

# 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 OF POPULATION DYNAMICS: INVERSE OPTIMIZATION MEETS JKO SCHEME

Anonymous authors

Paper under double-blind review

## ABSTRACT

Learning population dynamics involves recovering the underlying process that governs particle evolution, given evolutionary snapshots of samples at discrete time points. Recent methods frame this as an energy minimization problem in probability space and leverage the celebrated JKO scheme for efficient time discretization. In this work, we introduce  $iJKOnet$ , an approach that combines the JKO framework with inverse optimization techniques to learn population dynamics. Our method relies on a conventional *end-to-end* adversarial training procedure and does not require restrictive architectural choices, e.g., input-convex neural networks. We establish theoretical guarantees for our methodology and demonstrate improved performance over prior JKO-based methods.

## 1 INTRODUCTION

Modeling population dynamics is a fundamental challenge in many scientific domains, including biology (Schiebinger et al., 2019; Moon et al., 2019), ecology (Ayala et al., 1973), meteorology (Fisher et al., 2009; Sigrist et al., 2015; Verma et al., 2024; Price et al., 2025), transportation flows in urban networks (Medina-Salgado et al., 2022), and epidemiology (Wang et al., 2021; Kosma et al., 2023), among others. The task is to infer the underlying stochastic dynamics of a system – typically modeled by stochastic differential equations (SDEs) – from observed marginal distributions at discrete time points. While this problem has been studied extensively in settings where individual trajectories are available (Krishnan et al., 2017; Li et al., 2020; Brogat-Motte et al., 2024), such data are often unavailable in practice. In many real-world scenarios, we only observe *population-level data*, where it is infeasible to continuously track individual entities. Instead, we are forced to rely on temporally separated and mutually independent snapshots of the population.

In *single-cell genomics* (Macosko et al., 2015), for example, measuring the state of individual cells typically involves destructive sampling – a process where cells are destroyed during measurement, preventing any further observation of their future behavior. As a result, the data consist only of isolated profiles of cellular populations taken at discrete time points. Reconstructing the continuous developmental trajectories of cells from such fragmented data poses a challenge. A similar situation arises in *financial markets* (Gontis et al., 2010; Yang et al., 2023), where analysts often have access only to marginal distributions of asset prices at specific times, making it necessary to infer the underlying dynamics that govern these distributions. In *crowd dynamics* (Maury et al., 2010; 2011; Wan et al., 2023), modeling the temporal evolution of population densities is crucial for understanding and managing pedestrian flows. In this case, data on individual trajectories are also typically unavailable, and only aggregate distributions at various time points are observed.

A promising approach to address the absence of particle trajectories is the Jordan–Kinderlehrer–Otto (JKO) scheme (Jordan et al., 1998), which models the evolution of a particle system as a sequence of distributions that gradually approach the minimum of a total energy functional while remaining

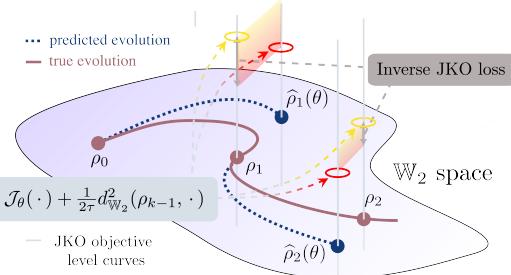


Figure 1:  $iJKOnet$  working scheme: our method minimizes the **gap** between the **optimal** values of parametric ( $\theta$ ) JKO functional and **suboptimal** values obtained at ground truth population measures.

close to the previous distributions. However, implementing the JKO scheme involves solving an optimization problem over the space of probability measures, which is computationally demanding. The first attempt to leverage JKO for learning population dynamics was introduced in (Bunne et al., 2022b, JKOnet). While innovative, this approach relies on a complex learning objective and is limited to potential energy functionals – meaning it cannot capture stochasticity in the dynamics. A recent work, (Terpin et al., 2024, JKOnet\*), proposes replacing the JKO optimization step with its first-order optimality conditions. This relaxation allows modeling more general energy functionals and reduces computational complexity, but it does not support end-to-end training. Instead, it requires precomputing optimal transport couplings (Cuturi, 2013) between subsequent time snapshots, which limits its scalability and generalization.

**Contributions.** In this work, we present a novel approach for recovering the system dynamics based on observed population-level data. Our key contributions are as follows:

1. **Methodology:** In §3.2, we cast the problem of reconstructing energy functionals within the JKO framework as an inverse optimization task. This perspective leads to a novel min-max optimization objective for population dynamics recovery.
2. **Algorithm:** We equip our population dynamics methodology with a conventional end-to-end adversarial learning procedure (§3.3). Importantly, our practical scheme does not pose restrictions on the architectures of the utilized neural networks, which contributes to the scalability.
3. **Theoretical Guarantees:** Under suitable assumptions, we show that our method can accurately recover the underlying energy functional governing the observed population dynamics (§ 3.4).

In §5 and Appendix B, we evaluate our approach on a range of synthetic and real-world datasets, including single-cell genomics. The results show that our method demonstrates improved performance over previous JKO-based approaches for learning population dynamics.

**Notation.** Let  $\mathcal{X}$  be a compact subset of  $\mathbb{R}^D$  equipped with the Euclidean norm  $\|\cdot\|_2$ . Let  $\mathcal{P}(\mathcal{X})$  denote the set of probability measures on  $\mathcal{X}$ , and let  $\mathcal{P}_{ac}(\mathcal{X})$  denote its subset of probability measures absolutely continuous with respect to the Lebesgue measure. For  $\rho \in \mathcal{P}_{ac}(\mathcal{X})$ , we use  $\rho$  to denote both the measure and its density function with respect to the Lebesgue measure. For measures  $\mu, \nu \in \mathcal{P}(\mathcal{X})$ , we denote the set of couplings (*transportation plans*) between them by  $\Pi(\mu, \nu)$ . For a measure  $\rho \in \mathcal{P}(\mathcal{X})$  and measurable map  $T : \mathcal{X} \rightarrow \mathcal{X}$ , we denote by  $T\#\rho$  the associated *push-forward* measure, and  $\text{id}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{X}$  is the identity mapping.  $\nabla \cdot F = \sum_{d=1}^D \frac{\partial F_d}{\partial x_d}$  denotes the *divergence* operator for a continuously differentiable vector field  $F : \mathcal{X} \rightarrow \mathcal{X}$ .

## 2 BACKGROUND

To describe the evolution of population measures, we use the theoretical framework of *Wasserstein gradient flows* (WGFs). In what follows, we first introduce the preliminary concept of Optimal Transport between probability measures (Villani et al., 2008; Santambrogio, 2015). Then we provide sufficient theoretical background on WGFs; see (Santambrogio, 2017; Figalli & Glaudo, 2023) for an overview and (Ambrosio et al., 2008) for a comprehensive study of WGFs theory. Finally, we get the reader acquainted with the JKO scheme, which is the cornerstone of our approach.

**Optimal Transport.** The (squared) *Wasserstein-2 distance*  $d_{\mathbb{W}_2}$  between two probability measures  $\mu, \nu \in \mathcal{P}(\mathcal{X})$  is defined as the solution to the *Kantorovich* problem (Kantorovich, 1942):

$$d_{\mathbb{W}_2}^2(\mu, \nu) \stackrel{\text{def}}{=} \min_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{X}} \|x - y\|_2^2 d\pi(x, y), \quad (1)$$

where the distribution  $\pi^*$  delivering the minimum to (1) is called the *optimal coupling* (or *plan*). If  $\mu \in \mathcal{P}_{ac}(\mathcal{X})$ , then the Brenier’s theorem (Brenier, 1991) establishes equivalence between the Kantorovich formulation (1) and *Monge’s problem* (Monge, 1781):

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

where the optimal transport map  $T^*$  is known as the *Monge map*. Moreover, there exists a *unique* (up to an additive constant) convex potential  $\psi^* : \mathcal{X} \rightarrow \mathbb{R}$  such that  $T^* = \nabla \psi^*$  and  $(\nabla \psi^*)\#\mu = \nu$  (McCann, 1995). In this setting, the optimal coupling in (1) is given by  $\pi^* = [\text{id}_{\mathcal{X}}, \nabla \psi^*]\#\mu$ .

**Wasserstein Gradient Flows.** For clarity, we first consider the concept of gradient flows in Euclidean space  $\mathbb{R}^D$  and then move to the space of probability measures  $\mathcal{P}(\mathcal{X})$ . For Euclidean space  $\mathbb{R}^D$ , a *gradient flow* is an absolutely continuous curve  $x(t)$  starting at  $x_0 \in \mathbb{R}^D$  that minimizes

108 an energy functional  $J : \mathbb{R}^D \rightarrow \mathbb{R}$  "as fast as possible". To find such a curve, one needs to  
 109 solve an ordinary differential equation (ODE) (Teschl, 2012) of the form  $x'(t) = -\nabla J(x(t))$  with  
 110 initial condition  $x(0) = x_0$ . The same idea is applicable to the space of probability measures  
 111  $\mathcal{P}(\mathcal{X})$ . If we equip  $\mathcal{P}(\mathcal{X})$  with the Wasserstein-2 distance  $d_{\mathbb{W}_2}$ , we obtain the *Wasserstein space*  
 112  $\mathbb{W}_2(\mathcal{X}) = (\mathcal{P}(\mathcal{X}), d_{\mathbb{W}_2})$  – a complete and separable metric space (Bogachev & Kolesnikov, 2012).  
 113 In this case, for an energy functional  $\mathcal{J} : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$ , the gradient flow in  $\mathbb{W}_2(\mathcal{X})$ , called the  
 114 *Wasserstein gradient flow* (WGF), is an absolutely continuous curve  $\rho_t : \mathbb{R}_+ \rightarrow \mathcal{P}(\mathcal{X})$  starting at  $\rho_0$   
 115 that follows the steepest descent direction of  $\mathcal{J}$ , i.e., solves

$$116 \quad \partial_t \rho_t = -\nabla_{\mathbb{W}_2} \mathcal{J}(\rho_t), \quad \rho_{(t=0)} = \rho_0, \quad (3)$$

117 where  $\nabla_{\mathbb{W}_2} \mathcal{J}(\rho_t)$  denotes the *Wasserstein gradient* in  $\mathbb{W}_2$  given by  $\nabla_{\mathbb{W}_2} \mathcal{J}(\rho) = -\nabla \cdot (\rho \nabla \frac{\delta \mathcal{J}}{\delta \rho})$  with  
 118  $\frac{\delta \mathcal{J}}{\delta \rho}$  denoting the first variation (Chewi et al., 2025) of energy  $\mathcal{J}$ . Thus, equation (3) can be rewritten  
 119 in the form of the *continuity equation*, expressing mass conservation under the velocity field  $v_t$ :  
 120

$$121 \quad \partial_t \rho_t + \nabla \cdot (\rho_t v_t) = 0, \quad v_t = -\nabla \frac{\delta \mathcal{J}}{\delta \rho}(\rho_t). \quad (4)$$

123 There exists an intriguing connection between WGFs and certain partial differential equations  
 124 (PDEs) (Evans, 2022). In particular, different energy functionals in the Wasserstein space corre-  
 125 spond to distinct PDEs (Santambrogio, 2015), some of which are associated with diffusion processes  
 126 appearing in practice (Gómez-Castro, 2024; Bailo et al., 2024b).

127 **Examples of PDEs as WGFs.** Consider the *free energy* functional (Carrillo et al., 2003, Eq. (1.3)):

$$128 \quad \mathcal{J}_{\text{FE}}(\rho) = \underbrace{\int_{\mathcal{X}} V(x) d\rho(x)}_{\mathcal{V}(\rho)} + \underbrace{\int_{\mathcal{X} \times \mathcal{X}} W(x-y) d\rho(x)d\rho(y)}_{\mathcal{W}(\rho)} + \underbrace{\int_{\mathcal{X}} U(\rho(x)) dx}_{\mathcal{U}(\rho)}, \quad (5)$$

132 where  $\mathcal{V}$ ,  $\mathcal{W}$ , and  $\mathcal{U}$  correspond to the system's *potential*, *interaction*, and *internal* energies, respec-  
 133 tively. The PDE corresponding to the WGF driven by this functional is known as the *aggregation-  
 134 diffusion* equation (ADE). It describes the evolution of density  $\rho_t$  under a corresponding velocity  
 135 field  $v_t$  (4). This evolution reflects a balance of three effects: *drift* – driven by the potential  $V(x)$ ,  
 136 modeling an external field (e.g., gravity, electric potential); *interaction* – governed by the symmetric  
 137 interaction kernel  $W(x-y)$ , accounting for non-local effects (e.g., particle interactions, long-range  
 138 forces); and *diffusion* – modeled by the internal energy  $\mathcal{U}(\rho)$ , representing the energy associated  
 139 with the local state of the system (e.g., thermal energy, chemical energy). A *weak* solution to the  
 140 equation exists if the *no-flux conditions*:  $\nabla V \cdot \vec{n} = 0$  on  $\partial \mathcal{X}$  holds, where  $\vec{n}$  is the outward normal  
 141 vector on the boundary, and if  $W \in \mathcal{C}(\mathcal{X} \times \mathcal{X})$ , i.e.  $W$  is a continuous function on  $\mathcal{X} \times \mathcal{X}$ . These  
 142 conditions ensure that no mass crosses the boundary of  $\mathcal{X}$ ; mass can only be redistributed within  $\mathcal{X}$ .  
 143 For a detailed discussion of the existence and classification of solutions, see (Gómez-Castro, 2024).

144 Such energy functionals are ubiquitous in real-world applications, including physics (Carrillo &  
 145 Gvalani, 2021), biology (Keller & Segel, 1971; Potts & Painter, 2024; Potts, 2024), economics (Fi-  
 146 aschi & Ricci, 2025), machine learning (Suzuki et al., 2023; Chizat et al., 2024; Nitanda, 2024), and  
 147 nonlinear optimization (Bailo et al., 2024a), to name a few. For further discussion and references,  
 148 see the recent surveys (Carrillo et al., 2019a; Gómez-Castro, 2024; Bailo et al., 2024b).

149 To demonstrate the flexibility of the energy formulation in (5), we examine a few representative  
 150 cases. When the system's energy is purely internal and given by the negative differential entropy,  
 151  $\mathcal{U}(\rho) = -\mathcal{H}(\rho) \stackrel{\text{def}}{=} \int \rho(x) \log \rho(x) dx$ , continuity equation (4) reduces to the classical heat equation:  
 152  $\partial_t \rho = \nabla^2 \rho$  (Vázquez, 2017). Alternatively, when the energy is  $\mathcal{J}_{\text{FP}}(\rho) = \mathcal{V}(\rho) - \beta \mathcal{H}(\rho)$ , the  
 153 resulting PDE is the *linear Fokker-Planck equation* with diffusion coefficient  $\beta$ :

$$154 \quad \partial_t \rho_t = \nabla \cdot (\nabla V(x) \rho_t) + \beta \nabla^2 \rho_t, \quad (6)$$

155 A fruitful branch of theoretical and practical research stems from the connection between the Fokker-  
 156 Planck PDE and stochastic differential equations (SDEs) (Risken & Frank, 1996; Bogachev et al.,  
 157 2022), the latter describes the stochastic evolution of *particles*. In particular, equation (6) is equiva-  
 158 lent (Weinan et al., 2021) to the following *Itô SDE*:

$$159 \quad dX_t = -\nabla V(X_t) dt + \sqrt{2\beta} dW_t, \quad (7)$$

160 where  $X = \{X_t\}_{t \geq 0}$  is a  $\mathbb{R}^D$ -valued stochastic process and  $W = \{W_t\}_{t \geq 0}$  is a standard Wiener  
 161 process (Särkkä & Solin, 2019). In other words, if  $X_t \sim \rho_t$  evolves according to (7), then the density  
 162  $\rho_t$  evolves according to the Fokker-Planck equation (6) in the space of probability measures.

162 WGFs provide a compelling framework for modeling PDEs and their associated SDEs across various  
 163 domains, giving rise to a range of methods for approximating solutions to WGFs. These include deep  
 164 learning-based approaches (Mokrov et al., 2021; Alvarez-Melis et al., 2022; Altekrüger et al., 2023),  
 165 particle-based methods such as Stein Variational Gradient Descent (SVGD) (Liu & Wang, 2016)  
 166 and its extensions (Das & Nagaraj, 2023; Tankala et al., 2025), as well as classical discretization  
 167 techniques in Wasserstein space (Carrillo et al., 2022, see §1.2). Many of these advances build on  
 168 the celebrated Jordan-Kinderlehrer-Otto (JKO) scheme (Jordan et al., 1998), which we review next.

169 **The JKO Scheme.** A classical method for solving ODEs in Euclidean space is the *implicit Euler scheme*. For an (Euclidean) gradient flow  $x'(t) = -\nabla J(x(t))$  and  $\tau > 0$ , this scheme  
 170 approximately solves this ODE by iteratively applying the proximal operator (Parikh et al., 2014,  
 171 §1.1):  $x_{k+1} = \text{prox}_{\tau J}(x_k) = \arg \min_{x \in \mathbb{R}^D} \{ J(x) + \frac{1}{2\tau} \|x - x_k\|_2^2 \}$ . Compared to the explicit Euler  
 172 scheme, this approach offers improved numerical stability (Butcher, 2016). Jordan, Kinderlehrer,  
 173 and Otto (Jordan et al., 1998) extended this idea to the space of probability measures, introducing a  
 174 variational time discretization of the Fokker-Planck equation (6), now known as the *JKO scheme*:

$$175 \rho_{k+1}^\tau = \arg \min_{\rho \in \mathcal{P}(\mathcal{X})} \left\{ \mathcal{J}(\rho) + \frac{1}{2\tau} d_{\mathbb{W}_2}^2(\rho, \rho_k^\tau) \right\} = \text{JKO}_{\tau \mathcal{J}}(\rho_k^\tau), \quad \rho_0^\tau = \rho_0, \quad (8)$$

176 where  $\tau > 0$  is the time step. As  $\tau \rightarrow 0$ , the sequence  $\rho_k^\tau, k \in \mathbb{N}$  converges to the continuous  
 177 solution  $\rho_t$  of the Fokker-Planck equation (6). Later, the JKO scheme's convergence was gener-  
 178 alized (Ambrosio et al., 2008, Thm. 4.0.4) for the free energy functional  $\mathcal{J}_{\text{FE}}$  (5). Note that since  
 179  $d_{\mathbb{W}_2}^2(\rho, \rho_k^\tau) \geq 0$ , the energy is non-increasing along the sequence:  $\mathcal{J}(\rho_{k+1}^\tau) \leq \mathcal{J}(\rho_k^\tau)$ . This mono-  
 180 tonicity property is often utilized in JKO-based algorithms (Salim et al., 2020) in Wasserstein space.

### 181 3 IJKONET METHOD

182 We begin by formally stating the problem addressed in our work (§3.1), then we develop our method-  
 183 ology (§3.2). Finally, we discuss the practical (§3.3) and theoretical aspects (§3.4) of our approach.

#### 184 3.1 PROBLEM STATEMENT

185 In our work, we address the problem of recovering the underlying energy functional  $\mathcal{J}^*$  that gov-  
 186 erns the evolution of a density  $\rho_t \in \mathcal{P}_{ac}(\mathcal{X})$  in (4) based on marginal population measures (Bunne  
 187 et al., 2022b; Lavenant et al., 2024). Specifically, we are given independent samples from marginals  
 188  $\{\rho_k\}_{k=0}^K$  of the evolving distribution  $\rho_t$  at corresponding time points  $t_0 < t_1 < \dots < t_K$ . Impor-  
 189 tantly, each distribution  $\rho_k$  may be represented by a different number of samples. As discussed in  
 190 §2, the JKO scheme approaches the continuous-time dynamics as the step sizes  $\Delta t_k = t_{k+1} - t_k$  are  
 191 small enough. This motivates our modeling assumption that the ground truth sequence of measures  
 192  $\{\rho_k\}_{k=0}^K$  follows the scheme  $\rho_{k+1} = \text{JKO}_{\Delta t_k \mathcal{J}^*}(\rho_k)$ .

193 Although the time intervals  $\Delta t_k$  between observations may vary, corresponding to non-uniform step  
 194 sizes in the JKO scheme, we assume equal spacing  $\Delta t_k = \tau$  in the remaining text for simplicity.

