

TracingFlow: A Simulation-Free Trajectory Inference Framework Based on Second-Order Dynamics

Yuhao Sun^{1*}, Zekun Wu^{2*}, Zixun Huang², Peijie Zhou^{1,3,4,5} [†]

¹Center for Machine Learning Research, Peking University

²School of Mathematical Sciences, Peking University

³Center for Quantitative Biology, Peking University

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

⁵AI for Science Institute, Beijing

Abstract

Inferring continuous system evolution from sparse temporal snapshots is a key challenge in generative modeling and single-cell omics. While Optimal Transport (OT) is popular, existing frameworks are largely restricted to first-order dynamics, assuming memoryless velocity fields. This limits expressiveness, as first-order systems fail to account for regulatory momentum and time-delayed responses inherent in processes like cell differentiation. Here, we introduce TracingFlow, a simulation-free Flow Matching framework generalizing to second-order dynamics. By using neural networks to regress the acceleration field, TracingFlow provides an exact, efficient solution to the Dynamical Optimal Acceleration Transport (DOAT) problem. Unlike first-order methods yielding over-smoothed trajectories, our second-order formulation captures high-curvature transitions and nonlinear evolutions by learning the underlying force fields. Evaluated on complex synthetic and large-scale scRNA-seq datasets, TracingFlow achieves superior accuracy in distributional reconstruction and trajectory faithfulness. Moreover, by integrating lineage tracing priors, it recovers dynamical structures that are both mathematically optimal and biologically plausible.

1 Introduction

Recovering underlying dynamics from discrete observations is a pivotal task in single-cell omics, known as Trajectory Inference (TI) [9, 72]. To identify the least costly evolution between distributions, Optimal Transport (OT) frameworks have emerged [7, 71]. These generally categorize into Static OT methods [22, 28, 51], which learn direct mappings, and Dynamical OT methods [24, 44, 57, 67], which capture continuous evolution via flow maps.

Early Dynamical OT frameworks relied on Neural ODEs [9], incurring high computational costs due to numerical integration. To address this, Flow Matching [35, 36, 48] was proposed as a simulation-free alternative, regressing velocity fields to push source distributions to targets. By decomposing transport costs into single-particle trajectories, Flow Matching efficiently solves Dynamical OT problems [27, 50, 58], including its various variants [8, 15, 45, 59].

However, most TI frameworks assume first-order dynamics ($\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}, t)$), inherently restricting $\mathbf{v}$ to be single-valued w.r.t $\mathbf{x}$. This limits expressiveness when modeling complex biological priors, such as lineage tracing [30, 38], potentially yielding dynamics that contradict with additional biological prior information. Furthermore, [19, 23, 34] pointed out that if more modalities

*Equal contribution

[†]Corresponding author: pjzhou@pku.edu.cn




Figure 1: An illustration figure for TracingFlow

such as chromatin accessibility or proteomics are considered, the biological kinetics may follow higher-order dynamics or could be modelled in augmented space. While 3MSBM [56] recently introduced a second-order Momentum Schrödinger Bridge for smoother trajectories, it relies on iterative retraining the acceleration field similar to Rectified Flow [36] to determine optimal couplings. Furthermore, it employs a heuristic method to estimate the initial velocity, rather than incorporating this estimation as part of the optimal control problem.

To overcome these limitations, we introduce **TracingFlow**, a simulation-free framework utilizing second-order dynamical systems. We propose the Dynamical Optimal Acceleration Transport (DOAT) problem, which determines evolution by minimizing acceleration costs. TracingFlow directly regresses the acceleration field and initial velocity, avoiding ODE simulation. Our experiments show that TracingFlow achieves competitive and often improved reconstruction accuracy relative to existing simulation-free frameworks [42, 58, 59, 69], while effectively incorporating biological priors. Our contributions include:

- We propose TracingFlow, a simulation-free TI framework based on second-order dynamics. It solves the proposed DOAT problem by regressing acceleration fields, reducing computational costs compared to ODE-based methods.
- We provide theoretical guarantees by decoupling the design of the optimal transport plan and single-particle control. We further introduce an iterative strategy to transform the real-world Velocity Missed-DOAT (VM-DOAT) problem into a standard DOAT problem, enabling approximate solutions.
- We demonstrate TracingFlow’s effectiveness on multi-time-point real world datasets. Results indicate superior distribution reconstruction accuracy and enhanced capability in preserving biological priors compared to state-of-the-art methods.

2 Related Works

Solving Optimal Transport via Flow Matching. Flow Matching [1, 35, 36, 48] transports probability distributions via learned flow maps, offering scalable simulation-free training. As an optimal control problem, Optimal Transport (OT) [47] decomposes into Optimal Coupling and single-particle geodesics, facilitating solutions via Flow Matching [15, 27, 50, 58]. This approach

extends to various variants and generalizations [4, 8, 13, 15, 26, 45, 46, 59, 61, 68]. However, these typically employ first-order dynamics. While second-order approaches have emerged [56], they require iterative training of the acceleration field to achieve optimal coupling. Moreover, they rely on heuristic methods for initial velocity estimation, rather than integrating it as a component of the optimal control problem. TracingFlow addresses this by decomposing the Dynamical Acceleration Optimal Transport (DOAT) problem into optimal coupling and single particle optimal trajectory, achieving a truly simulation-free process for DOAT.

Overcoming Single-Valued Velocity Constraint in Flow Matching. In standard Flow Matching, the velocity field is modeled as a function solely dependent on $\mathbf{x}$ and t . This single-valued dependency prevents the model from representing intersecting trajectories with distinct velocities, often resulting in curved inference paths and increased computational costs. To decouple the velocity from the current position, the velocity network can incorporate additional inputs like class labels [54, 73], initial positions [14], or hidden states [20]. Additionally, [69] employs hierarchical generation to mitigate this issue. However, these methods do not extend to second-order dynamics. While [42] utilized a second-order system, it is limited to constant-acceleration paths. In contrast, TracingFlow offers a flexible second-order setting. By designing minimal-cost paths, it effectively resolves the limitation of the single-valued velocity field.

Single-cell Trajectory Inference and Lineage Integration. Optimal Transport (OT) is a robust framework for inferring cellular dynamics from scRNA-seq data [2, 6, 10–12, 24, 25, 28, 29, 33, 37, 40, 41, 44, 51–53, 55, 57, 60, 65, 66, 70, 74]. However, transcriptomic similarity alone often fails to resolve complex trajectories. To mitigate this, studies incorporate lineage tracing priors such as clonal barcode information. Existing approaches integrate such info via regularization [17, 49], structural alignment [32], velocity mapping [62], or sparse transition modeling [18, 21, 63]. While improving accuracy, these generally focus on discrete couplings or static matrices. TracingFlow advances this by embedding priors into a continuous OT-based second-order flow matching framework, capturing complex dynamics consistent with lineage information.

3 Mathematical Background

In this section, we provide the mathematical formulation of the problem addressed by TracingFlow.

Dynamical Optimal Acceleration Transport (DOAT) Problem Let $\mathcal{X} \subset \mathbb{R}^d$ and $\mathcal{V} \subset \mathbb{R}^d$ denote the *position* and *velocity* spaces, respectively. We define the *augmented space* as $\mathcal{S} := \mathcal{X} \times \mathcal{V}$, so that the full system state is $\mathbf{s} = (\mathbf{x}, \mathbf{v}) \in \mathcal{S}$. Consider the second-order dynamic system:

$$\dot{\mathbf{x}} = \mathbf{v} \quad \dot{\mathbf{v}} = \mathbf{a}(\mathbf{x}, \mathbf{v}, t), \quad t \in [t_0, t_K]. \quad (1)$$

Modeling a large ensemble of such particles via a probability density ρ_t in $\mathcal{S}$, we follow [5, 11] to formulate the following optimal control problem:

$$\min_{\mathbf{a}, \rho} \mathcal{J}_{\text{DOAT}}(\mathbf{a}, \rho) = \int_{t_0}^{t_K} \int_{\mathcal{S}} \frac{1}{2} \|\mathbf{a}(\mathbf{x}, \mathbf{v}, t)\|^2 \rho_t(\mathbf{x}, \mathbf{v}) \, d\mathbf{x} d\mathbf{v} dt \quad (2)$$

$$\text{s.t.} \quad \partial_t \rho_t + \nabla_{\mathbf{x}} \cdot (\mathbf{v} \rho_t) + \nabla_{\mathbf{v}} \cdot (\mathbf{a} \rho_t) = 0, \quad (3)$$

$$\rho_{t_i}(\mathbf{x}, \mathbf{v}) = \mu_i(\mathbf{x}, \mathbf{v}), \quad i = 0, \dots, K. \quad (4)$$

Here, constraints are imposed at $K + 1$ timestamps t_i with given densities μ_i . Equation (3) represents the continuity equation in the augmented space. Analogous to standard Dynamical

Optimal Transport [47], we term this the *Dynamical Optimal Acceleration Transport (DOAT)* problem. To facilitate the subsequent numerical solution via particle methods, we assume the existence of an optimal flow map.

Velocity-Missed Dynamical Optimal Acceleration Transport (VM-DOAT) Problem

However, in practical applications such as single-cell datasets, only particle positions are observable, rendering velocities as unknown quantities. This presents a discrepancy with the DOAT formulation.

Consequently, TracingFlow addresses a relaxation of the problem:

$$\min_{\mathbf{a}, \rho} \mathcal{J}_{\text{VM-DOAT}}(\mathbf{a}, \rho) = \int_{t_0}^{t_K} \int_S \frac{1}{2} \|\mathbf{a}(\mathbf{x}, \mathbf{v}, t)\|^2 \rho_t(\mathbf{x}, \mathbf{v}) \, d\mathbf{x} d\mathbf{v} dt \quad (5)$$

$$\text{s.t.} \quad \partial_t \rho_t(\mathbf{x}, \mathbf{v}) + \nabla_{\mathbf{x}} \cdot (\mathbf{v} \rho_t) + \nabla_{\mathbf{v}} \cdot (\mathbf{a} \rho_t) = 0, \quad (6)$$

$$\int \rho_{t_i}(\mathbf{x}, \mathbf{v}) \, d\mathbf{v} = \mu_i^{(\text{pos})}(\mathbf{x}), \quad i = 0, 1, \dots, K. \quad (7)$$

Here, the constraints are imposed only on the marginal distribution of positions, $\int \rho_{t_j}(\mathbf{x}, \mathbf{v}) \, d\mathbf{v}$, at each time point, with no constraints on the velocity distribution. We refer to this problem as *Velocity Missed-Dynamical Optimal Acceleration Transport (VM-DOAT)*.

4 Methodology of TracingFlow

In this section, we first introduce the key properties of the DOAT problem that are essential for our algorithm design. Subsequently, drawing inspiration from Flow Matching approaches for standard Dynamical OT, we devise an algorithm to exactly solve the DOAT problem by combining Optimal Coupling with single-particle optimal trajectories. Finally, we introduce a global path-based velocity completion scheme that effectively bridges the gap between VM-DOAT and the standard DOAT framework. By inferring missing velocities through the lens of global trajectory consistency, our approach enables the application of optimal transport theory to solve the VM-DOAT problem in a principled manner.

4.1 Key Properties of DOAT Problem

Optimal Single Particle Trajectory The DOAT problem is an optimal control problem over distributions. We begin by considering the corresponding single-particle optimal control problem. Specifically, the optimal trajectory is a *cubic interpolation* determined by the boundary states.

Proposition 4.1. *Let $\gamma : [0, T] \rightarrow \mathbb{R}^d$ be a twice continuously differentiable curve satisfying the boundary conditions $\gamma(0) = \mathbf{x}_0$, $\gamma'(0) = \mathbf{v}_0$, $\gamma(T) = \mathbf{x}_T$, and $\gamma'(T) = \mathbf{v}_T$. The variational problem*

$$\min_{\gamma} \int_0^T \frac{1}{2} \|\gamma''(t)\|_2^2 \, dt \quad (8)$$

admits a unique minimizer, which is a cubic polynomial of the form

$$\gamma(t) = c_3 t^3 + c_2 t^2 + c_1 t + c_0, \quad (9)$$

where the coefficients are detailed in Section A.1. This result is got by Euler-Lagrange formula. See a detailed proof in Section A.1.

Optimal Coupling By utilizing the single-particle optimal trajectory, we can immediately obtain the minimum cost for transporting a particle from $(\mathbf{x}_0, \mathbf{v}_0)$ to $(\mathbf{x}_T, \mathbf{v}_T)$:

Corollary 4.2. *Substitute Equation (9) to Equation (8), we get the minimum cost :*

$$\mathcal{C}_{[0,T]} = \frac{2}{T^3} \left\{ (\|\mathbf{v}_0\|^2 + \langle \mathbf{v}_0, \mathbf{v}_T \rangle + \|\mathbf{v}_T\|^2) T^2 + 3\langle \mathbf{v}_0 + \mathbf{v}_T, \mathbf{x}_0 - \mathbf{x}_T \rangle T + 3\|\mathbf{x}_0 - \mathbf{x}_T\|^2 \right\} \quad (10)$$

The minimum cost allows us to define an optimal transport problem between $t = 0$ and $t = T$:

$$\pi_{[0,T]}^* = \min_{\pi_{[0,T]}} \int \mathcal{C}_{[0,T]}[(\mathbf{x}_0, \mathbf{v}_0), (\mathbf{x}_T, \mathbf{v}_T)] d\pi_{[0,T]} \quad (11)$$

where $\pi_{[0,T]}[(\mathbf{x}_0, \mathbf{v}_0), (\mathbf{x}_T, \mathbf{v}_T)]$ is any coupling satisfying the constraints $\int \pi_{[0,T]} d\mathbf{x}_0 d\mathbf{v}_0 = \mu_T(\mathbf{x}, \mathbf{v})$ and $\int \pi_{[0,T]} d\mathbf{x}_T d\mathbf{v}_T = \mu_0(\mathbf{x}, \mathbf{v})$, $\mu_0(\mathbf{x}, \mathbf{v})$ and $\mu_T(\mathbf{x}, \mathbf{v})$ are two distributions in augmented space. We term this formulation the **Static Optimal Acceleration Transport (SOAT)** problem.

It is worth noting that while DOAT permits crossings in the projected position space, it prohibits trajectory crossing in the augmented space, which guarantees the neatness of trajectories. The following remark provides the mathematical justification for this non-crossing property in the augmented space, demonstrating that any such intersection implies a strictly higher transport cost.

Remark 4.3. Consider two pairs of points in $\mathbb{R}^{2d}$, $(\mathbf{x}_0, \mathbf{v}_0), (\mathbf{x}_T, \mathbf{v}_T)$ and $(\tilde{\mathbf{x}}_0, \tilde{\mathbf{v}}_0), (\tilde{\mathbf{x}}_T, \tilde{\mathbf{v}}_T)$ with no points overlapping. Denote the cubic interpolation curves between each pair by $\gamma(t), \tilde{\gamma}(t) (t \in [0, T])$, respectively. If the curves cross with each other at time point $t_0 \in (0, T)$, then we have

$$\mathcal{C}_{[0,T]}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}_T, \mathbf{v}_T) + \mathcal{C}_{[0,T]}(\tilde{\mathbf{x}}_0, \tilde{\mathbf{v}}_0, \tilde{\mathbf{x}}_T, \tilde{\mathbf{v}}_T) > \mathcal{C}_{[0,T]}(\mathbf{x}_0, \mathbf{v}_0, \tilde{\mathbf{x}}_T, \tilde{\mathbf{v}}_T) + \mathcal{C}_{[0,T]}(\tilde{\mathbf{x}}_0, \tilde{\mathbf{v}}_0, \mathbf{x}_T, \mathbf{v}_T).$$

See a detailed proof in Section A.2.

Equation (2) represents a multi-marginal optimal transport problem over the interval $[t_0, t_K]$. To address this, we formulate sequential SOAT problems between adjacent time steps t_i and t_{i+1} by applying the substitutions $\mathbf{x}_0 \rightarrow \mathbf{x}_i$, $\mathbf{x}_T \rightarrow \mathbf{x}_{i+1}$, $\mathbf{v}_0 \rightarrow \mathbf{v}_i$, $\mathbf{v}_T \rightarrow \mathbf{v}_{i+1}$, and $T \rightarrow t_{i+1} - t_i$ to Equation (11). For simplicity, let $\mathcal{C}_{i \rightarrow i+1}$ denote the resulting pairwise cost and $\pi_{i \rightarrow i+1}^*$ the optimal plan. Due to the Markovian nature of the dynamics, the full multi-marginal plan decomposes into the sequence of these local plans $\{\pi_{i \rightarrow i+1}^*\}$ for $i = 0, \dots, K-1$.

4.2 Conditional and Marginal Probability Path

Obtaining Marginal Probability Path via Mixing Conditional Probability Paths

Similar to the approach in standard Flow Matching [35, 58], directly constructing the Marginal Probability Path is difficult. Therefore, we condition on $\mathbf{z}$ and define the Marginal Probability Path as a mixture of Conditional Probability Paths:

$$\rho_t(\mathbf{x}, \mathbf{v}) = \int \rho_t(\mathbf{x}, \mathbf{v} | \mathbf{z}) q(\mathbf{z}) d\mathbf{z} \quad (12)$$

The condition $\mathbf{z}$ consists of $K+1$ points selected respectively from the snapshot data at $K+1$ time points $t = 0, 1, \dots, K$. If $\rho_t(\mathbf{x}, \mathbf{v} | \mathbf{z})$ is generated by a conditional acceleration field $\mathbf{a}(\mathbf{x}, \mathbf{v}, t | \mathbf{z})$ starting from the initial condition $\rho_0(\mathbf{x}, \mathbf{v} | \mathbf{z})$, then the marginal acceleration field:

$$\mathbf{a}(\mathbf{x}, \mathbf{v}, t) := \mathbb{E}_{q(\mathbf{z})} \left\{ \frac{\mathbf{a}(\mathbf{x}, \mathbf{v}, t | \mathbf{z}) \rho_t(\mathbf{x}, \mathbf{v} | \mathbf{z})}{\rho_t(\mathbf{x}, \mathbf{v})} \right\} \quad (13)$$

will generate the marginal probability path $\rho_t(\mathbf{x}, \mathbf{v})$.

