
Variational Regularized Unbalanced Optimal Transport: Single Network, Least Action

Yuhao Sun^{1,*}, Zhenyi Zhang^{2,*}, Zihan Wang^{3,*}, Tiejun Li^{1,2,4,†} and Peijie Zhou^{1,3,4,5,†}

¹Center for Machine Learning Research, Peking University.

²LMAM and School of Mathematical Sciences, Peking University.

³Center for Quantitative Biology, Peking University.

⁴NELBDA, Peking University. ⁵AI for Science Institute, Beijing.

Emails: {2501111524, zhenyizhang, jackwzh}@stu.pku.edu.cn,
{tieli, pjzhou}@pku.edu.cn

Abstract

Recovering the dynamics from a few snapshots of a high-dimensional system is a challenging task in statistical physics and machine learning, with important applications in computational biology. Many algorithms have been developed to tackle this problem, based on frameworks such as optimal transport and the Schrödinger bridge. A notable recent framework is Regularized Unbalanced Optimal Transport (RUOT), which integrates both stochastic dynamics and unnormalized distributions. However, since many existing methods do not explicitly enforce optimality conditions, their solutions often struggle to satisfy the principle of least action and meet challenges to converge in a stable and reliable way. To address these issues, we propose Variational RUOT (Var-RUOT), a new framework to solve the RUOT problem. By incorporating the optimal necessary conditions for the RUOT problem into both the parameterization of the search space and the loss function design, Var-RUOT only needs to learn a scalar field to solve the RUOT problem and can search for solutions with lower action. We also examined the challenge of selecting a growth penalty function in the widely used Wasserstein-Fisher-Rao metric and proposed a solution that better aligns with biological priors in Var-RUOT. We validated the effectiveness of Var-RUOT on both simulated data and real single-cell datasets. Compared with existing algorithms, Var-RUOT can find solutions with lower action while exhibiting faster convergence and improved training stability. Our code is available at <https://github.com/ZeroVector/VarRUOT>.

1 Introduction

Inferring continuous dynamics from finite observations is crucial when analyzing systems with many particles (Chen et al. 2018). However, in many important applications such as single-cell RNA sequencing (scRNA-seq) experiments, only a few snapshot measurements are available, which makes recovering the underlying continuous dynamics a challenging task (Ding et al. 2022). Such a task of reconstructing dynamics from sparse snapshots is commonly referred to as *trajectory inference* in time-series scRNA-seq modeling (Zhang et al. 2025b; Ding et al. 2022; Heitz et al. 2024; Yeo et al. 2021a; Schiebinger et al. 2019a; Bunne et al. 2023b; Zhang et al. 2021) or the mathematical problem of *ensemble regression* (Yang et al. 2022).

A number of frameworks have been proposed to address this problem. For example, in dynamical optimal transport (OT), particles evolve according to the ordinary differential equations (ODEs) with

*These authors contributed equally. †Corresponding authors.

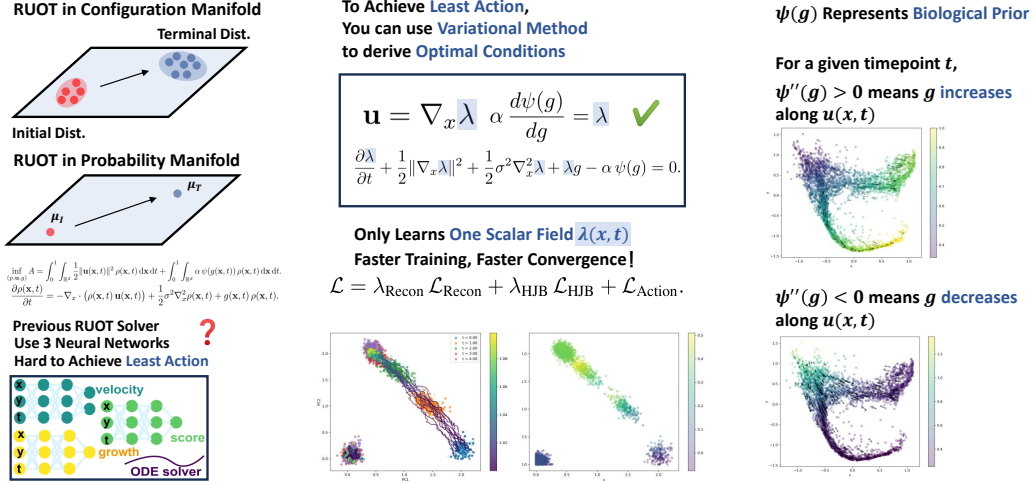


Figure 1: Overview of Variational RUOT (Var-RUOT).

the objective of minimizing the total action required to transport the initial distribution to the terminal distribution (Benamou and Brenier 2000). Unbalanced dynamical OT further extends this framework by adding a penalty term $\psi(g)$ on the particle growth or death processes in total transport energy (namely the Wasserstein–Fisher–Rao metric or WFR metric) to handle unnormalized distributions (Chizat et al. 2018a; Chizat et al. 2018b). Moreover, stochastic methods such as the Schrödinger Bridge adopt similar action principles while governing particle evolution via stochastic differential equations (SDEs) (Gentil et al. 2017; Léonard 2014). Recently, the Regularized Unbalanced Optimal Transport (RUOT) framework generalizes these ideas by incorporating both stochasticity and particle birth-death processes (Lavenant et al. 2024; Ventre et al. 2023; Chizat et al. 2022; Pariset et al. 2023; Zhang et al. 2025a). In machine learning, generative models such as diffusion models (Ho et al. 2020; Song et al. 2021; Sohl-Dickstein et al. 2015; Song et al. 2020) and flow matching techniques (Lipman et al. 2023; Tong et al. 2024a; Liu et al. 2022) have also been adapted to solve transport problems. However, these approaches face two major challenges: 1) They usually do not explicitly enforce optimality conditions, leading to solutions that violate the principle of least action, and they meet challenges to converge reliably; 2) Selecting an appropriate penalty function $\psi(g)$ that aligns with underlying biological priors remains challenging.

To overcome these challenges, we propose **Variational-RUOT (Var-RUOT)**. Our algorithm employs variational methods to derive the necessary conditions for action minimization within the RUOT framework. By parameterizing a single scalar field with a neural network and incorporating these optimality conditions directly into our loss design, Var-RUOT learns dynamics with lower action. Experiments on both simulated and real datasets demonstrate that our approach achieves competitive performance with fewer training epochs and improved stability. Furthermore, we show that different choices of the penalty function for the growth rate g yield distinct biologically relevant priors in single-cell dynamics modeling. Our contributions are summarized as follows:

- We introduce a new method for solving RUOT problems by incorporating the first-order optimality conditions directly into the solution parameterization. This reduces the learning task to a single scalar potential function, which significantly simplifies the model space.
- We show how incorporating these necessary conditions into the loss function and architecture enables Var-RUOT to consistently discover transport paths with lower action, providing a more efficient and stable training process for RUOT problem.
- We address a key limitation in the classical Wasserstein-Fisher-Rao metric, which can yield biologically implausible solutions due to its quadratic growth penalty term. We propose the criterion and practical solution to modify such a penalty term, therefore enhancing the more realistic modeling of single-cell dynamics.

2 Related Works

Deep Learning Solver for Trajectory Inference Problem There are a large number of deep learning-based trajectory inference problem solvers. For example, there are solvers for Optimal Transport using static OT solver, Neural ODE or Flow matching techniques (Tong et al. 2020; Huguët et al. 2022; Wan et al. 2023; Zhang et al. 2024a; Tong et al. 2024a; Albergo et al. 2023; Palma et al. 2025; Rohbeck et al. 2025; Petrović et al. 2025; Schiebinger et al. 2019b; Klein et al. 2025; Wang et al. 2025), as well as solvers for the Schrödinger bridge that utilize either its relation to entropy regularized OT or its optimal control formulation (Shi et al. 2024; De Bortoli et al. 2021; Gu et al. 2025; Koshizuka and Sato 2023; Neklyudov et al. 2024; Zhang et al. 2024b; Bunne et al. 2023a; Chen et al. 2022a; Zhou et al. 2024; Zhu et al. 2024; Maddu et al. 2024; Yeo et al. 2021b; Jiang and Wan 2024; Lavenant et al. 2024; Ventre et al. 2023; Chizat et al. 2022; Tong et al. 2024b; Atanackovic et al. 2025; Yang 2025; You et al. 2024). However, these methods typically **employ separate neural networks** to parameterize the velocity and growth functions, without leveraging their **optimality conditions** or the **inherent relationship** between them (Zhang et al. 2025d) with few exceptions such as Action Matching (Neklyudov et al. 2023; Neklyudov et al. 2024). This poses challenges in achieving optimal solutions that minimize the action energy. Our work generalizes Action Matching (Neklyudov et al. 2023; Neklyudov et al. 2024) to handle the setting where unbalanced distributions and stochastic effects coexist.

HJB equations in optimal transport Methods that leverage the optimality conditions (e.g., Hamilton-Jacobi-Bellman (HJB) equations) of dynamic OT and its variants, have been proposed (Neklyudov et al. 2024; Zhang et al. 2024b; Chen et al. 2016; Benamou and Brenier 2000; Neklyudov et al. 2023; Wu et al. 2025; Chow et al. 2020). However, these approaches typically do not simultaneously address both **unbalanced and stochastic** dynamics.

WFR metric in time-series scRNA-seq modeling In computational biology, several existing works model both cell state transitions and growth dynamics simultaneously in temporal scRNA-seq datasets by minimizing the action in the WFR metric i.e. solving the dynamical unbalanced optimal transport problem (Sha et al. 2024; Tong et al. 2023; Peng et al. 2024; Eyring et al. 2024b) or its variants (Pariset et al. 2023; Lavenant et al. 2024; Zhang et al. 2025a; Zhang et al. 2025c). However, these works usually adopt the default growth penalty function $\psi(g) = \frac{1}{2}g^2$ in the WFR metric and have not investigated the **biological implications of different choices** for $\psi(g)$.

3 Preliminaries and Backgrounds

Dynamical Optimal Transport The Dynamical Optimal Transport, also known as the Benamou–Brenier formulation, requires minimizing the following action functional (Benamou and Brenier 2000):

$$\inf_{\rho, \mathbf{u}} \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2} \|\mathbf{u}(\mathbf{x}, t)\|^2 \rho(\mathbf{x}, t) d\mathbf{x} dt,$$

where ρ and \mathbf{u} are subject to the continuity equation constraint:

$$\frac{\partial \rho(\mathbf{x}, t)}{\partial t} + \nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho(\mathbf{x}, t)) = 0, \quad \rho(\cdot, 0) = \mu_0, \quad \rho(\cdot, 1) = \mu_1.$$

Unbalanced Dynamical OT and Wasserstein–Fisher–Rao (WFR) metric In order to handle unnormalized probability densities in practical problems (for example, to account for cell proliferation and death in computational biology), one can modify the form of the continuity equation by adding a birth-death term, and accordingly include a corresponding penalty term in the action. This leads to the optimal transport problem under the Wasserstein–Fisher–Rao (WFR) metric (Chizat et al. 2018a; Chizat et al. 2018b).

$$\inf_{\rho, \mathbf{u}, g} \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}(\mathbf{x}, t)\|^2 + \alpha g^2(\mathbf{x}, t) \right) \rho(\mathbf{x}, t) d\mathbf{x} dt,$$

with ρ , \mathbf{u} , and g subject to the unnormalized continuity equation constraint

$$\frac{\partial \rho(\mathbf{x}, t)}{\partial t} + \nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho(\mathbf{x}, t)) - g(\mathbf{x}, t) \rho(\mathbf{x}, t) = 0, \quad \rho(\cdot, 0) = \mu_0, \quad \rho(\cdot, 1) = \mu_1.$$

Schrödinger Bridge Problem and Dynamical Formulation Schrödinger bridges aims to find the most likely way for a stochastic system to evolve from an initial distribution μ_0 to a terminal distribution μ_1 . Formally, let $\mu_{[0,1]}^{\mathbf{X}}$ denote the probability measure induced by the stochastic process $\mathbf{X}(t)$, $0 \leq t \leq 1$, and let $\mu_{[0,1]}^{\mathbf{Y}}$ denote the probability measure induced by a given reference process $\mathbf{Y}(t)$, $0 \leq t \leq 1$. The Schrödinger bridge seeks to solve $\min_{\mu_{[0,1]}^{\mathbf{X}}} \mathcal{D}_{\text{KL}}(\mu_{[0,1]}^{\mathbf{X}} \parallel \mu_{[0,1]}^{\mathbf{Y}})$. In particular, if \mathbf{X}_t follows the SDE $d\mathbf{X}_t = \mathbf{u}(\mathbf{X}_t, t) dt + \sigma(\mathbf{X}_t, t) d\mathbf{W}_t$, where $\mathbf{W}_t \in \mathbb{R}^d$ is a standard Brownian motion, $\sigma(\mathbf{x}, t) \in \mathbb{R}^{d \times d}$ is a given diffusion matrix, and the reference process is defined as $d\mathbf{Y}_t = \sigma(\mathbf{Y}_t, t) d\mathbf{W}_t$, then the Schrödinger bridge problem is equivalent to the following stochastic optimal control problem (Chen et al. 2016; Gentil et al. 2017):

$$\inf_{\rho, \mathbf{u}} \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \mathbf{u}^T(\mathbf{x}, t) \mathbf{a}^{-1}(\mathbf{x}, t) \mathbf{u}(\mathbf{x}, t) \right) \rho(\mathbf{x}, t) d\mathbf{x} dt,$$

where ρ and \mathbf{u} are subject to the Fokker–Planck equation constraint

$$\frac{\partial \rho(\mathbf{x}, t)}{\partial t} + \nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho(\mathbf{x}, t)) - \frac{1}{2} \nabla_{\mathbf{x}}^2 : (\mathbf{a}(\mathbf{x}, t) \rho(\mathbf{x}, t)) = 0, \quad \rho(\cdot, 0) = \mu_0, \quad \rho(\cdot, 1) = \mu_1.$$

Here, $\mathbf{a}(\mathbf{x}, t) = \sigma(\mathbf{x}, t) \sigma^T(\mathbf{x}, t)$ and $\nabla_{\mathbf{x}}^2 : (\mathbf{a}(\mathbf{x}, t) \rho(\mathbf{x}, t)) = \sum_{ij} \partial_{ij} (a_{ij} \rho(\mathbf{x}, t))$.

Regularized Unbalanced Optimal Transport If we consider both unnormalized probability densities and stochasticity simultaneously, we arrive at the Regularized Unbalanced Optimal Transport (RUOT) problem (Chen et al. 2022b; Baradat and Lavenant 2021; Zhang et al. 2025a).

Definition 3.1 (Regularized Unbalanced Optimal Transport (RUOT) Problem). *Consider minimizing the following action:*

$$\inf_{\rho, \mathbf{u}, g} \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \mathbf{u}^T(\mathbf{x}, t) \mathbf{a}^{-1}(\mathbf{x}, t) \mathbf{u}(\mathbf{x}, t) + \alpha \psi(g) \right) \rho(\mathbf{x}, t) d\mathbf{x} dt,$$

where $\psi : \mathbb{R} \rightarrow [0, +\infty)$ is a growth penalty function, α is a penalty coefficient, and the quantities ρ , \mathbf{u} and g are subject to the following constraint, which is an unnormalized continuity equation:

$$\frac{\partial \rho}{\partial t} + \nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho) - \frac{1}{2} \nabla_{\mathbf{x}}^2 : (\mathbf{a}(\mathbf{x}, t) \rho) - g(\mathbf{x}, t) \rho = 0, \quad \rho(\cdot, 0) = \mu_0, \quad \rho(\cdot, 1) = \mu_1.$$

Remark 3.1. When modeling the particle dynamics in these OT variants, one can use either autonomous systems (where v or g does not explicitly depend on t) or non-autonomous systems (where v or g explicitly depends on t). Here, we follow methods such as DeepRUOT (Zhang et al. 2025a) to choose non-autonomous systems. Indeed non-autonomous systems offer greater expressive power from a learning perspective and are less prone to underfitting the data.

4 Necessary Optimality Conditions for RUOT

To simplify our problem we adopt the assumption of isotropic time-invariant diffusion, i.e., $\mathbf{a}(\mathbf{x}, t) = \sigma^2 \mathbf{I}$. We refer to RUOT problem in this scenario as the *isotropic time-invariant RUOT problem*.

Definition 4.1 (Isotropic Time-Invariant (ITI) RUOT Problem). *Consider the following minimum-action problem with the action functional given by*

$$\inf_{(\rho, \mathbf{u}, g)} \mathcal{J} = \int_0^1 \int_{\mathbb{R}^d} \frac{1}{2} \|\mathbf{u}(\mathbf{x}, t)\|^2 \rho(\mathbf{x}, t) d\mathbf{x} dt + \int_0^1 \int_{\mathbb{R}^d} \alpha \psi(g(\mathbf{x}, t)) \rho(\mathbf{x}, t) d\mathbf{x} dt. \quad (1)$$

Here, $\psi : \mathbb{R} \rightarrow [0, +\infty)$ is the growth penalty function, and the triplet (ρ, \mathbf{u}, g) is subject to the constraint of the Fokker–Planck equation

$$\frac{\partial \rho(\mathbf{x}, t)}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\rho(\mathbf{x}, t) \mathbf{u}(\mathbf{x}, t)) + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \rho(\mathbf{x}, t) + g(\mathbf{x}, t) \rho(\mathbf{x}, t). \quad (2)$$

Additionally, ρ satisfies the initial and terminal conditions $\rho(\cdot, 0) = \mu_0$, $\rho(\cdot, 1) = \mu_1$.

In particular, if $\psi(g(\mathbf{x}, t)) = \frac{1}{2}g^2(\mathbf{x}, t)$, then this problem is referred to as the *unbalanced dynamic optimal transport with WFR metric*. We can derive the necessary conditions for the action functional to achieve a minimum using variational methods.

Theorem 4.1 (Necessary Conditions for Achieving the Optimal Solution in the ITI-RUOT Problem). *In the problem defined in Definition 4.1, the necessary conditions for the action \mathcal{T} to attain a minimum are*

$$\mathbf{u} = \nabla_{\mathbf{x}}\lambda, \quad \alpha \frac{d\psi(g)}{dg} = \lambda, \quad \frac{\partial\lambda}{\partial t} + \frac{1}{2}\|\nabla_{\mathbf{x}}\lambda\|^2 + \frac{1}{2}\sigma^2\nabla_{\mathbf{x}}^2\lambda + \lambda g - \alpha\psi(g) = 0. \quad (3)$$

Here, $\lambda(\mathbf{x}, t)$ is a scalar field. The proof of this theorem can be found in Appendix A.1.1.