#### 195 3.2 METHODOLOGY AND LOSS DERIVATION

196 In this section, we assume  $\rho_k \in \mathcal{P}_{ac}(\mathcal{X})$  and minimization is always taken over  $\rho \in \mathcal{P}_{ac}(\mathcal{X})$  in order  
 197 to utilize Brenier's theorem. The key idea of our method builds on *inverse* optimization techniques  
 198 (Chan et al., 2025) and is demonstrated in Figure 1. Thanks to assumption  $\rho_{k+1} = \text{JKO}_{\tau \mathcal{J}^*}(\rho_k)$ ,  
 199 we can derive an inequality that becomes an equality if a candidate functional  $\mathcal{J}$  matches the ground  
 200 truth functional  $\mathcal{J}^*$ , since  $\rho_{k+1}$  is obtained via the JKO step associated with  $\mathcal{J}^*$ :

$$201 \min_{\rho^k} \left\{ \mathcal{J}(\rho^k) + \frac{1}{2\tau} d_{\mathbb{W}_2}^2(\rho_k, \rho^k) \right\} \leq \mathcal{J}(\rho_{k+1}) + \frac{1}{2\tau} d_{\mathbb{W}_2}^2(\rho_k, \rho_{k+1}). \quad (9)$$

202 Moving the right-hand side to the left yields an expression that is always upper-bounded by zero,  
 203 regardless of the choice of  $\mathcal{J}$ . Maximizing the resulting gap with respect to  $\mathcal{J}$  encourages the  
 204 candidate functional to approximate the true functional  $\mathcal{J}^*$ . This yields the following objective:

$$205 \max_{\mathcal{J}} \sum_{k=0}^{K-1} \left[ \min_{\rho^k} \left\{ \mathcal{J}(\rho^k) + \frac{1}{2\tau} d_{\mathbb{W}_2}^2(\rho_k, \rho^k) \right\} - \mathcal{J}(\rho_{k+1}) - \underbrace{\frac{1}{2\tau} d_{\mathbb{W}_2}^2(\rho_k, \rho_{k+1})}_{\text{independent of } \mathcal{J}, \rho^k} \right] =$$

$$\max_{\mathcal{J}} \sum_{k=0}^{K-1} \min_{\rho^k} \left[ \mathcal{J}(\rho^k) - \mathcal{J}(\rho_{k+1}) + \frac{1}{2\tau} d_{\mathbb{W}_2}^2(\rho_k, \rho) \right] + \text{Const.} \quad (10)$$

Thanks to our assumption  $\rho_k \in \mathcal{P}_{ac}(\mathcal{X})$ , Brenier’s theorem (Brenier, 1991) ensures that each distribution  $\rho^k$  appearing in (10) can be written as a pushforward  $\rho^k = T^k \# \rho_k$  for some transport map  $T^k : \mathcal{X} \rightarrow \mathcal{X}$ . In addition, the Wasserstein distance (2) admits the upper-bound  $d_{\mathbb{W}_2}^2(\rho_k, \rho) \leq \int_{\mathcal{X}} \|x - T^k(x)\|^2 d\rho_k(x)$ . Since the minimizations over  $\rho^k$  (or equivalently over  $T^k$ ) are independent across  $k$ , the sum and the minimization can be interchanged. Applying these observations, we end up with our *final loss objective* equivalent to (10):

$$\max_{\mathcal{J}} \min_{T^k} \mathcal{L}(\mathcal{J}, T^k) \stackrel{\text{def}}{=} \max_{\mathcal{J}} \min_{T^k} \sum_{k=0}^{K-1} \left[ \mathcal{J}(T^k \# \rho_k) - \mathcal{J}(\rho_{k+1}) + \frac{1}{2\tau} \int_{\mathcal{X}} \|x - T^k(x)\|_2^2 \rho_k(x) dx \right]. \quad (11)$$

The key idea is that the inner minimization over maps  $T^k$  approximates a JKO step for a given energy functional  $\mathcal{J}$ ; that is, each optimal map  $T_{\mathcal{J}}^k = \arg \min_{T^k} \left[ \mathcal{J}(T^k \# \rho_k) + \frac{1}{2\tau} \int_{\mathcal{X}} \|x - T^k(x)\|_2^2 \rho_k(x) dx \right]$  pushes  $\rho_k$  to the JKO-updated distribution  $\hat{\rho}_{k+1} = \text{JKO}_{\tau\mathcal{J}}(\rho_k)$ . The outer maximization then drives  $\mathcal{J}$  toward the true functional  $\mathcal{J}^*$ , causing the pushforward distributions  $\hat{\rho}_{k+1} = T_{\mathcal{J}}^k \# \rho_k$  to approach  $\rho_{k+1}$ , see illustration in Figure 1.

### 3.3 PRACTICAL ASPECTS: METHOD PARAMETRIZATION AND LEARNING PROCEDURE

We denote by  $\theta \in \Theta$  and  $\varphi \in \Phi$  the parameters of sufficiently expressive function classes used to approximate the candidate functional  $\mathcal{J}$  and the transport maps  $T^k$ , respectively. Specifically,  $\mathcal{J}_{\theta}$  and  $T_{\varphi}^k$  are their parameterized counterparts.

**Mapping Parameterization.** Building upon (Benamou et al., 2016), prior works (Mokrov et al., 2021; Alvarez-Melis et al., 2022; Bunne et al., 2022b) typically parameterize the transport maps as  $T_{\varphi}^k = \nabla \psi_{\varphi}^k$ , where  $\psi_{\varphi}^k$  are modeled using input-convex neural networks (ICNNs) (Amos et al., 2017). However, ICNNs suffer from poor scalability in high-dimensional settings (Korotin et al., 2021b). In contrast, since our objective (11) imposes *no convexity constraints*, we parameterize  $T_{\varphi}^k$  directly using standard architectures like MLPs or ResNets (He et al., 2016). This relaxation simplifies optimization and yields improved empirical stability. *Alternatively, one can employ more task-specific parametrizations, such as triangular maps (Baptista et al., 2024).*

**Energy functional parametrization.** Following (Terpin et al., 2024), we parameterize the energy functional  $\mathcal{J}$  as a specific instance of the free energy formulation in (5), where the internal energy term is chosen to be scaled negative differential entropy:

$$\mathcal{J}_{\theta}(\rho) = \underbrace{\int_{\mathcal{X}} V_{\theta_1}(x) d\rho(x)}_{\mathcal{V}_{\theta_1}(\rho)} + \underbrace{\int_{\mathcal{X} \times \mathcal{X}} W_{\theta_2}(x - y) d\rho(x) d\rho(y)}_{\mathcal{W}_{\theta_2}(\rho)} + \underbrace{\theta_3 \int_{\mathcal{X}} \log \rho(x) d\rho(x)}_{\mathcal{U}_{\theta_3}(\rho) = -\theta_3 \mathcal{H}(\rho)}. \quad (12)$$

Here  $\theta = \{\theta_1, \theta_2, \theta_3\}$  represents the set of learnable parameters:  $\theta_1$  and  $\theta_2$  are parameters of neural networks that define the potential  $V_{\theta_1} : \mathcal{X} \rightarrow \mathbb{R}$  and the interaction kernel  $W_{\theta_2} : \mathcal{X} \rightarrow \mathbb{R}$ , respectively;  $\theta_3 \in \mathbb{R}$  is a learnable scalar diffusion coefficient, see (7).

**Entropy evaluation.** All the terms in the objective function (11) can be estimated using Monte Carlo integration (Metropolis & Ulam, 1949), except for the internal energy term  $\mathcal{U}_{\theta_3}$ . To handle this remaining component, we follow (Mokrov et al., 2021, **Theorem 1**), (Alvarez-Melis et al., 2022, Eq. (11)) and apply the change-of-variables formula, yielding:

$$\mathcal{U}_{\theta_3}(T_{\varphi}^k \# \rho_k) = \mathcal{U}_{\theta_3}(\rho_k) - \theta_3 \int_{\mathcal{X}} \log |\det \nabla_x T_{\varphi}^k(x)| d\rho_k(x), \quad (13)$$

where  $\nabla_x T_{\varphi}^k$  denotes the Jacobian of the *invertible* transport map  $T_{\varphi}^k$ . We note that, during training, the computation of  $\log \det$  – even for non-invertible mappings  $T_{\varphi}^k$  – can be effectively handled by modern programming tools; see Appendix C.3 for details. To efficiently compute the gradient of the  $\log \det$  term in (13) with respect to  $\varphi$ , one can use the Hutchinson trace estimator (Hutchinson, 1989), as proposed by (Finlay et al., 2020; Huang et al., 2021). However, in our experimental setting (§5), we find that computing the full Jacobian is tractable and sufficient. For estimating the negative

270 differential entropy  $\mathcal{H}(\rho_k)$  in (13), we employ the Kozachenko–Leonenko nearest-neighbor estimator  
 271 (Kozachenko, 1987; Berrett et al., 2019). Notably, entropy values for the population measures  
 272  $\{\rho_k\}_{k=0}^K$  can be precomputed prior to training.  
 273

274 **Learning Objective.** To facilitate gradient-based optimization, we employ Monte Carlo estimation  
 275 to approximate the loss in (11). At each time step  $t_k$  (for  $k = 0, \dots, K$ ), we draw a batch of  $B_k$   
 276 samples  $X_k = \{x_k^1, \dots, x_k^{B_k}\} \sim \rho_k$  and optimize the following *empirical loss objective*:

$$277 \quad 278 \quad 279 \quad 280 \quad \widehat{\mathcal{L}}(\theta, \varphi) = \sum_{k=0}^{K-1} \left[ \widehat{\mathcal{J}}_\theta(T_\varphi^k(X_k)) - \widehat{\mathcal{J}}_\theta(X_{k+1}) + \frac{1}{2\tau} \sum_{j=1}^{B_k} \|x_k^j - T_\varphi^k(x_k^j)\|_2^2 \right], \quad (14)$$

281 where  $T_\varphi^k(X_k)$  denotes the batched pushforward used to compute the empirical estimates  $\widehat{\mathcal{J}}_\theta$  of  
 282 the functional  $\mathcal{J}_\theta$ . We optimize the objective in (14) with respect to  $\theta$  and  $\varphi$  using a standard  
 283 gradient descent–ascent scheme; implementation details are provided in §5 and Appendix C, with  
 284 the complete training procedure summarized in Algorithm 1 therein.  
 285

### 286 3.4 THEORETICAL ASPECTS

287 The central idea behind our loss formulation (§3.2) is to minimize the *inverse gap* in (9), which  
 288 measures the discrepancy between the optimal value (left-hand side) and the expected value (right-  
 289 hand side) of the JKO step for a candidate functional  $\mathcal{J}$ . While this objective is intuitively justified,  
 290 it remains to be shown whether the minimizer  $\mathcal{J}_{\min}$  of (11) truly approximates the ground-truth  
 291 energy functional  $\mathcal{J}^*$ . Our Theorem 3.1 addresses this question, demonstrating that, under suitable  
 292 assumptions, optimizing (11) indeed recovers  $\mathcal{J}^*$  up to an additive constant that does not affect the  
 293 dynamics governed by (4). Full details and proofs are provided in Appendix D.

294 We state our theorem for  $K = 1$ . When considering  $K > 1$ , the statement of the theorem holds true  
 295 independently for each timestep  $k = 0, 1, \dots, K - 1$ , and the results extend straightforwardly to this  
 296 more general setting. We further assume that both the ground-truth functional  $\mathcal{J}^*$  and the candidate  
 297 functional  $\mathcal{J}$  are purely of potential energy form, i.e.,  $\mathcal{J}^*(\rho) = \mathcal{J}_{\text{PE}}^*(\rho) = \int_{\mathcal{X}} V^*(x) d\rho(x)$ ;  $\mathcal{J}(\rho) =$   
 298  $\mathcal{J}_{\text{PE}}(\rho) = \int_{\mathcal{X}} V(x) d\rho(x)$ . For notational simplicity, we denote the loss as  $\mathcal{L}(V, T)$  in place of  
 299  $\mathcal{L}(\mathcal{J}_{\text{PE}}, T)$ . For technical purposes, we introduce the modified version of a potential as  $V_q := \tau V +$   
 300  $\frac{1}{2} \|\cdot\|_2^2 : \mathcal{X} \rightarrow \mathbb{R}$ ; subscript  $q$  stands for “quadratic”.  
 301

302 **Theorem 3.1** (Quality bounds for recovered potential energy). *Let  $\varepsilon(V) \stackrel{\text{def}}{=} \mathcal{L}(V^*, T_{V^*}) - \mathcal{L}(V, T_V)$   
 303 be the gap between the optimal and optimized value of inverse JKO loss (11) with internal  $\min_T$   
 304 problem solved exactly, i.e.,  $T_V \stackrel{\text{def}}{=} \min_T \mathcal{L}(V, T)$ . Let  $\mathcal{X}$  be a convex set; (modified) potentials  $V_q$   
 305 be strictly convex and  $\frac{1}{\beta}$ -smooth (see the definition in Appendix D). Then there exists  $C = C(\tau, \beta)$   
 306 such that following inequality holds:*  
 307

$$308 \quad 309 \quad \int_{\mathcal{X}} \|\nabla V^*(y) - \nabla V(y)\|^2 d\rho_1(y) \leq C\varepsilon(V). \quad (15)$$

310 Notably, equation (15) compares *gradients* of the recovered and ground truth potentials, hiding the  
 311 appearance of a redundant additive constant. Importantly, the assumptions of the theorem are not  
 312 particularly restrictive in practice. Specifically, the smoothness of potentials can be ensured by  
 313 employing smooth activation functions such as CELU (Barron, 2017), SiLU (Hendrycks & Gimpel,  
 314 2016), SoftPlus, and others. Strict convexity, in turn, can be enforced through architectural  
 315 design choices (Amos et al., 2017). Furthermore, the strict convexity of the (modified) potentials  
 316  $V_q$  can often be assumed when the step size  $\tau$  is sufficiently small. In our experiments (§5), we  
 317 parameterize  $V$  using standard MLPs and observe that this approach performs adequately.

318 To the best of our knowledge, this work is the first to provide a quality analysis (Theorem 3.1) of  
 319 JKO-based solvers for population dynamics. At present, our analysis focuses exclusively on the  
 320 *potential* energy component. Extending this framework to incorporate *interaction* and *internal* ener-  
 321 gies presents an interesting direction for future research. In contrast, (Wu & Wang, 2025) analyzes  
 322 the JKO scheme in terms of convergence to the true density of the Fokker-Planck equation (6), ac-  
 323 counting for parameter uncertainty in iterative updates. Our work differs by focusing on the ability  
 to restore energy components from data, rather than reconstructing the full density.

324 **4 RELATED WORKS**

325

326 This section reviews research directions most relevant to our work. For an extended discussion of  
 327 related works, see Appendix A.1. First, we discuss existing methods for modeling the dynamics of  
 328  $\rho_k$  given a known energy functional  $\mathcal{J}$  (§4.1). Next, we focus on approaches closely related to our  
 329 setting – specifically, learning population dynamics from observed data (§4.2) via the JKO scheme.

330 **4.1 SOLVING AGGREGATION-DIFFUSION EQUATION WITH GIVEN ENERGY**

331

332 This line of research focuses on methods for *solving* aggregation-diffusion equations derived from  
 333 WGFs with given energy functionals of the form (5). A variety of numerical approaches exist,  
 334 including mesh-based schemes (Cancès et al., 2023; Jüngel et al., 2024), particle methods (Craig &  
 335 Bertozzi, 2016; Campos Pinto et al., 2018; Carrillo et al., 2019b), and JKO-based schemes (Carrillo  
 336 et al., 2022). For a comprehensive overview, see (Bailo et al., 2024b, §5). Recently, the research  
 337 field was empowered by deep learning techniques. In particular, (Alvarez-Melis et al., 2022; Mokrov  
 338 et al., 2021; Fan et al., 2022; Park et al., 2023; Lee et al., 2024) employ gradient-based optimization  
 339 and neural-network parametrization to solve WGFs with given energy  $\mathcal{J}$  as to special cases of (5).

340 **4.2 LEARNING ENERGY FOR AGGREGATION-DIFFUSION EQUATION VIA JKO SCHEME**

341

342 This line of research focuses on methods for *learning* population dynamics (§3.1) using the JKO  
 343 scheme, as in JKOnet (Bunne et al., 2022b) and JKOnet\* (Terpin et al., 2024). We discuss these  
 344 methods in detail below, as they are the most closely related to our approach. Additional methods  
 345 for learning population dynamics are reviewed in Appendix A.1.1.

346 **JKOnet** (Bunne et al., 2022b) formulates the task of population dynamics recovery as a bi-level  
 347 optimization problem aimed at minimizing the discrepancy between observed distributions  $\rho_k$  and  
 348 model predictions  $\hat{\rho}_k$ :

$$\begin{aligned} \mathcal{L}_{\text{JKOnet}}(\theta, \varphi) &= \sum_{k=0}^{K-1} d_{\mathbb{W}_2}^2(\hat{\rho}_k, \rho_k), \quad \text{s.t. } \hat{\rho}_0 = \rho_0, \quad \hat{\rho}_{k+1} = \nabla \psi_k^* \# \hat{\rho}_k, \\ \psi_k^* &\stackrel{\text{def}}{=} \arg \min_{\varphi: \psi_\varphi \in \text{CVX}} \mathcal{J}_\theta(\nabla \psi_\varphi \# \hat{\rho}_k) + \frac{1}{2\tau} \int_{\mathcal{X}} \|x - \nabla \psi_\varphi(x)\|^2 d\hat{\rho}_k(x), \end{aligned} \quad (16)$$

354 where CVX denotes the set of continuously differentiable convex functions from  $\mathcal{X}$  to  $\mathbb{R}$ . The trans-  
 355 port maps  $\psi_\varphi$  are parametrized either using ICNNs (Bunne et al., 2022a) or through the ‘Monge  
 356 gap’ approach (Uscidda & Cuturi, 2023). However, the energy functional is limited to the potential  
 357 energy term  $\mathcal{V}_\theta(\rho)$ , excluding interaction and internal components. This framework has two key  
 358 limitations: **(i)** solving the bi-level optimization problem in (16) is computationally challenging and  
 359 requires specific techniques like unrolling optimizer’s steps, and **(ii)** extending the method to more  
 360 realistic settings with richer energy structures is questionable due to computational complexity.

361 **JKOnet\*** (Terpin et al., 2024) addresses the limitations of the original JKOnet by replacing the full  
 362 JKO optimization problem (8) with its first-order optimality conditions. This reformulation allows  
 363 the use of more expressive energy functionals (12). The method minimizes the following objective,  
 364 utilizing a *precomputed* optimal transport plans  $\pi_k$  between  $\rho_k$  and  $\rho_{k+1}$ :

$$\begin{aligned} \mathcal{L}_{\text{JKOnet}^*}(\theta) &= \sum_{k=0}^{K-1} \int_{\mathcal{X} \times \mathcal{X}} \left\| \nabla V_{\theta_1}(x_{k+1}) + \int_{\mathcal{X}} \nabla U_{\theta_2}(x_{k+1} - y_{k+1}) d\rho_{k+1}(y_{k+1}) \right. \\ &\quad \left. + \theta_3 \frac{\nabla \rho_{k+1}(x_{k+1})}{\rho_{k+1}(x_{k+1})} + \frac{1}{\tau} (x_{k+1} - x_k) \right\|^2 d\pi_k(x_k, x_{k+1}). \end{aligned} \quad (17)$$

371 The authors emphasize several advantages over JKOnet, including reduced computational costs  
 372 and support for more general energy functionals. However, their method requires an additional op-  
 373 timization round to precompute optimal transport plans  $\pi_k$ , introducing extra sources of inaccuracy  
 374 and rendering the approach non-end-to-end. Moreover, the authors of (Terpin et al., 2024) employ  
 375 discrete OT solvers to compute  $\pi_k$ , which may fail to accurately represent OT maps between the  
 376 underlying distributions, particularly in high-dimensional settings (Deb et al., 2021).

377 **iJKOnet (Ours)** integrates the strengths of both JKOnet and JKOnet\*. From JKOnet, it inherits  
 378 parameterized models for the energy functional  $\mathcal{J}_\theta$  and the transport map  $T_\varphi$ , avoiding reliance on

378 precomputed discrete optimal transport plans  $\pi_k$ , which can hinder flexible mappings between  $\rho_k$   
 379 and  $\rho_{k+1}$ . From JKOnet\*, it adopts a more expressive formulation of the energy functional (12) –  
 380 capable of capturing complex structures. Unlike JKOnet\*, our formulation does not rely on ana-  
 381 lytical expressions of the JKO step (8) optimality conditions; instead, it solves the step directly via  
 382 the inner minimization in (11). This approach makes the method more readily extendable to general  
 383 energy forms, such as porous medium energies (Alvarez-Melis et al., 2022), although exploring such  
 384 extensions is beyond the scope of the current work.

## 385 5 EXPERIMENTS

386 Our implementation is primarily based on the publicly available code<sup>1</sup> from JKOnet\* (Terpin et al.,  
 387 2024), written in JAX (Bradbury et al., 2018), and will be made publicly available. Since our  
 388 approach builds on the JKO framework, in §5.1, we first reproduce the experimental setup from  
 389 JKOnet\* for learning a known potential energy. Although we found the original codebase well-  
 390 structured and accessible, we identified several inconsistencies in the data generation process, which  
 391 we discuss in detail there. Finally, in §5.2, we compare our method against popular non-JKO base-  
 392 lines on single-cell RNA-seq datasets. Extended comparisons can be found in Appendix B.

393 **JKO-based Models.** As demonstrated in Appendix B.1.2 and B.1.3, learning the interaction and  
 394 internal energy components directly from samples is challenging – likely because their estimation  
 395 requires computing integrals over functions that must themselves be estimated. Consequently, for  
 396 large-scale experiments, we restrict our model to the potential energy parametrization, imposing the  
 397 inductive bias  $\theta_2 = \theta_3 = 0$ . Following (Terpin et al., 2024), we refer to this variant as  $\text{iJKOnet}_V$ .  
 398 For moderate-dimensional tasks, we explore all combinations of parametrizations to illustrate the  
 399 challenges of full energy recovery. For a fair comparison, we also implemented a time-varying potential  
 400 energy parametrization, as described in Appendix B of the original paper, and denote this  
 401 variant by  $\text{iJKOnet}_{t,V}$ . However, as discussed in Appendix A.2, this parametrization does not cor-  
 402 respond to a single, consistent energy functional governing the dynamics. Instead, it actually solves  
 403 a sequence of  $K$  independent optimal transport problems between consecutive snapshots (Bunne  
 404 et al., 2023). Moreover, because JKOnet (Bunne et al., 2022b) is computationally expensive, we  
 405 limit our comparisons to experiments with moderate dimensionality, with the corresponding results  
 406 reported in Appendix B.2. Technical details of model training are provided in Appendix C.

407 **Metrics.** We evaluate the next-step distribution  $\hat{\rho}_{k+1}$  from  $\rho_k$  and compare it to the ground-truth  
 408 distribution  $\rho_{k+1}$  using the following metrics: **(a)** Earth Mover’s Distance (EMD) (or  $\mathcal{W}_1$ ) (Rub-  
 409 ner et al., 1998) and  $\mathcal{W}_2$ , **(b)** *Bures-Wasserstein Unexplained Variance Percentage* ( $\text{Bd}_{\mathcal{W}_2}^2$ -UVP)  
 410 (Korotin et al., 2021c), **(c)**  $\mathcal{L}_2$ -based Unexplained Variance Percentage ( $\mathcal{L}_2$ -UVP) (Korotin et al.,  
 411 2021a), which measures the discrepancy between the ground-truth functional  $F^*$  and its reconstruc-  
 412 tion  $\hat{F}$ , and **(d)** Maximum Mean Discrepancy (MMD) (Gretton et al., 2012). Formal definitions and  
 413 implementation details for all metrics are provided in Appendix C.1.

### 414 5.1 LEARNING POTENTIAL ENERGY

415 As discussed in §1, the population dynamics setting assumes that particle trajectories are not tracked  
 416 across time – i.e., each time snapshot should consist of independently sampled particles without  
 417 temporal correlation. This implies that data at each time step should be regenerated from scratch.  
 418 However, we identified an inconsistency in the original codebase<sup>1</sup>: particle trajectories are preserved  
 419 across time steps, resulting in temporally correlated samples. We refer to this as the **paired** setup,  
 420 where each particle  $x_k$  is directly linked to  $x_{k+1}$  along a trajectory. In contrast, the intended, tem-  
 421 porally uncorrelated setting is referred to as the **unpaired** setup. Table 3 in Appendix B.1.1 shows  
 422 that switching between setup has a substantial impact on performance. A detailed discussion of this  
 423 phenomenon is provided in Appendix A.3.

