

CellStream: Dynamical Optimal Transport Informed Embeddings for Reconstructing Cellular Trajectories from Snapshots Data

Yue Ling^{1*}, Peiqi Zhang^{1,5*}, Zhenyi Zhang^{1†}, Peijie Zhou^{2,3,4,5†}

¹ School of Mathematical Sciences, Peking University, China

² Center for Machine Learning Research, Peking University, China

³ Center for Quantitative Biology, Peking University, China

⁴ National Engineering Laboratory for Big Data Analysis and Applications, China

⁵ AI for Science Institute, Beijing, China

zhenyizhang@stu.pku.edu.cn; pjzhou@pku.edu.cn

Abstract

Single-cell RNA sequencing (scRNA-seq), especially temporally resolved datasets, enables genome-wide profiling of gene expression dynamics at single-cell resolution across discrete time points. However, current technologies provide only sparse, static snapshots of cell states and are inherently influenced by technical noise, complicating the inference and representation of continuous transcriptional dynamics. Although embedding methods can reduce dimensionality and mitigate technical noise, the majority of existing approaches typically treat trajectory inference separately from embedding construction, often neglecting temporal structure. To address this challenge, here we introduce CellStream, a novel deep learning framework that jointly learns embedding and cellular dynamics from single-cell snapshots data by integrating an autoencoder with unbalanced dynamical optimal transport. Compared to existing methods, CellStream generates dynamics-informed embeddings that robustly capture temporal developmental processes while maintaining high consistency with the underlying data manifold. We demonstrate CellStream’s effectiveness on both simulated datasets and real scRNA-seq data, including spatial transcriptomics. Our experiments indicate significant quantitative improvements over state-of-the-art methods in representing cellular trajectories with enhanced temporal coherence and reduced noise sensitivity. Overall, CellStream provides a new tool for learning and representing continuous streams from the noisy, static snapshots of single-cell gene expression.

Code — <https://github.com/PQ-Zhang/CellStream>

Introduction

Time-resolved single-cell RNA sequencing (scRNA-seq) captures gene expression at the genome scale in individual cells across discrete time points, enabling analysis of dynamic biological processes. However, because sequencing destroys cells, it yields sparse, noisy snapshots that obscure continuous dynamical trajectories (Schiebinger et al. 2019;

Klein et al. 2025; Zhang et al. 2025a). Dimensionality reduction through low-dimensional embeddings can partially alleviate sparsity and noise issues, and provide insights of the temporal inference through pseudotime analysis (Zhang et al. 2025a). However, existing embedding methods predominantly focus on preserving topological or geometric relationships (Lin, Troup, and Ho 2017; Mongia, Sengupta, and Majumdar 2019; Hou et al. 2020). They often neglect the temporal structure embedded in time-series snapshots data, potentially resulting in misinterpretation of cellular trajectories (Zheng et al. 2023; Chari and Pachter 2023). Thus, developing embeddings that explicitly incorporate temporal structure to faithfully reconstruct continuous cellular dynamics remains an open and critical challenge in the field.

For a single snapshot scRNA-seq data, many embedding methods are proposed to model and visualize the biological processes on a low-dimensional manifold. For example, embeddings derived from principal component analysis (PCA), t-distributed Stochastic Neighbor Embeddings (t-SNE), Uniform Manifold Approximation and Projection (UMAP) and diffusion maps are typical methods to decipher the topological structure of gene expression (van der Maaten and Hinton 2008; McInnes et al. 2018; Coifman and Lafon 2006). Canonical correlation analysis (CCA), Harmony and Partition-based graph abstraction (PAGA) shed light on preserving the statistical, geometric or topological structure (Wilks 2011; Korsunsky et al. 2019; Wolf et al. 2019). However, these embedding methods usually focus on retaining the geometric structure but are not designed to capture the cell-state transition dynamics and temporal relationships between snapshots. To address the problem, CellPath, Ocelli and VeloViz propose using RNA velocity results (Manno et al. 2018) to conduct velocity-informed visualization and pseudo-time trajectory inference (Zhang and Zhang 2021; Rutkowski and Tabaka 2025; Atta, Sahoo, and Fan 2022). However, the RNA velocity inference could encounter computational issues (Li et al. 2023; Zheng et al. 2023), especially in datasets with low unspliced counts, and does not consider cross-time point analysis explicitly. While deep generative model embeddings, like scVI, Geneformer and TarDis (Lopez et al. 2018; Theodoris et al. 2023; Inecik et al. 2025), attempt to address batch effects and library size

*These authors contributed equally.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

effects, they suffer from limited interpretability for the dynamical modeling.

To deal with the time-series scRNA-seq data, Optimal transport (OT) has emerged as a promising tool to model the cellular dynamics (Bunne et al. 2024; Zhang et al. 2025a). Its core idea relies on seeking the optimal transport plan between the population of cells among consecutive time points with minimal transportation cost (Schiebinger et al. 2019; Klein et al. 2025). To deal with cell proliferation effects, TIGON proposes the application of unbalanced dynamical OT in trajectory inference (Sha et al. 2024). Subsequent works further relate the regularized unbalanced optimal transport (RUOT) and mean-field terms to model stochastic effects and cell-cell communications, proposing a new framework with higher performance and efficiency (Zhang, Li, and Zhou 2025; Zhang et al. 2025b; Sun et al. 2025). Nevertheless, since direct inference of cellular dynamics in a high-dimensional gene expression space is usually unstable and demands high computational resources (Zhang et al. 2021), these methods usually rely on a pre-defined embedding as the input such as PCA or UMAP space, which is decoupled from the dynamical learning.

To address these challenges, we present CellStream, a deep learning framework that generates dynamics-informed embeddings based on unbalanced dynamical OT for time-series single-cell snapshots data. We evaluate the effectiveness of CellStream in modeling cellular dynamics using both simulated data and real data. Experiments demonstrate that CellStream embeddings can precisely reveal developmental trajectories from the noisy, static snapshots of single-cell gene expression. Overall, our method features the following contributions:

- We introduce CellStream that **jointly learns embedding and latent cellular dynamics** from cross-sectional snapshots of single-cell data by integrating an autoencoder with unbalanced dynamical optimal transport;
- We demonstrate that CellStream embeddings are capable of reconstructing the cellular dynamics **robustly against noise** through simulation benchmarking.
- We validate the effectiveness of CellStream on both simulated data and real snapshots data, showing its significant quantitative improvements over state-of-the-art methods in **representing cellular trajectories with enhanced temporal coherence**.

