,  $\mathcal{W}_1$ , and  $\mathcal{W}_2$  by 36.2%, 21.5%, and 13.6% on EB PHATE; by 16.1%, 26.4%, and 25.0% on iPSC-Liu; and by 31.2%, 25.5%, and 22.4% on Schiebinger2019.

Table 2. Time-averaged low-dimensional single-cell results. EB (Moon et al., 2019) uses  $t = 3$ , iPSC-Liu (Liu et al., 2020) uses  $t \in \{4, 16\}$ , and Schiebinger2019 (Schiebinger et al., 2019) uses  $t \in \{6, 11, 16\}$ .

Method	EB PHATE			iPSC-Liu			Schiebinger2019		
	MMD ↓	$\mathcal{W}_1$ ↓	$\mathcal{W}_2$ ↓	MMD ↓	$\mathcal{W}_1$ ↓	$\mathcal{W}_2$ ↓	MMD ↓	$\mathcal{W}_1$ ↓	$\mathcal{W}_2$ ↓
Action Matching	0.2436	0.4553	0.5273	0.5230	0.8005	0.9730	0.4057	0.4161	0.4754
Aligned CFM	0.1056	0.2657	0.3243	0.5428	0.6594	0.8320	0.3239	0.3175	0.3931
CURLY	0.1384	0.3223	0.4422	0.5613	0.7186	0.9216	0.7687	1.3684	1.4660
DMSB	0.1325	0.3053	0.4341	0.6456	1.0244	1.1023	0.3030	0.2774	0.3662
MFM	0.1518	0.3704	0.4451	0.5533	0.6657	0.8838	0.6990	1.5897	1.7942
OT-CFM	0.1069	0.2784	0.3622	0.5228	0.7166	0.9261	0.8072	1.9095	2.0235
<b>PACE(ours)</b>	<b>0.0674</b>	<b>0.2087</b>	<b>0.2803</b>	<b>0.4385</b>	<b>0.4853</b>	<b>0.6238</b>	<b>0.2084</b>	<b>0.2067</b>	<b>0.2840</b>

### 4.4. Higher-dimensional single-cell benchmarks

High-dimensional PCA tests whether PACE still helps when Euclidean distances lose temporal contrast. On iPSC-Liu, PACE is best on all time-averaged 10D/50D metrics in Table 3. On OP-Cite/OP-Multi 100D, it remains on the empirical Pareto front, leading OP-Cite MMD and OP-Multi  $\mathcal{W}_1/\mathcal{W}_2$  in Appendix Table A.9. Figure 4 shows the diagnostic behind this regime, with norm CV approaching 0.3 and inter-time separation approaching the within-time radius. The useful signal is therefore not global separation between time means, but local anisotropy within each snapshot; PACE turns that local geometry into coupling costs. This helps distinguish directionally consistent moves from distance-similar but geometrically implausible pairings. Appendix H and Figures A.6–A.7 provide the detailed concentration analysis.

Table 3. iPSC-Liu (Liu et al., 2020) time-averaged results over holdouts  $t \in \{4, 8, 12, 16\}$ .

Dim	Method	MMD ↓	$\mathcal{W}_1$ ↓	$\mathcal{W}_2$ ↓
10D	Action Matching	0.5555	3.7801	4.3743
	Aligned CFM	<u>0.4644</u>	<u>2.9405</u>	<u>3.3825</u>
	CURLY	0.5369	3.6128	4.2102
	DMSB	0.7713	6.3088	6.4593
	MFM	0.5273	3.4066	3.8382
	OT-CFM	0.5145	3.3850	3.9447
	<b>PACE (ours)</b>	<b>0.4095</b>	<b>2.9226</b>	<b>3.1879</b>
50D	Action Matching	0.3125	7.6126	8.4868
	Aligned CFM	<u>0.2624</u>	<u>7.3133</u>	<u>8.0552</u>
	CURLY	0.3144	8.1144	8.9096
	DMSB	0.5638	10.9368	11.2423
	MFM	0.2853	7.7924	8.5454
	OT-CFM	0.2898	7.8392	8.5954
	<b>PACE (ours)</b>	<b>0.2550</b>	<b>6.7537</b>	<b>7.2836</b>

### 4.5. Ablation study

We finally isolate the contribution of the main PACE components on the Schiebinger2019 holdouts used in Table A.7. Figure 5 compares the full model with variants that remove coupling rematching, set the metric penalty to the Euclidean case, use all neighbors rather than local neighborhoods, or ablate the spatial and temporal kernel structure used to smooth local geometry. Each simplification degrades at least one metric or time point, and the full PACE variant gives the most stable low errors across MMD,  $\mathcal{W}_1$ , and  $\mathcal{W}_2$ . This indicates that the gains do not come from a single implementation detail: the anisotropic metric, local-neighborhood construction, spatial-temporal geometry smoothing, and rematching step all contribute to the final trajectory reconstruction.




Figure 4. High-dimensional concentration diagnostics. Dashed lines mark norm CV = 0.3 and inter-time/within-time ratio = 1.0.




Figure 5. PACE ablations on Schiebinger2019 (Schiebinger et al., 2019) holdouts. Lower is better; variants remove rematching, use Euclidean action ( $\alpha = 0$ ), use all neighbors, or drop spatial/temporal kernels.

## 5. Conclusion

In this paper, we have introduced PACE, a geometry-aware framework for trajectory inference from destructive single-cell time-course snapshots. PACE replaces fixed Euclidean endpoint costs with an anisotropic path cost induced by a local spatiotemporal metric, then alternates bridge learning with coupling refinement and distills the resulting dynamics into a velocity field. Across controlled and biological benchmarks, the results support local geometry as a useful inductive bias for held-out snapshot reconstruction without paired cells, lineage tracing, or velocity supervision; the velocity diagnostics further show that the learned flows recover plausible directionality when external velocity references are available only for evaluation. Ablations show that this behavior depends on the combination of anisotropic local costs, neighborhood-restricted geometry, smoothing, and coupling rematching rather than a single implementation detail. Future work will focus on more robust metric estimation, uncertainty quantification, and partially supervised or multimodal time-course settings.

## References

Arvanitidis, G., Hansen, L. K., and Hauberg, S. Latent space oddity: on the curvature of deep generative models. *arXiv preprint arXiv:1710.11379*, 2017.

Bergen, V., Lange, M., Peidli, S., Wolf, F. A., and Theis, F. J. Generalizing rna velocity to transient cell states through dynamical modeling. *Nature biotechnology*, 38(12):1408–1414, 2020.

Beyer, K., Goldstein, J., Ramakrishnan, R., and Shaft, U.

When is “nearest neighbor” meaningful? In *International conference on database theory*, pp. 217–235. Springer, 1999.

Cacchiarelli, D., Qiu, X., Srivatsan, S., Manfredi, A., Ziller, M., Overbey, E., Grimaldi, A., Grimsby, J., Pokharel, P., Livak, K. J., et al. Aligning single-cell developmental and reprogramming trajectories identifies molecular determinants of myogenic reprogramming outcome. *Cell Systems*, 7(3):258–268, 2018.

Chassignet, E. P., Hurlburt, H. E., Smedstad, O. M., Halliwell, G. R., Hogan, P. J., Wallcraft, A. J., Baraille, R., and Bleck, R. The hycom (hybrid coordinate ocean model) data assimilative system. *Journal of Marine Systems*, 65(1-4):60–83, 2007.

Chen, R. T. Q. and Lipman, Y. Flow matching on general geometries. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=g7ohD1TITL>.

Chen, T., Liu, G.-H., Tao, M., and Theodorou, E. Deep momentum multi-marginal schrödinger bridge. *Advances in Neural Information Processing Systems*, 36:57058–57086, 2023.

De Bortoli, V., Thornton, J., Heng, J., and Doucet, A. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in neural information processing systems*, 34:17695–17709, 2021.

Farrell, J. A., Wang, Y., Riesenfeld, S. J., Shekhar, K., Regev, A., and Schier, A. F. Single-cell reconstruction of developmental trajectories during zebrafish embryogenesis. *Science*, 360(6392):eaar3131, 2018.

Griffiths, J. A., Scialdone, A., and Marioni, J. C. Using single-cell genomics to understand developmental processes and cell fate decisions. *Molecular systems biology*, 14(4):MSB178046, 2018.

Haghverdi, L., Büttner, M., Wolf, F. A., Buettner, F., and Theis, F. J. Diffusion pseudotime robustly reconstructs lineage branching. *Nature methods*, 13(10):845–848, 2016.

Hashimoto, T., Gifford, D., and Jaakkola, T. Learning population-level diffusions with generative rnns. In *International Conference on Machine Learning*, pp. 2417–2426. PMLR, 2016.

Hauberg, S., Freifeld, O., and Black, M. A geometric take on metric learning. *Advances in Neural Information Processing Systems*, 25, 2012.

