

Accelerating optimization over the space of probability measures

Shi Chen

*Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139, USA*

SCHEN636@MIT.EDU

Qin Li

*Department of Mathematics
University of Wisconsin-Madison
Madison, WI 53706, USA*

QINLI@MATH.WISC.EDU

Oliver Tse

*Department of Mathematics and Computer Science and Eindhoven Hendrik Casimir Institute
Eindhoven University of Technology
5600 MB Eindhoven, The Netherlands*

O.T.C.TSE@TUE.NL

Stephen J. Wright

*Department of Computer Sciences
University of Wisconsin-Madison
Madison, WI 53706, USA*

SWRIGHT@CS.WISC.EDU

Editor: Lam Nguyen

Abstract

The acceleration of gradient-based optimization methods is a subject of significant practical and theoretical importance, particularly within machine learning applications. While much attention has been directed towards optimizing within Euclidean space, the need to optimize over spaces of probability measures in machine learning motivates the exploration of accelerated gradient methods in this context, too. To this end, we introduce a Hamiltonian-flow approach analogous to momentum-based approaches in Euclidean space. We demonstrate that, in the continuous-time setting, algorithms based on this approach can achieve convergence rates of arbitrarily high order. We complement our findings with numerical examples.

Keywords: Acceleration methods, Momentum-based methods, Hamiltonian flows, Wasserstein gradient flows, Heavy-ball method.

1. Introduction

The search for a probability measure that minimizes an objective functional plays a significant role across many machine learning problems, encompassing areas such as generative modeling (Kingma and Welling, 2014; Goodfellow et al., 2014; Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021; Bengio et al., 2000), Bayesian inference (Jordan et al., 1999; Wainwright et al., 2008; Hoffman et al., 2013; Blei et al., 2017; Rezende and Mohamed, 2015; Lambert et al., 2022; Geffner and Domke, 2023) and reinforcement learning (Ziebart et al., 2008; Toussaint, 2009; Peters et al., 2010; Levine, 2018). These problems are stated as

$$\rho_* = \operatorname{argmin}_{\rho \in \mathcal{P}(\Omega)} E[\rho], \quad \text{and} \quad E_* := E[\rho_*], \quad (1)$$

where $\mathcal{P}(\Omega)$ is the collection of all probability measures supported on $\Omega \subset \mathbb{R}^d$ and $E : \mathcal{P}(\Omega) \rightarrow \mathbb{R}$ maps probability measures to \mathbb{R} . Throughout the paper, we use $[\cdot]$ to denote the dependence of a

Property	Convex		m -Strongly Convex	
	$f(x)$	$E[\rho]$	$f(x)$	$E[\rho]$
Gradient Flow	$O(1/t)$		$O(e^{-2mt})$	
Heavy-ball Method	$o(1/t)$		$O(e^{-\sqrt{m}t})$	
Variational Acceleration	$O(e^{-\beta t})$		-	

Table 1: Convergence rates of three momentum-based methods. All three methods share the same convergence rate to optimize a finite-dimensional function $f(x)$ and a functional $E[\rho]$.

functional on a function/probability measure. With a slight abuse of notation, we do not distinguish probability measures from their Lebesgue densities.

In this article, we focus on the continuous-time setting for solving (1). Analogous to the *gradient flow* in Euclidean space, a *gradient flow* on the space of probability measures can be used to find the minimizer. We define a gradient flow in $\mathcal{P}(\Omega)$ by making ρ depend on a “time” variable $t \geq 0$ (notation: ρ_t) and writing

$$\partial_t \rho_t = -\nabla_D E[\rho_t], \quad (\text{GF})$$

where ∇_D is an appropriately chosen metric on the space of probability measures. A commonly employed metric in the literature is the Wasserstein metric (Otto, 2001). The gradient flow determined by the Wasserstein gradient exhibits the same convergence rate as that of gradient flow in Euclidean spaces; see Table 1.

Given the numerous strategies developed to accelerate first-order (gradient-based) optimization methods in Euclidean space, it is natural to seek counterparts of these methods to optimization over the space of probability measures. Specifically, we are curious as to whether momentum-based acceleration methods on Euclidean space can be adapted to optimization problems of the form (1). Our focus lies on the convergence properties in continuous time, prompting the following questions:

- Is there a “*heavy-ball*” method applicable to (1), and is it provably faster than gradient flow (GF), akin to its superior performance for convex objective functions in Euclidean space?
- Is there a “*Nesterov acceleration*” method applicable to (1), does it exhibit provable speedup over gradient flow?
- What is the optimal convergence rate achievable by a first-order algorithm for (1) when employing a momentum strategy?

In addressing these questions, we design algorithms for minimizing over the space probability measures, resembling heavy-ball methods, variational acceleration methods (with Nesterov acceleration being one instance), and, more generally, Hamiltonian flows. These new algorithms demonstrate provable acceleration over regular gradient flow, mirroring their counterparts in Euclidean space. By careful design of the Hamiltonian, we achieve convergence of $e^{-\beta t}$ for almost all choices of β_t , including $\beta_t = t$ and $\log t^p$ for any power $p > 0$; see Table 1.

These developments hinge on two observations:

- In Euclidean space, many momentum-based acceleration methods rely on a carefully crafted Hamiltonian term $h_t : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$. It is shown, for various methods, that as long as particles follow trajectories defined by

$$\dot{x} = \nabla_v h_t(x, v), \quad \dot{v} = -\nabla_x h_t(x, v), \quad (2)$$

these particles can descend to the global minimum faster than the classical gradient flow. To obtain the acceleration effect in the space of probability measures, we need to conjure up the ‘‘Hamiltonian’’ concept for these spaces.

- The second observation is that a probability measure can be approximated by its samples:

$$\rho \approx \frac{1}{N} \sum_{i=1}^N \delta_{x_i} \in \mathcal{P}(\mathbb{R}^d) \quad N \in \mathbb{N}.$$

The evolution of ρ in the space of probability measures can be fully translated to the motion of its representative samples $\{x_1, x_2, \dots, x_N\}$. In the current context, methods such as heavy-ball and Nesterov prescribe the motion of particles by assigning dynamics to (x_i, v_i) . By deploying the relation between the motion of the particle sample and the evolution of its corresponding probability measure, we can lift these methods to define an evolution for ρ .

With these two observations in hand, and noting that each particle can be fully described by $(x_i, v_i) \in \mathbb{R}^d \times \mathbb{R}^d$, we define the empirical measure

$$\mu = \frac{1}{N} \sum_{i=1}^N \delta_{(x_i, v_i)} \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d), \quad N \in \mathbb{N}, \quad (3)$$

over this extended space—henceforth called the *phase space*. Deploying the relationship between the motion of particles and the empirical measure μ , and utilizing (2), we arrive at the *Hamiltonian flow* equation:

$$\partial_t \mu_t + \nabla_x \cdot \left(\mu_t \nabla_v \frac{\delta H_t}{\delta \mu} [\mu_t] \right) - \nabla_v \cdot \left(\mu_t \nabla_x \frac{\delta H_t}{\delta \mu} [\mu_t] \right) = 0, \quad (4)$$

where H_t is the counterpart of h_t lifted to the space of probability measures defined by

$$H_t[\mu] = \int_{\mathbb{R}^{2d}} h_t(x, v) d\mu(x, v),$$

and with $\frac{\delta H_t}{\delta \mu} [\mu_t]$ being its variational (Fréchet) derivative evaluated at μ_t^1 .

The formulation (4) serves as the foundation of our algorithm design. In particular, it is devised under the assumption that H_t and h_t are connected via a linear form, yet it remains valid for a general form of H_t . More precisely, it holds when

$$H_t[\mu] = K_t[\mu] + P_t[\mu], \quad (5)$$

with K_t and P_t representing the kinetic and potential energy functionals, respectively. The core of our strategy lies in carefully crafting the form of these functionals and connecting them to $E[\rho]$ to achieve acceleration.

1. The Fréchet derivative is the counterpart of Euclidean derivatives in a function space. The Euclidean derivative of a function $\nabla f(x)$ measures the first-order differentiation of this function at a point x : $f(x + \Delta x) - f(x) \approx \langle \nabla f(x), \Delta x \rangle$, with the bracket notation denoting the inner product defined on Euclidean space. Since $x \in \mathbb{R}^d$, we have $\nabla f(x) \in \mathbb{R}^d$. Similarly, the Fréchet derivative quantifies the first-order differentiation of a functional over the change in a function:

$$H[\mu + \delta\mu] - H[\mu] \approx \left\langle \frac{\delta H}{\delta \mu} [\mu], \delta\mu \right\rangle := \int \frac{\delta H}{\delta \mu} [\mu](z) \delta\mu(z) dz,$$

with the bracket notation denoting the inner product or duality pairing on the function space. In the following, we use $\frac{\delta}{\delta \rho}$ and $\frac{\delta}{\delta \mu}$ to denote the Fréchet derivatives of functionals that take in probability measures over \mathbb{R}^d and $\mathbb{R}^d \times \mathbb{R}^d$, respectively.

As in the Euclidean setting, the design of Hamiltonian induces a variety of convergence behaviors. By tailoring H_t to be consistent with the Hamiltonian term h_t crafted for the heavy-ball and Nesterov acceleration in Euclidean space, we can replicate these acceleration techniques over spaces of probability measures and establish the equivalence of convergence rates in these settings. Formal statements of our main results are given below.

Heavy-ball flow. The heavy-ball method was introduced by Polyak in 1960 (Polyak, 1964). In the continuous-time setting, the particle moves along the trajectory defined by

$$\dot{x} = v, \quad \dot{v} = -av - \nabla f(x), \quad (6)$$

where $a > 0$ is a user-defined parameter, independent of t . The method has $o(1/t)$ convergence for convex objective functions, and has faster convergence than gradient flow for strongly convex objectives, changing the rate from modulus of convexity m for gradient flow to \sqrt{m} for heavy ball, whenever $m \in (0, 1)$ is small.

As we discuss below in Section 4, analogous dynamics in the space of probability measures are captured by the following *heavy-ball flow* equation

$$\partial_t \mu_t + \nabla_x \cdot (v \mu_t) - \nabla_v \cdot \left(\left(av + \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] \right) \mu_t \right) = 0, \quad (\text{HBF})$$

where μ_t^X denotes the x -marginal of μ_t . Note the similarity between the coefficients of (6) and (HBF). The theoretical guarantee is also the same.

Formal Theorem A *Let $E : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ be convex along 2-Wasserstein geodesics (see Definition 2) and let μ_t be a solution to the heavy-ball flow equation (HBF). Then its x -marginal μ_t^X satisfies*

$$E[\mu_t^X] - E_* \leq o\left(\frac{1}{t}\right). \quad (7)$$

Furthermore, if $E : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ is m -strongly convex along 2-Wasserstein geodesics (see Definition 2) and we set $a = 2\sqrt{m}$, we have

$$E[\mu_t^X] - E_* \leq O(e^{-\sqrt{m}t}). \quad (8)$$

Note that the convergence rates in (7)-(8) exactly match those of the heavy-ball method (Attouch and Cabot, 2017; Wilson et al., 2016). The rigorous statement of this result is Theorem 7.

Variational acceleration flow. Variational acceleration methods (Wibisono et al., 2016) include Nesterov acceleration (Nesterov, 1983) as a special case. Each member of the class is defined by a triplet $(\alpha_t, \beta_t, \gamma_t)$ that satisfies certain requirements and follows the trajectory defined by the associated Hamiltonian:

$$\dot{x} = v, \quad \dot{v} = -(\dot{\gamma}_t - \dot{\alpha}_t)v - e^{2\alpha_t + \beta_t} \nabla f(x). \quad (9)$$

The method is known to converge with the rate $e^{-\beta_t}$ for convex objective functions f . Essentially, this means the method can converge at an arbitrary rate, given that β_t can be chosen to be any rapidly increasing function of t .

In Section 5, we analyze the counterpart of this approach in the space of probability measures, which we term as the *variational acceleration flow* equation:

$$\partial_t \mu_t + \nabla_x \cdot (v \mu_t) - \nabla_v \cdot \left(\left((\dot{\gamma}_t - \dot{\alpha}_t)v + e^{2\alpha_t + \beta_t} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] \right) \mu_t \right) = 0. \quad (\text{VAF})$$

Note once again the similarity between the coefficients of (VAF) and (9). We establish the following convergence result.

Formal Theorem B *Let $E : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ be convex along 2-Wasserstein geodesics (see Definition 2) and let μ_t be a solution to the Hamiltonian flow (VAF). If the optimal scaling conditions (44) hold, then the x -marginal μ_t^X satisfies*

$$E[\mu_t^X] - E_* \leq O(e^{-\beta t}). \quad (10)$$

As before, the rate of convergence in (10) exactly matches that of the corresponding class of methods in the Euclidean space (Wibisono et al., 2016). The rigorous statement can be found in Theorem 9.

1.1 Summary of Related Work

We identify two types of research results most relevant to the current paper: (1) Acceleration optimization methods (first-order momentum-based methods) on Euclidean space, and (2) accelerated methods on manifolds and for Bayesian sampling, a problem that shares many characteristics with ours.

Momentum-based type methods achieve acceleration by including an artificial *momentum* or *velocity* variable. Notable examples include the heavy-ball method (Polyak, 1964) and the Nesterov’s accelerated method (Nesterov, 1983), which has the optimal convergence rate for convex functions (Nesterov, 2003) and strongly convex functions (Nemirovskij and Yudin, 1983). Traditionally studied in the discrete-in-time setting, recent years have seen investigations of their continuous-time counterparts (Attouch and Alvarez, 2000; Cabot et al., 2009; Attouch and Cabot, 2017; Attouch et al., 2018; Su et al., 2014; Shi et al., 2021; Krichene et al., 2015; Wibisono and Wilson, 2015; Wilson et al., 2016; Betancourt et al., 2018; Muehlebach and Jordan, 2019; Diakonikolas and Jordan, 2021; Scieur et al., 2017; Moucer et al., 2023; Polyak and Shcherbakov, 2017; Allen-Zhu and Orecchia, 2017; Zhang et al., 2018; d’Aspremont et al., 2021; Maddison et al., 2018; França et al., 2020). Continuous-time analysis typically employ a Lyapunov function (Polyak and Shcherbakov, 2017). In Wibisono et al. (2016), the authors found that the introduction of the momentum variable allows one to achieve an arbitrarily high order of convergence, either through a special design of the Hamiltonian or through the time-dilation technique.

Accelerating convergence of first-order methods over the space of probability measures has yet to attract considerable interest, despite the evident importance of this optimization problem in machine learning applications. Topics related to this issue are discussed in Dwivedi et al. (2018); Cheng et al. (2018); Shen and Lee (2019); Lu et al. (2019); García Trillos and Morales (2022); Chow et al. (2020); Liu et al. (2019); Ma et al. (2021); Taghvaei and Mehta (2019); Wang and Li (2022); Zhang et al. (2023). Liu et al. (2019) proposes a framework for a class of accelerated Riemannian optimization algorithms over the probability manifold $\mathcal{P}(\mathbb{R}^d)$. Momentum-based acceleration methods are formulated as optimal control problems in Taghvaei and Mehta (2019), and a Lyapunov function is derived by drawing upon the analogy to classical methods. In the context of Bayesian sampling, Ma et al. (2021) adopts the perspective of extending probability measures to having support on the phase space and formulates the underdamped Langevin dynamics as a flow over the extended space. Convergence of the flows is proved under an assumption that the log-Sobolev inequality holds.

Among the papers referenced, we identify Chow et al. (2020) and Wang and Li (2022) as the ones related most closely to our work. Both papers leverage the second-order differential structure over the manifold of probability measures. Chow et al. (2020) directly formulate the second-order

differential equation, and Wang and Li (2022) adopt a strategy involving the introduction of a Hamiltonian flow across the tangent bundle of the probability measure space.

Both these works build on the definition of “Hamiltonian flow,” so on the surface, they are quite similar to ours. However, there are important differences with our work. Specifically, both studies develop their flow on the *physical space*, focusing on the quantity $\rho(t, x)$. In contrast, our approach introduces a distribution over the *phase space*, with the PDE spanning the entirety of $\mu(t, x, v)$. This shift in perspective results in two major consequences:

- There is no immediate well-posedness theory for the PDEs developed in Chow et al. (2020) and Wang and Li (2022): their PDEs may not have unique solutions. Stringent regularity assumptions were imposed in a different work by one of the authors of our article Carrillo et al. (2019a) to ensure the existence of a unique solution. On the contrary, the Hamiltonian flow PDE of our present paper is *guaranteed* to have a unique solution Ambrosio and Gangbo (2008). Similarly, while Wang and Li (2022) do provide a Lyapunov analysis of their flow, their proof relies on a smooth optimal transport map: It assumes that both the target distribution and the flow solution have Lebesgue densities. In comparison, we work directly on the transport plan and thus can circumvent the regularity assumption.

