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 kNN 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., 2025). 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),

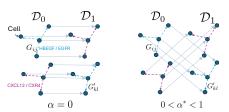


Figure 1: Augmenting feature distances with typed interactions improves alignment between snapshots \mathcal{D}_0 and \mathcal{D}_1 and encourages transport plans consistent with communication patterns.

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 Conditional Flow Matching (Lipman et al., 2022).

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 (ablating 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 $\psi: \mathbb{R}^d \times \mathbb{R}_+ \to \mathbb{R}^d$ 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.,

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

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

The static Optimal Transport problem. Optimal Transport (OT) provides a geometric framework for comparing probability distributions, enforcing the principle of least-action. In the context of cell trajectory inference, this principle assumes that the change between snapshots arises from the smallest rearrangement of cellular states consistent with biology: states evolve smoothly in expression 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 unit of mass from x_i to y_j , the discrete Kantorovich formulation seeks a coupling

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

where $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius dot product. The optimal coupling Γ^* therefore represents the most efficient mapping from a geometric standpoint, as it is defined based on the gene expression profiles. However, this formulation is limited as it only considers the cost of displacing individual points, ignoring any structural information encoded within the distributions.

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

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

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

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

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

3 RELATED WORK

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

Table 1: Comparison of trajectory methods. Legend: ✓ supported, ~ partial, X not supported.

			In-silico			
Method	Dynamic	Trajectories	Perturbation	Structure-Aware	scRNA Data Sufficient	Reference
PAGA (Scanpy)	Х	Х	Х	Х	✓	(Wolf et al., 2019)
Waddington-OT	✓	~	X	X	✓	(Schiebinger et al., 2019)
SCOT	X	X	X	~	X	(Demetci et al., 2022a)
CellOT	X	X	✓	X	X	(Bunne et al., 2023a)
OT-CFM	✓	✓	X	X	✓	(Tong et al., 2024)
TrajectoryNet	/	/	X	X	✓	(Tong et al., 2020)
Schrödinger Bridge	/	/	X	X	✓	(Hong et al., 2025)
scVelo	✓	X	X	X	✓	(Bergen et al., 2020)
IADOT (ours)	√	√	√	✓	✓	

Structure-aware alignments. Gromov–Wasserstein (GW) compares samples via their intrinsic geometry, and Fused GW (FGW) optimizes a joint feature+structure objective (Vayer et al., 2020a). In single-cell settings, GW/FGW pipelines (e.g., SCOT for multi-omics) typically rely on *undirected k*NN graphs that capture generic topology but lack *communication semantics* (Demetci et al., 2022b). We instead inject a *directed, typed* prior derived from ligand–receptor (LR) expression into an FGW objective. This encourages alignments that preserve signaling context and allows to probe the effect of specific LR interactions on the inferred dynamics. Orthogonal lines of work infer *directionality* from spliced/unspliced counts and propagate it on *k*NN graphs (Bergen et al., 2020). CellRank further combines velocity with transcriptomic similarity to estimate fate probabilities (La Manno et al., 2018; Bergen et al., 2020; Lange et al., 2022). Spatial OT approaches (e.g., SpaOTsc) instead exploit *physical proximity* to couple cells and infer possible communication, but their structure term remains geometric rather than *typed* signaling (Cang & Nie, 2020). In contrast, IADOT integrates *directed, typed* ligand–receptor pairs into a multi-dimensional FGW objective, and links the static coupling to *learned continuous dynamics* via conditional flow matching.

4 IADOT: INTERACTION-AWARE OPTIMAL TRANSPORT

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

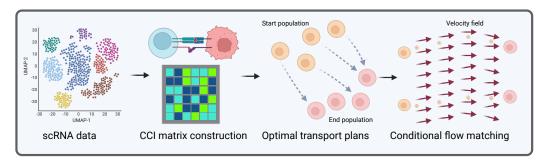


Figure 2: **Overview of IADOT.** From a ligand–receptor catalogue, we build directed, multi LR-pair CCI matrices. A structure-aware optimal transport problem balances feature similarity with interaction structure to produce a cross-snapshot coupling, used to train a time-continuous vector field learned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

4.1 INTERACTION-AWARE TRANSPORT VIA MULTI LR-PAIR FGW

Modeling cell-cell interactions from scRNA. Given a ligand-receptor (LR) catalog $\mathcal{P} = \{(l_k, r_k) \mid k \in [K]\}$ of K ligand-receptor pairs and a dataset of n cells, our aim is to construct 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 g=0.9 quantile of nonzero values in g=0.9 where g=0.9 is the normalized expression of gene g=0.9 for cell g=0.9 and exponent g=0.9 and g=0.9 and

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 $q_{i \to j}^{(p_k)} = s_{i\ell_k} \, s_{jr_k}$, 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 $G_{ijk} = q_{i \to j}^{(p_k)}$. 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 :

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

The structure term $S(\Gamma)$ 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 $S(\Gamma)$. Therefore, we optimize it using a conditional-gradient solver (Braun et al., 2022), with more details provided in Appendix D.4.

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

4.2 LEARNING CONTINUOUS DYNAMICS VIA CONDITIONAL FLOW MATCHING

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

Probability path. Let ρ_0 and ρ_1 denote the empirical distributions defined by \mathcal{D}_0 and \mathcal{D}_1 , respectively. Let $\Gamma^\star \in \mathbb{R}^{n_0 \times n_1}_+$ be the optimal coupling from Equation (5), with normalization constant $M = \sum_{i,j} \Gamma^\star_{ij}$. We define a joint distribution Π on $\mathcal{D}_0 \times \mathcal{D}_1$ by $\Pi = \sum_{i=1}^{n_0} \sum_{j=1}^{n_1} \frac{\Gamma^\star_{ij}}{M} \, \delta_{(x_i,y_j)}$, 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 . For $t \in [0,1]$, we then consider the affine interpolation $Z_t = (1-t)X + tY$, with $(X,Y) \sim \Pi$, and 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 and ρ_1 are the endpoints of this path.

