
Learning Velocity Prior-Guided Hamiltonian-Jacobi Flows with Unbalanced Optimal Transport

Amy Xiang Wang
Institut Imagine,
Paris, France
xw914@nyu.edu

Abstract

The connection between optimal transport (OT) and control theory is well established, most prominently in the Benamou–Brenier dynamic formulation. With quadratic cost, the OT problem can be reframed as a stochastic control problem in which a density ρ_t evolves under a controlled velocity field v_t subject to the continuity equation $\partial_t \rho_t + \nabla \cdot (\rho_t v_t) = 0$. In this work, we introduce a velocity prior into the continuity equation and derive a new Hamilton–Jacobi–Bellman (HJB) formulation to learn dynamical probability flows. We further extend the approach to the unbalanced setting by adding a growth term, capturing mass variation processes common in scientific domains such as cell proliferation and differentiation. Importantly, our method requires training only a single neural network to model v_t , without the need for a separate model for the growth term g_t . Finally, by decomposing the velocity field as $v_{\text{total}} = v_{\text{prior}} + v_{\text{corr}}$, our approach is able to capture complex transport patterns that prior methods struggle to learn due to the curl-free limitation.

1 Introduction

From flow matching (FM) to action matching (AM), learning transport maps between distributions has been widely explored in recent years [Lipman et al., 2022, Albergo and Vanden-Eijnden, 2022, Liu et al., 2022, Neklyudov et al., 2023a]. *Flow Matching (FM)* [Lipman et al., 2022] learns a time-dependent velocity field u_t that pushes ρ_0 to ρ_1 and can realize highly expressive transport paths; however, the original FM with independent coupling between source and target does not guarantee *least action* by minimizing the kinetic energy in the Benamou–Brenier sense. Instead, it trains u_t to match conditional expectations of displacement vectors under a chosen interpolation scheme, which may yield non-optimal flows.

Action Matching (AM) [Neklyudov et al., 2023a] addresses this by parameterizing a scalar potential s_t whose gradient ∇s_t induces the transport, aligning with the optimality conditions of OT and yielding lower kinetic energy than unconstrained FM. The price is reduced expressiveness: ∇s_t is *curl-free*, so AM cannot directly represent rotational or cyclic dynamics that are common in scientific domains. From the Helmholtz decomposition perspective [Neklyudov et al., 2023a], any vector field u_t^* can be written as $u_t^* = \nabla s_t^* + w_t$ with w_t divergence-free [Ambrosio et al., 2005, §8.4.2]. Under this lens, AM retains only the gradient component and discards w_t , explaining both its energy efficiency and its inability to encode rotations and cycles.

In this paper, we seek a middle ground – expressive like FM, energy-aware like AM – by introducing a velocity prior v_{prior} and learning only the residual potential. We note that even compared with energy-aware FM variants such as OT–CFM [Pooladian et al., 2023, Tong et al., 2023a], our approach achieves better energy efficiency, as demonstrated in upper Table 1. Specifically, we decompose the velocity

Table 1: Least Action Comparison for Balanced Gaussian Translation

Method	Mean error	Cov. error	W_2	Control action (∇s)	Total kinetic (v_{total})
Flow Matching	0.204	0.804	0.582	18.369	18.369
OT-FM	0.149	0.659	0.402	18.707	18.707
VP-HJF (Ours)	0.102	0.791	0.577	0.624	17.955
Prior-only ($\alpha=1$)	0.008	0.171	1.351	0	36.250

field as $v_{\text{total}}(t, x) = v_{\text{prior}}(t, x) + \nabla s_t(x)$. Here v_{prior} captures known rotational dynamics or domain-specific effects such as RNA velocity in single-cell biology, while ∇s_t accounts for the OT-consistent gradient component. We train s_t by minimizing a modified Hamilton–Jacobi residual that incorporates the prior, together with boundary terms that ensure $\rho_0 \rightarrow \rho_1$. This *residualized* design preserves OT optimality conditions for the learned component, improves interpretability, and injects inductive bias without paying the kinetic-energy cost of unconstrained original FM. We name our approach the Velocity Prior Hamiltonian–Jacobi Flow (VP-HJF).

2 Background

Dynamical Optimal Transport Beyond the classic static Monge–Kantorovich formulation in OT [Ambrosio et al., 2005, Villani et al., 2008], there exists a dynamical formulation known as the Benamou–Brenier problem which links OT with PDEs by representing the W_2 distance as the minimum kinetic energy where ρ_t is density and v_t is a velocity field with boundary conditions: $\rho|_{t=0} = \rho_0$, $\rho|_{t=1} = \rho_1$, [Benamou and Brenier, 2000]:

$$W_2^2(\rho_0, \rho_1) = \inf_{\rho_t, v_t} \int_0^1 \int \frac{1}{2} \|v_t(x)\|^2 \rho_t(x) dx dt, \quad \partial_t \rho_t + \nabla \cdot (\rho_t v_t) = 0. \quad (1)$$

Unbalanced Optimal Transport When total mass change over time such as following a growth-decay process in biology, we add a growth rate $g_t(x)$ term to the continuity equation to incorporate the weight changes [Chizat et al., 2018]:

$$\partial_t \rho_t(x) + \nabla \cdot (\rho_t(x) v_t(x)) = g_t(x) \rho_t(x), \quad \rho|_{t=0} = \rho_0, \quad \rho|_{t=1} = \rho_1, \quad (2)$$

The Wasserstein–Fisher–Rao distance with scale $\delta > 0$ is defined as the minimal action balancing transport cost and mass change: (Additional background see Appendix C)

$$\text{WFR}_\delta^2(\rho_0, \rho_1) = \inf_{\rho, v, g} \int_0^1 \int \left(\frac{1}{2} \|v_t(x)\|^2 + \frac{\delta^2}{2} g_t(x)^2 \right) \rho_t(x) dx dt, \quad \text{s.t. Eq. 2.} \quad (3)$$

3 Methodology

We introduce a velocity-prior guided approach, the *Velocity Prior Hamiltonian–Jacobi Flow (VP-HJF)*, to solve the unbalanced optimal transport problem under the Wasserstein–Fisher–Rao (WFR) metric (Eq. 2, 3). In contrast to prior approaches that fit two separate networks—one for transport and one for growth [Zhang et al., 2024, Wang et al., 2025], our method trains a single neural network. Following [Neklyudov et al., 2023a], we can represent both the transport velocity field and the growth term through a single scalar potential.