Remark 4.1. *Substituting the necessary conditions satisfied by \mathbf{u} and g into the Fokker–Planck equation, the evolution of the probability density $\rho(\mathbf{x}, t)$ is determined by $\frac{\partial\rho(\mathbf{x}, t)}{\partial t} = -\nabla_{\mathbf{x}} \cdot \left(\rho(\mathbf{x}, t)\nabla_{\mathbf{x}}\lambda(\mathbf{x}, t) \right) + \frac{1}{2}\sigma^2\nabla_{\mathbf{x}}^2\rho(\mathbf{x}, t) + (\psi')^{-1}\left(\frac{\lambda(\mathbf{x}, t)}{\alpha}\right)\rho(\mathbf{x}, t)$, where $\psi' = \frac{d\psi(g)}{dg}$, and $(\psi')^{-1}$ denotes the inverse function of ψ' .*

Remark 4.2. *If we choose the growth penalty function to take the form used in the WFR metric, i.e., $\psi(g) = \frac{1}{2}g^2$, and set $\alpha = 1$, $\sigma = 0$, then the above optimal necessary conditions immediately degenerate to $\mathbf{u} = \nabla_{\mathbf{x}}\lambda$, $g = \lambda$, $\frac{\partial\lambda}{\partial t} + \frac{1}{2}\|\nabla_{\mathbf{x}}\lambda\|^2 + \frac{1}{2}\lambda^2 = 0$, which is same as the form derived in (Neklyudov et al. 2024) under the WFR metric. If we let $g = 0$ and $\psi(0) = 0$ it becomes $\mathbf{u} = \nabla_{\mathbf{x}}\lambda$, $\frac{\partial\lambda}{\partial t} + \frac{1}{2}\|\nabla_{\mathbf{x}}\lambda\|^2 + \frac{1}{2}\sigma^2\nabla_{\mathbf{x}}^2\lambda = 0$, which is same as the form derived in (Neklyudov et al. 2024; Zhang et al. 2024b; Chen et al. 2016) under the Schrödinger Bridge problem.*

From Theorem 4.1 and Remark 4.1, the vector field $\mathbf{u}(\mathbf{x}, t)$ and the growth rate $g(\mathbf{x}, t)$ can be directly obtained from the scalar field $\lambda(\mathbf{x}, t)$. Moreover, since the initial density $\rho(\cdot, 0)$ is known, once the necessary conditions are satisfied, the evolution equation (i.e., the Fokker–Planck equation) is completely determined by $\lambda(\mathbf{x}, t)$. Thus, the scalar field $\lambda(\mathbf{x}, t)$ fully determines the system's evolution, and we only need to solve for one $\lambda(\mathbf{x}, t)$, which simplifies the problem.

However, these necessary conditions introduce a coupling between $\mathbf{u}(\mathbf{x}, t)$ and $g(\mathbf{x}, t)$, and this coupling could contradict biological prior knowledge. In biological data, it is generally believed that cells located at the upstream of a trajectory are stem cells with the highest proliferation and differentiation capabilities, and thus the corresponding g values should be maximal. Along the trajectory, as the cells gradually lose their "stemness," the g values would decrease. Under the necessary conditions, however, whether $g(\mathbf{x}, t)$ increases or decreases along $\mathbf{u}(\mathbf{x}, t)$ at a given time t depends on the form of the growth penalty function.

Theorem 4.2 (The relationship between \mathbf{u} and g ; Biological prior). *At a fixed time t , if $\frac{d\psi^2(g)}{dg^2} > 0$, then $g(\mathbf{x}, t)$ ascends in the direction of the velocity field $\mathbf{u}(\mathbf{x}, t)$ (i.e., $\mathbf{u}(\mathbf{x}, t)^T(\nabla_{\mathbf{x}}g(\mathbf{x}, t)) > 0$); otherwise, it descends.*

The proof is given in Appendix A.1.2. According to this theorem, to ensure the solution complies with biological prior, i.e., that at a given time the cells upstream in the trajectory exhibit the higher g value, it is necessary to ensure that that $\frac{d^2\psi(g)}{dg^2} < 0$.

5 Solving ITI ROUT Problem Through Neural Network

Given samples from distributions ρ_t at K discrete time points, $t \in \{T_1, \dots, T_K\}$, we aim to recover the continuous evolution process of the distributions by solving the ITI RUOT problem, that is, by minimizing the action functional while ensuring that $\rho(\mathbf{x}, t)$ matches the distributions ρ_t at the corresponding time points. Since the values of $\mathbf{u}(\mathbf{x}, t)$ and $g(\mathbf{x}, t)$, as well as the evolution of $\rho(\mathbf{x}, t)$ over time, are fully determined by the scalar field $\lambda(\mathbf{x}, t)$ in variational form (Appendix A.1.1), we approximate this scalar field using a single neural network. Specifically, we parameterize $\lambda(\mathbf{x}, t)$ as $\lambda_{\theta}(\mathbf{x}, t)$, where θ represents the neural network parameters.

5.1 Simulating SDEs Using the Weighted Particle Method

Directly solving the high-dimensional RUOT with PDE constraints is challenging. Therefore, we reformulate the problem by simulating the trajectories of a number of weighted particles.

Theorem 5.1. Consider a weighted particle system consisting of N particles, where the position of particle i at time t is given by $\mathbf{X}_i^t \in \mathbb{R}^d$ and its weight by $w_i(t) > 0$. The dynamics of each particle are described by

$$\begin{aligned} d\mathbf{X}_i^t &= \mathbf{u}(\mathbf{X}_i^t, t) dt + \sigma(t) d\mathbf{W}_t, \\ dw_i &= g(\mathbf{X}_i^t, t) w_i dt, \end{aligned} \quad (4)$$

where $\mathbf{u} : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ is a time-varying vector field, $g : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}$ is a growth rate function, $\sigma : [0, T] \rightarrow [0, +\infty)$ is a time-varying diffusion coefficient, and \mathbf{W}_t is an N -dimensional standard Brownian motion with independent components in each coordinate. The initial conditions are $\mathbf{X}_i^0 \sim \rho(\mathbf{x}, 0)$ and $w_i(0) = 1$. In the limit as $N \rightarrow \infty$, the empirical measure $\mu^N(\mathbf{x}, t) = \frac{1}{N} \sum_{i=1}^N w_i(t) \delta(\mathbf{x} - \mathbf{X}_i^t)$ converges to the solution of the following Fokker–Planck equation:

$$\frac{\partial \rho(\mathbf{x}, t)}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho(\mathbf{x}, t)) + \frac{1}{2} \sigma^2(t) \nabla_{\mathbf{x}}^2 \rho(\mathbf{x}, t) + g(\mathbf{x}, t) \rho(\mathbf{x}, t), \quad (5)$$

with the initial condition $\rho(\mathbf{x}, 0) = \rho_0(\mathbf{x})$.

The proof is provided in Appendix A.1.3. This theorem implies that we can approximate the evolution of $\rho(\mathbf{x}, t)$ by simulating N particles, where each particle’s weight w_i is governed by an ODE and its position \mathbf{X}_i is governed by an SDE. The evolution of the empirical measure $\mu^N(\mathbf{x}, t)$ thereby approximates the evolution of $\rho(\mathbf{x}, t)$.

5.2 Reformulating the Loss in Weighted Particle Form

The total loss function consists of three components such that $\mathcal{L} = \mathcal{L}_{\text{Recon}} + \gamma_{\text{HJB}} \mathcal{L}_{\text{HJB}} + \gamma_{\text{Action}} \mathcal{L}_{\text{Action}}$. Here, $\mathcal{L}_{\text{Recon}}$ ensures that the distribution generated by the model closely matches the true data distribution, \mathcal{L}_{HJB} enforces that the learned $\lambda_\theta(\mathbf{x}, t)$ satisfies the HJB equation in the necessary conditions, and $\mathcal{L}_{\text{Action}}$ minimizes the action as much as possible.

Reconstruction Loss Minimizing the reconstruction loss guarantees that the distribution generated by the model is consistent with the real data distribution. Since in the ITI RUOT problem the probability density $\rho(\mathbf{x}, t)$ is not normalized, we need to match both the total mass and the discrepancy between the two distributions. Our reconstruction loss is given by $\mathcal{L}_{\text{Recon}} = \gamma_{\text{Mass}} \mathcal{L}_{\text{Mass}} + \mathcal{L}_{\text{OT}}$, where, at time point k , the true mass is $\int_{\mathbb{R}^d} \rho(\mathbf{x}, T_k) d\mathbf{x} = M(T_k)$, and the weight of particle i is $w_i(T_k)$. Thus, the total mass of the model-generated distribution is $\hat{M}(T_k) = \frac{1}{N} \sum_{i=1}^N w_i(T_k)$. The mass reconstruction loss is then defined as $\mathcal{L}_{\text{Mass}} = \sum_{k=1}^K \left(M(T_k) - \hat{M}(T_k) \right)^2$. Let the true distribution at time point k be $\rho(\mathbf{x}, T_k)$. Its normalized version is given by $\tilde{\rho}(\mathbf{x}, T_k) = \frac{\rho(\mathbf{x}, T_k)}{\int_{\mathbb{R}^d} \rho(\mathbf{x}, T_k) d\mathbf{x}}$, while the normalized model-generated distribution is $\hat{\rho}(\mathbf{x}, T_k) = \frac{\frac{1}{N} \sum_{i=1}^N w_i(T_k) \delta(\mathbf{x} - \mathbf{X}_i)}{\hat{M}(T_k)}$. The distribution reconstruction loss is then defined as $\mathcal{L}_{\text{OT}} = \sum_{k=1}^K \mathcal{W}_2(\tilde{\rho}(\cdot, T_k), \hat{\rho}(\cdot, T_k))$, where γ_{Mass} is a hyperparameter that controls the importance of the mass reconstruction loss.

HJB Loss Minimizing the HJB loss ensures that the learned $\lambda_\theta(\mathbf{x}, t)$ obeys the HJB equation constraints specified in the necessary conditions. Since the gradient operator in the HJB equation is a local operator, we compute the HJB loss by integrating the extent to which $\lambda_\theta(\mathbf{x}, t)$ violates the HJB equation along the trajectories. When using N particles, the HJB loss is given by: $\mathcal{L}_{\text{HJB}}^N = \sum_{i=1}^N \left[\int_0^{T_K} \frac{w_i(t)}{\sum_{i=1}^N w_i(t)} \left(\frac{\partial \lambda_\theta(\mathbf{X}_i^t, t)}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda_\theta\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda_\theta + g_\theta(\mathbf{X}_i^t, t) - \alpha \psi(g_\theta) \right)^2 dt \right]$. Here g_θ is obtained from the necessary condition $\alpha \frac{d\psi(g_\theta)}{dg_\theta} = \lambda_\theta$.

Remark 5.1. The expectation of HJB Loss is $\mathbb{E}[\mathcal{L}_{\text{HJB}}^N] = \int_0^{T_K} \int_{\mathbb{R}^d} \hat{\rho}(\mathbf{x}, t) \left(\frac{\partial \lambda(\mathbf{x}, t)}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + g(\mathbf{x}, t) - \alpha \psi(g) \right)^2 d\mathbf{x} dt$ where $\hat{\rho}(\mathbf{x}, t) = \frac{\rho(\mathbf{x}, t)}{\int_{\mathbb{R}^d} \rho(\mathbf{x}, t) d\mathbf{x}}$ is the probability density by normalizing $\rho(\mathbf{x}, t)$. The proof is left in Appendix A.1.4.

Action Loss Since the variational method provides only the necessary conditions for achieving minimal action rather than sufficient ones, we need to incorporate the action into the loss so that the action is minimized as much as possible. The action loss is also computed via simulating weighted particles. When using N particles, it is given by $\mathcal{L}_{\text{Action}}^N = \frac{1}{N} \sum_{i=1}^N \left(\int_0^1 \frac{1}{2} \|\mathbf{u}_\theta(\mathbf{X}_i^t, t)\|^2 w_i(t) dt + \int_0^1 \alpha \psi(g_\theta(\mathbf{X}_i^t, t)) w_i(t) dt \right)$. Here \mathbf{u} and g is obtained from the necessary condition $\mathbf{u}_\theta = \nabla_{\mathbf{x}} \lambda_\theta(\mathbf{x}, t)$, $\alpha \frac{d\psi(g_\theta)}{dg_\theta} = \lambda_\theta$.

Remark 5.2. *The expectation of action loss is exactly the action defined in the ITI RUOT problem (Definition 4.1): $\mathbb{E}[\mathcal{L}_{\text{Action}}^N] = \mathcal{T}$. The proof is left in Appendix A.1.5.*

Overall, the training process of Var-RUOT involves minimizing the total of three loss terms described above to fit λ_θ . The training procedure is provided in Algorithm 1.

5.3 Adjusting the Growth Penalty Function to Match Biological Priors

As discussed in Theorem 4.2, the second-order derivative of $\psi(g)$ encodes the biological prior: if $\psi''(g) > 0$, then at any given time t , g increases in the direction of the velocity field, and vice versa. Therefore, we choose two representative forms of $\psi(g)$ for our solution. Given that $\psi(g)$ penalizes nonzero g , it should satisfy the following properties: (1) The further g deviates from 0, the larger $\psi(g)$ becomes, i.e., $\frac{d\psi(g)}{d|g|} > 0$. (2) Birth and death are penalized equally when prior knowledge is absent, i.e., $\psi(g) = \psi(-g)$.

Case 1: $\psi''(g) > 0$ In the case where $\psi''(g) > 0$, a typical form that meets the requirements is $\psi(g) = Cg^{2p}$, $p \in \mathbb{Z}^+$, $C > 0$. We select the form used in the WFR Metric, namely, $\psi_1(g) = \frac{1}{2}g^2$. The optimality conditions are presented in Appendix A.2.

Case 2: $\psi''(g) < 0$ For the case where $\psi''(g) < 0$, a typical form that meets the conditions is $\psi(g) = Cg^{(2p)/(2q+1)}$, where $p, q \in \mathbb{Z}^+$ and $2p < 2q + 1$. In order to obtain a smoother $g(\lambda)$ relationship from the necessary conditions, and as a illustrative example, we choose $\psi_2(g) = g^{2/15}$. The optimality conditions are also presented in Appendix A.2.

6 Numerical Results

In the experiments presented below, unless the use of the modified metric is explicitly stated, we utilize the standard WFR metric, namely, $\psi_1(g) = \frac{1}{2}g^2$.

6.1 Var-RUOT Minimizes Path Action

To evaluate the ability of Var-RUOT to capture the minimum-action trajectory, we first conducted experiments on a three-gene simulation dataset (Zhang et al. 2025a). The dynamics of the three-gene simulation data are governed by stochastic differential equations that incorporate self-activation, mutual inhibition, and external activation. The detailed specifications of the dataset are provided in Appendix B.1. The trajectories learned by DeepRUOT and Var-RUOT are illustrated in Fig. 2, and the \mathcal{W}_1 and \mathcal{W}_2 losses between the generated distributions and the ground truth distribution, as well as the corresponding action values, are reported in Table 1. In the table, we report the action of the method that utilizes the WFR Metric. The experimental results demonstrate that Var-RUOT achieves a lower path action while maintaining distribution matching accuracy. To further assess the performance of Var-RUOT on high-dimensional data, we also conducted experiments on an Epithelial Mesenchymal Transition (EMT) dataset (Sha et al. 2024; Cook and Vanderhyden 2020). This dataset was reduced to a 10-dimensional feature space, and the trajectories obtained after applying PCA for dimensionality reduction are shown in Fig. 3. Both Var-RUOT and DeepRUOT learn dynamics that can transform the distribution at $t = 0$ into the distributions at $t = 1, 2, 3$. Var-RUOT learns the nearly straight-line trajectory corresponding to the minimum action, whereas DeepRUOT learns a curved trajectory. The results of \mathcal{W}_1 , \mathcal{W}_2 distance and action, summarized in Table 2, Var-RUOT also learns trajectories with smaller action while achieving matching accuracy comparable to that of other algorithms.

Table 1: On the three-gene simulated dataset, the Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between the predicted distributions of each algorithm and the true distribution at various time points. Each experiment was run five times to compute the mean and standard deviation. (\mathcal{W}_1 and \mathcal{W}_2 are relative metrics and should therefore only be compared on the same dataset.)

Model	$t = 1$		$t = 2$		$t = 3$		$t = 4$		Path Action
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	
SF2M (Tong et al. 2024b)	0.1914 \pm 0.0051	0.3253 \pm 0.0059	0.4706 \pm 0.0200	0.7648 \pm 0.0059	0.7648 \pm 0.0200	1.0750 \pm 0.0267	2.1879 \pm 0.0451	2.8830 \pm 0.0741	—
PISDE (Jiang and Wan 2024)	0.1313 \pm 0.0023	0.3232 \pm 0.0013	0.2311 \pm 0.0015	0.5356 \pm 0.0015	0.4103 \pm 0.0006	0.7913 \pm 0.0035	0.5418 \pm 0.0015	0.9579 \pm 0.0037	—
MIO Flow (Huguet et al. 2022)	0.1290 \pm 0.0000	0.2087 \pm 0.0000	0.2963 \pm 0.0000	0.4565 \pm 0.0000	0.6461 \pm 0.0000	1.0165 \pm 0.0000	1.1473 \pm 0.0000	1.7827 \pm 0.0000	—
Action Matching (Neklyudov et al. 2023)	0.3801 \pm 0.0000	0.5033 \pm 0.0000	0.5028 \pm 0.0000	0.5637 \pm 0.0000	0.6288 \pm 0.0000	0.6822 \pm 0.0000	0.8480 \pm 0.0000	0.9034 \pm 0.0000	1.5491
OTCFM (Tong et al. 2024a)	0.1035 \pm 0.0000	0.3043 \pm 0.0000	0.2078 \pm 0.0000	0.4923 \pm 0.0000	0.2898 \pm 0.0000	0.5867 \pm 0.0000	0.3107 \pm 0.0000	0.4358 \pm 0.0000	—
UOTCFM (Eyring et al. 2024a)	0.1002 \pm 0.0000	0.2911 \pm 0.0000	0.1653 \pm 0.0000	0.3578 \pm 0.0000	0.1711 \pm 0.0000	0.2620 \pm 0.0000	0.4129 \pm 0.0000	0.7583 \pm 0.0000	—
WLF (Neklyudov et al. 2024)	0.4983 \pm 0.0000	0.5273 \pm 0.0000	0.8346 \pm 0.0000	0.8357 \pm 0.0000	0.8046 \pm 0.0000	0.8815 \pm 0.0000	0.4493 \pm 0.0000	0.8571 \pm 0.0000	—
TIGON (Sha et al. 2024)	0.0519 \pm 0.0000	0.0731 \pm 0.0000	0.0763 \pm 0.0000	0.1559 \pm 0.0000	0.1387 \pm 0.0000	0.2436 \pm 0.0000	0.1908 \pm 0.0000	0.2203 \pm 0.0000	1.2442
DeepRUOT (Zhang et al. 2025a)	0.0569 \pm 0.0019	0.1125 \pm 0.0033	0.0811 \pm 0.0037	0.1578 \pm 0.0079	0.1246 \pm 0.0040	0.2158 \pm 0.0081	0.1538 \pm 0.0056	0.2588 \pm 0.0088	1.4058
Var-RUOT (Ours)	0.0452 \pm 0.0024	0.1181 \pm 0.0064	0.0385 \pm 0.0022	0.1270 \pm 0.0121	0.0445 \pm 0.0033	0.1144 \pm 0.0160	0.0572 \pm 0.0034	0.2140 \pm 0.0067	1.1105 \pm 0.0515

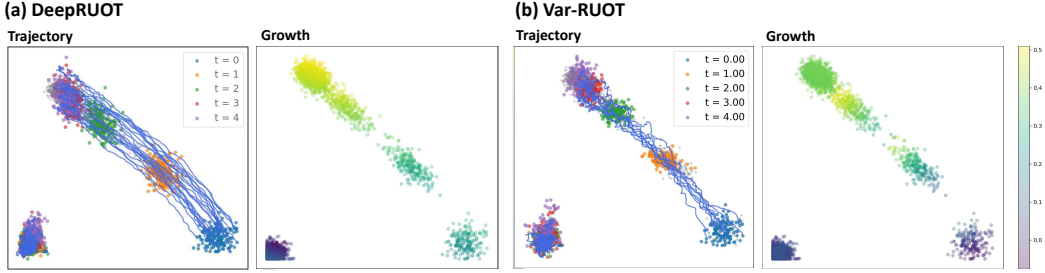


Figure 2: Comparison between DeepRUOT and Var-RUOT on dynamics reconstruction on three-gene simulation dataset. a) The trajectory and growth rate learned by DeepRUOT. b) The trajectory and growth rate learned by Var-RUOT on the same dataset. In the figure, the circles and '+' signs denote the generated data and the original data points, respectively.