Related Work

Manifold-learning Embedding

Numerous methods have been developed to project the complex, high-dimensional data onto low-dimensional manifolds. PCA, t-SNE and UMAP are widely used to reveal the spatial distribution of the given cellular trajectories (Ge et al. 2025; Gao, Qiao, and Huang 2022; Zhang, Li, and Zhou 2025). Diffusion map models the geometry of the dataset as a diffusion process on point clouds, capturing local similarity and global structure via the eigen-space of Markov chains (Coifman and Lafon 2006). Canonical correlation analysis (CCA) constructs the embedding space by

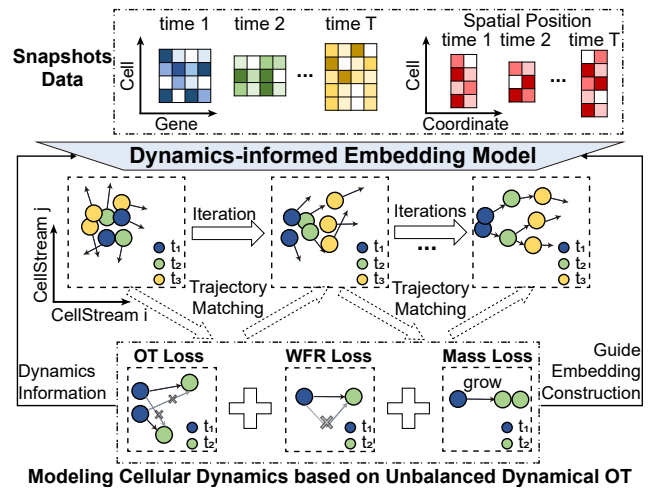


Figure 1: Overview of CellStream. CellStream is a dynamics-informed embedding model that jointly learns embeddings and continuous cellular streams from static, sparse snapshots data. CellStream uses an autoencoder and neural network to learn embeddings and continuous cellular streams from sparse snapshots data, guided by real-time trajectory feedback composed of \mathcal{L}_{OT} , \mathcal{L}_{WFR} and \mathcal{L}_{Mass} .

identifying linear transformations that maximize the correlations between multiple datasets to mitigate technical noise and preserve biological signals (Wilks 2011; Raslan et al. 2024). Similarly, Harmony performs iterative soft clustering and linear corrections in the latent space to remove batch effects (Korsunsky et al. 2019; Huang et al. 2025). Partition-based graph abstraction (PAGA) builds a topology-preserving embedding space through graph abstraction, coarsening single-cell neighborhood graphs to retain geometric structure (Wolf et al. 2019; Aissa et al. 2021). Although these approaches specialize in extracting topological traits from individual snapshots, they tend to overlook dynamic information and not fully consider temporal structure across different time points.

Dynamical Embedding

In recent years, several studies have explored how to manifest cellular dynamics in the embedding space. By integrating multiscale stochastic dynamics, Multiscale Transition Analysis (MuTrans) generates a dynamical manifold embedding that maps stable cell states to attractor basins and transient cells to saddle points, revealing continuous cell fate trajectories (Zhou et al. 2021). CellPath constructs meta-cell level embeddings based on gene expression similarity and RNA velocity coherence, employing a directed KNN graph for pseudo-time trajectory inference through path finding algorithms (Zhang and Zhang 2021). Integrating RNA velocity and multi-modal diffusion maps, Ocelli represents the cellular differentiation direction on a low-dimensional embedding from multimodal single-cell data (Rutkowski and Tabaka 2025). VeloViz constructs RNA velocity-informed embeddings based on a weighted k-nearest neighbor graph which integrates RNA velocity with gene expression (Atta,

Sahoo, and Fan 2022). However, the steady-state assumption of MuTrans may limit its application to non-stationary systems with rapid proliferation, and RNA velocity used in CellPath, Ocelli and VeloViz faces challenges to model complex temporal transitions across time points which have longer time scales than splicing dynamics (Zhou et al. 2024).

Deep Generative Model Embedding

Deep generative model embeddings leverage deep learning to generate robust, low-dimensional representations, enabling batch correction and regulatory inference. By integrating variational inference and conditional generative, scVI provides low-dimensional embeddings free from batch effects and library size effects (Lopez et al. 2018). GeneFormer employs a transformer architecture to tokenize rank-ordered gene expressions into sequences for self-supervised masked prediction, generating contextual cell embeddings that encode regulatory network dynamics (Theodoris et al. 2023). TarDis constructs embeddings in a variational autoencoder framework and leverages self-supervised disentanglement losses to distinguish covariate-specific variations from underlying biological processes (Inecik et al. 2025). However, the majority of existing deep generative model embeddings are usually established on VAE or Transformer architectures, encountering challenges in the interpretability issue of dynamical modeling.

Dynamical OT based Trajectory Inference

Dynamical OT proves to be an effective model for deciphering cellular dynamics (Zhang et al. 2025a; Gu, Chien, and Greenwald 2025). TrajectoryNet pioneers the application of dynamical OT in trajectory inference combined with continuous normalizing flows (Tong et al. 2020). MIOFlow introduces a Geodesic Autoencoder (GAE) to learn low-dimensional manifolds which enable trajectory inference from time-series snapshots based on dynamical optimal transport theory (Huguet et al. 2022). Leveraging unbalanced dynamical OT, TIGON infers differentiation trajectories and population growth from scRNA-seq data simultaneously (Sha et al. 2024). stVCR further reconstructs the interaction between gene expression and spatial migration from spatial transcriptome (ST) data (Peng, Zhou, and Li 2025; Zhang et al. 2025c). Focused on intercellular communications, CytoBridge relates the regularized unbalanced optimal transport (RUOT) to Mean-Field Schrödinger Bridge (MFSB), recovering unbalanced stochastic interaction dynamics (Zhang, Li, and Zhou 2025; Zhang et al. 2025b; Sun et al. 2025). There are also many flow matching-based methods to tackle this problem (Petrović et al. 2025; Tong et al. 2024b,a; Kapusniak et al. 2024; Wang et al. 2025) and also some methods are based on the Schrödinger Bridge (SB) problem (Lavenant et al. 2024; Pariset et al. 2023). However, to tackle sparsity and noise problems in original gene expression space, existing works often rely on a reduced dimension input as pre-computed embeddings. The separation of trajectory inference from embedding construction is limited when considering the temporal structure of datasets, impeding the reconstruction of spatiotemporal dynamics.

Methodology

Optimal Transport

Static Optimal Transport. Optimal transport(OT) seeks to minimize the transport cost between two populations, namely the Kantorovich problem (Kantorovich 1942). For $p \in [1, +\infty]$, the Wasserstein p -distance is defined as

$$\mathcal{W}_p(\mu, \nu) = \inf_{\gamma \in \Gamma(\mu, \nu)} (\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \gamma} [d(\mathbf{x}, \mathbf{y})^p])^{1/p},$$

where $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$, and a coupling γ is a joint probability measure whose marginals are μ and ν . $\Gamma(\mu, \nu)$ denotes the set of all couplings and $d : X \times Y \rightarrow \mathbb{R}^+$ is a given distance function. $\mathcal{W}_p(\mu, \nu)$ represents the optimal transport cost between two probability measures. Unless explicitly stated, we choose $p = 2$ by default in the calculation.