424 Following (Terpin et al., 2024, §4.1), we begin by examining how our method learns potentials in the  
 425 2D **paired** setup, which allows for a direct visual comparison with JKOnet\*. These experiments  
 426 are conducted on the synthetic dataset from (Terpin et al., 2024, Appendix B), with results presented  
 427 in Figure 4 and further details provided in Appendix B.1.1. We then repeat the same procedure in  
 428 the 2D **unpaired** setup, analyzing how performance scales with the number of samples (2K and  
 429 10K). The corresponding results are shown in Figure 2. They indicate that our method outperforms

430 <sup>1</sup><https://github.com/antonioterpin/jknet-star>

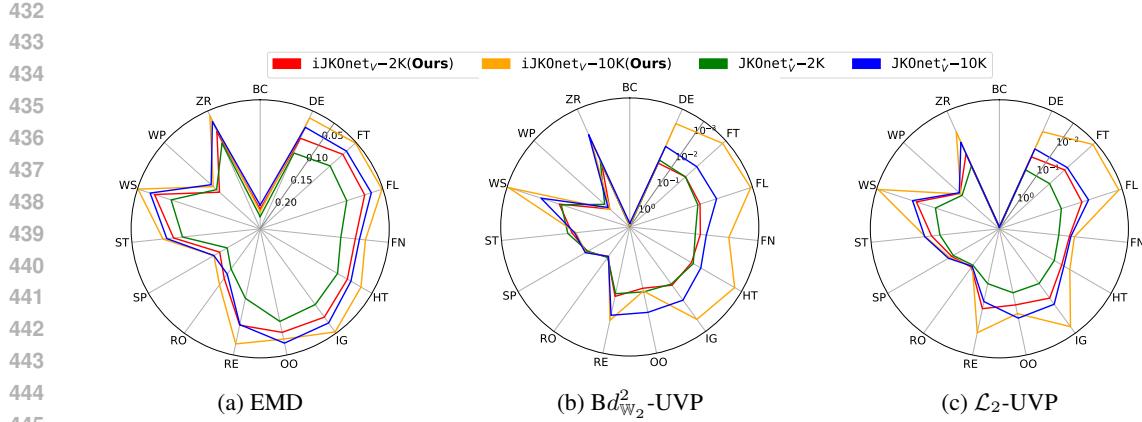


Figure 2: Numerical results from §5.1 for the **unpaired** setup. The reported absolute values show that, while increasing the number of samples generally improves performance across metrics, certain potentials remain challenging, highlighting the difficulty of this setup.

JK0net\* on nearly all potentials, and that for certain cases, increasing the number of samples does not improve performance, highlighting the difficulty of the unpaired setup.

## 5.2 LEARNING POPULATION DYNAMICS FOR SINGLE-CELL DATA

Following (Chen et al., 2023) and (Shen et al., 2025), in this section, we apply our method to modeling the dynamics of single-cell RNA data, with extended comparisons provided in Appendix B.2.

**Dataset.** We use the Embryoid Body (**EB**) dataset (Moon et al., 2019) and follow the preprocessing pipeline described in (Tong et al., 2020). The **EB** dataset comprises a collection of 5 timepoints sampled over a period of 30 days.

**Non-JKO Models.** We evaluate our method against classic algorithms for trajectory inference such as **TrajectoryNet** (Tong et al., 2020), **MIOFLOW** (Huguet et al., 2022), **DMSB** (Chen et al., 2023), **NLSB** (Koshizuka & Sato, 2023) and the recently proposed **MMSB** (Shen et al., 2025). For these methods, the reported metrics are taken from the corresponding referenced papers.

**Leave-two-out in 5D.** Following (Shen et al., 2025), we perform experiments using a leave-two-out setup. Since the **EB** dataset contains five timesteps, we remove the second ( $t_2$ ) and fourth ( $t_4$ ) timesteps **from training data**, and then evaluate how well our method can reconstruct the data from the remaining  $t_1$  and  $t_3$  timesteps. The results are demonstrated in Table 1. **One can see that our method with the “ $t, V$ ” parameterization yields the best metrics, whereas the use of alternative parameterizations (with interaction and internal energies) results in poorer/non-even metrics.** **Vanilla-SB** corresponds to the IPF (GP) method (Vargas et al., 2021); further details are provided in Appendix B.6 of (Shen et al., 2025).

Table 1: **5D** experiment.  $d_{W_2}$  distance ( $\downarrow$ ) comparison across  $t_2$  and  $t_4$ . Results for baselines (Shen et al., 2025).

Method	$t_2$	$t_4$
Vanilla-SB	$1.49 \pm 0.063$	$1.55 \pm 0.034$
DMSB	$1.13 \pm 0.082$	$1.45 \pm 0.16$
TrajectoryNet	$2.03 \pm 0.04$	$1.93 \pm 0.08$
MMSB	$1.27 \pm 0.028$	$1.57 \pm 0.048$
<b>Static</b>		
JK0net <sub>V</sub> *	$1.145 \pm 0.033$	$2.529 \pm 0.014$
JK0net <sub>V+U</sub> *	$1.099 \pm 0.119$	$2.537 \pm 0.054$
JK0net <sub>V+W</sub> *	$1.419 \pm 0.173$	$2.510 \pm 0.094$
JK0net <sub>W+U</sub> *	$1.887 \pm 0.017$	$1.739 \pm 0.037$
JK0net*	$1.361 \pm 0.257$	$2.557 \pm 0.042$
<b>Static (Ours)</b>		
iJK0net <sub>V</sub>	$1.082 \pm 0.011$	$1.147 \pm 0.001$
iJK0net <sub>V+U</sub>	$1.065 \pm 0.018$	$1.150 \pm 0.004$
iJK0net <sub>V+W</sub>	$2.865 \pm 0.166$	$1.376 \pm 0.015$
iJK0net <sub>W+U</sub>	$1.649 \pm 0.005$	$0.868 \pm 0.005$
iJK0net	$3.577 \pm 0.166$	$1.395 \pm 0.032$
<b>Time-varying</b>		
JK0net <sub>t,V</sub> *	$4.414 \pm 1.499$	$2.771 \pm 0.197$
iJK0net <sub>t,V</sub> (Ours)	$0.983 \pm 0.037$	$0.849 \pm 0.021$

Our method outperforms previous approaches both in terms of reconstruction metrics and execution time, as can be seen by comparing Table 12 in Appendix E of (Shen et al., 2025) with Table 7.

**Leave-one-out in 100D.** Following (Chen et al., 2023), we conduct experiments using a leave-one-out setup. One of the timesteps  $t_2$ ,  $t_3$ , or  $t_4$  is omitted, and we evaluate the method’s ability to reconstruct the distribution of the left-out timestep. The results are shown in Table 2. **Our method achieves performance comparable to DMSB in the all-time (w/o LO) setting while using a simpler, simulation-free optimization routine that requires no trajectory caching. Consequently, it outperforms DMSB in execution time, as shown by comparing Table 12 in Appendix E of (Shen et al., 2025) (for 5D) with Table 7 (for 100D).**

Table 2: **100D** experiment. MMD distance ( $\downarrow$ ). Comparison of methods across different leave-one-out splits. The results are averaged across 3 runs. Results for baseline methods: (Chen et al., 2023).

Method	LO- $t_2$	LO- $t_3$	LO- $t_4$	w/o LO
NLSB (Koshizuka & 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)	<b><math>0.042 \pm 0.020</math></b>	<b><math>0.033 \pm 0.003</math></b>	<b><math>0.040 \pm 0.020</math></b>	<b><math>0.032 \pm 0.003</math></b>
JKOnet $^*$ <sub>V</sub> (Terpin et al., 2024)	$0.220 \pm 0.025$	$0.293 \pm 0.018$	$0.235 \pm 0.006$	$0.229 \pm 0.052$
iJKOnet $V$ ( <b>Ours</b> )	$0.137 \pm 0.001$	$0.123 \pm 0.001$	$0.097 \pm 0.002$	$0.085 \pm 0.024$
JKOnet $^*$ <sub><math>t,V</math></sub> (Terpin et al., 2024)	$0.575 \pm 0.119$	$0.619 \pm 0.157$	$0.456 \pm 0.056$	$0.477 \pm 0.098$
iJKOnet $t,V$ ( <b>Ours</b> )	$0.848 \pm 0.043$	$0.370 \pm 0.233$	$0.055 \pm 0.007$	$0.124 \pm 0.243$

## 6 DISCUSSION

**Contributions.** We introduce the novel **iJKOnet** method, which provides a general framework for recovering any type of energy functional governing the evolution of a system; in this work, we focus on the free energy parametrization. Our work bridges the gap between inverse optimization theory and computational gradient flows. The proposed method outperforms previous JKO-based approaches in all comparisons and achieves comparable results to non-JKO methods on real-world tasks. We also establish theoretical guarantees for recovering dynamics driven solely by potential energy.

**Limitations.** Our method can not recover other types of internal energy functionals except for the Entropy, nor can it handle time-dependent interaction energies, although theoretical results for such cases exist (Ferreira & Valencia-Guevara, 2018). Furthermore, it does not account for birth-death dynamics as in recent trajectory inference methods (Zhang et al., 2025b). The reliance on entropy estimation may also lead to performance degradation in higher-dimensional settings, and we find that learning interaction energies remains particularly challenging. As a result, jointly optimizing all energy parameters ( $\theta_1, \theta_2, \theta_3$ ) can cause instability and convergence to inaccurate potential estimates. Finally, from a theoretical standpoint, our analysis currently provides guarantees only for potential energy (Theorem (3.1)); extending it to other energy types is an avenue for future work.

**LLM Usage.** Large Language Models (LLMs) were used exclusively to assist with sentence rephrasing and improving text clarity. All scientific content, results, and interpretations in this paper were produced solely by the authors.

## REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2623–2631, 2019.

Vitalii Aksenov and Martin Eigel. An eulerian approach to regularized jko scheme with low-rank tensor decompositions for bayesian inversion. *Journal of Scientific Computing*, 105(1):3, 2025.

Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=li7qeBbCR1t>.

540 Fabian Altekrüger, Johannes Hertrich, and Gabriele Steidl. Neural Wasserstein gradient flows for  
 541 discrepancies with riesz kernels. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara  
 542 Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International  
 543 Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*,  
 544 pp. 664–690. PMLR, 23–29 Jul 2023. URL [https://proceedings.mlr.press/v202/  
 545 altekruger23a.html](https://proceedings.mlr.press/v202/altekruger23a.html).

546 David Alvarez-Melis, Yair Schiff, and Youssef Mroueh. Optimizing functionals on the space of  
 547 probabilities with input convex neural networks. *Transactions on Machine Learning Research*,  
 548 2022. ISSN 2835-8856.

549 Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the  
 550 space of probability measures*. Springer Science & Business Media, 2008.

552 Brandon Amos, Lei Xu, and J Zico Kolter. Input convex neural networks. In *International conference  
 553 on machine learning*, pp. 146–155. PMLR, 2017.

555 Francisco Andrade, Gabriel Peyré, and Clarice Poon. Learning from samples: Inverse problems  
 556 over measures via sharpened fenchel-young losses. [arXiv preprint arXiv:2505.07124](https://arxiv.org/abs/2505.07124), 2025.

557 Lazar Atanackovic, Xi Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio,  
 558 Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the  
 559 wasserstein manifold. In *The Thirteenth International Conference on Learning Representations*,  
 560 2025. URL <https://openreview.net/forum?id=9SYczU3Qgm>.

562 Francisco J Ayala, Michael E Gilpin, and Joan G Ehrenfeld. Competition between species: theoretical  
 563 models and experimental tests. *Theoretical population biology*, 4(3):331–356, 1973.

564 Jean Baccou and Jacques Liandrat. Subdivision scheme for discrete probability measure-valued  
 565 data. *Applied Mathematics Letters*, 158:109233, 2024.

567 Rafael Bailo, Alethea Barbaro, Susana N Gomes, Konstantin Riedl, Tim Roith, Claudia Totzeck, and  
 568 Urbain Vaes. Cbx: Python and julia packages for consensus-based interacting particle methods.  
 569 *Journal of Open Source Software*, 9(98):6611, 2024a.

571 Rafael Bailo, José A Carrillo, and David Gómez-Castro. Aggregation-diffusion equations for col-  
 572 lective behaviour in the sciences. [arXiv preprint arXiv:2405.16679](https://arxiv.org/abs/2405.16679), 2024b.

573 Amartya Banerjee, Harlin Lee, Nir Sharon, and Caroline Moosmüller. Efficient trajectory infer-  
 574 ence in wasserstein space using consecutive averaging. In *International Conference on Artificial  
 575 Intelligence and Statistics*, pp. 2260–2268. PMLR, 2025.

577 Arindam Banerjee, Srujana Merugu, Inderjit S Dhillon, and Joydeep Ghosh. Clustering with breg-  
 578 man divergences. *Journal of machine learning research*, 6(Oct):1705–1749, 2005.

579 Ricardo Baptista, Youssef Marzouk, and Olivier Zahm. On the representation and learning of mono-  
 580 tone triangular transport maps. *Foundations of Computational Mathematics*, 24(6):2063–2108,  
 581 2024.

583 Jonathan T Barron. Continuously differentiable exponential linear units. [arXiv preprint  
 584 arXiv:1704.07483](https://arxiv.org/abs/1704.07483), 2017.

585 Nico Battich, Joep Beumer, Buys De Barbanson, Lenno Krenning, Chloé S Baron, Marvin E Tanen-  
 586 baum, Hans Clevers, and Alexander Van Oudenaarden. Sequencing metabolically labeled tran-  
 587 scripts in single cells reveals mrna turnover strategies. *Science*, 367(6482):1151–1156, 2020.

589 Amir Beck. *First-Order Methods in Optimization*, volume 25 of *MOS-SIAM Series on  
 590 Optimization*. SIAM, Philadelphia, 2017. ISBN 9781611974980. doi: 10.1137/1.  
 591 9781611974997.

592 Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-  
 593 kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

594 Jean-David Benamou, Guillaume Carlier, Quentin Mérigot, and Edouard Oudet. Discretization of  
 595 functionals involving the monge–ampère operator. *Numerische mathematik*, 134(3):611–636,  
 596 2016.

597 Jean-David Benamou, Thomas O Gallouët, and François-Xavier Vialard. Second-order models for  
 598 optimal transport and cubic splines on the wasserstein space. *Foundations of Computational  
 599 Mathematics*, 19(5):1113–1143, 2019.

600 Jules Berman and Benjamin Peherstorfer. CoLoRA: Continuous low-rank adaptation for reduced  
 601 implicit neural modeling of parameterized partial differential equations. In Ruslan Salakhutdinov,  
 602 Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix  
 603 Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume  
 604 235 of *Proceedings of Machine Learning Research*, pp. 3565–3583. PMLR, 21–27 Jul 2024.  
 605 URL <https://proceedings.mlr.press/v235/berman24b.html>.

606 Jules Berman, Tobias Blickhan, and Benjamin Peherstorfer. Parametric model reduction of mean-  
 607 field and stochastic systems via higher-order action matching. *Advances in Neural Information  
 608 Processing Systems*, 37:56588–56618, 2024.

609 Thomas B Berrett, Richard J Samworth, and Ming Yuan. Efficient multivariate entropy estimation  
 610 via k-nearest neighbour distances. *The Annals of Statistics*, 47(1):288–318, 2019.

611 Vladimir I Bogachev and Aleksandr V Kolesnikov. The monge–kantorovich problem: achievements,  
 612 connections, and perspectives. *Russian Mathematical Surveys*, 67(5):785, 2012.

613 Vladimir I Bogachev, Nicolai V Krylov, Michael Röckner, and Stanislav V Shaposhnikov.  
 614 *Fokker–Planck–Kolmogorov Equations*, volume 207. American Mathematical Society, 2022.

615 Clément Bonet, Christophe Vauthier, and Anna Korba. Flowing datasets with wasserstein over  
 616 wasserstein gradient flows. *arXiv preprint arXiv:2506.07534*, 2025.

617 James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal  
 618 Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao  
 619 Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/jax-ml/jax>.

620 Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions.  
 621 *Communications on pure and applied mathematics*, 44(4):375–417, 1991.

622 Luc Brodat-Motte, Riccardo Bonalli, and Alessandro Rudi. Learning controlled stochastic differen-  
 623 tial equations. *arXiv preprint arXiv:2411.01982*, 2024.

624 Charlotte Bunne, Andreas Krause, and Marco Cuturi. Supervised training of conditional monge  
 625 maps. *Advances in Neural Information Processing Systems*, 35:6859–6872, 2022a.

626 Charlotte Bunne, Laetitia Papaxanthos, Andreas Krause, and Marco Cuturi. Proximal optimal trans-  
 627 port modeling of population dynamics. In *International Conference on Artificial Intelligence and  
 628 Statistics*, pp. 6511–6528. PMLR, 2022b.

629 Charlotte Bunne, Stefan G Stark, Gabriele Gut, Jacobo Sarabia Del Castillo, Mitch Levesque,  
 630 Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-  
 631 cell perturbation responses using neural optimal transport. *Nature methods*, 20(11):1759–1768,  
 632 2023.

633 Daniel Burkhardt, Jonathan Bloom, Robrecht Cannoodt, Malte D Luecken, Smita Krishnaswamy,  
 634 Christopher Lance, Angela O Pisco, and Fabian J Theis. Multimodal single-cell integration across  
 635 time, individuals, and batches. *NeurIPS Competitions*, pp. 12, 2022.

636 Ivan Butakov, Aleksandr Tolmachev, Sofia Malanchuk, Anna Neopryatnaya, and Alexey Frolov.  
 637 Mutual information estimation via normalizing flows. *Advances in Neural Information Processing  
 638 Systems*, 37:3027–3057, 2024.

639 John Charles Butcher. *Numerical methods for ordinary differential equations*. John Wiley & Sons,  
 640 2016.

648 Martin Campos Pinto, José A Carrillo, Frédérique Charles, and Young-Pil Choi. Convergence of a  
 649 linearly transformed particle method for aggregation equations. *Numerische Mathematik*, 139:  
 650 743–793, 2018.

651 Clément Cancès, Maxime Herda, and Annamaria Massimini. Finite volumes for a generalized  
 652 poisson-ernst-planck system with cross-diffusion and size exclusion. In *International Conference  
 653 on Finite Volumes for Complex Applications*, pp. 57–73. Springer, 2023.

654 José A Carrillo and Rishabh S Gvalani. Phase transitions for nonlinear nonlocal aggregation-  
 655 diffusion equations. *Communications in mathematical physics*, 382(1):485–545, 2021.

656 José A Carrillo, Robert J McCann, and Cédric Villani. Kinetic equilibration rates for granular media  
 657 and related equations: entropy dissipation and mass transportation estimates. *Revista Matematica  
 658 Iberoamericana*, 19(3):971–1018, 2003.

659 José A Carrillo, Katy Craig, and Yao Yao. Aggregation-diffusion equations: dynamics, asymptotics,  
 660 and singular limits. *Active Particles, Volume 2: Advances in Theory, Models, and Applications*,  
 661 pp. 65–108, 2019a.

662 José A Carrillo, Katy Craig, Li Wang, and Chaozhen Wei. Primal dual methods for wasserstein  
 663 gradient flows. *Foundations of Computational Mathematics*, pp. 1–55, 2022.

664 José Antonio Carrillo, Katy Craig, and Francesco S Patacchini. A blob method for diffusion.  
 665 *Calculus of Variations and Partial Differential Equations*, 58:1–53, 2019b.

666 Timothy CY Chan, Rafid Mahmood, and Ian Yihang Zhu. Inverse optimization: Theory and appli-  
 667 cations. *Operations Research*, 73(2):1046–1074, 2025.

668 Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *The Twelfth  
 669 International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=g7ohDlTITL>.

670 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary  
 671 differential equations. *Advances in neural information processing systems*, 31, 2018a.

672 Tianrong Chen, Guan-Horng Liu, and Evangelos Theodorou. Likelihood training of schrödinger  
 673 bridge using forward-backward SDEs theory. In *International Conference on Learning  
 674 Representations*, 2022a. URL <https://openreview.net/forum?id=nioAdKCEdXB>.

675 Tianrong Chen, Guan-Horng Liu, Molei Tao, and Evangelos Theodorou. Deep momentum multi-  
 676 marginal schrödinger bridge. *Advances in Neural Information Processing Systems*, 36:57058–  
 677 57086, 2023.

678 Yongxin Chen, Giovanni Conforti, and Tryphon T Georgiou. Measure-valued spline curves: An  
 679 optimal transport viewpoint. *SIAM Journal on Mathematical Analysis*, 50(6):5947–5968, 2018b.

680 Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. The most likely evolution of diffusing and  
 681 vanishing particles: Schrodinger bridges with unbalanced marginals. *SIAM Journal on Control  
 682 and Optimization*, 60(4):2016–2039, 2022b.

683 Xiuyuan Cheng, Jianfeng Lu, Yixin Tan, and Yao Xie. Convergence of flow-based generative models  
 684 via proximal gradient descent in wasserstein space. *IEEE Transactions on Information Theory*,  
 685 2024.

686 Ziheng Cheng, Shiyue Zhang, Longlin Yu, and Cheng Zhang. Particle-based variational inference  
 687 with generalized wasserstein gradient flow. *Advances in Neural Information Processing Systems*,  
 688 36:66583–66609, 2023.

689 Andrei Chertkov and Ivan Oseledets. Solution of the fokker–planck equation by cross approximation  
 690 method in the tensor train format. *Frontiers in Artificial Intelligence*, 4:668215, 2021.

691 Sinho Chewi, Julien Clancy, Thibaut Le Gouic, Philippe Rigollet, George Stepaniants, and Austin  
 692 Stromme. Fast and smooth interpolation on wasserstein space. In *International Conference on  
 693 Artificial Intelligence and Statistics*, pp. 3061–3069. PMLR, 2021.

702 Sinho Chewi, Jonathan Niles-Weed, and Philippe Rigollet. *Statistical Optimal Transport:*  
 703 *École d’Été de Probabilités de Saint-Flour XLIX – 2019*, volume 2364 of *Lecture Notes in*  
 704 *Mathematics*. Springer Cham, 2025. ISBN 978-3-031-85159-9. doi: 10.1007/  
 705 978-3-031-85160-5.