Theorem 4.4. *The marginal acceleration field in Eq. (13) generates the marginal probability path. See the proof in Section A.3.*

Equivalence between Regressing Conditional Acceleration Fields and Marginal Acceleration Fields Assuming the marginal acceleration field $\mathbf{a}(\mathbf{x}, \mathbf{v}, t)$ is known and we can sample from the marginal probability path $\rho_t(\mathbf{x}, \mathbf{v})$, we can directly regress $\mathbf{a}(\mathbf{x}, \mathbf{v}, t)$ using a neural network. Let $\mathbf{a}_\theta(\cdot, \cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^1 \rightarrow \mathbb{R}^d$ be a time-dependent acceleration field parameterized by a neural network with parameters θ . The acceleration flow matching loss is

$$\mathcal{L}_{\text{AFM}} = \mathbb{E}_{(\mathbf{x}, \mathbf{v}) \sim \rho_t(\mathbf{x}, \mathbf{v}), t \sim \mathcal{U}[t_0, t_K]} \left\{ \|\mathbf{a}_\theta - \mathbf{a}(\mathbf{x}, \mathbf{v}, t)\|^2 \right\}$$

However, $\mathbf{a}(\mathbf{x}, \mathbf{v}, t)$ is generally intractable because it is defined via an expectation (integral), and its denominator $\rho_t(\mathbf{x}, \mathbf{v})$ also requires integration to compute. In contrast, the conditional acceleration field $\mathbf{a}(\mathbf{x}, \mathbf{v}, t|z)$ and the conditional probability path $\rho_t(\mathbf{x}, \mathbf{v}|z)$ have simple forms. Therefore, in actual training, we use the following Conditional Acceleration Flow Matching (CAFM) objective:

$$\mathcal{L}_{\text{CAFM}} = \mathbb{E}_{\substack{(\mathbf{x}, \mathbf{v}) \sim \rho_t(\mathbf{x}, \mathbf{v}|z), \\ t \sim \mathcal{U}[t_0, t_K], z \sim q(z)}} \left[\|\mathbf{a}_\theta(\mathbf{x}, \mathbf{v}, t) - \mathbf{a}(\mathbf{x}, \mathbf{v}, t|z)\|^2 \right]$$

These two objectives are equivalent in training, as described by the following theorem:

Theorem 4.5. *If $\rho_t(\mathbf{x}, \mathbf{v}) > 0$ for any $\mathbf{x} \in \mathbb{R}^d, \mathbf{v} \in \mathbb{R}^d, t \in [t_0, t_K]$, then $\mathcal{L}_{\text{AFM}}$ and $\mathcal{L}_{\text{CAFM}}$ differ only by a constant independent of θ . In other words, $\nabla_\theta \mathcal{L}_{\text{AFM}} = \nabla_\theta \mathcal{L}_{\text{CAFM}}$. The proof is left to Section A.4.*

Design of the Conditional Probability Path Consider the joint distributions $\mu_0(\mathbf{x}, \mathbf{v}), \mu_1(\mathbf{x}, \mathbf{v}), \dots, \mu_K(\mathbf{x}, \mathbf{v})$ at $K + 1$ time points. The optimal transport plans calculated via solving SOAT problem are $\pi_{0 \rightarrow 1}^*, \pi_{1 \rightarrow 2}^*, \dots, \pi_{K-1 \rightarrow K}^*$. Taking the condition variable $\mathbf{z} = [(\mathbf{x}_0, \mathbf{v}_0), \dots, (\mathbf{x}_K, \mathbf{v}_K)]$, and similar to standard Flow Matching, we choose a path with time-varying mean and invariant variance as the conditional probability path, i.e.:

$$\rho_t(\mathbf{x}, \mathbf{v}|z) = \mathcal{N}(\mathbf{x}|\mathbf{m}_x(t), \sigma_x^2 I) \cdot \mathcal{N}(\mathbf{v}|\mathbf{m}_v(t), \sigma_v^2 I) \quad (14)$$

Here, $\mathbf{m}_x, \mathbf{m}_v$ are given by the Minimum Cost Trajectory described in Proposition 4.1, where for $t \in [t_i, t_{i+1}]$:

$$\mathbf{m}_x(t) = c_3(t - t_i)^3 + c_2(t - t_i)^2 + c_1(t - t_i) + c_0 \quad \mathbf{m}_v(t) = 3c_3(t - t_i)^2 + 2c_2(t - t_i) + c_1$$

, the coefficients are identical to those given in Section A.1, with $\mathbf{x}_0, \mathbf{v}_0$ replaced by $\mathbf{x}_i, \mathbf{v}_i$, $\mathbf{x}_T, \mathbf{v}_T$ replaced by $\mathbf{x}_{i+1}, \mathbf{v}_{i+1}$, and T replaced by $t_{i+1} - t_i$. Based on the optimal transport plans, we define the optimal propagator (transition kernel) from time t_i to t_{i+1} as the probability of a particle at $(\mathbf{x}_i, \mathbf{v}_i)$ at time t_i being transported to $(\mathbf{x}_{i+1}, \mathbf{v}_{i+1})$ at time t_{i+1} :

$$\mathcal{K}_{i \rightarrow i+1}^*[(\mathbf{x}_i, \mathbf{v}_i), (\mathbf{x}_{i+1}, \mathbf{v}_{i+1})] = \frac{\pi_{i \rightarrow i+1}^*[(\mathbf{x}_i, \mathbf{v}_i), (\mathbf{x}_{i+1}, \mathbf{v}_{i+1})]}{\mu_i(\mathbf{x}_i, \mathbf{v}_i)} \quad (15)$$

Then,

$$\mathbf{z} \sim \mu_0(\mathbf{x}_0, \mathbf{v}_0) \cdot \prod_{i=0}^{K-1} \mathcal{K}_{i \rightarrow i+1}^*[(\mathbf{x}_i, \mathbf{v}_i), (\mathbf{x}_{i+1}, \mathbf{v}_{i+1})]$$

In the following, we demonstrate that the constructed marginal probability path recovers the true distribution at every given time point.

Theorem 4.6. *When $\sigma_x \rightarrow 0, \sigma_v \rightarrow 0$, the marginal probability density constructed by Equation (12) satisfies $\rho_{t_i}(\mathbf{x}, \mathbf{v}) \rightarrow \mu_i(\mathbf{x}, \mathbf{v})$ at the $K + 1$ given time points. Therefore, the Marginal Probability Path constructed by mixing Conditional Probability Paths can correctly reconstruct the marginal distributions. See the proof in Section A.5.*

Furthermore, we demonstrate that the design yields an exact solution to the DOAT problem.

Proposition 4.7. *As $\sigma_x \rightarrow 0$ and $\sigma_v \rightarrow 0$, the marginal probability path and the marginal acceleration field will converge to the solution to the DOAT problem. See the proof in Section A.6.*

4.3 Transform VM-DOAT to DOAT

Relations between DOAT and VM-DOAT The optimal transport cost of the DOAT problem is a functional of the sequence of joint distributions $\mu_i(\mathbf{x}, \mathbf{v})$ for $i = 0, \dots, K$. Formally, we denote this dependency as $\mathcal{J}_{\text{DOAT}}[\{\mu_i(\mathbf{x}, \mathbf{v})\}_{i=0}^K]$. In the VM-DOAT setting, we only have access to the marginal spatial distributions $\mu_i^{(\text{pos})}(\mathbf{x}) = \int \mu_i(\mathbf{x}, \mathbf{v}) d\mathbf{v}$. Consequently, the cost of VM-DOAT, denoted as $\mathcal{J}_{\text{VM-DOAT}}[\{\mu_i^{(\text{pos})}(\mathbf{x})\}_{i=0}^K]$, relates to $\mathcal{J}_{\text{DOAT}}$ through the following minimization:

$$\mathcal{J}_{\text{VM-DOAT}}[\{\mu_i^{(\text{pos})}(\mathbf{x})\}] = \min_{\omega_i(\mathbf{v}|\mathbf{x})} \mathcal{J}_{\text{DOAT}}[\{\mu_i^{(\text{pos})}(\mathbf{x}) \cdot \omega_i(\mathbf{v}|\mathbf{x})\}], \quad (16)$$

where $\omega_i(\mathbf{v}|\mathbf{x})$ represents the conditional distribution of velocity, such that the joint distribution factorizes as $\mu_i(\mathbf{x}, \mathbf{v}) = \mu_i^{(\text{pos})}(\mathbf{x}) \cdot \omega_i(\mathbf{v}|\mathbf{x})$. Thus, our objective is to determine the optimal $\omega_i(\mathbf{v}|\mathbf{x})$ that minimizes the expression above, thereby transforming the VM-DOAT problem into a solvable DOAT instance.

This transformation rests on two key properties: 1) The SOAT problem can be solved efficiently using the Sinkhorn method. Once the Transport Plan is computed, we can sample a trajectory $[(\mathbf{x}_0, \mathbf{v}_0), \dots, (\mathbf{x}_K, \mathbf{v}_K)]$; 2) Given the positions $\mathbf{x}_{0:K}$ of a trajectory, consider the following multi-marginal optimal control problem:

$$\min \int_{t_0}^{t_K} \frac{1}{2} \|\gamma''(t)\|^2 dt, \quad \text{s.t. } \gamma(t_i) = \mathbf{x}_i, \quad i = 0, \dots, K. \quad (17)$$

The velocities of the optimal trajectory, denoted as $V = [\gamma'(t_0), \dots, \gamma'(t_K)]$, can be obtained by solving a sparse linear system, as illustrated in the following proposition:

Proposition 4.8. *The velocities V defined above are the solution to the sparse linear system $VA = B$, where $A \in \mathbb{R}^{(K+1) \times (K+1)}$ is a tridiagonal matrix and $B \in \mathbb{R}^{d \times (K+1)}$. A detailed derivation and the explicit expressions for A and B are provided in Section A.7.*

It is worth noting that solving this linear system requires $\mathcal{O}(K)$ time, ensuring high computational efficiency.

Iterative Scheme to Perform the Transformation Based on this decomposition, we employ an iterative scheme to estimate the conditional velocity distributions across $t_0, \dots, t_K$. We initialize $\omega_i^{(0)}(\mathbf{v}|\mathbf{x}) = \delta(\mathbf{v} - \mathbf{0})$ and subsequently alternate between the following two steps:

- *Solve SOAT Plans:* Compute the K optimal SOAT plans between the distributions $\mu_i^{(\text{pos})}(\mathbf{x}) \cdot \omega_i^{(m)}(\mathbf{v}|\mathbf{x})$ for adjacent time points. This yields the plans $\pi_{i \rightarrow i+1}^{*(m)}[(\mathbf{x}_i, \mathbf{v}_i), (\mathbf{x}_{i+1}, \mathbf{v}_{i+1})]$ for $i = 0, \dots, K-1$, where m denotes the iteration index. This step effectively provides the multi-marginal optimal transport plan given the current $\omega_i^{(m)}(\mathbf{v}|\mathbf{x})$.
- *Update Velocity Distributions:* Update the conditional velocity distributions based on the computed SOAT plans:

$$\omega_i^{(m+1)}(\mathbf{v}|\mathbf{x}_i) = \frac{\int q^{(m)}(\mathbf{z}) \delta(\hat{\mathbf{v}}_i[\mathbf{x}_0 \dots \mathbf{x}_K] - \mathbf{v}) d\mathbf{z}_{\setminus i}}{\int q^{(m)}(\mathbf{z}) d\mathbf{z}_{\setminus i}}, \quad (18)$$

where $d\mathbf{z}_{\setminus i}$ denotes the volume element $d\mathbf{x}_0 d\mathbf{v}_0 \dots d\mathbf{x}_K d\mathbf{v}_K$ excluding $d\mathbf{x}_i$ (i.e., integrating over all variables except $\mathbf{x}_i$). The term $\hat{\mathbf{v}}_i[\mathbf{x}_0, \dots, \mathbf{x}_K]$ denotes the optimal velocity at time t_i derived from the sequence $\mathbf{x}_0, \dots, \mathbf{x}_K$ via Proposition 4.8. Intuitively, $\omega_i^{(m+1)}(\mathbf{v}|\mathbf{x}_i)$ is set to the distribution induced by the velocities of all optimal control trajectories passing through $\mathbf{x}_i$ at time t_i . Since the total transport cost is the sum of costs along individual trajectories, aligning $\mathbf{v}$ with the optimal velocity of each trajectory minimizes the global objective.

In the update step above, the trajectory variable $\mathbf{z} = [(\mathbf{x}_0, \mathbf{v}_0), \dots, (\mathbf{x}_K, \mathbf{v}_K)]$ follows the joint distribution $q^{(m)}(\mathbf{z})$ defined as:

$$q^{(m)}(\mathbf{z}) = \mu_0^{(\text{pos})}(\mathbf{x}_0) \cdot \omega_0^{(m)}(\mathbf{v}_0|\mathbf{x}_0) \times \prod_{i=0}^{K-1} \mathcal{K}_{i \rightarrow i+1}^{*(m)}[(\mathbf{x}_i, \mathbf{v}_i), (\mathbf{x}_{i+1}, \mathbf{v}_{i+1})],$$

where the propagator $\mathcal{K}_{i \rightarrow i+1}^{*(m)}$ is given by:

$$\mathcal{K}_{i \rightarrow i+1}^{*(m)}[(\mathbf{x}_i, \mathbf{v}_i), (\mathbf{x}_{i+1}, \mathbf{v}_{i+1})] = \frac{\pi_{i \rightarrow i+1}^{*(m)}[(\mathbf{x}_i, \mathbf{v}_i), (\mathbf{x}_{i+1}, \mathbf{v}_{i+1})]}{\mu_i^{(\text{pos})}(\mathbf{x}_i) \cdot \omega_i^{(m)}(\mathbf{v}_i|\mathbf{x}_i)}.$$

During the iteration process, the objective $\mathcal{L}_{\text{DOAT}}[\mu_i^{(\text{pos})}(\mathbf{x}) \cdot \omega_i^{(m)}(\mathbf{v}|\mathbf{x})]$ decreases monotonically. Since the cost is bounded below by 0, the iterative algorithm is guaranteed to converge.

5 TracingFlow Algorithm for Trajectory Inference

In the practical implementation of TracingFlow, we first perform an approximation of the conditional velocity distribution to reduce the VM-DOAT problem to a DOAT problem. Subsequently, we employ flow matching to solve for the acceleration field of the DOAT problem. Notably, for lineage tracing data, we incorporate biological priors into this framework.

5.1 Approximation of Conditional Velocity Distribution

Determining the full conditional velocity distribution $\omega_i(\mathbf{v}|\mathbf{x})$ to transform VM-DOAT to DOAT (Section 4.3) is computationally prohibitive, potentially requiring an auxiliary generative model. We therefore adopt a deterministic approximation, assigning a single velocity to each data point. Consequently, Equation (18) is modified to:

$$\omega_i^{(m+1)}(\mathbf{v}|\mathbf{x}_i) = \delta(\mathbf{v} - \mathbf{v}_{\text{exp}}^{(m+1)}(\mathbf{x}_i)), \quad \mathbf{v}_{\text{exp}}^{(m+1)}(\mathbf{x}_i) = \frac{\int q^{(m)}(\mathbf{z}) \hat{\mathbf{v}}_i[\mathbf{x}_0 \cdots \mathbf{x}_K] d\mathbf{z}_{\setminus i}}{\int q^{(m)}(\mathbf{z}) d\mathbf{z}_{\setminus i}}. \quad (19)$$

Intuitively, this assigns $\mathbf{v}_{\text{exp}}^{(m+1)}(\mathbf{x}_i)$ as the mean velocity of all optimal control trajectories passing through $\mathbf{x}_i$ at time t_i .

5.2 MiniBatch-OT

Solving SOAT repeatedly is computationally intensive for large-scale datasets. However, since the conditional velocity distribution $\omega_i(\mathbf{v}|\mathbf{x})$ is approximated as a Dirac delta, the joint empirical distribution $\mu_i(\mathbf{x}, \mathbf{v})$ effectively becomes a superposition of Dirac deltas. This allows us to employ a minibatch strategy: we partition μ_i and μ_{i+1} into B corresponding minibatches $\{\mu_i^{(n)}\}_{n=1}^B$ and $\{\mu_{i+1}^{(n)}\}_{n=1}^B$. By solving the local plans $\pi_{i \rightarrow i+1}^{*(n)}$ independently and aggregating them as $\pi_{i \rightarrow i+1}^* = \oplus_{n=1}^B \pi_{i \rightarrow i+1}^{*(n)}$, we significantly enhance computational efficiency.

5.3 Incorporation of Biological Priors

In lineage tracing applications, let indices l_i and l_{i+1} denote individual samples at time points t_i and t_{i+1} , respectively. Each data point $(\mathbf{x}_{i,l_i}, \mathbf{v}_{i,l_i})$ is associated with a barcode $b_{i,l_i} \in \mathbb{N}^+$. We incorporate this prior by modifying the cost matrix $\mathcal{C}_{i \rightarrow i+1}$. Specifically, when computing the cost between the l -th sample at t_i and the k -th sample at t_{i+1} , we compare their barcodes: if $b_{i,l_i} \neq b_{i+1,l_{i+1}}$, the standard cost is multiplied by a penalty factor p_0 ($p_0 > 1$). This soft constraint discourages transitions between distinct lineages, effectively embedding biological knowledge into the dynamics learning process.

Table 1: Average $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground truth distributions at various time points on the 2D Simulation, Cite 5D, and Cite 100D datasets for different algorithms.

Method	Dynamics	2D Simulation		Cite 5D		Cite 100D	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	0.5270 ± 0.0612	0.5922 ± 0.0813	0.8298 ± 0.0162	0.9155 ± 0.0160	10.5498 ± 0.0291	10.6220 ± 0.0310
SF ² M		0.9926 ± 0.1857	1.1318 ± 0.1734	0.7650 ± 0.0375	0.8956 ± 0.0408	12.3426 ± 0.0735	12.4757 ± 0.0783
3MSBM	2nd Order	1.1855 ± 0.2636	1.2767 ± 0.2215	3.6369 ± 0.3508	3.7703 ± 0.3178	15.7708 ± 0.1473	15.8833 ± 0.1832
MMFM		2.3689 ± 0.1830	2.6466 ± 0.1769	2.5346 ± 0.1520	2.7576 ± 0.1877	12.5353 ± 0.6268	12.7868 ± 0.7061
HRF		1.5390 ± 0.1130	1.6631 ± 0.1214	1.4016 ± 0.0437	1.5313 ± 0.0430	10.1532 ± 0.0643	10.2336 ± 0.0707
CAF		1.6385 ± 0.2777	1.6431 ± 0.2766	4.5589 ± 0.1073	4.7411 ± 0.1620	16.7429 ± 0.4371	16.9837 ± 0.4324
TF (Ours)		0.3748 ± 0.0505	0.5180 ± 0.0360	0.5861 ± 0.0688	0.6386 ± 0.0775	8.0700 ± 0.1537	8.2530 ± 0.1889

5.4 Algorithm of TracingFlow

The TracingFlow algorithm comprises three main steps: First, we apply the iterative approximation from Section 5.1 to determine initial velocities, effectively transforming the VM-DOAT problem into a DOAT problem. Second, we compute the Optimal Transport Plan for SOAT as defined in Equation (11). Third, we sample conditional probability paths to train the acceleration network $\mathbf{a}_\theta(\mathbf{x}, \mathbf{v}, t)$ via regression. Additionally, to provide initial conditions for inference, we parameterize the unique initial velocities derived in the first step using a neural network $\mathbf{v}_{0,\xi}(\mathbf{x})$, trained similarly via regression. The algorithm’s pseudocode is provided in Section D.

Theoretically, TracingFlow recovers the exact distribution if the acceleration and initial velocity are perfectly fitted. In practice, where losses are non-zero, we prove that the positional $\mathcal{W}_2$ distance between the generated and true distributions is bounded by the sum of the velocity and AFM losses, and the relevant theorem is detailed in Section A.8.

6 Experiment Results

To evaluate the effectiveness of our algorithm, TracingFlow, we conducted three categories of experiments: 1) verifying whether the acceleration field of TracingFlow can transport the initial data distribution $\mu_0^{(\text{pos})}(\mathbf{x}) = \int \rho_{t_0}(\mathbf{x}, \mathbf{v})d\mathbf{v}$ to the data distribution $\mu_i^{(\text{pos})}(\mathbf{x})$ at other given times t_i ; 2) verifying whether TracingFlow can effectively perform distribution interpolation and extrapolation; and 3) verifying whether TracingFlow can more effectively infer the dynamical laws in Lineage Tracing data given biological prior knowledge.