- 440 Huguet, G., Magruder, D. S., Tong, A., Fasina, O., Kuchroo,  
 441 M., Wolf, G., and Krishnaswamy, S. Manifold inter-  
 442 polating optimal-transport flows for trajectory inference.  
 443 *Advances in neural information processing systems*, 35:  
 444 29705–29718, 2022.
- 445 Kapuśniak, K., Potapchik, P., Reu, T., Zhang, L., Tong,  
 446 A., Bronstein, M., Bose, A. J., and Di Giovanni, F. Met-  
 447 ric flow matching for smooth interpolations on the data  
 448 manifold. *Advances in Neural Information Processing*  
 449 *Systems*, 37:135011–135042, 2024.
- 451 Kviman, O., Tamogashev, K., Branchini, N., Elvira, V.,  
 452 Lagergren, J., and Malkin, N. Multi-marginal flow match-  
 453 ing with adversarially learnt interpolants. In *The Four-*  
 454 *teenth International Conference on Learning Representations*,  
 455 2026. URL [https://openreview.net/](https://openreview.net/forum?id=AJls43yje7)  
 456 [forum?id=AJls43yje7](https://openreview.net/forum?id=AJls43yje7).
- 457 La Manno, G., Soldatov, R., Zeisel, A., Braun, E.,  
 458 Hochgerner, H., Petukhov, V., Lidschreiber, K., Kastri-  
 459 M. E., Lönnerberg, P., Furlan, A., et al. Rna velocity of  
 460 single cells. *Nature*, 560(7719):494–498, 2018.
- 462 Lance, C., Luecken, M. D., Burkhardt, D. B., Cannoodt, R.,  
 463 Rautenstrauch, P., Laddach, A., Ubingazhibov, A., Cao,  
 464 Z.-J., Deng, K., Khan, S., et al. Multimodal single cell  
 465 data integration challenge: results and lessons learned.  
 466 *BioRxiv*, pp. 2022–04, 2022.
- 467 Lange, M., Bergen, V., Klein, M., Setty, M., Reuter, B.,  
 468 Bakhti, M., Lickert, H., Ansari, M., Schniering, J.,  
 469 Schiller, H. B., et al. Cellrank for directed single-cell  
 470 fate mapping. *Nature methods*, 19(2):159–170, 2022.
- 472 Lavenant, H., Zhang, S., Kim, Y.-H., and Schiebinger, G.  
 473 Towards a mathematical theory of trajectory inference.  
 474 *arXiv preprint arXiv:2102.09204*, 2021.
- 475 Ledoux, M. *The concentration of measure phenomenon*.  
 476 Number 89. American Mathematical Soc., 2001.
- 478 Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and  
 479 Le, M. Flow matching for generative modeling. In *The*  
 480 *Eleventh International Conference on Learning Representations*,  
 481 2023. URL [https://openreview.net/](https://openreview.net/forum?id=PqvMRDCJT9t)  
 482 [forum?id=PqvMRDCJT9t](https://openreview.net/forum?id=PqvMRDCJT9t).
- 483 Liu, X., Ouyang, J. F., Rossello, F. J., Tan, J. P., Davidson,  
 484 K. C., Valdes, D. S., Schroeder, J., Sun, Y. B., Chen, J.,  
 485 Knaupp, A. S., et al. Reprogramming roadmap reveals  
 486 route to human induced trophoblast stem cells. *Nature*,  
 487 586(7827):101–107, 2020.
- 489 Moon, K. R., Van Dijk, D., Wang, Z., Gigante, S., Burkhardt,  
 490 D. B., Chen, W. S., Yim, K., Elzen, A. v. d., Hirn, M. J.,  
 491 Coifman, R. R., et al. Visualizing structure and transitions  
 492 in high-dimensional biological data. *Nature biotechnol-*  
 493 *ogy*, 37(12):1482–1492, 2019.
- 494 Neklyudov, K., Brekelmans, R., Severo, D., and Makhzani,  
 A. Action matching: Learning stochastic dynamics from  
 samples. In *International conference on machine learn-*  
*ing*, pp. 25858–25889. PMLR, 2023.
- Petrović, K., Atanackovic, L., Moro, V., Kapuśniak, K.,  
 Ceylan, I. I., Bronstein, M. M., Bose, J., and Tong, A.  
 Curly flow matching for learning non-gradient field dy-  
 namics. In *The Thirty-ninth Annual Conference on Neu-*  
*ral Information Processing Systems*, 2026. URL [https://](https://openreview.net/forum?id=7cqKVDgFZQ)  
[openreview.net/forum?id=7cqKVDgFZQ](https://openreview.net/forum?id=7cqKVDgFZQ).
- Peyré, G. and Cuturi, M. *Computational optimal transport:*  
*With applications to data science*. Now Foundations and  
 Trends, 2019.
- Saelens, W., Cannoodt, R., Todorov, H., and Saeys, Y. A  
 comparison of single-cell trajectory inference methods.  
*Nature biotechnology*, 37(5):547–554, 2019.
- Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subrama-  
 nian, V., Solomon, A., Gould, J., Liu, S., Lin, S., Berube,  
 P., et al. Optimal-transport analysis of single-cell gene  
 expression identifies developmental trajectories in repro-  
 gramming. *Cell*, 176(4):928–943, 2019.
- Setty, M., Kiseliovas, V., Levine, J., Gayoso, A., Mazutis, L.,  
 and Pe’Er, D. Characterization of cell fate probabilities  
 in single-cell data with palantir. *Nature biotechnology*,  
 37(4):451–460, 2019.
- Shen, Y., Berlinghieri, R., and Broderick, T. Multi-  
 marginal schrödinger bridges with iterative reference re-  
 finement. In *The 28th International Conference on Ar-*  
*tificial Intelligence and Statistics*, 2025. URL [https://](https://openreview.net/forum?id=VcwZ3gtYFY)  
[openreview.net/forum?id=VcwZ3gtYFY](https://openreview.net/forum?id=VcwZ3gtYFY).
- Shi, Y., De Bortoli, V., Campbell, A., and Doucet, A. Dif-  
 fusion schrödinger bridge matching. *Advances in neural*  
*information processing systems*, 36:62183–62223, 2023.
- Tanay, A. and Regev, A. Scaling single-cell genomics from  
 phenomenology to mechanism. *Nature*, 541(7637):331–  
 338, 2017.
- Tong, A., Huang, J., Wolf, G., Van Dijk, D., and Krish-  
 naswamy, S. Trajectorynet: A dynamic optimal transport  
 network for modeling cellular dynamics. In *International*  
*conference on machine learning*, pp. 9526–9536. PMLR,  
 2020.
- Tong, A., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks,  
 J., Fatras, K., Wolf, G., and Bengio, Y. Conditional flow  
 matching: Simulation-free dynamic optimal transport.  
*arXiv preprint arXiv:2302.00482*, 2(3), 2023.
- Trapnell, C., Cacchiarelli, D., Grimsby, J., Pokharel, P., Li,  
 S., Morse, M., Lennon, N. J., Livak, K. J., Mikkelsen,

- 495 T. S., and Rinn, J. L. The dynamics and regulators of cell  
496 fate decisions are revealed by pseudotemporal ordering of  
497 single cells. *Nature biotechnology*, 32(4):381–386, 2014.
- 498 Tritschler, S., Büttner, M., Fischer, D. S., Lange, M., Bergen,  
499 V., Lickert, H., and Theis, F. J. Concepts and limitations  
500 for learning developmental trajectories from single cell  
501 genomics. *Development*, 146(12):dev170506, 2019.
- 502  
503 Vershynin, R. *High-dimensional probability: An introduc-*  
504 *tion with applications in data science*, volume 47. Cam-  
505 bridge university press, 2018.
- 506  
507 Wagner, D. E., Weinreb, C., Collins, Z. M., Briggs, J. A.,  
508 Megason, S. G., and Klein, A. M. Single-cell mapping of  
509 gene expression landscapes and lineage in the zebrafish  
510 embryo. *Science*, 360(6392):981–987, 2018.
- 511  
512 Weinreb, C., Wolock, S., Tusi, B. K., Socolovsky, M., and  
513 Klein, A. M. Fundamental limits on dynamic inference  
514 from single-cell snapshots. *Proceedings of the National*  
515 *Academy of Sciences*, 115(10):E2467–E2476, 2018.
- 516  
517 Wolf, F. A., Hamey, F. K., Plass, M., Solana, J., Dahlin,  
518 J. S., Göttgens, B., Rajewsky, N., Simon, L., and Theis,  
519 F. J. Paga: graph abstraction reconciles clustering with  
520 trajectory inference through a topology preserving map  
521 of single cells. *Genome biology*, 20(1):59, 2019.

522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539  
540  
541  
542  
543  
544  
545  
546  
547  
548  
549

550 **Technical appendices and supplementary material**  
551  
552 A. Benchmark and Metric Details ..... 11  
553  
554 B. Baseline Details ..... 16  
555  
556 C. PACE Implementation Details ..... 16  
557  
558 D. Proof of Proposition 3.1 (Ill-posedness) ..... 17  
559  
560 E. Local Metric and Correspondence ..... 17  
561  
562 F. Monotonicity of Ideal Alternating Updates ..... 17  
563  
564 G. Stage 1 Regularizers ..... 18  
565  
566 H. High-Dimensional Concentration Effects ..... 18  
567  
568 I. iPSC-Liu Results ..... 19  
569  
570 J. Intuition behind alternating bridge and coupling opti-  
571 mization ..... 20  
572  
573 K. Adaptive tangent and normal projectors ..... 20  
574  
575 L. Adaptive bandwidth estimation ..... 20  
576  
577 M. Limitations and Future Work ..... 21

578 **A. Benchmark and Metric Details**

579 All datasets are cast as unpaired time-course point clouds.  
580 A point is a cell or synthetic particle state, and a time point  
581 is an empirical marginal distribution. Methods receive only  
582 the training snapshots; held-out snapshots are used only  
583 for evaluation. No method receives cell identities, cross-  
584 time correspondences, RNA velocity, or any ground-truth  
585 trajectory pairing.