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

We train v_{θ} by regressing to this drift along the path, yielding the following CFM objective:

$$\mathcal{L}_{\mathrm{CFM}}(\theta) = \mathbb{E}_{\substack{(X,Y) \sim \Pi \\ t \sim \mathrm{Unif}[0,1]}} \left[\left\| v_{\theta} \left(Z_{t}, t \right) - u_{t} \left(Z_{t} \mid X, Y \right) \right\|_{2}^{2} \right]$$
 (6)

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

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

5 EXPERIMENTS

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

5.1 SYNTHETIC SETUP

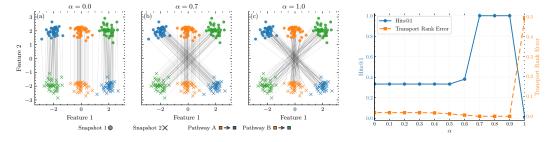


Figure 3: Structure-aware coupling recovers the ground-truth transport map. We show representative couplings (*left*) and matching metrics (*right*). Feature-only OT (α =0) ignores structure and misaligns clusters, structure-only (α =1) distorts geometry within the interaction types. A balanced trade-off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

Setup. We consider two 2D snapshots, each composed of three clusters. The second snapshot is obtained by translating each cluster by a distinct vector, inducing a known one-to-one ground-truth

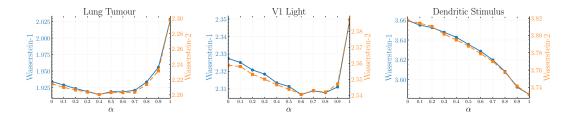


Figure 4: **Interpolation error.** We plot the W_1 and W_2 distances between the interpolated and empirical t_1 snapshots as α varies. Optimal performance occurs at dataset-specific $\alpha^* > 0$.

transport. We define an interaction structure with two types: the middle cluster points to the left ($Type\ I$) and to the right ($Type\ 2$), mirrored in the target snapshot (see Appendix C.1 for more details). We then obtain a coupling for each $\alpha \in \{0,0.1,\ldots,1\}$ by solving the FGW problem defined in Equation (5) with the ground-truth interaction structures.

Results. Representative couplings across α are shown in Figure 3 (left). With $\alpha=0$ (feature-only), the interaction structure is ignored and clusters are misaligned; with $\alpha=1$ (structure-only), interaction types are satisfied but geometry is distorted. An intermediate setting ($\alpha\approx0.7$) preserves the directed relations while maintaining within-interaction geometry. We quantify these observations by computing Hits@1, the fraction of source samples whose top-weighted target equals the ground-truth match, and $Transport\ Rank\ Error\ (TRE)$, the average fraction of targets ranked above the ground-truth match. Figure 3 (right) shows that Hits@1 peaks and TRE is minimized at mid-range α , indicating that a balanced mix of features and structure gives the most faithful transport map.

5.2 Cross-Snapshot Interpolation from Structure-Aware Couplings

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

Setup. We build CCI tensors by selecting dataset-specific ligand-receptor pairs via an automated procedure that accounts for stability of expression levels across snapshots (cf. Appendix D.6 for more 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 LR catalog $\mathcal P$ and hyperparameter $\alpha \in \{0,0.1,\ldots,1.0\}$, we obtain a coupling $\Gamma(\alpha,\mathcal P)$ by solving the OT problem defined in Equation (5). We define the marginal at t_1 by affine interpolation and 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} observed at t_1 , computing the Wasserstein-1 and Wasserstein-2 distances $W_1(\rho_{t_1}(\alpha,\mathcal P),\rho_{t_1})$ and $W_2(\rho_{t_1}(\alpha,\mathcal P),\rho_{t_1})$.

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

5.3 Cross-Snapshot Trajectory Inference via Flow Matching

Setup. Having shown that incorporating structure yields better couplings for these datasets, we now verify whether it also improves continuous-time dynamics learnt with IADOT. Starting from an optimal coupling Γ , we fit a time-conditional vector field v_{θ} using the conditional Flow matching loss (Equation (6)). We integrate v_{θ} to transport cells observed at t_0 to the held-out time t_1 , and compare the transported distribution to the empirical snapshot at t_1 using Wasserstein-1 and -2 distances. We repeat this for $\alpha \in \{0, 0.5, 1\}$. As baselines, we compare with *TrajectoryNet* (Tong

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.

		V1 co	ortex	Dendritic	Stimulus	Lung	tumor
Method	α	W_1	W_2	W_1	W_2	W_1	W_2
TRAJECTORYNET DSB		$3.022 (0.061) \\ 3.819 (0.152)$				2.712 (0.090) 3.700 (0.116)	3.056 (0.099) 3.967 (0.102)
IADOT	0 0.5 1	2.381 (0.004)	2.618 (0.003)	3.679 (0.009)	3.835 (0.010)	1.993 (0.004) 1.989 (0.004) 2.057 (0.005)	2.272 (0.005)

et al., 2020) (Neural ODE dynamics) and *Diffusion Schrödinger Bridges (DSB)* (De Bortoli et al., 2021) (diffusion-based trajectory inference).

Results. Table 2 reports W_1 and W_2 at the held-out time t_1 (lower is better). Results indicate that the procedure used to fit the velocity field affects performance: conditional flow matching yields consistently lower errors than TrajectoryNet and DSB across datasets. Second, within IADOT, structure helps: settings with $\alpha>0$ outperform the feature-only case $(\alpha=0)$, with the best results at $\alpha\in\{0.5,1\}$. These findings align with Section 5.2, indicating that CCI structure benefits both static alignment and the learned continuous-time dynamics.

5.4 PROBING TRAJECTORY SENSITIVITY TO LIGAND-RECEPTOR CATALOG EDITS

Setup. The previous sections showed that incorporating interaction structure can improve trajectory inference. We now examine how modifying these structures influences the inferred trajectories. Specifically, we focus on the *Lung Tumor* dataset and construct alternative ligand–receptor catalogs in which signaling for selected pathways is removed. From these modified catalogs, we recompute the corresponding CCI tensors and resolve the OT problem (Equation (5)) to obtain couplings. These couplings define mappings from source to target cells, which we then summarize using the 20 Hallmarks of Cancer gene sets (see Appendix D.9.1). We evaluate differences relative to the baseline (unmodified catalog) over a 24h interpolation window¹.

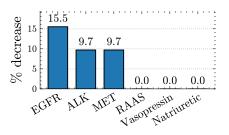


Figure 5: **In silico interventions.** Edits to the LR catalog (i.e. removing EGFR/ALK/MET interactions in the CCIs) reduce a Hallmark-based tumour progression proxy.

Results. Figure 5 shows the relative decrease in tumourassociated 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 per-

Table 3: Sensitivity analysis on the CCI construction.

Method	Tur	nor	Deno	lritic	Cor	tex
Method	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

¹Note this intervention edits the interaction prior used in the OT objective. We do not perturb gene expression or simulate drug pharmacodynamics.

