Numerical solution of the second boundary value problem for the Elliptic Monge-Ampère equation

Jean-David Benamou, Adam Oberman, Froese Britanny

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Numerical solution of the second boundary value problem for the Elliptic Monge-Ampère equation

Jean-David Benamou — Brittany Froese — Adam Oberman

N° ????

Juin 2012

Thème NUM

ISSN 0249-6399 ISRN INRIA/RR--????--FR+ENG

INSTITUT NATIONAL
DE RECHERCHE
EN INFORMATIQUE
ET EN AUTOMATIQUE

INRIA

centre de recherche
PARIS - ROCQUENCOURT

Numerical solution of the second boundary value problem for the Elliptic Monge-Ampère equation

Jean-David Benamou*, Brittany Froese† , Adam Oberman‡

Thème NUM — Systèmes numériques
Équipes-Projets POEMS

Rapport de recherche n° ???? — Juin 2012 — 37 pages

Abstract: This paper introduces a numerical method for the solution of the nonlinear elliptic Monge-Ampère equation. The boundary conditions correspond to the optimal transportation of measures supported on two domains, where one of these sets is convex. The new challenge is implementing the boundary conditions, which are implicit and non-local. These boundary conditions are reformulated as a nonlinear Hamilton-Jacobi PDE on the boundary. This formulation allows us to extend the convergent, wide stencil Monge-Ampère solvers proposed by Froese and Oberman to this problem. Several non-trivial computational examples demonstrate that the method is robust and fast.

Key-words: Optimal Transportation, Monge-Ampère Equation

* INRIA, Domaine de Voluceau Rocquencourt, 78153 Rocquencourt, France
† Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, Canada
‡ Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, Canada

Centre de recherche INRIA Paris – Rocquencourt
Domaine de Voluceau, Rocquencourt, BP 105, 78153 Le Chesnay Cedex
Téléphone : +33 1 39 63 55 11 — Télécopie : +33 1 39 63 53 30

Résolution numérique du deuxième problème aux limites pour l'équation de Monge Ampère

Résumé : Cet article présente une méthode de résolution numérique pour une équation de Monge Ampère elliptique non-linéaire. Les conditions aux limites particulières correspondent au problème du transport optimal entre deux mesures dont au moins un des supports est convexe. La difficulté posée tient au caractère implicite et non local de ces conditions aux limites. Nous proposons de les reformuler comme une équation de Hamilton-Jacobi sur le bord. Ceci permet d'étendre les schémas de type "wide-stencil" et résultats de convergence associés de Froese et Oberman à ce problème. Plusieurs cas tests, certains non triviaux, démontrent la rapidité et robustesse de la méthode.

Mots-clés : Transport Optimal, Equation de Monge-Ampère

NUMERICAL SOLUTION OF THE SECOND BOUNDARY VALUE PROBLEM FOR THE ELLIPTIC MONGE-AMPÈRE EQUATION

J.-D. BENAMOU(INRIA-ROCQUENCOURT, FRANCE),
B. D. FROESE (SFU, CANADA), AND A. OBERMAN (SFU, CANADA)

Abstract. This paper introduces a numerical method for the solution of the nonlinear elliptic Monge-Ampère equation, with boundary conditions corresponding to the optimal transportation of measures supported on two domains, X and Y, where one of these sets is convex. The new challenge is implementing the boundary conditions, which are implicit. These boundary conditions are reformulated as a nonlinear Hamilton-Jacobi PDE on the boundary. This formulation allows us to extend the convergent, wide stencil Monge-Ampère solvers proposed in [FO11a] to this problem. Several non-trivial computational examples demonstrate that the method is robust and fast.

Contents

Date: June 4, 2012.

1. Introduction

1.1. **Numerical Optimal transportation.** The paper introduces a numerical method for the solution of the fully nonlinear elliptic Monge-Ampère equation for a convex potential u :

$$\det(D^2u(x)) = \frac{\rho_X(x)}{\rho_Y(\nabla u(x))} \quad \text{for } x \in X. \tag{MA}$$

Here, D^2u is the Hessian and ∇u the gradient of the function $u : X \subset \mathbb{R}^d \to \mathbb{R}$. The equation arises from the Monge-Kantorovitch problem where the map ∇u, from the probability density ρ_X supported on X to the probability density ρ_Y supported on $Y \subset \mathbb{R}^d$, minimizes a transportation cost (see subsection 1.5). We require that X, Y are convex and bounded. Our method will allow ρ_X to vanish, but requires Lipschitz continuity of ρ_Y.

The usual boundary conditions are replaced by a state constraint on the gradient map from X to Y,

$$\nabla u(X) = Y \tag{BV2}$$

This condition is referred to in the literature as the *second boundary value problem* for the Monge-Ampère equation [Pog94].

The main idea in this paper is to replace (BV2) by a Hamilton-Jacobi equation on the boundary :

$$H(\nabla u(x)) = \operatorname{dist}(\nabla u(x), \partial Y) = 0, \quad \text{for } x \in \partial X \tag{HJ}$$

where we use the signed distance function to the boundary of the set Y as the Hamiltonian H.

The problem at hand now is (MA-HJ), and solving this combined problem requires several theoretical and numerical ideas. First, we use convexity of the solution to demonstrate obliqueness of the boundary condition. This obliqueness condition ensures that H can be discretized using information inside the domain (upwinding). The upwind discretization gives a monotone and consistent scheme on the boundary. Together with a consistent and monotone scheme for (MA) inside the domain, we obtain convergence for the combined discretization of (MA-HJ) using the approximation framework of Barles-Souganidis [BS91]. Using the work of [FO12], we also extend this framework to allow for more accurate (or nearly monotone) approximations.

1.2. **Discussion of numerical methods for Optimal Transportation.** To the best of our knowledge, there was previously no method available for implementing the condition (BV2) together with (MA) and, indeed, previous works have restricted their attention to geometries/boundary conditions where it can be simplified.

A first and widely used simplification is to work on the torus (with periodic densities) using the change of variable $u = Id + v$ (v is periodic). However, this severely restricts the range of solvable problems. Mass transfer between periodic cells may be optimal. There is no easy way to prevent it and capture the optimal transportation for the one cell non-periodic problem.

A second option for designing optimal transportation-compatible boundary conditions is to restrict attention to fixed simple geometries. For example, when dealing with square to square maps, e.g. $X = Y = [0,1] \times [0,1]$ and denoting by (x, y) the cartesian coordinates, the Neumann boundary condition $\partial_x u|_{\{x=0,1\}} = \{0, 1\}$, $\partial_y u|_{\{y=0,1\}} = \{0, 1\}$ yields a face to face mapping on the boundary, which is optimal when mass does not vanish. Note that Neumann BCs prescribe only one component of the gradient map, thus leaving the map free to stretch or contract in the tangential direction.

An idea to treat more general density geometries would be to extend the density support to a square and pad with zeros, but numerical methods for (MA) are not able to treat vanishing densities in the target space Y. Mollifying and adding constant mass everywhere can significantly modify the optimal map.

In [Fro12], a more general heuristic method is proposed, consisting in iteratively solving (MA) with Neumann boundary conditions, and projecting the resulting set onto the target set Y. The new projection is then used to derive new Neumann boundary conditions. This method required several iterations, and no convergence proof was available.

1.3. **Our approach to the numerical treatment of (HJ).** The current work treats the boundary condition naturally, admits a convergence proof, and allows the equation to be solved quickly: at the same cost as solving the Dirichlet problem, and at a cost equivalent to several linear solves on the same grid. The projection method proposed in [Fro12], mentioned above, can be interpreted as a gradient descent solution method for (HJ).

We use a reformulation of (BV2) for which a convergent monotone finite difference scheme, consistent with the treatment of the PDE (MA), can be easily built and incorporated into a fast Newton solver. It is based on the concept of *defining function* used in [Del91, Urb97] precisely to prove the existence of classical solutions for (MA-BV2). To summarize the idea: a convex function H defined on the whole space $\mathbb{R}^d$, which the target set Y, is a defining function for the set Y if the zero level set of H coincides with ∂Y. Such a function is easy to construct and compute; in this paper we use the signed distance function to the boundary of the set Y. As X, Y are convex, the second boundary value condition (BV2) is then equivalent to

the nonlinear but *local* boundary condition:

$$H(\nabla u(x)) = 0, \qquad \text{for } x \in \partial X. \tag{HJ}$$

The problem at hand now is (MA-HJ), and we show that it is tractable using the classical viscosity approximation theory.

1.4. **Discussion of OT and OT numerics.** We recall that the term "second boundary value" was proposed in the pioneering work of Pogorelov [Pog94]. Pogorelov's constructive method for solving Monge-Ampère assumes a target density ρ_Y in the form of a (possibly weighted) sum of Dirac masses at prescribed supporting points in the target space. He then constructs the convex potential solution u as a supremum of affine functions whose gradients/slopes necessarily take values in the finite set of supporting points. The support domain of u, X, is split into cells supporting the contribution of each affine function to the solution. Adjusting the constant/height of these affine functions, one can increase or decrease their sizes in order to satisfy the Jacobian equation. In this framework, the "second boundary value" problem arises as a natural and easy to treat constraint instead of the usual "first" Dirichlet boundary conditions. The theory for Neumann and more general oblique boundary conditions is more technical and was addressed later; see [LTU86]. Pogorelov solutions also provide examples of weak solutions, as the gradient map is discontinuous.

A similar idea has been pursued numerically in the contexts of meteorology [CP84], antenna design [GO03a], and more recently by in image processing [Mer11]. In this context, optimal mass transfer is a linear programming problem. When the initial density is also a sum of Diracs, the popular auction algorithm proposed by Bertsekas (see the survey paper [Ber92]) solves it with $\mathcal{O}(N^2 \log N)$ complexity. In [Bos10], the author compares different linear programming approaches and discusses the non-trivial issue of quantization (discretization of densities towards sum of Diracs), which is necessary to treat more general mass transfer problems.

The other (large) family of numerical iterative methods uses gradient flows combined with a linearization of the optimal mass transfer problem (which boils down to a linear second order elliptic equation). It can be solved either by deterministic [CDF11, DAT08, HRT10] or stochastic [MRCSV11] methods. These approaches lack a convergence theory. A provably convergent method is the CFD (Computation Fluid Dynamics) reformulated optimal transportation "time" problem [BB00], which relaxes the nonlinearity of the constraints at the cost of an additional virtual time dimension. Classical efficient numerical optimization methods may also be used: Augmented Lagrangian [BB00] and classical optimal control methods where the gradient of the functional is computed by a direct/adjoint system of PDEs [BB01, LST10]. Solving the Monge-Ampère equation using a Newton's method as in [LR05, LPS] may also be counted in this class of numerical methods.

The optimal transportation problem for general cost functions can be formulated as a Linear Programming problem. However, this requires doubling the dimension of the problem and using LP solvers, which scale poorly for large problems sizes, compared to PDE methods. An implementation of the LP solution method can be found in [RU00].

