

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LEARNING VELOCITY PRIOR-GUIDED HAMILTONIAN-JACOBI FLOWS WITH UNBALANCED OPTIMAL TRANS- PORT

Anonymous authors

Paper under double-blind review

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 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,

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

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).

Motivation in practice In domains like single-cell biology and physical systems with known drifts, accurate priors are available yet incomplete and and mass change such as cell proliferation and decay is ubiquitous. VP-HJF exploits these priors to encode hard-to-learn structure. The residual potential learns both the correction that the prior can not explain and the mass changes through the growth term. This yields a compact and interpretable alternative to fully free vector-field models, particularly effective when local supervision is noisy but prior knowledge is rich.

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:

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

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) equation $\partial_t V(t, x) + H(x, \nabla V(t, x), t) = 0$, where $V(1, x) = \psi(x)$ and the Hamiltonian is $H(x, p, t) := \inf_u \{L(x, u, t) + p^\top f(x, u, t)\}$ with $p = \nabla V(t, x)$.

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.

$$\mathcal{L}_{\text{uAM}} = \int_0^1 \mathbb{E}_{x \sim \rho_t} \left[\partial_t s_\theta(t, x) + \frac{1}{2} \|\nabla_x s_\theta(t, x)\|^2 + \frac{1}{2} s_\theta^2(t, x) \right] dt, \quad (4)$$

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

108 **3 METHODOLOGY**
 109

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

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

120 Building on Proposition 3.1, we reduce the WFR problem to learning a single model and incorpo-
 121 rate problem-specific dynamics through a simple velocity decomposition. Specifically, we decom-
 122 pose the velocity field into two parts: a known velocity prior and a learnable corrective velocity field
 123 component:

$$v_{\text{total}}(t, x) = v_{\text{prior}}(t, x) + v_{\text{corr}}(t, x), \quad (5)$$

124 where v_{prior} encodes domain knowledge (e.g. translations, rotations, RNA velocity), and v_{corr} is the
 125 data–driven corrective component. In this way, the prior captures coarse dynamics while the model
 126 focuses on refinements such as correcting the residual transport and learning mass imbalance that
 127 the prior cannot explain. In essence, our approach improves interpretability and reduces the learning
 128 complexity through adding prior knowledge of the velocity field v_{prior} – leaving the learnable veloc-
 129 ity field v_{corr} simpler learning tasks compared with other generative modeling methods of learning
 130 the entire velocity field v_{total} . Intuitively, our approach pays kinetic cost only for the *correction* to
 131 the prior drift and for the mass growth-decay component, making learning more efficient.

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

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

$$\mathcal{A}(\rho, v_{\text{corr}}, g) = \int_0^1 \int \left(\frac{1}{2} \|v(t, x)\|^2 + \frac{1}{2} g(t, x)^2 \right) \rho_t(x) dx dt, \quad (6)$$

$$\text{s.t. } \partial_t \rho_t = -\nabla \cdot (\rho_t (v_{\text{prior}} + v_{\text{corr}})) + g_t \rho_t, \quad \rho|_{t=0} = \rho_0, \quad \rho|_{t=1} = \rho_1. \quad (7)$$

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

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

$$\begin{aligned} \mathcal{A}(\rho, v_{\text{corr}}, g) &\geq \mathbb{E}_{\rho_0(x)}[s_0(x)] - \mathbb{E}_{\rho_1(x)}[s_1(x)] \\ &\quad - \int_0^1 \int \rho_t(x) \left(\partial_t s + \frac{1}{2} \|\nabla s\|^2 + \nabla s \cdot v_{\text{prior}} + \frac{1}{2} s^2 \right) dx dt. \end{aligned} \quad (8)$$

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

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

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

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

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

$$165 \quad r_\theta(t, x) := \partial_t s + \frac{1}{2} \|\nabla_x s\|^2 + \nabla_x s \cdot v_{\text{prior}} + \frac{1}{2} s^2. \quad (9)$$

166 *Then minimizing*

$$168 \quad \mathcal{L}_{\text{HJB}}(\theta) = \mathbb{E}_{\rho_0(x)}[s_0(x)] - \mathbb{E}_{\rho_1(x)}[s_1(x)] + \int_0^1 \mathbb{E}_{\rho_t(x)}[w(t, x)r_\theta(t, x)^2] dt \quad (10)$$

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

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

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

186 While the HJB residual enforces local optimality conditions, it does not guarantee that the terminal
 187 distribution ρ_1^θ matches with the target ρ_1 . To bridge this gap, we design a two-part reconstruction
 188 objective: (i) a density matching term through the sliced Wasserstein between the predicted $\hat{\rho}_1$ and
 189 ρ_1 , and (ii) a mass term aligning the global log-mass ratio. These two terms directly calibrates the
 190 terminal distribution's *shape* and *mass*, complementing the HJB residual.

