A BLOB METHOD FOR INHOMOGENEOUS DIFFUSION WITH APPLICATIONS TO MULTI-AGENT CONTROL AND SAMPLING

KATY CRAIG, KARTHIK ELAMVAZHUTHI, MATT HABERLAND, AND OLGA TURANOVA

ABSTRACT. As a counterpoint to classical stochastic particle methods for linear diffusion equations, such as Langevin dynamics for the Fokker-Planck equation, we develop a deterministic particle method for the weighted porous medium equation and prove its convergence on bounded time intervals. This generalizes related work on blob methods for unweighted porous medium equations. From a numerical analysis perspective, our method has several advantages: it is meshfree, preserves the gradient flow structure of the underlying PDE, converges in arbitrary dimension, and captures the correct asymptotic behavior in simulations.

The fact that our method succeeds in capturing the long time behavior of the weighted porous medium equation is significant from the perspective of related problems in quantization. Just as the Fokker-Planck equation provides a way to quantize a probability measure $\bar{\rho}$ by evolving an empirical measure $\rho^N(t) = \frac{1}{N} \sum_{i=1}^{N} \delta_{X^i(t)}$ according to stochastic Langevin dynamics so that $\rho^N(t)$ flows toward $\bar{\rho}$, our particle method provides a way to quantize $\bar{\rho}$ according to deterministic particle dynamics approximating the weighted porous medium equation. In this way, our method has natural applications to multi-agent coverage algorithms and sampling probability measures.

A specific case of our method corresponds exactly to the mean-field dynamics of training a two-layer neural network for a radial basis function activation function. From this perspective, our convergence result shows that, in the over parametrized regime and as the variance of the radial basis functions goes to zero, the continuum limit is given by the weighted porous medium equation. This generalizes previous results, which considered the case of a uniform data distribution, to the more general inhomogeneous setting. As a consequence of our convergence result, we identify conditions on the target function and data distribution for which convexity of the energy landscape emerges in the continuum limit.

CONTENTS

1.	Introduction	1
2.	Preliminaries	11
3.	Gradient flows of energies with regularization and confiment	16
4.	An H^1 bound on the mollified gradient flow of $\mathcal{F}_{\varepsilon,k}$	21
5.	Convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ to \mathcal{F}_k	26
6.	Convergence of the gradient flows of \mathcal{F}_k to \mathcal{F}	32
7.	Numerical Simulation	37
Appendices		45
References		

1. INTRODUCTION

Quantization is a fundamental problem throughout the sciences, in which one seeks to approximate a continuum distribution or signal by discrete objects [43]. Mathematically, the quantization problem may be modeled by fixing a *target probability measure* $\bar{\rho}$ on a subset Ω of \mathbb{R}^d and seeking locations $\{X^i\}_{i=1}^N$ in Ω so that the empirical measure $\rho^N = \frac{1}{N} \sum_{i=1}^N \delta_{X^i}$ approximates $\bar{\rho}$ in an appropriate sense. In statistics, this

March 2, 2022.

²⁰²⁰ Mathematics Subject Classification. Primary 35Q35, 35Q62, 35Q68, 35Q82, 65M12, 82C22, 93A16.

The work of K. Craig is supported by NSF DMS grant 1811012 and a Hellman Faculty Fellowship. K. Craig and O. Turanova gratefully acknowledge the support from the Simons Center for Theory of Computing, at which part of this work was completed. The work of K. Elamvazhuthi is supported by AFOSR grants FA9550-18-1-0502 and FA9550-18-1-0502.

The work of O. Turanova is supported by NSF DMS grant 1907221.

problem arises in the context of sampling, since the locations $\{X^i\}_{i=1}^N$ represent approximate samples drawn from $\bar{\rho}$. In control theory, this problem is relevant to multi-agent coverage algorithms [13,28], in which one seeks to control a fleet of robots to evolve from their current locations $\{X_0^i\}_{i=1}^N$ to terminal locations $\{X^i\}_{i=1}^N$ distributed according to $\bar{\rho}$.

There is a vast literature on different approaches to the quantization problem, arising from the many different criteria by which $\rho^N = \frac{1}{N} \sum_{i=1}^N \delta_{X^i}$ is considered a "good" approximation of $\bar{\rho}$. For example, if one seeks ρ^N to approximate $\bar{\rho}$ optimally in the Wasserstein metric of optimal transport, recent work has shown that this is closely related to the well-known Lloyd's algorithm and has a fascinating connection to weighted fast diffusion equations [11,12,17,46,47,60]. In the statistics literature, developing efficient sampling methods and quantifying their convergence is an active area of research, from classical methods based on Langevin dynamics to more recent developments, such as Hamiltonian Monte Carlo or Stein Variational Gradient Descent [10,54,77]. In the control theory literature, recent work has developed multiagent coverage algorithms based on stochastic and kernelized particle methods for linear diffusions, as well as theoretically explored the potential of nonlinear diffusions for the coverage task, via finite volume and graph-based methods [37–39,51,61]. Other authors have explored the role of different notions of optimality in designing coverage algorithms [5,6].

In each of these applications, quantization methods based on partial differential equations play an important role. A classical approach is given by evolving the locations of the particles by *Langevin dynamics*,

$$\begin{cases} dX_t^i = \sqrt{2}dB_t^i - \nabla \log \bar{\rho}(X_t^i)dt, \\ X^i(0) = X_0^i, \end{cases}$$

which is the stochastic particle discretization of the Fokker-Planck equation,

(FP)
$$\begin{cases} \partial_t \rho = \Delta \rho - \nabla \cdot (\rho \nabla \log \bar{\rho}) \\ \rho(0) = \rho_0. \end{cases}$$

In the present work, we continue in this line of PDE-principled methods for sampling and coverage algorithms. We introduce a new method based on the *weighted porous medium equation* (WPME). Given a bounded, convex domain $\Omega \subseteq \mathbb{R}^d$, a log-concave, strictly positive target $\bar{\rho} : \mathbb{R}^d \to \mathbb{R}$ with $\int_{\Omega} \bar{\rho} = 1$, and a fixed external potential $V \in C^2(\Omega)$, we consider the equation,

(WPME)
$$\begin{cases} \partial_t \rho = \nabla \cdot \left(\frac{\bar{\rho}}{2} \nabla \left(\frac{\rho^2}{\bar{\rho}^2}\right)\right) + \nabla \cdot (\nabla V \rho),\\ \rho(0) = \rho_0, \end{cases}$$

with no-flux boundary conditions on $\partial\Omega$. The initial conditions are chosen to satisfy $\rho_0 \ge 0$ and $\int_{\Omega} \rho_0 = 1$. (See Proposition 3.10 for the definition of weak solution.)

The dynamics of (WPME) arise in connection to quantization since, for V = 0, solutions of (WPME) converge as $t \to +\infty$ to $\bar{\rho}$ on Ω in the *Wasserstein metric*; see Proposition 3.14. Consequently, if one can approximate solutions $\rho(t)$ of (WPME) by an empirical measure $\rho^N(t) = \frac{1}{N} \sum_{i=1}^N \delta_{X^i(t)}$, this naturally leads to a method for flowing the empirical measure toward $\bar{\rho}$ on Ω in the long time limit.

The main goal of the present work is to develop a *deterministic particle method* for (WPME), constructing an empirical measure $\rho^N(t) = \frac{1}{N} \sum_{i=1}^N \delta_{X^i(t)}$ and a system of ordinary differential equations to govern the locations of the particles $X^i(t)$ so that $\rho^N(t)$ indeed converges, as $N \to +\infty$, to a solution $\rho(t)$ of (WPME) on bounded time intervals. In Sections 1.4-1.5 below, we describe the specific assumptions we impose and the precise statements of our results, including which of our results continue to hold for $\bar{\rho}$ not log-concave, on unbounded domains Ω , and for less regular V.

On one hand, (WPME) is of interest outside the context of quantization. Weighted porous medium equations arise throughout the sciences, from models of fluid flow to biological swarming [44,74]. From this perspective, Theorem 1.4 of the present work provides a new numerical method for simulating these phenomena. In particular, our work extends the *blob method* for the porous medium equation ($\bar{\rho} = 1$), which has been studied by Oelschläger [62], Lions and MasGallic [52], Carrillo, Craig, and Patacchini [20], and Burger and Esposito [15], to the case of weighted porous medium equations. (See below for a more detailed discussion of the relation with these results.) This provides a provably convergent numerical method for (WPME) in arbitrary dimensions, contributing to the substantial literature on numerical methods for such equations, including classical finite volume, finite element, and discontinuous Galerkin methods [8,14,19,71], as well as methods based on alternative deterministic particle methods in one spatial dimension [18,24,32,33, 58], Lagrangian evolution of the transport map along the flow [23,25,40,57,80], and many others [7,21,22,42]. From a numerical analysis perspective, the key benefits of our approach are that it is meshfree, deterministic, preserves the gradient flow structure and asymptotic behavior, and converges in arbitrary dimension.

On the other hand, we believe (WPME) is particularly interesting from the perspective of quantization for several reasons. First, as we describe below, there is a strong analogy between (WPME) and (FP), so that a quantization method based on (WPME) provides a counterpoint to classical Langevin dynamics.

A second reason for studying (WPME) in connection with quantization comes from applications in sampling. Over the past five years, Stein Variational Gradient Decent, originally introduced by Liu and Wang [54], has attracted attention in the statistics community as a novel method for sampling a target measure $\bar{\rho}$ via a deterministic interacting particle system, which has a formal Wasserstein gradient flow structure with respect to a *convex* mobility [50, 53, 55]. Recent work by Chewi et al. [26] identified that, when V = 0, Stein Variational Gradient Descent (SVGD) may be interpreted as a kernelized version of (WPME), which has a rigorous Wasserstein gradient flow structure, as we explain below. In this way, understanding properties of (WPME) and its discretizations has the potential shed light on behavior of SVGD more generally.

A third reason for interest in (WPME) from a quantization perspective comes from applications in control theory. This is due to the fact the particle method we succeed in developing for (WPME) is *deterministic*, an important attribute in the context of coverage algorithms, since the results of the algorithm wouldn't need to be averaged over many runs, and there is hope that future research could lead to quantitative convergence guarantees. This is in contrast to the case of classical quantization methods based on (FP), for which the natural Langevin particle approximation is stochastic.

A final reason for interest in (WPME) comes from a variant of the quantization problem arising in models of two-layer neural networks. As we will explain below, the particle method we develop to approximate solutions of (WPME) coincides exactly with the dynamics for training a two-layer neural network with a radial basis function activation function. In this way, our convergence result sheds light on the continuum limit of two-layer neural networks, showing that they converge to a solution of (WPME); see Corollary 1.6. This generalizes the previous convergence result of Javanmard, Mondelli, and Montanari [48] to the case of nonuniform data distributions. As a consequence of this result, we are able provide conditions on the target function and data distribution that guarantee that the continuum limit of the training dynamics of twolayer neural networks is the gradient flow of a *convex* energy, where the relevant notion of convexity along Wasserstein gradient flow is *displacement convexity* or *convexity along Wasserstein geodesics*; see Definition 2.6. This emergence of convexity in the continuum limit is relevant to the behavior of neural networks in practice, where researchers seek to explain why gradient descent dynamics sometimes converge to a global optimum, in spite of the fact that, at the discrete level, the energy landscape is nonconvex [27,81].

The remainder of the introduction proceeds as follows. In Section 1.1, we state fundamental properties of (WPME) and describe the analogy between (WPME) and (FP). In Section 1.2, we introduce our particle method for approximating solutions of (WPME). In Section 1.3, we describe the connection with two-layer neural networks. In Sections 1.4 and 1.5, we state our main assumptions and results. Finally, in Section 1.6, we outline our approach and describe directions for future work.

1.1. The weighted porous medium equation. A key feature of (WPME), which serves as a guiding principle of the present work, is that is it a *Wasserstein gradient flow* of the energy,

(1.1)
$$\mathcal{F}: \mathcal{P}(\mathbb{R}^d) \to \mathbb{R} \cup \{+\infty\}, \quad \mathcal{F}(\mu) = \mathcal{E}(\mu) + \mathcal{V}(\mu) + \mathcal{V}_{\Omega}(\mu)$$

where $\mathcal{P}(\Omega)$ denotes the set of Borel probability measures on Ω ; and the *internal energy* \mathcal{E} , *external potential* energy \mathcal{V}_{Ω} are given by,

(1.2)
$$\mathcal{E}(\mu) = \begin{cases} \frac{1}{2} \int_{\mathbb{R}^d} \frac{|\mu(x)|^2}{\bar{\rho}(x)} dx & \text{if } \mu \ll \bar{\rho}(x) dx \text{ and } d\mu(x) = \mu(x) dx, \\ +\infty & \text{otherwise,} \end{cases}$$

(1.3)
$$\mathcal{V}(\mu) = \int_{\mathbb{R}^d} V(x) d\mu(x),$$

(1.4)
$$\mathcal{V}_{\Omega}(\mu) = \begin{cases} 0 & \text{if supp } \mu \subseteq \overline{\Omega}, \\ +\infty & \text{otherwise.} \end{cases}$$

The internal energy \mathcal{E} induces the nonlinear diffusion term, the external potential \mathcal{V} induces the convection term, and the confining potential \mathcal{V}_{Ω} restricts the dynamics to Ω , with no flux boundary conditions on $\partial\Omega$. Our primary interest, and the main mathematical challenge in establishing our results, is in the nonlinear diffusion induced by \mathcal{E} and its approximation by a deterministic particle method. In Section 2, we provide detailed background on the Wasserstein metric W_2 and Wasserstein gradient flows. In Proposition 3.10, we recall the precise statement of the result that solutions of (WPME) are the gradient flow of \mathcal{F} .

The fact that (WPME) has a gradient flow structure is in close analogy with the (FP) equation: in their seminal work [49], Jordan, Kinderlehrer, and Otto established that (FP) is the Wasserstein gradient flow of the *Kullback-Leibler divergence*,

$$\operatorname{KL}(\mu, \bar{\rho}) = \int_{\Omega} \log(\mu/\bar{\rho}) d\mu, \quad \text{for } \mu \ll \bar{\rho}.$$

From this perspective, it is useful to notice that, when V = 0, (WPME) can also be thought of as the Wasserstein gradient flow of the χ^2 divergence [72],

$$\chi^{2}(\mu,\bar{\rho}) = \begin{cases} \frac{1}{2} \int \frac{|\mu(x)-\bar{\rho}(x)|^{2}}{\bar{\rho}(x)} dx, & \text{if } \mu \ll \mathcal{L}^{d}, \, d\mu(x) = \mu(x) dx, \text{ and supp } \mu \subseteq \overline{\Omega}, \\ +\infty & \text{otherwise.} \end{cases}$$

This can be seen by noticing $\int |\mu(x) - \bar{\rho}(x)|^2 / \bar{\rho}(x) dx = \int |\mu(x)|^2 / \bar{\rho}(x) dx - 1$, so that, when V = 0, our energy \mathcal{F} agrees with χ^2 , up to a constant that does not affect the dynamics of the gradient flow: $\mathcal{F}(\mu) + 1/2 = \chi^2(\mu, \bar{\rho})$. In what follows, we will always suppose that $\bar{\rho}$ is normalized to satisfy $\int_{\Omega} \bar{\rho} = 1$, and so that for $\mu \in \mathcal{P}(\mathbb{R}^d)$, the KL divergence and the χ^2 divergence measure the discrepancy between μ and $\bar{\rho}$ on Ω and vanish in the case that $\mu = \bar{\rho}$ on Ω .

The gradient flow structures of (WPME) and (FP) have important interpretations from the perspective of quantization, since they encode key information about how quickly solutions are flowing toward $\bar{\rho}$. The fact that solutions of the (FP) equation are the Wasserstein gradient flow of the KL divergence is equivalent to saying that they dissipate the KL divergence as quickly as possible, with respect to the Wasserstein structure. In the same way, solutions of the (WPME) equation dissipate the χ^2 divergence as quickly as possible, with respect to the Wasserstein structure.

Another important feature of (WPME) from the perspective of quantization are the available estimates quantifying its convergence to equilibrium. Chewi et al. [26], show that, if $\Omega = \mathbb{R}^d$, V = 0, and $\bar{\rho}$ satisfies a Poincaré inequality, then, along smooth solutions, the Kullback-Leibler divergence decreases exponentially:

(1.5)
$$\operatorname{KL}(\rho(t),\bar{\rho}) \leq e^{-C_{\bar{\rho}}t} \operatorname{KL}(\rho(0),\bar{\rho}), \quad \text{for } C_{\bar{\rho}} > 0.$$

If, in addition, $\bar{\rho}$ is strongly log-concave, then the χ^2 divergence decreases exponentially:

$$\chi^2(\rho(t),\bar{\rho}) \le e^{-C_{\bar{\rho}} t} \chi^2(\rho(0),\bar{\rho}), \quad \text{for } C_{\bar{\rho}} > 0,$$

This mirrors the theory for (FP), in which a Poincaré inquality ensures exponential decay of the χ^2 divergence and log-concavity ensures decay of the KL divergence. (See Matthes, McCann, and Savaré's *flow interchange method* for general results of this form [56]. In addition, see Grillo, Muratori, and Porzio [44], who rigorously proved exponential convergence to equilibrium of weak solutions in L^p spaces for all $p < +\infty$.) Furthermore, in the case of the (WPME) equation, if $\bar{\rho}$ merely satisfies a weaker condition, known as an $L^{2/3}$ -Poincaré inequality, then Dolbeault et al. [35] showed that the χ^2 divergence decreases polynomially. This raises the possibility that, for different choices of $\bar{\rho}$ and initial conditions ρ_0 , there may exist contexts in which solutions of (WPME) converge to $\bar{\rho}$ with stronger convergence guarantees than solutions of (FP). Since developing general conditions on the target $\bar{\rho}$ and the initialization ρ_0 that distinguish whether (WPME) or (FP) equilibrates more quickly remains an active area of research, we do not claim that the dynamics of (WPME) offer superior long time behavior to (FP). Instead, we merely observe that, at the continuum level, (WPME) provides competitive dynamics. Understanding when solutions to (WPME) or (FP) converge more quickly to equilibrium may, in the future, shed light on which quantization methods are superior in different contexts.

1.2. Particle approximation of (WPME). The aim of the present work is to design a deterministic particle method for approximating solutions of (WPME) that preserves its gradient flow structure. Since

solutions of (WPME) are gradient flows of the energy (1.1-1.4), we seek to approximate them by gradient flows of the *regularized energy*, defined by,

(1.6)
$$\mathcal{F}_{\varepsilon,k}: \mathcal{P}(\mathbb{R}^d) \to \mathbb{R} \cup \{+\infty\}, \quad \mathcal{F}_{\varepsilon,k}(\mu) = \mathcal{E}_{\varepsilon}(\mu) + \mathcal{V}_{\varepsilon}(\mu) + \mathcal{V}_k(\mu),$$

for the energies $\mathcal{E}_{\varepsilon}(\mu)$, $\mathcal{V}_{\varepsilon}(\mu)$, and $\mathcal{V}_{k}(\mu)$ given by,

(1.7)
$$\mathcal{E}_{\varepsilon}(\mu) = \frac{1}{2} \int_{\mathbb{R}^d} \frac{|\zeta_{\varepsilon} * \mu|^2(x)}{\bar{\rho}(x)} \, dx,$$

(1.8)
$$\mathcal{V}_{\varepsilon}(\mu) = \int (\zeta_{\varepsilon} * V)(x) d\mu(x),$$

(1.9)
$$\mathcal{V}_k(\mu) = \int V_k(x) d\mu(x).$$

Here $\zeta_{\varepsilon} \in C^{\infty}(\mathbb{R}^d)$ is a rapidly decreasing mollifier and $V_k \in C^2(\mathbb{R}^d)$, for $k \in \mathbb{N}$, is a convex function that vanishes on Ω and approaches $+\infty$ on Ω^c as $k \to +\infty$.

The energy $\mathcal{E}_{\varepsilon}(\mu)$ is an approximation, as $\varepsilon \to 0$, of \mathcal{E} . This regularized energy has superior differentiability properties along empirical measures, ensuring that the gradient flow starting at empirical measure initial data leads to a well-posed particle method. It also enjoys the property,

(1.10)
$$\mathcal{E}_{\varepsilon}(\rho) = \mathcal{E}(\zeta_{\varepsilon} * \rho),$$

which is a key element in our proof of an H^1 bound for $\zeta_{\varepsilon} * \rho_{\varepsilon}$ along solutions of the gradient flow; see Theorem 4.1. The energy $\mathcal{V}_{\varepsilon}$ is an approximation of \mathcal{V} . While many different methods of approximating \mathcal{V} would work well both numerically and theoretically, we focus our attention on $\mathcal{V}_{\varepsilon}$ due to the connection with two-layer neural networks. Finally, the energy \mathcal{V}_k is an approximation, as $k \to +\infty$, of \mathcal{V}_{Ω} .

While the main focus of our work is the analysis of how dynamics induced by $\mathcal{E}_{\varepsilon}$, for general initial data, approximate dynamics induced by \mathcal{E} (indeed, if Ω is the entire space \mathbb{R}^d and V is taken to be zero, then the energy $\mathcal{F}_{\varepsilon,k}$ is exactly $\mathcal{E}_{\varepsilon}$), our analysis of how the gradient flow dynamics induced by \mathcal{V}_k converge to those from \mathcal{V}_{Ω} as $k \to +\infty$ also generalizes existing results by Alasio, Bruna, and Carrillo to weighted porous medium equations [1]. (See also recent work by Patacchini and Slepčev, which uses a similar approach to study well-posedness of aggregation equations on compact manifolds [65].) Our motivations for considering this approximation of the confining potential \mathcal{V}_{Ω} are twofold. First, it simplifies the implementation of the particle method, obviating the need to implement reflection boundary conditions. Second, it allows for the most challenging aspect of the analysis — the relationship between the dynamics induced by $\mathcal{E}_{\varepsilon}$ and \mathcal{E} — to be carried out on \mathbb{R}^d , rather than on a domain with boundary.

Wasserstein gradient flows of the regularized energy $\mathcal{F}_{\varepsilon,k}$ are characterized by the equation,

(WPME_{$$\varepsilon,k$$})
$$\begin{cases} \partial_t \rho = \nabla \cdot \left(\rho \left(\nabla \zeta_{\varepsilon} * (\zeta_{\varepsilon} * \rho/\bar{\rho}) + \nabla \zeta_{\varepsilon} * V + \nabla V_k \right) \right) \\ \rho(0) = \rho_0, \end{cases}$$

defined on all of \mathbb{R}^d in the duality with $C_c^{\infty}(\mathbb{R}^d \times (0, +\infty))$; see Proposition 3.12. If the initial conditions are given by an empirical measure, $\rho_0 = \sum_{i=1}^N \delta_{X_0^i} m^i$, with $\sum_{i=1}^N m^i = 1$, then the solution remains an empirical measure for all time. Concretely, we have $\rho(t) = \sum_{i=1}^N \delta_{X_i(t)} m^i$, and the locations of the particles $\{X^i(t)\}_{i=1}^N$ are characterized as solutions of,

(1.11)
$$\begin{cases} \dot{X}^{i}(t) = -\sum_{j=1}^{N} f(X^{i}, X^{j})m^{j} - \nabla\zeta_{\varepsilon} * V(X^{i}) - \nabla V_{k}(X^{i}), \\ X^{i}(0) = X_{0}^{i}, \end{cases}$$

for,

(1.12)
$$f(x,y) := \int_{\mathbb{R}^d} \frac{\nabla \zeta_{\varepsilon}(x-z)\zeta_{\varepsilon}(y-z)}{\bar{\rho}(z)} \, dz;$$

see Proposition 3.13. In Section 7.1, we provide sufficient conditions on $\bar{\rho}$ for which the integral in f(x, y) has an analytic formula, in which case it can be precomputed exactly and does not contribute to the computational complexity of our method.

Based on the intuition that $\mathcal{F}_{\varepsilon,k}$ is an approximation of \mathcal{F} , it is natural to hope that gradient flows of $\mathcal{F}_{\varepsilon,k}$ approximate gradient flows of \mathcal{F} . Our main result is that this is indeed true. We show that the particle method defined by (1.11) converges to a solution of (WPME) on bounded time intervals, provided that

the initial conditions ρ_0 have bounded entropy and the number of particles N grows sufficiently quickly as $\varepsilon \to 0$; see Theorem 1.4. Note that this method formally extends to equations of the form (WPME) with an additional term $-\nabla \cdot (v\rho)$ on the right hand side, for general velocities v(x,t), by adding a term of the form $v(X^i(t), t)$ to the right hand side of (1.11).

Our work on the convergence of the $\varepsilon \to 0$ limit builds on several previous works studying the properties of (1.11) as $\varepsilon \to 0$. All previous works have considered the spatially homogeneous case $\bar{\rho} \equiv 1$. The first work in this direction was due to Oelschläger [62], who considered the case $V = V_k = 0$ and proved convergence to classical, strictly positive solutions of (WPME) in arbitrary dimensions and convergence to weak solutions in one dimension. Subsequently, Lions and Mas-Gallic [52], also in the case $V = V_k = 0$, proved convergence of (WPME_{ε,k}) as $\varepsilon \to 0$, provided that the initial conditions ρ_0 had uniformly bounded entropy, thereby excluding particle initial data required to connect (WPME_{$\varepsilon,k}) to the system of ODEs (1.11). The assumption$ $of bounded entropy played an important role in Lions and Mas-Gallic's proof of a <math>\dot{H}^1$ bound for regularized solutions to (WPME_{$\varepsilon,k}). (In fact, the analogous bound also plays an important role in the present work – see$ Theorem 4.1 for a generalization of this result to the spatially inhomogeneous setting.) Next, Carrillo, thefirst author, and Patacchini [20] generalized Lions and Mas-Gallic's approach to porous medium equationsof the form,</sub></sub>

$$\partial_t \rho = \Delta \rho^m + \nabla \cdot (\rho (\nabla V + \nabla W * \rho)).$$

In the case m = 2, they obtained convergence of the $\varepsilon \to 0$ limit under appropriate continuity and semiconvexity assumptions on V and W; for $1 \le m < 2$, they obtained Γ -convergence of the corresponding energies as $\varepsilon \to 0$; and for m > 2, they obtained conditional convergence of the $\varepsilon \to 0$ limit, as long as certain a priori estimates were preserved along the flow. Again, Carrillo, Craig, and Patacchini's work required the initial data to have bounded entropy, excluding particle solutions. Very recently, Burger and Esposito [15] continued the study of the m = 2 case for more general velocity fields v(x, t),

$$\partial_t \rho + \nabla \cdot (\rho v) = \Delta \rho^m,$$

and weaker regularity on the mollifier ζ .

Our work makes three contributions to this active area of research. We obtain true convergence of the particle method, relaxing the hypothesis that the initial data have bounded entropy by using stability properties of the regularized flow; see Theorem 1.4. Our result holds for spatially inhomogeneous porous medium equations, allowing general $\bar{\rho} \in C^1(\mathbb{R}^d)$ that are log-concave and bounded above and below on $\Omega \subseteq \mathbb{R}^d$. (See the next section for a discussion of where the log-concavity assumption may be weakened.) Finally, by allowing spatially inhomogenous equations, we identify a connection between our particle method and problems in sampling, control theory, and training of two-layer neural networks.

1.3. Application to two-layer neural networks. An additional reason for interest in the convergence of (1.11-1.12) to (WPME), aside from its utility as a particle approximation, is that the dynamics of (1.11-1.12) for $\Omega = \mathbb{R}^d$ coincide precisely with the training dynamics for mean field models of two-layer neural networks with a radial basis function activation function. In this context, one is given a data distribution ν , a nonnegative target function $f_0 \in L^2(\nu)$, and an activation function $\Phi_{\varepsilon}(x, z) = \zeta_{\varepsilon}(x - z)$, and one seeks to choose parameters, $\{X^i\}_{i=1}^N$, so that the empirical measure $\rho^N = \frac{1}{N} \sum_{i=1}^N \delta_{X^i}$ minimizes the following energy, known as the *population risk*:

(1.13)
$$\mathcal{R}_{\varepsilon}(\mu) = \frac{1}{2} \int \left| \int \Phi_{\varepsilon}(x, z) d\mu(x) - f_0(z) \right|^2 d\nu(z).$$

In several recent works, it was discovered that evolving the parameters $X^i(t)$ by gradient descent of the function $(X^1, \ldots, X^n) \mapsto \mathcal{R}_{\varepsilon}(\rho^N)$ is equivalent to evolving the empirical measure ρ^N by the Wasserstein gradient flow of $\mathcal{R}_{\varepsilon}$ [27, 48, 59, 66, 70, 79]. To see the connection with (1.11-1.12), note that, using the definition of Φ_{ε} , expanding the square, and applying Tonelli's theorem (see also the associativity property of convolution (2.1)), we obtain,

(1.14)
$$\mathcal{R}_{\varepsilon}(\mu) = \frac{1}{2} \int |\zeta_{\varepsilon} * \mu(z)|^2 d\nu(z) - \int \zeta_{\varepsilon} * \mu(z) f_0(z) d\nu(z) - \frac{1}{2} \int |f_0(z)|^2 d\nu(z)$$
$$= \mathcal{E}_{\varepsilon}(\mu) + \int (\zeta_{\varepsilon} * V)(x) d\mu(x) + C = \mathcal{E}_{\varepsilon}(\mu) + \mathcal{V}_{\varepsilon}(\mu) + C,$$

for,

(1.15)
$$\nu = 1/\bar{\rho} , \quad V = -f_0 \nu , \quad C = -\frac{1}{2} \int |f_0(z)|^2 d\nu(z).$$

This shows that, for $\Omega = \mathbb{R}^d$, which ensures $\mathcal{V}_k \equiv 0$, we have $\mathcal{R}_{\varepsilon}(\rho) = \mathcal{F}_{\varepsilon,k}(\rho) + C$. Since the constant does not effect the dynamics of the gradient flow, we see that the gradient flow of $\mathcal{R}_{\varepsilon}$ for general initial data ρ_0 is characterized by (WPME_{ε,k}), and the evolution for particle initial data corresponds to (1.11-1.12). Corollary 1.6, which follows from our convergence result for the gradient flows of $\mathcal{F}_{\varepsilon,k}$, states that, for well-behaved initial conditions, particle solutions of (1.11-1.12) converge to a gradient flow of,

(1.16)
$$\mathcal{R}(\mu) = \begin{cases} \frac{1}{2} \int |\mu(z) - f_0(z)|^2 d\nu(z) & \text{if } \mu \ll \mathcal{L}^d, \quad d\mu(z) = \mu(z) dz, \\ +\infty & \text{otherwise.} \end{cases}$$

This generalizes previous work due to Javanmard, Mondelli, and Montanari [48], which considered the limit $\varepsilon \to 0$ in the specific case of a uniform data distribution $\nu = \mathbb{1}_{\Omega}/|\Omega|$, smooth target function f, bounded domain Ω , and compactly supported radial basis function ζ , with the no flux boundary conditions from the continuum PDE imposed by reflecting boundary conditions for the particle ODEs. The fact that our result holds for general nonuniform data distributions ν is significant from the perspective of two-layer neural networks, since, as can be seen in Corollary 1.6, there is an interplay between the data distribution ν and the target function f to determine when convexity of the energy $\mathcal{F}_{\varepsilon,k}$ emerges in the continuum limit.

1.4. Assumptions. We now describe our assumptions. We consider a domain $\Omega \subseteq \mathbb{R}^d$ satisfying,

(D)
$$\Omega$$
 is nonempty, open, and convex.

We suppose our mollifier satisfies,

(M)

$$\zeta \in C^2(\mathbb{R}^d) \text{ is even, nonnegative, } \|\zeta\|_{L^1(\mathbb{R}^d)} = 1, D^2\zeta \in L^\infty(\mathbb{R}^d),$$

$$\zeta(x) \le C_\zeta |x|^{-q} \text{ and } |\nabla\zeta(x)| \le C_\zeta |x|^{-q'}, \text{ for } C_\zeta > 0, \quad q > d+1, \quad q' > d$$

This assumption is satisfied by both Gaussians and smooth functions with compact support. Note that this assumption ensures that ζ has finite first moment, $\int_{\mathbb{R}^d} |x| \zeta(x) dx < +\infty$.

We suppose the external potential V satisfies,

(V)
$$V \in C^2(\mathbb{R}^d) \cap L^1(\mathbb{R}^d) \cap L^\infty(\mathbb{R}^d)$$
, with $\nabla V \in L^\infty(\mathbb{R}^d)$ and D^2V uniformly bounded below

We are optimistic that our results may continue to hold under weaker regularity hypotheses on V, but we leave this question to future work, since our primary interest is the approximation of the diffusive dynamics arising from \mathcal{E} via the particle method induced by $\mathcal{E}_{\varepsilon}$.

We suppose that our approximation of the confining potential V_k , for $k \in \mathbb{N}$, satisfies,

(C)
$$V_k$$
 is nonnegative, convex, and twice differentiable with $D^2 V_k \in L^{\infty}(\mathbb{R}^d)$,

(Ck)
$$V_k = 0 \text{ on } \Omega \text{ and } \lim_{k \to \infty} \left(\inf_{x \in B} V_k(x) \right) = +\infty \text{ for any ball } B \subset \Omega^c$$

Note that assumption (C) ensures $V_k \in L^1(\mu)$ and $\nabla V_k \in L^2(\mu)$ for any $\mu \in \mathcal{P}(\mathbb{R}^d)$ with $\int |x|^2 d\mu(x) < +\infty$. These assumptions play the following role in our proof: Assumption (C) ensures well-posedness of the gradient flows, and Assumption (Ck) allows us to recover the correct limiting dynamics as $k \to +\infty$. In particular, note that (Ck) implies that, in the $k \to \infty$ limit, V_k approximates the hard cutoff potential V_{Ω} , which is given by,

(1.17)
$$V_{\Omega}(x) = \begin{cases} 0 & \text{for } x \in \Omega, \\ +\infty & \text{otherwise.} \end{cases}$$

Finally, we suppose that our target $\bar{\rho}$ satisfies the regularity assumption,

(T)
$$\bar{\rho} \in C^1(\mathbb{R}^d)$$
 and there exists $C > 0$ so that $1/C \leq \bar{\rho}(x) \leq C$, for all $x \in \mathbb{R}^d$.

Assumption (T) is sufficient to ensure that the energy $\mathcal{E}_{\varepsilon}$ is lower semicontinuous, convex, and subdifferentiable, so that gradient flows of $\mathcal{E}_{\varepsilon}$ are well posed. It also allows us to conclude that the energy \mathcal{E} is lower semicontinuous. However, in order to obtain convexity and subdifferentiability of \mathcal{E} , hence well-posedness of gradient flows, we leverage existing results due to Ambrosio, Gigli, and Savaré [3] and require $\bar{\rho}$ to be log-concave; that is,

$$x \mapsto \log(\bar{\rho}(x))$$
 is concave.

It is an open question whether well-posedness of the gradient flow of \mathcal{E} could be obtained under weaker assumptions on $\bar{\rho}$. Interestingly, the main estimates in our proof of the convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ to \mathcal{F} (Theorem 4.1, Theorem 5.1, and Proposition 5.2) do not require log-concavity of $\bar{\rho}$. Instead, log-concavity comes into play when we seek to identify that the limit as $\varepsilon \to 0$ of gradient flows of $\mathcal{E}_{\varepsilon}$ is indeed a gradient flow of \mathcal{E} , since log-concavity of $\bar{\rho}$ ensures that the metric slope of \mathcal{E} is a a strong upper gradient; see Section 2.3 and [3, Section 1.2]. For this reason, we are optimistic that, in future work, it will be possible to extend our results to $\bar{\rho}$ that are not log-concave, once the difficulty of obtaining well-posedness of the gradient flow of \mathcal{E} and characterization of its strong upper gradient are overcome.

1.5. Main Results. With these assumptions in hand, we now state our main results. In order to get convergence of (WPME_{ε,k}) to (WPME), we consider the limits $\varepsilon \to 0$ and $k \to +\infty$ separately: first, we show that gradient flows of $\mathcal{F}_{\varepsilon,k}$ converge, as $\varepsilon \to 0$, to a gradient flow of,

(1.18)
$$\mathcal{F}_k(\rho) = \mathcal{E}(\rho) + \mathcal{V}(\rho) + \mathcal{V}_k(\rho)$$

second, we establish that gradient flows of \mathcal{F}_k converge, as $k \to +\infty$, to a gradient flow of \mathcal{F} .

Our first theorem shows that, if the initial conditions $\rho_{\varepsilon,k}(0)$ of the gradient flow of $\mathcal{F}_{\varepsilon,k}$ have uniformly bounded entropy $\mathcal{S}(\mu)$ and second moment $M_2(\mu)$, given by,

(1.19)
$$\mathcal{S}(\mu) = \begin{cases} \int_{\mathbb{R}^d} \mu(x) \log \mu(x) d\mathcal{L}^d(x) & \text{if } \mu \ll \mathcal{L}^d \text{ and } d\mu(x) = \mu(x) dx, \\ +\infty & \text{otherwise,} \end{cases} \qquad M_2(\mu) = \int_{\mathbb{R}^d} |x|^2 d\mu(x),$$

and the initial conditions are "well-prepared," then the gradient flows of $\mathcal{F}_{\varepsilon,k}$ with initial data $\rho_{\varepsilon,k}(0)$ converge, as $\varepsilon \to 0$, to a gradient flow of \mathcal{F}_k with initial data $\rho_k(0)$.