587 **Shared split and rollout protocol.** For a held-out label  
588  $t_h$ , the evaluator finds the nearest observed training labels  
589  $t_a < t_h < t_b$ . The learned velocity field is rolled out  
590 from the empirical source cloud at  $t_a$  toward  $t_b$  on the same  
591 normalized training-time scale used during training. We  
592 use 101 Euler ODE steps and extract the intermediate frame  
593 at ratio  $(t_h - t_a)/(t_b - t_a)$ . The predicted cloud is then  
594 compared with the empirical held-out cloud at  $t_h$ . Thus, the  
595 reported numbers measure recovery of the missing marginal  
596 distribution, not recovery of unobserved cell identities.

598 **Ocean benchmark.** The Ocean benchmark follows the  
599 Gulf of Mexico vortex experiment of Shen et al. (Shen et al.,  
600 2025). The data are generated by first extracting a velocity  
601 field around a vortex feature from high-resolution HYbrid  
602 Coordinate Ocean Model (HYCOM) reanalysis data (Chas-  
603 signet et al., 2007), then simulating particles representing

buoys or ocean debris through that field. The original bench-  
mark contains approximately 1000 observations across five  
training times and four validation times; the released array  
used here has nine snapshots with 111 particles per snapshot.  
Unlike the reference-family model in Shen et al., PACE and  
all baselines in our comparison receive only the particle  
positions at training snapshots, with no velocity supervision.  
The HYCOM-derived velocity field is held out from training  
and used only for the velocity-alignment diagnostic in  
Table A.3.

**Velocity-alignment diagnostics.** Figure 3 summarizes the  
velocity-alignment checks used in the main experiments.  
The tables below report the numerical values behind the two  
panels: Ocean compares learned velocities with held-out  
simulator velocities, while EB PHATE compares learned 2D  
velocities with the RNA-velocity reference. These reference  
velocities are used only for diagnostics, not for training.

**Low-dimensional single-cell reconstruction results.** Ta-  
ble 2 reports the time-averaged summary used in the main  
text. The tables below give the corresponding per-timepoint  
results for each low-dimensional biological benchmark: EB  
PHATE has one held-out PHATE snapshot, iPSC-Liu has  
two held-out 2D PCA snapshots, and Schiebinger2019 has  
three held-out reprogramming days. These detailed tables  
show whether the averaged gains are consistent across indi-  
vidual held-out times.

**Subsampling and dimensionality.** For biological  
datasets, `samples_per_timepoint` is either an  
experiment-specific cap or `null`. When it is a cap,  
each selected time point is sampled without replacement  
using the experiment seed before whitening. When  
it is `null`, all available cells from the selected time  
points are used. For EB, the raw timepoint counts are  
2381, 4163, 3278, 3665, 3332 across labels  $0, \dots, 4$ . For  
iPSC-Liu and OP-Cite/OP-Multi, the dimensionality is  
chosen by taking the first  $d$  PCA coordinates from `X_pca`.  
The 10D and 50D iPSC-Liu tables aggregate separate  
single-holdout runs; in each run the held-out label is  
removed from the training label set.

**Metrics.** For every held-out time point, we compare the  
predicted point cloud with the observed point cloud us-  
ing three distributional metrics. MMD is computed with  
a Gaussian kernel and a median-heuristic bandwidth, and  
the reported value is the square root of the nonnegative  
MMD estimate.  $\mathcal{W}_1$  uses uniform empirical weights and  
the Euclidean ground cost between predicted and observed  
samples.  $\mathcal{W}_2$  uses uniform empirical weights and squared  
Euclidean ground cost, then reports the square root of the op-  
timal transport value. These metrics are distributional: they  
evaluate reconstruction of the held-out marginal distribution,

Table A.1. Dataset sources, representations, and held-out splits used in Section 4. Controlled 2D coordinates are used directly. Biological embeddings are standardized by fitting a `StandardScaler` on training snapshots only, then applying it to train and held-out snapshots.

Benchmark	Source and role	Representation and preprocessing	Train/test labels
Ocean (Gulf of Mexico vortex)	HYCOM-derived particle benchmark from Shen et al. (Shen et al., 2025), also used in CURLY (Petrović et al., 2026). Training uses only particle positions; velocities are not used for supervision.	<code>data/oceans/oceans.npz</code> ; <code>positions</code> has shape (9, 111, 2). No whitening; full frames are used.	Train {0, 2, 4, 6, 8}; test {1, 3, 5, 7}.
Two-Branch	In-house synthetic bifurcation generated by <code>src/data_preprocess/generate_two_branch_data.py</code> . It tests recovery of a trunk that splits into two branches.	Six 2D frames, 64 points per frame, with latent pseudotime and branch labels stored for diagnostics but not given to methods. No whitening; full frames are used.	Train {0, 2, 3, 5}; test {1, 4}.
EB PHATE	Embryoid body differentiation data from PHATE (Moon et al., 2019). It represents a low-dimensional differentiation trajectory.	<code>data/eb_velocity_v5.npz</code> ; 16,819 cells with <code>phate</code> ( $N, 2$ ) and <code>pcs</code> ( $N, 100$ ). The 2D experiment uses PHATE and fits whitening on training cells only.	Train {0, 1, 2, 4}; test {3}.
Schiebinger2019	Mouse reprogramming time course from Schiebinger et al. (Schiebinger et al., 2019), converted from the sc-NODE data package using serum-filtered FLE coordinates.	Local H5AD file; <code>X</code> stores 2D FLE coordinates and <code>obs["day"]</code> stores integer day labels. Whitening is fit on training cells only.	Train all integer days 0, ..., 18 except {6, 11, 16}; test {6, 11, 16}.
iPSC-Liu	Human reprogramming time course from Liu et al. (Liu et al., 2020). It is used to evaluate missing-timepoint reconstruction in PCA representations.	<code>data/iPSC_liu_pca.h5ad</code> ; <code>obsm["X_pca"]</code> provides PCA scores and <code>obs["timepoint"]</code> stores labels such as D0, parsed as integers. Whitening is fit on training cells only.	2D: train {0, 8, 12, 20, 24}, test {4, 16}. 10D/50D: separate single-holdout runs for $t \in \{4, 8, 12, 16\}$ .
OP-Cite/OP-Multi	NeurIPS Open Problems multimodal single-cell integration challenge data (Lance et al., 2022). OP-Cite uses CITE-seq features; OP-Multi uses Multiome features.	<code>op_cite_inputs_0.h5ad</code> and <code>op_train_multi_targets_0.h5ad</code> ; both use <code>obsm["X_pca"]</code> . Table A.9 uses all 100 PCs, with whitening fit on training cells only.	Train {2, 4, 7}; test {3}.

Table A.2. Branching toy (2D) per-timepoint results, holdout  $t \in \{1, 4\}$ .

Method	MMD ↓		$\mathcal{W}_1$ ↓		$\mathcal{W}_2$ ↓	
	$t = 1$	$t = 4$	$t = 1$	$t = 4$	$t = 1$	$t = 4$
Action Matching	0.6129	0.2905	0.3343	0.2732	0.3419	0.3048
Aligned CFM	0.6524	0.3439	0.3874	0.309	0.3944	0.3371
CURLY	0.2664	0.0767	0.1703	0.1797	0.2123	<b>0.2087</b>
DMSB	1.0232	1.1585	4.5832	2.8358	4.5845	2.8406
MFM	0.3817	0.1145	0.1686	0.1865	0.1832	0.2116
OT-CFM	0.5279	0.1244	0.2433	0.1841	0.2509	0.2119
<b>PACE(ours)</b>	<b>0.1968</b>	<b>0.0146</b>	<b>0.1386</b>	<b>0.1752</b>	<b>0.1686</b>	0.2236

Table A.3. Ocean (Petrović et al., 2026) velocity-alignment diagnostics on held-out time points. The simulator velocity is used only for evaluation. Cosine distance is  $1 - \cos(\hat{v}, v)$ , and normalized  $L_2$  compares unit-normalized velocity directions; lower is better.

