

THE BACK-AND-FORTH METHOD FOR WASSERSTEIN GRADIENT FLOWS

MATT JACOBS, WONJUN LEE, AND FLAVIEN LÉGER

ABSTRACT. We present a method to efficiently compute Wasserstein gradient flows. Our approach is based on a generalization of the back-and-forth method (BFM) introduced in [JL20] to solve optimal transport problems. We evolve the gradient flow by solving the dual problem to the JKO scheme. In general, the dual problem is much better behaved than the primal problem. This allows us to efficiently run large scale gradient flows simulations for a large class of internal energies including singular and non-convex energies.

1. INTRODUCTION

In this work, we are interested in simulating the evolution of parabolic equations of the form

$$(1.1) \quad \begin{aligned} \partial_t \rho - \nabla \cdot (\rho \nabla \phi) &= 0, \\ \phi &= \delta U(\rho). \end{aligned}$$

Equation (1.1), often referred to as Darcy's law or the generalized porous medium equation, describes the evolution of a mass density ρ flowing along a pressure gradient $\nabla \phi$ generated by an internal energy functional U . This class of equations models various physical phenomena such as fluid flow, heat transfer, aggregation-diffusion, and crowd motion [Váz07, San15]. In general, these equations are both stiff and non-linear making them challenging to solve numerically. For example, in the important special case where $U(\rho) = \frac{1}{m-1} \int \rho^m$ ($m > 1$), equation (1.1) becomes a non-linear version of the heat equation

$$\partial_t \rho - \Delta(\rho^m) = 0,$$

known as the porous medium equation (PME). When U is non-differentiable or non-convex, simulation of these equations becomes even more difficult. Thus, in this paper, our goal is to design a method that can efficiently and accurately simulate equation (1.1) for a wide variety of internal energies U .

Our approach to simulating Darcy's law is based on the celebrated interpretation of equation (1.1) as a gradient flow with respect to the Wasserstein metric [JKO98, Ott01]. This interpretation can be used to create a discrete-in-time approximation scheme known as the JKO scheme [JKO98]. The

scheme constructs approximate solutions by iterating

$$(1.2) \quad \rho^{(n+1)} := \operatorname{argmin}_{\rho} U(\rho) + \frac{1}{2\tau} W_2^2(\rho, \rho^{(n)}).$$

Here, τ plays the role of the time step in the scheme and $W_2(\cdot, \cdot)$ is the 2-Wasserstein metric from the theory of optimal transportation [San15] (see Section 2.1 for a brief overview of optimal transport and the 2-Wasserstein metric). Thanks to the variational structure of the scheme, the iterates are unconditionally energy stable and one can choose the time step τ independently from any spatial discretization. In addition, the JKO scheme retains many desirable properties of the continuum equation, such as comparison and contraction type principles [JKT20, DPMSV16, AKY14].

In light of the many favorable properties of the JKO scheme, there have been many works devoted to the computation of minimizers for problem (1.2), see [BCW10, CM10, Pey15, BCMO16, BCL16, CDPS17, CCWW19, CWXY20, LMSS20] to name just a few. Despite the amount of work on this problem, it remains a challenge to efficiently solve the JKO scheme at a high resolution. The main difficulty in solving problem (1.2) lies in the handling of the Wasserstein distance term. Indeed, there is not a simple formula that gives the variation of the Wasserstein distance with respect to the density ρ . As such, essentially all methods for solving (1.2) are adaptations of algorithms for computing the Wasserstein distance between two fixed densities.

In this paper, we solve problem (1.2) by adapting the back-and-forth method (BFM) introduced in [JL20]. BFM is a state-of-the-art algorithm for computing optimal transport maps between two fixed densities. Instead of directly solving Monge's optimal transportation problem, BFM finds optimal maps by solving the associated Kantorovich dual problem. Building on this approach, rather than directly solving problem (1.2), we instead compute solutions to its dual problem. The dual problem is a concave maximization problem that produces the pressure variable at the next time step $\phi^{(n+1)}$. The optimal density variable can then easily be recovered from the pressure via the duality relation $\phi^{(n+1)} = \delta U(\rho^{(n+1)})$.

There are several advantages to solving the dual problem rather than the original primal problem. The pressure variable ϕ has better regularity than the density variable ρ . Indeed, at worst, the pressure gradient must be square integrable. As a result, the pressure is better suited to discrete approximation schemes. In addition, there is an explicit formula to compute derivatives of the dual functional, hence one can apply gradient ascent to solve the dual problem (the corresponding gradient descent scheme for the primal problem is much more difficult). Finally, the dual approach is very convenient when U encodes hard constraints (such as incompressibility of the density), as the dual problem will be unconstrained.

Leveraging the advantages of the dual problem to (1.2) and the special gradient ascent structure of BFM, we are able to rapidly and accurately solve

the JKO scheme for a large class of internal energies U . We show that the algorithm increases the value of the dual problem at every step. In particular, this analysis holds even in cases where the Hessian of U is singular and our analysis has no dependence on the size of the computational grid. As a result, we are able to simulate equation (1.1) on a much larger scale than previous methods, and we are easily able to handle difficult cases like incompressible crowd motion models with obstacles and aggregation-diffusion equations.

1.1. Overall approach. The back-and-forth method for Wasserstein gradient flows is based on solving the dual problem associated to the JKO scheme. The starting point for this analysis is Kantorovich's dual formulation of optimal transport. Given two measures μ and ν , the dual formulation of the 2-Wasserstein distance is given by

$$(1.3) \quad \frac{1}{2\tau} W_2^2(\mu, \nu) = \sup_{(\phi, \psi) \in \mathcal{C}} \int_{\Omega} \psi(x) d\mu(x) - \int_{\Omega} \phi(y) d\nu(y),$$

where we maximize over the constraint

$$\mathcal{C} := \{(\phi, \psi) \in C(\Omega) \times C(\Omega) : \psi(x) - \phi(y) \leq \frac{1}{2\tau}|x - y|^2\}.$$

Using the dual formulation of optimal transport, we can rewrite problem (1.2) as

$$\inf_{\rho} \sup_{(\phi, \psi) \in \mathcal{C}} U(\rho) + \int_{\Omega} \psi(x) d\rho^{(n)}(x) - \int_{\Omega} \phi(y) d\rho(y).$$

When U is convex, we can interchange the inf and sup to get an equivalent dual problem to (1.2):

$$(1.4) \quad \sup_{(\phi, \psi) \in \mathcal{C}} \int_{\Omega} \psi(x) d\rho^{(n)}(x) - U^*(\phi),$$

where U^* is the convex conjugate of U ,

$$U^*(\phi) := \sup_{\rho} \int_{\Omega} \phi(y) d\rho(y) - U(\rho).$$

Problem (1.4) looks difficult due to the constraint encoded by \mathcal{C} . Nevertheless, there is a very convenient way to reformulate the problem. Because $\rho^{(n)}$ is a nonnegative measure, it is favorable to choose ψ to be pointwise as large as possible. If we fix ϕ , it then follows that the corresponding largest possible choice for ψ is given by

$$(1.5) \quad \phi^c(x) := \inf_{y \in \Omega} \phi(y) + \frac{1}{2\tau}|x - y|^2.$$

Conversely, U^* is increasing with respect to ϕ (see Section 2.1), therefore, we would like to choose ϕ to be pointwise as small as possible. Thus, if we fix ψ , then the corresponding smallest choice for ϕ is given by

$$(1.6) \quad \psi^{\bar{c}}(y) := \sup_{x \in \Omega} \psi(x) - \frac{1}{2\tau}|x - y|^2.$$

Formulas (2.1) and (1.6) are known as the backward- c -transform and forward- c -transform respectively. These transforms play an essential role in optimal transport and are integral to our method. Crucially, we can use these transforms to eliminate the constraint \mathcal{C} and either one of the variables ϕ or ψ . More explicitly, problem (1.4) is equivalent to maximizing either one of the following two *unconstrained* functionals:

$$(1.7) \quad J(\phi) := \int_{\Omega} \phi^c(x) d\rho^{(n)}(x) - U^*(\phi),$$

$$(1.8) \quad I(\psi) := \int_{\Omega} \psi(x) d\rho^{(n)}(x) - U^*(\psi^{\bar{c}}).$$

Indeed, if ϕ_* is a maximizer of J and ψ_* is a maximizer of I , then we must have the relations

$$\phi_*^c = \psi_*, \quad \psi_*^{\bar{c}} = \phi_*,$$

and (ϕ_*, ψ_*) is a maximizer of (1.4). The reformulations I and J genuinely simplify the task of finding maximizers. On a regular discrete grid, the c -transform can be computed very efficiently [Luc97, JL20]. As a result, it is much more tractable to maximize I and J , rather than trying work with (1.4) directly.

We will find the maximizers ϕ_* and ψ_* by building upon the BFM algorithm introduced in [JL20]. The original BFM gives a very efficient scheme for finding the maximizers in the special case where U^* is a linear functional. Rather than focusing on either I or J , BFM simultaneously maximizes both functionals. The method proceeds by hopping back-and-forth between gradient ascent updates on J in ϕ -space and gradient ascent updates on I in ψ -space (hence the name). In between gradient steps, information in one space (ϕ -space or ψ -space) is propagated back to the other by taking a forward/backward c -transform. As noted in [JL20], the advantage of the back-and-forth approach is that certain features of the optimal solution pair (ϕ_*, ψ_*) may be easier to build in one space compared to the other. As a result, the back-and-forth method converges far more rapidly than vanilla gradient ascent methods that operate only on ϕ -space or only on ψ -space.

In order to generalize BFM to the Wasserstein gradient flow case, we need to be able to guarantee the stability of gradient ascent steps on (1.7) and (1.8) when U^* is nonlinear. In fact, for many important cases, the Hessian of U^* may have a singular component. To overcome this difficulty, we perform the gradient ascent steps in an appropriately weighted Sobolev space. The Sobolev control allows us to use Stokes' Theorem to convert boundary integrals into integrals over the full space, thus taming the singularities of U^* (see Section 3.2). As a result of this continuous analysis, the discretized scheme will have a convergence rate that is independent of the grid size. The back-and-forth method is summarized in Algorithm 1, where H is the aforementioned weighted Sobolev space.

Algorithm 1: The back-and-forth scheme for solving (1.4)

Given $\rho^{(n)}$ and ϕ_0 , iterate:

$$\begin{aligned}\phi_{k+\frac{1}{2}} &= \phi_k + \nabla_H J(\phi_k) \\ \psi_{k+\frac{1}{2}} &= (\phi_{k+\frac{1}{2}})^c \\ \psi_{k+1} &= \psi_{k+\frac{1}{2}} + \nabla_H I(\psi_{k+\frac{1}{2}}) \\ \phi_{k+1} &= (\psi_{k+1})^{\bar{c}}\end{aligned}$$

Once we have solved the dual problem, we can recover the solution to the original problem (1.2). If U is convex, then the optimal dual variable ϕ_* is related to $\rho^{(n+1)}$ through the duality relation $\rho^{(n+1)} = \delta U^*(\phi_*)$ (see Theorem 2.7 in Section 2.2). When U is not convex, the connection between (1.2) and the dual problem becomes more tenuous. Luckily, we can circumvent this difficulty using a convexity splitting scheme [Eyr98]. Indeed, if we write $U = U_1 + U_0$ where U_1 is convex and U_0 is concave, then we can replace the JKO scheme (1.2) with the modified scheme

$$(1.9) \quad \rho^{(n+1)} = \operatorname{argmin}_{\rho} U_1(\rho) + U_0(\rho^{(n)}) + (\delta U_0(\rho^{(n)}), \rho - \rho^{(n)}) + \frac{1}{2\tau} W_2^2(\rho, \rho^{(n)}).$$

It is well-known that convexity splitting retains the energy stability of a fully implicit scheme. Crucially, the energy term $U_1(\rho) + U_0(\rho^{(n)}) + (\delta U_0(\rho^{(n)}), \rho - \rho^{(n)})$ in (1.9) is a convex function of the variable ρ , and thus, we can apply the duality approach. All together, our method gives an extremely rapid way to simulate the PDE (1.1) even when U is non-convex or irregular.

The remainder of the paper is organized as follows. In Section 2, we review important background information on optimal transport, convex analysis, and optimization. In Section 3, we present the back-and-forth algorithm and explain how to guarantee stability and choose step sizes. Lastly, in Section 4, we demonstrate the accuracy, speed, and versatility of the algorithm through a wide suite of numerical experiments. In particular, our experiments include many cases that are well-known to be numerically challenging.

2. BACKGROUND

In this section, we will rigorously establish the connection between the primal and dual formulations of the JKO scheme. Furthermore, we will review key concepts from optimal transport and convex analysis that are needed to compute the gradients $\nabla_H J, \nabla_H I$ and establish stability of Algorithm 1. Note that throughout the paper we shall assume that $\Omega \subset \mathbb{R}^d$ is a bounded open set.

2.1. The c -transform and optimal transport. Throughout this section the space of continuous functions over Ω will be denoted by $C(\Omega)$.

Definition 2.1. Given $\phi \in C(\Omega)$ its *backward c-transform* is

$$(2.1) \quad \phi^c(x) := \inf_{y \in \Omega} \phi(y) + \frac{1}{2\tau}|x - y|^2.$$

Given $\psi \in C(\Omega)$ its *forward c-transform* is

$$(2.2) \quad \psi^{\bar{c}}(y) := \sup_{x \in \Omega} \psi(x) - \frac{1}{2\tau}|x - y|^2.$$

Lemma 2.1 ([San15]). *Given $\phi, \psi \in C(\Omega)$, we have*

$$\phi^{\bar{c}\bar{c}} \leq \phi, \quad \psi \leq \psi^{\bar{c}\bar{c}},$$

and

$$\phi^{\bar{c}\bar{c}\bar{c}} = \phi^c, \quad \psi^{\bar{c}\bar{c}\bar{c}} = \psi^{\bar{c}}.$$

Definition 2.2. Given $\phi, \psi \in C(\Omega)$, we say that ϕ is *c-convex* if $\phi^{\bar{c}\bar{c}} = \phi$ and we say that ψ is *c-concave* if $\psi^{\bar{c}\bar{c}} = \psi$. Furthermore, if $\phi^c = \psi$ and $\psi^{\bar{c}} = \phi$, then we say the pair (ϕ, ψ) is *c-conjugate*.