Recall that a probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$ lies in the *domain* of an energy $\mathcal{G} : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R} \cup \{+\infty\}$ if $\mathcal{G}(\mu) < +\infty$. We denote this by $\mu \in D(\mathcal{G})$. We also write $\mathcal{P}_2(\mathbb{R}^d) = \mathcal{P}(\mathbb{R}^d) \cap D(M_2)$. Finally, we often use the notion of *narrow* convergence of probability measures; see Definition 2.2.

Theorem 1.1 (convergence of gradient flows as $\varepsilon \to 0$). Assume (D), (M), (V), (C), (T), and that $\bar{\rho}$ is log-concave. Fix T > 0 and $k \in \mathbb{N}$. For $\varepsilon > 0$, let $\rho_{\varepsilon,k} \in AC^2([0,T];\mathcal{P}_2(\mathbb{R}^d))$ be a gradient flow of $\mathcal{F}_{\varepsilon,k}$ satisfying,

(1.20)
$$\sup_{\varepsilon > 0} \mathcal{S}(\rho_{\varepsilon,k}(0)) < \infty \quad and \quad \sup_{\varepsilon > 0} M_2(\rho_{\varepsilon,k}(0)) < \infty.$$

Suppose there exists $\rho_k(0) \in D(\mathcal{F}_k) \cap \mathcal{P}_2(\mathbb{R}^d)$ such that,

(1.21)
$$\rho_{\varepsilon,k}(0) \xrightarrow{\varepsilon \to 0} \rho_k(0) \text{ narrowly} \text{ and } \lim_{\varepsilon \to 0} \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}(0)) = \mathcal{F}_k(\rho_k(0)).$$

Then $\rho_{\varepsilon,k}(t) \xrightarrow{\varepsilon \to 0} \rho_k(t)$ narrowly, for all $t \in [0,T]$, where $\rho_k \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ is the gradient flow of \mathcal{F}_k with initial data $\rho_k(0)$.

Next, we consider the limit as $k \to +\infty$, proving that gradient flows of \mathcal{F}_k with "well-prepared" initial data converge to a gradient flow of \mathcal{F} .

Theorem 1.2 (convergence of gradient flows as $k \to +\infty$). Assume (D), (V), (C), (Ck), (T), and that $\bar{\rho}$ is log-concave. Fix T > 0. For $k \in \mathbb{N}$, let $\rho_k \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ be a gradient flow of \mathcal{F}_k and suppose there exists $\rho(0) \in D(\mathcal{F}) \cap \mathcal{P}_2(\mathbb{R}^d)$ such that,

(1.22) $\rho_k(0) \xrightarrow{k \to +\infty} \rho(0) \text{ narrowly} \text{ and } \lim_{k \to +\infty} \mathcal{F}_k(\rho_k(0)) = \mathcal{F}(\rho(0)).$ Then $\rho_k(t) \xrightarrow{k \to \infty} \rho(t) \text{ narrowly, for all } t \in [0, T], \text{ where } \rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d)) \text{ is the gradient flow of } \mathcal{F}(t)$ with initial data $\rho(0)$.

Combining the previous two theorems, we immediately obtain the following corollary, which yields convergence of the gradient flows as $\varepsilon \to 0$ and $k \to +\infty$. For simplicity, we state this result in the specific case when the same initial conditions are used for all gradient flows, though the result continues to hold for varying initial data satisfying the hypotheses of Theorems 1.1 and 1.2.

Corollary 1.3 (convergence of gradient flows as $k \to +\infty$, $\varepsilon = \varepsilon(k) \to 0$). Assume (D), (M), (V), (C), (Ck), (T) and that $\bar{\rho}$ is log-concave. Fix T > 0 and $\rho(0) \in D(\mathcal{F}) \cap D(\mathcal{S}) \cap \mathcal{P}_2(\mathbb{R}^d)$.

For $\varepsilon > 0$ and $k \in \mathbb{N}$, let $\rho_{\varepsilon,k} \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ be a gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial data $\rho(0)$. Then, as $k \to +\infty$, there exists a sequence $\varepsilon = \varepsilon(k) \to 0$ so that $\rho_{\varepsilon,k}(t)$ narrowly converges to $\rho(t)$, for all $t \in [0,T]$, where $\rho \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ is the gradient flow of \mathcal{F} with initial data $\rho(0)$.

The preceding corollary requires that the initial conditions of the gradient flow of $\mathcal{F}_{\varepsilon,k}$ have bounded entropy, which explicitly excludes empirical measure initial data. However, we are able to extend this result to empirical measure initial data by leveraging stability properties of the gradient flow of $\mathcal{F}_{\varepsilon,k}$. In this way, we obtain the following convergence result for the deterministic particle method to weak solutions of (WPME), provided that the underlying continuum solution has initial data with bounded entropy. In Proposition 3.10, we state the precise notion of weak solution of (WPME) that we consider, and in Lemma A.4, we provide an explicit construction of $\rho_{\varepsilon,k}^{N}(0)$ satisfying condition (1.24).

Theorem 1.4 (convergence with particle initial data). Assume (D), (M), (V), (C), (Ck), (T), and that $\bar{\rho}$ is log-concave. Fix T > 0 and $\rho(0) \in D(\mathcal{F}) \cap D(\mathcal{S}) \cap \mathcal{P}_2(\mathbb{R}^d)$. For $k, N \in \mathbb{N}$, $\varepsilon > 0$, and $t \in [0, T]$, consider the evolving empirical measure,

$$\rho_{\varepsilon,k}^N(t) = \sum_{i=1}^N \delta_{X_{\varepsilon,k}^i}(t) m^i, \quad m^i \ge 0, \quad \sum_{i=1}^N m^i = 1,$$

where $X^i_{\varepsilon,k} \in C^1([0,T]; \mathbb{R}^d)$ solves,

(1.23)
$$\begin{cases} \dot{X}_{\varepsilon,k}^{i} = -\sum_{j=1}^{N} m^{j} \int_{\mathbb{R}^{d}} \nabla \zeta_{\varepsilon} (X_{\varepsilon,k}^{i} - z) \zeta_{\varepsilon} (z - X_{\varepsilon,k}^{j}) \frac{1}{\bar{\rho}(z)} dz - \nabla (\zeta_{\varepsilon} * V) (X_{\varepsilon,k}^{i}) - V_{k} (X_{\varepsilon,k}^{i}) + V_{k} (X_{\varepsilon,k}^{i$$

Suppose that as $\varepsilon \to 0$ there exist $N = N(\varepsilon) \to +\infty$, so that, for all $k \in \mathbb{N}$, $\rho_{\varepsilon,k}^N(0) = \sum_{i=1}^N \delta_{X_{0,\varepsilon}^i} m^i$ converges to $\rho(0)$ with the rate,

$$(1.24) \lim_{k \to \infty} e^{-\lambda_{\varepsilon} T} W_2(\rho_{\varepsilon,k}^N(0), \rho(0)) = 0, \text{ for } \lambda_{\varepsilon} = -\varepsilon^{-d-2} ||1/\bar{\rho}||_{L^{\infty}(\mathbb{R}^d)} ||D^2\zeta||_{L^{\infty}(\mathbb{R}^d)} + \inf_{\{x,\xi \in \mathbb{R}^d\}} \xi^t D^2 V(x)\xi.$$

Then, as $k \to +\infty$, there exist $\varepsilon = \varepsilon(k) \to 0$ and $N = N(\varepsilon) \to +\infty$ for which $\rho_{\varepsilon,k}^N(t) = \sum_{i=0}^N \delta_{X_{\varepsilon,k}^i(t)} m^i$ narrowly converges to $\rho(t)$ for all $t \in [0,T]$, where $\rho \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ is the unique weak solution of (WPME) with initial data $\rho(0)$.

The following corollary ensures that the particle method defined in the previous theorem indeed converges to $\bar{\rho}$ on Ω in the long time limit, as relevant for applications in quantization.

Corollary 1.5 (long time limit). Suppose the assumptions of Theorem 1.4 hold and again denote $\rho_{\varepsilon,k}^N(t) = \sum_{i=0}^N \delta_{X_{\varepsilon,k}^i(t)} m^i$. In addition, assume V = 0, Ω is bounded, and $\int_{\Omega} \bar{\rho} d\mathcal{L}^d = 1$.

Then there exist $k = k(T) \to +\infty$, $\varepsilon = \varepsilon(k) \to 0$, and $N = N(\varepsilon) \to +\infty$ so that $\rho_{\varepsilon,k}^N(\cdot,t)$ narrowly converges to $\mathbb{1}_{\overline{\Omega}}(\cdot)\overline{\rho}(\cdot)$ as $t \to \infty$.

The preceding theorems provide sufficient conditions to guarantee convergence of the particle method to (WPME) on bounded time intervals and convergence to the desired target distribution $\bar{\rho}$ on Ω when V = 0 and Ω is bounded. However, these results are purely qualitative, and it remains an open question to what extent they could be made quantitative in T, k, ε , and N. For example, an inspection of the construction in Lemma A.4 shows that, if the particles are initialized with uniform spacing on a bounded domain Ω , the number of particles is required to grow extremely quickly with respect to ε . In particular, it suffices to have

$$N(\varepsilon, k)^{-1} = o\left(e^{-1/\varepsilon^{d+2}}\right) \text{ as } \varepsilon \to 0.$$

On the other hand, we observe numerically that $N(\varepsilon) \sim \varepsilon^{-1.01}$ is sufficient for good performance in one dimension. We leave a finer quantitative convergence analysis to future work. For example, it would be interesting to investigate whether higher regularity of the initial data $\rho(0)$ could be used to decrease the rate at which N must grow with ε in our rigorous convergence results, as the numerical simulations suggest is possible. As a second corollary of our main convergence results, we identify the limit of the training dynamics of two-layer neural networks with a radial basis function activation function and quadratic loss, as described in equations (1.13-1.15). In particular, our result gives sufficient conditions under which the limit of these training dynamics is the gradient flow of a *convex* energy, in the sense that it is convex along Wasserstein geodesics; see Definition 2.6.

Corollary 1.6. Consider a radial basis function activation function $\Phi_{\varepsilon}(x,z) = \zeta_{\varepsilon}(x-z)$ satisfying (M), a data distribution $\nu = 1/\bar{\rho}$, for $\bar{\rho}$ satisfying (T) and log-concave, and a target function $f_0 = -V\bar{\rho}$, for Vsatisfying (V). Fix T > 0. For $\varepsilon > 0$, $N \in \mathbb{N}$, and $t \in [0,T]$, consider the training dynamics of a twolayer neural network corresponding to the quadratic loss $\mathcal{R}_{\varepsilon}$; that is, consider the evolution of the empirical measure of parameters,

$$\rho_{\varepsilon}^{N}(t) = \sum_{i=1}^{N} \delta_{X_{\varepsilon}^{i}}(t) m^{i}, \quad m^{i} \ge 0, \quad \sum_{i=1}^{N} m^{i} = 1,$$

where $X_{\varepsilon}^i \in C^1([0,T]; \mathbb{R}^d)$ solves,

(1.25)
$$\begin{cases} \dot{X}_{\varepsilon}^{i} = -\sum_{j=1}^{N} m^{j} \int_{\mathbb{R}^{d}} \nabla \zeta_{\varepsilon} (X_{\varepsilon}^{i} - z) \zeta_{\varepsilon} (z - X_{\varepsilon}^{j}) \nu(z) \, dz + \nabla (\zeta_{\varepsilon} * (f_{0}\nu)) (X_{\varepsilon}^{i}) \\ X_{\varepsilon}^{i}(0) = X_{0,\varepsilon}^{i}. \end{cases}$$

Suppose there exists $\rho(0) \in D(\mathcal{F}) \cap D(\mathcal{S}) \cap \mathcal{P}_2(\mathbb{R}^d)$, with $\Omega = \mathbb{R}^d$, so that, for all $\varepsilon > 0$, there exists $N = N(\varepsilon)$ so that $\rho^N(0)$ converges to $\rho(0)$ sufficiently quickly, according to the rate from equation (1.24). Then, $\rho_{\varepsilon}^N(t) = \sum_{i=0}^N \delta_{X_{\varepsilon}^i(t)} m^i$ narrowly converges as $\varepsilon \to 0$ to $\rho(t)$ for all $t \in [0,T]$, where $\rho \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ is the unique weak solution of (WPME) with initial data $\rho(0)$.

In particular, whenever ν is log-convex and $f_0\nu$ is concave, the limit of the training dynamics is the gradient flow of the convex energy \mathcal{R} .

Our last main result concerns the behavior of minimizers of the energies $\mathcal{F}_{\varepsilon,k}$ and \mathcal{F}_k . Our proofs of Theorems 1.1 and 1.2 on the convergence of gradient flows as $\varepsilon \to 0$ and $k \to +\infty$ leverage Serfaty's general metric space framework for Γ -convergence of gradient flows [69], which we recall in Section 2.4. As a consequence of this approach, we easily obtain that, under sufficient compactness assumptions on the approximation of our confining potential V_k , minimizers of $\mathcal{F}_{\varepsilon,k}$ converge to a minimizer of \mathcal{F}_k and minimizers of \mathcal{F}_k converge to a minimizer of \mathcal{F} .

Theorem 1.7 (minimizers converge to minimizers). Suppose Assumptions (T), (D), (M), (V) and (C) hold. Assume that for any $k \in \mathbb{N}$, the sublevel sets of V_k are compact.

- (i) Fix $k \in \mathbb{N}$. If $\rho_{\varepsilon,k} \in \mathcal{P}_2(\mathbb{R}^d)$ is a minimizer of $\mathcal{F}_{\varepsilon,k}$ for all $\varepsilon > 0$, then there exists a subsequence such that, as $\varepsilon \to 0$, $\rho_{\varepsilon,k}$ narrowly converges to $\rho_k \in \mathcal{P}_2(\mathbb{R}^d)$, where ρ_k is a minimizer of \mathcal{F}_k .
- (ii) Assume that $V_k(x) \ge V_1(x)$ for each $x \in \mathbb{R}^d$ and $k \in \mathbb{N}$. If $\rho_k \in \mathcal{P}_2(\mathbb{R}^d)$ is a minimizer of \mathcal{F}_k for each $k \in \mathbb{N}$, then there exists a subsequence such that, as $k \to +\infty$, ρ_k narrowly converges to $\rho \in \mathcal{P}_2(\mathbb{R}^d)$, where ρ is a minimizer of \mathcal{F} .

This theorem has the potential to shed light on the convergence of the gradient flows in the long time limit. In particular, while our main results on convergence of the gradient flows only hold on *bounded* time intervals, if one could show that a gradient flow $\rho_{\varepsilon,k}(t)$ of $\mathcal{F}_{\varepsilon,k}$ indeed converged as $t \to +\infty$ to a minimizer of $\mathcal{F}_{\varepsilon,k}$, uniformly in $\varepsilon > 0$ and $k \in \mathbb{N}$, then one could combine the above theorem with the preceding theorems to get convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ to \mathcal{F} globally in time. Proving these estimates remains an open question, closely related to our motivating applications in quantization.

1.6. Outline of approach and future directions. We now outline our approach to proving these results. We begin, in Section 2, by recalling preliminary information on optimal transport, including basic notation in Section 2.1, convolution and convergence of measures in Section 2.2, optimal transport and Wasserstein gradient flows in Section 2.3, and Serfaty's general framework for Γ -convergence of gradient flows in Section 2.4. In Section 3, we prove several fundamental properties of the energy $\mathcal{F}_{\varepsilon,k}$ and recall known properties of the energies \mathcal{F}_k and \mathcal{F} , including convexity and differentiability in Section 3.1. We give the PDE characterizations of gradient flows of these energies in Section 3.2 and address the long time behavior of gradient flows of the energy \mathcal{F} in Section 3.3.

With these results in hand, we move on to studying the behavior of gradient flows of $\mathcal{F}_{\varepsilon,k}$ as $\varepsilon \to 0$ and $k \to +\infty$. Section 4 is devoted to proving a key estimate for the analysis of the $\varepsilon \to 0$ limit, which shows that if the initial conditions of the gradient flow of $\mathcal{F}_{\varepsilon,k}$ have bounded entropy, then the mollified gradient flow $\zeta_{\varepsilon} * \rho_{\varepsilon,k}(t)$ satisfies an H^1 bound; see Theorem 4.1. In Section 4.1, we sketch our proof of this result, formally integrating by parts, and in Sections 4.2-4.3, we prove the result, using the flow interchange method developed by Matthes, McCann, and Savaré [56].

In Section 5, we use the results of Section 4 to study the $\varepsilon \to 0$ limit. In Section 5.1, we obtain Γ convergence of the energies $\mathcal{F}_{\varepsilon,k}$ as $\varepsilon \to 0$ and use this to prove convergence of minimizers, as in Theorem
1.7(i). In Section 5.2, we move on to considering Γ -convergence of the gradient flows as $\varepsilon \to 0$, ultimately
proving their convergence to a gradient flow of \mathcal{F}_k , as in Theorem 1.1, under the key hypothesis that the
initial conditions of the gradient flow has uniformly bounded entropy.

Section 6 considers the $k \to +\infty$ limit, obtaining Γ -convergence of the energies \mathcal{F}_k to \mathcal{F} , as $k \to +\infty$, as well as our main theorems on convergence of the minimizers, Theorem 1.7(ii), and convergence of the gradient flows, Theorem 1.2. With these results in hand, we turn in Section 6.1 to extending the preceding convergence results on the gradient flows as $\varepsilon \to 0$, $k \to +\infty$ to allow for gradient flows with particle initial data, thereby obtaining the proof of Theorem 1.4. We also prove Corollary 1.5 on the long time behavior of the particle method and Corollary 1.6 on the limit of two-layer neural networks.

We close in Section 7 with several numerical examples illustrating key properties of our method. We explore the dynamics and long time behavior of particle solutions, for targets $\bar{\rho}$ that satisfy the log-concavity assumptions of our main theorems, as well as targets that fail this assumption but satisfy a Poincaré inequality. In both cases, we observe that our particle discretization captures the behavior of the continuum PDE when V = 0 and flows toward $\bar{\rho}$ on Ω in the long-time limit. We also explore the effect of the confining potential V_k on the dynamics for various choices of k, observing the qualitative agreement with no-flux boundary conditions on Ω , as well as the quantitative effect on rate of convergence to (WPME) as $N \to +\infty$, $\varepsilon \to 0$. In the case of strong confinement ($k = 10^9$) and log-concave target $\bar{\rho}$, we observe first order convergence in N, with $\varepsilon = 4/N^{0.99}$ on $\Omega = (-1, 1)$, both for the rate of convergence of the particle method to solutions of (WPME) and for convergence of the particle method to the target $\bar{\rho}$ on Ω in the long time limit. Finally, as our scheme preserves the gradient flow structure of (WPME), it succeeds in capturing the exponential decay of the KL divergence along particle method solutions (see inequality (1.5)), up to discretization error and is energy decreasing for $\mathcal{F}_{\varepsilon,k}$ for all values of N, ε , and k.

There are several directions for future work. Many of our results only lightly use the assumption that $\bar{\rho}$ is log-concave, and it would be interesting to remove it. A key challenge in this direction is obtaining well-posedness of the gradient flow of \mathcal{F} in the absence of convexity of the energy and proving that the metric slope is a strong upper gradient. A second direction for future work would be to improve methods for computing or approximating f(x, y), as defined in (1.12), which drives the dynamics of our system of ODEs. To compute this exactly involves integrating the reciprocal of the target $\bar{\rho}$ against the mollifiers, which can be done analytically for a variety of targets $\bar{\rho}$, including piecewise constant $\bar{\rho}$; see appendix C). Better understanding of the minimal information required on $\bar{\rho}$ required to approximate (1.12) and the effect of this approximation on the dynamics would be important to applying this method in practice, especially when only partial information of $\bar{\rho}$ is known. A third interesting open question would be to obtain quantitative estimates could be combined with existing estimates on the long time behavior of (WPME) to provide convergence guarantees regarding the convergence of the particle method to the target $\bar{\rho}$ on Ω .

2. Preliminaries

2.1. **Basic notation.** For any r > 0 and $x \in \mathbb{R}^d$ we use $B_r(x)$ to denote the open ball of center x and radius r. We write $\mathbb{1}_S$ for the indicator function of a given subset S of \mathbb{R}^d . i.e.,

$$\mathbb{1}_{S}(x) = \begin{cases} 1 & \text{for } x \in S, \\ 0 & \text{otherwise.} \end{cases}$$

We denote the *d*-dimensional Lebesgue measure by \mathcal{L}^d . Given $\mu \in \mathcal{P}(\mathbb{R}^d)$, we write $\mu \ll \mathcal{L}^d$ if μ is absolutely continuous with respect to \mathcal{L}^d , in which case we will denote both the probability measure μ and its Lebesgue density by the same symbol, e.g. $d\mu(x) = \mu(x)dx$. Finally, we let $L^p(\mu; \Omega)$ denote the Lebesgue space of

functions f on Ω with $|f|^p$ being μ -integrable, and abbreviate $L^p(\Omega) = L^p(\mathcal{L}^d; \Omega)$. (We commit a slight abuse of notation by using the same notation for the Lebesgue spaces of real-valued and \mathbb{R}^d -valued functions.)

2.2. Convolution and convergence of measures. A fundamental aspect of our approach is the regularization of the energy (1.2) via convolution with a mollifier. We now recall some elementary results on convolution of probability measures. For any $\mu \in \mathcal{P}(\mathbb{R}^d)$ and $\phi \in L^{\infty}(\mathbb{R}^d)$, the convolution of ϕ with μ is defined by,

$$\phi * \mu(x) = \int_{\mathbb{R}^d} \phi(x - y) \, d\mu(y) \quad \text{for all } x \in \mathbb{R}^d.$$

Throughout, we use the fact that the definition of convolution allows us to move mollifiers from the measure to the integrand. In particular, for any f bounded below and $\phi \in L^1(\mathbb{R}^d)$ even, we have,

(2.1)
$$\int_{\mathbb{R}^d} f \, d(\phi * \mu) = \int_{\mathbb{R}^d} f * \phi \, d\mu$$

Likewise, we often use the following mollifier exchange lemma, which provides sufficient conditions for moving functions in and out of convolutions within integrals.

Lemma 2.1 (mollifier exchange lemma, [20, Lemma 2.2]). Let $f : \mathbb{R}^d \to \mathbb{R}$ be Lipschitz continuous with constant $L_f > 0$, and let σ and ν be finite, signed Borel measures on \mathbb{R}^d . There is p = p(q, d) > 0 so that,

$$\left| \int \zeta_{\varepsilon} * (f\nu) \, d\sigma - \int (\zeta_{\varepsilon} * \nu) f \, d\sigma \right| \le \varepsilon^p L_f \left(\int (\zeta_{\varepsilon} * |\nu|) \, d|\sigma| + C_{\zeta} |\sigma|(\mathbb{R}^d) |\nu|(\mathbb{R}^d) \right) \quad \text{for all } \varepsilon > 0$$

We will often use the following notion of convergence:

Definition 2.2 (narrow convergence). A sequence μ_n in $\mathcal{P}(\mathbb{R}^d)$ is said to *narrowly converge* to $\mu \in \mathcal{P}(\mathbb{R}^d)$ if $\int f d\mu_n \to \int f d\mu$ for all bounded and continuous functions f.

For fixed $\phi \in C_b(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$ even and any sequence μ_n narrowly converging to μ , we immediately obtain from the definition of narrow convergence that, for any $f \in C_b(\mathbb{R}^d)$,

(2.2)
$$\int f(\phi * \mu_n) = \int (f * \phi) d\mu_n = \int (f * \phi) d\mu = \int f(\phi * \mu) d\mu_n$$

so $\phi * \mu_n$ narrowly converges to $\phi * \mu$. Moreover, we have:

Lemma 2.3 (mollifiers and narrow convergence, [20, Lemma 2.3]). Suppose ζ_{ε} is a mollifier satisfying Assumption (M), and let μ_{ε} be a sequence in $\mathcal{P}(\mathbb{R}^d)$ converging narrowly to $\mu \in \mathcal{P}(\mathbb{R}^d)$. Then $\zeta_{\varepsilon} * \mu_{\varepsilon}$ narrowly converges to μ .

2.3. Optimal transport, the Wasserstein metric, and Wasserstein gradient flows. We now describe basic facts about optimal transport and the Wasserstein metric, which we will use in what follows. For further background, we refer the reader to one of the many excellent textbooks on the subject [2,3,41,68,75].

For a Borel measurable map $\mathbf{t} \colon \mathbb{R}^n \to \mathbb{R}^m$, we say that \mathbf{t} transports $\mu \in \mathcal{P}(\mathbb{R}^n)$ to $\nu \in \mathcal{P}(\mathbb{R}^m)$ if $\nu(A) = \mu(\mathbf{t}^{-1}(A))$ for all measurable sets A. We call \mathbf{t} a transport map and denote ν as $\mathbf{t}_{\#}\mu \in \mathcal{P}(\mathbb{R}^m)$, the pushforward of μ through \mathbf{t} . For $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$, the set of transport plans from μ to ν is given by,

$$\Gamma(\mu,\nu) := \{ \boldsymbol{\gamma} \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d) \mid \pi^1_{\#} \boldsymbol{\gamma} = \mu, \, \pi^2_{\#} \boldsymbol{\gamma} = \nu \},$$

where $\pi^1, \pi^2 \colon \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$ are the projections of $\mathbb{R}^d \times \mathbb{R}^d$ onto the first and second copy of \mathbb{R}^d , respectively. The Wasserstein distance [3, Chapter 7] between $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ is given by,

(2.3)
$$W_2(\mu,\nu) = \min_{\boldsymbol{\gamma}\in\Gamma(\mu,\nu)} \left(\int_{\mathbb{R}^d\times\mathbb{R}^d} |x-y|^2 d\boldsymbol{\gamma}(x,y) \right)^{1/2}$$

We say that a transport plan γ is *optimal* if it attains the minimum in (2.3). We denote the set of optimal transport plans by $\Gamma_0(\mu, \nu)$.

Convergence with respect to the Wasserstein metric is stronger than narrow convergence of probability measures [3, Remark 7.1.11]. In particular, if μ_n is a sequence in $\mathcal{P}_2(\mathbb{R}^d)$ and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, we have,

(2.4)
$$W_2(\mu_n, \mu) \to 0 \text{ as } n \to \infty \iff (\mu_n \to \mu \text{ narrowly and } M_2(\mu_n) \to M_2(\mu) \text{ as } n \to \infty)$$

In order to define Wasserstein gradient flows, we require the following notion of regularity in time with respect to the Wasserstein metric. **Definition 2.4** (absolutely continuous). We say $\mu : [0,T] \to \mathcal{P}(\mathbb{R}^d)$ is absolutely continuous on [0,T], and write $\mu \in AC^2_{loc}((0,T); P_2(\mathbb{R}^d))$, if there exists $f \in L^2_{loc}((0,T))$ so that,

(2.5)
$$W_2(\mu(t),\mu(s)) \le \int_s^t f(r) \, dr \quad \text{for all } t,s \in (0,T) \text{ with } s \le t.$$

Along such curves, we may define the metric derivative.

Definition 2.5 (metric derivative). Given $\mu \in AC^2_{loc}((0,T); \mathcal{P}_2(\mathbb{R}^d))$, the limit,

$$|\mu'|(t) := \lim_{s \to t} \frac{W_2(\mu(t), \mu(s))}{|t - s|}$$

exists for a.e. $t \in (0,T)$ and is called the *metric derivative* of μ .

In fact, the metric derivative is the minimal square integrable function satisfying (2.5): for any $\mu \in AC^2_{loc}((0,T); \mathcal{P}_2(\mathbb{R}^d))$, we have $|\mu'| \in L^2((0,T))$, and for any function f satisfying (2.5), we have $|\mu'|(t) \leq f(t)$ for a.e. $t \in (0,T)$ (see [3, Theorem 1.1.2]).

Geodesics form an important class of curves in the Wasserstein metric. Given $\mu_0, \mu_1 \in \mathcal{P}_2(\mathbb{R}^d)$, the geodesics connecting μ_0 to μ_1 are the curves of the form,

(2.6)
$$\mu_{\alpha} = ((1-\alpha)\pi^1 + \alpha\pi^2)_{\#} \boldsymbol{\gamma} \quad \text{for } \alpha \in [0,1], \, \boldsymbol{\gamma} \in \Gamma_0(\mu,\nu).$$

More generally, given $\mu_1, \mu_2, \mu_3 \in \mathcal{P}_2(\mathbb{R}^d)$, a generalized geodesic from μ_2 to μ_3 with base μ_1 is given by,

(2.7)
$$\mu_{\alpha}^{2 \to 3} = \left((1 - \alpha)\pi^2 + \alpha\pi^3 \right)_{\#} \gamma \quad \text{for } \alpha \in [0, 1] \text{ and } \gamma \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d)$$

such that
$$\pi^{1,2}_{\#} \boldsymbol{\gamma} \in \Gamma_0(\mu_1, \mu_2)$$
 and $\pi^{1,3}_{\#} \boldsymbol{\gamma} \in \Gamma_0(\mu_1, \mu_3)$

with $\pi^{1,i}: \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d \times \mathbb{R}^d$ the projection of onto the first and *i*th copies of \mathbb{R}^d . Note that when the base μ_1 coincides with one of the endpoints μ_2 or μ_3 , a generalized geodesic is a geodesic.

A key property for the uniqueness and stability of Wasserstein gradient flows is convexity, or more generally semiconvexity, along generalized geodesics.

Definition 2.6 (semiconvexity). A functional $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, \infty]$ is semiconvex along generalized geodesics if there exists $\lambda \in \mathbb{R}$ such that, for all $\mu_1, \mu_2, \mu_3 \in \mathcal{P}_2(\mathbb{R}^d)$, there exists a generalized geodesic from μ_2 to μ_3 with base μ_1 for which the following inequality holds:

(2.8)
$$\mathcal{G}(\mu_{\alpha}^{2\to3}) \leq (1-\alpha)\mathcal{G}(\mu_2) + \alpha \mathcal{G}(\mu_3) - \alpha(1-\alpha)\frac{\lambda}{2}W_{2,\gamma}^2(\mu_2,\mu_3) \text{ for all } \alpha \in [0,1],$$

where,

$$W_{2,\boldsymbol{\gamma}}^2(\mu_2,\mu_3) := \int_{\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d} |\pi^2 - \pi^3|^2 \, d\boldsymbol{\gamma}(x,y,z).$$

In this case, will sometimes say the functional is λ -convex. If a functional is 0-convex, we will say it is convex.

We recall the following sufficient condition for convexity, which is the Wasserstein analogue of the "above the tangent line" characterization of convexity from finite dimensional Euclidean space.

Lemma 2.7 (above the tangent line property [29, Proposition 2.8]). A functional $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, \infty]$ is λ -convex along generalized geodesics if and only if for all generalized geodesics $\mu_{\alpha}^{2\to3}$ connecting μ_2 to μ_3 with base μ_1 , the map $\alpha \mapsto \mathcal{G}(\mu_{\alpha}^{2\to3})$ is differentiable for all $\alpha \in [0,1]$ and,

$$\mathcal{G}(\mu_3) - \mathcal{G}(\mu_2) - \left. \frac{d}{d\alpha} \mathcal{G}(\mu_\alpha) \right|_{\alpha=0} \ge \frac{\lambda}{2} W_{2,\gamma}^2(\mu_2,\mu_3).$$

For any functional $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, +\infty]$, we denote its *domain* by $D(\mathcal{G}) = \{\mu \in \mathcal{P}_2(\mathbb{R}^d) \mid \mathcal{G}(\mu) < +\infty\}$, and say that \mathcal{G} is *proper* if $D(\mathcal{G}) \neq \emptyset$. For any measure μ in the domain of a functional \mathcal{G} , we may define the local slope of \mathcal{G} at μ as follows.

Definition 2.8 (local slope). Given $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, \infty]$, for any $\mu \in D(\mathcal{G})$, the local slope is,

$$|\partial \mathcal{G}|(\mu) = \limsup_{\nu \to \mu} \frac{(\mathcal{G}(\mu) - \mathcal{G}(\nu))_+}{W_2(\mu, \nu)},$$

where $(s)_{+} = \max\{s, 0\}$ denotes the positive part of s.

Next, we define the subdifferential of a functional $\mathcal{G} : \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, +\infty]$ that is lower semicontinuous with respect to Wasserstein convergence and λ -convex along generalized geodesics.¹

Definition 2.9 (subdifferential of λ -convex functional). Suppose $\mathcal{G} : \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, +\infty]$ is proper, lower semicontinuous, and λ -convex along geodesics. Let $\mu \in D(\mathcal{G})$ and $\boldsymbol{\xi} : \mathbb{R}^d \to \mathbb{R}^d$ with $\boldsymbol{\xi} \in L^2(d\mu)$. We say that $\boldsymbol{\xi}$ belongs to the *subdifferential* of \mathcal{G} at μ , and write $\boldsymbol{\xi} \in \partial \mathcal{G}(\mu)$, if for all $\nu \in \mathcal{P}_2(\mathbb{R}^d)$,

(2.9)
$$\mathcal{G}(\nu) - \mathcal{G}(\mu) \ge \int_{\mathbb{R}^d \times \mathbb{R}^d} \langle \boldsymbol{\xi}(x), y - x \rangle \, d\boldsymbol{\gamma}(x, y) + \frac{\lambda}{2} W_2^2(\mu, \nu) \quad \text{for all } \boldsymbol{\gamma} \in \Gamma_0(\mu, \nu).$$

Remark 2.10 (subdifferential of sum). Note that if \mathcal{G}_1 and \mathcal{G}_2 satisfy the hypotheses of Definition 2.9 and $\mu \in D(\mathcal{G}_1) \cap D(\mathcal{G}_2)$, then for any $\boldsymbol{\xi}_1 \in \partial \mathcal{G}_1(\mu)$ and $\boldsymbol{\xi}_2 \in \partial \mathcal{G}_2(\mu)$, we have $\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2 \in \partial (\mathcal{G}_1 + \mathcal{G}_2)(\mu)$.

The local slope and subdifferential are related by the following proposition, which is a direct adaptation of [3, Lemma 10.1.5] to the case of functionals which contain measures μ in their domain that are not necessarily absolutely continuous with respect to Lebesgue measure. We defer the proof to appendix A.

Proposition 2.11 (local slope and minimal subdifferential). Suppose $\mathcal{G} : \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, +\infty]$ is proper, lower semicontinuous, and λ -convex along generalized geodesics. Then for any $\mu \in D(|\partial \mathcal{G}|)$, we have,

(2.10)
$$|\partial \mathcal{G}|(\mu) \le \inf \left\{ \|\boldsymbol{\xi}\|_{L^2(\mu)} : \boldsymbol{\xi} \in \partial \mathcal{G}(\mu) \right\}.$$

If equality holds and $\boldsymbol{\xi}$ attains the infimum, we will write $\boldsymbol{\xi} = \partial^{\circ} \mathcal{G}(\mu)$. In this case, the element of the subdifferential attaining the infimum is unique.

We now turn to the definition of a gradient flow in the Wasserstein metric (c.f. [3, Definition 1.1.1, Proposition 8.3.1, Definition 11.1.1, Theorem 11.1.3]).

Definition 2.12 (gradient flow). Suppose $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, +\infty]$ is proper, lower semicontinous, and λ -convex along generalized geodesics. A curve $\mu(t) \in AC^2_{loc}((0, +\infty); \mathcal{P}_2(\mathbb{R}^d))$ is a gradient flow of \mathcal{G} in the Wasserstein metric if $\mu(t)$ is a weak solution of the continuity equation,

(2.11) $\partial_t \mu(t) + \nabla \cdot (v(t)\mu(t)) = 0$, in duality with $C_c^{\infty}((0, +\infty) \times \mathbb{R}^d)$,

and,

$$v(t) = -\partial^{\circ} \mathcal{G}(\mu(t))$$
 for \mathcal{L}^1 -a.e. $t > 0$.

Next, we recall sufficient conditions for well-posedness of the initial value problem for the gradient flow, when the initial condition $\mu(0)$ is in the closure of the domain of the energy $\overline{D(\mathcal{G})}$. We also recall equivalent characterizations of the gradient flow as a curve of maximal slope and evolution variational inequality. As the theorem is simply a collection of general results developed by Ambrosio, Gigli, and Savaré [3], we defer its proof to Appendix A.

Theorem 2.13 (well-posedness and characterization of gradient flow). Suppose $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \to (-\infty, +\infty]$ is proper, lower semicontinuous, and λ -convex along generalized geodesics and $\mu(0) \in \overline{D(\mathcal{G})}$. Then, there exists a unique gradient flow $\mu(t)$ of \mathcal{G} satisfying $\lim_{t\to 0^+} \mu(t) = \mu(0)$ in the Wasserstein metric.

Furthermore $\mu(t) \in AC^2_{loc}((0, +\infty); \mathcal{P}_2(\mathbb{R}^d))$ is the gradient flow of \mathcal{G} if and only if $\mu(t)$ satisfies one of the following equivalent conditions:

(i) Curve of Maximal Slope: For all $0 < s \le t$.

$$\frac{1}{2}\int_{s}^{t}|\mu'|^{2}(r)dr + \frac{1}{2}\int_{s}^{t}|\partial\mathcal{G}|^{2}(\mu(r))dr = \mathcal{G}(\mu(s)) - \mathcal{G}(\mu(t)).$$

(ii) Evolution Variational Inequality: For all $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and for \mathcal{L}^1 -a.e. $t \geq 0$,

$$\frac{1}{2}\frac{d^{+}}{dt}W_{2}^{2}(\mu(t),\nu) + \frac{\lambda}{2}W_{2}^{2}(\mu(t),\nu) + \mathcal{G}(\mu(t)) \leq \mathcal{G}(\nu).$$

¹Note that in Ambrosio, Gigli, and Savaré [3, Chapter 10] this is known as the *reduced subdifferential*, which is stronger than their notion of *extended subdifferential*: see Definition 10.3.1 of the extended subdifferential and equations (10.3.12)-(10.3.13) for the reduced subdifferential. The reduced subdifferential is sufficient for our purposes, due to the fact that our main Γ -convergence result considers gradient flow solutions that are absolutely continuous with respect to Lebesgue measure, and we extend the convergence to particle initial data separately.