We believe these improvements to have mathematical depth and to resonate with the comparison between the compressible Euler equation and the Boltzmann equation. The Euler equation serves as the counterpart of the Boltzmann equation on the fluid dynamics side. While Euler develops blow-up singularities, the Boltzmann equation is well-posed DiPerna and Lions (1989). By adding velocity to the unknowns, PDE solutions can span out the singularities to form a regular solution.

- Another important consequence of this difference in formulations lies in the particle representation. In deriving the equations for $\rho(t, x)$, Chow et al. (2020) and Wang and Li (2022) relied on the so-called *mono-kinetic ansatz*, implying that the velocity $v(t, x) = \nabla\phi(t, x)$ lives on the tangent bundle, and is a function of the space variable x . Our approach does not make this assumption, since v is an independent variable, allowing different particles at the same location x to have different velocities. Essentially, while Chow et al. (2020) and Wang and Li (2022) track only the bulk velocity, we allow particles the freedom to roam with individual velocities. Such a conceptual difference resonates with the improvement by the Underdamped Langevin Monte Carlo (ULMC) (Cao et al., 2023; Ma et al., 2021) over the overdamped Langevin Monte Carlo (LMC), where ULMC allows particles to adjust the velocity according to the Hamiltonian, and the PDE is formulated on phase space.

Tanaka (2023) explores an extension of Nesterov’s accelerated method over the space of probability measures. The results in that paper are based on a notion of convexity called *transport convexity*, which differs from *geodesic convexity* considered in our work. This notion can be difficult to verify for several commonly used geodesically convex functionals, including the KL divergence. Additionally, Tanaka (2023) does not provide a convergence rate for the heavy-ball method.

1.2 Organization of the paper

There are two main technical components of the paper. The first is the Wasserstein metric and its induced flow and convexity, while the second concerns Hamiltonian flow methods developed for accelerating optimization in Euclidean space. We review these techniques in Section 2. Section 3 presents our major contributions; We present the Hamiltonian flow PDE in its most general form, and describe the two examples: the heavy-ball method and the variational acceleration flow.

The convergence rates of these methods are discussed in Theorem 7 and Theorem 9, respectively. Section 4 and Section 5 are dedicated to the proof of the two theorems.

2. Background knowledge

This section outlines notions relevant to this paper from our two fundamental building blocks: the Wasserstein metric for quantifying distances between probability measures and its induced convexity, and the Hamiltonian flow that guides the dynamics of particles to achieve acceleration. (Readers familiar with these topics can skip this section.)

2.1 Hamiltonian flows

The idea of accelerating convergence in the space of probability measures arises from the fact that Hamiltonian flows accelerate classical optimization methods in the *Euclidean* space. For the latter, we consider the minimization problem

$$x_* \in \operatorname{argmin}_{x \in \mathbb{R}^d} f(x), \quad (11)$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a sufficiently smooth convex objective function. We denote the optimal value by $f_* = f(x_*)$.

A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is *m-strongly convex* if

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{m}{2} \|y - x\|^2, \quad \text{for all } x, y \in \mathbb{R}^d. \quad (12)$$

The parameter $m \geq 0$ is called the *modulus of convexity*. When $m = 0$, we recover the standard convexity condition. An equivalent definition is that for all $x, y \in \mathbb{R}^d$, we have

$$f(tx + (1-t)y) \leq tf(x) + (1-t)f(y) - \frac{m}{2}t(1-t)\|x - y\|^2, \quad \forall t \in [0, 1]. \quad (13)$$

The most basic first-order strategy for finding the optimal point is the gradient descent method, from which Gradient Flow equation is derived:

$$\dot{x} = -\nabla f(x). \quad (\text{GF})$$

It is well known (see for example Polyak and Shcherbakov, 2017) that (GF) converges with the rate

$$\begin{cases} f(x(t)) - f_* \leq O(t^{-1}) & \text{for convex } f, \\ f(x(t)) - f_* \leq O(e^{-2mt}) & \text{for } m\text{-strongly convex } f. \end{cases} \quad (14)$$

There are many ways to speed up these convergence rates, and Hamiltonian flows provide a path to do so. This approach adds to the position x a velocity v and evolves (x, v) according to a Hamiltonian trajectory. Defining the Hamiltonian $h_t : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ so that

$$h_t(x, v) = k_t(v) + p_t(x), \quad (15)$$

with k_t and p_t termed the kinetic and potential energy, respectively, the Hamiltonian trajectory is defined by (2), restated here:

$$\dot{x} = \nabla_v h_t(x, v), \quad \dot{v} = -\nabla_x h_t(x, v). \quad (16)$$

By selecting carefully h_t —specifically k_t and p_t —one can show that the sample following (16) converges to x_* with accelerated speed. We define the two most famous examples of methods in this class.

Example 1 (Heavy-ball ODE (Polyak, 1964)) *When we set*

$$k_t(v) = \frac{e^{-at}}{2} \|v\|^2, \quad p_t(x) = e^{at} f(x), \quad (17)$$

the Hamiltonian flow (16) becomes

$$\dot{x} = e^{-at} v, \quad \dot{v} = -e^{at} \nabla f(x). \quad (18)$$

Via a change of variable and definition of the scaled velocity $u = e^{-at} v$, we obtain

$$\dot{x} = u, \quad \dot{u} = -au - \nabla f(x). \quad (19)$$

Compared to (14), the heavy-ball method speeds up the convergence of the gradient flow (GF) in both convex and m -strongly convex cases. Specifically, we have (see Attouch and Cabot, 2017; Wilson et al., 2016) that

$$\begin{cases} f(x(t)) - f_* \leq o(t^{-1}) & \text{for convex } f, \text{ when we set } a > 0, \\ f(x(t)) - f_* \leq O(e^{-\sqrt{m}t}) & \text{for } m\text{-strongly convex } f, \text{ when we set } a = 2\sqrt{m}. \end{cases} \quad (20)$$

Example 2 (Variational acceleration (Wibisono et al., 2016)) *Variational acceleration methods give rise to a large class of algorithms proposed in Wibisono et al. (2016) that deploy the following kinetic and potential energy:*

$$k_t(v) = \frac{e^{\alpha_t - \gamma_t}}{2} \|v\|^2, \quad p_t(x) = e^{\alpha_t + \beta_t + \gamma_t} f(x), \quad (21)$$

where $\alpha_t, \beta_t, \gamma_t$ are time-dependent user-defined parameters. For this definition of the Hamiltonian, the flow is

$$\dot{x} = e^{\alpha_t - \gamma_t} v, \quad \dot{v} = -e^{\alpha_t + \beta_t + \gamma_t} \nabla f(x). \quad (22)$$

Defining the scaled velocity $u = e^{\alpha_t - \gamma_t} v$, these equations become

$$\dot{x} = u, \quad \dot{u} = (\dot{\alpha}_t - \dot{\gamma}_t)u - e^{2\alpha_t + \beta_t} \nabla f(x). \quad (23)$$

Under mild assumptions, it was proved in Wibisono et al. (2016) that the dynamics speed up the convergence of (GF) when f is convex, the new rate being

$$f(x(t)) - f_* \leq O(e^{-\beta_t}). \quad (24)$$

One special example within this framework is the Nesterov acceleration method, which chooses $\alpha_t = \log(2/t)$, $\beta_t = \log(t^2/4)$, and $\gamma_t = 2 \log(t)$ and yields

$$\dot{x} = u, \quad \dot{u} = -\frac{3}{t}u - \nabla f(x). \quad (25)$$

For this approach, we obtain

$$f(x(t)) - f_* \leq O(t^{-2}).$$

2.2 Wasserstein metrics and induced convexity

The set of probability measures forms a nonlinear manifold. To quantify the distance between two distributions, the standard L_2 norm inherited from the Hilbert space is insufficient. Instead, we use techniques developed for Riemannian metrics (Otto, 2001). We present the main concepts here, omitting details.

Denoting by $\mathcal{P}_2(\mathbb{R}^d)$ the collection of all probability measures supported on \mathbb{R}^d that have finite second moment, we have the following definition of the 2-Wasserstein distance.

Definition 1 *Given two probability measures $\rho_1, \rho_2 \in \mathcal{P}_2(\mathbb{R}^d)$, the 2-Wasserstein distance W_2 between them is defined by*

$$W_2^2(\rho_1, \rho_2) = \inf \left\{ \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 \gamma(dx dy) : \gamma \in \Gamma(\rho_1, \rho_2) \right\}, \quad (26)$$

where

$$\Gamma(\rho_1, \rho_2) = \left\{ \gamma \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d) : (\pi^1)_\# \gamma = \rho_1, (\pi^2)_\# \gamma = \rho_2 \right\}$$

denotes the collection of all couplings between ρ_1 and ρ_2 . Here, $(\pi^i)_\# \gamma, i = 1, 2$ denotes the i -th marginal of the coupling measure γ . We denote by $\Gamma_o(\rho_1, \rho_2) \subset \Gamma(\rho_1, \rho_2)$ the collection of optimal couplings that attain the minimum in (26).

Note that Γ_o is always non-empty (Villani et al., 2009). According to Brenier's theorem (Brenier, 1991), when the marginal measure ρ_1 (or ρ_2) has Lebesgue density, the optimal coupling γ_o is unique and is induced by a unique transport map $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$, that is, $\gamma_o = (id \times T)_\# \rho_1$.

The 2-Wasserstein distance induces a (formal) Riemannian structure (Otto, 2001) onto $\mathcal{P}_2(\mathbb{R}^d)$. On a Riemannian manifold, the notion of a gradient can be defined through the underlying metric, giving rise in our case to the 2-Wasserstein gradient: For any functional $E : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$,

$$\nabla_{W_2} E[\rho] = -\nabla_x \cdot \left(\rho \nabla_x \frac{\delta E}{\delta \rho} \right).$$

This notion of gradient allows us to define Wasserstein gradient flows in $\mathcal{P}_2(\mathbb{R}^d)$, resembling gradient flows in the Euclidean space. By guiding the evolution of a probability measure along the steepest descent direction, we define the Wasserstein gradient flow by

$$\partial_t \rho = -\nabla_{W_2} E[\rho] = \nabla_x \cdot \left(\rho \nabla_x \frac{\delta E}{\delta \rho} \right). \quad (\text{WGF})$$

With Definition 1 of the distance between probability measures, the concept of convexity needs to be rephrased accordingly.

Definition 2 *For $m \geq 0$, a functional $E : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ is called m -strongly convex if for every $\rho_1, \rho_2 \in \mathcal{P}_2(\mathbb{R}^d)$ and $\gamma_o \in \Gamma_o(\rho_1, \rho_2)$, we have*

$$E[\rho_2] \geq E[\rho_1] + \iint_{\mathbb{R}^d \times \mathbb{R}^d} \left\langle \nabla_x \frac{\delta E}{\delta \rho}[\rho_1](x), y - x \right\rangle d\gamma_o(x, y) + \frac{m}{2} W_2^2(\rho_1, \rho_2), \quad (27)$$

When E satisfies (27) with $m = 0$, we say that E is (geodesically) convex.

Note the resemblance between Definition 2 and strong convexity in Euclidean space (12). Similarly, extending from the equivalent formulation of strong convexity in (13), we should also expect that E evaluated at a point on an interpolation between ρ_1 and ρ_2 should satisfy similar conditions. Indeed, considering the geodesic curve

$$\rho_t^{1 \rightarrow 2} = ((1-t)\pi^1 + t\pi^2)_{\#}\gamma_o, \quad \gamma_o \in \Gamma_o(\rho_1, \rho_2), \quad (28)$$

that connects ρ_1 and ρ_2 , with $\rho_0^{1 \rightarrow 2} = \rho_1$ and $\rho_1^{1 \rightarrow 2} = \rho_2$, Definition 2 can be equivalently seen as requiring

$$E[\rho_t^{1 \rightarrow 2}] \leq (1-t)E[\rho_1] + tE[\rho_2] - \frac{m}{2}t(1-t)W_2^2(\rho_1, \rho_2), \quad \text{for all } t \in [0, 1]. \quad (29)$$

One interesting class of convex functionals is obtained by extending convex potentials. Given a potential function V that is (m -strongly) convex on Euclidean space, its associated potential energy $\mathcal{V} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ defined by

$$\mathcal{V}[\rho] = \int_{\mathbb{R}^d} V(x) \rho(dx), \quad (30)$$

is m -strongly convex on $\mathcal{P}_2(\mathbb{R}^d)$. Another class of convex functionals comes from measuring the KL divergence against a log-concave reference probability measure ρ_* , that is,

$$E[\rho] = KL(\rho || \rho_*) = \int_{\mathbb{R}^d} \rho(x) \log \frac{\rho(x)}{\rho_*(x)} dx, \quad (31)$$

If ρ_* is (m -strongly) log-concave, then E is (m -strongly) convex on $\mathcal{P}_2(\mathbb{R}^d)$. That is, if the reference measure takes the form $\rho_* \propto e^{-g}$ for some function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, then the (m -strong) log-concavity of ρ_* is equivalent to the (m -strong) convexity of g (Ambrosio et al., 2005).

In Euclidean space, the gradient flow finds a minimizer of a convex function. An analogous property holds for the Wasserstein gradient flow (WGF) whenever E is a (m -strongly) convex functional. The convergence behavior of (WGF), as shown by Ambrosio et al. (2005) and Cheng and Bartlett (2018), is as follows:

$$\begin{cases} E[\rho_t] - E_* \leq O(t^{-1}) & \text{for convex } E, \\ E[\rho_t] - E_* \leq O(e^{-2mt}) & \text{for } m\text{-strongly convex } E. \end{cases} \quad (32)$$

We note the exact match of the convergence rates in comparison to the gradient flow (GF) in Euclidean space (14).

3. Hamiltonian flows for optimizing in the space of probability measures

Building on the tools of the previous section, we are ready to define the Hamiltonian flow over the space of probability measures. We first collect all probability measures over the phase space that have finite second moment:

$$\mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d) = \left\{ \mu : \int |x|^2 + |v|^2 d\mu(x, v) < \infty \right\}.$$

For all $\mu \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$, denote by μ^V and μ^X the marginal distributions of x and v , respectively:

$$\mu^V(\cdot) = \int_{\mathbb{R}^d} \mu(dx, \cdot), \quad \mu^X(\cdot) = \int_{\mathbb{R}^d} \mu(\cdot, dv).$$

In the proofs, we use notation $\mu_{t,x}$ for the conditional distribution of $\mu_t \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$, following Ambrosio et al. (2005, Theorem 5.3.1):

$$\mu_{t,x}(v) := \mu_t(v|x).$$

Extending the Hamiltonian defined in (15), we define the Hamiltonian in the probability measure space $H_t : \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d) \rightarrow \mathbb{R}$, having the form

$$H_t[\mu] = K_t[\mu^V] + P_t[\mu^X], \quad (33)$$

where $K_t[\mu^V]$ and $P_t[\mu^X]$ represent the kinetic and potential energy, respectively.

Remark 3 *We note that the definition of the Hamiltonian separates kinetic and potential energy, each of which depends on just one of the v -marginal and the x -marginal of the distribution. It is also possible to define a Hamiltonian that depends on the joint distribution. One such possibility was deployed in the underdamped Langevin dynamics (Ma et al., 2021), where the Hamiltonian is the KL divergence between μ and the distribution $\rho_* \otimes \nu_*$. Here ρ_* is the target distribution and $\nu_* \propto \exp(-|v|^2/2)$ represents the standard Gaussian distribution over the velocity space. As elaborated in Ma et al. (2021), the underdamped Langevin dynamics can be viewed as a damped version of our Hamiltonian flow.*

We define the Hamiltonian flow on the space of probability measure as follows.

Definition 4 (Hamiltonian flow over probability measures) *Let $t \mapsto H_t$ be the time-dependent Hamiltonian over $\mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$. A Hamiltonian flow with respect to H_t is a curve $t \mapsto \mu_t$ that satisfies*

$$\partial_t \mu_t + \nabla_x \cdot \left(\mu_t \nabla_v \frac{\delta H_t}{\delta \mu} [\mu_t] \right) - \nabla_v \cdot \left(\mu_t \nabla_x \frac{\delta H_t}{\delta \mu} [\mu_t] \right) = 0 \quad \text{in the distributional sense,} \quad (34)$$

with initial condition $\mu_{t=0} = \mu_0 \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$.

This definition provides the evolution of measures $t \mapsto \mu_t$. Well-posedness and absolute continuity of (34) with geodesically convex Hamiltonian and general initial data for this equation have been studied in Ambrosio and Gangbo (2008). We note that our definition of Hamiltonian flow is different from the conventional one; see Chow et al. (2020). Specifically, our formulation expands μ^X to the phase space μ and allows each sample to take on different velocities. We argue that this flow is physically meaningful, intuitive, and gives a meaningful reason to deploy the Hamiltonian flow (34) to evolve the probability to minimize E as shown in the following result.

