

The boundary method for semi-discrete optimal transport partitions and Wasserstein distance computation[☆]

Luca Dieci

*School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332 U.S.A.
Tel.: +1 404-894-9209 Fax: +1 404-894-4409*

J.D. Walsh III

*Naval Surface Warfare Center, Panama City Division (X24), 110 Vernon Ave., Panama City, FL 32407 U.S.A.
Tel.: +1 850-234-4660 Fax: +1 850-235-5374*

Abstract

We introduce a new technique, which we call the *boundary method*, for solving semi-discrete optimal transport problems with a wide range of cost functions. The boundary method reduces the effective dimension of the problem, thus improving complexity. For cost functions equal to a p -norm with $p \in (1, \infty)$, we provide mathematical justification, convergence analysis, and algorithmic development. Our testing supports the boundary method with these p -norms, as well as other, more general cost functions.

Keywords: Optimal transport, Monge-Kantorovich, semi-discrete, Wasserstein distance, boundary method
2000 MSC: 65K10, 35J96, 49M25

1. Introduction

In this work, we consider a new solution method for optimal transport problems. Numerical optimal transport has applications in a wide range of fields, but the scaling properties and ground cost restrictions of current numerical methods make it difficult to find solutions for many applications.

The boundary method we propose focuses on a broad class of optimal transportation problems: *semi-discrete optimal transport*. Many other techniques assume semi-discrete transport, either implicitly or explicitly, as semi-discrete formulations can be used to approximate solutions to fully continuous problems, and the semi-discrete optimal transport problem is of practical relevance itself.

Key challenges in numerical optimal transport are: (a) the design of numerical methods capable of handling general ground costs, (b) efficient computation of the Wasserstein metric,

[☆]This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1650044. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Email addresses: dieci@math.gatech.edu (Luca Dieci), joseph.d.walsh@navy.mil (J.D. Walsh III)

Preprint submitted to Elsevier

Thursday 15th March, 2018

and (c) solutions of three (or higher) dimensional problems. The boundary method addresses these concerns by solving problems where the ground cost is a p -norm, $p \in (1, \infty)$, and by doing so in a way that reduces the effective dimension of the transport problem.

1.1. Description of optimal transport: the Monge-Kantorovich problem

The theory of optimal transport dates back to the work by Monge in 1781, [1]. In the 1940s, Kantorovich's papers [2, 3] relaxed Monge's requirement that no mass be split, creating what we now know as the Monge-Kantorovich problem.

Definition 1.1 (Monge-Kantorovich problem). *Let $X, Y \subseteq \mathbb{R}^d$, let μ and ν be probability densities defined on X and Y , and let $c(\mathbf{x}, \mathbf{y}) : X \times Y \rightarrow \mathbb{R}$ be a measurable ground cost function. Define the set of transport plans*

$$\Pi(\mu, \nu) := \left\{ \pi \in \mathcal{P}(X \times Y) \mid \begin{array}{l} \pi[A \times Y] = \mu[A], \pi[X \times B] = \nu[B], \\ \forall \text{ meas. } A \subseteq X, B \subseteq Y \end{array} \right\}, \quad (1.1)$$

where $\mathcal{P}(X \times Y)$ is the set of probability measures on the product space, and define the primal cost function $P : \Pi(\mu, \nu) \rightarrow \mathbb{R}$ as

$$P(\pi) := \int_{X \times Y} c(\mathbf{x}, \mathbf{y}) d\pi(\mathbf{x}, \mathbf{y}). \quad (1.2)$$

The Monge-Kantorovich problem is to find the optimal primal cost

$$P^* := \inf_{\pi \in \Pi(\mu, \nu)} P(\pi), \quad (1.3)$$

and an associated optimal transport plan

$$\pi^* := \arg \inf_{\pi \in \Pi(\mu, \nu)} P(\pi). \quad (1.4)$$

Kantorovich also identified the problem's dual formulation.