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

5.6 Does structure always improve cross-snapshot alignment?

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

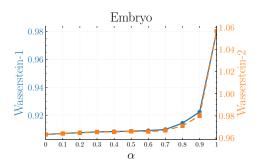


Figure 6: **Embryo dataset.** Incorporating CCIs does not improve performance over the feature-only baseline ($\alpha = 0$).

Results. We report the interpolation results for this dataset in Figure 6. In this setting, IADOT offers no additional gains over feature-only OT ($\alpha=0$), confirming that the CCI structure is not transferable across these days-long developmental intervals and thus becomes uninformative. This leads to an essentially flat curve with respect to α , with worse degradations at larger α . It also yields a practical guideline: when cross-snapshot interaction geometry does not persist over time, the structural term should be downweighted or omitted.

6 Discussion

CCI structure can improve cellular trajectory inference. Across synthetic and real scRNA–seq settings, we find that introducing typed, directed structure into the OT objective can lead to more faithful cross-snapshot couplings, and hence better continuous-time dynamics. Because structure is an explicit, editable prior, IADOT enables principled counterfactuals: we can quantify resulting shifts in inferred trajectories when perturbing pathway-specific LR libraries in the CCI construction. Finally, our ablation study shows that *biologically meaningful* interaction matrices are important and drive the observed performance gains.

Limitations. We stress that IADOT is not a silver bullet. Our approach assumes that key interaction structure is at least partly conserved between adjacent snapshots. In rapidly remodeling systems (e.g. embryo development), we have shown that incorporating the structure does not necesarily yield benefits. Furthermore, we evaluated IADOT on real scRNA-seq datasets from prior studies, but broader *scalability* to atlas level datasets is an interesting avenue for future work. Finally, future work might extend our OT formulation to an unbalanced setting (Séjourné et al., 2023) to explicitly account for cell proliferation and death over longer time scales.

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

REPRODUCIBILITY STATEMENT

We specify the OT objective in Section 4.1, and detail the multi-channel FGW solver as well as the normalization used to balance feature and structure terms in Appendices D.4 and D.5. Our continuous-time dynamics and Conditional Flow Matching objective are described in Section 4.2, with model architectures and all training hyperparameters listed in Appendix D.7. Datasets, sampling timepoints, and sizes are summarized in Appendix C, and the end-to-end preprocessing pipeline is documented in Appendix D.2. The construction of ligand—receptor catalogs is detailed in Appendix D.6. Baselines use authors' implementations with exact settings listed in Table 9 and Table 10. Software versions and key libraries are reported in Table 7. Code will be released upon acceptance.

REFERENCES

- Asif Adil, Vijay Kumar, Arif Tasleem Jan, and Mohammed Asger. Single-cell transcriptomics: Current methods and challenges in data acquisition and analysis. *Frontiers in Neuroscience*, 15, April 2021. ISSN 1662-453X. doi: 10.3389/fnins.2021.591122. URL http://dx.doi.org/10.3389/fnins.2021.591122.
- Prabhu S. Arunachalam, Madeleine K. D. Scott, Thomas Hagan, Chunfeng Li, Yupeng Feng, Florian Wimmers, Lilit Grigoryan, Meera Trisal, Venkata Viswanadh Edara, Lilin Lai, Sarah Esther Chang, Allan Feng, Shaurya Dhingra, Mihir Shah, Alexandra S. Lee, Sharon Chinthrajah, Sayantani B. Sindher, Vamsee Mallajosyula, Fei Gao, Natalia Sigal, Sangeeta Kowli, Sheena Gupta, Kathryn Pellegrini, Gregory Tharp, Sofia Maysel-Auslender, Sydney Hamilton, Hadj Aoued, Kevin Hrusovsky, Mark Roskey, Steven E. Bosinger, Holden T. Maecker, Scott D. Boyd, Mark M. Davis, Paul J. Utz, Mehul S. Suthar, Purvesh Khatri, Kari C. Nadeau, and Bali Pulendran. Systems vaccinology of the bnt162b2 mrna vaccine in humans. *Nature*, 596 (7872):410–416, July 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03791-x. URL http://dx.doi.org/10.1038/s41586-021-03791-x.
- Volker Bergen, Marius Lange, Stefan Peidli, F. Alexander Wolf, and Fabian J. Theis. Generalizing rna velocity to transient cell states through dynamical modeling. *Nature Biotechnology*, 38(12): 1408–1414, 2020. doi: 10.1038/s41587-020-0591-3.
- Gábor Braun, Alejandro Carderera, Cyrille W Combettes, Hamed Hassani, Amin Karbasi, Aryan Mokhtari, and Sebastian Pokutta. Conditional gradient methods. *arXiv preprint arXiv:2211.14103*, 2022.
- Robin Browaeys, Wouter Saelens, and Yvan Saeys. Nichenet: modeling intercellular communication by linking ligands to target genes. *Nature Methods*, 17(2):159–162, December 2019. ISSN 1548-7105. doi: 10.1038/s41592-019-0667-5. URL http://dx.doi.org/10.1038/s41592-019-0667-5.
- Charlotte Bunne, Stefan G. Stark, Gabriele Gut, Jacobo Sarabia del Castillo, Mitch Levesque, Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*, 20(11):1759–1768, September 2023a. ISSN 1548-7105. doi: 10.1038/s41592-023-01969-x. URL http://dx.doi.org/10.1038/s41592-023-01969-x.
- Charlotte Bunne, Stefan G Stark, Gabriele Gut, Jacobo Sarabia Del Castillo, Mitch Levesque, Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-cell perturbation responses using neural optimal transport. *Nature methods*, 20(11):1759–1768, 2023b.
- Charlotte Bunne, Geoffrey Schiebinger, Andreas Krause, Aviv Regev, and Marco Cuturi. Optimal transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1), August 2024. ISSN 2662-8449. doi: 10.1038/s43586-024-00334-2. URL http://dx.doi.org/10.1038/s43586-024-00334-2.