Proposition 3.1 (Neklyudov et al., 2023a, Prop. 3.3). *Suppose we have a continuous dynamic flow with density ρ_t . Under mild conditions, there exists a unique scalar potential function $\hat{s}_t(x)$ such that the unbalanced continuity equation (2) is satisfied, with the velocity field and growth function given by $v_t^*(x) = \nabla \hat{s}_t(x)$, $g_t^*(x) = \hat{s}_t(x)$.*

Building on Proposition 3.1, we reduce the WFR problem to learning a single model and incorporate problem-specific dynamics through a simple velocity decomposition. Specifically, we decompose the velocity field into two parts: a known velocity prior and a learnable corrective velocity field component: $v_{\text{total}}(t, x) = v_{\text{prior}}(t, x) + v_{\text{corr}}(t, x)$, where v_{prior} encodes domain knowledge (e.g. translations, rotations, RNA velocity), and v_{corr} is the data-driven corrective component. In this way,

the prior captures coarse dynamics while the model focuses on refinements such as correcting the residual transport and learning mass imbalance that the prior cannot explain. In essence, our approach improves interpretability and reduces the learning complexity through adding prior knowledge of the velocity field v_{prior} – leaving the learnable velocity field v_{corr} simpler learning tasks compared with other generative modeling methods of learning the entire velocity field v_{total} .

We define our velocity-prior guided unbalanced OT problem under the least-action principle as:

Definition 3.2. Consider the following least-action objective with $\delta=1$ and subject to the unbalanced velocity-prior guided continuity equation :

$$\mathcal{A}(\rho, v_{\text{corr}}, g) = \int_0^1 \int \left(\frac{1}{2} \|v(t, x)\|^2 + \frac{1}{2} g(t, x)^2 \right) \rho_t(x) dx dt, \quad (4)$$

$$\text{s.t. } \partial_t \rho_t = -\nabla \cdot (\rho_t (v_{\text{prior}} + v_{\text{corr}})) + g_t \rho_t, \quad \rho_{t=0} = \rho_0, \quad \rho_{t=1} = \rho_1. \quad (5)$$

Note that in our method we do *not* optimize over ρ directly. Instead, ρ_t is *induced* by a parametric flow Φ_t^θ via $\dot{x} = v_{\text{prior}}(t, x) + \nabla s_\theta(t, x)$ and defined as $\rho_t^\theta = (\Phi_t^\theta)_\# \rho_0$.

Prior-guided HJB residual Since solving for the minimum-action problem in primal form in Definition 3.2 is intractable, we turn to its dual formulation. The key derivation step is to introduce a scalar potential $s(t, x)$ as the Lagrange multiplier for the prior-guided continuity equation and applying the Fenchel–Young inequalities to the velocity field and growth term. We then obtain the following dual lower bound (see Appendix A for details):

$$\begin{aligned} \mathcal{A}(\rho, v_{\text{corr}}, g) &\geq \mathbb{E}_{\rho_0(x)}[s_0(x)] - \mathbb{E}_{\rho_1(x)}[s_1(x)] \\ &\quad - \int_0^1 \int \rho_t(x) \left(\partial_t s + \frac{1}{2} \|\nabla s\|^2 + \nabla s \cdot v_{\text{prior}} + \frac{1}{2} s^2 \right) dx dt. \end{aligned} \quad (6)$$

The bound is tight point-wise if and only if when we choose the primal variables as

$$v_{\text{corr}}(t, x) = \nabla_x s(t, x), \quad g(t, x) = s(t, x),$$

which shows that s can simultaneously control both the *corrective* transport $\nabla_x s$ and the local growth s . We can then plug these back to the continuity equation to get the optimal particle dynamics and their log-weights evolve as

$$\frac{d}{dt} x(t) = v_{\text{total}}(t, x) = v_{\text{prior}}(t, x) + \nabla_x s(t, x), \quad \frac{d}{dt} \log w_t(x(t)) = s(t, x(t)).$$

Corollary 3.3 (HJB residual objective). *Motivated by the duality form, we propose to parameterize $s_\theta(t, x)$ with a neural network and define the velocity-prior guided HJB residual as:*

$$r_\theta(t, x) := \partial_t s + \frac{1}{2} \|\nabla_x s\|^2 + \nabla_x s \cdot v_{\text{prior}} + \frac{1}{2} s^2. \quad (7)$$

Then minimizing

$$\mathcal{L}_{\text{HJB}}(\theta) = \mathbb{E}_{\rho_0(x)}[s_0(x)] - \mathbb{E}_{\rho_1(x)}[s_1(x)] + \int_0^1 \mathbb{E}_{\rho_t(x)}[w(t, x) r_\theta(t, x)^2] dt \quad (8)$$

drives s_θ toward dual feasibility. Note that in practice, we use squared residual to prevent positive and negative values from cancellation and add a importance weight $w(t)$ trick to reduce variance.

Importance Reweighting The squared HJB residual can be dominated by a few high-variance outliers (rare cells, sharp local flows), which destabilizes training. To ensure training stability and preventing these extreme high residual outliers, we adopt a simple batch-wise importance reweighting that down-weights large residuals. For a mini-batch $\{(t_i, x_i)\}_{i=1}^B$, let $r_i = |r_\theta(t_i, x_{t,i})| + \varepsilon$ and with a temperature $\tau > 0$. Then for each sample, the weight is inversely proportions to a temperature-shaped residual as $\tilde{w}_i \propto r_i^{-\tau}$. Thus, larger residuals get smaller weight, which reduces variance while keeping the update focused and stable.

