

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LEARNING CELLULAR DYNAMICS WITH CELL-CELL INTERACTION-AWARE OPTIMAL TRANSPORT

Anonymous authors

Paper under double-blind review

ABSTRACT

Inferring dynamics from population snapshots is a core challenge in machine learning and biology. In scRNA-sequencing (scRNA-seq), destructive measurements yield irregular, high-dimensional samples of cell states, obscuring how populations evolve. Existing trajectory inference methods either use graph heuristics or cast alignment as an Optimal Transport (OT) problem. However, they treat cells as independent points, ignoring intercellular interactions. In this work, we ask whether incorporating cell-cell interactions can improve the reconstruction of cellular dynamics from scRNA-seq snapshots. We introduce **IADOT** (*Interaction-Aware Dynamic Optimal Transport*), which integrates cell-cell interaction networks into an OT objective and then learns a time-continuous vector field via Conditional Flow Matching. Across a synthetic task and diverse scRNA-seq datasets, we find that incorporating interaction structure can improve snapshot alignment and inference of cellular dynamics versus feature-only baselines. **IADOT** also supports in-silico ligand-receptor perturbation analyses: we show on lung cancer data that inferred trajectories are sensitive to edits of the ligand-receptor catalog, consistent with known effects of targeted pathway inhibition.

1 INTRODUCTION

Single-cell technologies have turned the study of gene expression into a high-resolution, data-driven science (Picelli, 2016). By exposing cellular heterogeneity directly, these methods are reshaping how we approach complex biological systems (Cha & Lee, 2020). For instance, in embryonic development, they have traced lineage bifurcations that give rise to distinct tissues (Qiu et al., 2022). In oncology, they exposed how cancer populations branch and adapt (Yeo et al., 2022). More broadly, the capacity to measure cellular states at scale calls for computational methods that can recover the underlying dynamical rules of biology (Schiebinger et al., 2019). Importantly, such approaches hold major implications for pharmaceutical research, where experimental campaigns to explore disease mechanisms or evaluate therapeutic interventions are prohibitively costly and time-consuming (Sertkaya et al., 2024). By enabling *in silico* reconstruction and prediction of cellular dynamics, computational models can guide experiment design, prioritize drug targets, and reduce the need for exhaustive laboratory screening (Yue & Dutta, 2022).

Challenges of inferring cellular dynamics. Despite these advances, reconstructing cellular dynamics from single-cell measurements presents fundamental difficulties (Bunne et al., 2024). Measurements are destructive: the same cell cannot be followed over time, so there is no one-to-one correspondence between adjacent snapshots. Populations are imbalanced, with varying numbers of cells in each state, making one-to-one mappings ill-suited (Schiebinger et al., 2019). Gene expression measurements are noisy and sampled irregularly, and the ambient dimensionality of thousands of genes exacerbates statistical and computational difficulties (Adil et al., 2021). Reconstructing dynamics from such data means inferring smooth trajectories from noisy, unaligned population snapshots under partial observability.

Aligning snapshots with Optimal Transport. Classical trajectory inference constructs a cell-cell k NN graph in a low-dimensional embedding and then extracts pseudotime and branches via principal curves, diffusion distances (Haghverdi et al., 2016), or graph geodesics/spanning trees (Street et al., 2018). These locality-based heuristics implicitly assume geometric proximity within a snapshot reflects temporal adjacency and differentiation proceeds along geodesics of the learned mani-

fold. This often results in biased pseudotimes and spurious lineage structure (Saelens et al., 2022). More recent methods (Schiebinger et al., 2019; Bunne et al., 2023b) instead recast cell alignment as a global, uncertainty-aware coupling between *multiple* distributions via Optimal Transport (OT). This formulation has distinct advantages, as it produces soft correspondences, naturally handles unequal sample sizes, and encourages low-action trajectories via its prior.

Cell-cell interactions. Conventional OT-based alignment matches cells by minimizing distances in gene-expression space, effectively treating cells as independent points and ignoring the interaction networks that connect them. Consequently, they overlook potential smoothness and directionality in cell-cell interactions (CCIs) unless external structure (e.g., spatial coordinates) is provided (Klein et al., 2025a). This omission is at odds with the central role of directed CCIs in many applications, including pharmacological targeting (He & Xu, 2020; Liu et al., 2023). This motivates the following question: *Can structure derived from CCIs provide useful information to improve the reconstruction of cellular dynamics from scRNA-seq snapshots only?* To answer this, we introduce IADOT (*Interaction-Aware Dynamic Optimal Transport*), a framework that integrates gene-expression features and interaction networks into a single OT objective. IADOT constructs a directed CCI tensor from ligand–receptor expression at each snapshot, and optimizes an OT objective with two components: a feature cost in expression space and a structure cost that favors couplings *preserving specific CCI patterns across time*. The resulting soft couplings align snapshots while respecting the CCI structure, and we use them to learn continuous-time dynamics by training a velocity field via different flow matching techniques.

Contributions

Conceptually, we formalize trajectory inference as learning dynamics in interacting subsystems, introducing a structure-regularized OT objective where the persistence of directed, typed interaction networks (e.g., ligand–receptor signaling) serves as a prior. **Technically**, we propose IADOT which learns couplings between snapshots by integrating feature similarity with interaction structure in a multi-dimensional Fused-Gromov Wasserstein objective. Based on these couplings, we then learn continuous-time dynamics of cells by regressing velocity fields with Conditional Flow Matching. **Empirically**, we find on synthetic and diverse single-cell datasets that incorporating directed, typed CCI structure into the OT problem can improve trajectory inference. We also perform in-silico interventions on the LR prior (ablatting specific interactions) to assess the dependence of inferred dynamics on the structural assumption.

2 BACKGROUND

Problem formulation: cell trajectory inference. We consider k population snapshots $\{\mathcal{D}_i\}_{i=1}^k$, where each $\mathcal{D}_i \subset \mathbb{R}^d$ is a set of single-cell states measured at time t_i . The goal is to learn a time-continuous flow $\psi : \mathbb{R}^d \times \mathbb{R}_+ \rightarrow \mathbb{R}^d$ such that $\psi(x, t)$ returns the state obtained by evolving an initial state x to time t . Because scRNA-seq is *destructive*, the same cell cannot be observed at two times, so there is no one-to-one correspondence between cells in \mathcal{D}_i and \mathcal{D}_{i+1} . Classical time-series and ODE-fitting methods that require repeated observations of the same object are thus not directly applicable; trajectory inference must instead recover dynamics from *unaligned snapshots*.

Global alignment of snapshots. Rather than inferring trajectories from neighborhoods within a single snapshot (Haghverdi et al., 2016), recent work aligns *multiple snapshots at the population level* (Schiebinger et al., 2019), treating each snapshot as a probability distribution over cell states. For two timepoints $t_0 < t_1$ with datasets $\mathcal{D}_0 = \{x_i\}_{i=1}^{n_0}$ and $\mathcal{D}_1 = \{y_j\}_{j=1}^{n_1}$, where $x_i, y_j \in \mathbb{R}^d$ are gene-expression vectors, we form the empirical measures $\rho_0 = \sum_{i=1}^{n_0} a_i \delta_{x_i}$ and $\rho_1 = \sum_{j=1}^{n_1} b_j \delta_{y_j}$, with $a \in \Sigma_{n_0}$, $b \in \Sigma_{n_1}$, and $\Sigma_n := \{w \in \mathbb{R}_+^n : \sum_{k=1}^n w_k = 1\}$ (e.g., $a_i = 1/n_0$ for uniform weights). The alignment problem seeks a coupling Γ^* between ρ_0 and ρ_1 that respects the marginals, i.e.,

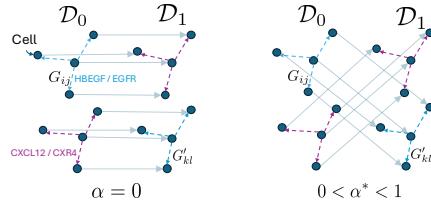


Figure 1: Augmenting feature distances with typed interactions improves alignment between snapshots \mathcal{D}_0 and \mathcal{D}_1 and encourages transport plans consistent with communication patterns.

$$108 \quad \Gamma^* \in \Pi(a, b) := \{ \Gamma \in \mathbb{R}_+^{n_0 \times n_1} \mid \Gamma \mathbf{1}_{n_1} = a, \Gamma^\top \mathbf{1}_{n_0} = b \}, \quad (1)$$

110 where $\mathbf{1}_n$ is the all-ones vector. Without additional structure, any $\Gamma \in \Pi(a, b)$ is admissible and
 111 the problem is underdetermined. Biological priors must therefore rule out implausible matchings.
 112 A widely used prior is the *principle of least action*: cell states change smoothly over time, making
 113 matchings that incur small feature-wise changes more likely. This recasts snapshot alignment as an
 114 Optimal Transport (OT) problem (Villani et al., 2008).

115 **The static Optimal Transport problem.** Optimal Transport (OT) provides a geometric frame-
 116 work for comparing probability distributions, enforcing *the principle of least-action*. In the context
 117 of cell trajectory inference, this principle assumes that the change between snapshots arises from the
 118 smallest rearrangement of cellular states consistent with biology: states evolve smoothly in expres-
 119 sion space. Given a cost matrix $C \in \mathbb{R}_+^{n_0 \times n_1}$, where $C_{ij} = c(x_i, y_j)$ is the cost of transporting a
 120 unit of mass from x_i to y_j , the discrete Kantorovich formulation seeks a coupling

$$121 \quad \Gamma^* \in \arg \min_{\Gamma \in \Pi(a, b)} \langle \Gamma, C \rangle_F \quad (2)$$

123 where $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius dot product. The optimal coupling Γ^* therefore represents the
 124 most efficient mapping from a geometric standpoint, as it is defined based on the gene expression
 125 profiles. **However, by focusing solely on intrinsic state changes, this formulation neglects the**
 126 **extrinsic cell-cell interactions that coordinate population dynamics.**

127 **Incorporating intra-snapshot structure.** Beyond inter-snapshot distances, it is frequent to have
 128 access to structural information in each snapshot. However, the optimization problem in Equa-
 129 tion (2) does not account for it, as it is purely based on inter-snapshot distances. The Gromov–
 130 Wasserstein (GW) problem extends OT to compare two distributions using their *pairwise*
 131 *relational* structure. We assume that this relational structure can be represented by two mat-
 132 rices $G^{(0)} \in \mathbb{R}^{n_0 \times n_0}$ (source) and $G^{(1)} \in \mathbb{R}^{n_1 \times n_1}$ (target). The GW problem seeks a coupling
 133 $\Gamma^* \in \Pi(a, b)$ that minimizes the distortion between the intra-domain structure matrices, $G^{(0)}$ and
 134 $G^{(1)}$. More precisely, the GW objective is the following quadratic program:

$$136 \quad \text{GW}(G^{(0)}, G^{(1)}, a, b) = \min_{\Gamma \in \Pi(a, b)} \sum_{i, k=1}^{n_0} \sum_{j, l=1}^{n_1} L(G_{ik}^{(0)}, G_{jl}^{(1)}) \Gamma_{ij} \Gamma_{kl} \quad (3)$$

138 where L denotes a pairwise distortion function. Finally, it is possible to compare distributions
 139 based on *both* their features and their relational structures, combining the Kantorovich and the GW
 140 formulations. For a given hyperparameter $\alpha \in [0, 1]$, the Fused Gromov-Wasserstein problem is
 141 defined by:

$$143 \quad \text{FGW}_\alpha(G^{(0)}, G^{(1)}, C, a, b) = \min_{\Gamma \in \Pi(a, b)} (1 - \alpha) \langle \Gamma, C \rangle_F + \alpha \sum_{i, k=1}^{n_0} \sum_{j, l=1}^{n_1} L(G_{ik}^{(0)}, G_{jl}^{(1)}) \Gamma_{ij} \Gamma_{kl} \quad (4)$$

145 The parameter α acts as a trade-off, balancing the importance of aligning individual cell features
 146 against preserving the structure between cells (controlled by the Gromov-Wasserstein term): setting
 147 $\alpha = 0$ recovers the Kantorovich problem, while $\alpha = 1$ recovers the GW problem.

3 RELATED WORK

152 **Distributional alignment for trajectory inference.** Classical trajectory-inference tools recon-
 153 struct cellular progressions from neighborhood graphs with pseudotime and branching heuristics
 154 (e.g., Monocle 2, DPT, Slingshot, PAGA) (Qiu et al., 2017; Haghverdi et al., 2016; Street et al.,
 155 2018; Wolf et al., 2019), typically within a *single* snapshot. Optimal transport (OT) (Villani et al.,
 156 2008; Peyré & Cuturi, 2019) provides an alternative that couples *distributions* across timepoints
 157 rather than stitching local paths. Waddington-OT (WOT) extends OT to sequences of time-labeled
 158 snapshots, estimating adjacent-time couplings (Schiebinger et al., 2019). Continuous-time coun-
 159 terparts such as TrajectoryNet learn neural ODE flows constrained by transport to interpolate dis-
 160 tributions over time (Tong et al., 2020). However, these families typically optimize match quality
 161 primarily in expression space, treating each cell as an isolated point and overlooking intercellular
 162 communication. Table 1 contrasts IADOT with other OT-based methods and an extended discussion
 163 is provided in Appendix A.

162 Table 1: **Comparison of trajectory methods.** Legend: ✓ supported, ~ partial, ✗ not supported.
163

164 Method	165 Dynamic	166 Trajectories	167 In-silico Perturbation	168 Structure-Aware	169 scRNA Data Sufficient	170 Reference
PAGA (Scanpy)	✗	✗	✗	✗	✓	(Wolf et al., 2019)
Waddington-OT	✓	~	✗	✗	✓	(Schiebinger et al., 2019)
SCOT	✗	✗	✗	~	✗	(Demetci et al., 2022a)
CellOT	✗	✗	✓	✗	✗	(Bunne et al., 2023a)
OT-CFM	✓	✓	✗	✗	✓	(Tong et al., 2024)
TrajectoryNet	✓	✓	✗	✗	✓	(Tong et al., 2020)
Schrödinger Bridge	✓	✓	✗	✗	✓	(Hong et al., 2025)
scVelo	✓	✗	✗	✗	✓	(Bergen et al., 2020)
IADOT (ours)	✓	✓	✓	✓	✓	—

171 **Structure-aware alignments.** Gromov–Wasserstein (GW) compares samples via their intrinsic
172 geometry, and Fused GW (FGW) optimizes a joint feature+structure objective (Vayer et al., 2020a).
173 In single-cell settings, GW/FGW pipelines typically rely on *undirected* k NN graphs that capture
174 generic topology but lack *communication semantics* (Demetci et al., 2022b; Lange et al., 2024).
175 We instead inject a *directed*, *typed* prior derived from ligand–receptor (LR) expression into an FGW
176 objective. This encourages alignments that preserve signaling context and allows to probe the effect
177 of specific LR interactions on the inferred dynamics. Orthogonal lines of work infer *directionality*
178 from spliced/unspliced counts and propagate it on k NN graphs (Bergen et al., 2020). CellRank
179 further combines velocity with transcriptomic similarity to estimate fate probabilities (La Manno
180 et al., 2018; Bergen et al., 2020; Lange et al., 2022). Spatial OT approaches instead exploit *physical*
181 *proximity* to couple cells and infer possible communication. For example, NicheFlow (Sakalyan
182 et al., 2025) models cells’ microenvironment but it assumes access to spatial transcriptomics data.
183 In general, the structure term in spatial OT remains geometric rather than *typed* signaling (Cang
184 & Nie, 2020). Contrasting these works, IADOT tackles the setting where only scRNA-seq data
185 is available. Meta Flow Matching (Atanackovic et al., 2024) learns an amortized vector field by
186 encoding the initial distribution with a graph neural network; however it assumes access to $N \geq 2$
187 datasets for training, which makes it inapplicable to the setting tackled by IADOT.
188

189 **Inductive biases in flow-based modeling.** Recent advances in flow matching have focused on
190 incorporating specific domain priors. MIOFlow (Huguet et al., 2022) and Metric Flow Matching
191 (Kapusniak et al., 2024) impose *geometric* inductive biases, restricting dynamics to the data mani-
192 fold or a learned Riemannian metric. Other approaches address biological mass conservation: UOT-
193 FM (Eyring et al., 2023) relaxes the exact mass constraint via unbalanced OT to model variable
194 population sizes, while VGF (Wang et al., 2025) explicitly incorporates *cellular growth* rates into
195 the generative flow. In contrast, IADOT integrates an *orthogonal* prior, based on *directed*, *typed*
196 ligand–receptor pairs.
197

4 IADOT: INTERACTION-AWARE OPTIMAL TRANSPORT

199 **Overview.** Our objective is to evaluate whether incorporating a structural prior on cell–cell interac-
200 tions (CCIs) (specifically, a bias toward transport maps that preserve CCI structure across snapshots)
201 can improve trajectory inference. Accordingly, we introduce *Interaction-Aware Dynamic Optimal*
202 *Transport* (IADOT), a framework that integrates gene-expression features and interaction networks
203 into a unified OT objective. Given source and target snapshots \mathcal{D}_0 and \mathcal{D}_1 , IADOT proceeds in
204 two stages. It first computes a *static* cross-snapshot coupling representing a probabilistic assign-
205 ment from source cells to target cells. IADOT enforces two desiderata regarding this coupling: **(D1)**
206 **Feature coherence**—the coupling should reflect smooth cell evolution in expression space; **(D2)**
207 **Communication preservation**—the coupling should capture the persistence of some directed CCI
208 geometry based on ligand/receptor expression. IADOT satisfies these two desiderata by optimizing
209 a *Fused Gromov–Wasserstein* objective balancing feature similarity and CCI preservation, yielding
210 a coupling Γ^* . In the second stage, IADOT fits a *continuous-time* velocity field from interpolants
211 derived from Γ^* using a *Conditional Flow Matching* loss. We can then integrate this velocity field
212 to obtain cell trajectories starting from any given initial state.
213

4.1 INTERACTION-AWARE TRANSPORT VIA MULTI LR-PAIR FGW

214 **Modeling cell–cell interactions from scRNA.** Given a ligand–receptor (LR) catalog $\mathcal{P} =$
215 $\{(l_k, r_k) \mid k \in [K]\}$ of K ligand–receptor pairs and a dataset of n cells, our aim is to construct

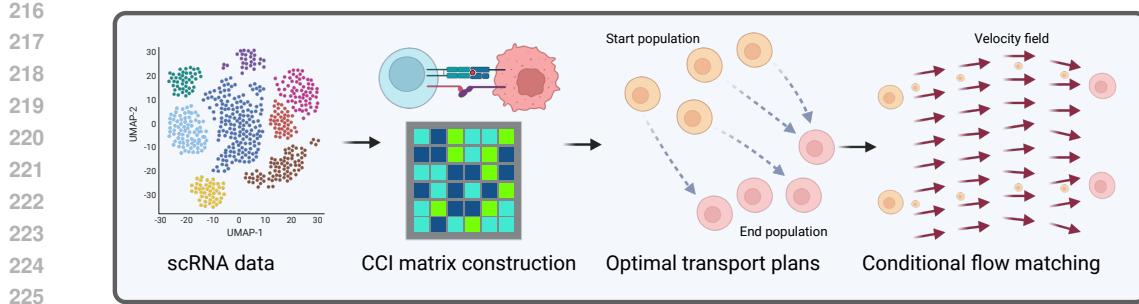


Figure 2: **Overview of IADOT.** From a ligand–receptor catalogue, we build directed, multi LR-pair CCI matrices. A structure-aware optimal transport problem balances feature similarity with interaction structure to produce a cross-snapshot coupling, used to train a time-continuous vector field learned via Conditional Flow Matching to recover cell trajectories.

a directed and nonnegative CCI tensor $G \in \mathbb{R}^{n \times n \times K}$ that summarizes potential signaling from any sender cell i to receiver cell j . Starting from raw expression counts, we first apply library-size normalization to make cells comparable. Rather than a $\log(1+\cdot)$ transform, which can suppress biologically meaningful high-expression events, we keep normalized counts and map each gene to $[0, 1]$ using a Hill-saturation function. For gene g , we then define $s_{cg} = x_{cg}^{h_g} / (x_{cg}^{h_g} + K_g^{h_g})$, with robust scale K_g (e.g., the $q = 0.9$ quantile of nonzero values in $\{x_{cg} \mid c \in [n]\}$ where x_{cg} is the normalized expression of gene g for cell c) and exponent h_g .

This gives bounded activations where near-saturating expression contributes strongly. For an LR pair $p_k = (l_k, r_k)$ and cells i (sender) and j (receiver), we score the interaction as $q_{i \rightarrow j}^{(p_k)} = s_{i l_k} s_{j r_k}$, capturing the intuitive requirement that ligand availability and receptor readiness must co-occur. We then define the value of G at (i, j, k) as $G_{ijk} = q_{i \rightarrow j}^{(p_k)}$. The CCI tensor G then serves as the directed structure we aim to preserve during cross-snapshot alignment.

Remark. For denoising purposes, cells can optionally be aggregated into metacells (e.g., by clustering in a low-dimensional embedding) before constructing the CCI tensors. We empirically evaluate this variant in Section 5.5.

