

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 to *confined* mean-field dynamics of training a two-layer neural network for a radial basis activation function. From this perspective, our convergence result shows that, in the overparametrized 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.

1. INTRODUCTION

Quantization is a fundamental problem throughout the sciences, in which one seeks to approximate a continuum distribution or signal by discrete objects [41].

Received by the editor March 9, 2022, and, in revised form, February 7, 2023, and February 14, 2023.

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

The work of the first author was supported by NSF DMS grants 1811012 and 2145900, as well as a Hellman Faculty Fellowship. The first and fourth authors were supported by the Simons Center for Theory of Computing, at which part of this work was completed. The work of the second author was supported by AFOSR grants FA9550-18-1-0502 and FA9550-18-1-0502. The work of the fourth author was supported by NSF DMS grant 1907221 and NSF DMS grant 2204722.

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 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* [12, 27], 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 [10, 11, 16, 44, 45, 58]. 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 [9, 52, 75]. In the control theory literature, recent work has developed multi-agent 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 [35–37, 49, 59]. Other authors have explored the role of different notions of optimality in designing coverage algorithms [4, 5].

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,

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

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 strictly positive target $\bar{\rho} : \mathbb{R}^d \rightarrow \mathbb{R}$ that is log-concave on Ω and satisfies $\int_{\Omega} \bar{\rho} = 1$, and a fixed external potential $V \in C^2(\Omega)$, we consider the equation,

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

with no-flux boundary conditions on $\partial\Omega$. The initial conditions are chosen to satisfy $\rho_0 \geq 0$ 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 $t \rightarrow +\infty$ 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 $N \rightarrow +\infty$, to a solution $\rho(t)$ of (WPME) on bounded time intervals. In Sections 1.4-1.5, 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 [42, 72]. From this perspective, Theorem 1.2 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 [60], Lions and MasGallic [50], Carrillo, Craig, and Patacchini [19], and Burger and Esposito [14], 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 [7, 13, 18, 69], as well as methods based on alternative deterministic particle methods in one spatial dimension [17, 23, 30, 31, 56], Lagrangian evolution of the transport map along the flow [22, 24, 38, 55, 78], and many others [6, 20, 21, 40]. 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 [52], 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 [48, 51, 53]. Recent work by Chewi et al. [25] 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 to 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 that 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 with *confined* dynamics for training a two-layer neural network with a radial basis 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) when confined to the domain Ω ; see Corollary 1.4. This generalizes the previous convergence result of Javanmard, Mondelli, and Montanari [46] to the case of nonuniform data distributions. As a consequence of this result, we are able to provide conditions on the target function and data distribution that guarantee that the continuum limit of the training dynamics of two-layer 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 [26, 79].

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 it is a *Wasserstein gradient flow* of the energy,

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

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

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

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

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

The internal energy $\mathcal{E}$ induces the nonlinear diffusion term, the external potential $\mathcal{V}$ induces the convection term, and the confining potential $\mathcal{V}_\Omega$ 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 [47], Jordan, Kinderlehrer, and Otto established that (FP) is the Wasserstein gradient flow of the *Kullback-Leibler divergence*,

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

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 [70],

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

where $\mathcal{L}^d$ denotes d -dimensional Lebesgue measure. 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 is the available estimates quantifying its convergence to equilibrium. Chewi et al. [25] 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:

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

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

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

This mirrors the theory for (FP), in which a Poincaré inequality 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 [54]. In addition, see Grillo, Muratori, and Porzio [42], 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. [33] 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,

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

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

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

$$(1.8) \quad \mathcal{V}_{\varepsilon}(\mu) = \int_{\mathbb{R}^d} (\zeta_{\varepsilon} * V)(x) d\mu(x),$$

$$(1.9) \quad \mathcal{V}_k(\mu) = \int_{\mathbb{R}^d} V_k(x) d\mu(x).$$

Here $\zeta_{\varepsilon}(x) = \varepsilon^{-d} \zeta(x/\varepsilon)$ for $\zeta \in C^{\infty}(\mathbb{R}^d)$ an even, 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 $k \rightarrow +\infty$.

The energy $\mathcal{E}_{\varepsilon}(\mu)$ is an approximation, as $\varepsilon \rightarrow 0$, 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,

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

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 $\mathcal{V}$ and likewise enjoys the analogous property $\mathcal{V}_{\varepsilon}(\rho) = \mathcal{V}(\zeta_{\varepsilon} * \rho)$. 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 $k \rightarrow +\infty$, of $\mathcal{V}_{\Omega}$.

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 $\mathcal{V}_{\Omega}$ as $k \rightarrow +\infty$ 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 [63].) Our motivations for considering this approximation of the confining potential $\mathcal{V}_{\Omega}$ 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,

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

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,

$$(1.11) \quad \begin{cases} \dot{X}^i(t) = -\sum_{j=1}^N f(X^i, X^j) m^j - \nabla \zeta_\varepsilon * V(X^i) - \nabla V_k(X^i), \\ X^i(0) = X_0^i, \end{cases}$$

for

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

see Proposition 3.13. Here, we again leverage the fact that ζ_ε is even. 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 of (WPME) have bounded entropy, the initial conditions $\rho_0^N(0) = \sum_{i=1}^N \delta_{X_0^i} m^i$ of the particle method converge sufficiently quickly to ρ_0 , and $\varepsilon \rightarrow 0$ and $k \rightarrow +\infty$ sufficiently rapidly; see Theorem 1.2. 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 $\varepsilon \rightarrow 0$, $k \rightarrow +\infty$ limit builds on several previous works. All previous works have considered the spatially homogeneous case $\bar{\rho} \equiv 1$. The first work in this direction was due to Oelschläger [60], 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 [50], also in the case $V = V_k = 0$, proved convergence of $(WPME_{\varepsilon,k})$ as $\varepsilon \rightarrow 0$, provided that the initial conditions ρ_0 of both $(WPME_{\varepsilon,k})$ and (WPME) 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 an $\dot{H}^1$ 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, the first author, and Patacchini [19] generalized Lions and Mas-Gallic's approach to porous medium equations of the form,

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

In the case $m = 2$, they obtained convergence of the $\varepsilon \rightarrow 0$ limit under appropriate continuity and semiconvexity assumptions on V and W ; for $1 \leq m < 2$, they obtained Γ -convergence of the corresponding energies as $\varepsilon \rightarrow 0$; and for $m > 2$, they obtained conditional convergence of the $\varepsilon \rightarrow 0$ limit, as long as certain a priori estimates were preserved along the flow. Again, Carrillo, Craig, and Patacchini's work required the initial data of both (WPME $_{\varepsilon,k}$) and (WPME) to have bounded entropy, excluding particle solutions. Very recently, Burger and Esposito [14] 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 of (WPME $_{\varepsilon,k}$) have bounded entropy by using stability properties of the regularized flow, so that only the initial data of (WPME) must have bounded entropy; see Theorem 1.2. Our result holds for spatially inhomogeneous porous medium equations, allowing general $\bar{\rho} \in C^1(\mathbb{R}^d)$ that are bounded above and below on $\mathbb{R}^d$ and log-concave on $\Omega \subseteq \mathbb{R}^d$. (See Section 1.4 for a discussion of where the log-concavity assumption may be weakened.) Finally, by allowing spatially inhomogeneous 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 represent a type of *confined* 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*:

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

In existing literature on this application, ζ_ε is typically assumed to be radial, and Φ_ε is referred to as a *radial basis function activation*. Our assumption on ζ_ε merely requires it to be even, so our results certainly hold under the conventional stronger assumption that ζ_ε is radial.

In several recent works, it was discovered that evolving the parameters $X^i(t)$ by gradient descent of the function $(X^1, \dots, X^n) \mapsto \mathcal{R}_\varepsilon(\rho^N) + \mathcal{V}_\Omega(\rho^N)$ is equivalent to evolving the empirical measure ρ^N by the Wasserstein gradient flow of $\mathcal{R}_\varepsilon$ restricted to Ω [26, 46, 57, 64, 68, 77]. Various methods for treating the boundary conditions are considered, including projection of the gradient descent direction into the convex hull of the domain Ω [26] or projection onto interior approximations of Ω [46].

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) for the even mollifier ζ_ε), we obtain,

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

for,

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

Moreover, the confining potential $\mathcal{V}_k$ provides an explicit method for projecting gradient descent dynamics onto an *exterior* approximation of Ω . In this way, since the addition of a constant to the energy does not affect the dynamics of the gradient flow, the *confined* training dynamics given by the gradient flow of $\mathcal{R}_\varepsilon + \mathcal{V}_k = \mathcal{F}_{\varepsilon,k} + C$ for general initial data ρ_0 is characterized by (WPME $_{\varepsilon,k}$), and the evolution for particle initial data corresponds to (1.11)-1.12. Corollary 1.4, 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,

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

This generalizes previous work due to Javanmard, Mondelli, and Montanari [46], which considered the limit $\varepsilon \rightarrow 0$ in the specific case of a uniform data distribution $\nu = \mathbb{1}_\Omega/|\Omega|$, smooth target function f , bounded convex domain Ω , and compactly supported radial basis function ζ . 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 1.4, 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,

$$(D) \quad \Omega \text{ is nonempty, open, and convex.}$$

We suppose our mollifier satisfies,

$$(M) \quad \begin{aligned} &\zeta \in C^2(\mathbb{R}^d) \text{ is even, nonnegative, } \|\zeta\|_{L^1(\mathbb{R}^d)} = 1, D^2\zeta \in L^\infty(\mathbb{R}^d), \\ &\zeta(x) \leq C_\zeta |x|^{-q} \text{ and } |\nabla \zeta(x)| \leq C_\zeta |x|^{-q'}, \text{ for } C_\zeta > 0, \quad q > d+1, \quad q' > d. \end{aligned}$$

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,

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

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 $D^2V_k \in L^\infty(\mathbb{R}^d)$,
 (Ck) $V_k = 0$ on Ω and $\lim_{k \rightarrow \infty} \left(\inf_{x \in B} V_k(x) \right) = +\infty$ for any ball $B \subset\subset \Omega^c$.

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 $k \rightarrow +\infty$. In particular, note that (Ck) implies that, in the $k \rightarrow \infty$ limit, V_k approximates the hard cutoff potential V_Ω , which is given by,

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

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

- (T) $\bar{\rho} \in C^1(\mathbb{R}^d)$, $\nabla \bar{\rho} \in L^\infty(\mathbb{R}^d)$, 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 $\mathcal{F}_{\varepsilon,k}$ is lower semicontinuous, convex, and subdifferentiable, so that gradient flows of $\mathcal{F}_{\varepsilon,k}$ are well-posed. It also allows us to conclude that the energy $\mathcal{F}$ is lower semicontinuous. However, in order to obtain convexity and subdifferentiability of $\mathcal{F}$, hence well-posedness of gradient flows, we require $\bar{\rho}$ to be log-concave on Ω ; that is,

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

It is an open question whether well-posedness of the gradient flows of $\mathcal{F}$ could be obtained under weaker assumptions on $\bar{\rho}$. Interestingly, the main estimates in our proof of the convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ as $\varepsilon \rightarrow 0$ (Theorem 4.1, Theorem 5.1, and Proposition 5.6) do not require log-concavity of $\bar{\rho}$. Instead, log-concavity comes into play when we seek to identify that the limit as $\varepsilon \rightarrow 0$ and $k \rightarrow \infty$ of gradient flows of $\mathcal{F}_{\varepsilon,k}$ is indeed a gradient flow of $\mathcal{F}$, since log-concavity of $\bar{\rho}$ ensures that the *metric slope* of $\mathcal{F}$ is a *strong upper gradient*; see Section 2.3 and [3, Section 1.2]. For this reason, we are optimistic that, in future work, it may 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 $\mathcal{F}$ and characterization of its strong upper gradient are overcome.

1.5. Main results. To state our main results, first we introduce some notation. Let the *entropy* $\mathcal{S}(\mu)$ and p -th *moment* $M_p(\mu)$, where $p \geq 1$, of a measure $\mu \in \mathcal{P}(\mathbb{R}^d)$

be given by

$$(1.18) \quad \begin{aligned} \mathcal{S}(\mu) &= \begin{cases} \int_{\mathbb{R}^d} \mu(x) \log \mu(x) d\mathcal{L}^d(x) & \text{if } \mu \ll \mathcal{L}^d \text{ and } d\mu(x) = \mu(x)dx, \\ +\infty & \text{otherwise,} \end{cases} \\ M_p(\mu) &= \int_{\mathbb{R}^d} |x|^p d\mu(x). \end{aligned}$$

Recall that a probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$ lies in the *domain* of an energy $\mathcal{G} : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R} \cup \{+\infty\}$ if $\mathcal{G}(\mu) < +\infty$. We denote this by $\mu \in D(\mathcal{G})$. We also write

$$(1.19) \quad \mathcal{P}_p(\mathbb{R}^d) = \mathcal{P}(\mathbb{R}^d) \cap D(M_p), \text{ for } p \geq 1.$$

Finally, we often use the notion of *narrow* convergence of probability measures (see Definition 2.2) and 2-absolutely continuous curves $AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ (see Definition 2.4).

Theorem 1.1 (Convergence of gradient flows as $k \rightarrow +\infty$, $\varepsilon = \varepsilon(k) \rightarrow 0$). *Assume (D), (M), (V), (C), (Ck), (T) and that $\bar{\rho}$ is log-concave on Ω . 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 the gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial data $\rho(0)$. Then, as $k \rightarrow +\infty$, there exists a sequence $\varepsilon = \varepsilon(k) \rightarrow 0$ so that

$$\lim_{k \rightarrow +\infty} W_1(\rho_{\varepsilon,k}(t), \rho(t)) = 0, \text{ uniformly for } t \in [0, T],$$

where $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ is the gradient flow of $\mathcal{F}$ with initial data $\rho(0)$.

Theorem 1.1 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.5, we provide an explicit construction of $\rho_{\varepsilon,k}^N(0)$ satisfying condition (1.21).

Theorem 1.2 (Convergence with particle initial data). *Assume (D), (M), (V), (C), (Ck), (T), and that $\bar{\rho}$ is log-concave on Ω . 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,*

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

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

$$(1.20) \quad \begin{cases} \dot{X}_{\varepsilon,k}^i = - \sum_{j=1}^N m^j \int_{\mathbb{R}^d} \nabla \zeta_{\varepsilon}(X_{\varepsilon,k}^i - z) \zeta_{\varepsilon}(z - X_{\varepsilon,k}^j) \frac{1}{\bar{\rho}(z)} dz \\ \quad - \nabla(\zeta_{\varepsilon} * V)(X_{\varepsilon,k}^i) - \nabla V_k(X_{\varepsilon,k}^i), \\ X_{\varepsilon,k}^i(0) = X_{0,\varepsilon}^i. \end{cases}$$

Suppose that as $\varepsilon \rightarrow 0$ there exist $N = N(\varepsilon) \rightarrow +\infty$, 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,

$$(1.21) \quad \lim_{k \rightarrow \infty} e^{-\lambda_\varepsilon T} W_2(\rho_{\varepsilon,k}^N(0), \rho(0)) = 0,$$

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

Then, as $k \rightarrow +\infty$, there exist $\varepsilon = \varepsilon(k) \rightarrow 0$ and $N = N(\varepsilon) \rightarrow +\infty$ for which $\rho_{\varepsilon,k}^N(t) = \sum_{i=0}^N \delta_{X_{\varepsilon,k}^i(t)} m^i$ satisfies

$$\lim_{k \rightarrow +\infty} W_1(\rho_{\varepsilon,k}^N(t), \rho(t)) \rightarrow 0, \text{ uniformly for } t \in [0, T],$$

where $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ is the unique weak solution of (WPME) with initial data $\rho(0)$.

Corollary 1.3 ensures that the particle method defined in Theorem 1.2 indeed converges to $\bar{\rho}$ on Ω in the long time limit, as relevant for applications in quantization. The idea behind the proof is as follows: given arbitrary $\delta > 0$, our log-concavity assumption for $\bar{\rho}$ on Ω allows us to prove that, for any exact solution $\rho(x, t)$ of (WPME), we may pick t sufficiently large so that $\rho(x, t)$ is within distance $\delta/2$ of $\mathbb{1}_{\bar{\Omega}} \bar{\rho}$ in the 1-Wasserstein metric, where $\mathbb{1}_{\bar{\Omega}}$ denotes the indicator function on $\bar{\Omega}$. For such a t , Theorem 1.2 ensures that for k sufficiently large (depending on t), ε sufficiently large (depending on k), and N sufficiently large (depending on ε), the particle solution $\rho_{\varepsilon,k}^N(t)$ is within distance $\delta/2$ of $\rho(x, t)$. In this way, there exists a choice of parameters so that the particle solution $\rho_{\varepsilon,k}^N(t)$ is arbitrarily close to $\mathbb{1}_{\bar{\Omega}} \bar{\rho}$.

Corollary 1.3 (Long time limit). *Suppose the assumptions of Theorem 1.2 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 $k = k(t) \rightarrow +\infty$, $\varepsilon = \varepsilon(k) \rightarrow 0$, and $N = N(\varepsilon) \rightarrow +\infty$ so that*

$$\lim_{t \rightarrow +\infty} W_1(\rho_{\varepsilon,k}^N(\cdot, t), \mathbb{1}_{\bar{\Omega}} \bar{\rho}) = 0.$$

Theorems 1.1 and 1.2 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.5 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

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

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 confined training dynamics of two-layer neural networks with a radial basis activation function and quadratic loss, as described in Section 1.3. 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 1.4 (Two-layer neural networks). *Consider a domain Ω satisfying (D), 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 on Ω , and a target function $f_0 = -V\bar{\rho}$, for V satisfying (V). For $k \in \mathbb{N}$, consider a confining potential V_k satisfying (C) and (Ck). Fix $T > 0$. For $\varepsilon > 0$, $N \in \mathbb{N}$, and $t \in [0, T]$, consider the confined training dynamics of a two-layer neural network corresponding to the energy $\mathcal{R}_\varepsilon + \mathcal{V}_k$; that is, consider the evolution of the empirical measure of parameters,*

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

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

$$(1.22) \quad \begin{cases} \dot{X}_\varepsilon^i = - \sum_{j=1}^N m^j \int_{\mathbb{R}^d} \nabla \zeta_\varepsilon(X_\varepsilon^i - z) \zeta_\varepsilon(z - X_\varepsilon^j) \nu(z) dz \\ \quad + \nabla(\zeta_\varepsilon * (f_0 \nu))(X_\varepsilon^i) - \nabla V_k(X_{\varepsilon,k}^i), \\ X_\varepsilon^i(0) = X_{0,\varepsilon}^i. \end{cases}$$

Suppose there exists $\rho(0) \in D(\mathcal{F}) \cap D(\mathcal{S}) \cap \mathcal{P}_2(\mathbb{R}^d)$ so that, for all $\varepsilon > 0$, there exists $N = N(\varepsilon)$ so that $\rho_\varepsilon^N(0)$ converges to $\rho(0)$ sufficiently quickly, according to the rate from equation (1.21). Then, as $k \rightarrow +\infty$, there exist $\varepsilon = \varepsilon(k) \rightarrow 0$ and $N = N(\varepsilon) \rightarrow +\infty$ for which $\rho_\varepsilon^N(t) = \sum_{i=1}^N \delta_{X_\varepsilon^i(t)} m^i$ satisfies

$$\lim_{k \rightarrow +\infty} W_1(\rho_{\varepsilon,k}^N(t), \rho(t)) \rightarrow 0, \quad \text{uniformly for } t \in [0, T],$$

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 on Ω and $f_0 \nu$ is concave on Ω , 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 $\mathcal{F}$. Our proof of Theorem 1.1 on the convergence of gradient flows as $k \rightarrow +\infty$ and $\varepsilon = \varepsilon(k) \rightarrow 0$ leverages the perspective of Serfaty's general metric space framework for Γ -convergence of gradient flows [67], 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 $\mathcal{F}$. Unlike in Theorem 1.1 on convergence of the gradient flows, the rate at which $\varepsilon \rightarrow 0$ does not depend on the rate $k \rightarrow +\infty$. Likewise, this result does not require $\bar{\rho}$ to be log-concave on Ω .

Theorem 1.5 (Minimizers converge to minimizers). *Suppose Assumptions (T), (D), (M), (V), (C), and (Ck) hold. Suppose further that $V_k \geq V_1$ for all $k \in \mathbb{N}$ and the sublevel sets of V_1 are compact. Then, if $\rho_{\varepsilon,k} \in \mathcal{P}_2(\mathbb{R}^d)$ is a minimizer of*

$\mathcal{F}_{\varepsilon,k}$ for all $\varepsilon > 0$, $k \in \mathbb{N}$, then as $\varepsilon \rightarrow 0$, $k \rightarrow +\infty$, $\rho_{\varepsilon,k}$ narrowly converges to $\rho \in \mathcal{P}(\mathbb{R}^d)$, where ρ is the unique minimizer of $\mathcal{F}$.

Theorem 1.5 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 $t \rightarrow +\infty$ to a minimizer of $\mathcal{F}_{\varepsilon,k}$, uniformly in $\varepsilon > 0$ and $k \in \mathbb{N}$, then one could combine Theorem 1.5 with Theorems 1.1 and 1.2 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 our variant of Serfaty’s 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 energy $\mathcal{F}$, including convexity and differentiability in Section 3.1. We give the PDE characterizations of the 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 $\varepsilon \rightarrow 0$ and $k \rightarrow +\infty$. Section 4 is devoted to proving a key estimate for the analysis of the $\varepsilon \rightarrow 0$ 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 [54].

In Section 5, we use the results of Section 4 to study the $\varepsilon \rightarrow 0$ limit. In Section 5.1, we obtain Γ -convergence of the energies $\mathcal{E}_\varepsilon + \mathcal{V}_\varepsilon$ as $\varepsilon \rightarrow 0$. In Section 5.2, we move on to considering convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ as $\varepsilon \rightarrow 0$ in Proposition 5.6, under the key hypothesis that the initial conditions of the gradient flow have uniformly bounded entropy. Ultimately, we prove the gradient flows converge to an “almost” curve of maximal slope (see Definition 5.2) of the intermediate energy $\mathcal{F}_k$, defined by

$$(1.23) \quad \mathcal{F}_k(\rho) = \mathcal{E}(\rho) + \mathcal{V}(\rho) + \mathcal{V}_k(\rho).$$

We introduce the notion of an “almost” curve of maximal slope as a weakening of the traditional notion of gradient flow. We need this weakening because we only suppose $\bar{\rho}$ is log-concave on Ω , instead of on all of $\mathbb{R}^d$. Heuristically, this causes the energy $\mathcal{F}_k$ to lack sufficient regularity to define its gradient flow. More precisely, our weak assumptions on $\bar{\rho}$ prevent us from characterizing the strong upper gradient of $\mathcal{F}_k$, which likewise prevent us from defining its curve of maximal slope.