Method	Cosine distance ↓					Normalized $L_2$ ↓				
	$t = 1$	$t = 3$	$t = 5$	$t = 7$	Avg.	$t = 1$	$t = 3$	$t = 5$	$t = 7$	Avg.
Action Matching	0.3535	0.0408	0.1974	0.1900	0.1954	0.8174	0.2317	0.5965	0.5993	0.5612
Aligned CFM	0.0440	0.0269	0.0101	0.0073	0.0221	0.2808	0.2118	0.1311	0.1122	0.1840
CURLY	0.0227	0.0068	0.0220	0.0231	0.0186	0.1997	0.0906	0.1709	0.1955	0.1642
MFM	<b>0.0029</b>	<b>0.0063</b>	0.0089	0.0043	<b>0.0056</b>	<b>0.0417</b>	0.0913	0.0965	0.0698	0.0748
OT-CFM	<u>0.0040</u>	0.0101	<u>0.0088</u>	0.0058	0.0072	<u>0.0738</u>	0.1147	<u>0.0731</u>	0.0801	0.0854
<b>PACE(ours)</b>	0.0041	<b>0.0030</b>	<b>0.0069</b>	<b>0.0031</b>	<b>0.0043</b>	0.0756	<b>0.0631</b>	<b>0.0671</b>	<b>0.0633</b>	<b>0.0673</b>

Table A.4. EB PHATE (Moon et al., 2019) velocity-alignment diagnostics at the held-out snapshot  $t = 3$ . The RNA-velocity reference is used only for evaluation. Cosine distance is  $1 - \cos(\hat{v}, v)$ , and normalized  $L_2$  compares unit-normalized velocity directions; lower is better.

Method	Cosine distance ↓	Normalized $L_2$ ↓
Action Matching	0.4041	0.7542
Aligned CFM	0.4273	0.7707
CURLY	0.6325	0.9693
MFM	0.6930	1.0261
OT-CFM	0.5571	0.8807
<b>PACE(ours)</b>	<b>0.3287</b>	<b>0.6630</b>

Table A.5. Per-timepoint results on EB PHATE (Moon et al., 2019) (2D), holdout  $t = 3$ .

Method	MMD ↓	$\mathcal{W}_1$ ↓	$\mathcal{W}_2$ ↓
	$t = 3$	$t = 3$	$t = 3$
Action Matching	0.2436	0.4553	0.5273
Aligned CFM	0.1056	0.2657	0.3243
CURLY	0.1384	0.3223	0.4422
DMSB	0.1325	0.3053	0.4341
MFM	0.1518	0.3704	0.4451
OT-CFM	0.1069	0.2784	0.3622
<b>PACE(ours)</b>	<b>0.0674</b>	<b>0.2087</b>	<b>0.2803</b>

Table A.6. Per-timepoint results on iPSC-Liu (Liu et al., 2020) (2D), holdout  $t \in \{4, 16\}$ .

Method	MMD ↓		$\mathcal{W}_1$ ↓		$\mathcal{W}_2$ ↓	
	$t = 4$	$t = 16$	$t = 4$	$t = 16$	$t = 4$	$t = 16$
Action Matching	0.7369	0.3090	0.9719	0.6290	1.0726	0.8733
Aligned CFM	0.8265	0.2590	0.7781	0.5407	0.7858	0.8781
CURLY	0.8156	0.3069	0.9030	0.5341	0.9292	0.9140
DMSB	0.8663	0.4248	0.7176	1.3311	0.7986	1.4060
MFM	0.8334	0.2731	0.7646	0.5667	0.7860	0.9815
OT-CFM	0.7787	0.2668	0.8695	0.5636	0.8983	0.9539
<b>PACE(ours)</b>	<b>0.6808</b>	<b>0.1962</b>	<b>0.5579</b>	<b>0.4126</b>	<b>0.5646</b>	<b>0.6830</b>

not cell-wise recovery of unobserved identities.

Table A.7. Per-timepoint results on Schiebinger2019 (Schiebinger et al., 2019), holdout  $t \in \{6, 11, 16\}$ .

Method	MMD ↓			$\mathcal{W}_1$ ↓			$\mathcal{W}_2$ ↓		
	$t = 6$	$t = 11$	$t = 16$	$t = 6$	$t = 11$	$t = 16$	$t = 6$	$t = 11$	$t = 16$
Action Matching	0.5056	0.4940	0.2174	0.3203	0.4982	0.4298	0.3283	0.5337	0.5642
Aligned CFM	0.5712	0.3405	<b>0.0601</b>	0.4126	0.3173	0.2228	0.4291	<u>0.3737</u>	<b>0.3764</b>
CURLY	0.9308	1.0201	0.3553	1.4737	1.9135	0.7182	1.5113	1.9386	0.9481
DMSB	0.5081	0.3196	0.0813	0.3118	0.3108	0.2096	0.3174	0.3760	0.4051
MFM	0.8676	0.8858	0.3437	2.6083	1.4602	0.7005	3.0265	1.4793	0.8768
OT-CFM	1.1550	1.0005	0.2661	3.1652	1.8933	0.6701	3.2608	1.9239	0.8856
<b>PACE(ours)</b>	<b>0.4219</b>	<b>0.1373</b>	<u>0.0660</u>	<b>0.2609</b>	<b>0.1590</b>	<b>0.2003</b>	<b>0.2669</b>	<b>0.2037</b>	<u>0.3813</u>

**Qualitative Stage-2 trajectory visualizations.** Figures A.1–A.5 complement the quantitative tables with selected 2D Stage-2 rollouts. For each dataset, the trajectory panel shows the learned rollout geometry for each baseline, while the prediction panel overlays the predicted held-out marginal with the empirical held-out snapshot. These plots are qualitative diagnostics only; all models are trained from unpaired snapshots without cell identities or ground-truth trajectories.

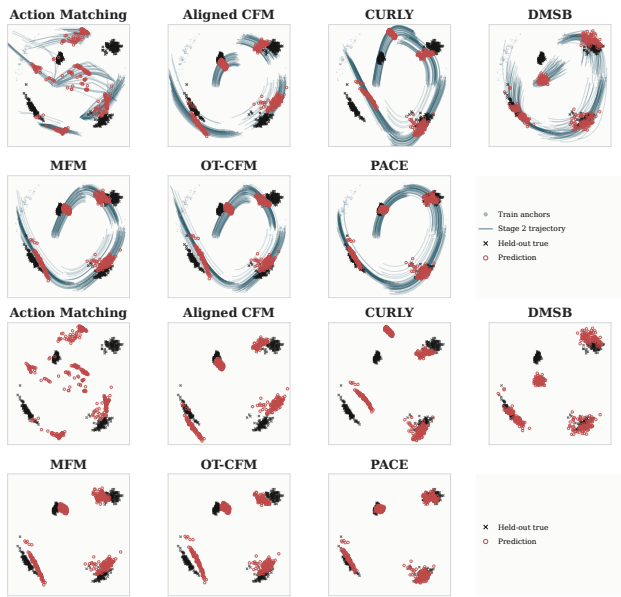


Figure A.1. Qualitative Stage-2 results on Ocean (Shen et al., 2025; Petrović et al., 2026). The top panel shows baseline rollouts; the bottom panel compares predicted and held-out point clouds.

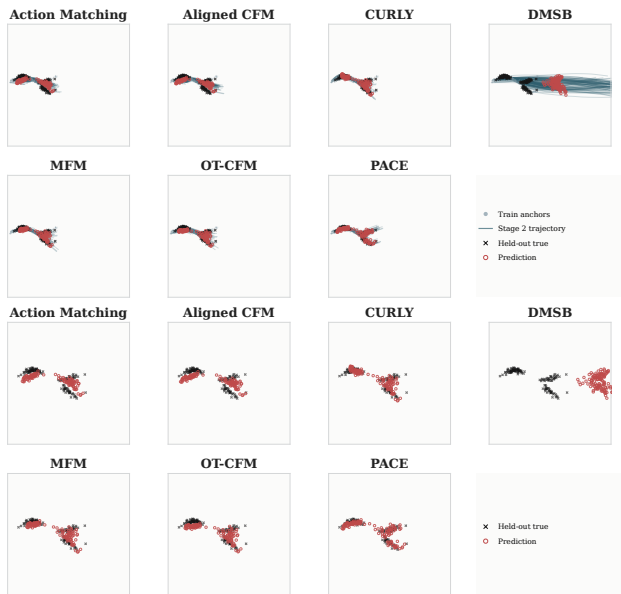


Figure A.2. Qualitative Stage-2 results on the Two-Branch toy benchmark. The top panel shows baseline rollouts; the bottom panel compares predicted and held-out point clouds.

770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809  
810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824

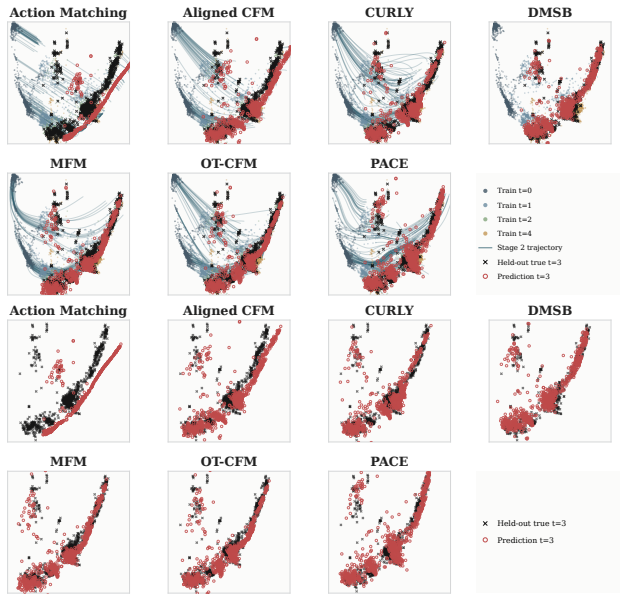


Figure A.3. Qualitative Stage-2 results on EB PHATE (Moon et al., 2019). The top panel shows baseline rollouts; the bottom panel compares predicted and held-out point clouds.