706 Lenaic Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. An interpolat-  
 707 ing distance between optimal transport and fisher-rao metrics. *Foundations of Computational*  
 708 *Mathematics*, 18(1):1–44, 2018a.

709 Lenaic Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced opti-  
 710 mal transport: Dynamic and kantorovich formulations. *Journal of Functional Analysis*, 274(11):  
 711 3090–3123, 2018b.

712 Lénaïc Chizat, Maria Colombo, Xavier Fernández-Real, and Alessio Figalli. Infinite-width limit  
 713 of deep linear neural networks. *Communications on Pure and Applied Mathematics*, 77(10):  
 714 3958–4007, 2024.

715 Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Scalable wasserstein gradient flow for gener-  
 716 ative modeling through unbalanced optimal transport. In *International Conference on Machine*  
 717 *Learning*, pp. 8629–8650. PMLR, 2024.

718 Julien Clancy and Felipe Suarez. Wasserstein-fisher-rao splines. *arXiv e-prints*, pp. arXiv–2203,  
 719 2022.

720 Katy Craig and Andrea Bertozzi. A blob method for the aggregation equation. *Mathematics of*  
 721 *computation*, 85(300):1681–1717, 2016.

722 Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural*  
 723 *information processing systems*, 26, 2013.

724 Marco Cuturi, Laetitia Meng-Papaxanthos, Yingtao Tian, Charlotte Bunne, Geoff Davis, and Olivier  
 725 Teboul. Optimal transport tools (ott): A jax toolbox for all things wasserstein. *arXiv preprint*  
 726 *arXiv:2201.12324*, 2022.

727 Aniket Das and Dheeraj Nagaraj. Provably fast finite particle variants of svgd via virtual particle  
 728 stochastic approximation. *Advances in Neural Information Processing Systems*, 36:49748–49760,  
 729 2023.

730 Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with  
 731 optimism. In *International Conference on Learning Representations*, 2018.

732 Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger  
 733 bridge with applications to score-based generative modeling. *Advances in neural information*  
 734 *processing systems*, 34:17695–17709, 2021.

735 Nabarun Deb, Promit Ghosal, and Bodhisattva Sen. Rates of estimation of optimal transport maps  
 736 using plug-in estimators via barycentric projections. *Advances in Neural Information Processing*  
 737 *Systems*, 34:29736–29753, 2021.

738 DeepMind, Igor Babuschkin, Kate Baumli, Alison Bell, Surya Bhupatiraju, Jake Bruce, Peter  
 739 Buchlovsky, David Budden, Trevor Cai, Aidan Clark, Ivo Danihelka, Antoine Dedieu, Clau-  
 740 dio Fantacci, Jonathan Godwin, Chris Jones, Ross Hemsley, Tom Hennigan, Matteo Hessel,  
 741 Shaobo Hou, Steven Kapturowski, Thomas Keck, Iurii Kemaev, Michael King, Markus Kunesch,  
 742 Lena Martens, Hamza Merzic, Vladimir Mikulik, Tamara Norman, George Papamakarios, John  
 743 Quan, Roman Ring, Francisco Ruiz, Alvaro Sanchez, Laurent Sartran, Rosalia Schneider, Eren  
 744 Sezener, Stephen Spencer, Srivatsan Srinivasan, Miloš Stanojević, Wojciech Stokowiec, Luyu  
 745 Wang, Guangyao Zhou, and Fabio Viola. The DeepMind JAX Ecosystem, 2020. URL <http://github.com/google-deepmind>.

746 Michael Ziyang Diao, Krishna Balasubramanian, Sinho Chewi, and Adil Salim. Forward-backward  
 747 gaussian variational inference via jko in the bures-wasserstein space. In *International Conference*  
 748 *on Machine Learning*, pp. 7960–7991. PMLR, 2023.

756 Nira Dyn and Nir Sharon. Subdivision schemes in metric spaces. [arXiv preprint arXiv:2509.08070](https://arxiv.org/abs/2509.08070),  
 757 2025.

758

759 Rossin Erbe, Genevieve Stein-O'Brien, and Elana J Fertig. Transcriptomic forecasting with neural  
 760 ordinary differential equations. [Patterns](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10208333/), 4(8), 2023.

761

762 Lawrence C Evans. [Partial differential equations](https://www.ams.org/publications/monographs/19), volume 19. American Mathematical Society,  
 763 2022.

764

765 Jiaojiao Fan, Qinsheng Zhang, Amirhossein Taghvaei, and Yongxin Chen. Variational wasserstein  
 766 gradient flow. In [International Conference on Machine Learning](https://proceedings.mlr.press/v162/), pp. 6185–6215. PMLR, 2022.

767

768 Jiaojiao Fan, Shu Liu, Shaojun Ma, Hao-Min Zhou, and Yongxin Chen. Neural monge map estimation  
 769 and its applications. [Transactions on Machine Learning Research](https://proceedings.mlr.press/v162/), 2023. ISSN 2835-8856.  
 770 URL <https://openreview.net/forum?id=2mZS1Qscj3>. Featured Certification.

771

772 Lucas CF Ferreira and Julio C Valencia-Guevara. Gradient flows of time-dependent functionals in  
 773 metric spaces and applications to pdes. [Monatshefte für Mathematik](https://www.degruyter.com/document/doi/10.1515/mf-2018-0022/html), 185(2):231–268, 2018.

774

775 Davide Fiaschi and Cristiano Ricci. The spatial evolution of economic activities and the emergence  
 776 of cities. [Journal of Economic Theory](https://www.sciencedirect.com/science/article/pii/S0898552622000607), pp. 105971, 2025.

777

778 Alessio Figalli and Federico Glaudo. [An Invitation to Optimal Transport, Wasserstein Distances,  
 779 and Gradient Flows](https://www.ems-ph.de/press/etb/258/). EMS Press, Berlin, 2nd edition, 2023. ISBN 978-3-98547-050-1. doi:  
 780 10.4171/ETB/25. URL <https://ems.press/books/etb/258>.

781

782 Chris Finlay, Jörn-Henrik Jacobsen, Levon Nurbekyan, and Adam Oberman. How to train your neu-  
 783 ral ode: the world of jacobian and kinetic regularization. In [International conference on machine  
 784 learning](https://proceedings.mlr.press/v162/), pp. 3154–3164. PMLR, 2020.

785

786 Mike Fisher, Jorge Nocedal, Yannick Trémolet, and Stephen J Wright. Data assimilation in weather  
 787 forecasting: a case study in pde-constrained optimization. [Optimization and Engineering](https://www.sciencedirect.com/science/article/pii/S0925005609000342), 10(3):  
 788 409–426, 2009.

789

790 Benyamin Ghojogh, Ali Ghodsi, Fakhri Karray, and Mark Crowley. Reproducing kernel hilbert  
 791 space, mercer's theorem, eigenfunctions, nyström method, and use of kernels in machine learn-  
 792 ing: Tutorial and survey. [arXiv preprint arXiv:2106.08443](https://arxiv.org/abs/2106.08443), 2021.

793

794 David Gómez-Castro. Beginner's guide to aggregation-diffusion equations. [SeMA Journal](https://www.sciencedirect.com/science/article/pii/S0898552622000607), 81(4):  
 795 531–587, 2024.

796

797 V Gontis, J Ruseckas, and A Kononovičius. A long-range memory stochastic model of the return in  
 798 financial markets. [Physica A: Statistical Mechanics and its Applications](https://www.sciencedirect.com/science/article/pii/S0378137110000937), 389(1):100–106, 2010.

799

800 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,  
 801 Aaron Courville, and Yoshua Bengio. Generative adversarial networks. [Communications of the  
 802 ACM](https://www.sciencedirect.com/science/article/pii/S0898552619300001), 63(11):139–144, 2020.

803

804 Erhard Gramer. Probability measures with given marginals and conditionals: I-projections and  
 805 conditional iterative proportional fitting. [Statistics & Risk Modeling](https://www.sciencedirect.com/science/article/pii/S0898552600000231), 18(3):311–330, 2000.

806

807 Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola.  
 808 A kernel two-sample test. [The journal of machine learning research](https://www.jmlr.org/papers/v13/gretton12a.html), 13(1):723–773, 2012.

809

810 Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Im-  
 811 proved training of wasserstein gans. [Advances in neural information processing systems](https://proceedings.neurips.cc/paper/7295_improved-training-of-wasserstein-gans.pdf), 30,  
 812 2017.

813

814 Fuqun Han, Stanley Osher, and Wuchen Li. Tensor train based sampling algorithms for ap-  
 815 proximating regularized wasserstein proximal operators. [SIAM/ASA Journal on Uncertainty  
 816 Quantification](https://www.siam.org/News/ViewNews.php?NID=470), 13(2):775–804, 2025.

817

818 Tatsunori Hashimoto, David Gifford, and Tommi Jaakkola. Learning population-level diffusions  
 819 with generative rnns. In [International Conference on Machine Learning](https://proceedings.mlr.press/v162/), pp. 2417–2426. PMLR,  
 820 2016.

810 Doron Haviv, Aram-Alexandre Pooladian, Dana Pe'er, and Brandon Amos. Wasserstein flow  
 811 matching: Generative modeling over families of distributions. In Forty-second International  
 812 Conference on Machine Learning, 2025. URL <https://openreview.net/forum?id=MRmI68k3gd>.

813

814 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recogni-  
 815 tion. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp.  
 816 770–778, 2016.

817

818 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). arXiv preprint  
 819 [arXiv:1606.08415](https://arxiv.org/abs/1606.08415), 2016.

820

821 Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth schrödinger  
 822 bridges. In Forty-second International Conference on Machine Learning, 2025. URL <https://openreview.net/forum?id=GHyvvWu1XC>.

823

824 Chin-Wei Huang, Ricky T. Q. Chen, Christos Tsirigotis, and Aaron Courville. Convex potential  
 825 flows: Universal probability distributions with optimal transport and convex optimization. In  
 826 International Conference on Learning Representations, 2021.

827

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

831

832 Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian  
 833 smoothing splines. Communications in Statistics-Simulation and Computation, 18(3):1059–1076,  
 1989.

834

835 Jan-Christian Hüttner and Philippe Rigollet. Minimax estimation of smooth optimal transport maps.  
 836 The Annals of Statistics, 49(2):1166 – 1194, 2021. doi: 10.1214/20-AOS1997. URL <https://doi.org/10.1214/20-AOS1997>.

837

838

839 Haoming Jiang, Zhehui Chen, Minshuo Chen, Feng Liu, Dingding Wang, and Tuo Zhao. On com-  
 840 putation and generalization of gans with spectrum control. arXiv preprint arXiv:1812.10912,  
 2018.

841

842 Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the fokker-  
 843 planck equation. SIAM journal on mathematical analysis, 29(1):1–17, 1998.

844

845 Ansgar Jüngel, Stefan Portisch, and Antoine Zurek. A convergent finite-volume scheme for nonlo-  
 846 cal cross-diffusion systems for multi-species populations. ESAIM: Mathematical Modelling and  
 847 Numerical Analysis, 58(2):759–792, 2024.

848

849 Jorge Justiniano, Martin Rumpf, and Matthias Erbar. Approximation of splines in wasserstein  
 850 spaces. ESAIM: Control, Optimisation and Calculus of Variations, 30:64, 2024.

851

852 Leonid V Kantorovich. On the translocation of masses. In Dokl. Akad. Nauk. USSR (NS), vol-  
 ume 37, pp. 199–201, 1942.

853

854 Kacper Kapusniak, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein,  
 855 Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the  
 856 data manifold. Advances in Neural Information Processing Systems, 37:135011–135042, 2024.

857

858 Keisuke Kawano, Takuro Kutsuna, Naoki Hayashi, Yasushi Esaki, and Hidenori Tanaka. Ct-ot  
 859 flow: Estimating continuous-time dynamics from discrete temporal snapshots. arXiv preprint  
[arXiv:2505.17354](https://arxiv.org/abs/2505.17354), 2025.

860

861 Evelyn F Keller and Lee A Segel. Model for chemotaxis. Journal of theoretical biology, 30(2):  
 225–234, 1971.

862

863 Diederik P Kingma. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980,  
 2014.

864 Dominik Klein, Théo Uscidda, Fabian Theis, and Marco Cuturi. Genot: Entropic (gromov) wasser-  
 865 stein flow matching with applications to single-cell genomics. *Advances in Neural Information*  
 866 *Processing Systems*, 37:103897–103944, 2024.

867

868 Nikita Kornilov, Petr Mokrov, Alexander Gasnikov, and Aleksandr Korotin. Optimal flow match-  
 869 ing: Learning straight trajectories in just one step. *Advances in Neural Information Processing*  
 870 *Systems*, 37:104180–104204, 2024.

871 Alexander Korotin, Vage Egiazarian, Arip Asadulaev, Alexander Safin, and Evgeny Burnaev.  
 872 Wasserstein-2 generative networks. In *International Conference on Learning Representations*,  
 873 2021a. URL [https://openreview.net/forum?id=bEoxzW\\_EXsa](https://openreview.net/forum?id=bEoxzW_EXsa).

874 Alexander Korotin, Lingxiao Li, Aude Genevay, Justin M Solomon, Alexander Filippov, and Evgeny  
 875 Burnaev. Do neural optimal transport solvers work? a continuous wasserstein-2 benchmark.  
 876 *Advances in neural information processing systems*, 34:14593–14605, 2021b.

877 Alexander Korotin, Lingxiao Li, Justin Solomon, and Evgeny Burnaev. Continuous wasserstein-2  
 878 barycenter estimation without minimax optimization. In *International Conference on Learning*  
 879 *Representations*, 2021c.

880 Alexander Korotin, Daniil Selikhanovich, and Evgeny Burnaev. Neural optimal transport. In  
 881 *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=d8CBR1WNkqH>.

882

883 Takeshi Koshizuka and Issei Sato. Neural lagrangian schr\”{o}dinger bridge: Diffusion modeling  
 884 for population dynamics. In *The Eleventh International Conference on Learning Representations*,  
 885 2023.

886

887 Chrysoula Kosma, Giannis Nikolentzos, George Panagopoulos, Jean-Marc Steyaert, and Michalis  
 888 Vazirgiannis. Neural ordinary differential equations for modeling epidemic spreading.  
 889 *Transactions on Machine Learning Research*, 2023.

890

891 Leonenko Kozachenko. Sample estimate of the entropy of a random vector. *Probl. Pered. Inform.*,  
 892 23:9, 1987.

893

894 Rahul Krishnan, Uri Shalit, and David Sontag. Structured inference networks for nonlinear state  
 895 space models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.

896

897 Marc Lambert, Sinho Chewi, Francis Bach, Silvère Bonnabel, and Philippe Rigollet. Variational  
 898 inference via wasserstein gradient flows. *Advances in Neural Information Processing Systems*,  
 899 35:14434–14447, 2022.

900

901 Jeffrey M Lane and Richard F Riesenfeld. A theoretical development for the computer genera-  
 902 tion and display of piecewise polynomial surfaces. *IEEE Transactions on Pattern Analysis and*  
*Machine Intelligence*, (1):35–46, 1980.

903

904 Hugo Lavenant, Stephen Zhang, Young-Heon Kim, Geoffrey Schiebinger, et al. Toward a math-  
 905 ematical theory of trajectory inference. *The Annals of Applied Probability*, 34(1A):428–500,  
 906 2024.

907

908 Wonjun Lee, Li Wang, and Wuchen Li. Deep jko: time-implicit particle methods for general non-  
 909 linear gradient flows. *Journal of Computational Physics*, 514:113187, 2024.

910

911 Ruikun Li, Jiazen Liu, Huandong Wang, Qingmin Liao, and Yong Li. Weightflow: Learning  
 912 stochastic dynamics via evolving weight of neural network. *arXiv preprint arXiv:2508.00451*,  
 913 2025a.

914

915 Xuechen Li, Ting-Kam Leonard Wong, Ricky TQ Chen, and David Duvenaud. Scalable gradients  
 916 for stochastic differential equations. In *International Conference on Artificial Intelligence and*  
*Statistics*, pp. 3870–3882. PMLR, 2020.

917

918 Zihao Li, Zhichen Zeng, Xiao Lin, Feihao Fang, Yanru Qu, Zhe Xu, Zhining Liu, Xuying Ning,  
 919 Tianxin Wei, Ge Liu, et al. Flow matching meets biology and life science: A survey. *arXiv*  
*e-prints*, pp. arXiv–2507, 2025b.

918 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching  
 919 for generative modeling. In 11th International Conference on Learning Representations, ICLR  
 920 2023, 2023.

921

922 Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, and  
 923 Ricky T. Q. Chen. Generalized schrödinger bridge matching. In The Twelfth International  
 924 Conference on Learning Representations, 2024. URL <https://openreview.net/forum?id=SoismgeX7z>.

925

926 Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference  
 927 algorithm. Advances in neural information processing systems, 29, 2016.

928

929 Xingchao Liu and Lemeng Wu. Learning diffusion bridges on constrained domains. In international  
 930 conference on learning representations (ICLR), 2023.

931

932 Xingchao Liu, Chengyue Gong, and qiang liu. Flow straight and fast: Learning to generate  
 933 and transfer data with rectified flow. In The Eleventh International Conference on Learning  
 934 Representations, 2023. URL <https://openreview.net/forum?id=XVjTT1nw5z>.

935

936 Zewen Liu, Xiaoda Wang, Bohan Wang, Zijie Huang, Carl Yang, and Wei Jin. Graph odes and be-  
 937 yond: A comprehensive survey on integrating differential equations with graph neural networks.  
 938 In Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining  
 939 V. 2, pp. 6118–6128, 2025.

940

941 Evan Z Macosko, Anindita Basu, Rahul Satija, James Nemesh, Karthik Shekhar, Melissa Goldman,  
 942 Itay Tirosh, Allison R Bialas, Nolan Kamitaki, Emily M Martersteck, et al. Highly parallel  
 943 genome-wide expression profiling of individual cells using nanoliter droplets. Cell, 161(5):1202–  
 944 1214, 2015.

945

946 Bertrand Maury, Aude Roudneff-Chupin, and Filippo Santambrogio. A macroscopic crowd motion  
 947 model of gradient flow type. Mathematical Models and Methods in Applied Sciences, 20(10):  
 948 1787–1821, 2010.

949

950 Bertrand Maury, Aude Roudneff-Chupin, Filippo Santambrogio, and Juliette Venel. Handling con-  
 951 gestion in crowd motion modeling. Networks and Heterogeneous Media, 6(3):485–519, 2011.

952

953 Robert J McCann. Existence and uniqueness of monotone measure-preserving maps. Duke  
 954 Mathematical Journal, 80(2):309, 1995.

955

956 Boris Medina-Salgado, Eddy Sánchez-DelaCruz, Pilar Pozos-Parra, and Javier E Sierra. Urban  
 957 traffic flow prediction techniques: A review. Sustainable Computing: Informatics and Systems,  
 958 35:100739, 2022.

959

960 Nicholas Metropolis and Stanislaw Ulam. The monte carlo method. Journal of the American  
 961 statistical association, 44(247):335–341, 1949.

962

963 Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for  
 964 generative adversarial networks. In International Conference on Learning Representations, 2018.  
 965 URL <https://openreview.net/forum?id=B1QRgziT->.

966

967 Petr Mokrov, Alexander Korotin, Lingxiao Li, Aude Genevay, Justin M Solomon, and Evgeny  
 968 Burnaev. Large-scale wasserstein gradient flows. Advances in Neural Information Processing  
 969 Systems, 34:15243–15256, 2021.

970

971 Gaspar Monge. Mémoire sur la théorie des déblais et des remblais. Mem. Math. Phys. Acad.  
 972 Royale Sci., pp. 666–704, 1781.

973

974 Kevin R Moon, David Van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen,  
 975 Kristina Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, et al. Visualizing  
 976 structure and transitions in high-dimensional biological data. Nature biotechnology, 37(12):1482–  
 977 1492, 2019.

972 Alex Morehead, Lazar Atanackovic, Akshata Hegde, Yanli Wang, Frimpong Boadu, Joel Sel-  
 973 varaj, Alexander Tong, Aditi Krishnapriyan, and Jianlin Cheng. How to go with the flow:  
 974 flow matching in bioinformatics and computational biology. *Authorea Preprints*, August 2025.  
 975 doi: 10.22541/au.175382408.89466370/v2. URL <http://dx.doi.org/10.22541/au.175382408.89466370/v2>.

977 Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learn-  
 978 ing stochastic dynamics from samples. In *International conference on machine learning*, pp.  
 979 25858–25889. PMLR, 2023.

981 Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza  
 982 Makhzani. A computational framework for solving wasserstein lagrangian flows. In *Proceedings*  
 983 of the 41st International Conference on Machine Learning, pp. 37461–37485, 2024a.

984 Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, qiang liu, and Alireza  
 985 Makhzani. A computational framework for solving wasserstein lagrangian flows. In *Forty-first*  
 986 *International Conference on Machine Learning*, 2024b. URL <https://openreview.net/forum?id=wwItuHdus6>.

989 Atsushi Nitanda. Improved particle approximation error for mean field neural networks. In  
 990 *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL  
 991 <https://openreview.net/forum?id=76NKidadct>.

992 Ivan V Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):  
 993 2295–2317, 2011.

995 F Otto. The geometry of dissipative evolution equations: the porous medium equation. *Comm.*  
 996 *Partial Differential Equations*, 26:101–174, 2001.

997 Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and trends® in Optimization*,  
 998 1(3):127–239, 2014.

1000 Min Sue Park, Cheolhyeong Kim, Hwijae Son, and Hyung Ju Hwang. The deep minimizing move-  
 1001 ment scheme. *Journal of Computational Physics*, 494:112518, 2023.

1002 A Paszke. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint*  
 1003 [arXiv:1912.01703](https://arxiv.org/abs/1912.01703), 2019.

1005 Stefano Peluchetti. Diffusion bridge mixture transports, schrödinger bridge problems and generative  
 1006 modeling. *Journal of Machine Learning Research*, 24(374):1–51, 2023a.

1008 Stefano Peluchetti. Non-denoising forward-time diffusions. *arXiv preprint arXiv:2312.14589*,  
 1009 2023b.

1010 Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data  
 1011 science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

1012 Simon Plazotta and Jonathan Zinsl. High-frequency limit of non-autonomous gradient flows of  
 1013 functionals with time-periodic forcing. *Journal of Differential Equations*, 261(12):6806–6855,  
 1014 2016.