- J. Gray Camp, Farhath Badsha, Marta Florio, Sabina Kanton, Tobias Gerber, Michaela Wilsch-Bräuninger, Eric Lewitus, Alex Sykes, Wulf Hevers, Madeline Lancaster, Juergen A. Knoblich, Robert Lachmann, Svante Pääbo, Wieland B. Huttner, and Barbara Treutlein. Human cerebral organoids recapitulate gene expression programs of fetal neocortex development. *Proceedings of the National Academy of Sciences*, 112(51):15672–15677, December 2015. ISSN 1091-6490. doi: 10.1073/pnas.1520760112. URL http://dx.doi.org/10.1073/pnas.1520760112.
- Zixuan Cang and Qing Nie. Inferring spatial and signaling relationships between cells from single-cell transcriptomic data. *Nature Communications*, 11(1):2084, 2020. doi: 10.1038/s41467-020-15968-5.
- Zixuan Cang, Yanxiang Zhao, Axel A. Almet, Adam Stabell, Raul Ramos, Maksim V. Plikus, Scott X. Atwood, and Qing Nie. Screening cell–cell communication in spatial transcriptomics via collective optimal transport. *Nature Methods*, 20(2):218–228, January 2023. ISSN 1548-7105. doi: 10.1038/s41592-022-01728-4. URL http://dx.doi.org/10.1038/s41592-022-01728-4.
- Junha Cha and Insuk Lee. Single-cell network biology for resolving cellular heterogeneity in human diseases. *Experimental & Molecular Medicine*, 52(11):1798–1808, November 2020. ISSN 2092-6413. doi: 10.1038/s12276-020-00528-0. URL http://dx.doi.org/10.1038/s12276-020-00528-0.
- Gan Chen, Chao Ren, Yao Xiao, Yujing Wang, Renqi Yao, Quan Wang, Guoxing You, Mingzi Lu, Shaoduo Yan, Xiaoyong Zhang, Jun Zhang, Yongming Yao, and Hong Zhou. Time-resolved single-cell transcriptomics reveals the landscape and dynamics of hepatic cells in sepsis-induced 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.
- Zhanlin Chen, William C. King, Aheyon Hwang, Mark Gerstein, and Jing Zhang. Deepvelo: Single-cell transcriptomic deep velocity field learning with neural ordinary differential equations. *Science Advances*, 8(48), December 2022. ISSN 2375-2548. doi: 10.1126/sciadv.abq3745. URL http://dx.doi.org/10.1126/sciadv.abq3745.
- Marco Cuturi, Michal Klein, and Pierre Ablin. Monge, Bregman and occam: Interpretable optimal transport in high-dimensions with feature-sparse maps. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 6671–6682. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/cuturi23a.html.
- Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in neural information processing systems*, 34:17695–17709, 2021.
- Pinar Demetci, Rebecca Santorella, Björn Sandstede, William Stafford Noble, and Ritambhara Singh. Scot: Single-cell multi-omics alignment with optimal transport. *Journal of Computational Biology*, 29(1):3–18, January 2022a. ISSN 1557-8666. doi: 10.1089/cmb.2021.0446. URL http://dx.doi.org/10.1089/cmb.2021.0446.
- Pınar Demetci, Rebecca Santorella, Björn Sandstede, William Stafford Noble, and Ritambhara Singh. Scot: Single-cell multi-omics alignment with optimal transport. *Journal of Computational Biology*, 29(1):3–18, 2022b. doi: 10.1089/cmb.2021.0446.
- Daniel Dimitrov, Philipp Sven Lars Schäfer, Elias Farr, Pablo Rodriguez-Mier, Sebastian Lobentanzer, Pau Badia-i Mompel, Aurelien Dugourd, Jovan Tanevski, Ricardo Omar Ramirez Flores, and Julio Saez-Rodriguez. Liana+ provides an all-in-one framework for cell-cell communication inference. *Nature Cell Biology*, 26(9):1613–1622, September 2024a. ISSN 1476-4679. doi: 10.1038/s41556-024-01469-w. URL http://dx.doi.org/10.1038/s41556-024-01469-w.
- Daniel Dimitrov, Philipp Sven Lars Schäfer, Elias Farr, Pablo Rodriguez-Mier, Sebastian Lobentanzer, Pau Badia-i Mompel, Aurelien Dugourd, Jovan Tanevski, Ricardo Omar Ramirez Flores,

- and Julio Saez-Rodriguez. Liana+ provides an all-in-one framework for cell–cell communication inference. *Nature Cell Biology*, 26(9):1613–1622, 2024b.
- Kalliopi Domvri, Paul Zarogoulidis, Kaid Darwiche, Robert F. Browning, Qiang Li, J. Francis Turner, Ioannis Kioumis, Dionysios Spyratos, Konstantinos Porpodis, Antonis Papaiwannou, Theodora Tsiouda, Lutz Freitag, and Konstantinos Zarogoulidis. Molecular targeted drugs and biomarkers in nsclc, the evolving role of individualized therapy. *Journal of Cancer*, 4(9):736–754, 2013. ISSN 1837-9664. doi: 10.7150/jca.7734. URL http://dx.doi.org/10.7150/jca.7734.
- Mirjana Efremova, Miquel Vento-Tormo, Sarah A. Teichmann, and Roser Vento-Tormo. Cellphonedb: inferring cell-cell communication from combined expression of multi-subunit ligand-receptor complexes. *Nature Protocols*, 15(4):1484–1506, February 2020. ISSN 1750-2799. doi: 10.1038/s41596-020-0292-x. URL http://dx.doi.org/10.1038/s41596-020-0292-x.
- Rossin Erbe, Genevieve Stein-O'Brien, and Elana J. Fertig. Transcriptomic forecasting with neural ordinary differential equations. *Patterns*, 4(8):100793, August 2023. ISSN 2666-3899. doi: 10.1016/j.patter.2023.100793. URL http://dx.doi.org/10.1016/j.patter.2023.100793.
- Nona Farbehi, Ralph Patrick, Aude Dorison, Munira Xaymardan, Vaibhao Janbandhu, Katharina Wystub-Lis, Joshua WK Ho, Robert E Nordon, and Richard P Harvey. Single-cell expression profiling reveals dynamic flux of cardiac stromal, vascular and immune cells in health and injury. *eLife*, 8, March 2019. ISSN 2050-084X. doi: 10.7554/elife.43882. URL http://dx.doi.org/10.7554/eLife.43882.
- Rémi Flamary, Cédric Vincent-Cuaz, Nicolas Courty, Alexandre Gramfort, Oleksii Kachaiev, Huy Quang Tran, Laurène David, Clément Bonet, Nathan Cassereau, Théo Gnassounou, Eloi Tanguy, Julie Delon, Antoine Collas, Sonia Mazelet, Laetitia Chapel, Tanguy Kerdoncuff, Xizheng Yu, Matthew Feickert, Paul Krzakala, Tianlin Liu, and Eduardo Fernandes Montesuma. Pot python optimal transport (version 0.9.5), 2024. URL https://github.com/PythonOT/POT.
- Aden Forrow and Geoffrey Schiebinger. Lineageot is a unified framework for lineage tracing and trajectory 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.
- Katherine C. Goldfarbmuren, Nathan D. Jackson, Satria P. Sajuthi, Nathan Dyjack, Katie S. Li, Cydney L. Rios, Elizabeth G. Plender, Michael T. Montgomery, Jamie L. Everman, Preston E. Bratcher, Eszter K. Vladar, and Max A. Seibold. Dissecting the cellular specificity of smoking effects and reconstructing lineages in the human airway epithelium. *Nature Communications*, 11 (1), May 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-16239-z. URL http://dx.doi.org/10.1038/s41467-020-16239-z.
- Laleh Haghverdi, Maren Büttner, F. Alexander Wolf, Florian Buettner, and Fabian J. Theis. Diffusion pseudotime robustly reconstructs lineage branching. *Nature Methods*, 13(10):845–848, 2016. doi: 10.1038/nmeth.3971.
- Douglas Hanahan and Robert A. Weinberg. Hallmarks of cancer: The next generation. *Cell*, 144 (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.
- Xing He and Chenqi Xu. Immune checkpoint signaling and cancer immunotherapy. *Cell Research*, 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.
- Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth schrödinger bridges. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=GHyvvWu1XC.
- Intekhab Hossain, Viola Fanfani, Jonas Fischer, John Quackenbush, and Rebekka Burkholz. Biologically informed neuralodes for genome-wide regulatory dynamics. *Genome Biology*, 25(1): 127, 2024.