Distribution Transport We evaluated TracingFlow on a 2-dimensional synthetic dataset and real single-cell omics datasets of varying dimensions. The evaluation metrics were the 1-Wasserstein ($\mathcal{W}_1$) and 2-Wasserstein ($\mathcal{W}_2$) distances, measuring the discrepancy between the fitted and ground truth distributions. The second-order dynamics model of TracingFlow provides stronger expressive power than first-order algorithms, yielding lower average $\mathcal{W}_1$ and $\mathcal{W}_2$ values across all datasets (Table 1). We visualized the evolutionary trajectories learned by OT-CFM and TracingFlow on the synthetic dataset in Figure 2. TracingFlow’s second-order model allows it to learn trajectories that cross in the position space $\mathbf{x}_1, \mathbf{x}_2$, whereas OT-CFM is unable to do so. In Section C.1, we present the distribution reconstruction accuracy at each time point, as well as experimental results on additional datasets (such as the Mexico Gulf dataset).

Interpolation and Extrapolation To verify TracingFlow’s interpolation and extrapolation capabilities, we conducted a Hold-One-Out experiment on the 5-dimensional EB dataset (time points $t = 0, \dots, 4$). In each iteration, we withheld one time point from the latter four for training and calculated the $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the interpolated and true distributions. As




Figure 2: On the Simulation 2D dataset: a) non-crossing paths learned by OT-CFM, and b) crossing paths learned by TracingFlow.

Table 2: Average $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the predicted and ground truth distributions at held-out time points on the EB 5D dataset for different algorithms.

Method	Dynamics	EB 5D	
		$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	5.1205 ± 0.2400	5.6910 ± 0.2561
SF ² M		4.5602 ± 0.1041	4.9905 ± 0.1262
3MSBM	2nd Order	7.1901 ± 0.8667	7.4697 ± 0.8588
MMFM		9.7777 ± 0.9006	10.5980 ± 0.9407
TF (Ours)		4.4974 ± 0.0950	4.7859 ± 0.1857

Table 3: Average lineage-weighted $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground truth distributions at various time points on the 3D Simulation Lineage and Hematopoiesis datasets for different algorithms.

Method	Dynamics	3D Sim-Lineage		Hematopoiesis	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	2.2383 ± 0.0205	2.3162 ± 0.0216	14.9945 ± 0.3442	15.3971 ± 0.3180
SF ² M		1.5255 ± 0.0513	1.5435 ± 0.0515	18.7170 ± 0.1124	18.9756 ± 0.1168
3MSBM	2nd Order	2.3126 ± 0.4756	2.4905 ± 0.4582	18.3313 ± 0.9834	18.6427 ± 0.9413
MMFM		2.6098 ± 0.3149	3.0309 ± 0.3265	27.0830 ± 2.9299	29.5421 ± 3.1420
HRF		1.7934 ± 0.4043	1.8085 ± 0.4022	15.4571 ± 0.1266	15.7138 ± 0.1282
CAF		2.6806 ± 0.2777	2.7547 ± 0.3521	15.2165 ± 0.0685	15.4974 ± 0.0576
TF w/o Bio-Prior		1.5292 ± 0.1714	1.5532 ± 0.2303	16.6166 ± 0.1229	17.0110 ± 0.1456
TF (Ours)		0.4538 ± 0.2039	0.5328 ± 0.2066	14.0470 ± 0.1530	14.4589 ± 0.1815

shown in Table 2, TracingFlow achieved the highest accuracy, demonstrating that models based on second-order dynamics are superior in capturing high-curvature and non-linear trajectories.

Lineage Tracing Data To verify whether TracingFlow better handles lineage tracing data given biological priors, we experimented on a 3D Simulation Lineage Dataset and the Hematopoiesis Dataset. Following Section 5.3, we introduced biological priors into TracingFlow by modifying the transport cost matrix, other baselines also incorporates priors via conditional velocity fields. We used lineage-weighted $\mathcal{W}_1$ and $\mathcal{W}_2$ distances to evaluate consistency with biological priors while learning dynamics. Results in Table 3 show that incorporating priors enables TracingFlow to outperform other algorithms. In Figure 3, we visualized dynamics on the simulation dataset to demonstrate how TracingFlow learns trajectories conforming to priors. In Figure 4, we plotted the PCA-reduced positions of two barcodes from the Hematopoiesis Dataset at $t = 1$; the positions inferred by TracingFlow are significantly closer to the ground truth than those by OT-CFM.

7 Conclusion and Limitation

In this work, we introduce TracingFlow, a Flow Matching framework for the Dynamical Acceleration Optimal Transport (DOAT) problem. By pre-solving SOAT to learn acceleration and initial velocity, it enables efficient, simulation-free solutions for large-scale DOAT problem. Experiments on simulated and real world data demonstrate that second-order dynamics enhance expressivity, yielding precise temporal distribution recovery. Additionally, incorporating biological priors better preserves intercellular lineages.

A limitation is the computationally intensive SOAT pre-calculation, though this is mitigable via minibatch-OT. Furthermore, explicitly modeling the theoretical conditional velocity distribution per data point is prohibitive, so we approximate it using its expectation. The objective




Figure 3: On the Sim-Lineage dataset: a) paths learned by OT-CFM without considering biological priors, and b) paths learned by TracingFlow that preserve biological priors.




Figure 4: Comparison of generated (circles) and ground truth ('x') positions at $t = 1$ for two barcodes on the hematopoiesis dataset [64]. (a) OT-CFM. (b) TracingFlow. Visualized via 2D PCA.

$\int_0^T \frac{1}{2} \|\mathbf{a}\|^2 dt$ also lacks a clear physical interpretation, as standard classical actions exclude second time derivatives (see Section E.2). Despite current sequencing measurements lacking velocity data, priors suggest biological systems may follow higher-order dynamics due to the existence of complex regulations. Future work will extend TracingFlow to further model these dynamics.

References

- [1] Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.
- [2] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- [3] Vladimir Igorevich Arnol'd. *Mathematical methods of classical mechanics*, volume 60. Springer Science & Business Media, 2013.
- [4] Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio, Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the wasserstein manifold. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [5] Jean-David Benamou, Thomas O Gallouët, and François-Xavier Vialard. Second-order models for optimal transport and cubic splines on the wasserstein space. *Foundations of Computational Mathematics*, 19(5):1113–1143, 2019.
- [6] Charlotte Bunne, Ya-Ping Hsieh, Marco Cuturi, and Andreas Krause. The schrödinger bridge between gaussian measures has a closed form. In *International Conference on Artificial Intelligence and Statistics*, pages 5802–5833. PMLR, 2023.
- [7] Charlotte Bunne, Geoffrey Schiebinger, Andreas Krause, Aviv Regev, and Marco Cuturi. Optimal transport for single-cell and spatial omics. *Nature Reviews Methods Primers*, 4(1): 58, 2024.
- [8] Zihan Cao, Yu Zhong, and Liang-Jian Deng. Taming flow matching with unbalanced optimal transport into fast pansharping. *arXiv preprint arXiv:2503.14975*, 2025.
- [9] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

- [10] Tianrong Chen, Guan-Horng Liu, and Evangelos Theodorou. Likelihood training of schrödinger bridge using forward-backward SDEs theory. In *International Conference on Learning Representations*, 2022.
- [11] Yongxin Chen, Giovanni Conforti, Tryphon T Georgiou, and Luigia Ripani. Multi-marginal schrödinger bridges. In *International Conference on Geometric Science of Information*, pages 725–732. Springer, 2019.
- [12] Lénaïc Chizat, Stephen Zhang, Matthieu Heitz, and Geoffrey Schiebinger. Trajectory inference via mean-field langevin in path space. *Advances in Neural Information Processing Systems*, 35:16731–16742, 2022.
- [13] Gabriele Corso, Vignesh Ram Somnath, Noah Getz, Regina Barzilay, Tommi Jaakkola, and Andreas Krause. Composing unbalanced flows for flexible docking and relaxation. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [14] Valentin De Bortoli, Guan-Horng Liu, Tianrong Chen, Evangelos A Theodorou, and Weilie Nie. Augmented bridge matching. *arXiv preprint arXiv:2311.06978*, 2023.
- [15] Luca Eyring, Dominik Klein, Théo Uscidda, Giovanni Palla, Niki Kilbertus, Zeynep Akata, and Fabian J Theis. Unbalancedness in neural Monge maps improves unpaired domain translation. In *The Twelfth International Conference on Learning Representations*, 2024.
- [16] Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z Alaya, Aurélie Boissunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, et al. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
- [17] Aden Forrow and Geoffrey Schiebinger. Lineageot is a unified framework for lineage tracing and trajectory inference. *Nature Communications*, 12(1), 2021.
- [18] Mingze Gao, Melania Barile, Shirom Chabra, Myriam Haltalli, Emily F. Calderbank, Yiming Chao, Weizhong Zheng, Nicola K. Wilson, Elisa Laurenti, Berthold Göttgens, and Yuanhua Huang. CLADES: a hybrid neuralode-gillespie approach for unveiling clonal cell fate and differentiation dynamics. *Nature Communications*, 16(1), 2025.
- [19] Gennady Gorin, Valentine Svensson, and Lior Pachter. Protein velocity and acceleration from single-cell multiomics experiments. *Genome biology*, 21(1):39, 2020.
- [20] Pengsheng Guo and Alexander G Schwing. Variational rectified flow matching. *arXiv preprint arXiv:2502.09616*, 2025.
- [21] Wenbo Guo, Zeyu Chen, Xinqi Li, Jingmin Huang, Qifan Hu, and Jin Gu. sctrace+: Enhancing cell fate inference by integrating the lineage-tracing and multi-faceted transcriptomic similarity information. *Cell Systems*, 16(9):101398, 2025.
- [22] Peter Halmos, Xinhao Liu, Julian Gold, Feng Chen, Li Ding, and Benjamin J Raphael. Dest-ot: Alignment of spatiotemporal transcriptomics data. *Cell Systems*, 16(2), 2025.
- [23] Ari Hong, Sangseon Lee, and Kwangsoo Kim. Multi-omic relay velocity modeling uncovers dynamic chromatin-transcription regulation across cell states. *Nature Communications*, 2025.
- [24] Guillaume Hugué, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in neural information processing systems*, 35:29705–29718, 2022.

- [25] Qi Jiang and Lin Wan. A physics-informed neural SDE network for learning cellular dynamics from time-series scRNA-seq data. *Bioinformatics*, 40:ii120–ii127, 09 2024. ISSN 1367-4811.
- [26] Kacper Kapusniak, Peter Potaptchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. *Advances in Neural Information Processing Systems*, 37:135011–135042, 2024.
- [27] Dominik Klein, Théo Uscidda, Fabian Theis, and Marco Cuturi. Genot: Entropic (gromov) wasserstein flow matching with applications to single-cell genomics. *Advances in Neural Information Processing Systems*, 37:103897–103944, 2024.
- [28] Dominik Klein, Giovanni Palla, Marius Lange, Michal Klein, Zoe Piran, Manuel Gander, Laetitia Meng-Papaxanthos, Michael Sterr, Lama Saber, Changying Jing, et al. Mapping cells through time and space with moscot. *Nature*, pages 1–11, 2025.
- [29] Takeshi Koshizuka and Issei Sato. Neural lagrangian schrödinger bridge: Diffusion modeling for population dynamics. In *The Eleventh International Conference on Learning Representations*, 2023.
- [30] Kai Kretzschmar and Fiona M. Watt. Lineage tracing. *Cell*, 148(1–2):33–45, 2012.
- [31] Christopher Lance, Malte D Luecken, Daniel B Burkhardt, Robrecht Cannoodt, Pia Rautenstrauch, Anna Laddach, Aidyn Ubiningazhibov, Zhi-Jie Cao, Kaiwen Deng, Sumeer Khan, et al. Multimodal single cell data integration challenge: results and lessons learned. *BioRxiv*, pages 2022–04, 2022.
- [32] Marius Lange, Zoe Piran, Michal Klein, Bastiaan Spanjaard, Dominik Klein, Jan Philipp Junker, Fabian J. Theis, and Mor Nitzan. Mapping lineage-traced cells across time points with moslin. *Genome Biology*, 25(1), 2024.
- [33] Hugo Lavenant, Stephen Zhang, Young-Heon Kim, Geoffrey Schiebinger, et al. Toward a mathematical theory of trajectory inference. *The Annals of Applied Probability*, 34(1A): 428–500, 2024.
- [34] Chen Li, Maria C Virgilio, Kathleen L Collins, and Joshua D Welch. Multi-omic single-cell velocity models epigenome–transcriptome interactions and improves cell fate prediction. *Nature biotechnology*, 41(3):387–398, 2023.
- [35] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.
- [36] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- [37] Suryanarayana Maddu, Victor Chardès, Michael Shelley, et al. Inferring biological processes with intrinsic noise from cross-sectional data. *arXiv preprint arXiv:2410.07501*, 2024.
- [38] Shanjun Mao, Chenyang Zhang, Runjiu Chen, Shan Tang, Xiaodan Fan, and Jie Hu. Cell lineage tracing: Methods, applications, and challenges. *Quantitative Biology*, 13(4), 2025.
- [39] 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.

- [40] Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *International conference on machine learning*, pages 25858–25889. PMLR, 2023.
- [41] Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A computational framework for solving wasserstein lagrangian flows. In *Forty-first International Conference on Machine Learning*, 2024.
- [42] Dogyun Park, Sojin Lee, Sihyeon Kim, Taehoon Lee, Youngjoon Hong, and Hyunwoo J Kim. Constant acceleration flow. *Advances in Neural Information Processing Systems*, 37: 90030–90060, 2024.
- [43] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [44] Qiangwei Peng, Peijie Zhou, and Tiejun Li. stVCR: Spatiotemporal dynamics of single cells. *bioRxiv*, pages 2024–06, 2024.
- [45] Qiangwei Peng, Zihan Wang, Junda Ying, Yuhao Sun, Qing Nie, Lei Zhang, Tiejun Li, and Peijie Zhou. WFR-FM: Simulation-free dynamic unbalanced optimal transport. *arXiv preprint arXiv:2601.06810*, 2026.
- [46] Katarina Petrović, Lazar Atanackovic, Kacper Kapusniak, Michael M. Bronstein, Joey Bose, and Alexander Tong. Curly flow matching for learning non-gradient field dynamics. In *Learning Meaningful Representations of Life (LMRL) Workshop at ICLR 2025*, 2025.
- [47] Gabriel Peyré and Marco Cuturi. *Computational optimal transport: With applications to data science*. Now Foundations and Trends, 2019.
- [48] Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.
- [49] Neha Prasad, Karren Yang, and Caroline Uhler. Optimal transport using gans for lineage tracing. *arXiv preprint arXiv:2007.12098*, 2020.
- [50] Martin Rohbeck, Edward De Brouwer, Charlotte Bunne, Jan-Christian Huetter, Anne Biton, Kelvin Y Chen, Aviv Regev, and Romain Lopez. Modeling complex system dynamics with flow matching across time and conditions. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [51] Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- [52] Yutong Sha, Yuchi Qiu, Peijie Zhou, and Qing Nie. Reconstructing growth and dynamic trajectories from single-cell transcriptomics data. *Nature Machine Intelligence*, 6(1):25–39, 2024.
- [53] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36, 2024.
- [54] Sagar Shrestha and Xiao Fu. Diversified flow matching with translation identifiability. *arXiv preprint arXiv:2511.05558*, 2025.

- [55] Yuhao Sun, Zhenyi Zhang, Zihan Wang, Tiejun Li, and Peijie Zhou. Variational regularized unbalanced optimal transport: Single network, least action. *arXiv preprint arXiv:2505.11823*, 2025.
- [56] Panagiotis Theodoropoulos, Augustinos D Saravanos, Evangelos A Theodorou, and Guan-Horng Liu. Momentum multi-marginal schrödinger bridge matching. *arXiv preprint arXiv:2506.10168*, 2025.
- [57] Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In *International conference on machine learning*, pages 9526–9536. PMLR, 2020.
- [58] 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*, 2024. ISSN 2835-8856. Expert Certification.
- [59] Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. In *International Conference on Artificial Intelligence and Statistics*, pages 1279–1287. PMLR, 2024.
- [60] Elias Ventre, Aden Forrow, Nitya Gadhiwala, Parijat Chakraborty, Omer Angel, and Geoffrey Schiebinger. Trajectory inference for a branching sde model of cell differentiation. *arXiv preprint arXiv:2307.07687*, 2023.
- [61] Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-growth flow matching for single-cell dynamics modeling. *arXiv preprint arXiv:2505.13413*, 2025.
- [62] Kun Wang, Liangzhen Hou, Xin Wang, Xiangwei Zhai, Zhaolian Lu, Zhike Zi, Weiwei Zhai, Xionglei He, Christina Curtis, Da Zhou, and Zheng Hu. Phylovelo enhances transcriptomic velocity field mapping using monotonically expressed genes. *Nature Biotechnology*, 42(5):778–789, 2023.
- [63] Shou-Wen Wang, Michael J. Herriges, Kilian Hurley, Darrell N. Kotton, and Allon M. Klein. Cospar identifies early cell fate biases from single-cell transcriptomic and lineage information. *Nature Biotechnology*, 40(7):1066–1074, 2022.
- [64] Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D Camargo, and Allon M Klein. Lineage tracing on transcriptional landscapes links state to fate during differentiation. *Science*, 367(6479):eaaw3381, 2020.
- [65] Maosheng Yang. Topological schrödinger bridge matching. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [66] Grace Hui Ting Yeo, Sachit D Saksena, and David K Gifford. Generative modeling of single-cell time series with prescient enables prediction of cell trajectories with interventions. *Nature communications*, 12(1):3222, 2021.
- [67] Jiaqi Zhang, Erica Larschan, Jeremy Bigness, and Ritambhara Singh. scNODE: generative model for temporal single cell transcriptomic data prediction. *Bioinformatics*, 40 (Supplement_2):ii146–ii154, 09 2024. ISSN 1367-4811.
- [68] Xi Nicole Zhang, Yuan Pu, Yuki Kawamura, Andrew Loza, Yoshua Bengio, Dennis Shung, and Alexander Tong. Trajectory flow matching with applications to clinical time series modelling. *Advances in Neural Information Processing Systems*, 37:107198–107224, 2024.

- [69] Yichi Zhang, Yici Yan, Alex Schwing, and Zhizhen Zhao. Towards hierarchical rectified flow. *arXiv preprint arXiv:2502.17436*, 2025.
- [70] Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through regularized unbalanced optimal transport. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [71] Zhenyi Zhang, Yuhao Sun, Qiangwei Peng, Tiejun Li, and Peijie Zhou. Integrating dynamical systems modeling with spatiotemporal scRNA-Seq data analysis. *Entropy*, 27(5), 2025. ISSN 1099-4300.
- [72] Grace XY Zheng, Jessica M Terry, Phillip Belgrader, Paul Ryvkin, Zachary W Bent, Ryan Wilson, Solongo B Ziraldo, Tobias D Wheeler, Geoff P McDermott, Junjie Zhu, et al. Massively parallel digital transcriptional profiling of single cells. *Nature communications*, 8(1):14049, 2017.
- [73] Qunxi Zhu and Wei Lin. Switched flow matching: Eliminating singularities via switching odes. *arXiv preprint arXiv:2405.11605*, 2024.
- [74] Qunxi Zhu, Bolin Zhao, Jingdong Zhang, Peiyang Li, and Wei Lin. Governing equation discovery of a complex system from snapshots. *arXiv preprint arXiv:2410.16694*, 2024.