1016 Jonathan R Potts. Aggregation-diffusion in heterogeneous environments. *arXiv preprint*  
 1017 [arXiv:2409.10147](https://arxiv.org/abs/2409.10147), 2024.

1019 Jonathan R Potts and Kevin J Painter. Distinguishing between long-transient and asymptotic states  
 1020 in a biological aggregation model. *Bulletin of Mathematical Biology*, 86(3):28, 2024.

1021 Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic  
 1022 Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Probabilistic  
 1023 weather forecasting with machine learning. *Nature*, 637(8044):84–90, 2025.

1025 Hannes Risken and Till Frank. *The Fokker-Planck Equation: Methods of Solution and Applications*,  
 volume 18. Springer Science & Business Media, 1996.

1026 R. Tyrrell Rockafellar. *Convex analysis*. Princeton Mathematical Series. Princeton University Press,  
 1027 Princeton, N. J., 1970.

1028

1029 Martin Rohbeck, Charlotte Bunne, Edward De Brouwer, Jan-Christian Huetter, Anne Biton,  
 1030 Kelvin Y. Chen, Aviv Regev, and Romain Lopez. Modeling complex system dynamics with flow  
 1031 matching across time and conditions. In *The Thirteenth International Conference on Learning  
 1032 Representations*, 2025. URL <https://openreview.net/forum?id=hwnObmOTrV>.

1033 Litu Rout, Alexander Korotin, and Evgeny Burnaev. Generative modeling with optimal transport  
 1034 maps. In *Proceedings of the 10th International Conference on Learning Representations (ICLR)*,  
 1035 2022. URL <https://openreview.net/forum?id=5JdLZg346Lw>.

1036

1037 Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. A metric for distributions with applica-  
 1038 tions to image databases. In *Sixth international conference on computer vision (IEEE Cat. No.  
 1039 98CH36271)*, pp. 59–66. IEEE, 1998.

1040 Adil Salim, Anna Korba, and Giulia Luise. The wasserstein proximal gradient algorithm. *Advances  
 1041 in Neural Information Processing Systems*, 33:12356–12366, 2020.

1042

1043 Filippo Santambrogio. *Optimal transport for applied mathematicians*, volume 87. Springer, 2015.

1044

1045 Filippo Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin  
 1046 of Mathematical Sciences*, 7:87–154, 2017.

1047

1048 Daniel Saragih, Deyu Cao, and Tejas Balaji. Flows and diffusions on the neural manifold. *arXiv  
 1049 e-prints*, pp. arXiv–2507, 2025.

1050

1051 Simo Särkkä and Arno Solin. *Applied stochastic differential equations*, volume 10. Cambridge  
 University Press, 2019.

1052

1053 Christopher Scarvelis and Justin Solomon. Riemannian metric learning via optimal transport. In  
 1054 *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=v3y68gz-WEz>.

1055

1056 Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon,  
 1057 Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell  
 1058 gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943,  
 2019.

1059

1060 Yutong Sha, Yuchi Qiu, Peijie Zhou, and Qing Nie. Reconstructing growth and dynamic trajectories  
 1061 from single-cell transcriptomics data. *Nature Machine Intelligence*, 6(1):25–39, 2024.

1062

1063 Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal schrödinger bridges with  
 1064 iterative reference refinement. In *The 28th International Conference on Artificial Intelligence and  
 1065 Statistics*, 2025. URL <https://openreview.net/forum?id=VcwZ3gtYFY>.

1066

1067 Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger  
 bridge matching. *Advances in Neural Information Processing Systems*, 36:62183–62223, 2023.

1068

1069 Fabio Sigrist, Hans R Künsch, and Werner A Stahel. Stochastic partial differential equation based  
 1070 modelling of large space–time data sets. *Journal of the Royal Statistical Society Series B:  
 1071 Statistical Methodology*, 77(1):3–33, 2015.

1072

1073 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 1074 Poole. Score-based generative modeling through stochastic differential equations. In *International  
 1075 Conference on Learning Representations*, 2021.

1076

1077 Yuhao Sun, Zhenyi Zhang, Zihan Wang, Tiejun Li, and Peijie Zhou. Variational regularized unbal-  
 1078 anced optimal transport: Single network, least action. *arXiv preprint arXiv:2505.11823*, 2025.

1079 Taiji Suzuki, Atsushi Nitanda, and Denny Wu. Uniform-in-time propagation of chaos for the  
 mean-field gradient langevin dynamics. In *The Eleventh International Conference on Learning  
 Representations*, 2023. URL [https://openreview.net/forum?id=\\_JScUk9TBUn](https://openreview.net/forum?id=_JScUk9TBUn).

1080 Chandan Tankala, Dheeraj M Nagaraj, and Anant Raj. Beyond propagation of chaos: A stochastic  
 1081 algorithm for mean field optimization. [arXiv preprint arXiv:2503.13115](https://arxiv.org/abs/2503.13115), 2025.

1082

1083 Antonio Terpin, Nicolas Lanzetti, Martín Gadea, and Florian Dorfler. Learning diffusion at light-  
 1084 speed. [Advances in Neural Information Processing Systems](https://proceedings.neurips.cc/paper/2024/file/6797-6832.pdf), 37:6797–6832, 2024.

1085

1086 Gerald Teschl. [Ordinary differential equations and dynamical systems](https://www.ams.org/quarterly/2012-140-01/01086), volume 140. American  
 1087 Mathematical Soc., 2012.

1088

1089 Panagiotis Theodoropoulos, Augustinos D Saravacos, Evangelos A Theodorou, and Guan-  
 1090 Horng Liu. Momentum multi-marginal schrödinger bridge matching. [arXiv preprint  
 1090 arXiv:2506.10168](https://arxiv.org/abs/2506.10168), 2025.

1091

1092 Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet:  
 1093 A dynamic optimal transport network for modeling cellular dynamics. In [International conference  
 1093 on machine learning](https://proceedings.neurips.cc/paper/2020/file/9526-9536.pdf), pp. 9526–9536. PMLR, 2020.

1094

1095 Alexander Tong, Manik Kuchroo, Shabarni Gupta, Aarthi Venkat, Beatriz P San Juan, Laura Rangel,  
 1096 Brandon Zhu, John G Lock, Christine L Chaffer, and Smita Krishnaswamy. Learning transcrip-  
 1097 tional and regulatory dynamics driving cancer cell plasticity using neural ode-based optimal trans-  
 1098 port. [bioRxiv](https://arxiv.org/abs/2303.02023), pp. 2023–03, 2023.

1099

1100 Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-  
 1101 Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models  
 1102 with minibatch optimal transport. [Transactions on Machine Learning Research](https://proceedings.mlr.press/v160/tong24a/tong24a.pdf), pp. 1–34, 2024a.

1103

1104 Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume  
 1105 Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow  
 1106 matching. In [The 27th International Conference on Artificial Intelligence and Statistics](https://proceedings.mlr.press/v160/tong24b/tong24b.pdf), pp. 1279–  
 1107 1287. Journal of Machine Learning Research-Proceedings Track, 2024b.

1108

1109 Théo Uscidda and Marco Cuturi. The monge gap: A regularizer to learn all transport maps. In  
 1110 [International Conference on Machine Learning](https://proceedings.mlr.press/v160/uscidda23/uscidda23.pdf), pp. 34709–34733. PMLR, 2023.

1111

1112 Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving schrödinger  
 1113 bridges via maximum likelihood. [Entropy](https://doi.org/10.3390/entropy-23-1134), 23(9):1134, 2021.

1114

1115 Juan Luis Vázquez. Asymptotic behaviour methods for the heat equation. convergence to the gaus-  
 1116 sian. [arXiv preprint arXiv:1706.10034](https://arxiv.org/abs/1706.10034), 2017.

1117

1118 Yogesh Verma, Markus Heinonen, and Vikas Garg. ClimODE: Climate and weather forecast-  
 1119 ing with physics-informed neural ODEs. In [The Twelfth International Conference on Learning  
 1119 Representations](https://openreview.net/forum?id=xuY33XhEGR), 2024. URL <https://openreview.net/forum?id=xuY33XhEGR>.

1120

1121 Alexander Vidal, Samy Wu Fung, Luis Tenorio, Stanley Osher, and Levon Nurbekyan. Taming  
 1122 hyperparameter tuning in continuous normalizing flows using the jko scheme. [Scientific reports](https://doi.org/10.1038/s41550-023-14501-1),  
 1123 13(1):4501, 2023.

1124

1125 Cédric Villani et al. [Optimal transport: old and new](https://www.springer.com/978-3-540-3610-0), volume 338. Springer, 2008.

1126

1127 Wei Wan, Yuejin Zhang, Chenglong Bao, Bin Dong, and Zuoqiang Shi. A scalable deep learning  
 1128 approach for solving high-dimensional dynamic optimal transport. [SIAM Journal on Scientific  
 1129 Computing](https://doi.org/10.1137/21m1437007), 2023.

1130

1131 Dongyi Wang, Yuanwei Jiang, Zhenyi Zhang, Xiang Gu, Peijie Zhou, and Jian Sun. Joint velocity-  
 1132 growth flow matching for single-cell dynamics modeling. [arXiv preprint arXiv:2505.13413](https://arxiv.org/abs/2505.13413), 2025.

1133

1134 Xiaomei Wang, Qi An, Zilong He, and Wei Fang. A literature review of social network analysis in  
 1135 epidemic prevention and control. [Complexity](https://doi.org/10.1002/cplx.28221), 2021(1):3816221, 2021.

1136

1137 E Weinan, Tiejun Li, and Eric Vanden-Eijnden. [Applied stochastic analysis](https://www.ams.org/quarterly/2021-199-01/01137), volume 199. American  
 1138 Mathematical Soc., 2021.

1139

1140 Shang Wu and Yazhen Wang. Computational and statistical asymptotic analysis of the jko scheme  
 1141 for iterative algorithms to update distributions. [arXiv preprint arXiv:2501.06408](https://arxiv.org/abs/2501.06408), 2025.

1134 Chen Xu, Xiuyuan Cheng, and Yao Xie. Normalizing flow neural networks by jko scheme.  
 1135 *Advances in Neural Information Processing Systems*, 36:47379–47405, 2023.  
 1136

1137 Luxuan Yang, Ting Gao, Yubin Lu, Jinqiao Duan, and Tao Liu. Neural network stochastic dif-  
 1138 ferential equation models with applications to financial data forecasting. *Applied Mathematical*  
 1139 *Modelling*, 115:279–299, 2023.

1140 Stephen Zhang, Suryanarayana Maddu, Xiaojie Qiu, and Victor Chardès. Inferring stochastic dy-  
 1141 namics with growth from cross-sectional data. *arXiv preprint arXiv:2505.13197*, 2025a.  
 1142

1143 Yanbo Zhang and Michael Levin. Equilibrium flow: From snapshots to dynamics. *arXiv preprint*  
 1144 *arXiv:2509.17990*, 2025.

1145 Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through  
 1146 regularized unbalanced optimal transport. In *The Thirteenth International Conference on Learning*  
 1147 *Representations*, 2025b. URL <https://openreview.net/forum?id=gQ1xd3Mtru>.  
 1148

1149 Zhenyi Zhang, Yuhao Sun, Qiangwei Peng, Tiejun Li, and Peijie Zhou. Integrating dynamical  
 1150 systems modeling with spatiotemporal scRNA-seq data analysis. *Entropy*, 27(5):453, 2025c.  
 1151

1152 Zhenyi Zhang, Zihan Wang, Yuhao Sun, Tiejun Li, and Peijie Zhou. Modeling cell dynamics and  
 1153 interactions with unbalanced mean field schrödinger bridge. *arXiv preprint arXiv:2505.11197*,  
 1154 2025d.

1155 Zhenyi Zhang, Zihan Wang, Yuhao Sun, Jiantao Shen, Qiangwei Peng, Tiejun Li, and Peijie Zhou.  
 1156 Deciphering cell-fate trajectories using spatiotemporal single-cell transcriptomic data. *Authorea*  
 1157 *Preprints*, September 2025e. doi: 10.22541/au.175735334.45794622/v1. URL <http://dx.doi.org/10.22541/au.175735334.45794622/v1>.  
 1158

1159 Qinqing Zheng, Matt Le, Neta Shaul, Yaron Lipman, Aditya Grover, and Ricky TQ Chen. Guided  
 1160 flows for generative modeling and decision making. *arXiv preprint arXiv:2311.13443*, 2023.  
 1161

1162 Huminhao Zhu, Fangyikang Wang, Chao Zhang, Hanbin Zhao, and Hui Qian. Neural sinkhorn  
 1163 gradient flow. *arXiv preprint arXiv:2401.14069*, 2024.  
 1164  
 1165  
 1166  
 1167  
 1168  
 1169  
 1170  
 1171  
 1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187

1188	CONTENTS	
1189		
1190		
1191	<b>A Extended Discussion</b>	<b>23</b>
1192	A.1 Related works . . . . .	23
1193	A.1.1 Learning population dynamics . . . . .	23
1194	A.1.2 JKO scheme meets deep learning . . . . .	25
1195	A.2 Time-varying potentials discussion . . . . .	26
1196	A.3 Paired vs. unpaired setups . . . . .	26
1197		
1198		
1199	<b>B Extended Comparisons</b>	<b>27</b>
1200	B.1 Synthetic comparisons . . . . .	27
1201	B.1.1 Learning potential energy . . . . .	28
1202	B.1.2 Learning interaction energy . . . . .	29
1203	B.1.3 Learning internal energy . . . . .	29
1204	B.2 Learning single-cell dynamics . . . . .	32
1205		
1206		
1207		
1208	<b>C Training and Model Details</b>	<b>33</b>
1209	C.1 Metric computation details . . . . .	33
1210	C.2 Training details . . . . .	34
1211	C.3 Entropy estimation details . . . . .	34
1212	C.4 Optimizer . . . . .	35
1213	C.5 Network architecture . . . . .	35
1214	C.6 Hardware . . . . .	36
1215	C.7 Functionals . . . . .	36
1216		
1217		
1218		
1219		
1220		
1221	<b>D Proofs</b>	<b>36</b>
1222		
1223		
1224	<b>A EXTENDED DISCUSSION</b>	
1225		
1226	We first review related work in §A.1 and then discuss the validity of time-varying potentials pro-	
1227	posed by Terpin et al. (2024) in §A.2 in contrast to the stationary setting considered in this paper.	
1228	Finally, in §A.3, we discuss potential reasons for the performance gap between the paired and un-	
1229	paired setups introduced in §5.1.	
1230		
1231	<b>A.1 RELATED WORKS</b>	
1232	In this section, we first provide an extended discussion of related work on learning population dy-	
1233	namics (also known as trajectory inference) in §A.1.1, followed by an overview of applications of	
1234	the JKO scheme (Jordan et al., 1998) in deep learning in §A.1.2.	
1235		
1236	<b>A.1.1 LEARNING POPULATION DYNAMICS</b>	
1237	In contrast to our approach, which models population dynamics as a Wasserstein gradient flow	
1238	(WGF) (Ambrosio et al., 2008) of the free-energy functional (5) (Gómez-Castro, 2024), several	
1239	other fruitful approaches have been proposed in the literature. We review these alternatives below.	
1240		
1241	<b>RNNs and Neural ODEs.</b> One of the first works on learning population dynamics is (Hashimoto et al., 2016), which used recurrent neural networks (RNNs) to learn the SDE (7) as a WGF of a poten-	

1242      trial energy functional. Later, (Chen et al., 2018a, CNF) introduced neural ODEs as the continuous-  
 1243      time limit of RNNs for learning ODE-based dynamics, and (Li et al., 2020) subsequently extended  
 1244      this approach to handle SDEs. Building on these ideas, (Erbe et al., 2023, RNAForecaster) ap-  
 1245      plied neural ODEs to predict future transcriptomic states of single cells (Battich et al., 2020). **More**  
 1246      **recently**, (Li et al., 2025a) proposed **WeightFlow**, which models stochastic population dynamics  
 1247      by learning the continuous evolution of neural network weights that parameterize distributions at  
 1248      each time point, rather than modeling the dynamics directly in the state space. Their key idea is  
 1249      to project the evolving distribution into the weight space of a backbone neural network trained to  
 1250      represent the probability distribution at each snapshot.

1251      **Static OT.** Another line of work models population dynamics by directly learning  $K$  *independent*  
 1252      transport maps from time-snapshot data, using the *static* optimal transport formulation (1) or (2),  
 1253      without attempting to recover the unifying energy functionals guiding the system evolution. For  
 1254      example, (Schiebinger et al., 2019, Waddington-OT) employs discrete transport maps, whereas  
 1255      (Bunne et al., 2023, CellOT) learns *neural* transport maps, as in (Fan et al., 2023, Eq.(6)), (Rout  
 1256      et al., 2022, Eq.(14)), and (Korotin et al., 2023, Eq.(15)), see Appendix A.2 for more details.

1257      **Dynamic OT.** (Tong et al., 2020, TrajectoryNet) models population dynamics using the dy-  
 1258      namic optimal transport formulation of (Benamou & Brenier, 2000), parameterized by a regularized  
 1259      CNF. The work (Huguet et al., 2022, MIOFlow) later extended this approach by incorporating a  
 1260      geodesic autoencoder to better capture manifold structure. More recently, (Wan et al., 2023) pro-  
 1261      posed computational techniques that make dynamic OT scalable to higher-dimensional settings.

1262      **Flow Matching.** The works in this paragraph extend Flow Matching (FM) techniques (Lipman  
 1263      et al., 2023; Albergo & Vanden-Eijnden, 2023; Liu et al., 2023) in several directions. Conditional  
 1264      Flow Matching (Tong et al., 2024a, CFM) introduce simulation-free objectives for learning deter-  
 1265      ministic flows. Building on the celebrated work of (Otto, 2001), the Wasserstein space can be nominally  
 1266      (i.e., heuristically) viewed as a Riemannian manifold. Motivated by this observation, Wasserstein  
 1267      Flow Matching (Haviv et al., 2025, WFM) applies the idea of Riemannian Flow Matching (Chen &  
 1268      Lipman, 2024, RFM) to the Wasserstein space, i.e., performing FM over a distribution of distribu-  
 1269      tions. Further developments, such as Meta Flow Matching (Atanackovic et al., 2025, Meta FM),  
 1270      embed the population of samples with a Graph Neural Network (GNN) (Liu et al., 2025) and use  
 1271      these embeddings as conditioning inputs for the learning vector field. Finally, Multi-Marginal Flow  
 1272      Matching (Rohbeck et al., 2025, MMFM) constructs a flow using smooth spline-based interpolation  
 1273      across time points and conditions, and regresses it with a neural network under the classifier-free  
 1274      guided Flow Matching framework (Zheng et al., 2023). For a broader overview of FM for learning  
 1275      population dynamics, see the recent surveys (Morehead et al., 2025; Li et al., 2025b).

1276      **Action Matching.** (Neklyudov et al., 2023, AM) introduced Action Matching (AM), a method that  
 1277      optimizes an action-gap objective. In contrast to Flow Matching, AM learns a vector field expressed  
 1278      as the gradient of an action,  $\nabla s_t^*$ , which uniquely defines the velocity field that transports particles  
 1279      optimally in the sense of optimal transport. Building on this idea, (Berman et al., 2024) conditioned  
 1280      the action  $s_{t,\mu}$  on physical parameters  $\mu$ , used (Berman & Peherstorfer, 2024, CoLoRa) for efficient  
 1281      action parametrization, and demonstrated its utility for surrogate modeling of classical numerical  
 1282      solvers, enabling fast prediction of system behavior across different physics parameter settings.  
 1283      Most recently, (Neklyudov et al., 2024b, WLF) proposed Wasserstein Lagrangian Flows, a unifying  
 1284      framework that minimizes Lagrangian action functionals over the space of probability densities  
 1285      rather than the ground space, thereby encapsulating AM as a special case.

1286      **Schrödinger bridges.** Two seminal works, Vargas et al. (2021, IPF (GP)) and De Bortoli et al.  
 1287      (2021, IPF (NN)), proposed solving the Schrödinger bridge problem using the Iterative Pro-  
 1288      portional Fitting (IPF) algorithm (Gramer, 2000) in application to generative modeling, which alter-  
 1289      nates between forward and backward processes. The first approach employed Gaussian processes  
 1290      as parametrization, whereas the second relied on score-based neural networks. Later, bridge match-  
 1291      ing methods was developed (Peluchetti, 2023b;a; Liu & Wu, 2023; Shi et al., 2023). Tong et al.  
 1292      (2024b, [SF]<sup>2</sup>M) introduced simulation-free objectives for both deterministic and stochastic flows,  
 1293      further reducing computational overhead. Additional contributions include generalized setups for  
 1294      Schrödinger bridges (Koshizuka & Sato, 2023; Liu et al., 2024), multi-marginal generalizations  
 1295      of the problem (Chen et al., 2023, DMSB), (Shen et al., 2025, MMSB), (Hong et al., 2025), and  
 1296      momentum-accelerated formulations (Theodoropoulos et al., 2025, 3MSBM).

**Unbalanced OT.** Several works jointly model marginal transitions and growth dynamics by minimizing the action in the Wasserstein-Fisher-Rao (WFR) metric (Chizat et al., 2018a), i.e., by solving the dynamical unbalanced optimal transport problem (Chizat et al., 2018b). (Tong et al., 2023) proposed `BEMIOflow`, an extension of (Huguet et al., 2022, `MIOflow`), that incorporates a neural network to predict cell growth and death rates continuously in time, thereby augmenting optimal transport with population-size dynamics. Later, (Chen et al., 2022b) formulated an unbalanced Schrödinger Bridge problem, which was applied in practice by (Zhang et al., 2025b, `DeepRUOT`), following the developments in (Sha et al., 2024, `TIGON`). `DeepRUOT` was further simplified into a single-network formulation, `Var-RUOT` (Sun et al., 2025), by exploiting optimality conditions; extended to include interaction modeling in (Zhang et al., 2025d, `CytoBridge`); and adapted to use two independent networks to separately model distributional drift and mass growth (Wang et al., 2025, `VGFM`), (Zhang et al., 2025a). In a different direction, (Klein et al., 2024, `GENOT`) generalized OT to simultaneously handle stochasticity, entropy regularization, and unbalanced mass transport, enabling applications to cross-modal and heterogeneous data. For a broader survey of trajectory inference methods and their applications to biology data, see (Zhang et al., 2025e;c).