2.4. Γ -convergence of energies and gradient flows. We now recall the general framework of Γ -convergence of energies, which is a classical tool in the Calculus of Variations, and Γ -convergence of gradient flows, as introduced by Serfaty [69]. The former provides sufficient conditions that, when combined with some compactness, ensure that minimizers of a sequence of energies converge to a minimizer of a limiting energy. The latter provides sufficient conditions that, again with sufficient compactness, ensure that gradient flows of a sequence of energies converge to a gradient flow of a limiting energy.

We begin by recalling the notion of Γ -convergence of energies, focusing in particular on the case of energies defined on $\mathcal{P}(\mathbb{R}^d)$, with respect to the narrow topology.

Definition 2.14 (Γ -convergence of energies). A sequence of functionals $\mathcal{G}_{\varepsilon} : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R} \cup \{+\infty\}$ is said to Γ -converge to $\mathcal{G} : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R} \cup \{+\infty\}$ if:

(2.12) For any sequence $\rho_{\varepsilon} \in \mathcal{P}(\mathbb{R}^d)$ converging narrowly to $\rho \in \mathcal{P}(\mathbb{R}^d)$, $\liminf_{\varepsilon \to 0} \mathcal{G}_{\varepsilon}(\rho_{\varepsilon}) \geq \mathcal{G}(\rho)$;

(2.13) For any $\rho \in \mathcal{P}_2(\mathbb{R}^d)$, there exists $\rho_{\varepsilon} \in \mathcal{P}(\mathbb{R}^d)$ converging narrowly to ρ , s.t. $\limsup \mathcal{G}_{\varepsilon}(\rho) \leq \mathcal{G}(\rho)$.

Next, we recall the notion of Γ -convergence of gradient flows. We state a version of [69, Theorem 2] that has been specialized to the present case, in which we consider functionals defined on $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ that are lower semicontinuous and semiconvex along generalized geodesics. In this case, the notions of "gradient flow" and "curve of maximal slope" are equivalent; see Theorem 2.13. Likewise, metric slopes are strong upper gradients [3, Corollary 2.4.10].

Theorem 2.15 (Γ -convergence of gradient flows, [69, Theorem 2]). Let $\mathcal{G}_{\varepsilon} : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R} \cup \{+\infty\}$ and $\mathcal{G} : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R} \cup \{+\infty\}$ be proper, lower semicontinuous functionals that are semiconvex along generalized geodesics. Furthermore, suppose that (2.12) holds. For all $\varepsilon > 0$, let $\rho_{\varepsilon} \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ be a gradient flow of $\mathcal{G}_{\varepsilon}$, and suppose that there exists a curve $\rho : [0,T] \to \mathcal{P}_2(\mathbb{R}^d)$ such that,

(2.14) $\rho_{\varepsilon}(t)$ narrowly converges to $\rho(t)$ for $t \in [0, T]$, and

(2.15)
$$\rho(0) \in D(\mathcal{G}) \text{ and } \lim_{\varepsilon \to 0} \mathcal{G}_{\varepsilon}(\rho_{\varepsilon}(0)) = \mathcal{G}(\rho(0)).$$

Moreover, suppose, for almost every $t \in [0, T]$,

(2.16)
$$\liminf_{\varepsilon \to 0} \int_0^t |\rho_\varepsilon'|^2(s) \, ds \ge \int_0^t |\rho'|(s) \, ds$$

(2.17)
$$\liminf_{\varepsilon \to 0} |\partial \mathcal{G}_{\varepsilon}|^2(\rho_{\varepsilon}(t)) \ge |\partial \mathcal{G}|^2(\rho(t))$$

Then $\rho \in AC^2([0,T], \mathcal{P}_2(\mathbb{R}^d))$, and ρ is a gradient flow of \mathcal{G} with initial data $\rho(0)$.

In fact, the existence of a subsequential limit of ρ_{ε} pointwise in time, as required in (2.14), as well as the lower semicontinuity of the metric derivatives, as required in (2.16) above, are both guaranteed under the following additional assumption.

Lemma 2.16 (existence of a narrowly convergent subsequence). Let $\mathcal{G}_{\varepsilon} : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R} \cup \{+\infty\}$ be proper, lower semicontinuous functionals that are semiconvex along generalized geodesics. For all $\varepsilon > 0$, let $\rho_{\varepsilon} \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ be a gradient flow of $\mathcal{G}_{\varepsilon}$. Suppose,

(2.18)
$$\sup_{\varepsilon>0} \mathcal{G}_{\varepsilon}(\rho_{\varepsilon}(0)) < +\infty \text{ and}$$

(2.19)
$$\inf_{\varepsilon > 0, t \in [0,T]} \mathcal{G}_{\varepsilon}(\rho_{\varepsilon}(t)) > -\infty.$$

Then there exists $\rho \in AC^2([0,T], \mathcal{P}_2(\mathbb{R}^d))$ such that, along a subsequence, (2.14) and (2.16) hold.

Proof. We shall produce the desired ρ and subsequence ρ_{ε} such that (2.14) holds. Once that is done, [31, Lemma 2.3] immediately implies that (2.16) holds.

To this end, we use Proposition A.3 to find that, for all $\varepsilon > 0$ and $t \in [0, T]$, there exists C > 0 so that $\rho_{\varepsilon}(t)$ belongs to the set $\{\rho : M_2(\rho) < C\}$, which is narrowly sequentially compact [3, Remark 5.1.5, Lemma 5.1.7]. Furthermore, item (i) of Theorem 2.13, together with assumptions (2.18) and (2.19), ensures,

(2.20)
$$\sup_{\varepsilon>0} \frac{1}{2} \int_0^T |\rho_{\varepsilon}'|^2(r) dr \le \sup_{\varepsilon>0} \left(\mathcal{G}_{\varepsilon}(\rho_{\varepsilon}(0)) - \mathcal{G}_{\varepsilon}(\rho_{\varepsilon}(T)) \right) < +\infty$$

From this, we deduce the equicontinuity:

$$\sup_{\varepsilon>0} W_2(\rho_\varepsilon(s),\rho_\varepsilon(t)) \le \sup_{\varepsilon>0} \int_s^t |\rho_\varepsilon'|(r)dr \le \sqrt{t-s} \left(\int_s^t |\rho_\varepsilon'|^2(r)dr\right)^{1/2} \le \sqrt{t-s} \left(\sup_{\varepsilon>0} \int_0^T |\rho_\varepsilon'|^2(r)\right)^{1/2}.$$

Therefore, the generalized Ascoli-Arzelá/Aubin-Lions theorem [3, Proposition 3.3.1] ensures that there exists $\rho : [0, +\infty) \to \mathcal{P}_2(\mathbb{R}^d)$ so that, up to a subsequence, $\rho_{\varepsilon}(t) \to \rho(t)$ narrowly, for all $t \ge 0$.

3. GRADIENT FLOWS OF ENERGIES WITH REGULARIZATION AND CONFIMENT

We now prove several fundamental properties of the internal energy \mathcal{E} and the regularized internal energy $\mathcal{E}_{\varepsilon}$, with the addition of external potential energies, \mathcal{V} and $\mathcal{V}_{\varepsilon}$, as well as the confining energies, \mathcal{V}_k and \mathcal{V}_{Ω} . In particular, will will characterize their lower semicontinuity, convexity, and subdifferentiability. Each of these properties provides information about the one-sided regularity of the energy functional, its first derivative, and its second derivative with respect to the Wasserstein metric. Since our study of gradient flows only considers well-posedness of the flow forward in time (which is natural given that our motivating equation is a diffusion equation), these one-sided estimates on the energy functionals' regularity are sufficient for our analysis. We will close the section by applying these properties to characterize the gradient flows of these energies in terms of partial differential equations.

3.1. Fundamental properties of energies. First, we recall that the functionals \mathcal{E} and $\mathcal{E}_{\varepsilon}$ are lower semicontinuous with respect to narrow convergence. Since narrow convergence is weaker than Wasserstein convergence, this in turn implies lower semicontinuity with respect to Wasserstein convergence. The proof of this result is standard, and we defer it to appendix B.

Lemma 3.1 (lower semicontinuity of \mathcal{E} and $\mathcal{E}_{\varepsilon}$). Suppose Assumptions (T) and (M) are satisfied. Then, for all $\varepsilon > 0$, the functionals \mathcal{E} and $\mathcal{E}_{\varepsilon}$ are lower semicontinuous with respect to narrow convergence.

The lower semicontinuity of the external potential energies, \mathcal{V} and $\mathcal{V}_{\varepsilon}$, and the confining energies, \mathcal{V}_k and \mathcal{V}_{Ω} , with respect to narrow convergence is an immediate consequence of the Portmanteau theorem, see e.g. [3, Lemma 5.1.7], since they all are obtained by integrating a function that is lower semicontinuous and bounded below against ρ .

Lemma 3.2 (lower semicontinuity of \mathcal{V} , $\mathcal{V}_{\varepsilon}$, \mathcal{V}_k , \mathcal{V}_{Ω}). Under Assumptions (M), (D), (V), and (C), the energies \mathcal{V} , $\mathcal{V}_{\varepsilon}$, \mathcal{V}_k , and \mathcal{V}_{Ω} are lower semicontinuous with respect to narrow convergence.

The convexity of the energies \mathcal{E} , \mathcal{V} , $\mathcal{V}_{\varepsilon}$, \mathcal{V}_k , and \mathcal{V}_{Ω} follows immediately from the theory developed by Ambrosio, Gigli, and Savaré [3]. We recall these results in the following proposition. The proof of this proposition is an immediate consequence of existing theory, so we defer it to appendix B.

Proposition 3.3 (convexity properties of of \mathcal{E} , \mathcal{V} , $\mathcal{V}_{\varepsilon}$, \mathcal{V}_k , and \mathcal{V}_{Ω}).

- (i) Suppose $\bar{\rho}$ satisfies Assumption (T) and is log-concave. Then \mathcal{E} is convex along generalized geodesics.
- (ii) Suppose V and V_k satisfy Assumptions (V) and (C). Then \mathcal{V} and $\mathcal{V}_{\varepsilon}$ are λ -convex along generalized geodesics, for $\lambda = \inf_{\{x,\xi \in \mathbb{R}^d\}} \xi^t D^2 V(x)\xi$, and \mathcal{V}_k is convex along generalized geodesics.
- (iii) Suppose $\Omega \subset \mathbb{R}^d$ satisfies Assumption (D). Then \mathcal{V}_{Ω} is convex along generalized geodesics.

We now aim to show that $\mathcal{E}_{\varepsilon}$ is also semiconvex for all $\varepsilon > 0$. In order to accomplish this, we begin by characterizing the directional derivative of $\mathcal{E}_{\varepsilon}$. For the reader's convenience, we also recall the directional derivatives of the external potential energies \mathcal{V} , $\mathcal{V}_{\varepsilon}$, and \mathcal{V}_k , which have been studied extensively in previous works; see, for example, [3, Proposition 10.4.2].

Proposition 3.4 (directional derivatives of $\mathcal{E}_{\varepsilon}$, \mathcal{V} , $\mathcal{V}_{\varepsilon}$, and \mathcal{V}_k). Suppose Assumptions (T), (M), (V) and (C) hold. Fix $\varepsilon > 0$, $\nu_1, \nu_2, \nu_3 \in \mathcal{P}_2(\mathbb{R}^d)$, and $\gamma \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d)$ with $\pi^i_{\#} \gamma = \nu_i$. Consider the curve,

$$\mu_{\alpha} = \left((1 - \alpha)\pi^2 + \alpha\pi^3 \right)_{\#} \boldsymbol{\gamma} \text{ for } \alpha \in [0, 1].$$

Then,

$$\begin{aligned} \left. \frac{d}{d\alpha} \mathcal{E}_{\varepsilon}(\mu_{\alpha}) \right|_{\alpha=0} &= \frac{1}{2} \int \frac{\zeta_{\varepsilon} * \nu_{2}(x)}{\bar{\rho}(x)} \int \left\langle \nabla \zeta_{\varepsilon} \left(x - y_{2} \right), y_{3} - y_{2} \right\rangle d\gamma(y_{1}, y_{2}, y_{3}) \ dx, \\ \left. \frac{d}{d\alpha} \mathcal{V}(\mu_{\alpha}) \right|_{\alpha=0} &= \int \left\langle \nabla V(y_{2}), y_{3} - y_{2} \right\rangle d\gamma(y_{1}, y_{2}, y_{3}), \\ \left. \frac{d}{d\alpha} \mathcal{V}_{\varepsilon}(\mu_{\alpha}) \right|_{\alpha=0} &= \int \left\langle \nabla (\zeta_{\varepsilon} * V)(y_{2}), y_{3} - y_{2} \right\rangle d\gamma(y_{1}, y_{2}, y_{3}), \\ \left. \frac{d}{d\alpha} \mathcal{V}_{k}(\mu_{\alpha}) \right|_{\alpha=0} &= \int \left\langle \nabla V_{k}(y_{2}), y_{3} - y_{2} \right\rangle d\gamma(y_{1}, y_{2}, y_{3}). \end{aligned}$$

Remark 3.5. Remark: Note that if $\gamma \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d)$ satisfies the hypotheses in the definition of generalized geodesic (2.7), then γ satisfies the assumptions of Proposition 3.4.

Proof. We begin with the characterization of the directional derivative $\frac{d}{d\alpha} \mathcal{E}_{\varepsilon}(\mu_{\alpha})|_{\alpha=0}$. As a first step in this direction, we estimate $\frac{d}{d\alpha} \zeta_{\varepsilon} * \mu_{\alpha}|_{\alpha=0}$. For all $x \in \mathbb{R}^d$ and $\alpha \in [0, 1]$,

(3.1)
$$\frac{1}{\alpha} \left(\zeta_{\varepsilon} * \mu_{\alpha}(x) - \zeta_{\varepsilon} * \mu_{0}(x) \right) = \int \frac{1}{\alpha} \left[\zeta_{\varepsilon}(x - ((1 - \alpha)y_{2} + \alpha y_{3})) - \zeta_{\varepsilon}(x - y_{2}) \right] d\gamma(y_{1}, y_{2}, y_{3}).$$

By the mean value theorem for ζ_{ε} , we may bound the integrand by,

(3.2)
$$\frac{1}{\alpha} \|\nabla \zeta_{\varepsilon}\|_{\infty} \left| ((1-\alpha)y_2 + \alpha y_3) - y_2 \right| \le \|\nabla \zeta_{\varepsilon}\|_{\infty} |y_3 - y_2| \in L^1(\gamma),$$

where the integrability holds since $M_1(\gamma) \leq M_2(\gamma)^{1/2} = (M_2(\nu_1) + M_2(\nu_2) + M_2(\nu_3))^{1/2} < +\infty$. Thus, by the dominated convergence theorem,

Now, we use this to compute $\left. \frac{d}{d\alpha} \mathcal{E}_{\varepsilon}(\mu_{\alpha}) \right|_{\alpha=0}$. First, note that we may express the difference quotient as,

$$(3.4) \qquad \frac{1}{\alpha} \left(\mathcal{E}_{\varepsilon}(\mu_{\alpha}) - \mathcal{E}_{\varepsilon}(\mu_{0}) \right) = \frac{1}{2\alpha} \int \left(\left(\zeta_{\varepsilon} * \mu_{\alpha} \right)^{2}(x) - \left(\zeta_{\varepsilon} * \mu_{0} \right)^{2}(x) \right) \bar{\rho}(x)^{-1} dx \\ = \int \frac{1}{2\alpha} \left[\left(\zeta_{\varepsilon} * \mu_{\alpha} \right)(x) + \left(\zeta_{\varepsilon} * \mu_{0} \right)(x) \right] \left[\left(\zeta_{\varepsilon} * \mu_{\alpha} \right)(x) - \left(\zeta_{\varepsilon} * \mu_{0} \right)(x) \right] \bar{\rho}(x)^{-1} dx.$$

By equations (3.1-3.2) and the fact that $\bar{\rho}$ is uniformly bounded below, the integrand is dominated by,

$$g_{\alpha}(x) := C\left[(\zeta_{\varepsilon} * \mu_{\alpha})(x) + (\zeta_{\varepsilon} * \mu_{0})(x) \right] \quad \text{for} \quad C = \|\nabla \varphi_{\varepsilon}\|_{\infty} \|\bar{\rho}^{-1}\|_{\infty} M_{1}(\gamma).$$

The narrow convergence of μ_{α} to μ_0 as $\alpha \to 0$, and the fact that ζ_{ε} is bounded and continuous ensures that $g_{\alpha}(x) \to 2C(\zeta_{\varepsilon} * \mu_0)(x)$ pointwise. Furthermore,

$$\lim_{\alpha \to 0} \int g_{\alpha}(x) dx = \lim_{\alpha \to 0} C \int \zeta_{\varepsilon} * \mu_{\alpha}(x) dx + C \int \zeta_{\varepsilon} * \mu_{0}(x) dx = 2C$$

Therefore, by the generalized dominated convergence theorem [67, Chapter 4, Theorem 19] and equations (3.3) and (3.4),

$$\lim_{\alpha \to 0} \frac{1}{\alpha} \left(\mathcal{E}_{\varepsilon}(\mu_{\alpha}) - \mathcal{E}_{\varepsilon}(\mu_{0}) \right) = \int \lim_{\alpha \to 0} \frac{1}{2\alpha} \left[\left(\zeta_{\varepsilon} * \mu_{\alpha} \right)(x) + \left(\zeta_{\varepsilon} * \mu_{0} \right)(x) \right] \left[\left(\zeta_{\varepsilon} * \mu_{\alpha} \right)(x) - \left(\zeta_{\varepsilon} * \mu_{0} \right)(x) \right] \bar{\rho}(x)^{-1} dx$$
$$= \int \frac{\left(\zeta_{\varepsilon} * \mu_{0} \right)(x)}{\bar{\rho}(x)} \int \left\langle \nabla \zeta_{\varepsilon}(x - y_{2}), y_{3} - y_{2} \right\rangle d\gamma(y_{1}, y_{2}, y_{3}) dx.$$

Next we consider the directional derivative $\frac{d}{d\alpha} \mathcal{V}(\mu_{\alpha})\Big|_{\alpha=0}$. By definition of \mathcal{V} and μ_{α} ,

$$\lim_{\alpha \to 0} \frac{1}{\alpha} \left(\mathcal{V}(\mu_{\alpha}) - \mathcal{V}(\mu_{0}) \right) = \lim_{\alpha \to 0} \int \frac{1}{\alpha} \left[V \left((1 - \alpha)y_{2} + \alpha y_{3} \right) - V(y_{2}) \right] d\gamma(y_{1}, y_{2}, y_{3}).$$

By the mean value theorem for V, we may bound the integrand by,

$$\frac{1}{\alpha} \|\nabla V\|_{\infty} |((1-\alpha)y_2 + \alpha y_3) - y_2| \le \|\nabla V\|_{\infty} |y_3 - y_2| \in L^1(\gamma)$$

Thus, by the dominated convergence theorem,

$$\lim_{\alpha \to 0} \frac{1}{\alpha} \left(\mathcal{V}(\mu_{\alpha}) - \mathcal{V}(\mu_{0}) \right) = \int \lim_{\alpha \to 0} \frac{1}{\alpha} \left[V \left((1 - \alpha)y_{2} + \alpha y_{3} \right) - V(y_{2}) \right] d\gamma(y_{1}, y_{2}, y_{3})$$
$$= \int \left\langle \nabla V(y_{2}), y_{3} - y_{2} \right\rangle d\gamma(y_{1}, y_{2}, y_{3}),$$

which gives the result. The result for $\mathcal{V}_{\varepsilon}$ follows exactly as above, replacing V with $(\zeta_{\varepsilon} * V)$.

Finally, we consider the directional derivative of \mathcal{V}_k . By definition of \mathcal{V}_k and μ_{α} and the assumption that $V_k \in C^2$ with $\|D^2 V_k\|_{\infty} < +\infty$, we may apply the Fundamental Theorem of Calculus to conclude,

$$\begin{split} \lim_{\alpha \to 0} \frac{1}{\alpha} \left(\mathcal{V}_k(\mu_{\alpha}) - \mathcal{V}_k(\mu_0) \right) \\ &= \lim_{\alpha \to 0} \frac{1}{\alpha} \int_{\mathbb{R}^d} \left[V_k \left((1 - \alpha) y_2 + \alpha y_3 \right) - V_k(y_2) \right] d\gamma(y_1, y_2, y_3), \\ &= \lim_{\alpha \to 0} \frac{1}{\alpha} \int_{\mathbb{R}^d} \int_0^\alpha \int_0^\beta (y_3 - y_2)^t D^2 V_k((1 - s) y_2 + s y_3)(y_3 - y_2) \, ds \, d\beta + \alpha \left\langle \nabla V_k(y_2), y_3 - y_2 \right\rangle d\gamma(y_1, y_2, y_3) \\ &= \int_{\mathbb{R}^d} \left\langle \nabla V_k(y_2), y_3 - y_2 \right\rangle d\gamma(y_1, y_2, y_3), \end{split}$$

where the first term vanishes since $D^2 V_k \in L^{\infty}(\mathbb{R}^d)$ and $\int |y_3 - y_2|^2 d\gamma \leq 2(M_2(\nu_1) + M_2(\nu_2)) < +\infty$.

Using this characterization of the directional derivative of $\mathcal{E}_{\varepsilon}$, we now prove that our energy $\mathcal{E}_{\varepsilon}$ is λ_{ε} -convex along generalized geodesics, where $\lambda_{\varepsilon} \xrightarrow{\varepsilon \to 0} -\infty$.

Proposition 3.6 (semiconvexity of $\mathcal{E}_{\varepsilon}$). Suppose Assumptions (T) and (M) hold. For all $\varepsilon > 0$, the functional $\mathcal{E}_{\varepsilon}$ is λ_{ε} -convex along generalized geodesics, where,

(3.5)
$$\lambda_{\varepsilon} = -\|1/\bar{\rho}\|_{L^{\infty}(\mathbb{R}^d)} \|D^2\zeta_{\varepsilon}\|_{L^{\infty}(\mathbb{R}^d)} = -\varepsilon^{-d-2} ||1/\bar{\rho}||_{L^{\infty}(\mathbb{R}^d)} ||D^2\zeta||_{L^{\infty}(\mathbb{R}^d)}$$

Proof. Let $(\mu_{\alpha}^{2\to 3})_{\alpha\in[0,1]}$ be a generalized geodesic with base $\mu_1 \in \mathcal{P}_2(\mathbb{R}^d)$ connecting two probability measures $\mu_2, \mu_3 \in \mathcal{P}_2(\mathbb{R}^d)$, and let $\gamma \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d)$ be the associated measure as defined in (2.7). Since $x \mapsto x^2$ is a convex function, using the above the tangent inequality for convex functions yields,

$$\mathcal{E}_{\varepsilon}(\mu_{3}) - \mathcal{E}_{\varepsilon}(\mu_{2}) = \frac{1}{2} \int \frac{\left(\zeta_{\varepsilon} * \mu_{3}(x)\right)^{2}}{\bar{\rho}(x)} dx - \frac{1}{2} \int \frac{\left(\zeta_{\varepsilon} * \mu_{2}(x)\right)^{2}}{\bar{\rho}(x)} dx$$
$$\geq \int \frac{\zeta_{\varepsilon} * \mu_{2}(x)}{\bar{\rho}(x)} \left(\zeta_{\varepsilon} * \mu_{3}(x)\right) - \zeta_{\varepsilon} * \mu_{2}(x)\right) dx$$
$$= \int \frac{\zeta_{\varepsilon} * \mu_{2}(x)}{\bar{\rho}(x)} \left(\iiint \zeta_{\varepsilon}(x - y_{3}) - \zeta_{\varepsilon}(x - y_{2})\right) d\gamma(y_{1}, y_{2}, y_{3}) dx$$

Therefore, by Proposition 3.4,

$$\begin{split} \mathcal{E}_{\varepsilon}(\mu_{3}) &- \mathcal{E}_{\varepsilon}(\mu_{2}) - \left. \frac{d}{d\alpha} \mathcal{E}_{\varepsilon}(\mu_{\alpha}^{2 \to 3}) \right|_{\alpha=0} \\ &\geq \int \frac{\zeta_{\varepsilon} * \mu_{2}(x)}{\bar{\rho}(x)} \left(\iiint \zeta_{\varepsilon}(x-y_{3}) - \zeta_{\varepsilon}(x-y_{2}) - \left\langle \nabla \zeta_{\varepsilon}(x-y_{2}), y_{2} - y_{3} \right\rangle d\gamma(y_{1}, y_{2}, y_{3}) \right) dx \\ &\geq - \frac{\|D^{2}\zeta_{\varepsilon}\|_{L^{\infty}(\mathbb{R}^{d})}}{2} \int \frac{\zeta_{\varepsilon} * \mu_{2}(x)}{\bar{\rho}(x)} dx \iiint |y_{2} - y_{3}|^{2} d\gamma(y_{1}, y_{2}, y_{3}) dx \\ &\geq - \frac{\|1/\bar{\rho}\|_{L^{\infty}(\mathbb{R}^{d})} \|D^{2}\zeta_{\varepsilon}\|_{L^{\infty}(\mathbb{R}^{d})}}{2} W_{2,\gamma}^{2}(\mu_{2}, \mu_{3}), \end{split}$$

where we have applied Young's inequality to conclude that $\|\zeta_{\varepsilon} * \mu_2\|_{L^1(\mathbb{R}^d)} = 1$. By Lemma 2.7, this gives the result.

19

The preceding results ensure that our energies $\mathcal{E}, \mathcal{E}_{\varepsilon}, \mathcal{V}, \mathcal{V}_{\varepsilon}, \mathcal{V}_{k}$, and \mathcal{V}_{Ω} are proper, lower semicontinuous, and semiconvex along generalized geodesics. Thus, the gradient flows of each of their energies, as well as the sum of any of the energies, is well posed, by Theorem 2.13, for any initial conditions in the closure of their domains. However, in order to characterize these gradient flows in terms of partial differential equations and prove our main Γ -convergence result, we must now characterize the minimal elements of their subdifferentials.

We begin with the following proposition, identifying elements in the subdifferential of $\mathcal{E}_{\varepsilon}$, \mathcal{V} , $\mathcal{V}_{\varepsilon}$, and \mathcal{V}_{k} . Note that the subdifferentials of $\mathcal{V}, \mathcal{V}_{\varepsilon}$, and \mathcal{V}_k were characterized in previous work [3, Proposition 10.4.2], and we recall key parts these results in item (ii) below for the reader's convenience.

Proposition 3.7 (subdifferentials of $\mathcal{E}_{\varepsilon}$, \mathcal{V} , $\mathcal{V}_{\varepsilon}$, and \mathcal{V}_k).

- (i) Suppose Assumptions (T) and (M) hold. For all $\varepsilon > 0$ and $\mu \in D(\mathcal{E}_{\varepsilon})$, we have $\nabla \frac{\delta \mathcal{E}_{\varepsilon}}{\delta \mu} \in \partial \mathcal{E}_{\varepsilon}(\mu)$,
- where $\frac{\delta \mathcal{E}_{\varepsilon}}{\delta \mu} = \zeta_{\varepsilon} * ((\zeta_{\varepsilon} * \mu) / \bar{\rho}).$ (ii) Suppose Assumptions (M), (V), and (C) hold. For all $\mu \in D(\mathcal{V})$, we have $\nabla V \in \partial \mathcal{V}(\mu)$. Similarly, for all $\mu \in D(\mathcal{V}_{\varepsilon})$, we have $\nabla (\zeta_{\varepsilon} * V) \in \partial \mathcal{V}_{\varepsilon}(\mu)$, and, for all $\mu \in D(\mathcal{V}_k)$, we have $\nabla V_k \in \partial \mathcal{V}_k(\mu)$.

Proof. We begin with the proof of (i). Fix $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\gamma \in \Gamma_0(\mu, \nu)$. Let $\mu_{\alpha} = ((1 - \alpha)\pi^1 + \alpha\pi^2)_{\#}\gamma$ be a geodesic from μ to ν . By Lemma 3.6, $\mathcal{E}_{\varepsilon}$ is λ_{ε} -convex along generalized geodesics, so in particular, it is convex along μ_{α} , and Lemma 2.7 ensures,

$$\mathcal{E}_{\varepsilon}(\nu) - \mathcal{E}_{\varepsilon}(\mu) - \left. \frac{d}{d\alpha} \mathcal{E}_{\varepsilon}(\mu_{\alpha}) \right|_{\alpha=0} \ge \frac{\lambda_{\varepsilon}}{2} W_2^2(\mu, \nu).$$

Rearranging and applying Proposition 3.4, with $\tilde{\gamma} = (\pi^1, \pi^1, \pi^2)_{\#} \gamma$, and Fubini's theorem, yields,

$$\begin{aligned} \mathcal{E}_{\varepsilon}(\nu) - \mathcal{E}_{\varepsilon}(\mu) &\geq \frac{1}{2} \int \frac{\zeta_{\varepsilon} * \mu(x)}{\bar{\rho}(x)} \int \left\langle \nabla \zeta_{\varepsilon} (x - y_2), y_2 - y_3 \right\rangle \, d\tilde{\gamma}(y_1, y_2, y_3) \, dx + \frac{\lambda_{\varepsilon}}{2} W_2^2(\mu, \nu) \\ &= \frac{1}{2} \int \frac{\zeta_{\varepsilon} * \mu(x)}{\bar{\rho}(x)} \int \left\langle \nabla \zeta_{\varepsilon} (x - y_1), y_1 - y_2 \right\rangle \, d\gamma(y_1, y_2) \, dx + \frac{\lambda_{\varepsilon}}{2} W_2^2(\mu, \nu) \\ &= \frac{1}{2} \int \left\langle \nabla \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \mu}{\bar{\rho}} \right) (y_1), y_2 - y_1 \right\rangle d\gamma(y_1, y_2) + \frac{\lambda_{\varepsilon}}{2} W_2^2(\mu, \nu) \\ &= \int \left\langle \nabla \frac{\delta \mathcal{E}_{\varepsilon}}{\delta \mu}(y_1), y_2 - y_1 \right\rangle d\gamma(y_1, y_2) + \frac{\lambda_{\varepsilon}}{2} W_2^2(\mu, \nu). \end{aligned}$$

This shows $\nabla \frac{\delta \mathcal{E}_{\varepsilon}}{\delta \mu} \in \partial \mathcal{E}_{\varepsilon}(\mu)$, by Definition 2.9 of the subdifferential.

For item (ii), we will show the result for \mathcal{V} , since the result for $\mathcal{V}_{\varepsilon}$ and \mathcal{V}_{k} follow from the same argument, simply via replacing V with $\zeta_{\varepsilon} * V$ and V_k , respectively. Let ν, μ, γ , and $\tilde{\gamma}$ be as in the proof of item (i). Applying Lemma 2.7, Proposition 3.3, Proposition 3.4, and rearranging, again as in the proof of (i), yields,

$$\begin{aligned} \mathcal{V}(\nu) - \mathcal{V}(\mu) &\geq \int \left\langle \nabla V(y_2), y_3 - y_2 \right\rangle d\tilde{\gamma}(y_1, y_2, y_3) + \frac{\lambda}{2} W_2^2(\nu, \mu) \\ &= \int \left\langle \nabla V(y_1), y_2 - y_1 \right\rangle d\gamma(y_1, y_2) + \frac{\lambda}{2} W_2^2(\nu, \mu), \end{aligned}$$

which shows $\nabla V \in \partial \mathcal{V}(\mu)$, by Definition 2.9 of the subdifferential.

Next, we characterize the minimal subdifferential of the energy $\mathcal{F}_{\varepsilon,k} = \mathcal{E}_{\varepsilon} + \mathcal{V}_{\varepsilon} + \mathcal{V}_k$ for all $\varepsilon > 0, k \in \mathbb{N}$. The proof is standard, and we defer it to Appendix B.

Proposition 3.8 (minimal subdifferential of $\mathcal{F}_{\varepsilon,k}$). Suppose Suppose Assumptions (T), (M), (V), and (C) hold. For all $\varepsilon > 0$ and $k \in \mathbb{N}$, $\mu \in D(\mathcal{F}_{\varepsilon,k})$,

(3.6)
$$\nabla \frac{\partial \mathcal{E}_{\varepsilon}}{\partial \mu} + \nabla (\zeta_{\varepsilon} * V) + \nabla V_{k} = \partial^{\circ} \mathcal{F}_{\varepsilon,k}(\mu).$$

Finally, we close by recalling Ambrosio, Gigli, and Savaré's characterization of the minimal subdifferentials of \mathcal{F}_k and \mathcal{F} [3, Theorems 10.4.9-10.4.13].

Proposition 3.9 (minimal subdifferentials of \mathcal{F}_k and \mathcal{F} , [3, Theorems 10.4.9-10.4.13]). Assume (T), (V), and $\bar{\rho}$ is log-concave.

(i) Suppose also (C) holds. Then, given $\mu \in D(\mathcal{F}_k)$, we have $|\partial \mathcal{F}_k|(\mu) < +\infty$ if and only if $(\mu/\bar{\rho})^2 \in W^{1,1}_{\text{loc}}(\mathbb{R}^d)$ and there exists $\boldsymbol{\xi} \in L^2(\mu)$ so that,

(3.7)
$$\boldsymbol{\xi}\mu = \frac{\bar{\rho}}{2}\nabla(\mu/\bar{\rho})^2 + \nabla V\mu + \nabla V_k\mu \quad on \quad \mathbb{R}^d.$$

In this case, $\boldsymbol{\xi} \in \partial^{\circ} \mathcal{F}_k(\mu)$.

(ii) Suppose also (D) holds. Given $\mu \in D(\mathcal{F})$, we have $|\partial \mathcal{F}|(\mu) < +\infty$ if and only if $(\mu/\bar{\rho})^2 \in W^{1,1}_{\text{loc}}(\Omega)$ and there exists $\boldsymbol{\xi} \in L^2(\mu)$ so that,

$$oldsymbol{\xi} \mu = rac{ar{
ho}}{2}
abla (\mu/ar{
ho})^2 +
abla V \mu \quad on \quad \Omega$$

In this case, $\boldsymbol{\xi} \in \partial^{\circ} \mathcal{F}(\mu)$.

3.2. Differential equation characterization of gradient flows. We close by identifying the differential equations that characterize gradient flows of $\mathcal{F}_{k,\varepsilon}$, \mathcal{F}_k and \mathcal{F} . These proofs are natural consequences of the properties of the energies proved in the previous section and the definition of gradient flow, so we defer them to Appendix B.

Proposition 3.10 (PDE characterization of GF of \mathcal{F}_k and \mathcal{F}). Assume (V), (T), and $\bar{\rho}$ is log-concave.

(i) Suppose also (C) holds. For every $\mu_0 \in \overline{D(\mathcal{F}_k)}$, we have that $\mu(t) \in AC^2_{loc}((0, +\infty); \mathcal{P}_2(\mathbb{R}^d))$ is the unique Wasserstein gradient flow of \mathcal{F}_k with initial data μ_0 if and only if $\mu(t)$ satisfies,

(3.8)
$$\begin{cases} \partial_t \mu - \nabla \cdot \left(\frac{\bar{\rho}}{2} \nabla \left(\frac{\mu^2}{\bar{\rho}^2}\right) + \nabla V \mu + \nabla V_k \mu\right) = 0, \text{ in duality with } C_c^{\infty}(\mathbb{R}^d \times (0, \infty)), \\ \lim_{t \to 0^+} \mu(t) = \mu_0 \quad \text{ in } W_2, \end{cases}$$

and satisfies,

(3.9)
$$\mu(t) \ll \mathcal{L}^d \text{ and } (\mu(t)/\bar{\rho})^2 \in W^{1,1}_{loc}(\mathbb{R}^d) \text{ for } \mathcal{L}^1\text{-a.e. } t > 0,$$

(3.10)
$$\int_{\mathbb{R}^d} \left| \bar{\rho} \,\nabla\left(\mu(t)^2/\bar{\rho}^2\right)/(2\mu) + \nabla V + \nabla V_k \right|^2 d\mu \in L^1_{loc}(0,\infty).$$

(ii) Suppose also (D) holds. For every $\mu_0 \in \overline{D(\mathcal{F})}$, we have that $\mu(t) \in AC^2_{loc}((0, +\infty); \mathcal{P}_2(\mathbb{R}^d))$ is the unique Wasserstein gradient flow of \mathcal{F} with initial data μ_0 if and only if $\mu(t)$ satisfies,

(3.11)
$$\begin{cases} \partial_t \mu - \nabla \cdot \left(\frac{\bar{\rho}}{2} \nabla \left(\frac{\mu^2}{\bar{\rho}^2}\right) + \nabla V \mu\right) = 0, \text{ in duality with } C_c^{\infty}(\mathbb{R}^d \times (0, \infty)), \\ \lim_{t \to 0^+} \mu(t) = \mu_0 \quad \text{ in } W_2, \end{cases}$$

and satisfies,

(3.12)
$$\mu(t) \ll \mathcal{L}^d, \mu = 0 \ \mathcal{L}^d \text{-a.e. on } \mathbb{R}^d \setminus \overline{\Omega}, \text{ and } (\mu(t)/\bar{\rho})^2 \in W^{1,1}_{loc}(\Omega) \text{ for } \mathcal{L}^1 \text{-a.e. } t > 0,$$

(3.13)
$$\int_{\mathbb{R}^d} \left| \bar{\rho} \, \nabla \left(\mu(t)^2 / \bar{\rho}^2 \right) / (2\mu) + \nabla V \right|^2 d\mu \in L^1_{loc}(0,\infty)$$

Remark 3.11 (relationship with existing work on nonlinear diffusion equations). First, note that if $\overline{\Omega}$ is compact, then the weak formulation of the PDE in equation (3.11) implies that the PDE also holds in the duality with $C^{\infty}((0, +\infty) \times \Omega)$, which is a weak formulation of the no flux boundary conditions,

(3.14)
$$\frac{\bar{\rho}}{2}\partial_{\mathbf{n}}\left(\frac{\mu^2}{\bar{\rho}^2}\right) + \partial_{\mathbf{n}}V\mu = 0 \text{ on } \partial\Omega,$$

since the test functions are merely required to be compactly supported $\mathbb{R}^d \times (0, +\infty)$, not $\Omega \times (0, +\infty)$. In particular, if μ is a smooth classical solution of (WPME) with no flux boundary conditions, it solves (3.11).