Proposition 5 *The motion of $\delta_{(x(t),v(t))}$, viewed as a probability measure to optimize E , agrees with that of $(x(t),v(t))$, viewed as a sample to optimize f , if E and f , H_t and h_t are related as follows:*

$$E[\rho] = \int_{\mathbb{R}^d} f \, d\rho, \quad H_t[\mu] = \int_{\mathbb{R}^{2d}} h_t \, d\mu. \quad (35)$$

More precisely, we have the following.

- (1) *If $t \mapsto (x(t),v(t))$ solves the Hamiltonian ODE in (16), then the curve of Dirac measure $t \mapsto \mu_t := \delta_{(x(t),v(t))}$ solves the Hamiltonian PDE (34).*
- (2) *If $t \mapsto x(t)$ in (16) converges to $x_* \in \operatorname{argmin}_x f$, then the Hamiltonian PDE (34) drives the x -marginal of μ_t towards δ_{x_*} , a minimizer of E .*

Proof To prove (1), we take an arbitrary $\phi \in C_c^\infty(\mathbb{R}^d \times \mathbb{R}^d)$ and test it on (34), showing that the result is zero. Test ϕ on the $\partial_t \mu_t$ term with $\mu_t = \delta_{(x(t), v(t))}$, we obtain

$$\begin{aligned} \frac{d}{dt} \int \phi d\mu_t &= \frac{d}{dt} \phi(x(t), v(t)) = \nabla_x \phi \cdot \dot{x} + \nabla_v \phi \cdot \dot{v} \\ &= \nabla_x \phi \cdot \nabla_v h_t(x(t), v(t)) - \nabla_v \phi \cdot \nabla_x h_t(x(t), v(t)), \end{aligned} \quad (36)$$

where we used chain rule and applied (16) in the last equation. Testing ϕ on the other two terms in (34), we obtain

$$\begin{aligned} \int \nabla_x \cdot \left(\mu_t \nabla_v \frac{\delta H_t}{\delta \mu} \right) \phi dx dv - \int \nabla_v \cdot \left(\mu_t \nabla_x \frac{\delta H_t}{\delta \mu} \right) \phi dx dv \\ = - \int \mu_t \left(\nabla_v \frac{\delta H_t}{\delta \mu} \cdot \nabla_x \phi - \nabla_x \frac{\delta H_t}{\delta \mu} \cdot \nabla_v \phi \right) dx dv \\ = - \nabla_x \phi \cdot \nabla_v \frac{\delta H_t}{\delta \mu}(x(t), v(t)) + \nabla_v \phi \cdot \nabla_x \frac{\delta H_t}{\delta \mu}(x(t), v(t)). \end{aligned} \quad (37)$$

The relation (35) implies that

$$H_t[\delta_{(x(t), v(t))}] = h_t(x(t), v(t)), \quad \frac{\delta H_t}{\delta \mu} = h_t.$$

Substituting into (37) and summing (36) and (37), we verify that the result is zero. Since ϕ is arbitrary, we conclude that $t \mapsto \delta_{(x(t), v(t))}$ solves (34) in the distributional sense.

To show item (2), we need only note that the relation (35) guarantees

$$E[\delta_{(x(t), v(t))}^X] = E[\delta_{x(t)}] = f(x(t)), \quad E_* = E[\delta_{x_*}] = f(x_*) = f_*,$$

thereby concluding the proof. ■

Building on the Hamiltonian flow of Definition 4, we provide two examples in the next two subsections. Both show an improvement in the convergence rate for the problem of finding an optimal ρ . Some other examples are collected in Appendix A.

3.1 Heavy-Ball Flow

The heavy-ball method is known to converge as $O(e^{-\sqrt{m}t})$ in Euclidean space, for m -strongly convex objectives and $o(1/t)$ for convex objectives. We find the corresponding rates for this algorithm in the probability measure space here.

By analogy to Example 1, we define the following Hamiltonian for any $\mu \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$:

$$H_t[\mu] = K_t[\mu^V] + P_t[\mu^X] = \frac{e^{-at}}{2} \int_{\mathbb{R}^d} \|v\|^2 d\mu^V + e^{at} E[\mu^X], \quad (38)$$

where $a > 0$ is a user-defined parameter. The Fréchet derivative is

$$\frac{\delta H_t}{\delta \mu}[\mu] = \frac{e^{-at}}{2} \|v\|^2 + e^{at} \frac{\delta E}{\delta \rho}[\mu^X], \quad (39)$$

so the Hamiltonian PDE (34) becomes

$$\partial_t \mu_t + \nabla_x \cdot (\mu_t e^{-at} v) - \nabla_v \cdot \left(\mu_t e^{at} \nabla_x \frac{\delta E}{\delta \rho}[\mu_t^X] \right) = 0. \quad (40)$$

In (19) we introduced a change of variables. Correspondingly, we define $u = e^{-at}v$ and denote by $\tilde{\mu}$ the probability measure over this new variable. Under this change of variable, (40) becomes (HBF). We spell out this claim in the following proposition.

Proposition 6 *Let μ_t solve the heavy-ball equation (40). Define the map $T_t(x, v) = (x, e^{-at}v)$ and set $\tilde{\mu}_t := (T_t)_\# \mu_t$ to be the pushforward of μ_t under T_t . Then $\tilde{\mu}_t$ solves the equation*

$$\partial_t \tilde{\mu}_t + \nabla_x \cdot (\tilde{\mu}_t u) - \nabla_u \cdot \left(\tilde{\mu}_t \left(au + \nabla_x \frac{\delta E}{\delta \rho} [\tilde{\mu}_t^X] \right) \right) = 0, \quad (41)$$

Proof To show that $\tilde{\mu}_t$ satisfies (41), we take an arbitrary $\phi \in C_c^\infty(\mathbb{R}^d \times \mathbb{R}^d)$, test it on $\tilde{\mu}_t$, and compute its time derivative to obtain

$$\begin{aligned} \frac{d}{dt} \int \phi(x, u) d\tilde{\mu}_t(x, u) &= \frac{d}{dt} \int \phi(x, e^{-at}v) d\mu_t(x, v) \\ &= \underbrace{\int \frac{d}{ds} \phi(x, e^{-as}v) \Big|_{s=t} d\mu_t(x, v)}_{\text{Term I}} + \underbrace{\int \phi(x, e^{-at}v) d\partial_t \mu_t(x, v)}_{\text{Term II}}, \end{aligned}$$

where in the first equality we use the definition of pushforward map. Term I can be written as an integral in $\tilde{\mu}_t$ by direct computation, that is,

$$\text{Term I} = -a \int e^{-at}v \cdot \nabla_v \phi(x, e^{-at}v) d\mu_t(x, v) = -a \int u \cdot \nabla_u \phi(x, u) d\tilde{\mu}_t(x, u),$$

where we use the pushforward definition again in the second equality. For Term II, we use the fact that μ_t solves (40) to obtain

$$\begin{aligned} \text{Term II} &= - \int \phi(x, e^{-at}v) d(\nabla_x \cdot (\mu_t e^{-at}v)) + \int \phi(x, e^{-at}v) d \left(\nabla_v \cdot \left(\mu_t e^{at} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] \right) \right) \\ &= \int e^{-at}v \cdot \nabla_x \phi(x, e^{-at}v) d\mu_t(x, v) - \int e^{at} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] \cdot e^{-at} \nabla_v \phi(x, e^{-at}v) d\mu_t(x, v) \\ &= \int u \cdot \nabla_x \phi(x, u) d\tilde{\mu}_t(x, u) - \int \nabla_x \frac{\delta E}{\delta \rho} [\tilde{\mu}_t^X] \cdot \nabla_u \phi(x, u) d\tilde{\mu}_t(x, u). \end{aligned}$$

In the last equality, we use $\tilde{\mu}_t^X = \mu_t^X$. By combining both terms, we arrive at (41). \blacksquare

The heavy-ball gradient-flow PDE (40) speeds up the convergence of (WGF) in the same way as the heavy-ball ODE (18) speeds up (GF). We state the result here and leave the proof to Section 4. In the following, we denote by $AC([0, \infty), \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d))$ the space of absolutely continuous curves over $\mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$ (Ambrosio et al., 2005).

Theorem 7 *Let $\mu \in AC([0, \infty), \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d))$ solve the heavy-ball flow (40) for $E : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$. (Equivalently, let $\tilde{\mu}$ solve (HBF).) If E and μ_t are sufficiently smooth, then the marginal distribution μ_t^X converges as follows:*

$$E[\mu_t^X] - E_* \leq o\left(\frac{1}{t}\right) \quad \text{for convex } E, \text{ when we set } a > 0, \quad (42a)$$

$$E[\mu_t^X] - E_* \leq O\left(e^{-\sqrt{m}t}\right) \quad \text{for } m\text{-strongly convex } E, \text{ when we set } a = 2\sqrt{m}. \quad (42b)$$

Note that convexity and m -strong convexity adopt the new geodesic convexity concept.

3.2 Variational Acceleration Flow

The second example generalizes the variational formulation method in Example 2, where the Hamiltonian is chosen to be

$$H_t[\mu] = K_t[\mu^V] + P_t[\mu^X] = \frac{e^{\alpha_t - \gamma_t}}{2} \int_{\mathbb{R}^d} \|v\|^2 d\mu^V + e^{\alpha_t + \beta_t + \gamma_t} E[\mu^X]. \quad (43)$$

Here, the functions $\alpha_t, \beta_t, \gamma_t$ are user-defined parameters that satisfy the optimal scaling conditions

$$\dot{\beta}_t \leq e^{\alpha_t}, \quad \dot{\gamma}_t = e^{\alpha_t} \quad (44)$$

By differentiating (43), we obtain

$$\frac{\delta H_t}{\delta \mu}[\mu] = \frac{e^{\alpha_t - \gamma_t}}{2} \|v\|^2 + e^{\alpha_t + \beta_t + \gamma_t} \frac{\delta E}{\delta \rho}[\mu^X], \quad (45)$$

so the associated Hamiltonian PDE is

$$\partial_t \mu_t + \nabla_x \cdot (\mu_t e^{\alpha_t - \gamma_t} v) - \nabla_v \cdot \left(\mu_t e^{\alpha_t + \beta_t + \gamma_t} \nabla_x \frac{\delta E}{\delta \rho}[\mu_t^X] \right) = 0. \quad (46)$$

Similar to the change of variables performed in (23), we set $u = e^{\alpha_t - \gamma_t} v$ and denote $\tilde{\mu}$ as the measure defined over the new variables (x, u) . Then, (46) becomes (VAF). We formalize this claim in the following proposition.

Proposition 8 *Let μ_t solve the variational acceleration flow equation (46). Define the map $T_t(x, v) = (x, e^{\alpha_t - \gamma_t} v)$ and set $\tilde{\mu}_t := (T_t)_\# \mu_t$ to be the pushforward of μ_t under T_t . Then $\tilde{\mu}_t$ solves the equation*

$$\partial_t \tilde{\mu}_t + \nabla_x \cdot (\tilde{\mu}_t u) - \nabla_v \cdot \left(\tilde{\mu}_t \left((\dot{\gamma}_t - \dot{\alpha}_t) u + e^{2\alpha_t + \beta_t} \nabla_x \frac{\delta E}{\delta \rho}[\tilde{\mu}_t^X] \right) \right) = 0. \quad (47)$$

Proof The derivation of (47) from (46) involves a computation similar to that of Proposition 6. We omit the details. \blacksquare

The following set of parameter choice satisfy (44):

$$\alpha_t = \log(2/t), \quad \beta_t = \log(t^2/4), \quad \gamma_t = 2 \log(t), \quad (48)$$

leading to the Nesterov flow:

$$\partial_t \tilde{\mu}_t + \nabla_x \cdot (\tilde{\mu}_t u) - \nabla_v \cdot \left(\tilde{\mu}_t \left(\frac{3}{t} u + \nabla_x \frac{\delta E}{\delta \rho}[\tilde{\mu}_t^X] \right) \right) = 0. \quad (49)$$

As for heavy-ball, we observe the speedup of this variational Hamiltonian flow (46) compared to the Wasserstein gradient-flow (WGF), with the improvement exactly matching that of variational acceleration method (22) over (GF). (The proof of this result is the subject of Section 5.)

Theorem 9 *Let the objective functional $E : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ be convex along 2-Wasserstein geodesics. Let $\mu \in AC([0, \infty), \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d))$ solve the Hamiltonian flow (46) for E . (Equivalently, let $\tilde{\mu}$ solve (VAF).) If E and μ_t are sufficiently smooth and the optimal scaling conditions (44) holds, then the marginal μ_t^X satisfies*

$$E[\mu_t^X] - E_* \leq O(e^{-\beta_t}). \quad (50)$$

Moreover, with coefficients configured as in (48), we have

$$E[\mu_t^X] - E_* \leq O\left(\frac{1}{t^2}\right). \quad (51)$$

4. The Heavy-Ball Flow

This section describes the convergence of the heavy-ball flow. Following some preliminary results, we treat the general convex case followed by the strongly convex case.

4.1 Preliminary Results

First, we consider $\nu \in AC([0, \infty), \mathcal{P}_2(\mathbb{R}^d))$ that solves

$$\partial_t \nu + \nabla \cdot (\nu \xi) = 0, \quad (52)$$

for a given bounded and sufficiently smooth vector field $\xi(t, x)$. Then for any $\sigma \in \mathcal{P}_2(\mathbb{R}^d)$, we have

$$\frac{1}{2} \frac{d}{dt} W_2^2(\nu_t, \sigma) = \int_{\mathbb{R}^d \times \mathbb{R}^d} \langle x - y, \xi_t(x) \rangle d\gamma_t(x, y), \quad (53)$$

where $\gamma_t \in \Gamma_o(\nu_t, \sigma)$ is an optimal coupling between ν_t and σ .

Let $\mu \in AC([0, \infty), \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d))$ solve

$$\partial_t \mu_t + \nabla_x \cdot (\mu_t F(t, v)) - \nabla_v \cdot (\mu_t G(t, \mu_t, x)) = 0. \quad (54)$$

Then under some smoothness requirement on F and G , we have

$$\begin{aligned} \frac{1}{2} \frac{d^+}{dt} \frac{d}{dt} W_2^2(\mu_t^X, \sigma) &\leq \int_{\mathbb{R}^{2d}} \|F(t, v)\|^2 d\mu_t \\ &+ \int_{\mathbb{R}^{3d}} \langle x - y, \partial_t F(t, v) - \nabla_v F(t, v) G(t, \mu_t, x) \rangle d\mu_{t,x}(v) d\gamma_t(x, y), \end{aligned} \quad (55)$$

where d^+/dt denotes the upper derivative at almost every $t \geq 0$, $\nabla_v F = (\partial_{v_j} F^i)_{ij}$ denotes the Jacobian, and $\mu_{t,x}$ is the conditional distribution of μ_t . The rigorous version of these statements and the smoothness requirements are presented in Appendix B.

For the heavy-ball flow (40), we obtain the following result.

Lemma 10 *Let $\mu \in AC([0, \infty), \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d))$ be a solution of the heavy-ball flow (40) and $\sigma \in \mathcal{P}_2(\mathbb{R}^d)$ be an arbitrary measure. If E and μ_t^X are sufficiently smooth, we obtain*

$$\frac{1}{2} \frac{d}{dt} W_2^2(\mu_t^X, \sigma) = \int_{\mathbb{R}^{3d}} \langle e^{-at} v, x - y \rangle d\mu_{t,x}(v) d\gamma_t(x, y), \quad (56)$$

and

$$\begin{aligned} \frac{1}{2} \frac{d^+}{dt} \frac{d}{dt} W_2^2(\mu_t^X, \sigma) &\leq \int_{\mathbb{R}^{2d}} \|e^{-at} v\|^2 d\mu_t(x, v) \\ &- \int_{\mathbb{R}^{3d}} \left\langle x - y, a e^{-at} v + \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\mu_{t,x}(v) d\gamma_t(x, y). \end{aligned} \quad (57)$$

where $\{\mu_{t,x}\}_{x \in \mathbb{R}^d} \subset \mathcal{P}_2(\mathbb{R}^d)$ is the conditional distribution of μ_t with respect to its x -marginal distribution μ_t^X , and $\gamma_t \in \Gamma_o(\mu_t^X, \sigma)$ is an optimal coupling between μ_t^X and σ .

Proof To prove (56), we apply (53). We first derive the continuity equation for μ_t^X analogously to (52) by integrating (40) over $v \in \mathbb{R}^d$, which gives

$$0 = \partial_t \mu_t^X + \nabla_x \cdot \left(\int_{\mathbb{R}^d} e^{-at} v \mu_t(\cdot, dv) \right) = \partial_t \mu_t^X + \nabla_x \cdot (e^{-at} \bar{v}_t(x) \mu_t^X), \quad (58)$$

where $\bar{v}_t(x) = \int_{\mathbb{R}^d} v \, d\mu_{t,x}(v)$. Applying (53) to (58) with $\xi(t, x) = e^{-at}\bar{v}_t(x)$ then gives (56).