Definition 1.2 (Dual formulation). *Define the set of functions*

$$\Phi_c(\mu, \nu) := \left\{ (\varphi, \psi) \in L^1(d\mu) \times L^1(d\nu) \mid \begin{array}{l} \varphi(\mathbf{x}) + \psi(\mathbf{y}) \leq c(\mathbf{x}, \mathbf{y}), \\ d\mu \text{ a.e. } \mathbf{x} \in X, d\nu \text{ a.e. } \mathbf{y} \in Y \end{array} \right\}. \quad (1.5)$$

Let the dual cost function, $D : \Phi_c(\mu, \nu) \rightarrow \mathbb{R}$, be defined as

$$D(\varphi, \psi) := \int_X \varphi d\mu + \int_Y \psi d\nu. \quad (1.6)$$

Then, the optimal dual cost is

$$D^* := \sup_{(\varphi, \psi) \in \Phi_c(\mu, \nu)} D(\varphi, \psi), \quad (1.7)$$

and an optimal dual pair is given by

$$(\varphi^*, \psi^*) := \arg \sup_{(\varphi, \psi) \in \Phi_c(\mu, \nu)} D(\varphi, \psi). \quad (1.8)$$

When the ground cost is a distance function (often but not necessarily Euclidean), Monge-Kantorovich solutions are related to the *Wasserstein metric*, a distance between probability distributions:

$$W_1(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \int_{X \times Y} c(\mathbf{x}, \mathbf{y}) d\pi(\mathbf{x}, \mathbf{y}). \quad (1.9)$$

We have $W_1(\mu, \nu) = P^* = D^*$, and hence, we may refer to any of these as the *Wasserstein distance*, the *optimal transport cost*, or simply the *optimal cost*.¹

Remark 1. $W_1(\mu, \nu)$ is often written as W_1 , with μ and ν implied. Furthermore, as Equation (1.9) makes clear, $W_1(\mu, \nu)$ also depends on the ground cost function $c(\mathbf{x}, \mathbf{y})$. In the literature, the Wasserstein distance formula often assumes the ground cost to be a specific pre-determined function, usually the Euclidean distance $\|\mathbf{x} - \mathbf{y}\|_2$.

Definition 1.3 (Monge problem). *In certain cases, there exists at least one solution to the semi-discrete Monge-Kantorovich problem that does not split transported masses. In other words, there exists some π^* such that*

$$\pi^*(\mathbf{x}, \mathbf{y}) = \pi_{T^*}^*(\mathbf{x}, \mathbf{y}) := \mu(\mathbf{x}) \delta[\mathbf{y} = T^*(\mathbf{x})], \quad (1.10)$$

where $T^* : X \rightarrow Y$ is a measurable map called the optimal transport map. When such a π^* exists, we say the solution also solves the Monge problem.

If the Monge-Kantorovich problem has a solution which solves the Monge problem, we can assume without loss of generality that every $\pi \in \Pi(\mu, \nu)$ satisfies

$$\pi(\mathbf{x}, \mathbf{y}) = \pi_T(\mathbf{x}, \mathbf{y}) := \mu(\mathbf{x}) \delta[\mathbf{y} = T(\mathbf{x})], \quad (1.11)$$

for some measurable transport map $T : X \rightarrow Y$, and that the primal cost can be written

$$P(\pi) := \int_X c(\mathbf{x}, T(\mathbf{x})) d\mu(\mathbf{x}). \quad (1.12)$$

1.2. Semi-discrete problem

The semi-discrete optimal transport problem we consider is the Monge-Kantorovich problem of Definition 1.1, with restrictions on μ and ν , and c .

- (1) Assume that μ satisfies the following:
 - (a) μ is absolutely continuous with respect to the Lebesgue measure.
 - (b) μ does not give positive measure to sets of less than d dimensions, and μ is nonatomic.
 - (c) The support of μ is given by the convex compact region $A \subseteq X$. (Since $A \subset \mathbb{R}^d$, it must also be the case that A is simply connected.)
- (2) Assume ν has exactly $n \geq 2$ non-zero values, located at $\{\mathbf{y}_i\}_{i=1}^n \subseteq Y$.
- (3) Assume c is a p -norm with $p \in (1, \infty)$.

As we will show, each of these conditions is required for one or more of the theorems given in Section 3. Condition (1)(a) ensures that the value of μ is bounded, which is required to show Wasserstein distance convergence in Theorem 3.24. Conditions (1)(b), (1)(c), (2), and (3) are all used to satisfy the conditions of Corollary 4² of [5], which we apply to show the μ -a.e. uniqueness of the solution in Theorem 3.7.

¹See also [4, p. 207], a definition of the Wasserstein metric W_p with $p \in [0, \infty)$.

²See Theorem 3.6, below, for a full statement of this result.