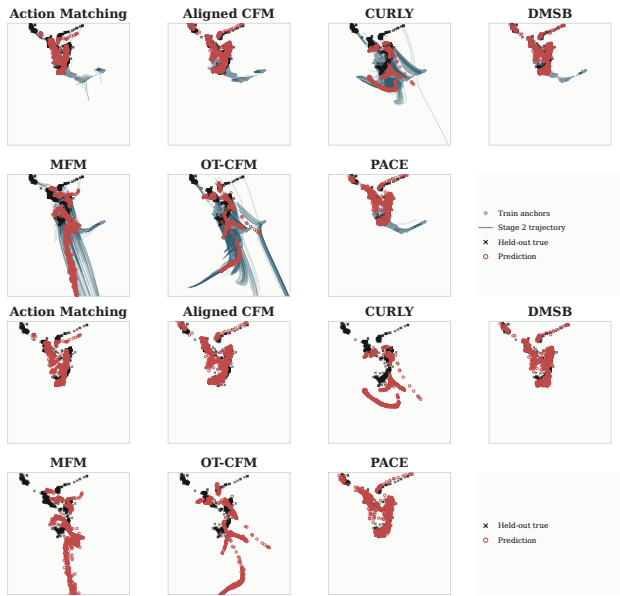


Figure A.4. Qualitative Stage-2 results on Schiebinger2019 (Schiebinger et al., 2019). The top panel shows baseline rollouts; the bottom panel compares predicted and held-out point clouds.

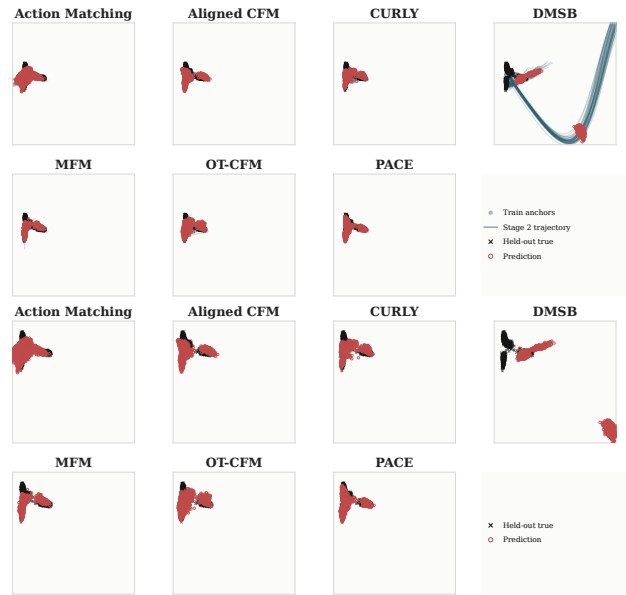


Figure A.5. Qualitative Stage-2 results on iPSC-Liu (Liu et al., 2020). The top panel shows baseline rollouts; the bottom panel compares predicted and held-out point clouds.

## B. Baseline Details

The result tables include only methods that were run for the corresponding experiment. All deterministic flow baselines use the same dataloaders, train/test labels, whitening rules, and held-out ODE rollout evaluator described in Appendix A. Unless an experiment-specific config overrides it, flow networks are MLP velocity fields trained with Adam/AdamW-style optimizers. The formulas below describe the implemented objectives under a segment  $[t_k, t_{k+1}]$ , with  $\Delta t = t_{k+1} - t_k$  and normalized segment time  $s = (t - t_k)/\Delta t \in [0, 1]$ .

**Action Matching.** Action Matching (Neklyudov et al., 2023) learns a scalar action network  $S_\theta(t, x)$ . For endpoint samples  $(x_0, x_1)$  and an interpolated point  $x_t = (1 - s)x_0 + sx_1$ , the implemented segment loss has the form

$$\begin{aligned} \mathcal{L}_{AM} = & \mathbb{E}[w(t_k)S_\theta(t_k, x_0) \\ & - w(t_{k+1})S_\theta(t_{k+1}, x_1) + \Delta t w(t)\partial_t S_\theta(t, x_t) \\ & + \frac{\Delta t}{2}w(t)\|\nabla_x S_\theta(t, x_t)\|^2 + \Delta t S_\theta(t, x_t)\partial_t w(t)]. \end{aligned} \quad (13)$$

At test time, the velocity is  $v_\theta(t, x) = \nabla_x S_\theta(t, x)$  and predictions are generated by solving  $\dot{x}_t = v_\theta(t, x_t)$ .

**OT-CFM.** OT-CFM (Tong et al., 2023) pairs mini-batches by an optimal-transport plan

$$\pi_k^* \in \arg \min_{\pi \in \Pi(\hat{\rho}_k, \hat{\rho}_{k+1})} \mathbb{E}_{(x_0, x_1) \sim \pi} \|x_0 - x_1\|^2. \quad (14)$$

For paired endpoints, it uses the straight conditional path

$$x_t = (1 - s)x_0 + sx_1, \quad u_t = \frac{x_1 - x_0}{\Delta t}, \quad (15)$$

and trains a velocity network with the flow-matching objective  $\mathbb{E}\|v_\theta(t, x_t) - u_t\|^2$ . The implementation uses the exact minibatch OT matcher from `torchcfm` and sets the conditional path noise to zero.

**MFM.** Metric Flow Matching (Kapuśniak et al., 2024) uses a two-stage pipeline. It first trains a geodesic correction network  $\psi_\eta$  under a data-induced metric, then trains a velocity field using the corrected interpolant

$$\mu_t = (1 - s)x_0 + sx_1 + \gamma(t)\psi_\eta(x_0, x_1, t), \quad (16)$$

where  $\gamma(t)$  vanishes at the segment endpoints. The flow-matching target is the time derivative

$$u_t = \frac{x_1 - x_0}{\Delta t} + \dot{\gamma}(t)\psi_\eta(x_0, x_1, t) + \gamma(t)\partial_t \psi_\eta(x_0, x_1, t), \quad (17)$$

with the last term present when the correction network depends explicitly on time. Our runs use the LAND-style data-manifold metric configuration.

**CURLY.** CURLY (Petrović et al., 2026) uses a related corrected interpolant, but with the normalized-time modulation  $s(1 - s)$ :

$$\mu_t = (1 - s)x_0 + sx_1 + s(1 - s)\psi_\eta(x_0, x_1, s). \quad (18)$$

The target velocity used for the second-stage flow model is

$$u_t = \frac{1}{\Delta t} [(x_1 - x_0) + (1 - 2s)\psi_\eta + s(1 - s)\partial_s \psi_\eta]. \quad (19)$$

The first stage fits  $\psi_\eta$  using pseudo-velocity supervision estimated from the observed temporal point clouds.

**Aligned CFM.** Aligned CFM follows the adversarially learned interpolant approach of ALI-CFM (Kviman et al., 2026). It trains a global interpolant  $I_\eta(x_0, x_1, t)$  from the first to the last training snapshot. Intermediate marginals of  $I_\eta$  are aligned with observed training snapshots by an adversarial objective of the schematic form

$$\begin{aligned} \min_{\eta} \max_D \sum_{\ell} & \left[ \mathbb{E}_{x \sim \hat{\rho}_\ell} \log D_\ell(x) \right. \\ & \left. + \mathbb{E}_{x_0, x_1} \log(1 - D_\ell(I_\eta(x_0, x_1, t_\ell))) \right]. \end{aligned} \quad (20)$$

After this stage, the interpolant is frozen and a velocity model is trained with

$$\mathcal{L}_{ALI-CFM} = \mathbb{E}\|v_\theta(t, I_\eta(x_0, x_1, t)) - \partial_t I_\eta(x_0, x_1, t)\|^2. \quad (21)$$

**DMSB.** DMSB (Chen et al., 2023) is a stochastic bridge baseline in joint position-velocity state space  $z_t = (x_t, v_t)$ . The implementation uses a momentum SDE discretization with learned control  $a_\theta(z_t, t)$ :

$$z_{n+1} = z_n + \begin{bmatrix} v_n \\ 0 \end{bmatrix} \Delta t + \sigma \begin{bmatrix} 0 \\ a_\theta \Delta t + \epsilon_n \end{bmatrix}, \quad (22)$$

with  $a_\theta = a_\theta(z_n, t_n)$  and  $\epsilon_n \sim \mathcal{N}(0, \Delta t I)$ . Forward and backward policies are trained on the temporal grid, and the learned dynamics are rolled out to predict the held-out marginal distributions. DMSB appears only in the tables where this baseline was run.

## C. PACE Implementation Details

PACE is implemented as a two-stage procedure. Stage 1 learns endpoint-preserving bridges and refines the cross-time matching; Stage 2 freezes the learned bridge correction and distills it into a global ODE velocity field for rollout. All reported experiments were run on a single NVIDIA A100 80GB GPU.

