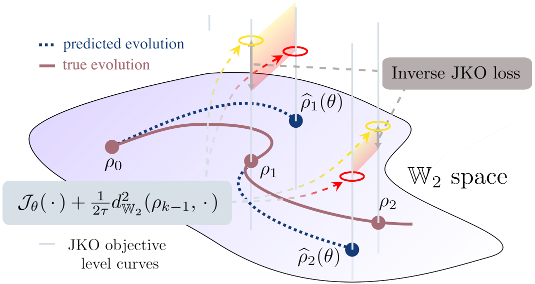

Learning of Population Dynamics: Inverse Optimization Meets JKO Scheme

---

Mikhail Persiianov[1] **Jiawei Chen**[1,2] **Petr Mokrov**[1] **Alexander Tyurin**[3,1]
Evgeny Burnaev[1,3] **Alexander Korotin**[1,3]

[1]Applied AI Institute [2]Moscow Center for Advanced Studies [3]AXXX

ICLR 2026

The Challenge: Unpaired Population Dynamics

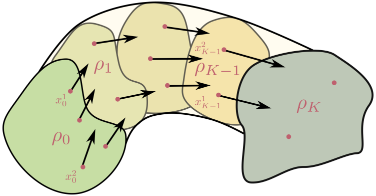

- **The Goal:** Infer underlying stochastic dynamics from marginal distributions ρ_t observed at discrete times $\{\rho_k\}_{k=0}^K$.
- **The Catch:** Individual particle trajectories are unavailable (e.g., single-cell genomics requires destructive sampling).
- **The JKO Scheme**[1]**:** Models evolution as a sequence $\{\rho_k^\tau\}_{k=0}^K$ of distributions approaching minimum energy $\mathcal{J} : \mathcal{P}(\mathcal{X}) \to \mathbb{R}$ while staying close to the previous state:

$$\rho_{k+1}^\tau = \arg\min_{\rho \in \mathcal{P}(\mathcal{X})} \left\{ \mathcal{J}(\rho) + \frac{1}{2\tau} d_{\mathbb{W}_2}^2(\rho, \rho_k^\tau) \right\} \stackrel{\text{def}}{=} \text{JKO}_{\tau\mathcal{J}}(\rho_k^\tau), \quad \rho_0^\tau = \rho_0, \tag{1}$$

where $d_{\mathbb{W}_2}$ is the *Wasserstein-2 distance* between two probability measures $\mu, \nu \in \mathcal{P}(\mathcal{X})$:

$$d_{\mathbb{W}_2}^2(\mu, \nu) = \min_{T : T\sharp\mu = \nu} \int_{\mathcal{X}} \|x - T(x)\|_2^2 \, \mathrm{d}\mu(x). \tag{2}$$

[1] Richard Jordan, David Kinderlehrer, and Felix Otto (1998). **"The variational formulation of the Fokker–Planck equation".** In: SIAM journal on mathematical analysis 29.1, pp. 1–17

Limitations of Prior JKO Methods (1/2). JKOnet[2]

- **Bi-level optimization (nested problem):**

$$\mathcal{L}_{\texttt{JKOnet}}(\theta) = \sum_{k=0}^{K-1} \widehat{d_{\mathbb{W}_2}}^2(\hat{\rho}_k, \rho_k), \quad \hat{\rho}_{k+1} = \nabla \psi_k^* \sharp \hat{\rho}_k \tag{3}$$

$$\psi_k^* = \arg\min_{\psi_\varphi \in \texttt{CVX}} \left[\mathcal{J}_\theta(\nabla \psi_\varphi \sharp \hat{\rho}_k) + \frac{1}{2\tau} \|x - \nabla \psi_\varphi(x)\|^2 \right] \tag{4}$$