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

- Sinisa Hrvatin, Daniel R Hochbaum, M Aurel Nagy, Marcelo Cicconet, Keiramarie Robertson, Lucas Cheadle, Rapolas Zilionis, Alex Ratner, Rebeca Borges-Monroy, Allon M Klein, et al. Singlecell analysis of experience-dependent transcriptomic states in the mouse visual cortex. *Nature neuroscience*, 21(1):120–129, 2018.
- Guillaume Huguet, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), Advances in Neural Information Processing Systems, 2022. URL https://openreview.net/forum?id=ahAEhOtVif.
- Erica L. Jackson, Nicholas Willis, Kim Mercer, Roderick T. Bronson, Denise Crowley, Raymond Montoya, Tyler Jacks, and David A. Tuveson. Analysis of lung tumor initiation and progression using conditional expression of oncogenic k-ras. *Genes & Development*, 15(24):3243–3248, December 2001. ISSN 1549-5477. doi: 10.1101/gad.943001. URL http://dx.doi.org/10.1101/gad.943001.
- Dominik Klein, Théo Uscidda, Fabian J Theis, and marco cuturi. GENOT: Entropic (gromov) wasserstein flow matching with applications to single-cell genomics. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=hjspWd7jvg.
- Dominik Klein, Giovanni Palla, Marius Lange, Michal Klein, Zoe Piran, Manuel Gander, Laetitia Meng-Papaxanthos, Michael Sterr, Lama Saber, Changying Jing, Aimée Bastidas-Ponce, Perla Cota, Marta Tarquis-Medina, Shrey Parikh, Ilan Gold, Heiko Lickert, Mostafa Bakhti, Mor Nitzan, Marco Cuturi, and Fabian J. Theis. Mapping cells through time and space with moscot. *Nature*, 638(8052):1065–1075, January 2025. ISSN 1476-4687. doi: 10.1038/s41586-024-08453-2. URL http://dx.doi.org/10.1038/s41586-024-08453-2.
- Gioele La Manno, Ruslan Soldatov, Amit Zeisel, and others. Rna velocity of single cells. *Nature*, 560(7719):494–498, 2018. doi: 10.1038/s41586-018-0414-6.
- Marius Lange, Volker Bergen, Michal Klein, Manu Setty, Bernhard Reuter, Mostafa Bakhti, Heiko Lickert, Meshal Ansari, Janine Schniering, Herbert B. Schiller, Dana Pe'er, and Fabian J. Theis. Cellrank for directed single-cell fate mapping. *Nature Methods*, 19(2):159–170, 2022. doi: 10. 1038/s41592-021-01346-6.
- Arthur Liberzon, Chet Birger, Helga Thorvaldsdóttir, Mahmoud Ghandi, Jill P. Mesirov, and Pablo Tamayo. The molecular signatures database hallmark gene set collection. *Cell Systems*, 1(6): 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.
- Zaikang Lin, Sei Chang, Aaron Zweig, Minseo Kang, Elham Azizi, and David A. Knowles. Interpretable neural odes for gene regulatory network discovery under perturbations, 2025. URL https://arxiv.org/abs/2501.02409.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2022.
- Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky TQ Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code. *arXiv* preprint arXiv:2412.06264, 2024.