Table 2: On the EMT dataset, the Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between the predicted distributions of each algorithm and the true distribution at various time points. Each experiment was run five times to compute the mean and standard deviation.

Model	$t = 1$		$t = 2$		$t = 3$		Path Action
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	
SF2M (Tong et al. 2024b)	0.2566 \pm 0.0016	0.2646 \pm 0.0016	0.2811 \pm 0.0016	0.2897 \pm 0.0012	0.2900 \pm 0.0010	0.3005 \pm 0.0010	—
PISDE (Jiang and Wan 2024)	0.2694 \pm 0.0016	0.2785 \pm 0.0016	0.2860 \pm 0.0013	0.2954 \pm 0.0012	0.2790 \pm 0.0015	0.2920 \pm 0.0016	—
MIO Flow (Huguet et al. 2022)	0.2439 \pm 0.0000	0.2529 \pm 0.0000	0.2665 \pm 0.0000	0.2770 \pm 0.0000	0.2841 \pm 0.0000	0.2984 \pm 0.0000	—
Action Matching (Neklyudov et al. 2023)	0.4723 \pm 0.0000	0.4794 \pm 0.0000	0.6382 \pm 0.0000	0.6454 \pm 0.0000	0.8453 \pm 0.0000	0.8524 \pm 0.0000	0.8583
OTCFM (Tong et al. 2024a)	0.2557 \pm 0.0000	0.2649 \pm 0.0000	0.2701 \pm 0.0000	0.2799 \pm 0.0000	0.2799 \pm 0.0000	0.2914 \pm 0.0000	—
UOTCFM (Eyring et al. 2024a)	0.2538 \pm 0.0000	0.2629 \pm 0.0000	0.2696 \pm 0.0000	0.2797 \pm 0.0000	0.2771 \pm 0.0000	0.2912 \pm 0.0000	—
WLF (Neklyudov et al. 2024)	0.3901 \pm 0.0000	0.3977 \pm 0.0000	0.4381 \pm 0.0000	0.4466 \pm 0.0000	0.2848 \pm 0.0000	0.2976 \pm 0.0000	—
TIGON (Sha et al. 2024)	0.2433 \pm 0.0000	0.2523 \pm 0.0000	0.2661 \pm 0.0000	0.2766 \pm 0.0000	0.2847 \pm 0.0000	0.2989 \pm 0.0000	0.4672
DeepRUOT (Zhang et al. 2025a)	0.2902 \pm 0.0009	0.2987 \pm 0.0012	0.3193 \pm 0.0006	0.3293 \pm 0.0008	0.3291 \pm 0.0008	0.3410 \pm 0.0023	0.4857
Var-RUOT(Ours)	0.2540 \pm 0.0016	0.2623 \pm 0.0017	0.2670 \pm 0.0013	0.2756 \pm 0.0014	0.2683 \pm 0.0014	0.2796 \pm 0.0015	0.3544 \pm 0.0019

6.2 Var-RUOT Stabilizes and Accelerates Training Process

To demonstrate that Var-RUOT converges faster and exhibits improved training stability, we further tested it on both the simulated and the EMT dataset. We trained the neural networks for the various algorithms using the same learning rate and optimizer, running each dataset five times. For each training, we recorded the number of epochs and wall-clock time required for the OT loss related to the distribution matching accuracy to decrease below a specified threshold (set to 0.30 in this study). Each training session was capped at a maximum of 500 epochs. If an algorithm's OT loss did not reach the threshold within 500 epochs, the required epoch was recorded as 500, and the wall-clock time was noted as the total duration of the training session. The experimental results are summarized in Table 3, which lists the mean and standard deviation of both the epochs and wall-clock times required for each algorithm on each dataset. The mean values reflect the convergence speed, while the standard deviations indicate the training stability. Our algorithm demonstrated both a faster convergence speed and better stability compared to the other methods. In Appendix C.1, we further demonstrate our training speed and stability by plotting the loss decay curves.

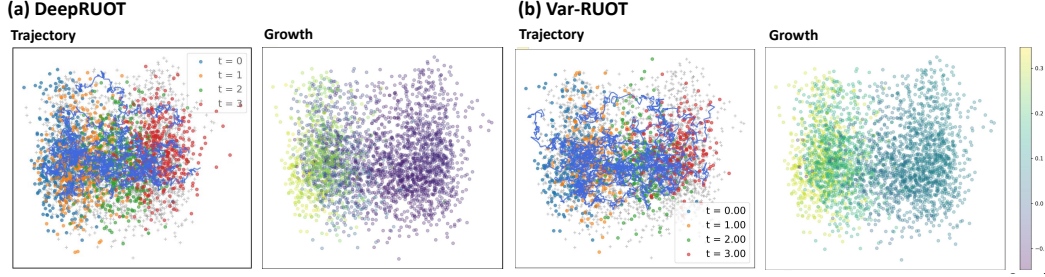


Figure 3: **Comparison between DeepRUOT and Var-RUOT on dynamics reconstruction on EMT dataset.** a) The trajectory and growth rate learned by DeepRUOT. b) The trajectory and growth rate learned by Var-RUOT on the same dataset. In the figure, the circles and '+' signs denote the generated data and the original data points, respectively.

Table 3: The number of epochs and wall time required for the OT Loss to drop below the threshold for each algorithm. We trained each algorithm five times to compute the mean and standard deviation.

Model	Simulation Gene		EMT	
	Epoch	Wall Time(Sec.)	Epoch	Wall Time(Sec.)
TIGON (Sha et al. 2024)	228.40 \pm 223.71	1142.79 \pm 1345.21	110.40 \pm 193.37	365.54 \pm 639.86
RUOT w/o Pretraining (Zhang et al. 2025a)	172.00 \pm 229.11	578.67 \pm 768.52	228.20 \pm 223.88	819.31 \pm 804.05
RUOT with 3 Epoch Pretraining (Zhang et al. 2025a)	204.40 \pm 238.29	653.33 \pm 761.35	221.60 \pm 226.52	801.18 \pm 819.46
Var-RUOT (Ours)	27.60 \pm 5.75	33.98 \pm 6.37	5.20 \pm 1.26	7.37 \pm 1.89

6.3 Different $\psi(g)$ Represents Different Biological Prior

To illustrate that the choice of $\psi(g)$ represents different biological priors, we present the learned dynamics under two selections of $\psi(g)$. We apply our algorithm on the Mouse Blood Hematopoiesis dataset (Weinreb et al. 2020; Sha et al. 2024). In Fig. 4(a), the standard WFR metric is applied, i.e., $\psi_1(g) = \frac{1}{2}g^2$, from which it can be observed that at time points $t = 0, 1, 2$, along the direction of the drift vector field $\mathbf{u}(\mathbf{x}, t)$, $g(\mathbf{x}, t)$ gradually increases. In Fig. 4(b), on the other hand, the alternative selection $\psi_2(g) = g^{2/15}$ mentioned in Section 5.3 is used, and it is evident that at each time point, $g(\mathbf{x}, t)$ gradually decreases along the direction of $\mathbf{u}(\mathbf{x}, t)$. The distribution matching accuracy and the action are reported in Table 4. When employing the modified metric, the corresponding action quantity is not directly comparable to those obtained using the WFR metric. Therefore we do not report its action here.

Table 4: Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between the predicted distributions of each algorithm and the true distribution at various time points on mouse blood hematopoiesis. Each experiment was run five times to compute the mean and standard deviation.

Model	$t = 1$		$t = 2$		Path Action
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	
Action Matching (Neklyudov et al. 2023)	0.4719 \pm 0.0000	0.5673 \pm 0.0000	0.8350 \pm 0.0000	0.8936 \pm 0.0000	4.3517
MIOFLOW (Huguet et al. 2022)	0.3546 \pm 0.0000	0.4083 \pm 0.0000	0.2772 \pm 0.0000	0.3459 \pm 0.0000	-
SF2M (Tong et al. 2024b)	0.1706 \pm 0.0043	0.2150 \pm 0.0062	0.1602 \pm 0.0029	0.2013 \pm 0.0039	-
PISDE (Jiang and Wan 2024)	0.3124 \pm 0.0065	0.3499 \pm 0.0062	0.2531 \pm 0.0142	0.2983 \pm 0.0175	-
OTCFM (Tong et al. 2024a)	0.3674 \pm 0.0000	0.4814 \pm 0.0000	0.3222 \pm 0.0000	0.3737 \pm 0.0000	-
UOTCFM (Eyring et al. 2024a)	0.3301 \pm 0.0000	0.3950 \pm 0.0000	0.2051 \pm 0.0000	0.2606 \pm 0.0000	-
WLF (Neklyudov et al. 2024)	0.3302 \pm 0.0000	0.3950 \pm 0.0000	0.2051 \pm 0.0000	0.2606 \pm 0.0000	-
TIGON (Sha et al. 2024)	0.4498 \pm 0.0000	0.5139 \pm 0.0000	0.4368 \pm 0.0000	0.4852 \pm 0.0000	3.7438
DeepRUOT (Zhang et al. 2025a)	0.1456 \pm 0.0016	0.1807 \pm 0.0019	0.1469 \pm 0.0046	0.1791 \pm 0.0061	5.5887
Var-RUOT (Standard WFR)	0.1200 \pm 0.0038	0.1459 \pm 0.0038	0.1431 \pm 0.0092	0.1764 \pm 0.0135	3.1491 \pm 0.0837
Var-RUOT (Modified Metric)	0.2953 \pm 0.0357	0.3117 \pm 0.0323	0.1917 \pm 0.0140	0.2226 \pm 0.0170	-

In addition to the three experiments presented in the main text, we conducted several ablation studies and supplementary experiments to further validate our method. First, we performed ablation studies on the weights of the HJB and action losses to verify their effectiveness in learning dynamics with a minimal action. Additionally, we explored alternative modeling choices, including a comparison

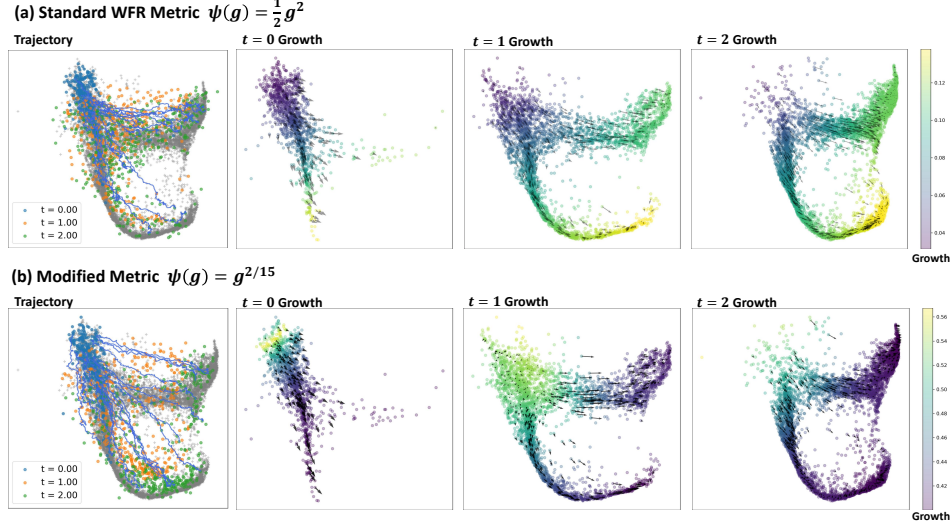


Figure 4: **Comparison of Var-RUOT using different growth metric on mouse blood hematopoiesis dataset.** a) The trajectory and growth at time points $t = 0, 1, 2$ learned using the standard WFR metric. b) The trajectory and growth at time points $t = 0, 1, 2$ learned using the modified metric. In the figure, the circles and '+' signs denote the generated data and the original data points, respectively.

between using L_1 and L_2 norms for the loss terms, and an analysis of the impact of different growth penalty functions beyond the $\psi_2(g) = g^{\frac{2}{15}}$ example where $\psi''(g) < 0$. The detailed results of these studies are provided in Appendix C.2. To assess the model's generalization capabilities, we performed hold-one-out and long-term extrapolation experiments. The results indicate that the Var-RUOT algorithm can effectively perform both interpolation and extrapolation, with the learned minimum-action dynamics leading to more accurate extrapolation outcomes (Appendix C.3). Furthermore, we carried out experiments on several high-dimensional datasets, which further validate the effectiveness of Var-RUOT in high-dimensional settings (Appendix C.4). Finally, we compared the path action computed by Var-RUOT against that from a traditional OT solver (Appendix C.5). The experimental results show that the path actions obtained by both methods are similar, which helps to validate our approach.

7 Conclusion

In this paper, we propose a new algorithm for solving the RUOT problem called Variational RUOT. By employing variational methods to derive the necessary conditions for the minimum action solution of RUOT, we solve the problem by learning a single scalar field. Compared to other algorithms, Var-RUOT can find solutions with lower action values while achieving the same level of fitting performance, and it offers faster training and convergence speeds. Finally, we emphasize that the selection of $\psi(g)$ in the action is crucial and directly linked to biological priors.

Although the Var-RUOT algorithm presented in this paper offers new insights for solving the RUOT problem, it is subject to several key limitations. Firstly, because the algorithm is based on neural network optimization, it can only converge to a local minimum rather than guaranteeing the attainment of the global minimum of the action. In practice, the HJB loss functions more as a regularization term than as a hard constraint that converges to zero. In addition, when using the modified metric, the goodness-of-fit for the distribution deteriorates, which warrants further investigation. Finally, the choice of the function $\psi(g)$ within the action is dependent on biological priors and is not automated. Future work could address these limitations by systematically investigating more canonical problems, which would provide clearer insight into the algorithm's fundamental behavior and limitations, and by automating the selection of $\psi(g)$ either by approximating it with a neural network or by deriving it from first-principle-based microscopic dynamics, such as a branching Wiener (Baradat and Lavenant 2021). We discussed more limitations of our work and potential broader impacts in Appendix E.1.

Acknowledgments and Disclosure of Funding

This work was supported by the National Key R&D Program of China (No. 2021YFA1003301 to T.L.), National Natural Science Foundation of China (NSFC No. 12288101 to T.L. & P.Z., and 8206100646, T2321001 to P.Z.) and The Fundamental Research Funds for the Central Universities, Peking University. We acknowledge the support from the High-performance Computing Platform of Peking University for computation.

References

- [1] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. “Stochastic interpolants: A unifying framework for flows and diffusions”. In: *arXiv preprint arXiv:2303.08797* (2023).
- [2] 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.
- [3] Ayméric Baradat and Hugo Lavenant. “Regularized unbalanced optimal transport as entropy minimization with respect to branching brownian motion”. In: *arXiv preprint arXiv:2111.01666* (2021).
- [4] Jean-David Benamou and Yann Brenier. “A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem”. In: *Numerische Mathematik* 84.3 (2000), pp. 375–393.
- [5] Jules Berman, Tobias Blickhan, and Benjamin Peherstorfer. *Parametric model reduction of mean-field and stochastic systems via higher-order action matching*. 2024. arXiv: [2410.12000 \[stat.ML\]](#).
- [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*. PMLR. 2023a, pp. 5802–5833.
- [7] Charlotte Bunne, Stefan G Stark, Gabriele Gut, Jacobo Sarabia Del Castillo, Mitch Levesque, Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. “Learning single-cell perturbation responses using neural optimal transport”. In: *Nature methods* 20.11 (2023b), pp. 1759–1768.
- [8] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. “Neural ordinary differential equations”. In: *Advances in neural information processing systems* 31 (2018).
- [9] Tianrong Chen, Guan-Hong Liu, and Evangelos Theodorou. “Likelihood Training of Schrödinger Bridge using Forward-Backward SDEs Theory”. In: *International Conference on Learning Representations*. 2022a.
- [10] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. “On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint”. In: *Journal of Optimization Theory and Applications* 169 (2016), pp. 671–691.
- [11] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. “The most likely evolution of diffusing and vanishing particles: Schrödinger bridges with unbalanced marginals”. In: *SIAM Journal on Control and Optimization* 60.4 (2022b), pp. 2016–2039.
- [12] Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. “An interpolating distance between optimal transport and Fisher–Rao metrics”. In: *Foundations of Computational Mathematics* 18 (2018a), pp. 1–44.
- [13] Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. “Unbalanced optimal transport: Dynamic and Kantorovich formulations”. In: *Journal of Functional Analysis* 274.11 (2018b), pp. 3090–3123.
- [14] Lénaïc Chizat, Stephen Zhang, Matthieu Heitz, and Geoffrey Schiebinger. “Trajectory inference via mean-field langevin in path space”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 16731–16742.
- [15] Shui-Nee Chow, Wuchen Li, and Haomin Zhou. “Wasserstein hamiltonian flows”. In: *Journal of Differential Equations* 268.3 (2020), pp. 1205–1219.
- [16] David P Cook and Barbara C Vanderhyden. “Context specificity of the EMT transcriptional response”. In: *Nature communications* 11.1 (2020), p. 2142.

- [17] Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. “Diffusion schrödinger bridge with applications to score-based generative modeling”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 17695–17709.
- [18] Jun Ding, Nadav Sharon, and Ziv Bar-Joseph. “Temporal modelling using single-cell transcriptomics”. In: *Nature Reviews Genetics* 23.6 (2022), pp. 355–368.
- [19] Luca Eyring, Dominik Klein, Théo Uscidda, Giovanni Palla, Niki Kilbertus, Zeynep Akata, and Fabian Theis. *Unbalancedness in Neural Monge Maps Improves Unpaired Domain Translation*. 2024a. arXiv: [2311.15100 \[cs.CV\]](#).
- [20] 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*. 2024b.
- [21] Nicolas Fournier and Arnaud Guillin. *On the rate of convergence in Wasserstein distance of the empirical measure*. 2013. arXiv: [1312.2128 \[math.PR\]](#).
- [22] Ivan Gentil, Christian Léonard, and Luigia Ripani. “About the analogy between optimal transport and minimal entropy”. In: *Annales de la Faculté des sciences de Toulouse: Mathématiques*. Vol. 26. 3. 2017, pp. 569–600.
- [23] Anming Gu, Edward Chien, and Kristjan Greenewald. “Partially Observed Trajectory Inference using Optimal Transport and a Dynamics Prior”. In: *The Thirteenth International Conference on Learning Representations*. 2025.
- [24] Matthieu Heitz, Yujia Ma, Sharvaj Kubal, and Geoffrey Schiebinger. “Spatial Transcriptomics Brings New Challenges and Opportunities for Trajectory Inference”. In: *Annual Review of Biomedical Data Science* 8 (2024).
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. “Denoising diffusion probabilistic models”. In: *Advances in neural information processing systems* 33 (2020), pp. 6840–6851.
- [26] Guillaume Hugué, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. “Manifold interpolating optimal-transport flows for trajectory inference”. In: *Advances in neural information processing systems* 35 (2022), pp. 29705–29718.
- [27] Qi Jiang and Lin Wan. “A physics-informed neural SDE network for learning cellular dynamics from time-series scRNA-seq data”. In: *Bioinformatics* 40 (2024), pp. ii120–ii127. ISSN: 1367-4811.
- [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”. In: *Nature* (2025), pp. 1–11.
- [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] Hugo Lavenant, Stephen Zhang, Young-Heon Kim, Geoffrey Schiebinger, et al. “Toward a mathematical theory of trajectory inference”. In: *The Annals of Applied Probability* 34.1A (2024), pp. 428–500.
- [31] Christian Léonard. “A survey of the Schrödinger problem and some of its connections with optimal transport”. In: *Discrete and Continuous Dynamical Systems-Series A* 34.4 (2014), pp. 1533–1574.
- [32] Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. “Scalable gradients for stochastic differential equations”. In: *International Conference on Artificial Intelligence and Statistics* (2020).
- [33] 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.
- [34] Xingchao Liu, Chengyue Gong, and Qiang Liu. “Flow straight and fast: Learning to generate and transfer data with rectified flow”. In: *arXiv preprint arXiv:2209.03003* (2022).