Theorem 3.4 (Prior-guided HJB optimality). *Suppose that the HJB residual defined in corollary.3.3 satisfies $r_\theta(t, x) = 0$ for ρ_t -a.e on $[0, 1] \times \mathbb{R}^d$, and the boundary constraints hold, then $(\rho_t, v_{\text{corr}}, g_\theta)$ satisfies the unbalanced continuity equation and the WFR optimality conditions in Definition 3.2. In particular, the learned corrective field $v_{\text{corr}}^* = \nabla_x s_\theta$ and growth $g^* = g_\theta$ satisfy the optimality conditions. (See Appendix B for proof).*




Figure 1: **Left** (first two): Ours correctly learned the rotating dynamic while AM failed. **Middle**: Vector field (red) bending away from v_{prior} (blue) to form petal shapes. **Right**: Predicted target distribution matches with ground truth

Table 2: Comparison on the EB dataset using SWD, MMD and W1, at the held-out marginals (t_1, t_3). Baseline results are taken from [Theodoropoulos et al., 2025]

Method	SWD t_1	SWD t_3	MMD t_1	MMD t_3	$W_1 t_1$	$W_1 t_3$
DeepRUOT	0.73	0.67	0.43	0.36	13.45	14.90
MIOFlow	0.84	0.94	1.01	0.92	13.20	13.57
SBIRR	0.80	0.91	0.71	0.73	15.09	20.39
MMFM	0.59	0.76	0.37	0.35	13.61	14.64
DMSB	0.58	0.54	0.38	0.36	14.08	15.22
3MSBM	0.48	0.38	0.14	0.18	13.89	13.11
VP-HJF (ours)	0.37	0.46	0.18	0.18	11.83	14.10

While the HJB residual enforces local optimality conditions, it does not directly align the terminal distribution with the target. To bridge this gap, we design a two-part reconstruction objective to align the learned terminal distribution for *shape* and *mass* with the ground truth distribution, see details in Appendix D. For related works, see Appendix E.

4 Experiments

Balanced case - Gaussian Translation In this experiment, we compare the least-action or energy costs across different methods. We define an affine prior drift as $v_{\text{prior}}(x) = \mu_1 - \mu_0$ with the parameters $\mu_0 = (0, 0)$, $\mu_1 = (0.5, 6.0)$, $\Sigma_0 = ((2.5, 0.0), (0.0, 0.3))$, $\Sigma_1 = ((0.4, 0.0), (0.0, 2.2))$. As shown in Table 1, our proposed VP-HJF achieves Wasserstein distance W_2 accuracy on par with FM Lipman et al. [2022] while OT-FM Tong et al. [2023a], Pooladian et al. [2023] attains the lowest Wasserstein distance. By contrast, both FM and OT-FM must learn the entire velocity field $u_t = \nabla s + w_t$, resulting in a much larger control action. This demonstrates that VP-HJF leverages the structured prior effectively, where the prior dynamics carry most of the transport, and the learned correction ∇s makes adjustments. To verify that our improvement is not solely due to a strong velocity prior v_{prior} itself, we also report a prior-only baseline. Additionally, Figure 1 highlights that incorporating the velocity prior is crucial in the rotational dynamics setting, as well as in a multi-trajectory example where paths diverge into a petal shape. See Appendix F for more.

EB scRNA-Seq data We evaluate cell-trajectory inference on the Embryoid Body (EB) dataset of Moon et al. [2019], using the preprocessed release from Koshizuka and Sato [2022], Tong et al. [2020]. The dataset comprises five snapshots over 27 days, grouped as $t_0 \in [0, 3]$, $t_1 \in [6, 9]$, $t_2 \in [12, 15]$, $t_3 \in [18, 21]$, $t_4 \in [24, 27]$. Leveraging RNA velocity as a prior v_{prior} at each snapshot, we train a local and shorter trajectory by adopting the *local per-interval* training: at each step we sample an adjacent pair (t_k, t_{k+1}) and learn only the transport and growth to move $\rho_{t_k} \rightarrow \rho_{t_{k+1}}$. This yields more stable gradients and low target variance than enforcing all time points jointly;. For details of the training algorithm, see Algorithm 1.

We test on 100-dim PCA components feature space and compare with recent works using the multi-marginal approach from 3MSBM[Theodoropoulos et al., 2025], SBIRR [Shen et al., 2024], MMFM [Rohbeck et al., 2025] and DMSB [Chen et al., 2023]. We follow the experiment setup from 3MSBM by having $t = 1, 3$ as the held out sets. Table 2 shows that our method outperforms others at $t = 1$

and remains competitive at other time snapshots. Notably, our method outperforms SBIRR and MMFM which solve a piecewise Schrödinger Bridges and OT coupling instead of a single global optimization with a joint coupling like 3MSBM and DMSB, indicating the benefits of using RNA velocity as a local prior with per-interval supervision.

Conclusion Our method decomposes the velocity field to embed domain knowledge for better energy efficiency and uses a single network to capture both unbalanced growth and transport. For discussions on limitations, see Appendix G.

Acknowledgments

The author thanks Takeshi Koshizuka for valuable discussions and feedback on the initial draft. We also thank the reviewers for detailed feedback, which helped improve this manuscript. Lastly, the author is grateful to Matt Ricci and Nathan Quiblier for the early-stage discussion.

References

- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *International conference on machine learning*, pages 25858–25889. PMLR, 2023a.
- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer, 2005.
- Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.
- Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023a.
- Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2008.
- Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.
- Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal transport: Dynamic and kantorovich formulations. *Journal of Functional Analysis*, 274(11): 3090–3123, 2018.
- Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through regularized unbalanced optimal transport. *arXiv preprint arXiv:2410.00844*, 2024.
- Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-growth flow matching for single-cell dynamics modeling. *arXiv preprint arXiv:2505.13413*, 2025.
- Panagiotis Theodoropoulos, Augustinos D Saravanos, Evangelos A Theodorou, and Guan-Hong Liu. Momentum multi-marginal schrödinger bridge matching. *arXiv preprint arXiv:2506.10168*, 2025.