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

$$196 \quad \text{SW}_2^2(\hat{\rho}_1, \rho_1) \approx \frac{1}{L} \sum_{\ell=1}^L W_2^2(\langle \theta_\ell, \hat{X}_1 \rangle, \langle \theta_\ell, Y \rangle), \quad (11)$$

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

200 To capture *global mass change*, we track the evolution of particle weights along the learned dy-
 201 namics. Assuming a WFR scale of $\delta = 1$ and a mini-batch of size B , we initialize log-weights as
 202 $\log w_i(0) = 0$. The weights evolve according to the potential s_θ through $\frac{d}{dt} \log w_i(t) = s_\theta(t, x_i(t))$,
 203 which yields terminal log-weights $\log w_i(1)$. At $t = 0$ the total mass is $M(0) = \sum_{i=1}^B w_i(0) = B$
 204 while at $t = 1$, it is $M(1) = \sum_{i=1}^B w_i(1) = \sum_{i=1}^B \exp(\log w_i(1))$.

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

$$212 \quad \log r_{\text{target}, k} = \log \frac{M(T_{k+1})}{M(T_k)} \approx \log \left(\frac{N_{k+1}}{N_k} \right), \quad (12)$$

214 The model's predicted ratios $\log \hat{r}_{\text{model}, k}$ is obtained from the weight evolution described above. We
 215 then penalize deviations from the ground truth ratio r in the logarithmic form to ensure stability and

216 to enforce WFR consistency: $\mathcal{L}_{\text{mass}} = (\log \hat{r} - \log r)^2$. Our final reconstruction loss combines the
 217 two components with tunable coefficients:
 218

$$219 \quad \mathcal{L}_{\text{recon}} = \lambda_{\text{sw}} \text{SW}_2^2(\hat{\rho}_1, \rho_1) + \lambda_{\text{mass}} (\log \hat{r} - \log r)^2 \quad (13)$$

222 with $\lambda_{\text{sw}} > 0$ and $\lambda_{\text{mass}} > 0$. In practice we use $L \in [64, 512]$ random projections, and choose
 223 $\lambda_{\text{mass}} \in [0.1, 1]$ to calibrate mass without overpowering other terms.
 224

225 **Total objective** Putting the pieces together, our total training loss is
 226

$$228 \quad \min_{\theta} \mathcal{L}(\theta) = \underbrace{\mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] + \lambda_{\text{hjb}} \int_0^1 \mathbb{E}_{x \sim \rho_t}[w(t, x) r_{\theta}(t, x)^2] dt}_{\mathcal{L}_{\text{HJB}}} \quad (14)$$

$$231 \quad + \underbrace{\lambda_{\text{sw}} \text{SW}_2^2(\hat{\rho}_1, \rho_1) + \lambda_{\text{mass}} (\log \hat{r} - \log r)^2}_{\mathcal{L}_{\text{recon}}}$$

235 where $r_{\theta}(t, x)$ is the HJB residual, $w(t, x)$ are the nonnegative weights, and $\hat{\rho}_1(\theta)$ is the terminal
 236 distribution obtained by pushing ρ_0 through the learned dynamics.
 237

239 4 RELATED WORKS

241 **Physics-constrained approaches** Existing works (Koshizuka and Sato, 2022; Neklyudov et al.,
 242 2023b; Tong et al., 2020) add a *potential-based prior* $V_t(x)$ to the HJ equation to incorporate prior
 243 knowledge for trajectory inference. In (Neklyudov et al., 2023b), this yields $\partial_t s + \frac{1}{2} \|\nabla s\|^2 + V_t +$
 244 $\frac{1}{2s^2} s^2 = 0$ and the conservative second-order law $\ddot{X}_t = -\nabla V_t(X_t)$. Meanwhile, our method is the
 245 *velocity-based prior* approach, the known velocity field v_{prior} enters as a drift inside the continuity
 246 equation, which produces the cross-term $\nabla s \cdot v_{\text{prior}}$. Since potential-based priors are curl-free, it
 247 can not represent rotational flows as in (Neklyudov et al., 2023a); By contrast, our data-driven
 248 formulation uses a flexible measured vector fields directly as v_{prior} , which can be considered as a
 249 “free drift” and learns only minimal optimal corrections and growth.
 250

251 Most recently, Curly FM (Petrović et al., 2025) proposes a two-stage pipeline: it first learns a smooth
 252 global velocity field from approximate velocities (e.g., RNA velocity), and then solves a Schrödinger
 253 bridge problem with this learned nonzero drift. In contrast, our VP–HJF framework does *not* require
 254 an explicit first-stage training step to learn a reference drift. We directly incorporate the velocity prior
 255 as the drift in a Hamilton–Jacobi dual formulation. Moreover, the dynamic prior in Gu et al. (2024)
 256 is assumed to be a clean, well-specified prior. Our setting is more flexible by using the corrective
 257 field v_{corr} to adjust noisy or misspecified priors.
 258

259 **Other trajectory inference approaches** Trajectory inference has also advanced through flow
 260 matching (Haviv et al., 2024; Kapusniak et al., 2024; Atanackovic et al., 2024; Eyring et al., 2023),
 261 and Schrödinger bridge methods, which scale effectively to high-dimensional data. Recent SB vari-
 262 ants further improve performance on single-cell datasets include (Huguet et al., 2022; Tong et al.,
 263 2023b; Shen et al., 2024; Hong et al., 2025; Pariset et al., 2023; Lavenant et al., 2024). For unbal-
 264 anced settings, variational and regularized UOT methods such as TIGON, DeepRUOT, Var-RUOT,
 265 VGF, and UMFSB directly learn transport dynamics and growth from snapshot data (Sha et al.,
 266 2024; Sun et al., 2025; Wang et al., 2025; Zhang et al., 2025). In particular, Var-RUOT also uses a
 267 single network to model both the velocity field and the growth term, but it did not incorporate prior
 268 known knowledge like ours. Furthermore, Var-RUOT trains a global-in-time trajectory and evaluates
 269 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.