- [35] Malte Luecken, Daniel Burkhardt, Robrecht Cannoodt, Christopher Lance, Aditi Agrawal, Hananeh Aliee, Ann Chen, Louise Deconinck, Angela Detweiler, Alejandro Granados, Shelly Huynh, Laura Isacco, Yang Kim, Dominik Klein, BONY DE KUMAR, Sunil Kuppasani, Heiko Lickert, Aaron McGeever, Joaquin Melgarejo, Honey Mekonen, Maurizio Morri, Michaela Müller, Norma Neff, Sheryl Paul, Bastian Rieck, Kaylie Schneider, Scott Steelman, Michael Sterr, Daniel Treacy, Alexander Tong, Alexandra-Chloe Villani, Guilin Wang, Jia Yan, Ce Zhang, Angela Pisco, Smita Krishnaswamy, Fabian Theis, and Jonathan M Bloom. “A sandbox for prediction and integration of DNA, RNA, and proteins in single cells”. In: *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Ed. by J. Vanschoren and S. Yeung. Vol. 1. 2021.
- [36] Suryanarayana Maddu, Victor Chardès, Michael Shelley, et al. “Inferring biological processes with intrinsic noise from cross-sectional data”. In: *arXiv preprint arXiv:2410.07501* (2024).
- [37] 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”. In: *Nature biotechnology* 37.12 (2019), pp. 1482–1492.
- [38] Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. “Action matching: Learning stochastic dynamics from samples”. In: *International conference on machine learning*. PMLR. 2023, pp. 25858–25889.
- [39] 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.
- [40] Alessandro Palma, Till Richter, Hanyi Zhang, Manuel Lubetzki, Alexander Tong, Andrea Dittadi, and Fabian J Theis. “Multi-Modal and Multi-Attribute Generation of Single Cells with CFGen”. In: *The Thirteenth International Conference on Learning Representations*. 2025.
- [41] Matteo Pariset, Ya-Ping Hsieh, Charlotte Bunne, Andreas Krause, and Valentin De Bortoli. “Unbalanced Diffusion Schrödinger Bridge”. In: *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*. 2023.
- [42] Qiangwei Peng, Peijie Zhou, and Tiejun Li. “stVCR: Reconstructing spatio-temporal dynamics of cell development using optimal transport”. In: *bioRxiv* (2024), pp. 2024–06.
- [43] 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.
- [44] Martin Rohbeck, Charlotte Bunne, Edward De Brouwer, 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.
- [45] 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”. In: *Cell* 176.4 (2019a), pp. 928–943.
- [46] 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”. In: *Cell* 176.4 (2019b), pp. 928–943.
- [47] Yutong Sha, Yuchi Qiu, Peijie Zhou, and Qing Nie. “Reconstructing growth and dynamic trajectories from single-cell transcriptomics data”. In: *Nature Machine Intelligence* 6.1 (2024), pp. 25–39.
- [48] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. “Diffusion Schrödinger bridge matching”. In: *Advances in Neural Information Processing Systems* 36 (2024).
- [49] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. “Deep unsupervised learning using nonequilibrium thermodynamics”. In: *International conference on machine learning*. PMLR. 2015, pp. 2256–2265.
- [50] Jiaming Song, Chenlin Meng, and Stefano Ermon. “Denoising diffusion implicit models”. In: *arXiv preprint arXiv:2010.02502* (2020).

- [51] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. “Score-Based Generative Modeling through Stochastic Differential Equations”. In: *International Conference on Learning Representations*. 2021.
- [52] Alexander Tong, Kilian FATRAS, Nikolay Malkin, Guillaume Huguét, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. “Improving and generalizing flow-based generative models with minibatch optimal transport”. In: *Transactions on Machine Learning Research* (2024a). Expert Certification. ISSN: 2835-8856.
- [53] Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. “Trajectory-net: A dynamic optimal transport network for modeling cellular dynamics”. In: *International conference on machine learning*. PMLR. 2020, pp. 9526–9536.
- [54] Alexander Tong, Manik Kuchroo, Shabarni Gupta, Aarthi Venkat, Beatriz P San Juan, Laura Rangel, Brandon Zhu, John G Lock, Christine L Chaffer, and Smita Krishnaswamy. “Learning transcriptional and regulatory dynamics driving cancer cell plasticity using neural ODE-based optimal transport”. In: *bioRxiv* (2023), pp. 2023–03.
- [55] Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguét, Guy Wolf, and Yoshua Bengio. “Simulation-Free Schrödinger Bridges via Score and Flow Matching”. In: *International Conference on Artificial Intelligence and Statistics*. PMLR. 2024b, pp. 1279–1287.
- [56] Elias Ventre, Aden Forrow, Nitya Gadhiwala, Parijat Chakraborty, Omer Angel, and Geoffrey Schiebinger. “Trajectory inference for a branching SDE model of cell differentiation”. In: *arXiv preprint arXiv:2307.07687* (2023).
- [57] Adrian Veres, Aubrey L Faust, Henry L Bushnell, Elise N Engquist, Jennifer Hyoje-Ryu Kenty, George Harb, Yeh-Chuin Poh, Elad Sintov, Mads Gürtler, Felicia W Pagliuca, et al. “Charting cellular identity during human in vitro β -cell differentiation”. In: *Nature* 569.7756 (2019), pp. 368–373.
- [58] Wei Wan, Yuejin Zhang, Chenglong Bao, Bin Dong, and Zuoqiang Shi. “A scalable deep learning approach for solving high-dimensional dynamic optimal transport”. In: *SIAM Journal on Scientific Computing* 45.4 (2023), B544–B563.
- [59] Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. “Joint Velocity-Growth Flow Matching for Single-Cell Dynamics Modeling”. In: *Advances in Neural Information Processing Systems* (2025).
- [60] Zihao Wang, Datong Zhou, Yong Zhang, Hao Wu, and Chenglong Bao. *Wasserstein-Fisher-Rao Document Distance*. 2019. arXiv: 1904.10294 [cs.LG].
- [61] Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D Camargo, and Allon M Klein. “Lineage tracing on transcriptional landscapes links state to fate during differentiation”. In: *Science* 367.6479 (2020), eaaw3381.
- [62] Hao Wu, Shu Liu, Xiaojing Ye, and Haomin Zhou. “Parameterized wasserstein hamiltonian flow”. In: *SIAM Journal on Numerical Analysis* 63.1 (2025), pp. 360–395.
- [63] Liu Yang, Constantinos Daskalakis, and George E Karniadakis. “Generative ensemble regression: Learning particle dynamics from observations of ensembles with physics-informed deep generative models”. In: *SIAM Journal on Scientific Computing* 44.1 (2022), B80–B99.
- [64] Maosheng Yang. “Topological Schrödinger Bridge Matching”. In: *The Thirteenth International Conference on Learning Representations*. 2025.
- [65] 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”. In: *Nature communications* 12.1 (2021a), p. 3222.
- [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”. In: *Nature communications* 12.1 (2021b), p. 3222.
- [67] Yuning You, Ruida Zhou, and Yang Shen. “Correlational Lagrangian Schrödinger Bridge: Learning Dynamics with Population-Level Regularization”. In: *arXiv preprint arXiv:2402.10227* (2024).
- [68] Jiaqi Zhang, Erica Larschan, Jeremy Bigness, and Ritambhara Singh. “scNODE: generative model for temporal single cell transcriptomic data prediction”. In: *Bioinformatics* 40.Supplement_2 (2024a), pp. ii146–ii154. ISSN: 1367-4811.

- [69] Peng Zhang, Ting Gao, Jin Guo, and Jinqiao Duan. “Action Functional as Early Warning Indicator in the Space of Probability Measures”. In: *arXiv preprint arXiv:2403.10405* (2024b).
- [70] Stephen Zhang, Anton Afanassiev, Laura Greenstreet, Tetsuya Matsumoto, and Geoffrey Schiebinger. “Optimal transport analysis reveals trajectories in steady-state systems”. In: *PLoS computational biology* 17.12 (2021), e1009466.
- [71] 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*. 2025a.
- [72] Zhenyi Zhang, Yuhao Sun, Qiangwei Peng, Tiejun Li, and Peijie Zhou. “Integrating Dynamical Systems Modeling with Spatiotemporal scRNA-Seq Data Analysis”. In: *Entropy* 27.5 (2025b). ISSN: 1099-4300.
- [73] Zhenyi Zhang, Zihan Wang, Yuhao Sun, Tiejun Li, and Peijie Zhou. “Modeling Cell Dynamics and Interactions with Unbalanced Mean Field Schrödinger Bridge”. In: *Advances in Neural Information Processing Systems* (2025c).
- [74] Zhenyi Zhang, Zihan Wang, Yuhao Sun, Jiantao Shen, Qiangwei Peng, Tiejun Li, and Peijie Zhou. “Deciphering cell-fate trajectories using spatiotemporal single-cell transcriptomic data”. In: *Preprint* (2025d).
- [75] Linqi Zhou, Aaron Lou, Samar Khanna, and Stefano Ermon. “Denoising Diffusion Bridge Models”. In: *The Twelfth International Conference on Learning Representations*. 2024.
- [76] Qunxi Zhu, Bolin Zhao, Jingdong Zhang, Peiyang Li, and Wei Lin. “Governing equation discovery of a complex system from snapshots”. In: *arXiv preprint arXiv:2410.16694* (2024).

A Technical Details

A.1 Proof of Theorems

A.1.1 Proof for Theorem 4.1

Theorem A.1. *The RUOT problem with isotropic and time-invariant diffusion intensity is formulated as:*

$$\min_{\rho, g} A = \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}(\mathbf{x}, t)\|^2 + \alpha \psi(g(\mathbf{x}, t)) \right) \rho(\mathbf{x}, t) d\mathbf{x} dt \quad (6)$$

$$s.t. \quad \frac{\partial \rho(\mathbf{x}, t)}{\partial t} = -\nabla \cdot \left(\rho(\mathbf{x}, t) \mathbf{u}(\mathbf{x}, t) \right) + \frac{1}{2} \sigma^2 \nabla^2 \rho(\mathbf{x}, t) + g(\mathbf{x}, t) \rho(\mathbf{x}, t) \quad (7)$$

In this problem, the necessary conditions for the action A to achieve a minimum are given by:

$$\mathbf{u} = \nabla_{\mathbf{x}} \lambda, \quad \alpha \frac{d\psi(g)}{dg} = \lambda, \quad \frac{\partial \lambda}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + \lambda g - \alpha \psi(g) = 0, \quad (8)$$

where $\lambda(\mathbf{x}, t)$ is a scalar field.

Proof. In order to incorporate the constraints of the Fokker–Planck equation, we construct an augmented action functional:

$$\begin{aligned} A^\dagger &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}(\mathbf{x}, t)\|^2 + \alpha \psi(g(\mathbf{x}, t)) \right) \rho(\mathbf{x}, t) d\mathbf{x} dt \\ &\quad + \int_0^1 \int_{\mathbb{R}^d} \lambda(\mathbf{x}, t) \left(\frac{\partial \rho(\mathbf{x}, t)}{\partial t} + \nabla \cdot \left(\rho(\mathbf{x}, t) \mathbf{u}(\mathbf{x}, t) \right) - \frac{1}{2} \sigma^2 \nabla^2 \rho(\mathbf{x}, t) - g(\mathbf{x}, t) \rho(\mathbf{x}, t) \right) d\mathbf{x} dt \end{aligned}$$

We take variations with respect to \mathbf{u} , g , and ρ . At the stationary point of the functional, the variation of the augmented action functional must vanish.

Step 1: Variation with respect to \mathbf{u} .

Let $\mathbf{u} \rightarrow \mathbf{u} + \delta \mathbf{u}$. The variation of the augmented action is

$$\begin{aligned} \delta A^\dagger &= \int_0^1 \int_{\mathbb{R}^d} (\mathbf{u}^T \delta \mathbf{u}) \rho + \lambda \nabla_{\mathbf{x}} \cdot (\rho \delta \mathbf{u}) d\mathbf{x} dt \\ &= \int_0^1 \int_{\mathbb{R}^d} (\mathbf{u}^T \delta \mathbf{u}) \rho + \int_0^1 \int_{\mathbb{R}^d} (\nabla_{\mathbf{x}} \cdot (\lambda \rho \delta \mathbf{u}) - \rho (\nabla_{\mathbf{x}} \lambda)^T \delta \mathbf{u}) d\mathbf{x} dt \\ &= \int_0^1 \int_{\mathbb{R}^d} (\mathbf{u}^T \delta \mathbf{u}) \rho + \int_0^1 \int_{S^\infty} \lambda \rho (\delta \mathbf{u})^T d\mathbf{S} - \int_0^1 \int_{\mathbb{R}^d} \rho (\nabla_{\mathbf{x}} \lambda)^T \delta \mathbf{u} d\mathbf{x} dt \\ &= \int_0^1 \int_{\mathbb{R}^d} (\mathbf{u}^T - (\nabla_{\mathbf{x}} \lambda)^T) \rho \delta \mathbf{u} d\mathbf{x} dt \end{aligned}$$

Here, S^∞ denotes the boundary at infinity in \mathbb{R}^d and $d\mathbf{S}$ is the surface element. Based on the assumption that

$$\int_{S^\infty} \lambda \rho (\delta \mathbf{u})^T d\mathbf{S} = 0,$$

and using the arbitrariness of $\delta \mathbf{u}$, we obtain the optimality condition

$$\mathbf{u} = \nabla_{\mathbf{x}} \lambda.$$

Step 2: Variation with respect to g .

Let $g \rightarrow g + \delta g$, then the variation of the augmented action becomes

$$\delta A^\dagger = \int_0^1 \int_{\mathbb{R}^d} \left(\alpha \frac{d\psi(g)}{dg} - \lambda \right) \rho \delta g d\mathbf{x} dt.$$

Since δg is arbitrary, we immediately obtain the optimality condition

$$\alpha \frac{d\psi(g)}{dg} = \lambda.$$

Step 3: Variation with respect to ρ .

Let $\rho \rightarrow \rho + \delta\rho$. Then the variation of the augmented action is given by

$$\begin{aligned} \delta A^\dagger &= \int_0^1 \int_{\mathbb{R}^d} \left[\left(\frac{1}{2} \|\mathbf{u}\|^2 + \alpha \psi(g) \right) \delta\rho + \lambda \left(\frac{\partial \delta\rho}{\partial t} + \nabla_{\mathbf{x}} \cdot (\mathbf{u} \delta\rho) - \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 (\delta\rho) - g \delta\rho \right) \right] d\mathbf{x} dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}\|^2 + \alpha \psi(g) - \lambda g \right) \delta\rho d\mathbf{x} dt + \int_{\mathbb{R}^d} \int_0^1 \left(\frac{\partial(\lambda \delta\rho)}{\partial t} - \delta\rho \frac{\partial \lambda}{\partial t} \right) dt d\mathbf{x} \\ &\quad + \int_0^1 \int_{\mathbb{R}^d} (\nabla_{\mathbf{x}} \cdot (\mathbf{u} \lambda \delta\rho) d\mathbf{x} dt - \mathbf{u}^T (\nabla_{\mathbf{x}} \lambda) \delta\rho) d\mathbf{x} dt \\ &\quad - \frac{1}{2} \sigma^2 \int_0^1 \int_{\mathbb{R}^d} (\nabla_{\mathbf{x}} \cdot (\lambda \nabla_{\mathbf{x}} \delta\rho) - (\nabla_{\mathbf{x}} \lambda)^T (\nabla_{\mathbf{x}} \delta\rho)) d\mathbf{x} dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}\|^2 + \alpha \psi(g) - \lambda g - \frac{\partial \lambda}{\partial t} - \mathbf{u}^T (\nabla_{\mathbf{x}} \lambda) \right) \delta\rho d\mathbf{x} dt + \cancel{\int_{\mathbb{R}^d} [\lambda \delta\rho]_{t=0}^{t=1} d\mathbf{x}} \\ &\quad + \cancel{\int_0^1 \int_{S^\infty} \nabla_{\mathbf{x}} \cdot (\mathbf{u} \lambda \delta\rho)^T d\mathbf{S} dt} - \cancel{\frac{1}{2} \sigma^2 \int_0^1 \int_{S^\infty} (\lambda \nabla_{\mathbf{x}} \delta\rho)^T d\mathbf{S} dt} \\ &\quad + \frac{1}{2} \sigma^2 \int_0^1 \int_{\mathbb{R}^d} (\nabla_{\mathbf{x}} \lambda)^T (\nabla_{\mathbf{x}} \delta\rho) d\mathbf{x} dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}\|^2 + \alpha \psi(g) - \lambda g - \frac{\partial \lambda}{\partial t} - \mathbf{u}^T (\nabla_{\mathbf{x}} \lambda) \right) \delta\rho d\mathbf{x} dt \\ &\quad + \frac{1}{2} \sigma^2 \int_0^1 \int_{\mathbb{R}^d} (\nabla_{\mathbf{x}} \cdot ((\nabla_{\mathbf{x}} \lambda) \delta\rho) - (\nabla_{\mathbf{x}}^2 \lambda) \delta\rho) d\mathbf{x} dt \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}\|^2 + \alpha \psi(g) - \lambda g - \frac{\partial \lambda}{\partial t} - \mathbf{u}^T (\nabla_{\mathbf{x}} \lambda) - \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda \right) \delta\rho d\mathbf{x} dt \\ &\quad - \cancel{\frac{1}{2} \sigma^2 \int_0^1 \int_{S^\infty} ((\nabla_{\mathbf{x}} \lambda) \delta\rho)^T d\mathbf{S} dt} \\ &= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}\|^2 + \alpha \psi(g) - \lambda g - \frac{\partial \lambda}{\partial t} - \mathbf{u}^T (\nabla_{\mathbf{x}} \lambda) - \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda \right) \delta\rho d\mathbf{x} dt. \end{aligned}$$

Since $\delta\rho$ is arbitrary, the corresponding optimality condition is

$$\frac{1}{2} \|\mathbf{u}\|^2 + \alpha \psi(g) - \lambda g - \frac{\partial \lambda}{\partial t} - \mathbf{u}^T (\nabla_{\mathbf{x}} \lambda) - \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda = 0.$$

Substituting the previously obtained condition $\mathbf{u} = \nabla_{\mathbf{x}} \lambda$, we arrive at the final optimality condition:

$$\frac{\partial \lambda}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + \lambda g - \alpha \psi(g) = 0.$$

□

A.1.2 Proof for Theorem 4.2

Theorem A.2. *The choice of $\psi(g)$ affects whether g ascends or descends along the direction of the velocity field \mathbf{u} at a given time. Specifically, at a fixed time t , if*

$$\frac{d\psi^2(g)}{dg^2} > 0, \quad (9)$$

then $g(\mathbf{x}, t)$ ascends in the direction of the velocity field $\mathbf{u}(\mathbf{x}, t)$ (i.e., $\mathbf{u}(\mathbf{x}, t)^T (\nabla_{\mathbf{x}} g(\mathbf{x}, t)) > 0$); otherwise, it descends.