Dynamical Optimal Transport. Equivalent to static OT, the framework of dynamical OT was formulated in the form of fluid dynamics equations, which modeled the transport as a continuous flow and optimized over smooth and time-dependent density $\rho(t, \mathbf{x})$ and velocity fields $\mathbf{v}(t, \mathbf{x})$ (Benamou and Brenier 2000). The spatiotemporal dynamics of the mass distribution was subject to the following continuity equation with the initial and final constraints:

$$\partial_t \rho + \nabla_{\mathbf{x}} \cdot (\mathbf{v}\rho) = 0, \quad \rho(0, \cdot) = \rho_0, \quad \rho(T, \cdot) = \rho_T$$

for any $t \in [0, T]$ and $\mathbf{x} \in \mathbb{R}^d$, where $\mathbf{v}(t, \mathbf{x}) \in \mathbb{R}^d$ determined the velocity field that governed the dynamics of the density. When defining the transport distance function as the Euclidean distance, namely $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2$, the cost function of dynamical OT can be expressed as

$$\mathcal{L}_{\text{dynamical OT}} = \int_0^T \int_{\mathbb{R}^d} \|\mathbf{v}(t, \mathbf{x})\|^2 \rho(t, \mathbf{x}) \, d\mathbf{x} \, dt.$$

Unbalanced Dynamical Optimal Transport. The mass conservation assumption of dynamical OT does not account for the total mass changing, which is an important phenomenon in biological processes due to cell apoptosis or divide. To model such dynamics, unbalanced dynamical OT introduces a growth term $g(t, \mathbf{x}) : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ to rewrite the continuity equation as

$$\partial_t \rho + \nabla_{\mathbf{x}} \cdot (\mathbf{v}\rho) = g\rho, \quad \rho(0, \cdot) = \rho_0, \quad \rho(T, \cdot) = \rho_T.$$

A transport cost function that considers both kinetic and growth energy simultaneously is imperative to constrain the transport dynamics. Wasserstein–Fisher–Rao (WFR) distance emerges as a desirable choice to describe the overall shipping cost (Chizat et al. 2018, 2019), and the corresponding cost function of unbalanced dynamical OT can be defined as:

$$\mathcal{L}_{\text{WFR}} = \int_0^T \int_{\mathbb{R}^d} (\|\mathbf{v}(t, \mathbf{x})\|^2 + \alpha g^2(t, \mathbf{x})) \rho(t, \mathbf{x}) \, d\mathbf{x} \, dt, \quad (1)$$

where integral terms involving $\|\mathbf{v}(t, \mathbf{x})\|^2$ and $g^2(t, \mathbf{x})$ are used to quantify the cost of transport and growth respectively. The hyperparameter α is introduced to explicitly control the relative weight of the above two terms when minimizing the cost function, and we set $\alpha = 1$ in our tasks. The sensitivity of our algorithm to α is presented in Appendix B.5.

Modeling the Single Cell Transcriptional Dynamics in the Low-dimensional Embedding Space

Despite the high-dimensionality of gene expression, it is commonly assumed that the dynamical biological processes underlying scRNA-seq data are indeed embedded on a low-dimensional manifold (Lin, Troup, and Ho 2017; Mongia, Sengupta, and Majumdar 2019; Hou et al. 2020). CellStream aims to model the dynamics of cellular trajectories and growth in an iteratively refined low-dimensional embedding space based on unbalanced dynamical OT. Formally, our goal is to jointly learn a suitable time-dependent low-dimensional embedding $f^{\text{enc}}(t, \mathbf{x})$, which maps the original D -dimensional data with time t to a d -dimensional space ($d < D$), together with the dynamical functions \mathbf{v} and g defined in this embedding space, such that the WFR distance $\mathcal{L}_{\text{WFR}}^{\text{emb}}$ in the dynamical optimal transport problem is minimized.

The constraint of $f^{\text{enc}}(t, \mathbf{x})$ is enforced by introducing an embedding loss $\mathcal{L}_{\text{embedding}}$, which promotes faithful reconstruction of the input in the latent space. For example, if f^{enc} is implemented as the encoder module of an autoencoder, a decoder $f^{\text{dec}} : \mathbb{R}^d \rightarrow \mathbb{R} \times \mathbb{R}^D$ can be jointly trained, and the embedding loss is defined as $\mathcal{L}_{\text{embedding}} = \text{MSE}(f^{\text{dec}} \circ f^{\text{enc}}([t, \mathbf{x}]), [t, \mathbf{x}])$. The specific form of the problem is as follows:

$$\min_{f^{\text{enc}}, f^{\text{dec}}, \mathbf{v}, g} \mathcal{L}_{\text{embedding}}(f^{\text{enc}}, f^{\text{dec}}) + \lambda \mathcal{L}_{\text{WFR}}^{\text{emb}}(f^{\text{enc}}, \mathbf{v}, g)$$

where $f^{\text{enc}}(t, \mathbf{x}) : \mathbb{R} \times \mathbb{R}^D \rightarrow \mathbb{R}^d$, $\mathbf{v}(t, \mathbf{z}) : \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, $g(t, \mathbf{z}) : \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$, with $\mathbf{z} = f^{\text{enc}}(t, \mathbf{x})$.

The system is constrained by $\rho(0, \cdot) = \rho_0$, $\rho(T, \cdot) = \rho_T$, and satisfies the continuity equation

$$\partial_t q(t, \mathbf{z}) + \nabla_{\mathbf{z}} \cdot (\mathbf{v}(t, \mathbf{z})q(t, \mathbf{z})) = g(t, \mathbf{z})q(t, \mathbf{z}), \quad (2)$$

where $q = f_{\#}^{\text{enc}} \rho$, i.e. $q(t, \mathbf{z}) = \int \rho(t, \mathbf{x}) \delta(\mathbf{z} - f^{\text{enc}}(\mathbf{x})) d\mathbf{x}$.

Deep Learning-based Solver in CellStream

We suppose that time series RNA-seq data are given by

$$(t_i, \mathbf{X}_i)_{i=0}^{T-1}, \quad \mathbf{X}_i = \{\mathbf{x}_i^j \in \mathbb{R}^D\}_{j=1}^{N_i}$$

where D indicates the dimension of gene expression, N_i represents the cell population at time t_i and we use index j for the single cells. We assume that the number of cells involved at time t_i is proportional to N_i if there is no prior information. In the embedding space, the input data is reduced to

$$(t_i, \mathbf{Z}_i, \mathbf{w}_i)_{i=0}^{T-1}, \quad \mathbf{Z}_i = \{\mathbf{z}_i^j = f^{\text{enc}}(\mathbf{x}_i^j) \in \mathbb{R}^d\}_{j=1}^{N_i}$$

where $\mathbf{w}_i = \frac{1}{N_i} \mathbf{1}_{N_i} \in \mathbb{R}^{N_i}$ implies the weight of each point. Driven by \mathbf{v} and g , we can also obtain predicted states:

$$\left(t_i, \hat{\mathbf{Z}}_i, \hat{\mathbf{w}}_i = \{\hat{w}_i^j(t_i)\}_{j=1}^{N_0} \right)_{i=0}^{T-1}, \quad \hat{\mathbf{Z}}_i = \{\hat{\mathbf{z}}^j(t_i) \in \mathbb{R}^d\}_{j=1}^{N_0}$$

initialized from $\hat{\mathbf{Z}}_0 = \mathbf{Z}_0$, $\hat{\mathbf{w}}_0 = \mathbf{w}_0$ by ODE such that

$$d\hat{\mathbf{z}}^j(t)/dt = \mathbf{v}(t, \hat{\mathbf{z}}^j), \quad d\log \hat{w}_i^j(t)/dt = g(t, \hat{\mathbf{z}}^j). \quad (3)$$