1.5. **Discussion of Optimal Transportation.** Our interest in (MA-BV2) is linked to the Monge Kantorovitch or *optimal transportation* or *optimal mass transfer* problem in the case of quadratic cost :

$$\inf_{M\in\mathbb{M}} \frac{1}{2}\int_X \|x - M(x)\|^2 \rho_X(x)dx \tag{1}$$

$$\mathbb{M} = \{M : X \mapsto Y,\ \ \rho_Y(M)\ \det(\nabla M) = \rho_X\}. \tag{2}$$

In [Bre91], it is proved that (1-2) is well posed and that the unique map M at which the minimum is reached is the gradient of a convex potential u, which is therefore also unique up to a constant. The correspondence with (MA-BV2) is then formally straightforward: just replace M by ∇u in the constraints (2). The mass transfer problem can be seen as a variational formulation of the second boundary value problem and has been a important tool in understanding well-posedness and regularity or lack of regularity of Monge-Ampère solutions. In particular, sensitive conditions for the existence of globally smooth solutions are the convexity of the sets X and Y and the strict positivity of the densities [Caf96].

1.6. **Applications of OT.** The problem of optimal mass transport arises in a large number applications including image registration [HZTA04], mesh generation [BW09], reflector design [GO03b], astrophysics (estimating the shape of the early universe) [FMMS02], and meteorology [CNP91]. Motivated by economic applications, Kantorovich contributed[1] to the understanding of optimal transport by reformulating the problem as a linear program and describing a simple dual formulation [Kan42, Kan48]. While this has made many theoretical questions easier to answer, this approach also effectively doubles the dimension of the problem. Consequently, computing the solution to even a small-scale problem is prohibitively expensive.

A large class of nonlinear continuity equations with confinement and/or possibly non local interaction potentials can be considered as semi-discrete gradient flows, known as JKO gradient flows [JKO98, Ott01], with respect to the Euclidean Wasserstein distance. The distance is the value function of the Optimal transportation problem. Again, the cost of its numerical resolution, let alone computing the gradient (with respect to one of the densities), has so far prevented the numerical use of this technique. In 1D the problem is trivial and [KW99] implements JKO gradient flow simulations for nonlinear diffusion. An interesting recent work [CM09] starts attacking the 2D case.

[1]He earned the '75 Nobel prize in economics.

The performance of our solver offers new perspectives for implementing JKO gradient flows.

2. Representation and approximation of (HJ)

We now describe the numerical approach to solving the second boundary value problem (BV2). We implement the boundary condition using the signed distance function to the boundary of the set Y. We are able to treat the boundary condition using a monotone finite difference scheme, which is consistent with the treatment of the PDE (MA).

Remark 2.1. *It is now common to use a distance function to determine a set, as is the case in the level set method. In this case, one solves a Hamilton-Jacobi equation* for *the distance function. What we are doing here is to use the distance function* as *the nonlinear PDE operator.*

The resulting boundary conditions are implicit, in contrast to, for example, Dirichlet or Neumann boundary conditions, which have been previously implemented for (MA)*. However, at each grid point the boundary condition can be treated as an implicit equation in a similar manner to how the interior grid points are treated: in our case, with a Newton solver.*

2.1. Representation and properties of the distance function.

The unsigned distance function can be represented as $\operatorname{dist}(y, \partial Y) = \inf_{y_0 \in \partial Y} \|y - y_0\|$, but since it is not convex, it is not suited to our purposes. For a bounded, convex set Y, the signed distance function,

$$H(y) = \begin{cases} +\operatorname{dist}(y, \partial Y), & y \text{ outside of } Y \\ -\operatorname{dist}(y, \partial Y), & y \text{ inside } Y \end{cases}$$

is convex and can be written in terms of the supporting hyperplanes to the convex set,

$$H(y) = \sup_{y_0 \in \partial Y} \{n(y_0) \cdot (y - y_0)\}, \tag{3}$$

where $n(y_0)$ is the outward normal to ∂Y at y_0; see Figure 1. Equivalently, since the image of the normals to ∂Y is the unit sphere, we can write

$$H(y) = \sup_{\|n\|=1} \{n \cdot (y - y(n))\} \tag{4}$$

where $y(n)$ is the point in ∂Y with normal n. This last representation is useful for computationally determining these values from a given representation of the target set; see subsection 2.2. The last representation is also valid in the case where the convex set is not smooth.

These statements follow from the *Supporting Hyperplane Theorem* [BV04, Section 2.5], which says that if $y_0 \in \partial Y$, for a convex set Y, then y_0 has a (possibly non-unique) supporting hyperplane,

$$P = \{A(y) = 0 \mid A(y) \equiv n \cdot (y - y_0)\},$$

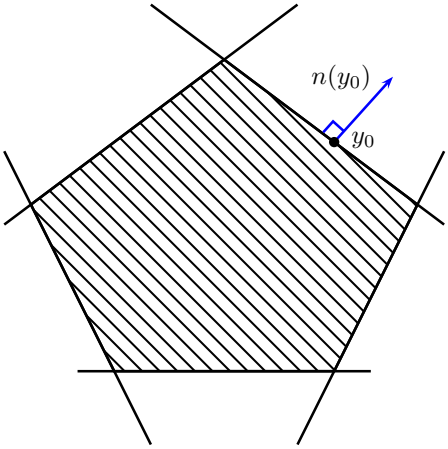

FIGURE 1. Polyhedral target set.

where $A(y) \leq 0$, for $y \in Y$. Without loss of generality, $\|n\| = 1$, and we can define n to be (an) outward normal to Y at y_0. Then, we can define

$$H^*(n) = n \cdot y(n) = \sup_{y \in \partial Y} n \cdot y,$$

where the equality follows from the supporting hyperplane result. We have proven the following lemma.

Lemma 2.2. *Let H be the signed distance function to a smooth, bounded set, Y. Then*

$$H(y) = \sup_{\|n\|=1} \{y \cdot n - H^*(n)\}$$

where

$$H^*(n) = \sup_{y_0 \in \partial Y} \{y_0 \cdot n\}. \tag{5}$$

2.2. **Representation of the target set.** The target set Y may be represented by its convex hull, or simply by scattered points. In the latter case, we wish to obtain the representation (3). This can be accomplished using computational geometry, or even convex optimization [BV04]. We now describe a method for obtaining the representation (3) from scattered points, which is convenient to our purposes, using the Legendre-Fenchel transform [BV04, section 3.3].

In the implementation, we will further discretize this Hamilton-Jacobi equation by computing the supremum over a finite subset of the admissible directions. These direction vectors are typically given by a uniform discretization of the directions, with discretization parameter $d\alpha$. We require only that $d\alpha \to 0$ for convergence.

We found this to be the simplest and computationally most efficient method to represent the geometry.

Even for an elliptical domain, where a large (hundreds, or on the order of the number of grid points on the perimeter) number of hyperplanes was needed to represent the convex set, this formulation was accurate and computationally inexpensive. This resulted from the fact that the hyperplanes (i.e. the values H^* below) were precomputed. Evaluating (5) is inexpensive.

2.3. **Obliqueness.** We recall here a fundamental property of maps characterized as the gradient of a convex potential. In, [Del91, Urb97] this *obliqueness* result is used to prove existence of classical solutions to (MA-BV2).

This condition, which leads to Lemma 2.4, will allow us to build an explicit monotone upwind discretization of (3), using points inside the domain.

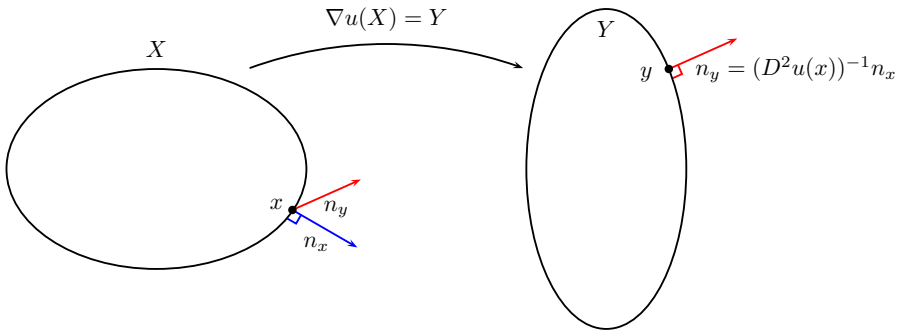

FIGURE 2. Illustration of the mapping $y = \nabla u(x)$ and the normal vectors.

Lemma 2.3. *Suppose X is a convex domain, and $Y = \nabla u(X)$ is the image of Y under the mapping ∇u, where u is a convex twice continuously differentiable function. Then the normal vectors make an acute angle,*

$$n_x \cdot n_y \geq 0. \tag{6}$$

See Figure 2.

Proof. Let $\nabla u(x) = y \in \partial Y$ and let $H(y) = \operatorname{dist}(y, \partial Y)$ be the signed distance function to Y, so that

$$Y = \{y \mid H(y) = 0\}$$

and

$$X = \{x \mid H(\nabla u(x)) = 0\}.$$

Then $n_y = \nabla H(y)$ and, by the chain rule for differentiation,

$$n_x = c\nabla D^2 u(x) \nabla H(\nabla u(x)) = cD^2 u(x) \nabla H(y)$$

for a normalization constant, c. Thus

$$n_x \cdot n_y = c(\nabla H(y))^T D^2 u(x) \nabla H(y) \geq 0,$$

since convexity of u means $D^2 u$ is positive definite. □

Lemma 2.4. *Let $u \in C^2(X)$ be any convex function that satisfies (BV2). For any $x \in \partial X$ with unit outward normal n_x, the supremum in (HJ) can be restricted to vectors making an acute angle with n_x:*

$$H(\nabla u(x)) = \sup_{\|n\|=1} \{\nabla u(x) \cdot n - H^*(n) \mid n \cdot n_x > 0\} = 0. \tag{7}$$

Proof. The supremum above will be attained for a value of $n = n^*$, which will be identical to the unit outward normal to the target at the point $\nabla u(x)$.

From Lemma 2.3, we know that $n^* \cdot n_x = n_y \cdot n_x \geq 0$. Consequently, it is only necessary to check values of n that make an acute angle with the boundary of the domain. □

2.4. **Monotone discretization of H.** In this section we explain how to build a monotone discretization of H using points at the boundary and on the inside of the domain X.

The expression (7), which comes from writing the convex set Y in terms of its tangent hyperplanes, leads to a natural convergent finite difference discretization. This expression can be understood as a Hamilton-Jacobi-Bellman equation arising from a control problem.

We recall (see [Obe06]) that an elliptic (monotone) discretization of $H(\nabla u)$ at the point x_i will take the form

$$H_i(u(x_i) - u(x_j))$$

and be non-decreasing in each of its arguments.

If we let $\Gamma_i = \{n \mid n \cdot n_{x_i} > 0, \|n\| = 1\}$, a simple way of writing an upwind discretization is to approximate the signed distance function by

$$H(\nabla u(x_i)) = \sup_{n \in \Gamma_i} \{\nabla u(x_i) \cdot n - H^*(n)\} \approx \sup_{n \in \Gamma_i} \left\{ \frac{u(x_i) - u(x_i - nh)}{h} - H^*(n) \right\}$$

where h is a small discretization parameter. As long as the domain is uniformly convex and h is sufficiently small, obliqueness ensures that the point $x_i - nh$ is inside the domain. However, this form does not lead to a simple, compact finite difference approximation.

Instead, for simplicity, we describe the discretization on a square domain. We note that a slightly more complicated discretization will generalize this to more general triangulated domains. However, by padding the source density ρ_X with zeros (see subsection 5.1), we can handle different geometries while still computing on a simple, square domain. We can also easily generalize the discretization to higher dimensions.