The following two propositions establish the fundamental relationship between optimal transport and the c-transform.

Proposition 2.2 ([Gan94, Gan95b, GM96]). *If $\phi: \Omega \rightarrow \mathbb{R}$ is c-convex and $\psi: \Omega \rightarrow \mathbb{R}$ is c-concave, then the maps*

$$(2.3) \quad T_\phi(x) := \operatorname{argmin}_{y \in \Omega} \phi(y) + \frac{1}{2\tau}|x - y|^2$$

and

$$(2.4) \quad S_\psi(y) := \operatorname{argmax}_{x \in \Omega} \psi(x) - \frac{1}{2\tau}|x - y|^2$$

are well-defined and unique almost everywhere. Furthermore, if $u \in C(\Omega)$, then for almost every $x, y \in \Omega$ we have the following perturbation formulas for the c-transform

$$(2.5) \quad \lim_{t \rightarrow 0^+} \frac{(\phi + tu)^c(x) - \phi^c(x)}{t} = u(T_\phi(x)),$$

$$(2.6) \quad \lim_{t \rightarrow 0^+} \frac{(\psi + tu)^{\bar{c}}(y) - \psi^{\bar{c}}(y)}{t} = u(S_\psi(y)).$$

Finally, if ϕ and ψ are c-conjugate, then

$$S_\psi(y) = y + \tau \nabla \phi(y),$$

$$T_\phi(x) = x - \tau \nabla \psi(x),$$

and $T_\phi(S_\psi(y)) = y$, $S_\psi(T_\phi(x)) = x$ almost everywhere.

Proposition 2.3 ([San15]). *If $\mu, \nu \in L^1(\Omega)$ are nonnegative densities with the same mass, then*

$$\begin{aligned} \frac{1}{2\tau} W_2^2(\mu, \nu) &= \sup_{\phi \in C(\Omega)} \int_{\Omega} \phi^c(x) \mu(x) dx - \int_{\Omega} \phi(y) \nu(y) dy, \\ \frac{1}{2\tau} W_2^2(\mu, \nu) &= \sup_{\psi \in C(\Omega)} \int_{\Omega} \psi(x) \mu(x) dx - \int_{\Omega} \psi^{\bar{c}}(y) \nu(y) dy. \end{aligned}$$

Now we can state the fundamental result that guarantees the existence and uniqueness of the optimal transport maps.

Theorem 2.4 ([Bre91, Gan95a, GM96]). *If $\mu, \nu \in L^1(\Omega)$ are nonnegative densities with the same mass, then there exists a c -conjugate pair (ϕ_*, ψ_*) such that*

$$\begin{aligned} \phi_* &\in \operatorname{argmax}_{\phi \in C(\Omega)} \int_{\Omega} \phi^c(x) \mu(x) dx - \int_{\Omega} \phi(y) \nu(y) dy, \\ \psi_* &\in \operatorname{argmax}_{\psi \in C(\Omega)} \int_{\Omega} \psi(x) \mu(x) dx - \int_{\Omega} \psi^{\bar{c}}(y) \nu(y) dy, \\ \frac{1}{2\tau} W_2^2(\mu, \nu) &= \int_{\Omega} \psi_*(x) \mu(x) dx - \int_{\Omega} \phi_*(y) \nu(y) dy, \end{aligned}$$

and T_{ϕ_*}, S_{ψ_*} are the unique optimal transport maps sending μ to ν and ν to μ respectively, i.e. $T_{\phi_*} \# \mu = \nu$ and $S_{\psi_*} \# \nu = \mu$.

2.2. Convex duality. Now that we have developed the basics of optimal transport, we are ready to return to the JKO scheme. To iterate the JKO scheme, one must be able to solve generalized optimal transport (GOT) problems of the form

$$(2.7) \quad \rho_* = \operatorname{argmin}_{\rho \in L^1(\Omega)} U(\rho) + \frac{1}{2\tau} W_2^2(\rho, \mu),$$

where $\mu \in L^1(\Omega)$ is a given nonnegative density. Our method solves the GOT problem by appealing to its dual formulation. In the rest of this subsection, we shall derive the dual problem and develop its basic properties. To obtain a well-behaved dual problem, we shall need the following assumptions on the energy U .

Assumption 1. The internal energy is given by a proper, convex, and lower semicontinuous functional $U: L^1(\Omega) \rightarrow \mathbb{R} \cup \{+\infty\}$ such that $U(\rho) = \infty$ if ρ is negative on a set of positive measure.

Assumption 2. There exists a function $s: \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$ with superlinear growth such that

$$U(\rho) \geq \int_{\Omega} s(\rho(y)) dy.$$

Remark 2.1. Assumption 1 encodes the fact that ρ must be a nonnegative density, while Assumption 2 guarantees that for each $B \in \mathbb{R}$ the sets $\{\rho \in L^1(\Omega) : U(\rho) < B\}$ are weakly compact.

Remark 2.2. Except for the convexity requirement, Assumptions 1 and 2 are very natural in the context of Wasserstein gradient flows. Note that we will eventually consider non-convex U in Section 3.3.

At the heart of duality is the notion of convex conjugation.

Definition 2.3. Given a functional $U: L^1(\Omega) \rightarrow \mathbb{R}$ its convex conjugate $U^*: L^\infty(\Omega) \rightarrow \mathbb{R}$ is defined by

$$U^*(\phi) := \sup_{\rho \in L^1(\Omega)} \int_{\Omega} \phi(x) \rho(x) dx - U(\rho),$$

Thanks to Assumption 1, U^* possess an important monotonicity property.

Lemma 2.5. U^* is monotonically increasing, i.e. if $\phi_0, \phi_1: \Omega \rightarrow \mathbb{R}$ are functions such that $\phi_0 \leq \phi_1$ pointwise everywhere, then

$$U^*(\phi_0) \leq U^*(\phi_1).$$

Proof. By Assumption 1 the internal energy is finite only over nonnegative densities, thus,

$$U^*(\phi) = \sup_{\rho \geq 0} \int_{\Omega} \phi(x) \rho(x) dx - U(\rho).$$

If we take some $\rho \in L^1(\Omega)$, with $\rho(x) \geq 0$ a.e., then we have

$$\int_{\Omega} \phi_0(x) \rho(x) dx - U(\rho) \leq \int_{\Omega} \phi_1(x) \rho(x) dx - U(\rho).$$

Taking a supremum over $\rho \geq 0$ finishes the proof. \square

Now we are ready to reintroduce the twin dual functionals I and J .

Proposition 2.6. Fix a nonnegative density $\mu \in L^1(\Omega)$. The functionals I, J given by

$$J(\phi) := \int_{\Omega} \phi^c(x) \mu(x) dx - U^*(\phi)$$

$$I(\psi) := \int_{\Omega} \psi(x) \mu(x) dx - U^*(\psi^{\bar{c}}),$$

are proper, weakly upper semicontinuous, concave and $\sup_{\phi \in C(\Omega)} J(\phi) = \sup_{\psi \in C(\Omega)} I(\psi)$. Furthermore, if ϕ is c -convex and ψ is c -concave, then J and I have first variations

$$\delta J(\phi) = T_{\phi \#} \mu - \delta U^*(\phi),$$

$$\delta I(\psi) = \mu - S_{\psi \#} \delta U^*(\psi^{\bar{c}}),$$

where δU^* is the first variation of U^* .

Proof. Following the logic in the proof of Lemma 2.5, we may write

$$U^*(\psi^{\bar{c}}) = \sup_{\rho \geq 0} \int_{\Omega} \psi^{\bar{c}}(y) \rho(y) dy - U(\rho).$$

Next, let $\mathcal{M}(\Omega \times \Omega)$ denote the space of nonnegative measures on $\Omega \times \Omega$, and for any given density $\rho \geq 0$ define

$$\Pi(\rho) := \left\{ \pi \in \mathcal{M}(\Omega \times \Omega) : \iint_{\Omega \times \Omega} f(y) d\pi(x, y) = \int_{\Omega} f(y) \rho(y) dy \text{ for all } f \in C(\Omega) \right\}.$$

Using the definition of the c -transform, we can then write

$$\int_{\Omega} \psi^{\bar{c}}(y) \rho(y) dy = \sup_{\pi \in \Pi(\rho)} \iint_{\Omega \times \Omega} \left(\psi(x) - \frac{1}{2\tau} |x - y|^2 \right) d\pi(x, y).$$

Therefore, we have

$$-U^*(\psi^{\bar{c}}) = \inf_{\rho \geq 0} \inf_{\pi \in \Pi(\rho)} U(\rho) - \iint_{\Omega \times \Omega} \left(\psi(x) - \frac{1}{2\tau} |x - y|^2 \right) d\pi(x, y).$$

Now it is clear that I can be written as the infimum over a family of linear functionals of ψ . Hence, I must be proper, concave and weakly upper semicontinuous. An essentially identical argument applies to J .

Since U^* is monotonically increasing, Lemma 2.5 implies that for any $\phi, \psi \in C(\Omega)$

$$J(\phi) \leq I(\phi^c), \quad I(\psi) \leq J(\psi^{\bar{c}}).$$

Therefore, we must have

$$\sup_{\psi \in C(\Omega)} I(\psi) = \sup_{\phi \in C(\Omega)} J(\phi).$$

When ϕ and ψ are c -convex/concave respectively, the formulas for the first variations follow directly from Proposition 2.2. \square

Finally, we conclude this subsection by stating the essential result linking the primal and dual generalized optimal transport problems. Crucially, this shows how to recover the solution to (2.7) from the maximizers of I and J .

Theorem 2.7 ([JKT20]). *If $\mu \in L^1(\Omega)$, U satisfies Assumptions 1, 2, and $\delta U(\mu)$ is not a constant function, then there exists a unique density ρ_* and a pair of c -conjugate functions (ϕ_*, ψ_*) such that*

$$\rho_* = \operatorname{argmin}_{\rho \in L^1(\Omega)} U(\rho) + \frac{1}{2\tau} W_2^2(\rho, \mu), \quad \phi_* \in \operatorname{argmax}_{\phi \in C(\Omega)} J(\phi), \quad \psi_* \in \operatorname{argmax}_{\psi \in C(\Omega)} I(\psi),$$

$$U(\rho_*) + \frac{1}{2\tau} W_2^2(\rho_*, \mu) = J(\phi_*) = I(\psi_*),$$

$$\rho_* \in \delta U^*(\phi_*), \quad \phi_* \in \delta U(\rho_*), \quad \rho_* = T_{\phi_*} \# \mu.$$

Remark 2.3. Note that if $\delta U(\mu)$ is constant, then $\mu = \operatorname{argmin}_{\rho \in L^1(\Omega)} U(\rho) + \frac{1}{2\tau} W_2^2(\rho, \mu)$. Thus, the excluded case is trivial.

2.3. Concave gradient ascent. Now that we see how to link the JKO scheme to the dual functionals I and J , it remains to develop a method to find the maximizers of I and J . To that end, in this subsection, we review classical unconstrained gradient ascent. Let us first recall the notion of *gradient*. This will require the structure of a real Hilbert space \mathcal{H} with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and norm $\|\cdot\|_{\mathcal{H}}$.

Definition 2.4. Given a point $\varphi \in \mathcal{H}$, we say that a bounded linear map $\delta F(\varphi): \mathcal{H} \rightarrow \mathbb{R}$ is the first variation (Fréchet derivative) of F at φ if

$$\lim_{\|h\|_{\mathcal{H}} \rightarrow 0} \frac{\|F(\varphi + h) - F(\varphi) - \delta F(\varphi)(h)\|_{\mathcal{H}}}{\|h\|_{\mathcal{H}}} = 0.$$

Definition 2.5. We say that a map $\nabla_{\mathcal{H}} F: \mathcal{H} \rightarrow \mathcal{H}$ is the \mathcal{H} -gradient of F (or simply gradient if there is no ambiguity about the space \mathcal{H}) if

$$\langle \nabla_{\mathcal{H}} F(\varphi), h \rangle_{\mathcal{H}} = \delta F(\varphi)(h)$$

for all $(\varphi, h) \in \mathcal{H} \times \mathcal{H}$.

The above identity highlights that gradients are intimately linked to the inner product of the Hilbert space, in contrast to first variations. Indeed, note that one can define the notion of a first variation over any normed vector space, while the notion of a gradient requires an inner product.

Gradient ascent method. Given a concave functional J over \mathcal{H} , consider the gradient ascent iterations

$$(2.8) \quad \phi_{k+1} = \phi_k + \nabla_{\mathcal{H}} J(\phi_k).$$

The gradient ascent scheme (2.8) can equivalently be written in the variational form

$$(2.9) \quad \phi_{k+1} = \operatorname{argmax}_{\phi} J(\phi_k) + \delta J(\phi_k)(\phi - \phi_k) - \frac{1}{2} \|\phi - \phi_k\|_{\mathcal{H}}^2.$$

Note that equations (2.8) and (2.9) typically include a step size parameter that controls how far one travels in the gradient direction. For reasons that will become clear shortly (see equation (2.13) and the subsequent discussion), we prefer to incorporate any parameters into the norm $\|\cdot\|_{\mathcal{H}}$ itself.