- Zhen-Ling Liu, Huan-Huan Chen, Li-Li Zheng, Li-Ping Sun, and Lei Shi. Angiogenic signaling pathways and anti-angiogenic therapy for cancer. *Signal Transduction and Targeted Therapy*, 8 (1), May 2023. ISSN 2059-3635. doi: 10.1038/s41392-023-01460-1. URL http://dx.doi.org/10.1038/s41392-023-01460-1.
- Kevin R Moon, David Van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen, Kristina Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, et al. Visualizing structure and transitions in high-dimensional biological data. *Nature biotechnology*, 37(12):1482–1492, 2019.
- Daniel J. Murphy, Melissa R. Junttila, Laurent Pouyet, Anthony Karnezis, Ksenya Shchors, Duyen A. Bui, Lamorna Brown-Swigart, Leisa Johnson, and Gerard I. Evan. Distinct thresholds govern myc's biological output in vivo. *Cancer Cell*, 14(6):447–457, December 2008. ISSN 1535-6108. doi: 10.1016/j.ccr.2008.10.018. URL http://dx.doi.org/10.1016/j.ccr.2008.10.018.
- Huize Pan, Chenyi Xue, Benjamin J. Auerbach, Jiaxin Fan, Alexander C. Bashore, Jian Cui, Dina Y. Yang, Sarah B. Trignano, Wen Liu, Jianting Shi, Chinyere O. Ihuegbu, Erin C. Bush, Jeremy Worley, Lukas Vlahos, Pasquale Laise, Robert A. Solomon, Edward S. Connolly, Andrea Califano, Peter A. Sims, Hanrui Zhang, Mingyao Li, and Muredach P. Reilly. Single-cell genomics reveals a novel cell state during smooth muscle cell phenotypic switching and potential therapeutic targets for atherosclerosis in mouse and human. *Circulation*, 142(21): 2060–2075, November 2020. ISSN 1524-4539. doi: 10.1161/circulationaha.120.048378. URL http://dx.doi.org/10.1161/CIRCULATIONAHA.120.048378.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019. doi: 10.1561/2200000073.
- Simone Picelli. Single-cell rna-sequencing: The future of genome biology is now. *RNA Biology*, 14 (5):637–650, July 2016. ISSN 1555-8584. doi: 10.1080/15476286.2016.1201618. URL http://dx.doi.org/10.1080/15476286.2016.1201618.
- Chengxiang Qiu, Junyue Cao, Beth K. Martin, Tony Li, Ian C. Welsh, Sanjay Srivatsan, Xingfan Huang, Diego Calderon, William Stafford Noble, Christine M. Disteche, Stephen A. Murray, Malte Spielmann, Cecilia B. Moens, Cole Trapnell, and Jay Shendure. Systematic reconstruction of cellular trajectories across mouse embryogenesis. *Nature Genetics*, 54(3):328–341, March 2022. ISSN 1546-1718. doi: 10.1038/s41588-022-01018-x. URL http://dx.doi.org/10.1038/s41588-022-01018-x.
- Chengxiang Qiu, Beth K. Martin, Ian C. Welsh, Riza M. Daza, Truc-Mai Le, Xingfan Huang, Eva K. Nichols, Megan L. Taylor, Olivia Fulton, Diana R. O'Day, Anne Roshella Gomes, Saskia Ilcisin, Sanjay Srivatsan, Xinxian Deng, Christine M. Disteche, William Stafford Noble, Nobuhiko Hamazaki, Cecilia B. Moens, David Kimelman, Junyue Cao, Alexander F. Schier, Malte Spielmann, Stephen A. Murray, Cole Trapnell, and Jay Shendure. A single-cell time-lapse of mouse prenatal development from gastrula to birth. *Nature*, 626(8001):1084–1093, February 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07069-w. URL http://dx.doi.org/10.1038/s41586-024-07069-w.
- Xiaojie Qiu, Qiaolin Mao, Yong Tang, Lihua Wang, Rakel Chawla, Hannah A. Pliner, and Cole Trapnell. Reversed graph embedding resolves complex single-cell trajectories. *Nature Methods*, 14(10):979–982, 2017. doi: 10.1038/nmeth.4402.
- Wouter Saelens, Robrecht Cannoodt, Helena Todorov, et al. A guide to trajectory inference and rna velocity. *bioRxiv*, 2022. doi: 10.1101/2021.12.22.473434.
- Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, Lia Lee, Jenny Chen, Justin Brumbaugh, Philippe Rigollet, Konrad Hochedlinger, Rudolf Jaenisch, Aviv Regev, and Eric S. Lander. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943.e22, 2019. doi: 10.1016/j.cell.2019.01.006.
- Thibault Séjourné, Gabriel Peyré, and François-Xavier Vialard. Unbalanced optimal transport, from theory to numerics. *Handbook of Numerical Analysis*, 24:407–471, 2023.

- Aylin Sertkaya, Trinidad Beleche, Amber Jessup, and Benjamin D. Sommers. Costs of drug development and research and development intensity in the us, 2000-2018. *JAMA Network Open*, 7 (6):e2415445, June 2024. ISSN 2574-3805. doi: 10.1001/jamanetworkopen.2024.15445. URL http://dx.doi.org/10.1001/jamanetworkopen.2024.15445.
 - Alex K Shalek, Rahul Satija, Joe Shuga, John J Trombetta, Dave Gennert, Diana Lu, Peilin Chen, Rona S Gertner, Jellert T Gaublomme, Nir Yosef, et al. Single-cell rna-seq reveals dynamic paracrine control of cellular variation. *Nature*, 510(7505):363–369, 2014.
 - Kelly Street, Davide Risso, Robert B. Fletcher, Diya Das, John Ngai, Nir Yosef, Elizabeth Purdom, and Sandrine Dudoit. Slingshot: cell lineage and pseudotime inference for single-cell transcriptomics. BMC Genomics, 19(1):477, 2018. doi: 10.1186/s12864-018-4772-0.
 - B. J. Strober, R. Elorbany, K. Rhodes, N. Krishnan, K. Tayeb, A. Battle, and Y. Gilad. Dynamic genetic regulation of gene expression during cellular differentiation. *Science*, 364(6447): 1287–1290, June 2019. ISSN 1095-9203. doi: 10.1126/science.aaw0040. URL http://dx.doi.org/10.1126/science.aaw0040.
 - Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In *Proceedings of the 37th International Conference on Machine Learning*, 2020. ICML 2020; arXiv:2002.04461.
- Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, pp. 1–34, 2024.
- Betsabeh Khoramian Tusi, Samuel L. Wolock, Caleb Weinreb, Yung Hwang, Daniel Hidalgo, Rapolas Zilionis, Ari Waisman, Jun R. Huh, Allon M. Klein, and Merav Socolovsky. Population snapshots predict early haematopoietic and erythroid hierarchies. *Nature*, 555(7694):54–60, February 2018. ISSN 1476-4687. doi: 10.1038/nature25741. URL http://dx.doi.org/10.1038/nature25741.
- Titouan Vayer, Laetitia Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Fused gromov–wasserstein distance for structured objects. *Algorithms*, 13(9):212, 2020a. doi: 10.3390/a13090212.
- Titouan Vayer, Laetitia Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Fused gromov-wasserstein distance for structured objects. *Algorithms*, 13(9):212, 2020b.
- Cédric Villani et al. Optimal transport: old and new, volume 338. Springer, 2008.
- Kathryn A Wierenga, Frank M Riemers, Bart Westendorp, Jack R Harkema, and James J Pestka. Single cell analysis of docosahexaenoic acid suppression of sequential lps-induced proinflammatory and interferon-regulated gene expression in the macrophage. *Frontiers in Immunology*, 13: 993614, 2022.
- F. Alexander Wolf, Fiona K. Hamey, Mireya Plass, Jordi Solana, Joakim S. Dahlin, Berthold Göttgens, Nikolaus Rajewsky, Lukas Simon, and Fabian J. Theis. Paga: graph abstraction reconciles clustering with trajectory inference through a topology preserving map of single cells. *Genome Biology*, 20(1):59, 2019. doi: 10.1186/s13059-019-1663-x.
- Alan T. Yeo, Shruti Rawal, Bethany Delcuze, Anthos Christofides, Agata Atayde, Laura Strauss, Leonora Balaj, Vaughn A. Rogers, Erik J. Uhlmann, Hemant Varma, Bob S. Carter, Vassiliki A. Boussiotis, and Al Charest. Single-cell rna sequencing reveals evolution of immune land-scape during glioblastoma progression. *Nature Immunology*, 23(6):971–984, May 2022. ISSN 1529-2916. doi: 10.1038/s41590-022-01215-0. URL http://dx.doi.org/10.1038/s41590-022-01215-0.
- Rongting Yue and Abhishek Dutta. Computational systems biology in disease modeling and control, review and perspectives. *npj Systems Biology and Applications*, 8(1), October 2022. ISSN 2056-7189. doi: 10.1038/s41540-022-00247-4. URL http://dx.doi.org/10.1038/s41540-022-00247-4.