We describe the discretization along the left side of the square domain, with normal $n_x = (-1, 0)$. Along this side, the set of admissible directions

will be given by

$$\{n = (n_1, n_2) \mid n_1 < 0, \|n\| = 1\}.$$

Then, letting h denote the spatial resolution of the grid, we can approximate the advection terms by

$$\nabla u(x_i) \cdot n \approx n_1 \frac{u(x_i + (h,0)) - u(x_i)}{h} + \max\{n_2, 0\} \frac{u(x_i) - u(x_i - (0,h))}{h} + \min\{n_2, 0\} \frac{u(x_i + (0,h)) - u(x_i)}{h}.$$

This scheme only relies on values inside the square and, because $n_1 < 0$, it is monotone. Taking the supremum of these monotone schemes over all admissible directions, we preserve monotonicity of the scheme.

We write a similar scheme on the other sides of the square. At corners, we take formal limits of the obliqueness constraint to limit the admissible directions to a single quadrant, which ensures that the required information will continue to reside inside the square.

3. Convergence

We begin with a review of background material that will be needed to construct and prove the convergence of our scheme for solving the second boundary value problem for the Monge-Ampère equation.

3.1. General discussion of numerical methods.

The viscosity approximation theory developed by Barles and Souganidis [BS91] provides criteria for the convergence of approximation schemes: schemes that are consistent, monotone, and stable converge to the unique viscosity solution of a degenerate elliptic equation. This general framework can be applied to a wide class of nonlinear second order equations and, in particular, when the PDE operator is defined on the closure of the support domain and may be discontinuous. For the second boundary value problem the equations can be written using the abstract operator

$$F(x, u(x), \nabla u(x), \nabla^2 u(x)) = 0, \ x \in \bar{X} \tag{8}$$

where F depends on ρ_X, ρ_Y and H:

$$F(x, u, p, M) = \begin{cases} \det(M) - \rho_X(x)/\rho_Y(p), & x \in X \\ H(p), & x \in \partial X. \end{cases} \tag{9}$$

Remark 3.1. *For the Dirichlet problem, for example, we would write $u - g$ instead of H for $x \in \partial X$.*

Note that H depends only on p; therefore, it will not affect the degenerate ellipticity of the whole problem cast in the discontinuous viscosity framework of [BS91]. This abstract framework does not indicate how to build such schemes, or how to produce fast solvers for the schemes. It is not obvious how to ensure that schemes satisfy the required comparison principle. A class of

schemes which for which this property holds was identified in [Obe06], and were called *degenerate elliptic*, by analogy with the structure condition for the PDE.

For uniformly elliptic PDEs, monotone schemes are not always necessary for convergence (for example, most higher order finite element methods are not monotone). However, for fully nonlinear or degenerate elliptic equations, the only convergence proof currently available requires that schemes be (nearly) monotone. One way to ensure monotonicity is to use wide stencil finite difference schemes; this has been done for the equation for motion by mean curvature [Obe04], for the Infinity Laplace equation [Obe05], for functions of the eigenvalues [Obe08b], for Hamilton-Jacobi-Bellman equations [BZ03], and for the convex envelope [Obe08a]. Even for linear elliptic equations, a wide stencil may be necessary to build a monotone scheme [MW53].

3.2. **Convergence of approximation schemes.** The techniques we use to prove the convergence of our scheme are based on the framework for approximating viscosity solutions that was introduced by Barles and Souganidis [BS91]. A framework for the construction of convergent schemes for degenerate elliptic equation in general, and the Monge-Ampère equation in particular, has been further developed in [Obe06, FO11a, FO12].

The key properties required to construct a scheme F^ϵ that converges to the viscosity solution of the underlying PDE F are *consistency* and *near ellipticity*; these terms are recalled below.

Definition 3.2 (Consistent)**.** *The scheme F^ϵ is* consistent *with the equation $F = 0$ if for any smooth function ϕ and $x \in \bar{\Omega}$,*

$$\limsup_{\epsilon \to 0, y \to x, \xi \to 0} F^\epsilon(y, \phi(y) + \xi, \phi(\cdot) + \xi) \leq F^*(x, \phi(x), \nabla\phi(x), D^2\phi(x)),$$

$$\liminf_{\epsilon \to 0, y \to x, \xi \to 0} F^\epsilon(y, \phi(y) + \xi, \phi(\cdot) + \xi) \geq F_*(x, \phi(x), \nabla\phi(x), D^2\phi(x)).$$

Definition 3.3 (Elliptic)**.** *The scheme F^ϵ is* elliptic *if it can be written*

$$F^\epsilon[v] = F^\epsilon(x, v(x), v(x) - v(\cdot)),$$

where F^ϵ is nondecreasing in its second and third arguments,

$$s \leq t,\ u(\cdot) \leq v(\cdot) \implies F^\epsilon(x, s, u(\cdot)) \leq F^\epsilon(x, t, v(\cdot)) \tag{10}$$

Definition 3.4 (Nearly Elliptic)**.** *The scheme F^ϵ is* nearly elliptic *if it can be written as*

$$F^\epsilon[v] = F_M[u] + F_P^\epsilon[u] \tag{11}$$

where F_M is a monotone (elliptic) scheme and F_P^ϵ is a perturbation, which satisfies

$$\lim_{\epsilon \to 0} \|F_P^\epsilon\| = 0.$$

Using these definitions, we now recall the main convergence theorem from [FO12].

Theorem 3.5 (Convergence of Approximation Schemes)**.** *Let u be the unique viscosity solution of the PDE (8) and let u^ϵ be a stable solution of the consistent, nearly elliptic approximation scheme (11). Then*

$$u^\epsilon \to u, \qquad \text{locally uniformly, as } \epsilon \to 0.$$

Moreover, if the non-monotone perturbation F_P^ϵ is continuous, u^ϵ exists and is stable.

In the same paper, the authors describe filtered approximation schemes, which combine a monotone (elliptic) scheme F_M with a more accurate, non-monotone scheme F_A. These schemes, which are convergent, have the form

$$F^\epsilon[u] = F_M[u] + r(\epsilon) S\left(\frac{F_A^\epsilon[u] - F_M^\epsilon[u]}{r(\epsilon)}\right) \tag{12}$$

where $r(\epsilon) \to 0$ as $\epsilon \to 0$. Here the function S, which is called a filter, can be defined, for example, by

$$S(x) = \begin{cases} x & \|x\| \le 1 \\ 0 & \|x\| \ge 2 \\ -x+2 & 1 \le x \le 2 \\ -x-2 & -2 \le x \le -1. \end{cases} \tag{13}$$

See Figure 3.

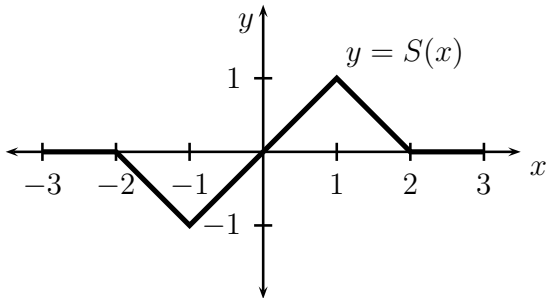

FIGURE 3. Filter function

3.3. **Discretization of Monge-Ampère equation.** The equation we want to solve is

$$\det(D^2 u(x)) = \rho_X(x)/\rho_Y(\nabla u(x)) + \langle u \rangle.$$

However, to reduce the cost of computations, we will replace the mean $\langle u \rangle$ with the value $u(x_0)$ for some point $x_0 \in X$, which should coincide with a grid point. The solution to the two problems is the same up to an additive constant.

We first describe the elliptic (monotone) scheme for the Monge-Ampère operator, which underlies the filtered scheme. This scheme was developed in [FO11a, Fro12].

To begin, we use Hadamard's inequality to represent the determinant of a positive definite matrix, $\det(M) \leq \Pi_i m_{ii}$, with equality when M is diagonal. Then, we can write

$$\det(M) = \min \left\{ \Pi_i (O^T M O)_{ii} \mid O^T O = I \right\}$$

where O is an orthogonal matrix. This last inequality, applied to the Hessian of a convex function, corresponds to taking products of second derivatives of the function along orthogonal directions,

$$\det(D^2 u) \equiv \min_{\{\nu_1 \ldots \nu_d\} \in V} \prod_{j=1}^{d} u_{\nu_j \nu_j},$$

where V is the set of all orthonormal bases for $\mathbb{R}^d$.

In the special case where the source density ρ_X vanishes, the Monge-Ampère operator reduces to the convex envelope operator [Obe07, OS09]. In this case, the operator enforces directional convexity in each direction, which is approximated by looking in grid directions [Obe08b, Obe08a]. In order to enforce the convexity constraint, we need to replace the derivatives with their positive part. As well, to guard against the possibility of non-convex solutions when the right-hand side vanishes, we will also subtract the negative parts of these second derivatives,

$$\det{}^+(D^2 u) \equiv \min_{\{\nu_1 \ldots \nu_d\} \in V} \left\{ \prod_{j=1}^{d} \max\{u_{\nu_j \nu_j}, 0\} + \sum_{j=1}^{d} \min\{u_{\nu_j \nu_j}, 0\} \right\}.$$

These modifications ensure that a non-convex function cannot solve our Monge-Ampère equation (with non-negative right-hand side) since it will lead to a negative value of the "convexified" Monge-Ampère operator.

To discretize this, we limit ourselves to considering a finite number of vectors ν that lie on the grid and have a fixed maximum length; this is the directional discretization, which gives us the angular resolution $d\theta$ of our stencil (see Figure 4). In this figure, values on the boundary are used to maintain the directional resolution $d\theta$ (at the expense of lower order accuracy in space because the distances from the reference point are not equal). Another option is to narrower stencils as the boundary is approached, which leads to lower angular resolution, but better spatial resolution. We denote the resulting set of orthogonal vectors by $\mathcal{G}$.

Each of the directional derivatives in the Monge-Ampère operator is then discretized using centered differences:

$$\mathcal{D}_{\nu\nu} u_i = \frac{1}{\|\nu\|^2 h^2} \left(u(x_i + \nu h) + u(x_i - \nu h) - 2u(x_i) \right).$$

In order to handle non-constant densities, we also need to discretize the gradient. A simple approach to this is akin to the Lax-Friedrichs scheme,

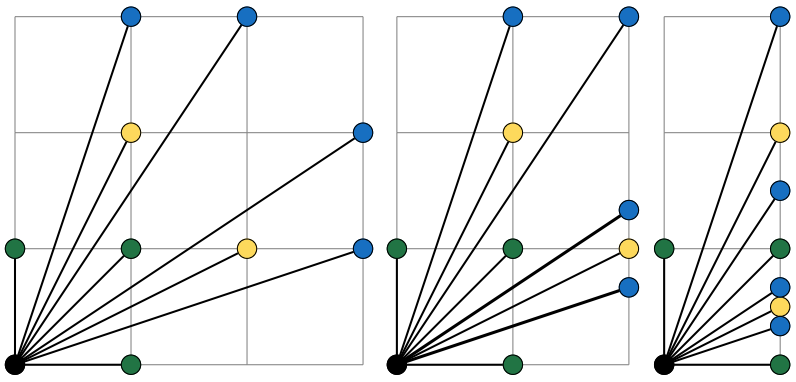

FIGURE 4. Neighboring grid points used for width one (green), two (yellow), and three (blue) stencils. The illustration shows the neighbors in the first quadrant. The modification near the boundary is illustrated in the second and third figures.

which involves adding a small multiple of the laplacian and discretizing with centered differences.

$$-\rho_X/\rho_Y(\nabla u) \approx -\rho_X/\rho_Y\left(\mathcal{D}_{x_1}u, \ldots, \mathcal{D}_{x_d}u\right) + \delta \sum_{j=1}^{d} \mathcal{D}_{x_j x_j} u.$$

The centered difference discretization of the first derivatives is

$$\mathcal{D}_\nu u_i = \frac{1}{2h}\left(u(x_i + \nu h) - u(x_i - \nu h)\right).$$

To preserve monotonicity, we require the parameter δ to satisfy $\delta > Kh$ where K is the Lipschitz constant (with respect to y) of $\rho_X(x)/\rho_Y(y)$.