To derive (57), we apply (55) to (40) with $F(t, v) = e^{-at}v$ and $G(t, \mu_t^X, x) = e^{at}\nabla\frac{\delta E}{\delta\rho}[\mu_t^X](x)$. ■

Furthermore, noting that $\frac{d}{dt}E[\mu_t^X] = \int \frac{\delta E}{\delta\mu_t^X} d\partial_t\mu_t^X$, we see that the evolution of $E[\mu_t^X]$ depends on the evolution of μ_t^X . Using (58), we then obtain

$$\begin{aligned} \frac{d}{dt}E[\mu_t^X] &= \int \frac{\delta E}{\delta\mu_t^X} \frac{d\mu_t^X}{dt} = \int_{\mathbb{R}^d} \left\langle \nabla \frac{\delta E}{\delta\rho}[\mu_t^X](x), e^{-at}\bar{v}_t(x) \right\rangle d\mu_t^X(x) \\ &= \int_{\mathbb{R}^{2d}} \left\langle \nabla \frac{\delta E}{\delta\rho}[\mu_t^X](x), e^{-at}v \right\rangle d\mu_t(x, v). \end{aligned} \quad (59)$$

The proofs below call for repeated use of (56), (57), and (59).

4.2 Convex case

We show here the convergence rate for the general convex case, stated in (42a). As for the Euclidean space analysis, we define a Lyapunov function as follows:

$$\mathcal{E}_t := \frac{1}{2} \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 d\mu_t(x, v) + E[\mu_t^X] - E_*. \quad (60)$$

It can be shown that \mathcal{E}_t decays in time. By differentiating (60), we obtain

$$\frac{d}{dt}\mathcal{E}_t = -a \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 d\mu_t + \underbrace{\frac{1}{2} \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 \frac{d\mu_t}{dt}}_{\text{Term II}} + \underbrace{\frac{d}{dt}E[\mu_t^X]}_{\text{Term III}}. \quad (61)$$

Noting that Term III is already expanded (59), we show that Term II cancels it. To see this, we recall (40) and use integration by parts in v to obtain

$$\frac{1}{2} \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 \frac{d\mu_t}{dt} = - \int_{\mathbb{R}^{2d}} \left\langle e^{-2at}v, e^{at}\nabla\frac{\delta E}{\delta\rho}[\mu_t^X](x) \right\rangle d\mu_t. \quad (62)$$

Thus (61) simplifies to

$$\frac{d}{dt}\mathcal{E}_t = -a \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 d\mu_t \leq 0, \quad (63)$$

showing monotonic decrease of \mathcal{E}_t . This feature implies that

$$\int_{\frac{t}{2}}^t \mathcal{E}_s \, ds \geq \frac{t}{2} \mathcal{E}_t, \quad \forall t \geq 0 \implies E[\mu_t^X] - E_* \leq \mathcal{E}_t \leq \frac{2}{t} \int_{\frac{t}{2}}^t \mathcal{E}_s \, ds. \quad (64)$$

Therefore, to show that $E[\mu_t^X] - E_* \leq o(t^{-1})$ requires showing that $\int_{t/2}^t \mathcal{E}_s \, ds \rightarrow 0$ as $t \rightarrow \infty$. In the current case, we can show further that $\mathcal{E}_t \in L^1([0, \infty))$, that is $\int_0^\infty \mathcal{E}_t \, dt < \infty$, implying that

$$\lim_{t \rightarrow \infty} \int_{t/2}^t \mathcal{E}_s \, ds \leq \lim_{t \rightarrow \infty} \int_{t/2}^\infty \mathcal{E}_s \, ds = 0.$$

To show $\int_0^\infty \mathcal{E}_t \, dt < \infty$, we note that

$$\int_0^\infty \mathcal{E}_t \, dt = \underbrace{\int_0^\infty \frac{1}{2} \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 d\mu_t \, dt}_{\text{Term 1}} + \underbrace{\int_0^\infty E[\mu_t^X] - E_* \, dt}_{\text{Term 2}}. \quad (65)$$

For Term 1, we integrate (63) over time to obtain

$$2 \times \text{Term 1} = \int_0^\infty \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 d\mu_t dt = \lim_{s \rightarrow \infty} \int_0^s \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 d\mu_t dt = \lim_{s \rightarrow \infty} \frac{\mathcal{E}_0 - \mathcal{E}_s}{a} \leq \frac{\mathcal{E}_0}{a}, \quad (66)$$

proving finiteness of this term. For Term 2, we show below that

$$\text{Term 2} = \int_0^\infty E[\mu_t^X] - E_* dt \leq a\mathcal{L}_0. \quad (67)$$

From (65), we thus have boundedness of $\int_0^\infty \mathcal{E}_t dt$ and thus the convergence of $E[\mu_t^X] - E_*$ with a rate of $o(t^{-1})$.

To show (67), we define the following Lyapunov function inspired by the proof of the regular heavy-ball method (Attouch and Cabot, 2017):

$$\mathcal{L}_t = \int_{\mathbb{R}^{3d}} \frac{1}{2} \left\| x + \frac{e^{-at}}{a}v - y \right\|^2 d\mu_{t,x}(v) d\gamma_t(x, y) + \frac{1}{a^2} (E[\mu_t^X] - E_*), \quad (68)$$

where, as above, $\{\mu_{t,x}\}_{x \in \mathbb{R}^d} \subset \mathcal{P}_2(\mathbb{R}^d)$ is the conditional distribution of μ_t at x and $\gamma_t \in \Gamma_o(\mu_t^X, \rho_*)$ is an optimal transport plan. By expanding the quadratic term, we obtain

$$\mathcal{L}_t = \underbrace{\int_{\mathbb{R}^{2d}} \frac{1}{2} \|x - y\|^2 d\gamma_t}_{\text{term A}} + \underbrace{\int_{\mathbb{R}^{3d}} \left\langle \frac{e^{-at}}{a}v, x - y \right\rangle d\mu_{t,x} d\gamma_t}_{\text{term B}} + \underbrace{\int_{\mathbb{R}^{2d}} \frac{1}{2} \left\| \frac{e^{-at}}{a}v \right\|^2 d\mu_t}_{\text{term C}} + \underbrace{\frac{1}{a^2} (E[\mu_t^X] - E_*)}_{\text{term D}}. \quad (69)$$

We now take the time derivative of each term to finally show that

$$\frac{d^+}{dt} \mathcal{L}_t \leq -\frac{1}{a} (E[\mu_t^X] - E_*). \quad (70)$$

Term A. Noticing that Term A is equivalent to $\frac{1}{2}W_2^2(\mu_t^X, \rho_*)$, we utilize (56) to obtain

$$\frac{d}{dt} \text{Term A} = \int_{\mathbb{R}^{3d}} \langle e^{-at}v, x - y \rangle d\mu_{t,x}(v) d\gamma_t(x, y). \quad (71)$$

Term B. From (56) with $\sigma = \rho_*$, we note that Term B can be written as

$$\text{Term B} = \frac{1}{a} \int_{\mathbb{R}^{3d}} \langle e^{-at}v, x - y \rangle d\mu_{t,x}(v) d\gamma_t(x, y) = \frac{1}{2a} \frac{d}{dt} W_2^2(\mu_t^X, \rho_*).$$

Thus, we have

$$\frac{d^+}{dt} \text{Term B} \leq \frac{1}{a} \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 d\mu_t - \int_{\mathbb{R}^{3d}} \left\langle x - y, e^{-at}v + \frac{1}{a} \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\mu_{t,x} d\gamma_t.$$

Term C. Differentiating this term in time gives

$$\frac{d}{dt} \text{Term C} = -\frac{1}{a} \int_{\mathbb{R}^{2d}} \|e^{-at}v\|^2 d\mu_t - \int_{\mathbb{R}^{2d}} \left\langle \frac{e^{-2at}}{a^2}v, e^{at} \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\mu_t, \quad (72)$$

where the two terms in this expression come from differentiating $\|e^{-at}v\|$ and μ_t .

Term D. By differentiating in time and following (59), we obtain

$$\frac{d}{dt} \text{Term D} = \frac{1}{a^2} \int_{\mathbb{R}^{2d}} \left\langle e^{-2at} v, e^{at} \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\mu_t.$$

By summing Terms A, B, C, and D, we obtain

$$\frac{d^+}{dt} \mathcal{L}_t \leq -\frac{1}{a} \int_{\mathbb{R}^{3d}} \left\langle x - y, \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\mu_{t,x} d\gamma_t \leq -\frac{1}{a} (E[\mu_t^X] - E_*), \quad (73)$$

where we used the convexity of E in the last inequality, concluding our proof of the claim (70).

By integrating (73) in time, we obtain

$$\int_0^\infty (E[\mu_t^X] - E_*) dt = \lim_{s \rightarrow \infty} \int_0^s (E[\mu_t^X] - E_*) dt \leq a \lim_{s \rightarrow \infty} (\mathcal{L}_0 - \mathcal{L}_s) \leq a \mathcal{L}_0. \quad (74)$$

showing (67) and concluding the proof.

4.3 Strongly Convex Case

This section is dedicated to showing the convergence rate (42b) for a strongly convex function f with modulus of convexity $m > 0$.

We define a Lyapunov function inspired by the one in Wilson et al. (2016):

$$\mathcal{L}_t = \left(\int_{\mathbb{R}^{3d}} \frac{m}{2} \left\| x + \frac{e^{-2\sqrt{m}t}}{\sqrt{m}} v - y \right\|^2 d\mu_{t,x}(v) d\gamma_t(x, y) + E[\mu_t^X] - E_* \right) e^{\sqrt{m}t}. \quad (75)$$

By expanding the quadratic term, using the same expansion as in (69), we can write

$$\mathcal{L}_t = e^{\sqrt{m}t} \left(\underbrace{\frac{m}{2} W_2^2(\mu_t^X, \rho_*)}_{\text{Term A}} + \underbrace{\frac{\sqrt{m}}{2} \frac{d}{dt} W_2^2(\mu_t^X, \rho_*)}_{\text{Term B}} + \underbrace{\frac{1}{2} \int_{\mathbb{R}^{2d}} \|e^{-2\sqrt{m}t} v\|^2 d\mu_t}_{\text{Term C}} + \underbrace{E[\mu_t^X] - E_*}_{\text{Term D}} \right). \quad (76)$$

We show below that $\frac{d^+}{dt} \mathcal{L}_t \leq 0$ by analyzing the contributions of the four terms in turn.

Term A. Using the same strategy as for Term A in Section 4.2 and applying (53), we have

$$\frac{d}{dt} \text{Term A} = \frac{m}{2} \frac{d}{dt} W_2^2(\mu_t^X, \rho_*) = m \times (56), \quad (77)$$

for $\gamma_t \in \Gamma_o(\mu_t^X, \rho_*)$.

Term B. By applying (55) to the heavy-ball flow PDE (40) with $a = 2\sqrt{m}$, we obtain an upper bound for the second-order derivative, as follows:

$$\frac{d^+}{dt} \text{Term B} \leq \frac{\sqrt{m}}{2} \frac{d^+}{dt} \frac{d}{dt} W_2^2(\mu_t^X, \rho_*) = \sqrt{m} \times (57).$$

Term C. By differentiating Term C with respect to t , we obtain

$$\frac{d}{dt} \text{Term C} = -2\sqrt{m} \int_{\mathbb{R}^{2d}} \|e^{-2\sqrt{m}t} v\|^2 d\mu_t - \int_{\mathbb{R}^{2d}} \left\langle e^{-2\sqrt{m}t} v, \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\mu_t.$$

Term D. Here we recall (59).

By assembling Terms A, B, C, and D, and substituting into (76), we obtain

$$\begin{aligned} \frac{d^+}{dt} \mathcal{L}_t &= \sqrt{m} e^{\sqrt{m}t} \left(\frac{m}{2} W_2^2(\mu_t^X, \rho_*) + E[\mu_t^X] - E_* - \int_{\mathbb{R}^{2d}} \left\langle x - y, \nabla \frac{\delta E}{\delta \rho}[\mu_t^X](x) \right\rangle d\gamma_t \right) \\ &\quad - \frac{\sqrt{m} e^{\sqrt{m}t}}{2} \int_{\mathbb{R}^{2d}} \|e^{-2\sqrt{m}t} v\|^2 d\mu_t \\ &\leq \underbrace{\sqrt{m} e^{\sqrt{m}t} \left(\frac{m}{2} W_2^2(\mu_t^X, \rho_*) + E[\mu_t^X] - E_* - \int_{\mathbb{R}^{3d}} \left\langle x - y, \nabla \frac{\delta E}{\delta \rho}[\mu_t^X](x) \right\rangle d\gamma_t \right)}_{\leq 0} \leq 0. \end{aligned} \quad (78)$$

The last inequality follows from the strong convexity of E (Definition 2).

Recalling (75), we obtain that

$$e^{\sqrt{m}t} (E[\mu_t^X] - E_*) \leq \mathcal{L}_t \leq \mathcal{L}_{t=0} \Rightarrow E[\mu_t^X] - E_* \leq O(e^{-\sqrt{m}t}), \quad (79)$$

concluding the proof.

5. Variational Acceleration Flow

Here we study the convergence rate of variational acceleration flow.

To prepare, we integrate (46) in v to obtain the evolution of the x -marginal μ_t^X :

$$0 = \partial_t \mu_t^X + \nabla_x \cdot \left(\int_{\mathbb{R}^d} e^{\alpha t - \gamma t} v \mu_t(\cdot, dv) \right) = \partial_t \mu_t^X + \nabla_x \cdot (e^{\alpha t - \gamma t} \bar{v}_t(x) \mu_t^X), \quad (80)$$

where $\bar{v}_t(x) = \int_{\mathbb{R}^d} v d\mu_{t,x}(v)$ and $\{\mu_{t,x}\}_{x \in \mathbb{R}^d} \subset \mathcal{P}_2(\mathbb{R}^d)$ is the conditional distribution, given μ_t^X .

5.1 Proof of Theorem 9

First, deploying the strategy of the previous section, we define a Lyapunov function

$$\mathcal{L}_t := \frac{1}{2} \int_{\mathbb{R}^{3d}} \|x + e^{-\gamma t} v - y\|^2 d\mu_{t,x}(v) d\bar{\gamma}_t(x, y) + e^{\beta t} (E[\mu_t^X] - E_*), \quad (81)$$

where $\bar{\gamma}_t \in \Gamma_o(\mu_t^X, \rho_*)$ is the optimal transport plan between μ_t^X and ρ_* . We show below that $\frac{d^+}{dt} \mathcal{L}_t \leq 0$, which implies that

$$e^{\beta t} (E[\mu_t^X] - E_*) \leq \mathcal{L}_t \leq \mathcal{L}_{t=0}, \quad \forall t \geq 0, \quad (82)$$

from which it follows that $E[\mu_t^X] - E_* \leq O(e^{-\beta t})$.

By expanding (81), we obtain

$$\mathcal{L}_t = \underbrace{\frac{1}{2} W_2^2(\mu_t^X, \rho_*)}_{\text{Term A}} + e^{-\alpha t} \underbrace{\frac{1}{2} \frac{d}{dt} W_2^2(\mu_t^X, \rho_*)}_{\text{Term B}} + \frac{1}{2} e^{-2\gamma t} \underbrace{\int_{\mathbb{R}^{2d}} \|v\|^2 d\mu_t}_{\text{Term C}} + e^{\beta t} \underbrace{(E[\mu_t^X] - E_*)}_{\text{Term D}}. \quad (83)$$

Following a familiar strategy, we analyze the derivatives of the four terms in turn.