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

APPENDIX

A Extended related works

Inferring 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., 2025). 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	$\min_{\Gamma} \langle \Gamma, C \rangle + \tau \text{ KL}(\Gamma 1 a) + \tau \text{ KL}(\Gamma^{\top} 1 b)$	Growth priors (no LR)	≥ 2	scRNA-seq (+growth)
SCOT (Demetci et al., 2022b)	static	Fused GW	$\min_{\Gamma} \langle \Gamma, C \rangle + \lambda \langle L(D^x, D^y), \Gamma \otimes \Gamma \rangle + \varepsilon H(\Gamma)$	k-NN geometry (untyped)	2	scRNA-seq
CellOT (Bunne et al., 2023b)	static	Kantorovich (dual/convex)	$\max_f \min_g C - \mathbb{E}_{g_s}[\langle x, \nabla g(x) \rangle - f(\nabla g(x))] - \mathbb{E}_{g_s}[f(y)]$	None	2	scRNA-seq
TrajectoryNet (Tong et al., 2020)	dynamic (continuous)	BB OT + neural ODE prior	$\min_{\theta} - \sum_{i} \log P_{t_i}(x_{t_i}) + \int v_{\theta} ^2 dt + \text{bio priors}$	None	≥ 2	scRNA-seq (often splicing)
OT-CFM (Tong et al., 2024)	dynamic (continuous)	Kantorovich + flow matching	$\min_{\Gamma} \langle \Gamma, C \rangle$; $\min_{\theta} \mathbb{E}_{(x,y) \sim \Gamma} v_{\theta} - v^{OT} ^2$	None	2	scRNA-seq
Schrödinger Bridge (DSB)	dynamic (stochastic)	SB/DSB	$\min_{p(x_{0:T})} KL(p \parallel ref-diffusion)$ s.t. $p_0 = \mu_0, p_T = \mu_T$	None	≥ 2	scRNA-seq (+noise model)
IADOT (ours)	dynamic (continuous)	Multi-channel FGW + CFM	$\min_{\Gamma} (1 - \alpha)\langle \Gamma, C \rangle + \alpha \langle \varphi(G^{(0)}, G^{(1)}), \Gamma \otimes \Gamma \rangle$; $\min_{\theta} CFM(\theta \mid \Gamma)$	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., 2025). 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).

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. (201:	
Cardiology	Heart/aorta scRNA-seq after myocardial infarction (e.g., D1, D3, D7) or atherosclerosis progression (early \rightarrow intermediate \rightarrow late).	Post-MI hear: Farbehi et al. (2019) Atherosclerosis: Pan et al. (2020)	
Immunology	Tissue + immune scRNA-seq across baseline \rightarrow active disease \rightarrow 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 (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.

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

Dataset	Reference	Timestamps	#Cells	#Genes
Tumour	_	0, 8, 24, 168 (h)	31,536	22,681
V1 Cortex	(Hrvatin et al., 2018)	0, 1, 4 (h)	6,505	17,008
Immune Stimulus	(Wierenga et al., 2022)	0, 1, 2, 4, 6 (h)	2,382	10,972
Mouse embryo	(Moon et al., 2019)	0, 6, 12, 18, 24 (d)	18,203	17,789
Macrophage Stimulus	(Shalek et al., 2014)	0, 3, 5 (h)	223	478

C.1 SYNTHETIC EXAMPLE

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

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

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 obtained by translating each cluster via

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

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

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

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

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

C.2 LUNG TUMOR

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

C.3 CORTEX—LIGHT STIMULATION

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

C.4 IMMUNE

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

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.

Table 7: Software stack (key Python packages).

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; Omega-Conf 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.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 Appendix 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 select the resolution ρ^* by scanning a small grid of resolutions and choosing the value whose *median* cluster size is closest to a target of $n^*=40$ cells.

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 to one-hot assignments). We obtain metacell-level activations by averaging the s_{cg} within clusters and form the metacell CCI in $\mathbb{R}^{g \times g}$ similarly as in the setting with individual cells.

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

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

This lifting operation ensures $S^{\top}\tilde{G}S = \bar{G}$. In contrast to G, the matrix \tilde{G} is constrained to lie 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 metacell–metacell interactions.

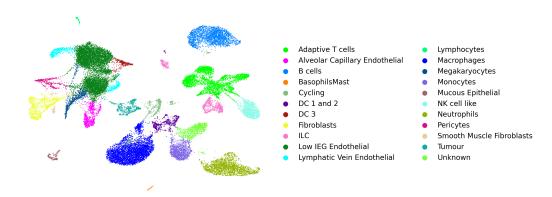


Figure 7: **Metacell construction example.** UMAP visualization of the single-cell RNA-seq data of the lung cancer dataset after Leiden clustering. Each point corresponds to an individual cell, colored by its assigned cluster and annotated with the corresponding cell type based on marker genes.

D.4 OPTIMAL TRANSPORT SOLVER

We extend POT's (Flamary et al., 2024) conditional-gradient (Frank–Wolfe) solver to handle multichannel 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 (uniform by default), and a matrix $\Sigma \succeq 0 \in \mathbb{R}^{d \times d}$, we measure discrepancies with the Mahalanobis norm $\|x\|_{\Sigma} = \sqrt{x^{\top} \Sigma^{-1} x}$.

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 quadratic term is

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

with

$$\operatorname{constC}_{ij} = \sum_{k} \|C_1[i, k]\|_{\Sigma}^2 p_k + \sum_{l} \|C_2[j, l]\|_{\Sigma}^2 q_l, \qquad \mathcal{B}(\Gamma) = \sum_{r=1}^{d} C_1^{(r)} \Gamma\left(C_2^{(r)}\right)^{\top}.$$

The gradient computed by the solver is

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

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

We minimize

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

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

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

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

Table 8: Flow Matching hyperparameters.