In spite of the fact that our hypotheses on $\bar{\rho}$ are too weak to identify the $\varepsilon \rightarrow 0$ limit of gradient flows of the energies $\mathcal{F}_{\varepsilon,k}$ as a true gradient flow, it turns out that our notion of “almost” curve of maximal slope is sufficient to identify the behavior as $k \rightarrow +\infty$. We consider this limit in Section 6, first obtaining Γ -convergence of the energies $\mathcal{F}_k$ to $\mathcal{F}$, as $k \rightarrow +\infty$, as well as our main theorem on convergence of the

minimizers, Theorem 1.5. We then prove that, as $k \rightarrow +\infty$, the “almost” curves of maximal slopes of $\mathcal{F}_k$ converge to the unique gradient flow of $\mathcal{F}$. Combining this with the $\varepsilon \rightarrow 0$ result from the previous section, we prove our main result, Theorem 1.1, showing that gradient flows of $\mathcal{F}_{\varepsilon,k}$ converge to a gradient flow of $\mathcal{F}$ as $k \rightarrow +\infty$ and $\varepsilon \rightarrow 0$ sufficiently quickly. With these results in hand, we turn in Section 6.1 to extending the preceding convergence results on the gradient flows as $\varepsilon \rightarrow 0$, $k \rightarrow +\infty$ to allow for gradient flows with particle initial data, thereby obtaining the proof of Theorem 1.2. We also prove Corollary 1.3 on the long time behavior of the particle method and Corollary 1.4 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 $N \rightarrow +\infty$, $\varepsilon \rightarrow 0$. 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 on Ω , 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 results on the rate of convergence depending on $N \in \mathbb{N}$, $\varepsilon > 0$, and $k \in \mathbb{N}$, particularly if these 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 $\Omega \subseteq \mathbb{R}^d$ 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,

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

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, [19, Lemma 2.2]). *Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ 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,*

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

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 $\int f d\mu_n \rightarrow \int f d\mu$ 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 $f \in C_b(\mathbb{R}^d)$,

$$(2.2) \quad \int f(\phi * \mu_n) = \int (f * \phi) d\mu_n \xrightarrow{n \rightarrow +\infty} \int (f * \phi) d\mu = \int f(\phi * \mu),$$

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

Lemma 2.3 (Mollifiers and narrow convergence, [19, 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, 39, 66, 73].

For a Borel measurable map $\mathbf{t}: \mathbb{R}^n \rightarrow \mathbb{R}^m$, 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 *push-forward* 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,

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

where $\pi^1, \pi^2: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ are the projections of $\mathbb{R}^d \times \mathbb{R}^d$ onto the first and second copies of $\mathbb{R}^d$, respectively. For $p \geq 1$, the *p-Wasserstein distance* [3, Chapter 7] between $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ is given by,

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

where the definitions of the p -th moment and the space $\mathcal{P}_p(\mathbb{R}^d)$ were recalled in (1.18) and (1.19). 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 $\Gamma_0(\mu, \nu)$.

We make the following observation: if γ is a transport plan from a measure μ to a Dirac mass δ_0 , then $\gamma = (\mathbf{id} \times 0)_\# \mu$, where 0 denotes the function $0: x \mapsto 0$. Using this in the definition (2.3), we obtain,

$$(2.4) \quad W_p^p(\mu, \delta_0) = \int_{\mathbb{R}^d \times \mathbb{R}^d} |x - y|^p d\gamma(x, y) = \int_{\mathbb{R}^d \times \mathbb{R}^d} |x - 0|^p d\mu(x) = M_p(\mu).$$

Note that applying Hölder's inequality in the definition of $W_p(\mu, \nu)$ yields,

$$(2.5) \quad W_p(\mu, \nu) \leq W_{p'}(\mu, \nu) \quad \text{for } 1 \leq p \leq p' \text{ and for all } \mu, \nu \in \mathcal{P}_{p'}(\mathbb{R}^d).$$

Convergence with respect to the p -Wasserstein metric is stronger than narrow convergence of probability measures [3, Remark 7.1.11]. In particular, for $p > 1$, if μ_n is a sequence in $\mathcal{P}_p(\mathbb{R}^d)$ and $\mu \in \mathcal{P}_p(\mathbb{R}^d)$, we have,

$$(2.6) \quad W_p(\mu_n, \mu) \rightarrow 0 \text{ as } n \rightarrow \infty \iff (\mu_n \rightarrow \mu \text{ narrowly and } M_p(\mu_n) \rightarrow M_p(\mu) \text{ as } n \rightarrow \infty).$$

When $p = 1$, the analogous result holds if convergence of first moments is replaced with the requirement that the first moments are *uniformly integrable* [3, Proposition 7.1.5].

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 $\mu: [0, T] \rightarrow \mathcal{P}(\mathbb{R}^d)$ is *2-absolutely continuous* on $[0, T]$, and write $\mu \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$, if there exists $f \in L^2([0, T])$ so that,

$$(2.7) \quad W_2(\mu(t), \mu(s)) \leq \int_s^t f(r) dr \quad \text{for all } t, s \in (0, T) \text{ with } s \leq t.$$

Along such curves, we may define the metric derivative.

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

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

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 (2.7): for any $\mu \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$, we have $|\mu'| \in L^2((0, T))$, and for any function f satisfying (2.7), 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,

$$(2.8) \quad \mu_\alpha = ((1 - \alpha)\pi^1 + \alpha\pi^2)_{\#}\gamma \quad \text{for } \alpha \in [0, 1], \gamma \in \Gamma_0(\mu, \nu).$$

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,

$$(2.9) \quad \mu_\alpha^{2 \rightarrow 3} = ((1 - \alpha)\pi^2 + \alpha\pi^3)_{\#}\gamma \quad \text{for } \alpha \in [0, 1] \text{ and } \gamma \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d) \\ \text{such that } \pi^{1,2}_{\#}\gamma \in \Gamma_0(\mu_1, \mu_2) \text{ and } \pi^{1,3}_{\#}\gamma \in \Gamma_0(\mu_1, \mu_3),$$

with $\pi^{1,i}: \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d \times \mathbb{R}^d$ the projection of $\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^d$ 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 $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \rightarrow (-\infty, \infty]$ 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:

$$(2.10) \quad \mathcal{G}(\mu_\alpha^{2 \rightarrow 3}) \leq (1 - \alpha)\mathcal{G}(\mu_2) + \alpha\mathcal{G}(\mu_3) - \alpha(1 - \alpha)\frac{\lambda}{2}W_{2,\gamma}^2(\mu_2, \mu_3) \quad \text{for all } \alpha \in [0, 1],$$

where,

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

In this case, we 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, [28, Proposition 2.8]). *A functional $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \rightarrow (-\infty, \infty]$ is λ -convex along generalized geodesics if, for all generalized geodesics $\mu_\alpha^{2 \rightarrow 3}$ connecting μ_2 to μ_3 with base μ_1 , the map $\alpha \mapsto \mathcal{G}(\mu_\alpha^{2 \rightarrow 3})$ is differentiable for all $\alpha \in [0, 1]$ and,*

$$\mathcal{G}(\mu_3) - \mathcal{G}(\mu_2) - \frac{d}{d\alpha}\mathcal{G}(\mu_\alpha^{2 \rightarrow 3}) \Big|_{\alpha=0} \geq \frac{\lambda}{2}W_{2,\gamma}^2(\mu_2, \mu_3).$$

For any functional $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \rightarrow (-\infty, +\infty]$, 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 $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \rightarrow (-\infty, \infty]$, for any $\mu \in D(\mathcal{G})$, the *local slope* is,

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

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

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

$$(2.11) \quad \mathcal{G}(\nu) - \mathcal{G}(\mu) \geq \int_{\mathbb{R}^d \times \mathbb{R}^d} \langle \xi(x), y - x \rangle d\gamma(x, y) + \frac{\lambda}{2} W_2^2(\mu, \nu) \quad \text{for all } \gamma \in \Gamma_0(\mu, \nu).$$

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 $\xi_1 \in \partial\mathcal{G}_1(\mu)$ and $\xi_2 \in \partial\mathcal{G}_2(\mu)$, we have $\xi_1 + \xi_2 \in \partial(\mathcal{G}_1 + \mathcal{G}_2)(\mu)$.

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

$$(2.12) \quad |\partial\mathcal{G}|(\mu) \leq \inf \{ \|\xi\|_{L^2(\mu)} : \xi \in \partial\mathcal{G}(\mu) \}.$$

If equality holds and ξ attains the infimum, we will write $\xi = \partial^\circ\mathcal{G}(\mu)$. 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 $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \rightarrow (-\infty, +\infty]$ is proper, lower semicontinuous, and λ -convex along generalized geodesics. A curve $\mu(t) \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ is a *gradient flow* of $\mathcal{G}$ in the Wasserstein metric if $\mu(t)$ is a weak solution of the continuity equation,

$$(2.13) \quad \partial_t \mu(t) + \nabla \cdot (v(t)\mu(t)) = 0, \text{ in duality with } C_c^\infty((0, T) \times \mathbb{R}^d),$$

and

$$v(t) = -\partial^\circ\mathcal{G}(\mu(t)) \text{ for } \mathcal{L}^1\text{-a.e. } t \in [0, T].$$

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.

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. Note that our notion of curve of maximal slope differs slightly from Ambrosio, Gigli, and Savaré, since we use a version that is integrated in time.

Theorem 2.13 (Well-posedness and characterization of gradient flow). *Suppose $\mathcal{G}: \mathcal{P}_2(\mathbb{R}^d) \rightarrow (-\infty, +\infty]$ 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 $\lim_{t \rightarrow 0^+} \mu(t) = \mu(0)$ in the Wasserstein metric.*

Furthermore $\mu(t) \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ 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:*

(2.14)

$$\frac{1}{2} \int_0^t |\mu'|^2(r) dr + \frac{1}{2} \int_0^t |\partial \mathcal{G}|^2(\mu(r)) dr \leq \mathcal{G}(\mu(0)) - \mathcal{G}(\mu(t)), \quad \text{for all } t \in [0, T].$$

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

$$\frac{1}{2} \frac{d^+}{dt} W_2^2(\mu(t), \nu) + \frac{\lambda}{2} W_2^2(\mu(t), \nu) + \mathcal{G}(\mu(t)) \leq \mathcal{G}(\nu).$$

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. This provides sufficient conditions that, when combined with some compactness, ensure minimizers of a sequence of energies converge to a minimizer of a limiting energy. Next, we introduce a variant of Serfaty's scheme of Γ -convergence of gradient flows [67] that is weak enough to accommodate our assumptions on $\bar{\rho}$. In particular, it will allow us to study the limiting behavior of both gradient flows of $\mathcal{F}_{\varepsilon, k}$ in Section 5, as well as “almost” curves of maximal slope in Section 6.

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 $\mathcal{G}_\alpha: \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R} \cup \{+\infty\}$ is said to Γ -converge to $\mathcal{G}: \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R} \cup \{+\infty\}$ if:

$$(2.15) \quad \begin{aligned} &\text{for any sequence } \rho_\alpha \in \mathcal{P}(\mathbb{R}^d) \text{ converging narrowly} \\ &\text{to } \rho \in \mathcal{P}(\mathbb{R}^d), \liminf_{\alpha \rightarrow 0} \mathcal{G}_\alpha(\rho_\alpha) \geq \mathcal{G}(\rho); \end{aligned}$$

$$(2.16) \quad \begin{aligned} &\text{for any } \rho \in \mathcal{P}_2(\mathbb{R}^d), \text{ there exists } \rho_\alpha \in \mathcal{P}(\mathbb{R}^d) \text{ converging} \\ &\text{narrowly to } \rho \text{ s.t. } \limsup_{\alpha \rightarrow 0} \mathcal{G}_\alpha(\rho) \leq \mathcal{G}(\rho). \end{aligned}$$

Next, we prove Lemma 2.15, which provides sufficient conditions for compactness of a sequence of absolutely continuous curves.

Lemma 2.15 (Compactness of absolutely continuous curves). *Fix $T > 0$. Suppose we have a sequence $\{\rho_\alpha\}_{\alpha>0} \subseteq AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ and*

$$(2.17) \quad \sup_{\alpha>0} \int_0^T |\rho'_\alpha|^2(r) dr < +\infty \quad \text{and} \quad \sup_{\alpha>0} M_2(\rho_\alpha(0)) < +\infty.$$

Then there exists $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ such that, along a subsequence $\alpha \rightarrow 0$, $W_1(\rho_\alpha(t), \rho(t)) \rightarrow 0$ uniformly in $t \in [0, T]$, and

$$(2.18) \quad \liminf_{\alpha \rightarrow 0} \int_0^t |\rho'_\alpha|^2(r) dr \geq \int_0^t |\rho'|^2(r) dr \text{ for every } t \in [0, T].$$

Proof. First, we shall produce $\rho : [0, T] \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ and a subsequence ρ_α such that $W_1(\rho_\alpha(t), \rho(t)) \rightarrow 0$ uniformly in $t \in [0, T]$. To this end, we use Proposition A.3, together with hypothesis (2.17) to find that there exists $C = C(T) > 0$, so that, for all $t \in [0, T]$ and $\alpha > 0$, $\rho_\alpha(t)$ belongs to the set $\{\rho : M_2(\rho) \leq C\}$. This set is narrowly sequentially compact [3, Remark 5.1.5, Lemma 5.1.7] and has uniformly integrable 1st moments [3, equation 5.1.20], so it is relatively compact in the 1-Wasserstein metric [3, Proposition 7.1.5]. In particular, $\{\rho_\alpha(t)\}_{\alpha>0}$ is relatively compact with respect to the 1-Wasserstein metric, pointwise in time.

Next, using inequality (2.5) and hypothesis (2.17), we deduce equicontinuity with respect to the 1-Wasserstein metric: for all $0 \leq s \leq t \leq T$,

$$\begin{aligned} \sup_{\alpha>0} W_1(\rho_\alpha(s), \rho_\alpha(t)) &\leq \sup_{\alpha>0} W_2(\rho_\alpha(s), \rho_\alpha(t)) \\ &\leq \sup_{\alpha>0} \int_s^t |\rho'_\alpha|(r) dr \\ &\leq \sqrt{t-s} \left(\sup_{\alpha>0} \int_0^T |\rho'_\alpha|^2(r) \right)^{1/2}. \end{aligned}$$

Therefore, the Ascoli-Arzelá theorem ensures that there exists $\rho : [0, T] \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ so that, up to a subsequence, $W_1(\rho_\alpha(t), \rho(t)) \rightarrow 0$ uniformly in $t \in [0, T]$.

It remains to show $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ and (2.18). To see this, note that hypothesis (2.17) ensures $\{|\rho'_\alpha|(r)\}_{\alpha>0}$ is bounded in $L^2([0, T])$. Thus, up to another subsequence, it is weakly convergent to some $\nu(r) \in L^2([0, T])$. Thus, for all $0 \leq s \leq t \leq T$, using the lower semicontinuity of the 2-Wasserstein metric with respect to narrow (hence 1-Wasserstein) convergence,

$$W_2(\rho(s), \rho(t)) \leq \liminf_{\alpha \rightarrow 0} W_2(\rho_\alpha(s), \rho_\alpha(t)) \leq \liminf_{\alpha \rightarrow 0} \int_s^t |\rho'_\alpha|(r) dr = \int_s^t \nu(r) dr.$$

This shows $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$. Furthermore, by [3, Theorem 1.1.2], it ensures $|\mu'(r)| \leq \nu(r)$ for a.e. $r \in [0, T]$. Thus, by lower semicontinuity of the $L^2([0, T])$ norm with respect to weak convergence, (2.18) holds. $\square$

With the preceding result in hand, we now introduce our variant of Γ -convergence of gradient flows. Our conditions strongly mirror Serfaty's framework [67, Theorem 2], in the context of functionals defined on $(\mathcal{P}_2(\mathbb{R}^d), W_2)$ that are lower semicontinuous and semiconvex along generalized geodesics. The main difference is that we do not identify either $\|\eta_\alpha\|_{L^2(\rho_\alpha)}$ or $\|\eta\|_{L^2(\rho)}$ in Proposition 2.16 as a strong upper gradient of $\mathcal{F}_\alpha$ or $\mathcal{F}$. For this reason, we cannot conclude that either ρ_α or ρ is a gradient flow of the respective energy. Our version of Serfaty's framework

allows us to accommodate our more general assumptions on $\bar{\rho}$, while still ultimately obtaining our main convergence result, Theorem 1.1.

Proposition 2.16. *Let $\mathcal{F}, \mathcal{F}_\alpha : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ be functionals that are proper and bounded from below uniformly in α , and suppose $\mathcal{F}_\alpha$ Γ -converges to $\mathcal{F}$ as $\alpha \rightarrow 0$. Fix $T > 0$. Suppose that, for all $\alpha > 0$, there exists $\rho_\alpha \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ and, for almost all $r \in [0, T]$, there exists $\eta_\alpha(r) \in L^2(\rho_\alpha(r))$, such that*

$$(2.19) \quad \frac{1}{2} \int_0^t |\rho'_\alpha|^2(r) dr + \frac{1}{2} \int_0^t \int_{\mathbb{R}^d} |\eta_\alpha(r)|^2 d\rho_\alpha(r) dr \leq \mathcal{F}_\alpha(\rho_\alpha(0)) - \mathcal{F}_\alpha(\rho_\alpha(t))$$

for all $0 \leq t \leq T$.

Suppose also that there exists $\rho(0) \in D(\mathcal{F}) \cap \mathcal{P}_2(\mathbb{R}^d)$ so that

$$(2.20) \quad \begin{aligned} \rho_\alpha(0) &\xrightarrow{\alpha \rightarrow 0} \rho(0) \text{ narrowly, } \quad \lim_{\alpha \rightarrow 0} \mathcal{F}_\alpha(\rho_\alpha(0)) = \mathcal{F}(\rho(0)), \\ &\text{and } \sup_{\alpha > 0} M_2(\rho_\alpha(0)) < +\infty. \end{aligned}$$

Then, there exists $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ so that, up to a subsequence in α ,

$$(2.21) \quad \lim_{\alpha \rightarrow 0} W_1(\rho_\alpha(r), \rho(r)) = 0 \text{ uniformly for } r \in [0, T].$$

Furthermore, we have

$$(2.22) \quad \frac{1}{2} \int_0^t |\rho'|^2(r) dr + \frac{1}{2} \int_0^t \left(\liminf_{\alpha \rightarrow 0} \int_{\mathbb{R}^d} |\eta_\alpha(r)|^2 d\rho_\alpha(r) \right) dr \leq \mathcal{F}(\rho(0)) - \mathcal{F}(\rho(t))$$

for all $0 \leq t \leq T$.

Proof. By (2.20), we may assume,

$$(2.23) \quad \sup_{\alpha} \mathcal{F}_\alpha(\rho_\alpha(0)) < +\infty.$$

Using this, together with our assumption that $\mathcal{F}_\alpha$ is bounded from below uniformly in α , we see that the right-hand side of (2.19) is bounded from above uniformly in α . From this we deduce $\sup_{\alpha} \int_0^t |\rho'_\alpha|^2(r) dr < +\infty$. Therefore, we may apply Lemma 2.15. We find that there exists $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ such that (2.21) holds along a subsequence $\alpha \rightarrow 0$ and we also have,

$$(2.24) \quad \liminf_{\alpha \rightarrow 0} \int_0^t |\rho'_\alpha|^2(r) dr \geq \int_0^t |\rho'|^2(r) dr \text{ for all } t \in [0, T].$$

Now, taking $\liminf_{\alpha \rightarrow 0}$ of (2.19), we find, for all $t \in [0, T]$,

$$(2.25) \quad \begin{aligned} \frac{1}{2} \int_0^t |\rho'|^2(r) dr + \liminf_{\alpha \rightarrow 0} \frac{1}{2} \int_0^t \int_{\mathbb{R}^d} |\eta_\alpha(r)|^2 d\rho_\alpha(r) dr \\ \leq \liminf_{\alpha \rightarrow 0} (\mathcal{F}(\rho_\alpha(0)) - \mathcal{F}(\rho_\alpha(t))). \end{aligned}$$

The Γ -convergence of $\mathcal{F}_\alpha$ to $\mathcal{F}$ and the narrow convergence of $\rho_\alpha(t)$ to $\rho(t)$ (which follows from the convergence in 1-Wasserstein), as well as the hypothesis (2.20), imply the following upper bound for the right-hand side of the previous line:

$$\mathcal{F}(\rho(0)) - \mathcal{F}(\rho(t)) \geq \limsup_{\alpha \rightarrow 0} \mathcal{F}(\rho_\alpha(0)) - \mathcal{F}(\rho_\alpha(t)) \geq \liminf_{\alpha \rightarrow 0} \mathcal{F}(\rho_\alpha(0)) - \mathcal{F}(\rho_\alpha(t)).$$

Finally, we use Fatou's Lemma to bound the second term on the left-hand side of (2.25) from below. Together with the previous line, this yields the desired estimate (2.22). $\square$

3. GRADIENT FLOWS OF ENERGIES WITH REGULARIZATION AND CONFINEMENT

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 $\mathcal{V}_\Omega$. In particular, we 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 $\mathcal{V}_\Omega$, 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 $\mathcal{V}$, $\mathcal{V}_\varepsilon$, $\mathcal{V}_k$, $\mathcal{V}_\Omega$). *Under Assumptions (M), (D), (V), and (C), the energies $\mathcal{V}$, $\mathcal{V}_\varepsilon$, $\mathcal{V}_k$, and $\mathcal{V}_\Omega$ are lower semicontinuous with respect to narrow convergence.*

The convexity of the energies $\mathcal{E}$, $\mathcal{V}$, $\mathcal{V}_\varepsilon$, $\mathcal{V}_k$, and $\mathcal{V}_\Omega$ follows immediately from the theory developed by Ambrosio, Gigli, and Savaré [3]. We recall these results in the following proposition. The proof of Proposition 3.3 is an immediate consequence of existing theory, so we defer it to Appendix B.

Proposition 3.3 (Convexity properties of $\mathcal{E} + \mathcal{V}_\Omega$, $\mathcal{V}$, $\mathcal{V}_\varepsilon$, and $\mathcal{V}_k$).