**Stage 1 anchor construction.** For each experiment, the selected training frames are first grouped by time label. PACE stacks the training frames into an anchor tensor of shape  $K \times N \times d$ , where  $K$  is the number of training labels and  $N$  is the minimum selected cell count across training time points. If time points have unequal counts, each frame is trimmed to this common  $N$  for the Stage 1 matching problem. This produces one adjacent matching problem per training interval.

**Initial and refined matchings.** The initial matching between adjacent anchors uses a cost combining normalized squared Euclidean distance with the local normal-direction penalty. After a fixed number of bridge-training epochs, PACE recomputes the matching by evaluating the learned path-action cost and solving the induced OT problem. We use the Python Optimal Transport (POT) package for OT computations. Rematching is repeated at a fixed experiment-specific interval (we usually choose 10 or 20), so bridge learning and correspondence estimation alternate throughout training.

**Bridge and flow distillation.** For endpoints  $(x_0, x_1)$  and local time  $s$ , the bridge has the endpoint-preserving form

$$\gamma_\theta(x_0, x_1, s) = (1-s)x_0 + sx_1 + s(1-s)\psi_\theta(x_0, x_1, s). \quad (23)$$

The bridge velocity  $\partial_s \gamma_\theta$  is obtained by automatic differentiation in 2D and by a Jacobian-vector product implementation in higher dimensions. After Stage 1,  $\psi_\theta$  is frozen. Stage 2 samples paired endpoints from the Stage 1 matchings and trains a global velocity network  $v_\phi(t, x)$  by regressing to the bridge velocity along generated points. Held-out predictions in the tables are produced by rolling out this Stage 2 velocity field.

## D. Proof of Proposition 3.1 (Ill-posedness)

*Proof of Proposition 3.1.* For one interval  $[t_k, t_{k+1}]$ , many couplings  $\pi_k \in \Pi(\hat{\rho}_k, \hat{\rho}_{k+1})$  can share the same endpoint marginals  $\hat{\rho}_k$  and  $\hat{\rho}_{k+1}$ . For any such coupling, each coupled endpoint pair  $(x, y)$  can be connected by infinitely many  $C^1$  paths  $\gamma : [0, 1] \rightarrow \mathbb{R}^d$  with  $\gamma(0) = x$  and  $\gamma(1) = y$ . Therefore, the observed snapshots determine neither a unique coupling nor a unique intermediate trajectory.  $\square$

## E. Local Metric and Correspondence

The metric used by PACE changes the matching criterion from endpoint proximity to path plausibility. In the locally constant idealization, suppose

$$G = I + \alpha P_N, \quad \alpha > 0,$$

where  $P_N$  is the orthogonal projector onto the local normal subspace. For a straight candidate displacement  $d = y - x$ ,

the metric action is

$$d^\top G d = \|d\|^2 + \alpha \|P_N d\|^2. \quad (24)$$

Thus, two candidate matches with the same Euclidean distance are not equivalent: the one whose displacement is more tangent-aligned has lower action. The implemented initial matching uses this principle through a normal-projection penalty, while the refined matching applies the same idea to the learned curved bridge by averaging

$$\dot{\gamma}_\theta(s)^\top G(\gamma_\theta(s), t(s)) \dot{\gamma}_\theta(s) \quad (25)$$

over a finite probe grid. This is the practical reason PACE can prefer a slightly longer but tangent-compatible path over a shorter chord that cuts across the inferred developmental geometry.

## F. Monotonicity of Ideal Alternating Updates

The alternating bridge-coupling updates can be viewed as block-coordinate descent for the discretized metric-action objective. For each adjacent time interval  $k$ , let

$$\Pi_k = \Pi(\hat{\rho}_k, \hat{\rho}_{k+1})$$

denote the feasible set of empirical couplings, and define the finite-dimensional objective

$$\mathcal{J}(\theta, \{\pi_k\}_{k=0}^{K-1}) = \sum_{k=0}^{K-1} \sum_{i,j} \pi_{ij}^{(k)} c_{ij}^{\text{path}}(\theta), \quad (26)$$

where  $c_{ij}^{\text{path}}(\theta)$  is the discretized path-action cost in Eq. (9). This objective is the full-batch version of the metric-action term optimized during the bridge update, using the same path-action quadrature as the coupling update.

**Proposition F.1** (Monotonicity under exact block updates). *Assume that, at iteration  $r$ , the bridge update computes a global minimizer*

$$\theta^{r+1} \in \arg \min_{\theta} \mathcal{J}(\theta, \{\pi_k^r\}_{k=0}^{K-1}),$$

*and the coupling update computes, for each  $k$ , a global minimizer*

$$\pi_k^{r+1} \in \arg \min_{\pi_k \in \Pi_k} \sum_{i,j} \pi_{ij}^{(k)} c_{ij}^{\text{path}}(\theta^{r+1}).$$

*Then the joint metric-action objective is non-increasing:*

$$\mathcal{J}(\theta^{r+1}, \{\pi_k^{r+1}\}_{k=0}^{K-1}) \leq \mathcal{J}(\theta^r, \{\pi_k^r\}_{k=0}^{K-1}). \quad (27)$$

*Moreover, the decrease is strict whenever either the bridge block or at least one coupling block achieves a strict improvement over the previous iterate.*

*Proof.* By the optimality of the bridge update with  $\{\pi_k^r\}$  fixed,

$$\mathcal{J}(\theta^{r+1}, \{\pi_k^r\}_{k=0}^{K-1}) \leq \mathcal{J}(\theta^r, \{\pi_k^r\}_{k=0}^{K-1}).$$

By the optimality of each coupling update with  $\theta^{r+1}$  fixed,

$$\begin{aligned} \sum_{i,j} \pi_{ij}^{(k),r+1} c_{ij}^{\text{path}}(\theta^{r+1}) \\ \leq \sum_{i,j} \pi_{ij}^{(k),r} c_{ij}^{\text{path}}(\theta^{r+1}) \quad \text{for every } k. \end{aligned}$$

Summing these inequalities over  $k$  gives

$$\mathcal{J}(\theta^{r+1}, \{\pi_k^{r+1}\}_{k=0}^{K-1}) \leq \mathcal{J}(\theta^{r+1}, \{\pi_k^r\}_{k=0}^{K-1}).$$

Combining the two inequalities proves monotonicity. If either inequality is strict, the combined decrease is strict.  $\square$

Because  $G_k(x, t) = I + \alpha C_N^{(k)}(x, t)$  is positive definite, every path-action cost is nonnegative; therefore the sequence of ideal objective values is bounded below by zero and hence converges as a sequence of numbers. This does not imply convergence to a global optimum, nor does it imply uniqueness of the bridge or coupling. Equality can occur when a block update returns an equivalent minimizer or when the current block is already optimal.

**Relation to the implemented algorithm.** The proposition applies to the ideal metric-action problem in Eq. (26). The implemented PACE training approximates this scheme with stochastic minibatches, a neural bridge optimized by finite gradient steps, periodic rather than continuous rematching, and optional stabilizing regularizers. These choices are used for scalability and numerical stability, but they do not provide a per-gradient-step strict decrease guarantee for the realized neural training trajectory. The monotonic result should therefore be read as the optimization principle behind the alternating updates, while the empirical sections evaluate the behavior of the practical stochastic implementation.

## G. Stage 1 Regularizers

PACE Stage 1 optimizes a metric action together with an optional regularization term:

$$\mathcal{L}_{\text{bridge}} = \lambda_{\text{metric}} \mathcal{L}_{\text{metric}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}. \quad (28)$$

The metric-action term is the discretized version of Eq. (8). The regularization term can include the implementation-level stabilizers below; in our experiments, this corresponds to a weighted combination

$$\lambda_{\text{reg}} \mathcal{L}_{\text{reg}} = \lambda_{\text{coh}} \mathcal{L}_{\text{coh}} + \lambda_{\text{orth}} \mathcal{L}_{\text{orth}}.$$

**Cross-segment velocity coherence.** Let  $(x_a, t_a, v_a)$  and  $(x_b, t_b, v_b)$  denote generated bridge points and velocities from different adjacent segments. PACE constructs Gaussian space-time weights

$$\begin{aligned} W_{ab} = \exp\left(-\frac{\|x_a - x_b\|^2}{\sigma_{x,a}\sigma_{x,b}}\right) \exp\left(-\frac{|t_a - t_b|^2}{\sigma_{t,a}\sigma_{t,b}}\right) \\ \times \mathbf{1}\{\text{seg}(a) \neq \text{seg}(b)\}, \end{aligned} \quad (29)$$

where the bandwidths are estimated from local projector variation. The coherence loss penalizes nearby generated points from different segments when their normalized velocities disagree:

$$\mathcal{L}_{\text{coh}} = \frac{\sum_{a,b} W_{ab} (1 - \hat{v}_a^\top \hat{v}_b)}{\sum_{a,b} W_{ab} + \varepsilon}, \quad \hat{v}_a = \frac{v_a}{\|v_a\| + \varepsilon}. \quad (30)$$

**Normal-motion suppression.** For generated bridge points, PACE estimates a local normal projector and penalizes the normal component of the generated velocity:

$$\mathcal{L}_{\text{orth}} = \frac{1}{|\mathcal{G}|} \sum_{(x,t,v) \in \mathcal{G}} \|P_N(x, t)v\|^2. \quad (31)$$

In 2D this reduces to a squared dot product with the locally estimated normal vector. In higher dimensions it uses the full normal-space projector  $P_N = I - P_T$ .