The concurrent work (Andrade et al., 2025, `iJKO`) also connects inverse optimization with the JKO scheme (8) but in the unbalanced setting. However, their focus is on sample complexity, the method addresses only potential energy, and experiments are restricted to low-dimensional cases.

**Riemannian perspective.** (Scarvelis & Solomon, 2023) proposed learning a metric tensor  $A(x)$  that minimizes the average 1-Wasserstein distance on the learned manifold between pairs of consecutive population snapshots. Furthermore, (Kapusniak et al., 2024, `MFM`) introduced Metric Flow Matching (`MFM`), which learns interpolants  $x_{t,\eta} = (1-t)x_0 + tx_1 + t(1-t)\varphi_{t,\eta}(x_0, x_1)$ , where  $\eta$  are the parameters of a neural network  $\varphi_{t,\eta}$  providing a nonlinear “correction” to straight-line interpolants (Albergo & Vanden-Eijnden, 2023). The resulting velocity field minimizes a data-dependent Riemannian metric, assigning lower transport cost to regions with higher data density.

**Splines in Wasserstein Space.** In the seminal work of (Schiebinger et al., 2019), a piecewise linear OT interpolation method was proposed to infer cell trajectories. Subsequent works introduced higher-order piecewise polynomials (e.g., cubic splines) in Wasserstein space: (Chen et al., 2018b; Benamou et al., 2019) formulated a global cubic spline minimization problem, which was later extended by (Chewi et al., 2021) to use Euclidean interpolation algorithm in  $\mathbb{R}^D$  after a finding optimal Monge map (2) between samples from consecutive measures, with further refinements by (Clancy & Suarez, 2022; Justiniano et al., 2024). More recently, (Banerjee et al., 2025, `WLR`) proposed the Wasserstein Lane–Riesenfeld method, leveraging the classical subdivision algorithm of (Lane & Riesenfeld, 1980). In parallel, (Dyn & Sharon, 2025) developed subdivision schemes for general metric spaces, including the Wasserstein space. (Baccou & Liandrat, 2024; Kawano et al., 2025) combines subdivision schemes with OT.

**Dynamics from stationary distributions.** A recent line of research aims to recover meaningful dynamical structure from static data alone. The Equilibrium Flow framework of (Zhang & Levin, 2025) addresses this problem by learning a vector field whose flow preserves the observed stationary distribution  $p$ . Their method enforces the invariance condition  $\nabla \cdot [p(x, t)v(x, t)] = 0$  using score-based estimates of  $\nabla \log p$ , enabling the recovery of plausible dynamics even without temporal information or trajectories. Despite this minimal supervision, the learned flows can be used for inverse design, where one constructs dynamics that yield a desired target distribution.

### A.1.2 JKO SCHEME MEETS DEEP LEARNING

**Generative Modeling.** (Vidal et al., 2023, `JKO-Flow`) and (Xu et al., 2023, `JKO-iFlow`) reinterpret the JKO scheme (8) through the framework of CNFs (Chen et al., 2018b), applying it to 2D synthetic data and image generation tasks, respectively. Subsequently, (Cheng et al., 2024) provided theoretical convergence guarantees for this approach. (Choi et al., 2024, `S-JKO`) further accelerated WGF-based generative modeling by leveraging a semi-dual unbalanced OT formulation. They constructed the WGF of an  $f$ -divergence  $D_f$  and used an (Song et al., 2021, `NCSN++`) backbone – similar to (Zhu et al., 2024, `NSGF`), who designed a generative model as the WGF with respect to the Sinkhorn divergence (Peyré et al., 2019).

**Tensor Train (TT) (Oseledets, 2011).** (Aksenov & Eigel, 2025) introduces a method for approximating probability distributions in Bayesian inversion by minimizing a KL-based functional via an

1350 entropically regularized JKO scheme. The approach discretizes the resulting coupled heat equations  
 1351 on a high-dimensional grid and solves them using accelerated fixed-point methods combined with  
 1352 low-rank TT representations. Related work includes (Chertkov & Oseledets, 2021), which applies  
 1353 TT approximations to the Fokker-Planck equation with drift and diffusion, and (Han et al., 2025),  
 1354 which employs entropy-regularized proximal steps with TT-cross for particle-based evolution.

1355 **Variational Inference.** (Lambert et al., 2022) and (Diao et al., 2023) study variational inference  
 1356 through the lens of Bures–Wasserstein gradient flows. (Cheng et al., 2023, GWG) introduce a gener-  
 1357 alized minimizing movement scheme on the space  $\mathcal{P}_{c_h}(\mathbb{R}^D)$ , where the transport cost is defined as  
 1358  $c_h(x, y) = g\left(\frac{x-y}{h}\right)h$  with  $g$  a continuously differentiable Young function, and apply this framework  
 1359 to particle-based variational inference.

1360 **Datasets/Weights Learning.** (Bonet et al., 2025) propose a explicit scheme that is computationally  
 1361 more efficient in practice than the implicit JKO scheme on the space  $\mathcal{P}_2(\mathcal{P}_2(\mathbb{R}^D))$ , and apply this  
 1362 approach to modeling flow datasets viewed as random measures. (Saragih et al., 2025) train a  
 1363 meta-model, based on JKO as well as other flow- and diffusion-based approaches, that generates  
 1364 dataset-conditioned classifiers by producing the neural network weights in a single forward pass.  
 1365

## 1366 A.2 TIME-VARYING POTENTIALS DISCUSSION

1367 In this section, we examine the validity of using time-varying potentials, as proposed in (Terpin et al.,  
 1368 2024, §4.4). We demonstrate that under this formulation, the original problem of reconstructing  
 1369 energy functionals (Bunne et al., 2022b) via JKO Scheme reduces to learning  $K$  independent neural  
 1370 optimal transport maps (Fan et al., 2023, Eq.(6)), (Rout et al., 2022, Eq.(14)) between the data  
 1371 snapshots  $\rho_k$  and  $\rho_{k+1}$ . To illustrate this, consider the loss in (11) and assume that, instead of a single  
 1372 functional  $\mathcal{J}$ , we now assign a separate only potential energy functional  $\mathcal{J}_{\text{PE}}^k = \int_{\mathcal{X}} V^k(x) d\rho_k(x)$   
 1373 for each time step  $k$ :

$$1375 \max_{\mathcal{J}_{\text{PE}}^k} \min_{T^k} \mathcal{L}(\mathcal{J}^k, T^k) \stackrel{\text{def}}{=} \max_{\mathcal{J}_{\text{PE}}^k} \min_{T^k} \sum_{k=0}^{K-1} \left[ \mathcal{J}_{\text{PE}}^k(T^k \# \rho_k) - \mathcal{J}_{\text{PE}}^k(\rho_{k+1}) + \frac{1}{2\tau} \int_{\mathcal{X}} \|x - T^k(x)\|_2^2 \rho_k(x) dx \right].$$

1378 It then becomes clear that the loss optimization decomposes into  $K$  independent terms  $\mathcal{L}^k(\mathcal{J}_{\text{PE}}^k, T^k)$ :

$$1380 \mathcal{L}^k(\mathcal{J}_{\text{PE}}^k, T^k) \stackrel{\text{def}}{=} \mathcal{J}_{\text{PE}}^k(T^k \# \rho_k) - \mathcal{J}_{\text{PE}}^k(\rho_{k+1}) + \frac{1}{2\tau} \int_{\mathcal{X}} \|x - T^k(x)\|_2^2 \rho_k(x) dx \rightarrow \max_{\mathcal{J}_{\text{PE}}^k} \min_{T^k}. \quad (18)$$

1383 By replacing  $V^k$  with the dual potential  $f_\eta$  and  $T^k$  with  $T_\theta$ , we recover Equation (6) from (Fan  
 1384 et al., 2023), which aligns with the objective proposed in (Bunne et al., 2023, Eq. (9)) for modeling  
 1385 single-cell dynamics. Thus, (Terpin et al., 2024) introduces a new strategy for minimizing this loss  
 1386 rather than directly addressing the recovery of ground-truth energy functionals (Bunne et al., 2022b).

1387 Fortunately, the theoretical framework developed by Ferreira & Valencia-Guevara (2018) and later  
 1388 extended by Plazotta & Zinsl (2016) establishes the validity of this approach even for time-varying  
 1389 free energy functionals (5), not only for the potential energy case, in the limit as the time discretiza-  
 1390 tion tends to zero. However, this connection is not acknowledged in (Terpin et al., 2024).

## 1392 A.3 PAIRED VS. UNPAIRED SETUPs

1394 In this section, we discuss the substantial performance gap observed between the paired and unpaired  
 1395 setups, as illustrated in Table 3. We suggest that this discrepancy stems from the intrinsic connection  
 1396 between the considered solvers,  $\text{iJKOnet}$  and  $\text{JKOnet}^*$ , and the underlying OT problem (2).

1397 Both solvers are designed to recover the optimal transport mappings  $T^{k,*}$  between consecutive dis-  
 1398 tributions  $\rho_k$  and  $\rho_{k+1}$ , satisfying the pushforward relation  $T^{k,*} \# \rho_k = \rho_{k+1}$ . In  $\text{JKOnet}^*$ , these  
 1399 mappings are obtained explicitly by precomputing the OT maps between  $\rho_k$  and  $\rho_{k+1}$ , whereas in  
 1400 our method, they are learned implicitly via neural networks  $T_\varphi^k$ .

1402 In practice, the continuous distributions  $\rho_k$  are replaced by their empirical counterparts, constructed  
 1403 from training samples  $\{x_k^1, x_k^2, \dots, x_k^N\} = X_k \sim \rho_k$ . This is precisely where the distinction be-  
 1404 tween the paired and unpaired setups becomes critical:

1404 • **Paired.** When samples are paired, they follow the ground-truth transport mappings, i.e.,  
 1405

$$1406 T^{0,*}(x_0^i) = x_1^i, \quad T^{1,*}(x_1^i) = x_2^i, \quad \dots, \quad T^{k,*}(x_k^i) = x_{k+1}^i, \quad \dots$$

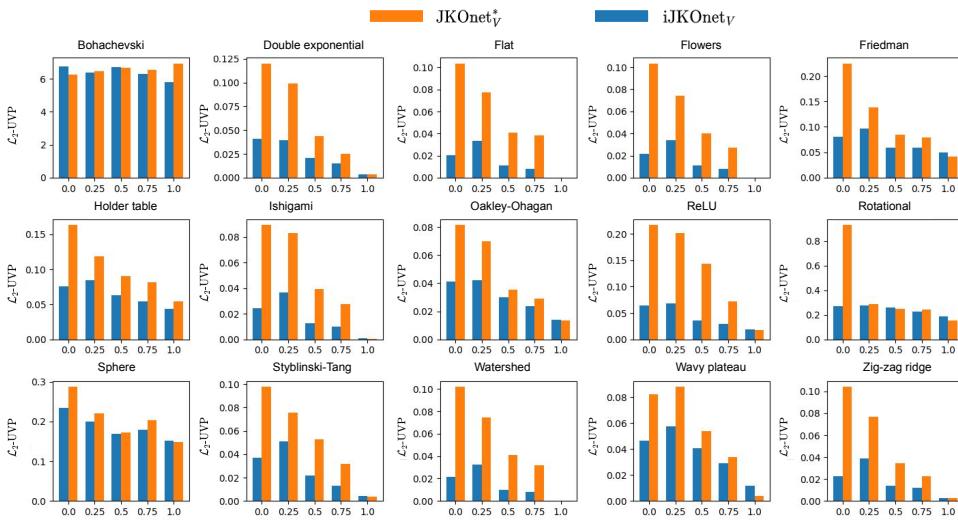
1408 In this case, estimating  $T^{k,*}$  becomes substantially simpler and can be viewed as a supervised  
 1409 regression problem over the given sample pairs.

1410 • **Unpaired.** In contrast, when the samples  $X_k, k \in \{0, \dots, K\}$ , are mutually independent, the  
 1411 quality of the estimated OT maps degenerate significantly. This degradation stems from the high  
 1412 sample complexity of Wasserstein-2 distances and maps (known to be relatively poor, see Hütter  
 1413 & Rigollet (2021)), as well as from the increased variance of the estimated transport mappings.  
 1414

1415 A deeper theoretical analysis of this gap between paired and unpaired setups remains an interesting  
 1416 direction for future work. To quantify how performance changes when transitioning from unpaired  
 1417 to paired data, we conducted the following *ablation study*.

1418 **Experimental Setup.** Following Appendix B.1.1, we used the same set of potentials and focused  
 1419 solely on the potential parametrization. In the 2D setting with  $K = 5$ , we trained for  $5K$  iterations  
 1420 using  $N = 2K$  training samples per time step and a 40% test split. The fraction of full trajectories  
 1421 in the training set was gradually increased from 0.0 (unpaired) to 1.0 (paired) in steps of 0.25.

1422 **Results.** Figure 3 shows the results. As the fraction of paired samples decreases, the performance  
 1423 of JKOnet\* degrades more rapidly than that of our method. For the Bohachevsky potential, both  
 1424 methods show reduced accuracy in both paired and unpaired scenarios.



1444 Figure 3: Ablation study comparing *paired* (1.0) and *unpaired* (0.0) setups. As the proportion  
 1445 of paired samples (i.e. full trajectories in the dataset) decreases, the performance of JKOnet\*  
 1446 degenerate more rapidly than that of our method.  
 1447

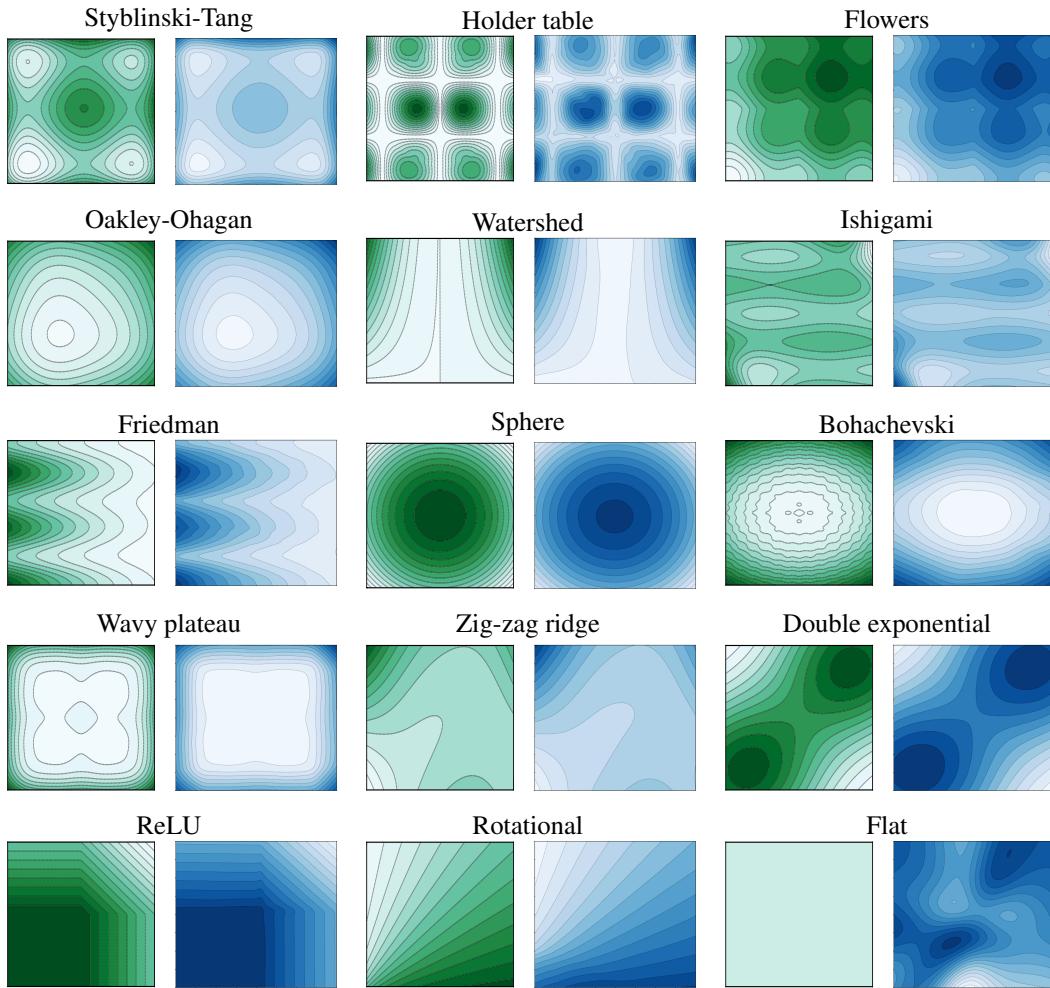
## 1448 B EXTENDED COMPARISONS

1452 In this section, we first present extended synthetic experiments in §B.1, followed by results on real-  
 1453 world single-cell data in §B.2.

### 1455 B.1 SYNTHETIC COMPARISONS

1456 In this section, we further investigate our method’s ability to learn potential energy (§B.1.1), inter-  
 1457 action energy (§B.1.2), and internal energy (§B.1.3).

1458 Since our approach builds on the JKO framework for modeling population dynamics, we adopt the  
 1459 experimental setup from JKOnet\* and compare our method against both JKOnet Bunne et al.  
 1460 (2022b) and JKOnet\* Terpin et al. (2024) on corresponding synthetic evaluation tasks.  
 1461



1494 Figure 4: Level curves of the true (green) and estimated (blue) potentials for the **paired** setup,  
 1495 following (Terpin et al., 2024, Appendix F). These results can be directly compared with those in  
 1496 (Terpin et al., 2024, Figure 6). Note that for the *flat* potential, the value range is near zero, as  
 1497 expected for the ground-truth potential.

### 1500 B.1.1 LEARNING POTENTIAL ENERGY

1502 **Paired Setup.** Following Terpin et al. (2024, §4.1), we evaluate our method in the paired setup (see  
 1503 details in §5.1), which enables a direct visual comparison with JKOnet\* (Terpin et al., 2024). The  
 1504 experiments are conducted on the synthetic dataset from Terpin et al. (2024, Appendix B), using a  
 1505 step size of  $\tau = 0.01$ ,  $K = 5$  time steps, and  $N = 2000$  samples per step with 40% of test samples.  
 1506 Figure 4 presents the ground-truth potentials  $V(x)$  (green), defined in Terpin et al. (2024, Appendix  
 1507 F, Eqs. (31)–(45)), alongside the reconstructed potentials  $V_\theta(x)$  (blue) learned by our method. As  
 1508 shown, our approach achieves performance comparable to JKOnet\* in the paired setup.

1509 **Unpaired Setup.** We repeat the experiments on a corrected version of the synthetic dataset, where  
 1510 samples are not temporally correlated. We select the six most challenging potentials in terms of  
 1511 convergence. Table 3 reports the ratio of the final metrics between the unpaired and paired setups.  
 As shown, switching to the unpaired setup significantly affects the method’s performance.

1512 Table 3: Comparison of paired and unpaired setups. Each value indicates the ratio of the final metric  
 1513 in the unpaired setup to that in the paired setup. Most ratios significantly exceed 1, highlighting the  
 1514 increased difficulty of the unpaired setting.

1516 1517 1518 1519 1520 1521 1522 1523 1524	1516 1517 1518 1519 1520 1521 1522 1523 1524	1516 1517 1518 1519 1520 1521 1522 1523 1524	iJKOnet <sub>V</sub>			JKOnet <sub>V</sub> <sup>*</sup>		
			Bd <sub>W<sub>2</sub></sub> <sup>2</sup> -UVP	EMD	L <sub>2</sub> -UVP	Bd <sub>W<sub>2</sub></sub> <sup>2</sup> -UVP	EMD	L <sub>2</sub> -UVP
FL	Flowers		8066	90	297	135325	305	5809
FN	Friedman		556	5	2	263	6	3
IG	Ishigami		2259	38	36	4276	64	210
WS	Watershed		18135	220	1509	512720	1014	102592
WP	Wavy plateau		2	2	7	1	2	29
ZR	Zigzag ridge		158	23	13	155	25	31

### B.1.2 LEARNING INTERACTION ENERGY

In this section, we evaluate how accurately our method and JKOnet<sup>\*</sup> (Terpin et al., 2024) can approximate the interaction energy in the unpaired setup. We exclude JKOnet<sup>†</sup> from comparison, as its design does not support learning interaction energies.

We identified an inconsistency in the theoretical formulation presented in the JKOnet<sup>\*</sup> paper (Terpin et al., 2024). Specifically, as discussed in §2 (see also (Santambrogio, 2015, §7.2)), the ground-truth interaction functional must be symmetric; that is, in (5), the interaction kernel should satisfy  $W(z) = W(-z)$ . Incorporating this correction, we conducted an experiment to assess the ability of both methods to recover the *interaction* energy in a 2D unpaired setting.

**Experimental Setup.** The ground-truth interaction kernels are defined as  $W(z) = \frac{1}{2}(W_b(z) + W_b(-z))$ , where  $W_b$  denotes the base functionals listed in Table 3. Both methods were restricted to use only the interaction energy component  $\mathcal{W}_{\theta_2}$  in the energy parameterization (12). The corresponding variants are denoted as iJKOnet<sub>W</sub> and JKOnet<sub>W</sub><sup>\*</sup>. Similar to §B.1.1, we use a step size of  $\tau = 0.01$ ,  $K = 5$  time steps, and two sample size settings: a small-scale setup with 12K total samples ( $N = 2K$  per step, with 40% reserved for testing) and a large-scale setup with 60K total samples ( $N = 10K$  per step, with 40% reserved for testing).

**Results.** Qualitative results are shown in Figure 5 for the paired setup and in Figure 6 for the unpaired setup, with quantitative  $\mathcal{L}_2$ -UVP metrics reported in Table 4 (see §C.1 for details). In the paired setup, most potentials are accurately restored, although Waby Plateau, Friedmann, and Flowers are not. In the unpaired setup, neither method successfully recovers the interaction energy, likely due to biases introduced by the batched estimation of the interaction term  $\mathcal{W}_{\theta_2}$  in (12), which involves squaring the input. This approximation seems to cause both models to converge to batch-specific minima.

1552 Table 4: Quantitative  $\mathcal{L}_2$ -UVP results for the **unpaired** setup in 2D interaction energy learning.  
 1553 Increasing the number of samples improves performance, though the effect is relatively limited.