Remark 3.6. *In practice, we do not add a multiple of the laplacian, but instead absorb the parameter δ into the second-derivative operators that are already present in the equation. The gradient can then be discretized using centered differences along rotated coordinate frames, which correspond to the directions that are active in the Monge-Ampère operator. This approach, which is described in* [Fro12]*, leads to sparser systems and improves the consistency error of the monotone scheme. These differences do not affect the convergence proof or the formal consistency error of the more accurate filtered scheme.*

Then an elliptic discretization of the Monge-Ampère equation is

$$MA_M^{h,d\theta,\delta}[u] = \min_{(\nu_1,\ldots,\nu_d)\in\mathcal{G}} G^{h,d\theta,\delta}_{(\nu_1,\ldots,\nu_d)}[u] \tag{14}$$

where each of the $G^{h,d\theta,\delta}_{(\nu_1,\ldots,\nu_d)}[u]$ is defined as

$$
\begin{aligned}
G^{h,d\theta,\delta}_{(\nu_1,\ldots,\nu_d)}[u] = &\prod_{j=1}^{d} \max\{\mathcal{D}_{\nu_j\nu_j}u, 0\} + \sum_{j=1}^{d} \min\{\mathcal{D}_{\nu_j\nu_j}u, 0\} \\
&- \rho_X(x)/\rho_Y\left(\mathcal{D}_{x_1}u, \ldots, \mathcal{D}_{x_d}u\right) + \delta \sum_{j=1}^{d} \mathcal{D}_{x_jx_j}u - u(x_0).
\end{aligned} \tag{15}
$$

This monotone scheme forms the basis of the filtered scheme (12). For improved accuracy on smooth solutions, we combine it with the accurate scheme F_A, which is simply a standard centered difference discretization of the (two-dimensional) equation

$$u_{x_1x_1}u_{x_2x_2} - u^2_{x_1x_2} - \rho_X(x)/\rho_Y(u_{x_1}, u_{x_2}) - u(x_0).$$

We denote the resulting discretization by $MA_S[u]$. A similar discretization is easily constructed in higher dimensions. Additional details can be found in [FO11b].

3.4. **Proof of convergence.** We combine the almost monotone schemes for (MA) with the upwind, monotone scheme for (HJ) into one equation, which we show converges to the unique convex viscosity solution of the system (MA),(HJ). The combined scheme is given as

$$F^{h,d\theta,d\alpha}[u_i] = \begin{cases} F^{h,d\theta}_F[u_i] & x_i \in X \\ H^{h,d\alpha}[u_i] & x_i \in \partial X. \end{cases} \tag{16}$$

In this definition, F_F is the filtered scheme for the Monge-Ampère equation (12), which relies on the discretizations described in subsection 3.3, and H is the upwind discretization of the boundary condition described in subsection 2.4.

Theorem 3.7 (Convergence)**.** *Let u be the unique convex viscosity solution of the PDE (MA) with boundary condition (HJ). Let $u^{h,d\theta,d\alpha}$ be a solution of the finite difference scheme (16). Then $u^{h,d\theta,d\alpha}$ converges uniformly to u as $h, d\theta, d\alpha \to 0$.*

Proof. From Theorem 3.5, we need only verify that the scheme is consistent and nearly elliptic (nearly monotone).

Consistency and near monotonicity of the scheme for the Monge-Ampère equation have been established in [FO11a, Fro12].

We recall the form of the boundary condition in (7),

$$H(\nabla u(x)) = \sup_{\|n\|=1} \{\nabla u(x) \cdot n - H^*(n) \mid n \cdot n_x > 0\}.$$

This is discretized using forward or backward differences for the gradient, which are consistent as $h \to 0$. The supremum is further approximated by restricting to a finite subset of directions, with resolution $d\alpha$. Since the

Legendre-Fenchel transform $H^*(n)$ is continuous, and since these admissible directions are approximated with an accuracy on the order of $d\alpha$, this approximation is consistent as $d\alpha \to 0$.

By exploiting the obliqueness property (Lemma 2.4), we were able to construct an upwind discretization of the boundary condition, which is monotone by construction. □

4. Solution methods

In the preceding sections, we have described a method for approximating the second boundary value problem for the Monge-Ampère equation by a system of nonlinear equations on a grid. Now we describe several different approaches for solving these equations.

4.1. **Explicit iteration.** The simplest approach to solving the almost-monotone system is to simply perform a forward Euler iteration on the parabolic equation

$$u_t = \det(D^2 u) = \rho_X / \rho_Y(\nabla u), \quad x \in X.$$

In order to enforce the boundary conditions, we can evolve the Hamilton-Jacobi equation

$$u_t = H(\nabla u), \quad x \in \partial X.$$

By allowing this system to evolve to steady state, we can obtain the solution of the discrete system.

While this explicit solution method will work, it is subject to a very restrictive nonlinear CFL condition [Obe06]. Consequently, this approach is very slow.

4.2. **Newton's method.** A much more efficient approach is to use an implicit solver, such as Newton's method. The filtered scheme for the Monge-Ampère equation has previously been solved with Newton's method [FO12] in the case of Dirichlet boundary conditions. Because we are now enforcing an implicit (nonlinear) boundary condition, the Newton solver must also be used to enforce the correct boundary values.

4.2.1. *The Monge-Ampère equation.* We begin by reviewing the form of a damped Newton's method for the filtered discretization of the Monge-Ampère equation in the interior of the computational domain. This involves performing an iteration of the form

$$u^{k+1} = u^k - \alpha(\nabla MA[u^k])^{-1} MA[u^k]$$

where the Jacobian is given by

$$\nabla MA[u] = \big(1 - S'[u]\big)\, \nabla MA_M[u] + S'[u] \nabla MA_S[u].$$

The derivative of the filter (13) is

$$S'(x) = \begin{cases} 1 & \|x\| < 1 \\ -1 & 1 < \|x\| < 2 \\ 0 & \|x\| > 2. \end{cases}$$

To prevent this from taking on negative values, which can lead to poorly conditioned or ill-posed linear systems, we approximate the Jacobian by

$$\tilde{\nabla} MA[u] = \left(1 - S'[u]\right) \nabla MA_M[u] + \max\{S'[u], 0\} \nabla MA_S[u].$$

The damping parameter α is chosen to ensure that the residual is decreasing.

This iteration requires the Jacobians of the monotone and accurate schemes. To simplify the expression, we define $F(x,p) = \rho_X(x)/\rho_Y(p)$. We also use $\mathbb{1}_0$ to denote the matrix that has entries equal to one in the column corresponding to the point x_0. We begin with the monotone scheme, recalling that this discretization has the form

$$MA_M[u] = \min_{(\nu_1,\ldots,\nu_d)\in\mathcal{G}} G_{(\nu_1,\ldots,\nu_d)}[u].$$

By Danskin's Theorem [Ber03], we can write the Jacobian of this as

$$\nabla MA_M[u] = \nabla G_{(\nu_1,\ldots,\nu_d)}[u],$$

where the $(\nu_1,\ldots,\nu_d)$ are the directions active in the minimum. We also have

$$\begin{aligned} \nabla_{u_i} G_{(\nu_1,\ldots,\nu_d)}[u] = \\ \sum_{m=1}^{d} \left[\left(\prod_{j\neq m} \max\{\mathcal{D}_{\nu_j\nu_j} u_i, 0\} \right) \mathbb{1}_{\mathcal{D}_{\nu_j\nu_j} u_i \geq 0} + \mathbb{1}_{\mathcal{D}_{\nu_j\nu_j} u_i < 0} \right] \mathcal{D}_{\nu_m\nu_m} \\ - \sum_{j=1}^{d} \frac{\partial F}{\partial p_j}\left(x, \mathcal{D}_{x_1} u_i, \ldots, \mathcal{D}_{x_d} u_i\right) \mathcal{D}_{\nu_j} + \delta \sum_{j=1}^{d} \mathcal{D}_{x_j x_j} - \mathbb{1}_0. \end{aligned}$$

Remark 4.1. *In fact, as mentioned in Remark 3.6, we use a slightly more sophisticated discretization. The resulting Jacobian is similar to this, though slightly more complicated; see* [Fro12] *for details.*

The Jacobian of the accurate scheme is (in two dimension)

$$\begin{aligned} \nabla_{u_i} MA_S[u] = (\mathcal{D}_{x_2x_2} u_i)\mathcal{D}_{x_1x_1} + (\mathcal{D}_{x_1x_1} u_i)\mathcal{D}_{x_2x_2} + 2(\mathcal{D}_{x_1x_2} u_i)\mathcal{D}_{x_1x_2} \\ - \frac{\partial F}{\partial p_1}(x, \mathcal{D}_{x_1} u_i, \mathcal{D}_{x_2} u_i)\mathcal{D}_{x_1} - \frac{\partial F}{\partial p_2}(x, \mathcal{D}_{x_1} u_i, \mathcal{D}_{x_2} u_i)\mathcal{D}_{x_2} - \mathbb{1}_0. \end{aligned}$$

4.2.2. *The boundary condition (BV2).* The Newton update is also used to enforce the correct boundary conditions.

We recall that the discrete form of the boundary condition obtained in section 3 can be written in compact form

$$H[u_i] = \max_{n\in\Lambda_i} \big\{ \max\{n_1,0\}\mathcal{D}_x^- u_i + \min\{n_1,0\}\mathcal{D}_x^+ u_i + \max\{n_2,0\}\mathcal{D}_y^- u_i + \min\{n_2,0\}\mathcal{D}_y^- u_i - H^*(n) \big\}.$$

Again we stress that at each point x_i on the source boundary, the admissible set Λ_i will only include directions n that allow for upwinding. We also recall that these directions, and the corresponding values of $H^*(n)$, are given or computed before the start of the Newton solve.

Now at boundary points, the Newton step will take the form

$$u^{k+1} = u^k - \alpha \left(\nabla H[u^k] \right)^{-1} H[u^k].$$

As with the Monge-Ampère operator itself, we can use Danskin's Theorem to compute the Jacobian at these points:

$$\nabla_{u_i} H[u] = \max\{n_1,0\}\mathcal{D}_x^- + \min\{n_1,0\}\mathcal{D}_x^+ + \max\{n_2,0\}\mathcal{D}_y^- + \min\{n_3,0\}\mathcal{D}_y^+$$

where again, the value of ν used is simply the direction that is active in the maximum.