- (i) *Suppose $\bar{\rho}$ is log-concave on Ω , where $\bar{\rho}$ satisfies Assumption (T) and Ω satisfies Assumption (D). Then $\mathcal{E} + \mathcal{V}_\Omega$ is convex along generalized geodesics.*
- (ii) *Suppose Assumptions (M), (V), and (C) hold. 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.*

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,*

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

Then,

$$\begin{aligned} \left. \frac{d}{d\alpha} \mathcal{E}_\varepsilon(\mu_\alpha) \right|_{\alpha=0} &= \int \frac{\zeta_\varepsilon * \nu_2(x)}{\bar{\rho}(x)} \int \langle \nabla \zeta_\varepsilon(x - y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3) dx, \\ \left. \frac{d}{d\alpha} \mathcal{V}(\mu_\alpha) \right|_{\alpha=0} &= \int \langle \nabla V(y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3), \\ \left. \frac{d}{d\alpha} \mathcal{V}_\varepsilon(\mu_\alpha) \right|_{\alpha=0} &= \int \langle \nabla(\zeta_\varepsilon * V)(y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3), \\ \left. \frac{d}{d\alpha} \mathcal{V}_k(\mu_\alpha) \right|_{\alpha=0} &= \int \langle \nabla V_k(y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3). \end{aligned}$$

Remark 3.5. 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 (2.9), then γ satisfies the assumptions of Proposition 3.4.

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

$$\begin{aligned} (3.1) \quad \frac{1}{\alpha} (\zeta_\varepsilon * \mu_\alpha(x) - \zeta_\varepsilon * \mu_0(x)) \\ = \int \frac{1}{\alpha} [\zeta_\varepsilon(x - ((1 - \alpha)y_2 + \alpha y_3)) - \zeta_\varepsilon(x - y_2)] d\gamma(y_1, y_2, y_3). \end{aligned}$$

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

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

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,

$$\begin{aligned} (3.3) \quad \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} (\zeta_\varepsilon * \mu_\alpha(x) - \zeta_\varepsilon * \mu_0(x)) \\ = \int \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} [\zeta_\varepsilon(x - ((1 - \alpha)y_2 + \alpha y_3)) - \zeta_\varepsilon(x - y_2)] d\gamma(y_1, y_2, y_3) \\ = \int \langle \nabla \zeta_\varepsilon(x - y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3). \end{aligned}$$

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

$$\begin{aligned}
 (3.4) \quad & \frac{1}{\alpha}(\mathcal{E}_\varepsilon(\mu_\alpha) - \mathcal{E}_\varepsilon(\mu_0)) \\
 &= \frac{1}{2\alpha} \int ((\zeta_\varepsilon * \mu_\alpha)^2(x) - (\zeta_\varepsilon * \mu_0)^2(x)) \bar{\rho}(x)^{-1} dx \\
 &= \int \frac{1}{2\alpha} [(\zeta_\varepsilon * \mu_\alpha)(x) + (\zeta_\varepsilon * \mu_0)(x)] [(\zeta_\varepsilon * \mu_\alpha)(x) - (\zeta_\varepsilon * \mu_0)(x)] \bar{\rho}(x)^{-1} dx.
 \end{aligned}$$

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

$$g_\alpha(x) := C [(\zeta_\varepsilon * \mu_\alpha)(x) + (\zeta_\varepsilon * \mu_0)(x)] \quad \text{for} \quad C = \|\nabla \zeta_\varepsilon\|_\infty \|\bar{\rho}^{-1}\|_\infty M_1(\gamma).$$

The narrow convergence of μ_α to μ_0 as $\alpha \rightarrow 0$, and the fact that ζ_ε is bounded and continuous, ensures that $g_\alpha(x) \rightarrow 2C(\zeta_\varepsilon * \mu_0)(x)$ pointwise. Furthermore,

$$\lim_{\alpha \rightarrow 0} \int g_\alpha(x) dx = \lim_{\alpha \rightarrow 0} C \int \zeta_\varepsilon * \mu_\alpha(x) dx + C \int \zeta_\varepsilon * \mu_0(x) dx = 2C.$$

Therefore, by the generalized dominated convergence theorem [65, Chapter 4, Theorem 19] and equations (3.3) and (3.4),

$$\begin{aligned}
 & \lim_{\alpha \rightarrow 0} \frac{1}{\alpha}(\mathcal{E}_\varepsilon(\mu_\alpha) - \mathcal{E}_\varepsilon(\mu_0)) \\
 &= \int \lim_{\alpha \rightarrow 0} \frac{1}{2\alpha} [(\zeta_\varepsilon * \mu_\alpha)(x) + (\zeta_\varepsilon * \mu_0)(x)] [(\zeta_\varepsilon * \mu_\alpha)(x) - (\zeta_\varepsilon * \mu_0)(x)] \bar{\rho}(x)^{-1} dx \\
 &= \int \frac{(\zeta_\varepsilon * \mu_0)(x)}{\bar{\rho}(x)} \int \langle \nabla \zeta_\varepsilon(x - y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3) dx.
 \end{aligned}$$

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

$$\lim_{\alpha \rightarrow 0} \frac{1}{\alpha}(\mathcal{V}(\mu_\alpha) - \mathcal{V}(\mu_0)) = \lim_{\alpha \rightarrow 0} \int \frac{1}{\alpha} [V((1-\alpha)y_2 + \alpha y_3) - V(y_2)] d\gamma(y_1, y_2, y_3).$$

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

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

Thus, by the dominated convergence theorem,

$$\begin{aligned}
 \lim_{\alpha \rightarrow 0} \frac{1}{\alpha}(\mathcal{V}(\mu_\alpha) - \mathcal{V}(\mu_0)) &= \int \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} [V((1-\alpha)y_2 + \alpha y_3) - V(y_2)] d\gamma(y_1, y_2, y_3) \\
 &= \int \langle \nabla V(y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3),
 \end{aligned}$$

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,

$$\begin{aligned}
 & \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} (\mathcal{V}_k(\mu_\alpha) - \mathcal{V}_k(\mu_0)) \\
 &= \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \int_{\mathbb{R}^d} [V_k((1-\alpha)y_2 + \alpha y_3) - V_k(y_2)] d\gamma(y_1, y_2, y_3), \\
 &= \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} \int_{\mathbb{R}^d} \int_0^\alpha \int_0^\beta (y_3 - y_2)^t D^2 V_k((1-s)y_2 + sy_3)(y_3 - y_2) ds d\beta \\
 &\quad + \alpha \langle \nabla V_k(y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3) \\
 &= \int_{\mathbb{R}^d} \langle \nabla V_k(y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3),
 \end{aligned}$$

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

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 $\lambda_\varepsilon \xrightarrow{\varepsilon \rightarrow 0} -\infty$.

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,*

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

Proof. Let $(\mu_\alpha^{2 \rightarrow 3})_{\alpha \in [0,1]}$ 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 (2.9). Since $x \mapsto x^2$ is a convex function, using the above the tangent inequality for convex functions yields,

$$\begin{aligned}
 \mathcal{E}_\varepsilon(\mu_3) - \mathcal{E}_\varepsilon(\mu_2) &= \frac{1}{2} \int \frac{(\zeta_\varepsilon * \mu_3(x))^2}{\bar{\rho}(x)} dx - \frac{1}{2} \int \frac{(\zeta_\varepsilon * \mu_2(x))^2}{\bar{\rho}(x)} dx \\
 &\geq \int \frac{\zeta_\varepsilon * \mu_2(x)}{\bar{\rho}(x)} (\zeta_\varepsilon * \mu_3(x)) - \zeta_\varepsilon * \mu_2(x) dx \\
 &= \int \frac{\zeta_\varepsilon * \mu_2(x)}{\bar{\rho}(x)} \left(\iiint \zeta_\varepsilon(x - y_3) - \zeta_\varepsilon(x - y_2) \right) d\gamma(y_1, y_2, y_3) dx.
 \end{aligned}$$

Therefore, by Proposition 3.4,

$$\begin{aligned}
 & \mathcal{E}_\varepsilon(\mu_3) - \mathcal{E}_\varepsilon(\mu_2) - \frac{d}{d\alpha} \mathcal{E}_\varepsilon(\mu_\alpha^{2 \rightarrow 3}) \Big|_{\alpha=0} \\
 &\geq \int \frac{\zeta_\varepsilon * \mu_2(x)}{\bar{\rho}(x)} \left(\iiint \zeta_\varepsilon(x - y_3) - \zeta_\varepsilon(x - y_2) \right. \\
 &\quad \left. - \langle \nabla \zeta_\varepsilon(x - y_2), y_2 - y_3 \rangle d\gamma(y_1, y_2, y_3) \right) dx \\
 &\geq -\frac{\|D^2 \zeta_\varepsilon\|_{L^\infty(\mathbb{R}^d)}}{2} \int \frac{\zeta_\varepsilon * \mu_2(x)}{\bar{\rho}(x)} dx \iiint |y_2 - y_3|^2 d\gamma(y_1, y_2, y_3) dx \\
 &\geq -\frac{\|1/\bar{\rho}\|_{L^\infty(\mathbb{R}^d)} \|D^2 \zeta_\varepsilon\|_{L^\infty(\mathbb{R}^d)}}{2} W_{2,\gamma}^2(\mu_2, \mu_3),
 \end{aligned}$$

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. $\square$

The preceding results ensure that our energies $\mathcal{E} + \mathcal{V}_\Omega, \mathcal{E}_\varepsilon, \mathcal{V}, \mathcal{V}_\varepsilon$, and $\mathcal{V}_k$ 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, are 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 Proposition 3.7, 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 of 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 $\nabla \mathcal{V} \in \partial \mathcal{V}(\mu)$. 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 Proposition 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,

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

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

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

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 $\mathcal{V}_k$ follows 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,

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

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

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

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

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

Proposition 3.9 (Minimal subdifferential of $\mathcal{F}$, [3, Theorems 10.4.9-10.4.13]). *Assume (T), (V), (D) hold and that $\bar{\rho}$ is log-concave on Ω . Given $\mu \in D(\mathcal{F})$, we have $|\partial \mathcal{F}|(\mu) < +\infty$ if and only if $(\mu/\bar{\rho})^2 \in W_{\text{loc}}^{1,1}(\Omega)$ and there exists $\xi \in L^2(\mu)$ so that,*

$$\xi \mu = \frac{\bar{\rho}}{2} \nabla(\mu/\bar{\rho})^2 + \nabla V \mu \quad \text{on } \Omega.$$

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

3.2. Differential equation characterization of gradient flows. We close by identifying the differential equations that characterize gradient flows of $\mathcal{F}_{\varepsilon,k}$ 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 $\mathcal{F}$). *Assume (V), (T), (D) hold and $\bar{\rho}$ is log-concave on Ω . For every $\mu_0 \in \overline{D(\mathcal{F})}$, we have that $\mu(t) \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ is the unique Wasserstein gradient flow of $\mathcal{F}$ with initial data μ_0 if and only if $\mu(t)$ satisfies,*

$$(3.7) \quad \begin{cases} \partial_t \mu - \nabla \cdot \left(\frac{\bar{\rho}}{2} \nabla \left(\frac{\mu^2}{\bar{\rho}^2} \right) + \nabla V \mu \right) = 0, & \text{in duality with } C_c^{\infty}(\mathbb{R}^d \times (0, \infty)), \\ \lim_{t \rightarrow 0^+} \mu(t) = \mu_0 & \text{in } W_2, \end{cases}$$

and satisfies,

$$(3.8) \quad \mu(t) \ll \mathcal{L}^d, \mu = 0 \text{ } \mathcal{L}^d\text{-a.e. on } \mathbb{R}^d \setminus \overline{\Omega}, \text{ and } (\mu(t)/\bar{\rho})^2 \in W_{\text{loc}}^{1,1}(\Omega) \text{ for } \mathcal{L}^1\text{-a.e. } t > 0,$$

$$(3.9) \quad \int_{\mathbb{R}^d} |\bar{\rho} \nabla(\mu(t)^2/\bar{\rho}^2)/(2\mu) + \nabla V|^2 d\mu \in L_{\text{loc}}^1(0, \infty).$$

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 (3.7) 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,

$$(3.10) \quad \frac{\bar{\rho}}{2} \partial_{\mathbf{n}} \left(\frac{\mu^2}{\bar{\rho}^2} \right) + \partial_{\mathbf{n}} V \mu = 0 \text{ on } \partial \Omega,$$

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 (3.7).

In [62], 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 [62, Definition 1], then it is a solution of (3.7), hence a gradient flow in the sense defined here.

More recently, Dolbeault et al. [33] and Grillo, Muratori, and Porzio [42] consider well-posedness of (WPME). If u is smooth enough, it is a solution to [42, 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 (3.7). More precisely comparing our notion of solution with [42, 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 $\mu(t) \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ is the unique Wasserstein gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial data μ_0 if and only if $\mu(t)$ satisfies,*

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

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 $\varepsilon > 0$, $N \in \mathbb{N}$, $\{X_0^1, \dots, X_0^N\} \in \mathbb{R}^d$, and $\{m^1, \dots, m^N\} \in \mathbb{R}_+$ satisfying $\sum_{i=1}^N m^i = 1$. Then, there exists a unique continuously differentiable function $X : [0, \infty) \rightarrow \mathbb{R}^{Nd}$, with components $(X^1(t), \dots, X^N(t))$, that satisfies the system,*

$$(3.12) \quad \begin{cases} \dot{X}^i = - \sum_{j=1}^N m^j \int_{\mathbb{R}^d} \nabla \zeta_\varepsilon(X^i - z) \zeta_\varepsilon(z - X^j) \frac{1}{\bar{\rho}(z)} dz - \nabla(\zeta_\varepsilon * V)(X^i) - \nabla V_k(X^i), \\ X^i(0) = X_0^i. \end{cases}$$

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) hold $V = 0$, Ω is bounded, and $\bar{\rho}$ is log-concave on Ω . 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,

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

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

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 $\varepsilon \rightarrow 0$ 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 on $\bar{\rho}$, V , and V_k , 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,*

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

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,

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

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,

$$(4.3) \quad \begin{aligned} & \frac{d}{dt} \mathcal{S}(\rho_{\varepsilon})(t) \\ &= \int_{\mathbb{R}^d} \log(\rho_{\varepsilon}) \partial_t \rho_{\varepsilon} d\mathcal{L}^d \\ &= - \int_{\mathbb{R}^d} \left\langle \nabla \rho_{\varepsilon}, \nabla \zeta_{\varepsilon} * \left(\frac{\zeta_{\varepsilon} * \rho_{\varepsilon}}{\bar{\rho}} \right) \right\rangle + \langle \nabla \rho_{\varepsilon}, \nabla (\zeta_{\varepsilon} * V) \rangle + \langle \nabla \rho_{\varepsilon}, \nabla V_k \rangle d\mathcal{L}^d \\ &= - \int_{\mathbb{R}^d} \frac{|\nabla \zeta_{\varepsilon} * \rho_{\varepsilon}|^2}{\bar{\rho}} + \left\langle \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon}), (\zeta_{\varepsilon} * \rho_{\varepsilon}) \nabla \left(\frac{1}{\bar{\rho}} \right) \right\rangle \\ & \quad + \langle \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon}), \nabla V \rangle - \rho_{\varepsilon} \Delta V_k d\mathcal{L}^d \\ &\leq -C_{\bar{\rho}} \|\nabla \zeta_{\varepsilon} * \rho_{\varepsilon}\|_{L^2(\mathbb{R}^d)}^2 - \int_{\mathbb{R}^d} \left\langle \nabla (\zeta_{\varepsilon} * \rho_{\varepsilon}), (\zeta_{\varepsilon} * \rho_{\varepsilon}) \nabla \left(\frac{1}{\bar{\rho}} \right) + \nabla V \right\rangle - \rho_{\varepsilon} \Delta V_k d\mathcal{L}^d, \end{aligned}$$

where the final inequality holds for some $C_{\bar{\rho}} > 0$ since $\bar{\rho}$ is uniformly bounded below.

We notice that the first term on the right-hand side is, up to a (good) sign, a constant, and an integration in time, exactly what appears on the left-hand side

of our desired inequality (4.1). To control the other two terms, we apply Cauchy's inequality and our hypotheses on $\bar{\rho}$, V and V_k to find,

$$\frac{d}{dt} \mathcal{S}(\rho_\varepsilon)(t) \leq -C_{\bar{\rho}} \|\nabla \zeta_\varepsilon * \rho_\varepsilon\|_{L^2(\mathbb{R}^d)} + \frac{C_{\bar{\rho}}}{2} \|\nabla \zeta_\varepsilon * \rho_\varepsilon\|_{L^2(\mathbb{R}^d)}^2 + C \|\zeta_\varepsilon * \rho_\varepsilon\|_{L^2(\mathbb{R}^d)}^2 + C',$$

where $C, C' > 0$ are constants depending on $\bar{\rho}$, V , and V_k . (We make the dependencies more precise in the actual proof). Integrating in time from 0 to T and rearranging, we obtain

$$\frac{C_{\bar{\rho}}}{2} \int_0^T \|\nabla \zeta_\varepsilon * \rho_\varepsilon\|_{L^2(\mathbb{R}^d)}^2 \leq S(\rho_\varepsilon(0)) - S(\rho_\varepsilon(T)) + C \int_0^T \|\zeta_\varepsilon * \rho_\varepsilon\|_{L^2(\mathbb{R}^d)}^2 + C'.$$

As we describe below, we may use a Carleman-type estimate to bound $S(\rho_\varepsilon(T))$ from below in terms of a dimension dependent constant plus $M_2(\rho_\varepsilon(T))$, which by Corollary A.4 can then be controlled by $M_2(\rho_\varepsilon(0))$, $\mathcal{F}_{\varepsilon,k}(\rho_\varepsilon(0))$, and $T > 0$. Likewise, we may bound $\|\zeta_\varepsilon * \rho_\varepsilon\|_{L^2(\mathbb{R}^d)}^2$ from above by $C'_{\bar{\rho}} \mathcal{F}_{\varepsilon,k}(\rho_\varepsilon(t))$ for $C'_{\bar{\rho}} > 0$, which can then be controlled by $C''_{\bar{\rho}} \mathcal{F}_{\varepsilon,k}(\rho_\varepsilon(0))$, since the energy decreases along the gradient flow. In this way, we obtain 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, Matthes, McCann, and Savaré introduced the *flow interchange method* [54]. This method strongly leverages the heuristic interpretation of the 2-Wasserstein metric as an infinite dimensional Riemannian manifold with inner product $\langle \cdot, \cdot \rangle_{W_2, \rho}$; see [73, Chapter 8.2] for more details on this interpretation, originally discovered by Otto [62]. 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:

$$\begin{aligned} \left. \frac{d}{dt} \mathcal{S}(\rho_\varepsilon) \right|_{t=0} &= \langle \nabla_{W_2} \mathcal{S}(\rho_\varepsilon), \partial_t \rho_\varepsilon \rangle_{W_2, \rho_\varepsilon} \Big|_{t=0} \\ &= - \langle \nabla_{W_2} \mathcal{S}(\rho_\varepsilon), \nabla_{W_2} \mathcal{F}_{\varepsilon,k}(\rho_\varepsilon) \rangle_{W_2, \rho_\varepsilon} \Big|_{t=0} \\ &= - \langle \nabla_{W_2} \mathcal{S}(\mu), \nabla_{W_2} \mathcal{F}_{\varepsilon,k}(\mu) \rangle_{W_2, \mu} \Big|_{t=0} \\ &= \langle \partial_t \mu, \nabla_{W_2} \mathcal{F}_{\varepsilon,k}(\mu) \rangle_{W_2, \mu} \Big|_{t=0} \\ &= \left. \frac{d}{dt} \mathcal{F}_{\varepsilon,k}(\mu) \right|_{t=0}. \end{aligned}$$

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 $\mathcal{S}$ 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 $\mathcal{S}$, then for all $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and for $\mathcal{L}^1$ -a.e. $t \geq 0$,

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

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 $\mathcal{S}(\rho_\varepsilon)$ 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 $J_{\tau,\varepsilon}^n \mu$ denote the n th step of the minimizing movement scheme of $\mathcal{F}_{\varepsilon,k}$ with time step τ and initial data $J_{\tau,\varepsilon}^0 \mu = \mu$.

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

Definition 4.3 (Heat flow semigroup). Given $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $h \geq 0$, 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,

$$(4.5) \quad \zeta_\varepsilon * (S_h(\mu)) = S_h(\zeta_\varepsilon * \mu).$$

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 $S_h(J_{\tau,\varepsilon}^n \mu)$). Assume (T), (M), (V), and (C) hold. Let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. We have,

$$(4.6) \quad \limsup_{h \rightarrow 0^+} \frac{\mathcal{E}_\varepsilon(J_{\tau,\varepsilon}^n \mu) - \mathcal{E}_\varepsilon(S_h(J_{\tau,\varepsilon}^n \mu))}{h} = - \int_{\mathbb{R}^d} \frac{1}{\rho} \Delta(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) (\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) d\mathcal{L}^d,$$

$$(4.7) \quad \limsup_{h \rightarrow 0^+} \frac{\mathcal{V}_\varepsilon(J_{\tau,\varepsilon}^n \mu) - \mathcal{V}_\varepsilon(S_h(J_{\tau,\varepsilon}^n \mu))}{h} = \int_{\mathbb{R}^d} \langle \nabla V, \nabla(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) \rangle d\mathcal{L}^d, \text{ and}$$

$$(4.8) \quad \limsup_{h \rightarrow 0^+} \frac{\mathcal{V}_k(J_{\tau,\varepsilon}^n \mu) - \mathcal{V}_k(S_h(J_{\tau,\varepsilon}^n \mu))}{h} = - \int_{\mathbb{R}^d} \Delta V_k dJ_{\tau,\varepsilon}^n \mu.$$

Our proof of Lemma 4.4 relies on two key facts, which we now recall. First, for any $\nu \in \mathcal{P}_2(\mathbb{R}^d)$,

$$(4.9) \quad \text{the map } h \mapsto S_h \nu \text{ is narrowly continuous;}$$

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 $\mathcal{S}$, is in $AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$, 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 $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\phi \in C_c^1(\mathbb{R}^d)$,

$$(4.10) \quad \int_{\mathbb{R}^d} \phi dS_h \nu - \int_{\mathbb{R}^d} \phi d\nu = - \int_0^h \int_{\mathbb{R}^d} \langle \nabla \phi(y), \nabla S_t \nu(y) \rangle dy dt.$$

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 [3, Lemma 8.1.2]. And, arguing as in [3, Example 11.1.9],

we have $S_h \nu \in W_{\text{loc}}^{1,1}(\mathbb{R}^d)$ for a.e. $h > 0$ and,

$$\begin{aligned} \int_0^T \int_{\mathbb{R}^d} |\nabla S_h \nu| &= \int_0^T \int_{\mathbb{R}^d} \frac{|\nabla S_h \nu|}{S_h \nu} S_h \nu \\ &\leq \left(\int_0^T \int_{\mathbb{R}^d} \frac{|\nabla S_h \nu|^2}{S_h \nu} \right)^{1/2} \left(\int_0^T \int_{\mathbb{R}^d} S_h \nu \right)^{1/2} \\ &= \sqrt{T} \left(\int_0^T \int_{\mathbb{R}^d} \frac{|\nabla S_h \nu|^2}{S_h \nu} \right)^{1/2}, \end{aligned}$$

where the quantity on the right-hand side is finite by equation (11.1.38) of [3].