Contents of Appendix

A	Proofs of Main Theorems	17
A.1	Proof of Proposition 4.1 : Optimal Trajectory of Single Particle	17
A.3	Proof of Remark 4.4 : No-Cross Property	18
A.4	Proof of Theorem 4.5 : Marginal Acceleration Field	19
A.5	Proof of Theorem 4.6 : The Relation Between AFM Loss and CAFM Loss	19
A.6	Proof of Theorem 4.7 : Marginal Probability Path	20
A.7	Proof of Proposition 4.8 : TracingFlow Exactly Solves DOAT Problem	21
A.8	Proof of Proposition 4.9 : Solve Linear System to Get Optimal Velocity	22
A.9	Proof of Theorem A.1 : $\mathcal{W}_2$ Distance can be bounded by $\mathcal{L}_{\text{AFM}}$ and $\mathcal{L}_{v_0}$	24
B	Datasets and Evaluation Metric	26
B.1	Experiment Setup	26
B.2	Evaluation Metrics	27
B.3	Datasets	27
C	Experiment Details	28
C.1	Experiment Details across Various Datasets	28
C.2	Training Time and Scalability of Tracing Flow	32
C.3	Sensitivity Analysis on Minibatch-OT	32
D	Pseudocode for Algorithm	34
E	Discussion	35
E.1	Relation with other Algorithms	35
E.2	Is $\mathcal{S} = \mathcal{X} \times \mathcal{V}$ a phase space?	35

A Proofs of Main Theorems

A.1 Proof of Proposition 4.1

Proof. First, we show that the unique minimizer is a cubic polynomial. Consider the energy functional defined on the interval $[0, T]$:

$$\mathcal{J}[\gamma] = \frac{1}{2} \int_0^T \|\gamma''(t)\|_2^2 dt. \quad (20)$$

The integrand $L(\gamma, \gamma', \gamma'') = \frac{1}{2} \|\gamma''\|_2^2$ depends on derivatives up to the second order. The necessary condition for optimality is given by the Euler–Lagrange equation:

$$\frac{\partial L}{\partial \gamma} - \frac{d}{dt} \left(\frac{\partial L}{\partial \gamma'} \right) + \frac{d^2}{dt^2} \left(\frac{\partial L}{\partial \gamma''} \right) = 0. \quad (21)$$

Since $\partial L / \partial \gamma = 0$, $\partial L / \partial \gamma' = 0$, and $\partial L / \partial \gamma'' = \gamma''$, the equation simplifies to:

$$\frac{d^2}{dt^2} \gamma''(t) = \gamma^{(4)}(t) = 0. \quad (22)$$

This implies that the optimal trajectory $\gamma(t)$ is a cubic polynomial:

$$\gamma(t) = c_3 t^3 + c_2 t^2 + c_1 t + c_0. \quad (23)$$

Imposing the boundary conditions $\gamma(0) = \mathbf{x}_0, \gamma'(0) = \mathbf{v}_0$ and $\gamma(T) = \mathbf{x}_T, \gamma'(T) = \mathbf{v}_T$ leads to the linear system:

$$\begin{cases} c_0 = \mathbf{x}_0, \\ c_1 = \mathbf{v}_0, \\ c_3 T^3 + c_2 T^2 + c_1 T + c_0 = \mathbf{x}_T, \\ 3c_3 T^2 + 2c_2 T + c_1 = \mathbf{v}_T. \end{cases} \quad (24)$$

Solving for the coefficients yields the unique solution:

$$\begin{cases} c_0 = \mathbf{x}_0, \\ c_1 = \mathbf{v}_0, \\ c_2 = \frac{3(\mathbf{x}_T - \mathbf{x}_0) - (2\mathbf{v}_0 + \mathbf{v}_T)T}{T^2}, \\ c_3 = \frac{(\mathbf{v}_0 + \mathbf{v}_T)T - 2(\mathbf{x}_T - \mathbf{x}_0)}{T^3}. \end{cases} \quad (25)$$

□

A.2 Proof of Remark 4.3

Proof. First we propose a lemma:

Lemma. Let $\mathcal{C}_{[0,t]}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}_t, \mathbf{v}_t)$ be the least acceleration cost through $[0, t]$, $(\mathbf{x}_0, \mathbf{v}_0), (\mathbf{x}_T, \mathbf{v}_T) \in \mathcal{X} \times \mathcal{V}$ are two different points. Then for arbitrary midpoint $\tau \in (0, 1)$ and state $(\mathbf{x}, \mathbf{v}) \in \mathcal{X} \times \mathcal{V}$, the following inequality holds:

$$\mathcal{C}_{[0,\tau]}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}, \mathbf{v}) + \mathcal{C}_{[\tau,T]}(\mathbf{x}, \mathbf{v}, \mathbf{x}_T, \mathbf{v}_T) \geq \mathcal{C}_{[0,T]}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}_T, \mathbf{v}_T) \quad (26)$$

The "=" holds if and only if $(\mathbf{x}, \mathbf{v}) = \gamma^*(\tau)$, where $\gamma^*(t) (t \in [0, T])$ is the cubic interpolation trajectory connecting $(\mathbf{x}_0, \mathbf{v}_0)$ and $(\mathbf{x}_T, \mathbf{v}_T)$.

Proof. Fix $\tau \in (0, T)$. Define objective function $F_\tau(\mathbf{x}, \mathbf{v}) = \mathcal{C}_{[0,\tau]}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}, \mathbf{v}) + \mathcal{C}_{[\tau,T]}(\mathbf{x}, \mathbf{v}, \mathbf{x}_T, \mathbf{v}_T)$. By Corollary 4.2, we can rewrite

$$\begin{aligned} F_\tau(\mathbf{x}, \mathbf{v}) &= \frac{2}{\tau^3} \left\{ (\|\mathbf{v}_0\|^2 + \langle \mathbf{v}_0, \mathbf{v} \rangle + \|\mathbf{v}\|^2)\tau^2 + 3\langle \mathbf{v}_0 + \mathbf{v}, \mathbf{x}_0 - \mathbf{x} \rangle\tau + 3\|\mathbf{x}_0 - \mathbf{x}\|^2 \right\} + \\ &\frac{2}{(T-\tau)^3} \left\{ (\|\mathbf{v}_T\|^2 + \langle \mathbf{v}_T, \mathbf{v} \rangle + \|\mathbf{v}\|^2)(T-\tau)^2 + 3\langle \mathbf{v}_T + \mathbf{v}, \mathbf{x} - \mathbf{x}_T \rangle(T-\tau) + 3\|\mathbf{x}_T - \mathbf{x}\|^2 \right\} \end{aligned} \quad (27)$$

Which is a quadratic form of $(\mathbf{x}, \mathbf{v})$ up to a constant. $F_\tau(\mathbf{x}, \mathbf{v}) \geq 0$ always holds, namely, the quadratic form has a finite lower bound, so it is semi-quadratic. To minimize $F_\tau(\mathbf{x}, \mathbf{v})$, we can take differentiations:

$$\nabla_{\mathbf{x}} F_\tau(\mathbf{x}, \mathbf{v}) = 3 \left(\frac{-(\mathbf{v}_0 + \mathbf{v})\tau + 2(\mathbf{x} - \mathbf{x}_0)}{\tau^2} + \frac{(\mathbf{v} + \mathbf{v}_T)(T-\tau) + 2(\mathbf{x} - \mathbf{x}_T)}{(T-\tau)^2} \right) = 0. \quad (28)$$

$$\nabla_{\mathbf{v}} F_\tau(\mathbf{x}, \mathbf{v}) = \frac{(\mathbf{v}_0 + 2\mathbf{v})\tau + 3(\mathbf{x}_0 - \mathbf{x})}{\tau^2} + \frac{(\mathbf{v}_T + 2\mathbf{v})(T-\tau) + 3(\mathbf{x} - \mathbf{x}_T)}{(T-\tau)^2} = 0, \quad (29)$$

Denote the cubic interpolation from $(\mathbf{x}_0, \mathbf{v}_0)$ to $(\mathbf{x}, \mathbf{v})$ on $[0, \tau]$ and $(\mathbf{x}, \mathbf{v})$ to $(\mathbf{x}_T, \mathbf{v}_T)$ on $[\tau, T]$ by $\gamma_{[0,\tau]}(t)$ and $\gamma_{[\tau,T]}(t)$, respectively. The coefficients are known according to Section A.1. Then we have:

$$\gamma''_{[0,\tau]}(\tau^-) = 6 \cdot \frac{(\mathbf{v}_0 + \mathbf{v})\tau - 2(\mathbf{x} - \mathbf{x}_0)}{\tau^2} + 2 \cdot \frac{3(\mathbf{x} - \mathbf{x}_0) - (2\mathbf{v}_0 + \mathbf{v})\tau}{\tau^2} = 2 \cdot \frac{3(\mathbf{x}_0 - \mathbf{x}) + (\mathbf{v}_0 + 2\mathbf{v})\tau}{\tau^2}, \quad (30)$$

$$\begin{aligned}\gamma''_{[\tau,T]}(\tau^+) &= 6 \cdot \frac{(\mathbf{v} + \mathbf{v}_T)\tau - 2(\mathbf{x}_T - \mathbf{x})}{\tau^2} + 2 \cdot \frac{3(\mathbf{x}_T - \mathbf{x}) - (2\mathbf{v} + \mathbf{v}_T)(T - \tau)}{(T - \tau)^2} \\ &= 2 \cdot \frac{3(\mathbf{x} - \mathbf{x}_T) + (\mathbf{v} + 2\mathbf{v}_T)(T - \tau)}{(T - \tau)^2}\end{aligned}\quad (31)$$

$$\gamma'''_{[0,\tau]}(\tau^-) = 6 \cdot \frac{(\mathbf{v}_0 + \mathbf{v})\tau - 2(\mathbf{x} - \mathbf{x}_0)}{\tau^2} \quad (32)$$

$$\gamma'''_{[\tau,T]}(\tau^+) = 6 \cdot \frac{(\mathbf{v} + \mathbf{v}_T)(T - \tau) - 2(\mathbf{x}_T - \mathbf{x})}{(T - \tau)^2} \quad (33)$$

Then we can find that 28 and 29 appear to be $\gamma'''_{[0,\tau]}(\tau^-) = \gamma'''_{[\tau,T]}(\tau^+)$, $\gamma''_{[0,\tau]}(\tau^-) = \gamma''_{[\tau,T]}(\tau^+)$. Because the two curves are cubic, and $\gamma^{(i)}_{[0,\tau]}(\tau^-) = \gamma^{(i)}_{[\tau,T]}(\tau^+)$, $i = 0, 1, 2, 3$, so we know that in the optimal case $\gamma_{[0,\tau]} = \gamma^*|_{[0,\tau]}$, $\gamma_{[\tau,T]} = \gamma^*|_{[\tau,T]}$. That is, the unique solution of 28 and 29 is $(\gamma^*(\tau), \gamma^{*'}(\tau))$, which completes the proof. $\square$

Back to our target, assume γ and $\tilde{\gamma}$ cross at $(\mathbf{x}_\tau, \mathbf{v}_\tau)$, then we have

$$\begin{aligned}&\mathcal{C}_{[0,T]}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}_T, \mathbf{v}_T) + \mathcal{C}_{[0,T]}(\tilde{\mathbf{x}}_0, \tilde{\mathbf{v}}_0, \tilde{\mathbf{x}}_T, \tilde{\mathbf{v}}_T) \\ &= \left\{ \mathcal{C}_{[0,\tau]}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}_\tau, \mathbf{v}_\tau) + \mathcal{C}_{[\tau,T]}(\mathbf{x}_\tau, \mathbf{v}_\tau, \mathbf{x}_T, \mathbf{v}_T) \right\} + \left\{ \mathcal{C}_{[0,\tau]}(\tilde{\mathbf{x}}_0, \tilde{\mathbf{v}}_0, \mathbf{x}_\tau, \mathbf{v}_\tau) + \mathcal{C}_{[\tau,T]}(\mathbf{x}_\tau, \mathbf{v}_\tau, \tilde{\mathbf{x}}_T, \tilde{\mathbf{v}}_T) \right\} \\ &= \left\{ \mathcal{C}_{[0,\tau]}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}_\tau, \mathbf{v}_\tau) + \mathcal{C}_{[\tau,T]}(\mathbf{x}_\tau, \mathbf{v}_\tau, \tilde{\mathbf{x}}_T, \tilde{\mathbf{v}}_T) \right\} + \left\{ \mathcal{C}_{[0,\tau]}(\tilde{\mathbf{x}}_0, \tilde{\mathbf{v}}_0, \mathbf{x}_\tau, \mathbf{v}_\tau) + \mathcal{C}_{[\tau,T]}(\mathbf{x}_\tau, \mathbf{v}_\tau, \mathbf{x}_T, \mathbf{v}_T) \right\} \\ &\geq \mathcal{C}_{[0,T]}(\mathbf{x}_0, \mathbf{v}_0, \tilde{\mathbf{x}}_T, \tilde{\mathbf{v}}_T) + \mathcal{C}_{[0,T]}(\tilde{\mathbf{x}}_0, \tilde{\mathbf{v}}_0, \mathbf{x}_T, \mathbf{v}_T).\end{aligned}\quad (34)$$

The equality holds if and only if γ and $\tilde{\gamma}$ coincide. Consequently, the strict inequality holds for distinct trajectories, which implies the non-crossing property. This completes the proof. $\square$

A.3 Proof of Theorem 4.4

Proof. Consider the time derivative of the marginal distribution:

$$\begin{aligned}\partial_t \rho_t(\mathbf{x}, \mathbf{v}) &= \partial_t \left\{ \int \rho_t(\mathbf{x}, \mathbf{v}|z) q(z) dz \right\} \\ &= \int \partial_t \rho_t(\mathbf{x}, \mathbf{v}|z) q(z) dz \\ &= - \int \nabla_{\mathbf{x}} \cdot (\rho_t(\mathbf{x}, \mathbf{v}|z) \mathbf{v}) q(z) dz - \int \nabla_{\mathbf{v}} \cdot (\rho_t(\mathbf{x}, \mathbf{v}|z) \mathbf{a}(\mathbf{x}, \mathbf{v}, t|z)) q(z) dz \\ &= - \nabla_{\mathbf{x}} \cdot \left(\int \rho_t(\mathbf{x}, \mathbf{v}|z) \mathbf{v} q(z) dz \right) - \nabla_{\mathbf{v}} \cdot \left(\int \rho_t(\mathbf{x}, \mathbf{v}|z) \mathbf{a}(\mathbf{x}, \mathbf{v}, t|z) q(z) dz \right) \\ &= - \nabla_{\mathbf{x}} \cdot (\rho_t(\mathbf{x}, \mathbf{v}) \mathbf{v}) - \nabla_{\mathbf{v}} \cdot (\mathbf{a}_t(\mathbf{x}, \mathbf{v}) \rho_t(\mathbf{x}, \mathbf{v}))\end{aligned}\quad (35)$$

Therefore, $\rho_t(\mathbf{x}, \mathbf{v})$ satisfies the continuity equation:

$$\partial_t \rho_t(\mathbf{x}, \mathbf{v}) = - \nabla_{\mathbf{x}} \cdot (\rho_t(\mathbf{x}, \mathbf{v}) \mathbf{v}) - \nabla_{\mathbf{v}} \cdot (\mathbf{a}_t(\mathbf{x}, \mathbf{v}) \rho_t(\mathbf{x}, \mathbf{v})) \quad (36)$$

This completes the proof. $\square$

A.4 Proof of Theorem 4.5

Proof. Consider the gradient of the AFM Loss and CAFM Loss:

$$\nabla_{\theta} \mathcal{L}_{\text{AFM}} = \nabla_{\theta} \mathbb{E}_{\rho_t(\mathbf{x}, \mathbf{v})} (\|\mathbf{a}_{\theta}(\mathbf{x}, \mathbf{v}, t)\|^2 - 2\mathbf{a}_{\theta}^T(\mathbf{x}, \mathbf{v}, t) \mathbf{a}(\mathbf{x}, \mathbf{v}, t)) \quad (37)$$

$$\nabla_{\theta} \mathcal{L}_{\text{CAFM}} = \nabla_{\theta} \mathbb{E}_{\rho_t(\mathbf{x}, \mathbf{v}|z) q(z)} (\|\mathbf{a}_{\theta}(\mathbf{x}, \mathbf{v}, t)\|^2 - 2\mathbf{a}_{\theta}^T(\mathbf{x}, \mathbf{v}, t) \mathbf{a}(\mathbf{x}, \mathbf{v}, t|z)) \quad (38)$$

For the first term of $\nabla_{\theta}\mathcal{L}_{\text{AFM}}$:

$$\begin{aligned}\mathbb{E}_{\rho_t(\mathbf{x},\mathbf{v})}\|\mathbf{a}_{\theta}(\mathbf{x},\mathbf{v},t)\|^2 &= \int \rho_t(\mathbf{x},\mathbf{v})\|\mathbf{a}_{\theta}(\mathbf{x},\mathbf{v},t)\|^2 d\mathbf{x}d\mathbf{v} \\ &= \int \rho_t(\mathbf{x},\mathbf{v}|z)q(z)\|\mathbf{a}_{\theta}(\mathbf{x},\mathbf{v},t)\|^2 d\mathbf{x}d\mathbf{v}dz \\ &= \mathbb{E}_{\rho_t(\mathbf{x},\mathbf{v}|z)q(z)}\|\mathbf{a}_{\theta}(\mathbf{x},\mathbf{v},t)\|^2\end{aligned}\tag{39}$$

For the second term:

$$\begin{aligned}\mathbb{E}_{\rho_t(\mathbf{x},\mathbf{v})}[\mathbf{a}_{\theta}^T(\mathbf{x},\mathbf{v},t)\mathbf{a}(\mathbf{x},\mathbf{v},t)] &= \int \rho_t(\mathbf{x},\mathbf{v})\mathbf{a}_{\theta}^T(\mathbf{x},\mathbf{v},t)\mathbf{a}(\mathbf{x},\mathbf{v},t)d\mathbf{x}d\mathbf{v} \\ &= \int \rho_t(\mathbf{x},\mathbf{v})\mathbf{a}_{\theta}^T(\mathbf{x},\mathbf{v},t)\left(\int q(z)\frac{\mathbf{a}(\mathbf{x},\mathbf{v},t|z)\rho_t(\mathbf{x},\mathbf{v}|z)}{\rho_t(\mathbf{x},\mathbf{v})}dz\right)d\mathbf{x}d\mathbf{v} \\ &= \int \int \mathbf{a}_{\theta}^T(\mathbf{x},\mathbf{v},t)\mathbf{a}(\mathbf{x},\mathbf{v},t|z)\rho_t(\mathbf{x},\mathbf{v}|z)q(z)d\mathbf{x}d\mathbf{v}dz \\ &= \mathbb{E}_{\rho_t(\mathbf{x},\mathbf{v}|z)q(z)}(\mathbf{a}_{\theta}^T(\mathbf{x},\mathbf{v},t)\mathbf{a}(\mathbf{x},\mathbf{v},t|z))\end{aligned}\tag{40}$$

□

A.5 Proof of Theorem 4.6

Proof. Consider a given time t_i . Based on the expressions for $\mathbf{m}_x, \mathbf{m}_v$, at this moment we exactly have:

$$\rho_{t_i}(\mathbf{x},\mathbf{v}|z) = \mathcal{N}(\mathbf{x}|\mathbf{x}_i, \sigma_x^2\mathbf{I}) \cdot \mathcal{N}(\mathbf{v}|\mathbf{v}_i, \sigma_v^2\mathbf{I})\tag{41}$$