4.3. **The projection method.** As an alternative solution method, we also consider the projection method described in [Fro12]. We show that it can be viewed as a simple iterative relaxation strategy for solving (MA-HJ).

The idea of this approach is that we wish to solve for the solution u to the Monge-Ampère equation in the domain, as well as its gradient $p = \nabla u$ on the boundary of the domain. We use a splitting approach that involves alternating between:

(i) the solution of a Monge-Ampère equation with Neumann boundary data obtained from the current estimate of the gradient map p at the boundary.
(ii) the solution of the Hamilton-Jacobi equation for the gradient p at the boundary.

More precisely, step (i) involves updating u^{k+1} by solving the following Monge-Ampère equation with Neumann boundary conditions.

$$\begin{cases} \det(D^2 u^{k+1}(x)) = \rho_X(x)/\rho_Y(\nabla u^{k+1}(x)) + \langle u^{k+1} \rangle & x \in X \\ \nabla u^{k+1} \cdot n_x = p_k \cdot n_x & x \in \partial X \\ u^{k+1} \text{ is convex.} \end{cases}$$

The gradient map ∇u^{k+1} obtained from solving this PDE may not solve the correct Hamilton-Jacobi boundary condition. Thus step (ii) involves updating the value of p^{k+1} to ensure that it does satisfy $H(p^{k+1}) = 0$. To ensure that we are making use of the information from the Monge-Ampère

equation, we define an intermediate value of the map $p^{k+1/2} = \nabla u^{k+1}$ at the boundary. Next, we update using a gradient descent approach.

$$p^{k+1} = p^{k+1/2} - H(p^{k+1/2})\nabla H(p^{k+1/2}).$$

This step has a geometric interpretation. Should we miss the target set so that $\nabla u^{k+1} \notin Y$, the new value of p^{k+1} is precisely its projection onto the boundary of the target (∂Y). Consequently, we recover the projection method of [Fro12].

5. Numerical study

We now provide the particular details of our numerical implementation of the second boundary value problem. This is followed by computational results for a number of challenging examples.

5.1. **Extension of densities.** When computing with finite difference methods, it is most convenient to work in rectangular domains. However, it is often desirable to solve the mass transport problem in more general domains. A simple solution, allowed by our solver, is to extend the density function ρ_X into a square, assigning it the value zero at points outside the set X that we are interested in. This will lead to a degenerate Monge-Ampère equation, but the filtered scheme is robust enough to handle this problem. As we will show in the numerics section, this approach allows to treat non-convex or non-connected domains, or even Dirac measures in the source density.

We emphasize that this trick cannot be used to extend the target density ρ_Y into a square. This is because the convergence and discretization error of the scheme are dependent on the Lipschitz constant of $\rho_X(x)/\rho_Y(y)$. Even if the target density is smoothly extended to have a small positive density ϵ outside the target set Y, this function will have a very large Lipschitz constant, which makes it impractical to obtain computational results with any reasonable accuracy.

However, it is still important to extend the target density so that it is defined in all space. This is because, while the given optimal transportation problem will require a density ρ_Y that is defined only in the set Y, we may be required to compute this density function at other points during the process of numerically solving the equation. As we have just noted, we cannot simply allow the density to vanish outside the target set. Instead, we must use a positve, Lipschitz continuous extension of ρ_Y into all space. In some cases, when ρ_Y is a given function, there is an obvious way of extending it into all space. Lipschitz extensions can always be obtained by using, for example, the method of [Obe05]. The resulting function $\rho_Y^*(y)$ can always be bounded away from zero by considering $\max\{\rho_Y^*(y), \rho_0\}$ for some $0 < \rho_0 \leq \min_{y \in Y} \rho_Y(y)$.

5.2. **Details of implementation.**

5.2.1. *Discretization of Monge-Ampère operator.* There are several sources of discretization error in this method. The first source is the discretization of the Monge-Ampère equation in the interior of the domain. This discretization includes three small parameters that contribute to the error: the spatial step size dx coming from the number of grid points N_X along each dimension, the angular resolution $d\theta$ coming from the width of the stencil, and the parameter δ that is used to bound the eigenvalues of the Hessian away from zero.

The precise accuracy we can expect to achieve will depend on the regularity of the solution. For sufficiently regular solutions, we expect the filtered scheme to reduce to the more accurate scheme, which has a formal accuracy of $\mathcal{O}(dx^2)$. For singular solutions, we cannot expect high accuracy. With this in mind, we simply use a narrow (9 point) stencil version of the scheme.

We previously observed that for convergence of general problems, the parameter δ should be $\mathcal{O}(dx)$. However, when the target density ρ_Y is constant, δ can be arbitrarily small. With this in mind, we set the parameter $\delta = dx^2$ so that it will not have an appreciable effect on the formal discretization error. The filtered scheme also requires us to define the maximum size of the non-monotone perturbation; in our computations we set this to $\sqrt{dx} + d\theta$.

5.2.2. *Discretization of transport equations.* The implementation of the boundary conditions requires the discretization of a number of transport equations. We use simple upwinding to accomplish this, which leads to a discretization error that is $\mathcal{O}(dx)$.

5.2.3. *Discretization of target set.* If the target set is polyhedral, we can determine exactly which vectors n are needed to represent the target, as well as the exact value of the Legendre-Fenchel transform of the distance function.

However, in the general case, our method leads us to approximate the target set by a polygon, which introduces additional error into the computations. In the discretization of the boundary condition (4), we will consider only N_Y different directions $n_j = 2\pi j/N_Y$, which represent the normal vectors to the polyhedral approximation. The boundary condition is the supremum of functions that are linear in n, so we expect this approximation to lead to error that is first order: $\mathcal{O}(1/N_Y)$.

We must also provide a discretization of the target boundary in order to estimate the Legendre-Fenchel transform, which is done using the formula (5). This requires us to compute the supremum of terms that are linear in the boundary points y_0 so that the discretization leads to first-order accuracy. In our computations, we discretize the boundary using $4N_X$ points, which is similar to the number of grid points along the boundary of the domain X.

5.3. **Computational examples.** We now provide several computational tests to validate the theoretical claims made in this paper.

As discussed in §5.1, we use a square domain in every example, setting the source density ρ_X to zero where necessary.

Whenever an exact solution is available, we provide the maximum norm of the distance between the mappings obtained from the exact and computed solutions (u_{ex} and u_{comp} respectively):

$$\max_{x \in X} \|\nabla u_{ex}(x) - \nabla u_{comp}(x)\|_2.$$

We also provide the number of Newton iterations and total computation time for computations done using the most refined target boundary (largest value of N_Y). We note that there is no appreciable difference in computation time as the value of N_Y is varied from 8 to 256.

5.3.1. *Mapping a square to a square.* We begin by recovering a mapping with an exact solution, which involves mapping a square onto a square. To set up this example, we define the function

$$q(z) = \left(-\frac{1}{8\pi}z^2 + \frac{1}{256\pi^3} + \frac{1}{32\pi}\right)\cos(8\pi z) + \frac{1}{32\pi^2}z\sin(8\pi z).$$

Now we map the density

$$f(x_1, x_2) = 1 + 4(q''(x_1)q(x_2) + q(x_1)q''(x_2)) + 16(q(x_1)q(x_2)q''(x_1)q''(x_2) - q'(x_1)^2 q'(x_2)^2)$$

in the square $(-0.5, 0.5) \times (-0.5, 0.5)$ onto a uniform density in the same square. This transport problem has the exact solution

$$u_{x_1}(x_1, x_2) = x_1 + 4q'(x_1)q(x_2), \quad u_{x_2}(x_1, x_2) = x_2 + 4q(x_1)q'(x_2).$$

This gradient map is picture in Figure 5.

Since the target set is a square, there is no need to discretize its boundary in this case.

Results are presented in Table 1. As anticipated in §5.2, we observe first-order accuracy.

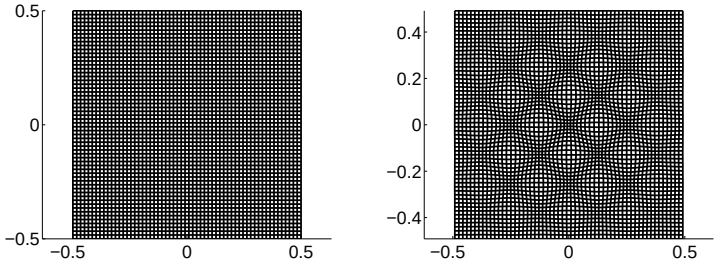

(a) (b)

FIGURE 5. (a) A cartesian mesh X and (b) its image under the gradient map ∇u (§5.3.1).

N_X	Maximum Error	L^2 Error	Iterations	CPU Time (s)
32	0.0220	0.0127	5	0.3
64	0.0110	0.0064	9	1.3
128	0.0055	0.0032	9	6.6
256	0.0028	0.0016	11	41.6
362	0.0020	0.0011	13	101.9

TABLE 1. Distance between exact and computed gradient maps for map from a square to a square. The number of Newton iterations and computation time are also given.

5.3.2. *Mapping an ellipse to an ellipse.* Next, we consider the problem of mapping an ellipse onto an ellipse. To describe the ellipses, we let M_x, M_y be symmetric positive definite matrices and let B_1 be the unit ball in $\mathbb{R}^d$. Now we take $X = M_x B_1$, $Y = M_y B_2$ to be ellipses with constant densities f, g in each ellipse.

In $\mathbb{R}^2$, the optimal map can be obtained explicitly [MO04] from

$$\nabla u(x) = M_y R_\theta M_x^{-1} x$$

where R_θ is the rotation matrix

$$R_\theta = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix},$$

the angle θ is given by

$$\tan(\theta) = \operatorname{trace}(M_x^{-1} M_y^{-1} J) / \operatorname{trace}(M_x^{-1} M_y^{-1}),$$

and the matrix J is equal to

$$J = R_{\pi/2} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

We use the particular example

$$M_x = \begin{pmatrix} 0.8 & 0 \\ 0 & 0.4 \end{pmatrix}, \quad M_y = \begin{pmatrix} 0.6 & 0.2 \\ 0.2 & 0.8 \end{pmatrix},$$

which is pictured in Figure 6.

Results are presented in Table 2, which demonstrates first order accuracy in both N_X and N_Y.

5.3.3. *Mapping from a disconnected region.* We now consider the problem of mapping the two half-circles

$$X = \{(x_1, x_2) \mid x_1 < -0.2, (x_1 + 0.2)^2 + x_2^2 < 0.85^2\} \cup \{(x_1, x_2) \mid x_1 > 0.1, (x_1 - 0.1)^2 + x_2^2 < 0.85^2\}$$

onto the circle

$$Y = \{(y_1, y_2) \mid y_1^2 + y_2^2 < 0.85^2\}.$$

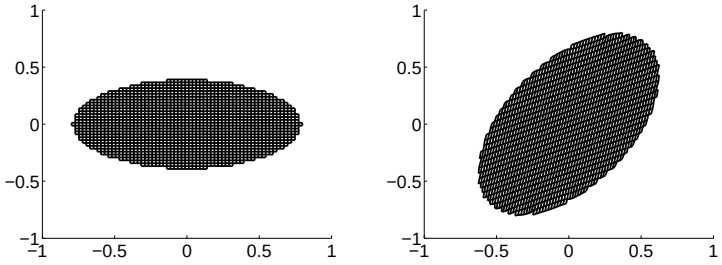

(a) (b)

FIGURE 6. (a) An ellipse X and (b) its image under the gradient map ∇u (§5.3.2).