In order to construct a feasible and efficient solver for CellStream, we implement f^{enc} as an autoencoder and parameterize these functions to be solved (including autoencoder

f^{enc} and f^{dec} , velocity field \mathbf{v} , and growth term g) based on neural networks as $f_{\theta}, \mathbf{v}_{\phi}, g_{\psi}$. Here θ, ϕ, ψ are parameters to be optimized in neural networks. The total loss function of neural networks (\mathcal{L}) can be divided into autoencoder loss (\mathcal{L}_{AE}), WFR distance loss (\mathcal{L}_{WFR}), and data matching loss ($\mathcal{L}_{\text{Match}}$):

$$\mathcal{L} = \lambda_{\text{AE}} \mathcal{L}_{\text{AE}} + \lambda_{\text{WFR}} \mathcal{L}_{\text{WFR}}^{\text{emb}} + \lambda_{\text{Match}} \mathcal{L}_{\text{Match}}. \quad (4)$$

The three loss terms refer to autoencoder reconstruction error, optimal transport distance, and matching error. $\lambda_{\text{AE}}, \lambda_{\text{WFR}}$ and λ_{Match} are hyperparameters controlling the relative weight between loss terms. In the following experiments, we set $\lambda_{\text{AE}} = 10$, $\lambda_{\text{WFR}} = 1$, $\lambda_{\text{Match}} = 5$. We will verify the robustness of the algorithm to those parameters in Appendix B.5.

Autoencoder Loss. To minimize the reconstruction error between the estimated single-cell data and the observed measurements, we introduce autoencoder loss \mathcal{L}_{AE} based on Mean Squared Error (MSE) :

$$\mathcal{L}_{\text{AE}}(\theta) = \text{MSE}(f_{\theta}^{\text{dec}} \circ f_{\theta}^{\text{enc}}(t, \mathbf{x}), [t, \mathbf{x}])$$

Wasserstein-Fisher-Rao Loss in Embedding Space. We can rewrite (1) under the circumstances of neural networks as:

$$\mathcal{L}_{\text{WFR}}^{\text{emb}}(\theta, \phi, \psi) = \int_0^T \int_{\mathbb{R}^d} (\|\mathbf{v}_{\phi}(t, \mathbf{z}_{\theta})\|^2 + \alpha g_{\psi}^2(t, \mathbf{z}_{\theta})) q(t, \mathbf{z}_{\theta}) d\mathbf{z}_{\theta} dt. \quad (5)$$

The integral form in (5) involves computing the density term $q(t, \mathbf{z})$ that satisfies (2), which is computationally prohibitive in practice since it requires solving complex partial differential equations (PDE) in high-dimensional space. To address this computational bottleneck, we leverage the fundamental equivalence between Eulerian and Lagrangian descriptions of fluid dynamics and further convert (5) to an equivalent dimensionless form (Sha et al. 2024; Zhang, Li, and Zhou 2025):

$$\mathcal{L}_{\text{WFR}}^{\text{emb, par}}(\theta, \phi, \psi) = \sum_{j=1}^{N_0} \int_0^T (\|\mathbf{v}_{\phi}(t, \hat{\mathbf{z}}^j(t))\|^2 + \alpha g_{\psi}^2(t, \hat{\mathbf{z}}^j(t))) \hat{w}_i^j(t) dt,$$

where $\hat{\mathbf{z}}^j(t)$ can be calculated efficiently using Neural ODE solver (Chen et al. 2018) for (3).

Data Matching Loss. Data matching loss aims to match the distribution of the estimated data at each time point with its corresponding ground truth. The Wasserstein distance is widely used to quantify such distribution difference. However, since we relax the mass conservation assumption to the unbalanced event (2), the combination of Mass Loss $\mathcal{L}_{\text{Mass}}$ and OT Loss \mathcal{L}_{OT} is calculated as $\mathcal{L}_{\text{Match}}(\theta, \phi, \psi) = \lambda_{\text{Mass}} \mathcal{L}_{\text{Mass}} + \lambda_{\text{OT}} \mathcal{L}_{\text{OT}}$, where $\mathcal{L}_{\text{Mass}}$ aims to ensure the consistency of cell population and \mathcal{L}_{OT} is used to evaluate the gap between different distributions. λ_{Mass} and λ_{OT} are hyperparameters to balance the relative weight, and we set $\lambda_{\text{Mass}} = 1$, $\lambda_{\text{OT}} = 1$ in the following experiments.

Data matching loss $\mathcal{L}_{\text{Match}}$ is intended to quantify the differences between the estimated distribution $(t_i, \hat{\mathbf{Z}}_i, \hat{\mathbf{w}}_i)_{i=0}^{T-1}$

and the ground truth $(t_i, \mathbf{Z}_i, \mathbf{w}_i)_{i=0}^{T-1}$. For Mass Loss $\mathcal{L}_{\text{Mass}}$, we use the global mass difference as the measurement:

$$\mathcal{L}_{\text{Mass}}(\theta, \phi, \psi) = \sum_{i=0}^{T-1} \left| \sum_{j=0}^{N_0-1} \hat{w}_i^j - \frac{N_i}{N_0} \right|.$$

For OT Loss \mathcal{L}_{OT} , we approximate it by using the Wasserstein 2-distance between two normalized distributions:

$$\mathcal{L}_{\text{OT}}(\theta, \phi, \psi) = \sum_{i=0}^{T-1} \mathcal{W}_2 \left(\frac{\hat{\mathbf{w}}_i}{\|\hat{\mathbf{w}}_i\|_1}, \frac{\mathbf{w}_i}{\|\mathbf{w}_i\|_1} \right).$$

Iterative Training Strategies

To improve the robustness and efficiency of convergence, we initialize the autoencoder using PCA. In practice, the encoder f_{θ}^{enc} and decoder f_{θ}^{dec} are optimized separately to minimize the loss function of the autoencoder \mathcal{L}_{AE} .

Due to the large number of parameters and the intricate calculation of loss functions, we employ Block Coordinate Descent (BCD), which involves separating the autoencoder from the dynamic components and alternately optimizing the parameters of both (Lyu and Li 2024). Instead of training them independently, when training the autoencoder to obtain cell embeddings, we not only take into account \mathcal{L}_{AE} but also incorporate the dynamics loss function (including \mathcal{L}_{WFR} , \mathcal{L}_{OT} and $\mathcal{L}_{\text{Mass}}$) proportionally. These training strategies facilitate the development of an interpretable embedding space. The specific training strategies are as follows:

Algorithm 1: Training CellStream

Input: Time Series scRNA data $(t_i, \mathbf{X}_i)_{i=0}^{T-1}$, the dimension of embedding space d , maximum iterations n_{max} .