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,

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

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,

$$(4.12) \quad \zeta_\varepsilon * (J_{\tau,\varepsilon}^n \mu) - S_h(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) = - \int_0^h \frac{d}{dt} S_t(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) dt = - \int_0^h \Delta S_t(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) dt.$$

Using this in (4.11), we obtain,

$$\begin{aligned} &\frac{\mathcal{E}_\varepsilon(J_{\tau,\varepsilon}^n \mu) - \mathcal{E}_\varepsilon(S_h(J_{\tau,\varepsilon}^n \mu))}{h} \\ &= \int_{\mathbb{R}^d} \frac{1}{2\bar{\rho}} \left(\frac{1}{h} \int_0^h -\Delta S_t(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) dt \right) (\zeta_\varepsilon * (J_{\tau,\varepsilon}^n \mu) + \zeta_\varepsilon * (S_h(J_{\tau,\varepsilon}^n \mu))) d\mathcal{L}^d. \end{aligned}$$

Classical elliptic regularity implies that $\|\Delta S_t(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu)\|_{L^\infty(\mathbb{R}^d)} \leq C_{\varepsilon,\tau,n}$ 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 $h \rightarrow 0+$, we find,

$$\limsup_{h \rightarrow 0+} \frac{\mathcal{E}_\varepsilon(J_{\tau,\varepsilon}^n \mu) - \mathcal{E}_\varepsilon(S_h(J_{\tau,\varepsilon}^n \mu))}{h} = - \int_{\mathbb{R}^d} \frac{1}{\bar{\rho}} \Delta(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) (\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) d\mathcal{L}^d.$$

We have again used that $S_t(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu)$ 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,

$$\begin{aligned} \frac{\mathcal{V}_\varepsilon(J_{\tau,\varepsilon}^n \mu) - \mathcal{V}_\varepsilon(S_h(J_{\tau,\varepsilon}^n \mu))}{h} &= \frac{1}{h} \left(\int_{\mathbb{R}^d} (\zeta_\varepsilon * V) dJ_{\tau,\varepsilon}^n \mu - \int_{\mathbb{R}^d} (\zeta_\varepsilon * V) dS_h(J_{\tau,\varepsilon}^n \mu) \right) \\ &= \frac{1}{h} \int_{\mathbb{R}^d} V (\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu - \zeta_\varepsilon * S_h(J_{\tau,\varepsilon}^n \mu)) d\mathcal{L}^d \\ &= \frac{1}{h} \int_{\mathbb{R}^d} V (\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu - S_h(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu)) d\mathcal{L}^d. \end{aligned}$$

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

$$\frac{\mathcal{V}_\varepsilon(J_{\tau,\varepsilon}^n \mu) - \mathcal{V}_\varepsilon(S_h(J_{\tau,\varepsilon}^n \mu))}{h} = - \int_{\mathbb{R}^d} V \frac{1}{h} \int_0^h \Delta S_t(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) dt d\mathcal{L}^d.$$

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

$$\limsup_{h \rightarrow 0+} \frac{\mathcal{V}_\varepsilon(J_{\tau,\varepsilon}^n \mu) - \mathcal{V}_\varepsilon(S_h(J_{\tau,\varepsilon}^n \mu))}{h} = - \int_{\mathbb{R}^d} V \Delta(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) d\mathcal{L}^d.$$

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,

$$\begin{aligned} \frac{\mathcal{V}_k(J_{\tau,\varepsilon}^n \mu) - \mathcal{V}_k(S_h(J_{\tau,\varepsilon}^n \mu))}{h} &= \frac{1}{h} \left(\int_{\mathbb{R}^d} V_k dJ_{\tau,\varepsilon}^n \mu - \int_{\mathbb{R}^d} V_k dS_h(J_{\tau,\varepsilon}^n \mu) \right) \\ &= \int_{\mathbb{R}^d} \frac{1}{h} \int_0^h \langle \nabla V_k(x), \nabla S_t J_{\tau,\varepsilon}^n \mu(x, t) \rangle dt dx \\ &= - \int_{\mathbb{R}^d} \frac{1}{h} \int_0^h \Delta V_k(x) S_t J_{\tau,\varepsilon}^n \mu(x, t) dt dx. \end{aligned}$$

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

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 $\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu$, since $J_{\tau,\varepsilon}^n \mu \in D(\mathcal{E}_\varepsilon)$.

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,*

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

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,

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

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

4.3. Proof of H^1 -type bound. We now apply Lemmas 4.4 and 4.5 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 $\mu \in D(\mathcal{F}_{\varepsilon,k})$,

$$\mathcal{F}_{\varepsilon,k}(J_{\tau,\varepsilon}^n \mu) - \mathcal{F}_{\varepsilon,k}(S_h(J_{\tau,\varepsilon}^n \mu)) \leq \frac{1}{2\tau} [W_2^2(S_h(J_{\tau,\varepsilon}^n \mu), J_{\tau,\varepsilon}^{n-1} \mu) - W_2^2(J_{\tau,\varepsilon}^n \mu, J_{\tau,\varepsilon}^{n-1} \mu)].$$

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

$$(4.13) \quad \limsup_{h \rightarrow 0^+} \frac{\mathcal{F}_{\varepsilon,k}(J_{\tau,\varepsilon}^n \mu) - \mathcal{F}_{\varepsilon,k}(S_h(J_{\tau,\varepsilon}^n \mu))}{h} \leq \frac{1}{2\tau} \frac{d^+}{dh} W_2^2(S_h(J_{\tau,\varepsilon}^n \mu), J_{\tau,\varepsilon}^{n-1} \mu) \Big|_{h=0} \\ \leq \frac{\mathcal{S}(J_{\tau,\varepsilon}^{n-1} \mu) - \mathcal{S}(J_{\tau,\varepsilon}^n \mu)}{\tau}.$$

The quantity on the right hand side will play the role of $-\frac{d}{dt}\mathcal{S}(\rho_\varepsilon)$ in the $\tau \rightarrow 0$ 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,

$$\limsup_{h \rightarrow 0^+} \frac{\mathcal{F}_{\varepsilon,k}(J_{\tau,\varepsilon}^n \mu) - \mathcal{F}_{\varepsilon,k}(S_h(J_{\tau,\varepsilon}^n \mu))}{h} \\ = - \int_{\mathbb{R}^d} \frac{1}{\bar{\rho}} \Delta(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) (\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) d\mathcal{L}^d \\ + \int_{\mathbb{R}^d} \langle \nabla V, \nabla(\zeta_\varepsilon * J_{\tau,\varepsilon}^n \mu) \rangle d\mathcal{L}^d - \int_{\mathbb{R}^d} \Delta V_k dJ_{\tau,\varepsilon}^n \mu.$$

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

$$\frac{\mathcal{S}(J_{\tau,\varepsilon}^0 \mu) - \mathcal{S}(J_{\tau,\varepsilon}^n \mu)}{\tau} = \sum_{i=1}^n \frac{\mathcal{S}(J_{\tau,\varepsilon}^{i-1} \mu) - \mathcal{S}(J_{\tau,\varepsilon}^i \mu)}{\tau} \\ \geq \sum_{i=1}^n - \int_{\mathbb{R}^d} \frac{1}{\bar{\rho}} \Delta(\zeta_\varepsilon * J_{\tau,\varepsilon}^i \mu) (\zeta_\varepsilon * J_{\tau,\varepsilon}^i \mu) d\mathcal{L}^d \\ + \int_{\mathbb{R}^d} \langle \nabla V, \nabla(\zeta_\varepsilon * J_{\tau,\varepsilon}^i \mu) \rangle d\mathcal{L}^d - \int_{\mathbb{R}^d} \Delta V_k dJ_{\tau,\varepsilon}^i \mu.$$

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

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

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 $\rho_\varepsilon(0)$,

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

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

$$\begin{aligned} & \int_{\mathbb{R}^d} -\frac{1}{\bar{\rho}} \Delta(\zeta_\varepsilon * \mu_{\tau,\varepsilon}(s)) (\zeta_\varepsilon * \mu_{\tau,\varepsilon}(s)) \, d\mathcal{L}^d \\ & \geq C_{\bar{\rho}} \|\nabla \zeta_\varepsilon * \mu_{\tau,\varepsilon}(s)\|_{L^2(\mathbb{R}^d)}^2 - C'_{\bar{\rho}} \|\zeta_\varepsilon * \mu_{\tau,\varepsilon}(s)\|_{L^2(\mathbb{R}^d)}^2. \end{aligned}$$

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,

$$\begin{aligned} & \int_0^T \int_{\mathbb{R}^d} -\frac{1}{\bar{\rho}} \Delta(\zeta_\varepsilon * \mu_{\tau,\varepsilon}(s)) (\zeta_\varepsilon * \mu_{\tau,\varepsilon}(s)) \, d\mathcal{L}^d \, ds \\ & \geq C_{\bar{\rho}} \int_0^T \|\nabla \zeta_\varepsilon * \mu_{\tau,\varepsilon}(s)\|_{L^2(\mathbb{R}^d)}^2 \, ds - TC'_{\bar{\rho}} (\mathcal{F}_{\varepsilon,k}(\rho_\varepsilon(0)) + \|V\|_{L^\infty(\mathbb{R}^d)}). \end{aligned}$$

(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,

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

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,

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

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

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

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

Thus, $\nabla(\zeta_\varepsilon * \mu_{\tau,\varepsilon})(s) \rightarrow \nabla(\zeta_\varepsilon * \rho_\varepsilon)(s)$ 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 $n \rightarrow +\infty$ in inequality (4.16) yields,

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

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], $\limsup_{n \rightarrow \infty} -\mathcal{S}(\mu_{\tau,\varepsilon}(T)) \leq -\mathcal{S}(\rho_\varepsilon(T))$. Thus, sending $n \rightarrow +\infty$ on the left hand side of (4.17), we estimate,

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

Next, we use a Carleman-type estimate [23, Lemma 4.1] to bound the entropy from below by a constant plus the second moment and apply Corollary A.4 to bound the second moment,

$$(4.19) \quad \limsup_{n \rightarrow +\infty} \mathcal{S}(\mu_{\tau,\varepsilon}(0)) - \mathcal{S}(\mu_{\tau,\varepsilon}(T)) \\ \leq \mathcal{S}(\rho_\varepsilon(0)) + (2\pi)^{d/2} + M_2(\rho_\varepsilon(T)) \\ \leq \mathcal{S}(\rho_\varepsilon(0)) + (2\pi)^{d/2} + (1 + Te^T)(M_2(\rho_\varepsilon(0)) + 2\mathcal{F}_{\varepsilon,k}(\rho_\varepsilon(0))).$$

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

$$\mathcal{S}(\rho_\varepsilon(0)) + (2\pi)^{d/2} + (1 + Te^T)(M_2(\rho_\varepsilon(0)) + 2\mathcal{F}_{\varepsilon,k}(\rho_\varepsilon(0))) \\ \geq \frac{C_{\bar{\rho}}}{2} \int_0^T \|\nabla \zeta_\varepsilon * \rho_\varepsilon(s)\|_{L^2(\mathbb{R}^d)}^2 ds - T(C'_\rho \mathcal{F}_{\varepsilon,k}(\rho_\varepsilon(0)) + C_V).$$

Rearranging then gives the result. $\square$

5. CONVERGENCE OF THE ENERGIES $\mathcal{F}_{\varepsilon,k}$ AND GRADIENT FLOWS AS $\varepsilon \rightarrow 0$

We now apply the properties of the energy $\mathcal{F}_{\varepsilon,k}$ and its gradient flows developed in the previous sections to study the behavior of minimizers and gradient flows as $\varepsilon \rightarrow 0$ for fixed $k \in \mathbb{N}$. In Subsection 5.1, we begin by proving the Γ -convergence of the energies $\mathcal{F}_{\varepsilon,k}$ to the energy $\mathcal{F}$. Next, in Subsection 5.2, we analyze the convergence of the gradient flows of $\mathcal{F}_{\varepsilon,k}$ as $\varepsilon \rightarrow 0$ with “well-prepared” initial data (bounded entropy and energy). Due to the fact that we only suppose $\bar{\rho}$ is log-concave on Ω , and not on all of $\mathbb{R}^n$, we are not able to conclude that the limit is a gradient flow of $\mathcal{F}_k$, which we recall is defined by

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

Instead we merely conclude it is an “almost” curve of maximal slope of $\mathcal{F}_k$, see Definition 5.2. Nevertheless, this weaker notion is still sufficient for our main convergence result, Theorem 1.1, studying the limits as $\varepsilon \rightarrow 0$, $k \rightarrow +\infty$.

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

Proof. We begin with the proof of (2.15). We first consider the energies $\mathcal{E}_\varepsilon + \mathcal{V}_\varepsilon$. Let ρ_ε narrowly converge to ρ . Lemma 2.3 implies,

$$(5.1) \quad \zeta_\varepsilon * \rho_\varepsilon \text{ narrowly converges to } \rho.$$

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 $\liminf_{\varepsilon \rightarrow 0}$ and using the lower semicontinuity of $\mathcal{E}$ with respect to narrow convergence, as well as (5.1), we obtain,

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

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,

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

This concludes the proof of (2.15) for $\mathcal{E}_\varepsilon + \mathcal{V}_\varepsilon$. Since $\mathcal{V}_k$ is lower semicontinuous, this likewise implies (2.15) holds for $\mathcal{F}_{\varepsilon,k}$.

Now we establish (2.16). Let $\rho \in \mathcal{P}(\mathbb{R}^d)$. Taking $\rho_\varepsilon = \rho$ for all $\varepsilon > 0$ in (5.2), we find that it suffices to prove $\limsup_{\varepsilon \rightarrow 0} \mathcal{E}_\varepsilon(\rho) \leq \mathcal{E}(\rho)$. 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,

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

We shall now prove:

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

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,

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

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)),

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

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,

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

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 $\varepsilon \rightarrow 0$. 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 theorem. $\square$

5.2. Convergence of the gradient flows. We seek to identify the limit of gradient flows of $\mathcal{F}_{\varepsilon,k}$ as $\varepsilon \rightarrow 0$. Heuristically, one may expect that they converge to a weak notion of gradient flow of $\mathcal{F}_k$, but in the absence of a log-concavity assumption on $\bar{\rho}$, the subdifferential of $\mathcal{F}_k$ lacks appropriate regularity for even a weak notion of gradient flow to be well-defined. However, inspired by Serfaty's approach for studying Γ -convergence of gradient flows, we are still able to identify a limit and show that it nearly satisfies the definition of a curve of maximal slope of $\mathcal{F}_k$. In order to simplify our exposition, we will call the limit an “almost” curve of maximal slope of $\mathcal{F}_k$.

Definition 5.2 (“Almost” curve of maximal slope of $\mathcal{F}_k$). A curve

$$\rho_k \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$$

is an “almost” curve of maximal slope of $\mathcal{F}_k$ if it satisfies,

$$\begin{aligned} (5.6) \quad \frac{1}{2} \int_0^t |\rho'_k|^2(r) dr + \frac{1}{2} \int_0^t \int_{\mathbb{R}^d} |\boldsymbol{\eta}_k(r)|^2 d\rho_k(r) dr \\ \leq \mathcal{F}_k(\rho_k(0)) - \mathcal{F}_k(\rho_k(t)) \text{ for all } t \in [0, T], \end{aligned}$$

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

$$(5.7) \quad \rho_k^2(t) \in W^{1,1}(\mathbb{R}^d) \text{ and } \boldsymbol{\eta}_k(t) \in L^2(\rho_k(t)) \text{ satisfies} \\ \boldsymbol{\eta}_k \rho_k = \frac{\bar{\rho}}{2} \nabla \left(\frac{\rho_k^2}{\bar{\rho}^2} \right) + \rho_k \nabla (V + V_k).$$

We emphasize that if $\bar{\rho}$ were log-concave on all of $\mathbb{R}^d$, $\rho_k \mapsto \int_{\mathbb{R}^d} |\boldsymbol{\eta}_k|^2 d\rho_k$ would be a strong upper gradient for $\mathcal{F}_k$ and any ρ_k satisfying Definition 5.2 would be a true curve of maximal slope of $\mathcal{F}_k$.

Our approach proceeds as follows. Inspired by Serfaty's framework for Γ -convergence of gradient flows, in Subsection 5.2.1, we first prove Proposition 5.3, which gives a weak notion of lower semicontinuity for the metric slopes along a sequence of gradient flows $\rho_\varepsilon(t)$: we show,

$$(5.8) \quad \liminf_{\varepsilon \rightarrow 0} |\partial \mathcal{F}_{\varepsilon,k}|^2(\rho_\varepsilon(t)) \geq \|\boldsymbol{\eta}_k(t)\|_{L^2(\rho_k(t))}^2,$$

where η_k is as in Definition 5.2. Next, in Subsection 5.2.2, we apply this to prove Proposition 5.6 on convergence of the gradient flows for “well-prepared” initial data. It is in this result that we employ the key estimate that we established in Theorem 4.1.

5.2.1. Limit of metric slopes. We begin by identifying sufficient conditions under which the limiting behavior of the metric slopes (5.8) holds.

Proposition 5.3 (Limiting behavior of metric slopes). *Assume (T), (M), (V), and (C) hold. Fix $k \in \mathbb{N}$. Consider a sequence ρ_ε in $\mathcal{P}_2(\mathbb{R}^d)$ satisfying,*

$$(5.9) \quad \sup_{\varepsilon > 0} \mathcal{F}_{\varepsilon, k}(\rho_\varepsilon) < +\infty,$$

$$(5.10) \quad \liminf_{\varepsilon \rightarrow 0} \|\nabla \zeta_\varepsilon * \rho_\varepsilon\|_{L^2(\mathbb{R}^d)} < +\infty, \text{ and}$$

$$(5.11) \quad \liminf_{\varepsilon \rightarrow 0} \int \left| \nabla \zeta_\varepsilon * \left(\frac{\zeta_\varepsilon * \rho_\varepsilon}{\bar{\rho}} \right) \right|^2 d\rho_\varepsilon < +\infty.$$

In addition, suppose there exists $\rho_k \in \mathcal{P}(\mathbb{R}^d)$ such that ρ_ε narrowly converges to ρ_k . Then $\rho_k^2 \in W^{1,1}(\mathbb{R}^d)$, and there exists $\eta_k \in L^2(\rho_k)$, with,

$$(5.12) \quad \eta_k \rho_k = \frac{\bar{\rho}}{2} \nabla \left(\frac{\rho_k^2}{\bar{\rho}^2} \right) + \rho_k \nabla (V + V_k),$$

and such that,

$$(5.13) \quad \liminf_{\varepsilon \rightarrow 0} \int \left| \nabla \zeta_\varepsilon * \left(\frac{\zeta_\varepsilon * \rho_\varepsilon}{\bar{\rho}} \right) + \nabla(\zeta_\varepsilon * V) + \nabla V_k \right|^2 d\rho_\varepsilon \geq \int |\eta_k|^2 d\rho_k.$$

The remainder of this subsection is devoted to the proof of Proposition 5.3. We begin with a preliminary lemma, showing that, under the assumptions of Proposition 5.3, we may upgrade the convergence of $\zeta_\varepsilon * \rho_\varepsilon$ to ρ from narrow convergence to convergence in $L^2_{\text{loc}}(\mathbb{R}^d)$.

Lemma 5.4 (Upgraded convergence of $\zeta_\varepsilon * \rho_\varepsilon$). *Assume (T), (M), (V), and (C) hold. Fix $k \in \mathbb{N}$. Consider any sequence ρ_ε in $\mathcal{P}(\mathbb{R}^d)$ and $\rho_k \in \mathcal{P}(\mathbb{R}^d)$ such that ρ_ε narrowly converges to ρ_k and (5.9) and (5.10) are satisfied. Then $\rho_k \in L^2(\mathbb{R}^d)$, and there exists a subsequence (still denoted ρ_ε) along which we have,*

$$(5.14) \quad \sup_{\varepsilon > 0} \|\zeta_\varepsilon * \rho_\varepsilon\|_{H^1(\mathbb{R}^d)} < +\infty \text{ and}$$

$$(5.15) \quad \zeta_\varepsilon * \rho_\varepsilon \text{ converges to } \rho_k \text{ in } L^2_{\text{loc}}(\mathbb{R}^d).$$

Proof of Lemma 5.4. By assumption (5.9) and the definition of $\mathcal{F}_{\varepsilon, k}$, we find,

$$(5.16) \quad \begin{aligned} +\infty &> \sup_{\varepsilon > 0} \mathcal{F}_{\varepsilon, k}(\rho_\varepsilon) + \|V\|_{L^\infty(\mathbb{R}^d)} \\ &\geq \sup_{\varepsilon > 0} \mathcal{E}_\varepsilon(\rho_\varepsilon) \\ &= \sup_{\varepsilon > 0} \frac{1}{2} \int_{\mathbb{R}^d} \frac{|\zeta_\varepsilon * \rho|^2}{\bar{\rho}} \\ &\geq \frac{1}{2\|\bar{\rho}\|_{L^\infty(\mathbb{R}^d)}} \sup_{\varepsilon > 0} \int_{\mathbb{R}^d} |\zeta_\varepsilon * \rho|^2. \end{aligned}$$

Similarly, since Theorem 5.1 ensures the Γ -convergence of $\mathcal{F}_{\varepsilon,k}$ to $\mathcal{F}_k$, statement (2.16) in Definition 2.14 of Γ -convergence ensures,

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

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

Combining assumption (5.10) with the estimate (5.16) we find that, up to a subsequence, (5.14) 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_k . The uniqueness of limits therefore implies (5.15). $\square$

A key step in studying the limiting behavior of the metric slopes of $\mathcal{F}_{\varepsilon,k}$, as in Proposition 5.3, is to identify the weak limit of $\nabla \zeta_\varepsilon * \left(\frac{1}{\rho}(\zeta_\varepsilon * \rho_\varepsilon)\right)$ in $L^1(\rho_\varepsilon)$. With this weak limit in hand, the desired result 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 Lemma 5.5, we characterize the weak limit.