Interaction-aware transport via multi LR-pair FGW. Given the two snapshots $\mathcal{D}_0 = \{x_i\}_{i=1}^{n_0}$ and $\mathcal{D}_1 = \{y_j\}_{j=1}^{n_1}$, we define a feature cost matrix $C \in \mathbb{R}_{\geq 0}^{n_0 \times n_1}$, such that for all i, j we have $C_{ij} = c(x_i, y_j)$, where c is typically the squared Euclidean distance. From the CCI construction described above, we obtain directed, nonnegative tensors $(G^{(0)}, G^{(1)})$ corresponding to the source and target snapshots respectively. Our objective is to find a coupling $\Gamma \in \mathbb{R}_{\geq 0}^{n_0 \times n_1}$ that aligns cells while respecting the CCI structures. To jointly account for feature distances and multi-LR pair CCIs, we optimize a *Fused Gromov–Wasserstein* objective that balances a feature term $\mathcal{F}(\Gamma)$ and a structure-preservation term $\mathcal{S}(\Gamma)$ defined with a similarity measure φ :

$$\min_{\Gamma \in \Pi(a, b)} (1 - \alpha) \underbrace{\langle \Gamma, C \rangle}_{\mathcal{F}(\Gamma)} + \alpha \underbrace{\sum_{i, k=1}^{n_0} \sum_{j, \ell=1}^{n_1} \varphi(G_{ik}^{(0)}, G_{j\ell}^{(1)}) \Gamma_{ij} \Gamma_{k\ell}}_{\mathcal{S}(\Gamma)}. \quad (5)$$

The structure term $\mathcal{S}(\Gamma)$ favours couplings that preserve the CCI patterns encoded in $G^{(0)}$ and $G^{(1)}$. Unlike the classical FGW setting (Vayer et al., 2020b), IADOT handles *multi-typed* interactions: each entry G_{ij} is a vector in \mathbb{R}^K rather than a scalar, allowing multiple LR pairs per cell–cell relation. We compare these interaction vectors with a similarity φ . By default we use the squared Euclidean norm $\varphi(u, v) = \|u - v\|^2$. Furthermore, Equation (5) is a non-linear and non-convex problem because of the structure term $\mathcal{S}(\Gamma)$. To efficiently solve this non-convex objective, we introduce a customized conditional-gradient solver adapted from (Braun et al., 2022). This tailored optimization routine, detailed in Section D.4, is specifically designed to handle the structural constraints of the interaction-aware coupling.

270 **Scale normalization.** To balance the feature and structure terms in Equation (5), we *normalize by*
 271 *endpoints*. Concretely, we first solve the feature-only problem ($\alpha = 0$) and the structure-only prob-
 272 lem ($\alpha = 1$), obtaining $\Gamma_{\alpha=0}^*$ and $\Gamma_{\alpha=1}^*$. We then rescale the feature cost and the CCI tensors using
 273 the corresponding objective values at these two optima, so that their magnitudes are comparable (see
 274 Section D.5 for more details).

275
 276 **Unbalanced setting.** The formalism above assumes that the coupling $\Gamma \in \Pi(a, b)$ preserves the
 277 marginals a and b . In practice, some developmental and perturbation settings exhibit net proliferation
 278 or apoptosis between snapshots. To account for that, IADOT can be extended to an unbalanced
 279 setting by relaxing the hard constraint $\Gamma \in \Pi(a, b)$ with divergence penalties on the row and column
 280 sums of Γ (more details are given in Section D.11). We provide results with this extension in
 281 Section 5.3 and Section E.4.

283 4.2 LEARNING CONTINUOUS DYNAMICS VIA CONDITIONAL FLOW MATCHING

284 **Objective.** The goal of IADOT is to learn a time-dependent velocity field that transports the source
 285 dataset \mathcal{D}_0 to the target dataset \mathcal{D}_1 , and can be integrated up to any time $t > 0$. We leverage the
 286 optimal coupling obtained from Equation (5) to align the two snapshots and convert this *static* cor-
 287 respondence into a *time-dependent* velocity field using Conditional Flow Matching (CFM) (Tong
 288 et al., 2024; Lipman et al., 2022). Concretely, we first construct a coupling-induced probability path
 289 $\{\rho_t\}_{t \in [0,1]}$ and then fit a velocity field to generate this probability path.

290 **Probability path.** Let ρ_0 and ρ_1 denote the empirical distributions defined by \mathcal{D}_0 and \mathcal{D}_1 , respec-
 291 tively. Let $\Gamma^* \in \mathbb{R}_+^{n_0 \times n_1}$ be the optimal coupling from Equation (5), with normalization constant
 292 $M = \sum_{i,j} \Gamma_{ij}^*$. We define a joint distribution Π on $\mathcal{D}_0 \times \mathcal{D}_1$ by $\Pi = \sum_{i=1}^{n_0} \sum_{j=1}^{n_1} \frac{\Gamma_{ij}^*}{M} \delta_{(x_i, y_j)}$,
 293 where $\delta_{(x_i, y_j)}$ denotes the Dirac measure at (x_i, y_j) . Therefore, the marginals of Π are ρ_0 and ρ_1 .
 294 For $t \in [0, 1]$, we then consider the affine interpolation $Z_t = (1-t)X + tY$, with $(X, Y) \sim \Pi$, and
 295 let $\rho_t = \mathcal{L}(Z_t)$ be the distribution of Z_t , yielding a probability path $\{\rho_t\}_{t \in [0,1]}$. By construction, ρ_0
 296 and ρ_1 are the endpoints of this path.

297
 300 **Learning the vector field with CFM.** Given the coupling-induced path $\{\rho_t\}_{t \in [0,1]}$, we learn a
 301 time-dependent velocity field $v_\theta : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ that generates it. For $(X, Y) \sim \Pi$ and
 302 $Z_t = (1-t)X + tY$, the interpolation implies a constant drift across time $u_t(Z_t | X, Y) = Y - X$,
 303 conditioned on (X, Y) .

304 We train v_θ by regressing to this drift along the path, yielding the following CFM objective:

$$307 \mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{\substack{(X, Y) \sim \Pi \\ t \sim \text{Unif}[0,1]}} \left[\| v_\theta(Z_t, t) - u_t(Z_t | X, Y) \|_2^2 \right] \quad (6)$$

$$310 \mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{\substack{(X, Y) \sim \Pi \\ t \sim \text{Unif}[0,1]}} \left[\| v_\theta(Z_t, t) - (Y - X) \|_2^2 \right]. \quad (7)$$

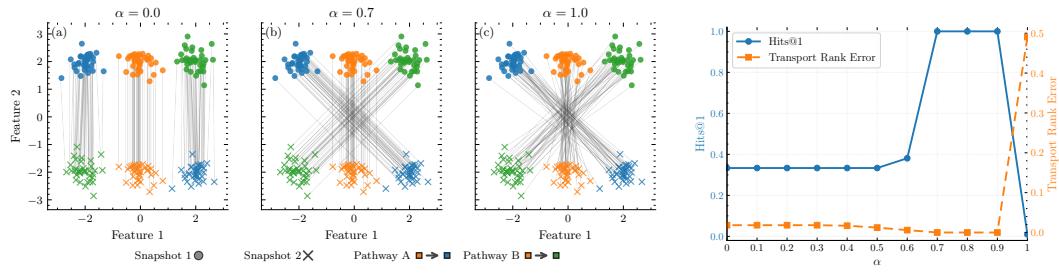
312
 313 Thus, converting the coupling to a velocity field reduces to supervised regression. As shown in
 314 (Lipman et al., 2024), the minimizer of this loss generates the probability path $\{\rho_t\}_{t \in [0,1]}$. After
 315 training, we can then sample trajectories starting from any point $x \in \mathbb{R}^d$ at time 0 by integrating the
 316 ODE $\dot{z}(t) = v_\theta(z(t), t)$ from 0 to $t > 0$, with the initial condition $z(0) = x$.

317 **Extensions.** Because our CCI prior is defined independently of how the velocity field is parameter-
 318 ized and regressed, it acts as an orthogonal component to the underlying flow-matching objective. In
 319 practice, this means that IADOT can be used as a plug-and-play prior on top of more advanced flow-
 320 matching methods without requiring any change to their architectures or training procedures. For
 321 example, we can directly combine IADOT with Metric Flow Matching (Kapusniak et al., 2024) by
 322 simply replacing the base CFM objective with its metric variant, while keeping the coupling derived
 323 from IADOT. We detail this in Section D.12 and we report results with this variant in Section 5.3
 and Section E.4.

324 5 EXPERIMENTS

326 We evaluate whether incorporating the cell-cell interaction (CCI) structure improves cross-snapshot
 327 alignment and continuous-time trajectory inference over feature-only baselines. In Section 5.1, we
 328 present a controlled synthetic study, showing that the solution to the structure-aware OT problem
 329 (Equation (5)) can exactly recover the ground-truth transport map. In Sections 5.2 and 5.3, we
 330 benchmark IADOT on three scRNA-seq datasets spanning diverse tissues and observe consistent
 331 gains over baselines in interpolation metrics, when incorporating the CCIs. We then provide biolog-
 332 ical insights by performing targeted edits to the ligand-receptor catalog, and quantify the resulting
 333 shifts in inferred dynamics in Section 5.4. Finally, we conduct a sensitivity analysis over the CCI
 334 construction choices in Section 5.5 and discuss potential failure modes of IADOT in Section 5.6.

335 5.1 SYNTHETIC SETUP



347 **Figure 3: Structure-aware coupling recovers the ground-truth transport map.** We show repre-
 348 sentative couplings (left) and matching metrics (right). Feature-only OT ($\alpha=0$) ignores structure and
 349 misaligns clusters, structure-only ($\alpha=1$) distorts geometry within the interaction types. A balanced
 350 trade-off ($\alpha \approx 0.7$) recovers the intended one-to-one mapping.

351 **Setup.** We consider two 2D snapshots, each composed of three clusters. The second snapshot is
 352 obtained by translating each cluster by a distinct vector, inducing a known one-to-one ground-truth
 353 transport. We define an interaction structure with two types : the middle cluster points to the left
 354 (*Type 1*) and to the right (*Type 2*), mirrored in the target snapshot (see Section C.1 for more details).
 355 We then obtain a coupling for each $\alpha \in \{0, 0.1, \dots, 1\}$ by solving the FGW problem defined in
 356 Equation (5) with the ground-truth interaction structures.

358 **Results.** Representative couplings across α are shown in Figure 3 (left). With $\alpha = 0$ (feature-
 359 only), the interaction structure is ignored and clusters are misaligned; with $\alpha = 1$ (structure-only),
 360 interaction types are satisfied but geometry is distorted. An intermediate setting ($\alpha \approx 0.7$) preserves
 361 the directed relations while maintaining within-interaction geometry. We quantify these observations
 362 by computing *Hits@1*, the fraction of source samples whose top-weighted target equals the ground-
 363 truth match, and *Transport Rank Error* (*TRE*), the average fraction of targets ranked above the
 364 ground-truth match. Figure 3 (right) shows that *Hits@1* peaks and *TRE* is minimized at mid-range
 365 α , indicating that a balanced mix of features and structure gives the most faithful transport map. **We**
 366 **refer to Section F.1 for a theoretical analysis of this synthetic setup.**

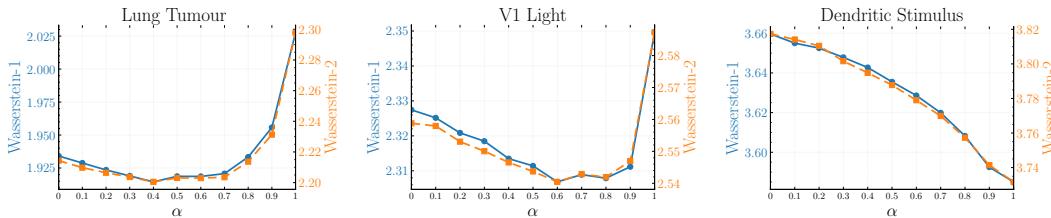
367 5.2 CROSS-SNAPSHOT INTERPOLATION FROM STRUCTURE-AWARE COUPLINGS

369 **Datasets.** We evaluate IADOT on *real-world datasets* whose characteristics are summarized in Ta-
 370 ble 6. We selected these datasets because their temporal coverage provides a favorable window in
 371 which ligand-receptor (LR) interactions are expected to remain approximately persistent. Following
 372 standard preprocessing, we project gene-expression profiles onto the top $d = 20$ principal compo-
 373 nents (Section D.2) and standardize them as in (Tong et al., 2024). Additional details on dataset
 374 collection are provided in Section C, and results on further datasets are in Section E.

375 **Setup.** We build CCI tensors by selecting dataset-specific ligand-receptor pairs via an automated
 376 procedure that accounts for stability of expression levels across snapshots (cf. Section D.6 for more
 377 details). We then assess the couplings produced by IADOT in an interpolation setup. Given three
 time points $t_0 < t_1 < t_2$, we hold out the snapshot at t_1 . Using only t_0 and t_2 , and for a chosen

378 LR catalog \mathcal{P} and hyperparameter $\alpha \in \{0, 0.1, \dots, 1.0\}$, we obtain a coupling $\Gamma(\alpha, \mathcal{P})$ by solving
 379 the OT problem defined in Equation (5). We define the marginal at t_1 by affine interpolation and
 380 denote it by $\rho_{t_1}(\alpha, \mathcal{P})$. For each α and \mathcal{P} , we compare $\rho_{t_1}(\alpha, \mathcal{P})$ with the empirical distribution ρ_{t_1}
 381 observed at t_1 , computing the Wasserstein-1 and Wasserstein-2 distances $W_1(\rho_{t_1}(\alpha, \mathcal{P}), \rho_{t_1})$ and
 382 $W_2(\rho_{t_1}(\alpha, \mathcal{P}), \rho_{t_1})$.

383 **Results.** We report these metrics in Figure 4. Across datasets, incorporating CCI structure improves
 384 alignment, with optimal performance at a dataset-specific $\alpha^* > 0$. We observe two regimes: a U-
 385 shaped curve with $0 < \alpha^* < 1$, indicating that combining CCI with feature-only OT is best, and an
 386 almost monotonic decrease with a minimum at $\alpha^* = 1$ for the *Dendritic Stimulus* dataset.
 387



396 **Figure 4: Interpolation error.** We plot the W_1 and W_2 distances between the interpolated and
 397 empirical t_1 snapshots as α varies. Optimal performance occurs at dataset-specific $\alpha^* > 0$.
 398

400 5.3 CROSS-SNAPSHOT TRAJECTORY INFERENCE VIA FLOW MATCHING

402 **Setup.** Having shown that incorporating structure yields better couplings for these datasets, we
 403 now verify whether it also improves continuous-time dynamics learnt with `IADOT`. Starting from an
 404 optimal coupling Γ , we fit a time-conditional vector field v_θ using the conditional Flow match-
 405 ing loss (Equation (6)). We integrate v_θ to transport cells observed at t_0 to the held-out time
 406 t_1 , and compare the transported distribution to the empirical snapshot at t_1 using Wasserstein-1
 407 and -2 distances. We repeat this for $\alpha \in \{0, 0.5, 1\}$. As baselines, we compare against neural
 408 ODE-based methods (`TrajectoryNet` (Tong et al., 2020), `MIOflow` (Huguet et al., 2022)),
 409 diffusion-based Schrödinger-bridge methods (`Diffusion Schrödinger Bridges` (DSB)
 410 (De Bortoli et al., 2021), `SF2M` (Tong et al., 2023)), flow-matching methods (`MFM` (Kapusniak et al.,
 411 2024), `UOT-FM` (Eyring et al., 2023), `VGFM` (Wang et al., 2025)), and `MOSCOT` (Klein et al., 2025b).
 412 Since `IADOT` provides a plug-and-play CCI prior that is complementary to the underlying OT or
 413 flow-matching objective, we also report “`IADOT+`” variants of several baselines (e.g. `IADOT+MFM`,
 414 `IADOT+UOT-FM`, `IADOT+SF2M`). We refer to Section D.12, Section D.11, and Section D.13 for
 415 details about these variants.

416 **Results.** Table 2 reports W_1 and W_2 at the held-out time t_1 (lower is better). Results indicate that
 417 the procedure used to fit the velocity field affects performance: conditional flow matching yields
 418 consistently lower errors than all the baselines across datasets. Second, within `IADOT`, structure
 419 helps: settings with $\alpha > 0$ outperform the feature-only case ($\alpha = 0$), with the best results at
 420 $\alpha \in \{0.5, 1\}$. These findings align with Section 5.2, indicating that CCI structure benefits both static
 421 alignment *and* the learned continuous-time dynamics. Furthermore, we see that plugging `IADOT`
 422 with other priors leads to improved performance, showing the wide benefits of incorporating CCI
 423 information.

424 5.4 PROBING TRAJECTORY SENSITIVITY TO LIGAND-RECEPTOR CATALOG EDITS

425 **Setup.** Having demonstrated that interaction structure guides trajectory inference, we now leverage
 426 `IADOT` to simulate intercellular perturbations. We focus on the *Lung Tumor* dataset and construct
 427 alternative ligand-receptor catalogs in which specific signaling pathways are ablated. From these
 428 modified catalogs, we recompute the CCI tensors and resolve the OT problem (Equation (5)). Crucially,
 429 this intervention modifies only the interaction prior governing the coupling, while holding the
 430 initial gene expression at $t = 0$ fixed. This design mimics a *pharmacological blockade*, where
 431 external signaling is inhibited before cells transcriptionally adapt (Lee et al., 2016), providing a
 432 mechanism-specific counterfactual that predicts how the loss of communication redirects the popu-

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

Table 2: Interpolation error for continuous time dynamics (lower is better). IADOT with varying structure weight α vs. baselines across the three datasets. We report mean \pm std over 5 runs.

Method	α	V1 Light		Dendritic Stimulus		Lung tumor	
		W_1	W_2	W_1	W_2	W_1	W_2
TrajectoryNet	—	3.022 (0.061)	3.338 (0.056)	4.410 (0.102)	4.607 (0.107)	2.712 (0.090)	3.056 (0.099)
DSB	—	3.819 (0.152)	3.875 (0.143)	4.099 (0.155)	4.249 (0.153)	3.700 (0.116)	3.967 (0.102)
OT-CFM	—	2.392 (0.005)	2.625 (0.007)	3.696 (0.007)	3.857 (0.009)	1.993 (0.004)	2.275 (0.005)
OT-MFM	—	2.401 (0.003)	2.636 (0.003)	3.714 (0.008)	3.880 (0.009)	1.984 (0.004)	2.285 (0.004)
UOT-FM	—	2.411 (0.005)	2.649 (0.006)	3.701 (0.006)	3.867 (0.007)	1.998 (0.004)	2.348 (0.004)
SF2M	—	3.254 (0.192)	3.368 (0.182)	4.333 (0.279)	4.436 (0.282)	3.826 (0.265)	3.974 (0.308)
VGFM	—	6.446 (0.114)	6.745 (0.102)	7.087 (0.022)	7.261 (0.026)	2.175 (0.017)	2.478 (0.019)
MIOFlow	—	6.360 (0.010)	6.655 (0.009)	6.970 (0.043)	7.159 (0.034)	2.001 (0.003)	2.316 (0.009)
Moscot	—	6.242 (0.000)	6.545 (0.000)	7.115 (0.000)	7.331 (0.000)	2.000 (0.000)	2.335 (0.000)
IADOT+SF2M	0.5	3.199 (0.117)	3.315 (0.110)	4.303 (0.213)	4.397 (0.205)	3.809 (0.302)	3.968 (0.374)
	1	3.226 (0.075)	3.339 (0.073)	4.289 (0.110)	4.387 (0.108)	3.638 (0.308)	3.739 (0.335)
IADOT+MFM	0.5	2.393 (0.007)	2.631 (0.008)	3.679 (0.007)	3.838 (0.009)	1.978 (0.004)	2.277 (0.003)
	1	2.363 (0.002)	2.606 (0.002)	3.668 (0.010)	3.824 (0.011)	2.013 (0.003)	2.304 (0.003)
IADOT+UOT-FM	0.5	2.377 (0.004)	2.619 (0.005)	3.688 (0.012)	3.854 (0.012)	1.971 (0.005)	2.322 (0.005)
	1	2.360 (0.002)	2.605 (0.001)	3.624 (0.004)	3.780 (0.002)	1.993 (0.004)	2.335 (0.005)
IADOT+CFM	0.5	2.381 (0.004)	2.618 (0.003)	3.679 (0.009)	3.835 (0.010)	1.989 (0.004)	2.272 (0.005)
	1	2.362 (0.003)	2.601 (0.005)	3.639 (0.021)	3.788 (0.021)	2.057 (0.005)	2.329 (0.005)

lation’s downstream trajectory. We quantify these shifts relative to the unperturbed baseline using the 20 Hallmarks of Cancer gene sets (see Section D.9.1) over a 24h interpolation window.

Results. Figure 5 shows the relative decrease in tumour-associated progression scores under different catalog edits. Attenuating signaling through EGFR, ALK, or MET produces measurable reductions (up to 15.5%), indicating that the inferred trajectories are sensitive to these pathways. This aligns with their established therapeutic relevance in non-small cell lung cancer, where EGFR inhibitors (e.g., gefitinib, osimertinib), ALK inhibitors (e.g., crizotinib, alectinib), and MET inhibitors (e.g., capmatinib, tepotinib) are used clinically (Domvri et al., 2013). By contrast, edits to unrelated cardio-renal pathways (RAAS, vasopressin, natriuretic peptides) yield negligible changes, suggesting that IADOT responds specifically to biologically relevant ligand–receptor structure rather than arbitrary perturbations.

5.5 SENSITIVITY TO CCI CONSTRUCTION CHOICES

Setup. Motivated by the previous observation that changing the CCI structure affects the learnt dynamics, we now conduct a sensitivity analysis via three controlled perturbations of the CCI construction process: *Random LR catalog*—replace the curated ligand–receptor (LR) catalog with a random subset of the same size as the initial catalog; *Shuffling*—randomly permute all entries of the CCI tensors, destroying coherent structure; *Metacells*—aggregate cells into metacells before constructing CCIs and then lift interactions back to the cell level (see Section D.3 for details), thereby smoothing the signal. We evaluate all variants under the interpolation protocol of Section 5.2 with $\alpha = 1$ (structure-only OT) to isolate structural effects.