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

compute the two endpoint couplings by solving the feature-only ($\alpha = 0$) and structure-only ($\alpha = 1$) problems, yielding $\Gamma_{\alpha=0}^{\star}$ and $\Gamma_{\alpha=1}^{\star}$. We then define the scaling factors as follows:

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

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

(11)

and rescale the feature cost matrix and the CCI tensors:

$$C \leftarrow \frac{C}{|\Delta \mathcal{F}|},\tag{12}$$

$$C \leftarrow \frac{C}{|\Delta \mathcal{F}|}, \tag{12}$$

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

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.

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/ $N_{
m sig\ edges}$ and retain only pairs with $0.10 \le coverage \le 0.40$ and at least two sources and two targets. We compute a standardized score s = 0.6z (mean_expr) + 0.4z (-spec_rank) and greedily select pairs in descending s while preventing repeated ligands or receptors. We keep the top 10 pairs for each dataset.

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.

Diffusion Shrodinger 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.









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	Hidden dimensions	64-64-64
	Blocks / layer type	1 block, concatsquash layers
	Activations	Softplus layers with tanh control
Regularization	$s_{ m L2int}$	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	$rtol = atol = 1 \times 10^{-5}$

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

Category	Hyperparameter	Setting
Model	Score network	Dual-encoder MLP; encoders [16, 32], decoder [64, 64,
		64], time embedding 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
	Learning rate	1×10^{-4}
Regularisation	γ schedule	$\gamma_{\rm min}=\gamma_{\rm max}=10^{-3}$, linear spacing
	Mean matching	Enabled
	Cache handling	2 048 cached particles, 2 048 plotting points, refresh every
		200 iters
	EMA	Disabled

D.9 LUNG CANCER DATA EXPERIEMENT

For the experiment described in Section 5.4, We annotated the lung cancer dataset using canonical lineage and state markers (Table 11); an overview of the full dataset is shown in Fig. 7. Because whole-lung profiling dilutes treatment effects (the tumour comprises only a small fraction of total cells), we constructed a focused tumour-niche subset to increase sensitivity and interpretability. Concretely, we retained all tumour cells and subsampled an equal number of T cells, B cells, fibroblasts, and endothelial cells from the same specimens to form a minimal viable tumour microenvironment. We then reused the analysis pipeline described earlier with matched timepoints at 0 h, 8 h, and 24 h. The only modification was to the ligand–receptor (LR) library: for pathway-specific probes, we toggled custom LR pairs to mimic the presence or absence of a given ligand (e.g., EGFR) and quantified the resulting changes in inferred communication and downstream dynamics. Marker definitions are provided in Table 11, and a dot-plot confirming marker specificity and minimal cross-lineage leakage is shown in Fig. 8.

TUMOUR PROGRESSION QUANTIFICATION USING HALLMARK GENE SETS

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

Cell Type	Positive Markers
Differentiated AT1	RTKN2, AGER
AT1	CLDN18
Tumour (AT2)	SFTPD, 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

Table 11: Curated panel of positive marker genes used for per-cell scoring and assignment in the lung cancer dataset.

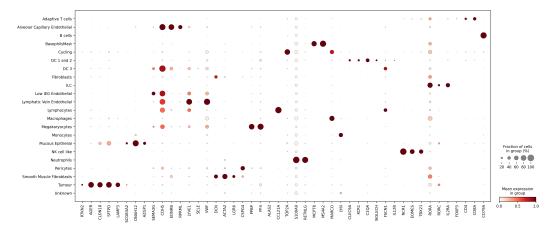


Figure 8: Dot-plot validation of curated marker genes across annotated cell types. Each column corresponds to a marker gene and each row to a cell-type labe. Dot size encodes the fraction of cells expressing that gene, while color intensity represents its standardized expression level.

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 12. 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 $0 \text{ h} \rightarrow 24 \text{ h}$; see Fig. 9.

Table 12: 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, UQCRQ, 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

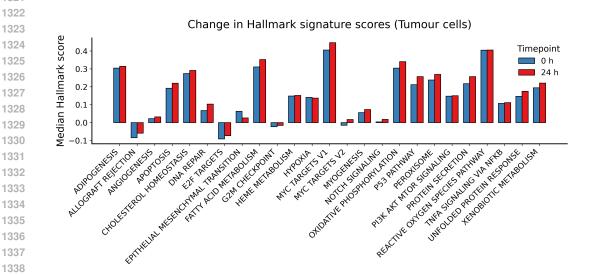


Figure 9: Hallmark changes. Changes in our dataset over 24 h following combined KRAS and MYC signalling across the 20 selected Hallmark gene sets.

E Additional results

STIMULUS DATASETS

We reproduce the experimental setup described in Section 5.2 and Section 5.3 with the macrophase stimulus-response dataset (Appendix C.6). We report the results in Figure 10 and Table 13, which are consistent with the findings on the other datasets.

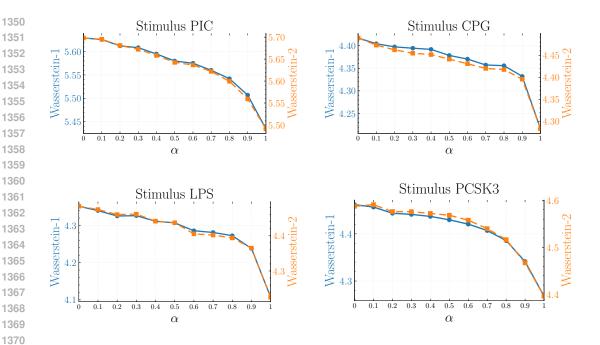


Figure 10: **Interpolation** results for the macrophage stimulus datasets.

Table 13: 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.

		Stimulus PIC		Stimulus CPG		Stimulus LPS		Stimulus PCSK3	
Method	α	W_1	W_2	W_1	W_2	W_1	W_2	W_1	W_2
TRAJECTORYNET DSB		5.628 (0.055) 5.796 (0.574)		${5.361 (0.085)}\atop 4.500 (0.128)$					
IADOT		5.544 (0.038) 5.490 (0.018) 5.446 (0.018)	5.555 (0.019)	4.427 (0.021)	4.502 (0.025)	4.380 (0.021)	4.517(0.021)	4.423 (0.020) 4.396 (0.023) 4.352 (0.020)	4.531 (0.019)

E.2 Sensitivity of couplings to catalog edits

The experiment presented in Section 5.4 involved perturbing the LR catalog by removing specific LR pairs. In Table 14, 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.

F LLM usage

We used large language models (LLMs) to assist with improving the clarity of writing and refining the formatting of tables and figures. LLMs were not used for research ideation, experimental design, analysis, or any substantive contributions that would merit authorship.

Table 14: 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%