Term A. Deploying (53) in the context of (46), we have

$$\begin{aligned}
 \frac{d}{dt} \text{Term A} &= \frac{1}{2} \frac{d}{dt} W_2^2(\mu_t^X, \rho_*) \\
 &= \frac{1}{2} \int_{\mathbb{R}^{2d}} \langle e^{\alpha_t - \gamma_t} \bar{v}_t(x), x - y \rangle d\bar{\gamma}_t(x, y) \\
 &= \frac{1}{2} \int_{\mathbb{R}^{3d}} \langle e^{\alpha_t - \gamma_t} v, x - y \rangle d\mu_{t,x}(v) d\bar{\gamma}_t(x, y).
 \end{aligned} \tag{84}$$

Term B. By deploying (55) in the context of (46), we have

$$\begin{aligned}
 \frac{d^+}{dt} \text{Term B} &= \frac{1}{2} \frac{d^+}{dt} \frac{d}{dt} W_2^2(\mu_t^X, \sigma) \\
 &\leq \int_{\mathbb{R}^{2d}} \|e^{\alpha_t - \gamma_t} v\|^2 d\mu_t \\
 &\quad - \int_{\mathbb{R}^{2d}} \left\langle e^{\alpha_t - \gamma_t} (x - y), (e^{\alpha_t} - \dot{\alpha}_t) \bar{v}_t(x) + e^{\alpha_t + \beta_t + \gamma_t} \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\bar{\gamma}_t(x, y).
 \end{aligned} \tag{85}$$

Term C. By differentiating Term C with respect to time and utilizing (46), we have

$$\frac{d}{dt} \text{Term C} = \int \|v\|^2 \frac{d\mu}{dt} = 2e^{\alpha_t + \beta_t + \gamma_t} \int \left\langle v, \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] \right\rangle d\mu_t.$$

Term D. By differentiating Term D with respect to time and utilizing (80), we obtain

$$\frac{d}{dt} \text{Term D} = \frac{d}{dt} E[\mu_t^X] = \int \frac{\delta E}{\delta \mu_t^X} \frac{d\mu_t^X}{dt} = \int \left\langle e^{\alpha_t - \gamma_t} \bar{v}_t(x), \nabla \frac{\delta E}{\delta \rho} [\mu_t^X] \right\rangle d\mu_t^X.$$

By assembling all terms and substituting into the limiting time derivative of (83), we obtain

$$\frac{d^+}{dt} \mathcal{L}_t \leq \dot{\beta}_t e^{\beta_t} (E[\mu_t^X] - E_*) - e^{\alpha_t + \beta_t} \int_{\mathbb{R}^{2d}} \left\langle x - y, \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\bar{\gamma}_t(x, y). \tag{86}$$

From the optimal scaling condition $\dot{\beta}_t \leq e^{\alpha_t}$, and using convexity of E , we obtain that

$$\frac{d^+}{dt} \mathcal{L}_t \leq e^{\alpha_t + \beta_t} \left[\underbrace{E[\mu_t^X] - E_* - \int_{\mathbb{R}^{2d}} \left\langle x - y, \nabla \frac{\delta E}{\delta \rho} [\mu_t^X](x) \right\rangle d\bar{\gamma}_t(x, y)}_{\leq 0} \right] \leq 0. \tag{87}$$

This completes the proof.

5.2 Time Dilation

We can show, using similar analysis to that of Wibisono et al. (2016), that the family of Hamiltonians defined by (43) is closed under time dilation. Given a smooth increasing function $\tau : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and a phase probability distribution $\mu_t(x, v)$, we consider the reparameterized curve $\tilde{\mu}_t = \mu_{\tau(t)}$. For clarity in the analysis of this section, we change the notation for the Hamiltonian of (43), denoting it by $H_{\alpha, \beta, \gamma}$ to emphasize its dependence on the time-dependent parameters α , β , and γ . We have the following result.

Theorem 11 *If μ_t satisfies the Hamiltonian flow equation (46) for the Hamiltonian $H_{\alpha,\beta,\gamma}$ of (43), then the reparameterized curve $\tilde{\mu}_t = \mu_{\tau(t)}$ satisfies the Hamiltonian flow equation for the rescaled Hamiltonian $H_{\tilde{\alpha},\tilde{\beta},\tilde{\gamma}}$, with the modified parameters defined as follows:*

$$\tilde{\alpha}_t = \alpha_{\tau(t)} + \log \dot{\tau}(t), \quad \tilde{\beta}_t = \beta_{\tau(t)}, \quad \tilde{\gamma}_t = \gamma_{\tau(t)}. \quad (88a)$$

Furthermore, α , β , and γ satisfy the optimal scaling conditions (44) if and only if $\tilde{\alpha}$, $\tilde{\beta}$, and $\tilde{\gamma}$ satisfy these conditions, respectively.

Proof Note that the time derivative of the reparameterized curve is given by $\partial_t \tilde{\mu}_t = \dot{\tau}(t) \partial_\tau \mu_\tau|_{\tau=\tau(t)}$. Thus, following (46), the following equation is satisfied by $\tilde{\mu}_t$:

$$\partial_t \tilde{\mu}_t = -\nabla_x \cdot \left(\mu_{\tau(t)} \dot{\tau}(t) e^{\alpha_{\tau(t)} - \gamma_{\tau(t)}} \right) + \nabla_v \cdot \left(\mu_{\tau(t)} \dot{\tau}(t) e^{\alpha_{\tau(t)} + \beta_{\tau(t)} + \gamma_{\tau(t)}} \nabla \frac{\delta E}{\delta \rho} \left[\mu_{\tau(t)}^X \right] \right). \quad (89)$$

With the definition (88) of the modified parameters $\tilde{\alpha}$, $\tilde{\beta}$, and $\tilde{\gamma}$, we can write this equation as

$$\partial_t \tilde{\mu}_t = -\nabla_x \cdot \left(\tilde{\mu}_t e^{\tilde{\alpha}_t - \tilde{\gamma}_t} \right) + \nabla_v \cdot \left(\tilde{\mu}_t e^{\tilde{\alpha}_t + \tilde{\beta}_t + \tilde{\gamma}_t} \nabla \frac{\delta E}{\delta \rho} \left[\tilde{\mu}_t^X \right] \right), \quad (90)$$

which is the Hamiltonian flow equation (46) for the rescaled Hamiltonian $H_{\tilde{\alpha},\tilde{\beta},\tilde{\gamma}}$. ■

The previous theorem is the analog of (Wibisono et al., 2016, Theorem 2.2). As observed in Wibisono et al. (2016), a major appeal of this theorem is that it links up a class of methods through time-dilation. In particular, set

$$\alpha_t = \log(2/t), \quad \beta_t = \log(t^2/4), \quad \gamma_t = 2 \log(t)$$

as the parameters for Nesterov acceleration, which achieves convergence of $E[\mu_t] - E_* = O(1/t^2)$. Then for any $p \geq 2$, setting $\tau(t) = t^{p/2}$, we have the accelerated rate $E[\tilde{\mu}_t] - E_* = O(1/t^p)$.

We note that in general, when we reparameterize time by a time-dilation function $\tau(t)$, the Hamiltonian functional transforms to $\tilde{H}_t[\mu] = \dot{\tau}(t) H_{\tau(t)}[\mu]$. Thus, the result of Theorem 11 can be written as

$$H_{\tilde{\alpha},\tilde{\beta},\tilde{\gamma}}(t) = \dot{\tau}(t) H_{\alpha,\beta,\gamma}(\tau(t)),$$

which can be verified by the definition of Hamiltonian (43) and the modified parameters (88).

6. Algorithms and Numerical Experiments

In this section, we report on numerical experiments with the Hamiltonian flows introduced above. In Section 6.1 we lay out the algorithms for running (HBF) and (VAF) using representative particles, while in Section 6.2, we showcase the application of the algorithms in two specific examples: potential energy and Bayesian sampling. We consider only continuous-time models in this section, deferring the development of discrete-in-time algorithms to future research.

6.1 Implementation of (HBF) and (VAF)

To study Hamiltonian flows, we find numerical solutions $\mu \in AC([0, \infty), \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d))$ of the equation

$$\partial_t \mu_t + \nabla_x \cdot \left(\mu_t \nabla_v \frac{\delta H_t}{\delta \mu} [\mu_t] \right) - \nabla_v \cdot \left(\mu_t \nabla_x \frac{\delta H_t}{\delta \mu} [\mu_t] \right) = 0. \quad (91)$$

Given that $(x, v) \in \mathbb{R}^d \times \mathbb{R}^d$, the classical numerical approach for computing this equation requires discretization over the domain $\mathbb{R}^d \times \mathbb{R}^d$, which is computationally prohibitive for nontrivial dimensions d . In this context, a Monte Carlo solver based on particle approximation can be more robust for higher values of d . We define an empirical distribution based on N particles (X^i, V^i) , as follows:

$$\mu_t \approx \bar{\mu}_t = \frac{1}{N} \sum_{i=1}^N \delta_{(X^i(t), V^i(t))}.$$

We denote by $\bar{\mu}_t^X = \frac{1}{N} \sum_{j=1}^N \delta_{X^j(t)}$ the x -marginal of the empirical measure. By substituting $\bar{\mu}_t$ into (91) and testing it on $\phi \in C_c^\infty(\mathbb{R}^d \times \mathbb{R}^d)$, we obtain that

$$\sum_{i=1}^N \nabla_x \phi(X^i, V^i) \left(\dot{X}^i - \nabla_v \frac{\delta H_t}{\delta \mu} [\bar{\mu}_t](X^i, V^i) \right) + \nabla_v \phi(X^i, V^i) \left(\dot{V}^i + \nabla_x \frac{\delta H_t}{\delta \mu} [\bar{\mu}_t](X^i, V^i) \right) = 0, \quad (92)$$

which suggests the following equations for evolution of the particles:

$$\dot{X}^i = \nabla_v \frac{\delta H_t}{\delta \mu} [\bar{\mu}_t](X^i, V^i), \quad \dot{V}^i = -\nabla_x \frac{\delta H_t}{\delta \mu} [\bar{\mu}_t](X^i, V^i), \quad i = 1, \dots, N. \quad (93)$$

By substituting the various definitions of H_t considered so far, we arrive at the following flows:

- Heavy-ball flow (HBF):

$$\dot{X}^i = V^i, \quad \dot{V}^i = -aV^i - \nabla_x \frac{\delta E}{\delta \rho} [\bar{\mu}_t^X](X^i), \quad i = 1, \dots, N. \quad (94)$$

- Variational-acceleration-flow (VAF) in its general form:

$$\dot{X}^i = V^i, \quad \dot{V}^i = -(\dot{\gamma}_t - \dot{\alpha}_t)V^i - e^{2\alpha_t + \beta_t} \nabla_x \frac{\delta E}{\delta \rho} [\bar{\mu}_t^X](X^i), \quad i = 1, \dots, N. \quad (95)$$

- Nesterov flow as an example of (VAF) using the coefficients (48):

$$\dot{X}^i = V^i, \quad \dot{V}^i = -\frac{3}{t}V^i - \nabla_x \frac{\delta E}{\delta \rho} [\bar{\mu}_t^X](X^i), \quad i = 1, \dots, N. \quad (96)$$

- Exponential convergence as an example of (VAF) using coefficients $[\alpha_t, \beta_t, \gamma_t] = [0, t, t]$:

$$\dot{X}^i = V^i, \quad \dot{V}^i = -V^i - e^t \nabla_x \frac{\delta E}{\delta \rho} [\bar{\mu}_t^X](X^i), \quad i = 1, \dots, N. \quad (97)$$

Note the similarity to deploying the particle method for solving the Wasserstein gradient flow (WGF). It is a standard technique to adopt the particle presentation:

$$\rho_t \approx \bar{\rho}_t = \frac{1}{N} \sum_{i=1}^N \delta_{X^i(t)}. \quad (98)$$

When this form is substituted into (WGF), we arrive at:

$$\dot{X}^i = -\nabla_x \frac{\delta E}{\delta \rho} [\bar{\rho}_t](X^i), \quad i = 1, \dots, N. \quad (99)$$

Remark 12 *The current paper focuses only on the convergence of the continuous-time dynamics. To make these approaches practical, we need to investigate the errors that arise in their numerical implementations. Our implementations make use of an adaptive solver (Section 6.2). We do not attempt an error analysis that is customized to the specific form of these differential equations. We next discuss the two sources of numerical error: time discretization and particle representation.*

- *Time discretization: When the time-step size h is small, the discrete solution should capture the fast decay rate of the continuous solution. In particular, the discretization error has to be compatible with the convergence rates of the PDE. This property is termed “rate-matching discretization” in Wibisono et al. (2016) in the Euclidean setting. Various techniques have been proposed, among which the symplectic integrator (Muehlebach and Jordan, 2021) appears to produce higher order accuracy. This observation is in line with the proposal presented in Ambrosio and Gangbo (2008), which does not discuss details. It would be interesting to adapt these techniques to the setting of probability measure space. However, we note that the nonlinear geometry (compared to Wibisono et al. (2016); Betancourt et al. (2018); Muehlebach and Jordan (2021) in Euclidean distance) makes the analysis significantly harder.*
- *Particle representation: It is widely known that Monte Carlo particle approximation suffers from the curse of dimensionality (Singh and Póczos, 2018; Niles-Weed and Berthet, 2022), since the number of particles needed to represent the underlying distribution scales exponentially in the dimension of the problem. However, the ultimate task is to find the minimizer of the energy functional. This is a “weak-sense” evaluation of the convergence, and there may still be a chance to achieve convergence without experiencing the curse of dimensionality. This intriguing possibility merits further investigation.*

6.2 Numerical Results

Here, we apply heavy-ball flow (94), Nesterov flow (96), and exponentially convergent variational acceleration flow (97) to three tasks. Example 1 involves minimization of a potential energy functional. Example 2 minimizes a KL divergence against a given target distribution, a problem from Bayesian sampling. Example 3 is the training of an infinitely-wide, single-layer neural network with ReLU activation.

In all these examples, the numerical integration over time is performed using the DIFFRAX library (Kidger, 2021). We use the Dormand-Prince 5/4 method with the default adaptive step size controller, setting both relative and absolute tolerances to 10^{-6} . The initial conditions of the particles are independently sampled from the standard Gaussian distribution: $X^i(0), V^i(0) \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$.

For strongly convex functions, the choice $a = 2\sqrt{m}$ used in the analysis is too small to produce optimal computational performance. Thus in all examples, heavy-ball flow is executed with $a = 0.5$.

Example 1: Potential Energy. We consider potential energy

$$E[\rho] = \mathcal{V}_\ell[\rho] = \int_{\mathbb{R}^d} V_\ell(x) \, d\rho, \quad \ell = 1, 2,$$

with two different forms of the potential functions:

$$V_1(x) = \frac{1}{2} \langle x - b, A(x - b) \rangle, \quad V_2(x) = h \log \left(\sum_{i=1}^M \exp \left(\frac{\langle w_i, x \rangle - q_i}{h} \right) \right). \quad (100)$$

For the potential \mathcal{V}_1 , we set spatial dimension to be $d = 500$, with $A \in \mathbb{R}^{500 \times 500}$ a random symmetric positive definite matrix and b is a random vector. The symmetric matrix is formed by setting $A = U^\top D U$ with D being a diagonal matrix and U being an orthogonal matrix. The diagonal elements in D are log-uniformly distributed between 10^{-5} and 1, and U is uniformly sampled from the Haar measure over the orthogonal group $O(d)$. The vector b is drawn from $\mathcal{N}(0, 100 \cdot I_d)$. This design ensures the objective functional \mathcal{V}_1 to be m -strongly convex with $m \approx 10^{-5}$. For potential \mathcal{V}_2 , we take $d = 200$ and choose $M = 1000$ and $h = 20$. Each $w_i \in \mathbb{R}^{200}$ is drawn from $\mathcal{N}(0, I_{200})$ and $q = (q_i) \in \mathbb{R}^{1000}$ is from $\mathcal{N}(0, I_{1000})$. Depending on the choice of w_i , \mathcal{V}_2 can be strongly convex, but its eigenvalues can be as small as one wishes. We use $N = 100$ particles for both \mathcal{V}_1 and \mathcal{V}_2 .

In both examples, $E[\rho_t]$ is estimated empirically from (98). To estimate the optimal E_* for \mathcal{V}_2 , we run all four methods for a long time and designate E_* the best value achieved over these four runs.

Figure 1 shows the decrease in optimality gap $E[\rho_t] - E_*$ over time for heavy-ball flow, Nesterov flow, exponentially convergent variational acceleration flow, and Wasserstein gradient flow. All three solvers demonstrate faster continuous-time convergence than the standard Wasserstein gradient flow, with exponentially convergent VAF being the fastest.




Figure 1: Optimality gap vs time for Wasserstein gradient flow (WGF), Heavy-Ball flow (HB), Nesterov flow (Nes) and exponentially convergent Hamiltonian flow (Exp), for the functionals \mathcal{V}_1 (left) and \mathcal{V}_2 (right). The functionals are evaluated at empirical measures; see (98).

In Figure 2 we show the total number of steps for the four methods, for different tolerance of the optimality gap. As the convergence tolerance is tightened (moving toward the right of the plot), the number of steps required grows. For any given problem, the total number of steps is an effective proxy for the actual computational cost, as it is proportional to the number of gradient evaluations. By comparing the two plots, we see that the Nesterov and Hamiltonian flows outperform the Wasserstein gradient flow for small tolerance.