In order to obtain convergence of the scheme

$$J(\phi_k) \xrightarrow[k \rightarrow \infty]{} \sup_{\phi} J(\phi),$$

with an efficient rate, it is essential to choose the norm $\|\cdot\|_{\mathcal{H}}$ properly. If the norm is too weak, then the algorithm may become unstable and fail to converge. On the other hand, if the norm is too strong, then very little change happens at each step and the algorithm converges slowly. The following theorem, one of the cornerstones of optimization, explains how to balance these competing considerations.

Theorem 2.8 ([Nes13]). *Let $J: \mathcal{H} \rightarrow \mathbb{R}$ be a twice Fréchet-differentiable concave functional with maximizer ϕ^* . If*

$$(2.10) \quad -\delta^2 J(\phi)(h, h) \leq \|h\|_{\mathcal{H}}^2,$$

for all $\phi, h \in \mathcal{H}$ (J is said to be “1-smooth”), then the gradient ascent scheme

$$\phi_{k+1} = \phi_k + \nabla_{\mathcal{H}} J(\phi_k)$$

starting at a point ϕ_0 satisfies the ascent property

$$(2.11) \quad J(\phi_{k+1}) \geq J(\phi_k) + \frac{1}{2} \|\nabla_{\mathcal{H}} J(\phi_k)\|_{\mathcal{H}}^2,$$

and has the convergence rate

$$(2.12) \quad J(\phi^*) - J(\phi_k) \leq \frac{\|\phi^* - \phi_0\|_{\mathcal{H}}^2}{2k}.$$

From Theorem 2.8, we can again see the competing interests of weakening or strengthening the norm $\|\cdot\|_{\mathcal{H}}$. A stronger norm makes it easier to satisfy equation (2.10), while a weaker norm gives a better convergence rate in (2.12). Putting these considerations together, we see that it is optimal to choose the weakest possible norm such that (2.10) holds.

Sobolev norm. Let Ω be an open bounded convex subset of \mathbb{R}^d . Our gradient ascent schemes use a norm H based on the Sobolev space $H^1(\Omega)$. For two constants $\Theta_1 > 0$ and $\Theta_2 > 0$ we define

$$(2.13) \quad \|h\|_H^2 = \int_{\Omega} \Theta_2 |\nabla h(x)|^2 + \Theta_1 |h(x)|^2 dx.$$

The precise value of Θ_1 and Θ_2 will depend on the functional being maximized (see for instance Theorem 3.3 in Section 3). In many instances, it will be optimal to choose Θ_1 and Θ_2 to have rather different values. For this reason, we do not wish to reduce these parameters to a single step size value. The next lemma describes how to compute gradients with respect to this inner product.

Lemma 2.9. *Suppose that $F = F(\phi)$ is a Fréchet-differentiable functional such that for any ϕ the first variation $\delta F(\phi)$ evaluated at any point h can be written as integration against a function f_{ϕ} , i.e.*

$$\delta F(\phi)(h) = \int_{\Omega} h(x) f_{\phi}(x) dx.$$

Define $\|\cdot\|_H$ by (2.13). Then the H -gradient of F can be written

$$\nabla_H F(\phi) = (\Theta_1 \text{Id} - \Theta_2 \Delta)^{-1} f_{\phi},$$

where Id is the identity operator and Δ is the Laplacian operator, taken together with zero Neumann boundary conditions.

Proof. Fix ϕ and consider the unique solution to the elliptic equation

$$\begin{cases} (\Theta_1 \text{Id} - \Theta_2 \Delta)g &= f_\phi & \text{in } \Omega, \\ n \cdot \nabla g &= 0 & \text{on } \partial\Omega. \end{cases}$$

Then we have the chain of equalities

$$\begin{aligned} \delta F(\phi)(h) &= \int_{\Omega} h(x) f_\phi(x) dx \\ &= \int_{\Omega} h(x) (\Theta_1 \text{Id} - \Theta_2 \Delta)g(x) dx \\ &= \int_{\Omega} \Theta_1 h(x) g(x) + \Theta_2 \nabla h(x) \cdot \nabla g(x) dx \\ &= \langle h, g \rangle_H. \end{aligned}$$

This shows that g is the H -gradient of F . \square

The above result can be restated as follows: the H -gradient of F is obtained by “preconditioning” δF with the inverse operator $(\Theta_1 \text{Id} - \Theta_2 \Delta)^{-1}$.

3. THE BACK-AND-FORTH METHOD

Our goal is to develop an efficient algorithm for solving the JKO scheme for a large class of interesting energies U . We begin in Section 3.1 with the case where U is convex with respect to ρ . In this case, the JKO scheme has an equivalent dual problem that we solve using an adaptation of the back-and-forth method from [JL20]. In Section 3.2, we show that the algorithm is gradient stable in a properly weighted H^1 space for convex energies of the form

$$U(\rho) = \int_{\Omega} u_m(\rho(x)) + V(x)\rho(x) dx,$$

where $V: \Omega \rightarrow [0, +\infty]$ is a fixed function, and

$$(3.1) \quad u_m(\rho) = \begin{cases} \frac{\gamma}{m-1}(\rho^m - \rho) & \text{if } \rho \geq 0, \\ +\infty & \text{otherwise,} \end{cases}$$

for some constants $\gamma > 0$ and $m > 1$. We shall also consider the two limiting cases $m \rightarrow 1$ and $m \rightarrow \infty$. Let us note that our analysis can easily be extended to more general functionals, however, we focus on the (important) special case above for clarity of exposition. After we have developed the method for convex energy functionals U , in Section 3.3 we show how to generalize the algorithm for non-convex U .

3.1. The back-and-forth method for convex U . To iterate the JKO scheme, we must be able to solve the generalized optimal transport (GOT) problem

$$(3.2) \quad \rho_* = \operatorname{argmin}_{\rho \in L^1(\Omega)} U(\rho) + \frac{1}{2\tau} W_2^2(\rho, \mu),$$

for any fixed nonnegative density $\mu \in L^1(\Omega)$. As we saw in Section 2 (see Theorem 2.7), when U is convex, the generalized optimal transport problem is in duality with the twin functionals I and J , i.e.

$$\inf_{\rho \in L^1(\Omega)} U(\rho) + \frac{1}{2\tau} W_2^2(\rho, \mu) = \sup_{\phi} J(\phi) = \sup_{\psi} I(\psi).$$

Recall I and J are given by

$$(3.3) \quad J(\phi) = \int_{\Omega} \phi^c(x) \mu(x) dx - U^*(\phi),$$

$$(3.4) \quad I(\psi) = \int_{\Omega} \psi(x) \mu(x) dx - U^*(\psi^{\bar{c}}).$$

Furthermore, the minimizer ρ_* of problem (3.2) is related to the maximizers ϕ_*, ψ_* through the relations

$$(3.5) \quad \rho_* = T_{\phi_* \#} \mu, \quad \rho_* \in \delta U^*(\phi_*), \quad \phi_*^c = \psi_*.$$

Both I and J are unconstrained concave functionals (see Proposition 2.6), therefore, it is now clear that one can find the maximizer of either functional via standard gradient ascent methods. On the other hand, choosing to work with solely I or solely J breaks the symmetry of the problem. Thus, rather than focusing on only one of the functionals, the back-and-forth method performs alternating gradient ascent steps on I and J . Although I and J use different variables, we can switch between ϕ and ψ by using the c -transform. As noted in [JL20], the alternating steps on I and J substantially accelerate the convergence rate of the method beyond standard gradient ascent.

We are now ready to introduce our approach to find the twin dual maximizers (ϕ_*, ψ_*) to problem (3.2). The method is outlined in Algorithm 2 and is based on two main ideas:

1. A back-and-forth update scheme, alternating between gradient ascent steps on I and J .
2. Gradient ascent steps in an H^1 -type norm H , with

$$\begin{aligned} \nabla_H J(\phi) &= (\Theta_1 \text{Id} - \Theta_2 \Delta)^{-1} \left[T_{\phi \#} \mu - \delta U^*(\phi) \right], \\ \nabla_H I(\psi) &= (\Theta_1 \text{Id} - \Theta_2 \Delta)^{-1} \left[\mu - S_{\psi \#}(\delta U^*(\psi^{\bar{c}})) \right]. \end{aligned}$$

Algorithm 2: The back-and-forth scheme for solving (3.3) and (3.4)

Given μ and ϕ_0 , iterate:

$$\begin{aligned} \phi_{k+\frac{1}{2}} &= \phi_k + \nabla_H J(\phi_k) \\ \psi_{k+\frac{1}{2}} &= (\phi_{k+\frac{1}{2}})^c \\ \psi_{k+1} &= \psi_{k+\frac{1}{2}} + \nabla_H I(\psi_{k+\frac{1}{2}}) \\ \phi_{k+1} &= (\psi_{k+1})^{\bar{c}} \end{aligned}$$

Our ultimate goal is to show that each step of Algorithm 2 increases the value of the functionals J and I . Thanks to Lemmas 2.1 and 2.5 it is easy to check that

$$J(\phi_{k+\frac{1}{2}}) \leq I((\phi_{k+\frac{1}{2}})^c), \quad I(\psi_{k+1}) \leq J((\psi_{k+1})^c).$$

Thus, we see that the alternating steps where we switch between the ϕ and ψ variables can only increase the values of the dual problems. To show that the gradient steps $\phi_{k+\frac{1}{2}} = \phi_k + \nabla_H J(\phi_k)$ and $\psi_{k+1} = \psi_{k+\frac{1}{2}} + \nabla_H I(\psi_{k+\frac{1}{2}})$ increase the values of J and I respectively requires a more detailed analysis, which will be the main focus of Section 3.2. As we shall see, the enhanced stability provided by the H^1 preconditioner $(\Theta_1 \text{Id} - \Theta_2 \Delta)^{-1}$ will be essential to ensure that the gradient steps have the ascent property.

Once the dual problems I and J have been solved to sufficient accuracy, one can recover the optimal density ρ_* in (3.2) through the duality relations in (3.5). In certain examples, such as incompressible flows, the subdifferential δU^* may be multivalued. When this happens, the relation $\rho_* \in \delta U^*(\phi_*)$ does not uniquely define ρ_* . In practice, however, δU^* is typically only multivalued on a single level set of ϕ_* which has zero measure. As a result, for numerical purposes, we can simply identify $\rho_* = \delta U^*(\phi_*)$. Note that it is advantageous to recover ρ_* in this way as opposed to the pushforward relation $\rho_* = T_{\phi_*} \# \mu$. Indeed, the formula $\rho_* = T_{\phi_*} \# \mu$ requires the computation of numerical derivatives of ϕ_* , while the duality relation $\rho_* \in \delta U^*(\phi_*)$ is derivative free.

Combining our work, we obtain an algorithm for evolving the JKO scheme.

Algorithm 3: Running the JKO scheme

Given initial data $\rho^{(0)}$, initialize $\phi^{(0)} = \delta U(\rho^{(0)})$.
for $n = 0, \dots, N$ **do**
 | $\phi^{(n+1)} \leftarrow$ Run Algorithm 2 with $\mu = \rho^{(n)}$ and $\phi_0 = \phi^{(n)}$.
 | $\rho^{(n+1)} = \delta U^*(\phi^{(n+1)})$.
end

3.2. H^1 gradient ascent. In order to ensure stability of the gradient ascent steps, the gradients of I and J are computed in a metric based on the H^1 Sobolev norm. Given two constants $\Theta_1 > 0$, $\Theta_2 > 0$, we define the Hilbert norm H by

$$(3.6) \quad \|h\|_H^2 = \int_{\Omega} \Theta_2 |\nabla h(x)|^2 + \Theta_1 |h(x)|^2 dx.$$

The main steps of the back-and-forth scheme are the gradient ascent steps in the H norm

$$\phi_{k+\frac{1}{2}} = \phi_k + \nabla_H J(\phi_k)$$

and

$$\psi_{k+1} = \psi_{k+\frac{1}{2}} + \nabla_H I(\psi_{k+\frac{1}{2}}).$$

In order to obtain convergence of our method, we want these steps to increase the values of the concave functionals J and I respectively. The so-called gradient ascent property

$$\begin{aligned} J(\phi_{k+\frac{1}{2}}) - J(\phi_k) &\geq \frac{1}{2} \|\nabla_H J(\phi_k)\|_H^2, \\ I(\psi_{k+1}) - I(\psi_{k+\frac{1}{2}}) &\geq \frac{1}{2} \|\nabla_H I(\psi_{k+\frac{1}{2}})\|_H^2, \end{aligned}$$

can be obtained when the Hessian bounds

$$(3.7) \quad \begin{aligned} -\delta^2 J(\phi)(h, h) &\leq \|h\|_H^2, \\ -\delta^2 I(\psi)(h, h) &\leq \|h\|_H^2 \end{aligned}$$

are satisfied (c.f. Theorem 2.8 in Section 2.3). When (3.7) holds, I and J are said to be “1-smooth” with respect to H .

We shall devote the rest of this subsection to obtaining inequalities of the form (3.7). Specifically, we shall show how to choose the constants Θ_1 and Θ_2 in equation (3.6) to ensure that I and J are 1-smooth (under regularity assumptions on ϕ and ψ) when U has the form

$$(3.8) \quad U(\rho) = \int_{\Omega} u_m(\rho(x)) dx + \int_{\Omega} V(x)\rho(x) dx,$$

where u_m is defined in (3.1) and $V: \Omega \rightarrow [0, +\infty]$ is some given function.