For simplicity, let $d\mathbf{s}_j = d\mathbf{x}_j d\mathbf{v}_j$. Therefore, the Marginal Probability Path is:

$$\begin{aligned}\rho_{t_i}(\mathbf{x},\mathbf{v}) &= \int \rho_{t_i}(\mathbf{x},\mathbf{v}|z)q(z)dz \\ &= \int \rho_{t_i}(\mathbf{x},\mathbf{v}|\mathbf{x}_i,\mathbf{v}_i)\mu_0(\mathbf{x}_0,\mathbf{v}_0) \cdot \prod_{j=0}^{K-1} \mathcal{K}_{j \rightarrow j+1}^*[(\mathbf{x}_j,\mathbf{v}_j),(\mathbf{x}_{j+1},\mathbf{v}_{j+1})]d\mathbf{s}_0 \cdots d\mathbf{s}_K\end{aligned}\tag{42}$$

Using the definition of the propagator, we know that:

$$\int \mathcal{K}_{j \rightarrow j+1}^*[(\mathbf{x}_j,\mathbf{v}_j),(\mathbf{x}_{j+1},\mathbf{v}_{j+1})]d\mathbf{x}_{j+1}d\mathbf{v}_{j+1} = 1\tag{43}$$

Thus, we can integrate out the future terms ($j > i$) first, and then iteratively collapse the past terms ($j < i$):

$$\begin{aligned}
\rho_{t_i}(\mathbf{x}, \mathbf{v}) &= \int \rho_{t_i}(\mathbf{x}, \mathbf{v} | (\mathbf{x}_i, \mathbf{v}_i)) \mu_0(\mathbf{x}_0, \mathbf{v}_0) \cdot \prod_{j=0}^{K-1} \mathcal{K}_{j \rightarrow j+1}^*[(\mathbf{x}_j, \mathbf{v}_j), (\mathbf{x}_{j+1}, \mathbf{v}_{j+1})] d\mathbf{s}_0 \cdots d\mathbf{s}_K \\
&= \int \rho_{t_i}(\mathbf{x}, \mathbf{v} | (\mathbf{x}_i, \mathbf{v}_i)) \mu_0(\mathbf{x}_0, \mathbf{v}_0) \cdot \prod_{j=0}^{i-1} \mathcal{K}_{j \rightarrow j+1}^*[(\mathbf{x}_j, \mathbf{v}_j), (\mathbf{x}_{j+1}, \mathbf{v}_{j+1})] d\mathbf{s}_0 \cdots d\mathbf{s}_i \\
&= \int \rho_{t_i}(\mathbf{x}, \mathbf{v} | (\mathbf{x}_i, \mathbf{v}_i)) \pi_{0 \rightarrow 1}^*[(\mathbf{x}_0, \mathbf{v}_0), (\mathbf{x}_1, \mathbf{v}_1)] \prod_{j=1}^{i-1} \mathcal{K}_{j \rightarrow j+1}^*[(\mathbf{x}_j, \mathbf{v}_j), (\mathbf{x}_{j+1}, \mathbf{v}_{j+1})] d\mathbf{s}_0 \cdots d\mathbf{s}_i \\
&= \int \rho_{t_i}(\mathbf{x}, \mathbf{v} | (\mathbf{x}_i, \mathbf{v}_i)) \mu_1(\mathbf{x}_1, \mathbf{v}_1) \prod_{j=1}^{i-1} \mathcal{K}_{j \rightarrow j+1}^*[(\mathbf{x}_j, \mathbf{v}_j), (\mathbf{x}_{j+1}, \mathbf{v}_{j+1})] d\mathbf{s}_1 \cdots d\mathbf{s}_i \\
&\vdots \\
&= \int \rho_{t_i}(\mathbf{x}, \mathbf{v} | (\mathbf{x}_i, \mathbf{v}_i)) \mu_i(\mathbf{x}_i, \mathbf{v}_i) d\mathbf{s}_i \\
&= \int (\mathcal{N}(\mathbf{x} | \mathbf{x}_i, \sigma_x^2 \mathbf{I}) \cdot \mathcal{N}(\mathbf{v} | \mathbf{v}_i, \sigma_v^2 \mathbf{I})) \mu_i(\mathbf{x}_i, \mathbf{v}_i) d\mathbf{s}_i \\
&= (\mu_i * g)(\mathbf{x}, \mathbf{v})
\end{aligned} \tag{44}$$

where $g(\mathbf{x}, \mathbf{v}) = (\mathcal{N}(\mathbf{x} | \mathbf{0}, \sigma_x^2 \mathbf{I}) \cdot \mathcal{N}(\mathbf{v} | \mathbf{0}, \sigma_v^2 \mathbf{I}))$. When $(\sigma_x, \sigma_v) \rightarrow \mathbf{0}$, we have $\rho_{t_i}(\mathbf{x}, \mathbf{v}) \rightarrow \mu_i(\mathbf{x}, \mathbf{v})$ in distribution.

If $\mu_i(\mathbf{x}, \mathbf{v})$ is absolutely continuous with respect to the Lebesgue measure and admits a density in $L^1(\mathbb{R}^{2d})$, we can guarantee strong convergence in the L^1 norm: $\lim_{\sigma \rightarrow 0} \|\rho_{t_i}(\mathbf{x}, \mathbf{v}) - \mu_i(\mathbf{x}, \mathbf{v})\|_{L^1} = 0$ for all $i = 0, 1, \dots, K$. □

A.6 Proof of Proposition 4.7

First, we prove that the Marginal Probability Path constructed by TracingFlow converges to the optimal solution of the DOAT problem when $\sigma_x \rightarrow 0, \sigma_v \rightarrow 0$.

We investigate directly in the augmented space. Let $\mathbf{s} = [\mathbf{x}, \mathbf{v}]^T \in \mathcal{X} \times \mathcal{V}$, and consider the cost of the DOAT problem. The subsequent position of a particle starting from the initial point $\mathbf{s}_0$ can be represented by the flow map $\phi_t(\mathbf{s}_0)$. Therefore, the cost can be written as:

$$\begin{aligned}
\mathcal{L}_{\text{DOAT}}(\mathbf{a}, \rho) &= \int_{t_0}^{t_k} \int_{\mathcal{X} \times \mathcal{V}} \frac{1}{2} \|\mathbf{a}(\mathbf{s}, t)\|^2 \rho_t(\mathbf{s}) d\mathbf{x} d\mathbf{v} dt \\
&= \int_{t_0}^{t_k} \int_{\mathcal{X} \times \mathcal{V}} \frac{1}{2} \|\mathbf{a}(\phi_t(\mathbf{s}_0), t)\|^2 \det\left(\frac{\partial \phi}{\partial \mathbf{s}_0}\right)^{-1} \rho_0(\mathbf{s}_0) \det\left(\frac{\partial \phi}{\partial \mathbf{s}_0}\right) d\mathbf{s}_0 dt \\
&= \int_{\mathcal{X} \times \mathcal{V}} \rho_0(\mathbf{s}_0) d\mathbf{s}_0 \int_{t_0}^{t_k} \frac{1}{2} \|\mathbf{a}(\phi_t(\mathbf{s}_0), t)\|^2 dt \\
&= \int_{\Omega} \left(\int_{t_0}^{t_k} \frac{1}{2} \|\gamma''(t)\|^2 dt \right) d\mathbb{P}[\gamma]
\end{aligned} \tag{45}$$

where $\mathbb{P}$ is a probability measure in the path space, satisfying the constraint $(e_{t_i})_{\#} \mathbb{P} = \mu_i(\mathbf{x}, \mathbf{v})$. Here, e_t is the evaluation map $e_t[\gamma] = \gamma(t)$, and $(e_{t_i})_{\#} \mathbb{P}$ denotes the push-forward of the probability measure by the evaluation map. According to Proposition 4.1, the minimum cost path connecting $\mathbf{s}_i$ and $\mathbf{s}_{i+1}$ is a cubic spline, and the minimum cost is $\mathcal{C}_{i \rightarrow i+1}[\mathbf{s}_i, \mathbf{s}_{i+1}]$. Therefore, for any path γ :

$$\int_{t_i}^{t_{i+1}} \frac{1}{2} \|\gamma''(t)\|^2 dt \geq \mathcal{C}_{i \rightarrow i+1}[\mathbf{s}_i, \mathbf{s}_{i+1}] \tag{46}$$

Thus, the DOAT problem cost has a lower bound:

$$\begin{aligned}
\mathcal{L}_{\text{DOAT}}(\rho, \mathbf{a}) &= \int_{\Omega} \sum_{i=0}^{K-1} \left(\int_{t_i}^{t_{i+1}} \frac{1}{2} \|\gamma''(t)\|^2 dt \right) d\mathbb{P}[\gamma] \\
&\geq \int_{\Omega} \sum_{i=0}^{K-1} \mathcal{C}_{i \rightarrow i+1}[\mathbf{s}_i, \mathbf{s}_{i+1}] d\mathbb{P}[\gamma] \\
&\geq \int \sum_{i=0}^{K-1} \mathcal{C}_{i \rightarrow i+1}[\mathbf{s}_i, \mathbf{s}_{i+1}] d\pi^*(\mathbf{s}_i, \mathbf{s}_{i+1})
\end{aligned} \tag{47}$$

where $\pi^*(\mathbf{s}_i, \mathbf{s}_{i+1})$ is the Optimal Transport Plan for the SOAT problem between $\mu(\mathbf{s}_i)$ and $\mu(\mathbf{s}_{i+1})$.

Consider the path constructed by TracingFlow. In the case where $\sigma_{\mathbf{x}}, \sigma_{\mathbf{v}} \rightarrow 0$, we have $\rho_t(\mathbf{x}, \mathbf{v} | \mathbf{z}) = \rho_t(\mathbf{s} | \mathbf{z}) \rightarrow \delta(\mathbf{s} - \gamma_{\mathbf{z}}^*(t))$, where $\gamma_{\mathbf{z}}^*$ is the optimal single-particle trajectory connecting $\mathbf{z} = [\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_K]$. We calculate the transport cost of the constructed path:

$$\begin{aligned}
\mathcal{L}_{\text{TF}} &= \int q(\mathbf{z}) d\mathbf{z} \sum_{i=0}^{K-1} \int_{t_i}^{t_{i+1}} \|\gamma_{\mathbf{z}}''(t)\|^2 dt \\
&= \int q(\mathbf{z}) d\mathbf{z} \sum_{i=0}^{K-1} \mathcal{C}_{i \rightarrow i+1}[\mathbf{s}_i, \mathbf{s}_{i+1}] \\
&= \int \mu_0(\mathbf{s}_0) \prod_{i=0}^{K-1} \mathcal{K}_{i \rightarrow i+1}^*[\mathbf{s}_i, \mathbf{s}_{i+1}] \left(\sum_{j=0}^{K-1} \mathcal{C}_{j \rightarrow j+1}[\mathbf{s}_j, \mathbf{s}_{j+1}] \right) d\mathbf{z}
\end{aligned} \tag{48}$$

Here, using the definition of the propagator:

$$\begin{aligned}
&\int \mu_0(\mathbf{s}_0) \prod_{i=0}^{K-1} \mathcal{K}_{i \rightarrow i+1}^*[\mathbf{s}_i, \mathbf{s}_{i+1}] (\mathcal{C}_{j \rightarrow j+1}[\mathbf{s}_j, \mathbf{s}_{j+1}]) d\mathbf{z} \\
&= \int \pi_{j \rightarrow j+1}^*[\mathbf{s}_j, \mathbf{s}_{j+1}] \mathcal{C}_{j \rightarrow j+1}[\mathbf{s}_j, \mathbf{s}_{j+1}] d\mathbf{s}_j d\mathbf{s}_{j+1}
\end{aligned} \tag{49}$$

Therefore:

$$\mathcal{L}_{\text{TF}} = \sum_{i=0}^{K-1} \int \pi_{i \rightarrow i+1}^*[\mathbf{s}_i, \mathbf{s}_{i+1}] \mathcal{C}_{i \rightarrow i+1}[\mathbf{s}_i, \mathbf{s}_{i+1}] d\mathbf{s}_i d\mathbf{s}_{i+1} \tag{50}$$

This coincides exactly with the lower bound of the DOAT problem cost derived above. Thus, the path constructed by TracingFlow is exactly the optimal solution to the DOAT problem in the case where $\sigma_{\mathbf{x}}, \sigma_{\mathbf{v}} \rightarrow 0$.

Next, we prove that under the construction of TracingFlow, the acceleration field $\mathbf{a}(\mathbf{s}, t)$ is single-valued with respect to $\mathbf{s}$ and t . By Remark 4.3, under the Optimal Transport Plan, the trajectories of flow maps starting from different points must not intersect. This implies that for a given time t , any point $\mathbf{s}$ in the augmented space is traversed by at most one flow map trajectory. Therefore, the acceleration field $\mathbf{a}(\mathbf{s}, t)$ possesses single-valuedness. In other words, the solution constructed by TracingFlow can indeed be expressed by a single optimal acceleration field $\mathbf{a}(\mathbf{s}, t)$.

A.7 Proof of Proposition 4.8

Proof. Let $\gamma : [0, T] \rightarrow \mathbb{R}^d$ be a path that is C^1 on $[0, T]$ and piecewise C^2 on each interval $[t_{j-1}, t_j]$.

The optimal solution is given by a piecewise cubic polynomial:

$$\begin{aligned}\gamma(t_{j-1} + \tau) &= \frac{(\mathbf{v}_{j-1} + \mathbf{v}_j)\Delta t_j - 2(\mathbf{x}_j - \mathbf{x}_{j-1})}{(\Delta t_j)^3} \tau^3 \\ &\quad + \frac{3(\mathbf{x}_j - \mathbf{x}_{j-1}) - (2\mathbf{v}_{j-1} + \mathbf{v}_j)\Delta t_j}{(\Delta t_j)^2} \tau^2 + \mathbf{v}_{j-1}\tau + \mathbf{x}_{j-1},\end{aligned}\tag{51}$$

for $\tau \in [0, \Delta t_j]$, where $\Delta t_j = t_j - t_{j-1}$,

According to Corollary 4.2, it suffices to minimize

$$\begin{aligned}\mathcal{J}(\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_K) &= 2 \sum_{j=1}^K \frac{1}{(\Delta t_j)^3} \left\{ (\|\mathbf{v}_{j-1}\|^2 + \langle \mathbf{v}_{j-1}, \mathbf{v}_j \rangle + \|\mathbf{v}_j\|^2)(\Delta t_j)^2 \right. \\ &\quad \left. + 3\langle \mathbf{v}_{j-1} + \mathbf{v}_j, \mathbf{x}_{j-1} - \mathbf{x}_j \rangle \Delta t_j + 3\|\mathbf{x}_{j-1} - \mathbf{x}_j\|^2 \right\}\end{aligned}\tag{52}$$

$\mathcal{J} \geq 0$ always holds, so the formula above is a semi-positive quadratic form.

Take differentiation, we have

$$\nabla_{\mathbf{v}_j} \mathcal{J} = 2 \left(\frac{\mathbf{v}_{j-1} + 2\mathbf{v}_j}{\Delta t_j} + \frac{2\mathbf{v}_j + \mathbf{v}_{j+1}}{\Delta t_{j+1}} \right) + 6 \left(\frac{\mathbf{x}_{j-1} - \mathbf{x}_j}{(\Delta t_j)^2} + \frac{\mathbf{x}_j - \mathbf{x}_{j+1}}{(\Delta t_{j+1})^2} \right), \quad \text{for } j = 1, \dots, K-1,\tag{53}$$

$$\nabla_{\mathbf{v}_0} \mathcal{J} = \frac{2(2\mathbf{v}_0 + \mathbf{v}_1)}{\Delta t_1} + \frac{6(\mathbf{x}_0 - \mathbf{x}_1)}{(\Delta t_1)^2}, \quad \nabla_{\mathbf{v}_K} \mathcal{J} = \frac{2(\mathbf{v}_{K-1} + 2\mathbf{v}_K)}{\Delta t_K} + \frac{6(\mathbf{x}_{K-1} - \mathbf{x}_K)}{(\Delta t_K)^2}.\tag{54}$$

Set the gradients to zero and rearranging the terms to separate the velocities $\mathbf{v}$ and positions $\mathbf{x}$, we obtain:

For $j = 0$:

$$\frac{2}{\Delta t_1} \mathbf{v}_0 + \frac{1}{\Delta t_1} \mathbf{v}_1 = \frac{3(\mathbf{x}_1 - \mathbf{x}_0)}{(\Delta t_1)^2}.\tag{55}$$

For $j = 1, \dots, K-1$:

$$\frac{1}{\Delta t_j} \mathbf{v}_{j-1} + 2 \left(\frac{1}{\Delta t_j} + \frac{1}{\Delta t_{j+1}} \right) \mathbf{v}_j + \frac{1}{\Delta t_{j+1}} \mathbf{v}_{j+1} = 3 \left(\frac{\mathbf{x}_j - \mathbf{x}_{j-1}}{(\Delta t_j)^2} + \frac{\mathbf{x}_{j+1} - \mathbf{x}_j}{(\Delta t_{j+1})^2} \right).\tag{56}$$

For $j = K$:

$$\frac{1}{\Delta t_K} \mathbf{v}_{K-1} + \frac{2}{\Delta t_K} \mathbf{v}_K = \frac{3(\mathbf{x}_K - \mathbf{x}_{K-1})}{(\Delta t_K)^2}.\tag{57}$$

Denote $V = [\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_K] \in \mathbb{R}^{d \times (K+1)}$ and writing this system in matrix form $VA = B$, where

$$A = \begin{bmatrix} \frac{2}{\Delta t_1} & \frac{1}{\Delta t_1} & 0 & \cdots & 0 & 0 \\ \frac{1}{\Delta t_1} & 2\left(\frac{1}{\Delta t_1} + \frac{1}{\Delta t_2}\right) & \frac{1}{\Delta t_2} & \cdots & 0 & 0 \\ 0 & \frac{1}{\Delta t_2} & 2\left(\frac{1}{\Delta t_2} + \frac{1}{\Delta t_3}\right) & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \frac{1}{\Delta t_{K-1}} & 0 \\ 0 & 0 & \cdots & \frac{1}{\Delta t_{K-1}} & 2\left(\frac{1}{\Delta t_{K-1}} + \frac{1}{\Delta t_K}\right) & \frac{1}{\Delta t_K} \\ 0 & 0 & \cdots & 0 & \frac{1}{\Delta t_K} & \frac{1}{\Delta t_K} \end{bmatrix} \in \mathbb{R}^{(K+1) \times (K+1)}\tag{58}$$

$$B = \left[\frac{3(\mathbf{x}_1 - \mathbf{x}_0)}{(\Delta t_1)^2}, 3 \left(\frac{\mathbf{x}_1 - \mathbf{x}_0}{(\Delta t_1)^2} + \frac{\mathbf{x}_2 - \mathbf{x}_1}{(\Delta t_2)^2} \right), \dots, \right. \\ \left. 3 \left(\frac{\mathbf{x}_{K-1} - \mathbf{x}_{K-2}}{(\Delta t_{K-1})^2} + \frac{\mathbf{x}_K - \mathbf{x}_{K-1}}{(\Delta t_K)^2} \right), \frac{3(\mathbf{x}_K - \mathbf{x}_{K-1})}{(\Delta t_K)^2} \right] \in \mathbb{R}^{d \times (K+1)} \quad (59)$$

The solution satisfies $\gamma'(t_j) = \mathbf{v}_j$ for all $j = 0, 1, \dots, K$.