Initialization: θ, ϕ, ψ

while the loss function remains decreasing **do**

Optimize θ **when fixing** ϕ, ψ :

 Calculate $(t_i, \mathbf{Z}_i, \mathbf{w}_i)_{i=0}^{T-1}$ based on the input and f_{θ}^{enc}

for iter = 1 to n_{max} **do**

 Estimate $(t_i, \hat{\mathbf{Z}}_i, \hat{\mathbf{w}}_i)_{i=0}^{T-1}$ via NeuralODE using (t_0, \mathbf{Z}_0) by (3)

 Calculate the loss function: $\mathcal{L} \leftarrow \lambda_{\text{AE}} \mathcal{L}_{\text{AE}} +$

$\lambda_{\text{WFR}} \mathcal{L}_{\text{WFR}}^{\text{emb, par}} + \lambda_{\text{Match}} \mathcal{L}_{\text{Match}}$

 Update θ

end for

Optimize ϕ, ψ **when fixing** θ :

 Calculate $(t_i, \mathbf{Z}_i, \mathbf{w}_i)_{i=0}^{T-1}$ based on the input and f_{θ}^{enc}

for iter = 1 to n_{max} **do**

 Estimate $(t_i, \hat{\mathbf{Z}}_i, \hat{\mathbf{w}}_i)_{i=0}^{T-1}$ via Neural ODE using (t_0, \mathbf{Z}_0) by (3)

 Calculate the loss function: $\mathcal{L} \leftarrow \lambda_{\text{WFR}} \mathcal{L}_{\text{WFR}}^{\text{emb, par}} +$

$\lambda_{\text{Match}} \mathcal{L}_{\text{Match}}$

 Update ϕ, ψ

end for

end while

Output:

 Autoencoder f_{θ}^{enc} and f_{θ}^{dec} ,

 Embedding space dynamical function \mathbf{v}_{ϕ} and g_{ψ} .

Evaluation Metrics

Besides the visual judgment, we employ several quantitative metrics to assess the inferred embeddings. In the simulated data, where the ground truth is known, we directly calculate the velocity accuracy (VA) between the estimated velocity and the true values. For the real datasets without ground truth, we develop two consistency metrics: velocity consistency (VC) and temporal consistency (TC) as follows.

Velocity Accuracy (VA) With known ground truth, we assess predictive accuracy by computing the distance correlation (dCor) between the predicted velocity $\{\mathbf{v}_i^{\text{pred}}\}_{i=1}^N$ in the cell embedding space and the ground truth velocity $\{\mathbf{v}_i^{\text{real}}\}_{i=1}^N$ in the same space:

$$\text{VA} = \text{dCor}(\mathbf{v}^{\text{real}}, \mathbf{v}^{\text{pred}}),$$

where the predicted velocity \mathbf{v}^{pred} can be directly derived from \mathbf{v}_{ϕ} , and the real value \mathbf{v}^{real} can be calculated by using coordinate transformation formula.

Velocity Consistency (VC) We propose velocity consistency (VC) to assess velocity coherence in the embedded latent space. A higher VC is expected since it indicates a smoother velocity field which boosts better robustness in downstream analysis. Specifically, the neighborhood of each cell is first identified, and then an individual average velocity similarity and overall velocity consistency are computed:

$$\text{VC}_r = \frac{\sum_{i,j} \text{vc}_r(\mathbf{z}_i^j)}{\sum_i N_i}, \quad \text{vc}_r(\mathbf{z}_i^j) = \frac{\sum_{\mathbf{z}_k^l \in \mathcal{N}_r(\mathbf{z}_i^j)} \cos\langle \mathbf{v}(\mathbf{z}_k^l), \mathbf{v}(\mathbf{z}_i^j) \rangle}{|\{\mathbf{z}_k^l \in \mathcal{N}_r(\mathbf{z}_i^j)\}|}$$

where $\mathcal{N}_r(\mathbf{z}_i^j)$ indicates the r -distance neighborhood of \mathbf{z}_i^j , and $\mathbf{v}(\mathbf{z}_k^l)$ is the velocity vector of \mathbf{z}_k^l in the embedding space. In the numerical experiment, we set $r = 0.05$.

Temporal Consistency (TC) We propose temporal consistency (TC) to assess the degree of temporal separation in the embedding space. For each cell from time point i , we calculate the proportion of temporally matched neighbors, namely the fraction of cells within its neighborhood that also belong to time point i . TC is then computed as the average of these proportions across all cells from all time points, providing a global measure of temporal coherence in the embedding:

$$\text{TC}_r = \frac{\sum_{i,j} \text{tc}_r(\mathbf{z}_i^j)}{\sum_i N_i}, \quad \text{tc}_r(\mathbf{z}_i^j) = \frac{|\{\mathbf{z}_k^l \in \mathcal{N}_r(\mathbf{z}_i^j) \cap \mathbf{Z}_i\}|}{|\{\mathbf{z}_k^l \in \mathcal{N}_r(\mathbf{z}_i^j)\}|}$$

Experiments

Experimental Setups

Datasets. We first benchmarked CellStream on the simulated data to demonstrate its capability in capturing cellular dynamics and its effectiveness of denoising in embedding construction. The cellular dynamics simulation involves a cell population with three time points and 6 genes, which gradually diverges into two cell subpopulations over time. Subsequently, three publicly available time-series datasets were utilized: a single cell RNA sequencing (scRNA-seq) dataset, a quantitative PCR (qPCR) dataset and a spatial transcriptome (ST) dataset. The EMT dataset describes an

A549 cancer cell line treated with TGFB1 to induce the epithelial-mesenchymal transition (EMT) response (Cook and Vanderhyden 2020). The iPSC dataset captures the differentiation of induced pluripotent stem cells (iPSCs) into specific cell types, featuring a bifurcation event (Bargaje et al. 2017). Finally, the mouse organogenesis spatiotemporal transcriptomic atlas (MOSTA) dataset illustrates the spatiotemporal transcriptomic dynamics during organogenesis in C57BL/6 mouse embryos (Chen et al. 2022). The data preprocessing details are left in Appendix A.2.

Baselines and benchmarks. To compare the embeddings, we use two quantitative metrics, velocity consistency (VC) and temporal consistency (TC), to evaluate the performance of CellStream compared with other mainstream methods. Since there are no previous approaches for dynamics-informed embeddings based on unbalanced dynamical OT, we compare our algorithm with 6 baseline methods which can be divided into two categories: 1) unbalanced dynamical OT model, TIGON, coupled with four geometric embeddings including PCA, t-SNE, UMAP and diffusion maps (Sha et al. 2024); 2) dynamical embeddings generated by VeloViz (RNA velocity based) and MIOFlow (dynamical OT based) (Atta, Sahoo, and Fan 2022; Huguet et al. 2022). Since RNA velocity analysis requires unspliced counts which is not available in some single-cell datasets, we use the velocity field estimated by TIGON as the input of VeloViz to approximate RNA velocity.

Implementation Details. CellStream is constructed with PyTorch 2.6.0 and Python 3.12.6. The autoencoder is set as an encoder-decoder architecture, comprising 3 hidden layers each with a dimension of 10 and RELU activations. The output dimension of the encoder is set to 2. The velocity network and growth network both employ a hidden dimension of 10 with Tanh activations, differing in layer number with 4 and 3 hidden layers respectively. The loss weights $\{\lambda_{AE}, \lambda_{WFR}, \lambda_{Match}, \lambda_{OT}, \lambda_{Mass}\}$ are set to $\{10, 1, 5, 1, 1\}$. α and r are set to 1 and 0.05 respectively. Our model is optimized using equation (4) with the Adam algorithm. The parameters of baseline methods remain as the default.