Crucially, we will give upper bounds on Θ_1 and Θ_2 that can be efficiently computed from the data. Obtaining tight bounds for Θ_1 and Θ_2 is important as they essentially control the step size of the algorithm (note that small values of Θ_1 and Θ_2 correspond to large gradient steps). As we explained in Section 2.3, it is optimal to choose the smallest values of Θ_1 and Θ_2 such that (3.7) holds. This analysis is actually practical, as our numerical experiments confirm that the convergence of BFM can be substantially accelerated by making good choices for Θ_1 and Θ_2 .

Those who are interested in the analysis of these bounds can continue reading this section, otherwise, one can immediately jump to the statements of Theorems 3.3 and 3.4, which give approximately optimal values of Θ_1 and Θ_2 for the functionals I and J .

3.2.1. Hessian bound analysis. It turns out that the Hessian bound analysis is nearly identical for I and J . Therefore, we will primarily focus on the analysis for J , and we will later explain how to deal with I in a similar fashion. To obtain Hessian bounds on $J(\phi) = \int_{\Omega} \phi^c \mu - U^*(\phi)$, we first derive bounds on the c -transform term

$$(3.9) \quad F(\phi) := \int_{\Omega} \phi^c(x) \mu(x) dx,$$

and then on the internal energy term $U^*(\phi)$. Let us begin by providing an expression for $\delta^2 F(\phi)$, the Hessian of F at a point ϕ that is c -convex.

Lemma 3.1 (Hessian bounds on the c -transform). *Let F be the functional defined in (3.9). If ϕ is a c -convex function, then the Hessian of F at ϕ can be written as*

$$\delta^2 F(\phi)(h, h) = -\tau \int_{\Omega} \nabla h(y) \cdot \operatorname{cof}(I_{d \times d} + \tau D^2 \phi(y)) \nabla h(y) \mu(y + \tau \nabla \phi(y)) dy,$$

where $\operatorname{cof}(I_{d \times d} + \tau D^2 \phi(y))$ denotes the cofactor matrix of $I_{d \times d} + \tau D^2 \phi(y)$. Furthermore, if the eigenvalues of $I_{d \times d} + \tau D^2 \phi(y)$ are bounded above by some constant Λ for every $y \in \Omega$, then we have the bound

$$(3.10) \quad -\delta^2 F(\phi)(h, h) \leq \tau \|\mu\|_{L^\infty} \Lambda^{d-1} \|\nabla h\|_{L^2}^2.$$

The proof of Lemma 3.1 can be found in the appendix. To gain some insight into the bound (3.10), note that given a positive definite symmetric matrix $M \in \mathbb{R}^{d \times d}$ with eigenvalues $\{\lambda_1, \dots, \lambda_d\}$, the eigenvalues of $\operatorname{cof}(M)$ are $\{\frac{\det(M)}{\lambda_1}, \dots, \frac{\det(M)}{\lambda_d}\}$. This produces the $d - 1$ degree scaling of Λ^{d-1} . To understand the meaning of Λ itself better, recall that the optimal primal variable ρ_* is given by $T_{\phi_*} \# \mu = \mu(y + \tau \nabla \phi_*(y)) \det(I_{d \times d} + \tau D^2 \phi(y))$. Hence, the eigenvalues of $I_{d \times d} + \tau D^2 \phi$ roughly measure how concentrated the mass of ρ_* is compared to μ . Since one expects the difference between ρ_* and μ to be on the order of τ , it is reasonable to expect that Λ will be close to 1.

We now turn our attention to bounding the Hessian of the internal energy term $U^*(\phi)$. When U takes the form (3.8), its convex conjugate can be written as

$$U^*(\phi) = \int_{\Omega} u_m^*(\phi(x) - V(x)) dx,$$

where

$$u_m^*(p) = \gamma^{-\frac{1}{m-1}} \left(\frac{(m-1)p + \gamma}{m} \right)_+^{\frac{m}{m-1}}$$

and $(\cdot)_+ = \max(\cdot, 0)$. Now it is clear that the Hessian of U^* is given by

$$(3.11) \quad \delta^2 U^*(\phi)(h, h) = \int_{\Omega} (u_m^*)''(\phi(x) - V(x)) |h(x)|^2 dx.$$

When $1 \leq m \leq 2$, the bounds are straightforward as $(u_m^*)''(p)$ is increasing with respect to p . Hence, in this case, we have

$$\delta^2 U^*(\phi)(h, h) = \int_{\Omega} (u_m^*)''(\phi(x) - V(x)) |h(x)|^2 dx \leq B \|h\|_{L^2(\Omega)}^2,$$

where $B = \sup_{x \in \Omega} (u_m^*)''(\phi(x) - V(x))$. It was shown in [JKT20] that the maximizer ϕ_* of J obeys a maximum type principle in the sense that

$$\phi_*(x) \leq M := \sup_{x \in \Omega} \delta U(\mu)(x).$$

It is therefore natural to assume that ϕ will be bounded above by M throughout the algorithm (the gradient steps tend to diffuse pressure in the regions of highest concentration). Assuming $V(x) \geq 0$ everywhere, it now follows that

$$\delta^2 U^*(\phi)(h, h) \leq (u_m^*)''(M) \|h\|_{L^2(\Omega)}^2.$$

The aforementioned maximum principle on the pressure, $\phi(x) \leq M$, can be used again to write the upper bound in terms of density instead of pressure. Indeed note that

$$\rho(x) = (u_m^*)'(\phi(x) - V(x)) \leq (u_m^*)'(\phi(x)) \leq (u_m^*)'(M).$$

Therefore the quantity

$$(3.12) \quad \rho_{\max} := (u_m^*)'(M)$$

acts a natural upper bound on the densities. Furthermore writing $(u_m^*)''(M) = (u_m^*)''(u_m'(\rho_{\max})) = u_m''(\rho_{\max})^{-1}$, we obtain

$$\delta^2 U^*(\phi)(h, h) \leq u_m''(\rho_{\max})^{-1} \|h\|_{L^2(\Omega)}^2.$$

The case $m > 2$ is substantially more complicated. When $m > 2$, $(u_m^*)''$ is singular at zero. Hence, the integrand may be unbounded near points where $\phi(x) = V(x)$. In this case, it may not be possible to bound (3.11) in terms of the L^2 norm of h . To understand this better, let us focus on the most difficult model we consider in this paper: the incompressible limit $m \rightarrow \infty$. When $m \rightarrow \infty$, the energy u_m encodes a hard ceiling constraint on the density values, i.e.

$$u_\infty(\rho) = \begin{cases} 0 & \text{if } 0 \leq \rho \leq 1, \\ +\infty & \text{otherwise.} \end{cases}$$

Hence, the dual energy u_∞^* is given by

$$u_\infty^*(p) = \begin{cases} 0 & \text{if } p < 0 \\ p & \text{if } p \geq 0. \end{cases}$$

We pause here to point out that u_∞^* has much better regularity than u_∞ , for instance u_∞^* is continuous over \mathbb{R} while u_∞ is discontinuous at 0 and 1. This illustrates once more the advantage of working with dual quantities. Nevertheless, u_∞^* is clearly not smooth in the convex sense, as there is a jump of derivative at 0. In fact, we have $(u_\infty^*)'' = \delta_0$, where δ_0 denotes the Dirac delta function at 0.

Luckily, even though U^* is built from u_∞^* which is not smooth, it is possible to bound the Hessian of U^* as long as the singularity only occurs on a small set. Indeed, if we make the assumption that $|\nabla\phi(x) - \nabla V(x)|$ stays away from zero on the surface $\{\phi = V\}$, i.e. there exists a constant $\Gamma_0 > 0$ such that

$$\sup_{\{x \in \Omega: \phi(x) = V(x)\}} \frac{1}{|\nabla\phi(x) - \nabla V(x)|} \leq \Gamma_0$$

(note this is a quantitative way of saying that $\{\phi = V\}$ is a lower dimensional set), then we can use the coarea formula to rewrite equation (3.11) as

$$(3.13) \quad \begin{aligned} \delta^2 U^*(\phi)(h, h) &= \int_{\{\phi=V\}} \frac{|h(x)|^2}{|\nabla\phi(x) - \nabla V(x)|} ds(x) \\ &\leq \Gamma_0 \int_{\{\phi=V\}} |h(x)|^2 ds(x), \end{aligned}$$

where ds is the usual surface measure. Due to the fact that the integration occurs over a surface, we cannot bound the right hand side of (3.13) in terms of $\|h\|_{L^2}$. However, we can use *trace inequalities* from PDE theory to bound surface integrals by volume integrals involving a higher derivative [Eva10] (this can be essentially viewed as an inequality version of Stokes' Theorem). More precisely, there exist constants C_1, C_2 depending on the surface $\{\phi = V\}$, but independent of h such that

$$\int_{\{\phi=V\}} |h(x)|^2 ds(x) \leq C_2 \|\nabla h\|_{L^2(\Omega)}^2 + C_1 \|h\|_{L^2(\Omega)}^2.$$

From there we can immediately deduce that U^* is H -smooth, since

$$\Gamma_0 \int_{\{\phi=V\}} |h(x)|^2 ds(x) \leq \|h\|_H^2$$

as long as we choose $\Theta_i \geq C_i \Gamma_0$, $i = 1, 2$.

Now that we have seen how to obtain Hessian bounds in the most singular case $m \rightarrow \infty$, we are ready to return to the case $2 < m < \infty$. Note that in this case, $(u_m^*)''(p)$ is zero if $p < 0$, singular at zero, and decreasing for $p > 0$. Hence, if we choose some value $\lambda > 0$ and let

$$A_\lambda = \{x \in \Omega : 0 \leq \phi(x) - V(x) \leq \lambda\},$$

then we immediately have the bound

$$\delta^2 U^*(\phi)(h, h) \leq (u_m^*)''(\lambda) \|h\|_{L^2(\Omega)}^2 + \int_{A_\lambda} (u_m^*)''(\phi(x) - V(x)) |h(x)|^2 dx.$$

To estimate the second term, we proceed along the same lines as the case $m = \infty$. Let $\{\phi - V = \lambda\} = \{x \in \Omega : \phi(x) - V(x) = \lambda\}$. As long as we have a constant Γ_λ and trace inequality constants $C_1(\alpha), C_2(\alpha)$ such that

$$(3.14) \quad \sup_{x \in A_\lambda} \frac{1}{|\nabla\phi(x) - \nabla V(x)|} \leq \Gamma_\lambda$$

and

$$(3.15) \quad \int_{\{\phi-V=\alpha\}} |h(x)|^2 ds(x) \leq C_2(\alpha) \|\nabla h\|_{L^2(\Omega)}^2 + C_1(\alpha) \|h\|_{L^2(\Omega)}^2,$$

then we can replicate the argument from above. Combining the coarea formula and trace inequality, we get the following string of inequalities

$$\int_{A_\lambda} (u_m^*)''(\phi(x) - V(x)) |h(x)|^2 dx$$

$$\begin{aligned} &\leq \Gamma_\lambda \int_0^\lambda (u_m^*)''(\alpha) \int_{\{\phi-V=\alpha\}} |h(x)|^2 ds(x) d\alpha \\ &\leq (u_m^*)'(\lambda) \Gamma_\lambda \left(C_{2,\lambda} \|\nabla h\|_{L^2(\Omega)}^2 + C_{1,\lambda} \|h\|_{L^2(\Omega)}^2 \right), \end{aligned}$$

where

$$(3.16) \quad C_{i,\lambda} = \max_{0 \leq \alpha \leq \lambda} C_i(\alpha).$$

Thus, $-\delta^2 U^*(h, h)$ is bounded by $\|h\|_H^2$ as long as we choose

$$\Theta_1 \geq (u_m^*)''(\lambda) + (u_m^*)'(\lambda) \Gamma_\lambda C_{1,\lambda}$$

and

$$\Theta_2 \geq (u_m^*)'(\lambda) \Gamma_\lambda C_{2,\lambda}$$

where we have the freedom to choose the precise value of λ .

Our above computations are now summarized in the following lemma.

Lemma 3.2 (Hessian bound on the internal energy). *Define ρ_{\max} , Γ_λ and $C_{i,\lambda}$ by (3.12), (3.14) and (3.16).*

1. *Case $1 \leq m \leq 2$. We have*

$$\delta^2 U^*(\phi)(h, h) \leq \frac{1}{\gamma m} (\rho_{\max})^{2-m} \|h\|_{L^2}^2.$$

2. *Case $2 < m < \infty$. For any $\lambda > 0$,*

$$\begin{aligned} \delta^2 U^*(\phi)(h, h) &\leq (\gamma m')^{1-m'} C_{2,\lambda} \Gamma_\lambda \|\nabla h\|_2^2 + \\ &\quad (\gamma m')^{1-m'} \left(C_{1,\lambda} \Gamma_\lambda \lambda^{m'-1} + (m'-1) \lambda^{m'-2} \right) \|h\|_2^2, \end{aligned}$$

where $m' = \frac{m}{m-1}$.

3. *Case $m = \infty$. We have*

$$\delta^2 U^*(\phi)(h, h) \leq C_{2,0} \Gamma_0 \|\nabla h\|_2^2 + C_{1,0} \Gamma_0 \|h\|_2^2.$$

Combining Lemma 3.1 and 3.2 we directly obtain the main theorem of this section.