Example 2: Bayesian Sampling. Next, we tackle the more challenging task of minimizing KL divergence between ρ and a target distribution ρ_* , defined by

$$E[\rho] = KL(\rho||\rho_*) = \int_{\mathbb{R}^d} \rho(x) \ln \frac{\rho(x)}{\rho_*(x)} dx = \int_{\mathbb{R}^d} \ln \rho(x) d\rho + \int_{\mathbb{R}^d} g(x) d\rho - \log C,$$

where the target measure has density $\rho_*(x) = C e^{-g(x)}$, with $C > 0$ being the normalizing constant. With a slight abuse of notation, we do not distinguish a probability measure from its Lebesgue density. We cannot apply the particle approximation directly to the KL divergence because the

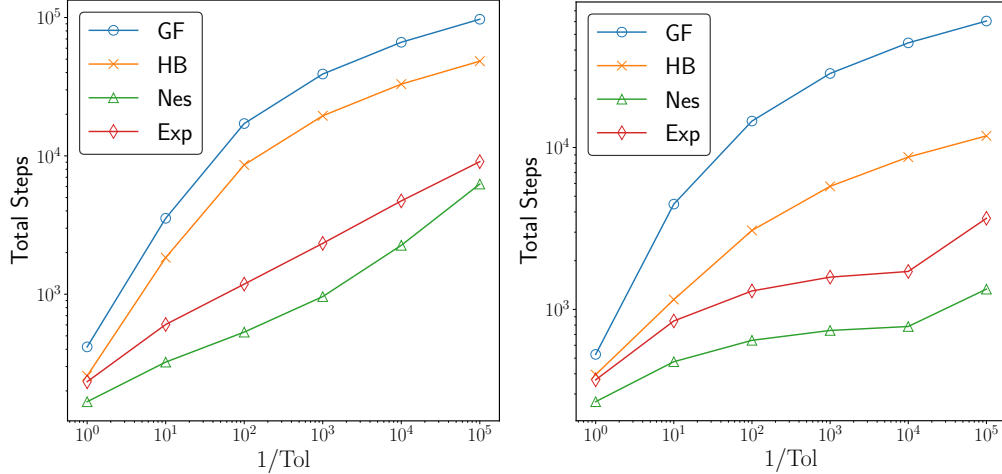


Figure 2: Total number of steps vs optimality gap (Tol) for Wasserstein gradient flow (WGF), Heavy-Ball flow (HB), Nesterov flow (Nes) and exponentially convergent Hamiltonian flow (Exp), for the functionals \mathcal{V}_1 (left) and \mathcal{V}_2 (right). The functionals are evaluated at empirical measures; see (98). The total number of steps includes those accepted and rejected in the adaptive step size controller.

empirical measure lacks a Lebesgue density. By analogy with the blob method for the Fokker-Planck equation (Carrillo et al., 2019b), we consider a regularized KL divergence

$$E^\varepsilon[\rho] = KL^\varepsilon(\rho||\rho_*) = \int_{\mathbb{R}^d} \ln K^\varepsilon * \rho \, d\rho + \int_{\mathbb{R}^d} g \, d\rho - \log C, \quad (101)$$

with the Gaussian convolution kernel $K^\varepsilon(x) = \frac{1}{(2\pi\varepsilon^2)^{d/2}} e^{-|x|^2/2\varepsilon^2}$ for some parameter $\varepsilon > 0$. This convolution allows E^ε being well-defined even for empirical measures. We note that E^ε , like E , is lower-bounded and is convex assuming an m -strongly convex $g(x)$, with m being sufficiently large, as outlined in Carrillo et al. (2019b).

We choose two different target measures ρ_* by specifying the log-density g in the same manner as the potential functions in (100), that is,

$$g_1(x) = \frac{1}{2} \langle x - b, A(x - b) \rangle, \quad g_2(x) = h \log \left(\sum_{i=1}^M \exp \left(\frac{\langle w_i, x \rangle - q_i}{h} \right) \right). \quad (102)$$

For g_1 , we take $d = 20$ so that $A \in \mathbb{R}^{20 \times 20}$ is a random positive definite symmetric matrix, and b is a random vector drawn from $\mathcal{N}(0, 10 \cdot I_{20})$. The symmetric matrix is formed by setting $A = U^\top D U$ with D being a diagonal matrix and U being an orthogonal matrix. The diagonal elements in D are log-uniformly distributed between 10^{-4} and 1, and U is uniformly sampled from the Haar measure over the orthogonal group $O(d)$. For g_2 , we set $d = 10$, $M = 200$, and $h = 10$. Each $w_i \in \mathbb{R}^{10}$ is drawn from $\mathcal{N}(0, I_{10})$, while $q = (q_i) \in \mathbb{R}^{200}$ is drawn from $\mathcal{N}(0, I_{200})$. In the numerical results below, we use $N = 1600$ particles. We choose $\varepsilon = 1$ (deferring the issue of choosing ε more optimally to future work).

To estimate the optimal E_*^ε , we run all four methods for a long time and designate E_*^ε the best value achieved over these four runs.

Figure 3 shows optimality gap as a function of t for heavy-ball flow, Nesterov flow, exponentially convergent variational acceleration flow, and Wasserstein gradient flow. In both examples, the

exponentially convergent Hamiltonian flow achieves the fastest convergence rate. It is noticeable that the Wasserstein gradient flow displays relatively lower errors during the initial stages, while the Hamiltonian flows exhibit slower decay due to oscillations. For larger t , the error of WGF saturates while the oscillations seen in the Hamiltonian flows taper off. Figure 4 shows the total number of steps as a function of the optimality gap. In both examples, the Nesterov and Hamiltonian flows outperforms the Wasserstein gradient flow for small tolerance. The performance of the Hamiltonian flows could be potentially enhanced by mitigating the oscillation through the incorporation of restarting strategies (O’donoghue and Candes, 2015; Su et al., 2014).




Figure 3: Optimality gap for minimization of regularized KL divergence with target g_1 (left) and g_2 (right) for four methods: Wasserstein gradient flow (WGF), Heavy-Ball flow (HB), Nesterov flow (Nes), and exponentially convergent Hamiltonian flow (Exp). The functionals are evaluated at empirical measures; see (98).




Figure 4: Total number of steps vs optimality gap (Tol) for minimization of regularized KL divergence with target g_1 (left) and g_2 (right) for four methods: Wasserstein gradient flow (WGF), Heavy-Ball flow (HB), Nesterov flow (Nes), and exponentially convergent Hamiltonian flow (Exp). The functionals are evaluated at empirical measures; see (98). The total number of steps includes those accepted and rejected in the adaptive step size controller.

Example 3: Neural network training. Our final example is related to the training of infinitely wide neural networks (Chizat and Bach, 2018; Mei et al., 2018; Sirignano and Spiliopoulos, 2020; Ding et al., 2021, 2022), where we have

$$E[\rho] = \frac{1}{2} \int_{\mathbb{R}^d} |f(x) - g(x, \rho)|^2 d\pi(x), \quad (103)$$

where π is a given distribution over the sampled data and $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is the target function. We take the function g to be a two-layer neural network: for every $x \in \mathbb{R}^d$ and $\rho \in \mathcal{P}(\mathbb{R}^{d+3})$

$$g(x, \rho) := \int_{\mathbb{R}^{d+3}} V(x, z) d\rho(z), \quad \text{with} \quad V(x, (\alpha, \beta, w, b)) = \alpha \sigma(w \cdot x + b) + \beta,$$

with σ being the ReLU function, which is positively 1-homogeneous, and $z = (\alpha, \beta, w, b) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}$. The functional $E[\rho]$ cannot be shown to be globally geodesic convex, but is locally geodesic convex, see Appendix C. We nevertheless tested the training with four methods (Wasserstein gradient flow (WGF), Heavy-Ball flow (HB), Nesterov flow (Nes), and exponentially convergent Hamiltonian flow (Exp)).

We set the spatial dimension to be $d = 1$ with the target function being $f(x) = \sin(\pi x)$. We choose the data distribution to be the uniform distribution over $[-1, 1]$ and 500 data points are sampled to evaluate the integration in π . In the numerical results below, we use $N = 200$ particles (neurons). The numerical results are presented in Figure 5. Note that the Nesterov and Hamiltonian flows outperform the Wasserstein gradient flow for small tolerances.




Figure 5: **Left:** Mean square errors for neural network training with target $f(x) = \sin(\pi x)$ for four methods: Wasserstein gradient flow (WGF), Heavy-Ball flow (HB), Nesterov flow (Nes), and exponentially convergent Hamiltonian flow (Exp). **Middle:** The target $f(x) = \sin(\pi x)$ and its neural network approximations obtained by running the four methods for $T = 14$. **Right:** The number of total steps as a function of the mean square error (Tol). The number of total steps includes those accepted and rejected in the adaptive step size controller.

Acknowledgments

Q.L. acknowledges support from Vilas Early Career award and ONR-N00014-21-1-214. The research of S.C. and Q.L. is supported in part by NSF-CAREER-1750488 and Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin Madison with funding from the Wisconsin Alumni Research Foundation. The work of S.C., Q.L., and S.W. is supported in part by NSF via grant DMS-2023239. S.W. also acknowledges support from NSF Award CCF-2224213

and AFOSR FA9550-21-1-0084. O.T. acknowledges support from NWO Vidi grant 016.Vidi.189.102 on *Dynamical-Variational Transport Costs and Application to Variational Evolutions*, and NWO grants OCENW.M.21.012 and NGF.1582.22.009.

Appendix A. Examples of Hamiltonian flows

We present several further examples of the Hamiltonian PDE.

Example 4: Kalman-Hamiltonian flow (Garbuno-Inigo et al., 2020; Wang and Li, 2022; Liu et al., 2022). Consider the Hamiltonian with weighted kinetic energy

$$H_t[\mu] = e^{-at} \frac{1}{2} \int_{\mathbb{R}^{2d}} \langle v, C^\lambda[\mu^X] v \rangle d\mu + e^{at} E[\mu^X], \quad (104)$$

where $a > 0$ and $C^\lambda[\nu] \in \mathbb{R}^{d \times d}$ is the covariance matrix defined by

$$C^\lambda[\nu] = \int_{\mathbb{R}^d} (x - \mathbb{E}_\nu[X]) \otimes (x - \mathbb{E}_\nu[X]) d\nu + \lambda I, \quad \lambda \geq 0. \quad (105)$$

The Kalman-Hamiltonian flow is then

$$\partial_t \mu_t + \nabla_x \cdot \left(\mu_t e^{-at} C^\lambda[\mu_t] v \right) - \nabla_v \cdot \left(\mu_t \left(e^{-at} \mathbb{E}_{\mu_t^V} [VV^\top] (x - \mathbb{E}_{\mu_t^X} [X]) + e^{at} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] \right) \right) = 0. \quad (106)$$

Particle dynamics are defined by

$$\begin{cases} \dot{X}_t &= e^{-at} C^\lambda[\mu_t^X] V_t \\ \dot{V}_t &= -e^{-at} \mathbb{E}_{\mu_t^V} [V_t V_t^\top] (X_t - \mathbb{E}_{\mu_t^X} [X_t]) - e^{at} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] (X_t). \end{cases} \quad (107)$$

Example 5: Stein-Hamiltonian flow (Duncan et al., 2023; Liu, 2017; Wang and Li, 2022). Consider the Hamiltonian with a kernel-weighted kinetic energy

$$H_t(\mu) = e^{-at} \int_{\mathbb{R}^{4d}} \frac{1}{2} w^\top K(x, y) v d\mu(x, v) d\mu(y, w) + e^{at} E[\mu^X], \quad (108)$$

where $K(x, y) \in \mathbb{R}^{d \times d}$ is a symmetric positive kernel function. The Stein-Hamiltonian flow is then

$$\begin{aligned} \partial_t \mu_t + \nabla_x \cdot \left(\mu_t e^{-at} \int_{\mathbb{R}^{2d}} K(x, y) w d\rho_t(y, w) \right) \\ - \nabla_v \cdot \left(\mu_t \left(e^{-at} \int_{\mathbb{R}^{2d}} \nabla_x [v^\top K(x, y) w] d\mu_t(y, w) + e^{at} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] \right) \right) = 0. \end{aligned} \quad (109)$$

Particle dynamics are defined by

$$\begin{cases} \dot{X}_t &= e^{-at} \int_{\mathbb{R}^{2d}} K(X_t, y) w d\mu_t(y, w) \\ \dot{V}_t &= -e^{-at} \int_{\mathbb{R}^{2d}} \nabla_x [V_t^\top K(X_t, y) w] d\mu_t(y, w) - e^{at} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] (X_t). \end{cases} \quad (110)$$

Example 6: Bregman-Hamiltonian flow (Wibisono et al., 2016; Wilson et al., 2016). The Bregman-Hamiltonian is defined by

$$H_t(\rho) = e^{\alpha t + \gamma t} \left(\int_{\mathbb{R}^{2d}} D_{\psi^*}(\nabla \psi(x) + e^{-\gamma t} v, \nabla \psi(x)) d\mu(x, v) + e^{\beta t} E[\mu^X] \right), \quad (111)$$

where $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex function of Legendre type (Rockafellar, 1997; Borwein and Lewis, 2006), and $\psi^* : \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the convex conjugate of ψ . (The Bregman divergence of a convex function ψ is defined by $D_\psi(y, x) := \psi(y) - \psi(x) - \langle \nabla \psi(x), y - x \rangle$, where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product on \mathbb{R}^d .) The Bregman-Hamiltonian flow is then

$$\begin{aligned} & \partial_t \mu_t + \nabla_x \cdot (\mu_t e^{\alpha_t} (\nabla \psi^*(\nabla \psi(x) + e^{-\gamma_t} v) - x)) \\ & - \nabla_v \cdot \left(\mu_t e^{\alpha_t + \gamma_t} \left[\nabla^2 \psi(x) (\nabla \psi^*(\nabla \psi(x) + e^{-\gamma_t} v) - x) - e^{-\gamma_t} v + e^{\beta_t} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X] \right] \right) = 0. \end{aligned} \quad (112)$$

Particle dynamics for this flow are defined by

$$\begin{aligned} \dot{X}_t &= e^{\alpha_t} (\nabla \psi^*(\nabla \psi(X_t) + e^{-\gamma_t} V_t) - X_t) \\ \dot{V}_t &= -e^{\alpha_t + \gamma_t} \left[\nabla^2 \psi(X_t) (\nabla \psi^*(\nabla \psi(X_t) + e^{-\gamma_t} V_t) - X_t) + e^{-\gamma_t} V_t - e^{\beta_t} \nabla_x \frac{\delta E}{\delta \rho} [\mu_t^X](X_t) \right]. \end{aligned} \quad (113)$$

We can define the mirror transform $M_t(x, v) = (x, z) = (x, \nabla \psi(x) + e^{-\gamma_t} v)$ and the pushforward measure $\nu_t = (M_t)_\# \mu_t \in \mathcal{P}_2(\mathbb{R}^{2d})$. Then $\nu_t^X = \mu_t^X$ and ν_t solves the equation

$$\partial_t \nu_t + \nabla_x \cdot (\nu_t e^{\alpha_t} (\nabla \psi^*(z) - x)) - \nabla_z \cdot \left(\nu_t e^{\alpha_t + \beta_t} \nabla_x \frac{\delta E}{\delta \rho} [\nu_t^X] \right) = 0, \quad (114)$$

where we have used $\dot{\gamma}_t = e^{\alpha_t}$ in the derivation. Under these transformations, the associated particle dynamics becomes

$$\begin{aligned} \dot{X}_t &= e^{\alpha_t} (\nabla \psi^*(Z_t) - X_t) \\ \dot{Z}_t &= -e^{\alpha_t + \beta_t} \nabla_x \frac{\delta E}{\delta \rho} [\nu_t^X](X_t). \end{aligned} \quad (115)$$

By choosing $\alpha_t = \log(r/t)$, $\beta_t = 2 \log(t/r)$ with $r > 0$, this method generalizes the accelerated mirror descent method (Krichene et al., 2015).

Appendix B. Rigorous statements and proofs for two properties in Section 4

Here we present rigorous statements regarding the time derivatives of the Wasserstein distance. The following theorem characterizes the first-order derivative of the Wasserstein distance.

Theorem 13 (Theorem 8.4.7 and Proposition 8.5.4 in Ambrosio et al. (2005)) *Let σ be a probability measure in $\mathcal{P}_2(\mathbb{R}^d)$ and $\nu \in C([0, \infty), \mathcal{P}_2(\mathbb{R}^d))$ be a solution to the continuity equation*

$$\partial_t \nu_t + \nabla \cdot (\nu_t \xi_t) = 0 \quad \text{in distribution,} \quad (116)$$

for locally Lipschitz vector fields ξ satisfying

$$\int_0^\infty \int_{\mathbb{R}^d} \|\xi_t\|^2 d\nu_t dt < \infty, \quad (117)$$

then $\nu \in \text{AC}([0, \infty), \mathcal{P}_2(\mathbb{R}^d))$ and for almost every $t \in (0, \infty)$, we have

$$\frac{1}{2} \frac{d}{dt} W_2^2(\nu_t, \sigma) = \int_{\mathbb{R}^{2d}} \langle x - y, \xi_t(x) \rangle d\gamma_t(x, y), \quad (118)$$

where $\gamma_t \in \Gamma_o(\nu_t, \sigma)$.

Before proceeding to characterize the second-order derivative of the Wasserstein distance, it is necessary to define the disintegration of the phase space probability measure $\mu \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ with respect to its x -marginal.

Theorem 14 (Theorem 5.3.1 in Ambrosio et al. (2005)) *Let $\mu \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$, and denote by $\mu^X \in \mathcal{P}(\mathbb{R}^d)$ the x -marginal distribution of μ . Then there exists a μ^X -a.e. uniquely determined family of Borel probability measure $\{\mu_x\}_{x \in \mathbb{R}^d} \in \mathcal{P}(\mathbb{R}^d)$ such that*

$$\int_{\mathbb{R}^{2d}} f(x, v) \, d\mu(x, v) = \int_{\mathbb{R}^d} \left(\int_{\mathbb{R}^d} f(x, v) \, d\mu_x(v) \right) d\mu^X(x),$$

for every Borel map $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, +\infty]$.