Results. Table 3 summarizes the results. Shuffling the CCI leads to a performance drop, confirming that the *structural organization* of LR interactions drives the gains. Using a random LR catalog also degrades the interpolation, highlighting the importance of *ligand–receptor specificity*. The CCI constructed with metacells yields intermediate results, as it improves results on two of the datasets. This can be attributed to its smoothing role, especially useful against dropout effect. However, it is

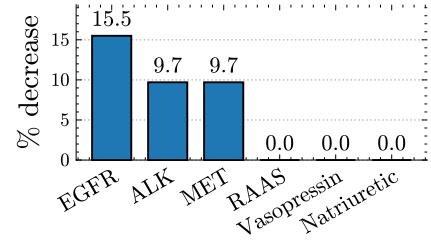


Figure 5: In silico interventions. Edits to the LR catalog (i.e. removing EGFR/ALK/MET interactions in the CCIs) reduce a Hallmark-based tumour progression proxy.

Table 3: Sensitivity analysis on the CCI construction.

Method	Tumor		Dendritic		Light	
	W_1	W_2	W_1	W_2	W_1	W_2
Shuffle	2.181	2.378	3.637	3.760	2.434	2.644
Random LR	2.186	2.408	3.587	3.722	2.441	2.646
Metacell	2.054	2.345	3.575	3.722	2.327	2.564
IADOT	2.028	2.298	3.585	3.732	2.350	2.587

486 not always optimal, as oversmoothing can bias the CCI and degrade performance, an observation
 487 consistent with previous spatiotemporal analyses (Klein et al., 2025a).
 488

489 5.6 DOES STRUCTURE ALWAYS IMPROVE CROSS-SNAPSHOT ALIGNMENT?

490
491 Setup. Our OT formulation penalizes couplings that do not preserve the CCI structure between
 492 two snapshots. Therefore, it assumes that this structure is at least approximately persistent across
 493 snapshots. When the system undergoes rapid and large-scale remodeling, this assumption can fail
 494 and the induced structure may no longer be informative. We illustrate this with a developing mouse
 495 embryo dataset (Moon et al., 2019), where tissue composition, size, and function change quickly
 496 during development (Qiu et al., 2024). Furthermore, the time interval between consecutive snapshots
 497 is substantial (6 days). As such, we expect the CCI structure at one stage to be poorly related to the
 498 next.

499 **Results.** We report the interpolation results for this
 500 dataset in Figure 6. In this setting, IADOT offers
 501 no additional gains over feature-only OT ($\alpha = 0$),
 502 confirming that the CCI structure is not transferable
 503 across these days-long developmental intervals and
 504 thus becomes uninformative. This leads to an essen-
 505 tially flat curve with respect to α , with worse de-
 506 gradations at larger α . It also yields a practical guide-
 507 line: when cross-snapshot interaction geometry does
 508 not persist over time, the structural term should be
 509 downweighted or omitted.

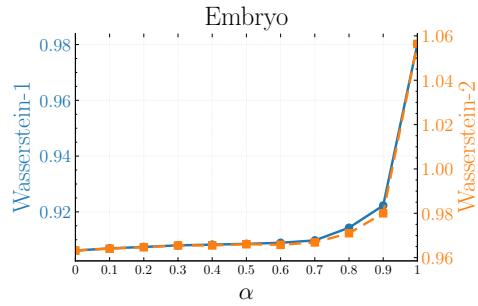
512 6 DISCUSSION

514 **Identifiability with biological structure.** Standard optimal transport relies solely on feature sim-
 515 ilarity to align populations, a formulation often insufficient for the high-dimensional, sparse, and
 516 non-linear nature of single-cell dynamics. In this work, we introduced IADOT, a framework that ad-
 517 dresses this fundamental identifiability problem by injecting a biologically grounded inductive bias:
 518 the persistence of ligand-receptor communication. By formulating alignment as a multi-channel
 519 Fused Gromov-Wasserstein problem, IADOT extracts directed signaling topologies and recovers
 520 trajectories that are not only geometrically smooth but also preserve CCIs. Because structure is an
 521 explicit, editable prior, IADOT enables principled counterfactuals: we can quantify resulting shifts
 522 in inferred trajectories when perturbing pathway-specific LR libraries in the CCI construction.

523 **A modular prior.** A key strength of IADOT is its modularity. Rather than being tied to a spe-
 524 cific dynamics model, it serves as a flexible "plug-and-play" cost function. Our extensive exper-
 525 iments demonstrate that the IADOT coupling consistently improves performance across diverse
 526 paradigms, including deterministic transport (CFM), geometric interpolation (Metric Flow Match-
 527 ing), stochastic bridges (SF2M), and unbalanced mass transport (UOT). This universality suggests
 528 that interaction-aware priors are orthogonal and complementary to recent advances in generative
 529 modeling, offering a generic recipe to refine any OT-based trajectory inference method. Finally, our
 530 ablation study shows that *biologically meaningful* interaction matrices are important and drive the
 531 observed performance gains.

532 **Limitations.** We stress that IADOT is not a silver bullet. Our approach assumes that key interaction
 533 structure is at least partly conserved between adjacent snapshots. In rapidly remodeling systems
 534 (e.g. embryo development), we have shown that incorporating the structure does not necessarily
 535 yield benefits. Furthermore, we evaluated IADOT on real scRNA-seq datasets from prior studies,
 536 but broader *scalability* to atlas level datasets is an interesting avenue for future work.

537 **Broader impact.** IADOT offers a simple recipe to inject typed interaction priors to disambiguate
 538 alignments. Beyond biology, IADOT offers a principled path to modeling dynamics in systems of
 539 interacting entities including financial markets, social networks and multi agent environments where
 structure aware couplings can improve alignments.



510
 511
 512 **Figure 6: Embryo dataset.** Incorporating
 513 CCIs does not improve performance over the
 514 feature-only baseline ($\alpha = 0$).
 515

540 REPRODUCIBILITY STATEMENT
541

542 We specify the OT objective in Section 4.1, and detail the multi-channel FGW solver as well as the
543 normalization used to balance feature and structure terms in Sections D.4 and D.5. Our continuous-
544 time dynamics and Conditional Flow Matching objective are described in Section 4.2, with model
545 architectures and all training hyperparameters listed in Section D.7. Datasets, sampling timepoints,
546 and sizes are summarized in Section C, and the end-to-end preprocessing pipeline is documented
547 in Section D.2. The construction of ligand–receptor catalogs is detailed in Section D.6. Baselines
548 use authors’ implementations with exact settings listed in Section D.8. Software versions and key
549 libraries are reported in Table 7. Code will be released upon acceptance.

550
551 REFERENCES
552

553 Asif Adil, Vijay Kumar, Arif Tasleem Jan, and Mohammed Asger. Single-cell transcriptomics:
554 Current methods and challenges in data acquisition and analysis. *Frontiers in Neuroscience*, 15,
555 April 2021. ISSN 1662-453X. doi: 10.3389/fnins.2021.591122. URL <http://dx.doi.org/10.3389/fnins.2021.591122>.

556 Prabhu S. Arunachalam, Madeleine K. D. Scott, Thomas Hagan, Chunfeng Li, Yupeng Feng, Flo-
557 rian Wimmers, Lilit Grigoryan, Meera Trisal, Venkata Viswanadh Edara, Lilin Lai, Sarah Es-
558 ther Chang, Allan Feng, Shaurya Dhingra, Mihir Shah, Alexandra S. Lee, Sharon Chinthrajah,
559 Sayantani B. Sindher, Vamsee Mallajosyula, Fei Gao, Natalia Sigal, Sangeeta Kowli, Sheena
560 Gupta, Kathryn Pellegrini, Gregory Tharp, Sofia Maysel-Auslender, Sydney Hamilton, Hadj
561 Aoued, Kevin Hrusovsky, Mark Roskey, Steven E. Bosinger, Holden T. Maecker, Scott D.
562 Boyd, Mark M. Davis, Paul J. Utz, Mehul S. Suthar, Purvesh Khatri, Kari C. Nadeau, and
563 Bali Pulendran. Systems vaccinology of the bnt162b2 mrna vaccine in humans. *Nature*, 596
564 (7872):410–416, July 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03791-x. URL
565 <http://dx.doi.org/10.1038/s41586-021-03791-x>.

566 Georgios Arvanitidis, Lars K Hansen, and Søren Hauberg. A locally adaptive normal distribution.
567 *Advances in Neural Information Processing Systems*, 29, 2016.

568 Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio,
569 Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the
570 wasserstein manifold. *arXiv preprint arXiv:2408.14608*, 2024.

571 Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-
572 kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

573 Volker Bergen, Marius Lange, Stefan Peidli, F. Alexander Wolf, and Fabian J. Theis. Generalizing
574 rna velocity to transient cell states through dynamical modeling. *Nature Biotechnology*, 38(12):
575 1408–1414, 2020. doi: 10.1038/s41587-020-0591-3.

576 Gábor Braun, Alejandro Carderera, Cyrille W Combettes, Hamed Hassani, Amin Karbasi,
577 Aryan Mokhtari, and Sebastian Pokutta. Conditional gradient methods. *arXiv preprint*
578 *arXiv:2211.14103*, 2022.

579 Robin Browaeys, Wouter Saelens, and Yvan Saeys. Nichenet: modeling intercellular communica-
580 tion by linking ligands to target genes. *Nature Methods*, 17(2):159–162, December 2019. ISSN
581 1548-7105. doi: 10.1038/s41592-019-0667-5. URL <http://dx.doi.org/10.1038/s41592-019-0667-5>.

582 Charlotte Bunne, Stefan G. Stark, Gabriele Gut, Jacobo Sarabia del Castillo, Mitch Levesque,
583 Kjøng-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-
584 cell perturbation responses using neural optimal transport. *Nature Methods*, 20(11):1759–1768,
585 September 2023a. ISSN 1548-7105. doi: 10.1038/s41592-023-01969-x. URL <http://dx.doi.org/10.1038/s41592-023-01969-x>.

586 Charlotte Bunne, Stefan G Stark, Gabriele Gut, Jacobo Sarabia Del Castillo, Mitch Levesque,
587 Kjøng-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-
588 cell perturbation responses using neural optimal transport. *Nature methods*, 20(11):1759–1768,
589 2023b.

594 Charlotte Bunne, Geoffrey Schiebinger, Andreas Krause, Aviv Regev, and Marco Cuturi. Optimal
 595 transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1), August
 596 2024. ISSN 2662-8449. doi: 10.1038/s43586-024-00334-2. URL <http://dx.doi.org/10.1038/s43586-024-00334-2>.

598 J. Gray Camp, Farhath Badsha, Marta Florio, Sabina Kanton, Tobias Gerber, Michaela Wilsch-
 599 Bräuninger, Eric Lewitus, Alex Sykes, Wulf Hevers, Madeline Lancaster, Juergen A. Knoblich,
 600 Robert Lachmann, Svante Pääbo, Wieland B. Huttner, and Barbara Treutlein. Human cerebral
 601 organoids recapitulate gene expression programs of fetal neocortex development. *Proceedings of
 602 the National Academy of Sciences*, 112(51):15672–15677, December 2015. ISSN 1091-6490. doi:
 603 10.1073/pnas.1520760112. URL <http://dx.doi.org/10.1073/pnas.1520760112>.

604 Zixuan Cang and Qing Nie. Inferring spatial and signaling relationships between cells from
 605 single-cell transcriptomic data. *Nature Communications*, 11(1):2084, 2020. doi: 10.1038/
 606 s41467-020-15968-5.

608 Zixuan Cang, Yanxiang Zhao, Axel A. Almet, Adam Stabell, Raul Ramos, Maksim V. Plikus,
 609 Scott X. Atwood, and Qing Nie. Screening cell-cell communication in spatial transcriptomics
 610 via collective optimal transport. *Nature Methods*, 20(2):218–228, January 2023. ISSN
 611 1548-7105. doi: 10.1038/s41592-022-01728-4. URL <http://dx.doi.org/10.1038/s41592-022-01728-4>.

613 Junha Cha and Insuk Lee. Single-cell network biology for resolving cellular heterogeneity in human
 614 diseases. *Experimental & Molecular Medicine*, 52(11):1798–1808, November 2020. ISSN
 615 2092-6413. doi: 10.1038/s12276-020-00528-0. URL <http://dx.doi.org/10.1038/s12276-020-00528-0>.

618 Gan Chen, Chao Ren, Yao Xiao, Yujing Wang, Renqi Yao, Quan Wang, Guoxing You, Mingzi Lu,
 619 Shaoduo Yan, Xiaoyong Zhang, Jun Zhang, Yongming Yao, and Hong Zhou. Time-resolved
 620 single-cell transcriptomics reveals the landscape and dynamics of hepatic cells in sepsis-induced
 621 acute liver dysfunction. *JHEP Reports*, 5(6):100718, June 2023. ISSN 2589-5559. doi: 10.1016/j.jhepr.2023.100718. URL <http://dx.doi.org/10.1016/j.jhepr.2023.100718>.

623 Zhanlin Chen, William C. King, Aheyon Hwang, Mark Gerstein, and Jing Zhang. Deepvelo :
 624 Single-cell transcriptomic deep velocity field learning with neural ordinary differential equations.
 625 *Science Advances*, 8(48), December 2022. ISSN 2375-2548. doi: 10.1126/sciadv.abq3745. URL
 626 <http://dx.doi.org/10.1126/sciadv.abq3745>.

627 Marco Cuturi, Michal Klein, and Pierre Ablin. Monge, Bregman and occam: Interpretable optimal
 628 transport in high-dimensions with feature-sparse maps. In Andreas Krause, Emma Brunskill,
 629 Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of
 630 the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine
 631 Learning Research*, pp. 6671–6682. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/cuturi23a.html>.

633 Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger
 634 bridge with applications to score-based generative modeling. *Advances in neural information
 635 processing systems*, 34:17695–17709, 2021.

637 Pinar Demetci, Rebecca Santorella, Björn Sandstede, William Stafford Noble, and Ritambhara
 638 Singh. Scot: Single-cell multi-omics alignment with optimal transport. *Journal of Computational
 639 Biology*, 29(1):3–18, January 2022a. ISSN 1557-8666. doi: 10.1089/cmb.2021.0446. URL
 640 <http://dx.doi.org/10.1089/cmb.2021.0446>.

641 Pinar Demetci, Rebecca Santorella, Björn Sandstede, William Stafford Noble, and Ritambhara
 642 Singh. Scot: Single-cell multi-omics alignment with optimal transport. *Journal of Computational
 643 Biology*, 29(1):3–18, 2022b. doi: 10.1089/cmb.2021.0446.

645 Daniel Dimitrov, Philipp Sven Lars Schäfer, Elias Farr, Pablo Rodriguez-Mier, Sebastian Loben-
 646 tanzer, Pau Badia-i Mompel, Aurelien Dugourd, Jovan Tanevski, Ricardo Omar Ramirez Flo-
 647 res, and Julio Saez-Rodriguez. Liana+ provides an all-in-one framework for cell-cell com-
 munication inference. *Nature Cell Biology*, 26(9):1613–1622, September 2024a. ISSN

648 1476-4679. doi: 10.1038/s41556-024-01469-w. URL <http://dx.doi.org/10.1038/s41556-024-01469-w>.

649

650

651 Daniel Dimitrov, Philipp Sven Lars Schäfer, Elias Farr, Pablo Rodriguez-Mier, Sebastian Loben-

652 tanzer, Pau Badia-i Mompel, Aurelien Dugourd, Jovan Tanevski, Ricardo Omar Ramirez Flores,

653 and Julio Saez-Rodriguez. Liana+ provides an all-in-one framework for cell-cell communication

654 inference. *Nature Cell Biology*, 26(9):1613–1622, 2024b.

655 Kalliopi Domvri, Paul Zarogoulidis, Kaid Darwiche, Robert F. Browning, Qiang Li, J. Francis

656 Turner, Ioannis Kioumis, Dionysios Spyros, Konstantinos Porpodis, Antonis Papaioannou,

657 Theodora Tsiodra, Lutz Freitag, and Konstantinos Zarogoulidis. Molecular targeted drugs and

658 biomarkers in nsclc, the evolving role of individualized therapy. *Journal of Cancer*, 4(9):736–754,

659 2013. ISSN 1837-9664. doi: 10.7150/jca.7734. URL <http://dx.doi.org/10.7150/jca.7734>.

660

661 Mirjana Efremova, Miquel Vento-Tormo, Sarah A. Teichmann, and Roser Vento-Tormo. Cell-

662 phonedb: inferring cell-cell communication from combined expression of multi-subunit

663 ligand–receptor complexes. *Nature Protocols*, 15(4):1484–1506, February 2020. ISSN

664 1750-2799. doi: 10.1038/s41596-020-0292-x. URL <http://dx.doi.org/10.1038/s41596-020-0292-x>.

665

666 Rossin Erbe, Genevieve Stein-O’Brien, and Elana J. Fertig. Transcriptomic forecasting with neural

667 ordinary differential equations. *Patterns*, 4(8):100793, August 2023. ISSN 2666-3899. doi:

668 10.1016/j.patter.2023.100793. URL <http://dx.doi.org/10.1016/j.patter.2023.100793>.

669

670 Luca Eyring, Dominik Klein, Théo Uscidda, Giovanni Palla, Niki Kilbertus, Zeynep Akata, and

671 Fabian Theis. Unbalancedness in neural monge maps improves unpaired domain translation.

672 *arXiv preprint arXiv:2311.15100*, 2023.

673

674 Nona Farbahi, Ralph Patrick, Aude Dorison, Munira Xaymardan, Vaibhao Janbandhu, Katharina

675 Wystub-Lis, Joshua WK Ho, Robert E Nordon, and Richard P Harvey. Single-cell expression

676 profiling reveals dynamic flux of cardiac stromal, vascular and immune cells in health and injury.

677 *eLife*, 8, March 2019. ISSN 2050-084X. doi: 10.7554/elife.43882. URL <http://dx.doi.org/10.7554/elife.43882>.

678

679 Kilian Fatras, Younes Zine, Szymon Majewski, Rémi Flamary, Rémi Gribonval, and Nicolas

680 Courty. Minibatch optimal transport distances; analysis and applications. *arXiv preprint*

681 *arXiv:2101.01792*, 2021.

682

683 Rémi Flamary, Cédric Vincent-Cuaz, Nicolas Courty, Alexandre Gramfort, Oleksii Kachaiev, Huy

684 Quang Tran, Laurène David, Clément Bonet, Nathan Cassereau, Théo Gnassounou, Eloi Tanguy,

685 Julie Delon, Antoine Collas, Sonia Mazelet, Laetitia Chapel, Tanguy Kerdoncuff, Xizheng Yu,

686 Matthew Feickert, Paul Krzakala, Tianlin Liu, and Eduardo Fernandes Montesuma. Pot python

687 optimal transport (version 0.9.5), 2024. URL <https://github.com/PythonOT/POT>.

688

689 Aden Forrow and Geoffrey Schiebinger. Lineageot is a unified framework for lineage tracing and tra-

690 jectory inference. *Nature Communications*, 12(1), August 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-25133-1. URL <http://dx.doi.org/10.1038/s41467-021-25133-1>.

691

692 Nicolas Fournier and Arnaud Guillin. On the rate of convergence in wasserstein distance of the

693 empirical measure. *Probability theory and related fields*, 162(3):707–738, 2015.

694

695 Katherine C. Goldfarbmuren, Nathan D. Jackson, Satria P. Sajuthi, Nathan Dyjack, Katie S. Li,

696 Cydney L. Rios, Elizabeth G. Plender, Michael T. Montgomery, Jamie L. Everman, Preston E.

697 Bratcher, Eszter K. Vladar, and Max A. Seibold. Dissecting the cellular specificity of smoking

698 effects and reconstructing lineages in the human airway epithelium. *Nature Communications*, 11

699 (1), May 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-16239-z. URL <http://dx.doi.org/10.1038/s41467-020-16239-z>.

700

701 Laleh Haghverdi, Maren Büttner, F. Alexander Wolf, Florian Buettner, and Fabian J. Theis. Dif-

fusion pseudotime robustly reconstructs lineage branching. *Nature Methods*, 13(10):845–848,

702 2016. doi: 10.1038/nmeth.3971.

702 Douglas Hanahan and Robert A. Weinberg. Hallmarks of cancer: The next generation. *Cell*, 144
 703 (5):646–674, March 2011. ISSN 0092-8674. doi: 10.1016/j.cell.2011.02.013. URL <http://dx.doi.org/10.1016/j.cell.2011.02.013>.

704

705 Xing He and Chenqi Xu. Immune checkpoint signaling and cancer immunotherapy. *Cell Research*,
 706 30(8):660–669, May 2020. ISSN 1748-7838. doi: 10.1038/s41422-020-0343-4. URL <http://dx.doi.org/10.1038/s41422-020-0343-4>.

707

708 Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth schrödinger
 709 bridges. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=GHyvvWu1XC>.

710

711 Intekhab Hossain, Viola Fanfani, Jonas Fischer, John Quackenbush, and Rebekka Burkholz. Bio-
 712 logically informed neuralodes for genome-wide regulatory dynamics. *Genome Biology*, 25(1):
 713 127, 2024.

714

715 Sinisa Hrvatin, Daniel R. Hochbaum, M. Aurel Nagy, Marcelo Cicconet, Keiramarie Robertson, Lu-
 716 cas Cheadle, Rapolas Zilionis, Alex Ratner, Rebeca Borges-Monroy, Allon M. Klein, Bernardo L.
 717 Sabatini, and Michael E. Greenberg. Single-cell analysis of experience-dependent transcriptomic
 718 states in the mouse visual cortex. *Nature Neuroscience*, 21(1):120–129, December 2017. ISSN
 719 1546-1726. doi: 10.1038/s41593-017-0029-5. URL <http://dx.doi.org/10.1038/s41593-017-0029-5>.

720

721 Sinisa Hrvatin, Daniel R Hochbaum, M Aurel Nagy, Marcelo Cicconet, Keiramarie Robertson, Lu-
 722 cas Cheadle, Rapolas Zilionis, Alex Ratner, Rebeca Borges-Monroy, Allon M Klein, et al. Single-
 723 cell analysis of experience-dependent transcriptomic states in the mouse visual cortex. *Nature
 724 neuroscience*, 21(1):120–129, 2018.