Lemma 5.5 (Weak limit of subdifferentials). *Assume (T), (M), (V), and (C) hold. Fix $k \in \mathbb{N}$. Consider any sequence ρ_ε in $\mathcal{P}(\mathbb{R}^d)$ and $\rho_k \in \mathcal{P}(\mathbb{R}^d)$ such that ρ_ε narrowly converges to ρ_k and (5.9), (5.10), and (5.11) are satisfied. For all $\varepsilon > 0$ and $f \in C_c^\infty(\mathbb{R}^d)$, define,*

$$\begin{aligned} (5.17) \quad L_\varepsilon(f) &= \int_{\mathbb{R}^d} f \left(\nabla \zeta_\varepsilon * \left(\frac{1}{\rho}(\zeta_\varepsilon * \rho_\varepsilon) \right) \right) d\rho_\varepsilon \\ \text{and } L(f) &= \int_{\mathbb{R}^d} -\frac{1}{2} \nabla \left(\frac{f}{\rho} \right) \rho_k^2 dx + \int_{\mathbb{R}^d} f \rho_k^2 \nabla \left(\frac{1}{\rho} \right) dx. \end{aligned}$$

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

$$(5.18) \quad \lim_{\varepsilon \rightarrow 0} L_\varepsilon(f) = L(f).$$

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

Proof. By Lemma 5.4, we may choose a subsequence, still denoted ρ_ε , along which (5.14) and (5.15) hold.

In order to characterize $\lim_{\varepsilon \rightarrow 0} L_\varepsilon(f)$, 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 $f \in C_c^\infty(\mathbb{R}^d)$,

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

We begin by showing,

$$(5.20) \quad \lim_{\varepsilon \rightarrow 0} J_\varepsilon(f) = \int_{\mathbb{R}^d} f \rho_k^2 \nabla \left(\frac{1}{\bar{\rho}} \right) d\mathcal{L}^d.$$

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,

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

By (5.14) of Lemma 5.4, the right-hand side converges to 0 as $\varepsilon \rightarrow 0$, which implies that (5.20) holds.

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

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

Note that the L^2 convergence of $\zeta_\varepsilon * \rho_\varepsilon$ to ρ_k established in (5.15) of Lemma 5.4 ensures that, for any $f \in C_c^\infty(\mathbb{R}^d)$,

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

Thus, to complete our proof that $\lim_{\varepsilon \rightarrow 0} L_\varepsilon(f) = L(f)$, it suffices to prove that, for any $f \in C_c^\infty(\mathbb{R}^d)$,

$$(5.21) \quad \lim_{\varepsilon \rightarrow 0} |I_\varepsilon(f) - \tilde{I}_\varepsilon(f)| = 0.$$

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

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

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,

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

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

$$(5.22) \quad |I_\varepsilon(f) - \tilde{I}_\varepsilon(f)| \leq \frac{\|\nabla f\|_\infty}{\inf \bar{\rho}} \iint |x-y| \rho_\varepsilon(y)\zeta_\varepsilon(x-y) |\nabla(\zeta_\varepsilon * \rho_\varepsilon)(x)| dy dx.$$

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

$$(5.23) \quad \zeta_\varepsilon(x-y)|x-y| \leq C\varepsilon^\delta \quad \text{for } |x-y| > \varepsilon^\gamma.$$

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,

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

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 (5.23) now follows by taking $\delta = 1 + \delta'/2$.

Thus, breaking up the integral on the right-hand side of (5.22) into two regions and using (5.23), we find,

$$\begin{aligned} |I_\varepsilon(f) - \tilde{I}_\varepsilon(f)| &\leq \frac{\|\nabla f\|_\infty}{\inf \bar{\rho}} \left(\varepsilon^\gamma \iint_{|x-y| < \varepsilon^\gamma} \rho_\varepsilon(y)\zeta_\varepsilon(x-y) |\nabla(\zeta_\varepsilon * \rho_\varepsilon)(x)| dy dx \right. \\ &\quad \left. + C\varepsilon^\delta \iint_{|x-y| > \varepsilon^\gamma} \rho_\varepsilon(y) |\nabla(\zeta_\varepsilon * \rho_\varepsilon)(x)| dy dx \right) \\ &\leq \frac{\|\nabla f\|_\infty}{\inf \bar{\rho}} \left(\varepsilon^\gamma \int (\rho_\varepsilon * \zeta_\varepsilon)(x) |\nabla(\zeta_\varepsilon * \rho_\varepsilon)(x)| dx + C\varepsilon^\delta \int |(\nabla \zeta_\varepsilon * \rho_\varepsilon)(x)| dx \right). \end{aligned}$$

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,

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

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 (5.14) 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 $\lim_{\varepsilon \rightarrow 0} |I_\varepsilon(f) - \tilde{I}_\varepsilon(f)| = 0$, which completes the proof that $\lim_{\varepsilon \rightarrow 0} L_\varepsilon(f) = L(f)$.

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

$$|L(f)| \leq C\|f\|_{L^2(\rho_k)}.$$

Indeed, since $\rho_k \in L^2(\mathbb{R}^d)$, ρ_k is a Radon measure, so $C_c^1(\mathbb{R}^d)$ is dense in $L^2(\rho_k)$ [8, Corollary 4.2.2], and there exists a unique extension of L to $L^2(\rho_k)$ enjoying the same bound.

Fix arbitrary $f \in C_c^\infty(\mathbb{R}^d)$. By definition of L_ε in equation (5.17) and Hölder's inequality,

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

Thus, by assumption (5.11), there exists $C > 0$ so that,

$$|L(f)| = \liminf_{\varepsilon \rightarrow 0} |L_\varepsilon(f)| \leq C \liminf_{\varepsilon \rightarrow 0} \|f\|_{L^2(\rho_\varepsilon)} = C \|f\|_{L^2(\rho_k)},$$

which gives the result. $\square$

We now apply Lemmas 5.4 and 5.5 to prove our result on the limit of the metric slopes.

Proof of Proposition 5.3. Choose a subsequence, still denoted by ρ_ε , so that,

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

It suffices to show $\rho_k^2 \in W^{1,1}(\mathbb{R}^d)$, there exists $\eta_k \in L^2(\rho)$ satisfying (5.12), and, up to a further subsequence,

$$(5.24) \quad \lim_{\varepsilon \rightarrow 0} \int_{\mathbb{R}^d} f \left(\nabla \zeta_\varepsilon * \left(\frac{1}{\bar{\rho}} (\zeta_\varepsilon * \rho_\varepsilon) \right) + \nabla(\zeta_\varepsilon * V) + \nabla V_k \right) d\rho_\varepsilon \\ = \int f \eta_k d\rho_k \quad \text{for all } f \in C_c^\infty(\mathbb{R}^d).$$

The estimate (5.13) 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,

$$(5.25) \quad \lim_{\varepsilon \rightarrow 0} \int_{\mathbb{R}^d} f(\nabla(\zeta_\varepsilon * V)) d\rho_\varepsilon = \lim_{\varepsilon \rightarrow 0} \int_{\mathbb{R}^d} \nabla V(\zeta_\varepsilon * (f\rho_\varepsilon)) d\mathcal{L}^d = \int f \nabla V d\rho_k,$$

$$(5.26) \quad \lim_{\varepsilon \rightarrow 0} \int_{\mathbb{R}^d} f(\nabla V_k) d\rho_\varepsilon = \int f \nabla V_k d\rho_k.$$

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 (5.18) of Lemma 5.5. Combining these with the Riesz Representation Theorem on $L^2(\rho_k)$ (which we can apply to the operator L due to, again, Lemma 5.5), we find that there exists $\tilde{\eta} \in L^2(\rho_k)$ such that,

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

Rearranging, we obtain,

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

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

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

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

$$\begin{aligned}\nabla\left(\frac{\rho_k^2}{\bar{\rho}^2}\right) &= \nabla(\rho_k^2)\frac{1}{\bar{\rho}^2} + \rho_k^2\nabla\left(\frac{1}{\bar{\rho}^2}\right) \\ &= \frac{1}{\bar{\rho}^2}\left(2\tilde{\eta}\rho_k\bar{\rho} - 2\bar{\rho}\rho_k^2\nabla\left(\frac{1}{\bar{\rho}}\right)\right) + \rho_k^2\nabla\left(\frac{1}{\bar{\rho}^2}\right) \\ &= 2\tilde{\eta}\frac{\rho_k}{\bar{\rho}}.\end{aligned}$$

Thus,

$$(5.27) \quad \tilde{\eta}\rho_k = \frac{\bar{\rho}}{2}\nabla\left(\frac{\rho_k^2}{\bar{\rho}^2}\right).$$

Finally, defining $\boldsymbol{\eta}_k = \tilde{\eta} + \nabla V + \nabla V_k$, the facts that $\nabla V \in L^\infty(\mathbb{R}^d)$ and $\nabla V_k \in L^2(\rho_k)$ (see sentence following Assumption (C)), ensure $\boldsymbol{\eta}_k \in L^2(\rho_k)$ and (5.24) holds. $\square$

5.2.2. Convergence of gradient flows. We now apply the result on the limiting behavior of the metric slopes, obtained in Proposition 5.3, as well as the Γ -convergence of the energies, obtained in Theorem 5.1, to show that gradient flows of $\mathcal{F}_{\varepsilon,k}$ with “well-prepared” initial data converge to an “almost curve of maximal slope” of $\mathcal{F}_k$. We emphasize that this result does not require a log-concavity assumption on $\bar{\rho}$.

Proposition 5.6. *Assume (M), (V), (C), (T) hold. 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,*

$$(5.28) \quad \sup_{\varepsilon > 0} \mathcal{S}(\rho_{\varepsilon,k}(0)) < \infty \quad \text{and} \quad \sup_{\varepsilon > 0} M_2(\rho_{\varepsilon,k}(0)) < \infty.$$

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

$$(5.29) \quad \rho_{\varepsilon,k}(0) \xrightarrow{\varepsilon \rightarrow 0} \rho_k(0) \text{ narrowly} \quad \text{and} \quad \lim_{\varepsilon \rightarrow 0} \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}(0)) = \mathcal{F}_k(\rho_k(0)).$$

Then, there exists $\rho_k \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ that is an “almost” curve of maximal slope of $\mathcal{F}_k$, in the sense of Definition 5.2, and a subsequence $\varepsilon_n^{(k)}$, depending on k , so that

$$(5.30) \quad \lim_{n \rightarrow +\infty} W_1(\rho_{\varepsilon_n^{(k)},k}(t), \rho_k(t)) = 0 \text{ uniformly for } t \in [0, T].$$

Proof of Proposition 5.6. By Theorem 2.13, $\rho_{\varepsilon,k}$ is a curve of maximal slope of $\mathcal{F}_{\varepsilon,k}$, so

$$(5.31) \quad \begin{aligned} \frac{1}{2} \int_0^t |\rho'_{\varepsilon,k}|^2(r) dr + \frac{1}{2} \int_0^t |\partial \mathcal{F}_{\varepsilon,k}|^2(\rho_{\varepsilon,k}(r)) dr \\ \leq \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}(0)) - \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}(t)), \quad \text{for all } 0 \leq t \leq T. \end{aligned}$$

We seek to apply Proposition 2.16. Theorem 5.1 ensures that $\mathcal{F}_{\varepsilon,k}$ Γ -converges to $\mathcal{F}_k$. Next, we note that, the previous line, together with the explicit characterization of $|\partial \mathcal{F}_{\varepsilon,k}|^2$ given in Proposition 3.8, yields that the hypothesis (2.19) holds with $\boldsymbol{\eta}_{\varepsilon,k}(r) \in L^2(\rho_{\varepsilon,k}(r))$ given by,

$$\boldsymbol{\eta}_{\varepsilon,k}(r) = \nabla \zeta_\varepsilon * \left(\frac{1}{\bar{\rho}} (\zeta_\varepsilon * \rho_{\varepsilon,k}(r)) \right) + \nabla(\zeta_\varepsilon * V) + \nabla V_k.$$

In addition, the hypotheses of the present proposition guarantee that (2.20) holds. Thus, Proposition 2.16 ensures that there exist $\rho_k \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ and a subsequence $\varepsilon_n^{(k)}$, depending on k , so that (5.30) holds and with, for all $t \in [0, T]$,

$$(5.32) \quad \frac{1}{2} \int_0^t |\rho'_k|^2(r) dr + \frac{1}{2} \int_0^t \left(\liminf_{n \rightarrow \infty} \int_{\mathbb{R}^d} |\boldsymbol{\eta}_{\varepsilon_n^{(k)}}(r)|^2 d\rho_{\varepsilon_n^{(k)}, k}(r) \right) dr \\ \leq \mathcal{F}_k(\rho_k(0)) - \mathcal{F}_k(\rho_k(t)).$$

In order to conclude, it suffices to establish that ρ_k satisfies the conditions of Definition 5.2, that is, for almost every $r \in [0, T]$, we have:

$$(5.33) \quad \rho_k(r)^2 \in W^{1,1}(\mathbb{R}^d),$$

$$(5.34) \quad \text{there exists } \boldsymbol{\eta}_k(r) \in L^2(\rho_k(R)) \text{ satisfying (5.7), and}$$

$$(5.35) \quad \liminf_{n \rightarrow \infty} \int_{\mathbb{R}^d} |\boldsymbol{\eta}_{\varepsilon_n^{(k)}}(r)|^2 d\rho_{\varepsilon_n^{(k)}, k}(r) \geq \int_{\mathbb{R}^d} |\boldsymbol{\eta}_k(r)|^2 d\rho_k(r).$$

Note that we may assume

$$(5.36) \quad \sup_{n \in \mathbb{N}} \mathcal{F}_{\varepsilon_n^{(k)}, k}(\rho_{\varepsilon_n^{(k)}, k}(0)) < +\infty.$$

Combining this with Theorem 4.1 and assumption (5.28) of the present proposition, we obtain,

$$\liminf_{n \rightarrow \infty} \int_0^T \|\nabla \zeta_{\varepsilon_n^{(k)}} * \rho_{\varepsilon_n^{(k)}, k}(r)\|_{L^2(\mathbb{R}^d)}^2 dr < +\infty.$$

Thus, by Fatou's lemma, for almost every $r \in [0, T]$, the above integrand must be finite. Likewise, inequality (5.32) ensures the left-hand side of (5.35) is finite for a.e. $r \in [0, T]$.

We seek to apply Proposition 5.3. Fix $r \in [0, T]$ such that

$$(5.37) \quad \liminf_{n \rightarrow \infty} \int_{\mathbb{R}^d} |\boldsymbol{\eta}_{\varepsilon_n^{(k)}}(r)|^2 d\rho_{\varepsilon_n^{(k)}, k}(r) < +\infty \\ \text{and } \liminf_{n \rightarrow \infty} \|\nabla \zeta_{\varepsilon_n^{(k)}} * \rho_{\varepsilon_n^{(k)}, k}(r)\|_{L^2(\mathbb{R}^d)} < +\infty.$$

Inequality (5.36) together with the fact that the energy $\mathcal{F}_{\varepsilon, k}$ decreases in time along the gradient flow $\rho_{\varepsilon, k}$ implies that (5.9) holds at time r . Thus, by our standing hypotheses on V and V_k , the hypotheses of Proposition 5.3 hold at time r . Consequently, the conclusion of Proposition 5.3 yields (5.33), (5.34), and (5.35) at time r . $\square$

6. CONVERGENCE OF ENERGIES $\mathcal{F}_k$ AND “ALMOST” CURVES OF MAXIMAL SLOPE AS $k \rightarrow +\infty$

The present section has two main goals. First, we show that the energies $\mathcal{F}_k$ Γ -converge to the energy $\mathcal{F}$ and use this to prove Theorem 1.5 that minimizers of $\mathcal{F}_{\varepsilon, k}$ converge to the unique minimizer of $\mathcal{F}$ as $\varepsilon \rightarrow 0$, $k \rightarrow +\infty$. Our second goal is to show that if $\bar{\rho}$ is log-concave on Ω , then as the confining potentials V_k approximate V_Ω , the “almost” curves of maximal slope of $\mathcal{F}_k$ (see Definition 5.2) converge to a gradient flow of $\mathcal{F}$ as $k \rightarrow +\infty$. We then use this to conclude our main result, Theorem 1.1, that gradient flows of $\mathcal{F}_{\varepsilon, k}$ converge to a gradient flow of $\mathcal{F}$ as $k \rightarrow +\infty$ and $\varepsilon = \varepsilon(k) \rightarrow 0$. Finally, in Section 6.1, we extend our result to cover particle initial data (Theorem 1.2), long time behavior (Corollary 1.3), and establish our results concerning two-layer neural networks (Corollary 1.4).

Our result on the $k \rightarrow +\infty$ limit generalizes work by Alasio, Bruna, and Carrillo [1] to the case of weighted porous medium equations. As in Proposition 5.6, which considered the $\varepsilon \rightarrow 0$ 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 $\mathcal{V}_k$ to $\mathcal{V}_\Omega$, in the sense of Definition 2.14.

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

Proof. We first establish item (2.15) for the energies $\mathcal{V}_k$ and $\mathcal{V}_\Omega$. Without loss of generality, we may assume $\liminf_{k \rightarrow +\infty} \mathcal{V}_k(\rho_k) < +\infty$ so, up to a subsequence,

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

To show inequality (2.15), it suffices to prove that $\text{supp } \rho \subseteq \overline{\Omega}$, since $\mathcal{V}_k(\rho_k)$ is nonnegative and $\mathcal{V}_\Omega(\rho)$ would equal zero. Suppose for the sake of contradiction that $\text{supp } \rho \not\subseteq \overline{\Omega}$, so that there exist $x \in \overline{\Omega}^c$ 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 $\rho_k \rightarrow \rho$ narrowly ensures $\liminf_{k \rightarrow +\infty} \rho_k(B) \geq \rho(B) > 0$. Thus, up to taking another subsequence, we may assume that there exists $\delta > 0$ so that $\rho_k(B) \geq \delta$ for all $k \in \mathbb{N}$. By definition of $\mathcal{V}_k$, this implies,

$$\begin{aligned} \liminf_{k \rightarrow +\infty} \int_{\mathbb{R}^d} V_k d\rho_k &\geq \liminf_{k \rightarrow +\infty} \int_B V_k d\rho_k \\ &\geq \liminf_{k \rightarrow +\infty} \left(\inf_{x \in B} V_k(x) \right) \rho_k(B) \\ &\geq \delta \liminf_{k \rightarrow +\infty} \left(\inf_{x \in B} V_k(x) \right) \\ &= +\infty, \end{aligned}$$

where the last inequality follows from Assumption (Ck) on V_k . This contradicts (6.1). Thus, we must have $\text{supp } \rho \subseteq \overline{\Omega}$, which completes the proof of item (2.15). Note that, since $\mathcal{E}$ and $\mathcal{V}$ are lower semicontinuous and bounded below, it follows immediately that inequality (2.15) holds for the energies $\mathcal{F}_k$ and $\mathcal{F}$.

It remains to prove item (2.16). 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,

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

and thus conclude by recalling the definitions of $\mathcal{V}_k$ and $\mathcal{V}$. Likewise, we also obtain (2.16) for $\mathcal{F}_k$ and $\mathcal{F}$. $\square$

As a corollary of Theorems 5.1 and 6.1, we obtain the result of Theorem 1.5: minimizers of $\mathcal{F}_{\varepsilon,k}$ converge to a minimizer of $\mathcal{F}$. The additional assumptions we add – that V_k are all greater than V_1 and the sublevel sets of V_1 are compact – are natural in the context of taking the V_k 's to be diverging to $+\infty$ off of $\overline{\Omega}$.

Proof of Theorem 1.5. First, we show that $\mathcal{F}$ has a unique minimizer. Suppose that ρ_0 and ρ_1 are both minimizers of $\mathcal{F}$. Since $\mathcal{F}$ is proper, we have $\mathcal{F}(\rho_0), \mathcal{F}(\rho_1) < +\infty$, so $\text{supp } \rho_1, \text{supp } \rho_2 \subseteq \overline{\Omega}$. Thus,

$$\begin{aligned}\mathcal{F}((1-\alpha)\rho_0 + \alpha\rho_1) &= \frac{1}{2} \int_{\Omega} \frac{|(1-\alpha)\rho_0 + \alpha\rho_1|^2}{\bar{\rho}} + \int_{\Omega} Vd((1-\alpha)\rho_0 + \alpha\rho_1) \\ &= (1-\alpha)\mathcal{F}(\rho_0) + \alpha\mathcal{F}(\rho_1) - \alpha(1-\alpha) \int_{\Omega} \frac{|\rho_0 - \rho_1|^2}{\bar{\rho}}.\end{aligned}$$

Since $\bar{\rho}$ is uniformly bounded above, we must have $\rho_0 = \rho_1$, else $\mathcal{F}((1-\alpha)\rho_0 + \alpha\rho_1) < \mathcal{F}(\rho_0)$ for $\alpha \in (0, 1)$, contradicting the choice of ρ_0 as a minimizer.