N_X	Maximum Error N_Y						Iterations	Time (s)
	8	16	32	64	128	256		
32	0.1163	0.0773	0.0693	0.0669	0.0665	0.0062	3	0.6
64	0.1188	0.0403	0.0302	0.0291	0.0282	0.0283	4	1.0
128	0.1214	0.0302	0.0201	0.0174	0.0168	0.0168	4	4.2
256	0.1206	0.0278	0.0116	0.0101	0.0092	0.0091	4	20.8
362	0.1175	0.0291	0.0098	0.0063	0.0057	0.0056	5	43.7

TABLE 2. Distance between exact and computed gradient maps for map from an ellipse to an ellipse. The number of Newton iterations and computation time is given for the largest number of directions (256).

This example, which is pictured in Figure 7, is extremely degenerate since the domain is not simply connected. Nevertheless, our method correctly computes the optimal map, as the results of Table 3 verify.

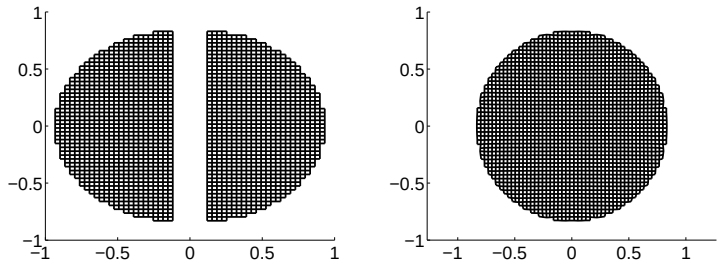

(a) (b)

FIGURE 7. (a) Two half-circles X and (b) its image under the gradient map ∇u (§5.3.3).

N_X	Maximum Error N_Y						Iterations	Time (s)
	8	16	32	64	128	256		
32	0.0453	0.0267	0.0255	0.0258	0.0259	0.0258	4	0.2
64	0.0397	0.0184	0.0158	0.0146	0.0144	0.0139	4	1,2
128	0.0392	0.0097	0.0063	0.0066	0.0065	0.0064	5	4.5
256	0.0432	0.0110	0.0084	0.0087	0.0086	0.0073	5	24.9
362	0.0448	0.0130	0.0070	0.0047	0.0045	0.0039	5	45.0

TABLE 3. Distance between exact and computed gradient maps for map from two semi-circles to a circle. The number of Newton iterations and computation time is given for the largest number of directions (256).

5.3.4. *Inverse mapping.* Next, we show that we can use our method to recover inverse mappings. In this particular example, we compute in the unit square using variable densities in both the source and the target set.

The target density is simply a gaussian in the center of the domain:

$$\rho_Y(y) = 2 + \frac{1}{0.2^2} \exp\left(-\frac{0.5|x|^2}{0.2^2}\right).$$

For the source density, we use four gaussians centered at the four corners of the domain. For example, in the quadrant $[-1, 0] \times [-1, 0]$, we use

$$\rho_X(x) = 2 + \frac{1}{0.2^2} \exp\left(-\frac{0.5|x - (-1, -1)|^2}{0.2^2}\right).$$

These density functions are pictured in Figure 8.

To visualize the mapping between these non-constant densities, we plot several curves in the domain X, as well as the image of these curves under the gradient map. These are also in Figure 8.

The optimal mapping is computed in two different ways:

(1) By solving the problem directly.
(2) By solving the inverse problem (mapping ρ_Y to ρ_X), the inverting the resulting gradient map.

In order to check that these two approaches produce the same result, we look at the distance between the two maps (Table 4). Even for this challenging example, which involves splitting the gaussian into several pieces or joining these pieces back together, the difference between the two computed maps depends linearly on the spatial resolution h.

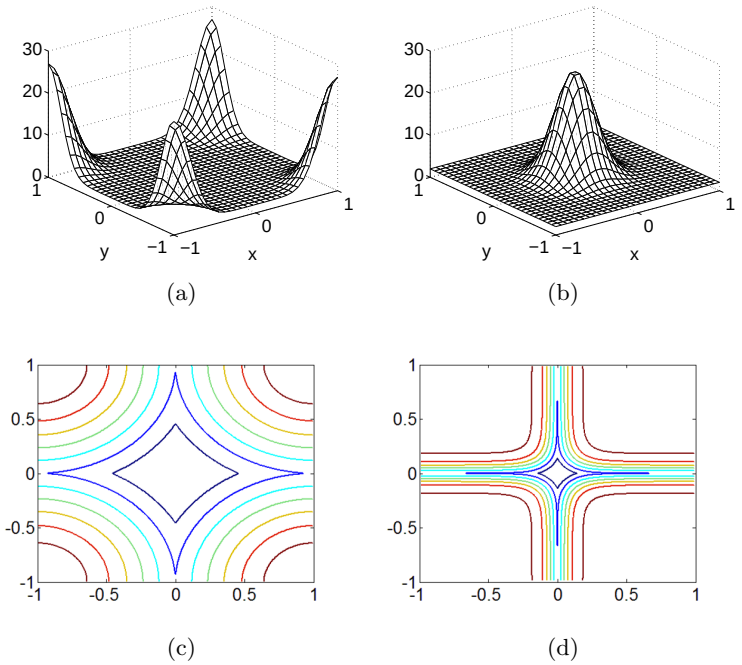

FIGURE 8. (a) The source density ρ_X and (b) target density ρ_Y. (c) Curves in the domain X and (d) their image under the gradient map.

5.3.5. *Mapping from a non-convex source.* As another challenging computational example, we consider the problem of mapping from a non-convex domain, which can lead to a breakdown in regularity. In this example, we choose a domain shaped like the letter "C", which is mapped into the unit circle. See Figure 9 for images of these sets, as well as the computed gradient map.

N_X	Max Distance	Iterations		CPU Time (s)	
		Forward	Inverse	Forward	Inverse
32	0.2337	6	3	0.3	0.2
64	0.1252	6	3	1.1	1.1
128	0.0649	8	3	6.3	5.2
256	0.0329	9	4	35.8	32.6
362	0.0233	11	4	95.2	52.2

TABLE 4. Distance between forward map and the inverse of the computed inverse map. The number of Newton iterations and computation time is also given for the two different approaches for computing the map.

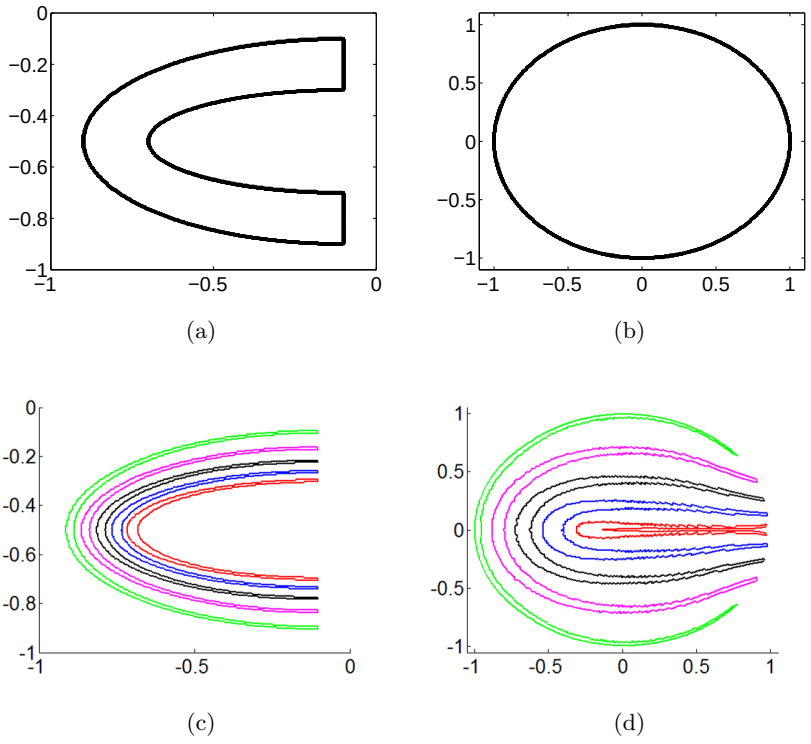

FIGURE 9. The boundaries of the (a) source X and (b) target Y sets. (c) Curves in the domain and (d) their image under the gradient map.

5.3.6. *Other geometries.* To give a flavor of the types of geometry that can be captured by solving (HJ), we provide several different maps in Figure 10. These were all obtained by mapping a square with uniform density onto a specified convex set, whose boundary could consist of a combination of

straight and curved edges. While no exact solutions are available, we can certainly verify that the computed gradient does map the square into the correct target set.

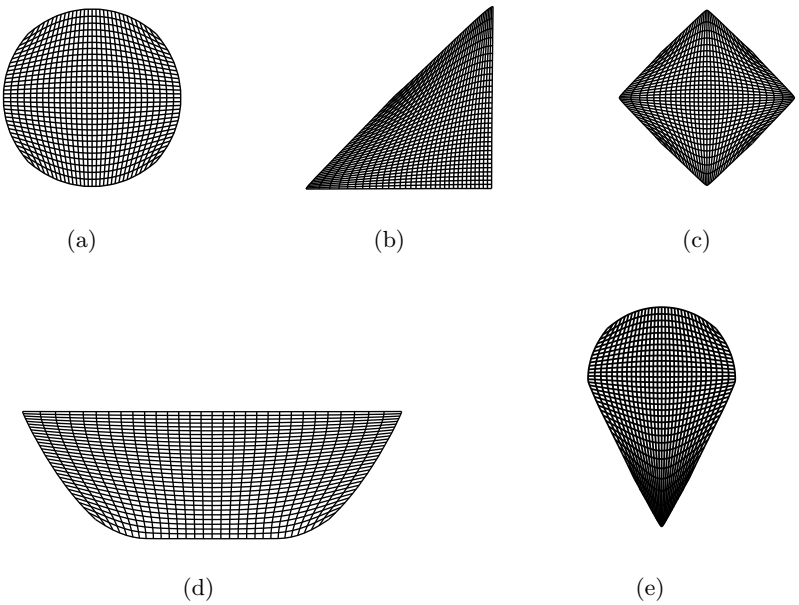

FIGURE 10. Maps computed using (HJ)

5.3.7. *Pogorelov solutions.* In this section we demonstrate the robustness of the solver when dealing with very singular solutions where the target density is a sum of possibly weighted Dirac masses positioned at points $\{y_j\}_{j=1..N_d}$:

$$\rho_Y = \sum_{j=1}^{j=Nd} q_j \, \delta_{y_j}.$$

The initial density is taken to be the Lebesgue measure on a closed convex subset of X, but any non-vanishing smooth probability density could be treated:

$$\rho_X = \mathbb{1}_X.$$

This configuration is actually the "reconstruction of the early shape of the universe" model used in [FMMS02].