Simulated Data with Bifurcation

We first demonstrated the performance of CellStream on simulated data with unbalanced bifurcating effects generated from a Stochastic Differential Equation (SDE) solver (Appendix A.1). We artificially added noise to mimic the technical noise of real snapshots data. CellStream successfully identifies two distinct cell populations and reconstructs their developmental trajectories from 3 artificial snapshots (Fig. 2(a)). In order to demonstrate the noise robustness of CellStream, we constructed 5 sets of simulated data with an increasing magnitude of noise, and compared the denoising effect of CellStream with other baseline methods. Fig. 2(b) shows that CellStream achieves high VA and TC despite the disturbance of high noise, while the performance of other methods is adversely affected. Though certain manifold learning embeddings like diffusion maps can attain relatively high VA scores under high noise conditions, there is a dramatic decrease in TC, indicating a challenge to interpret temporal structures accurately. A detailed comparison

is provided in Appendix B.3.

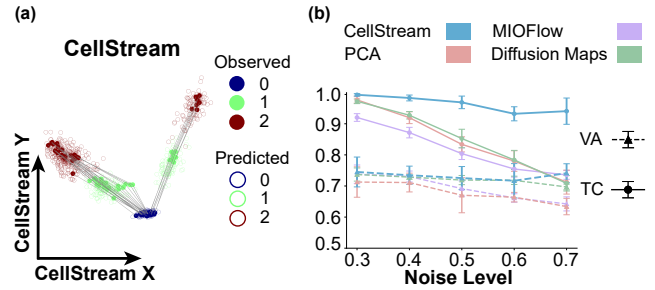


Figure 2: Application in the Simulated data. (a) The dynamics-informed embedding learned by CellStream. (b) Velocity accuracy and temporal consistency of embeddings from different methods across six runs on simulated data with increasing noise level.

Epithelial-to-mesenchymal Transition

We next applied our method to the epithelial-mesenchymal transition (EMT) dataset from an A549 cancer cell line treated with TGFB1 (Cook and Vanderhyden 2020) sampled across 4 time points. In Fig. 3(a), CellStream embeddings represent a clear cell-state transition temporal structure with highest VC and TC, aligned with the prior knowledge of EMT. In contrast, cell populations sampled from different time points overlap with each other in methods like VeloViz and MIOFlow (Fig. 3(b)(c)). While the embeddings generated by t-SNE or UMAP informed TIGON also capture such structure, their performances are inferior to CellStream in velocity consistency (Fig. 3(e)(f), Table 1). CellStream can also capture the growth dynamics consistent with relative cell population changes at each time point (Appendix B.2).

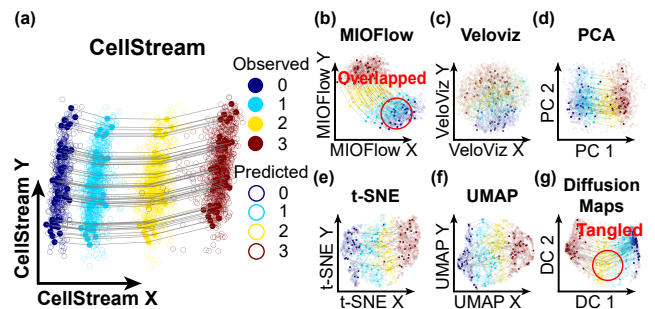


Figure 3: Application in the EMT dataset. (a) The dynamics-informed embedding learned by CellStream. (b) The embeddings learned by MIOFlow. (c) The embedding learned by VeloViz. (d-g) The embeddings from TIGON with PCA, t-SNE, UMAP and diffusion maps respectively.

Induced Pluripotent Stem Cells

We then used the induced pluripotent stem cells (iPSCs) differentiation dataset to validate the performance of CellStream in bifurcation event (Bargaje et al. 2017). Day 3 is

Method	CellStream(ours)		VeloViz		MIOFlow		TIGON							
	VC	TC	VC	TC	VC	TC	PCA		t-SNE		UMAP		Diffusion Maps	
dataset	VC	TC	VC	TC	VC	TC	VC	TC	VC	TC	VC	TC	VC	TC
EMT	0.97	0.99	0.88	0.70	0.96	0.77	0.66	0.59	0.70	0.94	0.67	0.91	0.68	0.57
iPSC	0.97	0.91	0.41	0.92	0.98	0.92	0.32	0.94	0.74	0.95	0.82	0.93	0.78	0.84
MOSTA	0.98	0.99	0.83	0.89	0.42	0.92	0.85	0.99	0.43	0.99	0.47	0.99	0.98	0.99

Table 1: Performance comparison of different trajectory inference methods on three real datasets. VC indicates velocity consistency and TC indicates temporal consistency.

the branching point when cells initially in PS states segregated into Mesodermal (M) states or Endodermal (En) states. We selected data at the first 6 time points that included the crucial bifurcation process. CellStream successfully reveals the cellular differentiation towards mesodermal or endodermal lineages which is not fully resolved by MIOFlow and TIGON with PCA (Fig. 4(a)(c)(e)). While VeloViz embeddings represent the bifurcation event, the velocity field is tangled at the branching point (Fig. 4(d)) and so do the other three embeddings(Appendix B.1).

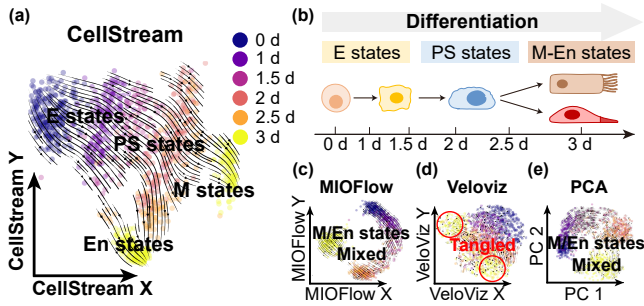


Figure 4: Application in the iPSC dataset. (a) The dynamics-informed embedding learned by CellStream. (b) Illustration of the bifurcation event in the iPSC dataset. (c) The embedding learned by MIOFlow. (d) The embedding learned by VeloViz. (e) The embedding learned by TIGON with PCA.

Spatiotemporal Analysis

Recently, the emergence of spatial transcriptomics (ST) technologies has enabled us to investigate the spatial distribution of gene expression. We further used the MOSTA (mouse organogenesis) dataset to assess the applicability of our method in ST data (Chen et al. 2022). We applied a pre-trained GNN (Dong and Zhang 2022) to integrate the gene expression matrix with spatial relationships between cells, using the resulting 50-dimensional space as the input of our model. In CellStream embedding (Fig. 5(a)), the cell population grew dramatically at the first time point and drifted forward consistently, agreeing with the prior spatiotemporal transcriptomic atlas (Appendix B.2). However, embeddings generated by MIOFlow and VeloViz do not accurately reflect the temporal structure. The four low-dimensional manifold embeddings also exhibit incoherent and disordered results compared with CellStream (Table 1, Fig. 5(d-g)).