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

271 **Require:** per time-interval t_k snapshots $\{x, t_k, v_k^{\text{prior}}\}_{k=0}^K$, network $s_\theta(t, x)$, coefficients
 272 $\alpha, \beta, \lambda_{\text{SW}}, \lambda_{\text{mass}}$, batch size B

273 1: **while** not converged **do**

274 2: Sample a global batch across all intervals $(x, t, v)_{b=1}^B \sim \{x, t, v_{\text{prior}}\}_{k=0}^K$

275 3: Compute $s_0 \leftarrow s_\theta(t_0, x_0)$, $s_1 \leftarrow s_\theta(t_1, x_1)$

276 4: **for** $k = 0$ to $K - 1$ **do**

277 5: Sample adjacent pairs $(x_k, v_k, x_{k+1}, v_{k+1})_{b=1}^B \sim \{x, t, v_{\text{prior}}\}_{k=0}^K$

278 6: **HJB residual loss for interval** $[t_k, t_{k+1}]$:

279 7: **for** $m = 1$ to M_{HJB} **do**

280 8: Sample normalized $u^{(m)} \sim \mathcal{U}(0, 1)$

281 9: Sample $x_{t,k}$ from the mixture $\rho_{t,k}$ of ρ_k and ρ_{k+1} , for $b = 1, \dots, B$

282 10: Compute $r_{t,k} \leftarrow \partial_t s_{t,k} + \frac{1}{2} \|\nabla_x s_{t,k}\|^2 + \nabla_x s_{t,k} \cdot v_{\text{prior}}(t_k^{(m)}, x_{t,k}) + \frac{1}{2} (s_{t,k})^2$

283 11: Compute importance weights $\tilde{w}_b \propto |r_{t,k}^{(b)}|^{-\tau}$

284 12: $\mathcal{L}_{\text{HJB},k} \leftarrow \mathcal{L}_{\text{HJB},k} + \frac{1}{B} \sum_{b=1}^B w_b (r_{t,k})^2$

285 13: **end for**

286 14: **Reconstruction loss:**

287 15: Define ODE rhs: $\dot{x} = \nabla_x s_\theta(t, x) + v_{\text{prior}}(t, x)$, $\dot{\log w} = s_\theta(t, x)$

288 16: Compute $(x_{k+1}^{\text{pred}}, \log w_{k+1}) \leftarrow \text{odeint}(\text{ode rhs}, (x_k, \mathbf{0}), t \in [t_k, t_{k+1}])$

289 17: $\mathcal{L}_{\text{recon}} \leftarrow \mathcal{L}_{\text{recon}} + \lambda_{\text{SW}} \frac{1}{L} \sum_{\ell=1}^L W_2^2(x_{k+1}^{\text{pred}}, x_{k+1}) + \lambda_{\text{mass}} (\log \hat{r}_k - \log r_k^2)$

290 18: **end for**

291 19: **Total loss:** $\mathcal{L}(\theta) \leftarrow \alpha \mathcal{L}_{\text{HJB}} + \beta \mathcal{L}_{\text{recon}}$

292 20: Update $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}(\theta)$

293 21: **end while**

5 EXPERIMENTS

5.1 SYNTHETIC DATASET

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

312 **Diverging Petal** We created a curved and rotated petal-shape dataset to test our method on di-
 313 verging 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-
 314 dependent term that push outward the points along the radius with different speed depend on the
 315 angle $\theta = \text{atan2}(y, x)$. We have

$$v_{\text{petal}}(x) = s(\theta) \hat{r}, \quad \hat{r} = \frac{x}{\|x\|}, \quad s(\theta) = \max(0, b + a \cos(k\theta)) \quad (15)$$

316 Compared with the petal shape appeared in AM and MIOFlow (Huguet et al., 2022), the underlying
 317 dynamic flow in our example is harder to learn, where the former one has a straight-axis aligned
 318 radial expansion as $r(x) = |x_1| + |x_2|$ with the gradient of $r(x)$ being a piece-wise constant and
 319 curl-free. Figure 1 middle shows that our petal shape matches with the target shape. Figure 1 right
 320 shows that the vector field (red arrows) are bending away from pure rotation (blue arrows) to align
 321 strongly with the petal shape by pushing the mass outward.