1.2.1. Semi-discrete transport and the Monge problem

Because c is continuous and μ is nonatomic, at least one solution to the semi-discrete Monge-Kantorovich problem also satisfies the Monge problem, described in Definition 1.3; see Theorem B in [6]. Thus, by applying Equation (1.11), we can assume without loss of generality that any transport plan π partitions A into n sets A_i , where A_i is the set of points in A that are transported by the map T to \mathbf{y}_i . Using this partitioning scheme in combination with Equation (1.12) allows us to rewrite the primal cost function for the semi-discrete problem as

$$P(\pi) := \sum_{i=1}^n \int_{A_i} c(\mathbf{x}, \mathbf{y}_i) d\mu(\mathbf{x}). \quad (1.13)$$

1.3. Shift characterization for semi-discrete optimal transport

Using this idea of sets A_i , we are ready to describe the shift characterization of the semi-discrete optimal transport problem. The definition of the characterization, which follows, is based on one given by Rüschendorf and Uckelmann in [7, 8].

Definition 1.4 (Shift characterization). Let $\{a_i\}_{i=1}^n$ be a set of n finite values, referred to as shifts. Define

$$F(\mathbf{x}) := \max_{1 \leq i \leq n} \{a_i - c(\mathbf{x}, \mathbf{y}_i)\}. \quad (1.14)$$

For $i \in \mathbb{N}_n$, where $\mathbb{N}_n = \{1, \dots, n\}$, let

$$A_i := \{\mathbf{x} \in A \mid F(\mathbf{x}) = a_i - c(\mathbf{x}, \mathbf{y}_i)\}. \quad (1.15)$$

Note that $\cup_{i=1}^n A_i = A$. The problem of determining an optimal transport plan π^* is equivalent to determining shifts $\{a_i\}_{i=1}^n$ such that for all $i \in \mathbb{N}_n$, the total mass transported from A_i to \mathbf{y}_i equals $\nu(\mathbf{y}_i)$.

1.4. Numerical approaches to the MK problem

Applications of optimal transport are found in many areas of research, including medicine, economics, image processing, machine learning, physics, and many others; e.g., see [9, 10, 11, 12, 13]. For that reason, many people have focused their research on numerical methods for the Monge-Kantorovich problem.

The solution to a semi-discrete problem can be approximated by treating the problem as fully discrete, and the solution to a fully continuous problem can be approximated by treating it as either semi- or fully discrete. By “treating,” we refer primarily to assumptions about continuity: in practice, nearly every approach fully discretizes the problem, and the complexity of such approaches is relative to the measure of the discretization.

The semi-discrete problem has received significant attention in its role as a discretization of the continuous problem (where continuity assumptions are employed over X but not Y). Substantial effort has been taken to quantify the extent to which solutions to such semi-discrete problems approximate the solution to the original continuous problem; for example, see [14]. However, the semi-discrete problem has interesting applications in its own right. Recent developments include works in economics [15, 16, 17], image processing [18], and optics [19, 20]. In addition, there are obvious similarities between the semi-discrete problem and generalized Voronoi diagrams, which drive ongoing research in physics and other fields.

When the ground cost for the semi-discrete problem is the squared 2-norm, $\|\cdot\|_2^2$, significant numerical progress has been achieved. In 1988, Oliker and Prussner introduced what came to be called the Oliker-Prussner algorithm for nonlinear Monge-Ampère-type equations in \mathbb{R}^2 ; see [21]. Oliker and Prussner were significantly ahead of their time. A 1992 paper by Aurenhammer et al., [22], while applying a different algorithm (Newton’s method), explicitly connected the Oliker and Prussner’s approach to semi-discrete transport and its resulting “Voronoi-type diagrams.” In 1998 Aurenhammer et al. published [23], a revision that clarified important details, and incorporated an argument from [5] to guarantee that the sets A_i partition A μ -a.e.³

When sets A_i and A_j share a boundary, for some $i \neq j$, there is a monotone relationship between the volume of A_i and the difference of shifts, $a_i - a_j$. The Oliker-Prussner approach and the boundary method both exploit this relationship, though in very different ways. Whether applying the Oliker-Prussner algorithm or some variation such as Newton’s method, the Oliker-Prussner approach begins with approximated sets \tilde{A}_i , and directly perturbs the approximated shift difference $\tilde{a}_i - \tilde{a}_j$ in order to bring $\mu(\tilde{A}_i)$ closer to $\nu(y_i)$. This approach is extended over all the shift differences,⁴ making it, in essence, a method for solving the Monge-Kantorovich dual problem with $c = \|\cdot\|_2^2$. Because the squared 2-norm is strictly convex, and it ensures that the boundary for each adjacent A_i and A_j is a hyperplane, algorithms based on the Oliker-Prussner approach are generally able to quantify convergence behavior and guarantee termination after a finite number of refinement steps.