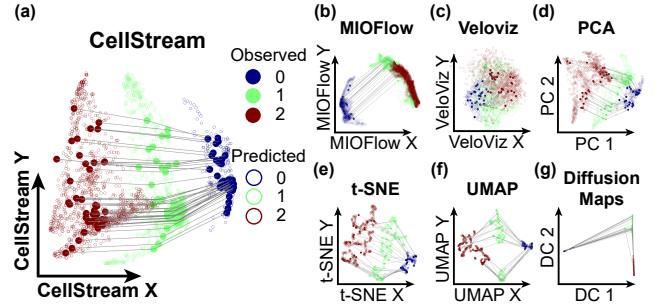


Figure 5: Application in the MOSTA dataset. (a) The dynamics-informed embedding learned by CellStream. (b) The embedding learned by MIOFlow. (c) The embedding learned by VeloViz. (d-g) The embeddings from TIGON with PCA, t-SNE, UMAP and diffusion maps respectively.

Ablation studies and hyperparameter sensitivity

We conduct ablation experiments to evaluate the contribution of different components. Specifically, we first decouple AE from the dynamics model to assess the necessity of jointly learning embeddings and cellular dynamics. Next, we dismiss the growth term to demonstrate the impact of unbalanced dynamical OT. Ablation study results are provided in Appendix B.4. Furthermore, we examine CellStream’s sensitivity to parameter variance by tuning hyperparameters including α , λ_{AE} , λ_{WFR} , λ_{Match} , λ_{OT} , λ_{Mass} and r . We include the detailed parameter sensitivity in Appendix B.5.

Conclusion

In this paper, we introduce CellStream, a new algorithm that jointly learns low-dimensional embeddings and continuous cellular stream from static, sparse snapshots data. We have demonstrated on both simulated data and real snapshots data that the dynamical embeddings derived from CellStream yield enhanced temporal coherence and improved denoising effect in revealing the underlying cellular dynamics. However, a key limitation of CellStream is that a sufficiently potent decoder architecture is yet to be developed to project dynamics back into the original gene expression space with high reconstruction fidelity. Future directions also involve providing a deeper insight into gene regulation networks from embedding space and extending the applications to various biological scenarios, like multi-omics data or cell-to-cell communication (Zhang et al. 2025b).

Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC No. 12288101, 8206100646, T2321001 to P.Z.) and The Fundamental Research Funds for the Central Universities, Peking University. We acknowledge the support from the High-performance Computing Platform of Peking University for computation.

References

- Aissa, A. F.; Islam, A. B. M. M. K.; Ariss, M. M.; Go, C. C.; Rader, A. E.; Conrardy, R. D.; and et al. 2021. Single-cell transcriptional changes associated with drug tolerance and response to combination therapies in cancer. *Nat Commun*, 12: 1628.
- Atta, L.; Sahoo, A.; and Fan, J. 2022. VeloViz: RNA velocity-informed embeddings for visualizing cellular trajectories. *Bioinformatics*, 38(2): 391–396.
- Bargaje, R.; Trachana, K.; Shelton, M.; McGinnis, C.; Zhou, J.; Chadick, C.; Cook, S.; Cavanaugh, C.; and Hood, S. H. L. 2017. Cell population structure prior to bifurcation predicts efficiency of directed differentiation in human induced pluripotent cells. *Proceedings of the National Academy of Sciences*, 114(9): 2271–2276.
- Benamou, J.-D.; and Brenier, Y. 2000. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numer.Math*, 84: 375–393.
- Bunne, C.; Schiebinger, G.; Krause, A.; Regev, A.; and Cuturi, M. 2024. Optimal transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1): 58.
- Chari, T.; and Pachter, L. 2023. The specious art of single-cell genomics. *PLOS Computational Biology*, 19(8): e1011288.
- Chen, A.; Liao, S.; Cheng, M.; Ma, K.; Wu, L.; Lai, Y.; and et al. 2022. Spatiotemporal transcriptomic atlas of mouse organogenesis using DNA nanoball-patterned arrays. *Cell*, 185(10): 1777–1792.e21.
- Chen, R. T.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural ordinary differential equations. *Advances in neural information processing systems*, 31.
- Chizat, L.; Peyré, G.; Schmitzer, B.; and Vialard, F.-X. 2018. An Interpolating Distance Between Optimal Transport and Fisher–Rao Metrics. *Found. Comput. Math*, 18: 1–44.
- Chizat, L.; Peyré, G.; Schmitzer, B.; and Vialard, F.-X. 2019. Unbalanced Optimal Transport: Dynamic and Kantorovich Formulation. arXiv:1508.05216.
- Coifman, R. R.; and Lafon, S. 2006. Diffusion maps. *Applied and Computational Harmonic Analysis*, 5–30.
- Cook, D. P.; and Vanderhyden, B. C. 2020. Context specificity of the EMT transcriptional response. *Nature Communications*, 11: 2142.
- Dong, K.; and Zhang, S. 2022. Deciphering spatial domains from spatially resolved transcriptomics with an adaptive graph attention auto-encoder. *Nature Communications*, 13: 1739.
- Gao, M.; Qiao, C.; and Huang, Y. 2022. UniTVelo: temporally unified RNA velocity reinforces single-cell trajectory inference. *Nat Commun*, 13: 6586.
- Ge, M.; Miao, J.; Qi, J.; Zhou, X.; and Lin, Z. 2025. TIVelo: RNA velocity estimation leveraging cluster-level trajectory inference. *Nat Commun*, 16: 6258.
- Gu, A.; Chien, E.; and Greenewald, K. 2025. Partially Observed Trajectory Inference using Optimal Transport and a Dynamics Prior. In *The Thirteenth International Conference on Learning Representations*.
- Hou, W.; Ji, Z.; Hongkai, J.; and Hicks, S. 2020. A systematic evaluation of single-cell RNA-sequencing imputation methods. *Genome biology*, 21: 1–30.
- Huang, Y.-A.; Li, Y.-C.; You, Z.-H.; Hu, L.; Hu, P.-W.; Wang, L.; and et al. 2025. Consensus representation of multiple cell-cell graphs from gene signaling pathways for cell type annotation. *BMC Biol*, 23: 23.
- Huguet, G.; Magruder, D. S.; Tong, A.; Fasina, O.; Kuchroo, M.; Wolf, G.; and Krishnaswamy, S. 2022. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in neural information processing systems*, 35: 29705–29718.
- Inecik, K.; Kara, A.; Rose, A.; Haniffa, M.; and Theis, F. J. 2025. TarDis: Achieving Robust and Structured Disentanglement of Multiple Covariates. In Sankararaman, S., ed., *Research in Computational Molecular Biology*, 285–289. Cham: Springer Nature Switzerland.
- Kantorovich, L. 1942. On the transfer of masses (in russian). *Doklady Akademii Nauk*, 37(2): 227–229.
- Kapusniak, K.; Potaptchik, P.; Reu, T.; Zhang, L.; Tong, A.; Bronstein, M.; Bose, J.; and Di Giovanni, F. 2024. Metric flow matching for smooth interpolations on the data manifold. *Advances in Neural Information Processing Systems*, 37: 135011–135042.
- Klein, D.; Palla, G.; Lange, M.; Klein, M.; Piran, Z.; and et al., M. G. 2025. Mapping cells through time and space with moscot. *Nature*, 1–11.
- Korsunsky, I.; Millard, N.; Fan, J.; Slowikowski, K.; Zhang, F.; Wei, K.; and et al. 2019. Fast, sensitive and accurate integration of single-cell data with Harmony. *Nat Methods*, 16: 1289–1296.
- Lavenant, H.; Zhang, S.; Kim, Y.-H.; Schiebinger, G.; et al. 2024. Toward a mathematical theory of trajectory inference. *The Annals of Applied Probability*, 34(1A): 428–500.
- Li, T.; Wang, Y.; Yang, G.; and Zhou, P. 2023. On the Mathematics of RNA Velocity II: Algorithmic Aspects. *bioRxiv*.
- Lin, P.; Troup, M.; and Ho, J. 2017. CIDR: Ultrafast and accurate clustering through imputation for single-cell RNA-seq data. *Genome biology*, 18: 1–11.
- Lopez, R.; Regier, J.; Cole, M. B.; Jordan, M. I.; and Yosef, N. 2018. Deep generative modeling for single-cell transcriptomics. *Nat Methods*, 15: 1053–1058.
- Lyu, H.; and Li, Y. 2024. Block majorization-minimization with diminishing radius for constrained nonsmooth nonconvex optimization. arXiv:2012.03503.