Classical solution methods are based on the Pogorelov remark that the solution (a convex potential) is necessarily of the form

$$u(x) = \max_{j=1..N_d} \{x \cdot y_j - v_j\}$$

and a necessary and sufficient condition on the $\{v_j\}_{j=1..N_d}$ to solve the Monge-Kantorovitch problem is shown to be

$$\rho_X(V_j) = q_j, \ j = 1..N_d \tag{17}$$

where $V_j = \{x \in X, \ j = \text{argmax}_k \left[x \cdot y_k - v_k\right]\}$ are the support cells of each of the affine functions. Modern methods rely on the quantization of the density ρ_X to a sum of Diracs. It reduces the problem to a classical assignment problem which can be solved by the Auction algorithm in $O(N^2 \log N)$ operations, where N is the number of points used in the quantization (assuming this is bigger than N_d). See [Bos10] for a review and also [Mer11] for a multiscale approach.

The Monge-Ampère solver used in this paper can deal with a singular right-hand side (ρ_X), but needs smooth target densities (ρ_Y). It does not rely on quantization, just on the discretization of the X and Y spaces using uniform grids. There is no theoretical proof of convergence, but the numerical experiments indicate that the cost of the Newton's method is linear in terms of the number of grid points (N^2). The method will therefore not depend on the (delicate) quantization step. This means that the cost is the same for 1 or 1000 Dirac masses.

In order to solve the problem above, we simply remark that the optimal $\{v_j\}$ are actually

$$v_j = u^*(y_j), \ j = 1, \ldots, N_d \tag{18}$$

where u^* is the Legendre-Fenchel transform of u and the solution of the dual problem of mapping ρ_Y (singular) to ρ_X (smooth). We first solve this dual problem using our approach, then reconstruct the potential u and the cells $\{V_J\}$s from the solution[2]. Note that it is an extremely singular right-hand side for the Monge-Ampère equation. Given M, the mass of the target density, we brutally set the weight at grid points representing the Dirac masses to $\frac{M}{N_d h^2}$ and 0 elsewhere. As we will see, the solution necessarily loses accuracy.

Figures (11) to (13) show randomly positioned Dirac masses embedded in a square, which are mapped to a uniform density on a ball. We plot the cells, and the colormap indicates the error with respect to the optimality condition (17). We also show the convex potential, noting that there is a one to one correspondence between the gradient of the affine facets and the positions of the target Dirac masses. We use a 256×256 discretization. All computation are done in Matlab on a 2.2 GHz intel Core I7 Laptop with 8 GB of RAM.

Table (5) shows the maximum and L^2 percentage error, together with the (modest) runtime. We recall that the convex potential computed by our method is the dual conjugate of the one represented in the figures.

[2]We use the MPT Matlab toolbox http://control.ee.ethz.ch/~mpt/.

This potential will be singular at Dirac locations[3] and the solution loses accuracy and causes deterioration in the values we compute for (18). Note, however, that the solution has the correct form of a convex piecewise affine potential. We get the correct number of cells, and the optimal map *ordering* must be correct. These solutions could be used for initialization of exact combinatorial optimization methods.

Table (5) also shows that, with this dual approach, the number of Dirac masses has no impact on the run time. In Table (6) we look at the run times and accuracy for 300 Dirac masses when we increase resolution up to 4 million points (remember the discretisation of the initial square domain is $N_X \times N_X$.) The cost of the method is linear until $N_X = 1024$ because of out of core memory overheads.

N_d	L^∞ Error	L^2 err	CPU Time (s)
3	0.05	0.02	10.48
30	0.48	0.21	10.56
300	0.56	0.18	11.8

TABLE 5. Normalized errors (percentage) for $N_X = 256$

N_X	L^∞ Error	L^2 err	CPU Time (s)
64	0.93	0.28	1.48
128	0.96	0.24	2.6
256	0.83	0.21	11.7
512	0.66	0.20	46.76
1024	0.64	0.20	281.53

TABLE 6. Run time and error (percentage) for increasing grid resolution with 300 Dirac masses.

6. Conclusion and Perspectives

We presented a new algorithm for solving the Monge Kantorovich Optimal transportation problem. The method is robust to source densities that vanish, are singular, or have non-convex support. The target density

[3]Keep in mind that these computations are being done using a compact (9 point) filtered scheme. While this is enough to observe convergence and good accuracy for smooth or even moderately singular solutions, a wide stencil can become necessary for more singular examples. This is what we observed when we computed the cone solution in [FO12]. Additionally, these Pogorolev solutions (and the cone solution) are not actually viscosity solutions, so it is a bit remarkable if we can compute them at all.

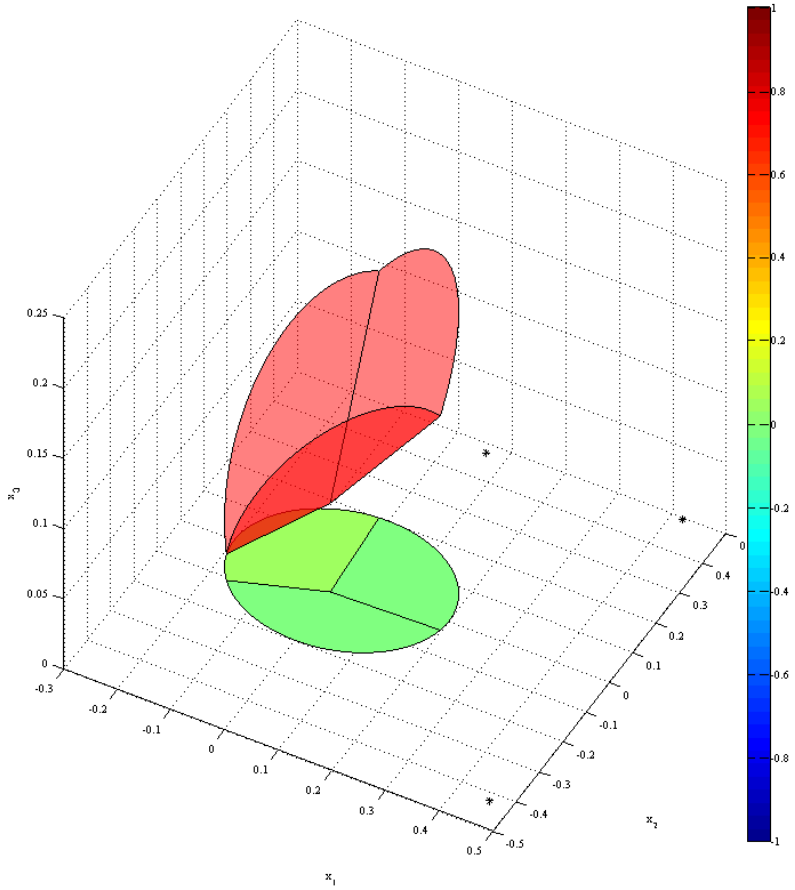

FIGURE 11. The convex potential surface (red) and cell projections V_j, with the colormap indicating the percentage error in cell area, for 3 randomly positioned Dirac masses (stars) mapped to a uniform density on a ball with a 256×256 grid.

is Lipschitz continuous and non-vanishing, with support on a convex set. (Of course, the numerical method allows us to choose the target and source densities, so the restriction is on only one of the densities).

The algorithm run time, using Newton's Method, is experimentally linear in the number of grid points.

Extensions and perspectives to the work presented in this paper include:

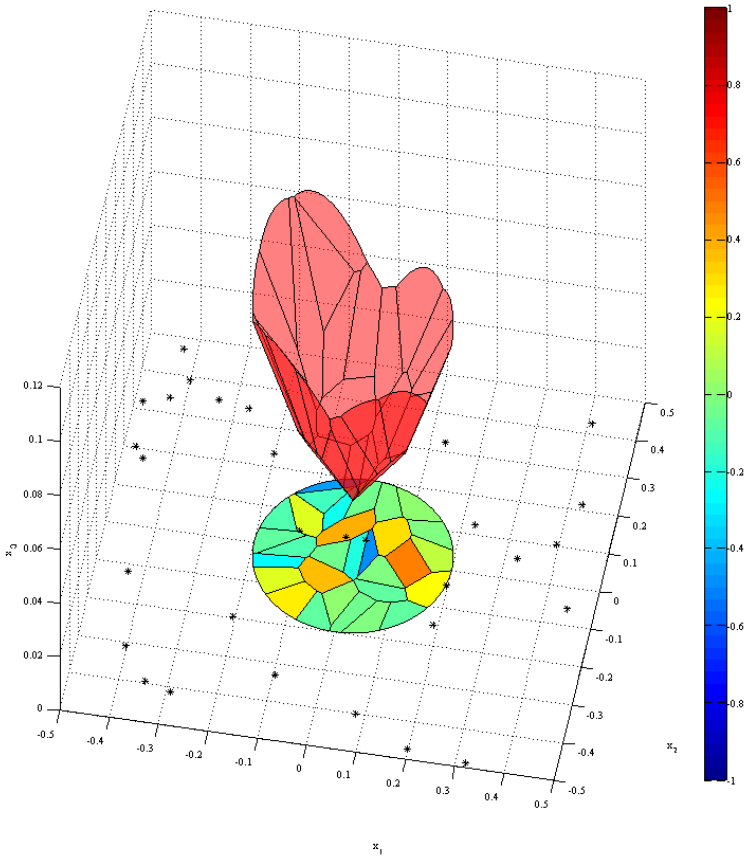

FIGURE 12. The convex potential surface (red) and cell projections V_j, with the colormap indicating the percentage error in cell area, for 30 randomly positioned Dirac masses (stars) mapped to a uniform density on a ball with a 256×256 grid.

- writing a 3D code. The 3D version of the Monge-Ampere solver for classical boundary conditions exist. The extension of the Hamilton-Jacobi solver on the boundary will be straightforward.
- studying the structure of the linear problem in the Newton iterate. We currently use the Matlab backslash operator, which utilizes a direct LU solver. Ellipticity suggests that an iterative Multigrid may be more efficient.
- extending the method as a tool for JKO gradient seems straightforward since the linearized equations will directly give the gradient.

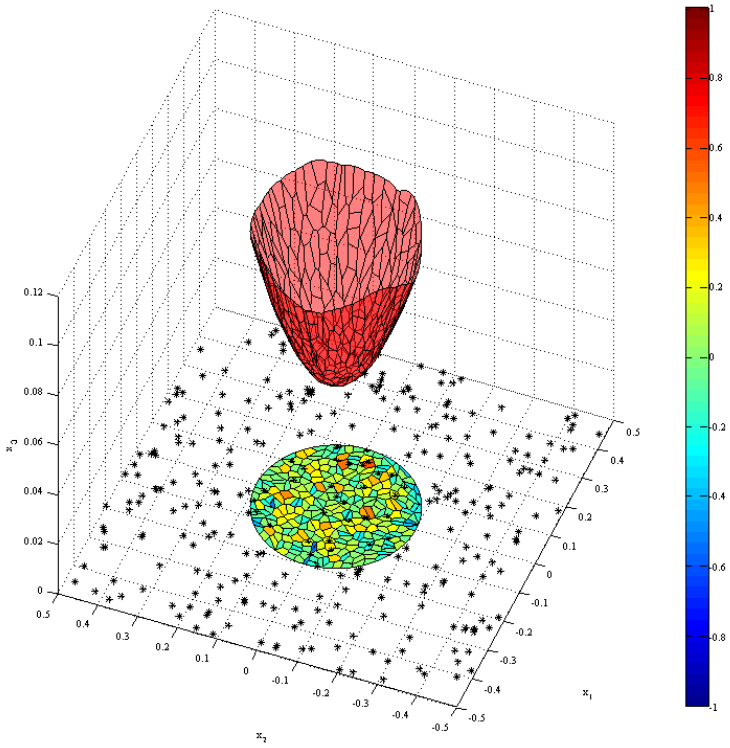

FIGURE 13. The convex potential surface (red) and cell projections V_j, with the colormap indicating the percentage error in cell area, for 300 randomly positioned Dirac masses (stars) mapped to a uniform density on a ball with a 256×256 grid.