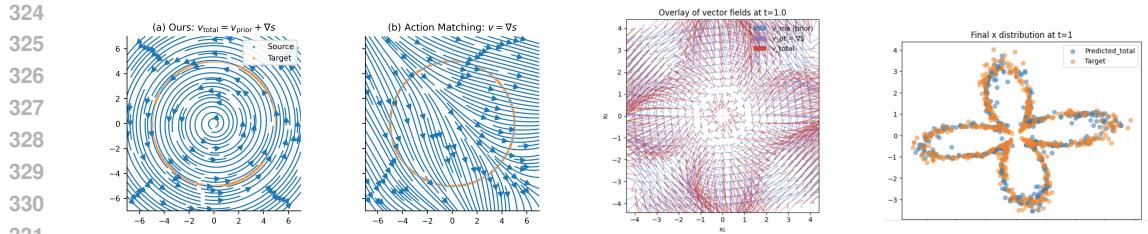


Figure 1: **Left** (first two): Ours correctly learned the rotating dynamic while AM failed. **Middle**: Vector field (red) bending away from v_{prior} (blue) to form petal shapes. **Right**: Predicted target distribution matches with ground truth

Table 2: Effect of adding $\mathcal{L}_{\text{mass}}$ for Lokta-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

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. The prior alone shows moderate accuracy but with high kinetic cost, whereas our method balances both accuracy and energy efficiency.

Lotka–Volterra with growth. We model the prey and predator densities $x_1(t), x_2(t)$ by a first-order nonlinear ODE, $\dot{x}_1(t) = \alpha x_1(t) - \beta x_1(t) x_2(t)$, $\dot{x}_2(t) = -\gamma x_2(t) + \delta x_1(t) x_2(t)$, 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 (Goel et al., 1971). To model the population expansion and decay dynamics, we use a simple scalar growth field and evolve local mass via the weight dynamics $g(x(t)) = \kappa(x_1(t) - x_2(t))$, $\frac{d}{dt} \log w(t) = g(x(t))$. 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 prior e oracle LV drift with added Gaussian noise with no growth term.

Table 2 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.

5.2 REAL-WORLD DATASET

In this section, we evaluate our method on two single-cell RNA-seq datasets. Both provide RNA velocity, which we use as the velocity prior v_{prior} . Such priors are common in biological and scientific applications beyond single-cell data. Incorporating them introduces an inductive bias that reduces

378 Table 3: Comparison on the EB dataset using SWD, MMD and W1, at the held-out marginals (t_1 ,
 379 t_3). Baseline results other than * 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
Var-RUOT*	0.37	0.24	0.25	0.06	10.28	11.92
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	0.47	0.18	0.17	11.83	13.98

390 Table 4: Robustness analysis of VP–HJF to perturbations of the velocity prior on the EB dataset of
 391 100 dim. We report mean \pm std over 5 seeds

	Clean	Gaussian noise η		scale c	
		0.25	0.75	0.5	1.5
$W_1 t_2$	13.26 ± 0.08	13.19 ± 0.08	13.19 ± 0.08	13.22 ± 0.08	13.21 ± 0.08
$W_1 t_4$	14.59 ± 0.09	14.58 ± 0.10	14.57 ± 0.09	14.56 ± 0.10	14.62 ± 0.09

400 learning complexity — our model needs only to learn a corrective flow and growth rather than the
 401 full dynamics from scratch.

402 **EB scRNA-Seq data** We evaluate cell-trajectory inference on the Embryoid Body (EB) dataset
 403 of Moon et al. (2019), using the preprocessed release from Koshizuka and Sato (2022); Tong et al.
 404 (2020). The dataset comprises five snapshots over 27 days, grouped as $t_0 \in [0, 3]$, $t_1 \in [6, 9]$,
 405 $t_2 \in [12, 15]$, $t_3 \in [18, 21]$, $t_4 \in [24, 27]$. Leveraging RNA velocity as a prior v_{prior} at each snapshot,
 406 we train a local and shorter trajectory by adopting the multi-marginal *local per-interval* training: at
 407 each step we sample an adjacent pair (t_k, t_{k+1}) and learn only the transport and growth to move
 408 $\rho_{t_k} \rightarrow \rho_{t_{k+1}}$. This yields more stable gradients and low target variance than enforcing all time points
 409 jointly. For details of the training algorithm, see Algorithm 1.

410 We test on 100-dim PCA components feature space and compare with recent works using the
 411 multi-marginal approach from 3MSBM (Theodoropoulos et al., 2025), SBIRR (Shen et al., 2024),
 412 MMFM (Rohbeck et al., 2025) and DMSB (Chen et al., 2023) as well as other methods using
 413 global-in-time joint training or unbalanced optimal transport DeepRUOT (Zhang et al., 2024), Var-
 414 RUOT (Sun et al., 2025) and MIOFlow (Huguet et al., 2022). We follow the experiment setup from
 415 3MSBM by having $t = 1, 3$ as the held-out sets and evaluating on various metrics. Table 3 shows
 416 that our method outperforms most methods and remains competitive with Var-RUOT and 3MSBM.
 417 Notably, our method outperforms SBIRR and MMFM, which solve piecewise Schrödinger bridges
 418 and OT couplings whereas methods like 3MSBM and DMSB solve a single global optimization
 419 with a joint coupling, indicating the benefits of using RNA velocity as a local prior with per-interval
 420 supervision. For additional results and comparison with Var-RUOT see Appendix C.

421 We also conducted robustness analysis on the mis-specification of v_{prior} of our approach. Specifi-
 422 cally, we perturb the reference field by (i) adding Gaussian noise, $v_{\text{prior}} = v_{\text{prior}} + \eta$, and (ii)
 423 rescaling its magnitude, $v_{\text{prior}} = c v_{\text{prior}}$. In Table 4, W_1 at t_2 and t_4 changes only marginally
 424 under both noise levels ($\alpha \in \{0.25, 0.75\}$) and scaling factors ($c \in \{0.5, 1.5\}$). This indicates that
 425 VP–HJF is robust under mild to moderate prior perturbations, with the learned corrective field v_{corr}
 426 adapting to and compensating for mis-specification in v_{prior} .