Theorem 3.3 (1-smoothness of J). *Let $1 \leq m \leq \infty$, $\gamma > 0$ and $U(\rho) = \int_\Omega u_m(\rho(x)) + V(x)\rho(x) dx$, where u_m is defined by (3.1). Then $J(\phi) := \int_\Omega \phi^c(x) \mu(x) dx - U^*(\phi)$ satisfies the Hessian bound*

$$-\delta^2 J(\phi)(h, h) \leq \Theta_2 \|\nabla h\|_{L^2}^2 + \Theta_1 \|h\|_{L^2}^2,$$

where Θ_1 and $\Theta_2 > 0$ are given by the table below. As in Lemma 3.1, Λ is an upper bound on the eigenvalues of $I_{d \times d} + \tau D^2 \phi(y)$ uniformly in y . Additionally $\lambda > 0$ is a parameter to choose and ρ_{\max} , Γ_λ and $C_{i,\lambda}$ are defined by (3.12), (3.14) and (3.16).

m	Θ_1	Θ_2
$m = 1$	$\frac{\rho_{\max}}{\gamma}$	$\tau\Lambda^{d-1}\ \mu\ _{L^\infty}$
$1 < m < 2$	$\frac{\rho_{\max}^{2-m}}{\gamma m}$	$\tau\Lambda^{d-1}\ \mu\ _{L^\infty}$
$m = 2$	$\frac{1}{2\gamma}$	$\tau\Lambda^{d-1}\ \mu\ _{L^\infty}$
$m > 2$	$(\gamma m')^{1-m'} \left(\lambda^{m'-1} C_{1,\lambda} \Gamma_\lambda + \frac{m'-1}{\lambda^{2-m'}} \right)$	$(\gamma m')^{1-m'} C_{2,\lambda} \Gamma_\lambda + \tau\Lambda^{d-1}\ \mu\ _{L^\infty}$
$m = \infty$	$C_{1,0} \Gamma_0$	$C_{2,0} \Gamma_0 + \tau\Lambda^{d-1}\ \mu\ _{L^\infty}$

In order to use Theorem 3.3 in the case $m > 2$, we need to be able to compute Γ_λ and $C_{i,\lambda}$ and we need to choose a value for λ when $m \in (2, \infty)$. On a discrete grid with n points, one can easily compute Γ_λ for all λ in $O(n)$ operations. On the other hand, it requires $O(n)$ operations to compute $C_1(\alpha)$ and $C_2(\alpha)$ for a single value of α (c.f. Section 4.1). Thus, for the case $m = \infty$, we can compute the constants explicitly in $O(n)$ operations. The case $2 < m < \infty$ is harder, since we cannot efficiently compute $C_{i,\lambda} = \max_{0 \leq \alpha \leq \lambda} C_i(\alpha)$. To overcome this difficulty, we typically choose λ by minimizing

$$\lambda^* = \operatorname{argmin}_{\lambda \geq 0} (\gamma m')^{1-m'} \left(\lambda^{m'-1} \Gamma_\lambda C_1(0) + \frac{m'-1}{\lambda^{2-m'}} \right),$$

which gives a reasonable estimate for the optimal choice of λ to make Θ_1 as small as possible. We then estimate $\max_{0 \leq \alpha \leq \lambda^*} C_i(\alpha)$ by simply taking the max over $C_i(0)$ and $C_i(\lambda^*)$, which appears to work well in practice.

To conclude this discussion we turn our attention to the other functional I for which a similar analysis can be made. Define first

$$p(x) = (\psi^{\bar{c}} - V)(T_{\psi^{\bar{c}}}(x)).$$

Then for $\lambda > 0$ we consider

$$(3.17) \quad \tilde{\Gamma}_\lambda = \sup_{x: 0 \leq p(x) \leq \lambda} \frac{1}{|\nabla p(x)|}.$$

The trace constants $\tilde{C}_{i,\lambda}$ are such that

$$(3.18) \quad \int_{\{p=\alpha\}} |h(x)|^2 ds(x) \leq \tilde{C}_{2,\lambda} \|\nabla h\|_{L^2}^2 + \tilde{C}_{1,\lambda} \|h\|_{L^2}^2,$$

uniformly over $\alpha \in (0, \lambda)$.

Theorem 3.4. *Let $I(\psi) = \int_\Omega \psi(x) \mu(x) dx - U^*(\psi^{\bar{c}})$, with $U(\rho) = \int_\Omega u_m(\rho(x)) + V(x)\rho(x) dx$, u_m is defined by (3.1) and $1 \leq m \leq \infty$. The Hessian of I can*

be written

$$-\delta^2 I(\psi)(h, h) = \delta^2 U^*(\psi^c)(h \circ S_\psi, h \circ S_\psi) + \tau \int_{\Omega} \nabla h(x) \cdot \operatorname{cof}(I_{d \times d} - \tau D^2 \psi(x)) \nabla h(x) \delta U^*(\psi^c)(x - \tau \nabla \psi(x)) dx.$$

It satisfies the bound

$$-\delta^2 I(\psi)(h, h) \leq \Theta_2 \|\nabla h\|_{L^2}^2 + \Theta_1 \|h\|_{L^2}^2,$$

where Θ_1 and $\Theta_2 > 0$ are given by the table below. Here Λ is an upper bound on the eigenvalues of $I_{d \times d} - \tau D^2 \psi(x)$ uniformly in x . Additionally $\lambda > 0$ is a parameter to choose and ρ_{\max} is defined by (3.12), $\tilde{\Gamma}_\lambda$ by (3.17) and $\tilde{C}_{i,\lambda}$ by (3.18).

m	Θ_1	Θ_2
$m = 1$	$\frac{\Lambda^d \rho_{\max}}{\gamma}$	$\tau \Lambda^{d-1} \rho_{\max}$
$1 < m < 2$	$\frac{\Lambda^d (\rho_{\max})^{2-m}}{\gamma m}$	$\tau \Lambda^{d-1} \rho_{\max}$
$m = 2$	$\frac{\Lambda^d}{2\gamma}$	$\tau \Lambda^{d-1} \rho_{\max}$
$m > 2$	$\Lambda^d (\gamma m')^{1-m'} \left(\tilde{C}_{1,\lambda} \tilde{\Gamma}_\lambda \lambda^{m'-1} + \frac{m'-1}{\lambda^{2-m'}} \right)$	$\Lambda^d (\gamma m')^{1-m'} \tilde{C}_{2,\lambda} \tilde{\Gamma}_\lambda + \tau \Lambda^{d-1} \rho_{\max}$
$m = \infty$	$\Lambda^d \tilde{C}_{1,0} \tilde{\Gamma}_0$	$\Lambda^d \tilde{C}_{2,0} \tilde{\Gamma}_0 + \tau \Lambda^{d-1} \rho_{\max}$

3.3. Back-and-forth for non-convex U . In this section, we will discuss how to extend our method when U is not convex with respect to ρ . The trick is to appeal to convexity splitting [Eyr98], a well-known technique for simulating gradient flows with non-convex energies. The idea behind convexity splitting is to write U as a sum of a convex function and a concave function, i.e.

$$U(\rho) = U_1(\rho) + U_0(\rho),$$

where U_1 is convex and U_0 is concave. Thanks to the concavity of U_0 , given any fixed density $\bar{\rho}$, we have the inequality

$$(3.19) \quad U(\rho) \leq U_1(\rho) + U_0(\bar{\rho}) + (\delta U_0(\bar{\rho}), \rho - \bar{\rho}).$$

Crucially, the right-hand-side of equation (3.19) is a convex function. As such, if we replace the JKO scheme with the relaxed scheme

$$(3.20) \quad \rho^{(n+1)} = \operatorname{argmin}_{\rho} U_1(\rho) + U_0(\rho^{(n)}) + (\delta U_0(\rho^{(n)}), \rho - \rho^{(n)}) + \frac{1}{2\tau} W_2^2(\rho, \rho^{(n)}),$$

then we obtain a convex variational problem. The beauty of convexity splitting is that the relaxed scheme is still unconditionally energy stable. Combining (3.19) and (3.20) we have the string of inequalities

$$\begin{aligned} U(\rho^{(n+1)}) + \frac{1}{2\tau} W_2^2(\rho^{(n+1)}, \rho^{(n)}) &\leq \\ U_1(\rho^{(n+1)}) + U_0(\rho^{(n)}) + (\delta U_0(\rho^{(n)}), \rho^{(n+1)} - \rho^{(n)}) + \frac{1}{2\tau} W_2^2(\rho^{(n+1)}, \rho^{(n)}) &\leq \\ \inf_{\rho} U_1(\rho) + U_0(\rho^{(n)}) + (\delta U_0(\rho^{(n)}), \rho - \rho^{(n)}) + \frac{1}{2\tau} W_2^2(\rho, \rho^{(n)}) &. \end{aligned}$$

By choosing $\rho = \rho^{(n)}$ in the last line, we can conclude that

$$U(\rho^{(n+1)}) + \frac{1}{2\tau} W_2^2(\rho^{(n+1)}, \rho^{(n)}) \leq U(\rho^{(n)}).$$

Thus, we see that the energy is still decreasing along the iterates of the relaxed scheme.

Now let us turn to solving the relaxed problem (3.20). Since the energy term in (3.20) is convex, we can solve the problem using the dual approach outlined above. The twin dual problems associated to (3.20), which we shall denote as \tilde{J} and \tilde{I} , are given by

$$(3.21) \quad \tilde{J}(\phi) := \int_{\Omega} \phi^c(x) \rho^{(n)}(x) dx - \tilde{U}^*(\phi),$$

$$(3.22) \quad \tilde{I}(\psi) := \int_{\Omega} \psi(x) \rho^{(n)}(x) dx - \tilde{U}^*(\psi^c),$$

where

$$\tilde{U}^*(\phi) := U_1^*(\phi - \delta U_0(\rho^{(n)})) + (\delta U_0(\rho^{(n)}), \rho^{(n)}) - U_0(\rho^{(n)})$$

is the convex conjugate of $U_1(\rho) + U_0(\rho^{(n)}) + (\delta U_0(\rho^{(n)}), \rho - \rho^{(n)})$. We can then find the dual maximizers $(\phi^{(n+1)}, \psi^{(n+1)})$ of (3.21) and (3.22) using Algorithm 2 along with the Hessian bounds developed in the previous subsection. As before, one can recover the solution $\rho^{(n+1)}$ of (3.20) through the duality relation $\rho^{(n+1)} = \delta \tilde{U}^*(\phi^{(n+1)})$.

4. NUMERICAL IMPLEMENTATION AND EXPERIMENTS

4.1. Implementation details. In this section, we use the back-and-forth method to numerically simulate equation (1.1) for a wide variety of internal energies U . Throughout this section we will assume that the domain $\Omega = [-1/2, 1/2]^2$ is the unit square in \mathbb{R}^2 , discretized using a regular rectangular grid. The numerical simulations in this section were coded in C++ and ran on 2019 MacBook Pro with 2.6 GHz 6-core and 16 GB RAM.

Following the approach in [JL20], we will compute the forward and backward c -transforms using the fast Legendre transform (FLT) algorithm [Luc97]. On a regular rectangular grid with n points, the FLT algorithm can be used to compute either the forward or backward c -transform in $O(n)$ operations.

See [JL20] for more detail on the equivalence of the c -transform and the Legendre transform.

When computing gradients with respect to the weighted norm (2.13), we will need to solve a Poisson equation with zero Neumann boundary condition. We will solve this equation numerically via the fast Fourier transform (FFT). All FFTs were calculated using the free FFTW C++ library.

To compute the gradients of I and J , we will also need to compute pushforwards. Given a density μ and an invertible map $Z : \Omega \rightarrow \Omega$ we can compute the pushforward $Z_{\#}\mu$ via the Jacobian formula

$$Z_{\#}\mu(x) = \frac{\mu(Z^{-1}(x))}{|\det(DZ(Z^{-1}(x)))|} = \mu(Z^{-1}(x))|\det(D(Z^{-1})(x))|.$$

In our case, we will only need to compute pushforwards with respect to the maps T_{ϕ} and S_{ψ} that are induced by the forward and backward c -transforms respectively. Thanks to the structure of BFM, we only need to compute $T_{\phi\#}\rho^{(n)}$ and $S_{\psi\#}\delta U^*(\psi^{\bar{c}})$ when ϕ and ψ are c -convex and c -concave respectively. As a result, we have the simple formulas $T_{\phi}^{-1}(y) = y + \tau\nabla\phi(y)$ and $S_{\psi}^{-1}(x) = x - \tau\nabla\psi(x)$. Therefore,

$$T_{\phi\#}\rho^{(n)}(y) = \rho^{(n)}(y + \tau\nabla\phi(y)) \det(I_{d \times d} + \tau D^2\phi(y)),$$

and

$$S_{\psi\#}\delta U^*(\psi^{\bar{c}})(x) = \left(\delta U^*(\psi^{\bar{c}}) \circ (x - \tau\nabla\psi(x)) \right) \det(I_{d \times d} - \tau D^2\psi(x)).$$

When implementing our algorithm, we compute these quantities using a simple centered difference scheme.

Finally, let us briefly explain how to compute the trace inequality constants $C_i(\alpha)$ defined in equation (3.15). From Lemma A.1 and Corollary A.2 in the Appendix, we see that $C_i(\alpha)$ can be computed from the solution u to the Eikonal equation

$$\begin{cases} |\nabla u(x)| = 1 & \text{if } \phi(x) - V(x) \neq \alpha, \\ u(x) > 0 & \text{if } \phi(x) - V(x) < \alpha, \\ u(x) < 0 & \text{if } \phi(x) - V(x) > \alpha. \end{cases}$$

Note that

$$|u(x)|^2 = \min_{\{y: \phi(y) - V(y) = \alpha\}} |x - y|^2,$$

which is nothing but a c -transform of the indicator function

$$\chi_{\alpha}(y) = \begin{cases} 0 & \text{if } \phi(y) - V(y) = \alpha, \\ +\infty & \text{else.} \end{cases}$$

Therefore, $|u|^2$ can be computed in $O(n)$ operations using the Fast Legendre transform, and from there one can recover u . Once one has u , it is straightforward to compute the constants in Corollary A.2 in $O(n)$ operations.