## H. High-Dimensional Concentration Effects

The high-dimensional experiments in Section 4 operate in a regime where Euclidean probability mass and pairwise costs are strongly concentrated. This section explains the diagnostics used in Figure 4 and why concentration affects OT, MMD, and nearest-neighbor-based trajectory objectives.

**Norm concentration.** For  $X \sim \mathcal{N}(0, I_d)$ , the map  $f(x) = \|x\|_2$  is 1-Lipschitz. The Gaussian concentration inequality therefore gives

$$\Pr(\|X\|_2 - \mathbb{E}\|X\|_2 \geq t) \leq 2 \exp(-t^2/2), \quad (32)$$

and  $\mathbb{E}\|X\|_2 \asymp \sqrt{d}$  implies

$$\frac{\text{Std}(\|X\|_2)}{\mathbb{E}\|X\|_2} = O(d^{-1/2}). \quad (33)$$

This is the standard Lipschitz concentration phenomenon (Ledoux, 2001): after whitening or PCA scaling, most points are pushed toward a thin shell rather than spreading across many radial scales.

**Direction and nearest-neighbor concentration.** For independent isotropic sub-Gaussian vectors  $X, Y \in \mathbb{R}^d$ , normalized directions satisfy

$$\left\langle \frac{X}{\|X\|_2}, \frac{Y}{\|Y\|_2} \right\rangle = O_p(d^{-1/2}), \quad (34)$$

so random directions become nearly orthogonal as  $d$  grows (Vershynin, 2018). A related nearest-neighbor result states that if the relative variance of distances vanishes,

$$\frac{\text{Var}(\|Q - X_1\|_2)}{\mathbb{E}[\|Q - X_1\|_2]^2} \rightarrow 0, \quad (35)$$

then

$$\Pr\left(\frac{D_{\max}^{(d)} - D_{\min}^{(d)}}{D_{\min}^{(d)}} \leq \epsilon\right) \rightarrow 1 \quad \text{for every } \epsilon > 0, \quad (36)$$

where  $D_{\min}^{(d)}$  and  $D_{\max}^{(d)}$  are the nearest and farthest distances from the query point (Beyer et al., 1999). Thus, when distance contrast collapses, the nearest neighbor and farthest neighbor become less distinguishable in relative terms.

**Effect on OT, MMD, and flow-matching losses.** If two approximately isotropic samples  $X, Y \in \mathbb{R}^d$  have coordinate-wise fluctuations with comparable scale, then

$$\|X - Y\|_2^2 = \sum_{j=1}^d (X_j - Y_j)^2 \quad (37)$$

is a sum of many coordinate-level contributions. Under standard independence or weak-dependence assumptions, a law-of-large-numbers or sub-Gaussian concentration argument implies that  $\|X - Y\|_2^2/d$  concentrates around its mean, with relative fluctuations that typically scale as  $O(d^{-1/2})$ . For transport objectives,

$$W_2^2(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \mathbb{E}_{(x,y) \sim \pi} \|x - y\|_2^2, \quad (38)$$

so concentration directly reduces the dynamic range of the cost matrix used by OT and flow-matching couplings. For kernel metrics,

$$\text{MMD}^2(\mu, \nu) = \mathbb{E}[k(x, x') + k(y, y') - 2k(x, y)], \quad (39)$$

and a distance-based kernel such as  $k(x, y) = \exp(-\|x - y\|_2^2/(2\sigma^2))$  also loses contrast when most  $\|x - y\|_2^2$  values occupy a narrow interval. Consequently, different predicted distributions can receive similar numerical losses even when their local geometry differs. This is the metric-degeneration issue encountered by high-dimensional distance-based trajectory methods, including metric and OT flow-matching objectives (Kapuśniak et al., 2024).

**Reference thresholds in Figure 4.** The first diagnostic is

$$\text{CV}_{\text{norm}} = \frac{\text{Std}(\|X\|_2)}{\mathbb{E}\|X\|_2}. \quad (40)$$

The horizontal reference level 0.3 is a practical warning threshold: below it, one standard deviation of radial variation is less than 30% of the mean radius, so most cells lie in

Table A.8. iPSC-Liu (Liu et al., 2020) (10D/50D) per-timepoint results across holdouts  $t \in \{4, 8, 12, 16\}$ .

Dim	Method	MMD ↓				W <sub>1</sub> ↓				W <sub>2</sub> ↓			
		t = 4	t = 8	t = 12	t = 16	t = 4	t = 8	t = 12	t = 16	t = 4	t = 8	t = 12	t = 16
10D	Action Matching	0.6504	0.6959	0.5132	0.3626	3.8458	4.1608	3.1014	4.0123	4.2076	5.9396	3.4279	3.9222
	Aligned CFM	0.6245	0.5936	0.4051	<b>0.2345</b>	3.7937	3.4433	2.5372	<b>1.9877</b>	4.2084	4.0761	2.9973	<b>2.2482</b>
	CURLY	<b>0.5461</b>	0.6054	0.4936	0.5024	3.8694	3.9773	3.1387	3.4660	4.3490	5.5607	3.4974	3.4336
	DMSB	0.9250	0.7057	0.7448	0.7098	10.5837	4.1953	5.5040	4.9522	10.6553	4.5341	5.6287	5.0190
	MF	0.3864	0.6674	0.3033	0.3521	3.5972	4.0239	3.2515	2.7539	4.0205	5.0654	3.5968	2.6699
	OT-CFM	0.6039	0.6339	0.4951	0.3251	3.6543	4.0040	3.2932	2.5884	4.0624	5.5807	3.6325	2.5030
	PACE(ours)	<b>0.5244</b>	<b>0.5153</b>	<b>0.3113</b>	<b>0.2871</b>	<b>3.5918</b>	<b>3.2507</b>	<b>2.3185</b>	<b>2.5294</b>	<b>3.7485</b>	<b>3.5755</b>	<b>2.7868</b>	<b>2.6406</b>
	50D	Action Matching	<b>0.3517</b>	0.3506	0.3278	0.2197	7.2261	8.4554	7.9593	6.8096	7.8273	10.3927	8.5704
Aligned CFM	0.3519	<b>0.2917</b>	0.2397	0.1663	7.0562	7.5723	7.5947	7.0302	7.6855	8.2791	8.2483	8.0080	
CURLY	0.3824	0.4189	0.2235	0.2327	8.2766	9.7010	7.6085	6.8715	8.8476	11.2946	8.2990	7.1972	
DMSB	0.7593	0.5730	0.4787	0.4441	14.5899	10.4588	9.7531	8.9456	14.7121	10.7933	10.1192	9.3447	
MF	0.3539	0.3467	0.2679	0.1728	<b>7.0197</b>	9.2505	8.5746	<b>6.3246</b>	<b>7.6665</b>	10.5394	9.2502	<b>6.7254</b>	
OT-CFM	0.3530	0.3458	0.2802	0.1803	7.1278	9.0608	8.7395	6.4288	7.7653	10.4423	9.3687	6.8054	
PACE(ours)	0.3609	<b>0.3004</b>	<b>0.2193</b>	<b>0.1396</b>	7.0656	<b>6.8194</b>	<b>6.8718</b>	<b>6.2581</b>	7.6896	<b>7.0798</b>	<b>7.5736</b>	<b>6.7914</b>	

a relatively thin shell and pairwise Euclidean costs are dominated by small angular or local fluctuations. This threshold is not a theorem-specific cutoff; it is a scale marker chosen to make the  $O(d^{-1/2})$  collapse visible on the empirical PCA representations.

The second diagnostic compares time separation with within-time dispersion. For snapshot centers  $c_i, c_j$ , define

$$R_{\text{time}} = \text{median}_{i < j} \frac{\|c_i - c_j\|_2}{\bar{r}_{ij}}, \quad (41)$$

$$\bar{r}_{ij} = \frac{1}{2} \left( \mathbb{E}_{X \sim \mu_i} \|X - c_i\|_2 + \mathbb{E}_{Y \sim \mu_j} \|Y - c_j\|_2 \right).$$

The reference level 1.0 marks the point where the median displacement between time-point centers is comparable to the typical within-time radius. When  $R_{\text{time}} \leq 1$ , the Euclidean shift between snapshots is no larger than the spread of the snapshots themselves, so time labels are difficult to separate using raw pairwise distances alone. Figure 4 summarizes these diagnostics in the main text. The appendix diagnostics in Figures A.6 and A.7 show the corresponding norm-concentration behavior for iPSC-Liu, OP-Cite, and OP-Multi.


**High-dimensional benchmark results.** Table A.8 reports the per-timepoint reconstruction metrics for iPSC-Liu (Liu et al., 2020) in 10D and 50D PCA representations, evaluated on holdout time points  $t \in \{4, 8, 12, 16\}$ . Table A.9 gives the corresponding results for the multimodal OP-Cite and OP-Multi datasets (Lance et al., 2022) in 100D. These tables complement the time-averaged summary in Table 3 and the concentration diagnostics in Figure A.6.