427 **Bone marrow scRNA-Seq data** We evaluate our approach on a real scRNA-seq bone-marrow
 428 atlas with multiple hematopoietic fates from scVelo (Bergen et al., 2020). Figure 2 (left) shows
 429 trajectories that emanates from early progenitor regions at $t=0$ (dark blue) and spread out to other
 430 branches by closely following the UMAP reference (gray). Figure 2 (right) shows the learned growth

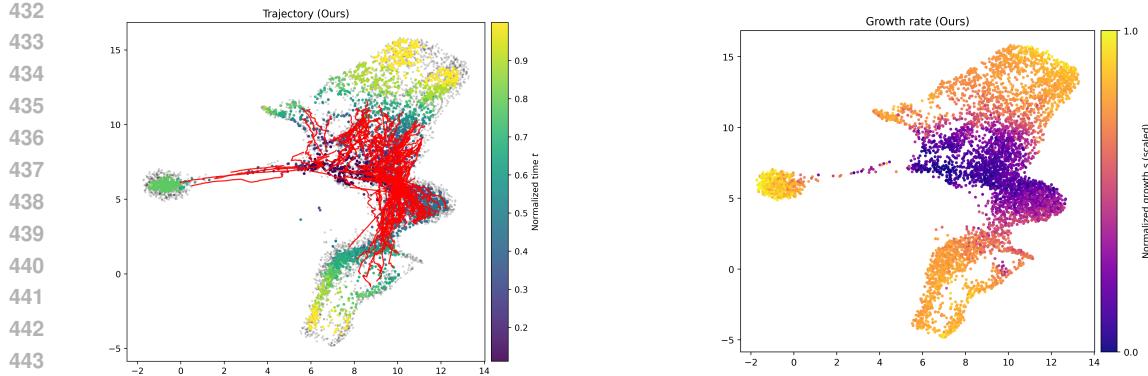


Figure 2: **Left:** Bone marrow trajectories. Colored points (ours) show inferred cell trajectories overlaid on the reference manifold (gray). **Right:** Learned growth field — orange means high growth, purple means low growth.

field $g_\theta(t, x) = s_\theta(t, x)$ that governs local mass dynamics. This map 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.

6 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 the learning complexity and improving on sample efficiency. A mis-specified 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 mis-specification are corrected by s_θ , whereas severe mis-specification 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 under-determined 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.

7 CONCLUSION

Our method decomposes the velocity field by using domain knowledge as prior and using a single network to capture both growth and transport. This decomposition yields robust performance even under mild to moderate prior mis-specification, indicating the flexibility of this framework to incorporate priors, making it a promising direction for modeling complex cellular dynamics.

486

REFERENCES

487

488 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
489 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.490 Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic inter-
491 polants. *arXiv preprint arXiv:2209.15571*, 2022.492 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
493 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.494
495 Kirill Neklyudov, Rob Brekelsmans, Daniel Severo, and Alireza Makhzani. Action matching: Learn-
496 ing stochastic dynamics from samples. In *International conference on machine learning*, pages
497 25858–25889. PMLR, 2023a.498 Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the*
499 *space of probability measures*. Springer, 2005.
500501 Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lip-
502 man, and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch cou-
503 plings. *arXiv preprint arXiv:2304.14772*, 2023.504 Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-
505 Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models
506 with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023a.507
508 Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2008.509 Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-
510 kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.
511512 Lenaic Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced opti-
513 mal transport: Dynamic and kantorovich formulations. *Journal of Functional Analysis*, 274(11):
514 3090–3123, 2018.515 Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through
516 regularized unbalanced optimal transport. *arXiv preprint arXiv:2410.00844*, 2024.
517518 Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-
519 growth flow matching for single-cell dynamics modeling. *arXiv preprint arXiv:2505.13413*, 2025.520 Takeshi Koshizuka and Issei Sato. Neural lagrangian schrödinger bridge: Diffusion modeling for
521 population dynamics. *arXiv preprint arXiv:2204.04853*, 2022.522
523 Kirill Neklyudov, Rob Brekelsmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza
524 Makhzani. A computational framework for solving wasserstein lagrangian flows. *arXiv preprint*
525 *arXiv:2310.10649*, 2023b.526 Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet:
527 A dynamic optimal transport network for modeling cellular dynamics. In *International conference*
528 *on machine learning*, pages 9526–9536. PMLR, 2020.529
530 Katarina Petrović, Lazar Atanackovic, Viggo Moro, Kacper Kapuśnian, Ismail İlkan Ceylan,
531 Michael Bronstein, Avishek Joey Bose, and Alexander Tong. Curly flow matching for learning
532 non-gradient field dynamics. *arXiv preprint arXiv:2510.26645*, 2025.533 Anming Gu, Edward Chien, and Kristjan Greenewald. Partially observed trajectory inference using
534 optimal transport and a dynamics prior. *arXiv preprint arXiv:2406.07475*, 2024.535 Doron Haviv, Aram-Alexandre Pooladian, Dana Pe'er, and Brandon Amos. Wasserstein flow match-
536 ing: Generative modeling over families of distributions. *arXiv preprint arXiv:2411.00698*, 2024.
537538 Kacper Kapuśnian, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein,
539 Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the
data manifold. *Advances in Neural Information Processing Systems*, 37:135011–135042, 2024.