- Kevin R Moon, David Van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen, Kristina Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, et al. Visualizing structure and transitions in high-dimensional biological data. *Nature biotechnology*, 37(12): 1482–1492, 2019.
- Takeshi Koshizuka and Issei Sato. Neural lagrangian schrödinger bridge: Diffusion modeling for population dynamics. *arXiv preprint arXiv:2204.04853*, 2022.
- Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. TrajectoryNet: A dynamic optimal transport network for modeling cellular dynamics. In *International conference on machine learning*, pages 9526–9536. PMLR, 2020.
- Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal schrödinger bridges with iterative reference refinement. *arXiv preprint arXiv:2408.06277*, 2024.
- Martin Rohbeck, Edward De Brouwer, Charlotte Bunne, Jan-Christian Huetter, Anne Biton, Kelvin Y Chen, Aviv Regev, and Romain Lopez. Modeling complex system dynamics with flow matching across time and conditions. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Tianrong Chen, Guan-Hong Liu, Molei Tao, and Evangelos Theodorou. Deep momentum multi-marginal schrödinger bridge. *Advances in Neural Information Processing Systems*, 36:57058–57086, 2023.
- Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A computational framework for solving wasserstein lagrangian flows. *arXiv preprint arXiv:2310.10649*, 2023b.
- Katarina Petrović, Lazar Atanackovic, Viggo Moro, Kacper Kapuśniak, Ismail Ilkan Ceylan, Michael Bronstein, Avishek Joey Bose, and Alexander Tong. Curly flow matching for learning non-gradient field dynamics. *arXiv preprint arXiv:2510.26645*, 2025.
- Anming Gu, Edward Chien, and Kristjan Greenewald. Partially observed trajectory inference using optimal transport and a dynamics prior. *arXiv preprint arXiv:2406.07475*, 2024.
- Guillaume Hugué, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in neural information processing systems*, 35:29705–29718, 2022.
- Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Hugué, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. *arXiv preprint arXiv:2307.03672*, 2023b.
- Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth schrödinger bridges. *arXiv preprint arXiv:2503.00530*, 2025.
- Yuhao Sun, Zhenyi Zhang, Zihan Wang, Tiejun Li, and Peijie Zhou. Variational regularized unbalanced optimal transport: Single network, least action. *arXiv preprint arXiv:2505.11823*, 2025.
- Zhenyi Zhang, Zihan Wang, Yuhao Sun, Tiejun Li, and Peijie Zhou. Modeling cell dynamics and interactions with unbalanced mean field schrödinger bridge. *arXiv preprint arXiv:2505.11197*, 2025.
- Volker Bergen, Marius Lange, Stefan Peidli, F. Alexander Wolf, and Fabian J. Theis. Generalizing rna velocity to transient cell states through dynamical modeling. *Nature Biotechnology*, 38(12):1408–1414, August 2020. ISSN 1546-1696. doi: 10.1038/s41587-020-0591-3. URL <http://dx.doi.org/10.1038/s41587-020-0591-3>.

A Prior-guided HJB Inequality Derivation

To derive the velocity-prior guided HJB residual inequality in Eq.6.

$$\begin{aligned} \mathcal{A}(\rho, v, g) &\geq \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \\ &\quad - \int_0^1 \int \rho_t(x) \left(\partial_t s + \frac{1}{2} \|\nabla s\|^2 + \nabla s \cdot v_{\text{prior}} + \frac{1}{2} s^2 \right) dx dt. \end{aligned} \quad (9)$$

We want to start by minimizing Definition 3.2. Since v_{prior} is constant and known, we only need to minimize the "learned" action portion v_{corr} , then the problem becomes the following:

Consider the velocity-prior guided WFR action:

$$\begin{aligned} \min_{\rho, v_{\text{ot}}, g} \quad \mathcal{A}(\rho, v_{\text{corr}}, g) &:= \int_0^1 \int_{\Omega} \left(\frac{1}{2} \|v_{\text{corr}}(t, x)\|^2 + \frac{1}{2} g(t, x)^2 \right) \rho_t(x) dx dt \quad (10) \\ \text{s.t.} \quad \partial_t \rho_t + \nabla \cdot (\rho_t (v_{\text{prior}} + v_{\text{corr}})) &= g_t \rho_t, \quad \rho|_{t=0} = \rho_0, \quad \rho|_{t=1} = \rho_1. \end{aligned}$$

Step 1: Lagrangian formulation

First, we introduce a scalar multiplier $s(t, x)$, the Lagrangian becomes:

$$\begin{aligned} \mathcal{L} &= \int_0^1 \int_{\Omega} \left(\frac{1}{2} \|v_{\text{corr}}\|^2 + \frac{1}{2} g^2 \right) \rho dx dt \quad (11) \\ &\quad + \int_0^1 \int_{\Omega} s \left(\partial_t \rho + \nabla \cdot (\rho (v_{\text{prior}} + v_{\text{corr}})) - g \rho \right) dx dt. \end{aligned}$$

Step 2: Integration by parts

Integration by parts in time on $\int s \partial_t \rho$

$$\begin{aligned} \int_0^1 \int_{\Omega} s \partial_t \rho dx dt &= \left[\int_{\Omega} s \rho dx \right]_{t=0}^{t=1} - \int_0^1 \int_{\Omega} \rho \partial_t s dx dt \quad (12) \\ &= \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] - \int_0^1 \int_{\Omega} \rho \partial_t s dx dt. \end{aligned}$$

Integration by parts in space on $\int s \nabla \cdot (\rho w)$ with $w := v_{\text{prior}} + v_{\text{ot}}$

$$\int_{\Omega} s \nabla \cdot (\rho w) dx = \int_{\partial \Omega} s \rho w \cdot n d\sigma - \int_{\Omega} \nabla s \cdot (\rho w) dx. \quad (13)$$

Assuming zero boundary flux or fast decay, the boundary term vanishes:

$$\int_0^1 \int_{\Omega} s \nabla \cdot (\rho w) dx dt = - \int_0^1 \int_{\Omega} \rho w \cdot \nabla s dx dt. \quad (14)$$

Then, the Lagrangian becomes

$$\begin{aligned} \mathcal{L} &= \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \quad (15) \\ &\quad + \int_0^1 \int_{\Omega} \rho_t(x) \left[\frac{1}{2} \|v_{\text{corr}}\|^2 - v_{\text{corr}} \cdot \nabla s + \frac{1}{2} g^2 - s g - \partial_t s - \nabla s \cdot v_{\text{prior}} \right] dx dt. \end{aligned}$$