In [64], Otto pioneered the connection between PDEs and Wasserstein gradient flows, characterizing solutions to homogeneous porous medium equations ($\bar{\rho} = 1$) without boundary ($\Omega = \mathbb{R}^d$) as gradient flows of the internal energy $\mathcal{F}(\rho) = \frac{1}{2} \int \rho^2$. The notion of solution used in this previous work is stronger than the one in Proposition 3.10. In particular, if ρ is a solution to the porous medium equation in this previous sense [64, Definition 1], then it is a solution of (3.11), hence a gradient flow in the sense defined here.

More recently, Dolbeault, et al. [35] and Grillo, Muratori, and Porzio [44] consider well-posedness of (WPME). If u is smooth enough, it is a solution to [44, equation (1.1)] (with $\rho_{\nu} = \rho_{\mu} = \bar{\rho}$, and with $\Omega = \mathbb{R}^d$) if and only if $\mu := \bar{\rho}u$ satisfies (3.11). More precisely comparing our notion of solution with [44, Definition

3.5], we observe that our definition requires the same regularity in space, stronger regularity in time, and we employ a smaller class of test functions.

Next, we provide a PDE characterization of the gradient flow of $\mathcal{F}_{\varepsilon,k}$, the proof of which we again defer to Appendix B.

Proposition 3.12 (PDE characterization of GF of $\mathcal{F}_{\varepsilon,k}$). Suppose Assumptions (T), (M), (V), and (C) hold. For every $\mu_0 \in \overline{D(\mathcal{F}_{\varepsilon,k})}$, we have that $\mu(t) \in AC^2_{loc}((0, +\infty); \mathcal{P}_2(\mathbb{R}^d))$ is the unique Wasserstein gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial data μ_0 if and only if $\mu(t)$ satisfies,

(3.15)
$$\begin{cases} \partial_t \mu - \nabla \cdot \left(\mu \left(\nabla \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \mu}{\bar{\rho}} \right) + \nabla (\zeta_{\varepsilon} * V) + \nabla V_k \right) \right) = 0, \text{ in duality with } C_c^{\infty}(\mathbb{R}^d \times (0, \infty)), \\ \lim_{t \to 0^+} \mu(t) = \mu_0 \text{ in } W_2. \end{cases}$$

Finally, we characterize the dynamics of the gradient flow of $\mathcal{F}_{\varepsilon,k}$ when the initial data is given by an empirical measure. We show that it remains an empirical measure for all time, that is, "particles remain particles", and we explicitly state the ODE that characterizes the empirical measure's evolution. The proof is in Appendix B.

Proposition 3.13 (particle evolution for $\mathcal{F}_{\varepsilon,k}$). Suppose Assumptions (T), (M), (V), and (C) hold. Fix $\varepsilon > 0, N \in \mathbb{N}, \{X_0^1, \ldots, X_0^N\} \in \mathbb{R}^d$, and $\{m^1, \ldots, m^N\} \in \mathbb{R}_+$ satisfying $\sum_{i=1}^N m^i = 1$. Then, there exists a unique continuously differentiable function $X : [0, \infty) \to \mathbb{R}^{Nd}$, with components $(X^1(t), \ldots, X^N(t))$, that satisfies the system,

(3.16)
$$\begin{cases} \dot{X}^i = -\sum_{j=1}^N m^j \int_{\mathbb{R}^d} \nabla \zeta_{\varepsilon} (X^i - z) \zeta_{\varepsilon} (z - X^j) \frac{1}{\bar{\rho}(z)} dz - \nabla (\zeta_{\varepsilon} * V) (X^i) - \nabla V_k (X^i) \\ X^i(0) = X_0^i. \end{cases}$$

Moreover, $\mu(t) := \sum_{i=1}^{N} \delta_{X^{i}(t)} m^{i}$ is the unique Wasserstein gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial conditions $\mu(0)$.

3.3. Long-time behavior. We conclude this section by recalling known properties of the long time behavior of (WPME) or, equivalently, gradient flows of \mathcal{F} , which motivate its connection to quantization.

Proposition 3.14 (long time behavior, [3]). Assume (D), (T), V = 0, Ω is bounded, and $\bar{\rho}$ is log-concave. Let $\rho_0 \in D(\mathcal{F})$ and let $\rho(t)$ be the gradient flow ρ of \mathcal{F} with initial data ρ_0 . Then we have,

$$\lim_{t \to +\infty} W_2\left(\rho(t), \frac{\mathbb{1}_{\overline{\Omega}}\bar{\rho}}{\int_{\Omega} \bar{\rho} \, d\mathcal{L}^d}\right) = 0.$$

Proof. This is an immediate consequence of [3, Corollary 4.0.6].

4. An H^1 bound on the mollified gradient flow of $\mathcal{F}_{\varepsilon,k}$

A key element in our proof of the convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ to a gradient flow of \mathcal{F}_k as $\varepsilon \to 0$ is the following H^1 -type bound on $\zeta_{\varepsilon} * \rho_{\varepsilon}(t)$ (the mollified gradient flow of $\mathcal{F}_{\varepsilon,k}$) in terms of the energy, second moment, and entropy of the initial data. We remark that this bound holds without a log-concavity assumption on $\bar{\rho}$.

Theorem 4.1 (H^1 bound on mollified GF of $\mathcal{F}_{\varepsilon,k}$). Assume (T), (M), (V), and (C) hold. There exist positive constant $C_{\bar{\rho}}$ and C_V , depending only on $\bar{\rho}$ and V, respectively, so that, for all T > 0, $k \in \mathbb{N}$, and $\varepsilon > 0$ and for any gradient flow $\rho_{\varepsilon} \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ of $\mathcal{F}_{\varepsilon,k}$, we have,

$$(4.1) \quad \int_0^T ||\nabla \zeta_{\varepsilon} * \rho_{\varepsilon}(t)||^2_{L^2(\mathbb{R}^d)} dt \le C_{\bar{\rho}} \left(\mathcal{S}(\rho_{\varepsilon}(0)) + \sqrt{2\pi} + (1 + T + Te^T) \left(M_2(\rho_{\varepsilon}(0)) + \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}(0)) + C_V \right) \right).$$

4.1. **Proof sketch.** First, we describe a formal argument to obtain inequality (4.1), and then we explain how to make the argument rigorous. By Proposition 3.12, $\rho_{\varepsilon}(t)$ is a weak solution of the PDE,

(4.2)
$$\partial_t \rho_{\varepsilon} = \nabla \cdot \left(\rho_{\varepsilon} \nabla \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \rho_{\varepsilon}}{\bar{\rho}} \right) + \rho_{\varepsilon} \nabla (\zeta_{\varepsilon} * V) + \rho_{\varepsilon} \nabla V_k \right),$$

in the duality with $C_c^{\infty}(\mathbb{R}^d \times (0, \infty))$. Thus, formally evaluating the entropy $\mathcal{S}(\rho)$ along the gradient flow, differentiating in time, and integrating by parts, we obtain,

$$(4.3) \quad \frac{d}{dt}\mathcal{S}(\rho_{\varepsilon})(t) = \int_{\mathbb{R}^{d}} \log(\rho_{\varepsilon})\partial_{t}\rho_{\varepsilon} \, d\mathcal{L}^{d} \\ = -\int_{\mathbb{R}^{d}} \left\langle \nabla\rho_{\varepsilon}, \nabla\zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \rho_{\varepsilon}}{\bar{\rho}}\right) \right\rangle + \left\langle \nabla\rho_{\varepsilon}, \nabla(\zeta_{\varepsilon} * V) \right\rangle + \left\langle \nabla\rho_{\varepsilon}, \nabla V_{k} \right\rangle d\mathcal{L}^{d} \\ = -\int_{\mathbb{R}^{d}} \frac{|\nabla\zeta_{\varepsilon} * \rho_{\varepsilon}|^{2}}{\bar{\rho}} + \left\langle \nabla(\zeta_{\varepsilon} * \rho_{\varepsilon}), (\zeta_{\varepsilon} * \rho_{\varepsilon})\nabla\left(\frac{1}{\bar{\rho}}\right) \right\rangle + \left\langle \nabla(\zeta_{\varepsilon} * \rho_{\varepsilon}), \nabla V \right\rangle - \rho_{\varepsilon}\Delta V_{k} \, d\mathcal{L}^{d}.$$

Integrating in time and estimating the terms on the right hand side then leads to inequality (4.1).

The key difficulty in making the above argument rigorous is justifying the time differentiation of the entropy, in the absence of relevant a priori estimates for ρ_{ε} . In order to overcome this difficulty, McCann, Matthes, and Savaré introduced the *flow interchange method* [56]. Suppose that $\rho_{\varepsilon}(t)$ and $\mu(t)$ are, respectively, the gradient flows of the energy $\mathcal{F}_{\varepsilon,k}$ and the entropy \mathcal{S} , and we have $\rho_{\varepsilon}(0) = \mu(0)$. The flow interchange method is based on the following formal observation, with ∇_{W_2} denoting the Wasserstein gradient and $\langle \cdot, \cdot \rangle_{W_2}$ denoting the Wasserstein inner product at ρ :

$$\frac{d}{dt}\mathcal{S}(\rho_{\varepsilon})\Big|_{t=0} = \left\langle \nabla_{W_{2}}\mathcal{S}(\rho_{\varepsilon}), \partial_{t}\rho_{\varepsilon} \right\rangle_{W_{2},\rho_{\varepsilon}}\Big|_{t=0} = -\left\langle \nabla_{W_{2}}\mathcal{S}(\rho_{\varepsilon}), \nabla_{W_{2}}\mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}) \right\rangle_{W_{2},\rho_{\varepsilon}}\Big|_{t=0} = -\left\langle \nabla_{W_{2}}\mathcal{S}(\mu), \nabla_{W_{2}}\mathcal{F}_{\varepsilon,k}(\mu) \right\rangle_{W_{2},\mu}\Big|_{t=0} = \left\langle \partial_{t}\mu, \nabla_{W_{2}}\mathcal{F}_{\varepsilon,k}(\mu) \right\rangle_{W_{2},\mu}\Big|_{t=0} = \frac{d}{dt}\mathcal{F}_{\varepsilon,k}(\mu)\Big|_{t=0}.$$

Consequently, at a fixed time, differentiating $\mathcal{F}_{\varepsilon,k}$ along the gradient flow of \mathcal{S} should give the same result as equation (4.3). The former is much easier to justify in practice, since the gradient flow $\mu(t)$ of \mathcal{S} with initial data $\mu(0)$ is precisely the solution of the heat equation on \mathbb{R}^d with initial data $\mu(0)$ [3, Examples 11.2.7], for which we have robust a priori estimates.

Note that, since the entropy S is a 0-convex energy [3, Proposition 9.3.9], the evolution variational inequality characterization of gradient flows, recalled in Theorem 2.13, ensures that, if $\mu(t)$ is the gradient flow of S, then for all $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and for \mathcal{L}^1 -a.e. $t \ge 0$,

(4.4)
$$\frac{1}{2}\frac{d^+}{dt}W_2^2(\mu(t),\nu) + \mathcal{S}(\mu(t)) \le \mathcal{S}(\nu).$$

4.2. Preliminaries for the proof. Now, we introduce the machinery we need for our rigorous argument, following the outline described above. To avoid differentiating $S(\rho_{\varepsilon})$ in time, we work with the discrete time analogue of the gradient flow of $\mathcal{F}_{\varepsilon,k}$, given by the minimizing movement scheme (see Definition A.1).

Definition 4.2 (minimizing movement scheme for $\mathcal{F}_{\varepsilon,k}$). Given $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, let $J^n_{\tau,\varepsilon}\mu$ denote the *n*th step of the minimizing movement scheme of $\mathcal{F}_{\varepsilon,k}$ with time step τ and initial data $J^0_{\tau,\varepsilon}\mu = \mu$.

Due to the robust a priori estimates available for solutions of the heat equation, we will work with continuous time gradient flow of S.

Definition 4.3 (heat flow semigroup). Gven $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $h \ge 0$, we will let $S_h\mu$ denote the (continuous time) gradient flow of \mathcal{S} with initial data μ at time h; in other words, S_h is the heat flow semigroup operator.

We will use the fact that, for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, we have,

(4.5)
$$\zeta_{\varepsilon} * (\mathbf{S}_h(\mu)) = \mathbf{S}_h(\zeta_{\varepsilon} * \mu).$$

A key step in the proof is computing the derivatives in h of $\mathcal{E}_{\varepsilon}(S_h(J_{\tau,\varepsilon}^n\mu)), \mathcal{V}_{\varepsilon}(S_h(J_{\tau,\varepsilon}^n\mu))$, and $\mathcal{V}_k(S_h(J_{\tau,\varepsilon}^n\mu))$ at h = 0. We separate this step into a separate lemma:

Lemma 4.4 (derivatives along $S_h(\mathcal{J}^n_{\tau,\varepsilon}\mu)$). Assume (T), (M), (V), and (C) hold. Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. We have,

(4.6)
$$\limsup_{h \to 0^+} \frac{\mathcal{E}_{\varepsilon}(J^n_{\tau,\varepsilon}\mu) - \mathcal{E}_{\varepsilon}(S_h(J^n_{\tau,\varepsilon}\mu))}{h} = -\int_{\mathbb{R}^d} \frac{1}{\bar{\rho}} \Delta(\zeta_{\varepsilon} * J^n_{\tau,\varepsilon}\mu) \left(\zeta_{\varepsilon} * J^n_{\tau,\varepsilon}\mu\right) \, d\mathcal{L}^d,$$

(4.7)
$$\limsup_{h \to 0^+} \frac{\mathcal{V}_{\varepsilon}(\mathbf{J}_{\tau,\varepsilon}^n \mu) - \mathcal{V}_{\varepsilon}(\mathbf{S}_h(\mathbf{J}_{\tau,\varepsilon}^n \mu))}{h} = \int_{\mathbb{R}^d} \left\langle \nabla V, \nabla(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^n \mu) \right\rangle d\mathcal{L}^d, \text{ and}$$

(4.8)
$$\limsup_{h \to 0^+} \frac{\mathcal{V}_k(\mathbf{J}_{\tau,\varepsilon}^n \mu) - \mathcal{V}_k(\mathbf{S}_h(\mathbf{J}_{\tau,\varepsilon}^n \mu))}{h} = -\int_{\mathbb{R}^d} \Delta V_k \, d\mathbf{J}_{\tau,\varepsilon}^n \mu.$$

Our proof of this lemma relies on two key facts, which we now recall. First, for any $\nu \in \mathcal{P}_2(\mathbb{R}^d)$,

(4.9) the map
$$h \mapsto S_h \nu$$
 is narrowly continuous

that is, $h \mapsto \int f \, dS_h \nu$ is continuous for any bounded and continuous function f. This holds since $S_h \nu$, by virtue of being the gradient flow of S, is in $AC_{loc}^2((0, +\infty); \mathcal{P}_2(\mathbb{R}^d))$, hence $h \mapsto S_h \nu$ is continuous with respect to W_2 , which implies narrow continuity.

The second fact we will use is that, for any for any $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\phi \in C_c^1(\mathbb{R}^d)$,

(4.10)
$$\int_{\mathbb{R}^d} \phi \, d\mathbf{S}_h \nu - \int_{\mathbb{R}^d} \phi \, d\nu = -\int_0^h \int_{\mathbb{R}^d} \left\langle \nabla \phi(y), \nabla \mathbf{S}_t \nu(y) \right\rangle \, dy \, dt.$$

Notice that, at a formal level, the integrand on the left-hand side is exactly $\int_0^h \frac{d}{dt} S_t \nu dt$, which, upon using the fact that $S_t \nu$ satisfies the heat equation, and integrating by parts, yields the desired equality. More rigorously, one may obtain (4.10) as a consequence of [4, Lemma 8.1.2]. And, arguing as in [4, Example 11.1.9], we have $S_h \nu \in W_{loc}^{1,1}(\mathbb{R}^d)$ for a.e. h > 0 and,

$$\int_0^T \int_{\mathbb{R}^d} |\nabla S_h \nu| = \int_0^T \int_{\mathbb{R}^d} \frac{|\nabla S_h \nu|}{S_h \nu} S_h \nu \le \left(\int_0^T \int_{\mathbb{R}^d} \frac{|\nabla S_h \nu|^2}{S_h \nu} \right)^{1/2} \left(\int_0^T \int_{\mathbb{R}^d} S_h \nu \right)^{1/2} = \sqrt{T} \left(\int_0^T \int_{\mathbb{R}^d} \frac{|\nabla S_h \nu|^2}{S_h \nu} \right)^{1/2}$$
where the quantity on the right hand side is finite by equation (11.1.28) of [4].

where the quantity on the right-hand side is finite by equation (11.1.38) of [4].

With these facts in hand, we now turn to the proof of Lemma 4.4.

Proof of Lemma 4.4. We begin by proving equation (4.6). For all h > 0, using the definition of $\mathcal{E}_{\varepsilon}$ and the commutativity relation (4.5), we find,

$$\frac{\mathcal{E}_{\varepsilon}(J_{\tau,\varepsilon}^{n}\mu) - \mathcal{E}_{\varepsilon}(S_{h}(J_{\tau,\varepsilon}^{n}\mu))}{h} = \frac{1}{2h} \int_{\mathbb{R}^{d}} \frac{\left|\zeta_{\varepsilon} * (J_{\tau,\varepsilon}^{n}\mu)\right|^{2}}{\bar{\rho}} d\mathcal{L}^{d} - \frac{1}{2h} \int_{\mathbb{R}^{d}} \frac{\left|\zeta_{\varepsilon} * (S_{h}(J_{\tau,\varepsilon}^{n}\mu))\right|^{2}}{\bar{\rho}} d\mathcal{L}^{d}
= \int_{\mathbb{R}^{d}} \frac{1}{2\bar{\rho}} \left(\frac{\zeta_{\varepsilon} * (J_{\tau,\varepsilon}^{n}\mu) - \zeta_{\varepsilon} * (S_{h}(J_{\tau,\varepsilon}^{n}\mu))}{h}\right) \left(\zeta_{\varepsilon} * (J_{\tau,\varepsilon}^{n}\mu) + \zeta_{\varepsilon} * (S_{h}(J_{\tau,\varepsilon}^{n}\mu))\right) d\mathcal{L}^{d}
(4.11) = \int_{\mathbb{R}^{d}} \frac{1}{2\bar{\rho}} \left(\frac{\zeta_{\varepsilon} * (J_{\tau,\varepsilon}^{n}\mu) - S_{h}(\zeta_{\varepsilon} * J_{\tau,\varepsilon}^{n}\mu)}{h}\right) \left(\zeta_{\varepsilon} * (J_{\tau,\varepsilon}^{n}\mu) + \zeta_{\varepsilon} * (S_{h}(J_{\tau,\varepsilon}^{n}\mu))\right) d\mathcal{L}^{d}.$$

Recalling that $\zeta_{\varepsilon} * (J_{\tau,\varepsilon}^n \mu)$ is a smooth function, and using that $S_h(\zeta_{\varepsilon} * J_{\tau,\varepsilon}^n \mu)$ satisfies the heat equation in the classical sense, we find,

(4.12)
$$\zeta_{\varepsilon} * (\mathbf{J}_{\tau,\varepsilon}^{n}\mu) - \mathbf{S}_{h}(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu) = -\int_{0}^{h} \frac{d}{dt} \mathbf{S}_{t}(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu) dt = -\int_{0}^{h} \Delta \mathbf{S}_{t}(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu) dt$$

Using this in (4.11), we obtain,

$$\frac{\mathcal{E}_{\varepsilon}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) - \mathcal{E}_{\varepsilon}(\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu))}{h} = \int_{\mathbb{R}^{d}} \frac{1}{2\bar{\rho}} \left(\frac{1}{h} \int_{0}^{h} -\Delta \mathbf{S}_{t}(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu) \, dt\right) \left(\zeta_{\varepsilon} * (\mathbf{J}_{\tau,\varepsilon}^{n}\mu) + \zeta_{\varepsilon} * (\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu))\right) d\mathcal{L}^{d}.$$

Classical elliptic regularity implies that $\|\Delta S_t(\zeta_{\varepsilon} * J^n_{\tau,\varepsilon}\mu)\|_{L^{\infty}(\mathbb{R}^d)} \leq C_{\varepsilon,\tau,n}$ holds for all t. Hence, the integrand on the right-hand side of the previous line is bounded in $L^1(\mathbb{R}^d)$, independently of h. Thus, upon applying the dominated convergence theorem to take the limit $h \to 0+$, we find,

$$\limsup_{h\to 0+} \frac{\mathcal{E}_{\varepsilon}(\mathrm{J}^n_{\tau,\varepsilon}\mu) - \mathcal{E}_{\varepsilon}(\mathrm{S}_h(\mathrm{J}^n_{\tau,\varepsilon}\mu))}{h} = -\int_{\mathbb{R}^d} \frac{1}{\bar{\rho}} \Delta(\zeta_{\varepsilon} * \mathrm{J}^n_{\tau,\varepsilon}\mu) \left(\zeta_{\varepsilon} * \mathrm{J}^n_{\tau,\varepsilon}\mu\right) d\mathcal{L}^d.$$

We have again used that $S_t(\zeta_{\varepsilon} * J^n_{\tau,\varepsilon})$ satisfies the heat equation in the classical sense, and is therefore continuous in t. This completes the proof of equation (4.6).

Next we establish equation (4.7). For all h > 0, using the definition of $\mathcal{V}_{\varepsilon}$, followed by (4.5), we obtain,

$$\frac{\mathcal{V}_{\varepsilon}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) - \mathcal{V}_{\varepsilon}(\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu))}{h} = \frac{1}{h} \left(\int_{\mathbb{R}^{d}} \left(\zeta_{\varepsilon} * V \right) d\mathbf{J}_{\tau,\varepsilon}^{n}\mu - \int_{\mathbb{R}^{d}} \left(\zeta_{\varepsilon} * V \right) d\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) \right)$$
$$= \frac{1}{h} \int_{\mathbb{R}^{d}} V \left(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu - \zeta_{\varepsilon} * \mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) \right) d\mathcal{L}^{d}$$
$$= \frac{1}{h} \int_{\mathbb{R}^{d}} V \left(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu - \mathbf{S}_{h}(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu) \right) d\mathcal{L}^{d}.$$

As in the computation for $\mathcal{E}_{\varepsilon}$, we now use (4.12) to find,

$$\frac{\mathcal{V}_{\varepsilon}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) - \mathcal{V}_{\varepsilon}(\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu))}{h} = -\int_{\mathbb{R}^{d}} V\frac{1}{h} \int_{0}^{h} \Delta \mathbf{S}_{t}(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu) \, dt \, d\mathcal{L}^{d}.$$

Assumption (V) implies $V \in L^1(\mathbb{R}^d)$, so we can pass to the limit in h (again, as above), and find,

$$\limsup_{h \to 0+} \frac{\mathcal{V}_{\varepsilon}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) - \mathcal{V}_{\varepsilon}(\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu))}{h} = -\int_{\mathbb{R}^{d}} V\Delta(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu) \, d\mathcal{L}^{d}$$

Integrating by parts yields (4.7).

Finally, we establish (4.8). For all h > 0, using the definition of \mathcal{V}_k , followed by (4.10), and an integration by parts, yields,

$$\frac{\mathcal{V}_{k}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) - \mathcal{V}_{k}(\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu))}{h} = \frac{1}{h} \left(\int_{\mathbb{R}^{d}} V_{k} \, d\mathbf{J}_{\tau,\varepsilon}^{n}\mu - \int_{\mathbb{R}^{d}} V_{k} \, d\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) \right)$$
$$= \int_{\mathbb{R}^{d}} \frac{1}{h} \int_{0}^{h} \left\langle \nabla V_{k}(x), \nabla \mathbf{S}_{t} \mathbf{J}_{\tau,\varepsilon}^{n}\mu(x,t) \right\rangle dt \, dx$$
$$= -\int_{\mathbb{R}^{d}} \frac{1}{h} \int_{0}^{h} \Delta V_{k}(x) \mathbf{S}_{t} \mathbf{J}_{\tau,\varepsilon}^{n}\mu(x,t) \, dt \, dx.$$

Since $\|\Delta V_k\|_{L^{\infty}(\mathbb{R}^d)}$ is bounded, we use the dominated convergence theorem, as well as the narrow continuity of $S_t J^n_{\tau,\varepsilon} \mu$ in t (see (4.9)), to pass to the limit in h and obtain the desired result.

Before proceeding to the main result of the section, we estimate the right-hand side of (4.6). Notice that the hypotheses on ϕ in the statement are satisfied by $\zeta_{\varepsilon} * J^n_{\tau,\varepsilon}\mu$, since $J^n_{\tau,\varepsilon}\mu \in D(\mathcal{E}_{\varepsilon})$.

Lemma 4.5. Let $\phi \in C^{\infty}(\mathbb{R}^d) \cap L^1(\mathbb{R}^d) \cap L^2(\mathbb{R}^d)$. Then we have,

$$-\int_{\mathbb{R}^d} \frac{1}{\bar{\rho}} (\Delta\phi)(\phi) \, d\mathcal{L} \ge C_{\bar{\rho}} \|\nabla\phi\|_{L^2(\mathbb{R}^d)}^2 - C_{\bar{\rho}}' \|\phi\|_{L^2(\mathbb{R}^d)}^2,$$

where $C_{\bar{\rho}}$ and $C'_{\bar{\rho}}$ depend only on $\bar{\rho}$.

Proof. Integrating by parts, using the product rule, and the fact that $\bar{\rho}$ is bounded uniformly away from zero, we find,

$$-\int_{\mathbb{R}^d} \frac{1}{\bar{\rho}} (\Delta \phi)(\phi) = \int_{\mathbb{R}^d} \left\langle \nabla \phi, \nabla \left(\frac{1}{\bar{\rho}}\phi\right) \right\rangle = \int_{\mathbb{R}^d} \frac{|\nabla \phi|^2}{\bar{\rho}} + \phi \left\langle \nabla \phi, \nabla \left(\frac{1}{\bar{\rho}}\right) \right\rangle$$
$$\geq C_{\bar{\rho}} \int_{\mathbb{R}^d} |\nabla \phi|^2 - C'_{\bar{\rho}} \int_{\mathbb{R}^d} |\nabla \phi| |\phi| \geq \frac{C_{\bar{\rho}}}{2} \int_{\mathbb{R}^d} |\nabla \phi|^2 - C'_{\bar{\rho}} \int |\phi|^2,$$

where the last estimate follows from the Cauchy-Schwartz inequality, and $C'_{\bar{\rho}}$ changes from line to line (but depends only on $\bar{\rho}$).

4.3. **Proof of** H^1 -**type bound.** We now apply the previous lemmas to prove the main result of the section. *Proof of Theorem 4.1.* By definition of the minimizing movement scheme (see Definition A.1), for any $\mu \in D(\mathcal{F}_{\varepsilon,k})$,

$$\mathcal{F}_{\varepsilon,k}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu) - \mathcal{F}_{\varepsilon,k}(\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu)) \leq \frac{1}{2\tau} \left[W_{2}^{2}(\mathbf{S}_{h}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu), \mathbf{J}_{\tau,\varepsilon}^{n-1}\mu) - W_{2}^{2}(\mathbf{J}_{\tau,\varepsilon}^{n}\mu, \mathbf{J}_{\tau,\varepsilon}^{n-1}\mu) \right].$$

Dividing by h, taking the limit as $h \to 0$, and applying the evolution variational inequality characterization of the gradient flow of S, inequality (4.4), we obtain,

$$(4.13) \quad \limsup_{h \to 0^+} \frac{\mathcal{F}_{\varepsilon,k}(\mathbf{J}_{\tau,\varepsilon}^n \mu) - \mathcal{F}_{\varepsilon,k}(\mathbf{S}_h(\mathbf{J}_{\tau,\varepsilon}^n \mu))}{h} \le \frac{1}{2\tau} \left. \frac{d^+}{dh} W_2^2(\mathbf{S}_h(\mathbf{J}_{\tau,\varepsilon}^n \mu), \mathbf{J}_{\tau,\varepsilon}^{n-1} \mu) \right|_{h=0} \le \frac{\mathcal{S}(\mathbf{J}_{\tau,\varepsilon}^{n-1} \mu) - \mathcal{S}(\mathbf{J}_{\tau,\varepsilon}^n \mu)}{\tau}.$$

The quantity on the right hand side will play the role of $-\frac{d}{dt}S(\rho_{\varepsilon})$ in the $\tau \to 0$ limit. Thus, in order to obtain (4.1), we aim to bound it from below by estimating the left hand side of (4.13).

Recalling that $\mathcal{F}_{\varepsilon,k} = \mathcal{E}_{\varepsilon} + \mathcal{V}_{\varepsilon} + \mathcal{V}_k$ and applying Lemma 4.4, we find,

$$\begin{split} \limsup_{h \to 0^+} \frac{\mathcal{F}_{\varepsilon,k}(\mathbf{J}_{\tau,\varepsilon}^n \mu) - \mathcal{F}_{\varepsilon,k}(\mathbf{S}_h(\mathbf{J}_{\tau,\varepsilon}^n \mu))}{h} \\ &= -\int_{\mathbb{R}^d} \frac{1}{\bar{\rho}} \Delta(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^n \mu) \left(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^n \mu\right) \, d\mathcal{L}^d + \int_{\mathbb{R}^d} \left\langle \nabla V, \nabla(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^n \mu) \right\rangle d\mathcal{L}^d - \int_{\mathbb{R}^d} \Delta V_k \, d\mathbf{J}_{\tau,\varepsilon}^n \mu. \end{split}$$

Combining this with (4.13), and and summing over n, we obtain,

$$\frac{S(\mathbf{J}_{\tau,\varepsilon}^{0}\mu) - S(\mathbf{J}_{\tau,\varepsilon}^{n}\mu)}{\tau} = \sum_{i=1}^{n} \frac{S(\mathbf{J}_{\tau,\varepsilon}^{i-1}\mu) - S(\mathbf{J}_{\tau,\varepsilon}^{i}\mu)}{\tau}$$
$$\geq \sum_{i=1}^{n} -\int_{\mathbb{R}^{d}} \frac{1}{\bar{\rho}} \Delta(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{i}\mu) \left(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{i}\mu\right) \, d\mathcal{L}^{d} + \int_{\mathbb{R}^{d}} \left\langle \nabla V, \nabla(\zeta_{\varepsilon} * \mathbf{J}_{\tau,\varepsilon}^{n}\mu) \right\rangle d\mathcal{L}^{d} - \int_{\mathbb{R}^{d}} \Delta V_{k} \, d\mathbf{J}_{\tau,\varepsilon}^{i}\mu$$

Take $\tau = T/n$, and let $\mu_{\tau,\varepsilon}(t)$ denote the piecewise constant interpolation of the minimizing movement scheme $J^n_{\tau,\varepsilon}\mu$; see equation (A.2). Then the above line implies,

$$(4.14) \quad \mathcal{S}(\mu_{\tau,\varepsilon}(0)) - \mathcal{S}(\mu_{\tau,\varepsilon}(T)) \\ \geq \int_0^T \int_{\mathbb{R}^d} -\frac{1}{\bar{\rho}} \Delta(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)) \left(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\right) + \langle \nabla V, \nabla(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)) \rangle \, d\mathcal{L}^d \, ds - \int_0^T \int_{\mathbb{R}^d} \Delta V_k \, d\mu_{\tau,\varepsilon}(s) \, ds.$$

We consider the right-hand side. The first term on the right-hand side is the most important one, since this is where the derivative we seek to estimate will come from. First, we note, using the definition of $\mathcal{E}_{\varepsilon}$, the properties of $\bar{\rho}$, the fact that the energy $\mathcal{F}_{\varepsilon,k}$ decreases along the minimizing movements scheme (see inequality (A.1)), and the fact that the minimizing movements scheme is initialized at $\rho_{\varepsilon}(0)$,

$$(4.15) \|\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\|_{L^{2}(\mathbb{R}^{d})}^{2} \leq \|\bar{\rho}^{-1}\|_{L^{\infty}(\mathbb{R}^{d})} \mathcal{E}_{\varepsilon}(\mu_{\tau,\varepsilon}(s)) \leq \|\bar{\rho}^{-1}\|_{L^{\infty}(\mathbb{R}^{d})} \left(\mathcal{F}_{\varepsilon,k}(\mu_{\tau,\varepsilon}(s)) + \|V\|_{L^{\infty}(\mathbb{R}^{d})}\right) \\ \leq \|\bar{\rho}^{-1}\|_{L^{\infty}(\mathbb{R}^{d})} \left(\mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}(0)) + \|V\|_{L^{\infty}(\mathbb{R}^{d})}\right) < +\infty.$$

Thus, for each fixed s, we may apply Lemma 4.5 to find,

$$\int_{\mathbb{R}^d} -\frac{1}{\bar{\rho}} \Delta(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)) \left(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\right) d\mathcal{L}^d \ge C_{\bar{\rho}} \|\nabla\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\|_{L^2(\mathbb{R}^d)}^2 - C_{\bar{\rho}}' \|\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\|_{L^2(\mathbb{R}^d)}^2$$

Using (4.15) to bound the second term on the right-hand side of the previous line from below, and integrating in time, we find,

$$\int_{0}^{T} \int_{\mathbb{R}^{d}} -\frac{1}{\bar{\rho}} \Delta(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)) \left(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\right) d\mathcal{L}^{d} ds \geq C_{\bar{\rho}} \int_{0}^{T} \|\nabla\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\|_{L^{2}(\mathbb{R}^{d})}^{2} ds - TC_{\bar{\rho}}' \left(\mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}(0)) + \|V\|_{L^{\infty}(\mathbb{R}^{d})}\right).$$

(Here $C'_{\bar{\rho}}$ is allowed to change from line to line, but only depends on $\bar{\rho}$.)