540 Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio,
 541 Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the
 542 wasserstein manifold. *arXiv preprint arXiv:2408.14608*, 2024.

543

544 Luca Eyring, Dominik Klein, Théo Uscidda, Giovanni Palla, Niki Kilbertus, Zeynep Akata, and
 545 Fabian Theis. Unbalancedness in neural monge maps improves unpaired domain translation.
 546 *arXiv preprint arXiv:2311.15100*, 2023.

547

548 Guillaume Huguet, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik
 549 Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows
 550 for trajectory inference. *Advances in neural information processing systems*, 35:29705–29718,
 551 2022.

552

553 Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume
 554 Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and
 555 flow matching. *arXiv preprint arXiv:2307.03672*, 2023b.

556

557 Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal schrödinger bridges
 558 with iterative reference refinement. *arXiv preprint arXiv:2408.06277*, 2024.

559

560 Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth schrödinger
 561 bridges. *arXiv preprint arXiv:2503.00530*, 2025.

562

563 Matteo Pariset, Ya-Ping Hsieh, Charlotte Bunne, Andreas Krause, and Valentin De Bortoli. Unbalanced
 564 diffusion schrödinger bridge. *arXiv preprint arXiv:2306.09099*, 2023.

565

566 Hugo Lavenant, Stephen Zhang, Young-Heon Kim, Geoffrey Schiebinger, et al. Toward a mathematical
 567 theory of trajectory inference. *The Annals of Applied Probability*, 34(1A):428–500, 2024.

568

569 Yutong Sha, Yuchi Qiu, Peijie Zhou, and Qing Nie. Reconstructing growth and dynamic trajectories
 570 from single-cell transcriptomics data. *Nature Machine Intelligence*, 6(1):25–39, 2024.

571

572 Yuhao Sun, Zhenyi Zhang, Zihan Wang, Tiejun Li, and Peijie Zhou. Variational regularized unbalanced
 573 optimal transport: Single network, least action. *arXiv preprint arXiv:2505.11823*, 2025.

574

575 Zhenyi Zhang, Zihan Wang, Yuhao Sun, Tiejun Li, and Peijie Zhou. Modeling cell dynamics and
 576 interactions with unbalanced mean field schrödinger bridge. *arXiv preprint arXiv:2505.11197*,
 577 2025.

578

579 Narendra S Goel, Samaresh C Maitra, and Elliott W Montroll. On the volterra and other nonlinear
 580 models of interacting populations. *Reviews of modern physics*, 43(2):231, 1971.

581

582 Kevin R Moon, David Van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen,
 583 Kristina Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, et al. Visualizing
 584 structure and transitions in high-dimensional biological data. *Nature biotechnology*, 37(12):1482–
 585 1492, 2019.

586

587 Panagiotis Theodoropoulos, Augustinos D Saravacos, Evangelos A Theodorou, and Guan-Horng
 588 Liu. Momentum multi-marginal schrödinger bridge matching. *arXiv preprint arXiv:2506.10168*, 2025.

589

590 Martin Rohbeck, Edward De Brouwer, Charlotte Bunne, Jan-Christian Huetter, Anne Biton,
 591 Kelvin Y Chen, Aviv Regev, and Romain Lopez. Modeling complex system dynamics with flow
 592 matching across time and conditions. In *The Thirteenth International Conference on Learning
 593 Representations*, 2025.

594

595 Tianrong Chen, Guan-Horng Liu, Molei Tao, and Evangelos Theodorou. Deep momentum multi-
 596 marginal schrödinger bridge. *Advances in Neural Information Processing Systems*, 36:57058–
 597 57086, 2023.

598

599 Volker Bergen, Marius Lange, Stefan Peidli, F. Alexander Wolf, and Fabian J. Theis. Generalizing
 600 rna velocity to transient cell states through dynamical modeling. *Nature Biotechnology*, 38(12):
 601 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>.

594 **A PRIOR-GUIDED HJB INEQUALITY DERIVATION**
 595

596 To derive the velocity-prior guided HJB residual inequality in Eq.8.

597
$$\begin{aligned} 598 \quad \mathcal{A}(\rho, v, g) &\geq \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \\ 599 &\quad - \int_0^1 \int_{\Omega} \rho_t(x) \left(\partial_t s + \frac{1}{2} \|\nabla s\|^2 + \nabla s \cdot v_{\text{prior}} + \frac{1}{2} s^2 \right) dx dt. \end{aligned} \quad (16)$$

 600

601 We want to start by minimizing Definition 3.2. Since v_{prior} is constant and known, we only need to
 602 minimize the "learned" action portion v_{corr} , then the problem becomes the following:
 603