Now, we show minimizers of $\mathcal{F}_{\varepsilon,k}$ converge to the unique minimizer of $\mathcal{F}$. Again, using that $\mathcal{F}$ is proper, take $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ such that $\mathcal{F}(\nu) < +\infty$, so in particular, $\mathcal{V}_{\Omega}(\nu) = 0 = \mathcal{V}_k(\nu)$. Since $\rho_{\varepsilon,k}$ minimizes $\mathcal{F}_{\varepsilon,k}$, using the fact from Theorem 5.1 that $\mathcal{E}_{\varepsilon} + \mathcal{V}_{\varepsilon}$ Γ -converges to $\mathcal{E} + \mathcal{V}$, we have

$$(6.2) \quad \limsup_{\varepsilon \rightarrow 0, k \rightarrow +\infty} \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}) \leq \limsup_{\varepsilon \rightarrow 0, k \rightarrow +\infty} \mathcal{F}_{\varepsilon,k}(\nu) = \limsup_{\varepsilon \rightarrow 0, k \rightarrow +\infty} \mathcal{E}_{\varepsilon}(\nu) + \mathcal{V}_{\varepsilon}(\nu) + 0 \\ = \mathcal{E}(\nu) + \mathcal{V}(\nu) + 0 = \mathcal{F}(\nu) < +\infty.$$

Thus, we may assume that, up to a subsequence, $\mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k})$ is uniformly bounded above in $\varepsilon > 0$, $k \in \mathbb{N}$. Since $\mathcal{E}_{\varepsilon}$ and $\mathcal{V}_{\varepsilon}$ are bounded below uniformly in $\varepsilon > 0$, $\mathcal{V}_k(\rho_{\varepsilon,k})$ must be bounded above uniformly in $\varepsilon > 0$ and $k \in \mathbb{N}$. Next, the assumption $V_k \geq V_1$ implies,

$$\sup_{\varepsilon > 0, k \in \mathbb{N}} \int V_1 d\rho_{\varepsilon,k} \leq \sup_{\varepsilon > 0, k \in \mathbb{N}} \int V_k d\rho_{\varepsilon,k} \leq \sup_{\varepsilon > 0, k \in \mathbb{N}} \mathcal{V}_k(\rho_{\varepsilon,k}) < +\infty.$$

Together with the fact that the sublevel sets of V_1 are compact, this guarantees that the sequence $\rho_{\varepsilon,k}$ is tight; see [3, Remark 5.1.5]. Thus, up to another subsequence, there exists $\rho \in \mathcal{P}(\mathbb{R}^d)$ so that $\rho_{\varepsilon,k} \rightarrow \rho$. By inequality (6.2), Theorem 5.1, and Theorem 6.1, we have, for any $\nu \in D(\mathcal{F})$,

$$\begin{aligned}\mathcal{F}(\nu) &= \mathcal{E}(\nu) + \mathcal{V}(\nu) + \mathcal{V}_{\Omega}(\nu) \\ &= \lim_{\varepsilon \rightarrow 0, k \rightarrow +\infty} \mathcal{E}_{\varepsilon}(\nu) + \mathcal{V}_{\varepsilon}(\nu) + \mathcal{V}_k(\nu) \\ &= \lim_{\varepsilon \rightarrow 0, k \rightarrow +\infty} \mathcal{F}_{\varepsilon,k}(\nu) \\ &\geq \liminf_{\varepsilon \rightarrow 0, k \rightarrow +\infty} \mathcal{F}_{\varepsilon,k}(\rho_{\varepsilon,k}) \\ &\geq \liminf_{\varepsilon \rightarrow 0, k \rightarrow +\infty} \left(\mathcal{E}_{\varepsilon}(\rho_{\varepsilon,k}) + \mathcal{V}_{\varepsilon}(\rho_{\varepsilon,k}) \right) + \liminf_{\varepsilon \rightarrow 0, k \rightarrow +\infty} \mathcal{V}_k(\rho_{\varepsilon,k}).\end{aligned}$$

Next, we choose subsequences that attain the $\liminf$, and then apply Theorem 6.1, to find,

$$\begin{aligned}\mathcal{F}(\nu) &\geq \lim_{n \rightarrow +\infty} \left(\mathcal{E}_{\varepsilon_n}(\rho_{\varepsilon_n, k_n}) + \mathcal{V}_{\varepsilon_n}(\rho_{\varepsilon_n, k_n}) \right) + \lim_{m \rightarrow +\infty} \mathcal{V}_{k_m}(\rho_{\varepsilon_m, k_m}) \\ &\geq \mathcal{E}(\rho) + \mathcal{V}(\rho) + \mathcal{V}_{\Omega}(\rho) = \mathcal{F}(\rho).\end{aligned}$$

Since ν was an arbitrary measure in the domain of $\mathcal{F}$, this shows ρ is the unique minimizer of $\mathcal{F}$. Finally, since the above argument shows that every subsequence of $\rho_{\varepsilon,k}$ has a further subsequence that converges to ρ , the original sequence $\rho_{\varepsilon,k}$ must converge to ρ . $\square$

We now turn our attention from minimizers to gradient flows and prove that “almost” curves of maximal slopes of $\mathcal{F}_k$ converge to a gradient flow of $\mathcal{F}$ as $k \rightarrow +\infty$.

Proposition 6.2. *Assume (D), (V), (C), (Ck), (T) hold and that $\bar{\rho}$ is log-concave on Ω . Fix $T > 0$. For $k \in \mathbb{N}$, let $\rho_k \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ be an “almost” curve of maximal slope of $\mathcal{F}_k$, in the sense of Definition 5.2, and suppose there exists $\rho(0) \in D(\mathcal{F}) \cap \mathcal{P}_2(\mathbb{R}^d)$ such that*

$$(6.3) \quad \begin{aligned} \rho_k(0) &\xrightarrow{k \rightarrow +\infty} \rho(0) \text{ narrowly,} \quad \lim_{k \rightarrow +\infty} \mathcal{F}_k(\rho_k(0)) = \mathcal{F}(\rho(0)), \\ &\text{and} \quad \sup_{k \in \mathbb{N}} M_2(\rho_k(0)) < +\infty. \end{aligned}$$

Then

$$(6.4) \quad \lim_{k \rightarrow +\infty} W_1(\rho_k(t), \rho(t)) = 0, \text{ uniformly for } t \in [0, T],$$

where $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ is the unique gradient flow of $\mathcal{F}$ with initial conditions $\rho(0)$.

Proof. Recall from Theorem 6.1 that $\mathcal{F}_k$ Γ -converges to $\mathcal{F}$. Thus, by Proposition 2.16, we find that there exists $\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ so that, up to a subsequence, (6.4) holds and

$$(6.5) \quad \begin{aligned} \frac{1}{2} \int_0^t |\rho'|^2(r) dr + \frac{1}{2} \int_0^t \left(\liminf_{k \rightarrow \infty} \int_{\mathbb{R}^d} |\boldsymbol{\eta}_k(r)|^2 d\rho_k(r) \right) dr \\ \leq \mathcal{F}(\rho(0)) - \mathcal{F}(\rho(t)) \quad \text{for all } 0 \leq t \leq T. \end{aligned}$$

Furthermore, by Definition 5.2 of an “almost” curve of maximal slope, we see that $\mathcal{F}_k(\rho_k(t)) \leq \mathcal{F}_k(\rho_k(0))$ for all $k \in \mathbb{N}$, $t \in [0, T]$. Combining this with (6.3) yields

$$(6.6) \quad \sup_{t \in [0, T], k \in \mathbb{N}} \mathcal{F}_k(\rho_k(t)) \leq \sup_{k \in \mathbb{N}} \mathcal{F}_k(\rho_k(0)) < +\infty.$$

Furthermore, since $\rho_k(t) \xrightarrow{W_1} \rho(t)$, hence narrowly, for all $t \in [0, T]$, Theorem 6.1 implies,

$$(6.7) \quad \sup_{t \in [0, T]} \mathcal{F}(\rho(t)) \leq \sup_{t \in [0, T]} \liminf_{k \rightarrow +\infty} \mathcal{F}_k(\rho_k(t)) \leq \sup_{t \in [0, T], k \in \mathbb{N}} \mathcal{F}_k(\rho_k(t)) < +\infty.$$

Let us use A_k to denote the measure zero subset of $[0, T]$ on which condition (5.7) of Definition 5.2 fails, and note that $\cup_{k=1}^{+\infty} A_k$ is likewise a set of measure zero, so that for almost every $t \in [0, T]$, condition (5.7) holds for all $k \in \mathbb{N}$. The remainder of the proof will be devoted to establishing that, for almost every $t \in [0, T]$, we have

$$(6.8) \quad \liminf_{k \rightarrow +\infty} \int_{\mathbb{R}^d} |\boldsymbol{\eta}_k(t)|^2 d\rho_k(t) \geq \int_{\mathbb{R}^d} |\boldsymbol{\eta}(t)|^2 d\rho(t),$$

for ρ and $\boldsymbol{\eta}$ satisfying

$$(6.9) \quad (\rho(t)/\bar{\rho})^2 \in W_{\text{loc}}^{1,1}(\Omega),$$

$$(6.10) \quad \boldsymbol{\eta}(t)\rho(t) = \frac{\bar{\rho}}{2} \nabla(\rho(t)/\bar{\rho})^2 + \nabla V \rho(t) \text{ on } \Omega.$$

Indeed, combining inequalities (6.5) and (6.8) yields that ρ satisfies

$$\begin{aligned} \frac{1}{2} \int_0^t |\rho'|^2(r) dr + \frac{1}{2} \int_0^t \int_{\mathbb{R}^d} |\boldsymbol{\eta}(r)|^2 d\rho(r) dr \\ \leq \mathcal{F}(\rho(0)) - \mathcal{F}(\rho(t)) , \quad \text{for all } 0 \leq t \leq T. \end{aligned}$$

Proposition 3.9, which characterizes the metric slope of $\mathcal{F}$, and Theorem 2.13 imply that ρ is the unique gradient flow of $\mathcal{F}$ with initial data $\rho(0)$. Finally, we remark that this argument shows that every subsequence of ρ_k has a further subsequence that converges to ρ in the sense (6.4), implying that the original sequence must also converge to ρ .

For almost every $t \in [0, T]$, inequality (6.5) ensures that the left-hand side of (6.8) is finite. Fix such a $t \in [0, T]$. Since the left-hand side of (6.8) is finite, passing to a subsequence in k , we may assume

$$(6.11) \quad \sup_{k \in \mathbb{N}} \|\boldsymbol{\eta}_k(t)\|_{L^2(\rho_k(t))} < +\infty.$$

To conclude the proof, it remains to show that (6.8), (6.9), and (6.10) hold at this time. From now on, we will suppress dependence on t , for simplicity of notation.

Since $\mathcal{V}$ and $\mathcal{V}_k$ are bounded below, inequality (6.6) implies,

$$(6.12) \quad \left(\frac{1}{2 \sup \bar{\rho}} \right) \sup_{k \in \mathbb{N}} \|\rho_k\|_2^2 \leq \sup_{k \in \mathbb{N}} \frac{1}{2} \int_{\mathbb{R}^d} \frac{|\rho_k|^2}{\bar{\rho}} = \sup_{k \in \mathbb{N}} \mathcal{E}(\rho_k) < +\infty.$$

Likewise, inequality (6.7) and the definition of $\mathcal{F}$ imply $\text{supp } \rho(t) \subseteq \bar{\Omega}$.

Next, we note that, since Assumption (V) ensures $\nabla V \in L^\infty$, applying the triangle inequality and inequality (6.11) yields,

$$(6.13) \quad \begin{aligned} \sup_k \left| \int \frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k \right| &= \sup_k \|\boldsymbol{\eta}_k - \nabla V\|_{L^1(\rho_k)} \\ &\leq \sup_k \|\boldsymbol{\eta}_k - \nabla V\|_{L^2(\rho_k)} \\ &< +\infty. \end{aligned}$$

By [3, Theorem 5.4.4], provided we have (6.9) and (6.10), in order to show (6.8), it suffices to show

$$(6.14) \quad \begin{aligned} \liminf_{k \rightarrow +\infty} \int_{\mathbb{R}^d} f \left(\frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V \rho_k + \nabla V_k \rho_k \right) \\ = \int_{\Omega} f \left(\frac{\bar{\rho}}{2} \nabla(\rho/\bar{\rho})^2 + \nabla V \rho \right) \quad \text{for all } f \in C_c^\infty(\mathbb{R}^d), \end{aligned}$$

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,

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

Thus, (6.14) is equivalent to the claim that, for all $f \in C_c^\infty(\mathbb{R}^d)$,

$$(6.15) \quad \liminf_{k \rightarrow +\infty} \int_{\mathbb{R}^d} f \left(\frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k \right) = \int_{\Omega} f \frac{\bar{\rho}}{2} \nabla(\rho/\bar{\rho})^2.$$

We will establish (6.15) 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 (6.15).

First, we consider the region Ω . By Assumption (Ck), which ensures V_k vanishes on Ω for all k , inequality (6.13) implies,

$$(6.16) \quad \left(\frac{\inf \bar{\rho}}{2} \right) \sup_k \int_{\Omega} |\nabla(\rho_k/\bar{\rho})^2| \leq \sup_k \int_{\Omega} \left| \frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 \right| < +\infty.$$

Since $(\rho_k/\bar{\rho})^2 \in W^{1,1}(\mathbb{R}^d)$, combining (6.12) and (6.16), 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 $g \geq 0$. 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 $\rho_k/\bar{\rho} \rightarrow \sqrt{g}$ in $L^1(\Omega)$. Combining this with the fact that $\rho_k \rightarrow \rho$ 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 $\bar{\Omega}$ ensures,

$$\begin{aligned} (6.17) \quad \liminf_{k \rightarrow +\infty} \int_{\Omega} f \left(\frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k \right) &= \liminf_{k \rightarrow +\infty} \int_{\Omega} f \frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 \\ &= - \liminf_{k \rightarrow +\infty} \int_{\Omega} \nabla \left(f \frac{\bar{\rho}}{2} \right) (\rho_k/\bar{\rho})^2 \\ &= - \int_{\Omega} \nabla \left(f \frac{\bar{\rho}}{2} \right) (\rho/\bar{\rho})^2. \end{aligned}$$

By inequality (6.13), the left hand side of the equation may be bounded above by,

$$\sup_k \|f\|_{\infty} \left\| \frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k \right\|_{L^1(\mathbb{R}^d)} \leq \|f\|_{\infty} (\sup_k \|\eta_k\|_{L^2(\rho_k)} + \|\nabla V\|_{\infty}),$$

which is finite by (6.11). Thus, we conclude $(\rho/\bar{\rho})^2 \in BV(\Omega)$. Thus, for all $f \in C_c^\infty(\mathbb{R}^d)$,

$$(6.18) \quad - \int_{\Omega} \nabla \left(f \frac{\bar{\rho}}{2} \right) (\rho/\bar{\rho})^2 = \int_{\Omega} f \frac{\bar{\rho}}{2} \nabla(\rho/\bar{\rho})^2.$$

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 (6.17) and (6.18) to find,

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

Recalling the definition of η_k , then using Hölder's inequality, and using the boundedness of ∇V , we obtain,

$$\begin{aligned} \int_{\Omega} \left| f \left(\frac{\bar{\rho}}{2} \nabla(\rho_k/\bar{\rho})^2 + \nabla V_k \rho_k \right) \right| &= \int_{\Omega} |f(\eta_k \rho_k - \nabla V \rho_k)| \\ &\leq \|f\|_{L^2(\rho_k; \Omega)} (\sup_k \|\eta_k\|_{L^2(\rho_k)} + \|\nabla V\|_{L^\infty(\mathbb{R}^d)}). \end{aligned}$$

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:

$$|L(f)| \leq \|f\|_{L^2(\rho;\Omega)} \left(\sup_k \|\boldsymbol{\eta}_k\|_{L^2(\rho_k)} + \|\nabla V\|_{L^\infty(\mathbb{R}^d)} \right),$$

where we use the estimate (6.11) to see that the first term in the parentheses is finite.

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

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

Since $\|\mathbf{w}\rho\|_{L^1(\mathcal{L}^d;\Omega)} = \|\mathbf{w}\|_{L^1(\rho;\Omega)} \leq \|\mathbf{w}\|_{L^2(\rho;\Omega)}$, this shows that $(\rho/\bar{\rho})^2 \in W^{1,1}(\Omega)$, so (6.9) holds. Likewise, $\boldsymbol{\eta} := \mathbf{w} + \nabla V \in L^2(\rho)$ satisfies the conditions of (6.10). Finally, integrating by parts on the right hand side of (6.17) gives (6.15) for all $f \in C_c^\infty(\Omega)$.

It remains to show that (6.15) 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 $0 \leq \phi \leq 1$ that is compactly supported in Ω , we have,

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

To estimate I_1 , note that,

$$\left| f\phi \frac{\bar{\rho}}{2} \nabla(\rho/\bar{\rho})^2 \right| \leq \frac{\|f\|_\infty \|\bar{\rho}\|_\infty}{2} |\mathbf{w}\rho| \in L^1(\mathcal{L}^d;\Omega).$$

To estimate I_2 , note that,

$$\begin{aligned} I_2 &\leq \liminf_{k \rightarrow +\infty} \|f(1-\phi)\|_{L^2(\rho_k)} (\|\boldsymbol{\eta}_k\|_{L^2(\rho_k)} + \|\nabla V\|_\infty) \\ &\leq \|f(1-\phi)\|_{L^2(\rho)} \left(\sup_k \|\boldsymbol{\eta}_k\|_{L^2(\rho_k)} + \|\nabla V\|_\infty \right), \end{aligned}$$

where,

$$|f(1-\phi)|^2 \rho \leq |f|^2 \rho \in L^1(\mathcal{L}^d;\Omega).$$

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

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

Since $\delta > 0$ was arbitrary, this completes the proof of (6.15). $\square$

We conclude with the proof of Theorem 1.1.

Proof of Theorem 1.1. As in the statement of the theorem, let

$$\rho \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$$

be the unique gradient flow of $\mathcal{F}$ with initial condition $\rho(0)$, the existence of which is guaranteed by Proposition 3.10. By Theorem 5.1, for all $k \in \mathbb{N}$,

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

so by Proposition 5.6, there exists an “almost” curve of maximal slope $\rho_k \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ and a subsequence $\{\varepsilon_j^{(k)}\}_{j=1}^\infty$, depending on k , so that

$$(6.19) \quad \lim_{j \rightarrow +\infty} W_1(\rho_{\varepsilon_j^{(k)}, k}(t), \rho_k(t)) = 0 \text{ uniformly for } t \in [0, T].$$

In particular, for each $k \in \mathbb{N}$, there exists $\varepsilon_k > 0$ so that $\lim_{k \rightarrow +\infty} \varepsilon_k = 0$ and

$$(6.20) \quad W_1(\rho_{\varepsilon_k, k}(t), \rho_k(t)) < \frac{1}{k} \text{ for all } t \in [0, T].$$

Furthermore, since Theorem 5.1 ensures

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

Proposition 6.2 implies

$$(6.21) \quad \lim_{k \rightarrow +\infty} W_1(\rho_k(t), \rho(t)) = 0, \text{ uniformly for } t \in [0, T].$$

Fix $\delta > 0$. Choose $K_\delta > 0$ so that, for all $k \geq K_\delta$, $W_1(\rho_k(t), \rho(t)) < \delta/2$ for all $t \in [0, T]$. Then, for all $k \geq \max\{2/\delta, K_\delta\}$,

$$W_1(\rho_{\varepsilon_k, k}(t), \rho(t)) \leq W_1(\rho_k(t), \rho(t)) + W_1(\rho_{\varepsilon_k, k}(t), \rho_k(t)) \leq \frac{\delta}{2} + \frac{1}{k} \leq \delta,$$

for all $t \in [0, T]$. This gives the result. $\square$

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 $k \rightarrow +\infty$, $\varepsilon = \varepsilon(k) \rightarrow 0$. Unfortunately, our assumption of “well-preparedness” requires that the initial data of $\mathcal{F}_{\varepsilon, k}$ have bounded entropy (5.28), 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 5.3). 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 1.2. This is based on the elementary fact 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.5. (In fact, our construction of the empirical measure in the proof of Lemma A.5 closely parallels what we employ in our numerical method.) It can be seen from the proof of Lemma A.5 that if $\text{supp } \mu \subseteq [-R, R]^d$, then to approximate μ by an empirical measure $\mu^N = \sum_{i=1}^N \delta_{X_i} m^i$ to accuracy $\delta > 0$ in the 2-Wasserstein metric, N can be taken to be the smallest integer larger than $(2\sqrt{d}R/\delta)^d$. 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 $O(1/\delta^d)$ [32, 34]. 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 [76].

Once we have extended our result to particle initial data, in Theorem 1.2, we are then able to quickly obtain our two main corollaries. Corollary 1.3 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 1.4 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 1.2.

Proof of Theorem 1.2. First, let $\rho_{\varepsilon,k}(t)$ be the gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial data $\rho(0)$. By Theorem 1.1, as $k \rightarrow +\infty$, $\varepsilon = \varepsilon(k) \rightarrow 0$,

$$(6.22) \quad \lim_{k \rightarrow +\infty} W_1(\rho_{\varepsilon,k}(t), \rho(t)) = 0, \text{ uniformly for } t \in [0, T],$$

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,

$$(6.23) \quad \lambda_\varepsilon = -\varepsilon^{-d-2} \|1/\bar{\rho}\|_{L^\infty(\mathbb{R}^d)} \|D^2\zeta\|_{L^\infty(\mathbb{R}^d)} + \inf_{\{x, \xi \in \mathbb{R}^d\}} \xi^t D^2V(x)\xi,$$

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 1.2, is the unique gradient flow of $\mathcal{F}_{\varepsilon,k}$ with initial data $\rho_{\varepsilon,k}^N(0)$. By (6.22), it suffices to show that, as $k \rightarrow +\infty$, $\varepsilon = \varepsilon(k) \rightarrow 0$, $N = N(\varepsilon) \rightarrow +\infty$,

$$(6.24) \quad W_1(\rho_{\varepsilon,k}^N(t), \rho_{\varepsilon,k}(t)) \rightarrow 0, \text{ uniformly for } t \in [0, T].$$

Since $\rho_{\varepsilon,k}^N(t)$ and $\rho_{\varepsilon,k}(t)$ 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 $t \in [0, T]$,

$$W_1(\rho_{\varepsilon,k}^N(t), \rho_{\varepsilon,k}(t)) \leq W_2(\rho_{\varepsilon,k}^N(t), \rho_{\varepsilon,k}(t)) \leq e^{-\lambda_\varepsilon t} W_2(\rho_{\varepsilon,k}^N(0), \rho(0)).$$

By hypothesis (1.21), the right hand side goes to zero uniformly in $t \in [0, T]$, which completes the proof. $\square$

We now apply this to obtain the proof of Corollary 1.3.

Proof of Corollary 1.3. Let $\rho(t)$ be the solution of (WPME) with initial data $\rho(0)$, as in Theorem 1.2. By Proposition 3.10, ρ is the unique gradient flow of $\mathcal{F}$ with initial data $\rho(0)$, so by Proposition 3.14,

$$\lim_{t \rightarrow +\infty} W_1(\rho(t), \mathbb{1}_{\bar{\Omega}} \bar{\rho}) \leq \lim_{t \rightarrow +\infty} W_2(\rho(t), \mathbb{1}_{\bar{\Omega}} \bar{\rho}) = 0.$$

Combining this with Theorem 1.2 gives the result. $\square$

We conclude with the proof of Corollary 1.4.

Proof of Corollary 1.4. The evolving empirical measure $\rho_\varepsilon^N(t)$, as defined in the statement of Corollary 1.4, coincides with the evolving empirical measure in Theorem 1.2. 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 $\mathcal{F}$ coincides with $\mathcal{R}$, 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. $\square$

7. NUMERICAL SIMULATION

We now implement the particle method described in Theorem 1.2, 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 $N \rightarrow +\infty$, $\varepsilon \rightarrow 0$, for fixed $k \gg 1$, 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 [43, 61, 71, 74].

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,

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

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

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

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

Unless otherwise specified, we choose,

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

Note that this relationship between ε and N is better than expected from our rigorous results; see the discussion after Corollary 1.3. 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

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

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 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.5.