Step 3: Fenchel-Young inequality

The Fenchel-Young inequality states that for any vectors a and p , we have:

$$\frac{1}{2} \|a\|^2 \geq p \cdot a - \frac{1}{2} \|p\|^2 \quad (16)$$

We set $a = v_{\text{corr}}$ and $p = \nabla s$, then we have:

$$\frac{1}{2} \|v_{\text{corr}}\|^2 - v_{\text{corr}} \cdot \nabla s \geq -\frac{1}{2} \|\nabla s\|^2 \quad (17)$$

Similarly, we set $a = g$ and $p = s$, then we have:

$$\frac{1}{2} g^2 - s g \geq -\frac{1}{2} s^2, \quad (18)$$

with equality iff $v_{\text{corr}} = \nabla s$ and $g = s$. Thus, putting pieces together, we have:

$$\begin{aligned} \mathcal{A}(\rho, v_{\text{corr}}, g) &\geq \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \\ &\quad - \int_0^1 \int_{\Omega} \rho_t(x) \left(\partial_t s + \frac{1}{2} \|\nabla s\|^2 + \nabla s \cdot v_{\text{prior}} + \frac{1}{2} s^2 \right) dx dt. \end{aligned} \quad (19)$$

B Proof of Theorem 3.4

Theorem B.1 (Prior-guided HJB optimality). *Suppose that the HJB residual defined in corollary.3.3 satisfies $r_\theta(t, x) = 0$ for ρ_t -a.e on $[0, 1] \times \mathbb{R}^d$, and the boundary constraints hold, then $(\rho_t, v_{\text{corr}}, g_\theta)$ satisfies the unbalanced continuity equation and the WFR optimality conditions in Definition 3.2. In particular, the learned corrective field $v_{\text{corr}}^* = \nabla_x s_\theta$ and growth $g^* = g_\theta$ satisfy the optimality conditions.*

Proof sketch The proof of this theorem builds upon the derivations from the previous proof. Recall, from step 3 Fenchel–Young inequality above, we have:

$$\mathcal{A}(\rho, v_{\text{corr}}, g) \geq \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] - \int_0^1 \int_\Omega \rho_t(x) r_s(t, x) dx dt. \quad (20)$$

where the HJB residual is defined as:

$$r_s(t, x) := \partial_t s + \frac{1}{2} \|\nabla_x s\|^2 + \nabla_x s \cdot v_{\text{prior}} + \frac{1}{2} s^2. \quad (21)$$

Step 1 (Dual constraint and lower bound)

To ensure a finite dual lower bound, we need to restrict to potentials s satisfying:

$$r_s(t, x) \geq 0 \quad \forall (t, x). \quad (22)$$

Then under this constraint, since the last term is nonpositive, we have:

$$\mathcal{A}(\rho, v_{\text{corr}}, g) \geq \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \quad \forall s \quad (23)$$

This inequality also holds when we take the infimum over all feasible $(\rho, v_{\text{corr}}, g)$, which gives the dual problem

$$\inf_{\rho, v_{\text{corr}}, g} \mathcal{A}(\rho, v_{\text{corr}}, g) \geq \sup_{s: r_s \geq 0} \left(\mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \right). \quad (24)$$

Step 2 (Residual Optimality)

Now suppose there exists a potential s_θ and a feasible triplet $(\rho_t, v_{\text{corr}}, g)$ such that

$$r_{s_\theta}(t, x) = 0 \quad \text{for } \rho_t\text{-a.e.}, \quad (25)$$

$$v_{\text{corr}} = \nabla_x s_\theta, \quad g = s_\theta, \quad (26)$$

and the boundary constraints $\rho|_{t=0} = \rho_0, \rho|_{t=1} = \rho_1$ hold.

Then the Fenchel–Young inequalities are equalities and equation 20 becomes:

$$\mathcal{A}(\rho, v_{\text{corr}}, g) = \mathbb{E}_{x \sim \rho_0}[s_\theta(0, x)] - \mathbb{E}_{x \sim \rho_1}[s_\theta(1, x)] - \int_0^1 \int_\Omega \rho_t(x) r_{s_\theta}(t, x) dx dt. \quad (27)$$

Since $r_{s_\theta} = 0$ ρ -a.e., then we have:

$$\mathcal{A}(\rho, v_{\text{corr}}, g) = \mathbb{E}_{x \sim \rho_0}[s_\theta(0, x)] - \mathbb{E}_{x \sim \rho_1}[s_\theta(1, x)]. \quad (28)$$

Combining this with the dual lower bound, we have:

$$\inf_{\rho, v_{\text{corr}}, g} \mathcal{A}(\rho, v_{\text{corr}}, g) < \mathcal{A}(\rho, v_{\text{corr}}, g) \quad (29)$$

$$\mathbb{E}_{x \sim \rho_0}[s_\theta(0, x)] - \mathbb{E}_{x \sim \rho_1}[s_\theta(1, x)] < \sup_{s: r_s \geq 0} \left(\mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \right). \quad (30)$$

So we have:

$$\mathcal{A}(\rho, v_{\text{corr}}, g) = \inf_{\rho, v_{\text{corr}}, g} \mathcal{A}(\rho', v'_{\text{corr}}, g') = \sup_{s: r_s \geq 0} \left(\mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \right), \quad (31)$$

Finally, we conclude that $(\rho_t, v_{\text{corr}}, g)$ is primal optimal. In particular,

$$v_{\text{corr}}^* = \nabla_x s_\theta, \quad g^* = s_\theta,$$

and $(\rho_t, v_{\text{corr}}^*, g^*)$ satisfies the unbalanced continuity equation and the WFR optimality conditions in Definition 3.2.