- **Heavy computation:**
  - Inner JKO optimization must be solved at each time step k
  - **Optimizer unrolling**: all iterates $\varphi_k^0 \to \cdots \to \varphi_k^N$ need to be stored for backpropagation to θ
  - Each loss $\widehat{d_{\mathbb{W}_2}}^2$ requires iterative Sinkhorn (discrete OT) computations and gradients
- **Limited modeling:**
  - Only potential energy $\mathcal{V}_\theta(\rho) = \int V_\theta(x)\, d\rho(x)$
  - Missing interaction / internal (e.g., entropy) terms
  - Restricted to convex architectures (ICNNs)

[2] Charlotte Bunne et al. (2022). **"Proximal optimal transport modeling of population dynamics".** In: International Conference on Artificial Intelligence and Statistics. PMLR, pp. 6511–6528.

Limitations of Prior JKO Methods (2/2). `JKOnet*`[3]

- **First-order objective:**

$$\mathcal{L}_{\texttt{JKOnet}^*}(\theta) = \sum_{k=0}^{K-1} \int_{\mathcal{X}\times\mathcal{X}} \Big\| \nabla V_{\theta_1}(x_{k+1}) + \int_{\mathcal{X}} \nabla W_{\theta_2}(x_{k+1} - y_{k+1})\,\mathrm{d}\rho_{k+1}(y_{k+1}) + \theta_3 \frac{\nabla \rho_{k+1}(x_{k+1})}{\rho_{k+1}(x_{k+1})} + \frac{1}{\tau}(x_{k+1} - x_k) \Big\|^2 \mathrm{d}\pi_k(x_k, x_{k+1}). \tag{5}$$

- **Non-End-to-End:** Requires *precomputed* OT plans π_k; cannot learn plans jointly with the energy
- **Discretization Error:** Computing π_k with discrete OT solvers is challenging in high-dimensional settings, which can degrade accuracy
- **Rigid Structure:** Strongly depends on analytical first-order optimality conditions; limits flexibility in modeling

Goal: End-to-end JKO learning with expressive energies and no precomputed plans.

[3]Antonio Terpin et al. (2024). **"Learning diffusion at lightspeed".** In: Advances in Neural Information Processing Systems 37, pp. 6797–6832.

`iJKOnet`: Methodology & Loss Derivation

Goal: Learn energy $\mathcal{J}_\theta$ such that $\rho_{k+1} \approx \text{JKO}_{\tau\mathcal{J}_\theta}(\rho_k)$.

- **The Inverse Gap:** For any candidate $\mathcal{J}$ and ground-truth snapshots ρ_k, ρ_{k+1}:

$$\min_{\rho} \left\{ \mathcal{J}(\rho) + \frac{1}{2\tau} d^2_{\mathbb{W}_2}(\rho_k, \rho) \right\} \leq \mathcal{J}(\rho_{k+1}) + \frac{1}{2\tau} d^2_{\mathbb{W}_2}(\rho_k, \rho_{k+1}) \tag{6}$$

  - LHS = optimal JKO step under $\mathcal{J}$
  - RHS = observed transition under $\mathcal{J}^*$ ($\rho_k \to \rho_{k+1}$)
  - Maximizing the difference between the LHS and RHS of (6) drives $\mathcal{J} \to \mathcal{J}^*$:

`iJKOnet` objective (adversarial training)

$$\mathcal{L}(\theta, \phi) = \sum_{k=0}^{K-1} \left[\mathcal{J}_\theta(T^k_\varphi \sharp \rho_k) - \mathcal{J}_\theta(\rho_{k+1}) + \frac{1}{2\tau} \int_{\mathcal{X}} \|x - T^k_\varphi(x)\|_2^2 \rho_k(x) dx \right] \to \max_{\mathcal{J}_\theta} \min_{T^k_\varphi} \tag{7}$$

Parametrization:

- T^k_φ: learns the **best competing transport** (approximate JKO minimizer)
- $\mathcal{J}_\theta$: adjusts energy to make observed transitions appear optimal
- **MLPs used for T^k_φ**, unlike prior work: no ICNNs or precomputed OT plans

Theoretical Aspects: Quality Bounds

Question: Does minimizing the inverse JKO loss recover the true energy $\mathcal{J}^*$?