725

726 Guillaume Huguet, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik
 727 Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows
 728 for trajectory inference. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun
 729 Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=ahAEhOtVif>.

730

731 Erica L. Jackson, Nicholas Willis, Kim Mercer, Roderick T. Bronson, Denise Crowley, Raymond
 732 Montoya, Tyler Jacks, and David A. Tuveson. Analysis of lung tumor initiation and progression
 733 using conditional expression of oncogenic k-ras. *Genes & Development*, 15(24):3243–3248,
 734 December 2001. ISSN 1549-5477. doi: 10.1101/gad.943001. URL <http://dx.doi.org/10.1101/gad.943001>.

735

736 Kacper Kapusniak, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein,
 737 Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the
 738 data manifold. *Advances in Neural Information Processing Systems*, 37:135011–135042, 2024.

739

740 Dominik Klein, Théo Uscidda, Fabian J Theis, and marco cuturi. GENOT: Entropic (gromov)
 741 wasserstein flow matching with applications to single-cell genomics. In *The Thirty-eighth Annual
 742 Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=hjspWd7jvg>.

743

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

749

750 Dominik Klein, Giovanni Palla, Marius Lange, Michal Klein, Zoe Piran, Manuel Gander, Laetitia
 751 Meng-Papaxanthos, Michael Sterr, Lama Saber, Changying Jing, et al. Mapping cells through
 752 time and space with moscot. *Nature*, 638(8052):1065–1075, 2025b.

753

754 Gioele La Manno, Ruslan Soldatov, Amit Zeisel, and others. Rna velocity of single cells. *Nature*,
 755 560(7719):494–498, 2018. doi: 10.1038/s41586-018-0414-6.

756 Marius Lange, Volker Bergen, Michal Klein, Manu Setty, Bernhard Reuter, Mostafa Bakhti, Heiko
 757 Lickert, Meshal Ansari, Janine Schniering, Herbert B. Schiller, Dana Pe'er, and Fabian J. Theis.
 758 Cellrank for directed single-cell fate mapping. *Nature Methods*, 19(2):159–170, 2022. doi: 10.
 759 1038/s41592-021-01346-6.

760
 761 Marius Lange, Zoe Piran, Michal Klein, Bastiaan Spanjaard, Dominik Klein, Jan Philipp Junker,
 762 Fabian J. Theis, and Mor Nitzan. Mapping lineage-traced cells across time points with moslin.
 763 *Genome Biology*, 25(1), October 2024. ISSN 1474-760X. doi: 10.1186/s13059-024-03422-4.
 764 URL <http://dx.doi.org/10.1186/s13059-024-03422-4>.

765 Ju Yeon Lee, Hyun Tae Lee, Woori Shin, Jongseok Chae, Jaemo Choi, Sung Hyun Kim, Heejin
 766 Lim, Tae Won Heo, Kyeong Young Park, Yeon Ji Lee, Seong Eon Ryu, Ji Young Son, Jee Un
 767 Lee, and Yong-Seok Heo. Structural basis of checkpoint blockade by monoclonal antibodies in
 768 cancer immunotherapy. *Nature Communications*, 7(1), October 2016. ISSN 2041-1723. doi:
 769 10.1038/ncomms13354. URL <http://dx.doi.org/10.1038/ncomms13354>.

770
 771 Arthur Liberzon, Chet Birger, Helga Thorvaldsdóttir, Mahmoud Ghandi, Jill P. Mesirov, and Pablo
 772 Tamayo. The molecular signatures database hallmark gene set collection. *Cell Systems*, 1(6):
 773 417–425, December 2015. ISSN 2405-4712. doi: 10.1016/j.cels.2015.12.004. URL <http://dx.doi.org/10.1016/j.cels.2015.12.004>.

774
 775 Zaikang Lin, Sei Chang, Aaron Zweig, Minseo Kang, Elham Azizi, and David A. Knowles. In-
 776 terpretable neural odes for gene regulatory network discovery under perturbations, 2025. URL
 777 <https://arxiv.org/abs/2501.02409>.

778
 779 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow match-
 780 ing for generative modeling. In *The Eleventh International Conference on Learning Representa-
 781 tions*, 2022.

782
 783 Yaron Lipman, Marton Havasi, Peter Holderith, Neta Shaul, Matt Le, Brian Karrer, Ricky TQ
 784 Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code. *arXiv
 785 preprint arXiv:2412.06264*, 2024.

786
 787 Zhen-Ling Liu, Huan-Huan Chen, Li-Li Zheng, Li-Ping Sun, and Lei Shi. Angiogenic signaling
 788 pathways and anti-angiogenic therapy for cancer. *Signal Transduction and Targeted Therapy*, 8
 789 (1), May 2023. ISSN 2059-3635. doi: 10.1038/s41392-023-01460-1. URL <http://dx.doi.org/10.1038/s41392-023-01460-1>.

790
 791 Kevin R Moon, David Van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen,
 792 Kristina Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, et al. Visualizing
 793 structure and transitions in high-dimensional biological data. *Nature biotechnology*, 37(12):1482–
 794 1492, 2019.

795
 796 Daniel J. Murphy, Melissa R. Junttila, Laurent Pouyet, Anthony Karnezis, Ksenya Shchors,
 797 Duyen A. Bui, Lamorna Brown-Swigart, Leisa Johnson, and Gerard I. Evan. Distinct thresh-
 798 olds govern myc's biological output in vivo. *Cancer Cell*, 14(6):447–457, December 2008.
 799 ISSN 1535-6108. doi: 10.1016/j.ccr.2008.10.018. URL <http://dx.doi.org/10.1016/j.ccr.2008.10.018>.

800
 801 Huize Pan, Chenyi Xue, Benjamin J. Auerbach, Jiaxin Fan, Alexander C. Bashore, Jian Cui,
 802 Dina Y. Yang, Sarah B. Trignano, Wen Liu, Jianting Shi, Chinyere O. Ihuegbu, Erin C. Bush,
 803 Jeremy Worley, Lukas Vlahos, Pasquale Laise, Robert A. Solomon, Edward S. Connolly, An-
 804 drea Califano, Peter A. Sims, Hanrui Zhang, Mingyao Li, and Muredach P. Reilly. Single-
 805 cell genomics reveals a novel cell state during smooth muscle cell phenotypic switching and
 806 potential therapeutic targets for atherosclerosis in mouse and human. *Circulation*, 142(21):
 807 2060–2075, November 2020. ISSN 1524-4539. doi: 10.1161/circulationaha.120.048378. URL
 808 <http://dx.doi.org/10.1161/CIRCULATIONAHA.120.048378>.

809
 810 Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in
 811 Machine Learning*, 11(5-6):355–607, 2019. doi: 10.1561/2200000073.

810 Simone Picelli. Single-cell rna-sequencing: The future of genome biology is now. *RNA Biology*, 14
 811 (5):637–650, July 2016. ISSN 1555-8584. doi: 10.1080/15476286.2016.1201618. URL <http://dx.doi.org/10.1080/15476286.2016.1201618>.

812

813 Chengxiang Qiu, Junyue Cao, Beth K. Martin, Tony Li, Ian C. Welsh, Sanjay Srivatsan, Xingfan
 814 Huang, Diego Calderon, William Stafford Noble, Christine M. Disteche, Stephen A. Murray,
 815 Malte Spielmann, Cecilia B. Moens, Cole Trapnell, and Jay Shendure. Systematic reconstruction
 816 of cellular trajectories across mouse embryogenesis. *Nature Genetics*, 54(3):328–341, March
 817 2022. ISSN 1546-1718. doi: 10.1038/s41588-022-01018-x. URL <http://dx.doi.org/10.1038/s41588-022-01018-x>.

818

819 Chengxiang Qiu, Beth K. Martin, Ian C. Welsh, Riza M. Daza, Truc-Mai Le, Xingfan Huang, Eva K.
 820 Nichols, Megan L. Taylor, Olivia Fulton, Diana R. O’Day, Anne Roshella Gomes, Saskia Il-
 821 cisin, Sanjay Srivatsan, Xinxian Deng, Christine M. Disteche, William Stafford Noble, Nobuhiko
 822 Hamazaki, Cecilia B. Moens, David Kimelman, Junyue Cao, Alexander F. Schier, Malte Spiel-
 823 mann, Stephen A. Murray, Cole Trapnell, and Jay Shendure. A single-cell time-lapse of mouse
 824 prenatal development from gastrula to birth. *Nature*, 626(8001):1084–1093, February 2024. ISSN
 825 1476-4687. doi: 10.1038/s41586-024-07069-w. URL <http://dx.doi.org/10.1038/s41586-024-07069-w>.

826

827 Xiaojie Qiu, Qiaolin Mao, Yong Tang, Lihua Wang, Rakel Chawla, Hannah A. Pliner, and Cole
 828 Trapnell. Reversed graph embedding resolves complex single-cell trajectories. *Nature Methods*,
 829 14(10):979–982, 2017. doi: 10.1038/nmeth.4402.

830

831 Wouter Saelens, Robrecht Cannoodt, Helena Todorov, et al. A guide to trajectory inference and rna
 832 velocity. *bioRxiv*, 2022. doi: 10.1101/2021.12.22.473434.

833

834 Kristiyan Sakalyan, Alessandro Palma, Filippo Guerranti, Fabian J Theis, and Stephan Günemann.
 835 Modeling microenvironment trajectories on spatial transcriptomics with nicheflow. In *The Thirty-
 836 ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=5ofJyjgrth>.

837

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

843

844 Aylin Sertkaya, Trinidad Beleche, Amber Jessup, and Benjamin D. Sommers. Costs of drug devel-
 845 opment and research and development intensity in the us, 2000-2018. *JAMA Network Open*, 7
 846 (6):e2415445, June 2024. ISSN 2574-3805. doi: 10.1001/jamanetworkopen.2024.15445. URL
<http://dx.doi.org/10.1001/jamanetworkopen.2024.15445>.

847

848 Alex K Shalek, Rahul Satija, Joe Shuga, John J Trombetta, Dave Gennert, Diana Lu, Peilin Chen,
 849 Rona S Gertner, Jellert T Gaublomme, Nir Yosef, et al. Single-cell rna-seq reveals dynamic
 850 paracrine control of cellular variation. *Nature*, 510(7505):363–369, 2014.

851

852 Kelly Street, Davide Risso, Robert B. Fletcher, Diya Das, John Ngai, Nir Yosef, Elizabeth Purdom,
 853 and Sandrine Dudoit. Slingshot: cell lineage and pseudotime inference for single-cell transcrip-
 854 toomics. *BMC Genomics*, 19(1):477, 2018. doi: 10.1186/s12864-018-4772-0.

855

856 B. J. Strober, R. Elorbany, K. Rhodes, N. Krishnan, K. Tayeb, A. Battle, and Y. Gilad. Dy-
 857 namic genetic regulation of gene expression during cellular differentiation. *Science*, 364(6447):
 858 1287–1290, June 2019. ISSN 1095-9203. doi: 10.1126/science.aaw0040. URL <http://dx.doi.org/10.1126/science.aaw0040>.

859

860 Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. Trajectorynet:
 861 A dynamic optimal transport network for modeling cellular dynamics. In *Proceedings of the 37th
 862 International Conference on Machine Learning*, 2020. ICML 2020; arXiv:2002.04461.

863

864 Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume
 865 Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and
 866 flow matching. *arXiv preprint arXiv:2307.03672*, 2023.

864 Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-
 865 Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models
 866 with minibatch optimal transport. *Transactions on Machine Learning Research*, pp. 1–34, 2024.
 867

868 Betsabeh Khoramian Tusi, Samuel L. Wolock, Caleb Weinreb, Yung Hwang, Daniel Hidalgo, Rapo-
 869 las Zilionis, Ari Waisman, Jun R. Huh, Allon M. Klein, and Merav Socolovsky. Population snap-
 870 shots predict early haematopoietic and erythroid hierarchies. *Nature*, 555(7694):54–60, February
 871 2018. ISSN 1476-4687. doi: 10.1038/nature25741. URL <http://dx.doi.org/10.1038/nature25741>.
 872

873 Titouan Vayer, Laetitia Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Fused
 874 gromov–wasserstein distance for structured objects. *Algorithms*, 13(9):212, 2020a. doi:
 875 10.3390/a13090212.

876 Titouan Vayer, Laetitia Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Fused
 877 gromov-wasserstein distance for structured objects. *Algorithms*, 13(9):212, 2020b.
 878

879 Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2008.

880 Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-
 881 growth flow matching for single-cell dynamics modeling. In *The Thirty-ninth Annual Confer-
 882 ence on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=aXAkNlbGa>.
 883

884 Kathryn A Wierenga, Frank M Riemers, Bart Westendorp, Jack R Harkema, and James J Pestka.
 885 Single cell analysis of docosahexaenoic acid suppression of sequential lps-induced proinflamma-
 886 tory and interferon-regulated gene expression in the macrophage. *Frontiers in Immunology*, 13:
 887 993614, 2022.

888 F. Alexander Wolf, Fiona K. Hamey, Mireya Plass, Jordi Solana, Joakim S. Dahlin, Berthold
 889 Göttgens, Nikolaus Rajewsky, Lukas Simon, and Fabian J. Theis. Paga: graph abstraction rec-
 890 onciles clustering with trajectory inference through a topology preserving map of single cells.
 891 *Genome Biology*, 20(1):59, 2019. doi: 10.1186/s13059-019-1663-x.

892 Alan T. Yeo, Shruti Rawal, Bethany Delcuze, Anthos Christofides, Agata Atayde, Laura Strauss,
 893 Leonora Balaj, Vaughn A. Rogers, Erik J. Uhlmann, Hemant Varma, Bob S. Carter, Vassi-
 894 liki A. Boussiotis, and Al Charest. Single-cell rna sequencing reveals evolution of immune land-
 895 scape during glioblastoma progression. *Nature Immunology*, 23(6):971–984, May 2022. ISSN
 896 1529-2916. doi: 10.1038/s41590-022-01215-0. URL <http://dx.doi.org/10.1038/s41590-022-01215-0>.
 897

898 Rongting Yue and Abhishek Dutta. Computational systems biology in disease modeling and con-
 899 trol, review and perspectives. *npj Systems Biology and Applications*, 8(1), October 2022. ISSN
 900 2056-7189. doi: 10.1038/s41540-022-00247-4. URL <http://dx.doi.org/10.1038/s41540-022-00247-4>.
 901

902 Fabio Zanini, Makeda L. Robinson, Derek Croote, Malaya Kumar Sahoo, Ana Maria Sanz, Eliana
 903 Ortiz-Lasso, Ludwig Luis Albornoz, Fernando Rosso, Jose G. Montoya, Leslie Goo, Benjamin A.
 904 Pinsky, Stephen R. Quake, and Shirit Einav. Virus-inclusive single-cell rna sequencing reveals
 905 the molecular signature of progression to severe dengue. *Proceedings of the National Academy
 906 of Sciences*, 115(52), December 2018. ISSN 1091-6490. doi: 10.1073/pnas.1813819115. URL
 907 <http://dx.doi.org/10.1073/pnas.1813819115>.
 908

909

910

911

912

913

914

915

916

917

APPENDIX

A Extended related works

Infering cellular trajectories. Methods for trajectory inference differ in both their assumptions and the temporal scales they target. Pseudotime approaches order cells along low-dimensional embeddings (such as UMAP), capturing smooth expression trends but relying on manifold geometry rather than explicit dynamical models (Erbe et al., 2023). RNA velocity augments expression with spliced and unspliced counts to estimate short-term directional change (typically minutes to hours), but its accuracy degrades over longer horizons and in the presence of sparse signals (Chen et al., 2022). More recent formulations use continuous-time models to interpolate between snapshots: Neural ODEs and dynamic optimal transport (OT) learn flows across cell states, with variants tailored to gene-regulatory-network dynamics or intervention-aware recovery in scRNA-seq (Lin et al., 2025), and others incorporating biological priors to regularize the inferred trajectories (e.g., PHOENIX) (Hossain et al., 2024). DeepVelo applies Neural ODEs to high-dimensional, sparse measurements, yielding predictive flows without committing to a mechanistic model (Chen et al., 2022). Related work couples dynamics with OT on learned manifolds to better respect transcriptomic geometry during alignment (Huguet et al., 2022). A closely related line casts dynamics as a Schrödinger bridge used for trajectory inference and generative modeling in single-cell RNA data (Hong et al., 2025).

Optimal transport in biology. Optimal transport (OT) is widely used for static alignment of cellular populations in biology and has been extended to dynamic settings for modeling complex scRNA-seq trajectories (Tong et al., 2020). High dimensionality is a central challenge, and remedies include dimensionality reduction, regularization, and scalable solvers (Cuturi et al., 2023), often operating in learned or low-dimensional representations (e.g. PCA or manifold embeddings) to better respect transcriptomic geometry (Huguet et al., 2022). OT-based frameworks have been applied to recover cellular trajectories in development (Schiebinger et al., 2019), and incorporating inductive biases (such as lineage information) can further improve identifiability and accuracy (Forrow & Schiebinger, 2021). Recent work relates OT and continuous-time dynamics via flow matching Klein et al. (2024) and explores multi-modal integration directly within the OT formulation (Klein et al., 2025a). An overview of different OT-based trajectory methods and their mathematical fomulation is provided in Table 4.

Table 4: **OT-based methods: objectives and assumptions.** BB = Benamou–Brenier; FGW = Fused Gromov–Wasserstein; SB/DSB = (Diffusion) Schrödinger Bridge.

Method	Static / dynamic	OT formulation	Optimization objective (schematic)	Structure used	Timepoints	Data assumed
Waddington-OT (Schiebinger et al., 2019)	dynamic (discrete)	Unbalanced entropic Kantorovich	$\min_{\Gamma} \langle \Gamma, C \rangle + \tau \text{KL}(\Gamma \ q) + \tau \text{KL}(\Gamma^\top \ b)$	Growth priors (no LR)	≥ 2	scRNA-seq (+growth)
SCOT (Demetsi et al., 2022b)	static	Fused GW	$\min_{\Gamma} \langle \Gamma, C \rangle + \lambda \langle \Gamma, D^x, D^y, \Gamma^\top \rangle + \text{H}(\Gamma)$	k-NN geometry (untyped)	2	scRNA-seq
CellOT (Bunne et al., 2023b)	static	Kantorovich (dual/convex)	$\max_{\Gamma} \min_{\mu} \mathbb{E}_{\mu} [\langle x, \nabla g(x) \rangle - f(\nabla g(x)) - \mathbb{E}_{\mu} [f(y)]$	None	2	scRNA-seq
TrajectoryNet (Tong et al., 2020)	dynamic (continuous)	BB OT + neural ODE prior	$\min_{\Gamma} \min_{\mu} - \sum_t \log P_t(x_t) + \int \ u\ ^2 dt + \text{bio priors}$	None	≥ 2	scRNA-seq (often splicing)
OT-CFM (Tong et al., 2024)	dynamic (continuous)	Kantorovich + flow matching	$\min_{\Gamma} \langle \Gamma, C \rangle : \min_{\mu} \mathbb{E}_{(\mu, y) \sim \Gamma} \ y - \mu\ ^2$	None	2	scRNA-seq
Schrödinger Bridge (DSB)	dynamic (stochastic)	SB/DSB	$\min_{\mu \sim p(\mu \mid \Gamma)} \text{KL}(\mu \parallel \text{ref-diffusion}) \text{ s.t. } \mu_0 = \mu_0, \mu_T = \mu_T$	None	≥ 2	scRNA-seq (+noise model)
IADOT (ours)	dynamic (continuous)	Multi-channel FGW + CFM	$\min_{\Gamma} (1 - \alpha) \langle \Gamma, C \rangle + \alpha \langle \varphi(G^{(0)}, G^{(1)}), \Gamma \otimes \Gamma \rangle : \min_{\theta} \text{CFM}(\theta \mid \Gamma)$	Typed, directed LR CCI	2 (extendable)	scRNA-seq

Interaction modeling between cells A line of work seeks to infer cell–cell communication directly from single-cell gene expression, using curated ligand–receptor (LR) knowledge to score putative interactions between sender–receiver pairs (Browaeys et al., 2019). Tools such as Cell-PhoneDB systematically enumerate LR co-expression across cell types (Efremova et al., 2020), and related approaches have been extended to spatial transcriptomics to incorporate physical proximity as an additional constraint on feasible communications (Cang et al., 2023). Beyond purely geometric priors, multi-modal OT frameworks like MOSCOT can integrate diverse structure (e.g. spatial adjacency) into the coupling itself (Klein et al., 2025a). Finally, meta-frameworks like LIANA+ unify and standardize CCI scoring across multiple LR resources and methods, facilitating method-agnostic comparisons and consensus analyses (Dimitrov et al., 2024a).

Comparison with Related Works Table 1 compares our proposed framework against existing state-of-the-art methods across five key capabilities essential for modeling complex cellular dynamics. We define these criteria as follows:

- **Dynamic** indicates whether the method explicitly models temporal evolution across multiple experimental timepoints, as opposed to inferring dynamics or trajectories from a single static snapshot.
- **Trajectories** distinguishes methods that recover a continuous smooth path enabling predictions at unobserved intermediate timepoints from those that solely compute discrete couplings or transport maps between timepoints.
- **In-silico Perturbation** refers to the capability to perform principled interventions, allowing users to simulate and predict the system’s response to specific stimuli or perturbations.
- **Structure-Aware** assesses whether the optimization objective explicitly models interactions between cells (e.g., via cell-cell communication or topological constraints) rather than treating cells as independent, isolated entities.
- **scRNA Data Sufficient** confirms whether the method can operate effectively using standard single-cell RNA sequencing inputs alone, without requiring auxiliary spatial transcriptomics data or multi-modal integration that are often unavailable.