1555 1556 1557 1558 1559 1560 1561 1562 1563 1564 1565	1555 1556 1557 1558 1559 1560 1561 1562 1563 1564 1565	Sample size: 2K		Sample size: 10K	
		iJKOnet <sub>W</sub> <b>(Ours)</b>	JKOnet <sub>W</sub> <sup>*</sup>	iJKOnet <sub>W</sub> <b>(Ours)</b>	JKOnet <sub>W</sub> <sup>*</sup>
Flowers		0.0032	0.0101	0.0009	0.0069
Friedman		0.0087	0.0010	0.0024	0.0093
Ishigami		0.0046	0.0087	0.0008	0.0066
Watershed		0.0043	0.0086	0.0009	0.0081
Zigzag Ridge		0.0060	0.0104	0.0017	0.0034

### B.1.3 LEARNING INTERNAL ENERGY

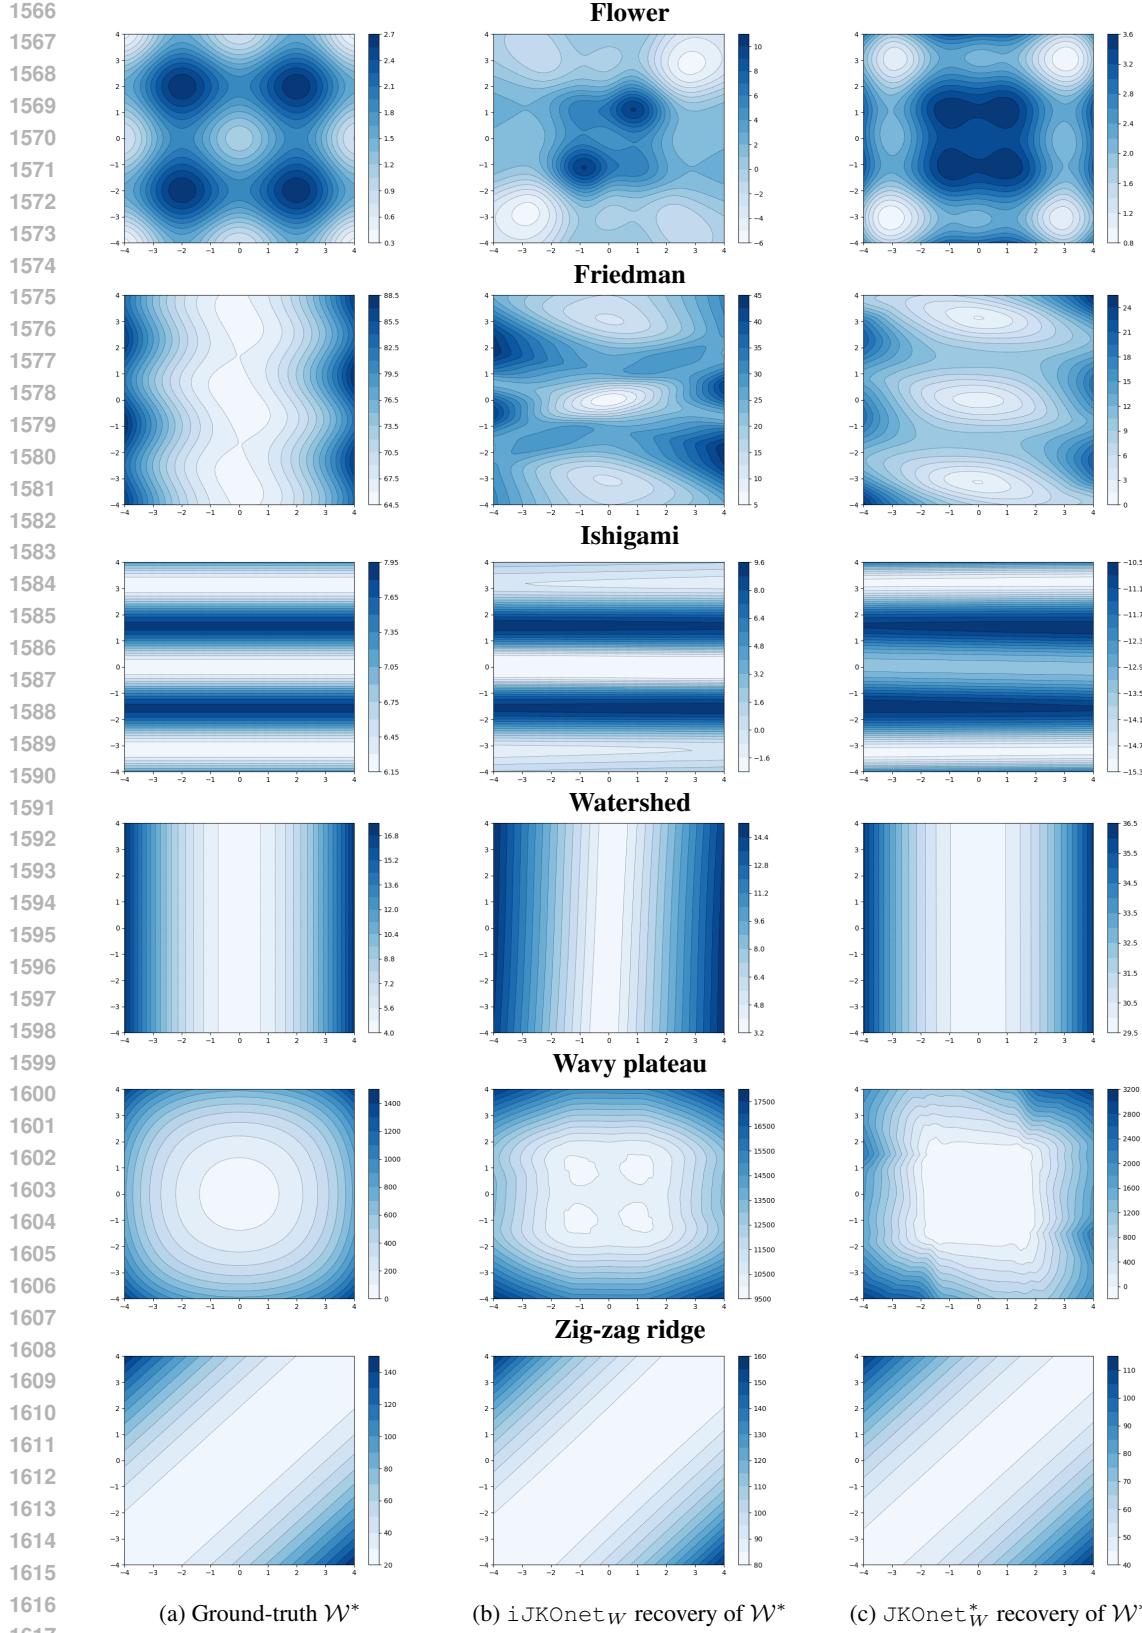


Figure 5: Qualitative results from §B.1.2 in the **paired** setup, using 10K samples per time step (with 40% reserved for testing). Compared to Figure 6, which shows the **unpaired** setup, the model performs noticeably better: nearly all potentials are accurately reconstructed, highlighting the relative simplicity of the paired scenario.

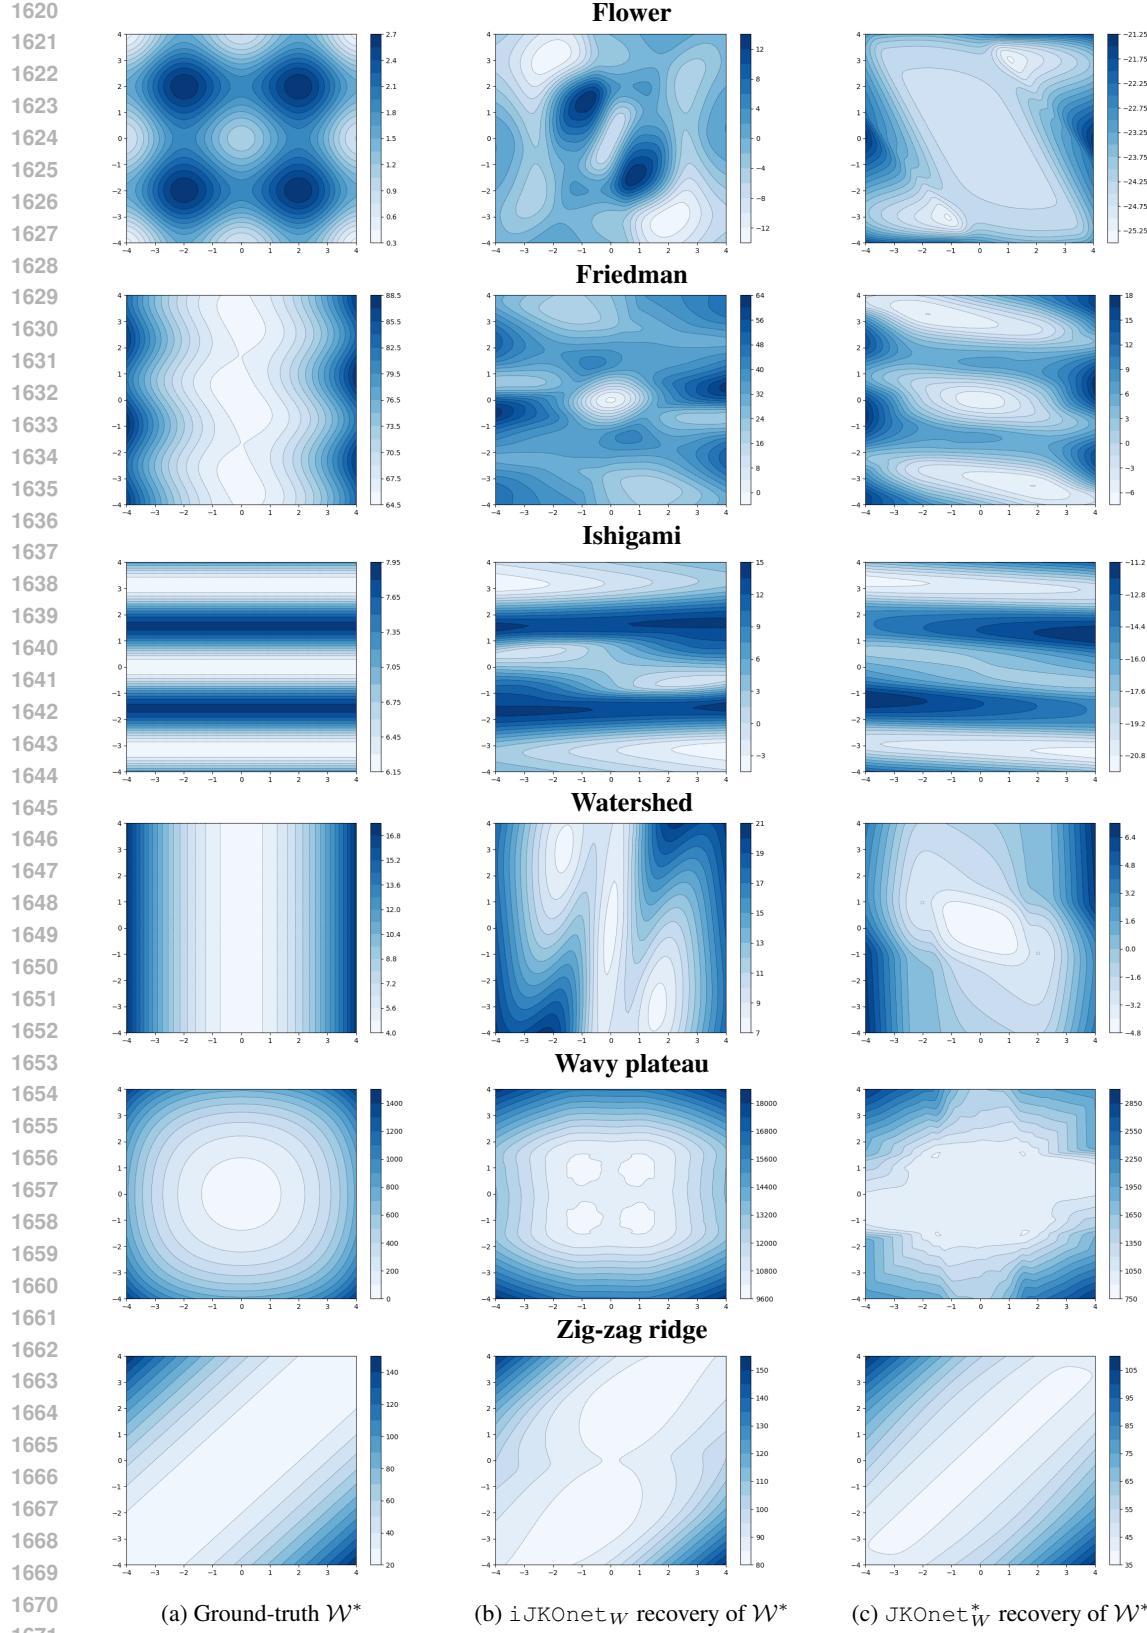


Figure 6: Qualitative results from §B.1.2 in the **unpaired** setup with 10K samples in total (40% reserved for testing). Neither model is able to accurately reconstruct the ground-truth interaction energy  $\mathcal{W}^*$ , likely due to biases introduced in the estimation of the integral in (12).

In this section, we assess how accurately our method and JKOnet\* (Terpin et al., 2024) can estimate the diffusion coefficient  $\theta_3$  in (12) corresponding to the internal energy term under the **unpaired** setup. We exclude JKOnet from comparison, as its design does not support learning internal energy.

**Experimental Setup.** Following the protocol of Terpin et al. (2024, §4.3), we estimate the diffusion coefficient  $\theta_3$  in 2D and 20D settings using both iJKOnet and JKOnet\*, i.e., the full parameterization in (12). The same functional from Appendix C.7 is used for both potential and interaction energies. The ground-truth diffusion levels are set to  $\beta^* \in \{0.0, 0.1, 0.2\}$ .

**Results.** Figure 7 shows that iJKOnet fails to recover the ground-truth  $\beta^*$  values, while JKOnet\* provides a closer approximation. However, all predicted  $\theta_3$  values tend to converge toward 0, deviating from the true levels  $\beta^* \in \{0.0, 0.1, 0.2\}$ . This indicates that accurately learning the internal energy remains challenging for both models. We assume that such behavior stems from difficulties of entropy estimation developed methods.

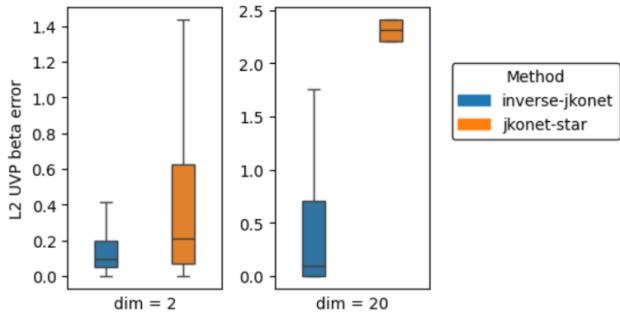


Figure 7: Estimation error of the diffusion coefficient  $\theta_3$  relative to the ground-truth values  $\beta^*$ . Blue bars correspond to iJKOnet and orange bars to JKOnet\*. The Y-axis indicates the absolute deviation between the estimated  $\theta_3$  and true  $\beta^*$  values.

## B.2 LEARNING SINGLE-CELL DYNAMICS

In this section, we present extended comparisons on real-world single-cell datasets, starting with 5D results and then considering 50D and 100D.

**Datasets.** Following §5.2, we consider the Embryoid Body (**EB**) dataset (Moon et al., 2019) and additionally the Multiome (**Multi**) dataset (Burkhardt et al., 2022). For both datasets, we apply the preprocessing pipeline of (Tong et al., 2020). The Multi dataset contains single-cell measurements across four time points (days 2, 3, 4, and 7), which we use for high-dimensional experiments.

**5D.** Following (Neklyudov et al., 2024a, Table 2) and (Terpin et al., 2024, §4.4), we conducted experiments on the **EB** dataset in 5D. At each time step, we train for 1000 epochs on 60% of the data and compute the  $d_{W_1}$  distance between the observed distribution  $\rho_k$  (remaining 40% of the data) and the one-step-ahead prediction  $\hat{\rho}_k$ . The results are shown in Table 5. Our method outperforms several previous approaches when using a time-varying potential parametrization.

**50D and 100D data.** Following Neklyudov et al. (2024a), we conducted experiments on the **Multi** dataset in 50D and 100D using a leave-one-out setup, averaged over three runs. Models were trained on marginals from timepoint partitions [2,4,7] and [2,3,7], and evaluated on the corresponding left-out marginals at timepoints [3] and [4]. Training was performed for 5000 epochs.

The results are presented in Table 6. Our method achieves comparable performance for the left-out [3] marginal. However, performance on the [4] marginal is poor, leading to a lower overall average. This is likely due to the longer time intervals between the learned marginals, which violate our assumption that consecutive marginals are obtained from JKO steps (see §3.1), thereby reducing the accuracy of predicting the [4] marginal.

Table 5: The results for the **EB** dataset in 5D are obtained by training on all time steps, evaluated for next-time-step prediction using the  $d_{\mathbb{W}_1}$  metric. Results for non-JKO methods are taken from (Neklyudov et al., 2024a, Table 2):

Model	$t_1$	$t_2$	$t_3$	$t_4$	Mean
Neural SDE (Li et al., 2020)	0.69	0.91	0.85	0.81	0.82
TrajectoryNet (Tong et al., 2020)	0.73	1.06	0.90	1.01	0.93
SB-FBSDE (Chen et al., 2022a)	0.56	0.80	1.00	1.00	0.84
NLSB (Koshizuka & Sato, 2023)	0.68	0.84	0.81	0.79	0.78
OT-CFM (Tong et al., 2024a)	0.78	0.76	0.77	0.75	0.77
WLF-OT (Neklyudov et al., 2024a)	0.65	0.78	0.76	0.75	0.74
WLF-SB (Neklyudov et al., 2024a)	0.63	0.79	0.77	0.74	0.73
JKOnet (Bunne et al., 2022b)	1.53	1.27	1.13	1.41	1.34
JKOnet* <sub>V</sub> Terpin et al. (2024)	0.99	1.11	1.06	1.30	1.12
<b>iJKOnet<sub>V</sub> (Ours)</b>	0.92	1.11	0.95	1.21	1.05
JKOnet* <sub>t,V</sub> (Terpin et al., 2024)	0.69	0.77	0.69	0.78	0.73
<b>iJKOnet<sub>t,V</sub> (Ours)</b>	0.51	0.58	0.57	0.64	<b>0.58</b>

Table 6: The results for **Multi** dataset for direct comparison with from (Neklyudov et al., 2024a, Table 1), averaged for 3 runs for leave-one-out setup,  $d_{\mathbb{W}_2}$  metric.

Dimension	50			100		
	Metric	BW-UVP	EMD	MMD	BW-UVP	EMD
JKOnet* <sub>V</sub>	$128.729 \pm 78.307$	$68.406 \pm 6.055$	0.0003	$92.512 \pm 38.075$	$72.639 \pm 3.161$	0.0002
<b>iJKOnet<sub>V</sub> (Ours)</b>	$38.318 \pm 18.325$	$50.560 \pm 7.187$	0.0003	$38.991 \pm 16.918$	$59.216 \pm 6.869$	0.0002
JKOnet* <sub>t,V</sub>	$89.264 \pm 38.372$	$68.600 \pm 9.846$	0.0003	$93.429 \pm 34.332$	$78.674 \pm 8.236$	0.0002
JKOnet <sub>t,V</sub> (Ours)	$36.156 \pm 16.515$	$50.026 \pm 7.727$	0.0003	$38.387 \pm 18.525$	$59.318 \pm 7.365$	0.0002

## C TRAINING AND MODEL DETAILS

### C.1 METRIC COMPUTATION DETAILS

In this section, we provide details on the evaluation metrics used in our experiments.

**$\mathcal{L}_2$ -UVP.** When the ground-truth functional  $F^*$  (e.g.,  $V$  or  $W$ ) is available, we assess the discrepancy between it and its reconstruction  $\hat{F}$  using the *backward  $\mathcal{L}_2$ -based Unexplained Variance Percentage* ( $\mathcal{L}_2$ -UVP) metric introduced by Korotin et al. (2021a), defined as follows:

$$\mathcal{L}_2\text{-UVP}(F^*, \hat{F}) = 100 \cdot \frac{\tau^2 \|\nabla \hat{F} - \nabla F^*\|_{\rho_{k+1}}^2}{\text{Var}(\rho_k)} \%, \quad (19)$$

where the norm  $\|\cdot\|_{\rho_{k+1}}$  is computed with respect to the ground-truth distribution  $\rho_{k+1}$ , and  $\tau$  denotes the time step size. Values close to 0% indicate that  $\nabla \hat{F}$  closely approximates  $\nabla F^*$ .

**$Bd_{\mathbb{W}_2}^2$ -UVP.** The *Bures-Wasserstein UVP* introduced by (Korotin et al., 2021c) is defined as

$$Bd_{\mathbb{W}_2}^2\text{-UVP}(\rho_k, \hat{\rho}_k) \stackrel{\text{def}}{=} 100 \cdot \frac{Bd_{\mathbb{W}_2}^2(\rho_k, \hat{\rho}_k)}{\frac{1}{2}\text{Var}(\rho_k)} \%, \quad (20)$$

where the *Bures-Wasserstein distance* is given by

$$Bd_{\mathbb{W}_2}^2(\mathbb{P}, \mathbb{Q}) \stackrel{\text{def}}{=} d_{\mathbb{W}_2}^2(\mathcal{N}(\mu_{\mathbb{P}}, \Sigma_{\mathbb{P}}), \mathcal{N}(\mu_{\mathbb{Q}}, \Sigma_{\mathbb{Q}})), \quad (21)$$

with  $\mu_{\mathbb{P}}$  and  $\Sigma_{\mathbb{P}}$  denoting the mean and covariance of distribution  $\mathbb{P}$ , respectively.