Theorem 3.1 (informal for $K = 1$)

Let $\varepsilon(V) \stackrel{\text{def}}{=} \mathcal{L}(V^*, T_{V^*}) - \mathcal{L}(V, T_V)$ be the gap between the optimal and optimized value of inverse JKO loss (7) with internal $\min_T$ problem solved exactly, i.e., $T_V \stackrel{\text{def}}{=} \min_T \mathcal{L}(V, T)$. Assume $\mathcal{X}$ is convex and V satisfies mild convexity and smoothness conditions. Then there exists a constant $C = C(\tau, \beta)$ such that

$$\int_{\mathcal{X}} \left\|\nabla V^*(y) - \nabla V(y)\right\|^2 \mathrm{d}\rho_1(y) \leq C\varepsilon(V). \tag{8}$$

- **Meaning:** Small gap $\varepsilon(V) \Rightarrow \nabla V \approx \nabla V^*$ (V recovered up to an additive constant)
- **Scope:** Applies to **potential energy only**

Takeaway: Minimizing the inverse JKO loss recovers the <u>driving force</u> ∇V^* of the dynamics.

Experimental Results: Synthetic 2D Benchmark

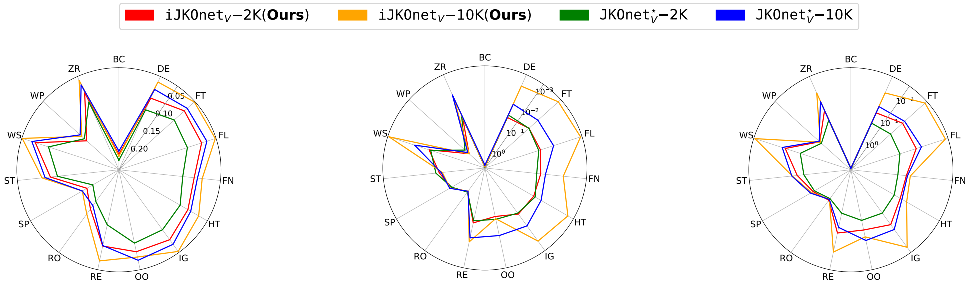

(a) EMD　　**(b)** B$d^2_{\mathbb{W}_2}$-UVP　　**(c)** $\mathcal{L}_2$-UVP

Setup: Recovery of known potential energies (V) in a 2D space.

- **Unpaired vs. Paired Setup:**
  - Population setting requires *independent samples* $x_k \sim \rho_k$ at each time step k
  - Prior code uses **paired trajectories** ($x_k \leftrightarrow x_{k+1}$), introducing temporal correlation
  - We correct this by regenerating samples (**unpaired** setup)
- **Impact:** Setup choice significantly affects performance
- `iJKOnet` outperforms `JKOnet`*[3] across nearly all tested potentials.

Single-Cell Dynamics: 5D Leave-Two-Out

- **Dataset:** Embryoid Body (EB) single-cell data[a], 5 timesteps $\mathbf{t_0} - \mathbf{t_4}$
- **Setup:**
  - Train on snapshots $\mathbf{t_0}, \mathbf{t_2}, \mathbf{t_4}$
  - Predict missing timesteps $\mathbf{t_1}, \mathbf{t_3}$ (leave-two-out)
- **Metric:** $d_{\mathbb{W}_2}$ distance (↓)

Result

`iJKOnet`$_{t,V}$ achieves the best performance on both missing timepoints t_2 and t_4, outperforming all JKO-based and non-JKO baselines.