The second-order derivative of the W_2 -distance can be computed by the following theorem.

Theorem 15 (Modification of Theorem 1 in Carrillo et al. (2019a)) *Let $\sigma \in \mathcal{P}_2(\mathbb{R}^d)$ and $\mu \in AC([0, \infty), \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d))$ be a solution to the Hamiltonian flow*

$$\partial_t \mu_t + \nabla_x \cdot (\mu_t F(t, v)) - \nabla_v \cdot (\mu_t G(t, \mu_t, x)) = 0, \quad \forall t \in [0, T] \quad (119)$$

with locally-in- t and globally-in- (x, v) Lipschitz vector fields $(t, x, v) \mapsto F(t, v)$, $G(t, \mu_t, x)$ satisfying

$$t \mapsto \|F(t, v)\|_{L^2(\mu_t)}, \|\partial_t F(t, v)\|_{L^2(\mu_t)}, \|\nabla_v F(t, v)\|_{L^2(\mu_t)}, \|G(t, \mu_t, x)\|_{L^2(\mu_t)} \in C([0, \infty)) \cap L^2([0, \infty)). \quad (120)$$

then for any $T > 0$, the following inequality holds:

$$\begin{aligned} \frac{1}{2} \frac{d^+}{dt} \frac{d}{dt} W_2^2(\mu_t^X, \sigma) &\leq \int_{\mathbb{R}^{2d}} \|F(t, v)\|^2 d\mu_t(x, v) \\ &+ \int_{\mathbb{R}^{3d}} \langle x - y, \partial_t F(t, v) - \nabla_v F(t, v) G(t, \mu_t, x) \rangle d\mu_{t,x}(v) d\gamma_t(x, y), \end{aligned} \quad (121)$$

where d^+/dt denotes the upper derivative in almost every $t > 0$ and the Jacobian $\nabla_v F = (\partial_{v^j} F^i)_{ij}$. Here $\mu_{t,x}$ is the disintegration of μ_t with respect to its x -marginal μ_t^X .

Proof We start by defining the following flow: For fixed $t \in (0, \infty)$, let $\Phi_\tau = (\Phi_\tau^X, \Phi_\tau^V)$ satisfy the following equations:

$$\begin{aligned} \partial_\tau \Phi_\tau^X(x, v) &= F(t + \tau, \Phi_\tau^V(x, v)), \\ \partial_\tau \Phi_\tau^V(x, v) &= -G(t + \tau, \mu_{t+\tau}, \Phi_\tau^X(x, v)), \end{aligned} \quad (\Phi_0^X, \Phi_0^V) = (x, v), \text{ for } \mu_t\text{-a.e. } (x, v). \quad (122)$$

These formulas define the corresponding Lipschitz flow for $\tau \in (-t, \infty)$. Set

$$\mu_{t \pm h} = (\Phi_{\pm h})_\# \mu_t. \quad (123)$$

By defining $J(t, \mu, x, v) := \partial_t F(t, v) - \nabla_v F(t, v) G(t, \mu, x)$, one can compute

$$\partial_\tau^2 \Phi_\tau^X(x, v) = J(t + \tau, \mu_{t+\tau}, \Phi_\tau(x, v)). \quad (124)$$

Let $\gamma_t \in \Gamma_o(\mu_t^X, \sigma)$ be an optimal transport plan between μ_t^X and $\sigma \in \mathcal{P}_2(\mathbb{R}^d)$. To prove (121), we use an approximation argument based on finite differences. To this end, we evolve γ_t in time, using the map Φ_τ , such that it remains a transport plan between the x -marginal $\mu_{t+\tau}^X$ and the

measure σ . However, to execute the flow (122), both the initial coordinate x and the velocity v are required, whereas the velocity information is absent in γ_t . To address this issue, we define an extended coupling plan over \mathbb{R}^{3d} by appending the velocity information into γ_t . Specifically, we define the extended coupling plan γ_t^E as the probability measure over $\mathcal{P}(\mathbb{R}^{3d})$ that satisfies the equation:

$$\int \phi(x, v, y) d\gamma_t^E(x, v, y) = \int \phi(x, v, y) d\mu_{t,x}(v) d\gamma_t(x, y), \quad \forall \phi \in C_b(\mathbb{R}^{3d}).$$

To evolve the coupling plan γ_t^E , we define the map $Q_\tau : \mathbb{R}^{3d} \rightarrow \mathbb{R}^{2d}$ as follows:

$$Q_\tau(x, v, y) = (\Phi_\tau^X(x, v), y), \quad \tau \in (-t, \infty).$$

The evolved transport plan is then defined as the pushforward measure: $\gamma_t^\tau := (Q_\tau)_\# \gamma_t^E \in \mathcal{P}(\mathbb{R}^{2d})$.

We observe that γ_t^τ defines a transport plan between $\mu_{t+\tau}^X$ and σ . Indeed, for any $\phi^X \in C_b(\mathbb{R}^d)$, we integrate it against the plan γ_t^τ and obtain

$$\begin{aligned} \int \phi^X(x) d\gamma_t^\tau(x, v, y) &= \int \phi^X(\Phi_\tau^X(x, v)) d\mu_{t,x}(v) d\gamma_t(x, y) \\ &= \int \phi^X(\Phi_\tau^X(x, v)) d\mu_t(x, v) \\ &= \int \phi^X(x) d\mu_{t+\tau}^X(x), \end{aligned}$$

where in the first equality, we use the definition of γ_t^τ , and in the second equality we employ the definition of the transport plan γ_t and the disintegration $\mu_{t,x}$. Similarly, for any $\phi^Y \in C_b(\mathbb{R}^d)$, we integrate it against the plan γ_t^τ and obtain

$$\int \phi^Y(y) d\gamma_t^\tau(x, v, y) = \int \phi^Y(y) d\mu_{t,x}(v) d\gamma_t(x, y) = \int \phi^Y(y) d\sigma(y).$$

Thus, γ_t^τ is a transport plan between $\mu_{t+\tau}^X$ and σ for any $\tau \in (-t, \infty)$.

The rest of the proof follows the same strategy as that of Carrillo et al. (2019a, Theorem 5.3.1). We outline it here for completeness.

For some fixed $t \in (0, \infty)$ and $h \in (0, t)$, consider the finite difference

$$\Delta_h \mathcal{K}(\mu_t^X, \sigma) := (D_{h/2}^2 W_2^2)(\mu_t^X, \sigma) = \frac{1}{h^2} (W_2^2(\mu_{t+h}, \sigma) - 2W_2^2(\mu_t, \sigma) + W_2^2(\mu_{t-h}, \sigma)), \quad (125)$$

where D_τ denotes the symmetric difference operator with step $\tau > 0$, that is,

$$(D_\tau W_2^2)(\mu_t^X, \sigma) := \frac{1}{2\tau} (W_2^2(\mu_{t+\tau}^X, \sigma) - W_2^2(\mu_{t-\tau}^X, \sigma)). \quad (126)$$

Recalling that γ_t^τ is a coupling plan between $\mu_{t+\tau}^X$ and σ for any $\tau \in [-h, h]$, we obtain

$$W_2^2(\mu_{t\pm\tau}^X, \sigma) \leq \int_{\mathbb{R}^{2d}} \|x - y\|^2 d\gamma_t^{\pm\tau}(x, y) = \int_{\mathbb{R}^{3d}} \|\Phi_{\pm\tau}^X(x, v) - y\|^2 d\mu_{t,x}(v) d\gamma_t(x, y). \quad (127)$$

By making use of (127), for any $h \in (0, t)$, we have

$$\begin{aligned}
\Delta_h \mathcal{K}(\mu_t^X, \sigma) &\leq \frac{1}{h^2} \int_{\mathbb{R}^{3d}} \|\Phi_\tau^X(x, v) - y\|^2 - 2\|x - y\|^2 + \|\Phi_{-\tau}^X(x, v) - y\|^2 d\mu_{t,x}(v) d\gamma_t(x, y) \\
&\leq \frac{1}{h} \int_{-h}^h \int_{\mathbb{R}^{2d}} \|\partial_\tau \Phi_\tau^X(x, v)\|^2 d\mu_t(x, v) d\tau \\
&\quad + \frac{1}{h^2} \int_0^h \int_{\mathbb{R}^{3d}} (h - \tau) \langle x - y, \partial_\tau^2 \Phi_\tau^X(x, v) \rangle d\mu_{t,x}(v) d\gamma_t(x, y) d\tau \\
&\quad + \frac{1}{h^2} \int_{-h}^0 \int_{\mathbb{R}^{3d}} (h + \tau) \langle x - y, \partial_\tau^2 \Phi_\tau^X(x, v) \rangle d\mu_{t,x}(v) d\gamma_t(x, y) d\tau \\
&= \int_{-1}^1 \int_{\mathbb{R}^{2d}} \|F(t + sh, \Phi_{sh}^V(x, v))\|^2 d\mu_{t+sh} ds \\
&\quad + \int_0^1 (1 - s) \int_{\mathbb{R}^{3d}} \langle x - y, J(t + sh, \mu_{t+sh}, \Phi_{sh}(x, v)) \rangle d\mu_{t,x}(v) d\gamma_t(x, y) ds \\
&\quad + \int_{-1}^0 (1 + s) \int_{\mathbb{R}^{3d}} \langle x - y, J(t + sh, \mu_{t+sh}, \Phi_{sh}(x, v)) \rangle d\mu_{t,x}(v) d\gamma_t(x, y) ds,
\end{aligned} \tag{128}$$

where in the second inequality we use the fundamental theorem of calculus and Jensen's inequality, and in the last equality, we use (124).

For a fixed $T > 0$, we choose an integer $N > 0$ such that $h = T/N$. Let $\{\mu_{nh}\}_{n=0}^N$ be recursively defined by $\mu_{(n+1)h} = (\Phi_h^n)_\# \mu_{nh}$ for $n = 0, \dots, N$ where $\Phi_h^n = (\Phi_h^{n,X}, \Phi_h^{n,V})$ satisfies

$$\begin{cases} \partial_\tau \Phi_\tau^{n,X}(x, v) &= F(nh + \tau, \Phi_\tau^{n,V}(x, v)), \\ \partial_\tau \Phi_\tau^{n,V}(x, v) &= -G(nh + \tau, \mu_{nh+\tau}, \Phi_\tau^{n,X}(x, v)), \end{cases} \quad \Phi_0^n(x, v) = (x, v) \quad \text{for } \mu_{nh}\text{-a.e. } (x, v), \tag{129}$$

with initial condition $(\Phi_0^{n,X}, \Phi_0^{n,V}) = (x, v)$ and $\tau \in (-h, h)$. Then, for $n = 0, \dots, N$, (128) provides the inequality

$$\begin{aligned}
\Delta_h \mathcal{K}(\mu_{nh}^X, \sigma) &\leq \int_{-1}^1 \int_{\mathbb{R}^{2d}} \|F((n + s)h, \Phi_{sh}^V(x, v))\|^2 d\mu_{(n+s)h} ds \\
&\quad + \int_0^1 (1 - s) \int_{\mathbb{R}^{3d}} \langle x - y, J((n + s)h, \mu_{(n+s)h}, \Phi_{sh}(x, v)) \rangle d\mu_{nh,x}(v) d\gamma_{nh}(x, y) ds \\
&\quad + \int_{-1}^0 (1 + s) \int_{\mathbb{R}^{3d}} \langle x - y, J((n + s)h, \mu_{(n+s)h}, \Phi_{sh}(x, v)) \rangle d\mu_{nh,x}(v) d\gamma_{nh}(x, y) ds \\
&=: (A) + (B) + (C).
\end{aligned} \tag{130}$$

Multiplying the inequality with h and summing over $n = 1, \dots, N - 1$ yields for the LHS

$$\begin{aligned}
\sum_{n=1}^{N-1} h \Delta_h \mathcal{K}(\mu_{nh}^X, \sigma) &= (D_{h/2} W_2^2)(\mu_{(N-1/2)h}, \sigma) - (D_{h/2} W_2^2)(\mu_{h/2}, \sigma) \\
&= \frac{1}{h} \left(\int_{(N-1)h}^{Nh} - \int_0^h \right) \frac{d}{d\tau} W_2^2(\mu_\tau^X, \sigma) d\tau \\
&= 2 \int_0^1 \int_{\mathbb{R}^{3d}} \langle F((N - 1 + s)h, v), x - y \rangle d\mu_{(N-1+s)h,x} d\gamma_{(N-1+s)h}(x, y) ds \\
&\quad - 2 \int_0^1 \int_{\mathbb{R}^{3d}} \langle F(sh, v), x - y \rangle d\mu_{sh,x} d\gamma_{sh}(x, y) ds.
\end{aligned} \tag{131}$$

Passing to the limit $h \rightarrow 0$ with $Nh = T$ gives

$$\begin{aligned}
 \lim_{h \rightarrow 0} \sum_{n=1}^{N-1} h \Delta_h \mathcal{K}(\mu_{nh}, \sigma) &= 2 \int_{\mathbb{R}^{3d}} \langle F(T, v), x - y \rangle d\mu_{T,x} d\gamma_T(x, y) \\
 &\quad - 2 \int_{\mathbb{R}^{3d}} \langle F(0, v), x - y \rangle d\mu_{0,x} d\gamma_0(x, y) \\
 &= \frac{d}{dt} W_2^2(\mu_T^X, \sigma) - \frac{d}{dt} W_2^2(\mu_0^X, \sigma),
 \end{aligned} \tag{132}$$

which holds due to the dominated convergence theorem. On the other hand, the following convergences hold for the terms on the RHS of (130):

$$\begin{aligned}
 \sum_{n=1}^{N-1} hA &\longrightarrow 2 \int_0^T \int_{\mathbb{R}^{2d}} \|F(t, v)\|^2 d\mu_t dt \\
 \sum_{n=1}^{N-1} hB &\longrightarrow \frac{1}{2} \int_0^T \int_{\mathbb{R}^{3d}} \langle x - y, J(t, \mu_t, x, v) \rangle d\mu_{t,x}(v) d\gamma_t(x, y) dt \\
 \sum_{n=1}^{N-1} hC &\longrightarrow \frac{3}{2} \int_0^T \int_{\mathbb{R}^{3d}} \langle x - y, J(t, \mu_t, x, v) \rangle d\mu_{t,x}(v) d\gamma_t(x, y) dt
 \end{aligned} \tag{133}$$

by the definition of Riemann integrable functions and the assumed regularity (120). ■

Appendix C. Local convexity of neural network training

For the neural network architecture proposed above, the loss functional (103) is not geodesically convex over $\mathcal{P}(\mathbb{R}^d)$, but we claim it is locally convex along geodesics satisfying certain conditions (cf. (134) below). We discuss the argument explicitly in this section.