Proof. Let

$$\frac{d\psi(g)}{dg} = \psi'(g) \quad \text{and} \quad \frac{d^2\psi(g)}{dg^2} = \psi''(g).$$

Using the optimality condition for g from Appendix A.1.1,

$$\alpha \frac{d\psi(g)}{dg} = \lambda,$$

taking the gradient with respect to x on both sides yields

$$\nabla_{\mathbf{x}} g(\mathbf{x}, t) = \frac{1}{\alpha \psi''(g)} \nabla_{\mathbf{x}} \lambda(\mathbf{x}, t).$$

The condition for g to increase along the velocity field is that the inner product between $\nabla_{\mathbf{x}} g(\mathbf{x}, t)$ and $\mathbf{u}(\mathbf{x}, t)$ is positive everywhere. Using the optimality condition for the velocity,

$$\mathbf{u}(\mathbf{x}, t) = \nabla_{\mathbf{x}} \lambda(\mathbf{x}, t),$$

we have

$$\mathbf{u}(\mathbf{x}, t)^T \nabla_{\mathbf{x}} g(\mathbf{x}, t) = \frac{1}{\alpha \psi''(g)} \|\nabla_{\mathbf{x}} \lambda(\mathbf{x}, t)\|^2.$$

Since $\alpha > 0$, the condition

$$\mathbf{u}(\mathbf{x}, t)^T \nabla_{\mathbf{x}} g(\mathbf{x}, t) > 0$$

is equivalent to requiring that

$$\psi''(g) > 0 \quad \forall g.$$

□

A.1.3 Proof for Theorem 5.1

Theorem A.3. *Consider a weighted particle system consisting of N particles, where the position of particle i is given by $\mathbf{X}_i^t \in \mathbb{R}^d$ and its weight by $w_i(t) > 0$. The dynamics of each particle are described by*

$$\begin{aligned} \frac{d\mathbf{X}_i^t}{dt} &= \mathbf{u}(\mathbf{X}_i^t, t) dt + \sigma(t) d\mathbf{W}_t, \\ \frac{dw_i}{dt} &= g(\mathbf{X}_i^t, t) w_i dt, \end{aligned} \quad (10)$$

where $\mathbf{u} : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ is a time-varying vector field, $g : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}$ is a growth rate function, $\sigma : [0, T] \rightarrow [0, +\infty)$ is a time-varying diffusion coefficient, and \mathbf{W}_t is an N -dimensional standard Brownian motion with independent components in each coordinate. The initial conditions are $\mathbf{X}_i^0 \sim \rho(\mathbf{x}, 0)$ and $w_i(0) = 1$. In the limit as $N \rightarrow \infty$, the empirical measure

$$\mu^N(\mathbf{x}, t) = \frac{1}{N} \sum_{i=1}^N w_i(t) \delta(\mathbf{x} - \mathbf{X}_i^t) \quad (11)$$

converges to the solution of the following Fokker–Planck equation:

$$\frac{\partial \rho(\mathbf{x}, t)}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho(\mathbf{x}, t)) + \frac{1}{2} \sigma^2(t) \nabla_{\mathbf{x}}^2 \rho(\mathbf{x}, t) + g(\mathbf{x}, t) \rho(\mathbf{x}, t), \quad (12)$$

with the initial condition $\rho(\mathbf{x}, 0) = \rho_0(\mathbf{x})$.

Proof. Consider a smooth test function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$. We study the evolution of the expectation

$$\int_{\mathbb{R}^d} \phi(\mathbf{x}) \mu^N(\mathbf{x}, t) d\mathbf{x} = \frac{1}{N} \sum_{i=1}^N w_i(t) \phi(\mathbf{X}_i^t).$$

By applying Itô's formula, we have

$$d\left(w_i(t) \phi(\mathbf{X}_i^t)\right) = w_i(t) d\phi(\mathbf{X}_i^t) + \phi(\mathbf{X}_i^t) dw_i(t) + dw_i(t) d\phi(\mathbf{X}_i^t).$$

Using Itô's formula to compute $d\phi(\mathbf{X}_i^t)$, we obtain

$$d\phi(\mathbf{X}_i^t) = (\nabla_{\mathbf{x}} \phi(\mathbf{X}_i^t))^T d\mathbf{X}_i^t + \frac{1}{2} \nabla_{\mathbf{x}}^2 \phi(\mathbf{X}_i^t) \sigma^2(t) dt.$$

Since $dw_i(t) = g(\mathbf{X}_i^t, t) w_i(t) dt$ contains no stochastic term (i.e., there is no $d\mathbf{W}$), the term $dw_i(t) d\phi(\mathbf{X}_i^t)$ is of higher order and can be neglected. Therefore, we have:

$$\begin{aligned} d\left(w_i(t) \phi(\mathbf{X}_i^t)\right) &= \phi(\mathbf{X}_i^t) g(\mathbf{X}_i^t, t) w_i(t) dt \\ &\quad + w_i(t) (\nabla_{\mathbf{x}} \phi(\mathbf{X}_i^t))^T \left(\mathbf{u}(\mathbf{X}_i^t, t) dt + \sigma(t) d\mathbf{W}_t \right) \\ &\quad + \frac{1}{2} \nabla_{\mathbf{x}}^2 \phi(\mathbf{X}_i^t) \sigma^2(t) dt. \end{aligned}$$

Next, we compute

$$\begin{aligned} \mathbb{E} \left[\frac{d}{dt} \left(\frac{1}{N} \sum_{i=1}^N w_i(t) \phi(\mathbf{X}_i^t) \right) \right] &= \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \left(\phi(\mathbf{X}_i^t) g(\mathbf{X}_i^t, t) w_i(t) \right. \right. \\ &\quad \left. \left. + w_i(t) (\nabla_{\mathbf{x}} \phi(\mathbf{X}_i^t))^T \mathbf{u}(\mathbf{X}_i^t, t) \right. \right. \\ &\quad \left. \left. + \frac{1}{2} \nabla_{\mathbf{x}}^2 \phi(\mathbf{X}_i^t) \sigma^2(t) \right) \right]. \end{aligned}$$

Thus, in the limit as $N \rightarrow \infty$, and let $\rho(\mathbf{x}, t) = \mu^\infty(\mathbf{x}, t)$, we have

$$\begin{aligned} \frac{d}{dt} \int_{\mathbb{R}^d} \phi(\mathbf{x}) \rho(\mathbf{x}, t) d\mathbf{x} &= \int_{\mathbb{R}^d} \left(g(\mathbf{x}, t) \rho(\mathbf{x}, t) \phi(\mathbf{x}) + \rho(\mathbf{x}, t) (\nabla_{\mathbf{x}} \phi(\mathbf{x}))^T \mathbf{u}(\mathbf{x}, t) \right. \\ &\quad \left. + \frac{1}{2} \sigma^2(t) \rho(\mathbf{x}, t) \nabla_{\mathbf{x}}^2 \phi(\mathbf{x}) \right) d\mathbf{x}. \end{aligned}$$

By integrating by parts, we obtain

$$\int_{\mathbb{R}^d} \rho(\mathbf{x}, t) (\nabla_{\mathbf{x}} \phi(\mathbf{x}))^T \mathbf{u}(\mathbf{x}, t) d\mathbf{x} = - \int_{\mathbb{R}^d} \phi(\mathbf{x}, t) \nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho(\mathbf{x}, t)) d\mathbf{x},$$

and

$$\int_{\mathbb{R}^d} \rho(\mathbf{x}, t) \nabla_{\mathbf{x}}^2 \phi(\mathbf{x}) d\mathbf{x} = \int_{\mathbb{R}^d} \phi(\mathbf{x}, t) \nabla_{\mathbf{x}}^2 \rho(\mathbf{x}, t) d\mathbf{x}.$$

Hence, we deduce that

$$\frac{d}{dt} \int_{\mathbb{R}^d} \phi(\mathbf{x}, t) \rho(\mathbf{x}, t) d\mathbf{x} = \int_{\mathbb{R}^d} \phi(\mathbf{x}, t) \left[-\nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho(\mathbf{x}, t)) + \frac{1}{2} \sigma^2(t) \nabla_{\mathbf{x}}^2 \rho(\mathbf{x}, t) + g(\mathbf{x}, t) \rho(\mathbf{x}, t) \right] d\mathbf{x}.$$

Since $\phi(\mathbf{x})$ is arbitrary, we obtain the Fokker–Planck equation:

$$\frac{\partial \rho(\mathbf{x}, t)}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\mathbf{u}(\mathbf{x}, t) \rho(\mathbf{x}, t)) + \frac{1}{2} \sigma^2(t) \nabla_{\mathbf{x}}^2 \rho(\mathbf{x}, t) + g(\mathbf{x}, t) \rho(\mathbf{x}, t).$$

□

A.1.4 Proposition : the Expectation of HJB Loss

Proposition A.1. Consider the following HJB loss:

$$\mathcal{L}_{HJB}^N = \sum_{i=1}^N \left[\int_0^{T_K} \frac{w_i(t)}{\sum_{i=1}^N w_i(t)} \left(\frac{\partial \lambda(\mathbf{X}_i^t, t)}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + g(\mathbf{X}_i^t, t) - \alpha \psi(g) \right)^2 dt \right] \quad (13)$$

where

$$d\mathbf{X}_i^t = \mathbf{u}(\mathbf{X}_i^t, t) dt + \sigma(t) d\mathbf{W}_t, \quad dw_i = g(\mathbf{X}_i^t, t) w_i dt,$$

The expectation of HJB loss is

$$\mathbb{E}[\mathcal{L}_{HJB}^N] = \int_0^{T_K} \int_{\mathbb{R}^d} \hat{\rho}(\mathbf{x}, t) \left(\frac{\partial \lambda(\mathbf{x}, t)}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + g(\mathbf{x}, t) - \alpha \psi(g) \right)^2 d\mathbf{x} dt \quad (14)$$

where $\hat{\rho}(\mathbf{x}, t) = \frac{\rho(\mathbf{x}, t)}{\int_{\mathbb{R}^d} \rho(\mathbf{x}, t) d\mathbf{x}}$ is the normalized probability density.

Proof. Taking the expectation of \mathcal{L}_{HJB}^N is equivalent to drawing N particles each time to obtain \mathcal{L}_{HJB}^N , repeating this process infinitely many times, and computing the average of the \mathcal{L}_{HJB}^N obtained in each instance. Since the particles are independent, this operation is directly equivalent to taking the number of particles $N \rightarrow \infty$, thus:

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{HJB}^N] &= \lim_{N \rightarrow \infty} \sum_{i=1}^N \left[\int_0^{T_K} \frac{w_i(t)}{\sum_{j=1}^N w_j(t)} \left(\frac{\partial \lambda}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + g(\mathbf{X}_i^t, t) - \alpha \psi(g) \right)^2 dt \right] \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \left[\int_0^{T_K} \frac{w_i(t)}{\frac{1}{N} \sum_{j=1}^N w_j(t)} \left(\frac{\partial \lambda}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + g - \alpha \psi(g) \right)^2 dt \right] \\ &= \lim_{N \rightarrow \infty} \left[\int_0^{T_K} \int_{\mathbb{R}^d} \frac{1}{\int_{\mathbb{R}^d} \mu^N d\mathbf{x}} \mu^N(\mathbf{x}, t) \left(\frac{\partial \lambda}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + g - \alpha \psi(g) \right)^2 d\mathbf{x} dt \right] \\ &= \int_0^{T_K} \int_{\mathbb{R}^d} \frac{1}{\lim_{N \rightarrow \infty} \int_{\mathbb{R}^d} \mu^N d\mathbf{x}} \lim_{N \rightarrow \infty} \left(\mu^N \left(\frac{\partial \lambda}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + g - \alpha \psi(g) \right)^2 d\mathbf{x} dt \right) \\ &= \int_0^{T_K} \int_{\mathbb{R}^d} \hat{\rho}(\mathbf{x}, t) \left(\frac{\partial \lambda(\mathbf{x}, t)}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + g(\mathbf{x}, t) - \alpha \psi(g) \right)^2 d\mathbf{x} dt \end{aligned}$$

In the final equality of the proof, we employed the previously proven Appendix A.1.3. \square

A.1.5 Proposition : the Expectation of Action Loss

Proposition A.2. Consider the following action loss:

$$\mathcal{L}_{Action}^N = \frac{1}{N} \sum_{i=1}^N \left(\int_0^1 \frac{1}{2} \|\mathbf{u}(\mathbf{X}_i^t, t)\|^2 w_i(t) dt + \int_0^1 \alpha \psi(g(\mathbf{X}_i^t, t)) w_i(t) dt \right) \quad (15)$$

where

$$d\mathbf{X}_i^t = \mathbf{u}(\mathbf{X}_i^t, t) dt + \sigma(t) d\mathbf{W}_t, \quad dw_i = g(\mathbf{X}_i^t, t) w_i dt,$$

The expectation of action loss is equal to the action defined in the RUOT formulation, namely,

$$\mathbb{E}[\mathcal{L}_{Action}^N] = \int_0^{T_K} \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}(\mathbf{x}, t)\|^2 + \alpha \psi(g(\mathbf{x}, t)) \right) \rho(\mathbf{x}, t) d\mathbf{x} dt. \quad (16)$$

Proof. Taking the expectation of $\mathcal{L}_{\text{Action}}^N$ is equivalent to drawing N particles each time to obtain $\mathcal{L}_{\text{Action}}^N$, repeating this process infinitely many times, and computing the average of the $\mathcal{L}_{\text{Action}}^N$ obtained in each instance. Since the particles are independent, this operation is directly equivalent to taking the number of particles $N \rightarrow \infty$, thus:

$$\begin{aligned}\mathbb{E}[\mathcal{L}_{\text{Action}}^N] &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \left[\int_0^{T_K} \frac{1}{2} \|\mathbf{u}(\mathbf{X}_i^t, t)\|^2 w_i(t) dt + \int_0^{T_K} \alpha \psi(g(\mathbf{X}_i^t, t)) w_i(t) dt \right] \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \left[\int_0^{T_K} \int_{\mathbb{R}^d} \frac{1}{2} \|\mathbf{u}(\mathbf{x}, t)\|^2 \mu^N(\mathbf{x}) dt \right. \\ &\quad \left. + \int_0^{T_K} \int_{\mathbb{R}^d} \alpha \psi(g(\mathbf{x}, t)) \mu^N(\mathbf{x}) dt \right] \\ &= \int_0^{T_K} \int_{\mathbb{R}^d} \left(\frac{1}{2} \|\mathbf{u}(\mathbf{x}, t)\|^2 + \alpha \psi(g(\mathbf{x}, t)) \right) \rho(\mathbf{x}, t) d\mathbf{x} dt.\end{aligned}$$

In the final equality of the proof, we employed the previously proven Appendix A.1.3. \square

A.2 Optimality Conditions Under Different $\psi(g)$

In our Experiment, we use two different $\psi(g)$ as examples. When $\psi_1(g) = \frac{1}{2}g^2$, the optimality conditions are :

$$\mathbf{u} = \nabla_{\mathbf{x}} \lambda, \quad g = \frac{\lambda}{\alpha}, \quad \frac{\partial \lambda}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda + \frac{1}{2} \frac{\lambda^2}{\alpha} = 0.$$

When $\psi_2(g) = g^{2/15}$, the optimality conditions are :

$$\mathbf{u} = \nabla_{\mathbf{x}} \lambda, \quad g = \left(\frac{2\alpha}{15\lambda} \right)^{\frac{15}{13}}, \quad \frac{\partial \lambda}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda - \frac{13}{15} \alpha \left(\frac{2\alpha}{15\lambda} \right)^{\frac{2}{13}} = 0.$$

Note that in this case the function $g(\lambda)$ exhibits a singularity at $\lambda = 0$. In fact, given the two properties we imposed on $\psi(g)$ ($\frac{d\psi(g)}{d|g|} > 0$ and $\psi(g) = \psi(-g)$) along with the constraint $\psi''(g) < 0$, it follows that $\psi'(g)$ must be discontinuous at 0, and hence $g(\lambda) = (\psi')^{-1}\left(\frac{\lambda(\mathbf{x}, t)}{\alpha}\right)$ necessarily has a singularity at $\lambda = 0$. For the sake of training stability, we slightly modify $g(\lambda)$ to remove this singularity. We redefine $g(\lambda)$ as:

$$g^\dagger(\lambda) = \begin{cases} g(\lambda), & \lambda \geq \delta, \\ \frac{1}{2\delta} (g(\delta) - g(-\delta))(\lambda + \delta) + g(-\delta), & -\delta \leq \lambda < \delta, \\ g(-\lambda), & \lambda < -\delta, \end{cases}$$

where δ is a small positive constant. In our computations, we set $\delta = 0.1$. Another possible treatment for this singularity involves fitting the reciprocal of λ , μ , thereby canceling the singularity at $\lambda = 0$. In this case, $g = \left(\frac{2\alpha\mu}{15} \right)^{\frac{15}{13}}$, and thus the singularity near zero points no longer exists. We leave how to eliminate this singularity more systematically in numerical computation as the future work.

A.3 Training Algorithm and Computation of Losses

For the algorithms mentioned above, we use the particle method to compute the losses. In Appendix D.2, we discuss the convergence rates of these losses: both \mathcal{L}_{HJB} and $\mathcal{L}_{\text{Action}}$ have a convergence rate of $\mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$, while for data with high dimensionality ($d \gg 1$), the convergence

Algorithm 1 Training Var-RUOT

Require: Datasets D_1, \dots, D_K , batch size N , training epochs N_{Epoch} , initialized network $\lambda_\theta(\mathbf{x}, t)$.

Ensure: Trained scalar field $\lambda_\theta(\mathbf{x}, t)$.

```
1: for  $i = 1$  to  $N_{\text{Epoch}}$  do
2:   From the data at the first time point,  $D_1$ , sample  $N$  particles and set all their weights to
      $w_i(0) = 1$ , for  $i = \{1, 2, \dots, N\}$ .
3:   for  $j = 1$  to  $K$  do
4:     use optimality conditions  $\mathbf{u}_\theta(\mathbf{x}, t) = \nabla_{\mathbf{x}} \lambda_\theta(\mathbf{x}, t)$ ,  $\alpha \frac{d\psi(g_\theta)}{dg_\theta} = \lambda_\theta(\mathbf{x}, t)$  to calculate
      $\mathbf{u}_\theta(\mathbf{x}, t)$  and  $g_\theta(\mathbf{x}, t)$ , where  $t \in [T_i, T_{i+1})$ 
      $\mathcal{L}_{\text{Action}}^N \leftarrow \mathcal{L}_{\text{Action}}^N +$ 
5:        $\frac{1}{N} \sum_{i=1}^N \left( \int_{T_i}^{T_{i+1}} \frac{1}{2} \|\mathbf{u}_\theta(\mathbf{X}_i^t, t)\|^2 w_i(t) dt + \int_{T_i}^{T_{i+1}} \alpha \psi(g_\theta(\mathbf{X}_i^t, t)) w_i(t) dt \right)$ 
      $\mathcal{L}_{\text{HJB}}^N \leftarrow \mathcal{L}_{\text{HJB}}^N + \sum_{i=1}^N \left[ \int_{T_i}^{T_{i+1}} \frac{w_i(t)}{\sum_{j=1}^N w_j(t)} \right.$ 
6:        $\times \left( \frac{\partial \lambda_\theta(\mathbf{X}_i^t, t)}{\partial t} + \frac{1}{2} \|\nabla_{\mathbf{x}} \lambda_\theta\|^2 + \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}}^2 \lambda_\theta + g_\theta(\mathbf{X}_i^t, t) - \alpha \psi(g_\theta) \right)^2 dt \Big]$ 
7:    $\mathcal{L}_{\text{Recon}} \leftarrow \mathcal{L}_{\text{Recon}} + (M(T_{i+1}) - \hat{M}(T_{i+1}))^2 + \mathcal{W}_2(\tilde{\rho}(\cdot, T_{i+1}), \hat{\rho}(\cdot, T_{i+1}))$ 
8:    $\mathcal{L}_{\text{Total}} = \gamma_{\text{Action}} \mathcal{L}_{\text{Action}} + \gamma_{\text{HJB}} \mathcal{L}_{\text{HJB}} + \mathcal{L}_{\text{Recon}}$ 
9:   update  $\lambda_\theta(\mathbf{x}, t)$  w.r.t.  $\mathcal{L}_{\text{Total}}$ 
```

rate of \mathcal{L}_{OT} is $\mathcal{O}(N^{-\frac{2}{d}})$. To balance training speed, memory consumption, and the algorithm's convergence rate, we recommend using approximately $N = 10^3$ to 10^4 particles to compute \mathcal{L}_{HJB} , $\mathcal{L}_{\text{Action}}$, and $\mathcal{L}_{\text{Recon}}$. Furthermore, to ensure high accuracy in the integration for \mathcal{L}_{HJB} and $\mathcal{L}_{\text{Action}}$, we advise using a high-order integrator (Berman et al. 2024) (e.g., a Stochastic Runge-method (Li et al. 2020)) and setting the integration step size to approximately $\Delta t = 0.1$. $\mathcal{L}_{\text{Mass}}$ measures the discrepancy between the mass learned by the model and the true mass. The true mass at time point i (for $i > 1$) is defined as the ratio of the number of data points at that time, N_i , to the number of data points at the initial time, N_1 .