- Manno, G. L.; Soldatov, R.; Zeisel, A.; Braun, E.; Hochgerner, H.; Petukhov, V.; and et al. 2018. RNA velocity of single cells. *Nature*, 560: 494–498.
- McInnes, L.; Healy, J.; Saul, N.; and Großberger, L. 2018. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 3: 861.
- Mongia, A.; Sengupta, D.; and Majumdar, A. 2019. McImpute: matrix completion based imputation for single cell RNA-seq data. *Frontiers in genetics*, 10: 9.
- Pariset, M.; Hsieh, Y.-P.; Bunne, C.; Krause, A.; and Bortoli, V. D. 2023. Unbalanced Diffusion Schrödinger Bridge. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*.
- Peng, Q.; Zhou, P.; and Li, T. 2025. stVCR: Spatiotemporal dynamics of single cells. *bioRxiv*.
- Petrović, K.; Atanackovic, L.; Kapusniak, K.; Bronstein, M. M.; Bose, J.; and Tong, A. 2025. Curly Flow Matching for Learning Non-gradient Field Dynamics. In *Learning Meaningful Representations of Life (LMRL) Workshop at ICLR 2025*.
- Raslan, A. A.; Pham, T. X.; Lee, J.; Kontodimas, K.; Tilston-Lunel, A.; Schmottlach, J.; and et al. 2024. Lung injury-induced activated endothelial cell states persist in aging-associated progressive fibrosis. *Nat Commun*, 15: 5449.
- Rutkowski, P.; and Tabaka, M. 2025. Ocelli: an open-source tool for the analysis and visualization of developmental multimodal single-cell data. *NAR Genomics and Bioinformatics*, 7: lqaf040.
- Schiebinger, G.; Shu, J.; Tabaka, M.; Cleary, B.; Subramanian, V.; Solomon, A.; and et al. 2019. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4): 928–943.
- Sha, Y.; Qiu, Y.; Zhou, P.; and Nie, Q. 2024. Reconstructing growth and dynamic trajectories from single-cell transcriptomics data. *Nature Machine Intelligence*, 6(1): 25–39.
- Sun, Y.; Zhang, Z.; Wang, Z.; Li, T.; and Zhou, P. 2025. Variational Regularized Unbalanced Optimal Transport: Single Network, Least Action. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Theodoris, T. V.; Xiao, L.; Chopra, A.; Chaffin, M. D.; Sayed, Z. R. A.; Hill, M. C.; and et al. 2023. Transfer learning enables predictions in network biology. *Nature*, 618: 616–624.
- Tong, A.; FATRAS, K.; Malkin, N.; Huguét, G.; Zhang, Y.; Rector-Brooks, J.; Wolf, G.; and Bengio, Y. 2024a. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*. Expert Certification.
- Tong, A.; Huang, J.; Wolf, G.; Dijk, D. V.; and Krishnaswamy, S. 2020. TrajectoryNet: A Dynamic Optimal Transport Network for Modeling Cellular Dynamics. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 9526–9536. PMLR.
- Tong, A.; Malkin, N.; Fatras, K.; Atanackovic, L.; Zhang, Y.; Huguét, G.; Wolf, G.; and Bengio, Y. 2024b. Simulation-Free Schrödinger Bridges via Score and Flow Matching. In *International Conference on Artificial Intelligence and Statistics*, 1279–1287. PMLR.
- van der Maaten, L.; and Hinton, G. E. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9: 2579–2605.
- Wang, D.; Jiang, Y.; Zhang, Z.; Gu, X.; Zhou, P.; and Sun, J. 2025. Joint Velocity-Growth Flow Matching for Single-Cell Dynamics Modeling. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Wilks, D. 2011. Chapter 13 - Canonical Correlation Analysis (CCA). In Wilks, D. S., ed., *Statistical Methods in the Atmospheric Sciences*, volume 100 of *International Geophysics*, 563–582. Academic Press.
- Wolf, F. A.; Hamey, F. K.; Plass, M.; Solana, J.; Dahlin, J. S.; Göttgens, B.; Rajewsky, N.; Simon, L. M.; and Theis, F. J. 2019. PAGA: graph abstraction reconciles clustering with trajectory inference through a topology preserving map of single cells. *Genome Biol*, 20: 59.
- Zhang, S.; Afanassiev, A.; Greenstreet, L.; Matsumoto, T.; and Schiebinger, G. 2021. Optimal transport analysis reveals trajectories in steady-state systems. *PLoS Comput Biol*, (12).
- Zhang, Z.; Li, T.; and Zhou, P. 2025. Learning stochastic dynamics from snapshots through regularized unbalanced optimal transport. In *The Thirteenth International Conference on Learning Representations*.
- Zhang, Z.; Sun, Y.; Peng, Q.; Li, T.; and Zhou, P. 2025a. Integrating Dynamical Systems Modeling with Spatiotemporal scRNA-Seq Data Analysis. *Entropy*, 27(5).
- Zhang, Z.; Wang, Z.; Sun, Y.; Li, T.; and Zhou, P. 2025b. Modeling Cell Dynamics and Interactions with Unbalanced Mean Field Schrödinger Bridge. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Zhang, Z.; Wang, Z.; Sun, Y.; Shen, J.; Q. Peng, Q.; Li, T.; and Pei, Z. 2025c. Deciphering cell-fate trajectories using spatiotemporal single-cell transcriptomic data. *Preprint*.
- Zhang, Z.; and Zhang, X. 2021. Inference of high-resolution trajectories in single-cell RNA-seq data by using RNA velocity. *Cell Reports Methods*, 1.
- Zheng, S. C.; Stein-O’Brien, G.; Boukas, L.; Goff, L. A.; and Hansen, K. D. 2023. Pumping the brakes on RNA velocity by understanding and interpreting RNA velocity estimates. *Genome biology*, 24(1): 246.
- Zhou, P.; Bocci, F.; Li, T.; and Nie, Q. 2024. Spatial transition tensor of single cells. *Nature Methods*, 21(6): 1053–1062.
- Zhou, P.; Wang, S.; Li, T.; and Nie, Q. 2021. Dissecting transition cells from single-cell transcriptome data through multiscale stochastic dynamics. *Nat Commun*, 12: 5609.