B Potential applications of IADOT

Snapshots of cellular systems using single-cell RNA sequencing are now pervasive across diverse areas of biology and medicine. A few representative longitudinal datasets are summarized in Table 5. IADOT provides a principled framework to analyze such data by combining snapshot measurements with biologically typed ligand–receptor structure. This enables the reconstruction of coherent cell-state trajectories through optimal transport couplings and a learned continuous flow, as well as the exploration of counterfactual scenarios by selectively re-weighting interaction channels. The resulting outputs (shifts in lineage fate, changes in pathway usage, and differences in progression timing) offer interpretable readouts that can guide mechanistic hypotheses and help prioritize therapeutic strategies before experimental validation.

Table 5: Public longitudinal single-cell datasets. Each row lists an application area, a brief description, and representative studies/accessions (not exhaustive).

Area	Dataset description (≥ 3 timepoints)	References / accessions
Virology	PBMC/tissue scRNA-seq across acute, peak critical or challenge series (D0, D1–3, D7+).	Dengue virus: (Zanini et al., 2018) Influenza: (Arunachalam et al., 2021)
Neurology	Brain single-cell timecourses including immune infiltration and glial responses.	Brain organoids: (Camp et al., 2015)
Cardiology	Heart/aorta scRNA-seq after myocardial infarction (e.g., D1, D3, D7) or atherosclerosis progression (early → intermediate → late).	Post-MI heart: (Farbahi et al., 2019) Atherosclerosis: (Pan et al., 2020)
Immunology	Tissue + immune scRNA-seq across baseline → active disease → remission/recovery in model systems.	Lung: (Goldfarbmuren et al., 2020)
Development	Human iPSC/hPSC differentiation series (e.g., D0, D4, D8, D12/15), tracking lineage commitment and maturation.	Cardiomyocytes: (Strober et al., 2019) Blood cells: (Tusi et al., 2018)
Regeneration	Liver/kidney/muscle injury timecourses (e.g., 0h, 24h, 48h/96h; or 0d, 2d, 5d, 7d) capturing repair trajectories.	Liver injury: (Chen et al., 2023)

C Datasets

In addition to the synthetic dataset, we used 5 real-world scRNA datasets to showcase the effectiveness and limitations of our method. Details on the number of genes and the number of cells in each dataset can be found in Table 6.

1026 **Table 6: Datasets used in our experiments.** Counts reflect the preprocessed objects used by
 1027 IADOT. “Timestamps (h)” lists observed hours.
 1028

1029 <i>Dataset</i>	1030 <i>Reference</i>	1031 <i>Timestamps</i>	1032 <i>#Cells</i>	1033 <i>#Genes</i>
1031 Tumour	1032 –	1033 0, 8, 24, 168 (h)	1034 31,536	1035 22,681
1032 V1 Cortex	1033 (Hrvatin et al., 2018)	1034 0, 1, 4 (h)	1035 6,505	17,008
1033 Immune Stimulus	1034 (Wierenga et al., 2022)	1035 0, 1, 2, 4, 6 (h)	1036 2,382	10,972
1034 Mouse embryo	1035 (Moon et al., 2019)	1036 0, 6, 12, 18, 24 (d)	1037 18,203	1038 17,789
1035 Macrophage Stimulus	1036 (Shalek et al., 2014)	1037 0, 3, 5 (h)	1038 223	1039 478

1036 C.1 SYNTHETIC EXAMPLE

1039 In this section we detail the synthetic setup used in Section 5.1. We construct \mathcal{D}_0 as three 2D
 1040 Gaussian clusters,

$$1041 \mathcal{D}_0 = \bigcup_{k=0}^2 \mathcal{S}_k, \quad \mathcal{S}_k = \{X_i^{(k)}\}_{i=1}^{35}, \quad X_i^{(k)} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_k, 0.1 I_2),$$

1044 with centers $\mu_0 = (-2, 2)$, $\mu_1 = (0, 2)$, and $\mu_2 = (2, 2)$. The target snapshot $\mathcal{D}_1 = \bigcup_{k=0}^2 \mathcal{S}'_k$ is
 1045 obtained by translating each cluster via

$$1046 T_0(x) = x + (4, -4), \quad T_1(x) = x + (0, -4), \quad T_2(x) = x + (-4, -4),$$

1047 so that $\mathcal{S}'_k = \{T_k(X) : X \in \mathcal{S}_k\}$.

1049 For structure, we define two-channel, directed relation tensors $G, G' \in \{0, 1\}^{105 \times 105 \times 2}$ over \mathcal{D}_0
 1050 and \mathcal{D}_1 , respectively. Writing $G^{(c)}$ for channel c , we set

$$1051 G_{ij}^{(1)} = \mathbf{1}\{X_i \in \mathcal{S}_1, X_j \in \mathcal{S}_0\}, \quad G_{ij}^{(2)} = \mathbf{1}\{X_i \in \mathcal{S}_1, X_j \in \mathcal{S}_2\},$$

1053 with G' defined analogously on \mathcal{D}_1 . Thus, channel 1 encodes $\mathcal{S}_1 \rightarrow \mathcal{S}_0$ and channel 2 encodes
 1054 $\mathcal{S}_1 \rightarrow \mathcal{S}_2$.

1055 C.2 LUNG TUMOR

1058 We use a scRNA-seq dataset to study rapid tumour progression driven by RAS–MYC signalling
 1059 using a *Kras*^{G12D} lung tumour model with tamoxifen-inducible MycER. Samples were collected
 1060 at 0 h (vehicle), 8 h, 24 h ($n = 8$ biological replicates per condition; 0 h is time zero). Lungs
 1061 from LSL-*Kras*^{G12D} (Jackson et al., 2001) and LSL-*Rosa26*^{MIE/MIE} (*MycERT2*) mice (Murphy et al.,
 1062 2008) were dissociated to single cells, red blood cells removed, filtered (70 μm), and 6,000 cells per
 1063 sample were loaded for 10x Chromium 3' v3 libraries. Libraries were sequenced on a NovaSeq 6000
 1064 and processed with Cell Ranger v6.1.1 against mm10. All animal work complied with institutional
 1065 ethical regulations.

1066 C.3 V1 CORTEX—LIGHT STIMULATION

1068 Adult (6–8 week) mice were dark-adapted for 7 days, then either euthanized in darkness (0h, con-
 1069 trol) or exposed to ambient light for 1h or 4h (Hrvatin et al., 2017). The visual cortex was profiled
 1070 by scRNA-seq to capture early transcriptional responses to sensory input. We treat 0h as the source
 1071 snapshot, 4h as the target snapshot, and use 1h as an intermediate timepoint for interpolation/validation.
 1072 After filtering and subsampling we are left with 6505 cells.

1073 C.4 IMMUNE

1075 To probe innate immune modulation, we use scRNA-seq of murine fetal liver-derived macrophages
 1076 exposed to LPS with or without 24 h pre-treatment by docosahexaenoic acid (DHA, 25 μM)
 1077 (Wierenga et al., 2022). Cells were collected at 0 h (vehicle), 1 h, and 4 h after LPS (20 ng/mL)
 1078 and sequenced on the 10x Chromium platform. We use 0h as source, 4h as target, and 1h for inter-
 1079 polation/validation; when comparing conditions, we stratify by DHA vs. vehicle and subsample to
 balance groups.

1080
1081
1082 **Table 7: Software stack (key Python packages).**
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

Core scientific Python	NumPy 2.2.6; SciPy 1.16.1; pandas 2.3.1; scikit-learn 1.7.1; numba 0.61.2; matplotlib 3.10.5; seaborn 0.13.2
Deep learning / training	PyTorch 2.8.0; PyTorch Lightning 2.5.3; torchmetrics 1.8.1; Triton 3.4.0; Hydra-core 1.3.2; OmegaConf 2.3.0
Optimal transport / geometry	POT 0.9.5; GeomLoss 0.2.6; GraphTools 1.5.3
Single-cell analysis	Scanpy 1.11.4; anndata 0.12.2; scVelo 0.3.3; harmonypy 0.0.10; UMAP-learn 0.5.9.post2; PHATE 1.0.11; igraph 0.11.9; leidenalg 0.10.2; networkx 3.5; OmniPath 1.0.12; pypath-omnipath 0.16.20

C.5 EMBRYO DEVELOPMENT

In Section 5.6, we analyze a human embryoid body (EB) differentiation time course used in [Moon et al. \(2019\)](#), which profiles human embryonic stem cells differentiating toward germ layers over 27 days by scRNA-seq. We use the first (Day 0) and third (Day 12) snapshot to infer the cellular dynamics, reserving data at Day 6 for interpolation/validation.

C.6 MACROPHAGE STIMULUS

To evaluate robustness across experimental platforms, we deliberately included datasets generated with multiple scRNA-seq technologies: 10x Chromium (droplet-based, whole-transcriptome) and BD Rhapsody (microwell-based, often targeted or lower-depth whole-transcriptome). This cross-platform design allows us to test whether our method generalizes despite differences in capture chemistry, library preparation, and typical read depth, which can affect UMI yield and the number of detected genes per cell. All datasets were processed through a consistent downstream pipeline to ensure comparability. We use a macrophage stimulus-response time series that profiled single-cell dynamics across three polarization states (M0, M1 via IFN_γ , M2 via IL-4) responding to six immune ligands (LPS, poly(I:C), CpG, PCSK3) ([Shalek et al., 2014](#)). Cells were sampled at 0 h (baseline) and multiple post-stimulation time points (15/30 min, 1 h, 3 h, 5 h, 8 h), (BD Rhapsody). For our alignment tasks we treat 0 h as the source snapshot, 5 h as the target, and use 3 h as intermediate validation points. Experimental results for this dataset can be found in Section E.1.

D Experimental details

In what follows, we provide details about our experiments presented in Section 5. Code will be released upon acceptance.

D.1 SOFTWARE AND LIBRARIES USED

We provide in Table 7 the main Python packages we used.

D.2 DATA PRE-PROCESSING

Raw scRNA-seq files for all datasets were converted to AnnData to standardize processing. We applied basic QC, removing cells with < 300 detected genes and genes expressed in < 3 cells. Counts were library-size normalized per cell (fixed total). We then selected the 2000 highly variable genes and computed a 20-component PCA on these features. Finally, we performed Harmony batch correction in PCA space (retaining both corrected and uncorrected embeddings for downstream analyses).

D.3 CONSTRUCTING CCIs USING METACELLS

We detail how we construct CCIs using metacells in the ablation presented in Section 5.5. Without loss of generality and to keep the presentation simple (with matrix multiplications), we assume $K = 1$ (i.e., one LR pair) reducing the CCI tensors to matrices. Before constructing the CCI matrices, we cluster the cells in each snapshot using Leiden community detection on a k -nearest-neighbour (kNN) graph built from the PCA representations with Euclidean distances and $k = 10$. An example of the Leiden clustering with subsequent cell annotations is provided in Figure 7. We

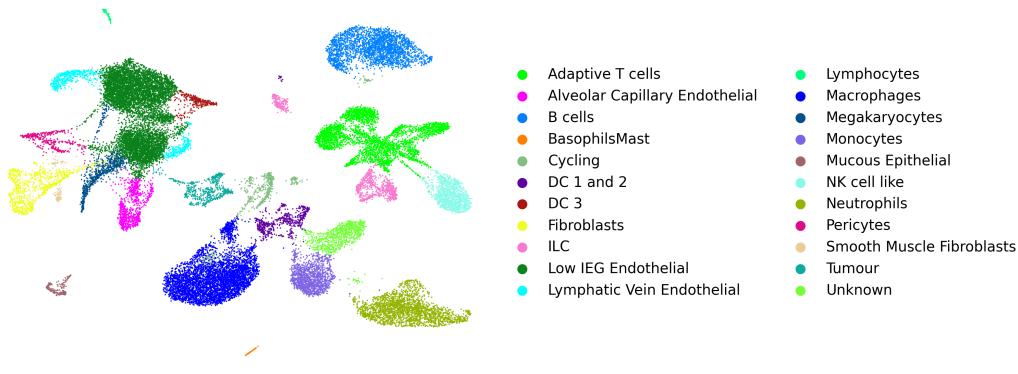
1134 select the resolution ρ^* by scanning a small grid of resolutions and choosing the value whose *median*
 1135 cluster size is closest to a target of $n^* = 40$ cells.
 1136

1137 Let $S \in \mathbb{R}_{\geq 0}^{n \times g}$ be the membership matrix of the resulting g clusters (rows sum to 1, are correspond
 1138 to one-hot assignments). We obtain metacell-level activations by averaging the s_{cg} within clusters
 1139 and form the metacell CCI in $\mathbb{R}^{g \times g}$ similarly as in the setting with individual cells.

1140 Having constructed the metacell CCI matrix \tilde{G} , we lift it back to the cell level via
 1141

$$\tilde{G} = S(S^\top S)^{-1} \tilde{G} (S^\top S)^{-1} S^\top,$$

1143 This lifting operation ensures $S^\top \tilde{G} S = \tilde{G}$. In contrast to G , the matrix \tilde{G} is constrained to lie
 1144 in the subspace $\{SMS^\top \mid M \in \mathbb{R}^{g \times g}\}$, i.e., cell-cell interactions in \tilde{G} are entirely mediated by
 1145 metacell-metacell interactions.



1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160 **Figure 7: Metacell construction example.** UMAP visualization of the single-cell RNA-seq data of
 1161 the lung cancer dataset after Leiden clustering. Each point corresponds to an individual cell, colored
 1162 by its assigned cluster and annotated with the corresponding cell type based on marker genes.

1163 D.4 OPTIMAL TRANSPORT SOLVER

1164 We extend POT’s (Flamary et al., 2024) conditional-gradient (Frank–Wolfe) solver to handle multi-
 1165 channel interactions. Given structure tensors $C_1 \in \mathbb{R}^{n_s \times n_s \times d}$, $C_2 \in \mathbb{R}^{n_t \times n_t \times d}$, marginals p, q
 1166 (uniform by default), and a matrix $\Sigma \succeq 0 \in \mathbb{R}^{d \times d}$, we measure discrepancies with the Mahalanobis
 1167 norm $\|x\|_\Sigma = \sqrt{x^\top \Sigma^{-1} x}$.

1168 Let $\langle A, B \rangle = \sum_{i,j} A_{ij} B_{ij}$ and write $C^{(r)}$ for the r -th channel of a structure tensor C . The GW
 1169 quadratic term is

$$1170 \mathcal{Q}(\Gamma) = \sum_{i,k,j,l} \|C_1[i, k] - C_2[j, l]\|_\Sigma^2 \Gamma_{ij} \Gamma_{kl} = \langle \text{constC}, \Gamma \rangle - \langle \mathcal{B}(\Gamma), \Gamma \rangle,$$

1171 with

$$1172 \text{constC}_{ij} = \sum_k \|C_1[i, k]\|_\Sigma^2 p_k + \sum_l \|C_2[j, l]\|_\Sigma^2 q_l, \quad \mathcal{B}(\Gamma) = \sum_{r=1}^d C_1^{(r)} \Gamma (C_2^{(r)})^\top.$$

1173 The gradient computed by the solver is

$$1174 \nabla \mathcal{Q}(\Gamma) = 2(\text{constC} - \mathcal{B}(\Gamma)) \tag{8}$$

1175 We keep POT’s CG loop, stopping criteria, and line-search options unchanged.

1176 We minimize

$$1177 \min_{\Gamma \in \Pi(p, q)} (1 - \alpha) \langle M, \Gamma \rangle + \alpha \mathcal{Q}(\Gamma),$$

1178 with the same CG loop, where this objective is linearized using the gradient in Equation (8).

1179 When $d = 1$ (scalar edges), the method reduces to the original POT solver.

Table 8: Flow Matching hyperparameters.

Category	Hyperparameter	Setting / Notes
Model	Architecture	Velocity MLP (hidden dim 64, depth 3, no dropout)
Training schedule	Time conditioning	Sinusoidal embedding (dim 16), concatenated to inputs
	Epochs	500
	Minibatch size	128 (train loader), 2048 for validation batches
	Optimizer	AdamW, lr = 10^{-3} , weight decay 10^{-4} , betas (0.9, 0.999)

D.5 NORMALIZATION

To balance the contributions of the feature term and the structure term in the objective described at Equation (5), we rescale the feature cost matrix C and the CCI tensors $G^{(0)}$ and $G^{(1)}$. We first compute the two endpoint couplings by solving the feature-only ($\alpha = 0$) and structure-only ($\alpha = 1$) problems, yielding $\Gamma_{\alpha=0}^*$ and $\Gamma_{\alpha=1}^*$. We then define the scaling factors as follows:

$$\Delta\mathcal{F} := \mathcal{F}(\Gamma_{\alpha=0}^*) - \mathcal{F}(\Gamma_{\alpha=1}^*), \quad (9)$$

$$\Delta\mathcal{S} := \mathcal{S}(\Gamma_{\alpha=1}^*) - \mathcal{S}(\Gamma_{\alpha=0}^*), \quad (10)$$

$$(11)$$

and rescale the feature cost matrix and the CCI tensors:

$$C \leftarrow \frac{C}{|\Delta\mathcal{F}|}, \quad (12)$$

$$G^{(0)} \leftarrow \frac{G^{(0)}}{\sqrt{\Delta\mathcal{S}}}, \quad G^{(1)} \leftarrow \frac{G^{(1)}}{\sqrt{\Delta\mathcal{S}}}. \quad (13)$$

This places the terms on comparable scales so that α meaningfully reflects the feature/structure trade-off, and increasing α from 0 to 1 smoothly interpolates between the Kantorovich and the Gromov–Wasserstein problems.

D.6 SELECTION OF LIGAND / RECEPTOR PAIRS

We apply LIANA’s (Dimitrov et al., 2024b) consensus rank aggregation with `expr_prop` = 0.1 to obtain per-cell-type interaction scores. We retain interactions with `cellphone_pvals` ≤ 0.05 and `lr_logfc` ≥ 0 , then keep ligand–receptor pairs whose `expr_prod` exceeds the median within that significant set. We require the same significance criteria in each snapshot. For every surviving pair, we aggregate LIANA results across significant edges to compute the mean expression product, average specificity ranks, counts of significant source→target edges, and the numbers of unique source and target cell types. We define coverage as `coverage` = `n.edges`/ $N_{\text{sig edges}}$ and retain only pairs with $0.10 \leq \text{coverage} \leq 0.40$ and at least two sources and two targets. We compute a standardized score $s = 0.6 z(\text{mean_expr}) + 0.4 z(-\text{spec_rank})$ and greedily select pairs in descending s while preventing repeated ligands or receptors. We keep the top 10 pairs for each dataset.

D.7 CONDITIONAL FLOW MATCHING HYPERPARAMETERS

We detail the hyperparameters used for the CFM stage in Table 8, which we kept fixed across the datasets. Given a 0.9/0.1 train/val split, we keep the checkpoint that minimizes the validation loss over the run.

D.8 BASELINES

TrajectoryNet. We use the implementation from the authors (Tong et al., 2020) available at <https://github.com/KrishnaswamyLab/TrajectoryNet>. We summarize the hyperparameters used in Table 9.

Table 9: TrajectoryNet hyperparameters.

Category	Hyperparameter	Setting
Optimization	Training iterations	1,000
	Batch size	1,000
	Learning rate	1×10^{-3}
	Weight decay	1×10^{-5}
Model	Blocks / layer type	1 block, concatsquash layers Hidden dimensions: 64-64-64
	Activations	Softplus layers with tanh control
Regularization	$s_{L2\text{int}}$	1×10^{-3}
	k_{top} regularizer	1×10^{-2}
	Training noise	0.1
ODE solver	Time scale	0.4 (five integration points)
	Solver	dopri5
	Tolerances	$\text{rtol} = \text{atol} = 1 \times 10^{-5}$

Diffusion Schrodinger Bridges. We use the implementation from the authors (De Bortoli et al., 2021) available at https://github.com/JTT94/diffusion_schrodinger_bridge. We summarize the hyperparameters used in Table 10.

MIOFlow We use the implementation from the authors (Huguet et al., 2022) available at <https://github.com/KrishnaswamyLab/MIOFlow>. We summarize the hyperparameters used in Table 11.

Moscot We use the implementation from the authors (Klein et al., 2025b) available at <https://github.com/theislab/moscot>. We summarize the hyperparameters used in Table 12.

VGFM We use the implementation from the authors (Wang et al., 2025) available at <https://github.com/DongyiWang-66/VGFM>. We summarize the hyperparameters used in Table 13.

MFM We use the implementation from the authors (Kapusniak et al., 2024) available at <https://github.com/kkapusniak/metric-flow-matching>. We summarize the hyperparameters used in Table 14.

SF2M We developed a custom implementation of the SF2M framework (Tong et al., 2023) to enable the integration of the IADOT structural prior into the simulation-free training objective. We summarize the hyperparameters used in Table 15.

UOT-FM To incorporate the interaction-aware coupling in an unbalanced setting, we utilized a custom implementation of UOT-FM based on the original formulation (Eyring et al., 2023). We summarize the hyperparameters used in Table 16.

Table 10: Diffusion Schrödinger Bridge (DSB) baseline hyperparameters.