In particular, computing \mathcal{L}_{HJB} requires second-order derivatives to calculate the Laplacian $\nabla_x^2 \lambda(\mathbf{x}, t)$. This step can be computationally expensive. However, we only need the trace of the Hessian (i.e., the Laplacian), not the full matrix. This can be efficiently estimated using Hutchinson's randomized trace estimator. For a Hessian matrix H , the trace is given by $\text{trace}(H) = \mathbb{E}_{\mathbf{v}}[\mathbf{v}^T H \mathbf{v}]$, where the vector \mathbf{v} is drawn from a distribution with zero mean and unit covariance, such as $\mathcal{N}(0, I)$.

B Experiential Details

B.1 Additional Information for Datasets

Simulation Dataset In the main text, we utilize a simulated dataset derived from a three-gene regulatory network (Zhang et al. 2025a). The dynamics of this system are governed by stochastic differential equations that incorporate self-activation, mutual inhibition, and external activation. The dynamics of the three genes are described by the following equations:

$$\begin{aligned} \frac{dX_1^i}{dt} &= \frac{\alpha_1 (X_1^i)^2 + \beta}{1 + \gamma_1 (X_1^i)^2 + \alpha_2 (X_2^i)^2 + \gamma_3 (X_3^i)^2 + \beta} - \delta_1 X_1^i + \eta_1 \xi_t, \\ \frac{dX_2^i}{dt} &= \frac{\alpha_2 (X_2^i)^2 + \beta}{1 + \gamma_1 (X_1^i)^2 + \alpha_2 (X_2^i)^2 + \gamma_3 (X_3^i)^2 + \beta} - \delta_2 X_2^i + \eta_2 \xi_t, \\ \frac{dX_3^i}{dt} &= \frac{\alpha_3 (X_3^i)^2}{1 + \alpha_3 (X_3^i)^2} - \delta_3 X_3^i + \eta_3 \xi_t, \end{aligned}$$

where $\mathbf{X}^i(t)$ represents the gene expression levels of the i th cell at time t . The coefficients α_i , γ_i , and β control the strengths of self-activation, inhibition, and the external stimulus, respectively.

The parameters δ_i indicate the rates of gene degradation, and the terms $\eta_i \xi_t$ account for stochastic influences using additive white noise.

The probability of cell division is linked to the expression level of X_2 and is given by

$$g = \alpha_g \frac{X_2^2}{1 + X_2^2}.$$

When a cell divides, the resulting daughter cells are created with each gene perturbed by an independent random noise term, $\eta_d N(0, 1)$, around the parent cell's gene expression profile $(X_1(t), X_2(t), X_3(t))$. Detailed hyper-parameters are provided in Table 5. The initial population of cells is independently sampled from two normal distributions, $\mathcal{N}([2, 0.2, 0], 0.1)$ and $\mathcal{N}([0, 0, 2], 0.1)$. At every time step, any negative expression values are set to zero.

Table 5: Simulation parameters for the gene regulatory network.

Parameter	Value	Description
α_1	0.5	Self-activation strength for X_1 .
γ_1	0.5	Inhibition strength exerted by X_3 on X_1 .
α_2	1	Self-activation strength for X_2 .
γ_2	1	Inhibition strength exerted by X_3 on X_2 .
α_3	1	Self-activation strength for X_3 .
γ_3	10	Half-saturation constant in the inhibition term.
δ_1	0.4	Degradation rate for X_1 .
δ_2	0.4	Degradation rate for X_2 .
δ_3	0.4	Degradation rate for X_3 .
η_1	0.05	Noise intensity for X_1 .
η_2	0.05	Noise intensity for X_2 .
η_3	0.01	Noise intensity for X_3 .
η_d	0.014	Noise intensity for perturbations during cell division.
β	1	External signal activating X_1 and X_2 .
dt	1	Time step size.
Time Points	[0, 8, 16, 24, 32]	Discrete time points when data is recorded.

Other Datasets Used in Main Text In addition to the three-gene simulated dataset, our main text also utilizes the EMT dataset and the Mouse Blood Hematopoiesis dataset. The EMT dataset is sourced from (Sha et al. 2024; Cook and Vanderhyden 2020) and is derived from A549 cancer cells undergoing TGF β 1-induced epithelial-mesenchymal transition (EMT). It comprises data from four distinct time points, containing a total of 3133 cells, with each cell represented by 10 features obtained through PCA dimensionality reduction. Meanwhile, the Mouse Blood Hematopoiesis dataset covers 3 time points and includes 10,998 cells in total (Weinreb et al. 2020; Sha et al. 2024) and was reduced to 2-dimensional space using nonlinear dimensionality reduction.

Gaussian Datasets To examine the differences between the optimal transport paths solved by our algorithm and those obtained from traditional OT solvers, and to validate the scalability of our model (i.e., its learning capability on high-dimensional data), we constructed two 2D Gaussian datasets (a balanced version and an unbalanced version) and three high-dimensional Gaussian datasets. The 2D Gaussian dataset consists of two time points, where the distribution at each time point is a mixture of multiple Gaussian distributions, as illustrated in Figure 5. The three high-dimensional Gaussian datasets have dimensions of 50, 100, and 150, respectively. The 100-dimensional dataset was adapted from the experiments in DeepRUOT (Zhang et al. 2025a). The results of these high-dimensional datasets, after being reduced to two dimensions using PCA for visualization, are shown in Figure 5.

Other High Dimensional Datasets In addition, we employed four real-world datasets for our evaluation. The first is the Mouse Blood Hematopoiesis dataset from (Weinreb et al. 2020), which comprises 49,302 cells collected at three time points. We reduced its dimensionality to 50 using PCA, and a subset of this processed data was used for the experiments in our main text. The second is the Pancreatic β -cell Differentiation dataset from (Veres et al. 2019), consisting of 51,274 cells sampled across eight time points. We subsequently reduced its dimensionality to 30 via PCA. The third dataset, sourced from (Moon et al. 2019), profiles 16,819 cells sampled at five time points from human embryoid bodies (EBs), which developed via the spontaneous aggregation of stem cells in




Figure 5: Diagram of the Gaussian mixture datasets. PCA is used to reduce the high dimensional data to 2 dimensions.

3D culture throughout a 27-day differentiation process. We reduced its dimensionality to 100 using PCA and then standardized each feature. The fourth dataset is from the NeurIPS Challenge (Luecken et al. 2021). From this, we utilized the gene expression data for the donor with DonorID 28483. We extracted the 1,079 cells from this donor that were annotated with pseudotime values. To create discrete time points, we performed K-Means clustering on these cells based on their pseudotime. The resulting clusters were then ordered by their centroid pseudotime values and assigned discrete labels $t = 0, 1, 2, 3, 4$. This data was also standardized on a per-feature basis.

B.2 Evaluation Metrics

To assess the fitting accuracy of the learned dynamics to the data distribution, we compute the \mathcal{W}_1 and \mathcal{W}_2 distances between the data points generated by the model and the real data points. They are defined as

$$\mathcal{W}_1(p, q) = \min_{\pi \in \Pi(p, q)} \int \|\mathbf{x} - \mathbf{y}\|_2 d\pi(\mathbf{x}, \mathbf{y}),$$

and

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

We compute these two metrics using the emd function from the pot library.

To evaluate the action of the dynamics learned by the model, we directly compute the action loss. Appendix A.1.5 guarantees that the expectation of the loss is equal to the action defined in the RUOT problem. The action loss is:

$$\mathcal{L}_{\text{Action}}^N = \frac{1}{N} \sum_{i=1}^N \left(\int_0^1 \frac{1}{2} \|\mathbf{u}(\mathbf{X}_i^t, t)\|^2 w_i(t) dt + \int_0^1 \alpha \psi(g(\mathbf{X}_i^t, t)) w_i(t) dt \right)$$

where

$$d\mathbf{X}_i^t = \mathbf{u}(\mathbf{X}_i^t, t) dt + \sigma(t) d\mathbf{W}_t,$$

$$dw_i = g(\mathbf{X}_i^t, t) w_i dt,$$

with initial conditions $\mathbf{X}_i^0 \sim \rho(\mathbf{x}, 0)$ and $w_i(0) = 1$. We run our model 5 times on each dataset, to calculate the mean and standard deviation of $\mathcal{W}_1, \mathcal{W}_2$ and action.

To evaluate the training speed of the model, we use the SamplesLoss class from the geomloss library to compute the OT loss at each epoch during the training process for each method, with the blur parameter set to 0.10. We sum the OT losses at all time points to obtain the total OT loss. For each model, we perform 5 training runs, recording the number of epochs and the time required for the OT loss to drop below a specified threshold. We then compute the mean and standard deviation of these values, with the mean reflecting the training/convergence speed and the standard deviation reflecting the training stability.

For models whose dynamics are governed by stochastic differential equations, the choice of σ directly affects the results (both the OT loss and the path action). Therefore, when running the RUOT and our Var-RUOT models on each dataset, σ is set to 0.10.

C Additional Experiment Results

C.1 Additional Results on Training Speed and Stability

We plotted the average loss per epoch across five training runs in Fig. 6. Experimental results show that on the Simulation Gene dataset, our algorithm converges approximately 10 times faster than the fastest among the other algorithms (RUOT with 3-epoch pretraining), and on the EMT dataset, our algorithm converges roughly 20 times faster than the fastest alternative (TIGON).




Figure 6: Loss curves of each algorithm over five training runs. The curves are obtained by averaging the losses from the five runs.

C.2 Hyperparameter Selection and Ablation Study

Hyperparameter Selection We used NVIDIA A100 GPUs (with 40G memory) and 128-core CPUs to conduct the experiments described in this paper. The neural network used to fit $\lambda(\mathbf{x}, t)$ is a fully connected network augmented with layer normalization and residual connections. It consists of 2 hidden layers, each with 512 dimensions. In our algorithm, the main hyperparameters that need tuning include the penalty coefficient α for growth in the action, and the weights γ_{HJB} and γ_{Action} for the two regularization losses, \mathcal{L}_{HJB} and $\mathcal{L}_{\text{Action}}$, respectively. Here, α represents our prior regarding the strength of cell birth and death in the data: a larger α imposes a greater penalty on cell birth and death, thereby making it easier for the model to learn solutions with lower birth and death intensities. Meanwhile, the HJB loss and the action loss, serving as regularizers, are both designed to ensure that the solution obtained by the algorithm has as low an action as possible—the HJB equation being a necessary condition for the action to reach its minimum, and the inclusion of the action loss ensuring that the model learns a solution with an even smaller action under those necessary conditions.

To ensure that our algorithm generalizes well across a wide range of real-world datasets, we only used two sets of parameters: one for the standard WFR Metric ($\psi(g) = \frac{1}{2}g^2$) and one for the Modified Metric ($\psi(g) = g^{2/15}$). The parameters used in each case are listed in Table 6. The primary reason for using two sets is that different metrics yield different scales for the HJB loss.

Table 6: Parameter settings for standard WFR metric and modified metric.

Parameter	γ_{HJB}	γ_{Action}	α	Learning Rate	Optimizer
Standard WFR Metric ($\psi_1(g) = \frac{1}{2}g^2$)	6.25×10^{-2}	6.25×10^{-2}	2.00	1×10^{-4}	AdamW
Modified Metric ($\psi_2(g) = g^{2/15}$)	6.25×10^{-3}	6.25×10^{-2}	7.00	2×10^{-5}	AdamW

Sensitivity Analysis of α To demonstrate the robustness of our algorithm with respect to hyperparameter selection, we first varied the penalty coefficient α for growth and examined the resulting

changes in model performance. This sensitivity analysis was conducted on the 2D Mouse Blood Hematopoiesis dataset, and we performed the analysis for both the Standard WFR metric and the modified metric. The performance of the model under different values of α is shown in Table 7. The experimental results indicate that our algorithm is not sensitive to α , as similar performance can be achieved across multiple different values of α . In comparison with the Standard WFR metric, however, the algorithm appears to be somewhat more sensitive to α when the modified metric is used.

Table 7: On the Mouse Blood Hematopoiesis dataset, the Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between the predicted distributions of Var-RUOT with different α and the true distribution at each time points. Each experiment was run five times to compute the mean and standard deviation.

Model	$t = 1$		$t = 2$	
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2
Var-RUOT (Standard WFR, $\alpha = 1$)	0.1622 ± 0.0072	0.2027 ± 0.0097	0.1280 ± 0.0123	0.1522 ± 0.0178
Var-RUOT (Standard WFR, $\alpha = 2$)	0.1203 ± 0.0060	0.1498 ± 0.0043	0.1389 ± 0.0068	0.1701 ± 0.0096
Var-RUOT (Standard WFR, $\alpha = 3$)	0.1402 ± 0.0054	0.1704 ± 0.0077	0.1350 ± 0.0100	0.1655 ± 0.0132
Var-RUOT (Modified Metric, $\alpha = 5$)	0.3783 ± 0.0194	0.3326 ± 0.0128	0.2110 ± 0.0164	0.2226 ± 0.0219
Var-RUOT (Modified Metric, $\alpha = 7$)	0.2953 ± 0.0357	0.3117 ± 0.0323	0.1917 ± 0.0140	0.2226 ± 0.0170
Var-RUOT (Modified Metric, $\alpha = 9$)	0.2737 ± 0.0095	0.3116 ± 0.0072	0.1970 ± 0.0072	0.2224 ± 0.0075

Sensitivity Analysis of Growth Penalty Function

In this work, we select $\psi_2(g) = g^{2/15}$ to illustrate a case where the conditions $\psi'(g) > 0$ and $\psi''(g) < 0$ hold. As established in Theorem 4.2, these conditions ensure that the growth rate $g(\mathbf{x}, t)$ increases along the direction of the velocity field $\mathbf{u}(\mathbf{x}, t)$ at each time step.

The specific choice of $\psi_2(g) = g^{2/15}$ is motivated by its simplicity and, more importantly, by considerations for training stability. For a general penalty of the form $\psi(g) = g^{(2p)/(2q+1)}$, the first-order optimality condition $\alpha \frac{d\psi(g)}{dg} = \lambda$ yields an analytical solution for g in terms of λ :

$$g(\lambda) = \left(\frac{\lambda}{\alpha} \frac{2q+1}{2p} \right)^{\frac{2q+1}{2p-(2q+1)}} \propto \left(\frac{1}{\lambda} \right)^{\frac{2q+1}{(2q+1)-2p}}.$$

Since g is computed directly from λ during training, we must avoid an excessively steep $g(\lambda)$ curve, especially near $\lambda = 0$, to ensure stability. This requires the exponent $\frac{2q+1}{(2q+1)-2p}$ to be small. We achieve this by setting $p = 1$ (the smallest positive integer) and choosing a large value for q . Selecting $q = 7$ results in our final penalty function, $\psi_2(g) = g^{2/15}$, which balances simplicity and stability.

To empirically validate this choice, we conducted a sensitivity study on the 2D Mouse Blood Hematopoiesis dataset using different parameter choices for $\psi(g)$. The results, shown in Table 8, compare the Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) at different time points.

Table 8: Sensitivity analysis on the choice of $\psi(g)$.

Penalty Function $\psi(g)$	$t = 1$		$t = 2$	
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2
$\psi(g) = g^{2/9}$ ($p = 1, q = 4$)	0.9476 ± 0.0835	1.0824 ± 0.0887	1.2197 ± 0.1450	1.3414 ± 0.1449
$\psi(g) = g^{2/15}$ ($p = 1, q = 7$)	0.2953 ± 0.0357	0.3117 ± 0.0323	0.1917 ± 0.0140	0.1764 ± 0.0135
$\psi(g) = g^{2/21}$ ($p = 1, q = 10$)	0.4239 ± 0.0181	0.5009 ± 0.0162	0.2638 ± 0.0146	0.2194 ± 0.0175

The results demonstrate that $\psi(g) = g^{2/15}$ achieves the best performance. Its performance is comparable to that of $\psi(g) = g^{2/21}$, whereas the choice of $\psi(g) = g^{2/9}$ (with a smaller q) leads to poorer reconstruction quality due to training instability. This confirms that selecting an exponent within a reasonable range (i.e., with a sufficiently large q) is effective and stable.

Ablation Study of γ_{HJB} and γ_{Action} In order to verify whether \mathcal{L}_{HJB} and $\mathcal{L}_{\text{Action}}$ facilitate the algorithm find solutions with lower action, we conducted ablation studies. These experiments were carried out on the EMT data and 2D Mouse Blood Hematopoiesis Data, since in this dataset the transition from the initial distribution to the terminal distribution can be achieved through relatively

simple dynamics (each particle moving in a roughly straight line). Therefore, if the HJB loss and the action loss are effective, the model will learn this simple dynamics rather than a more complex one. We varied the HJB loss weight γ_{HJB} to the following values: $[0, 6.25 \times 10^{-3}, 3.125 \times 10^{-2}, 6.25 \times 10^{-2}, 6.25 \times 10^{-1}, 3.125]$, while keeping the action loss weight γ_{Action} fixed at 1. We then shown both the mean \mathcal{W}_1 distances between the predicted and true distributions at four different time points and the trajectory action (as shown in Table 9). Similarly, we fixed $\gamma_{\text{HJB}} = 1$ and varied λ_{Action} over the same set of values: $[0, 6.25 \times 10^{-3}, 3.125 \times 10^{-2}, 6.25 \times 10^{-2}, 6.25 \times 10^{-1}, 3.125]$, with the corresponding results illustrated in Table 10. The figures indicate that as γ_{HJB} and γ_{Action} increase, the action of the learned trajectories decreases monotonically, demonstrating that both loss terms are effective. However, as these weights increase, the model’s ability to fit the distribution deteriorates. Therefore, we recommend that in practical applications, both γ_{HJB} and γ_{Action} should be set to values below 0.1, as configured in this paper.

Table 9: Comparison of \mathcal{W}_1 and Path Action for different values of the parameter γ_{HJB} on a specific dataset.

Dataset	Metric	Value of γ_{HJB}					
		0	6.25×10^{-3}	3.125×10^{-2}	6.25×10^{-2}	6.25×10^{-1}	3.125
Mouse Blood Hematopoiesis	\mathcal{W}_1	0.1573	0.1609	0.1550	0.1316	0.2109	0.2539
	Path Action	3.3002	3.2165	3.2397	3.1491	3.1182	3.1200
EMT	\mathcal{W}_1	0.2702	0.2701	0.2686	0.2719	0.2722	0.2748
	Path Action	0.4512	0.4519	0.3867	0.3641	0.2865	0.2796

Table 10: Comparison of \mathcal{W}_1 and Path Action for different values of the parameter γ_{Action} on a specific dataset.