Since A is strictly diagonally dominant, the solution V is unique. $\square$

A.8 Proof of Theorem A.1

Theorem A.1. Let $\mathbf{a}_\theta(\mathbf{x}, \mathbf{v}, t)$ be L_θ -Lipschitz continuous of $(\mathbf{x}, \mathbf{v})$. The Wasserstein-2 distance between the generated distribution $\hat{\rho}_{t_i}^{(pos)}(\mathbf{x}) = \int \hat{\rho}_{t_i}(\mathbf{x}, \mathbf{v}) d\mathbf{v}$ and the true distribution $\mu_i^{(pos)}(\mathbf{x})$ in the position space is bounded by

$$\mathcal{W}_2^2(\hat{\rho}_{t_i}^{(pos)}, \mu_i^{(pos)}) \leq 2e^{\kappa(t_i - t_0)} \left(\mathcal{L}_{v_0} + (t_i - t_0) \mathcal{L}_{AFM}|_{[t_0, t_i]} \right) \quad (60)$$

where the cumulative AFM loss and the initial velocity loss are defined as:

$$\mathcal{L}_{AFM}|_{[t_0, t_i]} = \mathbb{E}_{t, (\mathbf{x}, \mathbf{v}) \sim \rho_t} \|\mathbf{a}_\theta(\mathbf{x}, \mathbf{v}, t) - \mathbf{a}(\mathbf{x}, \mathbf{v}, t)\|^2, \mathcal{L}_{v_0} = \mathbb{E}_{(\mathbf{x}_0, \mathbf{v}_0) \sim \rho_{t_0}} \|\mathbf{v}_0 - \mathbf{v}_0\|^2.$$

and κ is a constant only depending on L_θ .

Proof. We address the problem directly within the augmented space $\mathcal{X} \times \mathcal{V}$. Let the system state be denoted by $\mathbf{s} = [\mathbf{x}, \mathbf{v}]^T$, and let $\mu_i(\mathbf{s})$ represent the data distribution in this space. We define a flow map $\phi_t : [t_0, t_K] \times \mathcal{X} \times \mathcal{V} \rightarrow \mathcal{X} \times \mathcal{V}$, which is governed by the exact marginal acceleration field $\mathbf{a}(\mathbf{s}, t)$ defined in Equation (13).

Using $(\phi_t)_\#$ to denote the push-forward operator, the time-dependent density evolves as $\rho_t = (\phi_t)_\# \rho_{t_0}$, where $\rho_{t_0} = \mu_0$. Similarly, let ϕ_t^θ be the flow map induced by the learned acceleration field $\mathbf{a}_\theta(\mathbf{s}, t)$. The generated distribution $\hat{\rho}_t$ then evolves according to $\hat{\rho}_t = (\phi_t^\theta)_\# \hat{\rho}_{t_0}$. It is important to note that $\hat{\rho}_{t_0} \neq \rho_{t_0}$, as the conditional distribution of the initial velocity is approximated via a neural network.

We introduce a bridging measure $\nu_t = (\phi_t^\theta)_\# \rho_{t_0}$, then the Wasserstein-2 distance between joint distributions can be decomposed as

$$\mathcal{W}_2^2(\rho_t, \hat{\rho}_t) \leq (\mathcal{W}_2(\rho_t, \nu_t) + \mathcal{W}_2(\nu_t, \hat{\rho}_t))^2 \leq 2 \left(\mathcal{W}_2^2(\rho_t, \nu_t) + \mathcal{W}_2^2(\nu_t, \hat{\rho}_t) \right) \quad (61)$$

Estimation of Term 1: Define $Q_t = \int \mu_0(\mathbf{s}_0) \|\phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0)\|^2 d\mathbf{s}_0$, here $\mu_0 = \rho_{t_0}$ according to definition. Noting that $Q_{t_0} = 0$ since the starting points are identical. Define $A = \begin{bmatrix} \mathbf{0} & I \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$.

Taking the time derivative, we have:

$$\begin{aligned} \frac{dQ_t}{dt} &= 2 \int \mu_0(\mathbf{s}_0) (\phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0))^T \left(\frac{d\phi_t^\theta(\mathbf{s}_0)}{dt} - \frac{d\phi_t(\mathbf{s}_0)}{dt} \right) d\mathbf{s}_0 \\ &= 2 \int \mu_0(\mathbf{s}_0) (\phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0))^T \left(A\phi_t^\theta(\mathbf{s}_0) - A\phi_t(\mathbf{s}_0) + \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_\theta(\phi_t^\theta(\mathbf{s}_0), t) \end{bmatrix} - \begin{bmatrix} \mathbf{0} \\ \mathbf{a}(\phi_t(\mathbf{s}_0), t) \end{bmatrix} \right) d\mathbf{s}_0 \\ &= 2 \int \mu_0(\mathbf{s}_0) \left[(\phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0))^T A (\phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0)) \right. \\ &\quad + (\phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0))^T \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_\theta(\phi_t^\theta(\mathbf{s}_0), t) - \mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t) \end{bmatrix} \\ &\quad \left. + (\phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0))^T \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t) - \mathbf{a}(\phi_t(\mathbf{s}_0), t) \end{bmatrix} \right] d\mathbf{s}_0 \end{aligned} \quad (62)$$

We bound the three terms inside the integral separately. Let $\Delta\phi_t(\mathbf{s}_0) = \phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0)$.

1. Since $\|A\| \leq 1$, we have $\Delta\phi_t(\mathbf{s}_0)^T A \Delta\phi_t(\mathbf{s}_0) \leq \|\Delta\phi_t(\mathbf{s}_0)\|^2$.

2. Using the Lipschitz continuity of the network $\mathbf{a}_\theta$ with constant L_θ :

$$\begin{aligned} \Delta\phi_t(\mathbf{s}_0)^T \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_\theta(\phi_t^\theta(\mathbf{s}_0), t) - \mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t) \end{bmatrix} &\leq \|\Delta\phi_t(\mathbf{s}_0)\| \cdot \|\mathbf{a}_\theta(\phi_t^\theta(\mathbf{s}_0), t) - \mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t)\| \\ &\leq L_\theta \|\Delta\phi_t(\mathbf{s}_0)\|^2. \end{aligned} \quad (63)$$

3. Using Cauchy-Schwarz and the inequality $2xy \leq x^2 + y^2$:

$$\begin{aligned} \Delta\phi_t(\mathbf{s}_0)^T \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t) - \mathbf{a}(\phi_t(\mathbf{s}_0), t) \end{bmatrix} &\leq \|\Delta\phi_t(\mathbf{s}_0)\| \cdot \|\mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t) - \mathbf{a}(\phi_t(\mathbf{s}_0), t)\| \\ &\leq \frac{1}{2} \|\Delta\phi_t(\mathbf{s}_0)\|^2 + \frac{1}{2} \|\mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t) - \mathbf{a}(\phi_t(\mathbf{s}_0), t)\|^2. \end{aligned} \quad (64)$$

Substituting these bounds back into the derivative:

$$\begin{aligned} \frac{dQ_t}{dt} &\leq 2 \int \mu_0(\mathbf{s}_0) \left[(1 + L_\theta + \frac{1}{2}) \|\Delta\phi_t(\mathbf{s}_0)\|^2 + \frac{1}{2} \|\mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t) - \mathbf{a}(\phi_t(\mathbf{s}_0), t)\|^2 \right] d\mathbf{s}_0 \\ &= (2L_\theta + 3)Q_t + \int \mu_0(\mathbf{s}_0) \|\mathbf{a}_\theta(\phi_t(\mathbf{s}_0), t) - \mathbf{a}(\phi_t(\mathbf{s}_0), t)\|^2 d\mathbf{s}_0 \end{aligned} \quad (65)$$

By Gronwall's inequality, since $Q_{t_0} = 0$:

$$Q_t \leq e^{(2L_\theta+3)(t-t_0)} \int_{t_0}^t \left(\int \mu_0(\mathbf{s}_0) \|\mathbf{a}_\theta(\phi_\tau(\mathbf{s}_0), \tau) - \mathbf{a}(\phi_\tau(\mathbf{s}_0), \tau)\|^2 d\mathbf{s}_0 \right) d\tau \quad (66)$$

By the definition of Wasserstein-2 distance, we have:

$$\begin{aligned} t_0 \mathcal{W}_2^2(\rho_t, \nu_t) &= \min_{\pi} \int \|\mathbf{s}_1 - \mathbf{s}_2\|^2 d\pi \quad \text{s.t.} \quad \int \pi d\mathbf{s}_1 = \rho_t, \int \pi d\mathbf{s}_2 = \nu_t \\ &\leq \int \mu_0(\mathbf{s}_0) \|\phi_t^\theta(\mathbf{s}_0) - \phi_t(\mathbf{s}_0)\|^2 d\mathbf{s}_0 \\ &= Q_t \leq e^{(2L_\theta+3)(t-t_0)} \int_{t_0}^t \mathbb{E}_{\mathbf{s} \sim \mu_\tau} \|\mathbf{a}_\theta(\mathbf{s}, \tau) - \mathbf{a}(\mathbf{s}, \tau)\|^2 d\tau \end{aligned} \quad (67)$$

$$= e^{(2L_\theta+3)(t-t_0)} (t - t_0) \mathcal{L}_{\text{AFM}}|_{[t_0, t]}. \quad (68)$$

Estimation of Term 2: Take $\pi_0(\mathbf{s}_0, \mathbf{s}_0^\xi)$ as the optimal coupling that attains $\mathcal{W}_2^2(\nu_{t_0}, \hat{\rho}_{t_0}) = \int \|\mathbf{s}_0 - \mathbf{s}_0^\xi\|^2 d\pi_0(\mathbf{s}_0, \mathbf{s}_0^\xi)$. Note that $\nu_{t_0} = \mu_0$ (the true data distribution).

From a pair of starting points $(\mathbf{s}_0, \mathbf{s}_0^\xi) \sim \pi_0$, define the square distance between trajectories under the same approximate flow as $\Delta(t) = \|\phi_t^\theta(\mathbf{s}_0) - \phi_t^\theta(\mathbf{s}_0^\xi)\|^2$. Taking the time derivative, we have:

$$\begin{aligned} \frac{d\Delta(t)}{dt} &= 2(\phi_t^\theta(\mathbf{s}_0) - \phi_t^\theta(\mathbf{s}_0^\xi))^T \left(\frac{d\phi_t^\theta(\mathbf{s}_0)}{dt} - \frac{d\phi_t^\theta(\mathbf{s}_0^\xi)}{dt} \right) \\ &= 2(\phi_t^\theta(\mathbf{s}_0) - \phi_t^\theta(\mathbf{s}_0^\xi))^T \left(A(\phi_t^\theta(\mathbf{s}_0) - \phi_t^\theta(\mathbf{s}_0^\xi)) + \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_\theta(\phi_t^\theta(\mathbf{s}_0), t) - \mathbf{a}_\theta(\phi_t^\theta(\mathbf{s}_0^\xi), t) \end{bmatrix} \right) \\ &\leq 2 \left(\|A\| \cdot \|\phi_t^\theta(\mathbf{s}_0) - \phi_t^\theta(\mathbf{s}_0^\xi)\|^2 + \|\phi_t^\theta(\mathbf{s}_0) - \phi_t^\theta(\mathbf{s}_0^\xi)\| \cdot \|\mathbf{a}_\theta(\phi_t^\theta(\mathbf{s}_0), t) - \mathbf{a}_\theta(\phi_t^\theta(\mathbf{s}_0^\xi), t)\| \right) \\ &\leq 2(1 + L_\theta)\Delta(t). \end{aligned} \quad (69)$$

Again by Gronwall's inequality, we have:

$$\Delta(t) \leq e^{2(1+L_\theta)(t-t_0)} \Delta(t_0) = e^{2(1+L_\theta)(t-t_0)} \|\mathbf{s}_0 - \mathbf{s}_0^\xi\|^2. \quad (70)$$

Since $\nu_t = (\phi_t^\theta)_\# \nu_{t_0}$ and $\hat{\rho}_t = (\phi_t^\theta)_\# \hat{\rho}_{t_0}$, the push-forward of the optimal coupling $\pi_t = (\phi_t^\theta, \phi_t^\theta)_\# \pi_0$ is a valid coupling for $(\nu_t, \hat{\rho}_t)$. Integrate the squared distance over all starting pairs with respect to π_0 , we obtain:

$$\begin{aligned} \mathcal{W}_2^2(\nu_t, \hat{\rho}_t) &\leq \int \|\phi_t^\theta(\mathbf{s}_0) - \phi_t^\theta(\mathbf{s}_0^\xi)\|^2 d\pi_0(\mathbf{s}_0, \mathbf{s}_0^\xi) \\ &\leq \int e^{2(1+L_\theta)(t-t_0)} \|\mathbf{s}_0 - \mathbf{s}_0^\xi\|^2 d\pi_0(\mathbf{s}_0, \mathbf{s}_0^\xi) \\ &= e^{2(1+L_\theta)(t-t_0)} \mathcal{W}_2^2(\nu_{t_0}, \hat{\rho}_{t_0}) \end{aligned} \quad (71)$$

Let $\tilde{\pi}(\mathbf{s}_0, \mathbf{s}_0^\xi)$ be the coupling between ν_0 and $\hat{\rho}_0$ induced by the identity mapping on the position space, i.e., pairing $\mathbf{s}_0 = (\mathbf{x}_0, \mathbf{v}_0)$ with $\mathbf{s}_0^\xi = (\mathbf{x}_0, \mathbf{v}_0^\xi)$. We can bound Wasserstein distance by L_2 loss:

$$\begin{aligned} \mathcal{W}_2^2(\nu_{t_0}, \hat{\rho}_{t_0}) &\leq \int \|\mathbf{s}_0 - \mathbf{s}_0^\xi\|^2 d\tilde{\pi}(\mathbf{s}_0, \mathbf{s}_0^\xi) \\ &= \int (\|\mathbf{x}_0 - \mathbf{x}_0^\xi\|^2 + \|\mathbf{v}_0 - \mathbf{v}_0^\xi\|^2) d\tilde{\pi}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}_0^\xi, \mathbf{v}_0^\xi) \\ &= \int \|\mathbf{v}_0 - \mathbf{v}_0^\xi\|^2 d\tilde{\pi}(\mathbf{x}_0, \mathbf{v}_0, \mathbf{x}_0^\xi, \mathbf{v}_0^\xi) = \mathcal{L}_{v_0} \end{aligned} \quad (72)$$

Thus $\mathcal{W}_2^2(\nu_t, \hat{\rho}_t) \leq e^{2(1+L_\theta)(t-t_0)} \mathcal{L}_{v_0}$.

Combining the bounds above, we have

$$\begin{aligned} \mathcal{W}_2^2(\hat{\rho}_{t_i}, \rho_{t_i}) &\leq 2 \left[\mathcal{W}_2^2(\hat{\rho}_{t_i}, \nu_{t_i}) + \mathcal{W}_2^2(\nu_{t_i}, \rho_{t_i}) \right] \\ &\leq 2 \left[e^{(2L_\theta+3)(t_i-t_0)} \cdot (t_i - t_0) \mathcal{L}_{\text{AFM}}|_{[t_0, t_i]} + e^{(2L_\theta+2)(t_i-t_0)} \mathcal{L}_{v_0} \right]. \end{aligned} \quad (73)$$

Define $\kappa = (2L_\theta + 3)$, then the final bound is obtained by:

$$\begin{aligned} \mathcal{W}_2^2(\hat{\rho}_{t_i}^{(\text{pos})}, \rho_{t_i}^{(\text{pos})} = \mu_i^{(\text{pos})}) &= \min_{\pi_{\mathbf{x}}} \int \|\mathbf{x}_1 - \mathbf{x}_2\|^2 d\pi_{\mathbf{x}}(\mathbf{x}_1, \mathbf{x}_2) \quad \text{s.t.} \quad \int \pi_{\mathbf{x}} d\mathbf{x}_2 = \rho_{t_i}^{(\text{pos})}, \quad \int \pi_{\mathbf{x}} d\mathbf{x}_1 = \hat{\rho}_{t_i}^{(\text{pos})} \\ &\leq \min_{\pi} \int \|\mathbf{x}_1 - \mathbf{x}_2\|^2 d\pi(\mathbf{s}_1, \mathbf{s}_2) \quad \text{s.t.} \quad \int \pi d\mathbf{s}_2 = \rho_{t_i}, \quad \int \pi d\mathbf{s}_1 = \hat{\rho}_{t_i} \\ &\leq \min_{\pi} \int \|\mathbf{s}_1 - \mathbf{s}_2\|^2 d\pi(\mathbf{s}_1, \mathbf{s}_2) \end{aligned} \quad (74)$$

$$= \mathcal{W}_2^2(\hat{\rho}_{t_i}, \rho_{t_i}) \leq 2e^{\kappa(t_i-t_0)} \left((t_i - t_0) \mathcal{L}_{\text{AFM}}|_{[t_0, t_i]} + \mathcal{L}_{v_0} \right). \quad (75)$$

This completes the proof. Note that our derivation holds even for multi-valued or stochastic cases, since the upper bound on the Wasserstein distance is established via a valid coupling, which remains well-defined regardless of whether the mapping is deterministic. $\square$

B Datasets and Evaluation Metric

B.1 Experiment Setup

The experiments were conducted on a shared high-performance computing (HPC) cluster. All experimental runs utilized GPUs equipped with 80GB of memory. Both the network $\mathbf{a}_\theta(\mathbf{x}, \mathbf{v}, t)$, employed to approximate acceleration, and the network $\mathbf{v}_\phi(\mathbf{x}, \mathbf{v}, t)$, employed to

approximate velocity, utilize residual architectures. Specifically, $\mathbf{a}_\theta(\mathbf{x}, \mathbf{v}, t)$ consists of 4 hidden layers, while $\mathbf{v}_\phi(\mathbf{x}, \mathbf{v}, t)$ comprises 2 hidden layers. The Optimal Transport (OT) calculations were implemented using the Python Optimal Transport (POT) library [16], and the training process of neural networks were implemented using Pytorch [43].

In all experiments, TracingFlow employs a completely consistent set of hyperparameters. When solving for OT using the Sinkhorn algorithm, the regularization coefficient is set to $\epsilon = 1 \times 10^{-3}$. When incorporating biological priors, the penalty coefficient for barcode mismatches is set to $p_0 = 25$. Since the number of cells at each time point in the datasets processed in this paper is on the order of 10^3 , we do not utilize mini-batch OT; instead, we compute the Transport Plan on the full dataset.