604 Consider the velocity-prior guided WFR action:

605
$$\begin{aligned} 606 \quad \min_{\rho, v_{\text{ot}}, g} \quad \mathcal{A}(\rho, v_{\text{corr}}, g) &:= \int_0^1 \int_{\Omega} \left(\frac{1}{2} \|v_{\text{corr}}(t, x)\|^2 + \frac{1}{2} g(t, x)^2 \right) \rho_t(x) dx dt \\ 607 \quad \text{s.t.} \quad \partial_t \rho_t + \nabla \cdot (\rho_t(v_{\text{prior}} + v_{\text{corr}})) &= g_t \rho_t, \quad \rho|_{t=0} = \rho_0, \quad \rho|_{t=1} = \rho_1. \end{aligned} \quad (17)$$

 608

609 **Step 1: Lagrangian formulation**

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

611
$$\begin{aligned} 612 \quad \mathcal{L} &= \int_0^1 \int_{\Omega} \left(\frac{1}{2} \|v_{\text{corr}}\|^2 + \frac{1}{2} g^2 \right) \rho dx dt \\ 613 &\quad + \int_0^1 \int_{\Omega} s \left(\partial_t \rho + \nabla \cdot (\rho(v_{\text{prior}} + v_{\text{corr}})) - g \rho \right) dx dt. \end{aligned} \quad (18)$$

 614

615 **Step 2: Integration by parts**

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

617
$$\begin{aligned} 618 \quad \int_0^1 \int_{\Omega} s \partial_t \rho dx dt &= \left[\int_{\Omega} s \rho dx \right]_{t=0}^{t=1} - \int_0^1 \int_{\Omega} \rho \partial_t s dx dt \\ 619 &= \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] - \int_0^1 \int_{\Omega} \rho \partial_t s dx dt. \end{aligned} \quad (19)$$

 620

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

622
$$\int_{\Omega} s \nabla \cdot (\rho w) dx = \int_{\partial \Omega} s \rho w \cdot n d\sigma - \int_{\Omega} \nabla s \cdot (\rho w) dx. \quad (20)$$

 623

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

625
$$\int_0^1 \int_{\Omega} s \nabla \cdot (\rho w) dx dt = - \int_0^1 \int_{\Omega} \rho w \cdot \nabla s dx dt. \quad (21)$$

 626

627 Then, the Lagrangian becomes

628
$$\begin{aligned} 629 \quad \mathcal{L} &= \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \\ 630 &\quad + \int_0^1 \int_{\Omega} \rho_t(x) \left[\frac{1}{2} \|v_{\text{corr}}\|^2 - v_{\text{corr}} \cdot \nabla s + \frac{1}{2} g^2 - s g - \partial_t s - \nabla s \cdot v_{\text{prior}} \right] dx dt. \end{aligned} \quad (22)$$

 631

632 **Step 3: Fenchel–Young inequality**

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

634
$$\frac{1}{2} \|a\|^2 \geq p \cdot a - \frac{1}{2} \|p\|^2 \quad (23)$$

 635

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

637
$$\frac{1}{2} \|v_{\text{corr}}\|^2 - v_{\text{corr}} \cdot \nabla s \geq -\frac{1}{2} \|\nabla s\|^2 \quad (24)$$

 638

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

640
$$\frac{1}{2} g^2 - s g \geq -\frac{1}{2} s^2, \quad (25)$$

 641

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

643
$$\begin{aligned} 644 \quad \mathcal{A}(\rho, v_{\text{corr}}, g) &\geq \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \\ 645 &\quad - \int_0^1 \int_{\Omega} \rho_t(x) \left(\partial_t s + \frac{1}{2} \|\nabla s\|^2 + \nabla s \cdot v_{\text{prior}} + \frac{1}{2} s^2 \right) dx dt. \end{aligned} \quad (26)$$

 646

648 B PROOF OF THEOREM 3.4

649
650 **Theorem B.1** (Prior-guided HJB optimality). *Suppose that the HJB residual defined in corollary 3.3
651 satisfies $r_\theta(t, x) = 0$ for ρ_t -a.e on $[0, 1] \times \mathbb{R}^d$, and the boundary constraints hold, then $(\rho_t, v_{\text{corr}}, g_\theta)$
652 satisfies the unbalanced continuity equation and the WFR optimality conditions in Definition 3.2. In
653 particular, the learned corrective field $v_{\text{corr}}^* = \nabla_x s_\theta$ and growth $g^* = g_\theta$ satisfy the optimality
654 conditions.*

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

$$658 \quad \mathcal{A}(\rho, v_{\text{corr}}, g) \geq \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] - \int_0^1 \int_{\Omega} \rho_t(x) r_s(t, x) dx dt. \quad (27)$$

661 where the HJB residual is defined as:

$$662 \quad r_s(t, x) := \partial_t s + \frac{1}{2} \|\nabla_x s\|^2 + \nabla_x s \cdot v_{\text{prior}} + \frac{1}{2} s^2. \quad (28)$$

664 Step 1 (Dual constraint and lower bound)

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

$$666 \quad r_s(t, x) \geq 0 \quad \forall (t, x). \quad (29)$$

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

$$669 \quad \mathcal{A}(\rho, v_{\text{corr}}, g) \geq \mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \quad \forall s \quad (30)$$

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

$$672 \quad \inf_{\rho, v_{\text{corr}}, g} \mathcal{A}(\rho, v_{\text{corr}}, g) \geq \sup_{s: r_s \geq 0} \left(\mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \right). \quad (31)$$