Next, we apply the Cauchy-Schwartz inequality to the second term on the right-hand side of (4.14) to obtain,

$$\int_0^T \int_{\mathbb{R}^d} \left\langle \nabla V, \nabla(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)) \right\rangle d\mathcal{L}^d \, ds \ge -\frac{C_{\bar{\rho}}}{2} \int_0^T \|\nabla\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\|_{L^2(\mathbb{R}^d)}^2 \, ds - C'_{\bar{\rho}} T \|\nabla V\|_{L^2(\mathbb{R}^d)}.$$

Finally, for third term on the right-hand side of (4.14), we bound it from below simply by $T \|\Delta V_k\|_{L^{\infty}(\mathbb{R}^d)}$. which is finite by assumption. Using this, along with the two previous estimates, we find,

(4.16)
$$\mathcal{S}(\mu_{\tau,\varepsilon}(0)) - \mathcal{S}(\mu_{\tau,\varepsilon}(T)) \geq \frac{C_{\bar{\rho}}}{2} \int_0^T \|\nabla\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)\|_{L^2(\mathbb{R}^d)}^2 \, ds - TC'_{\bar{\rho}} \left(\mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}(0)) + C_V\right),$$

where $C_V = \|V\|_{L^\infty(\mathbb{R}^d)} + \|\nabla V\|_{L^2(\mathbb{R}^d)} + \|\Delta V_k\|_{L^\infty(\mathbb{R}^d)}.$

We now aim to send $n \to +\infty$ in inequality (4.16), using the fact that $\mu_{\tau,\varepsilon}(t) \to \rho_{\varepsilon}(t)$ narrowly for all $t \geq 0$; see Theorem A.2. Note that, for any $f \in L^2(\mathbb{R}^d)$ and $s \in [0, T]$,

$$\int_{\mathbb{R}^d} f\nabla(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon}(s)) = -\int_{\mathbb{R}^d} (\nabla\zeta_{\varepsilon} * f) \mu_{\tau,\varepsilon}(s) \xrightarrow{n \to +\infty} - \int_{\mathbb{R}^d} (\nabla\zeta_{\varepsilon} * f) \rho_{\varepsilon}(s) = \int_{\mathbb{R}^d} f\nabla(\zeta_{\varepsilon} * \rho_{\varepsilon}(s)).$$

Thus, $\nabla(\zeta_{\varepsilon} * \mu_{\tau,\varepsilon})(s) \to \nabla(\zeta_{\varepsilon} * \rho_{\varepsilon})(s)$ weakly in $L^2(\mathbb{R}^d)$ for all $s \in [0,T]$. By the lower semicontinuity of the $L^2(\mathbb{R}^d)$ norm with respect to weak convergence and Fatou's lemma, sending $n \to +\infty$ in inequality (4.16) yields,

(4.17)
$$\limsup_{n \to \infty} \mathcal{S}(\mu_{\tau,\varepsilon}(0)) - \mathcal{S}(\mu_{\tau,\varepsilon}(T)) \ge \frac{C_{\bar{\rho}}}{2} \int_0^T \|\nabla \zeta_{\varepsilon} * \rho_{\varepsilon}(s)\|_{L^2(\mathbb{R}^d)}^2 \, ds - T\left(C'_{\bar{\rho}}\mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}(0)) + C_V\right).$$

For the left hand side of (4.17), note that the choice of initial data for the minimizing movement scheme ensures $\mathcal{S}(\mu_{\tau,\varepsilon}(0)) = \rho_{\varepsilon}(0)$ for all $\tau > 0$ and, by the lower semicontinuity of the entropy with respect to narrow convergence [3, Remark 9.3.8], $\limsup_{n\to\infty} -\mathcal{S}(\mu_{\tau,\varepsilon}(T)) \leq -\mathcal{S}(\rho_{\varepsilon}(T))$. Thus, sending $n \to +\infty$ on the left hand side of (4.17), we estimate,

(4.18)
$$\limsup_{n \to +\infty} \mathcal{S}(\mu_{\tau,\varepsilon}(0)) - \mathcal{S}(\mu_{\tau,\varepsilon}(T)) \leq \mathcal{S}(\rho_{\varepsilon}(0)) - \mathcal{S}(\rho_{\varepsilon}(T))$$

Finally, using a Carleman-type estimate [24, Lemma 4.1] to bound the entropy below by a constant plus the second moment and applying Proposition A.3 to bound the second moment, we obtain,

(4.19)
$$\lim_{n \to +\infty} \sup \mathcal{S}(\mu_{\tau,\varepsilon}(0)) - \mathcal{S}(\mu_{\tau,\varepsilon}(T)) \leq \mathcal{S}(\rho_{\varepsilon}(0)) + \sqrt{2\pi} + M_2(\rho_{\varepsilon}(T)) \leq \mathcal{S}(\rho_{\varepsilon}(0)) + \sqrt{2\pi} + (1 + Te^T) (M_2(\rho_{\varepsilon}(0)) + \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}(0))).$$

Thus, combining inequalities (4.17) and (4.19), we obtain

$$\mathcal{S}(\rho_{\varepsilon}(0)) + \sqrt{2\pi} + (1 + Te^{T}) \left(M_{2}(\rho_{\varepsilon}(0)) + \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}(0)) \right) \geq \frac{C_{\bar{\rho}}}{2} \int_{0}^{T} \|\nabla\zeta_{\varepsilon} * \rho_{\varepsilon}(s)\|_{L^{2}(\mathbb{R}^{d})}^{2} ds - T \left(C_{\rho}' \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}(0)) + C_{V} \right)$$

Bearranging then gives the result.

iging

5. Convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ to \mathcal{F}_k

We now apply the properties of the energy $\mathcal{F}_{\varepsilon,k}$ and its gradient flows developed in the previous sections to prove two of our main results: Thorems 1.1 and 1.7(i). Our strategy to prove these three theorems is to leverage the framework of Γ -convergence.

In Subsection 5.1, we begin by proving the Γ -convergence of the energies $\mathcal{F}_{\varepsilon,k}$ to the energy \mathcal{F}_k , in the sense of Definition 2.14. It is then a routine application of standard Calculus of Variations techniques to conclude that, as long as the external potentials offer sufficient compactness, minimizers of $\mathcal{F}_{\varepsilon,k}$ converge to a minimizers of \mathcal{F}_k as $\varepsilon \to 0$. This yields Theorem 1.7(i).

Next, in Subsection 5.2.1, we prove the Γ -convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ to the the gradient flow of \mathcal{F}_k , in the sense of Serfaty [69]; see Theorem 2.15 in which we recall her approach. This ensures that, given a sequence of gradient flows of $\mathcal{F}_{\varepsilon,k}$ with "well-prepared" initial data (bounded entropy and energy), the gradient flows of $\mathcal{F}_{\varepsilon,k}$ converge to a gradient flow of \mathcal{F}_k as $\varepsilon \to 0$, completing the proof of Theorem 1.1.

5.1. Γ -convergence of the energies and convergence of minimizers. We now turn to the proof of Theorem 1.7(i), that minimizers of $\mathcal{F}_{\varepsilon,k}$ converge to a minimizer of \mathcal{F}_k . We begin by showing the Γ convergence of the energies $\mathcal{F}_{\varepsilon,k}$ to the energy \mathcal{F}_k , in the sense of Definition 2.14.

Theorem 5.1 (Γ -convergence of energies). Assume (T), (M), (V), and (C) hold. Fix $k \in \mathbb{N}$. Then the energies $\mathcal{F}_{\varepsilon,k}$ Γ -converge to \mathcal{F}_k as $\varepsilon \to 0$. In particular, for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $\lim_{\varepsilon \to 0} \mathcal{F}_{\varepsilon,k}(\mu) = \mathcal{F}_k(\mu)$.

Proof. We begin with the proof of (2.12). Let ρ_{ε} narrowly converge to ρ . Lemma 2.3 implies,

(5.1)
$$\zeta_{\varepsilon} * \rho_{\varepsilon}$$
 narrowly converges to ρ .

By definition of $\mathcal{E}_{\varepsilon}$ and \mathcal{E} , we have, as in (1.10), $\mathcal{E}_{\varepsilon}(\rho_{\varepsilon}) = \mathcal{E}(\zeta_{\varepsilon} * \rho_{\varepsilon})$. Taking $\liminf_{\varepsilon \to 0}$ and using the lower semicontinuity of \mathcal{E} with respect to narrow convergence, as well as (5.1), we obtain,

$$\liminf_{\varepsilon \to 0} \mathcal{E}_{\varepsilon}(\rho_{\varepsilon}) = \liminf_{\varepsilon \to 0} \mathcal{E}(\zeta_{\varepsilon} * \rho_{\varepsilon}) \ge \mathcal{E}(\rho).$$

For the $\mathcal{V}_{\varepsilon}$ term, we first use the properties of convolution, followed by the assumption $V \in C_b(\mathbb{R}^d)$ and (5.1), to find,

(5.2)
$$\int_{\mathbb{R}^d} (\zeta_{\varepsilon} * V) \, d\rho_{\varepsilon} = \int_{\mathbb{R}^d} V(\zeta_{\varepsilon} * \rho_{\varepsilon}) \, d\mathcal{L}^d \to \int_{\mathbb{R}^d} V \, d\rho.$$

Finally, Lemma 3.1 ensures \mathcal{V}_k is lower semicontinuous with respect to narrow convergence. Together with the definitions of $\mathcal{F}_{\varepsilon,k}$ and \mathcal{F}_k , this concludes the proof of (2.12).

Now we establish (2.13). Let $\rho \in \mathcal{P}(\mathbb{R}^d)$. Taking $\rho_{\varepsilon} = \rho$ for all $\varepsilon > 0$ in (5.2), and using the fact that \mathcal{V}_k is independent of ε , we find that it suffices to prove $\limsup_{\varepsilon \to 0} \mathcal{E}_{\varepsilon}(\rho) \leq \mathcal{E}(\rho)$. Without loss of generality, we assume ρ is such that $\mathcal{E}(\rho) < +\infty$, otherwise, the desired inequality is trivially true. Together with the definition of \mathcal{E} and our assumption (T) that $\bar{\rho}$ is bounded uniformly above and below, we deduce $\rho \in L^2(\mathbb{R}^d)$. We use the definition of $\mathcal{E}_{\varepsilon}$ to find,

$$2\mathcal{E}_{\varepsilon}(\rho) = \int_{\mathbb{R}^d} |\zeta_{\varepsilon} * \rho|^2(x) \frac{1}{\bar{\rho}(x)} \, dx = \int_{\mathbb{R}^d} \left| \int_{\mathbb{R}^d} \zeta_{\varepsilon}(x-y)\rho(y) \, dy \right|^2 \frac{1}{\bar{\rho}(x)} \, dx$$

Next we use Jensen's inequality, followed by Fubini's Theorem, to obtain,

(5.3)
$$2\mathcal{E}_{\varepsilon}(\rho) \leq \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \zeta_{\varepsilon}(x-y)\rho(y)^2 \frac{1}{\bar{\rho}(x)} \, dy \, dx = \int_{\mathbb{R}^d} \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}}\right)(y)\rho^2(y) \, dy.$$

We shall now prove:

(5.4)
$$\lim_{\varepsilon \to 0} \left| \int_{\mathbb{R}^d} \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}} \right) (y) \rho^2(y) \, dy - 2\mathcal{E}(\rho) \right| = 0.$$

Together with (5.3), this will yield the desired result.

In order to establish (5.4), we first use the definition of $\mathcal{E}(\rho)$ to write,

(5.5)
$$\left| \int_{\mathbb{R}^d} \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}} \right) \rho^2 \, d\mathcal{L}^d - 2\mathcal{E}(\rho) \right| = \left| \int_{\mathbb{R}^d} \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}} \right) \rho^2 \, d\mathcal{L}^d - \int_{\mathbb{R}^d} \frac{\rho^2}{\bar{\rho}} \, d\mathcal{L}^d \right| \le \int_{\mathbb{R}^d} \left| \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}} \right) - \frac{1}{\bar{\rho}} \right| \rho^2 \, d\mathcal{L}^d.$$

Fix $\delta > 0$ arbitrary. Since $\rho \in L^2(\mathbb{R}^d)$, there exists R > 0 such that $\int_{B_R^c} \rho^2 \leq \delta$. Moreover, since $1/\bar{\rho}$ is uniformly bounded (see Assumption (T)),

$$\int_{B_R^c} \left| \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}} \right) - \frac{1}{\bar{\rho}} \right| \rho^2 \, d\mathcal{L}^d \le C \int_{B_R^c} \rho^2 \le C\delta,$$

where C is independent of δ and ε . Now, splitting the integral in (5.5) into integrals over B_R and B_R^c , we find,

$$\left| \int_{\mathbb{R}^d} \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}} \right) \rho^2 \, d\mathcal{L}^d - 2\mathcal{E}(\rho) \right| \le \int_{B_R} \left| \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}} \right) - \frac{1}{\bar{\rho}} \right| \rho^2 \, d\mathcal{L}^d + C\delta \le \left\| \left(\zeta_{\varepsilon} * \frac{1}{\bar{\rho}} \right) - \frac{1}{\bar{\rho}} \right\|_{L^{\infty}(B_R)} \|\rho\|_{L^2(\mathbb{R}^d)} + C\delta.$$

Since $1/\bar{\rho}$ is continuous, $\zeta_{\varepsilon} * \frac{1}{\bar{\rho}}$ converges to $1/\bar{\rho}$ uniformly on compact subsets of \mathbb{R}^d as $\varepsilon \to 0$. In particular, we may choose $\varepsilon > 0$ small enough so that the value of the right-hand side of the previous line is no larger than δ . Since $\delta > 0$ was arbitrary, this completes the proof of estimate (5.4) and therefore of the lemma. \Box

We now apply this result to complete the proof of Theorem 1.7(i), which ensures that, under the assumption that the sublevel sets of V_k are compact, minimizers of $\mathcal{F}_{\varepsilon,k}$ converge to a minimizer of \mathcal{F}_k .

Proof of Theorem 1.7(i). The proof is classical, following directly from Theorem 5.1 and the fact that, since the sublevel sets of V_k are compact, a sequence of minimizers of $\mathcal{F}_{\varepsilon,k}$ is tight. For more details, see a similar argument in previous work by the first author [20, Theorem 4.5].

5.2. Γ -convergence of the gradient flows. We now turn to the proof of Theorem 1.1. We seek to show that gradient flows of $\mathcal{F}_{\varepsilon,k}$ with "well-prepared" initial data converge to a gradient flow of \mathcal{F}_k . As previously described, our approach relies on Serfaty's framework of Γ -convergence of gradient flows, so we seek to verify the hypotheses in Theorem 2.15. In general, the most challenging hypothesis to verify in this framework is hypothesis (2.17), which ensures lower semicontinuity of the metric slopes of $\mathcal{F}_{\varepsilon,k}$ along a sequence of gradient flows $\rho_{\varepsilon}(t)$,

(5.6)
$$\liminf_{\varepsilon \to 0} |\partial \mathcal{F}_{\varepsilon,k}|^2(\rho_{\varepsilon}(t)) \ge |\partial \mathcal{F}_k|^2(\rho(t)).$$

In Subsection 5.2.1, we prove Proposition 5.2, which provides sufficient conditions to ensure this lower semicontinuity. In Subsection 5.2.2, we apply this to conclude convergence of the gradient flows for "well-prepared" initial data.

5.2.1. Lower semicontinuity of metric slopes. We begin by stating sufficient conditions under which the lower semicontinuity inequality for the metric slopes holds.

Proposition 5.2 (lower semicontinuity of metric slopes). Assume (T), (M), (V), and (C) hold. Consider a sequence ρ_{ε} in $\mathcal{P}_2(\mathbb{R}^d)$ satisfying,

(5.7)
$$\sup_{\varepsilon>0} \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}) < +\infty,$$

(5.8)
$$\liminf_{\varepsilon \to 0} \|\nabla \zeta_{\varepsilon} * \rho_{\varepsilon}\|_{L^{2}(\mathbb{R}^{d})} < +\infty, \text{ and}$$

(5.9)
$$\liminf_{\varepsilon \to 0} \int \left| \nabla \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \rho_{\varepsilon}}{\bar{\rho}} \right) \right|^2 d\rho_{\varepsilon} < +\infty.$$

In addition, suppose there exists $\rho \in \mathcal{P}(\mathbb{R}^d)$ such that ρ_{ε} narrowly converges to ρ . Then $\rho^2 \in W^{1,1}(\mathbb{R}^d)$, and there exists $\eta \in L^2(\rho)$, with,

(5.10)
$$\boldsymbol{\eta}\rho = \frac{\bar{\rho}}{2}\nabla\left(\frac{\rho^2}{\bar{\rho}^2}\right) + \rho\nabla(V+V_k),$$

and such that,

(5.11)
$$\liminf_{\varepsilon \to 0} \int \left| \nabla \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \rho_{\varepsilon}}{\bar{\rho}} \right) + \nabla (\zeta_{\varepsilon} * V) + \nabla V_k \right|^2 d\rho_{\varepsilon} \ge \int |\boldsymbol{\eta}|^2 d\rho$$

Remark 5.3 (lower semicontinuity of metric slopes and log-concavity). Note that, under assumptions (T), (M), (V), and (C), if we impose the additional assumption that $\bar{\rho}$ is log-concave, then, by Propositions 3.8-3.9, we have (5.11) is equivalent to (5.6). However, we emphasize that the conclusion of Proposition 5.7 holds even for $\bar{\rho}$ not log-concave.

The remainder of this subsection is devoted to the proof of Proposition 5.2. We begin with a preliminary lemma, showing that, under the assumptions of Proposition 5.2, we may upgrade the convergence of $\zeta_{\varepsilon} * \rho_{\varepsilon}$ to ρ from narrow convergence to convergence in $L^2_{loc}(\mathbb{R}^d)$.

Lemma 5.4 (upgraded convergence of $\zeta_{\varepsilon} * \rho_{\varepsilon}$). Assume (T), (M), (V), and (C) hold. Consider any sequence ρ_{ε} in $\mathcal{P}(\mathbb{R}^d)$ and $\rho \in \mathcal{P}(\mathbb{R}^d)$ such that ρ_{ε} narrowly converges to ρ and (5.7) and (5.8) are satisfied. Then $\rho \in L^2(\mathbb{R}^d)$, and there exists a subsequence (still denoted ρ_{ε}) along which we have,

(5.12)
$$\sup_{\varepsilon>0} ||\zeta_{\varepsilon} * \rho_{\varepsilon}||_{H^{1}(\mathbb{R}^{d})} < +\infty \text{ and,}$$

(5.13)
$$\zeta_{\varepsilon} * \rho_{\varepsilon} \text{ converges to } \rho \text{ in } L^{2}_{loc}(\mathbb{R}^{d}).$$

Proof of Lemma 5.4. By assumption (5.7) and the definition of $\mathcal{F}_{\varepsilon,k}$, we find,

$$(5.14) + \infty > \sup_{\varepsilon > 0} \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}) + \|V\|_{L^{\infty}(\mathbb{R}^d)} \ge \sup_{\varepsilon > 0} \mathcal{E}_{\varepsilon}(\rho_{\varepsilon}) = \sup_{\varepsilon > 0} \frac{1}{2} \int_{\mathbb{R}^d} \frac{|\zeta_{\varepsilon} * \rho|^2}{\bar{\rho}} \ge \frac{1}{2\|\bar{\rho}\|_{L^{\infty}(\mathbb{R}^d)}} \sup_{\varepsilon > 0} \int_{\mathbb{R}^d} |\zeta_{\varepsilon} * \rho|^2.$$

Similarly, since Theorem 5.1 ensures the Γ -convergence of $\mathcal{F}_{\varepsilon,k}$ to \mathcal{F}_k , statement (2.13) in Definition 2.14 of Γ -convergence ensures,

$$+\infty > \sup_{\varepsilon > 0} \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon}) + \|V\|_{L^{\infty}(\mathbb{R}^{d})} \ge \mathcal{F}_{k}(\rho) + \|V\|_{L^{\infty}(\mathbb{R}^{d})} \ge \mathcal{E}(\rho) \ge \frac{1}{2\|\bar{\rho}\|_{L^{\infty}(\mathbb{R}^{d})}} \int_{\mathbb{R}^{d}} \rho^{2} dx,$$

so $\rho \in L^2(\mathbb{R}^d)$.

Combining assumption (5.8) with the estimate (5.14) we find that, up to a subsequence, (5.12) holds. Therefore, by the Rellich-Kondrachov embedding theorem, we find that, up to another subsequence, $\zeta_{\varepsilon} * \rho_{\varepsilon}$ converges in $L^2_{loc}(\mathbb{R}^d)$. On the other hand, Lemma 2.3, implies that $\zeta_{\varepsilon} * \rho_{\varepsilon}$ narrowly converges to ρ . The uniqueness of limits therefore implies (5.13).

A key step in proving lower semicontinuity of the metric slopes of $\mathcal{F}_{\varepsilon,k}$, as in Proposition 5.2, is to identify the weak limit of $\nabla \zeta_{\varepsilon} * \left(\frac{1}{\bar{\rho}} \left(\zeta_{\varepsilon} * \rho_{\varepsilon}\right)\right)$ in $L^{1}(\rho_{\varepsilon})$. With this weak limit in hand, lower semicontinuity of the metric slopes will then follow from general results due to Ambrosio, Gigli, and Savaré on lower semicontinuity of integral functions with varying measures [3, Theorem 5.4.4 (ii)]. In the following lemma, we characterize the weak limit.

Lemma 5.5 (weak limit of subdifferentials). Assume (T), (M), (V), and (C) hold. Consider any sequence ρ_{ε} in $\mathcal{P}(\mathbb{R}^d)$ and $\rho \in \mathcal{P}(\mathbb{R}^d)$ such that ρ_{ε} narrowly converges to ρ and (5.7), (5.8), and (5.9) are satisfied. For all $\varepsilon > 0$ and $f \in C_c^{\infty}(\mathbb{R}^d)$, define,

(5.15)
$$L_{\varepsilon}(f) = \int_{\mathbb{R}^d} f\left(\nabla\zeta_{\varepsilon} * \left(\frac{1}{\bar{\rho}}\left(\zeta_{\varepsilon} * \rho_{\varepsilon}\right)\right)\right) d\rho_{\varepsilon} \quad and \quad L(f) = \int_{\mathbb{R}^d} -\frac{1}{2}\nabla\left(\frac{f}{\bar{\rho}}\right)\rho^2 dx + \int_{\mathbb{R}^d} f\rho^2\nabla\left(\frac{1}{\bar{\rho}}\right) dx.$$

There exists a subsequence, still denoted by ε , so that, for any $f \in C_c^{\infty}(\mathbb{R}^d)$, we have,

(5.16)
$$\lim_{\varepsilon \to 0} L_{\varepsilon}(f) = L(f)$$

Furthermore, L is a bounded linear operator on $L^2(\rho)$.

Proof. By Lemma 5.4, we may choose a subsequence, still denoted ρ_{ε} , along which (5.12) and (5.13) hold.

In order to characterize $\lim_{\varepsilon \to 0} L_{\varepsilon}(f)$, we begin by breaking up the expression for $L_{\varepsilon}(f)$ into two terms, which we will estimate separately. Using the definition of L_{ε} and properties of convolution, we find that, for any $f \in C_c^{\infty}(\mathbb{R}^d)$,

$$L_{\varepsilon}(f) = \int_{\mathbb{R}^d} f\left(\zeta_{\varepsilon} * \left(\frac{1}{\bar{\rho}}\nabla\left(\zeta_{\varepsilon} * \rho_{\varepsilon}\right)\right)\right) d\rho_{\varepsilon} + \int_{\mathbb{R}^d} f\left(\zeta_{\varepsilon} * \left(\nabla\left(\frac{1}{\bar{\rho}}\right)\left(\zeta_{\varepsilon} * \rho_{\varepsilon}\right)\right)\right) d\rho_{\varepsilon}$$
$$= \int_{\mathbb{R}^d} \left(\left(f\rho_{\varepsilon}\right) * \zeta_{\varepsilon}\right) \left(\frac{1}{\bar{\rho}}\nabla\left(\zeta_{\varepsilon} * \rho_{\varepsilon}\right)\right) d\mathcal{L}^d + \int_{\mathbb{R}^d} \left(\left(f\rho_{\varepsilon}\right) * \zeta_{\varepsilon}\right) \left(\left(\nabla\frac{1}{\bar{\rho}}\right)\left(\zeta_{\varepsilon} * \rho_{\varepsilon}\right)\right) d\mathcal{L}^d$$
$$=: I_{\varepsilon}(f) + J_{\varepsilon}(f).$$

We begin by showing,

(5.18)
$$\lim_{\varepsilon \to 0} J_{\varepsilon}(f) = \int_{\mathbb{R}^d} f \rho^2 \nabla \left(\frac{1}{\bar{\rho}}\right) d\mathcal{L}^d$$

To this end, we apply Lemma 2.1, with $\sigma = \zeta_{\varepsilon} * \rho_{\varepsilon} \nabla \left(\frac{1}{\bar{\rho}}\right) d\mathcal{L}^d$ and $\nu = \rho_{\varepsilon}$ to find, for $C_{\zeta} > 0$ as in assumption (M), there exist $p, L_f > 0$, and $C_{\bar{\rho}} > 0$ so that,

$$\begin{aligned} \left| J_{\varepsilon}(f) - \int_{\mathbb{R}^d} f(\zeta_{\varepsilon} * \rho_{\varepsilon})^2 \nabla\left(\frac{1}{\bar{\rho}}\right) d\mathcal{L}^d \right| &\leq \varepsilon^p L_f\left(\int_{\mathbb{R}^d} (\zeta_{\varepsilon} * \rho_{\varepsilon})^2 |\nabla(1/\bar{\rho})| \, d\mathcal{L}^d + C_{\zeta} \int (\zeta_{\varepsilon} * \rho_{\varepsilon}) \left| \nabla \frac{1}{\bar{\rho}} \right| \, d\mathcal{L}^d \right) \\ &\leq \varepsilon^p L_f C_{\bar{\rho}} \left(\int_{\mathbb{R}^d} (\zeta_{\varepsilon} * \rho_{\varepsilon})^2 \, d\mathcal{L}^d + C_{\zeta} \right). \end{aligned}$$

By (5.12) of Lemma 5.4, the right-hand side converges to 0 as $\varepsilon \to 0$, which implies that (5.18) holds.

Next, we consider $\lim_{\varepsilon \to 0} I_{\varepsilon}(f)$. For any $f \in C_c^{\infty}(\mathbb{R}^d)$, define,

$$\tilde{I}_{\varepsilon}(f) = \frac{1}{2} \int_{\mathbb{R}^d} \frac{f}{\bar{\rho}} \nabla \left(\left(\zeta_{\varepsilon} * \rho_{\varepsilon} \right)^2 \right) \, d\mathcal{L}^d$$

Note that the L^2 convergence of $\zeta_{\varepsilon} * \rho_{\varepsilon}$ to ρ established in (5.13) of Lemma 5.4 ensures that, for any $f \in C_c^{\infty}(\mathbb{R}^d)$,

$$\lim_{\varepsilon \to 0} \tilde{I}_{\varepsilon}(f) = \lim_{\varepsilon \to 0} -\frac{1}{2} \int_{\mathbb{R}^d} \nabla\left(\frac{f}{\bar{\rho}}\right) (\zeta_{\varepsilon} * \rho_{\varepsilon})^2 \, d\mathcal{L}^d = -\frac{1}{2} \int_{\mathbb{R}^d} \nabla\left(\frac{f}{\bar{\rho}}\right) \rho^2 \, d\mathcal{L}^d.$$

Thus, to complete our proof that $\lim_{\varepsilon \to 0} L_{\varepsilon}(f) = L(f)$, it suffices to prove that, for any $f \in C_c^{\infty}(\mathbb{R}^d)$, (5.19) $\lim_{\varepsilon \to 0} |I_{\varepsilon}(f) - \tilde{I}_{\varepsilon}(f)| = 0.$

Using the definitions of $I_{\varepsilon}(f)$ and $\tilde{I}_{\varepsilon}(f)$, followed by some rearranging, we find,

$$\begin{aligned} |I_{\varepsilon}(f) - \tilde{I}_{\varepsilon}(f)| &= \left| \int_{\mathbb{R}^d} \left((f\rho_{\varepsilon}) * \zeta_{\varepsilon} \right) (x) \frac{1}{\bar{\rho}(x)} \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon})(x) \, dx - \int_{\mathbb{R}^d} f(x) (\zeta_{\varepsilon} * \rho_{\varepsilon})(x) \frac{1}{\bar{\rho}(x)} \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon})(x) \, dx \right|, \\ &= \left| \int_{\mathbb{R}^d} \left[\left((f\rho_{\varepsilon}) * \zeta_{\varepsilon} \right) (x) - f(x) (\zeta_{\varepsilon} * \rho_{\varepsilon})(x) \right] \frac{1}{\bar{\rho}(x)} \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon})(x) \, dx \right|. \end{aligned}$$

We have, for all $x \in \mathbb{R}^d$,

$$((f\rho_{\varepsilon})*\zeta_{\varepsilon})(x) - f(x)(\zeta_{\varepsilon}*\rho_{\varepsilon})(x) = \int_{\mathbb{R}^d} (f(y) - f(x))\rho_{\varepsilon}(y)\zeta_{\varepsilon}(x-y)\,dy$$

Thus, using this and Fubini's Theorem we find,

$$|I_{\varepsilon}(f) - \tilde{I}_{\varepsilon}(f)| \le \iint |f(y) - f(x)| \,\rho_{\varepsilon}(y)\zeta_{\varepsilon}(x - y)\frac{1}{\bar{\rho}(x)} \,|\nabla(\zeta_{\varepsilon} * \rho_{\varepsilon})(x)| \,dy\,dx$$

Since $\bar{\rho}$ is bounded uniformly from below and $f \in C_c^{\infty}(\mathbb{R}^d)$,

(5.20)
$$|I_{\varepsilon}(f) - \tilde{I}_{\varepsilon}(f)| \leq \frac{\|\nabla f\|_{\infty}}{\inf \bar{\rho}} \iint |x - y| \rho_{\varepsilon}(y)\zeta_{\varepsilon}(x - y) |\nabla(\zeta_{\varepsilon} * \rho_{\varepsilon})(x)| \, dy \, dx$$

Next, we claim that there exist $C > 0, \gamma \in (0, 1)$ and $\delta > 1$, all depending only on ζ , such that,

(5.21)
$$\zeta_{\varepsilon}(x-y)|x-y| \le C\varepsilon^{\delta} \quad \text{for } |x-y| > \varepsilon^{\gamma}.$$

Indeed, let q be as in Assumption (M), define $\delta' = q - (d+1) > 0$ and $\gamma = \delta'/2(d+\delta')$. The definition of ζ_{ε} and assumption (M) imply,

$$\zeta_{\varepsilon}(z)|z| = \zeta\left(\frac{z}{\varepsilon}\right)\frac{|z|}{\varepsilon^d} \le C|z|^{-(d+1+\delta')}\varepsilon^{d+1+\delta'}|z|\varepsilon^{-d} = C|z|^{-d-\delta'}\varepsilon^{1+\delta'}.$$

Thus, for $|z| > \varepsilon^{\gamma}$ we obtain, $\zeta_{\varepsilon}(z)|z| \leq C\varepsilon^{-(d+\delta')\gamma}\varepsilon^{1+\delta'} = C\varepsilon^{1+\delta'/2}$. The inequality (5.21) now follows by taking $\delta = 1 + \delta'/2$.

Thus, breaking up the integral on the right-hand side of (5.20) into two regions and using (5.21), we find, $|I_{\varepsilon}(f) - \tilde{I}_{\varepsilon}(f)|$

$$\leq \frac{\|\nabla f\|_{\infty}}{\inf \bar{\rho}} \left(\varepsilon^{\gamma} \iint_{|x-y|<\varepsilon^{\gamma}} \rho_{\varepsilon}(y) \zeta_{\varepsilon}(x-y) \left| \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon})(x) \right| \, dy \, dx + C \varepsilon^{\delta} \iint_{|x-y|>\varepsilon^{\gamma}} \rho_{\varepsilon}(y) \left| \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon})(x) \right| \, dy \, dx \right)$$

$$\leq \frac{\|\nabla f\|_{\infty}}{\inf \bar{\rho}} \left(\varepsilon^{\gamma} \int (\rho_{\varepsilon} * \zeta_{\varepsilon})(x) \left| \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon})(x) \right| \, dx + C \varepsilon^{\delta} \int \left| (\nabla \zeta_{\varepsilon} * \rho_{\varepsilon})(x) \right| \, dx \right).$$

Now we use Hölder's inequality for the first term on the right-hand side, and Young's inequality for the second term to obtain,

$$|I_{\varepsilon}(f) - \tilde{I}_{\varepsilon}(f)| \le C_f \left(\varepsilon^{\gamma} ||\zeta_{\varepsilon} * \rho_{\varepsilon}||_{L^2(\mathbb{R}^d)} ||\nabla \zeta_{\varepsilon} * \rho_{\varepsilon}||_{L^2(\mathbb{R}^d)} + \varepsilon^{\delta} ||\nabla \zeta_{\varepsilon}||_{L^1(\mathbb{R}^d)}\right).$$

To bound the first term on the right-hand side we recall that $\zeta_{\varepsilon} * \rho_{\varepsilon}$ is bounded in $H^1(\mathbb{R}^d)$ uniformly in ε (see the estimate (5.12) from Lemma 5.4). For the second term, we note $\varepsilon^{\delta} ||\nabla \zeta_{\varepsilon}||_{L^1(\mathbb{R}^d)} = \varepsilon^{\delta-1} ||\nabla \zeta||_{L^1(\mathbb{R}^d)}$. Since $\gamma > 0$ and $\delta - 1 > 0$, this ensures $\lim_{\varepsilon \to 0} |I_{\varepsilon}(f) - \tilde{I}_{\varepsilon}(f)| = 0$, which completes the proof that $\lim_{\varepsilon \to 0} L_{\varepsilon}(f) = L(f)$.

It remains to show that L is a bounded linear operator on $L^2(\rho)$. We will show that, for any $f \in C_c^{\infty}(\mathbb{R}^d)$,

$$|L(f)| \le C ||f||_{L^2(\rho)}$$

Indeed, since $\rho \in L^2(\mathbb{R}^d)$, ρ is a Radon measure, so $C_c^1(\mathbb{R}^d)$ is dense in $L^2(\rho)$ [9, Corollary 4.2.2], and there exists a unique extension of L to $L^2(\rho)$ enjoying the same bound.

Fix arbitrary $f \in C_c^{\infty}(\mathbb{R}^d)$. By definition of L_{ε} in equation (5.15) and Hölder's inequality,

$$|L_{\varepsilon}(f)| \leq ||f||_{L^{2}(\rho_{\varepsilon})} \left\| \nabla \zeta_{\varepsilon} * \left(\frac{1}{\bar{\rho}} \left(\zeta_{\varepsilon} * \rho_{\varepsilon} \right) \right) \right\|_{L^{2}(\rho_{\varepsilon})}.$$

Thus, by assumption (5.9), there exists C > 0 so that,

$$|L(f)| = \liminf_{\varepsilon \to 0} |L_{\varepsilon}(f)| \le C \liminf_{\varepsilon \to 0} ||f||_{L^{2}(\rho_{\varepsilon})} = C||f||_{L^{2}(\rho)},$$

which gives the result.

We now apply the previous lemmas to prove our main proposition, ensuring lower semicontinuity of the metric slopes.

Proof of Proposition 5.2. Choose a subsequence, still denoted by ρ_{ε} , so that,

$$\lim_{\varepsilon \to 0} |\partial \mathcal{F}_{\varepsilon,k}|(\rho_{\varepsilon}) = \liminf_{\varepsilon \to 0} |\partial \mathcal{F}_{\varepsilon,k}|(\rho_{\varepsilon}).$$

It suffices to show $\rho^2 \in W^{1,1}(\mathbb{R}^d)$, there exists $\eta \in L^2(\rho)$ satisfying (5.10), and, up to a further subsequence,

(5.22)
$$\lim_{\varepsilon \to 0} \int_{\mathbb{R}^d} f\left(\nabla\zeta_\varepsilon * \left(\frac{1}{\bar{\rho}}\left(\zeta_\varepsilon * \rho_\varepsilon\right)\right) + \nabla(\zeta_\varepsilon * V) + \nabla V_k\right) d\rho_\varepsilon = \int f\boldsymbol{\eta} \, d\rho \quad \text{for all } f \in C_c^\infty(\mathbb{R}^d).$$

The estimate (5.11) then follows by applying [3, Theorem 5.4.4 (ii)], completing the proof.

Notice that, for any $f \in C_c^{\infty}(\mathbb{R}^d)$, the fact that ∇V and ∇V_k are continuous and Lemma 2.3 ensure,

(5.23)
$$\lim_{\varepsilon \to 0} \int_{\mathbb{R}^d} f\left(\nabla(\zeta_{\varepsilon} * V)\right) d\rho_{\varepsilon} = \lim_{\varepsilon \to 0} \int_{\mathbb{R}^d} \nabla V\left(\zeta_{\varepsilon} * (f\rho_{\varepsilon})\right) d\mathcal{L}^d = \int f \nabla V d\rho$$

(5.24)
$$\lim_{\varepsilon \to 0} \int_{\mathbb{R}^d} f(\nabla V_k) \, d\rho_{\varepsilon} = \int f \nabla V_k \, d\rho$$

Next, we use the definitions of $L_{\varepsilon}(f)$ and L(f), as well as the convergence of $L_{\varepsilon}(f)$ to L(f) established in (5.16) of Lemma 5.5. Combining these with the Riesz Representation Theorem on $L^2(\rho)$ (which we can apply to the operator L due to, again, Lemma 5.5), we find that there exists $\tilde{\eta} \in L^2(\rho)$ such that,

$$\lim_{\varepsilon \to 0} \int_{\mathbb{R}^d} f\left(\nabla \zeta_{\varepsilon} * \left(\frac{1}{\bar{\rho}} \left(\zeta_{\varepsilon} * \rho_{\varepsilon}\right)\right)\right) \, d\rho_{\varepsilon} = \int_{\mathbb{R}^d} -\frac{1}{2} \nabla \left(\frac{f}{\bar{\rho}}\right) \rho^2 \, dx + \int_{\mathbb{R}^d} f \rho^2 \nabla \left(\frac{1}{\bar{\rho}}\right) \, dx = \int f \tilde{\eta} \, d\rho.$$

Rearranging, we obtain,

$$-\frac{1}{2}\int_{\mathbb{R}^d} \nabla\left(\frac{f}{\bar{\rho}}\right)\rho^2 \, dx = \int_{\mathbb{R}^d} f\tilde{\eta}\rho - f\rho^2 \nabla\left(\frac{1}{\bar{\rho}}\right) \, dx = \int_{\mathbb{R}^d} \frac{f}{\bar{\rho}}\left(\tilde{\eta}\rho\bar{\rho} - \bar{\rho}\rho^2 \nabla\left(\frac{1}{\bar{\rho}}\right)\right) \, dx.$$

Since the previous line holds for all $f \in C_c^{\infty}(\mathbb{R}^d)$, we deduce $\rho^2 \in W^{1,1}(\mathbb{R}^d)$ and

$$\nabla\left(\frac{\rho^2}{2}\right) = \tilde{\eta}\rho\bar{\rho} - \bar{\rho}\rho^2\nabla\left(\frac{1}{\bar{\rho}}\right).$$

Finally, by the chain rule for $W^{1,1}(\mathbb{R}^d)$ functions and the previous line, we have,

$$\nabla\left(\frac{\rho^2}{\bar{\rho}^2}\right) = \nabla(\rho^2)\frac{1}{\bar{\rho}^2} + \rho^2\nabla\left(\frac{1}{\bar{\rho}^2}\right) = \frac{1}{\bar{\rho}^2}\left(2\tilde{\eta}\rho\bar{\rho} - 2\bar{\rho}\rho^2\nabla\left(\frac{1}{\bar{\rho}}\right)\right) + \rho^2\nabla\left(\frac{1}{\bar{\rho}^2}\right) = 2\tilde{\eta}\frac{\rho}{\bar{\rho}}$$

Thus,

(5.25)
$$\tilde{\boldsymbol{\eta}}\rho = \frac{\bar{\rho}}{2}\nabla\left(\frac{\rho^2}{\bar{\rho}^2}\right)$$

Finally, defining $\eta = \tilde{\eta} + \nabla V + \nabla V_k$, the facts that $\nabla V \in L^{\infty}(\mathbb{R}^d)$ and $\nabla V_k \in L^2(\rho)$ (see sentence following Assumption (C), ensure $\eta \in L^2(\rho)$ and (5.22) holds.