B.2 Evaluation Metrics

We employ the 1-Wasserstein Distance ($\mathcal{W}_1$) and the 2-Wasserstein Distance ($\mathcal{W}_2$) to quantify the discrepancy between the predicted distribution and the ground truth distribution. These metrics are defined as follows:

$$\mathcal{W}_1(\mu, \nu) = \min_{\pi \in \Pi(\mu, \nu)} \int \|\mathbf{x} - \mathbf{y}\| d\pi(\mathbf{x}, \mathbf{y}) \quad (76)$$

$$\mathcal{W}_2(\mu, \nu) = \left(\min_{\pi \in \Pi(\mu, \nu)} \int \|\mathbf{x} - \mathbf{y}\|_2^2 d\pi(\mathbf{x}, \mathbf{y}) \right)^{1/2} \quad (77)$$

For the evaluation of datasets containing lineage information, we utilize the lineage-weighted $\mathcal{W}_1$ and $\mathcal{W}_2$ distances. These are defined based on the decomposition of the distributions. Consider the distribution generated by the model, denoted as $\mu(\mathbf{x})$, and the reference distribution provided by the dataset, denoted as $\nu(\mathbf{x})$. Both consist of L lineages and can be decomposed as:

$$\mu(\mathbf{x}) = \sum_{i=1}^L w_i \mu^{(i)}(\mathbf{x}), \quad \nu(\mathbf{x}) = \sum_{i=1}^L v_i \nu^{(i)}(\mathbf{x}) \quad (78)$$

where $\mu^{(i)}(\mathbf{x})$ and $\nu^{(i)}(\mathbf{x})$ represent the distributions of the i -th lineage within $\mu(\mathbf{x})$ and $\nu(\mathbf{x})$, respectively, satisfying $\int \mu^{(i)}(\mathbf{x}) d\mathbf{x} = 1$ and $\int \nu^{(i)}(\mathbf{x}) d\mathbf{x} = 1$. The terms w_i and v_i denote the weights of the i -th lineage in $\mu(\mathbf{x})$ and $\nu(\mathbf{x})$, respectively, subject to the constraint $\sum_{i=1}^L w_i = \sum_{i=1}^L v_i = 1$. The lineage-weighted $\mathcal{W}_1$ and $\mathcal{W}_2$ distances are defined as:

$$\mathcal{W}_1^{(L)}(\mu \| \nu) = \sum_{i=1}^L v_i \mathcal{W}_1(\mu^{(i)}, \nu^{(i)}), \quad \mathcal{W}_2^{(L)}(\mu \| \nu) = \sum_{i=1}^L v_i \mathcal{W}_2(\mu^{(i)}, \nu^{(i)}) \quad (79)$$

In essence, we first calculate the $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the distributions of corresponding lineages in the two datasets. Subsequently, we compute a weighted average of these distances using the proportions of each lineage given by the reference distribution $\nu(\mathbf{x})$. It is important to note that the lineage-weighted $\mathcal{W}_1$ and $\mathcal{W}_2$ metrics serve solely as indicators of the model's ability to preserve lineage prior information while learning the dynamics. They are not true distance metrics in the mathematical sense, as they are evidently asymmetric with respect to μ and ν .

B.3 Datasets

3D Simulation Lineage Data. Consider the following dynamics with three lineage barcodes and a twist structure:

$$\begin{cases} dX_t^{(0)} = v dt + \sigma_x dB_t^X \\ dY_t^{(0)} = \sigma_y dB_t^Y \\ dZ_t^{(0)} = \sigma_z dB_t^Z \end{cases} \quad \begin{cases} dX_t^{(1)} = v dt + \sigma_x dB_t^X \\ dY_t^{(1)} = -\mu Z_t^{(1)} dt + \sigma_y dB_t^Y \\ dZ_t^{(1)} = \mu Y_t^{(1)} dt + \sigma_z dB_t^Z \end{cases} \quad \begin{cases} dX_t^{(2)} = v dt + \sigma_x dB_t^X \\ dY_t^{(2)} = \mu Z_t^{(2)} dt + \sigma_y dB_t^Y \\ dZ_t^{(2)} = -\mu Y_t^{(2)} dt + \sigma_z dB_t^Z \end{cases} \quad (80)$$

We take $v = 3.0, \mu = 3.14, \sigma_x = \sigma_y = \sigma_z = 0.1$ and time points $t = 0, 1, 2$. Initialize $(X_0^{(0)}, Y_0^{(0)}, Z_0^{(0)}) \sim \mathcal{N}([0, 0, 0], 0.5I)$, $(X_0^{(1)}, Y_0^{(1)}, Z_0^{(1)}) \sim \mathcal{N}([0, -2, 0], 0.5I)$, $(X_0^{(2)}, Y_0^{(2)}, Z_0^{(2)}) \sim \mathcal{N}([0, 2, 0], 0.5I)$ and sample 500 particles for each.




Figure 5: Dynamics of simulation datasets. (a) 2D simulation data, (b) 3D simulation data and (c) 3D simulation data projected to XY plane.

Cite Data. We utilized the Cite-seq dataset introduced in [31], which comprises 31,240 cells collected across 4 distinct time points. To evaluate the performance of TracingFlow on real-world biological data, we applied Principal Component Analysis (PCA) to reduce the dimensionality of the gene expression profiles to 100 and 5 dimensions, respectively.

Embryoid Bodies Data. We employed the Embryoid Bodies (EB) dataset from [39], consisting of 16,819 cells sampled at 5 time points. We utilized PCA to reduce the dimensionality of the gene expression data to 5 dimensions.

Hematopoiesis Data. We employed the Hematopoiesis dataset from [64], consisting of 130,887 cells sampled at 3 time points. Part of the cells are recorded with barcodes. We selected 22,329 cells with barcodes from the dataset and utilized PCA to reduce the dimensionality of gene expression data to 50 dimensions. When processing this dataset, we filtered for data where barcodes were present at all three time points. The barcodes for the remaining data were set to undefined, incurring no additional transport cost during the incorporation of biological priors (Section 5.3).

Mexico Gulf Data To assess the capability and robustness of Tracing Flow in capturing long-term system dynamics, we utilize the Mexico Gulf dataset, following [56]. This two-dimensional dataset comprises 9 time points.

C Experiment Details

C.1 Experiment Details across various datasets

Simulation 2D Data Table 4 presents the performance of TracingFlow and other algorithms on the 2D Simulation dataset. On average, TracingFlow reconstructs the data distribution at each time point with higher accuracy.

Cite 5D & Cite 100D Data We present the distribution reconstruction performance of each algorithm on the Cite 5D and Cite 100D datasets in Table 5 and Table 6, respectively. On these

Table 4: $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground-truth distributions at each time point on the 2D Simulation dataset.

Method	Dynamics	$t = 1$		$t = 2$		$t = 3$		$t = 4$	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	0.1215 $\pm$ 0.0164	0.1441 $\pm$ 0.0164	0.3644 $\pm$ 0.1177	0.4143 $\pm$ 0.1168	0.5863 $\pm$ 0.2615	0.6238 $\pm$ 0.2657	1.0356 $\pm$ 0.2325	1.1867 $\pm$ 0.2890
SF ² M		0.2389 $\pm$ 0.1217	0.2668 $\pm$ 0.1305	0.5700 $\pm$ 0.2574	0.6108 $\pm$ 0.2654	1.0630 $\pm$ 0.2769	1.3557 $\pm$ 0.2770	2.0986 $\pm$ 1.3198	2.2938 $\pm$ 1.3193
3MSBM	2nd Order	1.4029 $\pm$ 0.6892	1.4459 $\pm$ 0.6596	0.9318 $\pm$ 0.2553	1.0181 $\pm$ 0.2050	1.8079 $\pm$ 0.7761	1.9162 $\pm$ 0.7529	0.5993 $\pm$ 0.6176	0.7265 $\pm$ 0.6935
MMFM		1.3226 $\pm$ 0.1774	1.9287 $\pm$ 0.3002	1.8387 $\pm$ 0.5624	1.1779 $\pm$ 0.5114	1.8158 $\pm$ 0.1794	1.8740 $\pm$ 0.1701	4.4984 $\pm$ 0.2937	4.6059 $\pm$ 0.2977
HRF		1.2036 $\pm$ 0.0425	1.2096 $\pm$ 0.0244	0.9955 $\pm$ 0.0921	1.0009 $\pm$ 0.0873	1.9829 $\pm$ 0.1311	2.0040 $\pm$ 0.1582	1.9740 $\pm$ 0.1873	2.4379 $\pm$ 0.2157
CAF		0.4051 $\pm$ 0.1712	0.5117 $\pm$ 0.1738	0.8857 $\pm$ 0.3630	0.8996 $\pm$ 0.3689	1.3274 $\pm$ 0.1194	1.3731 $\pm$ 0.1107	3.9358 $\pm$ 0.4572	3.7880 $\pm$ 0.4530
TF (Ours)		0.1609 $\pm$ 0.0232	0.1787 $\pm$ 0.0258	0.1753 $\pm$ 0.0299	0.2045 $\pm$ 0.0277	0.4091 $\pm$ 0.0523	0.4821 $\pm$ 0.0698	0.6404 $\pm$ 0.2661	1.1436 $\pm$ 0.1577

datasets TracingFlow achieves superior reconstruction accuracy at all time steps except for $t = 3$. The trajectories learned by TracingFlow on these two datasets are illustrated in Figure 6.

Table 5: $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground-truth distributions at each time point on the Cite 5D dataset

Method	Dynamics	$t = 1$		$t = 2$		$t = 3$	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	0.5615 $\pm$ 0.0156	0.6484 $\pm$ 0.0268	0.8472 $\pm$ 0.0463	0.9416 $\pm$ 0.0443	1.0536 $\pm$ 0.0534	1.1566 $\pm$ 0.0573
SF ² M		0.5932 $\pm$ 0.0249	0.6797 $\pm$ 0.0424	0.7768 $\pm$ 0.0451	0.8670 $\pm$ 0.0480	0.9251 $\pm$ 0.0279	1.0221 $\pm$ 0.0352
3MSBM	2nd Order	5.3735 $\pm$ 0.3961	5.4244 $\pm$ 0.3582	4.3118 $\pm$ 0.4589	4.5458 $\pm$ 0.4141	1.2254 $\pm$ 0.2998	1.3427 $\pm$ 0.3307
MMFM		0.5335 $\pm$ 0.0134	0.5911 $\pm$ 0.0057	2.2290 $\pm$ 0.2151	2.4557 $\pm$ 0.1613	4.8413 $\pm$ 0.2538	5.2261 $\pm$ 0.4061
HRF		0.4225 $\pm$ 0.0152	0.4622 $\pm$ 0.0185	0.6227 $\pm$ 0.0421	0.6814 $\pm$ 0.0392	0.7638 $\pm$ 0.0738	0.8394 $\pm$ 0.0713
CAF		5.1724 $\pm$ 0.0452	5.6242 $\pm$ 0.0763	5.4333 $\pm$ 0.0914	5.4637 $\pm$ 0.1425	3.0710 $\pm$ 0.1853	3.1354 $\pm$ 0.2672
TF (Ours)		0.2682 $\pm$ 0.0241	0.2916 $\pm$ 0.0312	0.6094 $\pm$ 0.0615	0.6516 $\pm$ 0.0684	0.8808 $\pm$ 0.1208	0.9725 $\pm$ 0.1329

Table 6: $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground-truth distributions at each time point on the Cite 100D dataset

Method	Dynamics	$t = 1$		$t = 2$		$t = 3$	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	10.1136 $\pm$ 0.0104	10.1873 $\pm$ 0.0135	10.2504 $\pm$ 0.0209	10.3145 $\pm$ 0.0201	11.2682 $\pm$ 0.0938	11.3642 $\pm$ 0.0997
SF ² M		10.3598 $\pm$ 0.0051	10.4448 $\pm$ 0.0053	11.4647 $\pm$ 0.0384	11.5456 $\pm$ 0.0402	15.2032 $\pm$ 0.1813	15.4367 $\pm$ 0.1949
3MSBM	2nd Order	19.9677 $\pm$ 0.1373	20.1015 $\pm$ 0.1105	15.3218 $\pm$ 0.1232	15.4089 $\pm$ 0.1343	12.0229 $\pm$ 0.2895	12.1395 $\pm$ 0.3621
MMFM		9.8726 $\pm$ 0.0383	9.9305 $\pm$ 0.0361	12.1374 $\pm$ 0.2422	12.3076 $\pm$ 0.2495	15.5960 $\pm$ 1.4972	16.1222 $\pm$ 2.0195
HRF		9.9660 $\pm$ 0.0215	10.0580 $\pm$ 0.0341	9.9451 $\pm$ 0.0632	9.9986 $\pm$ 0.1526	10.5485 $\pm$ 0.1082	10.6442 $\pm$ 0.1254
CAF		18.1182 $\pm$ 0.3214	18.1971 $\pm$ 0.2845	17.4947 $\pm$ 0.4267	17.5767 $\pm$ 0.4912	14.6158 $\pm$ 0.5632	15.1773 $\pm$ 0.5215
TF (Ours)		2.8955 $\pm$ 0.1429	3.2763 $\pm$ 0.1776	9.8235 $\pm$ 0.1489	9.8773 $\pm$ 0.1827	11.4910 $\pm$ 0.1693	11.6054 $\pm$ 0.2064

EB 5D Data Table 7 displays the interpolation accuracy on the EB 5D dataset, evaluated by holding out one time point during training. TracingFlow outperformed other methods in reconstruction accuracy for all time points except $t = 1$. Figure 7 illustrates the interpolated distributions generated by TracingFlow across the four time points.

Table 7: $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between predicted and ground-truth distributions at held-out time points on the EB 5D dataset.

Method	Dynamics	$t = 1$		$t = 2$		$t = 3$		$t = 4$	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	3.4335 $\pm$ 0.1465	3.8462 $\pm$ 0.1449	4.8404 $\pm$ 0.1496	5.5792 $\pm$ 0.1358	5.0713 $\pm$ 0.2839	5.6448 $\pm$ 0.2861	7.1368 $\pm$ 0.3799	7.6938 $\pm$ 0.4574
SF ² M		3.4204 $\pm$ 0.0637	3.6713 $\pm$ 0.0589	3.5511 $\pm$ 0.0394	3.9441 $\pm$ 0.0418	4.3175 $\pm$ 0.2517	4.8738 $\pm$ 0.3247	6.9516 $\pm$ 0.0616	7.4727 $\pm$ 0.0795
3MSBM	2nd Order	6.1086 $\pm$ 1.1215	6.5434 $\pm$ 1.2227	7.3626 $\pm$ 0.8043	8.0596 $\pm$ 1.0252	7.5862 $\pm$ 1.1170	8.1734 $\pm$ 0.7900	7.7029 $\pm$ 0.4241	7.1022 $\pm$ 0.3971
MMFM		10.7803 $\pm$ 0.6718	11.8376 $\pm$ 0.7823	5.7932 $\pm$ 0.2162	6.2242 $\pm$ 0.1828	9.3326 $\pm$ 0.5617	9.9837 $\pm$ 0.4994	13.2046 $\pm$ 2.1528	14.3464 $\pm$ 2.2982
TF (Ours)		3.5611 $\pm$ 0.0362	3.8930 $\pm$ 0.0243	3.5002 $\pm$ 0.0983	3.8035 $\pm$ 0.1086	3.5643 $\pm$ 0.1140	3.7947 $\pm$ 0.4123	7.3639 $\pm$ 0.1316	7.6523 $\pm$ 0.1976

3D Simulation Lineage Dataset Table 8 demonstrates the ability of various algorithms to preserve biological priors on the 3D Simulation Lineage dataset. Since the trajectories learned




Figure 6: Trajectories learned by TracingFlow on the (a) Cite 5D and (b) Cite 100D datasets, projected to 2D using PCA.




Figure 7: Predicted distributions by TracingFlow at four held-out time points, projected to 2D using PCA.

by TracingFlow align with biological priors (see Figure 3), it achieves lower lineage-weighted $\mathcal{W}_1$ and $\mathcal{W}_2$ scores on average.

Table 8: lineage-weighted $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground-truth distributions at each time point on the 3D Simulation Lineage dataset

Method	Dynamics	$t = 1$		$t = 2$	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	2.6080 ± 0.0029	2.6869 ± 0.0030	1.8686 ± 0.0382	1.9456 ± 0.0403
SF ² M		2.8262 ± 0.0409	2.8373 ± 0.0462	0.2247 ± 0.0617	0.2533 ± 0.0568
3MSBM	2nd Order	4.0141 ± 0.8788	4.2262 ± 0.8596	0.7548 ± 0.0726	0.6751 ± 0.0573
MMFM		1.8585 ± 0.0350	2.1618 ± 0.0496	3.5695 ± 0.5692	4.1028 ± 0.6037
HRF		2.9150 ± 0.2595	2.9265 ± 0.3901	0.6717 ± 0.5491	0.6905 ± 0.4143
CAF		1.6284 ± 0.2249	1.6623 ± 0.3726	3.7328 ± 0.3305	3.8471 ± 0.3316
TF w/o Bio. Prior		2.8267 ± 0.2896	2.8376 ± 0.3714	0.2317 ± 0.0512	0.2687 ± 0.0892
TF(Ours)		0.4282 ± 0.1606	0.4989 ± 0.1585	0.4795 ± 0.2680	0.5667 ± 0.2654

Hematopoiesis Dataset Table 9 presents the performance of different algorithms in retaining biological priors on real-world datasets, measured by the lineage-weighted $\mathcal{W}_1, \mathcal{W}_2$ distance. TracingFlow surpassed other algorithms at all time points. We also evaluated the algorithms

without introducing biological priors by Vanilla $\mathcal{W}_1$, $\mathcal{W}_2$ distance, and as shown in Table 10, TracingFlow maintained superior performance. Figure 8 visualizes the data points generated by TracingFlow on the real-world dataset under both conditions (with and without biological priors).

Table 9: lineage-weighted $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground-truth distributions at each time point on the Hematopoiesis dataset

Method	Dynamics	$t = 1$		$t = 2$	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	1st Order	13.2972 ± 0.0804	13.7134 ± 0.0831	16.6918 ± 0.1973	17.0807 ± 0.2025
SF ² M		14.1427 ± 0.0341	14.4199 ± 0.0303	23.2913 ± 0.2590	23.5314 ± 0.2640
3MSBM	2nd Order	17.6487 ± 0.5250	17.8710 ± 0.5204	19.0139 ± 1.4418	19.4144 ± 1.3622
MMFM		18.4040 ± 2.8129	19.8538 ± 2.8906	36.0412 ± 3.0429	40.2282 ± 3.3934
HRF		14.1007 ± 0.0804	14.3122 ± 0.0831	16.8135 ± 0.1973	17.1154 ± 0.2025
CAF		13.1098 ± 0.0628	13.3263 ± 0.0742	17.3232 ± 0.0580	17.6685 ± 0.0572
TF w/o Bio. Prior		14.1260 ± 0.0526	14.5958 ± 0.0537	19.1073 ± 0.1932	19.4262 ± 0.2375
TF		11.5203 ± 0.0695	12.0151 ± 0.0809	16.5736 ± 0.2603	16.9027 ± 0.2999

Table 10: $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground-truth distributions at each time point on the Hematopoiesis dataset

Method	Dynamics	$t = 1$		$t = 2$		Average	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	First-order	9.8780	10.6132	10.7621	11.4466	10.3201	11.0299
SF ² M		10.2352	10.8017	13.2507	13.6879	11.7430	12.2428
3MSBM	Second-order	12.8214	13.2501	20.0262	20.6374	16.4238	16.9438
MMFM		10.4305	10.8695	22.3269	23.5322	16.3787	17.2008
HRF		10.6233	11.4510	10.9570	11.6201	10.7902	11.5356
CAF		21.8716	22.3309	17.1251	18.3914	19.4984	20.3612
TF w/o Bio. Prior		7.0139	7.4140	10.1642	10.5925	8.5891	9.0032




Figure 8: Original data (a), and data points generated by TracingFlow on the Hematopoiesis dataset without biological priors (b) and with biological priors (c), projected to 2D using UMAP.

Mexico Gulf Dataset Table 11 presents the performance of various algorithms on the Mexico Gulf dataset. Tracing Flow achieves the best distribution reconstruction accuracy at all time points, yielding smooth trajectories.