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,

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

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 :

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

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 $\varepsilon \rightarrow 0$, then the kernel density estimator $\tilde{\rho}_{\varepsilon,k}^N$ also narrowly converges to μ as $\varepsilon \rightarrow 0$. 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,

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

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 $\bar{\Omega}$, the restriction of $\tilde{\rho}_{\varepsilon,k}^N$ to $\bar{\Omega}$ is not a probability measure and $\text{KL}(\tilde{\rho}_{\varepsilon,k}^N(t), \bar{\rho})$ can be negative. On the other hand, $\tilde{\rho}_{\varepsilon,k}^N/C_{\varepsilon,k,N}$ is always a probability measure on $\bar{\Omega}$, 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:

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

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 $\text{supp } \zeta_{\varepsilon} \subset \subset B_{R\varepsilon}(0)$, 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. A second issue one would have to contend with in the case of compactly supported mollifiers is that, as described above, we expect quantitative accuracy of our method to degrade if there is not sufficient overlap between mollifiers. This degradation could be particularly sudden in the case of compactly supported mollifiers, and we expect that some “remeshing” (that

is, resampling the locations of the particles) might be needed when the particles get too far apart, in order to maintain sufficient overlap in the long time limit.

In spite of the above-described potential benefits of a compactly supported mollifier, we are still able to achieve good computational speed in practice with a strictly positive mollifier 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 [61]. 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 [19]. Namely, we performed the same simulations as those used to generate Figure 1 of [19] (the evolution of density over time, starting from Barenblatt initial data), both using the code of [19], as well as with our implementation. In Figure 1, we record the resulting improvement in terms of computational time.

Time	Carrillo et al. [19]	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 [19] (left column). Here the target is $\bar{\rho}_{\text{uni}}$, we take $k = 0$, $t = 0.15$, and the initial condition is the Barenblatt profile; see equation (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,

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

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

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

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




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). Left Column: Snapshots of $\tilde{\rho}_{\varepsilon,k}^N(t)$ for the indicated times t . Right Column: Evolution of corresponding particle trajectories $X^i(t)$.

$\tilde{\rho}_{\varepsilon,k}^N(x,t)$, 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: $\bar{\rho}_{\text{uni}}$ (left), $\bar{\rho}_{\text{log-con}}$ (middle), and $\bar{\rho}_{\text{pw-const}}$ (right). In all cases, our initial condition is given by a *Barenblatt profile* $\psi_\tau(x)$, with $\tau = 0.0625$:

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

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 $\bar{\rho}_{\text{uni}}$ and $\bar{\rho}_{\text{log-con}}$, this provides numerical verification of Corollary 1.3, since these targets $\bar{\rho}$ are log-concave. On the other hand, while $\bar{\rho}_{\text{pw-const}}$ 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 [25] and weak [33, 42] solutions of (WPME) ensures that exact solutions of the continuum PDE converge to the target $\bar{\rho}_{\text{pw-const}}$ 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 $\bar{\rho}_{\text{pw-const}}$ 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 [9], and stochastic methods in the control theory literature [35, 59], 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.

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 $\bar{\rho}_{\text{pw-const}}$, 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 $\bar{\Omega} = [-1, 1]$. 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 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 $\bar{\Omega}$, 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 $\bar{\Omega}$. 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 $\bar{\Omega}$, the kernel density estimate will satisfy,

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




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). Top: No confinement ($k = 0$). Middle: Medium confinement ($k = 100$). Bottom: Strong confinement ($k = 10^9$).

However, in the limit $N \rightarrow +\infty$, $\varepsilon \rightarrow 0$, and $k \rightarrow +\infty$, the support of $\tilde{\rho}_{\varepsilon,k}^N$ will be contained in $\bar{\Omega}$. 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 $\bar{\rho}_{\text{pw-const}}$ on Ω , in agreement with Corollary 1.3.




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, $\bar{\rho}_{\text{uni}}$ (left), $\bar{\rho}_{\text{log-con}}$ (middle), and $\bar{\rho}_{\text{pw-const}}$ (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 $N \rightarrow +\infty$, $\varepsilon \rightarrow 0$, and $k \rightarrow +\infty$, as in Theorem 1.2.

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, $\bar{\rho}_{\text{uni}}$ (left), $\bar{\rho}_{\text{log-con}}$ (middle), and $\bar{\rho}_{\text{pw-const}}$ (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, $\bar{\rho}_{\text{uni}}$ and $\bar{\rho}_{\text{log-con}}$, 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 $\bar{\rho}_{\text{pw-const}}$.

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 $k \rightarrow +\infty$, $\varepsilon \rightarrow 0$, and $N \rightarrow +\infty$. 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,

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

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




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 top and middle plots. We take $k = 10^9$, and the initial data is the Barenblatt profile (7.13).

when the solution of the underlying weighted porous medium equation is sufficiently regular, we restrict our attention to the smooth targets $\bar{\rho}_{\text{uni}}$ and $\bar{\rho}_{\text{log-con}}$.

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




FIGURE 6. The effect of k on the rate of convergence in N of the L^1 error (7.14) between $\bar{\rho}_{\varepsilon,k}^N$ 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).

Barenblatt initial data. We choose values of N from $N = 20$ 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 strength 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 3. On the other hand, selecting the confinement to be too strong can lead in a degradation of the rate of convergence as $\varepsilon \rightarrow 0$, $N \rightarrow +\infty$, as more particles would be required for a given degree of accuracy.




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).

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 conditions, we now consider uniform initial conditions,

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

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 $\bar{\rho}_{\text{uni}}$ (left) and $\bar{\rho}_{\text{log-con}}$ (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 [29] 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).




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).

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 $\bar{\rho}_{\text{uni}}$ and $\bar{\rho}_{\text{log-con}}$. Furthermore, as illustrated in Figure 3, since strong confinement is necessary to obtain convergence to the target as $t \rightarrow +\infty$, 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,

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

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

For both $\bar{\rho}_{\text{uni}}$ and $\bar{\rho}_{\text{log-con}}$ 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 1.3, in which we show that there exist parameters $T \rightarrow +\infty$, $k \rightarrow +\infty$, $\varepsilon \rightarrow 0$, $N \rightarrow +\infty$ for which our particle method indeed provides a way to approximate $\bar{\rho}$ on Ω , as relevant for applications in quantization.

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 $\gamma \in \Gamma_0(\mu, \nu)$,

$$\begin{aligned} |\partial\mathcal{G}|(\mu) &= \limsup_{\nu \rightarrow \mu} \frac{(\mathcal{G}(\mu) - \mathcal{G}(\nu))_+}{W_2(\mu, \nu)} \\ &\leq \limsup_{\nu \rightarrow \mu} \frac{1}{W_2(\mu, \nu)} \left(\int_{\mathbb{R}^d \times \mathbb{R}^d} \langle \xi(x), x - y \rangle d\gamma(x, y) - \frac{\lambda}{2} W_2^2(\mu, \nu) \right)_+ \\ &\leq \limsup_{\nu \rightarrow \mu} \left(\frac{\|\xi\|_{L^2(\mu)} W_2(\mu, \nu)}{W_2(\mu, \nu)} + \frac{\lambda_-}{2} W_2(\mu, \nu) \right) \\ &= \|\xi\|_{L^2(\mu)}, \end{aligned}$$

where $\lambda_- = \max\{-\lambda, 0\}$. This shows inequality (2.12). Uniqueness of the minimal subdifferential follows from the strict convexity of $\|\cdot\|_{L^2(\mu)}$. $\square$

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, suppose $\mu(t) \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ is a curve of maximal slope in the sense of inequality (2.14). By the definition of strong upper gradient [3, Definition 1.2.1], Young's inequality, and the fact that, under the assumptions of the theorem, the local slope $|\partial\mathcal{G}|$ is a strong upper gradient [3, Corollary 2.4.19], $t \mapsto \mathcal{G}(\mu(t))$ is absolutely continuous, and

$$-\frac{d}{dt}\mathcal{G}(\mu(t)) \leq \frac{1}{2}|\partial\mathcal{G}|^2(\mu(t)) + \frac{1}{2}|\mu'|^2(t), \text{ for almost every } t \geq 0.$$

Define $f(t) = \frac{d}{dt}\mathcal{G}(\mu(t)) + \frac{1}{2}|\partial\mathcal{G}|^2(\mu(t)) + \frac{1}{2}|\mu'|^2(t)$. Then we must have $f(t) \geq 0$ for a.e. $t \geq 0$ and inequality (2.14) ensures

$$\int_0^t f(r)dr \leq 0 \text{ for all } t \geq 0.$$

Therefore, we must have $f(t) = 0$ for a.e. $t \geq 0$ and, integrating f from $(a, b) \subseteq [0, +\infty)$, we see $\mathcal{G}(\mu(t))$ must be decreasing. This shows $\mu(t)$ is a curve of maximal slope in the pointwise sense of Ambrosio, Gigli, and Savaré [3, Definition 1.3.2]. Finally, [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.)

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]. $\square$

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 [47].

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,

$$J_\tau \mu = \operatorname{argmin}_{\nu \in \mathcal{P}_2(\mathbb{R}^d)} \left\{ \frac{1}{2\tau} W_2^2(\mu, \nu) + \mathcal{G}(\nu) \right\},$$

and define the *minimizing movement scheme* $J_\tau^n \mu$ by,

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

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

$$(A.1) \quad \mathcal{G}(J_\tau^n \mu) \leq \mathcal{G}(J_\tau^{n-1} \mu).$$

We recall Theorem A.2 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,*

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

Then for all $t \in [0, T]$, we have $\lim_{n \rightarrow +\infty} \mu_\tau(t) = \mu(t)$ 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]. $\square$

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 bound for 2-absolutely continuous curves). *Suppose $\rho(t) \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$. Then we have*

$$(A.3) \quad M_2(\rho(t)) \leq (1 + te^t) \left(M_2(\rho(0)) + \int_0^t |\rho'|^2(r) dr \right) \quad \text{for all } t \in [0, T].$$

Proof. We shall show that, for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$,

$$(A.4) \quad W_2^2(\rho(t), \mu) \leq (1 + te^t) \left[W_2^2(\rho(0), \mu) + \int_0^t |\rho'|^2(r) dr \right] \quad \text{for all } t \in [0, T].$$

The desired estimate then follows from taking $\mu = \delta_0$, as in (2.4).

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 $|\partial\mathcal{W}|(\rho)$ is a strong upper gradient for $\mathcal{W}$ (see [3, Definition 1.2.1, Corollary 2.4.10]), which implies,

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

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

$$\begin{aligned}
 |\partial\mathcal{W}|(\rho) &= \limsup_{\nu \rightarrow \rho} \frac{W_2^2(\nu, \mu) - W_2^2(\rho, \mu)}{2W_2(\rho, \nu)} \\
 &= \limsup_{\nu \rightarrow \rho} \frac{(W_2(\nu, \mu) - W_2(\rho, \mu))(W_2(\nu, \mu) + W_2(\rho, \mu))}{2W_2(\rho, \nu)} \\
 (A.6) \quad &\leq \limsup_{\nu \rightarrow \rho} \frac{W_2(\rho, \nu)(W_2(\nu, \mu) + W_2(\rho, \mu))}{2W_2(\rho, \nu)} = W_2(\rho, \mu).
 \end{aligned}$$

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

$$\begin{aligned}
 \frac{1}{2} [W_2^2(\rho(t), \mu) - W_2^2(\rho(0), \mu)] &\leq |\mathcal{W}(\rho(t)) - \mathcal{W}(\rho(0))| \\
 &\leq \int_0^t W_2(\rho(s), \mu) |\rho'(s)| ds \\
 &\leq \|W_2(\rho(s), \mu)\|_{L^2([0, t])} \|\rho'(s)\|_{L^2([0, t])} \\
 &\leq \frac{1}{2} \int_0^t W_2^2(\rho(s), \mu) ds + \frac{1}{2} \int_0^t |\rho'|^2(s) ds.
 \end{aligned}$$

By Gronwall's inequality, this implies inequality (A.4). $\square$

As a corollary of Proposition A.3, we obtain a bound on the second moment for gradient flows in terms of the energy driving the gradient flow.

Corollary A.4 (M_2 bound for curves of maximal slope). *Suppose*

$$\rho(t) \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$$

is a curve of maximal slope of a nonnegative functional $\mathcal{G} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow [0, +\infty]$, that is, $\rho(t)$ satisfies Theorem 2.13(i). Suppose further that $\rho_0 \in D(\mathcal{G})$. Then we have,

$$(A.7) \quad M_2(\rho(t)) \leq (1 + te^t) (M_2(\rho(0)) + 2\mathcal{G}(\rho(0))).$$

Proof. This is an immediate consequence of Proposition A.3, using the definition of a curve of maximal slope, Theorem 2.13(i), and the fact that $\mathcal{G}$ is assumed to be nonnegative. $\square$

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.5 (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 $\mu^N = \sum_{i=1}^N \delta_{X^i} m^i$ satisfies $W_2(\mu, \mu^N) \leq \delta$.*

Proof of Lemma A.5. 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, $Q_r(0) = [-\frac{r}{2}, \frac{r}{2}]^d$. 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 approximate μ by a compactly supported measure μ_R . To this end, note that since $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, there exists $R > 0$ so that $\int_{Q_R(0)^c} |x|^2 d\mu \leq (\frac{\delta}{2})^2$. 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,

$$(A.8) \quad W_2(\mu, \mu_R) \leq \left(\int |\mathbf{t}_R(x) - x|^2 d\mu \right)^{1/2} \leq \left(\int_{Q_R(0)^c} |x|^2 d\mu \right)^{1/2} \leq \frac{\delta}{2}.$$

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

$$(A.9) \quad K \geq \frac{2\sqrt{d}R}{\delta},$$

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 $\{X^i\}_{i=1}^{K^d}$ 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:

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

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

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

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, $\frac{\sqrt{d}R}{K}$. Finally, we conclude by using the definition of K in (A.9), together with the estimate (A.8), and the triangle inequality:

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

□

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 Buttazzo [15, Corollary 3.4.2]: given $g : \mathbb{R}^d \times \mathbb{R} \rightarrow$

$[0, +\infty]$, consider the functional $\mathcal{G} : \mathcal{P}(\mathbb{R}^d) \rightarrow [0, +\infty]$ defined by,

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

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 $\theta : \mathbb{R} \rightarrow \mathbb{R}$ with $\lim_{t \rightarrow \infty} \frac{\theta(t)}{t} = \infty$ and $g(x, y) \geq \theta(|y|)$ 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 $\mathcal{E}$. $\square$

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. First we show item (i). Define the log-concave extension of $\bar{\rho}$ by

$$\tilde{\rho} = e^{-W}, \quad W(x) = \begin{cases} -\log \bar{\rho}(x) & \text{if } x \in \Omega, \\ +\infty & \text{otherwise.} \end{cases}$$

In this way, $\tilde{\rho} = \bar{\rho}$ on Ω , but $\tilde{\rho}$ is log-concave on all of $\mathbb{R}^d$. Furthermore, for all $\rho \in \mathcal{P}_2(\mathbb{R}^d)$,

$$\mathcal{E}(\rho) + \mathcal{V}_\Omega(\rho) = \begin{cases} \frac{1}{2} \int_{\mathbb{R}^d} \frac{|\rho(x)|^2}{\tilde{\rho}(x)} d\mathcal{L}^d(x) & \text{if } \rho \ll \tilde{\rho}, \\ +\infty & \text{otherwise.} \end{cases}$$

Finally, [3, Theorem 9.4.12] ensures the energy on the right hand side is convex along generalized geodesics.

Item (ii) is a consequence of the fact that, for any potential $W : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ 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 $V \in C^2(\mathbb{R}^d)$ with Hessian bounded below implies $D^2V \geq \lambda I_{d \times d}$ for $\lambda = \inf_{\{x, \xi \in \mathbb{R}^d\}} \xi^t D^2V(x) \xi$, hence we also have $D^2(\zeta_\varepsilon * V) \geq \lambda I_{d \times d}$ 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. $\square$

Next we prove Proposition 3.8, characterizing the minimal element of the subdifferential of $\mathcal{F}_{\varepsilon, k}$.

Proof of Proposition 3.8. For simplicity of notation, denote,

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

Note that Proposition 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,

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

By definition of the metric slope,

$$\begin{aligned} \text{(B.3)} \quad |\partial\mathcal{F}_{\varepsilon,k}|(\mu) &= \limsup_{\nu \rightarrow \mu} \frac{(\mathcal{F}_{\varepsilon,k}(\mu) - \mathcal{F}_{\varepsilon,k}(\nu))_+}{W_2(\mu, \nu)} \geq \limsup_{\alpha \rightarrow 0} \frac{(\mathcal{F}_{\varepsilon,k}(\mu) - \mathcal{F}_{\varepsilon,k}(\mu_\alpha))_+}{W_2(\mu, \mu_\alpha)} \\ &\geq \frac{1}{\|\nabla\psi\|_{L^2(\mu)}} \limsup_{\alpha \rightarrow 0} \frac{(\mathcal{F}_{\varepsilon,k}(\mu) - \mathcal{F}_{\varepsilon,k}(\mu_\alpha))_+}{\alpha}. \end{aligned}$$

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,

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

we obtain,

$$\begin{aligned} |\partial\mathcal{F}_{\varepsilon,k}|(\mu)\|\nabla\psi\|_{L^2(\mu)} &\geq \lim_{\alpha \rightarrow 0} \frac{\mathcal{E}_\varepsilon(\mu) - \mathcal{E}_\varepsilon(\mu_\alpha)}{\alpha} + \frac{\mathcal{V}_\varepsilon(\mu) - \mathcal{V}_\varepsilon(\mu_\alpha)}{\alpha} + \frac{\mathcal{V}_k(\mu) - \mathcal{V}_k(\mu_\alpha)}{\alpha} \\ &= - \int \frac{\zeta_\varepsilon * \mu(x)}{\bar{\rho}(x)} \int \langle \nabla\zeta_\varepsilon(x - y_2), y_3 - y_2 \rangle d\gamma(y_1, y_2, y_3) \\ &\quad + \int \langle \nabla(\zeta_\varepsilon * V)(y_2) + \nabla V_k(y_2), y_3 - y_2 \rangle dx \\ &= - \int \left\langle \left(\nabla\zeta_\varepsilon * \left(\frac{\zeta_\varepsilon * \mu}{\bar{\rho}} \right) \right) + \nabla(\zeta_\varepsilon * V) + \nabla V_k, \nabla\psi \right\rangle d\mu. \end{aligned}$$

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

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

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

We now turn to a proof of Proposition 3.10, which characterizes the gradient flow of $\mathcal{F}$ in terms of a partial differential equation.

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

$$\text{(B.4)} \quad \int_0^t |\partial\mathcal{F}|^2(\mu(r)) dr \leq \mathcal{F}(\mu_0) + \|V\|_\infty < +\infty.$$

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

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

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 $t \geq 0$. Furthermore, Proposition 3.9 implies that, for almost every $t \geq 0$, $(\mu(t)/\bar{\rho})^2 \in W_{\text{loc}}^{1,1}(\Omega)$ and that there exists $\xi(t) \in \partial^\circ \mathcal{F}(\mu)$ with,

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

By Definition 2.12 of gradient flow, we obtain that μ satisfies the continuity equation (2.13) with $v(t) = -\xi(t)$. Using the expression (B.6) for ξ therefore yields (3.7). Finally, the containment (3.9) follows from inequality (B.4) and equation (B.6).

On the other hand, suppose μ solves (3.7) and satisfies (3.8)-3.9. Then, defining ξ on the support of μ via (B.6) implies that the hypotheses of Proposition 3.9 are satisfied, so $\xi \in \partial^\circ \mathcal{F}(\mu)$. From this we find that (3.7) is exactly the continuity equation in Definition 2.12 of the gradient flow, with $v(t) = -\xi(t)$ satisfying $\|v(t)\|_{L^2(\mu(t))} \in L_{\text{loc}}^1(0, +\infty)$. Thus, we have that $\mu \in AC^2([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ [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. $\square$

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 (3.11) follows directly from Definition 2.12, Proposition 3.8, and Theorem 2.13.

Now suppose that $\mu(t)$ satisfies (3.11). Then, the fact that the velocity field in the continuity equations is uniformly bounded ensures, by [3, 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. $\square$

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 $(X^1, \dots, X^N)$ that appears on the right-hand side of (3.12) is Lipschitz continuous, and therefore the ODE system (3.12) is well-posed. Suppose $X^i(t)$ solves (3.12). We claim that it suffices to show that $\mu(t) = \sum_{i=1}^N \delta_{X^i(t)} m^i$ solves (3.11). Proposition 3.12 then ensures that $\mu(t)$ is the unique solution of the gradient flow.