5.2.2. Convergence of gradient flows. We now apply the lower semicontinuity of the metric slopes, obtained in Proposition 5.2, as well as the Γ -convergence of the energies, obtained in Theorem 5.1, to prove Theorem 1.1: that gradient flows of $\mathcal{F}_{\varepsilon,k}$ with "well-prepared" initial data converge to a gradient flow of \mathcal{F}_k .

Proof of Theorem 1.1. First, we will show that, up to a subsequence, the hypotheses of Serfaty's Γ-convergence of gradient flows framework, recalled in Theorem 2.15, are satisfied. The results of Section 3.1 ensure that $\mathcal{F}_{\varepsilon,k}$ and \mathcal{F}_k are proper, lower semicontinuous, and semiconvex along generalized geodesics. Hypothesis (2.12) holds by Theorem 5.1. Hypothesis (2.15) holds by assumption.

31

Next, we consider hypotheses (2.14) and (2.16). First, note that, up to a subsequence, we may assume,

(5.26)
$$\sup_{\varepsilon>0} \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}(0)) < +\infty.$$

Likewise, $\mathcal{F}_{\varepsilon,k}(\mu) \geq -\|V\|_{L^{\infty}(\mathbb{R}^d)}$ for all $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Thus, applying Lemma 2.16 with $\mathcal{G}_{\varepsilon} = \mathcal{F}_{\varepsilon,k}$, we obtain that there exists $\rho_k \in AC^2([0,T], \mathcal{P}_2(\mathbb{R}^d))$ so that (2.14) and (2.16) hold, up to a subsequence.

It remains to show (2.17). By Theorem 4.1, inequality (5.26), and assumption (1.20) of the present theorem, we obtain,

$$\liminf_{\varepsilon \to 0} \int_0^T ||\nabla \zeta_\varepsilon * \rho_{\varepsilon,k}(t)||_{L^2(\mathbb{R}^d)}^2 dt < +\infty.$$

Thus, by Fatou's lemma, for almost every $t \in [0, T]$,

(5.27)
$$\liminf_{\varepsilon \to 0} ||\nabla \zeta_{\varepsilon} * \rho_{\varepsilon,k}(t)||_{L^2(\mathbb{R}^d)} < +\infty.$$

Let us fix a time t so that (5.27) holds. We now apply Proposition 5.2 with $\rho_{\varepsilon,k}(t)$ and $\rho_k(t)$ instead of ρ_{ε} and ρ . We find $\rho_k(t)^2 \in W^{1,1}(\mathbb{R}^d)$, and that there exists $\eta_k(t) \in L^2(\rho_k(t))$ satisfying,

$$\boldsymbol{\eta}_k(t)\rho_k(t) = \frac{\bar{\rho}}{2}\nabla\left(\frac{\rho_k(t)^2}{\bar{\rho}^2}\right) + \nabla(\zeta_{\varepsilon} * V)\rho_k(t) + \nabla V_k\rho_k(t),$$

and such that,

(5.28)
$$\liminf_{\varepsilon \to 0} \int_{\mathbb{R}^d} \left| \nabla \zeta_{\varepsilon} * \left(\frac{1}{\bar{\rho}} \left(\zeta_{\varepsilon} * \rho_{\varepsilon,k}(t) \right) \right) + \nabla (\zeta_{\varepsilon} * V) + \nabla V_k \right|^2 d\rho_{\varepsilon,k}(t) \ge \int_{\mathbb{R}^d} |\boldsymbol{\eta}_k(t)|^2 d\rho_k(t)$$

Finally, we notice that, according to Propositions 3.8 and 3.9, the previous line is exactly (2.17).

Thus, we can apply Theorem 2.15, from which the conclusion of Theorem 1.1 follows immediately. Finally, we remark that, although we have established this for a subsequence of $\varepsilon \to 0$, the uniqueness of gradient flows of \mathcal{F}_k (Proposition 3.10) shows that the limit is the same along any subsequence, which implies that the whole sequence $\rho_{\varepsilon,k}$ converges to ρ_k .

	_	
	_	

6. Convergence of the gradient flows of \mathcal{F}_k to \mathcal{F}

We now show that, if the confining potentials V_k approximate the hard cutoff function V_{Ω} , which is defined in (1.17), on the convex domain Ω , then gradient flows of \mathcal{F}_k converge to gradient flows of \mathcal{F} as $k \to +\infty$. Our result generalizes work by Alasio, Bruna, and Carrillo [1] to the case of weighted porous medium equations. As in our previous theorem on the $\varepsilon \to 0$ limit, we use an approach based on Γ -convergence of gradient flows, which is different from the approach used in the aforementioned work [1]. We are optimistic this new approach will be more easily generalizable to a range of Wasserstein gradient flows.

We begin by showing Γ -convergence of the energies \mathcal{F}_k to \mathcal{F} , in the sense of Definition 2.14.

Theorem 6.1 (Γ convergence of energies \mathcal{F}_k to \mathcal{F}). Assume (T), (D), (V), (C), and (Ck). Then, the energies $\mathcal{F}_k \Gamma$ -converge to \mathcal{F} as $k \to \infty$. In particular, $\lim_{k\to+\infty} \mathcal{F}_k(\mu) = \mathcal{F}(\mu)$ for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$.

Proof. We first establish item (2.12). Without loss of generality, we may assume $\liminf_{k\to+\infty} \mathcal{F}_k(\rho_k) < +\infty$ and, up to a subsequence, that $\sup_{k\in\mathbb{N}} \mathcal{F}_k(\rho_k) < +\infty$. Using that $\mathcal{F}_k = \mathcal{E} + \mathcal{V} + \mathcal{V}_k$ and the fact that $\mathcal{E}(\rho_k)$ and $\mathcal{V}(\rho_k)$ are bounded below uniformly in k, this likewise gives,

(6.1)
$$\sup_{k \in \mathbb{N}} \mathcal{V}_k(\rho_k) < +\infty.$$

We claim that it suffices to show,

(6.2)
$$\lim_{k \to +\infty} \mathcal{V}_k(\rho_k) \ge \mathcal{V}_{\Omega}(\rho)$$

In particular, if this is true, then by the lower semicontinuity of \mathcal{E} and \mathcal{V} , see Lemmas 3.1 and 3.2, we have,

$$\liminf_{k \to +\infty} \mathcal{F}_k(\rho_k) = \liminf_{k \to +\infty} \mathcal{E}(\rho_k) + \mathcal{V}(\rho_k) + \mathcal{V}_k(\rho_k) \ge \mathcal{E}(\rho) + \mathcal{V}(\rho) + \mathcal{V}_{\Omega}(\rho) = \mathcal{F}(\rho)$$

which gives item (2.12).

To show inequality (6.2), (2.13), by definition of \mathcal{V}_k and \mathcal{V}_{Ω} , it suffices to prove that supp $\rho \subseteq \overline{\Omega}$, since then the left hand side of the inequality is nonnegative and the right hand side is zero. Suppose, for the sake of contradiction that supp $\rho \not\subseteq \overline{\Omega}$, so that there exists $x \in \overline{\Omega}^c$ and an open ball B containing x so that $B \subset \subset \overline{\Omega}^c$ and $\rho(B) > 0$. By the Portmanteau theorem, the fact that $\rho_k \to \rho$ narrowly ensures $\liminf_{k \to +\infty} \rho_k(B) \ge \rho(B) > 0$. Thus, up to taking another subsequence, we may assume that there exists $\delta > 0$ so that $\rho_k(B) \ge \delta$ for all $k \in \mathbb{N}$. By definition of \mathcal{V}_k , this implies,

$$\liminf_{k \to +\infty} \int_{\mathbb{R}^d} V_k d\rho_k \ge \liminf_{k \to +\infty} \int_B V_k d\rho_k \ge \liminf_{k \to +\infty} \left(\inf_{x \in B} V_k(x) \right) \rho_k(B) \ge \delta \liminf_{k \to +\infty} \left(\inf_{x \in B} V_k(x) \right) = +\infty,$$

where the last inequality follows from Assumption (Ck) on V_k . This contradicts (6.1). Thus, we must have supp $\rho \subseteq \overline{\Omega}$, which completes the proof of item (2.12).

It remains to prove item (2.13). To this end, we note that we may write $\mathcal{V}_{\Omega}(\rho) = \int V_{\Omega} d\rho$, where $V_{\Omega}(x)$ is given by (1.17). Assumption Ck on V_k implies $V_k(x) \leq V_{\Omega}(x)$ for all $x \in \mathbb{R}^d$. Therefore we find,

$$\limsup_{k \to +\infty} \mathcal{V}_k(\rho) = \limsup_{k \to +\infty} \int V_k d\rho \le \int V_\Omega d\rho = \mathcal{V}_\Omega(\rho),$$

and thus conclude by recalling the definitions of \mathcal{F}_k and \mathcal{F} .

As a corollary of Theorem 6.1, we obtain Theorem 1.7(ii): minimizers of \mathcal{F}_k converge to minimizers of \mathcal{F} . The additional assumption we add – that the V_k are all greater than V_1 – is natural in the context of taking the V_k 's to be diverging to $+\infty$ off of Ω .

Proof of Theorem 1.7(ii). Since \mathcal{F} is proper, take $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ such that $\mathcal{F}(\nu) < +\infty$. We now use that the ρ_k minimize \mathcal{F}_k , followed by the Γ -convergence of the \mathcal{F}_k to \mathcal{F} , to find,

$$\liminf_{k \to \infty} \mathcal{F}_k(\rho_k) \le \limsup_{k \to \infty} \mathcal{F}_k(\nu) < \mathcal{F}(\nu) < \infty.$$

Thus, up to taking a subsequence, we may assume $\mathcal{F}_k(\rho_k)$ is uniformly bounded. Next, the assumption $V_k \geq V_1$ implies,

$$\sup_{k\in\mathbb{N}}\int V_1\,d\rho_k\leq \sup_{k\in\mathbb{N}}\int V_k\,d\rho_k\leq \sup_{k\in\mathbb{N}}\mathcal{F}_k<+\infty.$$

Together with the assumption that the sublevel sets of V_1 are compact, this guarantees that the sequence ρ_k is tight; see [4, Remark 5.1.5]. The remainder of the argument is classical. See, for example, a similar argument in previous work by the first author [20, Theorem 4.5].

With the previous theorem on the Γ -convergence of the energies \mathcal{F}_k to \mathcal{F} , we now prove Theorem 1.2 on convergence of the gradient flows.

Proof of Theorem 1.2. We seek to apply Serfaty's framework for Γ-convergence of gradient flows, recalled in Theorem 2.15. We note that the hypotheses on semicontinuity and convexity are satisfied due to the results in Subsection 3.1, and Theorem 6.1 guarantees that \mathcal{F}_k and \mathcal{F} satisfy (2.12). Assumption (2.15) is satisfied due to the analogous assumptions (1.22).

We now apply Lemma 2.16 to demonstrate that, up to a subsequence, hypotheses (2.14) and (2.16) are satisfied. First, (1.22) ensures,

(6.3)
$$\sup_{k \in \mathbb{N}} \mathcal{F}_k(\rho_k(0)) < +\infty,$$

so assumption (2.18) of Lemma 2.16 is satisfied. Next, note that since \mathcal{E} and \mathcal{V}_k are non-negative and \mathcal{V} is bounded below, assumption (2.19) is also satisfied. Thus, Lemma 2.16 guarantees the existence of a subsequence of ρ_k along which (2.14) and (2.16) hold.

It remains to verify assumption (2.17). In particular, we will show that for all $t \ge 0$,

(6.4)
$$\liminf_{k \to +\infty} |\partial \mathcal{F}_k|^2(\rho_k(t)) \ge |\partial \mathcal{F}|^2(\rho(t)).$$

Theorem 2.15 then ensures that, up to a subsequence, the ρ_k converge to a gradient flow of \mathcal{F} with initial data $\rho(0)$. Finally, the uniqueness of gradient flows of \mathcal{F} (Proposition 3.10) ensures that the limit is the same along any subsequence, which implies that the whole sequence ρ_k converges to the unique gradient flow ρ with initial data $\rho(0)$.

First, combining inequality (6.3) with Theorem 2.13, we obtain,

(6.5)
$$\sup_{k\in\mathbb{N}}\sup_{t\geq 0}\mathcal{F}_k(\rho_k(t))\leq \sup_{k\in\mathbb{N}}\mathcal{F}_k(\rho_k(0))<+\infty$$

Furthermore, since $\rho_k(t) \rightarrow \rho(t)$ narrowly, Theorem 6.1 implies,

(6.6)
$$\sup_{t>0} \mathcal{F}(\rho(t)) < +\infty$$

Since \mathcal{V} and \mathcal{V}_k are bounded below, inequality (6.5) implies,

(6.7)
$$\left(\frac{\inf\bar{\rho}}{2}\right)\sup_{k\in\mathbb{N},t\geq0}\|\rho_k\|_2^2 \le \sup_{k\in\mathbb{N},t\geq0}\frac{1}{2}\int\frac{|\rho_k(t)|^2}{\bar{\rho}} = \sup_{k\in\mathbb{N},t\geq0}\mathcal{E}(\rho_k(t)) < +\infty$$

Likewise, inequality (6.6) implies $\rho(t) \in L^2(\mathbb{R}^d)$ and supp $\rho(t) \subseteq \overline{\Omega}$ for all $t \ge 0$. In particular, $\rho(t) = 0$ a.e. on Ω^c for all $t \ge 0$.

Fix $t \in [0, T]$. We seek to show (6.4). From now on, we will suppress dependence on t, for simplicity of notation. Without loss of generality, we may suppose that the left hand side of inequality (6.4) is finite and that, up to a subsequence, we have,

(6.8)
$$\sup_{k\in\mathbb{N}} |\partial \mathcal{F}_k|^2(\rho_k) < +\infty.$$

By Proposition 3.9, we have $(\rho_k/\bar{\rho})^2 \in W^{1,1}_{\text{loc}}(\mathbb{R}^d)$ and there exists $\boldsymbol{\xi}_k \in L^2(\rho_k)$ satisfying,

(6.9)
$$\boldsymbol{\xi}_k \rho_k = \frac{\bar{\rho}}{2} \nabla (\rho_k / \bar{\rho})^2 + \nabla V \rho_k + \nabla V_k \rho_k,$$

and such that $\|\boldsymbol{\xi}_k\|_{L^2(\rho_k)} = |\partial \mathcal{F}_k|(\rho_k)$. In particular, by (6.8), we have,

(6.10)
$$\sup_{k\in\mathbb{N}} \|\boldsymbol{\xi}_k\|_{L^2(\rho_k)} < +\infty.$$

Furthermore, since Assumption (V) ensures $\nabla V \in L^{\infty}$, applying the triangle inequality yields,

(6.11)
$$\sup_{k} \int \left| \frac{\rho}{2} \nabla (\rho_{k}/\bar{\rho})^{2} + \nabla V_{k} \rho_{k} \right| = \sup_{k} \left\| \boldsymbol{\xi}_{k} - \nabla V \right\|_{L^{1}(\rho_{k})} \leq \sup_{k} \left\| \boldsymbol{\xi}_{k} - \nabla V \right\|_{L^{2}(\rho_{k})} \\ \leq \sup_{k} \left| \partial \mathcal{F}_{k} \right| (\rho_{k}) + \left\| \nabla V \right\|_{L^{2}(\rho_{k})} \leq \sup_{k} \left| \partial \mathcal{F}_{k} \right| (\rho_{k}) + \left\| \nabla V \right\|_{\infty} < +\infty.$$

By a second application of Proposition 3.9, combined with [3, Theorem 5.4.4], in order to show (6.4), it suffices to show $(\rho/\bar{\rho})^2 \in W^{1,1}_{\text{loc}}(\Omega)$, and that there exists $\boldsymbol{\xi} \in L^2(\rho)$ satisfying:

(6.12)
$$\boldsymbol{\xi}\rho = \frac{\bar{\rho}}{2}\nabla(\rho/\bar{\rho})^2 + \nabla V\rho$$
 holds on Ω , and

(6.13)
$$\liminf_{k \to +\infty} \int_{\mathbb{R}^d} f\left(\frac{\bar{\rho}}{2} \nabla (\rho_k/\bar{\rho})^2 + \nabla V \rho_k + \nabla V_k \rho_k\right) = \int_{\Omega} f\left(\frac{\bar{\rho}}{2} \nabla (\rho/\bar{\rho})^2 + \nabla V \rho\right) \quad \text{for all } f \in C_c^{\infty}(\mathbb{R}^d),$$

where we use that $\rho = 0$ a.e. on Ω^c . By Assumption (V) on V, we have $f \nabla V \in C_b(\mathbb{R}^d)$, so since ρ_k narrowly converges to ρ , we find,

$$\liminf_{k \to +\infty} \int_{\mathbb{R}^d} f \nabla V \rho_k = \int_{\Omega} f \nabla V \rho , \quad \text{for a.e. } t \in [0, T].$$

Thus, (6.13) is equivalent to the claim that, for all $f \in C_c^{\infty}(\mathbb{R}^d)$,

(6.14)
$$\lim_{k \to +\infty} \inf_{\mathbb{R}^d} f\left(\frac{\bar{\rho}}{2} \nabla (\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k\right) = \int_{\Omega} f\frac{\bar{\rho}}{2} \nabla (\rho/\bar{\rho})^2$$

We will establish (6.14) for test functions $f \in C_c^{\infty}(\Omega)$. Then, we will extend to the general case of $f \in C_c^{\infty}(\mathbb{R}^d)$ via a cutoff function to obtain (6.14).

First, we consider the region Ω . By Assumption Ck, which ensures V_k vanishes on Ω for all k, inequality (6.11) implies,

(6.15)
$$\left(\frac{\inf\bar{\rho}}{2}\right)\sup_{k}\int_{\Omega}\left|\nabla(\rho_{k}/\bar{\rho})^{2}\right| \leq \sup_{k}\int_{\Omega}\left|\frac{\bar{\rho}}{2}\nabla(\rho_{k}/\bar{\rho})^{2}\right| < +\infty$$

Since $(\rho_k/\bar{\rho})^2 \in W^{1,1}_{\text{loc}}(\mathbb{R}^d)$, combining (6.7) and (6.15), we obtain that $(\rho_k/\bar{\rho})^2$ is bounded in $W^{1,1}(\Omega)$. Thus, up to a subsequence, $(\rho_k/\bar{\rho})^2$ converges in $L^1(\Omega)$ and almost everywhere to some $g \in L^1(\Omega)$ with $g \ge 0$.

Furthermore,

$$\begin{aligned} \left\| \frac{\rho_k}{\bar{\rho}} - \sqrt{g} \right\|_{L^1(\Omega)} &\leq \sqrt{|\Omega|} \left\| \frac{\rho_k}{\bar{\rho}} - \sqrt{g} \right\|_{L^2(\Omega)} \leq \sqrt{|\Omega|} \left(\int_{\Omega} \left| \frac{\rho_k}{\bar{\rho}} - \sqrt{g} \right| \left(\frac{\rho_k}{\bar{\rho}} + \sqrt{g} \right) \right)^{1/2} \\ &= \sqrt{|\Omega|} \left(\int_{\Omega} \left| \left(\frac{\rho_k}{\bar{\rho}} \right)^2 - g \right| \right)^{1/2}, \end{aligned}$$

so $\rho_k/\bar{\rho} \to \sqrt{g}$ in $L^1(\Omega)$. Combining this with the fact that $\rho_k \to \rho$ narrowly, we obtain $\sqrt{g} = \rho/\bar{\rho}$ a.e. on Ω . Therefore, for all $f \in C_c^{\infty}(\Omega)$, the fact that V_k vanishes on $\overline{\Omega}$ ensures,

(6.16)
$$\liminf_{k \to +\infty} \int_{\Omega} f\left(\frac{\bar{\rho}}{2} \nabla (\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k\right) = \liminf_{k \to +\infty} \int_{\Omega} f\frac{\bar{\rho}}{2} \nabla (\rho_k/\bar{\rho})^2 = -\liminf_{k \to +\infty} \int_{\Omega} \nabla \left(f\frac{\bar{\rho}}{2}\right) (\rho_k/\bar{\rho})^2 \\ = -\int_{\Omega} \nabla \left(f\frac{\bar{\rho}}{2}\right) (\rho/\bar{\rho})^2.$$

By inequality (6.11), the left hand side of the equation may be bounded above by,

$$\sup_{k} \|f\|_{\infty} \left\| \frac{\bar{\rho}}{2} \nabla (\rho_{k}/\bar{\rho})^{2} + \nabla V_{k}\rho_{k} \right\|_{L^{1}(\mathbb{R}^{d})} \leq \|f\|_{\infty} (\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty}).$$

so we conclude $(\rho/\bar{\rho})^2 \in BV(\mathbb{R}^d)$. Thus, for all $f \in C_c^{\infty}(\mathbb{R}^d)$,

(6.17)
$$-\int_{\Omega} \nabla \left(f \frac{\bar{\rho}}{2} \right) (\rho/\bar{\rho})^2 = \int_{\Omega} f \frac{\bar{\rho}}{2} \nabla (\rho/\bar{\rho})^2.$$

Next, we seek to apply the Riesz Representation Theorem to the operator,

$$L(f) = \int_{\Omega} f \frac{\bar{\rho}}{2} \nabla (\rho/\bar{\rho})^2.$$

We first verify the boundedness of this operator on $L^2(\rho; \Omega)$. To this end, we use the definition of L and the equalities (6.16) and (6.17) to find,

$$L(f) = \liminf_{k \to +\infty} \int_{\Omega} f\left(\frac{\bar{\rho}}{2} \nabla (\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k\right).$$

Recalling the definition of $\boldsymbol{\xi}_k$ in (6.9), then using Hölder's inequality, and finally using the estimate (6.10) and the boundedness of ∇V , we obtain,

$$\begin{split} \int_{\Omega} \left| f\left(\frac{\bar{\rho}}{2} \nabla (\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k\right) \right| &= \int_{\Omega} \left| f(\boldsymbol{\xi}_k \rho_k - \nabla V \rho_k) \right| \\ &\leq \|f\|_{L^2(\rho_k;\Omega^\circ)} (\|\boldsymbol{\xi}_k\|_{L^2(\rho_k)} + \|\nabla V\|_{L^2(\rho_k;\Omega)}) \\ &\leq \|f\|_{L^2(\rho_k;\Omega)} \left(\sup_k |\partial \mathcal{F}_k|(\rho_k) + \|\nabla V\|_{L^\infty(\mathbb{R}^d)} \right). \end{split}$$

Finally, taking the limit in k, and using the narrow convergence of ρ_k to ρ , we find that the desired bound on L holds:

$$|L(f)| \le ||f||_{L^2(\rho;\Omega)} \left(\sup_k |\partial \mathcal{F}_k|(\rho_k) + ||\nabla V||_{L^{\infty}(\mathbb{R}^d)} \right).$$

Thus, by the Riesz Representation theorem, there exists $\mathbf{w} \in L^2(\rho; \Omega)$ so that,

$$\int f \frac{\bar{\rho}}{2} \mathbf{w} \rho = \int f \frac{\bar{\rho}}{2} \nabla (\rho/\bar{\rho})^2, \quad \text{for all } f \in C_c^{\infty}(\Omega).$$

Since $\|\mathbf{w}\rho\|_{L^1(\lambda^d;\Omega)} = \|\mathbf{w}\|_{L^1(\rho;\Omega)} \le \|\mathbf{w}\|_{L^2(\rho;\Omega)}$, this shows that $(\rho/\bar{\rho})^2 \in W^{1,1}_{\text{loc}}(\Omega)$. Likewise, $\boldsymbol{\xi} := \mathbf{w} + \nabla V \in L^2(\rho)$ satisfies the conditions of (6.12). Finally, integrating by parts on the right hand side of (6.16) gives (6.14) for all $f \in C_c^{\infty}(\Omega)$.

It remains to show that (6.14) holds for all $f \in C_c^{\infty}(\mathbb{R}^d)$. By the fact that we just showed it holds for test functions in $C_c^{\infty}(\Omega)$, for any smooth cutoff function $0 \leq \eta \leq 1$ that is compactly supported in Ω , we have,

$$\liminf_{k \to +\infty} \int_{\mathbb{R}^d} f\left(\frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k\right) = \liminf_{k \to +\infty} \int_{\mathbb{R}^d} (f\eta + f(1-\eta)) \left(\frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k\right)$$
$$= \underbrace{\int_{\Omega} f\eta \frac{\bar{\rho}}{2} \nabla(\rho/\bar{\rho})^2}_{I_1} + \underbrace{\liminf_{k \to +\infty} \int_{\mathbb{R}^d} f(1-\eta) \left(\frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k\right)}_{I_2}.$$

To estimate I_1 , note that,

$$|f\eta\frac{\bar{\rho}}{2}\nabla(\rho/\bar{\rho})^2| \le \frac{\|f\|_{\infty}\|\bar{\rho}\|_{\infty}}{2}|\mathbf{w}\rho| \in L^1(\lambda^d;\Omega).$$

To estimate I_2 , note that,

$$I_{2} \leq \liminf_{k \to +\infty} \|f(1-\eta)\|_{L^{2}(\rho_{k})} \left(|\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \right) \|_{L^{2}(\rho)} \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) \right) \|_{L^{2}(\rho)} \leq \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^{2}(\rho)} \left(\sup_{k} |\partial \mathcal{F}_{k}|(\rho_{k}) + \|\nabla V\|_{\infty} \right) + \|f(1-\eta)\|_{L^$$

where,

$$|f(1-\eta)|^2 \rho \le |f|^2 \rho \in L^1(\lambda^d; \Omega).$$

Thus, by the dominated convergence theorem, for all $\delta > 0$, choosing η sufficiently close to 1 pointwise on $\overline{\Omega}$, we obtain,

$$\left|\liminf_{k \to +\infty} \int_{\mathbb{R}^d} f\left(\frac{\bar{\rho}}{2} \nabla (\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k\right) - \int_{\Omega} f\frac{\bar{\rho}}{2} \nabla (\rho/\bar{\rho})^2 \right| \le \left|I_1 - \int_{\Omega} f\frac{\bar{\rho}}{2} \nabla (\rho/\bar{\rho})^2\right| + I_2 \le \delta.$$

Since $\delta > 0$ was arbitrary, this completes the proof of (6.14).

We conclude with the proof of Corollary 1.3.

Proof of Corollary 1.3. By Theorem 5.1,

$$\lim_{\varepsilon \to 0} \mathcal{F}_{\varepsilon,k}(\rho(0)) = \mathcal{F}_k(\rho(0))$$

so by Theorem 1.1, $\rho_{\varepsilon,k}(t) \xrightarrow{\varepsilon \to 0} \rho_k(t)$ narrowly, for all $t \in [0,T]$, where $\rho_k \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ is the gradient flow of \mathcal{F}_k with initial data $\rho(0)$. Furthermore, by Theorem 6.1,

$$\lim_{k \to +\infty} \mathcal{F}_k(\rho(0)) = \mathcal{F}(\rho(0)),$$

so, by Theorem 1.2, $\rho_k(t) \xrightarrow{k \to \infty} \rho(t)$ narrowly, for all $t \in [0,T]$, where $\rho \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ is the gradient flow of \mathcal{F} with initial data $\rho(0)$. Thus, by first choosing $k \in \mathbb{N}$ sufficiently large and then $\varepsilon > 0$ sufficiently small, we may find $\rho_{\varepsilon,k}(t)$ arbitrarily close to $\rho(t)$.

6.1. Extension to particle initial data and application to two-layer neural networks. In the previous sections, we have shown that gradient flows of $\mathcal{F}_{\varepsilon,k}$ with "well-prepared" initial data converge to a gradient flow of \mathcal{F} , as $k \to +\infty$, $\varepsilon = \varepsilon(k) \to 0$. Unfortunately, our assumption of "well-preparedness" requires that the initial data of $\mathcal{F}_{\varepsilon,k}$ have bounded entropy (1.20), which is a crucial assumption for obtaining the H^1 -type bound on the mollified gradient flow (Theorem 4.1) and the lower semicontinuity of the metric slopes (Proposition 5.2). This assumption explicitly excludes initial data given by an empirical measure.

We now use stability of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ to extend the convergence result to initial data given by an empirical measure, obtaining the proof of our third major theorem, Theorem 1.4. This is based on the elementary fact that that any measure $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ can be approximated to arbitrary accuracy by an empirical measure. For lack of a reference, we recall this in Lemma A.4. (In fact, our construction of the empirical measure in the proof of Lemma A.4 closely parallels what we employ in our numerical method.) It can be seen from the proof of Lemma A.4 that, if supp $\mu \subseteq [-R, R)^d$, then N can be taken to be the smallest integer larger than $(2\sqrt{dR}/\delta)^d$. More generally, in order for an empirical measure constructed from N i.i.d. samples of a measure $\mu \ll \mathcal{L}^d$ to converge to μ in the Wasserstein metric, N must scale like $O(1/\delta^d)$ [34,36]. Our requirement that the initial conditions of $\mathcal{F}_{\varepsilon,k}$ have bounded entropy implies $\mu \ll \mathcal{L}^d$, so this scaling requirement is sharp in our case. However, if μ were permitted to concentrate on lower dimensional sets, recent work by Weed and Bach has shown these requirements can be weakened [78].

Once we have extended our result to particle initial data, in Theorem 1.4, we are then able to quickly obtain our two main corollaries. Corollary 1.5 shows that, on bounded domains Ω and in the absence of an external potential V, the particle solution indeed approximates the target $\bar{\rho}$ in the long time limit. Next, Corollary 1.6 shows that the overparametrized limit of two-layer neural networks converges, as the variance of the radial basis function goes to zero, to a solution of (WPME), which is the gradient flow of a convex energy.

We begin with the proof of Theorem 1.4.

Proof of Theorem 1.4. First, let $\rho_{\varepsilon,k}(t)$ be the gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial data $\rho(0)$. By Corollary 1.3, as $k \to +\infty, \varepsilon = \varepsilon(k) \to 0$,

(6.18)
$$\rho_{\varepsilon,k}(t)$$
 narrowly converges to $\rho(t)$ for all $t \in [0,T]$

where $\rho(t)$ is the gradient flow of \mathcal{F} with initial data $\rho(0)$. By Proposition 3.10, ρ is the unique weak solution of (WPME). Recall from Lemmas 3.1-3.2 and Propositions 3.3 and 3.6 that $\mathcal{F}_{\varepsilon,k}$ is lower semicontinuous and λ_{ε} -convex along generalized geodesics with,

(6.19)
$$\lambda_{\varepsilon} = -\varepsilon^{-d-2} ||1/\bar{\rho}||_{L^{\infty}(\mathbb{R}^d)} ||D^2\zeta||_{L^{\infty}(\mathbb{R}^d)} + \inf_{\{x,\xi\in\mathbb{R}^d\}} \xi^t D^2 V(x)\xi,$$

and note that $-\infty < \lambda_{\varepsilon} \leq 0$.

By Proposition 3.13, the evolving empirical measure $\rho_{\varepsilon,k}^N(t)$, as defined in the statement of Theorem 1.4, is the unique gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial data $\rho_{\varepsilon,k}^N(0)$. By (6.18), it suffices to show that, as $k \to +\infty$, $\varepsilon = \varepsilon(k) \to 0$, $N = N(\varepsilon) \to +\infty$,

(6.20)
$$(\rho_{\varepsilon,k}^{N}(t) - \rho_{\varepsilon,k}(t)) \text{ narrowly converges to 0 for all } t \in [0,T]$$

Since $\rho_{\varepsilon,k}^N(t)$ and $\rho_{\varepsilon,k}$ are both gradient flows of the λ_{ε} -convex energy $\mathcal{F}_{\varepsilon,k}$, classical stability estimates for gradient flows [3, Theorem 11.2.1] ensure that, for all $t \in [0, T]$,

$$W_2(\rho_{\varepsilon,k}^N(t),\rho_{\varepsilon,k}(t)) \le e^{-\lambda_{\varepsilon}t} W_2(\rho_{\varepsilon,k}^N(0),\rho(0))$$

By hypothesis (1.24), the right hand side goes to zero, which by (2.4) completes the proof.

We now apply this to obtain the proof of Corollary 1.5.

Proof of Corollary 1.5. Let $\rho(t)$ be the solution of (WPME) with initial data $\rho(0)$, as in Theorem 1.4. By Proposition 3.10, ρ is the unique gradient flow of \mathcal{F} with initial data $\rho(0)$, so by Proposition 3.14,

$$\lim_{t \to +\infty} W_2\left(\rho(t), \mathbb{1}_{\overline{\Omega}}\bar{\rho}\right) = 0.$$

Thus, choosing t sufficiently large so that $\rho(t)$ is sufficiently close to $\mathbb{1}_{\overline{\Omega}}\bar{\rho}$, Theorem 1.4 gives the result. \Box

We conclude with the proof of Corollary 1.6.

Proof of Corollary 1.6. The evolving empirical measure $\rho_{\varepsilon}^{N}(t)$, as defined in the statement of Corollary 1.6, coincides with the evolving empirical measure in Theorem 1.4, in the case $\Omega = \mathbb{R}^{d}$, hence $V_{k} \equiv 0$. Thus, the convergence of $\rho_{\varepsilon,k}^{N}(t)$ to $\rho(t)$ is an immediate consequence of this theorem.

Furthermore, by Proposition 3.10, $\rho(t)$ is the unique gradient flow of \mathcal{F} . Expanding the square in the definition of \mathcal{R} and applying Tonelli's theorem, as in equation (1.14), we see that \mathcal{F} equals \mathcal{R} , up to a constant. By Definitions 2.9 and 2.12, the gradient flows of two energies coincide. Thus, $\rho(t)$ is the gradient flow of \mathcal{R} . Similarly, from Definition 2.6, we see that adding or subtracting a constant from an energy does not affect its convexity properties. Thus, Proposition 3.3 ensures that \mathcal{R} is convex.

7. NUMERICAL SIMULATION

We now implement the particle method described in Theorem 1.4, demonstrating how the system of deterministic ordinary differential equations (1.11-1.12) can be used to numerically approximate solutions of the diffusive partial differential equation (WPME). Our numerical examples explore long time behavior of solutions, the effect of the confining potential V_k on the dynamics, the decay of the KL divergence along particle method solutions, and the rate of convergence as $N \to +\infty$, $\varepsilon \to 0$, for fixed k >> 1, both to

solutions of (WPME) at intermediate times and to the target $\bar{\rho}$ on Ω in the long time limit. Our simulations are conducted in Python using the NumPy, SciPy, CuPy, and Matplotlib libraries [45, 63, 73, 76].

7.1. Details of numerical approach. We now describe the details of our numerical approach. Since the main goal of our study is to illustrate how nonlocal particle dynamics can approximate local diffusion equations, we consider the external potential V = 0. We take the dimension d = 1, a Gaussian mollifier,

(7.1)
$$\zeta_{\varepsilon}(x) = e^{-x^2/2\varepsilon^2} / (\sqrt{2\pi\varepsilon^2}).$$

and choose the underlying domain as $\Omega = (-1, 1)$. We approximate no flux boundary conditions on Ω via the confining potential,

(7.2)
$$V_k(x) = \begin{cases} \frac{k}{2}(x+1)^2 & \text{if } x < -1, \\ \frac{k}{2}(x-1)^2 & \text{if } x > 1, \\ 0 & \text{otherwise,} \end{cases}$$

where $k \in \mathbb{N}$ controls the strength of the confinement.

Unless otherwise specified, we choose,

(7.3)
$$\varepsilon = 4/N^{0.99}.$$

Note that this relationship between ε and N is better than expected from our rigorous results; see the discussion after Corollary 1.5. As will be seen from our choice of initial conditions $\{X_0^i\}_{i=1}^N$ below, the choice of ε in (7.3) ensures that the mollifiers have sufficient overlap and that different particles can "sense" each other through the function $f(X^i, X^j)$.

Similarly, unless otherwise specified, we choose

(7.4)
$$k = 10^9$$
.

Our choice of k, corresponding to strong confinement, is motivated by the desire to more closely approximate the dynamics of (WPME) on the bounded domain Ω with no flux boundary conditions. We anticipate that different choices of dimension, mollifier, underlying domain, and confining potential may affect the rate of convergence of our method, but, as our main convergence theorems are not quantitative, we leave a detailed numerical analysis of the these effects to future work.