4.2. Experiments. We present four sets of numerical experiments. In the first set of experiments, we demonstrate the speed and accuracy of our method by comparing to the so-called Barenblatt solutions, a special case of equation (1.1) where closed-form solutions are available. In the next set of experiments, we simulate the porous media equation $\partial_t \rho = \Delta(\rho^m) + \nabla \cdot (\rho \nabla V)$ for various interesting functions $V : \Omega \rightarrow \mathbb{R} \cup \{+\infty\}$ and values of m . Note that if V takes the value $+\infty$ on some closed set $E \subset \Omega$, then ρ can never enter E . Hence, this is equivalent to solving (1.1) on the more complicated domain $\Omega \setminus E$. In the third set of experiments, we use the splitting scheme from Section 3.3 to simulate (1.1) when U is nonconvex. In this case, the non-convexity will come from an interaction energy of the form $\mathcal{W}(\rho) = \int_{\Omega} \int_{\Omega} W(x-y) \rho(x) \rho(y) dy dx$. Finally, in the last set of experiments, we study incompressible flows where U encodes the hard constraint $\rho \leq 1$ everywhere. In this case, the dual energy U^* will have a very singular Hessian at the boundary of the support of ρ . Nonetheless, we are still able to simulate the evolution even on very fine grids.

4.2.1. Accuracy: Barenblatt solutions. In this experiment, we use our back-and-forth algorithm to solve the PME,

$$(4.1) \quad \partial_t \rho = \gamma \Delta(\rho^m),$$

with the initial data

$$\rho(0, x) = M \delta_0(x).$$

Here, $\gamma > 0$ is a constant that controls the speed of the diffusion, $M > 0$ is the total initial mass and δ_0 is the standard Dirac distribution centered at zero. When $m > 1$, this equation is the Wasserstein gradient flow of the energy $U(\rho) = \int_{\Omega} \frac{\gamma}{m-1} \rho(x)^m dx$. Thanks to the simplicity of the initial data, on the domain \mathbb{R}^2 the equation has a closed form solution, known as the Barenblatt solution [Bar96, Bar03],

$$(4.2) \quad \rho(t, x) = \left(\left(\frac{M}{4\pi m t \gamma} \right)^{\frac{m-1}{m}} - \frac{(m-1)}{4m^2 t \gamma} |x|^2 \right)_+^{\frac{1}{m-1}},$$

where $(\cdot)_+ = \max(\cdot, 0)$. The Barenblatt solution is compactly supported, therefore, it agrees with the solution on the square $[-1/2, 1/2]^2$ up until the time $t_c = \frac{m-1}{16m^2\gamma} \left(\frac{\pi(m-1)}{4mM} \right)^{m-1}$ when the mass hits the boundary of the square.

Using the Barenblatt solution as a benchmark, we can test the accuracy and efficiency of our scheme. We will simulate the equation for the exponents $m = 2, 4, 6$. Since the Dirac delta function is challenging to work with numerically, we shall instead fix a height $h_0 > 0$ and start the flow at a time $t_0 > 0$, where t_0 is chosen so that $\|\rho(t_0, \cdot)\|_{L^\infty} = h_0$. Note that the value of t_0 will depend on the exponent m , and can be found explicitly from equation (4.2). In addition, we will only consider the flow within in the time interval $[t_0, t_c]$, since the Barenblatt solution is only valid on the unit square up to time t_c .

In all of our benchmark experiments, we shall set $M = 0.5$ $h_0 = 15$ and $\gamma = 10^{-3}$. Note that the small value of γ is just a time rescaling to ensure that the flow occurs on a macroscopic time interval. We will compute the evolution between the times $t_0 \leq t \leq 2 + t_0$ with different step sizes $\tau = 0.4, 0.2, 0.1, 0.05, 0.025$ (one can check that with our parameter choices $t_0 + 2 < t_c$ for $m = 2, 4, 6$). Running the experiments with various time step sizes allows us to verify that the scheme becomes more accurate as the time step is decreased. We shall measure the accuracy of the solution using the L^1 norm, which is very natural in the context of Wasserstein gradient flows (see for instance [JKT20]). The precise formula for our error estimate is

$$(4.3) \quad \text{error} = \frac{1}{N_\tau} \sum_{n=0}^{N_\tau} \int_{\Omega} |\rho(n\tau + t_0, x) - \rho^{(n)}(x)| dx,$$

where $N_\tau = \lfloor \frac{2}{\tau} \rfloor$, $\rho(n\tau + t_0, x)$ is the Barenblatt solution and $\rho^{(n)}$ is the n^{th} JKO iterate starting from the initial data $\rho^{(0)}(x) = \rho(t_0, x)$. When solving for $\phi^{(n+1)}$, we will run Algorithm 2 until the residual $\|T_\phi - \delta U^*(\phi)\|_{L^1(\Omega)}$ is less than $\epsilon = 10^{-3}$.

The results of these experiments are displayed in Table 1 and Figure 1. Table 1 displays the error (4.3) and the total computation time for all of the aforementioned experiments. In Figure 1, we plot a cross section of our solutions and the exact solution at various time snapshots. The cross section is taken along the horizontal line $\{(x_1, 0) : x_1 \in [-1/2, 1/2]\}$. One can see that as the time step is decreased, our solution is in excellent agreement with the exact solution for all exponents $m = 2, 4, 6$. Figure 1 also shows that our method correctly captures the discontinuity of $\nabla \rho$ at the boundary of the support of ρ . This is notable as most other numerical methods smooth out the discontinuity. The reason that we are able to correctly capture the discontinuity is due to the fact that we recover the density through the duality relation $\rho^{(n+1)} = \delta U^*(\phi^{(n+1)}) = \left(\frac{m-1}{m\gamma} \max(\phi, 0)\right)^{\frac{1}{m-1}}$. The function $s(x) = \max(x, 0)^{\frac{1}{m-1}}$ has discontinuous derivatives at zero, therefore even when $\phi^{(n+1)}$ is smooth, $\nabla \rho$ will still have a discontinuity at the boundary of its support.

4.2.2. *Slow diffusion with drifts and obstacles.* In our next set of experiments, we add spatially varying potentials to the energy functional. The resulting equations are a type of drift-diffusion equations. The energy takes the specific form

$$U(\rho) = \int_{\Omega} \frac{\gamma}{m-1} \rho^m(x) + V(x)\rho(x) dx,$$

where V is a given function.

In the first set of experiments, we consider an example where the initial density is the characteristic function of a star shaped region normalized to

TABLE 1. Barenblatt solution test case

τ	N_τ	$m = 2$		$m = 4$		$m = 6$	
		Error	Time (s)	Error	Time (s)	Error	Time (s)
0.4	5	6.35×10^{-2}	14.54	1.19×10^{-1}	23.11	1.13×10^{-1}	22.02
0.2	10	3.72×10^{-2}	22.16	7.95×10^{-2}	30.34	7.48×10^{-1}	30.41
0.1	20	2.08×10^{-2}	36.57	5.03×10^{-2}	48.41	4.74×10^{-2}	43.95
0.05	40	1.18×10^{-2}	55.64	3.06×10^{-2}	77.03	2.90×10^{-2}	80.10
0.025	80	8.26×10^{-3}	77.67	1.89×10^{-2}	140.38	1.79×10^{-2}	164.89

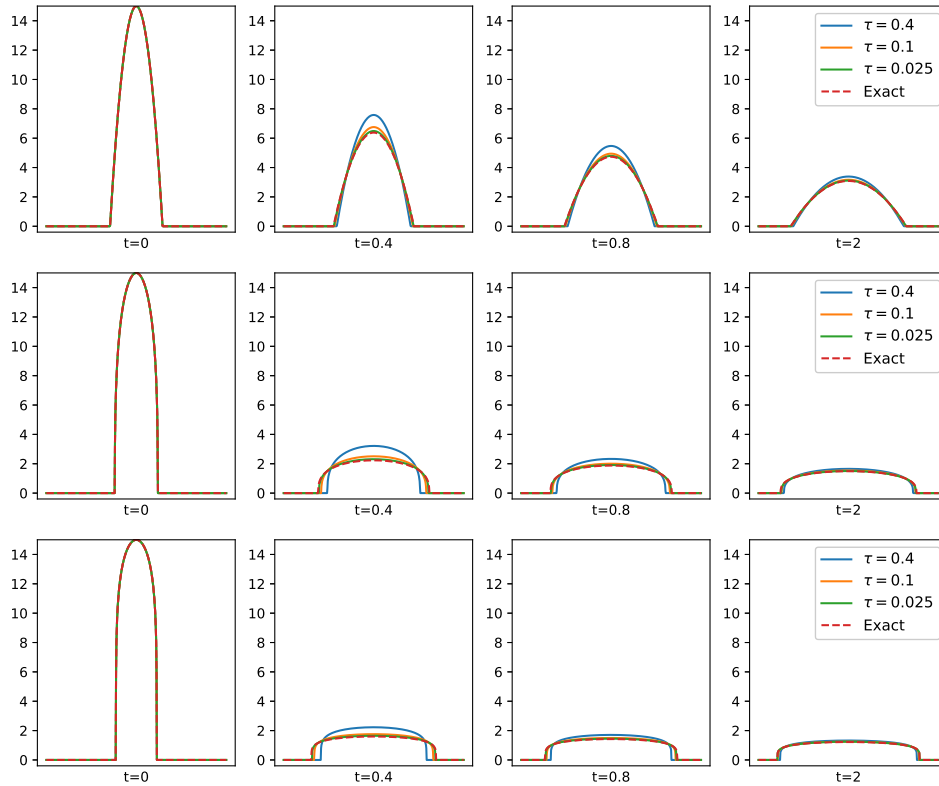


FIGURE 1. Cross sections of our computed solutions and the exact Barenblatt solution at times $t = t_0, t_0 + 0.4, t_0 + 0.8, t_0 + 2$ along the horizontal line $\{(x_1, 0) : x_1 \in [-1/2, 1/2]\}$. Row 1: $m = 2$, Row 2: $m = 4$, Row 3: $m = 6$.

have mass 1, and we use the fixed potential function

$$(4.4) \quad V_1(x) = 1 - \sin(5\pi x_1) \sin(3\pi x_2).$$

The initial data and the potential V_1 are shown in Figure 2.




FIGURE 2. Higher values are depicted with brighter pixels.

Using this setup, we run two different experiments, one where $m = 2$ and another where $m = 4$. In both cases, we set $\gamma = 0.1$ and use the time step $\tau = 0.001$. We run the equations until we reach a state that is essentially stationary. The flow for $m = 2$ is run from time $t = 0$ to time $t = 5$, and the flow for $m = 4$ is run from time $t = 0$ to time $t = 2$. The flow for the $m = 2$ case is shown in Figure 3 and the $m = 4$ case is shown in Figure 4. The solutions show the density is drawn to regions where the potential is small, while avoiding concentration due to the ρ^m term. Notice that the steady state for $m = 4$ is much more diffuse than the steady state for $m = 2$, this is because ρ^4 penalizes concentration much more than ρ^2 .




FIGURE 3. PME with exponent $m = 2$ and potential given by (4.4). The images show the evolution from time $t = 0$ to $t = 5$ (top left to bottom right). The final image is the approximate steady state. Images are 512×512 pixels. Brighter pixels indicate larger density values.

Next, we consider a different potential function:

$$(4.5) \quad V_2(x) = 10 \left((x_1 - 0.4)^2 + (x_2 - 0.4)^2 \right) + \iota_{\Omega \setminus E}(x)$$




FIGURE 4. PME with exponent $m = 4$ and potential given by (4.4). The images show the evolution from time $t = 0$ to $t = 2$ (top left to bottom right). The final image is the approximate steady state. Images are 512×512 pixels. Brighter pixels indicate larger density values.

where E is a given subset of Ω and $\iota_{\Omega \setminus E} : \Omega \rightarrow \mathbb{R} \cup \{+\infty\}$ is the indicator function

$$\iota_{\Omega \setminus E}(x) = \begin{cases} 0 & \text{if } x \in \Omega \setminus E \\ +\infty & \text{if } x \in E. \end{cases}$$

With this setup, the set E represents an obstacle that the density is not allowed to penetrate. During the flow, the density diffuses and drifts towards the lower level sets of V_2 , all while avoiding the set E .

In Figure 5 and Figure 6, we display two different experiments with different obstacles E , but the same diffusion exponent $m = 4$. In both experiments, the starting density is the characteristic function of a square centered at $(-0.3, -0.3)$ with side length 0.2 renormalized to have unit mass. In Figure 5, the obstacle is a disc with radius 0.2 centered at the origin, and in Figure 6, the obstacle is a star shaped region centered at the origin. In both experiments, we set $\tau = 0.001$, $\gamma = 0.0075$ and we run the flow until time $t = 2$. An interesting difference between the two flows is that the non-convexity of the star shaped obstacle results in some mass being trapped between the arms of the star. It is not entirely clear if the mass eventually escapes as time goes to infinity. This is because the PME allows for compactly supported solutions (in contrast to say the behavior of the heat equation).

4.2.3. *Non-convex U (aggregation-diffusion)*. In this experiment, we simulate (1.1) with an energy functional U that is not a convex with respect to ρ . Specifically, we consider the energy

$$(4.6) \quad U(\rho) = \mathcal{W}(\rho) + \int_{\Omega} \frac{1}{60} \rho^3(x) dx,$$




FIGURE 5. PME with exponent $m = 4$, $\gamma = .0075$ and potential given by (4.5). The obstacle E is represented by the white region. The images show the evolution from time $t = 0$ to $t = 2$ (top left to bottom right). Images are 512×512 pixels. With the exception of the obstacle, brighter pixels indicate larger density values.