Category	Hyperparameter	Setting
Model	Score network	Encoder [16, 32] Decoder [64, 64, 64], dim =16
Training schedule	IPF rounds	10 outer IPF iterations
	Optimisation steps	10 000 gradient updates
	Langevin steps	12 steps per bridge trajectory
	Batch size	128
Regularisation	Learning rate	1×10^{-4}
	γ schedule	$\gamma_{\min} = \gamma_{\max} = 10^{-3}$, linear spacing
	Mean matching	Enabled
	EMA	Disabled

1296 Table 11: MioFlow Baseline Hyperparameters.
1297

1298 Category	1299 Hyperparameter	1300 Setting
1301 Model Architecture	1302 Network layers	1303 [64, 64, 64]
1304 Optimization & Network Training	1305 Learning rate (LR)	1306 1×10^{-4}
	1307 Total epochs (N_{epochs})	20
	1308 Local epochs per stage (N_{local})	5
	1309 Post-local epochs ($N_{\text{post_local}}$)	5
	1310 Sample size (Batch size)	256
	1311 Number of batches per epoch (N_{batches})	100

1308 Table 12: Moscot Baseline Hyperparameters.
1309

1310 Category	1311 Hyperparameter	1312 Setting
1313 Optimal Transport Parameters	1314 Epsilon (ϵ , regularization)	0.001
	1315 τ_a (Source regularization)	1.0
	1316 τ_b (Target regularization)	1.0

1315 Table 13: VGFM Baseline Hyperparameters.
1316

1317 Category	1318 Hyperparameter	1319 Setting
1320 Model Architecture	1321 Hidden dimension (hidden_dim)	64
	1322 Number of hidden layers (n_hiddens)	3
	1323 Activation function	Tanh
1324 Optimization	1325 Pre-train epochs	300
	1326 Training epochs	50
	1327 Batch size	256
	1328 Initial learning rate (learning_rate1)	1×10^{-3}
	1329 Second learning rate (learning_rate2)	1×10^{-4}
	1330 Stepsize (for solver)	0.01

1329 D.9 LUNG CANCER DATA EXPERIMENT

1331 For the experiment described in Section 5.4, We annotated the lung cancer dataset using canonical
 1332 lineage and state markers (Table 17); an overview of the full dataset is shown in Fig. 7. Because
 1333 whole-lung profiling dilutes treatment effects (the tumour comprises only a small fraction of total
 1334 cells), we constructed a focused *tumour-niche* subset to increase sensitivity and interpretability. Con-
 1335 cretely, we retained all tumour cells and subsampled an equal number of T cells, B cells, fibroblasts,
 1336 and endothelial cells from the same specimens to form a minimal viable tumour microenvironment.
 1337 We then reused the analysis pipeline described earlier with matched timepoints at 0 h, 8 h, and 24 h.
 1338 The only modification was to the ligand–receptor (LR) library: for pathway-specific probes, we tog-
 1339 gled custom LR pairs to mimic the presence or absence of a given ligand (e.g., EGFR) and quantified
 1340 the resulting changes in inferred communication and downstream dynamics. Marker definitions are
 1341 provided in Table 17, and a dot-plot confirming marker specificity and minimal cross-lineage leak-
 1342 age is shown in Fig. 8.

1343 D.9.1 TUMOUR PROGRESSION QUANTIFICATION USING HALLMARK GENE SETS
1344

1345 There is no single, universally accepted definition of tumour progression. Clinical assessments typ-
 1346 ically use lesion size, extent of metastasis, and histopathology. While we observe distinct cellular
 1347 changes and invasion over our 24 h window, these measures are not applicable at single-cell reso-
 1348 lution. Instead, we construct an approximate *tumour differentiation* score based on the Hallmarks
 1349 of Cancer (Hanahan & Weinberg, 2011), using the MSigDB Hallmark gene sets (Liberzon et al.,
 2015).

Table 14: MFM hyperparameters.

Category	Hyperparameter	Setting / Notes
Velocity network (Flow Matching)		
Model	Architecture	Velocity MLP (hidden dim 64, depth 3, no dropout)
	Time conditioning	Sinusoidal embedding (dim 16), concatenated to inputs
Training	Epochs	500
	Batch size	128 (train), 2 048 (validation)
Optimizer		AdamW, lr = 10^{-3} , wd = 10^{-4} , $\beta = (0.9, 0.999)$
	Gradient clipping	1.0
GeoPath network (Riemannian Correction)		
Model	Architecture	GeoPath MLP (hidden dim 128, depth 3)
	Time conditioning	Concatenated directly (no embedding)
Training	Activation	SELU
	Epochs	100
Optimizer	Batch size	128 (train)
		Adam, lr = 10^{-4} , wd = 0.0, $\beta = (0.9, 0.999)$
Metric	Time sampler	Uniform
	Type	LAND
Metric	LAND parameters	$\gamma = 0.2$, $\rho = 10^{-3}$, $\alpha = 1.0$
	Max samples	4 096 for metric computation

Table 15: SF2M hyperparameters.

Category	Hyperparameter	Setting / Notes
Model	Architecture	Velocity and Score MLP (hidden dim 64, depth 3)
	Time conditioning	Sinusoidal embedding (dim 16), concat. to inputs
	Activation	SELU
Distribution parameters	σ_{bridge}	1.0
	σ_{sample}	1.0
Training schedule	Epochs	500
	Minibatch size	128 (train loader), 2 048 for validation batches
Optimizer	Optimizer	AdamW,
	lr	10^{-3}
	weight decay	10^{-4}
	betas	(0.9, 0.999)
Gradient clipping		1.0

For each hallmark, we compute a per-cell score as the *median* expression across its member genes (chosen over the mean for robustness to sparsity and outliers). The overall progression score is then the mean across the 20 retained hallmarks. The full hallmark definitions are available in MSigDB (Liberzon et al., 2015); the selected hallmarks, their gene counts, and five example genes each are listed in Table 18. Hallmarks not applicable to our tumour context (e.g., hormonal signalling for breast/prostate, long-term metabolic programs) were excluded.

As a baseline check, we verify that tumour cells exhibit coherent changes along the selected hallmarks over 0 h → 24 h; see Fig. 9.

1404
1405
1406 Table 16: UOT-FM hyperparameters.
1407
1408
1409
1410
1411
1412
1413

Category	Hyperparameter	Setting / Notes
Model	Architecture	Velocity MLP (hidden dim 64, depth 3, no dropout)
Training schedule	Time conditioning	Sinusoidal embedding (dim 16), concatenated to inputs
	Epochs	500
	Minibatch size	128 (train loader), 2 048 for validation batches
	Optimizer	AdamW, lr = 10^{-3} , weight decay 10^{-4} , betas (0.9, 0.999)
Optimal transport	Convergence tol.	10^{-9} (relative and absolute)
	Marginal reg.	1.0

Cell Type	Positive Markers
Differentiated AT1	RTKN2, AGER
AT1	CLDN18
Tumour (AT2)	SFTP, LAMP3, SCGB3A2
Mucous Epithelial	DNAH12, AZGP1
Endothelial	SEMA3G
Low IEG Endothelial	CDH5
Alveolar Capillary Endothelial	EDNRB, RPRML
Lymphatic Vein Endothelial	LYVE1, SELE, VWF
Fibroblasts	COL1A2, PDGFRA
Smooth Muscle Fibroblasts	ACTA2, LGR6
Fibroblast Subset	DCN
Pericytes	CSPG4
Megakaryocytes	PPBP, PF4
Erythrocytes	ALAS2
Lymphocytes	CCL21A
Cycling	TOP2A
Neutrophils	S100A9, RETNLG
Basophils & Mast cells	MCPT8, MS4A2
Macrophages	MARCO
Monocytes	LY6I
DC 1 and 2	CLEC9A, XCR1, C1QA, SIGLECH
DC 3	FSCN1, IL12B
NK cell like	NCR1, EOMES, TBX21
ILC	RORA, RORC, IL2RA
Adaptive T cells	FOXP3, CD4, CD8A
B cells	CD79A

1439
1440 Table 17: Curated panel of positive marker genes used for per-cell scoring and assignment in the
1441 lung cancer dataset.
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

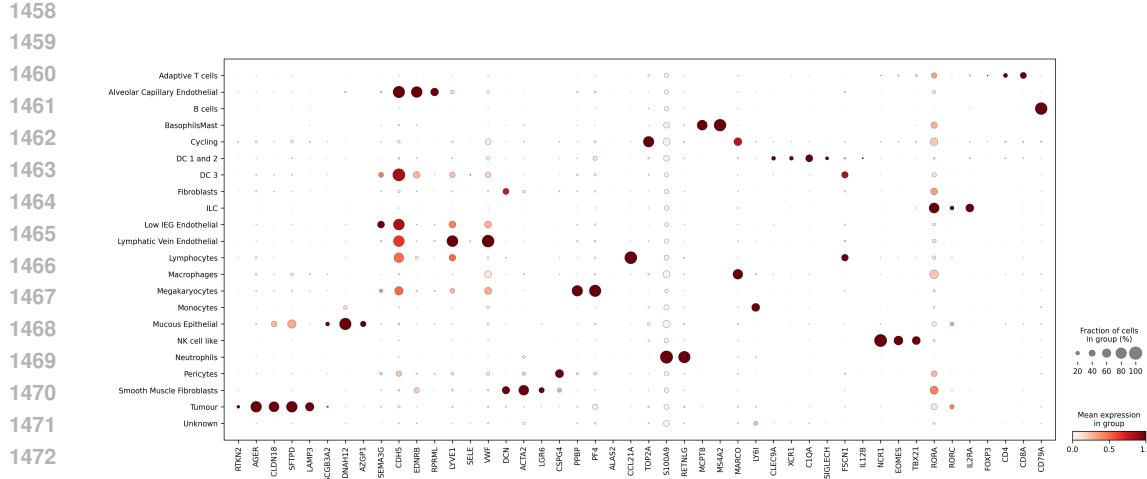


Figure 8: Dot-plot validation of curated marker genes across annotated cell types. Each column corresponds to a marker gene and each row to a cell-type label. Dot size encodes the fraction of cells expressing that gene, while color intensity represents its standardized expression level.

Table 18: **Hallmark gene sets used for trajectory summarization.** We list each set’s size and five randomly sampled member genes.

Gene set	# Genes	Random gene examples (5)
Angiogenesis	36	TIMP1, POSTN, VTN, THBD, NRP1
Apoptosis	161	ERBB2, IL1B, DPYD, NEDD9, MADD
DNA Repair	150	GTF2B, RAE1, ADCY6, POLA2, TAF1C
E2F Targets	200	MCM7, PCNA, MCM4, RFC2, GINS1
Epithelial–Mesenchymal Transition	200	SPP1, GPX7, LOX, THBS1, SLC6A8
G2M Checkpoint	200	RBM14, AMD1, CDC27, UCK2, NDC80
Glycolysis	200	SPAG4, PKP2, SLC25A13, PRPS1, ZNF292
Hypoxia	200	S100A4, CSRP2, DTNA, PIM1, TPST2
KRAS Signaling v1	200	FSHB, YPEL1, BARD1, SLC6A3, ATP6V1B1
KRAS Signaling v2	200	CIDEA, KIF5C, LAT2, PDCD1LG2, PIGR
MYC Targets v1	200	RAD23B, USP1, NAP1L1, NDUFAB1, SNRPA1
MYC Targets v2	58	PRMT3, AIMP2, SRM, EXOSC5, SUPV3L1
Myogenesis	200	EIF4A2, PDE4DIP, ANKRD2, EPHB3, ATP6AP1
Notch Signaling	32	SKP1, MAML2, HES1, FBXW11, DTX1
Oxidative Phosphorylation	200	NDUFS8, VDAC1, UQCRC, NDUFB3, NDUFB2
p53 Pathway	200	TNNI1, SLC35D1, BTG1, FDXR, JAG2
Peroxisome	104	IDH2, FIS1, EPHX2, SLC23A2, SLC25A4
Reactive Oxygen Species Pathway	49	PRNP, OXSR1, SOD1, PDLIM1, TXN
TNF α Signaling via NF κ B	200	DUSP2, CEBPB, OLR1, CCL20, IL1A
Xenobiotic Metabolism	200	SSR3, HACL1, ARPP19, AHCY, GSR

D.10 COMPUTATIONAL AND MEMORY COSTS

Complexity of the full OT solver. We solve Equation (5) with a custom conditional-gradient (Frank–Wolfe) solver detailed in Section D.4. Let n_0 and n_1 be the numbers of cells in the two snapshots and K the number of ligand–receptor (LR) pairs (interaction channels).

Each Frank–Wolfe iteration consists of two main steps:

1. **Gradient computation.** This yields a per-iteration cost

$$\mathcal{O}(K n_0 n_1 (n_0 + n_1)).$$

since it requires performing the matrix multiplication of $C_1^{(r)} \Gamma$ and $(C_1^{(r)} \Gamma) (C_2^{(r)})^\top$ for each channel r .

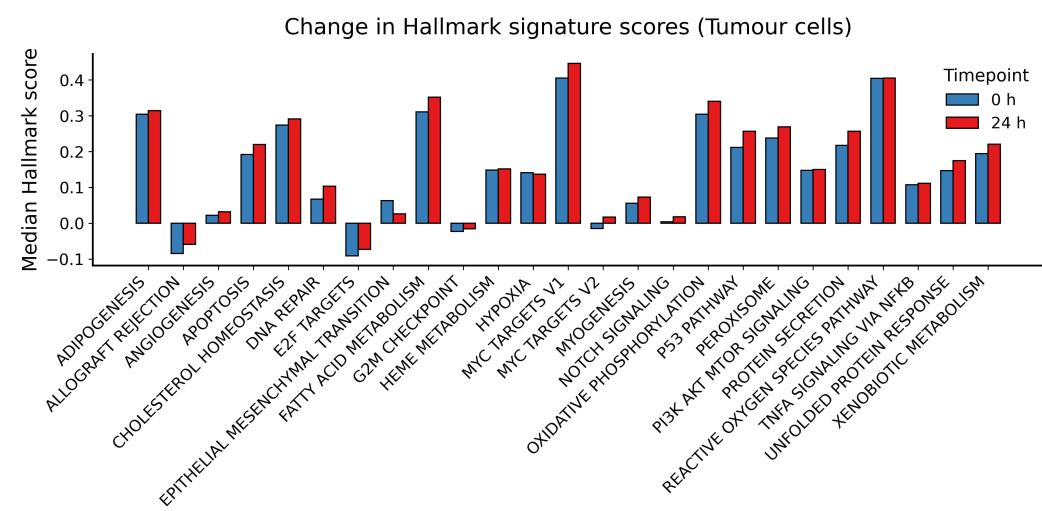


Figure 9: **Hallmark changes.** Changes in our dataset over 24 h following combined KRAS and MYC signalling across the 20 selected Hallmark gene sets.

Table 19: **Wall-clock runtime (in seconds) of IADOT, decomposing into the OT part and the Flow matching part.**

	Lung Tumour	V1 Light	Dendritic Stimulus
OT [s]	211.3	107.0	4.3
FM [s]	189.1	186.0	24.2

2. **Linear OT subproblem.** Given the linearized objective, we solve a linear OT problem over $\Pi(a, b)$ using POT’s (Flamary et al., 2024) existing OT routine. Its complexity is

$$\mathcal{O}(\text{cost}_{\text{OT}}(n_0, n_1)),$$

(e.g. cubic in n for a network-simplex LP, or $\mathcal{O}(T_{\text{Sinkhorn}} n_0 n_1)$ for entropic OT).

If T_{CG} denotes the number of Frank–Wolfe iterations required to reach the desired tolerance, the total complexity of the IADOT OT stage is

$$\mathcal{O}\left(T_{\text{CG}} [K n_0 n_1 (n_0 + n_1) + \text{cost}_{\text{OT}}(n_0, n_1)]\right).$$

Wall-clock runtimes. We report the wall-clock runtimes (seconds) in Table 19, decomposing it into the OT part (finding the coupling Γ^*) and the Flow matching part (fitting the velocity model).

Memory footprint of the OT stage. The dominant memory costs come from: (i) the coupling $\Gamma \in \mathbb{R}^{n_0 \times n_1}$, (ii) the feature cost matrix $C \in \mathbb{R}^{n_0 \times n_1}$, (iii) the multi-channel structure tensors $C_1 \in \mathbb{R}^{n_0 \times n_0 \times K}$ and $C_2 \in \mathbb{R}^{n_1 \times n_1 \times K}$ (corresponding to the CCI tensors $G^{(0)}$ and $G^{(1)}$), and (iv) a small number of auxiliary matrices of size $n_0 \times n_1$ (e.g. constC , $B(\Gamma)$, and the gradient). Crucially, we never construct the full tensor of pairwise structure discrepancies in Equation (5). Instead, the structure term is implemented through the matrix products $C_1^{(r)} \Gamma (C_2^{(r)})^\top$. As a result, the memory complexity of the OT solver scales as

$$\mathcal{O}(K(n_0^2 + n_1^2) + n_0 n_1),$$

Hence it is quadratic in the number of cells per snapshot and linear in the number of LR pairs K . In comparison, a feature-only OT solver ($\alpha = 0$) needs C and Γ , with memory $\mathcal{O}(n_0 n_1)$.

Using IADOT with large-scale datasets. While the computational and memory cost remained reasonable across the datasets we used, for very large datasets it can be mitigated using standard scalable techniques that are orthogonal to our formulation:

- adding entropic regularization (Peyré & Cuturi, 2019) and using Sinkhorn-type solvers, which make the problem easier to optimize and reduce memory at the price of a small, controllable bias
- employing mini-batch optimization (Fatras et al., 2021), where the CCI prior is estimated from couplings computed on minibatches instead of the whole dataset
- constructing metacells, with details provided in Section D.3. Using metacells reduces the effective sample size. We evaluate this variant in Section 5.5.

These strategies preserve the form of the IADOT prior while substantially improving scalability for large scale datasets.

D.11 EXTENDING IADOT TO THE UNBALANCED SETTING

As mentioned in Section 4.1, we can take into account cell proliferation or apoptosis between snapshots with an *unbalanced* OT formulation. Instead of searching the coupling over the space $\Pi(a, b)$ of couplings with marginals a and b , we instead optimize an objective that relaxes this hard constraint. More specifically, we adopt a two-step procedure:

Step 1: inferring non-uniform marginals via unbalanced OT. We first ignore CCI structure and solve an unbalanced feature-only problem

$$\Gamma^u \in \arg \min_{\Gamma \in \mathbb{R}_{\geq 0}^{n_0 \times n_1}} \left\{ \langle \Gamma, C \rangle + \lambda_0 \text{KL}(\Gamma \mathbf{1} \| a) + \lambda_1 \text{KL}(\Gamma^\top \mathbf{1} \| b) \right\}, \quad (14)$$

where $\text{KL}(\cdot \| \cdot)$ denotes the Kullback–Leibler divergence and $\mathbf{1}$ is the all-ones vector. The penalty terms are soft constraints on the marginals of Γ , allowing deviations from (a, b) that capture net cell birth or death between snapshots. From the optimal coupling Γ^u we extract the reweighted marginals

$$\tilde{a} = \Gamma^u \mathbf{1}, \quad \tilde{b} = (\Gamma^u)^\top \mathbf{1},$$

which are renormalized to sum to one. We note that this step is α - agnostic, which allows to keep fixed marginals across the different α . In practice, we use a solver based on L-BFGS-B implemented in (Flamary et al., 2024).

Step 2: FGW with frozen unbalanced marginals. In a second step, we fix \tilde{a} and \tilde{b} and solve the interaction-aware FGW problem of Section 4.1 with these new marginals:

$$\min_{\Gamma \in \Pi(\tilde{a}, \tilde{b})} (1 - \alpha) \mathcal{F}(\Gamma) + \alpha \mathcal{S}(\Gamma), \quad (15)$$

where \mathcal{F} and \mathcal{S} are defined as in Equation (5). This second step preserves the multi-LR-pair CCI structure while respecting the unequal total mass at the two snapshots inferred in Step 1.

D.12 COMBINING IADOT WITH MFM

Here we detail how our CCI-based prior can be combined with Metric Flow Matching (MFM) (Kapusniak et al., 2024). MFM generalizes Conditional Flow Matching by learning interpolants $x_{t,\eta}$ that approximate geodesics of a data–dependent Riemannian metric g on the ambient space. Given a coupling q between p_0 and p_1 , MFM first trains a network $\phi_{t,\eta}$ to minimize the geodesic loss

$$L_g(\eta) = \mathbb{E}_{(x_0, x_1) \sim q, t} [\dot{x}_{t,\eta}^\top G(x_{t,\eta}; \mathcal{D}) \dot{x}_{t,\eta}],$$

where $G(\cdot; \mathcal{D})$ is the coordinate representation of the metric and $x_{t,\eta} = (1 - t)x_0 + tx_1 + t(1 - t)\phi_{t,\eta}(x_0, x_1)$.

In our experiments in Section 5.3 and Section E.4, we instantiate MFM with the LAND metric g_{LAND} (Arvanitidis et al., 2016).

Once the interpolant parameters η^* have been fitted via L_g , we can optimize the following MFM velocity–regression loss using the coupling found with IADOT:

1620

1621

$$\mathcal{L}_{\text{MFM}}(\theta, \eta) = \mathbb{E}_{\substack{(X, Y) \sim \Pi \\ t \sim \text{Unif}[0, 1]}} \left[\| v_\theta(Z_{t, \eta}, t) - u_{t, \eta}(Z_{t, \eta} \mid X, Y) \|_{g(Z_{t, \eta})}^2 \right] \quad (16)$$

1624

1625

$$= \mathbb{E}_{\substack{(X, Y) \sim \Pi \\ t \sim \text{Unif}[0, 1]}} \left[\| v_\theta(Z_{t, \eta}, t) - \dot{x}_{t, \eta}(X, Y) \|_{g(Z_{t, \eta})}^2 \right], \quad (17)$$

1626

where $\|\cdot\|_{g(Z_{t, \eta})}$ is the norm induced by the Riemannian metric g at $Z_{t, \eta}$, i.e.

1628

1629

$$\|w\|_{g(Z_{t, \eta})}^2 = w^\top G(Z_{t, \eta}; \mathcal{D}) w.$$

1630

1631

D.13 COMBINING IADOT WITH SF2M

1633

1634

Here we describe how our CCI-based prior can be combined with SF2M (Tong et al., 2023) to obtain a Schrödinger-bridge-type dynamics that uses the IADOT coupling.