C Additional Background

Hamilton–Jacobi–Bellman (HJB) We recall the classical connection between optimal control and Hamilton–Jacobi (HJ) theory. Consider a deterministic control system with state $x(t) \in \mathbb{R}^d$, control $u(t)$, dynamics $\dot{x} = f(x, u, t)$, running cost $L(x, u, t)$, and terminal cost $\psi(x)$. The *value function*

$$V(t, x) = \inf_{u(\cdot)} \left\{ \int_t^1 L(x(s), u(s), s) ds + \psi(x(1)) \right\}$$

gives the minimal cost-to-go from (t, x) under admissible controls. It is well known that V solves the Hamilton–Jacobi–Bellman (HJB) PDE

$$\partial_t V(t, x) + H(x, \nabla V(t, x), t) = 0, \quad V(1, x) = \psi(x),$$

where the Hamiltonian is

$$H(x, p, t) := \inf_u \left\{ L(x, u, t) + p^\top f(x, u, t) \right\}, \quad p = \nabla V(t, x).$$

Action Matching (AM) AM fits a scalar potential s_θ to learn a energy-minimizing flow between distributions by minimizing the (un)balanced HJB residuals.

$$\mathcal{L}_{\text{uAM}} = \int_0^1 \mathbb{E}_{x \sim \rho_t} \left[\partial_t s_\theta(t, x) + \frac{1}{2} \|\nabla_x s_\theta(t, x)\|^2 + \frac{1}{2} s_\theta^2(t, x) \right] dt, \quad (32)$$

with boundary constraints as: $\mathbb{E}_{x \sim \rho_0} [s_0(x)] - \mathbb{E}_{x \sim \rho_1} [s_1(x)]$.

D Reconstruction Loss

Reconstruction loss To align the terminal distribution in *shape*, we use a sliced Wasserstein objective. Let $\hat{\rho}_1$ be the empirical terminal distribution learned from our model, and ρ_1 be the ground truth target distribution. θ_ℓ is a random projection sampled from $\theta_\ell \sim \text{Unif}(\mathbb{S}^{d-1})$ for $\ell = 1, \dots, L$, \hat{X}_1 are predicted samples, and Y are ground-truth samples, the sliced Wasserstein loss is defined as:

$$\text{SW}_2^2(\hat{\rho}_1, \rho_1) \approx \frac{1}{L} \sum_{\ell=1}^L W_2^2(\langle \theta_\ell, \hat{X}_1 \rangle, \langle \theta_\ell, Y \rangle), \quad (33)$$

where W_2^2 on \mathbb{R} is the 1D Wasserstein distance computed by sorting projections.

To capture *global mass change*, we track the evolution of particle weights along the learned dynamics. Assuming a WFR scale of $\delta = 1$ and a mini-batch of size B , we initialize log-weights as $\log w_i(0) = 0$. The weights evolve according to the potential s_θ through $\frac{d}{dt} \log w_i(t) = s_\theta(t, x_i(t))$, which yields terminal log-weights $\log w_i(1)$. At $t = 0$ the total mass is $M(0) = \sum_{i=1}^B w_i(0) = B$ while at $t = 1$, it is $M(1) = \sum_{i=1}^B w_i(1) = \sum_{i=1}^B \exp(\log w_i(1))$.

To calculate the ground truth mass, we neither have access to the full density ρ_t nor to the absolute scale of $M(T)$. Instead, we use relative mass changes between time points, estimated from the number of observed particles at each time T_k , which is the k -th time point in a multi-snapshot setting. For instance, on the single cell datasets with interval $[T_k, T_{k+1}]$, let N_k be the number of cells observed at time T_k . We can approximate the ground-truth log mass ratio by

$$\log r_{\text{target}, k} = \log \frac{M(T_{k+1})}{M(T_k)} \approx \log \left(\frac{N_{k+1}}{N_k} \right), \quad (34)$$

The model’s predicted ratios $\log \hat{r}_{\text{model}, k}$ is obtained from the weight evolution described above. We then penalize deviations from the ground truth ratio r in the logarithmic form to ensure stability and to enforce WFR consistency: $\mathcal{L}_{\text{mass}} = (\log \hat{r} - \log r)^2$. Our final reconstruction loss combines the two components with tunable coefficients:

$$\mathcal{L}_{\text{recon}} = \lambda_{\text{sw}} \text{SW}_2^2(\hat{\rho}_1, \rho_1) + \lambda_{\text{mass}} (\log \hat{r} - \log r)^2 \quad (35)$$

with $\lambda_{\text{sw}} > 0$ and $\lambda_{\text{mass}} > 0$. In practice we use $L \in [64, 512]$ random projections, and choose $\lambda_{\text{mass}} \in [0.1, 1]$ to calibrate mass without overpowering other terms.

Total objective Putting the pieces together, our total training loss is

$$\min_{\theta} \mathcal{L}(\theta) = \underbrace{\mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] + \lambda_{\text{hjb}} \int_0^1 \mathbb{E}_{x \sim \rho_t}[w(t, x) r_{\theta}(t, x)^2] dt}_{\mathcal{L}_{\text{HJB}}} \quad (36)$$

$$+ \underbrace{\lambda_{\text{sw}} \text{SW}_2^2(\hat{\rho}_1, \rho_1) + \lambda_{\text{mass}} (\log \hat{r} - \log r)^2}_{\mathcal{L}_{\text{recon}}}$$

where $r_{\theta}(t, x)$ is the HJB residual, $w(t, x)$ are the nonnegative weights, and $\hat{\rho}_1(\theta)$ is the terminal distribution obtained by pushing ρ_0 through the learned dynamics.

E Related Works

Physics-constrained approaches Existing works [Koshizuka and Sato, 2022, Neklyudov et al., 2023b, Tong et al., 2020] add a *potential-based prior* $V_t(x)$ to the HJ equation to incorporate prior knowledge for trajectory inference. In [Neklyudov et al., 2023b], this yields $\partial_t s + \frac{1}{2} \|\nabla s\|^2 + V_t + \frac{1}{2\delta^2} s^2 = 0$ and the conservative second-order law $\ddot{X}_t = -\nabla V_t(X_t)$. Meanwhile, our method is the *velocity-based prior* approach, the known velocity field v_{prior} enters as a drift inside the continuity equation, which produces the cross-term $\nabla s \cdot v_{\text{prior}}$. Since potential-based priors is curl-free, it can not represent rotational flows as in [Neklyudov et al., 2023a]; By contrast, our data-driven formulation uses a flexible measured vector fields directly as v_{prior} , which can be considered as a "free drift" and learns only minimal optimal corrections and growth.