Dataset	Metric	Value of γ_{Action}					
		0	6.25×10^{-3}	3.125×10^{-2}	6.25×10^{-2}	6.25×10^{-1}	3.125
Mouse Blood Hematopoiesis	\mathcal{W}_1	0.1454	0.1344	0.1674	0.1316	0.4001	0.7914
	Path Action	3.4043	3.4497	3.2737	3.1491	2.4716	1.2059
EMT	\mathcal{W}_1	0.2706	0.2681	0.2656	0.2774	0.3043	0.3832
	Path Action	0.4192	0.4171	0.3934	0.3288	0.1686	0.0709

The Effectiveness of different types of Loss Norms

In the experiments above, we used the L_2 loss for computing both \mathcal{L}_{HJB} and $\mathcal{L}_{\text{Action}}$. To investigate the effect of the loss function type, we present the results of using the L_1 loss. On the Simulation and EMT datasets, the experimental results with the L_1 loss are shown in Tables 11 and 12. The results indicate that with our hyperparameter values ($\gamma_{\text{HJB}} = 6.25 \times 10^{-2}$, $\gamma_{\text{Action}} = 6.25 \times 10^{-2}$), using either the L_1 loss or the L_2 loss when computing the HJB loss achieves similar reconstruction performance and Path Action values.

Table 11: Performance comparison of Var-RUOT with L_2 and L_1 loss functions on Simulation Dataset. The table shows \mathcal{W}_1 and \mathcal{W}_2 distances, as well as the final path action.

Model	$t = 1$		$t = 2$		$t = 3$		$t = 4$		Path Action
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	
Var-RUOT (L_2 Loss)	0.0452 ± 0.0024	0.1181 ± 0.0064	0.0385 ± 0.0022	0.1270 ± 0.0121	0.0445 ± 0.0033	0.1144 ± 0.0160	0.0572 ± 0.0034	0.2140 ± 0.0067	1.1105 ± 0.0515
Var-RUOT (L_1 Loss)	0.0489 ± 0.0024	0.1112 ± 0.0135	0.0556 ± 0.0038	0.1541 ± 0.0175	0.0589 ± 0.0062	0.1372 ± 0.0226	0.0664 ± 0.0040	0.2181 ± 0.0130	1.1249 ± 0.0035

Table 12: Performance comparison of Var-RUOT with L_2 and L_1 loss functions on EMT Dataset for time steps $t = 1, 2, 3$.

Model	$t = 1$		$t = 2$		$t = 3$		Path Action
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	
Var-RUOT (L_2 Loss)	0.2540 ± 0.0016	0.2623 ± 0.0017	0.2670 ± 0.0013	0.2756 ± 0.0014	0.2683 ± 0.0014	0.2796 ± 0.0015	0.3544 ± 0.0019
Var-RUOT (L_1 Loss)	0.2631 ± 0.0016	0.2717 ± 0.0014	0.2796 ± 0.0013	0.2886 ± 0.0014	0.2809 ± 0.0025	0.2916 ± 0.0025	0.3528 ± 0.0016

C.3 Hold-one-out Experiment

Table 13: Performance comparison of Var-RUOT with L_2 and L_1 loss functions on 2D Mouse Blood Hematopoiesis Dataset for time steps $t = 1, 2$.

Model	$t = 1$		$t = 2$		Path Action
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	
Var-RUOT (L_2 Loss)	0.1200 ± 0.0038	0.1459 ± 0.0038	0.1431 ± 0.0092	0.1764 ± 0.0135	3.1491 ± 0.0837
Var-RUOT (L_1 Loss)	0.1747 ± 0.0095	0.2097 ± 0.0172	0.1690 ± 0.0156	0.2102 ± 0.0242	3.1889 ± 0.0946

In order to validate whether our algorithm can learn the correct dynamical equations from a limited set of snapshot data, we conducted hold-one-out experiments on the three-gene simulated data, the EMT data, and the 2D Mouse Blood Hematopoiesis data. This experiment is designed to test the interpolation and extrapolation capabilities of the algorithm. For a dataset with n time points, we perform n experiments. In each experiment, one time point is removed from the n time points, and the model is trained using the remaining time points. Afterwards, we compute the \mathcal{W}_1 and \mathcal{W}_2 distances between the predicted distribution and the true distribution at the missing time point. When a time point from $\{1, 2, \dots, n-1\}$ is removed, the model is performing interpolation; when the time point n is removed, the model is performing extrapolation. The results of these experiments are shown in Table 14, Table 15, and Table 16. The experimental results indicate that our model’s interpolation performance is superior to that of TIGON and comparable to that of DeepRUOT: in the EMT data and the Mouse Blood Hematopoiesis data, our model achieves the superior extrapolation performance. To further assess the long-term extrapolation capabilities of Var-RUOT and investigate whether minor errors from early time points are amplified over the longer time period, we conducted an additional extrapolation experiment. For this, we extended the three-gene simulation dataset to nine time points. The initial five time points ($t = 0, 1, 2, 3, 4$) were designated for training, while the final four ($t = 5, 6, 7, 8$) were reserved for evaluating extrapolation performance. We compared the performance of TIGON, DeepRUOT, and Var-RUOT, with the results presented in Table 17. These results indicate that Var-RUOT maintains performance superior to TIGON and comparable to DeepRUOT, even in this long-term extrapolation setting. A degradation in performance over longer extrapolation horizons is expected. This is primarily because we employ a non-autonomous system, where the parameter λ is a function of both the state \mathbf{x} and time t . Consequently, the model learns the mapping $\lambda(\mathbf{x}, t)$ only within the temporal domain of the training data, and its behavior beyond this range cannot be accurately inferred.

From a physical viewpoint, the dynamical equations governing the biological processes of cells can be formulated in the form of a minimum action principle (in this work, ITI RUOT Problem is a surrogate model whose action is not the true action derived from studying the biological process, but rather a simple and numerically convenient form of action). Compared to other algorithms, our method can find trajectories with lower action, i.e., it is more capable of learning dynamics that conform to the prior prescribed by the action functional. These dynamics yield better extrapolation performance, which indicates that the design of the action in the RUOT problem is at least partially reasonable. From a machine learning perspective, forcing the model to learn trajectories corresponding to the minimum action serves as a form of regularization that enhances the model’s generalization capability.

In addition, we separately illustrate the learned trajectories and growth profiles on the three-gene simulated dataset after removing four different time points, as shown in Fig. 7 and Fig. 8, respectively. The consistency in the learned results indirectly demonstrates that the model is still able to learn the correct dynamics and perform effective interpolation and extrapolation, even when snapshots at certain time points are missing. We further illustrate the interpolated and extrapolated trajectories of both the DeepRUOT and Var-RUOT algorithms on the Mouse Blood Hematopoiesis dataset, as shown in Fig. 9 and Fig. 10, respectively. This dataset comprises only three time points, $t = 0, 1, 2$. When one time point is removed, Var-RUOT tends to favor a straight-line trajectory connecting the remaining two time points (since such a trajectory represents the minimum-action path), which serves as an effective prior and leads to a reasonably accurate interpolation. In contrast, because DeepRUOT does not explicitly incorporate the minimum-action objective into its model, the trajectories it learns tend to be more intricate and curved. These more complex trajectories might present challenges for generalization, making accurate interpolation or extrapolation more difficult.




Figure 7: Trajectories learned on the three-gene simulated dataset after individually removing $t = 1, 2, 3, 4$.




Figure 8: Growth learned on the three-gene simulated dataset after individually removing $t = 1, 2, 3, 4$.

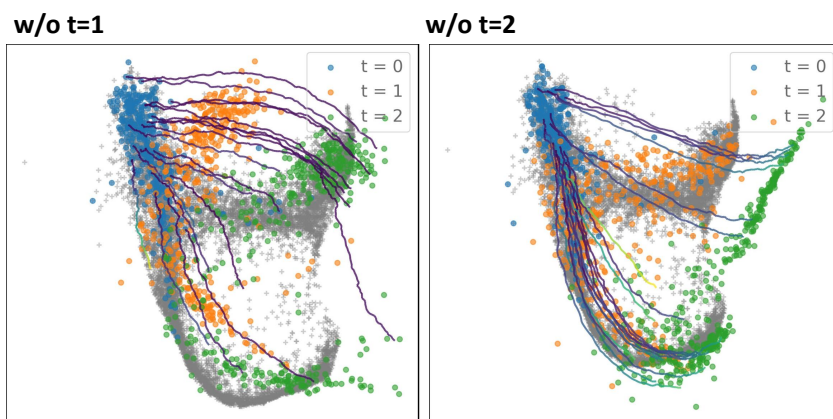


Figure 9: The results of the DeepRUOT algorithm on the 2D Mouse Blood Hematopoiesis dataset for interpolation (with $t=1$ removed) and extrapolation (with $t=2$ removed)

Table 14: On the three-gene simulated dataset, after removing the data of each time point in turn and training on the remaining data, the Wasserstein distances (i.e., \mathcal{W}_1 and \mathcal{W}_2 distances) between the model-predicted data for the missing time points and the ground truth are computed.

Model	$t = 1$		$t = 2$		$t = 3$		$t = 4$	
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2
TIGON	0.1205 \pm 0.0000	0.1679 \pm 0.0000	0.0931 \pm 0.0000	0.1919 \pm 0.0000	0.2390 \pm 0.0000	0.3369 \pm 0.0000	0.2403 \pm 0.0000	0.3616 \pm 0.0000
RUOT	0.0960 \pm 0.0027	0.1505 \pm 0.0018	0.0887 \pm 0.0069	0.1501 \pm 0.0062	0.1184 \pm 0.0058	0.1704 \pm 0.0079	0.1428 \pm 0.0062	0.2179 \pm 0.0135
Var-RUOT(Ours)	0.0880 \pm 0.0036	0.1210 \pm 0.0066	0.1043 \pm 0.0035	0.2293 \pm 0.0045	0.0943 \pm 0.0029	0.1769 \pm 0.0092	0.1401 \pm 0.0047	0.3382 \pm 0.0045

Table 15: On the EMT dataset, after removing the data of each time point in turn and training on the remaining data, the Wasserstein distances (i.e., \mathcal{W}_1 and \mathcal{W}_2 distances) between the model-predicted data for the missing time points and the ground truth are computed.

Model	$t = 1$		$t = 2$		$t = 3$	
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2
TIGON	0.3457 \pm 0.0000	0.3560 \pm 0.0000	0.3733 \pm 0.0000	0.3849 \pm 0.0000	0.5260 \pm 0.0000	0.5424 \pm 0.0000
RUOT	0.3107 \pm 0.0017	0.3201 \pm 0.0016	0.3344 \pm 0.0024	0.3445 \pm 0.0021	0.4947 \pm 0.0019	0.5074 \pm 0.0019
Var-RUOT(Ours)	0.3018 \pm 0.0030	0.3104 \pm 0.0031	0.3375 \pm 0.0027	0.3460 \pm 0.0028	0.4082 \pm 0.0027	0.4189 \pm 0.0027

Table 16: On the Mouse Blood Hematopoiesis dataset, after removing the data of each time point in turn and training on the remaining data, the Wasserstein distances (i.e., \mathcal{W}_1 and \mathcal{W}_2 distances) between the model-predicted data for the missing time points and the ground truth are computed.

Model	$t = 1$		$t = 2$	
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2
TIGON	0.5838 \pm 0.0000	0.6726 \pm 0.0000	1.3264 \pm 0.0000	1.3928 \pm 0.0000
RUOT	0.6235 \pm 0.0014	0.6971 \pm 0.0012	1.0723 \pm 0.0096	1.1397 \pm 0.0120
Var-RUOT(Ours)	0.2696 \pm 0.0054	0.3279 \pm 0.0044	0.2594 \pm 0.0069	0.3016 \pm 0.0095

Table 17: Long-term extrapolation performance comparison on the three-gene simulation dataset. All models were trained on three-gene simulation data from the first five time points ($t = 0, 1, 2, 3, 4$) and evaluated on four subsequent, unseen time points ($t = 5, 6, 7, 8$). The table reports performance metrics (e.g., \mathcal{W}_1 and \mathcal{W}_2 distances) against the ground truth.

Method	$t = 5$		$t = 6$		$t = 7$		$t = 8$	
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2
TIGON	0.1932 \pm 0.0	0.2496 \pm 0.0	0.3437 \pm 0.0	0.4011 \pm 0.0	0.5209 \pm 0.0	0.5643 \pm 0.0	0.6913 \pm 0.0	0.7155 \pm 0.0
DeepRUOT	0.1027 \pm 0.0048	0.1459 \pm 0.0010	0.1940 \pm 0.0069	0.2192 \pm 0.0055	0.3359 \pm 0.0088	0.3644 \pm 0.0074	0.5047 \pm 0.0076	0.5362 \pm 0.0059
VarRUOT	0.1057 \pm 0.0175	0.2191 \pm 0.0067	0.1806 \pm 0.0223	0.2604 \pm 0.0073	0.2890 \pm 0.0032	0.3787 \pm 0.0083	0.4260 \pm 0.0032	0.5557 \pm 0.0067

C.4 Experiments on High Dimensional Dataset

High Dimensional Gaussian Dataset To evaluate the effectiveness of our method on high-dimensional datasets, we first tested it on Gaussian datasets of 50 and 100 dimensions. We learned the dynamics of the data using the standard WFR metric ($\psi(g) = \frac{1}{2}g^2$) as well as a modified growth penalty function, $\psi(g) = g^{2/15}$, which ensures $\psi''(g) < 0$. The learned trajectories and growth rates are illustrated in Fig. 11. Under both choices of $\psi(g)$, our method captures reasonable dynamics: the Gaussian distribution centered on the left shifts upward and downward, while the Gaussian distribution on the right exhibits growth without displacement. We also tested the performance of VarRUOT and other algorithms on a 150-dimensional Gaussian Mixture dataset, as shown in Table 18.

50D Mouse Blood Hematopoiesis and Pancreatic β -cell Differentiation Dataset We tested our method on two high-dimensional real scRNA-seq datasets including 50D Mouse Blood Hematopoiesis dataset and Pancreatic β -cell Differentiation dataset. We used UMAP to reduce the dimensionality of the datasets to 2 (only for visualization), plotted the growth of each data point, and visualized the vector fields $\mathbf{u}(\mathbf{x}, t)$ on the reduced dimensions using the scvelo library. The results for the two datasets are shown in Fig. 12 and Fig. 13, respectively. As can be seen from the figures, the reduced

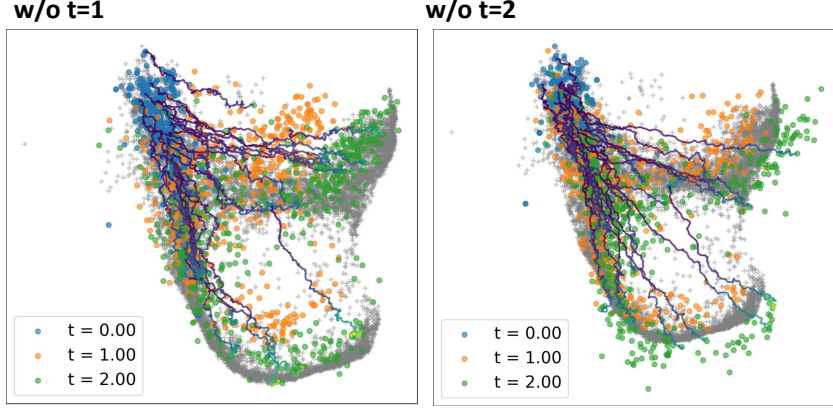


Figure 10: The results of the Var-RUOT algorithm on the 2D Mouse Blood Hematopoiesis dataset for interpolation (with $t=1$ removed) and extrapolation (with $t=2$ removed)

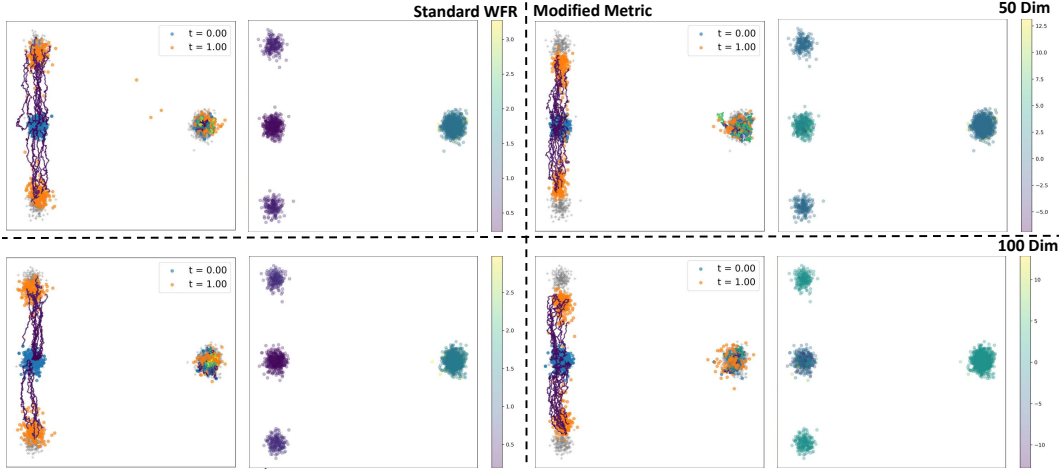


Figure 11: Trajectories and growth learned on the 50-dimensional and 100-dimensional Gaussian datasets using the Standard WFR Metric and Modified Metric.

Table 18: Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between the predicted distributions of each algorithm and the true distribution at various time points on 150D Gaussian Data. Each experiment was run five times to compute the mean and standard deviation.

Model	$t = 1$		$t = 2$		Path Action
	\mathcal{W}_1	\mathcal{W}_2	\mathcal{W}_1	\mathcal{W}_2	
SF2M	7.286 ± 0.002	7.345 ± 0.002	9.457 ± 0.001	9.558 ± 0.002	
PISDE	6.051 ± 0.002	6.061 ± 0.002	7.546 ± 0.001	7.597 ± 0.001	
MIO Flow	6.494 ± 0.000	6.521 ± 0.000	7.812 ± 0.000	7.866 ± 0.000	
TIGON	6.157 ± 0.000	6.169 ± 0.000	7.615 ± 0.000	7.665 ± 0.000	13.4913
DeepRUOT	6.202 ± 0.001	6.213 ± 0.001	7.648 ± 0.003	7.696 ± 0.003	13.5992
Var-RUOT	6.072 ± 0.008	6.089 ± 0.009	7.544 ± 0.005	7.596 ± 0.005	10.1033 ± 0.5677

velocity field points from points with smaller t to those with larger t , which indicates that our model can correctly learn a vector field that transfers the distribution even in high-dimensional data.

EB Dataset and NeurIPS Challenge Dataset

To further quantitatively evaluate the performance of VarRUOT on high-dimensional data, we employed the 100-dimensional EB dataset and the 50-dimensional NeurIPS Challenge dataset. The

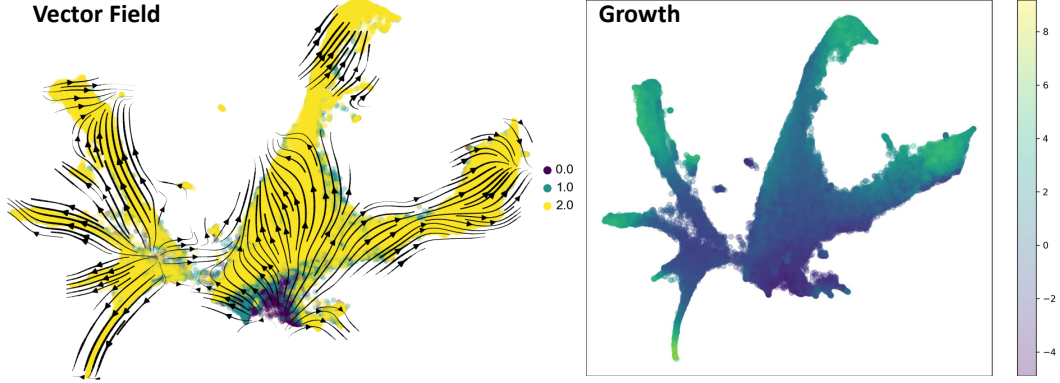


Figure 12: Learned vector field $\mathbf{u}(\mathbf{x}, t)$ and growth on the 50D Mouse Blood Hematopoiesis dataset (data reduced using UMAP and vector field stream plots generated with the scvelo library).