The fact that $\lim_{t \rightarrow 0^+} \mu(t) = \mu(0)$ in W_2 follows immediately from the definition of $\mu(t)$ and $\mu(0)$. Next, note that,

$$(B.7) \quad \begin{aligned} & - \int_{\mathbb{R}^d} \nabla \zeta_\varepsilon(X^i(t) - z) \frac{1}{\bar{\rho}(z)} \sum_{j=1}^N m^j \zeta(z - X^j(t)) dz \\ & = - \int_{\mathbb{R}^d} \nabla \zeta_\varepsilon(X^i(t) - z) \frac{1}{\bar{\rho}(z)} \zeta(z - y) d\mu(y) dz \\ & = - \left(\nabla \zeta_\varepsilon * \left(\frac{\zeta_\varepsilon * \mu}{\bar{\rho}} \right) \right) (X^i(t)). \end{aligned}$$

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.7), for each $1 \leq i \leq N$,

$$\begin{aligned} 0 &= \int_0^\infty \frac{d}{dt} f(X^i(t), t) dt \\ &= \int_0^\infty \left\langle \nabla f(X^i(t), t), \dot{X}^i(t) \right\rangle + \partial_t f(X^i(t), t) dt \\ &= \int_0^\infty \left\langle \nabla f(X^i(t), t), \left(- \left(\nabla \zeta_\varepsilon * \left(\frac{\zeta_\varepsilon * \mu}{\bar{\rho}} \right) \right)(X^i(t)) - \nabla(\zeta_\varepsilon * V)(X^i(t)) \right. \right. \\ &\quad \left. \left. - \nabla V_k(X^i(t)) \right) \right\rangle + \partial_t f(X^i(t), t) dt. \end{aligned}$$

Multiplying by m^i , summing over i , and recalling the definition of μ yields,

$$\begin{aligned} 0 &= \int_0^\infty \int_{\mathbb{R}^d} \left\langle \nabla f(x, t), \left(- \left(\nabla \zeta_\varepsilon * \left(\frac{\zeta_\varepsilon * \mu}{\bar{\rho}} \right) \right)(x) - \nabla(\zeta_\varepsilon * V)(x) \right. \right. \\ &\quad \left. \left. - \nabla V_k(x) \right) \right\rangle + \partial_t f(x, t) d\mu(x, t) dt. \end{aligned}$$

Thus, μ is a distributional solution of the continuity equation (3.11). $\square$

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,

$$\begin{aligned} f(x_i, x_j) &= \frac{-2\varepsilon^2 x_i - 6\varepsilon^2 x_j + x_i^3 + x_j^3 - x_i x_j^2 + 4x_i - x_j^3 - 4x_j}{16\sqrt{\pi}\varepsilon^3} C_{\bar{\rho}} e^{-(x_i - x_j)^2 / (4\varepsilon^2)}, \\ g(x_i, x_j) &= [\psi_{ij}(+\infty) - \psi_{ij}(-\infty)], \\ \psi_{ij}(z) &= \frac{C_{\bar{\rho}}}{8} \varepsilon e^{\frac{-(x_i^2 + x_j^2 + 2z^2)}{2\varepsilon^2}} \left(-\sqrt{\pi}(2\varepsilon^2 + x_i^2 + 2x_i x_j + z^2 + 4) \right. \\ &\quad \left. e^{\frac{(x_i^2 + x_j^2 + 2x_i x_j)}{4\varepsilon^2}} \operatorname{erf}\left(\frac{x_i + x_j - 2z}{2}\right) - 2\varepsilon(x_i + x_j + 2z)\varepsilon^{\frac{z(x_i + x_j)}{\varepsilon^2}} \right). \end{aligned}$$

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

$$f(x_i, x_j) = \sum_{k=1}^N c_k^{-1} [\varphi_{ij}(b_{k+1}) - \varphi_{ij}(b_k)],$$

$$\varphi_{ij}(z) = -\frac{e^{-(x_i^2 + 2z^2 + x_j^2)/(2\varepsilon^2)}}{8\pi\varepsilon^3} \left(2\varepsilon e^{z(x_i + x_j)/\varepsilon^2} \right. \\ \left. - \sqrt{\pi}(x_i - x_j) e^{((x_i + x_j)^2 + 4z^2)/(4\varepsilon^2)} \operatorname{erf}\left(\frac{x_i - 2z + x_j}{2\varepsilon}\right) \right) \\ + \frac{e^{-(x_i^2 + x_j^2)/(2\varepsilon^2)}}{8\pi\varepsilon^3} \left(2\varepsilon - \sqrt{\pi}(x_i - x_j) e^{(x_i + x_j)^2/(4\varepsilon^2)} \operatorname{erf}\left(\frac{x_i + x_j}{2\varepsilon}\right) \right),$$

$$g(x_i, x_j) = \sum_{k=1}^N [\psi_{ij}(b_{k+1}) - \psi_{ij}(b_k)],$$

$$\psi_{ij} = \frac{-1}{4\sqrt{\pi}} e^{\frac{-(x_i - x_j)^2}{4\varepsilon^2}} \operatorname{erf}\left(\frac{x_i + x_j - 2z}{2\varepsilon}\right).$$

REFERENCES

- [1] L. Alasio, M. Bruna, and J. A. Carrillo, *The role of a strong confining potential in a nonlinear Fokker-Planck equation*, Nonlinear Anal. **193** (2020), 111480, 28, DOI 10.1016/j.na.2019.03.003. MR4062975
- [2] L. Ambrosio, E. Brué, and D. Semola, *Lectures on Optimal Transport*, Unitext, vol. 130, Springer, Cham, [2021] ©2021. La Matematica per il 3+2, DOI 10.1007/978-3-030-72162-6. MR4294651
- [3] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient Flows in Metric Spaces and in the Space of Probability Measures*, 2nd ed., Lectures in Mathematics ETH Zürich, Birkhäuser Verlag, Basel, 2008. MR2401600
- [4] B. Anderson, E. Loeser, M. Gee, F. Ren, S. Biswas, O. Turanova, M. Haberland, and A. L. Bertozzi, *Quantitative assessment of robotic swarm coverage*, Proceedings of the 15th International Conference on Informatics in Control, Automation, and Robotics, vol. 2, 2018, pp. 91–101.
- [5] B. G. Anderson, E. Loeser, M. Gee, F. Ren, S. Biswas, O. Turanova, M. Haberland, and A. L. Bertozzi, *Quantifying robotic swarm coverage*, Informatics in Control, Automation and Robotics, Lect. Notes Electr. Eng., vol. 613, Springer, Cham, [2020] ©2020, pp. 276–301, DOI 10.1007/978-3-030-31993-9_13. MR4328240
- [6] J.-D. Benamou, G. Carlier, Q. Mérigot, and É. Oudet, *Discretization of functionals involving the Monge-Ampère operator*, Numer. Math. **134** (2016), no. 3, 611–636, DOI 10.1007/s00211-015-0781-y. MR3555350
- [7] M. Bessemoulin-Chatard and F. Filbet, *A finite volume scheme for nonlinear degenerate parabolic equations*, SIAM J. Sci. Comput. **34** (2012), no. 5, B559–B583, DOI 10.1137/110853807. MR3023729
- [8] V. I. Bogachev, *Measure Theory. Vol. I, II*, Springer-Verlag, Berlin, 2007, DOI 10.1007/978-3-540-34514-5. MR2267655
- [9] N. Bou-Rabee and A. Eberle, *Markov chain Monte Carlo methods*, Lecture notes, 2020 <https://uni-bonn.sciebo.de/s/kzTUff5FrWGAay>.
- [10] D. P. Bourne and R. Cristoferi, *Asymptotic optimality of the triangular lattice for a class of optimal location problems*, Comm. Math. Phys. **387** (2021), no. 3, 1549–1602, DOI 10.1007/s00220-021-04216-6. MR4324385
- [11] D. P. Bourne and S. M. Roper, *Centroidal power diagrams, Lloyd’s algorithm, and applications to optimal location problems*, SIAM J. Numer. Anal. **53** (2015), no. 6, 2545–2569, DOI 10.1137/141000993. MR3419889

- [12] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*, Princeton Series in Applied Mathematics, Princeton University Press, Princeton, NJ, 2009. A mathematical approach to motion coordination algorithms, DOI 10.1515/9781400831470. MR2524493
- [13] M. Burger, J. A. Carrillo, and M.-T. Wolfram, *A mixed finite element method for nonlinear diffusion equations*, Kinet. Relat. Models **3** (2010), no. 1, 59–83, DOI 10.3934/krm.2010.3.59. MR2580954
- [14] M. Burger and A. Esposito, *Porous medium equation as limit of nonlocal interaction*, Preprint, [arXiv:2202.05030](https://arxiv.org/abs/2202.05030), 2022.
- [15] G. Buttazzo, *Semicontinuity, Relaxation and Integral Representation in the Calculus of Variations*, Pitman Research Notes in Mathematics Series, vol. 207, Longman Scientific & Technical, Harlow; copublished in the United States with John Wiley & Sons, Inc., New York, 1989. MR1020296
- [16] E. Caglioti, F. Golse, and M. Iacobelli, *A gradient flow approach to quantization of measures*, Math. Models Methods Appl. Sci. **25** (2015), no. 10, 1845–1885, DOI 10.1142/S0218202515500475. MR3358447
- [17] J. A. Carrillo, Y. Huang, F. S. Patacchini, and G. Wolansky, *Numerical study of a particle method for gradient flows*, Kinet. Relat. Models **10** (2017), no. 3, 613–641, DOI 10.3934/krm.2017025. MR3591126
- [18] J. A. Carrillo, A. Chertock, and Y. Huang, *A finite-volume method for nonlinear nonlocal equations with a gradient flow structure*, Commun. Comput. Phys. **17** (2015), no. 1, 233–258, DOI 10.4208/cicp.160214.010814a. MR3372289
- [19] J. A. Carrillo, K. Craig, and F. S. Patacchini, *A blob method for diffusion*, Calc. Var. Partial Differential Equations **58** (2019), no. 2, Paper No. 53, 53, DOI 10.1007/s00526-019-1486-3. MR3913840
- [20] J. A. Carrillo, K. Craig, L. Wang, and C. Wei, *Primal dual methods for Wasserstein gradient flows*, Found. Comput. Math. **22** (2022), no. 2, 389–443, DOI 10.1007/s10208-021-09503-1. MR4407747
- [21] J. A. Carrillo, K. Craig, and Y. Yao, *Aggregation-diffusion equations: dynamics, asymptotics, and singular limits*, Active Particles. Vol. 2. Advances in theory, models, and applications, Model. Simul. Sci. Eng. Technol., Birkhäuser/Springer, Cham, 2019, pp. 65–108. MR3932458
- [22] J. A. Carrillo and J. S. Moll, *Numerical simulation of diffusive and aggregation phenomena in nonlinear continuity equations by evolving diffeomorphisms*, SIAM J. Sci. Comput. **31** (2009/10), no. 6, 4305–4329, DOI 10.1137/080739574. MR2566595
- [23] J. A. Carrillo, F. S. Patacchini, P. Sternberg, and G. Wolansky, *Convergence of a particle method for diffusive gradient flows in one dimension*, SIAM J. Math. Anal. **48** (2016), no. 6, 3708–3741, DOI 10.1137/16M1077210. MR3566905
- [24] J. A. Carrillo, H. Ranetbauer, and M.-T. Wolfram, *Numerical simulation of nonlinear continuity equations by evolving diffeomorphisms*, J. Comput. Phys. **327** (2016), 186–202, DOI 10.1016/j.jcp.2016.09.040. MR3564334
- [25] S. Chewi, T. L. Gouic, C. Lu, T. Maunu, and P. Rigollet, *SVGD as a kernelized Wasserstein gradient flow of the chi-squared divergence*, Preprint, [arXiv:2006.02509](https://arxiv.org/abs/2006.02509), 2020.
- [26] L. Chizat and F. Bach, *On the global convergence of gradient descent for over-parameterized models using optimal transport*, Preprint, [arXiv:1805.09545](https://arxiv.org/abs/1805.09545), 2018.
- [27] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, *Coverage control for mobile sensing networks*, IEEE Trans. Robot. Autom. **20** (2004), no. 2, 243–255.
- [28] K. Craig, *Nonconvex gradient flow in the Wasserstein metric and applications to constrained nonlocal interactions*, Proc. Lond. Math. Soc. (3) **114** (2017), no. 1, 60–102, DOI 10.1112/plms.12005. MR3653077
- [29] K. Craig and A. L. Bertozzi, *A blob method for the aggregation equation*, Math. Comp. **85** (2016), no. 300, 1681–1717, DOI 10.1090/mcom3033. MR3471104
- [30] S. Daneri, E. Radici, and E. Runa, *Deterministic particle approximation of aggregation-diffusion equations on unbounded domains*, J. Differential Equations **312** (2022), 474–517, DOI 10.1016/j.jde.2021.12.019. MR4361838
- [31] M. Di Francesco and M. D. Rosini, *Rigorous derivation of nonlinear scalar conservation laws from follow-the-leader type models via many particle limit*, Arch. Ration. Mech. Anal. **217** (2015), no. 3, 831–871, DOI 10.1007/s00205-015-0843-4. MR3356989
- [32] V. Dobrić and J. E. Yukich, *Asymptotics for transportation cost in high dimensions*, J. Theoret. Probab. **8** (1995), no. 1, 97–118, DOI 10.1007/BF02213456. MR1308672

- [33] J. Dolbeault, I. Gentil, A. Guillin, and F.-Y. Wang, *L^q -functional inequalities and weighted porous media equations*, Potential Anal. **28** (2008), no. 1, 35–59, DOI 10.1007/s11118-007-9066-0. MR2366398
- [34] R. M. Dudley, *The speed of mean Glivenko-Cantelli convergence*, Ann. Math. Statist. **40** (1968), 40–50, DOI 10.1214/aoms/1177697802. MR236977
- [35] K. Elamvazhuthi, C. Adams, and S. Berman, *Coverage and field estimation on bounded domains by diffusive swarms*, 2016 IEEE 55th Conference on Decision and Control (CDC), IEEE, 2016, pp. 2867–2874.
- [36] K. Elamvazhuthi and S. Berman, *Nonlinear generalizations of diffusion-based coverage by robotic swarms*, 2018 IEEE Conference on Decision and Control (CDC), IEEE, 2018, pp. 1341–1346.
- [37] U. Eren and B. Açıkmese, *Velocity field generation for density control of swarms using heat equation and smoothing kernels*, IFAC-PapersOnLine **50** (2017), no. 1, 9405–9411.
- [38] L. C. Evans, O. Savin, and W. Gangbo, *Diffeomorphisms and nonlinear heat flows*, SIAM J. Math. Anal. **37** (2005), no. 3, 737–751, DOI 10.1137/04061386X. MR2191774
- [39] A. Figalli and F. Glaudo, *An Invitation to Optimal Transport, Wasserstein Distances, and Gradient Flows*, EMS Textbooks in Mathematics, EMS Press, Berlin, [2021] ©2021, DOI 10.4171/ETB/22. MR4331435
- [40] T. O. Gallouët, Q. Mérigot, and A. Natale, *Convergence of a Lagrangian discretization for barotropic fluids and porous media flow*, SIAM J. Math. Anal. **54** (2022), no. 3, 2990–3018, DOI 10.1137/21M1422756. MR4421481
- [41] S. Graf and H. Luschgy, *Foundations of Quantization for Probability Distributions*, Lecture Notes in Mathematics, vol. 1730, Springer-Verlag, Berlin, 2000, DOI 10.1007/BFb0103945. MR1764176
- [42] 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 Contin. Dyn. Syst. **33** (2013), no. 8, 3599–3640, DOI 10.3934/dcds.2013.33.3599. MR3021373
- [43] J. D. Hunter, *Matplotlib: a 2D graphics environment*, Comput. Sci. Eng. **9** (2007), no. 3, 90–95.
- [44] M. Iacobelli, *A Gradient Flow Perspective on the Quantization Problem*, PDE Models for Multi-Agent Phenomena, Springer INdAM Ser., vol. 28, Springer, Cham, 2018, pp. 145–165. MR3888971
- [45] M. Iacobelli, *Asymptotic analysis for a very fast diffusion equation arising from the 1D quantization problem*, Discrete Contin. Dyn. Syst. **39** (2019), no. 9, 4929–4943, DOI 10.3934/dcds.2019201. MR3986316
- [46] A. Javanmard, M. Mondelli, and A. Montanari, *Analysis of a two-layer neural network via displacement convexity*, Ann. Statist. **48** (2020), no. 6, 3619–3642, DOI 10.1214/20-AOS1945. MR4185822
- [47] R. Jordan, D. Kinderlehrer, and F. Otto, *The variational formulation of the Fokker-Planck equation*, SIAM J. Math. Anal. **29** (1998), no. 1, 1–17, DOI 10.1137/S0036141096303359. MR1617171
- [48] A. Korba, A. Salim, M. Arbel, G. Luise, and A. Gretton, *A non-asymptotic analysis for Stein variational gradient descent*, Adv. Neural Inf. Process. Syst. **33** (2020), 4672–4682.
- [49] V. Krishnan and S. Martínez, *Distributed optimal transport for the deployment of swarms*, 2018 IEEE Conference on Decision and Control (CDC), IEEE, 2018, pp. 4583–4588.
- [50] P.-L. Lions and S. Mas-Gallic, *Une méthode particulière déterministe pour des équations diffusives non linéaires* (French, with English and French summaries), C. R. Acad. Sci. Paris Sér. I Math. **332** (2001), no. 4, 369–376, DOI 10.1016/S0764-4442(00)01795-X. MR1821479
- [51] Q. Liu, *Stein variational gradient descent as gradient flow*, Adv. Neural Inf. Process. Syst. **30** (2017).
- [52] Q. Liu and D. Wang, *Stein variational gradient descent: a general purpose Bayesian inference algorithm*, Adv. Neural Inf. Process. Syst. **29** (2016).
- [53] J. Lu, Y. Lu, and J. Nolen, *Scaling limit of the Stein variational gradient descent: the mean field regime*, SIAM J. Math. Anal. **51** (2019), no. 2, 648–671, DOI 10.1137/18M1187611. MR3919409
- [54] D. Matthes, R. J. McCann, and G. Savaré, *A family of nonlinear fourth order equations of gradient flow type*, Comm. Partial Differential Equations **34** (2009), no. 10-12, 1352–1397, DOI 10.1080/03605300903296256. MR2581977

- [55] D. Matthes and H. Osberger, *Convergence of a variational Lagrangian scheme for a nonlinear drift diffusion equation*, ESAIM Math. Model. Numer. Anal. **48** (2014), no. 3, 697–726, DOI 10.1051/m2an/2013126. MR3177862
- [56] D. Matthes and B. Söllner, *Convergent Lagrangian discretization for drift-diffusion with nonlocal aggregation*, Innovative Algorithms and Analysis, Springer INdAM Ser., vol. 16, Springer, Cham, 2017, pp. 313–351. MR3644286
- [57] S. Mei, A. Montanari, and P.-M. Nguyen, *A mean field view of the landscape of two-layer neural networks*, Proc. Natl. Acad. Sci. USA **115** (2018), no. 33, E7665–E7671, DOI 10.1073/pnas.1806579115. MR3845070
- [58] Q. Mérigot, *A multiscale approach to optimal transport*, Comput. Graph. Forum 30 (2011), 1583–1592. Wiley Online Library.
- [59] A. R. Mesquita, J. P. Hespanha, and K. Åström, *Optimotaxis: a stochastic multi-agent optimization procedure with point measurements*, Hybrid Systems: Computation and Control, Lecture Notes in Comput. Sci., vol. 4981, Springer, Berlin, 2008, pp. 358–371, DOI 10.1007/978-3-540-78929-1_26. MR2728891
- [60] K. Oelschläger, *Large systems of interacting particles and the porous medium equation*, J. Differential Equations **88** (1990), no. 2, 294–346, DOI 10.1016/0022-0396(90)90101-T. MR1081251
- [61] 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.
- [62] F. Otto, *The geometry of dissipative evolution equations: the porous medium equation*, Comm. Partial Differential Equations **26** (2001), no. 1-2, 101–174, DOI 10.1081/PDE-100002243. MR1842429
- [63] F. S. Patacchini and D. Slepčev, *The nonlocal-interaction equation near attracting manifolds*, Discrete Contin. Dyn. Syst. **42** (2022), no. 2, 903–929, DOI 10.3934/dcds.2021142. MR4369035
- [64] G. M. Rotskoff and E. Vanden-Eijnden, *Trainability and accuracy of artificial neural networks: an interacting particle system approach*, Comm. Pure Appl. Math. **75** (2022), no. 9, 1889–1935. MR4465905
- [65] H. L. Royden, *Real Analysis*, The Macmillan Company, New York; Collier Macmillan Ltd., London, 1963. MR0151555
- [66] F. Santambrogio, *Optimal Transport for Applied Mathematicians*, Progress in Nonlinear Differential Equations and their Applications, vol. 87, Birkhäuser/Springer, Cham, 2015. Calculus of variations, PDEs, and modeling, DOI 10.1007/978-3-319-20828-2. MR3409718
- [67] S. Serfaty, *Gamma-convergence of gradient flows on Hilbert and metric spaces and applications*, Discrete Contin. Dyn. Syst. **31** (2011), no. 4, 1427–1451, DOI 10.3934/dcds.2011.31.1427. MR2836361
- [68] J. Sirignano and K. Spiliopoulos, *Mean field analysis of neural networks: a law of large numbers*, SIAM J. Appl. Math. **80** (2020), no. 2, 725–752, DOI 10.1137/18M1192184. MR4074020
- [69] Z. Sun, J. A. Carrillo, and C.-W. Shu, *A discontinuous Galerkin method for nonlinear parabolic equations and gradient flow problems with interaction potentials*, J. Comput. Phys. **352** (2018), 76–104, DOI 10.1016/j.jcp.2017.09.050. MR3717129
- [70] A. B. Tsybakov, *Introduction to Nonparametric Estimation*, Springer Series in Statistics, Springer, New York, 2009. Revised and extended from the 2004 French original; Translated by Vladimir Zaiats, DOI 10.1007/b13794. MR2724359
- [71] S. Van Der Walt, S. C. Colbert, and G. Varoquaux, *The NumPy array: a structure for efficient numerical computation*, Comput. Sci. Eng. 13 (2011), no. 2, 22–30.
- [72] J. L. Vázquez, *The mathematical theories of diffusion: nonlinear and fractional diffusion*, Nonlocal and Nonlinear Diffusions and Interactions: New Methods and Directions, Lecture Notes in Math., vol. 2186, Springer, Cham, 2017, pp. 205–278. MR3588125
- [73] C. Villani, *Topics in Optimal Transportation*, Graduate Studies in Mathematics, vol. 58, American Mathematical Society, Providence, RI, 2003, DOI 10.1090/gsm/058. MR1964483
- [74] 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*, Nat. Methods 17 (2020), no. 3, 261–272.
- [75] N. K. Vishnoi, *An introduction to Hamiltonian Monte Carlo method for sampling*, Preprint, arXiv:2108.12107, 2021.

- [76] J. Weed and F. Bach, *Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance*, Bernoulli **25** (2019), no. 4A, 2620–2648, DOI 10.3150/18-BEJ1065. MR4003560
- [77] W. E, C. Ma, and L. Wu, *Machine learning from a continuous viewpoint, I*, Sci. China Math. **63** (2020), no. 11, 2233–2266, DOI 10.1007/s11425-020-1773-8. MR4170870
- [78] M. Westdickenberg and J. Wilkening, *Variational particle schemes for the porous medium equation and for the system of isentropic Euler equations*, M2AN Math. Model. Numer. Anal. **44** (2010), no. 1, 133–166, DOI 10.1051/m2an/2009043. MR2647756
- [79] S. Wojtowytsch, *On the convergence of gradient descent training for two-layer ReLU-networks in the mean field regime*, Preprint, [arXiv:2005.13530](https://arxiv.org/abs/2005.13530), 2020.

DEPARTMENT OF MATHEMATICS, SOUTH HALL, ROOM 6607, UNIVERSITY OF CALIFORNIA, SANTA BARBARA, CALIFORNIA 93106-3080

Email address: kcraig@math.ucsb.edu

DEPARTMENT OF MECHANICAL ENGINEERING, BOURNS HALL A342, 900 UNIVERSITY AVE., RIVERSIDE, CALIFORNIA 92521

Email address: kelamvazhuthi@engr.ucr.edu

BIORESOURCE AND AGRICULTURAL ENGINEERING DEPT., AGRICULTURAL ENGINEERING BLDG. (08), ROOM 101, SAN LUIS OBISPO, CALIFORNIA 93407

Email address: mhaberla@calpoly.edu

DEPARTMENT OF MATHEMATICS, MICHIGAN STATE UNIVERSITY, 619 RED CEDAR ROAD, C212 WELLS HALL, EAST LANSING, MI 48824

Email address: turanova@msu.edu