FIGURE 6. PME with exponent $m = 4$, $\gamma = .0075$ and potential given by (4.5). The obstacle E is represented by the white region. The images show the evolution from time $t = 0$ to $t = 2$ (top left to bottom right). Images are 512×512 pixels. With the exception of the obstacle, brighter pixels indicate larger density values.

where

$$\mathcal{W}(\rho) := \frac{1}{2} \int_{\Omega} \int_{\Omega} |x - y|^2 \rho(x) \rho(y) dy dx.$$

By separating out the square, one can check \mathcal{W} is concave with respect to ρ .

While convex energies U encourage mass diffusion, non-convex energies allow for both aggregation and diffusion phenomena. Indeed, one can see that $\mathcal{W}(\rho)$ encourages the density to concentrate while the ρ^3 term encourages the density to diffuse. Due to the convolution, \mathcal{W} can be viewed as a “lower order” term as compared to ρ^3 . However, since the coefficients in front of the

convolution is much larger than the coefficient in front of the ρ^3 term, the aggregation effect will dominate until the density reaches a certain saturation level.

Here we run a single experiment starting with an initial density that is the sum of the characteristic function of four squares with side lengths 0.2 centered at each combination of $(\pm 0.3, \pm 0.3)$ and renormalized to have total mass equal to one. We set $\tau = .005$ and run the flow from time $t = 0$ to $t = 10$, at which time the evolution appears to have reached a steady state.

The results of the experiment are displayed in Figures 7 and 8. Figure 7 displays a heat map of the density evolution, while Figure 8 gives a 3 dimensional plot showing the height of the density. Throughout the evolution, one can see the competing effects of aggregation and diffusion. The heights of the four densities decrease due to diffusion, however aggregation pulls the four separate components together towards the center of the domain.




FIGURE 7. Aggregation-diffusion equation with an energy given by (4.6). The images show the evolution from time $t = 0$ to $t = 10$ (top left to bottom right). The final image is the approximate steady state. Images are 512×512 pixels. Brighter pixels indicate larger density values.

4.2.4. *Incompressible projections and flows.* In our last set of experiments, we consider incompressible flows, which have applications to crowd motion models and fluid mechanics. Here the energy takes the form

$$(4.7) \quad U(\rho) = s_\infty(\rho) + \int_{\Omega} V(x)\rho(x) dx,$$

where

$$s_\infty(\rho) = \begin{cases} 0 & \text{if } 0 \leq \rho(x) \leq 1 \text{ for a.e. } x \in \Omega, \\ \infty & \text{otherwise,} \end{cases}$$

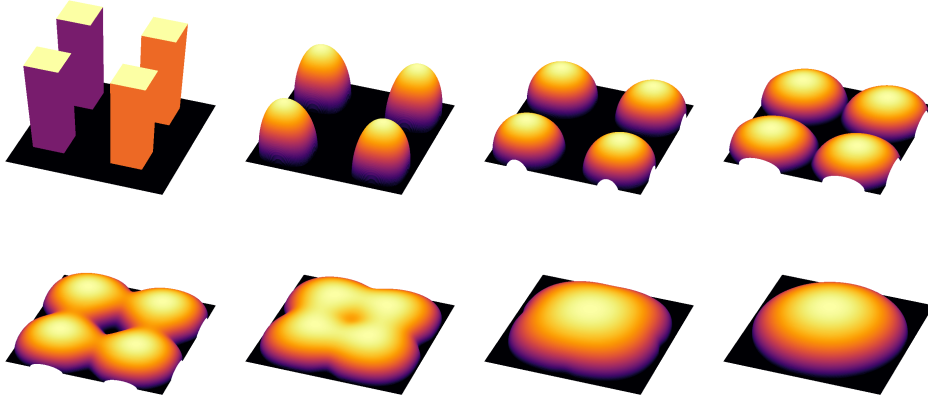


FIGURE 8. Aggregation-diffusion equation with an energy given by (4.6). The images show a 3-d surface plot of the evolution from time $t = 0$ to $t = 10$ (top left to bottom right). The final image is the approximate steady state. Images are 512×512 pixels.

and V is a fixed potential function. Note that $s_\infty(\rho)$ can be seen as the limit of the energy

$$s_m(\rho) = \frac{1}{m-1} \int_{\Omega} \rho^m(x) dx$$

as $m \rightarrow \infty$.

We will run our experiments, using the potential energy

$$(4.8) \quad V(x) = \frac{1}{2} \left((x_1 - \frac{3}{10})^2 + (x_2 - \frac{3}{10})^2 \right) + \iota_{\Omega \setminus E}(x)$$

where E is a closed set that represents an impenetrable obstacle. We run two simulations using two different obstacles

$$E_1 = B_{\frac{1}{4}}(\frac{1}{5}, -\frac{1}{5}) \cup B_{\frac{1}{4}}(-\frac{1}{5}, \frac{1}{5})$$

and

$$E_2 = B_{\frac{1}{10}}(0, \frac{1}{5}) \cup B_{\frac{1}{10}}(0, -\frac{1}{5}) \cup B_{\frac{1}{10}}(\frac{1}{5}, 0) \cup B_{\frac{1}{10}}(-\frac{1}{5}, 0),$$

where $B_r(x_1, x_2)$ denotes the closed ball of radius r centered at (x_1, x_2) . In both experiments, we choose an initial density $\rho^{(0)}$, which equals 1 on a ball of a radius 0.15 centered at $(-0.3, -0.3)$ and is equal to 0 elsewhere.

The results of our experiments are displayed in Figures 9 and 10. Figure 9 uses the obstacle E_1 , while Figure 10 uses the obstacle E_2 . In the figures, the yellow pixels represent the density $\rho^{(n)}$ and white pixels represents the obstacle. In both experiments we use a time step $\tau = 0.05$ and run the evolution from time $t = 0$ to time $t = 20$. Both experiments are conducted on 1024×1024 pixel grids.

Notably, in both of the simulations depicted in Figures 9 and 10, there is a sharp interface separating the regions $\rho = 1$ and $\rho = 0$. This matches the expected behavior of the flow with our chosen potentials. In general, it is

difficult for numerical methods to correctly capture sharp interfaces. Again, the reason that our method is able to do so is because of our dual approach. By recovering the density through the duality relation $\rho^{(n+1)} \in \delta U^*(\phi^{(n+1)})$ we automatically produce a discontinuity at the level set $\{y \in \Omega : \phi^{(n+1)}(y) - V(y) = 0\}$.




FIGURE 9. Incompressible flow with the energy (4.7), potential (4.8), and obstacle E_1 . The images show the evolution from time $t = 0$ to $t = 20$ (top left to bottom right). The final image is the approximate steady state. Images are 1024×1024 pixels. Yellow pixels represents the density and white pixels represents the obstacle.




FIGURE 10. Incompressible flow with the energy (4.7), potential (4.8), and obstacle E_2 . The images show the evolution from time $t = 0$ to $t = 20$ (top left to bottom right). The final image is the approximate steady state. Images are 1024×1024 pixels. Yellow pixels represents the density and white pixels represents the obstacle.

APPENDIX A. PROOFS

Proof of Lemma 3.1.

1. In order to obtain the Hessian of F let us start with its first derivative. We have

$$F(\phi + h) - F(\phi) = \int_{\Omega} [(\phi + h)^c(x) - \phi^c(x)] \mu(x) dx.$$

Assume that ϕ is c -convex. Then Proposition ?? tells us how to differentiate the c -transform, so that we can write

$$\int_{\Omega} (\phi + h)^c(x) - \phi^c(x) \mu(x) dx = \int_{\Omega} h(T_{\phi}(x)) \mu(x) dx + o(h).$$

Therefore $\delta F(\phi)(h) = \int_{\Omega} h(T_{\phi}(x)) \mu(x) dx$. To derive the Hessian of F we similarly compute

$$\delta F(\phi + h)(h) - \delta F(\phi)(h) = \int_{\Omega} [h(T_{\phi+h}(x)) - h(T_{\phi}(x))] \mu(x) dx.$$

We must now differentiate the maps T_{ϕ} with respect to ϕ . By Proposition 2.2 we know that $T_{\phi}(x) = x - \tau \nabla \phi^c(x)$. As a consequence

$$\begin{aligned} T_{\phi+h}(x) - T_{\phi}(x) &= -\tau \nabla [(\phi + h)^c - \phi^c](x) \\ &= -\tau \nabla (h \circ T_{\phi})(x) + o(h) \\ &= -\tau DT_{\phi}(x)^T \nabla h(T_{\phi}(x)) + o(h). \end{aligned}$$

Note that $DT_{\phi} = I_{d \times d} - \tau D^2 \phi^c$ is a symmetric matrix. We deduce from the above computations that

$$\begin{aligned} \delta F(\phi + h)(h) - \delta F(\phi)(h) &= \\ &= \int_{\Omega} \nabla h(T_{\phi}(x)) \cdot (-\tau) DT_{\phi}(x) \nabla h(T_{\phi}(x)) \mu(x) dx + o(h), \end{aligned}$$

from which we conclude that

$$(A.1) \quad \delta^2 F(\phi)(h, h) = -\tau \int_{\Omega} \nabla h(T_{\phi}(x)) \cdot DT_{\phi}(x) \nabla h(T_{\phi}(x)) \mu(x) dx.$$

Since our goal is to bound this Hessian by a norm of h we do the change of variable $y = T_{\phi}(x)$, or equivalently $x = S_{\phi^c}(y)$ since S_{ϕ^c} is the inverse of T_{ϕ} . Then

$$\delta^2 F(\phi)(h, h) = -\tau \int_{\Omega} \nabla h(y) \cdot DT_{\phi}(S_{\phi^c}(y)) \nabla h(y) \mu(S_{\phi^c}(y)) \det DS_{\phi^c}(y) dy.$$

Note that DS is a positive semi-definite matrix and therefore no absolute value is needed on the determinant term. Moreover we have $DT_{\phi}(S_{\phi^c}(y)) = DS_{\phi^c}(y)^{-1}$ and putting this term together with the determinant we can form the cofactor matrix $\text{cof}(DS) = \det(DS) DS^{-1}$.

As a result we obtain the expression

$$\delta^2 F(\phi)(h, h) = -\tau \int_{\Omega} \nabla h(y) \cdot \text{cof}(DS_{\phi^c}(y)) \nabla h(y) \mu(S_{\phi^c}(y)) dy.$$

2. Since ϕ is c -convex, $\phi = \phi^{c\bar{c}}$ and therefore $S_{\phi^c}(y) = y + \tau \nabla \phi(y)$. The c -convexity of ϕ also implies that the symmetric matrix $DS_{\phi^c}(y) = I_{d \times d} + \tau D^2 \phi(y)$ is positive semi-definite. Assume now that $I_{d \times d} + \tau D^2 \phi(y) \leq \Lambda I$ for all $y \in \Omega$. Then $I_{d \times d} + \tau D^2 \phi(y)$ is a symmetric matrix with eigenvalues between 0 and Λ . By general properties of the cofactor matrix the eigenvalues of $\text{cof}(DS_{\phi^c}(y))$ lie between 0 and Λ^{d-1} where d is the space dimension. We immediately deduce

$$-\delta^2 F(\phi)(h, h) \leq \tau \Lambda^{d-1} \|\mu\|_{L^\infty} \int_{\Omega} |\nabla h(y)|^2 dy.$$

□

Lemma A.1. *Suppose that $E \subset \mathbb{R}^d$ is a bounded set with C^2 boundary and let $R := \text{Reach}(\partial E)$. Let $u_0 : \mathbb{R}^d \rightarrow \mathbb{R}$ be a solution to the Eikonal equation $|\nabla u_i| = 1$ where $u_0 < 0$ inside E and $u_0 > 0$ outside E and set $u_1 = -u_0$. Let*

$$E_r^0 = \{x \in \mathbb{R}^d : u_0(x) \in (0, r)\}.$$

and

$$E_r^1 = \{x \in \mathbb{R}^d : u_1(x) \in (0, r)\}.$$

If $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth function, then for $i = 0, 1$

$$\int_{\partial E} |g(x)| ds(x) \leq \inf_{0 < r' < R} \left(\int_{E_r^i} |\nabla g(x)| dx + C_i(E, r) \int_{E_r^i} |g(x)| dx \right)$$

where

$$C_i(E, r) = \inf_{0 < r' < r} \frac{1}{r'} + \sup_{x \in E_r^i} (\Delta u_i(x))_+$$

Remark A.1. The reach of ∂E is the largest number r such that the characteristics of u_0 do not cross in $E_r^0 \cup E_r^1$. When ∂E is C^2 , the reach must be strictly positive and the Laplacian Δu must be bounded on $E_r^0 \cup E_r^1$ for all r smaller than the reach of ∂E .

Remark A.2. If E is a convex set, then $C_1(E, r) = \frac{1}{r}$.