Numerous efforts have been made to extend the approach proposed by Oliker and Prussner. An application-focused paper by Caffarelli et al. extends the Oliker-Prussner algorithm to \mathbb{R}^3 , assuming special geometries [24]. Lévy presents a parallelized Newton’s method for surfaces embedded in three dimensions, one which scales well when Y consists of large numbers of Dirac masses [25]. Other works, such as [26], attempt to integrate the Oliker-Prussner approach with the Wide Stencil methods developed for continuous Monge-Ampère problems; see, e.g., [27, 28]. All of these assume $c = \|\cdot\|_2^2$.

A few authors have attempted to develop approaches for ground costs other than the squared 2-norm. Most of these do not employ Oliker-Prussner. In [8], Rüschemdorf and Uckelmann report on numerical experiments with ground costs given by the Euclidean distance taken to the powers 2, 3, 4, and 10. They assume that μ is the uniform distribution, and test various weights and placements for the set $\{y_i\}_{i=1}^n$. When an exact solution cannot be directly determined, they fully discretize the problem and use a linear programming solver.

In [29], Schmitzer works with cost functions $c = \|\cdot\|_2^p$ for $p \in (1, \infty)$, and applies a form of adaptive scaling done by “shielding” regions: his method attempts to determine points of influence in order to solve primarily local problems. He restricts his examples to \mathbb{R}^2 .

Solving the semi-discrete problem for the 2-norm is discussed in [30].⁵ Starting with an alternative form of Equation (1.16), taken from [31], Barrett and Prigozhin develop a mixed formulation of the Monge-Kantorovich problem, which they solve using a standard finite element discretization.

Kitagawa’s 2014 paper, [32], offers a potentially broad generalization of the Oliker-Prussner algorithm, which works for ground costs other than $\|\cdot\|_2^2$, provided those ground costs satisfy strict

³When the ground cost is the squared 2-norm, a shift-characterized Monge transport map that partitions A μ -a.e. is called a *power* or *Laguerre diagram*, and the sets A_i are referred to as *Laguerre cells*.

⁴They refer to a set of shift differences $\{a_i - a_j \mid i, j \in \mathbb{N}_n, i < j\}$ as a *weight vector*.

⁵In [30], the partition of A is called an “optimal coupling.”

conditions, including Strong Ma-Trudinger-Wang; see also [33]. His proposals, while densely theoretical, do not include numerics or an explicit iterative scheme.

As [34] states, the special case $c = \|\cdot\|_2^2$ has two methods specifically designed for solving semi-discrete problems directly: the Oliker-Prussner algorithm and the damped Newton methods proposed in papers like [23]. Both rely on some variant of what we call the Oliker-Prussner approach, described above. However, approaches developed for fully discrete or continuous transport can also be applied to the semi-discrete problems, though with varying degrees of effectiveness. Rüschemdorf and Uckelmann apply a discrete linear program solver in [8], and the solver Barrett and Prigozhin use in [30] was developed for continuous transport.

Discrete methods assume a fully discrete (X, μ) and (Y, ν) , and solve the resulting minimization problem using network flow minimization techniques. As described in [35], there are over 20 established methods for solving such problems, and at least seven software packages capable of handling one or more of these methods.

Most approaches to the fully continuous Monge-Kantorovich problem assume specific ground costs and solve using techniques developed for elliptic partial differential equations, particularly those of the Monge-Ampère-type:

$$-\nabla \cdot (a\nabla u) = f, \text{ where } |\nabla u| \leq 1, a \geq 0, \text{ and } |\nabla u| < 1 \implies a = 0. \quad (1.16)$$

If the ground cost function is strictly convex, or otherwise satisfies the Ma-Trudinger-Wang regularity conditions described in [33], such problems are well-posed. To date, the requirements of well-posedness have largely restricted the application of such continuous methods to well-behaved cost functions such as $\|\cdot\|_2^2$ or a regularized Euclidean distance. Continuous methods currently in use apply finite difference, gradient descent, or Bregman's method, all attempting to map X to a fully discretized Y [36, 37, 38].

As we will show, the boundary method offers a new approach to solving semi-discrete transport, distinct from all of those described above. By and large, the solution methods described above only work for a specific fixed cost, usually $c = \|\cdot\|_2^2$. The boundary method quickly solves problems with more general ground costs. When the ground cost is a p -norm, with $p \in (1, \infty)$, the boundary method provides a global rate of convergence that is proportional to the volume of A .

2. Boundary Method

At a high level, the idea behind the boundary method is simple: track only the boundaries between regions, without resolving the regions' interiors. To do this in practice and obtain an efficient technique, we must account for the interplay between discretization, a mechanism for discarding interior regions, and a fast solver.