Table 11: $\mathcal{W}_1$ and $\mathcal{W}_2$ distances between the generated and ground-truth distributions at each time point on the Mexico Gulf dataset

Method	Dynamics	$t = 1$		$t = 2$		$t = 3$		$t = 4$	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	First-order	0.0336	0.0375	0.0791	0.0887	0.1367	0.1415	0.1409	0.1474
SF ² M		0.0768	0.0917	0.0802	0.0964	0.0819	0.0968	0.0809	0.0987
3MSBM	Second-order	0.4407	0.6305	0.5565	0.7617	0.5887	0.9113	0.5231	1.0714
MMFM		0.0346	0.0364	0.2961	0.3008	1.0538	1.0571	1.8755	1.8766
HRF		0.6205	0.6217	1.2004	1.2011	1.5334	1.5337	1.4986	1.5013
CAF		1.7659	1.7663	2.4713	2.4714	2.0742	2.0755	1.3040	1.3071
TF		0.0128	0.0145	0.0301	0.0339	0.0407	0.0458	0.0348	0.0456

Method	Dynamics	$t = 5$		$t = 6$		$t = 7$		$t = 8$		Average	
		$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
OT-CFM	First-order	0.1318	0.1370	0.1364	0.1425	0.1610	0.1799	0.1871	0.2236	0.1258	0.1373
SF ² M		0.1030	0.1191	0.0851	0.0956	0.1401	0.1494	0.1965	0.2068	0.1056	0.1193
3MSBM	Second-order	0.2185	1.2783	0.5362	2.0436	0.8474	4.5953	1.3606	13.2434	0.6340	3.0669
MMFM		2.3926	2.3928	2.7214	2.7215	2.9191	2.9199	2.9930	2.9945	1.7857	1.7875
HRF		1.3022	1.3075	0.9476	0.9518	0.6480	0.6535	0.3296	0.3341	1.0101	1.0131
CAF		0.7086	0.7124	0.4669	0.4751	0.6672	0.6914	0.9641	0.9914	1.3028	1.3113
TF		0.0545	0.0657	0.0778	0.0824	0.1039	0.1133	0.1206	0.1309	0.0594	0.0665

C.2 Training time and scalability of TracingFlow

To evaluate the scalability of TracingFlow, we report the training time required by each method across various datasets, as presented in Table 12. Experimental results indicate that Flow Matching algorithms based on second-order dynamics (e.g., 3 MSBM and our TracingFlow) generally require more training time compared to first-order dynamics algorithms (e.g., OT-CFM and SF²M). Notably, the training time of TracingFlow is slightly lower than that of 3 MSBM, demonstrating that our method does not introduce additional computational overhead.

Table 12: Training time comparison (in seconds) of different methods across various datasets.

Method	Simulation	Cite 5D	Cite 100D	SimLineage-3D	RealLineage
OT-CFM	12 s	175 s	218 s	13 s	196 s
SF ² M	26 s	79 s	84 s	34 s	61 s
MMFM	42 s	96 s	179 s	68 s	201 s
3 MSBM	220 s	495 s	910 s	244 s	855 s
CAF	94 s	274 s	382 s	256 s	264 s
HRF	247 s	271 s	359 s	289 s	347 s
TF	225 s	319 s	806 s	269 s	742 s

C.3 Sensitivity Analysis on Minibatch-OT

In all experiments presented in this paper, the full dataset was utilized to compute the OT plan. However, as the scale of the data increases, the computational cost of the OT plan grows rapidly, making the adoption of the minibatch computation method described in Section 5.2 inevitable. We conducted a sensitivity analysis on the Cite-5D dataset to evaluate the impact of batch size in minibatch-OT computation on the accuracy of distribution reconstruction. The results are presented in Table 13. Experimental results demonstrate that TracingFlow consistently and accurately reconstructs the distributions at each time point when the batch size is adjusted within a certain range (batch size $\sim 10^3$).

Table 13: Sensitivity analysis of batch size for Minibatch-OT on the Cite-5D dataset.

Batch Size	$t = 1$		$t = 2$		$t = 3$		Average	
	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$	$\mathcal{W}_1$	$\mathcal{W}_2$
1000	0.4344	0.4798	0.6522	0.7183	0.7658	0.8479	0.6175	0.6820
2000	0.4196	0.4624	0.5640	0.6251	0.7972	0.8820	0.5936	0.6565
Full Dataset	0.2868	0.3185	0.5487	0.5997	0.7689	0.8499	0.5348	0.5894

D Pseudocode for Algorithm

The pseudocode for the training process of Tracing Flow is provided in Algorithm 1.

Algorithm 1 Tracing-Flow Matching Training

Require: Snapshot datasets $\mathcal{D}_k = \{\mathbf{x}_k^{(i)}\}_{i=1}^{N_k}$ at times $t_0 < t_1 < \dots < t_K$;
Acceleration Network $\mathbf{a}_\theta$; Velocity Network $\mathbf{v}_\xi$;
Hyperparameters: batch size B_{traj} , learning rates η_a, η_v .

- 1: **Initialize:** Estimated velocity sets $\mathcal{V}_k = \{\mathbf{v}_k^{(i)}\}_{i=1}^{N_k} \leftarrow \{\mathbf{0}\}$ for all k .
- 2: **while** Velocity Estimates Not Converged **do**
- 3: **Stage 1: Velocity Assignment (Iterative SOAT)**
- 4: Initialize velocity accumulators $\hat{\mathcal{V}}_k^{(i)} \leftarrow \emptyset$ for all k, i .
- 5: *// Step 1.1: Solve Optimal Transport*
- 6: **for** $k = 0$ to $K - 1$ **do**
- 7: Compute cost matrix $\mathcal{C}_{i \rightarrow i+1}$ with / w.o. biological prior.
- 8: Compute optimal coupling $\pi_{k \rightarrow k+1}^*$ between $(\mathbf{x}_k, \mathbf{v}_k)$ and $(\mathbf{x}_{k+1}, \mathbf{v}_{k+1})$ by solving the SOAT Problem.
- 9: Compute propagator (transition) matrix $\mathcal{K}_{k \rightarrow k+1}^*$ by normalizing $\pi_{k \rightarrow k+1}^*$:
- 10:
$$\mathcal{K}_{k \rightarrow k+1}^*(i, j) = \frac{\pi_{k \rightarrow k+1}^*(i, j)}{\sum_{j'} \pi_{k \rightarrow k+1}^*(i, j')}$$
- 11: **end for**
- 12: *// Step 1.2: Trajectory Sampling & Velocity Update*
- 13: Sample B_{traj} discrete trajectories $T_m = \{(\mathbf{x}_k^{(s_{m,k})}, \mathbf{v}_k^{(s_{m,k})})\}_{k=0}^K$ for $m = 1 \dots B_{\text{traj}}$.
- 14: · Initial indices $s_{m,0}$ are sampled uniformly.
- 15: · Subsequent indices $s_{m,k+1} \sim \mathcal{K}_{k \rightarrow k+1}^*(s_{m,k}, \cdot)$.
- 16: **for** $m = 1$ to B_{traj} **do**
- 17: Compute optimal path velocities $\{\hat{\mathbf{v}}_k^{(s_{m,k})}\}_{k=0}^K$ for trajectory T_m using Proposition 4.8.
- 18: Accumulate: $\hat{\mathcal{V}}_k^{(s_{m,k})} \leftarrow \hat{\mathcal{V}}_k^{(s_{m,k})} \cup \{\hat{\mathbf{v}}_k^{(s_{m,k})}\}$ for $k = 0 \dots K$.
- 19: **end for**
- 20: *// Step 1.3: Average Accumulated Velocities*
- 21: **for** $k = 0$ to K , $i = 1$ to N_k **do**
- 22: **if** $\hat{\mathcal{V}}_k^{(i)} \neq \emptyset$ **then**
- 23:
$$\mathbf{v}_k^{(i)} \leftarrow \frac{1}{|\hat{\mathcal{V}}_k^{(i)}|} \sum_{\mathbf{v} \in \hat{\mathcal{V}}_k^{(i)}} \mathbf{v}$$
- 24: **end if**
- 25: **end for**
- 26: **end while**
- 27: **Stage 2: Final Transport Plan Computation**
- 28: Compute optimal couplings π^* and propagators $\mathcal{K}^*$ using the converged velocities $\mathcal{V}_k$.
- 29: **Stage 3: Neural Network Training**
- 30: **while** Not Converged **do**
- 31: Sample random time $t \sim \mathcal{U}[t_0, t_K]$ and coupling $\mathbf{z} \sim q(\mathbf{z})$.
- 32: Compute Acceleration Matching Loss:
- 33:
$$\mathcal{L}_{\text{CAFM}} = \mathbb{E}_{(\mathbf{x}, \mathbf{v}) \sim \rho_t(\cdot|z), \mathbf{z} \sim q(\mathbf{z}), t \sim \mathcal{U}[t_0, t_K]} [\|\mathbf{a}_\theta(\mathbf{x}, \mathbf{v}, t) - \mathbf{a}(\mathbf{x}, \mathbf{v}, t|z)\|^2]$$
- 34: Compute Initial Velocity Loss:
- 35:
$$\mathcal{L}_{v_0} = \mathbb{E}_{(\mathbf{x}, \mathbf{v}) \sim \rho_0} [\|\mathbf{v}_\xi(\mathbf{x}) - \mathbf{v}\|^2]$$
- 36: Update parameters (Gradient Descent):
- 37: $\theta \leftarrow \theta - \eta_a \nabla_\theta \mathcal{L}_{\text{CAFM}}$
- 38: $\xi \leftarrow \xi - \eta_v \nabla_\xi \mathcal{L}_{v_0}$
- 39: **end while**

E Discussion

E.1 Relation to other Algorithms

Relation to 3 MSBM 3 MSBM considers the following stochastic optimal control problem:

$$\min \mathbb{E}_{\rho_t} \int [\|\mathbf{a}_t\|^2] dt \quad (81)$$

$$d\mathbf{m}_t = A\mathbf{m}_t dt + \mathbf{u}_t dt + g d\mathbf{W}_t, \quad \mathbf{x}_n \sim q_n = \int \pi_n(\mathbf{x}, \mathbf{v}) d\mathbf{v}_n \quad (82)$$

where $\mathbf{m}_t = [\mathbf{x}_t, \mathbf{v}_t]^T$, $A = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$, and $g = \begin{bmatrix} 0 & 0 \\ 0 & \sigma \end{bmatrix}$. This problem can be viewed as a “stochastic” version of the VM-DOAT problem proposed in this paper. 3 MSBM also provides a simulation-free training method to regress the acceleration field. However, 3 MSBM does not explicitly solve this stochastic optimal control problem using the form of optimal coupling plus optimal single-particle trajectories. Its optimal coupling requires iterative solutions similar to Rectified Flow [36]: a process of repeatedly selecting the coupling and training a_θ . Furthermore, it relies on a heuristic method to estimate the initial velocity (via normal initialization and iteratively solving forward and backward SDEs), rather than incorporating this estimation as part of the optimal control problem. Experiments show that Tracing Flow achieves better distribution reconstruction accuracy than 3MSBM.

Relation to MMFM In MMFM, the single-particle path is similarly specified by:

$$\int \|\gamma''(t)\|^2 dt, \quad \mathbf{x}_k = \gamma(t_k) \quad (83)$$

which corresponds to a natural spline. However, MMFM does not solve an optimal control problem over distributions. Furthermore, the velocity field in MMFM remains single-valued with respect to $\mathbf{x}$, which prevents it from learning trajectories that cross in the position space.

Relation to OAT-FM OAT-FM introduces a DOAT problem formulation consistent with this paper; however, the objective of OAT-FM is not to genuinely solve the DOAT problem within the augmented space $\mathcal{X} \times \mathcal{V}$, but rather to fine-tune a pre-trained FM model. It employs the loss:

$$\begin{aligned} \mathcal{L}_{\text{CAFM}} = \mathbb{E}_{\pi^*[(x_0, v_0), (x_1, v_1)]} & \left[\alpha \left\| \frac{\mathbf{x}_t - \mathbf{x}_0}{t} - \frac{\mathbf{v}_0 + \mathbf{v}_\theta}{2} \right\|^2 + (1 - \alpha) \|\mathbf{v}_\theta(\mathbf{x}_t, t) - \mathbf{v}_0\|^2 \right. \\ & \left. + \alpha \left\| \frac{\mathbf{x}_1 - \mathbf{x}_t}{1 - t} - \frac{\mathbf{v}_\theta + \mathbf{v}_1}{2} \right\|^2 + (1 - \alpha) \|\mathbf{v}_1 - \mathbf{v}_\theta(\mathbf{x}_t, t)\|^2 \right] \end{aligned} \quad (84)$$

to enforce the velocity field along the line segment connecting any two samples to be as parallel to that line as possible. As a fine-tuning approach, it effectively still addresses a distribution transport problem involving only two time points within the position space $\mathcal{X}$, rather than tackling a multi-marginal optimal control problem. Furthermore, OAT-FM does not establish a connection between its loss function and the cost along each trajectory $\frac{1}{2} \int_0^1 \|\ddot{\gamma}\|^2 dt$, nor does it explicitly learn the acceleration field $\mathbf{a}(\mathbf{x}, \mathbf{v}, t)$.

E.2 Is $\mathcal{S} = \mathcal{X} \times \mathcal{V}$ a Phase Space?

In the Conclusion and Limitation section, we noted that if $\mathbf{x}$ is regarded as the generalized coordinate and $\mathbf{v}$ as its time derivative, the term $\int \frac{1}{2} \|\mathbf{a}\|^2 dt$ cannot be interpreted as a physical action, as the action in classical mechanics typically does not involve second-order time derivatives. Here, we offer an alternative physical interpretation of this control objective.

Proposition E.1. Consider a d -dimensional second-order dynamical system governed by the equations:

$$\dot{\mathbf{x}} = \mathbf{v} \quad (85)$$

$$\dot{\mathbf{v}} = \mathbf{a} \quad (86)$$

where $\mathbf{x}, \mathbf{v}, \mathbf{a} \in \mathbb{R}^d$. The optimal control objective is given by:

$$\min \frac{1}{2} \int_0^T \|\mathbf{a}\|^2 dt \quad (87)$$

The evolution of this system under the optimal control law is equivalent to a Hamiltonian system in $4d$ -dimensional phase space with the following Hamiltonian:

$$H = \mathbf{p}_1^T \mathbf{q}_2 + \frac{1}{2} \|\mathbf{p}_2\|^2 \quad (88)$$

where $\mathbf{q}_1, \mathbf{q}_2 \in \mathbb{R}^d$ are the canonical coordinates, and $\mathbf{p}_1, \mathbf{p}_2 \in \mathbb{R}^d$ are the canonical momenta (note that $\mathbf{q}_1$ does not appear in the Hamiltonian). Here, $\mathbf{q}_1$ and $\mathbf{q}_2$ correspond to $\mathbf{x}$ and $\mathbf{v}$ in the optimal control problem, respectively, while $\mathbf{p}_2$ corresponds to $\mathbf{a}$.

Proof. We derive the equations of motion directly using Hamilton's canonical equations:

$$\dot{\mathbf{q}}_1 = \frac{\partial H}{\partial \mathbf{p}_1} = \mathbf{q}_2, \quad \dot{\mathbf{q}}_2 = \frac{\partial H}{\partial \mathbf{p}_2} = \mathbf{p}_2, \quad \dot{\mathbf{p}}_1 = -\frac{\partial H}{\partial \mathbf{q}_1} = \mathbf{0}, \quad \dot{\mathbf{p}}_2 = -\frac{\partial H}{\partial \mathbf{q}_2} = -\mathbf{p}_1 \quad (89)$$

This implies that $\mathbf{p}_1$ is constant. Solving these equations sequentially yields:

$$\mathbf{p}_2(t) = \mathbf{c}'_2 - \mathbf{p}_1 t \quad (90)$$

$$\mathbf{q}_2(t) = \mathbf{c}'_1 + \mathbf{c}'_2 t - \frac{1}{2} \mathbf{p}_1 t^2 \quad (91)$$

$$\mathbf{q}_1(t) = \mathbf{c}'_0 + \mathbf{c}'_1 t + \frac{1}{2} \mathbf{c}'_2 t^2 - \frac{1}{6} \mathbf{p}_1 t^3 \quad (92)$$

Comparing this with the results in Equation (9), we observe that $\mathbf{q}_1(t)$ and $\mathbf{q}_2(t)$ satisfy the same equations as $\mathbf{x}(t)$ and $\mathbf{v}(t)$, with the undetermined constants determined by the boundary conditions. This completes the proof. $\square$

Consequently, we observe that although the augmented space $\mathcal{S} = \mathcal{X} \times \mathcal{V}$ concatenates position $\mathbf{x}$ and velocity (momentum) $\mathbf{v}$, it **cannot** be interpreted as a **phase space**. Instead, it should be viewed as the **configuration space** of the classical mechanical system described by the Hamiltonian above. We can attempt to recover the Lagrangian from this Hamiltonian. The canonical equations established that $\dot{\mathbf{q}}_1 = \mathbf{q}_2$ and $\dot{\mathbf{q}}_2 = \mathbf{p}_2$. Applying the Legendre transformation:

$$\begin{aligned} L &= \mathbf{p}_1^T \dot{\mathbf{q}}_1 + \mathbf{p}_2^T \dot{\mathbf{q}}_2 - H \\ &= \mathbf{p}_1^T (\dot{\mathbf{q}}_1 - \mathbf{q}_2) + \frac{1}{2} \|\dot{\mathbf{q}}_2\|^2 \end{aligned} \quad (93)$$

We find that the term $\mathbf{p}_1^T$ cannot be eliminated; the Lagrangian cannot be expressed solely as a function of the canonical coordinates and their first derivatives. This is a characteristic of constrained systems (in Dirac's theory of constraints, a constrained system is typically defined as one where the Hessian matrix of the Lagrangian is singular. To see this intuitively, if we regard $\mathbf{p}_1$ in the Lagrangian as d new canonical coordinates, their corresponding canonical momenta are zero, implying the motion is confined to a submanifold of the phase space).

In Equation (10), we derived the minimum cost for the system to transition from $(\mathbf{x}_1, \mathbf{v}_1)$ to $(\mathbf{x}_2, \mathbf{v}_2)$. While this cost is a symmetric positive-definite quadratic form with respect to $\mathbf{x}$

and $\mathbf{v}$, it does not constitute a metric on the augmented space $\mathcal{S} = \mathcal{X} \times \mathcal{V}$. (If it were a metric, the augmented space would be flat, and the optimal control trajectories—geodesics—would be linear functions of t). Maupertuis’ principle suggests a close relationship between the metric in the configuration space and the system’s Lagrangian [3]: if the Lagrangian takes the form:

$$L(\mathbf{q}, \dot{\mathbf{q}}) = \frac{1}{2} m_{ij}(\mathbf{q}) \dot{q}^i \dot{q}^j - V(\mathbf{q}) \quad (94)$$

then the configuration space possesses the metric:

$$g_{ij}(\mathbf{q}) = 2(E - V(\mathbf{q})) m_{ij}(\mathbf{q}) \quad (95)$$

where E is the total energy of the particle. For the optimal control problem discussed herein, we have seen that the Lagrangian cannot be written in such a form; therefore, it is fundamentally impossible to equip the configuration space $\mathcal{S} = \mathcal{X} \times \mathcal{V}$ with such a metric.

E.3 Broader Impacts

This paper presents work whose goal is to advance the field of machine learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here. However, as our algorithm is applied to single-cell data analysis, the fidelity of the generated trajectories is heavily dependent on the quality of the input data. Consequently, when utilizing our method for biological or medical research purposes, it is critical to employ high-quality datasets, incorporate established and accurate biological priors, and ensure that the algorithmic outputs are rigorously validated by domain experts.