- extending to c-convexity. A popular generalization of the Monge-Kantorovitch problem is to replace the Euclidean distance $\|x - y\|^2$ by a convex function $c(x, y)$ in the transportation cost. The theory is well developed and one can write a version of the nonlinear Monge-Ampère PDE for c-convex potential solutions. The numerical solution of this equation is an open problem.

REFERENCES

[BB00] Jean-David Benamou and Yann Brenier, *A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem*, Numer. Math. **84** (2000), no. 3, 375–393. MR1738163 (2000m:65111)

[BB01] J. D. Benamou and Y. Brenier, *Mixed L^2-Wasserstein optimal mapping between prescribed density functions*, J. Optim. Theory Appl. **111** (2001), no. 2, 255–271. MR1865668 (2002h:49069)

[Ber03] Dimitri P. Bertsekas, *Convex analysis and optimization*, Athena Scientific, Belmont, MA, 2003. With Angelia Nedić and Asuman E. Ozdaglar. MR2184037 (2006j:90001)

[Ber92] ———, *Auction algorithms for network flow problems: a tutorial introduction*, Comput. Optim. Appl. **1** (1992), no. 1, 7–66. MR1195629 (93h:90033)

[Bos10] D. Bosc, *Numerical approximation of optimal transport maps*, SSRN (2010).

[Bre91] Yann Brenier, *Polar factorization and monotone rearrangement of vector-valued functions*, Comm. Pure Appl. Math. **44** (1991), no. 4, 375–417. MR1100809 (92d:46088)

[BS91] Guy Barles and Panagiotis E. Souganidis, *Convergence of approximation schemes for fully nonlinear second order equations*, Asymptotic Anal. **4** (1991), no. 3, 271–283. MR92d:35137

[BV04] Stephen Boyd and Lieven Vandenberghe, *Convex optimization*, Cambridge University Press, Cambridge, 2004. MR2061575 (2005d:90002)

[BW09] C. J. Budd and J. F. Williams, *Moving mesh generation using the parabolic Monge-Ampère equation*, SIAM J. Sci. Comput. **31** (2009), no. 5, 3438–3465. MR2538864

[BZ03] J. Frédéric Bonnans and Housnaa Zidani, *Consistency of generalized finite difference schemes for the stochastic HJB equation*, SIAM J. Numer. Anal. **41** (2003), no. 3, 1008–1021. MR2005192 (2004i:49061)

[Caf96] Luis A. Caffarelli, *Boundary regularity of maps with convex potentials. II*, Ann. of Math. (2) **144** (1996), no. 3, 453–496. MR1426885 (97m:35027)

[CDF11] L. Chacón, G. L. Delzanno, and J. M. Finn, *Robust, multidimensional mesh-motion based on Monge-Kantorovich equidistribution*, J. Comput. Phys. **230** (2011), no. 1, 87–103. MR2734282 (2011i:65233)

[CM09] 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. MR2566595 (2011b:65200)

[CNP91] M. J. P. Cullen, J. Norbury, and R. J. Purser, *Generalised Lagrangian solutions for atmospheric and oceanic flows*, SIAM J. Appl. Math. **51** (1991), no. 1, 20–31. MR1089128 (92g:76081)

[CP84] M. J. P. Cullen and R. J. Purser, *An extended Lagrangian theory of semi-geostrophic frontogenesis*, J. Atmospheric Sci. **41** (1984), no. 9, 1477–1497. MR881109 (87k:86011)

[DAT08] Ayelet Dominitz, Sigurd Angenent, and Allen Tannenbaum, *On the computation of optimal transport maps using gradient flows and multiresolution analysis*, Recent advances in learning and control, 2008, pp. 65–78. MR2409075 (2009j:49091)

[Del91] P. Delanoë, *Classical solvability in dimension two of the second boundary-value problem associated with the Monge-Ampère operator*, Ann. Inst. H. Poincaré Anal. Non Linéaire **8** (1991), no. 5, 443–457. MR1136351 (92g:35070)

[FMMS02] Uriel Frisch, Sabino Matarrese, Roya Mohayaee, and Andrei Sobolevski, *A reconstruction of the initial conditions of the universe by optimal mass transportation*, Nature **417** (2002).

[FO11a] Brittany D. Froese and Adam M. Oberman, *Convergent finite difference solvers for viscosity solutions of the elliptic Monge-Ampère equation in dimensions two and higher*, SIAM J. Numer. Anal. **49** (2011), no. 4, 1692–1714. MR2831067

[FO11b] ———, *Fast finite difference solvers for singular solutions of the elliptic Monge-Ampère equation*, J. Comput. Phys. **230** (2011), no. 3, 818–834. MR2745457

[FO12] ———, *Convergent almost monotone finite difference approximations for viscosity solutions of the elliptic monge-ampère partial differential equation*, submitted (2012).

[Fro12] Brittany D. Froese, *A numerical method for the elliptic Monge-Ampère equation with transport boundary conditions*, SIAM J. Sci. Comput. **34** (2012), no. 3, A1432–A1459.

[GO03a] T. Glimm and V. Oliker, *Optical design of single reflector systems and the Monge-Kantorovich mass transfer problem*, J. Math. Sci. (N. Y.) **117** (2003), no. 3, 4096–4108. Nonlinear problems and function theory. MR2027449 (2004k:49101)

[GO03b] ———, *Optical design of single reflector systems and the Monge-Kantorovich mass transfer problem*, J. Math. Sci. (N. Y.) **117** (2003), no. 3, 4096–4108. Nonlinear problems and function theory. MR2027449 (2004k:49101)

[HRT10] Eldad Haber, Tauseef Rehman, and Allen Tannenbaum, *An efficient numerical method for the solution of the L_2 optimal mass transfer problem*, SIAM J. Sci. Comput. **32** (2010), no. 1, 197–211. MR2599774 (2011b:49120)

[HZTA04] Steven Haker, Lei Zhu, Allen Tannenbaum, and Sigurd Angenent, *Optimal mass transport for registration and warping*, Int. J. Comput. Vision **60** (2004), no. 3, 225–240.

[JKO98] Richard Jordan, David Kinderlehrer, and Felix Otto, *The variational formulation of the Fokker-Planck equation*, SIAM J. Math. Anal. **29** (1998), no. 1, 1–17. MR1617171 (2000b:35258)

[Kan42] L. V. Kantorovich, *On the transfer of masses*, Dokl. Akad. Nauk. SSSR **37** (1942), no. 7–8, 227–229.

[Kan48] ———, *On a problem of Monge*, Uspekhi Mat. Nauk. **3** (1948), no. 2, 225–226.

[KW99] David Kinderlehrer and Noel J. Walkington, *Approximation of parabolic equations using the Wasserstein metric*, M2AN Math. Model. Numer. Anal. **33** (1999), no. 4, 837–852. MR1726488 (2000k:65173)

[LPS] Boualem Khouider Louis-Philippe Saumier Martial Agueh, *An efficient numerical algorithm for the l2 optimal transport problem with applications to image processing.*

[LR05] Grégoire Loeper and Francesca Rapetti, *Numerical solution of the Monge-Ampère equation by a Newton's algorithm*, C. R. Math. Acad. Sci. Paris **340** (2005), no. 4, 319–324. MR2121899

[LST10] Aime Lachapelle, Julien Salomon, and Gabriel Turinici, *Computation of mean field equilibria in economics*, Math. Models Methods Appl. Sci. **20** (2010), no. 4, 567–588. MR2647032 (2011e:91027)

[LTU86] P.-L. Lions, N. S. Trudinger, and J. I. E. Urbas, *The Neumann problem for equations of Monge-Ampère type*, Comm. Pure Appl. Math. **39** (1986), no. 4, 539–563. MR840340 (87j:35114)

[Mer11] Q. Merigot, *A multiscale approach to optimal transport*, Comp. Graph. Forum **5** (2011), no. 30, 1583–1592.

[MO04] Robert J. McCann and Adam M. Oberman, *Exact semi-geostrophic flows in an elliptical ocean basin*, Nonlinearity **17** (2004), no. 5, 1891–1922. MR2086155 (2005g:37164)

[MRCSV11] Bertrand Maury, Aude Roudneff-Chupin, Filippo Santambrogio, and Juliette Venel, *Handling congestion in crowd motion modeling*, Netw. Heterog. Media **6** (2011), no. 3, 485–519. MR2826756

[MW53] Theodore S. Motzkin and Wolfgang Wasow, *On the approximation of linear elliptic differential equations by difference equations with positive coefficients*, J. Math. Physics **31** (1953), 253–259. MR14,693i

[Obe04] Adam M. Oberman, *A convergent monotone difference scheme for motion of level sets by mean curvature*, Numer. Math. **99** (2004), no. 2, 365–379.

[Obe05] ———, *A convergent difference scheme for the infinity Laplacian: construction of absolutely minimizing Lipschitz extensions*, Math. Comp. **74** (2005), no. 251, 1217–1230 (electronic). MR2137000 (2006h:65165)

[Obe06] ———, *Convergent difference schemes for degenerate elliptic and parabolic equations: Hamilton-Jacobi equations and free boundary problems*, SIAM J. Numer. Anal. **44** (2006), no. 2, 879–895 (electronic). MR2218974 (2007a:65173)

[Obe07] ———, *The convex envelope is the solution of a nonlinear obstacle problem*, Proc. Amer. Math. Soc. **135** (2007), no. 6, 1689–1694 (electronic).

[Obe08a] ———, *Computing the convex envelope using a nonlinear partial differential equation*, Math. Models Methods Appl. Sci. **18** (2008), no. 5, 759–780. MR2413037 (2009d:35102)

[Obe08b] ———, *Wide stencil finite difference schemes for the elliptic Monge-Ampère equation and functions of the eigenvalues of the Hessian*, Discrete Contin. Dyn. Syst. Ser. B **10** (2008), no. 1, 221–238. MR2399429 (2009f:35101)

[OS09] Adam M. Oberman and Luis Silvestre, *The dirichlet problem for the convex envelope*, Trans. Amer. Math. Soc. (2009), (16 pages). accepted.

[Ott01] Felix Otto, *The geometry of dissipative evolution equations: the porous medium equation*, Comm. Partial Differential Equations **26** (2001), no. 1-2, 101–174. MR1842429 (2002j:35180)

[Pog94] A. V. Pogorelov, *Generalized solutions of Monge-Ampère equations of elliptic type*, A tribute to Ilya Bakelman (College Station, TX, 1993), 1994, pp. 47–50. MR1423367

[RU00] Ludger Rüschendorf and Ludger Uckelmann, *Numerical and analytical results for the transportation problem of Monge-Kantorovich*, Metrika **51** (2000), no. 3, 245–258 (electronic). MR1795372 (2002c:60021)

[Urb97] John Urbas, *On the second boundary value problem for equations of Monge-Ampère type*, J. Reine Angew. Math. **487** (1997), 115–124. MR1454261 (98f:35057)

Centre de recherche INRIA Paris – Rocquencourt
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399