1636

1637

SF2M learns a drift v_θ and score s_θ by regressing to the conditional flow and score of a mixture of Brownian bridges between endpoints (x_0, x_1) . For a single bridge with diffusion σ , the conditional marginal at time $t \in (0, 1)$ is

1639

1640

$$p_t(x \mid x_0, x_1) = \mathcal{N}(x; \mu_t(x_0, x_1), \sigma^2 t(1-t)I_d), \quad \mu_t(x_0, x_1) = (1-t)x_0 + tx_1,$$

1641

with closed-form drift $u_t^\circ(x \mid x_0, x_1)$ and score $\nabla_x \log p_t(x \mid x_0, x_1)$ given in (Tong et al., 2024).

1642

1643

To combine SF2M with IADOT, we simply instantiate the endpoint coupling with the IADOT coupling Π instead of the entropic OT plan. Training triples are then sampled as

1644

1645

$$t \sim \text{Unif}[0, 1], \quad (X, Y) \sim \Pi, \quad Z_t \sim p_t(\cdot \mid X, Y),$$

1646

1647

where $p_t(\cdot \mid X, Y)$ is the Brownian-bridge marginal above with $\mu_t(X, Y) = (1-t)X + tY$. The SF2M objective specialized to IADOT is

1648

1649

1650

$$\mathcal{L}_{\text{SF2M}}(\theta) = \mathbb{E}_{\substack{(X, Y) \sim \Pi \\ t \sim \text{Unif}[0, 1] \\ Z_t \sim p_t(\cdot \mid X, Y)}} \left[\| v_\theta(t, Z_t) - u_t^\circ(Z_t \mid X, Y) \|_2^2 \right] \quad (18)$$

1651

1652

$$+ \lambda(t)^2 \| s_\theta(t, Z_t) - \nabla_z \log p_t(Z_t \mid X, Y) \|_2^2 \right], \quad (19)$$

1653

1654

where $\lambda(t) = 2\sqrt{t(1-t)}/\sigma$ is the time-dependent weighting.

1655

E Additional results

1656

E.1 STIMULUS DATASETS

1657

1658

We reproduce the experimental setup described in Section 5.2 and Section 5.3 with the macrophase stimulus-response dataset (Section C.6). We report the results in Figure 10 and Table 20, which are consistent with the findings on the other datasets.

1663

1664

E.2 SENSITIVITY OF COUPLINGS TO CATALOG EDITS

1665

1666

1667

The experiment presented in Section 5.4 involved perturbing the LR catalog by removing specific LR pairs. In Table 21, we show how the coupling changes, by computing the fraction of source cells whose target argmax differs between "active" vs. "inactive" LR libraries for each pathway.

1668

1669

E.3 COMPARING CELL INTERACTION TYPES

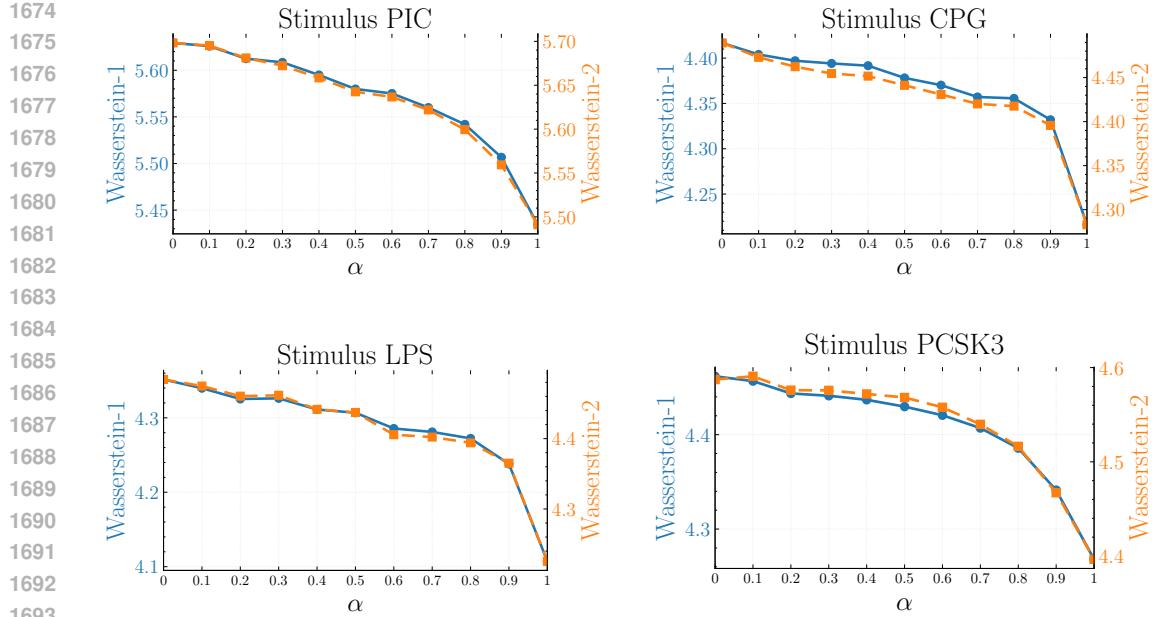
1670

1671

1672

1673

We further examined how different classes of molecular interactions influence the resulting transport couplings. Using our automated selection procedure (Section D.6), we identified a top-ranking set of 10 ligand-receptor pairs for each of the datasets. We contrasted this against a matched set of 10 canonical long-range soluble cytokines and growth factors: (CXCL12-CXCR4, VEGFA-KDR,

Figure 10: **Interpolation** results for the macrophage stimulus datasets.Table 20: **Interpolation error for continuous time dynamics (lower is better).** IADOT with varying structure weight α vs. baselines for the macrophase stimulus datasets. We report mean \pm std over 5 runs.

Method	α	Stimulus PIC		Stimulus CPG		Stimulus LPS		Stimulus PCSK3	
		W_1	W_2	W_1	W_2	W_1	W_2	W_1	W_2
TrajectoryNet	—	5.628 (0.055)	5.930 (0.049)	5.361 (0.085)	5.826 (0.080)	5.087 (0.109)	5.589 (0.078)	5.033 (0.051)	5.434 (0.051)
DSB	—	5.796 (0.574)	5.833 (0.571)	4.500 (0.128)	4.594 (0.114)	4.685 (0.533)	4.815 (0.528)	4.648 (0.328)	4.749 (0.324)
OT-CFM	—	5.544 (0.038)	5.614 (0.036)	4.430 (0.019)	4.512 (0.020)	4.434 (0.052)	4.582 (0.058)	4.423 (0.020)	4.551 (0.018)
OT-MFM	—	5.501 (0.020)	5.566 (0.022)	4.464 (0.013)	4.556 (0.015)	4.440 (0.033)	4.592 (0.038)	4.384 (0.008)	4.505 (0.008)
UOT-FM	—	5.414 (0.027)	5.492 (0.024)	4.572 (0.033)	4.733 (0.045)	4.570 (0.122)	4.785 (0.143)	4.332 (0.007)	4.483 (0.008)
SF2M	—	6.132 (0.059)	6.212 (0.058)	4.959 (0.044)	5.053 (0.043)	4.930 (0.051)	5.057 (0.054)	5.048 (0.031)	5.155 (0.032)
VGFM	—	9.796 (0.658)	9.863 (0.683)	8.242 (0.192)	8.362 (0.205)	8.026 (0.135)	8.158 (0.1469)	8.205 (0.110)	8.297 (0.110)
MIOFlow	—	9.365 (0.382)	9.440 (0.404)	16899 (22805)	22007 (31264)	7.807 (0.038)	7.927 (0.034)	8.179 (0.150)	8.281 (0.166)
Moscat	—	7.213 (0.000)	7.244 (0.000)	6.983 (0.000)	7.087 (0.000)	7.663 (0.000)	7.786 (0.000)	6.734 (0.000)	6.832 (0.000)
IADOT +SF2M	0.5	6.109 (0.081)	6.193 (0.086)	4.948 (0.066)	5.027 (0.063)	4.846 (0.046)	4.973 (0.046)	4.960 (0.058)	5.084 (0.058)
	1	6.110 (0.062)	6.197 (0.063)	4.951 (0.056)	5.037 (0.058)	4.874 (0.068)	5.016 (0.065)	5.016 (0.029)	5.173 (0.039)
IADOT +MFM	0.5	5.483 (0.012)	5.544 (0.014)	4.448 (0.007)	4.527 (0.006)	4.442 (0.035)	4.589 (0.041)	4.386 (0.016)	4.520 (0.013)
	1	5.376 (0.021)	5.440 (0.022)	4.460 (0.079)	4.543 (0.082)	4.477 (0.050)	4.641 (0.062)	4.329 (0.023)	4.507 (0.029)
IADOT +UOT-FM	0.5	5.355 (0.021)	5.451 (0.018)	4.544 (0.033)	4.700 (0.035)	4.489 (0.054)	4.692 (0.061)	4.343 (0.025)	4.522 (0.033)
	1	5.346 (0.024)	5.487 (0.038)	4.547 (0.035)	4.755 (0.040)	4.489 (0.038)	4.716 (0.041)	4.325 (0.038)	4.531 (0.038)
IADOT +CFM	0.5	5.490 (0.018)	5.555 (0.019)	4.427 (0.021)	4.502 (0.025)	4.380 (0.021)	4.517 (0.021)	4.396 (0.023)	4.531 (0.019)
	1	5.446 (0.018)	5.512 (0.016)	4.440 (0.041)	4.518 (0.044)	4.431 (0.126)	4.577 (0.141)	4.352 (0.020)	4.530 (0.033)

Table 21: **Coupling changes (argmax) at $\alpha = 1.0$.** Fraction of source cells whose target argmax differs between “active” vs. “inactive” LR libraries for each pathway; $N=2195$ source cells. Targeted pathways (EGFR/ALK/MET) show large shifts, while cardio–renal controls (RAAS, Vasopressin, Natriuretic) show little or moderate effect, as expected.

Pathway / System	Coupling changed (count / N)	Percent
EGFR (targeted)	2071/2195	94.35%
ALK (targeted)	2164/2195	98.59%
MET (targeted)	2154/2195	98.13%
RAAS (control)	0/2195	0.00%
Vasopressin (control)	0/2195	0.00%
Natriuretic (control)	1582/2195	72.07%

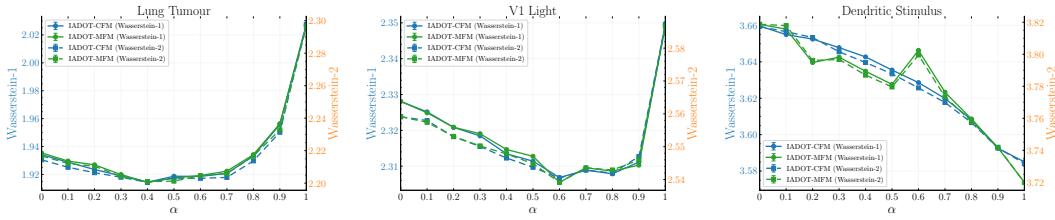


Figure 11: **Incorporating the CCI prior with MFM.** We plot the W_1 and W_2 distances between the interpolated and empirical t_1 snapshots as α varies.

1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
CCL5-CCR5, TGFB1-TGFB2, IL6-IL6R, EGF-EGFR, TNF-TNFRSF1A, IGF1-IGF1R, CSF1-CSF1R, IFNG-IFNLR1). As shown in Table 22, the cytokine pairs exhibit slightly higher Wasserstein (W_1 and W_2) distances compared to the results obtained previously with our selection procedure. This suggests that the specific interaction modes we keep have more informative topological constraints on the transport map than generic diffusive signaling, effectively recovering structure-aware couplings that reflect the physical tissue architecture.

Table 22: **Wasserstein distances for pure structural alignment ($\alpha = 1$).** Comparison of interpolation performance using generic Long Range priors versus our automated selection procedure. Lower values indicate better alignment.

Dataset	Interaction Prior	$W_1 \downarrow$	$W_2 \downarrow$
V1 Light	Long range	2.42	2.63
	Dataset-specific	2.35	2.59
Immune	Long range	3.58	3.73
	Dataset-specific	3.59	3.73
Lung Cancer	Long range	2.10	2.33
	Dataset-specific	2.02	2.30

E.4 COMBINING IADOT WITH OTHER PRIORS

1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
A key advantage of IADOT is its modularity: the CCI-derived prior only depends on the CCI tensors ($G^{(0)}, G^{(1)}$) and on a coupling Γ , and is therefore largely orthogonal to how Γ is obtained. As a consequence, the CCI prior can be combined with a wide range of existing priors or architectural choices for trajectory inference. Here, we illustrate this flexibility by extending IADOT to two settings: (i) unbalanced OT, which explicitly accounts for cell birth and death between snapshots, and (ii) metric flow matching, which replaces the standard Euclidean flow-matching objective with a geometry-aware variant. Details on both of these implementations can be found in Section D.11 and Section D.12.

We reproduce the experiment in Section 5.2 with these IADOT variants, and report the results in Figure 11 and Figure 12. We notice the following:

- $\alpha > 0$ remains optimal. For all datasets and the two IADOT variants, the best W_1/W_2 values occur at a non-zero structure weight α , mirroring the behavior observed in Section 5.2.
- Complementary to other priors. The fact that $\alpha > 0$ remains optimal shows that adding the CCI prior on top of MFM or UOT-FM yields consistent improvements over the corresponding feature-only baselines, highlighting that IADOT’s gains are not tied to a specific OT or flow-matching objective, but rather come from the biological prior.

E.5 SENSITIVITY ANALYSIS ON THE LR EXPRESSIONS

1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
In this section, we study the sensitivity of IADOT to measurement noise in the LR expressions. We inject this noise in LR genes expressions by adding zero-mean Gaussian noise to the gene

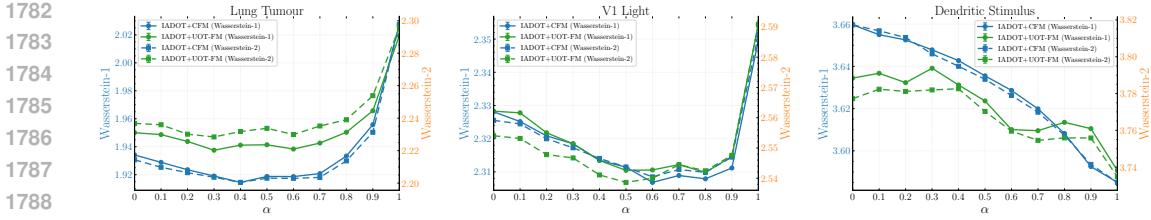


Figure 12: **Incorporating the CCI prior with UOT-CFM.** We plot the W_1 and W_2 distances between the interpolated and empirical t_1 snapshots as α varies.

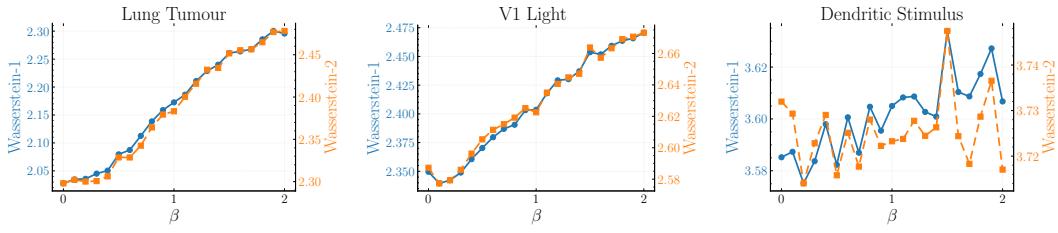


Figure 13: **Perturbation of the LR expressions with Gaussian noise.** We plot the W_1 and W_2 distances between the interpolated and empirical t_1 snapshots as the scaling factor of the noise β increases.

expressions before applying the Hill transform and clipping below by 0, i.e. we define $\tilde{x}_{cg} = \max(0, x_{cg} + \epsilon_{cg})$ where $\epsilon_{cg} \sim \mathcal{N}(0, \sigma_g^2)$. σ_g^2 denotes a gene-specific noise variance, defined as $\sigma_g = \beta\hat{\sigma}_g$, with $\hat{\sigma}_g$ the empirical standard deviation of $\{x_{cg} \mid c \in [n]\}$ (to take into account per-gene variance) and β a scaling factor. From these perturbed expressions, we compute the activations \tilde{s}_{cg} and we construct the CCI tensors with the entries $\tilde{q}_{i \rightarrow j}^{(pk)} = \tilde{s}_{i\ell_k} \tilde{s}_{jr_k}$ and obtain the couplings by solving Equation (5).

We report the results in Figure 13, where we sweep β for different values across the interval $[0, 2]$, with $\alpha = 1$.

As the noise scale β increases, both W_1 and W_2 gradually deteriorate across all three datasets. This non-zero sensitivity is expected and desirable: if the CCI prior was irrelevant, corrupting the LR expressions would leave the interpolation error unchanged. Instead, adding noise worsens alignment, showing the benefits of the prior. Furthermore, the performance is relatively robust to small level of noises for the Lung tumour and V1 Light datasets. Interestingly, for the V1 Light dataset, we see that $\beta \in \{0.1, 0.2\}$ improves the results upon $\beta = 0$, which we attribute to a small regularization / denoising effect. Adding a small amount of centered Gaussian noise before the Hill transform and clipping makes low-intensity ligand or receptor expressions become zero while leaving strongly expressed pairs essentially unchanged. The results are noisier for the Dendritic Stimulus dataset, which we attribute to the smaller size of the dataset.

E.6 SENSITIVITY WITH RESPECT TO K_g AND h_g

In this section, we conduct a sensitivity analysis on the hyperparameters K_g and h_g , used to define the interaction scores in Section 4.1 as $s_{cg} = x_{cg}^{h_g} / (x_{cg}^{h_g} + K_g^{h_g})$. We consider different values of the percentile level $p \in \{80, 90, 99\}$ (with $K_g = Q_g(p)$ denoting the p -th percentile of $\{x_{cg} \mid c \in [n]\}$) and $h_g \in \{1, 2, 4\}$, for $\alpha = 0.5$. We report the interpolation results in Figure 14. We observe that the performance is largely insensitive to the specific choice of these parameters. This stability justifies the use of standard default values (90th percentile and $h_g = 1$) across our experiments without the need for extensive per-dataset tuning.

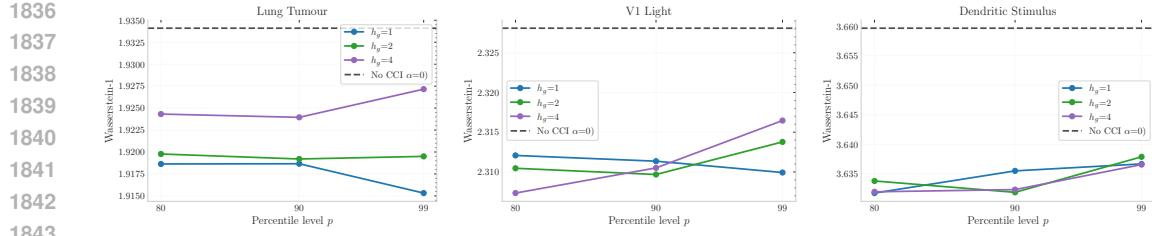


Figure 14: **Sensitivity analysis on the hyperparameters of the Hill transform.** We plot the W_1 distances between the interpolated and empirical t_1 snapshots for different values of K_g and h_g .

E.7 PATH CURVATURE ANALYSIS

To empirically demonstrate that `IADOT` learns non-linear interaction effects, we measure the average path length ratio (displacement divided by path length) of the inferred trajectories:

$$S(z_0, v_\theta) = \frac{\|z_1 - z_0\|_2}{\int_0^1 \|v_\theta(z_t, t)\|_2 dt} \quad (20)$$

where $z_t = z_0 + \int_0^t v_\theta(z_t, t) dt$ for $t \in [0, 1]$ and z_0 denotes the initial point.

A ratio of 1.0 indicates a straight line, while values < 1.0 indicate curvature.

As shown in Table 23, increasing the interaction weight α leads to significantly higher curvature (lower ratios), confirming that incorporating interactions prevents the model from simply learning independent straight lines.

Table 23: Path length ratio comparison across different datasets.