## I. iPSC-Liu Results

Figure A.7 shows the norm-concentration diagnostics for iPSC-Liu across increasing PCA dimensions, illustrating the  $O(d^{-1/2})$  norm-concentration behavior discussed in Appendix H. These panels correspond to the high-dimensional benchmark experiments reported in Table A.8 and summarized in the main text in Table 3.


Table A.9. OP-Cite and OP-Multi (Lance et al., 2022) (100D) results.

Method	OP-CITE (100D)			OP-MULTI (100D)		
	MMD ↓	$\mathcal{W}_1$ ↓	$\mathcal{W}_2$ ↓	MMD ↓	$\mathcal{W}_1$ ↓	$\mathcal{W}_2$ ↓
Action Matching	0.1753	<b>9.9388</b>	<b>10.0448</b>	<b>0.1622</b>	10.7711	10.8160
Aligned CFM	<u>0.1434</u>	10.7009	10.8094	0.1720	10.8330	10.8826
CURLY	0.2529	11.5163	11.6015	0.2786	12.0015	12.0440
MFM	0.1435	<u>10.4301</u>	<u>10.5409</u>	0.1699	10.6274	10.6702
OT-CFM	0.1448	10.4447	10.5491	<u>0.1662</u>	<u>10.6128</u>	<u>10.6564</u>
<b>PACE(ours)</b>	<b>0.1409</b>	10.4574	10.5605	0.1693	<b>10.5824</b>	<b>10.6293</b>


 Figure A.6. Norm-concentration diagnostics for iPSC-Liu (Liu et al., 2020) and OP-Cite/OP-Multi (Lance et al., 2022) representations. The panels illustrate the empirical thin-shell behavior summarized by  $CV_{\text{norm}}$ , with the 0.3 reference level used as a practical concentration warning threshold.

## J. Intuition behind alternating bridge and coupling optimization

**Why coupling matters for the bridge.** The neural bridge is trained on endpoint pairs sampled from the current coupling. If the coupling is biologically implausible (for example, pairing an early stem cell with a terminally differentiated cell), the bridge must learn a path that traverses the entire developmental manifold in a single segment. Such a path is likely to pass through regions where the metric strongly penalizes normal motion, yielding high action and making it difficult for the network to find a low-action deformation. Conversely, when the coupling pairs cells at nearby pseudotime stages, the displacement is small and largely tangent-aligned, so the bridge can easily learn a geometry-consistent interpolant. In short, the coupling determines the training set for the bridge; wrong pairs give the network an impossible learning problem.


 Figure A.7. Norm-concentration diagnostics for iPSC-Liu (Liu et al., 2020) representations. Increasing PCA dimension reduces relative radial variation, consistent with the  $O(d^{-1/2})$  norm-concentration behavior described in Appendix H.

**Why the bridge matters for the coupling.** The coupling update solves an OT problem whose cost matrix is the path action evaluated on the neural bridge (Eq. (9)). If the bridge is untrained, its paths are nearly straight lines and the cost reduces to Euclidean distance; the OT solver then cannot distinguish a biologically plausible pairing from an implausible one. Once the bridge has learned metric-aware paths, the action cost becomes geometry-sensitive. Pairs whose displacement aligns with the local developmental tangent incur low action, while pairs that would require motion across normal (non-developmental) directions incur high action. The bridge therefore transforms an uninformative Euclidean cost into a geometry-aware cost, allowing OT to select biologically meaningful pairings.

**Why amortize with a shared neural network.** Without amortization, every candidate endpoint pair would require solving an independent boundary-value problem (the geodesic equation) under the state-dependent metric  $G_k$ . With  $N_k N_{k+1}$  candidate pairs, this is computationally infeasible. The neural bridge amortizes this cost by learning a single parametric family  $\gamma_\theta(x, y, \tau)$  that approximates the minimum-action path for all pairs. The approximation need not be perfect at initialization; it only needs to be good enough to provide a more informative cost than Euclidean distance, and it improves as the coupling improves.

## K. Adaptive tangent and normal projectors

The weighted local covariance at each anchor point yields eigenvalues  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d$  and orthonormal eigenvectors  $v_1, \dots, v_d$ . The tangent subspace is spanned by the leading (maximum-variance) eigenvectors. Its effective dimension  $q_r$  is chosen as the smallest integer such that the cumulative explained variance reaches a prescribed threshold  $\tau \in (0, 1)$  (e.g., 0.95):

$$q_r = \min \left\{ q : \frac{\sum_{j=1}^q \lambda_{d-j+1}}{\sum_{j=1}^d \lambda_j} \geq \tau \right\}.$$

The tangent projector is  $P_T^{(r)} = \sum_{j=1}^{q_r} v_{d-j+1} v_{d-j+1}^\top$  and the normal projector is its orthogonal complement  $P_N^{(r)} = I - P_T^{(r)}$ . In two dimensions this reduces to  $P_N^{(r)} = n_r n_r^\top$  with  $n_r = v_1$  and  $q_r = 1$ .

## L. Adaptive bandwidth estimation

For each segment  $k$ , the spatial and temporal bandwidths  $h_x^{(k)}$  and  $h_t^{(k)}$  are estimated from the anchor normal bank rather than set by hand. The estimation aggregates information over a small window of anchor snapshots centered on segment  $k$  (by default one snapshot on each side). The same local eigendecomposition also yields the tangent subspace

$P_T^{(r)}$  (Appendix K), which is used for bandwidth estimation below but does not enter the metric tensor directly.

**Spatial geometry per snapshot.** For each anchor snapshot  $a$  in the window, PACE computes three quantities from the  $k$  nearest neighbors of every cell (excluding the cell itself):

- **Local spatial spacing**  $\Delta_x^{(a)}$ : the median tangential step size  $\|\langle x_j - x_i, t_i \rangle\|$  between each cell  $i$  and its neighbors  $j$ , where  $t_i$  is the local tangent direction. This captures how densely cells are spaced along the developmental manifold.
- **Normal spatial rate**  $\rho_{N,x}^{(a)}$ : the median Frobenius-norm difference  $\|P_N^{(j)} - P_N^{(i)}\|_F / \Delta_x^{(a)}$  between the normal projectors of neighboring cells. This measures how rapidly the local geometry changes in space.
- **Tangent spatial rate**  $\rho_{T,x}^{(a)}$ : the analogous quantity for tangent projectors.

**Cross-snapshot temporal rates.** For each consecutive pair  $(a, a+1)$  in the window, PACE matches cells by nearest-neighbor correspondence and computes:

- **Normal temporal rate**  $\rho_{N,t}^{(a,a+1)}$ : the Frobenius-norm difference  $\|P_N^{(nn(i))} - P_N^{(i)}\|_F / |t_{a+1} - t_a|$  between matched cells, divided by the experimental time gap. This measures how rapidly the local geometry evolves over time.
- **Tangent temporal rate**  $\rho_{T,t}^{(a,a+1)}$ : the analogous quantity for tangent projectors.

**Segment bandwidth assembly.** All per-snapshot and cross-snapshot quantities are concatenated and their positive medians are taken, yielding four segment-level scalars:  $\Delta_x^{(k)}$ ,  $\rho_{N,x}^{(k)}$ ,  $\rho_{N,t}^{(k)}$ , and  $\rho_{T,t}^{(k)}$ . The metric bandwidths are then

$$h_x^{(k)} = \max\left(\Delta_x^{(k)}, \frac{1}{\rho_{N,x}^{(k)}}\right), \quad (42)$$

$$c_N^{(k)} = \frac{\rho_{N,t}^{(k)}}{\rho_{N,x}^{(k)}}, \quad (43)$$

$$h_t^{(k)} = \frac{h_x^{(k)}}{c_N^{(k)}}. \quad (44)$$

The spatial bandwidth  $h_x$  is the larger of the local spacing and the reciprocal spatial variation scale, preventing the kernel from being narrower than either the point density or

the geometry variation. The ratio  $c_N^{(k)}$  converts spatial scale into equivalent temporal scale through the speed at which the normal geometry changes in time relative to space. A similar pair  $(\sigma_x^{(k)}, \sigma_t^{(k)})$  is computed from the tangent rates for the coupling-refinement kernel, but the metric field uses  $(h_x^{(k)}, h_t^{(k)})$ .

## M. Limitations and Future Work

PACE has several limitations. First, its local metric depends on nearest-neighbor covariance estimates within each snapshot. When data are sparse, noisy, or highly heterogeneous, the estimated tangent and normal subspaces may be unstable; adaptive neighborhood selection may improve robustness.

Second, PACE is optimized with stochastic neural training and periodic rematching. Although Appendix F states a monotonicity property for ideal exact block-coordinate updates, the implemented algorithm is only a scalable approximation and does not guarantee global optimality.

Third, destructive snapshot data do not provide ground-truth correspondences or continuous cell histories. The reported metrics evaluate held-out marginal reconstruction rather than individual trajectory correctness. PACE-inferred paths should be interpreted as plausible geometry-consistent reconstructions, not directly observed cellular histories.

Future work will extend PACE beyond closed-population transport by incorporating birth-death or growth terms (Schiebinger et al., 2019). Another direction is to learn representations that preserve local trajectory geometry in high-dimensional settings, where distance concentration can make coupling costs less discriminative.