Consider the training of an infinitely wide 2-layer neural network with the loss functional

$$E[\rho] = \frac{1}{2} \iint_{\mathbb{R} \times \mathbb{R}^d} |y - g(x, \rho)|^2 d\pi(x, y),$$

where π is a given distribution over the sampled data and the function g is a two-layer neural network defined according to:

$$g(x, \rho) := \int_{\mathbb{R} \times \mathbb{R}^d} V(x, z) d\rho(z), \quad \text{with} \quad V(x, (\alpha, w)) = \alpha \sigma(w \cdot x),$$

for all $(x, \rho) \in \mathbb{R}^d \times \mathcal{P}(\mathbb{R} \times \mathbb{R}^d)$. Here x is the input to the neural network, ρ is the probability measure according to which the neuron weights are drawn, and σ is the positively 1-homogeneous ReLU function. We slightly modify the representation of V . Noting that for $\alpha \in \mathbb{R}$, $\alpha = \alpha 1_{\{\alpha > 0\}} - |\alpha| 1_{\{\alpha < 0\}}$, so we rewrite:

$$\begin{aligned}
 V(x, (\alpha, w)) &= \sigma(\alpha 1_{\{\alpha > 0\}} w \cdot x) - \sigma(|\alpha| 1_{\{\alpha < 0\}} w \cdot x) \\
 &=: \sigma(\omega_1 \cdot x) - \sigma(\omega_2 \cdot x) =: \widehat{V}(x, \omega),
 \end{aligned}$$

where we defined $\omega_1 = \alpha 1_{\{\alpha > 0\}} w$ and $\omega_2 = |\alpha| 1_{\{\alpha < 0\}} w$. This relation forms the definition:

$$\mathbb{R}^d \times \mathbb{R}^d \ni \omega = \mathbf{r}(\alpha, w) = (\omega_1, \omega_2),$$

and correspondingly, we define: $\hat{\rho} = r_{\#}\rho \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$. We thus obtain

$$\begin{aligned} g(x, \rho) &= \iint_{\mathbb{R} \times \mathbb{R}^d} V(x, (\alpha, w)) \, d\rho(\alpha, w) = \iint_{\mathbb{R} \times \mathbb{R}^d} [\sigma(r_1(\alpha, w) \cdot x) - \sigma(r_2(\alpha, w) \cdot x)] \, d\rho(\alpha, w) \\ &= \int_{\mathbb{R}^{2d}} [\sigma(\omega_1 \cdot x) - \sigma(\omega_2 \cdot x)] \, d\hat{\rho}(\omega) = \int_{\mathbb{R}^{2d}} \hat{V}(x, \omega) \, d\hat{\rho}(\omega) =: \hat{g}(x, \hat{\rho}). \end{aligned}$$

Consequently, we relax the training of ρ using $E[\rho]$ into a different problem: training $\hat{\rho}$ using the following functional:

$$E[\rho] = \frac{1}{2} \iint_{\mathbb{R} \times \mathbb{R}^d} |y - \hat{g}(x, \hat{\rho})|^2 \, d\pi(x, y) =: \hat{E}[\hat{\rho}].$$

We now show that \hat{E} is *locally* geodesically convex when σ is the ReLU function. To show geodesically convexity at a probability measure $\hat{\rho}$ amounts to show the objective functional is convex along any geodesics whose origin is at $\hat{\rho}$. To do so, we set $\hat{\eta}$ to be any probability measure in $\mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$, and denote by \mathbb{T} the optimal transport map between $\hat{\rho}$ and $\hat{\eta}$. Then, along the geodesics with a constant speed (McCann, 1997), $\hat{\gamma}_t := [(1-t)id + t\mathbb{T}]_{\#}\hat{\rho}$, for $t \in [0, 1]$, we have

$$\hat{g}(x, \hat{\gamma}_t) = \int_{\mathbb{R}^{2d}} \hat{V}(x, \omega) \, d\hat{\gamma}_t(\omega) = \int_{\mathbb{R}^{2d}} \hat{V}(x, (1-t)\omega + t\mathbb{T}(\omega)) \, d\hat{\rho}(\omega).$$

If we can successfully rewrite as

$$\begin{aligned} \hat{g}(x, \hat{\gamma}_t) &= (1-t) \int_{\mathbb{R}^{2d}} \hat{V}(x, \omega) \, d\rho(\omega) + t \int_{\mathbb{R}^{2d}} \hat{V}(x, \mathbb{T}(\omega)) \, d\hat{\rho}(\omega) \\ &= (1-t) \int_{\mathbb{R}^{2d}} \hat{V}(x, \omega) \, d\rho(\omega) + t \int_{\mathbb{R}^{2d}} \hat{V}(x, \omega') \, d\eta(\omega') \\ &= (1-t)\hat{g}(x, \hat{\rho}) + t\hat{g}(x, \hat{\eta}), \end{aligned} \tag{134}$$

we are showing $\hat{\rho} \mapsto \hat{g}(x, \hat{\rho})$ is geodesically linear for every $x \in \mathbb{R}^d$. Owing to the convexity of $r \mapsto |r|^2$, we obtain

$$\begin{aligned} \hat{E}[\hat{\gamma}_t] &= \frac{1}{2} \iint_{\mathbb{R} \times \mathbb{R}^d} |y - \hat{g}(x, \hat{\gamma}_t)|^2 \, d\pi(x, y) = \frac{1}{2} \iint_{\mathbb{R} \times \mathbb{R}^d} |(1-t)(y - \hat{g}(x, \hat{\rho})) + t(y - \hat{g}(x, \hat{\eta}))|^2 \, d\pi(x, y) \\ &\leq (1-t)\frac{1}{2} \iint_{\mathbb{R} \times \mathbb{R}^d} |y - \hat{g}(x, \hat{\rho})|^2 \, d\pi(x, y) + t\frac{1}{2} \iint_{\mathbb{R} \times \mathbb{R}^d} |y - \hat{g}(x, \hat{\eta})|^2 \, d\pi(x, y) = (1-t)\hat{E}[\hat{\rho}] + t\hat{E}[\hat{\eta}], \end{aligned}$$

thus implying the (local) geodesic convexity of \hat{E} when (134) holds.

Generally speaking, however, (134) does not hold for any given $\hat{\eta}$. It *would* hold if the optimal transport map \mathbb{T} between $\hat{\rho}$ and $\hat{\eta}$ were to satisfy

$$\text{sign}(\omega_i \cdot x) = \text{sign}(\mathbb{T}_i(\omega) \cdot x), \quad \text{for } \hat{\rho}\text{-almost every } \omega,$$

with the convention that $\text{sign}(0) = 0$. When this happens, we would have

$$\text{sign}(((1-t)\omega_i \cdot x + t\mathbb{T}_i(\omega) \cdot x)) = \text{sign}(\omega_i \cdot x) = \text{sign}(\mathbb{T}_i(\omega) \cdot x), \quad \text{for every } t \in [0, 1],$$

and noticing $\sigma(\alpha + \beta) = \sigma(\alpha) + \sigma(\beta)$ for $\alpha, \beta \geq 0$, we can split

$$\hat{V}(x, \omega) = \hat{V}(x, (1-t)\omega + t\mathbb{T}(\omega)) = (1-t)\hat{V}(x, \omega) + t\hat{V}(x, \mathbb{T}(\omega)),$$

ensuring (134).

References

- Zeyuan Allen-Zhu and Lorenzo Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. In *8th Innovations in Theoretical Computer Science Conference, ITCS*, volume 67 of *LIPICs*, pages 3:1–3:22, 2017.
- Luigi Ambrosio and Wilfrid Gangbo. Hamiltonian ODEs in the Wasserstein space of probability measures. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 61(1):18–53, 2008.
- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.
- Hedy Attouch and Felipe Alvarez. The heavy ball with friction dynamical system for convex constrained minimization problems. In *Optimization: Proceedings of the 9th Belgian-French-German Conference on Optimization Namur, September 7–11, 1998*, pages 25–35. Springer, 2000.
- Hedy Attouch and Alexandre Cabot. Asymptotic stabilization of inertial gradient dynamics with time-dependent viscosity. *Journal of Differential Equations*, 263(9):5412–5458, 2017.
- Hedy Attouch, Zaki Chbani, Juan Peypouquet, and Patrick Redont. Fast convergence of inertial dynamics and algorithms with asymptotic vanishing viscosity. *Mathematical Programming*, 168: 123–175, 2018.
- Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. A neural probabilistic language model. *Advances in neural information processing systems*, 13, 2000.
- Michael Betancourt, Michael I Jordan, and Ashia C Wilson. On symplectic optimization. *arXiv preprint arXiv:1802.03653*, 2018.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.
- Jonathan Borwein and Adrian Lewis. *Convex Analysis*. Springer, 2006.
- Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.
- Alexandre Cabot, Hans Engler, and Sébastien Gadat. On the long time behavior of second order differential equations with asymptotically small dissipation. *Transactions of the American Mathematical Society*, 361(11):5983–6017, 2009.
- Yu Cao, Jianfeng Lu, and Lihan Wang. On explicit L^2 -convergence rate estimate for underdamped Langevin dynamics. *Archive for Rational Mechanics and Analysis*, 247(5):90, 2023.
- José A Carrillo, Young-Pil Choi, and Oliver Tse. Convergence to equilibrium in Wasserstein distance for damped Euler equations with interaction forces. *Communications in Mathematical Physics*, 365:329–361, 2019a.
- José Antonio Carrillo, Katy Craig, and Francesco S Patacchini. A blob method for diffusion. *Calculus of Variations and Partial Differential Equations*, 58:1–53, 2019b.
- Xiang Cheng and Peter Bartlett. Convergence of Langevin MCMC in KL-divergence. In *Algorithmic Learning Theory*, pages 186–211. PMLR, 2018.

- Xiang Cheng, Niladri S Chatterji, Peter L Bartlett, and Michael I Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. In *Conference on learning theory*, pages 300–323. PMLR, 2018.
- Lenaïc Chizat and Francis Bach. On the global convergence of gradient descent for overparameterized models using optimal transport. *Advances in neural information processing systems*, 31, 2018.
- Shui-Nee Chow, Wuchen Li, and Haomin Zhou. Wasserstein Hamiltonian flows. *Journal of Differential Equations*, 268(3):1205–1219, 2020.
- Jelena Diakonikolas and Michael I Jordan. Generalized momentum-based methods: A Hamiltonian perspective. *SIAM Journal on Optimization*, 31(1):915–944, 2021.
- Zhiyan Ding, Shi Chen, Qin Li, and Stephen J. Wright. On the global convergence of gradient descent for multi-layer ResNets in the mean-field regime. *arXiv preprint arXiv:2110.02926*, 2021.
- Zhiyan Ding, Shi Chen, Qin Li, and Stephen J. Wright. Overparameterization of deep ResNet: zero loss and mean-field analysis. *Journal of machine learning research*, 2022.
- R. J. DiPerna and P. L. Lions. On the cauchy problem for boltzmann equations: Global existence and weak stability. *Annals of Mathematics*, 130(2):321–366, 1989.
- Andrew B. Duncan, Nikolas Nüsken, and Lukasz Szpruch. On the geometry of Stein variational gradient descent. *Journal of machine learning research*, 24:56:1–56:39, 2023.
- Raaz Dwivedi, Yuansi Chen, Martin J Wainwright, and Bin Yu. Log-concave sampling: Metropolis-Hastings algorithms are fast! In *Conference on learning theory*, pages 793–797. PMLR, 2018.
- Alexandre d’Aspremont, Damien Scieur, and Adrien Taylor. Acceleration methods. *Foundations and Trends® in Optimization*, 5(1-2):1–245, 2021.
- Guilherme França, Jeremias Sulam, Daniel Robinson, and René Vidal. Conformal symplectic and relativistic optimization. *Advances in Neural Information Processing Systems*, 33:16916–16926, 2020.
- Alfredo Garbuno-Inigo, Franca Hoffmann, Wuchen Li, and Andrew M Stuart. Interacting Langevin diffusions: Gradient structure and ensemble Kalman sampler. *SIAM Journal on Applied Dynamical Systems*, 19(1):412–441, 2020.
- Nicolás García Trillos and Javier Morales. Semi-discrete optimization through semi-discrete optimal transport: a framework for neural architecture search. *Journal of Nonlinear Science*, 32(3):27, 2022.
- Tomas Geffner and Justin Domke. Langevin diffusion variational inference. In *International Conference on Artificial Intelligence and Statistics*, pages 576–593. PMLR, 2023.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 2013.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37:183–233, 1999.
- Patrick Kidger. *On Neural Differential Equations*. PhD thesis, University of Oxford, 2021.
- Diederik P. Kingma and Max Welling. Auto-Encoding variational Bayes. In *2nd International Conference on Learning Representations, ICLR*, 2014.
- Walid Krichene, Alexandre Bayen, and Peter L Bartlett. Accelerated mirror descent in continuous and discrete time. *Advances in neural information processing systems*, 28, 2015.
- Marc Lambert, Sinho Chewi, Francis R. Bach, Silvère Bonnabel, and Philippe Rigollet. Variational inference via Wasserstein gradient flows. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems, 2022*.
- Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- Chang Liu, Jingwei Zhuo, Pengyu Cheng, Ruiyi Zhang, and Jun Zhu. Understanding and accelerating particle-based variational inference. In *International Conference on Machine Learning*, pages 4082–4092. PMLR, 2019.
- Qiang Liu. Stein variational gradient descent as gradient flow. *Advances in neural information processing systems*, 30, 2017.
- Ziming Liu, Andrew M Stuart, and Yixuan Wang. Second order ensemble Langevin method for sampling and inverse problems. *arXiv preprint arXiv:2208.04506*, 2022.
- Yulong Lu, Jianfeng Lu, and James Nolen. Accelerating Langevin sampling with birth-death. *arXiv preprint arXiv:1905.09863*, 2019.
- Yi-An Ma, Niladri S. Chatterji, Xiang Cheng, Nicolas Flammarion, Peter L. Bartlett, and Michael I. Jordan. Is there an analog of Nesterov acceleration for gradient-based MCMC? *Bernoulli*, 27(3): 1942 – 1992, 2021.
- Chris J Maddison, Daniel Paulin, Yee Whye Teh, Brendan O’Donoghue, and Arnaud Doucet. Hamiltonian descent methods. *arXiv preprint arXiv:1809.05042*, 2018.
- Robert J McCann. A convexity principle for interacting gases. *Advances in mathematics*, 128(1): 153–179, 1997.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Céline Moucer, Adrien B. Taylor, and Francis R. Bach. A systematic approach to Lyapunov analyses of continuous-time models in convex optimization. *SIAM J. Optim.*, 33(3):1558–1586, 2023.
- Michael Muehlebach and Michael Jordan. A dynamical systems perspective on Nesterov acceleration. In *International Conference on Machine Learning*, pages 4656–4662. PMLR, 2019.

- Michael Muehlebach and Michael I Jordan. Optimization with momentum: Dynamical, control-theoretic, and symplectic perspectives. *Journal of Machine Learning Research*, 22(73):1–50, 2021.
- Arkadij Semenovič Nemirovskij and David Borisovich Yudin. *Problem complexity and method efficiency in optimization*. Wiley-Interscience, 1983.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- Yurii Evgen’evich Nesterov. A method of solving a convex programming problem with convergence rate $O(k^2)$. In *Doklady Akademii Nauk*, volume 269 (3), pages 543–547. Russian Academy of Sciences, 1983.
- Jonathan Niles-Weed and Quentin Berthet. Minimax estimation of smooth densities in Wasserstein distance. *The Annals of Statistics*, 50(3):1519–1540, 2022.
- Felix Otto. The geometry of dissipative evolution equations: The porous medium equation. *Communications in Partial Differential Equations*, 26(1-2):101–174, 2001.
- Brendan O’donoghue and Emmanuel Candes. Adaptive restart for accelerated gradient schemes. *Foundations of computational mathematics*, 15:715–732, 2015.
- Jan Peters, Katharina Mulling, and Yasemin Altun. Relative entropy policy search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24 (1), pages 1607–1612, 2010.
- Boris Polyak and Pavel Shcherbakov. Lyapunov functions: An optimization theory perspective. *IFAC-PapersOnLine*, 50(1):7456–7461, 2017.
- Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR computational mathematics and mathematical physics*, 4(5):1–17, 1964.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- R Tyrrell Rockafellar. *Convex Analysis*, volume 11. Princeton university press, 1997.
- Damien Scieur, Vincent Roulet, Francis Bach, and Alexandre d’Aspremont. Integration methods and optimization algorithms. *Advances in Neural Information Processing Systems*, 30, 2017.
- Ruoqi Shen and Yin Tat Lee. The randomized midpoint method for log-concave sampling. *Advances in Neural Information Processing Systems*, 32, 2019.
- Bin Shi, Simon S Du, Michael I Jordan, and Weijie J Su. Understanding the acceleration phenomenon via high-resolution differential equations. *Mathematical Programming*, pages 1–70, 2021.
- Shashank Singh and Barnabás Póczos. Minimax distribution estimation in Wasserstein distance. *arXiv preprint arXiv:1802.08855*, 2018.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752, 2020.

- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *9th International Conference on Learning Representations, ICLR*, 2021.
- Weijie Su, Stephen Boyd, and Emmanuel Candes. A differential equation for modeling Nesterov’s accelerated gradient method: theory and insights. *Advances in neural information processing systems*, 27, 2014.
- Amirhossein Taghvaei and Prashant Mehta. Accelerated flow for probability distributions. In *International Conference on Machine Learning*, pages 6076–6085. PMLR, 2019.
- Ken’ichiro Tanaka. Accelerated gradient descent method for functionals of probability measures by new convexity and smoothness based on transport maps. *arXiv preprint arXiv:2305.05127*, 2023.
- Marc Toussaint. Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning*, pages 1049–1056, 2009.
- Cédric Villani et al. *Optimal Transport: Old and New*, volume 338. Springer, 2009.
- Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- Yifei Wang and Wuchen Li. Accelerated information gradient flow. *Journal of Scientific Computing*, 90:1–47, 2022.
- Andre Wibisono and Ashia C Wilson. On accelerated methods in optimization. *arXiv preprint arXiv:1509.03616*, 2015.
- Andre Wibisono, Ashia C Wilson, and Michael I Jordan. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, 113(47):E7351–E7358, 2016.
- Ashia C Wilson, Benjamin Recht, and Michael I Jordan. A Lyapunov analysis of momentum methods in optimization. *arXiv preprint arXiv:1611.02635*, 2016.
- Jingzhao Zhang, Aryan Mokhtari, Suvrit Sra, and Ali Jadbabaie. Direct Runge-Kutta discretization achieves acceleration. *Advances in neural information processing systems*, 31, 2018.
- Shunshi Zhang, Sinho Chewi, Mufan Li, Krishna Balasubramanian, and Murat A Erdogdu. Improved discretization analysis for underdamped Langevin Monte Carlo. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 36–71. PMLR, 2023.
- Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.