**EMD or  $d_{\mathbb{W}_1}$ .** The *Earth Mover's Distance* (EMD) (Rubner et al., 1998) is defined as

$$\text{EMD}(\rho_k, \hat{\rho}_k) \equiv d_{\mathbb{W}_1}(\rho_k, \hat{\rho}_k) \stackrel{\text{def}}{=} \min_{\pi \in \Pi(\rho_k, \hat{\rho}_k)} \int_{\mathcal{X} \times \mathcal{X}} \|x - y\| d\pi(x, y), \quad (22)$$

1782 where  $\rho_k$  and  $\hat{\rho}_k$  denote the ground truth and predicted distributions at time step  $k$ .  
 1783

1784 **MMD.** *Maximum Mean Discrepancy* (MMD) (Gretton et al., 2012) between two distributions  $\rho_k$   
 1785 and  $\hat{\rho}_k$  is defined as

$$1786 \quad \text{MMD}^2(\rho_k, \hat{\rho}_k) = \|\mathbb{E}_{x \sim \rho_k}[h(x)] - \mathbb{E}_{\hat{x} \sim \hat{\rho}_k}[h(\hat{x})]\|_{\mathcal{H}}^2, \quad (23)$$

1788 where  $h$  is the feature map associated with the reproducing kernel  $\mathbf{H}$ , and  $\|\cdot\|_{\mathcal{H}}$  denotes the norm in  
 1789 the corresponding RKHS (Ghojogh et al., 2021). Here,  $\rho_k$  and  $\hat{\rho}_k$  are defined as in the EMD metric.  
 1790

## 1791 C.2 TRAINING DETAILS

1793 **Energy.** To ensure stability, we accumulate gradients across all time steps  $k = 0, \dots, K$ , using  
 1794 mini-batch estimates  $\hat{\mathcal{J}}_{\theta}$  of the energy function  $\mathcal{J}_{\theta}$  (see Eq. (12)). For each time step  $t_k$ , we sample  
 1795 a mini-batch  $X_k = \{x_k^1, \dots, x_k^B\} \sim \rho_k$ , with a fixed batch size  $B$ . The energy is then estimated as:  
 1796

$$1797 \quad \mathcal{J}_{\theta}(\rho_k) \approx \hat{\mathcal{J}}_{\theta}(X_k) = \frac{1}{B} \sum_{i=0}^B \left[ V_{\theta_1}(x_k) + \frac{1}{B} \sum_{j=0}^B W_{\theta_2}(x_i - x_j) \right] - \theta_3 \hat{\mathcal{H}}(X_k), \quad (24)$$

1800 where  $\hat{\mathcal{H}}(X_k)$  denotes the estimated entropy of  $\rho_k$  discussed in the following section.  
 1801

1802 **Map.** In practice, we use two strategies to parameterize each transport map  $T_{\varphi}^k$ : for large-scale  
 1803 tasks, we assign a separate network to each time step; for moderate-dimensional tasks, we encode  
 1804 the time index  $k$  as an additional input, i.e.,  $T_{\varphi}^k(x) = T_{\varphi}(x, k)$ , which helps prevent overfitting. This  
 1805 approach performs better than using a network without any time embedding. Thanks to Optax  
 1806 (DeepMind et al., 2020), the inference of all  $T_{\varphi}^k$  can be performed in parallel, treating them as a  
 1807 single ‘generator’ step. We also experimented with initializing the maps during the early training  
 1808 epochs by aligning them with discrete OT maps computed using the OTT-JAX (Cuturi et al., 2022).  
 1809

1810 **Training.** We aggregate gradients for  $\mathcal{J}_{\theta}$  and  $T_{\varphi}^k$  across all time steps. We also experimented with  
 1811 alternative aggregation strategies, but found that equal aggregation across all time steps was the  
 1812 most effective. During optimization, we perform multiple update steps for the transport maps  $T_{\varphi}^k$ ,  
 1813 parameterized by  $\varphi$ , while updating the energy function  $\mathcal{J}_{\theta}$ , parameterized by  $\theta$ , only once. This  
 1814 follows standard practice in min–max optimization (Goodfellow et al., 2020). The overall procedure  
 1815 is summarized in the pseudocode in Algorithm 1.  
 1816

1817 **Stability.** The main challenge was avoiding suboptimal energy functionals that did not converge  
 1818 to zero. We found that the simplest setup (combined with careful hyperparameter tuning using the  
 1819 Optuna framework (Akiba et al., 2019)) was the most effective. We additionally experimented  
 1820 with common stabilization techniques from GAN training, including gradient penalties (Gulrajani  
 1821 et al., 2017), spectral normalization (Miyato et al., 2018; Jiang et al., 2018), and extragradient up-  
 1822 dates (Daskalakis et al., 2018). Overall, we perform optimization for single energy functional, this  
 1823 naturally acts as a regularizer, improving the stability of the training.  
 1824

1825 **Scalability.** Our method scales well with the number of time points. Leveraging JAX (Bradbury  
 1826 et al., 2018) and Optax (DeepMind et al., 2020), we avoid memory issues even for large batches and  
 1827 many time steps, since we do not backpropagate through time as in JKOnet (Bunne et al., 2022b).  
 1828 Each step is processed independently, and the loss is computed by averaging outputs from separate  
 1829 networks, avoiding large computation graphs. Training time on GPU is comparable to JKOnet\*, but  
 1830 our approach eliminates the costly precomputation of OT couplings. In high-dimensional settings  
 1831 (50D–100D), JKOnet\* exceeds GPU memory limits and requires CPU-based precomputations,  
 1832 making it substantially slower.  
 1833

## 1834 C.3 ENTROPY ESTIMATION DETAILS

1835 Prior to training, we estimated  $\hat{\mathcal{H}}(\rho_k)$  for each  $k = 0, \dots, K$  using the Kozachenko–Leonenko  
 1836 nearest-neighbor estimator (Kozachenko, 1987; Berrett et al., 2019) with 5 nearest neighbors. We  
 1837 used the publicly available implementation from (Butakov et al., 2024)<sup>2</sup>. As discussed in §3.3, we  
 1838

<sup>2</sup><https://github.com/VanessB/mutinfo>

1836

**Algorithm 1:** iJKOnet Training

1837

---

**Input:**

Sequence of measures  $\{\rho_k\}_{k=0}^K$  (accessible via samples);  
 Mapping networks  $T_\varphi^k : \mathcal{X} \rightarrow \mathcal{X}$ ;  
 Potential network  $V_{\theta_1} : \mathcal{X} \rightarrow \mathbb{R}$ , interaction network  $W_{\theta_2} : \mathcal{X} \rightarrow \mathbb{R}$ ;  
 Diffusion coefficient  $\theta_3 \in \mathbb{R}^+$ , time step  $\tau$ , max inner iterations  $I_T$ ;

**Output:**

Learned parameters  $\theta = (\theta_1, \theta_2, \theta_3)$  and  $\varphi$ ;

**while** not converged **do**

**for**  $i \leftarrow 1$  to  $I_T$  **do**

// Update transport maps  $T_\varphi^k$

**for**  $k \leftarrow 0$  to  $K - 1$  **do**

Sample batch  $\bar{X}_k \sim \rho_k$ ,  $X_{k+1} \sim \rho_{k+1}$ ;  
 $X_{k+1}^{pred} \leftarrow T_\varphi^k(X_k)$ ;  
 $\mathcal{L}_\varphi^k \leftarrow \mathcal{J}_\theta(X_{k+1}^{pred}) + \frac{1}{2\tau} \|X_{k+1}^{pred} - X_{k+1}\|_2^2$ ;

$\varphi \leftarrow \varphi - \nabla_\varphi \sum_k \mathcal{L}_\varphi^k$ ;

// Update energy parameters  $\theta$

**for**  $k \leftarrow 0$  to  $K - 1$  **do**

Sample batch  $\bar{X}_k \sim \rho_k$ ;  
 $X_{k+1}^{pred} \leftarrow T_\varphi^k(X_k)$ ;  
 $\mathcal{L}_\theta^k \leftarrow -\mathcal{J}_\theta(X_{k+1}^{pred}) + \mathcal{J}_\theta(\textcolor{orange}{X}_{k+1})$ ;

$\theta \leftarrow \theta - \nabla_\theta \sum_k \mathcal{L}_\theta^k$ ;

---

1860

1861

1862

1863

use the following equation for estimating  $\widehat{\mathcal{H}}(T_\varphi^k \# \rho_k)$ :

1864

1865

$$\widehat{\mathcal{H}}(T_\varphi^k \# \rho_k) = \widehat{\mathcal{H}}(\rho_k) - \int_{\mathcal{X}} \log |\det \nabla_x T_\varphi^k(x)| d\rho_k(x). \quad (25)$$

1866

1867

1868

1869

Since our map is not required to be invertible during training, as is the case for gradient-based ICNN (Amos et al., 2017) parameterizations (Bunne et al., 2022b), we compute the sign and the natural logarithm of the absolute value of the Jacobian determinant, both of which are supported by modern computational frameworks (Bradbury et al., 2018; Paszke, 2019).

1870

1871

1872

1873

1874

1875

**C.4 OPTIMIZER**

1876

1877

1878

1879

1880

1881

We optimize two loss functions with respect to  $\mathcal{J}_\theta$  and  $T_\varphi^k$  using the Adam optimizer (Kingma, 2014). For  $\mathcal{J}_\theta$ , we use parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\varepsilon = 1 \times 10^{-8}$ , and a constant learning rate of  $5 \times 10^{-4}$ , with gradient clipping applied using a maximum global norm of 10. For  $T_\varphi^k$ , we use parameters  $\beta_1 = 0.5$ ,  $\beta_2 = 0.9$ , and  $\varepsilon = 1 \times 10^{-8}$ , with a learning rate of  $1 \times 10^{-3}$ . Training is performed with mini-batches of size 500.

1882

1883

1884

**C.5 NETWORK ARCHITECTURE**

1885

1886

1887

1888

1889

The neural networks for the potential  $V_{\theta_1}$  and interaction  $W_{\theta_2}$  energies are multi-layer perceptrons (MLPs) with two hidden layers of size 64, using `softplus` activation functions, and a one-dimensional output layer. The neural network for the optimal transport maps  $T_\varphi^k$  is also an MLP with two hidden layers of size 64. The time step  $k$  is concatenated with the input  $x_k^i$ , and the network uses `selu` activations, with an output layer that matches the input dimension of  $x_k^i$ .

35

1890  
1891

## C.6 HARDWARE

1892  
1893  
1894

Experiments were conducted on a CentOS Linux 7 (Core) system with NVIDIA A100 GPUs. Most of the computation time was spent on metric evaluation. Leveraging JAX’s just-in-time compilation, our method completed 100 training epochs in approximately one minute.

1895

1896

Table 7: Training time (in hours) for **EB** (Moon et al., 2019) dataset for 5 and 100 dimensions.

1897

1898

1899

1900

1901

1902

Solver	Dimension	mean	std
iJKOnet <sub>V</sub>	100	0.333	0.020
iJKOnet <sub>t,V</sub>	100	0.338	0.023
iJKOnet <sub>V</sub>	5	0.321	0.005
iJKOnet <sub>t,V</sub>	5	0.332	0.001

1903

1904

1905

## C.7 FUNCTIONALS

1906

1907

1908

1909

1910

1911

For easier representation of our experimental results, we use abbreviations for each potential, each abbreviation being shown in parentheses and the definitions can be found in (Terpin et al., 2024, Appendix F): Wavy Plateau (**WP**), Double Exponential (**DE**), Rotational (**RO**), ReLU (**RE**), Flat (**FT**), Friedman (**FN**), Watershed (**WS**), Ishigami (**IG**), Flowers (**FL**), Bohachevsky (**BC**), Sphere (**SP**), Styblinski-Tang (**ST**), Oakley–Ohagan (**OO**), Zigzag Ridge (**ZR**), and Holder Table (**HT**).

1912

1913

## D PROOFS

1914

1915

1916

To begin with, we recall the notion of strong smoothness, see (Beck, 2017, §5.1). A functional  $F : \mathcal{X} \rightarrow \mathbb{R}$  is called  $\frac{1}{\beta}$  - *strongly smooth* if it is continuously differentiable on  $\mathcal{X}$  and it holds:

1917

$$\beta \|\nabla F(x) - \nabla F(y)\|_2 \leq \|x - y\|_2 \quad \forall x, y \in \mathcal{X}.$$

1918

1919

Now we proceed to some auxiliary results.

1920

1921

1922

**Lemma D.1** (Solution to the JKO problem with potential energy is unique). *Let  $\mathcal{J}^*(\rho) = \int_{\mathcal{X}} V^*(x) d\rho(x)$ ,  $\rho_0 \in \mathcal{P}_{2,ac}(\mathcal{X})$  and  $\rho_1 = \text{JKO}_{\tau, \mathcal{J}^*}(\rho_0)$ . Then  $T_{V^*} = \arg \min_{T: \mathcal{X} \rightarrow \mathcal{X}} \mathcal{L}(V^*, T)$  is the (unique) optimal transport map between  $\rho_0$  and  $\rho_1$ . In particular,  $T_{V^*} \# \rho_0 = \rho_1$ .*

1923

1924

1925

*Proof.* Consider the functional  $\rho \mapsto d_{\mathbb{W}_2}^2(\rho_0, \rho)$ , where  $\rho \in \mathcal{P}_2(\mathcal{X})$ . In what follows, we prove that this functional is strictly convex.

1926

1927

Consider  $\rho_a, \rho_b \in \mathcal{P}_2(\mathcal{X})$ ;  $\alpha > 0$ . Let  $\rho = \alpha \rho_a + (1 - \alpha) \rho_b$ . Note that  $\rho_0$  is absolutely continuous. By the Brenier’s theorem, it holds:

1928

1929

1930

$$d_{\mathbb{W}_2}^2(\rho_0, \rho_a) = \int \|\rho_0(x) - T_a(x)\|_2^2 d\rho_0(x); \quad d_{\mathbb{W}_2}^2(\rho_0, \rho_b) = \int \|\rho_0(x) - T_b(x)\|_2^2 d\rho_0(x)$$

1931

1932

1933

for the corresponding (unique) OT maps  $T_a : T_a \# \rho_0 = \rho_a$ ;  $T_b : T_b \# \rho_0 = \rho_b$ .

Now consider the weighted sum:

1934

1935

1936

1937

1938

1939

$$\begin{aligned} \alpha d_{\mathbb{W}_2}^2(\rho_0, \rho_a) + (1 - \alpha) d_{\mathbb{W}_2}^2(\rho_0, \rho_b) &= \int_{\mathcal{X}} (\alpha \|x - T_a(x)\|_2^2 + (1 - \alpha) \|x - T_b(x)\|_2^2) d\rho_0(x) \\ &= \int_{\mathcal{X}} \|x - y\|_2^2 d\pi_{ab}(x, y), \end{aligned}$$

where plan  $\pi_{ab} \in \Pi(\rho_0, \rho)$  has the following conditionals:

1940

1941

1942

$$\pi(y|x) = \alpha \delta(y = T_a(x)) + (1 - \alpha) \delta(y = T_b(x)).$$

Since the OT map between  $\rho_0$  and  $\rho$  is unique (and deterministic), we have

1943

$$d_{\mathbb{W}_2}^2(\rho_0, \rho) = \min_{T: \# \rho_0 = \rho} \int_{\mathcal{X}} \|x - T(x)\|_2^2 d\rho(x) < \int_{\mathcal{X}} \|x - y\|_2^2 d\pi_{ab}(x, y),$$

1944 which yields the strict convexity of  $\rho \mapsto d_{\mathbb{W}_2}^2(\rho_0, \rho)$ .  
 1945

1946 The main statement of the Lemma follows from the strict convexity of  
 1947  $\rho \mapsto \int_{\mathcal{X}} V^*(x) d\rho(x) + d_{\mathbb{W}_2}^2(\rho_0, \rho)$  and Brenier's theorem.  $\square$   
 1948

1949 Now we are ready to prove our main quality bound theorem.  
 1950

1951 *Proof.* (Theorem 3.1).

1952 *Preliminary Note:* The facts from convex analysis (properties of convex functions and their conju-  
 1953 gates) which we use below could be found in (Kornilov et al., 2024, Lemma 1).

1954 At first, we simplify the expression for the JKO loss:  
 1955

$$\begin{aligned} 1956 \quad \mathcal{L}(V, T_V) &= \int V(T_V(x)) d\rho_0(x) - \int V(y) d\rho_1(y) + \frac{1}{2\tau} \int \|x - T_V(x)\|_2^2 d\rho_0(x) \\ 1957 &= \frac{1}{\tau} \int V_q(T_V(x)) d\rho_0(x) - \frac{1}{\tau} \int V_q(y) d\rho_1(y) - \frac{1}{\tau} \int \langle x, T_V(x) \rangle d\rho_0(x) \\ 1958 &\quad + \underbrace{\int \frac{\|y\|_2^2}{2\tau} d\rho_1(y) + \int \frac{\|x\|_2^2}{2\tau} d\rho_0(x)}_{\stackrel{\text{def}}{=} C(\rho_0, \rho_1)}, \\ 1960 & \\ 1961 & \\ 1962 & \\ 1963 & \end{aligned} \quad (26)$$

1964 where  $V_q = \tau V + \frac{1}{2} \|\cdot\|_2^2$ .  
 1965

1966 From (26) we note that  $T_V = \arg \min_T \int (V_q(T_V(x)) - \langle x, T_V(x) \rangle) d\rho_0(x)$ . From the convex  
 1967 analysis, it follows that  $T_V(x) = \nabla \bar{V}_q(x)$  delivers the minimum;  $\bar{V}_q$  is the (Fenchel) conjugate of  
 1968  $V_q$ . Note that  $T_V$  is measurable since it is continuous almost surely (Rockafellar, 1970, Theorem  
 1969 25.5). Substituting  $\nabla \bar{V}_q$  into (26) on par with Fenchel-Young equality yields:  
 1970

$$1971 \quad \mathcal{L}(V, T_V) = -\frac{1}{\tau} \int \bar{V}_q(x) d\rho_0(x) - \frac{1}{\tau} \int V_q(y) d\rho_1(y) + C(\rho_0, \rho_1). \quad (27)$$

1973 Note that  $T_{V^*} \sharp \rho_0 = \rho_1$  (Lemma D.1). Therefore, eq. (26) for  $\mathcal{L}(V^*, T_{V^*})$  could be simplified:  
 1974

$$\begin{aligned} 1975 \quad \mathcal{L}(V^*, T_{V^*}) &= \\ 1976 \quad \frac{1}{\tau} \underbrace{\int V_q^*(T_{V^*}(x)) d\rho_0(x)}_{= \int V_q^*(y) d\rho_1(y)} - \frac{1}{\tau} \int V_q^*(y) d\rho_1(y) - \frac{1}{\tau} \int \langle x, T_{V^*}(x) \rangle d\rho_0(x) + C(\rho_0, \rho_1) &= \\ 1977 & \\ 1978 & \\ 1979 & \\ 1980 & \\ 1981 & \\ 1982 & \\ 1983 & \\ 1984 & \end{aligned} \quad (28)$$

Now we analyze the gap  $\varepsilon(V)$  between optimal and optimized JKO losses. Leveraging (27) and (28) yields:  
 1985

$$1986 \quad \tau \varepsilon(V) = \tau \mathcal{L}(V^*, T_{V^*}) - \tau \mathcal{L}(V, T_V) = \int \bar{V}_q(x) d\rho_0(x) + \int V_q(y) d\rho_1(y) - \int \langle x, T_{V^*}(x) \rangle d\rho_0(x). \\ 1987 \quad (29)$$

1988 Now we note that  $T_{V^*} = \nabla \bar{V}_q^* : \mathcal{X} \rightarrow \mathbb{R}^D$ . From the properties of convex functions it follows that  
 1989  $\nabla V_q^* : \mathcal{X} \rightarrow \mathbb{R}^D$  defines the *inverse* Optimal Transport mapping. In particular (we assume  $V_q^*$  to be  
 1990 convex)  $\nabla V_q^* \sharp \rho_1 = \rho_0$ . Then, changing the variables in (29) results in:  
 1991

$$\begin{aligned} 1992 \quad \tau \varepsilon(V) &= \int \bar{V}_q(\nabla V_q^*(y)) d\rho_1(y) + \int V_q(y) d\rho_1(y) - \int \langle \nabla V_q^*(y), y \rangle d\rho_1(y) \\ 1993 &= \int_{\mathcal{X}} [\bar{V}_q(\nabla V_q^*(y)) + V_q(y) - \langle \nabla V_q^*(y), y \rangle] d\rho_1(y). \quad (30) \\ 1994 & \\ 1995 & \\ 1996 & \\ 1997 & \end{aligned}$$

$$= \int_{\mathcal{X}} \mathcal{D}_{\bar{V}_q}(\nabla V_q^*(y), \nabla V_q(y)) d\rho_1(y), \quad (31)$$

1998 where  $\mathcal{D}_{\overline{V}_q}(\cdot, \cdot)$  is the *Bregman* divergence, see (Banerjee et al., 2005, Def. 1) for the definition  
 1999 and (Banerjee et al., 2005, Appendix A) for the “Dual Divergences” property used in the transition  
 2000 from (30) to (31).  
 2001

2002 Since  $V_q$  is  $\frac{1}{\beta}$ -smooth, then  $\overline{V}_q$  is  $\beta$ -strongly convex. Therefore, by the property of strongly convex  
 2003 functions (we pick  $y \in \text{supp}(\rho_1)$ ):  
 2004

$$\begin{aligned} 2005 \quad \frac{\beta}{2} \|\nabla V_q^*(y) - \nabla V_q(y)\|_2^2 &\leq \overline{V}_q(\nabla V_q^*(y)) - \overline{V}_q(\nabla V_q(y)) - \langle \nabla \overline{V}_q(\nabla V_q(y)), \nabla \overline{V}_q^*(y) - \nabla V_q(y) \rangle \\ 2006 \quad &= \mathcal{D}_{\overline{V}_q}(\nabla V_q^*(y), \nabla V_q(y)). \end{aligned} \quad (32)$$

2008 Combining (31) and (32) yields:  
 2009

$$2010 \quad \tau \varepsilon(V) \geq \frac{\beta}{2} \int_{\mathcal{X}} \|\nabla V_q^*(y) - \nabla V_q(y)\|_2^2 d\rho_1(y). \\ 2011$$

2012 To complete the proof, we are left to note that  $\nabla V_q^*(y) - \nabla V_q(y) = \tau(\nabla V^*(y) - \nabla V(y))$ .  $\square$   
 2013

2014  
 2015  
 2016  
 2017  
 2018  
 2019  
 2020  
 2021  
 2022  
 2023  
 2024  
 2025  
 2026  
 2027  
 2028  
 2029  
 2030  
 2031  
 2032  
 2033  
 2034  
 2035  
 2036  
 2037  
 2038  
 2039  
 2040  
 2041  
 2042  
 2043  
 2044  
 2045  
 2046  
 2047  
 2048  
 2049  
 2050  
 2051