Most recently, Curly FM [Petrović et al., 2025] proposes a two-stage pipeline: it first learns a smooth global velocity field from approximate velocities (e.g., RNA velocity), and then solves a Schrödinger bridge problem with this learned nonzero drift. In contrast, our VP-HJF framework does *not* require an explicit first-stage training step to learn a reference drift. We directly incorporate the velocity prior as the drift in a Hamilton–Jacobi dual formulation. Moreover, the dynamic prior in Gu et al. [2024] is assumed to be a clean, well-specified prior. Our setting is more flexible by using the corrective field v_{corr} to adjust noisy or misspecified priors.

Other trajectory inference approaches Trajectory inference has also advanced through flow matching and Schrödinger bridge methods, which scale effectively to high-dimensional data. Recent SB variants further improve performance on single-cell datasets include [Huguet et al., 2022, Tong et al., 2023b, Shen et al., 2024, Hong et al., 2025]. For unbalanced settings, variational and regularized UOT methods such as DeepRUOT, Var-RUOT, VGFM, and UMFSB directly learn transport dynamics and growth from snapshot data [Zhang et al., 2024, Sun et al., 2025, Wang et al., 2025, Zhang et al., 2025]. In particular, Var-RUOT also uses a single network to model both the velocity field and the growth term, but it did not incorporate prior known knowledge like ours. Furthermore, Var-RUOT trains a global-in-time trajectory and evaluates its objective by integrating the dynamics over the entire time horizon, which requires simulating the full trajectory on a fine time grid. In contrast, we use a local per interval training scheme combined with a mixture-based sampling strategy to interpolate data between each interval.

F Additional Experiments

Balanced case - Rotating Ring First, we show a case when utilizing the velocity prior is crucial in learning the correct velocity field where curl-free methods like AM fails. We tested on a 2D rotating ring dataset where the points on the ring (source) are rotated by a fixed angle θ (target). The velocity prior is defined as $v_{\text{prior}} = \omega Jx$, where J is the skew-symmetric rotation matrix and $x \in \mathbb{R}^2$, so this becomes $Jx = \begin{bmatrix} 0 & 1; & -1 & 0 \end{bmatrix} \begin{bmatrix} x_1 & x_2 \end{bmatrix}^T = (-x_2, x_1)^T$. The task for our model v_{ot} is to learn the residual correction after given the prior rotation knowledge, such as ensuring the boundary condition by aligning the mismatched source and target density or correcting the radial drift by push the points inwards or outwards, etc. In Figure 1 left (b), we show that since AM has the curl-free limitation, without a prior, its model $\nabla s_t(x)$ failed to represent pure rotation where the streamlines cut through the circle. In Figure 1 left (a), the streamlines from our method form a circular flow indicating that v_{prior} gives the model the correct inductive bias.

Table 3: Effect of adding $\mathcal{L}_{\text{mass}}$ for Lotka-Volterra. r_{pred} is the predicted mass ratio, the absolute log error $|\Delta \log r| = |\log r_{\text{pred}} - \log r_{\text{true}}|$, the relative error with $r_{\text{true}} = 1.418$ and the velocity RMSE.

Method	r_{pred}	$ \Delta \log r $	Rel. % err.	Vel. RMSE
VP-HJF (ours)	1.356	0.044	4.30%	0.137
VP-HJF (w/o $\mathcal{L}_{\text{mass}}$)	0.899	0.455	36.56%	0.158
Unbalanced AM	0.938	0.413	33.80%	0.302
Prior-only	1.000	0.348	29.40%	0.050

Diverging Petal We created a curved and rotated petal-shape dataset to test our method on diverging multi-trajectory paths. The source is a gaussian distribution concentrated in the center and $v_{\text{prior}} = \omega Jx$ the rotation dynamic defined as before. Our task is to learn v_{petal} , a radial and angle-dependent term that push outward the points along the radius with different speed depend on the angle $\theta = \text{atan2}(y, x)$. We have $v_{\text{petal}}(x) = s(\theta) \hat{r}$, $\hat{r} = \frac{x}{\|x\|}$, $s(\theta) = \max(0, b + a \cos(k\theta))$. Compared with the petal shape appeared in AM and MioFlow [Huguet et al., 2022], the underlying dynamic flow in our example is harder to learn, where the former one has a straight-axis aligned radial expansion as $r(x) = |x_1| + |x_2|$ with the gradient of $r(x)$ being a piece-wise constant and curl-free. Figure 1 middle shows that our petal shape matches with the target shape. Figure 1 right shows that the vector field (red arrows) are bending away from pure rotation (blue arrows) to align strongly with the petal shape by pushing the mass outward.

Lotka-Volterra with growth. We model the prey and predator densities $x_1(t), x_2(t)$ by a first-order nonlinear ODE,

$$\dot{x}_1(t) = \alpha x_1(t) - \beta x_1(t) x_2(t), \quad \dot{x}_2(t) = -\gamma x_2(t) + \delta x_1(t) x_2(t),$$

where α is the prey’s intrinsic growth rate, β is the predation rate, γ is the predator’s mortality rate, and δ is the predator’s growth rate from consuming prey. To model the population expansion and decay dynamics, we use a simple scalar growth field and evolve local mass via the weight dynamics

$$g(x(t)) = \kappa (x_1(t) - x_2(t)), \quad \frac{d}{dt} \log w(t) = g(x(t))$$

The total mass $M(t) = \mathbb{E}[w(t)]$ and the ground-truth mass ratio is calculated as $r_{\text{true}} = M(t)/M(0)$. We define the v_{prior} as the oracle LV drift with added Gaussian noise with no growth term.

Table 3 shows that adding the explicit mass term $\mathcal{L}_{\text{mass}}$ enables our method to closely match r_{true} with 4% of relative error while both unbalanced AM and our method without the $\mathcal{L}_{\text{mass}}$ term suffer from a much higher relative error of over 30%. Although AM aligns transport but leaves the *scale* of the scalar potential s_θ unconstrained, so the integrated growth $\int_0^t g(x(t)) dt$ is miscalibrated. By contrast, $\mathcal{L}_{\text{mass}}$ provides endpoint constraint on the mass – yielding better mass dynamics and more aligned with the ground truth mass ratio.