[a] Kevin R Moon et al. (2019). **"Visualizing structure and transitions in high-dimensional biological data".** In: Nature biotechnology 37.12, pp. 1482–1492

Method	$\mathbf{t_1}$	$\mathbf{t_3}$
Vanilla-SB Vargas et al. 2021	1.49 ± 0.063	1.55 ± 0.034
DMSB Chen et al. 2023	1.13 ± 0.082	1.45 ± 0.16
TrajectoryNet Tong et al. 2020	2.03 ± 0.04	1.93 ± 0.08
MMSB Shen et al. 2025	1.27 ± 0.028	1.57 ± 0.048
Static		
JKOnet$^*_V$	1.145 ± 0.033	2.529 ± 0.014
JKOnet$^*_{V+U}$	1.099 ± 0.119	2.537 ± 0.054
JKOnet$^*_{V+W}$	1.419 ± 0.173	2.510 ± 0.094
JKOnet$^*_{W+U}$	1.887 ± 0.017	1.739 ± 0.037
JKOnet*	1.361 ± 0.257	2.557 ± 0.042
Static (Ours)		
iJKOnet$_V$	1.082 ± 0.011	1.147 ± 0.001
iJKOnet$_{V+U}$	*1.065 ± 0.018*	*1.150 ± 0.004*
iJKOnet$_{V+W}$	2.865 ± 0.166	1.376 ± 0.015
iJKOnet$_{W+U}$	1.649 ± 0.005	0.868 ± 0.005
iJKOnet	3.577 ± 0.166	1.395 ± 0.032
Time-varying		
JKOnet$^*_{t,V}$	4.414 ± 1.499	2.771 ± 0.197
iJKOnet$_{t,V}$ **(Ours)**	$\mathbf{0.983 \pm 0.037}$	$\mathbf{0.849 \pm 0.021}$

Single-Cell Dynamics: 100D Leave-One-Out

- **Dataset:** EB single-cell data (100D)
- **Setup (leave-one-out):**
  - Remove one timestep (t_1, t_2, or t_3)
  - Reconstruct the omitted distribution
- **Metric:** MMD distance (↓)

Result

iJKOnet$_V$ achieves performance on par with **DMSB** in the full-data setting (w/o LO), while using a **simpler, simulation-free** training procedure (no trajectory caching) and faster execution.

Method	LO-$\mathbf{t_1}$	LO-$\mathbf{t_2}$	LO-$\mathbf{t_3}$	w/o LO
NLSB Koshizuka and Sato 2023	0.38	0.37	0.37	0.66
MIOFLOW Huguet et al. 2022	0.23	0.90	0.23	0.23
DMSB Chen et al. 2023	$\mathbf{0.042 \pm 0.020}$	$\mathbf{0.033 \pm 0.003}$	$\mathbf{0.040 \pm 0.020}$	$\mathbf{0.032 \pm 0.003}$
JKOnet$^{\star}_{V}$ Terpin et al. 2024	0.220 ± 0.025	0.293 ± 0.018	0.235 ± 0.006	0.229 ± 0.052
iJKOnet$_V$ **(Ours)**	*0.137 ± 0.001*	*0.123 ± 0.001*	0.097 ± 0.002	*0.085 ± 0.024*
JKOnet$^{\star}_{t,V}$ Terpin et al. 2024	0.575 ± 0.119	0.619 ± 0.157	0.456 ± 0.056	0.477 ± 0.098
iJKOnet$_{t,V}$ **(Ours)**	0.848 ± 0.043	0.370 ± 0.233	*0.055 ± 0.007*	0.124 ± 0.243

Conclusion

- `iJKOnet` bridges inverse optimization and computational gradient flows.
- Provides a scalable, end-to-end framework for recovering energy functionals.
- Establishes theoretical bounds for JKO-based potential energy recovery.
- Outperforms existing JKO-based methods on complex synthetic and real-world biological datasets.

Thank You!

Source code available at:
`https://github.com/MuXauJl11110/iJKOnet`