Proof. Note that if $x \in \partial E$ and $n(x)$ is the outward facing normal at x , then $\nabla u_0(x) = n(x)$. Therefore,

$$\int_{\partial E} |g(x)| ds(x) = \int_{\partial E} |g(x)| \nabla u_0(x) \cdot n(x) ds(x)$$

For some $r \in (0, R)$ let $\alpha_r : \mathbb{R} \rightarrow \mathbb{R}$ be a function such that

$$\alpha_r'(t) = \begin{cases} 1 & \text{if } t \geq 0, \\ 1 + \frac{t}{r} & \text{if } t \in (-r, 0), \\ 0 & \text{if } t \leq -r. \end{cases}$$

We then have

$$\int_{\partial E} |g(x)| \nabla u_0(x) \cdot n(x) ds(x) = \int_{\partial E} |g(x)| \nabla (\alpha_r(u_0(x))) \cdot n(x) ds(x) =$$

$$\int_E \nabla \cdot \left(|g(x)| \nabla \left(\alpha_r(u_0(x)) \right) \right) dx$$

where the last equality follows from Stokes Theorem. Expanding out the derivatives and noting that $\alpha_r'(t) \in [0, 1]$, $\alpha_r''(t) \in [0, \frac{1}{r}]$ and $\alpha'(u_0(x)), \alpha''(u_0(x))$ both vanish for x outside of E_r^0 , we get

$$\begin{aligned} \int_{\partial E} |g(x)| ds(x) &\leq \int_{E_r^0} |\nabla g(x)| dx + \int_{E_r^0} |g(x)| \left((\Delta u_0(x))_+ + \frac{1}{r} \right) dx \leq \\ &\int_{E_r^0} |\nabla g(x)| dx + C_0(E, r) \int_{E_r^0} |g(x)| dx \end{aligned}$$

Our choice of r was arbitrary, thus we can take an inf over $r \in (0, R)$ to conclude the result when $i = 0$.

To tackle the case $i = 1$, we will employ a nearly identical argument, except we will use Stokes Theorem to convert the boundary integral into an integral over $\mathbb{R}^d \setminus E$. Since $\nabla u_1(x) \cdot n(x) = -1$ for $x \in \partial E$, we have

$$\begin{aligned} \int_{\partial E} |g(x)| ds(x) &= - \int_{\partial E} |g(x)| \nabla \left(\alpha_r(u_1(x)) \right) \cdot n(x) ds(x) = \\ &\int_{\mathbb{R}^d \setminus E} \nabla \cdot \left(|g(x)| \nabla \left(\alpha_r(u_1(x)) \right) \right) dx \end{aligned}$$

Now an identical argument to the one above gives the bound for the case $i = 1$. \square

Corollary A.2. *Suppose that $E \subset \Omega$ is a set with C^2 boundary and let $R := \min(\text{Reach}(\partial E), \text{dist}(E, \partial \Omega))$. Define u_i , E_r^i , and $C_i(E, r)$ as in Lemma A.1, and let*

$$C(E, \Omega) = \min_{i \in \{0, 1\}} \inf_{0 < r < R} C_i(E, r).$$

If $h : \Omega \rightarrow \mathbb{R}$ is an H^1 function, then

$$\int_{\partial E} |h(x)|^2 dx \leq \frac{1}{C} \int_{\Omega} |\nabla h(x)|^2 + 2C \int_{\Omega} |h(x)|^2 dx,$$

where

$$C = \max(1, C(E, \Omega)).$$

Proof. Suppose first that $h : \Omega \rightarrow \mathbb{R}$ is a smooth function. By Lemma A.1, we have

$$\int_{\partial E} |h(x)|^2 ds(x) \leq \inf_{0 < r < R} \left(\int_{E_r^i} 2|h(x)| \nabla h(x) dx + C_i(E, r) \int_{E_r^i} |h(x)|^2 dx \right)$$

for $i = 0, 1$. Clearly this is bounded from above by

$$\int_{\Omega} 2|h(x)| \nabla h(x) dx + \inf_{0 < r < R} C_i(E, r) \int_{\Omega} |h(x)|^2 dx$$

Taking a minimum over $i = 0, 1$, we can conclude that

$$\int_{\partial E} |h(x)|^2 ds(x) \leq \int_{\Omega} 2|h(x)| \nabla h(x) dx + C \int_{\Omega} |h(x)|^2 dx.$$

We can then use Cauchy-Schwarz to get

$$\int_{\partial E} |h(x)|^2 ds(x) \leq \frac{1}{C} \int_{\Omega} |\nabla h(x)| dx + 2C \int_{\Omega} |h(x)|^2 dx.$$

The result extends to H^1 functions thanks to the continuity of the trace operator over H^1 . \square

Proof of Theorem 3.4. Recall that $I(\psi) = \int_{\Omega} \psi(x) \mu(x) dx - U^*(\psi^{\bar{c}})$.

Step 1: formula for the Hessian of I . The derivation of the Hessian of I is similar to the one of J (see for instance the proof of Lemma 3.1). Using the formulas for the first variation of the c -transform (see Lemma ??) we can check that

$$\delta I(\psi)h = -\delta U^*(\psi^{\bar{c}})(h \circ S_{\psi}),$$

for any test function h . To obtain the Hessian of I , we need to differentiate S_{ψ} . As in the proof of Lemma 3.1 we can show that $S_{\psi+h}(y) - S_{\psi}(y) = \tau DS_{\psi}(y)^T \nabla h(S_{\psi}(y)) + o(h)$. This implies

$$\begin{aligned} \delta^2 I(\psi)(h, h) &= -\delta^2 U^*(\psi^{\bar{c}})(h \circ S_{\psi}, h \circ S_{\psi}) - \\ &\quad \tau \int_{\Omega} \eta(y) \nabla h(S_{\psi}(y)) \cdot DS_{\psi}(y) \nabla h(S_{\psi}(y)) dy, \end{aligned}$$

where we let $\eta = \delta U^*(\psi^{\bar{c}})$. Thus as in the case of J , the Hessian of I $\delta^2 I(\psi)(h, h) = -(I) - (II)$ contains two terms which we can bound separately.

Step 2: Bound on (II). Do the change of variables $x = S_{\psi}(y)$, i.e. $y = T_{\psi^{\bar{c}}}(x)$ in (II). We obtain

$$(II) = \tau \int_{\Omega} \eta(T_{\psi^{\bar{c}}}(x)) \nabla h(x) \cdot \text{cof } DT_{\psi^{\bar{c}}}(x) \nabla h(x) dx,$$

which can be bounded above by $\tau \|\eta\|_{L^\infty} \Lambda^{d-1} \|\nabla h\|_{L^2}$ in the same spirit as in the proof of Lemma 3.1. Moreover $\|\eta\|_{L^\infty} \leq \rho_{\max}$. Indeed, assuming $V(x) \geq 0$ we have for all $x \in \Omega$

$$\eta(x) = \delta U^*(\psi^{\bar{c}})(x) = (u_m^*)'(\psi^{\bar{c}}(x) - V(x)) \leq (u_m^*)'(\psi^{\bar{c}}(x)) \leq \rho_{\max},$$

by monotonicity of $(u_m^*)'$ and by definition of ρ_{\max} . As a consequence

$$(II) \leq \tau \rho_{\max} \Lambda^{d-1} \|\nabla h\|_{L^2}^2.$$

Step 3: Bound on (I). We have

$$(I) = \delta U^*(\psi^{\bar{c}})(h \circ S_{\psi}, h \circ S_{\psi}) = \int_{\Omega} (u_m^*)''(\psi^{\bar{c}}(y) - V(y)) |h(S_{\psi})|^2 dy.$$

Do again the change of variables $y = T_{\psi^{\bar{c}}}(x)$ to obtain

$$(I) = \int_{\Omega} (u_m^*)''(p(x)) |h(x)|^2 \det(DT_{\psi^{\bar{c}}}(x)) dx,$$

where we recall that $p(x) = \psi^{\bar{c}}(T_{\psi^{\bar{c}}}(x)) - V(T_{\psi^{\bar{c}}}(x))$. We bound the determinant term by Λ^d . Then, to go further we must distinguish between the three cases $1 \leq m \leq 2$, $2 < m < \infty$ and $m = \infty$.

When $1 \leq m \leq 2$, the function $(u_m^*)''$ is increasing and therefore

$$(u_m^*)''(p(x)) \leq (u_m^*)''(M) = u_m''(\rho_{\max})^{-1},$$

where $M = \sup_x \delta U(\mu)(x)$ (see the maximum principle and the related discussion when ρ_{\max} is defined in equation (3.12)). To sum up,

$$(I) \leq u''(\rho_{\max})^{-1} \Lambda^d \|h\|_{L^2}^2.$$

When $2 < m \leq \infty$, one can follow the same line of proof as in the case of J , using now the function $p(x)$ instead of $\phi(x) - V(x)$ which modifies the related constants accordingly. □

REFERENCES

- [AKY14] Damon Alexander, Inwon Kim, and Yao Yao. Quasi-static evolution and congested crowd transport. *Nonlinearity*, 27(4):823–858, mar 2014. [2](#)
- [Bar96] Grigory Isaakovich Barenblatt. *Scaling, self-similarity, and intermediate asymptotics*, volume 14 of *Cambridge Texts in Applied Mathematics*. Cambridge University Press, Cambridge, 1996. With a foreword by Ya. B. Zel'dovich. [24](#)
- [Bar03] Grigory Isaakovich Barenblatt. *Scaling*. Cambridge Texts in Applied Mathematics. Cambridge University Press, Cambridge, 2003. With a foreword by Alexandre Chorin. [24](#)
- [BCL16] Jean-David Benamou, Guillaume Carlier, and Maxime Laborde. An augmented Lagrangian approach to Wasserstein gradient flows and applications. In *Gradient flows: from theory to application*, volume 54 of *ESAIM Proc. Surveys*, pages 1–17. EDP Sci., Les Ulis, 2016. [2](#)
- [BCMO16] Jean-David Benamou, Guillaume Carlier, Quentin Mérigot, and Édouard Oudet. Discretization of functionals involving the Monge–Ampère operator. *Numer. Math.*, 134(3):611–636, 2016. [2](#)
- [BCW10] Martin Burger, José A. Carrillo, and Marie-Therese Wolfram. A mixed finite element method for nonlinear diffusion equations. *Kinet. Relat. Models*, 3(1):59–83, 2010. [2](#)
- [Bre91] Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Comm. Pure Appl. Math.*, 44(4):375–417, 1991. [7](#)
- [CCWW19] Jose A Carrillo, Katy Craig, Li Wang, and Chaozhen Wei. Primal dual methods for Wasserstein gradient flows. *arXiv preprint arXiv:1901.08081*, 2019. [2](#)
- [CDPS17] Guillaume Carlier, Vincent Duval, Gabriel Peyré, and Bernhard Schmitzer. Convergence of entropic schemes for optimal transport and gradient flows. *SIAM J. Math. Anal.*, 49(2):1385–1418, 2017. [2](#)
- [CM10] 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(6):4305–4329, 2009/10. [2](#)
- [CWXY20] Jose A Carrillo, Li Wang, Wuzhe Xu, and Ming Yan. Variational asymptotic preserving scheme for the Vlasov–Poisson–Fokker–Planck system. *arXiv preprint arXiv:2007.01969*, 2020. [2](#)

- [DPMSV16] Guido De Philippis, Alpár Richárd Mészáros, Filippo Santambrogio, and Bozhidar Velichkov. BV estimates in optimal transportation and applications. *Arch. Ration. Mech. Anal.*, 219(2):829–860, 2016. [2](#)
- [Eva10] Lawrence C. Evans. *Partial differential equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, second edition, 2010. [18](#)
- [Eyr98] David J. Eyre. Unconditionally gradient stable time marching the Cahn–Hilliard equation. *MRS Proceedings*, 529:39, 1998. [5](#), [21](#)
- [Gan94] Wilfrid Gangbo. An elementary proof of the polar factorization of vector-valued functions. *Archive for rational mechanics and analysis*, 128(4):381–399, 1994. [6](#)
- [Gan95a] Wilfrid Gangbo. *Quelques problèmes d’analyse non convexe. Habilitation à diriger des recherches en mathématiques*. Habilitation, Université de Metz, January 1995. [7](#)
- [Gan95b] Wilfrid Gangbo. Quelques problèmes d’analyse non convexe. *Habilitation à diriger des recherches en mathématiques. Université de Metz (Janvier 1995)*, 1995. [6](#)
- [GM96] Wilfrid Gangbo and Robert J. McCann. The geometry of optimal transportation. *Acta Math.*, 177(2):113–161, 1996. [6](#), [7](#)
- [JKO98] Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker–Planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998. [1](#)
- [JKT20] Matt Jacobs, Inwon Kim, and Jiajun Tong. The L^1 -contraction principle in optimal transport. *arXiv preprint arXiv:2006.09557*, 2020. [2](#), [9](#), [16](#), [25](#)
- [JL20] Matt Jacobs and Flavien Léger. A fast approach to optimal transport: the back-and-forth method. *Numerische Mathematik*, pages 1–32, Oct 2020. [1](#), [2](#), [4](#), [12](#), [13](#), [22](#), [23](#)
- [LMSS20] Hugo Leclerc, Quentin Mérigot, Filippo Santambrogio, and Federico Stra. Lagrangian discretization of crowd motion and linear diffusion. *SIAM Journal on Numerical Analysis*, 58(4):2093–2118, 2020. [2](#)
- [Luc97] Yves Lucet. Faster than the fast Legendre transform, the linear-time Legendre transform. *Numer. Algorithms*, 16(2):171–185, 1997. [4](#), [22](#)
- [Nes13] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013. [11](#)
- [Ott01] Felix Otto. The geometry of dissipative evolution equations: the porous medium equation. *Comm. Partial Differential Equations*, 26(1-2):101–174, 2001. [1](#)
- [Pey15] Gabriel Peyré. Entropic approximation of Wasserstein gradient flows. *SIAM J. Imaging Sci.*, 8(4):2323–2351, 2015. [2](#)
- [San15] Filippo Santambrogio. *Optimal transport for applied mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser/Springer, Cham, 2015. Calculus of variations, PDEs, and modeling. [1](#), [2](#), [6](#), [7](#)
- [Váz07] Juan Luis Vázquez. *The porous medium equation: mathematical theory*. Oxford University Press, 2007. [1](#)

UCLA, LOS ANGELES, CA, USA
Email address: majaco@math.ucla.edu

UCLA, LOS ANGELES, CA, USA

ENS, PSL UNIVERSITY, PARIS, FRANCE
Email address: flavienleger@nyu.edu