The first step in our method is to approximate the initial condition ρ_0 in (WPME) by an empirical measure $\sum_{i=1}^{N} \delta_{X_0^i} m^i$ with locations $\{X_0^i\}_{i=1}^N \subseteq \mathbb{R}^d$ and weights $\{m^i\}_{i=1}^N \subseteq [0, +\infty)$ satisfying $\sum_{i=1}^{N} m^i = 1$. In practice, we do this by dividing the domain $\Omega = (-1, 1)$ into N intervals of equal measure. The location X_0^i is chosen to be the center of the i^{th} interval, and the weight m^i is chosen to approximate the integral of ρ_0 over the interval:

$$m^{i} = h\rho_{0}(X_{0}^{i}) \approx \int_{X_{0}^{i} - h/2}^{X_{0}^{i} + h/2} \rho_{0}(x)dx, \quad \text{for } h = |\Omega|/N$$

See Lemma A.4.

With the initial conditions in hand, the next step is to solve the system of ODEs (1.11-1.12). For general $\bar{\rho}$, this is an integral equation, which would be expensive to compute. In the present work, we consider $\bar{\rho}$ for which the integral in equation (1.12) can be pre-computed exactly, yielding to a closed form, analytic expression for $f(X^i, X^j)$ and reducing (1.11) to a standard system of ODEs. In Appendix C, we provide explicit formulas for $f(X^i, X^j)$ in the case $\bar{\rho}$ is piecewise constant or $\bar{\rho} = C/(1 + |x|^2)$, the latter being a prototypical example of a log-concave target. While it will not be possible to obtain a closed form expression for $f(X^i, X^j)$ for all choices of $\bar{\rho}$, we are optimistic that taking sufficiently accurate piecewise constant approximations would yield good results. We leave a detailed analysis of the convergence of our method under various approximations of the target $\bar{\rho}$ to future work. Once a closed form expression for $f(X^i, X^j)$ is obtained, the system of ODEs (1.11) may then be solved using a standard numerical integrator. In the present work, we use the SciPy implementation of the backward differentiation formula (BDF) with a maximum time step of 10^{-5} .

Finally, we seek to understand qualitative properties of the *particle solution*, that is, the evolving empirical measure,

(7.5)
$$\rho_{\varepsilon,k}^{N}(t) = \sum_{i=1}^{N} \delta_{X^{i}(t)} m^{i},$$

as well as its relation to the solution $\rho(t)$ of (WPME) and the target $\bar{\rho}$. To visually depict $\rho_{\varepsilon,k}^N(t)$ and compute its difference from $\rho(t)$ and $\bar{\rho}$ with respect to classical L^p norms and statistical divergences, we will often consider the following kernel density estimate, given by convolving $\rho_{\varepsilon,k}^N(t)$ with the mollifier ζ_{ε} :

(7.6)
$$\tilde{\rho}_{\varepsilon,k}^{N}(x,t) = (\rho_{\varepsilon,k}^{N}(t) * \zeta_{\varepsilon})(x) = \sum_{i=1}^{N} \zeta_{\varepsilon}(X^{i}(t) - x)m^{i}$$

According to Lemma 2.3, if there exists $\mu \in \mathcal{P}(\mathbb{R}^d)$ so that $\rho_{\varepsilon,k}^N$ narrowly converges to μ as $\varepsilon \to 0$, then the kernel density estimator $\tilde{\rho}_{\varepsilon,k}^N$ also narrowly converges to μ as $\varepsilon \to 0$. Thus our main results that guarantee convergence of $\rho_{\varepsilon,k}^N$ also ensure convergence of $\tilde{\rho}_{\varepsilon,k}^N$.

Furthermore, when the target $\bar{\rho}$ is normalized to satisfy $\int_{\Omega} \bar{\rho} = 1$, solutions of (WPME) dissipate the Kullback-Leibler (KL) divergence with respect to $\bar{\rho}$ on Ω exponentially quickly in time (see inequality (1.5)). We will numerically illustrate that this key property is preserved by our approximate solutions $\tilde{\rho}_{\varepsilon,k}^N$. We compute the KL divergence on $\bar{\Omega}$ via,

(7.7)
$$\operatorname{KL}\left(\frac{\tilde{\rho}_{\varepsilon,k}^{N}(t)}{C_{\varepsilon,k,N}(t)},\bar{\rho}\right) = \int_{\Omega}\left(\frac{\tilde{\rho}_{\varepsilon,k}^{N}(x,t)}{C_{\varepsilon,k,N}(t)}\right)\log\left(\frac{\tilde{\rho}_{\varepsilon,k}^{N}(x,t)}{C_{\varepsilon,k,N}(t)\bar{\rho}(x)}\right)dx, \text{ for } C_{\varepsilon,k,N}(t) = \int_{\Omega}\tilde{\rho}_{\varepsilon,k}^{N}(x,t)dx,$$

where the constant $C_{\varepsilon,k,N}(t)$ allows us to compensate for the fact that, since $\tilde{\rho}_{\varepsilon,k}^N$ is not in general supported on $\overline{\Omega}$, the restriction of $\tilde{\rho}_{\varepsilon,k}^N$ to $\overline{\Omega}$ is not a probability measure and $\operatorname{KL}(\tilde{\rho}_{\varepsilon,k}^N(t), \bar{\rho})$ can be negative. On the other hand, $\tilde{\rho}_{\varepsilon,k}^N/C_{\varepsilon,k,N}$ is always a probability measure on $\overline{\Omega}$, so that equation (7.7) gives a well-defined, nonnegative statistical divergence. We compute the integrals in (7.7) numerically, using the SciPy library's quad function.

A final key quantity of our numerical scheme is the value of the energy $\mathcal{F}_{\varepsilon,k}$ along the solution of the gradient flow $\rho_{\varepsilon,k}^N$. At the continuous time level, the gradient flow structure ensures that $\mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}^N(t))$ is always decreasing in time; see Theorem 2.13 and Proposition 3.13. To investigate the rate of decrease numerically, we first obtain the following expression for $\mathcal{F}_{\varepsilon}$ in this setting:

$$(7.8) \qquad \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}^{N}(t)) = \mathcal{E}_{\varepsilon}(\rho_{\varepsilon,k}^{N}(t)) + \mathcal{V}_{k}(\rho_{\varepsilon,k}^{N}(t)) = \frac{1}{2} \int_{\mathbb{R}^{d}} \frac{|\zeta_{\varepsilon} * \rho_{\varepsilon,k}^{N}(t)|^{2}}{\bar{\rho}} d\mathcal{L}^{d} + \int_{\mathbb{R}^{d}} V_{k} d\rho_{\varepsilon,k}^{N}(t)
= \frac{1}{2} \int_{\mathbb{R}^{d}} \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \rho_{\varepsilon,k}^{N}}{\bar{\rho}}\right) (t) d\rho_{\varepsilon,k}^{N}(t) + \int_{\mathbb{R}^{d}} V_{k} d\rho_{\varepsilon,k}^{N}(t)
= \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} g(X^{i}(t), X^{j}(t)) m^{i} m^{j} + \sum_{i=1}^{N} m^{i} V_{k}(X^{i}(t)), \text{ where}
(7.9) \qquad g(x,y) := \int_{\mathbb{R}} \frac{\zeta_{\varepsilon}(x-z)\zeta_{\varepsilon}(y-z)}{\bar{\rho}(z)} dz.$$

We note that g(x, y) is related to the function f(x, y) defined in equation (1.12) by $f = \nabla_x g$, and the integral in the definition of g may be likewise computed explicitly for our choices of $\bar{\rho}$, as we describe in Appendix C.

We close this discussion of the details of our numerical method with a few remarks on its efficient implementation in Python. As an interacting particle system, computing the evolution of the particle trajectories (1.11-1.12) is inherently an $O(N^2)$ computation for a strictly positive mollifier ζ_{ε} . The expectation is that the computational effort would decrease for a compactly supported mollifier: indeed, if supp $\zeta_{\varepsilon} \subset B_{R\varepsilon}(0)$, then $f(X^i, X^j)$ would vanish for $|X^i - X^j| > 2R\varepsilon$. However, rigorously proving that the computational effort indeed decreases to O(mN), where *m* represented the average number of particles lying within the radius of a given mollifier, would require careful estimates on the repulsive forces between the particles and is left for future work. Nevertheless, even for a strictly positive mollifier, we are able to achieve good computational speed in practice by using the following techniques. First, we provide an analytical Jacobian to the ODE solver rather than relying on finite difference approximations. Second, we leverage the structure of the integrand to compute these partial derivatives efficiently. Finally, we parallelize these computations using the CuPy library for GPU-accelerated computing [63]. These elements of our implementation allow us to speed up our calculations by two orders of magnitude compared to previous work by the first author [20]. Namely, we performed the same simulations as those used to generate Figure 1 of [20] (the evolution of density over time, starting from Barenblatt initial data), both using the code of [20], as well as with our implementation. In Figure 7.1, we record the resulting improvement in terms of computational time.

Time	Carrillo, et. al. [20]	Present Work
N = 100	0.04s	0.05s
N = 200	0.41s	0.08s
N = 400	3.35s	0.14s
N = 800	38.96s	0.35s
N = 1600	461.96s	5.73s

FIGURE 1. Computational time for simulation of $\rho_{\varepsilon,k}^N(t)$ using our numerical method and implementation (right column) and that of [20] (left column). Here the target is $\bar{\rho}_{uni}$, we take k = 0, t = 0.15, and the initial condition is the Barenblatt profile (7.13.

These simulations were performed on a standard desktop PC (Intel Core i7-10700 CPU @ 2.9 GHz, 16 GB RAM) with a consumer-level GPU (NVIDIA GeForce RTX 2060 Super). This improvement demonstrates how recent advances in open source scientific computing methods, even in high level languages like Python, are making computing interacting particle systems tractable, even for large numbers of particles.

7.2. Simulation Results. We now turn to several numerical examples that illustrate key properties of our method. In the following simulations, we consider three main choices of target: uniform, log-concave, and piecewise-constant, given by,

(7.10)
$$\bar{\rho}_{\text{uni}}(x) = \frac{1}{2},$$

(7.11)
$$\bar{\rho}_{\log\text{-con}}(x) = \frac{2}{\pi(1+|x|^2)},$$

(7.12)
$$\bar{\rho}_{\text{pw-const}}(x) = \begin{cases} 1/3 & \text{for } x \in (-\infty, -0.75) \cup [-0.25, 0.25) \cup [0.75, +\infty), \\ 2/3 & \text{for } x \in [-0.75, -0.25) \cup [0.25, 0.75). \end{cases}$$

7.2.1. Evolution of density and particle trajectories. In Figure 2, we illustrate qualitative properties of numerical solutions by plotting the kernel density estimate $\tilde{\rho}_{\varepsilon,k}^N(x,t)$, defined in equation (7.6), in the top row and the trajectories of the particles $X^i(t)$ in the bottom row. We conduct our simulation for N = 101 particles, of which 20 are plotted in the bottom row. We consider three choices of target: $\bar{\rho}_{uni}$ (left), $\bar{\rho}_{log-con}$ (middle), and $\bar{\rho}_{pw-const}$ (right). In all cases, our initial condition is given by a *Barenblatt profile* $\psi_{\tau}(x)$, with $\tau = 0.0625$:

(7.13)
$$\psi_{\tau}(x) = \frac{\tau^{-1/3}}{12} \left(3^{4/3} - \frac{|x|^2}{\tau^{2/3}} \right)_+.$$

In the top row of Figure 2, we observe that, for all choices of target $\bar{\rho}$, the kernel density estimate of the solution $\tilde{\rho}_{\varepsilon,k}^{N}(x,t)$ flows toward $\bar{\rho}$ on Ω . For $\bar{\rho}_{uni}$ and $\bar{\rho}_{\log-con}$, this provides numerical verification of Corollary 1.5, since these targets $\bar{\rho}$ are log-concave. On the other hand, while $\bar{\rho}_{pw-const}$ is not log-concave, and thus falls outside the scope of our theoretical results, it does satisfy a Poincaré inequality, so previous work on asymptotic behavior on smooth [26] and weak [35,44] solutions of (WPME) ensure that exact solutions of the continuum PDE converge to the target $\bar{\rho}_{pw-const}$ exponentially quickly in time; see, for example, inequality (1.5). Consequently, although this case lies outside the realm of our rigorous results, it is not surprising that we observe convergence of $\rho_{\varepsilon,k}^{N}$ to $\bar{\rho}_{pw-const}$ in the long-time limit numerically.

In the bottom row of Figure 2, we observe that the particles evolve relatively quickly to their steady state, with most stopping by time t = 0.3. This stands in stark contrast to classical stochastic approaches for sampling, such as Langevin dynamics [10], and stochastic methods in the control theory literature [37,61], in which particles remain in perpetual motion, complicating the choice of an appropriate *stopping time*, beyond which continued evolution doesn't lead to improved accuracy.



FIGURE 2. Simulation of the evolution of the density $\tilde{\rho}_{\varepsilon,k}^{N}(t)$ for the three targets defined in (7.10)-(7.12), with N = 101, $k = 10^9$, and initial data the Barenblatt profile (7.13). Top Row: Snapshots of $\tilde{\rho}_{\varepsilon,k}^{N}(t)$ for the indicated times t. Bottom Row: Evolution of corresponding particle trajectories $X^{i}(t)$.

7.2.2. Effect of confining potential on evolution of density. In Figure 3, we consider the effect of the confining potential on the dynamics. For a fixed number of particles N = 200 and initial conditions given by $\bar{\rho}_{pw-const}$, we plot the evolution of the kernel density estimate $\tilde{\rho}_{\varepsilon,k}^{N}(x,t)$ as the strength of the confining potential V_k is increased, from k = 0 (left, no confinement) to k = 100 (middle, moderate confinement) and $k = 10^9$ (right, strong confinement). All simulations are conducted with Barenblatt initial data, as in equation (7.13).

In the k = 0 plot in Figure 3, we observe that the support of $\tilde{\rho}_{\varepsilon,k}^N(x,t)$ quickly spreads outside the closure of the domain $\overline{\Omega} = [-1,1]$. This is due to fact that k = 0 implies $V_0 = 0$, by equation (7.2), so there is no confining potential, which is equivalent to taking $\Omega = \mathbb{R}^d$. In this case, Theorem 1.1 ensures that, for $\varepsilon > 0$ small and $N \in \mathbb{N}$ large, the particle method approximates solutions of the (WPME) equation on \mathbb{R}^d without boundary. The diffusive effect of this equation causes the particles to spread.

In the k = 100 plot, we observe that even a weak confining potential causes the support of the kernel density estimate to remain mostly inside of $\overline{\Omega}$, with only a small amount of mass leaking out the sides of the domain. And, in the $k = 10^9$ plot, when the confinement effect is very strong, we observe that the support of the kernel density estimate is even closer to $\overline{\Omega}$. In general, we expect the support of the kernel density estimate $\tilde{\rho}_{\varepsilon,k}^N(t)$ to always be slightly larger than the domain, since even when all particles are confined to $\overline{\Omega}$, the kernel density estimate will satisfy,

supp
$$\tilde{\rho}_{\varepsilon,k}^{N}(t) = \{X^{1}(t), \dots, X^{n}(t)\} + \text{supp } \varphi_{\varepsilon}.$$



FIGURE 3. Comparison of how the strength of the confining potential affects the evolution of the density. Here, $\bar{\rho} = \bar{\rho}_{\text{pw-const}}$, N = 200, and the initial data is the Barenblatt profile (7.13). Left: no confinement (k = 0). Middle: medium confinement (k = 100). Right: strong confinement $(k = 10^9)$.

However, in the limit $N \to +\infty$, $\varepsilon \to 0$, and $k \to +\infty$, the support of $\tilde{\rho}_{\varepsilon,k}^N$ will be contained in $\overline{\Omega}$. Finally, note that, by preventing mass from leaking out of the domain, strong confinement gives the best agreement between the long time behavior (t = 1) of the kernel density estimate and the desired target $\bar{\rho}_{pw-const}$ on Ω , in agreement with Corollary 1.5.



FIGURE 4. Evolution of KL divergence between $\tilde{\rho}_{\varepsilon,k}^N(t)$ and $\bar{\rho}$ for three choices of target (7.10)-(7.12) and three choices of N (solid lines). We plot the line of best fit for $t \in [0, 0.25]$ (dashed line). We take $k = 10^9$, t = 2, and initial data is the Barenblatt profile (7.13).

7.2.3. Decay of KL divergence. In Figure 4, we examine the decay of KL divergence between the kernel density estimate $\tilde{\rho}_{\varepsilon,k}^N(t)$ and the target $\bar{\rho}$ on Ω , as computed via equation (7.7). We consider three choices of target, $\bar{\rho}_{uni}$ (left), $\bar{\rho}_{log-con}$ (middle), and $\bar{\rho}_{pw-const}$ (right), and varying numbers of particles N. All simulations are conducted with Barenblatt initial data. Since each of the three targets $\bar{\rho}$ satisfies a Poincaré inequality, the inequality (1.5) implies that the KL divergence between $\bar{\rho}$ and smooth solutions $\rho(t)$ of the (WPME) equation decays exponentially quickly in time. We seek to observe to what extent this property is preserved by the numerical solution $\tilde{\rho}_{\varepsilon,k}^N(t)$, which approximates $\rho(t)$ in the limit $N \to +\infty$, $\varepsilon \to 0$, and $k \to +\infty$, as in Theorem 1.4.

For all three choices of target, we indeed observe an initial regime in which the KL divergence decays exponentially, as indicated by linear decay on the semilog plots in Figure 4. We estimate the rate of decay by plotting the line of best fit on the time interval $t \in [0, 0.25]$, as shown by the dashed line. After the initial period of exponential decay, the KL divergence often appears to level off, particularly for smaller numbers of particles. For larger numbers of particles, the period of exponential decay lasts longer. This indicates that, for smaller numbers of particles, the discretization error in the approximation of (WPME) becomes dominant sooner, slowing the decay of the KL divergence.

The fact that our numerical approximation $\tilde{\rho}_{\varepsilon,k}^{N}(t)$ preserves, up to discretization error, the key property of exponential decay of the KL divergence testifies to the benefit of structure-preserving numerical schemes—in our case, designing a numerical scheme that preserves the continuum PDE's gradient flow structure also succeeds in capturing asymptotic behavior at the level of the particle method.

7.2.4. Decay of energy. In Figure 5, we examine the decay of the energy $\mathcal{F}_{\varepsilon,k}$ along the particle method solution $\rho_{\varepsilon,k}^N(t)$, as computed via equations (7.8-7.9). We consider three choices of target, $\bar{\rho}_{uni}$ (left), $\bar{\rho}_{log-con}$ (middle), and $\bar{\rho}_{pw-const}$ (right), and varying numbers of particles N. All simulations are conducted with Barenblatt initial data.

In all three cases, we observe that the energy decreases along the flow. This is expected since (up to the time discretization error of the ODE solver) our particle method solution $\rho_{\varepsilon,k}^N(t)$ is exactly a gradient flow of the energy $\mathcal{F}_{\varepsilon,k}$. For both of the log-concave energies, $\bar{\rho}_{uni}$ and $\bar{\rho}_{\log-con}$, we observe an initial period of exponential decay, for $t \in [0, 0.5]$, which we approximate by a line of best fit, shown by the dashed line. We do not observe a corresponding period of exponential decay for the non-log-concave energy $\bar{\rho}_{pw-const}$.



FIGURE 5. Evolution of $\mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}^N)$ for three choices of target (7.10)-(7.12) and three choices of N (solid lines). We include the line of best fit for $t \in [0, 0.5]$ (dashed line) on the left and middle plots. We take $k = 10^9$, and the initial data is the Barenblatt profile (7.13).

7.2.5. Convergence to weighted porous medium equation. In Figures 6 and 7, we examine the rate of convergence of the kernel density estimate $\tilde{\rho}_{\varepsilon,k}^N(t)$ as $k \to +\infty$, $\varepsilon \to 0$, and $N \to +\infty$. Given that, for general $\bar{\rho}$, we lack an analytic expression for the solution $\rho(t)$ of (WPME) to which we expect the solutions to converge, we instead compare our numerical solution with N particles at time t = 0.1 to the numerical solution with $N_{\text{max}} = 1,280$ particles at time t = 0.1 via,

(7.14)
$$L^{1} \operatorname{error} = \int_{\Omega} \left| \tilde{\rho}_{\varepsilon(N),k}^{N}(x,t) - \tilde{\rho}_{\varepsilon(N_{\max}),k}^{N_{\max}}(x,t) \right| dx, \qquad t = 0.1,$$

where $\varepsilon(N)$ is as in equation (7.3) and the integral is evaluated using the SciPy library's **quad** function. Furthermore, since we only expect good convergence rates when the solution of the underlying weighted porous medium equation is sufficiently regular, we restrict our attention to the smooth targets $\bar{\rho}_{uni}$ and $\bar{\rho}_{log-con}$.

In Figure 6, we consider how the presence of a confining potential affects the rate of convergence, for both $\bar{\rho}_{\text{uni}}$ and $\bar{\rho}_{\text{log-con}}$. All simulations are conducted with Barenblatt initial data. We choose values of Nfrom N = 20 to N = 640, with logarithmic spacing. In the top row, for no confinement (k = 0), we observe second order convergence. In the middle row, for moderate confinement (k = 100), we observe slightly less than second order convergence. Finally, in the bottom row, for strong confinement $(k = 10^9)$, we observe less than first order convergence.

This example illustrates that there is a delicate balance underlying the choice of the stegnth of the confining potential. On one hand, the confinement must be selected to be sufficiently strong to prevent mass from leaking out of the domain and to ensure that the long time limit agrees well with the desired target; see



FIGURE 6. The effect of k on the rate of convergence in N of the L^1 error (7.14) between $\tilde{\rho}_{\varepsilon,k}^N$ and the numerical solution. In the left-hand column the target is $\bar{\rho}_{uni}$ and in the right-hand column the target is $\bar{\rho}_{log-con}$. Here t = 0.1, and the initial condition is the Barenblatt profile (7.13).

Figure 3. On the other hand, selecting the confinement to be too strong can lead in a degradation of the rate of convergence as $\varepsilon \to 0$, $N \to +\infty$, as more particles would be required for a given degree of accuracy.

In Figure 7, we consider the role the initial conditions play in determining the rate of convergence of the method. In particular, unlike the previous simulation, which was conducted with Barenblatt initial





FIGURE 7. The effect of the initial condition on the rate of convergence in N of the L^1 error (7.14) between $\tilde{\rho}_{\varepsilon,k}^N$ and the numerical solution for two choices of target $\bar{\rho}$. Here $k = 10^9$, t = 0.1, and we take the uniform initial condition (7.15).

conditions, we now consider uniform initial conditions,

(7.15)
$$\mu_0(x) = \begin{cases} \frac{1}{2} & \text{if } x \in [-1,1], \\ 0 & \text{otherwise.} \end{cases}$$

We consider the case of no confinement, k = 0, since the previous figure showed the fastest rate of convergence, of approximately second order, in this case; see Figure 6, top row. We compute the L^1 error as in equation (7.14) with $N_{\text{max}} = 1,280$ and N from N = 20 to N = 640 logarithmically spaced.

Unlike in the previous case, in which we observed near second order convergence in the absence of confinement, in this case we observe closer to first order convergence for both $\bar{\rho}_{uni}$ (left) and $\bar{\rho}_{log-con}$ (right). We believe this is due to the fact that the continuum solution $\rho(t)$ of (WPME) with uniform initial conditions, as above, has worse regularity than the solution for Barenblatt initial conditions. In previous work by the first author and Bertozzi [30] on a regularized particle method for the related aggregation equation, which also has a gradient flow structure in the Wasserstein metric, it was shown that the rate of convergence of the particle method depended strongly on the regularity of the solution of the underlying PDE, in the sense that lower regularity of the continuum solution led to a slower rate of convergence of the numerical solution. While the convergence results in the present paper are purely qualitative, it appears that there may a similar dependence on regularity for the rate of convergence of our particle method to (WPME).

7.2.6. Convergence to Steady State. In Figure 8, we conclude our study of properties of the numerical method by examining the rate of convergence of the kernel density estimate $\tilde{\rho}_{\varepsilon,k}^{N}(t)$ to the target $\bar{\rho}$ in the long time limit, as the number of particles N increases. As we only expect good convergence rates when the target is sufficiently regular, we restrict our attention to the smooth targets $\bar{\rho}_{uni}$ and $\bar{\rho}_{log-con}$. Furthermore, as illustrated in Figure 3, since strong confinement is necessary to obtain convergence to the target as $t \to +\infty$, we choose $k = 10^9$. We consider Barenblatt initial conditions and values of N from N = 20 to N = 720, logarithmically spaced. We compute the L^1 error via,

(7.16)
$$L^{1} \operatorname{error} = \int_{\Omega} \left| \tilde{\rho}_{\varepsilon,k}^{N}(x,T) - \bar{\rho}(x) \right| dx, \qquad T = 2.0,$$

where the integral is evaluated using the SciPy library's quad function.

For both $\bar{\rho}_{uni}$ and $\bar{\rho}_{log-con}$ we observe nearly first order convergence of our particle approximation to the target $\bar{\rho}$. This provides a quantitative numerical result to complement our qualitative result from Corollary 1.5, in which we show that there exist parameters $T \to +\infty$, $k \to +\infty$, $\varepsilon \to 0$, $N \to +\infty$ for which our particle method indeed provides a way to approximate $\bar{\rho}$ on Ω , as relevant for applications in quantization.



FIGURE 8. The rate of convergence in N of the L^1 error (7.16) between $\tilde{\rho}_{\varepsilon,k}^N$ and the target $\bar{\rho}$ for two choices of target. Here t = 2, $k = 10^9$, and we take the uniform initial condition (7.15).

APPENDIX A. WASSERSTEIN GRADIENT FLOWS

We begin with a proof of Proposition 2.11, relating the metric slope and subdifferential.

Proof of Proposition 2.11. By definition of the subdifferential and local slope, for all $\gamma \in \Gamma_0(\mu, \nu)$,

$$\begin{split} |\partial \mathcal{G}|(\mu) &= \limsup_{\nu \to \mu} \frac{(\mathcal{G}(\mu) - \mathcal{G}(\nu))_{+}}{W_{2}(\mu, \nu)} \leq \limsup_{\nu \to \mu} \frac{1}{W_{2}(\mu, \nu)} \left(\int_{\mathbb{R}^{d} \times \mathbb{R}^{d}} \langle \boldsymbol{\xi}(x), x - y \rangle \, d\boldsymbol{\gamma}(x, y) - \frac{\lambda}{2} W_{2}^{2}(\mu, \nu) \right)_{+} \\ &\leq \limsup_{\nu \to \mu} \left(\frac{\|\boldsymbol{\xi}\|_{L^{2}(\mu)} W_{2}(\mu, \nu)}{W_{2}(\mu, \nu)} + \frac{\lambda_{-}}{2} W_{2}(\mu, \nu) \right) = \|\boldsymbol{\xi}\|_{L^{2}(\mu)}, \end{split}$$

where $\lambda_{-} = \max\{-\lambda, 0\}$. This shows inequality (2.10). Uniqueness of the minimal subdifferential follows from the strict convexity of $\|\cdot\|_{L^{2}(\mu)}$.

We now describe the proof of Theorem 2.13, which is a collection of results due to Ambrosio, Gigli, and Savaré that ensure well-posedness of Wasserstein gradient flows, as well as their characterization via curves of maximal slope.

Proof of Theorem 2.13. Existence and uniqueness of the gradient flow, as well as the fact that the gradient flow is a curve of maximal slope, follows from [3, Theorem 11.2.1]. Conversely, if $\mu(t)$ is a curve of maximal slope, then [3, Theorem 11.1.3] ensures it is a gradient flow of \mathcal{G} . (This theorem applies since functionals that are λ -convex are *regular*, in the sense required by the theorem, and functionals that are λ -convex along generalized geodesics satisfy the required coercivity assumption in [3, equation 11.1.13b]: see [3, Lemma 10.3.8, Definition 10.3.9] for regularity and [3, Assumption 4.0.1, Lemma 4.1.1] for coercivity. Furthermore, the λ -convexity and lower semicontinuity of \mathcal{G} ensure that its local slope is a strong upper gradient [3, Corollary 2.4.10].) Finally, the fact that $\mu(t)$ is a gradient flow of \mathcal{G} if and only if it satisfies the Evolution Variational Inequality follows from [3, Theorem 11.1.4].

Next, we define a discrete time approximation of a Wasserstein gradient flow, known as a *minimizing* movement scheme, which was famously introduced in the Wasserstein context by Jordan, Kinderlehrer, and Otto [49].

Definition A.1 (minimizing movement scheme). Suppose \mathcal{G} is proper, lower semicontinuous, and λ -convex along generalized geodesics. Define the *proximal operator* J_{τ} by,

$$\mathbf{J}_{\tau}\mu = \operatorname*{argmin}_{\nu \in \mathcal{P}_{2}(\mathbb{R}^{d})} \left\{ \frac{1}{2\tau} W_{2}^{2}(\mu,\nu) + \mathcal{G}(\nu) \right\},$$

and define the minimizing movement scheme $J^n_{\tau}\mu$ by,

$${}^{n}_{\tau}(\mu) = \underbrace{\mathbf{J}_{\tau} \circ \mathbf{J}_{\tau} \circ \cdots \circ \mathbf{J}_{\tau}}_{n \text{ times}}(\mu).$$

Note that, by definition, the energy decreases along the minimizing movement scheme:

(A.1)
$$\mathcal{G}(\mathbf{J}_{\tau}^{n}\mu) \leq \mathcal{G}(\mathbf{J}_{\tau}^{n-1}\mu).$$

We recall the following theorem on the convergence, due to Ambrosio, Gigli, and Savaré.

Theorem A.2 (convergence of minimizing movement scheme, [3, Theorem 4.0.9]). Suppose \mathcal{G} is proper, lower semicontinuous, and λ -convex along generalized geodesics and $\mu \in D(\mathcal{G})$. Fix T > 0, and take a piecewise constant interpolation of the minimizing movement scheme,

(A.2)
$$\mu_{\tau}(s) = \mathbf{J}_{\tau}^{n} \mu \quad \text{for } s \in ((n-1)\tau, n\tau]$$

Then for all $t \in [0,T]$, we have $\lim_{n \to +\infty} \mu_{\tau}(t) = \mu(t)$ narrowly, where $\mu(t)$ is the gradient flow of \mathcal{G} with initial data μ .

Proof. This theorem is an immediate consequence of [3, Theorem 4.0.9].

We continue with an elementary result bounding the Wasserstein distance between a curve of maximal slope and a fixed reference measure.

Proposition A.3 $(M_2 \text{ bound for curves of maximal slope)}$. Suppose $\rho(t) \in AC^2_{loc}((0, +\infty); \mathcal{P}_2(\mathbb{R}^d))$ is a curve of maximal slope of a nonnegative functional $\mathcal{G} : \mathcal{P}_2(\mathbb{R}^d) \to [0, +\infty]$, that is, $\rho(t)$ satisfies Theorem 2.13(i). Suppose further that $\rho_0 \in D(\mathcal{G})$. Then we have,

(A.3)
$$M_2(\rho(t)) \le (1 + te^t) (M_2(\rho(0)) + \mathcal{G}(\rho(0))).$$

Proof. We shall first establish, for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, the estimate,

(A.4)
$$W_2^2(\rho(t),\mu) \le \left(1 + te^t\right) \left[W_2^2(\rho(0),\mu) + \mathcal{G}(\rho(0))\right].$$

To this end, fix some $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, and define $\mathcal{W}(\rho) = -\frac{1}{2}W_2^2(\rho,\mu)$. Since \mathcal{W} is (-1)-convex and lower semicontinuous [3, Proposition 9.3.12], the local slope $|\partial \mathcal{W}|(\mu)$ is a strong upper gradient for \mathcal{W} (see [3, Definition 1.2.1, Corollary 2.4.10]), which implies,

(A.5)
$$|\mathcal{W}(\rho(t)) - \mathcal{W}(\rho(0))| \le \int_0^t |\partial \mathcal{W}|(\rho(s))|\rho'|(s)ds.$$

Furthermore, using the definition of local slope, rearranging, and applying the triangle inequality, yields,

$$\begin{aligned} |\partial \mathcal{W}|(\rho) &= \limsup_{\nu \to \rho} \frac{W_2^2(\nu, \mu) - W_2^2(\rho, \mu)}{2W_2(\rho, \nu)} = \limsup_{\nu \to \rho} \frac{(W_2(\nu, \mu) - W_2(\rho, \mu))(W_2(\nu, \mu) + W_2(\rho, \mu))}{2W_2(\rho, \nu)} \\ \leq \limsup_{\nu \to \rho} \frac{W_2(\rho, \nu)(W_2(\nu, \mu) + W_2(\rho, \mu))}{2W_2(\rho, \nu)} = W_2(\rho, \mu). \end{aligned}$$

Thus, combining (A.5) and (A.6), we obtain,

$$\frac{1}{2} \left[W_2^2(\rho(t),\mu) - W_2^2(\rho(0),\mu) \right] \le |\mathcal{W}(\rho(t)) - \mathcal{W}(\rho(0))| \le \int_0^t W_2(\rho(s),\mu) |\rho'|(s) ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} |||\rho'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2^2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |\rho'|(s)^2 ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} |||\rho'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2^2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |\rho'|(s)^2 ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} |||\rho'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |\rho'|(s)^2 ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} |||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |\rho'|(s)^2 ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |\rho'|(s)^2 ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |\rho'|(s)^2 ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |\rho'|(s)^2 ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |\rho'|(s)^2 ds \\
\le ||W_2(\rho(s),\mu)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |P'|(s)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |P'|(s)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |P'|(s)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |P'|(s)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{1}{2} \int_0^t |P'|(s)||_{L^2([0,t])} ||P'|(s)||_{L^2([0,t])} \le \frac{1}{2} \int_0^t W_2(\rho(s),\mu) ds + \frac{$$

Since $\rho(s)$ is a curve of maximal slope of \mathcal{G} and \mathcal{G} is nonnegative,

$$\frac{1}{2}\int_0^t |\rho'|(s)^2 ds \le \mathcal{G}(\rho(0)).$$

Combining the previous two inequalities gives,

$$\frac{1}{2} \left[W_2^2(\rho(t),\mu) - W_2^2(\rho(0),\mu) \right] \le \frac{1}{2} \int_0^t W_2^2(\rho(s),\mu) ds + \mathcal{G}(\rho(0)).$$

By Gronwall's inequality, this implies,

$$W_2^2(\rho(t),\mu) \le \left[W_2^2(\rho(0),\mu) + \mathcal{G}(\rho(0))\right] (1 + te^t).$$

This shows inequality (A.4).

To obtain (A.3), it suffices to recall the definition of the Wasserstein metric in equation (2.3) in terms of transport plans γ . If γ is a transport plan from a measure μ to a Dirac mass δ_0 , then for all measurable sets A and B,

$$\gamma(A \times \mathbb{R}^d) = \mu(A) \text{ and } \gamma(\mathbb{R}^d \times B) = \delta_0(B),$$

so $\gamma = (\mathbf{id} \times 0)_{\#} \mu$, where 0 denotes the function $0: x \mapsto 0$. Therefore, applying equation (2.3), we obtain,

$$W_2^2(\mu,\nu) = \int_{\mathbb{R}^d \times \mathbb{R}^d} |x-y|^2 d\gamma(x,y) = \int_{\mathbb{R}^d \times \mathbb{R}^d} |x-0|^2 d\mu(x) = M_2(\mu).$$

Combining this with inequality (A.4) gives inequality (A.3).

We close this section by providing the construction of an empirical measure approximating any measure $\mu \in \mathcal{P}_2(\mathbb{R}^d)$.

Lemma A.4 (approximation via empirical measures). For all $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\delta > 0$, there exists $N \in \mathbb{N}$, $\{X^i\}_{i=1,\dots,N} \subseteq \mathbb{R}^d$, and $\{m^i\}_{i=1,\dots,N} \subseteq \mathbb{R}^+$ with $\sum_{i=1}^N m^i = 1$, such that $\mu^N = \sum_{i=i}^N \delta_{X^i} m^i$ satisfies $W_2(\mu, \mu^N) \leq \delta$.

Proof of Lemma A.4. Throughout this proof, we shall use $Q_r(0)$ to denote a cube in \mathbb{R}^d centered at 0 and with side length r > 0; namely, $Q_r(0) = \left[-\frac{r}{2}, \frac{r}{2}\right]^d$. For $x \in \mathbb{R}^d$, let $Q_r(x) = Q_r(0) + x$. Fix $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\delta > 0$. First, we reduce to the case of approximating a compactly supported measure.