Figure 13: Learned vector field $\mathbf{u}(\mathbf{x}, t)$ and growth on the Pancreatic β -cell Differentiation dataset (data reduced using UMAP and vector field stream plots generated with the scvelo library).

quantitative results for VarRUOT against other algorithms are presented in Table 19 and Table 20, respectively. On the EB dataset, Var-RUOT demonstrates distribution reconstruction performance comparable to that of TIGON and DeepRUOT, while surpassing simulation-free approaches (OTCFM, UOTCFM) and Wasserstein Gradient Flow. For the NeurIPS Challenge dataset, Var-RUOT outperforms all other prior methods evaluated. We also report the results for VarRUOT and competing algorithms on the unnormalized 5-dimensional EB dataset in Table 21.

Table 19: Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between the predicted distributions of each algorithm and the true distribution at various time points on EB Data. Each experiment was run five times to compute the mean and standard deviation.

Method	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
SF2M	9.6732 ± 0.0006	9.7446 ± 0.0005	11.0711 ± 0.0006	11.1649 ± 0.0024	12.5436 ± 0.0041	12.6476 ± 0.0047	14.7100 ± 0.0077	14.8127 ± 0.0006
PISDE	8.8317 ± 0.0007	8.9075 ± 0.0008	9.4283 ± 0.0015	9.5312 ± 0.0016	9.7872 ± 0.0013	9.8952 ± 0.0016	11.1102 ± 0.0018	11.2071 ± 0.0022
MIOFLOW	8.7182 ± 0.0000	8.7924 ± 0.0000	9.4139 ± 0.0000	9.5180 ± 0.0000	9.7547 ± 0.0000	9.8671 ± 0.0000	11.0080 ± 0.0000	11.1065 ± 0.0000
TIGON	8.7992 ± 0.0000	8.8769 ± 0.0000	9.6497 ± 0.0000	9.7533 ± 0.0000	10.0130 ± 0.0000	10.1209 ± 0.0000	11.3284 ± 0.0000	11.2452 ± 0.0000
RUOT	8.8029 ± 0.0023	8.8803 ± 0.0022	9.6518 ± 0.0071	9.7555 ± 0.0065	10.0365 ± 0.0126	10.1424 ± 0.0118	11.3555 ± 0.0133	11.4501 ± 0.0132
OTCFM	9.4459 ± 0.0000	9.5280 ± 0.0000	11.3786 ± 0.0000	11.5893 ± 0.0000	14.0529 ± 0.0000	14.9498 ± 0.0000	19.7489 ± 0.0000	19.7489 ± 0.0000
UOTCFM	9.4794 ± 0.0000	9.5596 ± 0.0000	11.6060 ± 0.0000	11.8336 ± 0.0000	14.6715 ± 0.0000	15.6972 ± 0.0000	21.7653 ± 0.0000	25.5683 ± 0.0000
Action Matching	12.7145 ± 0.0000	12.9813 ± 0.0000	16.7809 ± 0.0000	16.8206 ± 0.0000	13.4663 ± 0.0000	13.7058 ± 0.0000	18.1165 ± 0.0000	18.1657 ± 0.0000
WGF	9.9847 ± 0.0000	10.0482 ± 0.0000	11.0562 ± 0.0000	11.1344 ± 0.0000	11.9719 ± 0.0000	12.0508 ± 0.0000	13.2417 ± 0.0000	13.3035 ± 0.0000
Var-RUOT	8.9158 ± 0.0026	8.9996 ± 0.0006	9.7955 ± 0.0203	9.9651 ± 0.0245	10.2647 ± 0.0214	10.4849 ± 0.0316	11.4927 ± 0.0585	11.6323 ± 0.0629

Table 20: Wasserstein distances (\mathcal{W}_1 and \mathcal{W}_2) between the predicted distributions of each algorithm and the true distribution at various time points on NeurIPS Challenge Data. Each experiment was run five times to compute the mean and standard deviation.

Method	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
SF2M	6.1876 \pm 0.0085	6.2565 \pm 0.0081	12.6170 \pm 0.0602	13.6426 \pm 0.0707	9.1830 \pm 0.0041	10.9930 \pm 0.1741	12.5831 \pm 0.6213	13.2691 \pm 0.4054
PISDE	5.6642 \pm 0.0078	5.7069 \pm 0.0079	5.2358 \pm 0.0073	5.2779 \pm 0.0077	6.2793 \pm 0.0134	6.4015 \pm 0.0134	8.9086 \pm 0.0053	9.0298 \pm 0.0060
MIOFLOW	6.4267 \pm 0.0000	6.4817 \pm 0.0000	6.2715 \pm 0.0000	6.3141 \pm 0.0000	7.1079 \pm 0.0000	7.2222 \pm 0.0000	8.9035 \pm 0.0000	9.0160 \pm 0.0000
TIGON	5.9481 \pm 0.0000	6.0003 \pm 0.0000	5.7599 \pm 0.0000	5.8008 \pm 0.0000	6.5539 \pm 0.0000	6.6787 \pm 0.0000	8.9573 \pm 0.0000	9.0587 \pm 0.0000
RUOT	6.0166 \pm 0.0020	6.0579 \pm 0.0018	5.6972 \pm 0.0048	5.7415 \pm 0.0050	6.5099 \pm 0.0034	6.6419 \pm 0.0062	8.8047 \pm 0.0069	8.9103 \pm 0.0069
OTCFM	5.5082 \pm 0.0000	5.5645 \pm 0.0000	5.4476 \pm 0.0000	5.5151 \pm 0.0000	7.0935 \pm 0.0000	7.3473 \pm 0.0000	10.1416 \pm 0.0000	10.7654 \pm 0.0000
UOTCFM	5.5526 \pm 0.0000	5.6089 \pm 0.0000	5.4349 \pm 0.0000	5.5086 \pm 0.0000	6.7482 \pm 0.0000	6.6849 \pm 0.0000	9.6782 \pm 0.0000	10.3230 \pm 0.0000
Action Matching	7.3701 \pm 0.0000	7.4312 \pm 0.0000	8.5156 \pm 0.0000	8.5485 \pm 0.0000	9.2842 \pm 0.0000	9.3269 \pm 0.0000	9.9665 \pm 0.0000	10.0428 \pm 0.0000
WGF	7.1785 \pm 0.0000	7.2411 \pm 0.0000	7.7135 \pm 0.0000	7.7479 \pm 0.0000	8.2521 \pm 0.0000	8.3044 \pm 0.0000	8.6629 \pm 0.0000	8.7749 \pm 0.0000
Var-RUOT	4.7961 \pm 0.0763	4.9015 \pm 0.0844	4.1816 \pm 0.0145	4.2887 \pm 0.0151	6.2312 \pm 0.0130	6.4587 \pm 0.0144	9.1058 \pm 0.0511	9.3636 \pm 0.0537

Table 21: Comparison of Wasserstein-1 distances (\mathcal{W}_1) between the predicted distributions and the ground truth on the unnormalized 5-dimensional EB dataset.

Model	Wasserstein-1 Distance (\mathcal{W}_1) at time t			
	$t = 1$	$t = 2$	$t = 3$	$t = 4$
MMFM	0.477	0.554	0.781	0.872
Metric FM	0.449	0.552	0.583	0.597
SF2M	0.556	0.715	0.750	0.650
MIOFlow	0.442	0.585	0.651	0.670
TIGON	0.386	0.502	0.602	0.600
DeepRUOT	0.386	0.497	0.591	0.585
UOT-FM	0.544	0.670	0.729	0.852
VARRUOT	0.416	0.486	0.509	0.511

C.5 Comparison with Traditional OT Solvers

To verify whether the Var-RUOT algorithm can find the correct minimum action paths for both Optimal Transport (OT) and Unbalanced Optimal Transport (UOT) problems, we constructed two normalized (balanced) and two unbalanced Gaussian Mixture datasets on a 2D plane. We then compared the action computed by the Var-RUOT method (with $\sigma = 0$ to eliminate stochasticity) against the minimum action values calculated by traditional solvers.

To compute the ground truth for the minimum action on the normalized datasets, we used the `pot` library. For the unbalanced datasets, we could not find an open-source solver for dynamic unbalanced optimal transport (apart from deep learning-based methods like TIGON). Therefore, we adopted a static Wasserstein-Fisher-Rao (WFR) distance solver, designed as a variant of the Sinkhorn algorithm as described in (Wang et al. 2019), to serve as our baseline. The comparison results are presented in Table 22.

Table 22: Comparison of Path Action values on 2D Gaussian Mixture datasets.

Dataset	Normalized		Unbalanced	
	Single Gaussian	Gaussian Mixture	Single Gaussian	Gaussian Mixture
Path Action (Traditional Solver)	0.5046	1.9062	1.3776	2.4464
Path Action (Var-RUOT)	0.4998	1.8754	1.0102	2.3909

The results show that for the two normalized datasets, Var-RUOT finds Path Action values comparable to those computed by the `pot` library. On the two unbalanced datasets, the Path Action found by Var-RUOT is even lower than that of the traditional solver. This outcome might be attributable to either (1) suboptimal solution of the baseline solver or (2) numerical roundoff errors. Nevertheless, our review of the literature did not identify alternative open-source baselines addressing this specific problem. Overall, these results demonstrate that Var-RUOT can achieve performance comparable to traditional OT solvers or WFR distance solvers for the task of finding the minimum action path, which validates its effectiveness.

D Further Discussions

D.1 Comparison with Relevant Algorithms

Several methods have been proposed to address the problem of trajectory inference for temporal snapshots of single-cell data. We list the differences between our method and previous trajectory inference algorithms in Table 23.

Table 23: Comparison of trajectory inference methods. Properties assessed are: handling **Unbalanced Distributions** (unequal mass); modeling **Stochastic Dynamics**; addressing both **Unbalanced + Stochastic** properties simultaneously; and formulation based on the **Least Action** principle via first-order conditions, not just as a loss penalty. The ‘+’ indicates the method possesses the property, while a ‘−’ indicates it does not.

Method	Unbalanced Distribution	Stochastic Dynamics	Unbalanced + Stochastic	Least Action
SF2M	−	+	−	−
PISDE	−	+	−	+
MIOFlow	−	−	−	−
Action Matching	+	+	−	+
TIGON	+	−	−	−
DeepRUOT	+	+	+	−
Var-RUOT (Ours)	+	+	+	+

Difference from DeepRUOT

Like DeepRUOT, Var-RUOT addresses the critical challenge of handling both unbalanced distributions and stochastic dynamics, which is essential for single-cell omics data.

The primary difference is architectural. DeepRUOT employs three interdependent neural networks to model the velocity field $\mathbf{v}(\mathbf{x}, t)$, growth rate $g(\mathbf{x}, t)$, and log-density $\log \rho(\mathbf{x}, t)$. It enforces physical constraints as soft penalties in the loss function. This complex setup necessitates a multi-stage pre-training procedure to ensure convergence.

In contrast, Var-RUOT uses a single network to parameterize a scalar field $\lambda_\theta(\mathbf{x}, t)$. Both the velocity $\mathbf{u}(\mathbf{x}, t)$ and growth rate $g(\mathbf{x}, t)$ are analytically derived from $\lambda_\theta(\mathbf{x}, t)$ via first-order optimality conditions (Theorem 4.1). This simpler, end-to-end design eliminates the need for pre-training and improves training stability.

Difference from Action Matching

The differences between Var-RUOT and Action Matching could be summarized from three perspectives. Firstly, Var-RUOT handles both unbalanced distributions and stochastic dynamics, which is crucial for single-cell omics.

Secondly, Action Matching’s training requires sampling from intermediate distributions $q_t(x)$. These are typically challenging to obtain in single-cell omics, where data exists only as sparse time-point snapshots. This data limitation can potentially hinder the ability to capture complex dynamics, whereas Var-RUOT is designed for such snapshot data, and empirical results suggest Var-RUOT’s desirable performance.

Lastly, the methods employ different loss formulations. Action Matching minimizes an Action Gap loss:

$$\mathcal{L}_{\text{AM}} = \mathbb{E}_{q_0(x)}[s_0(x)] - \mathbb{E}_{q_1(x)}[s_1(x)] + \int_0^1 \mathbb{E}_{q_t(x)} \left[\frac{1}{2} \|\nabla s_t(x)\|^2 + \frac{\partial s}{\partial t}(x) \right] dt.$$

On the other hand, Var-RUOT directly minimizes the residual of the Hamilton-Jacobi-Bellman (HJB) equation. By directly enforcing this first-order optimality condition in its loss, Var-RUOT can find solutions with low action robustly in empirical tests.

D.2 Convergence of the Particle Method and Loss Computation

The VarRUOT algorithm involves three loss terms: \mathcal{L}_{HJB} , $\mathcal{L}_{\text{Action}}$, and $\mathcal{L}_{\text{Recon}}$. Among these, \mathcal{L}_{HJB} and $\mathcal{L}_{\text{Action}}$ are computed using the Monte Carlo method. This involves sampling a set of particles,

integrating along their trajectories, and then averaging the results over all particles. The convergence rate of the Monte Carlo method is $\mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$, where N is the number of particles used.

The reconstruction loss, $\mathcal{L}_{\text{Recon}}$, is composed of two parts: \mathcal{L}_{OT} and $\mathcal{L}_{\text{Mass}}$. \mathcal{L}_{OT} , calculated using the `geomLoss` library, represents the \mathcal{W}_2 distance between the normalized empirical distribution and the true distribution. According to the triangle inequality, $\mathcal{W}_2(\mu^N, \nu) \leq \mathcal{W}_2(\mu^N, \mu) + \mathcal{W}_2(\mu, \nu)$, the convergence of \mathcal{L}_{OT} as the number of particles increases depends on the rate at which the empirical measure $\mu^N = \frac{1}{N} \sum_{i=1}^N \delta(\mathbf{x} - \mathbf{x}_i)$ (with $\mathbf{x}_i \sim \mu$) approximates the true measure μ . The expected error is bounded as follows (Fournier and Guillin 2013):

$$\mathbb{E}[\mathcal{W}_2(\mu^N, \mu)] \leq C \begin{cases} N^{-\frac{1}{2}} & \text{if } d = 2, 3 \\ N^{-\frac{1}{2}} \log(1 + N) & \text{if } d = 4 \\ N^{-\frac{2}{d}} & \text{if } d > 4 \end{cases}$$

E Limitations and Broader Impacts

E.1 Limitations

The algorithm presented in this paper offers new insights for solving the RUOT problem; however, it still has several limitations. Firstly, although Var-RUOT parameterizes \mathbf{u} and g with a single neural network, and designs the loss function based on the necessary conditions for the minimal action solution, since neural network optimization only finds local minima, there is still no guarantee that the solution found is indeed the one with minimal action. Moreover, in practice, the HJB loss does not necessarily converge to zero (A typical training curve for the HJB loss is shown in Fig. 14). Thus, it effectively acts as a regularization term. This could be addressed by conducting a more detailed analysis on simpler versions of the RUOT problem (for instance, Gaussian-to-Gaussian transport or Dirac-to-Dirac Transport).

Furthermore, when using the modified metric, the goodness-of-fit in the distribution deteriorates, which may suggest that the \mathbf{u} and g satisfying the optimal necessary conditions derived via the variational method are limited in transporting the initial distribution to the terminal distribution. This might reflect a controllability issue in control theory that warrants further investigation.

Finally, the choice of $\psi(g)$ in the action is dependent on biological priors. To automate it, one could approximate ψ with a neural network or derive it from microscopic or mesoscopic dynamics, such as the branching Wiener process to model cell division for a more physically grounded action (Baradat and Lavenant 2021).

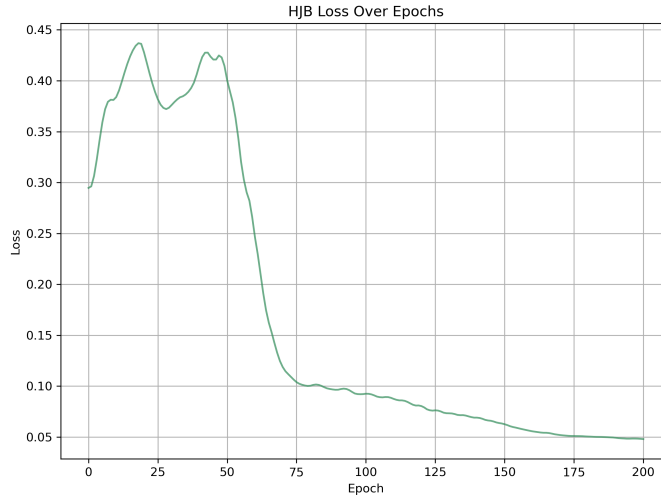


Figure 14: Training curve of HJB Loss on three-gene simulation dataset.

E.2 Broader Impacts

Var-RUOT explicitly incorporates the first-order optimality conditions of the RUOT problem into both the parameterization process and the loss function. This approach enables our algorithm to find solutions with a smaller action while maintaining excellent distribution fitting accuracy. Compared to previous methods, Var-RUOT employs only a single network to approximate a scalar field, which results in a faster and more stable training process. Additionally, we observe that the selection of the growth penalty function $\psi(g)$ within the WFR metric is highly correlated with the underlying biological priors. Consequently, our new algorithm provides a novel perspective on the RUOT problem.

Our approach can be extended to other analogous systems. For example, in the case of simple mesoscopic particle systems where the action can be explicitly formulated such as in diffusion or chemical reaction processes, our framework can effectively infer the evolution of particle trajectories and distributions. This capability makes it applicable to tasks such as experimental data processing and interpolation. In the biological or medical field, our method can be employed to predict cellular fate and to provide quantitative diagnostic results or treatment plans for certain diseases.

It should be noted that the performance of Var-RUOT largely depends on the quality of the data. Datasets containing significant noise may lead the model to produce results with a slight bias. Moreover, the particular form of the action can have a substantial impact on the model’s outcomes, potentially affecting important biological priors. These factors could present challenges for subsequent biological analyses or clinical decision-making, and care must be taken in the use and dissemination of the model-generated interpolation results to avoid data contamination.

When applying our method in biological or medical contexts, it is crucial to train the model using high-quality experimental data, select an action formulation that is well-aligned with the relevant domain-specific priors, and ensure that the results are validated by domain experts. Furthermore, there is a need to enhance the interpretability of the model and to further improve training speed through methods such as simulation-free techniques. These directions represent important avenues for our future work.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction accurately reflect our claim. Our [Appendix A.1](#) provides the proof of the important theorem of the claim; our numerical results are presented in [Section 6](#).

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss our limitations in [Section 7](#) and [Appendix E.1](#).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide our proof for our theorems in Appendix A.1.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide our training procedure in Appendix A.3, and we provide our code in a public repository <https://github.com/ZerooVector/VarRUOT>.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We use public datasets for our experiment, and we provide our code in a public repository <https://github.com/ZerooVector/VarRUOT>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide our experimental details in Appendix B and Appendix C. In Appendix C, we also provide some additional results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide the mean and standard deviation values of our metrics in each experiment.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the description of compute resources in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have all reviewed and confirmed this research conforms with the NeurIPS Code of Ethics. We have striven to maintain and preserve anonymity.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss broader impacts of our work in Appendix E.2.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All datasets are publicly available, and we provide details of datasets in Appendix B.1.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new datasets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The method development does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.