α	Lung Tumour	Dendritic Stimulus	V1 Light
0	0.974 ± 0.001	0.982 ± 0.002	0.954 ± 0.004
0.5	0.952 ± 0.002	0.979 ± 0.003	0.907 ± 0.011
1	0.862 ± 0.011	0.969 ± 0.005	0.635 ± 0.009

1890 **F Theoretical analysis**
 1891

1892 **F.1 SYNTHETIC SETUP**
 1893

1894 In this section, we provide a theoretical guarantee for the synthetic setup of Section C.1.
 1895

1896 **Theorem 1.** *Let \mathcal{D}_0 and \mathcal{D}_1 be the source and target datasets defined by the synthetic clusters, and
 1897 let $G^{(0)}, G^{(1)}$ be the associated directed interaction tensors.*

1898 *Consider the two candidate couplings:*

1899 1. Γ_{GT} : *The transport plan corresponding to the true translation vectors (preserving cluster
 1900 identity).*
 1901

1902 2. Γ_{FO} : *The transport plan corresponding to the feature-only map.*
 1903

1904 *As in Section D.5, define the (unnormalized) feature and structure gaps between these two couplings
 1905 as*

$$1906 \Delta\mathcal{F} := \mathcal{F}(\Gamma_{FO}) - \mathcal{F}(\Gamma_{GT}), \quad \Delta\mathcal{S} := \mathcal{S}(\Gamma_{FO}) - \mathcal{S}(\Gamma_{GT}),$$

1907 *and the corresponding normalized feature and structure terms*

$$1908 \tilde{\mathcal{F}}(\Gamma) := \frac{\mathcal{F}(\Gamma)}{|\Delta\mathcal{F}|}, \quad \tilde{\mathcal{S}}(\Gamma) := \frac{\mathcal{S}(\Gamma)}{|\Delta\mathcal{S}|}.$$

1910 *Let*

$$1911 J(\Gamma, \alpha) = (1 - \alpha) \tilde{\mathcal{F}}(\Gamma) + \alpha \tilde{\mathcal{S}}(\Gamma)$$

1912 *be the normalized FGW objective function.*

1913 *Let N denote the size of each cluster. Then, for sufficiently large N , there exists a critical threshold
 1914 $\alpha^* \in (0, 1)$ such that for all $\alpha > \alpha^*$, the ground truth coupling strictly minimizes the objective
 1915 relative to the feature-only alternative: $J(\Gamma_{GT}, \alpha) < J(\Gamma_{FO}, \alpha)$.*

1916 *Proof.* Let the source measure be μ and the target measure be ν , with the variance of the Normal
 1917 distributions set to $\sigma^2 = 0.1$. The centroids are located at $\mu_0 = (-2, 2), \mu_1 = (0, 2), \mu_2 = (2, 2)$
 1918 and $\mu'_0 = (2, -2), \mu'_1 = (0, -2), \mu'_2 = (-2, -2)$. The interaction tensors $G^{(0)}$ and $G^{(1)}$ encode
 1919 directed edges from the middle cluster ($k = 1$) to the left ($k = 0$) via Channel 1, and to the right
 1920 ($k = 2$) via Channel 2. In what follows, we first compute the *unnormalized* feature and structure
 1921 costs. We then incorporate the normalization scheme of Section D.5 in the final threshold derivation.
 1922

1923 **1. Analysis of the feature cost \mathcal{F}**

1924 For the ground truth coupling Γ_{GT} , each cluster k maps to its true image μ'_k . The cost is the mean
 1925 squared norm of the translation vectors $v_0 = (4, -4), v_1 = (0, -4)$, and $v_2 = (-4, -4)$:
 1926

$$1927 \mathcal{F}(\Gamma_{GT}) = \frac{1}{3} \sum_{k=0}^2 \|v_k\|^2 = \frac{1}{3}(32 + 16 + 32) = \frac{80}{3}. \quad (21)$$

1928 For the feature-only coupling Γ_{FO} , \mathcal{S}_0 maps to \mathcal{S}'_2 , \mathcal{S}_1 to \mathcal{S}'_1 , and \mathcal{S}_2 to \mathcal{S}'_0 . In the finite sample
 1929 regime with N points, the optimal transport cost between two empirical Gaussian distributions with
 1930 identical covariance matrices converges to the squared Euclidean distance between their means. We
 1931 denote the finite-sample deviation by δ_N :
 1932

$$1933 \mathcal{F}(\Gamma_{FO}) = \frac{1}{3} (\|\mu_0 - \mu'_2\|^2 + \|\mu_1 - \mu'_1\|^2 + \|\mu_2 - \mu'_0\|^2) + \delta_N. \quad (22)$$

1934 *Hence:*

$$1935 \mathcal{F}(\Gamma_{FO}) = 16 + \delta_N. \quad (23)$$

1936 The term δ_N represents the error between the empirical measures and their population counterparts.
 1937 For distributions in dimension $d = 2$, this error decays at a rate of $\delta_N = O(N^{-1/2})$ (Fournier &
 1938 Guillen, 2015). Provided N is sufficiently large, standard OT ($\alpha = 0$) prefers the incorrect mapping
 1939 since $16 < \frac{80}{3}$.

1940 **2. Analysis of structure cost \mathcal{S}**

1944 The structure cost is the Gromov-Wasserstein cost:
 1945
 1946
 1947

$$\mathcal{S}(\Gamma) = \sum_{i,k} \sum_{j,l} \|G_{ik}^{(0)} - G_{jl}^{(1)}\|^2 \Gamma_{ij} \Gamma_{kl}. \quad (24)$$

1948 Since Γ_{GT} maps every source cluster k to the target cluster with the same index k , and $G^{(1)}$ is
 1949 defined to preserve the index-based structure of $G^{(0)}$, we have:
 1950

$$\mathcal{S}(\Gamma_{GT}) = 0. \quad (25)$$

1952 For Γ_{FO} , the mapping permutes indices as $\pi(0) = 2, \pi(1) = 1, \pi(2) = 0$. We evaluate the cost for
 1953 the two active interactions in $G^{(0)}$:
 1954

- **Edge 1 → 0 (Channel 1):** The source relation is $[1, 0]^\top$ and the target relation (from $\pi(1) = 1$ to $\pi(0) = 2$) is Channel 2 ($[0, 1]^\top$), which yields a squared difference of 2 with mass weight $1/9$.
- **Edge 1 → 2 (Channel 2):** The source relation is $[0, 1]^\top$ and the target relation (from $\pi(1) = 1$ to $\pi(2) = 0$) is Channel 1 ($[1, 0]^\top$), which yields a squared difference of 2 with mass weight $1/9$.

1962 The total structure cost is:
 1963
 1964

$$\mathcal{S}(\Gamma_{FO}) = \frac{1}{9} \times 2 + \frac{1}{9} \times 2 = \frac{4}{9}. \quad (26)$$

3. Threshold derivation with normalization

1967 We now incorporate the normalization scheme of Section D.5. Define the (unnormalized) feature
 1968 and structure gaps between the two couplings as
 1969

$$\Delta\mathcal{F} := \mathcal{F}(\Gamma_{FO}) - \mathcal{F}(\Gamma_{GT}), \quad \Delta\mathcal{S} := \mathcal{S}(\Gamma_{FO}) - \mathcal{S}(\Gamma_{GT}). \quad (27)$$

1971 From the computations above,

$$\Delta\mathcal{F} = 16 + \delta_N - \frac{80}{3}, \quad \Delta\mathcal{S} = \frac{4}{9}. \quad (28)$$

1972 For sufficiently large N , we have $\mathcal{F}(\Gamma_{GT}) > \mathcal{F}(\Gamma_{FO})$, so $|\Delta\mathcal{F}| = \mathcal{F}(\Gamma_{GT}) - \mathcal{F}(\Gamma_{FO}) > 0$ and
 1973 the normalization is well-defined. Normalizing, we get:
 1974

$$\tilde{\mathcal{F}}(\Gamma) = \frac{\mathcal{F}(\Gamma)}{|\Delta\mathcal{F}|}, \quad \tilde{\mathcal{S}}(\Gamma) = \frac{\mathcal{S}(\Gamma)}{\Delta\mathcal{S}}. \quad (29)$$

1975 The normalized FGW objective can therefore be written as
 1976

$$J(\Gamma, \alpha) = (1 - \alpha)\tilde{\mathcal{F}}(\Gamma) + \alpha\tilde{\mathcal{S}}(\Gamma). \quad (30)$$

1977 For the two couplings of interest, we obtain
 1978

$$\tilde{\mathcal{S}}(\Gamma_{GT}) = \frac{\mathcal{S}(\Gamma_{GT})}{\Delta\mathcal{S}} = 0, \quad \tilde{\mathcal{S}}(\Gamma_{FO}) = \frac{\mathcal{S}(\Gamma_{FO})}{\Delta\mathcal{S}} = 1, \quad (31)$$

1979 and
 1980

$$\tilde{\mathcal{F}}(\Gamma_{GT}) - \tilde{\mathcal{F}}(\Gamma_{FO}) = \frac{\mathcal{F}(\Gamma_{GT}) - \mathcal{F}(\Gamma_{FO})}{|\Delta\mathcal{F}|} = \frac{|\Delta\mathcal{F}|}{|\Delta\mathcal{F}|} = 1. \quad (32)$$

1981 We seek α such that $J(\Gamma_{GT}, \alpha) < J(\Gamma_{FO}, \alpha)$ under this normalization, i.e.
 1982

$$(1 - \alpha)\tilde{\mathcal{F}}(\Gamma_{GT}) < (1 - \alpha)\tilde{\mathcal{F}}(\Gamma_{FO}) + \alpha. \quad (33)$$

1983 Using $\tilde{\mathcal{F}}(\Gamma_{GT}) - \tilde{\mathcal{F}}(\Gamma_{FO}) = 1$, this inequality becomes
 1984

$$(1 - \alpha) < \alpha \iff \alpha > \frac{1}{2}. \quad (34)$$

1985 Hence, under the normalization of Section D.5 and in the asymptotic regime, a critical threshold
 1986 $\alpha^* = 1/2$ exists above which the ground truth coupling strictly improves the normalized objective
 1987 relative to the feature-only alternative: $J(\Gamma_{GT}, \alpha) < J(\Gamma_{FO}, \alpha)$ for all $\alpha > 1/2$. \square
 1988

1998 **Remarks.** In theory, $\alpha^* = 0.5$ comes from an idealized analysis of the normalized objective that
 1999 only compares the feature-only and structure-only couplings in the *population limit*. Finite-sample
 2000 effects, approximate normalization, and the existence of many 'almost-correct' couplings break the
 2001 symmetry of the idealized setting and make the optimal α slightly bigger than 0.5.

2002 Second, Theorem 1 shows that the ground truth coupling Γ_{GT} is better than Γ_{FO} at $\alpha = 1$. However,
 2003 it is not the *only* one. The interaction tensors $G^{(0)}$ and $G^{(1)}$ are constant for all points within a
 2004 cluster. Therefore, the structure cost $\mathcal{S}(\Gamma)$ depends only on which clusters are matched, not on how
 2005 individual points are mapped within them.

2007 Any coupling that correctly maps source clusters to their corresponding target clusters yields a struc-
 2008 ture cost of 0. This includes the ground truth coupling Γ_{GT} , but also any coupling that correctly
 2009 matches clusters while randomly permuting points inside them. This explains the results observed
 2010 in Figure 3 : at $\alpha = 1$, the solver returns a solution that is structurally perfect but fails to recover the
 2011 exact point-to-point correspondence.

2012 **F.2 DYNAMIC INTERPRETATION OF IADOT**

2014 We provide a dynamic viewpoint on IADOT, showing that it can be seen as the solution of a joint
 2015 static-dynamic energy minimization problem combining kinetic energy in expression space and a
 2016 structure-preserving term.

2017 As before, let

$$\Pi(a, b) := \left\{ \Gamma \in \mathbb{R}_+^{n_0 \times n_1} : \Gamma \mathbf{1}_{n_1} = a, \Gamma^\top \mathbf{1}_{n_0} = b \right\}.$$

2019 We further consider the common choice of feature cost

$$C_{ij} = \|x_i - y_j\|^2, \quad 1 \leq i \leq n_0, 1 \leq j \leq n_1. \quad (35)$$

2022 **Admissible processes for a fixed coupling.** Let $\Gamma \in \Pi(a, b)$ and define the associated joint law on
 2023 endpoints

$$\Pi_\Gamma := \sum_{i=1}^{n_0} \sum_{j=1}^{n_1} \Gamma_{ij} \delta_{(x_i, y_j)}. \quad (36)$$

2025 In the balanced case $\sum_{i,j} \Gamma_{ij} = 1$, so Π_Γ is a probability measure with marginals ρ_0, ρ_1 .

2028 We consider continuous-time processes $(X_t)_{t \in [0, 1]}$ taking values in \mathbb{R}^d and satisfying:

- X . has almost surely absolutely continuous paths
- the joint law of its endpoints is $(X_0, X_1) \sim \Pi_\Gamma$

2033 We write $\mathcal{A}(\Pi_\Gamma)$ for the class of all such processes. For any $X \in \mathcal{A}(\Pi_\Gamma)$, define the kinetic energy
 2034 as:

$$\mathcal{K}(X) := \mathbb{E} \left[\int_0^1 \|\dot{X}_t\|^2 dt \right]. \quad (37)$$

2038 The following lemma is standard but we include it for completeness.

2039 **Lemma 1.** Let $x, y \in \mathbb{R}^d$ and let $\gamma : [0, 1] \rightarrow \mathbb{R}^d$ be absolutely continuous with $\gamma(0) = x, \gamma(1) = y$.
 2040 Then

$$\int_0^1 \|\dot{\gamma}(t)\|^2 dt \geq \|y - x\|^2, \quad (38)$$

2043 with equality if and only if $\gamma(t) = (1 - t)x + ty$ for all $t \in [0, 1]$.

2045 *Proof.* By Cauchy–Schwarz inequality,

$$\left\| \int_0^1 \dot{\gamma}(t) dt \right\|^2 \leq \int_0^1 \|\dot{\gamma}(t)\|^2 dt, \quad (39)$$

2049 with equality if and only if $\dot{\gamma}(t)$ is constant in t . Since $\gamma(1) - \gamma(0) = y - x$, this yields

$$\|y - x\|^2 = \left\| \int_0^1 \dot{\gamma}(t) dt \right\|^2 \leq \int_0^1 \|\dot{\gamma}(t)\|^2 dt, \quad (40)$$

2052 Equality holds if and only if $\dot{\gamma}(t) = y - x$, i.e. $\gamma(t) = (1 - t)x + ty$. \square
 2053

2054 **Proposition 1.** *Let $\Gamma \in \Pi(a, b)$ and Π_Γ be as above. Consider the admissible class $\mathcal{A}(\Pi_\Gamma)$ and the
 2055 kinetic energy \mathcal{K} . Then:*

2056 1. *The energy $\mathcal{K}(X_\cdot)$ is minimized over $\mathcal{A}(\Pi_\Gamma)$ by the process*

$$2058 \quad X_t^{\text{lin}} := (1 - t)X + tY, \quad (X, Y) \sim \Pi_\Gamma. \quad (41)$$

2059 2. *The minimal value of the kinetic energy is*

$$2061 \quad \inf_{X_\cdot \in \mathcal{A}(\Pi_\Gamma)} \mathcal{K}(X_\cdot) = \mathbb{E}_{(X, Y) \sim \Pi_\Gamma} [\|X - Y\|^2] = \sum_{i,j} \Gamma_{ij} C_{ij}. \quad (42)$$

2064 *Proof.* Any $X_\cdot \in \mathcal{A}(\Pi_\Gamma)$ satisfies $(X_0, X_1) \sim \Pi_\Gamma$. Condition on the endpoints:

$$2066 \quad \mathcal{K}(X_\cdot) = \mathbb{E}_{(X, Y) \sim \Pi_\Gamma} \left[\mathbb{E} \left[\int_0^1 \|\dot{X}_t\|^2 dt \mid (X_0, X_1) = (X, Y) \right] \right]. \quad (43)$$

2068 For each fixed pair $(X, Y) = (x, y)$, Lemma 1 shows that the conditional energy is minimized by
 2069 the straight-line path $t \mapsto (1 - t)x + ty$, with minimal value $\|x - y\|^2$. Thus the global minimizer
 2070 over $\mathcal{A}(\Pi_\Gamma)$ is the straight-line process X_t^{lin} , and

$$2071 \quad \inf_{X_\cdot \in \mathcal{A}(\Pi_\Gamma)} \mathcal{K}(X_\cdot) = \mathbb{E}_{(X, Y) \sim \Pi_\Gamma} [\|X - Y\|^2] = \sum_{i,j} \Gamma_{ij} \|x_i - y_j\|^2 = \sum_{i,j} \Gamma_{ij} C_{ij}. \quad (44)$$

2074 \square

2075 **Joint static–dynamic energy and reduction to FGW.** We can view IADOT as minimizing over
 2076 both couplings and dynamics the joint energy functional
 2077

$$2078 \quad \mathcal{E}_\alpha(\Gamma, X_\cdot) := (1 - \alpha) \mathcal{K}(X_\cdot) + \alpha S(\Gamma), \quad (45)$$

2079 subject to $\Gamma \in \Pi(a, b)$ and $X_\cdot \in \mathcal{A}(\Pi_\Gamma)$.

2080 **Proposition 2.** *Fix $\alpha \in [0, 1]$. Consider the optimization problem*

$$2082 \quad \inf_{\Gamma \in \Pi(a, b)} \inf_{X_\cdot \in \mathcal{A}(\Pi_\Gamma)} \mathcal{E}_\alpha(\Gamma, X_\cdot) = \inf_{\Gamma \in \Pi(a, b)} \inf_{X_\cdot \in \mathcal{A}(\Pi_\Gamma)} [(1 - \alpha) \mathcal{K}(X_\cdot) + \alpha S(\Gamma)]. \quad (46)$$

2083 *Then:*

2084 1. *For any fixed Γ , the inner infimum over $X_\cdot \in \mathcal{A}(\Pi_\Gamma)$ is attained by the straight-line process
 2085 $X_t^{\text{lin}} = (1 - t)X + tY$, $(X, Y) \sim \Pi_\Gamma$, and*

$$2088 \quad \inf_{X_\cdot \in \mathcal{A}(\Pi_\Gamma)} \mathcal{E}_\alpha(\Gamma, X_\cdot) = (1 - \alpha) \sum_{i,j} \Gamma_{ij} C_{ij} + \alpha S(\Gamma). \quad (47)$$

2089 2. *Consequently, the joint static–dynamic problem reduces to the purely static FGW problem*

$$2092 \quad \inf_{\Gamma \in \Pi(a, b)} \inf_{X_\cdot \in \mathcal{A}(\Pi_\Gamma)} \mathcal{E}_\alpha(\Gamma, X_\cdot) = \inf_{\Gamma \in \Pi(a, b)} [(1 - \alpha) \langle \Gamma, C \rangle_F + \alpha S(\Gamma)], \quad (48)$$

2093 *whose minimizers are exactly the FGW-optimal couplings used by IADOT.*

2094 *Proof.* Point (1) follows directly from Proposition 1: for any Γ ,

$$2097 \quad \inf_{X_\cdot \in \mathcal{A}(\Pi_\Gamma)} \mathcal{E}_\alpha(\Gamma, X_\cdot) = (1 - \alpha) \inf_{X_\cdot \in \mathcal{A}(\Pi_\Gamma)} \mathcal{K}(X_\cdot) + \alpha S(\Gamma) = (1 - \alpha) \sum_{i,j} \Gamma_{ij} C_{ij} + \alpha S(\Gamma). \quad (49)$$

2099 Taking the infimum over $\Gamma \in \Pi(a, b)$ yields (2), which coincides with the static FGW objective. \square
 2100

2101 Intuitively, Proposition 2 shows that IADOT does not use linear interpolations between matched
 2102 cells as a heuristic, but as the *unique* minimal-action choice once the coupling Γ^* is fixed. The
 2103 static FGW step therefore selects an interaction-aware coupling that trades off feature displacement
 2104 and CCI preservation, and the subsequent dynamic step realizes this coupling by approximating the
 2105 lowest-kinetic-energy flow in expression space. When $\alpha = 0$, this recovers the classical OT–CFM
 (Tong et al., 2024)/ Benamou–Brenier (Benamou & Brenier, 2000) interpolation.

2106 F.3 CONNECTION TO THE VELOCITY FIELD LEARNT WITH CFM.
21072108 In practice, `IADOT` does not explicitly construct the process X_t^{lin} but instead uses CFM to learn a
2109 time-dependent vector field v_θ that generates the same probability path.
21102111 In the infinite-capacity and optimization limit, the minimizer v^* of \mathcal{L}_{CFM} coincides with the velocity
2112 field of X_t^{lin} constructed in Section F.2, in the sense that
2113

2114
$$v^*(z, t) = \mathbb{E}[Y - X \mid Z_t = z],$$

2115 and the ODE
2116

2117
$$\dot{z}_t = v^*(z_t, t)$$

2118 generates exactly the probability path $\{\rho_t\}_{t \in [0, 1]}$ induced by Γ^* .
2119

2120 We can relate v^* to the kinetic energy of the straight-line process, following a similar technique as
2121 in (Lipman et al., 2024). For a time-dependent vector field $w : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ that generates
2122 $\{\rho_t\}$, define its kinetic energy along this path by
2123

2124
$$\mathcal{K}_{\text{Eul}}(w) := \int_0^1 \mathbb{E}_{Z_t \sim \rho_t} \left[\|w(Z_t, t)\|_2^2 \right] dt.$$

2125

2126 Using the formula of v^* above and Jensen's inequality, we obtain
2127

2128
$$\begin{aligned} \mathcal{K}_{\text{Eul}}(v^*) &= \int_0^1 \mathbb{E} \left[\left\| \mathbb{E}[Y - X \mid Z_t] \right\|_2^2 \right] dt \\ 2129 &\leq \int_0^1 \mathbb{E} \left[\mathbb{E}[\|Y - X\|_2^2 \mid Z_t] \right] dt \\ 2130 &= \int_0^1 \mathbb{E}[\|Y - X\|_2^2] dt \\ 2131 &= \mathbb{E}_{(X, Y) \sim \Pi} [\|Y - X\|_2^2]. \end{aligned}$$

2132

2133 By Proposition 1 we have
2134

2135
$$\mathbb{E}_{(X, Y) \sim \Pi_{\Gamma^*}} [\|Y - X\|_2^2] = \sum_{i,j} \Gamma_{ij}^* C_{ij} = K(X^{\text{lin}}),$$

2136

2137 Hence
2138

2139
$$\mathcal{K}_{\text{Eul}}(v^*) \leq \sum_{i,j} \Gamma_{ij}^* C_{ij} = K(X^{\text{lin}}).$$

2140

2141 In other words, for a fixed coupling Γ , the feature term
2142

2143
$$F(\Gamma) = \langle \Gamma, C \rangle_F = \sum_{i,j} \Gamma_{ij} C_{ij}$$

2144

2145 provides an explicit upper bound on the kinetic energy of the velocity field recovered by CFM from
2146 the corresponding straight-line dynamics. Combined with the joint static-dynamic formulation in
2147 Equation (45), this shows that the `IADOT` objective
2148

2149
$$(1 - \alpha) \langle \Gamma, C \rangle_F + \alpha S(\Gamma)$$

2150

2151 can be viewed as selecting a coupling that balances CCI preservation with a surrogate upper bound
2152 on the kinetic energy of the flow that CFM learns from that coupling.
21532154 G LLM usage
21552156 We used large language models (LLMs) to assist with improving the clarity of writing and refining
2157 the formatting of tables and figures. LLMs were not used for research ideation, experimental design,
2158 analysis, or any substantive contributions that would merit authorship.
2159