Fix $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\delta > 0$. First, we reduce to the case of approximating a compactly supported measure. To this end, note that since $\mu \in M_2(\mathbb{R}^d)$, there exists R > 0 so that $\int_{Q_R(0)^c} |x|^2 d\mu \leq \left(\frac{\delta}{2}\right)^2$. Consider the transport map,

$$\mathbf{t}_R(x) = \begin{cases} x & \text{if } x \in Q_R(0), \\ 0 & \text{otherwise,} \end{cases}$$

and define μ_R via $\mu_R = (\mathbf{t}_R)_{\#}\mu$. Then we find,

(A.7)
$$W_2(\mu,\mu_R) \le \left(\int |\mathbf{t}_R(x) - x|^2 \, d\mu\right)^{1/2} \le \left(\int_{Q_R(0)^c} |x|^2 \, d\mu\right)^{1/2} \le \frac{\delta}{2},$$

We are now ready to define the approximating measure μ^N . Choose $K \in \mathbb{N}$ large enough so that,

(A.8)
$$K \ge \frac{2\sqrt{dR}}{\delta},$$

and consider a grid on $Q_R(0)$ where each cell has side length R/K, so that we have $Q_R(0) = \bigcup_{i=1}^{K^d} Q_{R/K}(X^i)$, where the centers $\{X^i\}_{i=1}^{m^d}$ are chosen such that the above union is disjoint. Let $N = K^d$, and define μ^N to be the sum of Dirac masses at the centers of the cells, with weights given by the mass of μ_R in each cell:

$$\mu^N = \sum_{i=1}^N \delta_{X^i} m^i, \quad \text{with } m^i = \mu_R \left(Q_{R/K}(X^i) \right).$$

To estimate $W_2(\mu_R, \mu^N)$, we consider the transport map $\mathbf{t} : \mathbb{R}^d \to \mathbb{R}^d$ which, for i = 1, ..., N, moves all the mass in cell $Q_{R/K}(X^i)$ to X^i . Then $\mu^N = \mathbf{t}_{\#}\mu_R$ and,

$$W_2^2(\mu_R, \mu^N) \le \int |\mathbf{t}(x) - x|^2 \, d\mu_R = \sum_{i=1}^N \int_{Q_{R/K}(X^i)} |\mathbf{t}(x) - x|^2 \, d\mu_R \le \sum_{i=1}^N \int_{Q_{R/K}(X^i)} \left(\frac{\sqrt{dR}}{K}\right)^2 \, d\mu_R = \left(\frac{\sqrt{dR}}{K}\right)^2$$

where the second inequality follows from the fact that mass in the *i*th cell stays in the *i*th cell, so the largest distance mass could be moved is the diagonal length of the cell, $\frac{\sqrt{dR}}{K}$. Finally, we conclude by using the definition of K in (A.8), together with the estimate (A.7), and the triangle inequality:

$$W_2(\mu, \mu^N) \le W_2(\mu, \mu_R) + W_2(\mu_R, \mu^N) \le \frac{\delta}{2} + \frac{\delta}{2}.$$

Appendix B. Further properties of energies and gradient flows with regularization and confinement

We provide the proof of Lemma 3.1, which ensures that the energies \mathcal{E} and $\mathcal{E}_{\varepsilon}$ are lower semicontinuous with respect to narrow convergence.

Proof of Lemma 3.1. First we consider \mathcal{E} . For this energy, lower semicontinuity follows from the following result of Buttazo [16, Corollary 3.4.2]: given $g : \mathbb{R}^d \times \mathbb{R} \to [0, +\infty]$, consider the functional $\mathcal{G} : \mathcal{P}(\mathbb{R}^d) \to [0, +\infty]$ defined by,

(B.1)
$$\mathcal{G}(\mu) = \begin{cases} \int_{\mathbb{R}^d} g(x, \mu(x)) dx & \text{if } \mu \ll \mathcal{L}^d, \\ +\infty & \text{otherwise.} \end{cases}$$

Then if (i) g is lower semicontinuous, (ii) for every $x \in \mathbb{R}^d$, the function $g(x, \cdot)$ is convex on \mathbb{R} , and (iii) there exists $\theta : \mathbb{R} \to \mathbb{R}$ with $\lim_{t\to\infty} \frac{\theta(t)}{t} = \infty$ and $g(x, y) \ge \theta(|y|)$ for every $x \in \mathbb{R}^d, y \in \mathbb{R}$, then the functional \mathcal{G} is lower semicontinuous with respect to narrow convergence.

We now verify these hypotheses: our energy \mathcal{E} is of the form (B.1), for $g(x,y) = \frac{y^2}{2\bar{\rho}(x)}$, which satisfies (i) and (ii). Furthermore, by setting $\theta(t) = Ct^2$, where $C = (\max_{x \in \mathbb{R}^d} 2\bar{\rho}(x))^{-1}$, we see that g satisfies (iii). Thus, \mathcal{E} is lower semicontinuous with respect to narrow convergence.

The lower semicontinuity of $\mathcal{E}_{\varepsilon}$ follows directly from the definition of $\mathcal{E}_{\varepsilon}(\mu) = \mathcal{E}(\zeta_{\varepsilon} * \mu)$, Lemma 2.3, and the lower semicontinuity of \mathcal{E} .

We now prove Proposition 3.3 by applying the general results of Ambrosio, Gigli, and Savaré [3] to immediately characterize the convexity of \mathcal{E} , \mathcal{V} , \mathcal{V}_k , and $\mathcal{V}_{\varepsilon}$.

Proof of Proposition 3.3. Item (i) is an direct consequence of [3, Theorem 9.4.12]. Item (ii) is a consequence of the fact that, for any potential $W : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ that is proper, lower semicontinuous, bounded below, and λ -convex, the corresponding energy $\rho \mapsto \int W\rho$ is λ -convex along generalized geodesics [3, Proposition 9.3.2]. Next, recall that that $V \in C^2(\mathbb{R}^d)$ with Hessian bounded below implies $D^2 V \ge \lambda I_{d\times d}$ for $\lambda = \inf_{\{x,\xi\in\mathbb{R}^d\}}\xi^t D^2 V(x)\xi$, hence we also have $D^2(\zeta_{\varepsilon} * V) \ge \lambda I_{d\times d}$ for all $\varepsilon > 0$. In particular, both V and $(\zeta_{\varepsilon} * V)$ are λ -convex, which implies \mathcal{V} and $\mathcal{V}_{\varepsilon}$ are λ -convex along generalized geodesics. likewise, since V_k is continuous, bounded below, and convex, \mathcal{V}_k is convex along generalized geodesics. Finally, for item ((iii)), we first note that for $\mu \in \mathcal{P}_2(\mathbb{R}^D)$, we may write $\mathcal{V}_{\Omega}(\rho) = \int \mathcal{V}_{\Omega} d\rho$, where V_{Ω} is given by (1.17). Assumption (D) implies V_{Ω} is a lower-semicontinous, nonnegative, convex function. Hence, the result follows once again from [3, Proposition 9.3.2].

Next we prove Proposition 3.6, characterizing the minimal element of the subdifferential of $\mathcal{F}_{\varepsilon,k}$.

Proof of Proposition 3.6. For simplicity of notation, denote,

(B.2)
$$\boldsymbol{\xi} = \nabla \frac{\delta \mathcal{E}_{\varepsilon}}{\delta \mu} + \nabla (\zeta_{\varepsilon} * V) + \nabla V_k.$$

Note that Lemma 3.7 and Remark 2.10 on the additivity of the subdifferential ensure that $\boldsymbol{\xi} \in \partial \mathcal{F}_{\varepsilon,k}(\mu)$. In order to conclude $\boldsymbol{\xi} \in \partial^{\circ} \mathcal{F}_{\varepsilon,k}$, it remains to show that $\|\boldsymbol{\xi}\|_{L^{2}(\mu)} \leq |\partial \mathcal{F}_{\varepsilon,k}|(\mu)$. Proposition 2.11 will then give the result.

Fix $\psi \in C^1(\mathbb{R}^d)$ satisfying $\nabla \psi \in L^2(\mu)$, and define $\mu_{\alpha} = (\mathbf{id} + \alpha \nabla \psi)_{\#} \mu$. By definition of the Wasserstein distance from μ to μ_{α} in terms of minimizing over all transport plans from μ to μ_{α} , equation (2.3), and the fact that $(\mathbf{id} \times (\mathbf{id} + \alpha \nabla \psi))_{\#} \mu$ is such a plan,

$$W_2(\mu_{\alpha},\mu) \le \|(\mathbf{id} + \alpha \nabla \psi) - \mathbf{id}\|_{L^2(\mu)} = \alpha \|\nabla \psi\|_{L^2(\mu)}$$

By definition of the metric slope,

(B.3)
$$|\partial \mathcal{F}_{\varepsilon,k}|(\mu) = \limsup_{\nu \to \mu} \frac{(\mathcal{F}_{\varepsilon,k}(\mu) - \mathcal{F}_{\varepsilon,k}(\nu))_+}{W_2(\mu,\nu)} \ge \limsup_{\alpha \to 0} \frac{(\mathcal{F}_{\varepsilon,k}(\mu) - \mathcal{F}_{\varepsilon,k}(\mu_\alpha))_+}{W_2(\mu,\mu_\alpha)}$$
$$\ge \frac{1}{\|\nabla \psi\|_{L^2(\mu)}} \limsup_{\alpha \to 0} \frac{(\mathcal{F}_{\varepsilon,k}(\mu) - \mathcal{F}_{\varepsilon,k}(\mu_\alpha))_+}{\alpha}.$$

We now apply inequality (B.3) to complete the proof. Recall from the sentence following assumption (C) that $V_k \in L^1(\nu)$ and $\nabla V_k \in L^2(\nu)$ for all $\nu \in \mathcal{P}_2(\mathbb{R}^d)$. Hence, $\mu_{\alpha} \in D(\mathcal{F}_{\varepsilon,k})$ for all $\alpha \geq 0$. Thus, combining inequality (B.3) with Proposition 3.4, which characterizes the directional derivatives of $\mathcal{E}_{\varepsilon}$, $\mathcal{V}_{\varepsilon}$, and \mathcal{V}_k , applied with,

$$\boldsymbol{\gamma} = (\mathbf{id}, \mathbf{id}, \mathbf{id} + \nabla \psi)_{\#} \mu,$$

we obtain,

$$\begin{split} |\partial \mathcal{F}_{\varepsilon,k}|(\mu) \| \nabla \psi \|_{L^{2}(\mu)} &\geq \lim_{\alpha \to 0} \frac{\mathcal{E}_{\varepsilon}(\mu) - \mathcal{E}_{\varepsilon}(\mu_{\alpha})}{\alpha} + \frac{\mathcal{V}_{\varepsilon}(\mu) - \mathcal{V}_{\varepsilon}(\mu_{\alpha})}{\alpha} + \frac{\mathcal{V}_{k}(\mu) - \mathcal{V}_{k}(\mu_{\alpha})}{\alpha} \\ &= -\frac{1}{2} \int \frac{\zeta_{\varepsilon} * \mu(x)}{\bar{\rho}(x)} \int \left\langle \nabla \zeta_{\varepsilon} \left(x - y_{2} \right), y_{3} - y_{2} \right\rangle d\gamma(y_{1}, y_{2}, y_{3}) \\ &+ \int \left\langle \nabla (\zeta_{\varepsilon} * V)(y_{2}) + \nabla V_{k}(y_{2}), y_{3} - y_{2} \right\rangle dx \\ &= -\int \left\langle \frac{1}{2} \left(\nabla \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \mu}{\bar{\rho}} \right) \right) + \nabla V + \nabla V_{k}, \nabla \psi \right\rangle d\mu. \end{split}$$

Since the above inequality holds for any $\psi \in C^1$ with $\nabla \psi \in L^2(\mu)$, taking,

$$\psi = -\frac{1}{2} \left(\zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \mu}{\bar{\rho}} \right) \right) - \left(\zeta_{\varepsilon} * V \right) - V_k, \text{ so that } \nabla \psi = -\boldsymbol{\xi},$$

we obtain $|\partial \mathcal{F}_{\varepsilon,k}|(\mu) \| \nabla \psi \|_{L^2(\mu)} \ge \| \nabla \psi \|_{L^2(\mu)}^2$. Dividing through by $\| \nabla \psi \|_{L^2(\mu)} = \| \boldsymbol{\xi} \|_{L^2(\mu)}$ gives the result.

We now turn to a proof of Proposition 3.10, which characterizes the gradient flows of \mathcal{F}_k and \mathcal{F} in terms of partial differential equations.

Proof of Proposition 3.10. Note that, for either $\mathcal{G} = \mathcal{F}_k$ or $\mathcal{G} = \mathcal{F}$, if μ is a gradient flow of \mathcal{G} , with initial data $\mu_0 \in \overline{D(\mathcal{G})}$, then, according to Theorem 2.13, μ is unique and is a curve of maximal slope for \mathcal{G} . Since $\mathcal{G} \geq -\|V\|_{\infty}$, this implies that for any t > 0,

(B.4)
$$\int_0^t |\partial \mathcal{G}|^2(\mu(r)) \, dr \le \mathcal{G}(\mu_0) + \|V\|_{\infty} < +\infty.$$

This ensures that $|\partial \mathcal{G}|(\mu(t)) < +\infty$ for \mathcal{L}^1 almost every t > 0, and since $D(|\partial \mathcal{G}|) \subseteq D(\mathcal{G})$, we also have,

(B.5)
$$\mathcal{G}(\mu(t)) < +\infty$$
 for a.e. $t > 0$.

First, consider the case $\mathcal{G} = \mathcal{F}_k$. Suppose μ is a gradient flow of \mathcal{F}_k with initial data $\mu_0 \in \overline{D(\mathcal{F}_k)}$. By inequality (B.5), $\mu(t) \ll \mathcal{L}^d$ for almost every $t \ge 0$. Furthermore, Proposition 3.9 implies that, for almost every $t \ge 0$, we have $(\mu(t)/\bar{\rho})^2 \in W^{1,1}_{loc}(\mathbb{R}^d)$ and there exists $\boldsymbol{\xi}(t) \in \partial^\circ \mathcal{F}_k(\mu)$ with,

(B.6)
$$\boldsymbol{\xi}(t)\mu(t) = \frac{\rho}{2}\nabla(\mu(t)^2/\bar{\rho}^2) + \nabla V\mu(t) + \nabla V_k\mu(t) \quad \text{and} \quad |\partial \mathcal{F}_k|(\mu) = \|\boldsymbol{\xi}(t)\|_{L^2(\mu(t))}.$$

By Definition 2.12 of gradient flow, we obtain that μ satisfies the continuity equation (2.11) with $v(t) = -\boldsymbol{\xi}(t)$. Using the expression (B.6) for $\boldsymbol{\xi}$ therefore yields (3.8). Finally, the containment (3.10) follows from inequality (B.4) and equation (B.6).

On the other hand, suppose μ solves (3.8) and satisfies (3.9-3.10). Then, defining $\boldsymbol{\xi}$ on the support of μ via (B.6) implies that the hypotheses of Proposition 3.9 are satisfied, so $\boldsymbol{\xi} \in \partial^{\circ} \mathcal{F}(\mu)$. From this we find that (3.8) is exactly the continuity equation in Definition 2.12 of the gradient flow, with $v(t) = -\boldsymbol{\xi}(t)$ satisfying $\|v(t)\|_{L^2(\mu(t)} \in L^1_{loc}(0, +\infty)$. Thus, we have that $\mu \in AC^2_{loc}(0, +\infty)$; $\mathcal{P}_2(\mathbb{R}^d)$) [3, Theorem 8.3.1], hence $\mu(t)$ is the unique gradient flow of \mathcal{F}_k with initial data μ_0 , completing the proof of the first part of the proposition.

Now, consider the case $\mathcal{G} = \mathcal{F}$. Suppose μ is a gradient flow of \mathcal{F} with initial data $\mu_0 \in \overline{D(\mathcal{F})}$. By inequality (B.5), $\mu(t) \ll \mathcal{L}^d$ and $\mu = 0$ a.e. on $\mathbb{R}^d \setminus \overline{\Omega}$ for almost every $t \ge 0$. Furthermore, Proposition 3.9 implies that, for almost every $t \ge 0$, $(\mu(t)/\bar{\rho})^2 \in W^{1,1}_{loc}(\Omega)$ and that there exists $\boldsymbol{\xi}(t) \in \partial^{\circ} \mathcal{F}(\mu)$ with,

(B.7)
$$\boldsymbol{\xi}(t)\mu(t) = \frac{\rho}{2}\nabla(\mu(t)^2/\bar{\rho}^2) + \nabla V\mu(t) \quad \text{on } \Omega \quad \text{and} \quad |\partial \mathcal{F}|(\mu) = \|\boldsymbol{\xi}(t)\|_{L^2(\mu(t))}$$

By Definition 2.12 of gradient flow, we obtain that μ satisfies the continuity equation (2.11) with $v(t) = -\boldsymbol{\xi}(t)$. Using the expression (B.7) for $\boldsymbol{\xi}$ therefore yields (3.11). Finally, the containment (3.13) follows from inequality (B.4) and equation (B.7).

On the other hand, suppose μ solves (3.11) and satisfies (3.12-3.13). Then, defining $\boldsymbol{\xi}$ on the support of μ via (B.7) implies that the hypotheses of Proposition 3.9 are satisfied, so $\boldsymbol{\xi} \in \partial^{\circ} \mathcal{F}(\mu)$. From this we find that (3.11) is exactly the continuity equation in Definition 2.12 of the gradient flow, with $v(t) = -\boldsymbol{\xi}(t)$ satisfying $\|v(t)\|_{L^2(\mu(t)} \in L^1_{\text{loc}}(0, +\infty)$. Thus, we have that $\mu \in AC^2_{\text{loc}}(0, +\infty)$; $\mathcal{P}_2(\mathbb{R}^d)$) [3, Theorem 8.3.1], hence $\mu(t)$ is the unique gradient flow of \mathcal{F} with initial data μ_0 , completing the proof of the proposition.

The next result is a proof of Proposition 3.12, which characterizes the gradient flow of $\mathcal{F}_{\varepsilon,k}$ in terms of a partial differential equation.

Proof of Proposition 3.12. Suppose that $\mu(t)$ is the gradient flow of $\mathcal{F}_{\varepsilon,k}$. Then the fact that $\mu(t)$ satisfies (3.15) follows directly from Definition 2.12, Proposition 3.8, and Theorem 2.13.

Now suppose that $\mu(t)$ satisfies (3.15). Then, the fact that the velocity field in the continuity equations is uniformly bounded ensures, by [4, Theorem 8.3.1], that $\mu \in AC^2([0,T]; \mathcal{P}_2(\mathbb{R}^d))$. Thus, the fact that μ is the gradient flow of $\mathcal{F}_{\varepsilon,k}$ is again a consequence of Definition 2.12, Proposition 3.8, and Theorem 2.13.

We now consider the proof of Proposition 3.13, which shows that the gradient flow of $\mathcal{F}_{\varepsilon,k}$ beginning at an empirical measure remains an empirical measure for all time and characterizes the ODE governing the evolution of the locations of the Dirac masses.

Proof of Proposition 3.13. First note that, for all $\varepsilon > 0$ fixed, the function of $(X^1, ..., X^N)$ that appears on the right-hand side of (3.16) is Lipschitz continuous, and therefore the ODE system (3.16) is well-posed. Suppose $X^{i}(t)$ solves (3.16). We claim that it suffices to show that $\mu(t) = \sum_{i=1}^{N} \delta_{X^{i}(t)} m^{i}$ solves (3.15). Proposition 3.12 then ensures that $\mu(t)$ is the unique solution of the gradient flow.

The fact that $\lim_{t\to 0^+} \mu(t) = \mu(0)$ in W_2 follows immediately from the definition of $\mu(t)$ and $\mu(0)$. Next, note that,

$$-\int_{\mathbb{R}^d} \nabla \zeta_{\varepsilon} (X^i(t) - z) \frac{1}{\bar{\rho}(z)} \sum_{j=1}^N m^j \zeta(z - X^j(t)) \, dz = -\int_{\mathbb{R}^d} \nabla \zeta_{\varepsilon} (X^i(t) - z) \frac{1}{\bar{\rho}(z)} \zeta(z - y) \, d\mu(y) dz$$
$$= -\left(\nabla \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \mu}{\bar{\rho}}\right)\right) (X^i(t)).$$

Now, fix a test function $f \in C_c^{\infty}(\mathbb{R}^d \times (0, +\infty))$. By the Fundamental Theorem of Calculus and equation (B.8), for each $1 \leq i \leq N$,

$$0 = \int_0^\infty \frac{d}{dt} f(X^i(t), t) dt = \int_0^\infty \left\langle \nabla f(X^i(t), t), \dot{X}^i(t) \right\rangle + \partial_t f(X^i(t), t) dt$$
$$= \int_0^\infty \left\langle \nabla f(X^i(t), t), \left(-\left(\nabla \zeta_\varepsilon * \left(\frac{\zeta_\varepsilon * \mu}{\bar{\rho}}\right)\right) (X^i(t)) - \nabla (\zeta_\varepsilon * V) (X^i(t)) - \nabla V_k(X^i(t))\right) \right\rangle + \partial_t f(X^i(t), t) dt.$$

Multiplying by m^i , summing over *i*, and recalling the definition of μ yields,

$$0 = \int_0^\infty \int_{\mathbb{R}^d} \left\langle \nabla f(x,t), \left(-\left(\nabla \zeta_\varepsilon * \left(\frac{\zeta_\varepsilon * \mu}{\bar{\rho}} \right) \right)(x) - \nabla (\zeta_\varepsilon * V)(x) - \nabla V_k(x) \right) \right\rangle + \partial_t f(x,t) \, d\mu(x,t) \, dt.$$
us, μ is a distributional solution of the continuity equation (3.15).

Thus, μ is a distributional solution of the continuity equation (3.15).

(B.8)

APPENDIX C. EXPLICIT FORMULAS FOR NUMERICAL METHOD

In this section, we collect a few explicit formulas that we use in the implementation of our numerical method. For our choices of uniform (7.10), log-concave (7.11), and piecewise constant (7.12) target, we have explicit formulas for the functions f(x,y) and g(x,y) defined in section 7.1: see equations (1.12) and (7.9). For the log-concave target measure, we obtain,

For the uniform and piecewise constant targets, note that both may be expressed as,

$$\bar{\rho}(x) = \sum_{k=1}^{N} c_k \mathbb{1}_{[b_k, b_{k+1}]}(x),$$

where $\{c_k\}_{k=1}^N$ are positive constants chosen so that $\int_{\Omega} \bar{\rho} = 1$, $\{b_k\}_{k=1}^{N+1} \subseteq \mathbb{R}$. For any target of this form, we obtain

$$\begin{split} f(x_i, x_j) &= \sum_{k=1}^{N} c_k^{-1} \left[\varphi_{ij}(b_{k+1}) - \varphi_{ij}(b_k) \right] \\ \varphi_{ij}(z) &= -\frac{e^{-(x_i^2 + 2z^2 + x_j^2)/(2\varepsilon^2)}}{8\pi\varepsilon^3} \left(2\varepsilon e^{z(x_i + x_j)/\varepsilon^2)} - \sqrt{\pi}(x_i - x_j)e^{((x_i + x_j)^2 + 4z^2)/(4\varepsilon^2)} \operatorname{erf}\left(\frac{x_i - 2z + x_j}{2\varepsilon}\right) \right) \\ &+ \frac{e^{-(x_i^2 + x_j^2)/(2\varepsilon^2)}}{8\pi\varepsilon^3} \left(2\varepsilon - \sqrt{\pi}(x_i - x_j)e^{(x_i + x_j)^2/(4\varepsilon^2)} \operatorname{erf}\left(\frac{x_i + x_j}{2\varepsilon}\right) \right) \\ g(x_i, x_j) &= \sum_{k=1}^{N} [\psi_{ij}(b_{k+1}) - \psi_{ij}(b_k)] \\ \psi_{ij} &= \frac{-1}{4\sqrt{\pi}} e^{\frac{-(x_i - x_j)^2}{4\varepsilon^2}} \operatorname{erf}\left(\frac{x_i + x_j - 2z}{2\varepsilon}\right). \end{split}$$

References

- L. Alasio, M. Bruna, and J. A. Carrillo. The role of a strong confining potential in a nonlinear Fokker-Planck equation. Nonlinear Analysis, 193:111480, 2020.
- [2] L. Ambrosio, E. Brué, and D. Semola. Lectures on optimal transport. Springer, 2021.
- [3] L. Ambrosio, N. Gigli, and G. Savaré. Gradient flows: in metric spaces and in the space of probability measures. Springer Science & Business Media, 2008.
- [4] L. Ambrosio, N. Gigli, and G. Savaré. Gradient flows in metric spaces and in the space of probability measures. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, second edition, 2008.
- [5] B. Anderson, E. Loeser, M. Gee, F. Ren, S. Biswas, O. Turanova, M. Haberland, and A. L. Bertozzi. Quantitative assessment of robotic swarm coverage. Proc. 15th Int. Conf. on Informatics in Control, Automation, and Robotics, 2:91–101, 2018.
- [6] B. Anderson, E. Loeser, M. Gee, F. Ren, S. Biswas, O. Turanova, M. Haberland, and A. L. Bertozzi. Quantifying swarm coverage. In Springer Lecture Notes in Electrical Engineering, volume 613. Springer, 2020.
- [7] J.-D. Benamou, G. Carlier, Q. Mérigot, and E. Oudet. Discretization of functionals involving the monge-ampère operator. Numerische mathematik, 134(3):611-636, 2016.
- [8] M. Bessemoulin-Chatard and F. Filbet. A finite volume scheme for nonlinear degenerate parabolic equations. SIAM Journal on Scientific Computing, 34(5):B559–B583, 2012.
- [9] V. I. Bogachev. Measure theory, volume 1. Springer Science & Business Media, 2007.
- [10] N. Bou-Rabee and A. Eberle. Markov Chain Monte Carlo Methods. lecture notes, https://uni-bonn.sciebo.de/s/ kzTUFff5FrWGAay, 2020.
- [11] D. P. Bourne and R. Cristoferi. Asymptotic optimality of the triangular lattice for a class of optimal location problems. Communications in Mathematical Physics, 387(3):1549–1602, 2021.
- [12] D. P. Bourne and S. M. Roper. Centroidal power diagrams, lloyd's algorithm, and applications to optimal location problems. SIAM Journal on Numerical Analysis, 53(6):2545–2569, 2015.
- [13] F. Bullo, J. Cortés, and S. Martinez. Distributed control of robotic networks. Princeton University Press, 2009.
- [14] M. Burger, J. A. Carrillo, and M.-T. Wolfram. A mixed finite element method for nonlinear diffusion equations. *Kinetic & Related Models*, 3(1):59, 2010.
- [15] M. Burger and A. Esposito. Porous medium equation as limit of nonlocal interaction. arXiv preprint arXiv:2202.05030, 2022.

- [16] G. Buttazzo. Semicontinuity, relaxation and integral representation in the calculus of variations, volume 207. Longman, 1989.
- [17] E. Caglioti, F. Golse, and M. Iacobelli. A gradient flow approach to quantization of measures. Mathematical Models and Methods in Applied Sciences, 25(10):1845–1885, 2015.
- [18] J. Carrillo, Y. Huang, F. Patacchini, and G. Wolansky. Numerical study of a particle method for gradient flows. *Kinetic and Related Models*, 10(3), 2017.
- [19] J. A. Carrillo, A. Chertock, and Y. Huang. A finite-volume method for nonlinear nonlocal equations with a gradient flow structure. *Communications in Computational Physics*, 17(1):233–258, 2015.
- [20] J. A. Carrillo, K. Craig, and F. S. Patacchini. A blob method for diffusion. Calculus of Variations and Partial Differential Equations, 2019.
- [21] J. A. Carrillo, K. Craig, L. Wang, and C. Wei. Primal dual methods for wasserstein gradient flows. Foundations of Computational Mathematics, pages 1–55, 2021.
- [22] J. A. Carrillo, K. Craig, and Y. Yao. Aggregation-diffusion equations: dynamics, asymptotics, and singular limits. In Active Particles, Volume 2, pages 65–108. Springer, 2019.
- [23] J. A. Carrillo and J. S. Moll. Numerical simulation of diffusive and aggregation phenomena in nonlinear continuity equations by evolving diffeomorphisms. SIAM Journal on Scientific Computing, 31(6):4305–4329, 2010.
- [24] J. A. Carrillo, F. S. Patacchini, P. Sternberg, and G. Wolansky. Convergence of a particle method for diffusive gradient flows in one dimension. SIAM Journal on Mathematical Analysis, 48(6):3708–3741, 2016.
- [25] J. A. Carrillo, H. Ranetbauer, and M.-T. Wolfram. Numerical simulation of nonlinear continuity equations by evolving diffeomorphisms. *Journal of Computational Physics*, 327:186–202, 2016.
- [26] S. Chewi, T. L. Gouic, C. Lu, T. Maunu, and P. Rigollet. Svgd as a kernelized wasserstein gradient flow of the chi-squared divergence. arXiv preprint arXiv:2006.02509, 2020.
- [27] L. Chizat and F. Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. arXiv preprint arXiv:1805.09545, 2018.
- [28] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on robotics and Automation*, 20(2):243–255, 2004.
- [29] K. Craig. Nonconvex gradient flow in the wasserstein metric and applications to constrained nonlocal interactions. Proceedings of the London Mathematical Society, 114(1):60–102, 2017.
- [30] K. Craig and A. L. Bertozzi. A blob method for the aggregation equation. Math. Comp., 85(300):1681–1717, 2016.
- [31] K. Craig and I. Topaloglu. Aggregation-diffusion to constrained interaction: Minimizers & gradient flows in the slow diffusion limit. arXiv preprint arXiv:1806.07415, 2018.
- [32] S. Daneri, E. Radici, and E. Runa. Deterministic particle approximation of aggregation-diffusion equations on unbounded domains. *Journal of Differential Equations*, 312:474–517, 2022.
- [33] M. Di Francesco and M. D. Rosini. Rigorous derivation of nonlinear scalar conservation laws from follow-the-leader type models via many particle limit. Archive for rational mechanics and analysis, 217(3):831–871, 2015.
- [34] V. Dobrić and J. E. Yukich. Asymptotics for transportation cost in high dimensions. Journal of Theoretical Probability, 8(1):97–118, 1995.
- [35] J. Dolbeault, I. Gentil, A. Guillin, and F.-Y. Wang. L^q-functional inequalities and weighted porous media equations. Potential Analysis, 28:35–59, 2008.
- [36] R. M. Dudley. The speed of mean Glivenko-Cantelli convergence. The Annals of Mathematical Statistics, 40(1):40–50, 1969.
- [37] K. Elamvazhuthi, C. Adams, and S. Berman. Coverage and field estimation on bounded domains by diffusive swarms. In 2016 IEEE 55th Conference on Decision and Control (CDC), pages 2867–2874. IEEE, 2016.
- [38] K. Elamvazhuthi and S. Berman. Nonlinear generalizations of diffusion-based coverage by robotic swarms. In 2018 IEEE Conference on Decision and Control (CDC), pages 1341–1346. IEEE, 2018.
- [39] U. Eren and B. Açıkmeşe. Velocity field generation for density control of swarms using heat equation and smoothing kernels. IFAC-PapersOnLine, 50(1):9405–9411, 2017.
- [40] L. C. Evans, O. Savin, and W. Gangbo. Diffeomorphisms and nonlinear heat flows. SIAM journal on mathematical analysis, 37(3):737–751, 2005.
- [41] A. Figalli and F. Glaudo. An Invitation to Optimal Transport, Wasserstein Distances, and Gradient Flows. EMS Textbooks in Mathematics, 2021.
- [42] T. Gallouët, Q. Merigot, and A. Natale. Convergence of a lagrangian discretization for barotropic fluids and porous media flow. arXiv preprint arXiv:2105.12605, 2021.
- [43] S. Graf and H. Luschgy. Foundations of quantization for probability distributions. Springer, 2007.
- [44] G. Grillo, M. Muratori, and M. M. Porzio. Porous media equations with two weights: Smoothing and decay properties of energy solutions via poincaré inequalities. Discrete & Continuous Dynamical Systems, 33(8):3599–3640, 2013.
- [45] J. D. Hunter. Matplotlib: A 2D graphics environment. Computing in science & engineering, 9(03):90–95, 2007.
- [46] M. Iacobelli. A gradient flow perspective on the quantization problem. In PDE Models for Multi-Agent Phenomena, pages 145–165. Springer, 2018.
- [47] M. Iacobelli. Asymptotic analysis for a very fast diffusion equation arising from the 1d quantization problem. Discrete & Continuous Dynamical Systems, 39(9):4929, 2019.
- [48] A. Javanmard, M. Mondelli, and A. Montanari. Analysis of a two-layer neural network via displacement convexity. The Annals of Statistics, 48(6):3619–3642, 2020.
- [49] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker–Planck equation. SIAM journal on mathematical analysis, 29(1):1–17, 1998.

- [50] A. Korba, A. Salim, M. Arbel, G. Luise, and A. Gretton. A non-asymptotic analysis for Stein variational gradient descent. Advances in Neural Information Processing Systems, 33:4672–4682, 2020.
- [51] V. Krishnan and S. Martínez. Distributed optimal transport for the deployment of swarms. In 2018 IEEE Conference on Decision and Control (CDC), pages 4583–4588. IEEE, 2018.
- [52] P.-L. Lions and S. Mas-Gallic. Une méthode particulaire déterministe pour des équations diffusives non linéaires. Comptes Rendus de l'Académie des Sciences-Series I-Mathematics, 332(4):369–376, 2001.
- [53] Q. Liu. Stein variational gradient descent as gradient flow. Advances in neural information processing systems, 30, 2017.
- [54] Q. Liu and D. Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. Advances in neural information processing systems, 29, 2016.
- [55] J. Lu, Y. Lu, and J. Nolen. Scaling limit of the Stein variational gradient descent: The mean field regime. SIAM Journal on Mathematical Analysis, 51(2):648–671, 2019.
- [56] D. Matthes, R. J. McCann, and G. Savaré. A family of nonlinear fourth order equations of gradient flow type. Communications in Partial Differential Equations, 34(11):1352–1397, 2009.
- [57] D. Matthes and H. Osberger. Convergence of a variational lagrangian scheme for a nonlinear drift diffusion equation. ESAIM: Mathematical Modelling and Numerical Analysis, 48(3):697-726, 2014.
- [58] D. Matthes and B. Söllner. Convergent lagrangian discretization for drift-diffusion with nonlocal aggregation. In *Innovative algorithms and analysis*, pages 313–351. Springer, 2017.
- [59] S. Mei, A. Montanari, and P.-M. 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.
- [60] Q. Mérigot. A multiscale approach to optimal transport. In Computer Graphics Forum, volume 30, pages 1583–1592. Wiley Online Library, 2011.
- [61] A. R. Mesquita, J. P. Hespanha, and K. Åström. Optimotaxis: A stochastic multi-agent optimization procedure with point measurements. In International workshop on hybrid systems: Computation and control, pages 358–371. Springer, 2008.
- [62] K. Oelschläger. Large systems of interacting particles and the porous medium equation. Journal of differential equations, 88(2):294–346, 1990.
- [63] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis. Cupy: A numpy-compatible library for nvidia gpu calculations. 31st Conference on Neural Information Processing Systems, 2017.
- [64] F. Otto. The geometry of dissipative evolution equations: the porous medium equation. Comm. Partial Differential Equations, 26(1-2):101–174, 2001.
- [65] F. S. Patacchini and D. Slepčev. The nonlocal-interaction equation near attracting manifolds. arXiv preprint arXiv:2106.01823, 2021.
- [66] G. M. Rotskoff and E. Vanden-Eijnden. Trainability and accuracy of neural networks: An interacting particle system approach. arXiv preprint arXiv:1805.00915, 2018.
- [67] H. L. Royden and P. Fitzpatrick. *Real analysis*, volume 32. Macmillan New York, 1988.
- [68] F. Santambrogio. Optimal transport for applied mathematicians. Birkäuser, NY, 55(58-63):94, 2015.
- [69] S. Serfaty. Gamma-convergence of gradient flows on hilbert and metric spaces and applications. Discrete Contin. Dyn. Syst, 31(4):1427–1451, 2011.
- [70] J. Sirignano and K. Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. SIAM Journal on Applied Mathematics, 80(2):725–752, 2020.
- [71] Z. Sun, J. A. Carrillo, and C.-W. Shu. A discontinuous galerkin method for nonlinear parabolic equations and gradient flow problems with interaction potentials. *Journal of Computational Physics*, 352:76–104, 2018.
- [72] A. B. Tsybakov. Introduction to nonparametric estimation. Springer Series in Statistics, 2009.
- [73] S. Van Der Walt, S. C. Colbert, and G. Varoquaux. The numpy array: a structure for efficient numerical computation. Computing in science & engineering, 13(2):22–30, 2011.
- [74] J. L. Vázquez. The mathematical theories of diffusion: Nonlinear and fractional diffusion. In Nonlocal and nonlinear diffusions and interactions: new methods and directions, pages 205–278. Springer, 2017.
- [75] C. Villani. Topics in optimal transportation, volume 58. American Mathematical Soc., 2003.
- [76] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.
- [77] N. K. Vishnoi. An introduction to Hamiltonian Monte Carlo method for sampling. arXiv preprint arXiv:2108.12107, 2021.
- [78] J. Weed and F. Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. Bernoulli, 25(4A):2620–2648, 2019.
- [79] E. Weinan, C. Ma, and L. Wu. Machine learning from a continuous viewpoint, i. Science China Mathematics, 63(11):2233– 2266, 2020.
- [80] M. Westdickenberg and J. Wilkening. Variational particle schemes for the porous medium equation and for the system of isentropic euler equations. ESAIM: Mathematical Modelling and Numerical Analysis, 44(1):133-166, 2010.
- [81] S. Wojtowytsch. On the convergence of gradient descent training for two-layer relu-networks in the mean field regime. arXiv preprint arXiv:2005.13530, 2020.

UNIVERSITY OF CALIFORNIA, SANTA BARBARA, DEPARTMENT OF MATHEMATICS *Email address*: kcraig@math.ucsb.edu

University of California, Los Angeles, Department of Mathematics $Email \ address: {\tt karthikevaz@math.ucla.edu}$

 $\label{eq:california} California \ Polytechnic \ State \ University, \ BioResource \ and \ Agricultural \ Engineering \ Department \ Email \ address: \ maberla@calpoly.edu$

MICHIGAN STATE UNIVERSITY, DEPARTMENT OF MATHEMATICS *Email address*: turanova@msu.edu