675 Step 2 (Residual Optimality)

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

$$677 \quad r_{s_\theta}(t, x) = 0 \quad \text{for } \rho_t\text{-a.e.}, \quad (32)$$

$$678 \quad v_{\text{corr}} = \nabla_x s_\theta, \quad g = s_\theta, \quad (33)$$

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

680 Then the Fenchel–Young inequalities are equalities and equation 27 becomes:

$$683 \quad \mathcal{A}(\rho, v_{\text{corr}}, g) = \mathbb{E}_{x \sim \rho_0}[s_\theta(0, x)] - \mathbb{E}_{x \sim \rho_1}[s_\theta(1, x)] - \int_0^1 \int_{\Omega} \rho_t(x) r_{s_\theta}(t, x) dx dt. \quad (34)$$

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

$$687 \quad \mathcal{A}(\rho, v_{\text{corr}}, g) = \mathbb{E}_{x \sim \rho_0}[s_\theta(0, x)] - \mathbb{E}_{x \sim \rho_1}[s_\theta(1, x)]. \quad (35)$$

689 Combining this with the dual lower bound, we have:

$$690 \quad \inf_{\rho, v_{\text{corr}}, g} \mathcal{A}(\rho, v_{\text{corr}}, g) < \mathcal{A}(\rho, v_{\text{corr}}, g) \quad (36)$$

$$692 \quad \mathbb{E}_{x \sim \rho_0}[s_\theta(0, x)] - \mathbb{E}_{x \sim \rho_1}[s_\theta(1, x)] < \sup_{s: r_s \geq 0} \left(\mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \right). \quad (37)$$

694 So we have:

$$696 \quad \mathcal{A}(\rho, v_{\text{corr}}, g) = \inf_{\rho, v_{\text{corr}}, g} \mathcal{A}(\rho', v'_{\text{corr}}, g') = \sup_{s: r_s \geq 0} \left(\mathbb{E}_{x \sim \rho_0}[s_0(x)] - \mathbb{E}_{x \sim \rho_1}[s_1(x)] \right), \quad (38)$$

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

$$699 \quad v_{\text{corr}}^* = \nabla_x s_\theta, \quad g^* = s_\theta,$$

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

702 Table 5: Comparison on the EB dataset using W_2 at the held-out marginals (t_1, t_3) . Baseline results
 703 other than * are taken from (Theodoropoulos et al., 2025).

Method	W_2	t_1	W_2	t_3
DeepRUOT	13.64		15.10	
Var-RUOT*	10.34		12.02	
MIOFlow	13.66		14.05	
SBIRR	15.42		20.98	
MMFM	14.68		14.83	
DMSB	14.83		15.49	
3MSBM	14.51		13.26	
VP-HJF (ours)*	11.94		12.28	

715 Table 6: Comparison on the EB dataset using weighted W_1 at marginals t_1-t_4 for Var-RUOT and
 716 VP-HJF.

Method	W_1^{weighted}			
	t_1	t_2	t_3	t_4
Var-RUOT*	10.28	11.58	11.90	13.28
VP-HJF (ours)*	11.14	12.45	13.14	14.45

C ADDITIONAL EXPERIMENTS

727 We conducted further comparison on the EB dataset with 100-dim using the \mathcal{W}_2 metric at the held-
 728 out marginals at t_1, t_3 . Table 5 shows that our approach remains competitive and outperforms most
 729 methods. We then compare more directly with VAR-RUOT in Table 6 using the weighted \mathcal{W}_1 ,
 730 since both methods are in the unbalanced optimal transport setting. In this evaluation, we assign
 731 non-uniform particle weights by integrating the learned dynamics through ODE integration and use
 732 these predicted weights when computing \mathcal{W}_1 , instead of uniform masses.

733 On this weighted \mathcal{W}_1 metric, Var-RUOT achieves slightly lower values than VP-HJF. This gap
 734 could partly due to the use of noisy RNA-velocity as v_{prior} in our framework, which can trade a
 735 small increase in transport cost for better agreement with the measured dynamics. In addition, Var-
 736 RUOT optimizes a single global-in-time trajectory via SDE simulations, whereas VP-HJF relies on
 737 deterministic ODE rollouts with local per-interval supervision.

738 **further training details on single-cell datasets** For both the EB and bone marrow datasets, we
 739 use a 4-layer MLP with Swish activation. The MLP outputs follows the Action Matching implemen-
 740 tation, where the output $h_\theta(t, x)$ needs to multiply by the original data input x so the scalar output
 741 becomes $s = (h \times x).\text{sum}()$ Moreover, We optimize the model with Adam and set the learning rate
 742 to $2e-4$ for both datasets. We follow Algorithm 1 to train with 300 epochs and 256 batch size,
 743 using dopri5 with 16 steps for the ode integration. We set the HBJ loss coefficients to $\lambda_{\text{hbj}} = 0.01$,
 744 sliced Wasserstein loss coefficient $\lambda_{\text{sw}} = 10$ and the mass loss coefficient to $\lambda_{\text{mass}} = 0.01$.

745 For the Var-RUOT baseline on EB, we use the authors' publicly released implementation and con-
 746 figuration, changing only the training epoch to 500 epochs.

D LLM USAGE

750 We used LLM for improve on writing, mainly for checking grammar. We also used LLM for finding
 751 relevant and related works.