Bone marrow scRNA-Seq data We evaluate our approach on a real scRNA-seq bone-marrow atlas with multiple hematopoietic fates from scVelo [Bergen et al., 2020]. Figure 2 (left) shows trajectories that emanates from early progenitor regions at $t=0$ (dark blue) and spread out to other branches by closely following the UMAP reference (gray). Figure 2 (right) shows the learned growth field $g_\theta(t, x) = s_\theta(t, x)$ that governs local mass dynamics. This maps shows that our model successfully captured low growth rate in early progenitors and increases as cells enter the active cycling and amplification stage. However, we also observe that the high growth rate near some terminal regions. While this indicates the model can assign higher growth to specific cell types, a strong terminal-phase growth is biologically implausible. For the bone marrow case, mature or exiting cells should have near-zero or negative growth. This likely reflect the objective imbalance where transport terms dominating mass calibration, suggesting mild regularization such as time-smoothness on s_θ or branch-wise boundary constraints to better align growth with biology.

G Discussion and Limitations

On the velocity prior quality and assumptions The velocity prior v_{prior} indeed plays a constructive—but double-edged role in our method. A good prior captures coarse dynamics - reducing




Figure 2: **Left:** Bone marrow trajectories. Colored points (ours) show inferred cell trajectories overlaid on the reference manifold (gray) **Right:** Learned growth field — orange means high growth, purple means low growth.

the learning complexity and improving on sample efficiency. A misspecified prior can bias the learned corrective field $s_\theta(t, x)$ and slow or destabilize training. Hence, the *quality* of v_{prior} strongly influences both optimization and generalization. In practice, mild perturbation of the prior through noise, scale or misspecification are corrected by s_θ , whereas severe misspecification such as overly large or structurally wrong drifts can bias the learned corrective flow. Moreover, v_{prior} does *not* require divergence-free assumption. Our dual objective explicitly includes the cross term $\nabla_x s_\theta \cdot v_{\text{prior}}$ in the HJB residual avoiding hidden orthogonality requirements.

Limitations In Fig. 2 (right) we observe high growth near terminal regions, which is biologically implausible for mature or cell-cycle-exiting states. Without additional biological constraints such as cell-cycle markers, branch-terminal boundary conditions or proliferation markers, growth-transport disentanglement may remain underdetermined in some regions. A promising direction is to add weak supervision on the learned growth model to improve identifiability.

For single-cell datasets we currently use local supervision—training on adjacent pairs with a time-continuous shared network. This choice is simple and scalable and induces a globally smooth field, but it does not jointly enforce all marginals as in recent multi-marginal methods, which may limit long-range trajectory coherence. In future works, extending VP-HJF with global consistency could improve on long-range trajectory inference.

H Algorithm

Algorithm 1 Training VP-HJF (Velocity-Prior Hamiltonian–Jacobi Flow)

Require: per time-interval t_k snapshots $\{x, t_k, v_k^{\text{prior}}\}_{k=0}^K$, network $s_\theta(t, x)$, coefficients $\alpha, \beta, \lambda_{\text{SW}}, \lambda_{\text{mass}}$, batch size B

- 1: **while** not converged **do**
- 2: Sample a global batch across all intervals $(x, t, v)_{b=1}^B \sim \{x, t, v_{\text{prior}}\}_{k=0}^K$
- 3: Compute $s_0 \leftarrow s_\theta(t_0, x_0), s_1 \leftarrow s_\theta(t_1, x_1)$
- 4: **for** $k = 0$ to $K - 1$ **do**
- 5: Sample adjacent pairs $(x_k, v_k, x_{k+1}, v_{k+1})_{b=1}^B \sim \{x, t, v_{\text{prior}}\}_{k=0}^K$
- 6: **HJB residual loss for interval** $[t_k, t_{k+1}]$:
- 7: **for** $m = 1$ to M_{HJB} **do**
- 8: Sample normalized $u^{(m)} \sim \mathcal{U}(0, 1)$
- 9: Sample $x_{t,k}$ from the mixture $\rho_{t,k}$ of ρ_k and ρ_{k+1} , for $b = 1, \dots, B$
- 10: Compute $r_{t,k} \leftarrow \partial_t s_{t,k} + \frac{1}{2} \|\nabla_x s_{t,k}\|^2 + \nabla_x s_{t,k} \cdot v_{\text{prior}}(t_k^{(m)}, x_{t,k}) + \frac{1}{2} (s_{t,k})^2$
- 11: Compute importance weights $\tilde{w}_b \propto |r_{t,k}^{(b)}|^{-\tau}$
- 12: $\mathcal{L}_{\text{HJB},k} \leftarrow \mathcal{L}_{\text{HJB},k} + \frac{1}{B} \sum_{b=1}^B w_b (r_{t,k})^2$
- 13: **end for**
- 14: **Reconstruction loss:**
- 15: Define ODE rhs: $\dot{x} = \nabla_x s_\theta(t, x) + v_{\text{prior}}(t, x), \quad \dot{\log w} = s_\theta(t, x)$
- 16: Compute $(x_{k+1}^{\text{pred}}, \log w_{k+1}) \leftarrow \text{odeint}(\text{ode rhs}, (x_k, \mathbf{0}), t \in [t_k, t_{k+1}])$
- 17: $\mathcal{L}_{\text{recon}} \leftarrow \mathcal{L}_{\text{recon}} + \lambda_{\text{SW}} \frac{1}{L} \sum_{\ell=1}^L W_2^2(x_{k+1}^{\text{pred}}, x_{k+1}) + \lambda_{\text{mass}} (\log \hat{r}_k - \log r_k^2)$
- 18: **end for**
- 19: **Total loss:** $\mathcal{L}(\theta) \leftarrow \alpha \mathcal{L}_{\text{HJB}} + \beta \mathcal{L}_{\text{recon}}$
- 20: Update $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}(\theta)$
- 21: **end while**