At its heart, the boundary method can be viewed as an adaptive refinement technique, one which focuses on the shared region boundaries. The method discards interior regions, but a well-chosen initial discretization prevents any corresponding loss of accuracy. The necessary conditions for a well-chosen initial discretization are presented in Theorem 3.20 and discussed in Remark 3.

2.1. Boundary identity and system of equations

For all $i, j \in \mathbb{N}_n$ such that $i \neq j$, let

$$A_{ij} := A_i \cap A_j. \quad (2.1)$$

The *boundary set* is defined as

$$B := \bigcup_{1 \leq i < n} \bigcup_{i < j \leq n} A_{ij}, \quad (2.2)$$

and for each $i \in \mathbb{N}_n$, let the *strict interior* of A_i be defined as

$$\mathring{A}_i := A_i \setminus B. \quad (2.3)$$

For all $i, j \in \mathbb{N}_n$ such that $i \neq j$, define $g_{ij} : X \rightarrow \mathbb{R}$ as

$$g_{ij}(\mathbf{x}) := c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{x}, \mathbf{y}_j). \quad (2.4)$$

By Corollary 3.11 below, $B \neq \emptyset$ and for each $\mathbf{x} \in B$ there exist $i, j \in \mathbb{N}_n$, $i \neq j$, such that $\mathbf{x} \in A_{ij}$. Because $\mathbf{x} \in A_i$, we have $F(\mathbf{x}) = a_i - c(\mathbf{x}, \mathbf{y}_i)$, and because $\mathbf{x} \in A_j$, we have $F(\mathbf{x}) = a_j - c(\mathbf{x}, \mathbf{y}_j)$. Combining and rearranging these, we get

$$g_{ij}(\mathbf{x}) = a_i - a_j, \quad \forall \mathbf{x} \in A_{ij}. \quad (2.5)$$

Thus, Equation (2.5) implies that A_{ij} is a subset of a level set of g_{ij} ; the value $a_i - a_j$ is constant, regardless of which $\mathbf{x} \in A_{ij}$ is chosen. Using this information, for each $i, j \in \mathbb{N}_n$, $i \neq j$, such that $A_{ij} \neq \emptyset$, we can define the constant *shift difference*

$$a_{ij} := g_{ij}(\mathbf{x}_{ij}) \quad \forall \mathbf{x}_{ij} \in A_{ij}. \quad (2.6)$$

Given a sufficiently large set of linearly independent equations of the form given in Equation (2.6), one could determine most or all of the shifts $\{a_i\}_{i=1}^n$. As we show in Theorem 3.13, it is possible to obtain exactly $(n - 1)$ linearly independent equations of the desired form, but a set of n such independent equations does not exist.

Since we know that the set of shifts allows exactly one degree of freedom, the boundary method's approach is to obtain $(n - 1)$ well-chosen a_{ij} values, fix one a_i , and use linearly independent equations of the form given in Equation (2.5) to solve for the remaining $(n - 1)$ shifts. The crucial observation is that for the a_i 's, there is no need to retain information about interior of the regions.

The Wasserstein distance can also be computed without saving region interiors. Once we have determined that $R \subset A_i$ for some region R , the (partial) Wasserstein distance corresponding to R is equal to

$$P_R := v(\mathbf{y}_i) \int_R c(\mathbf{x}, \mathbf{y}_i) d\mu(\mathbf{x}), \quad (2.7)$$

and the total Wasserstein distance P^* is equal to the sum of all such partial distances P_R , computed over every A_i .

Recognizing these facts, inherent in the shift characterization, inspired both the boundary method's name and its guiding principles, summarized below:

Do *not* solve for the entire transport plan;
rather, identify region boundaries.

2.2. The boundary method

As described below, the boundary method generates a grid A^r over the unevaluated region of A , and uses it to determine the subgrid B^r containing the boundary set B . This subgrid is determined by finding an optimal transport solution from the grid A^r to the point set $\{\mathbf{y}_i\}_{i=1}^n$.

Although not strictly necessary, we will restrict ourselves to $A = [0, 1]^d$ and apply a Cartesian grid over that region. At the r -th refinement level of the algorithm, the grid will thus consist of a collection of boxes with width w_r in each dimension of our discretization. By a slight abuse of notation, we use \mathbf{x}^r to refer to such a box, centered at the point \mathbf{x} . Thus, $\mu(\mathbf{x}^r)$ refers to the μ -measure of the box of width w_r centered at \mathbf{x} .

Neighboring boxes are those with center points that differ by no more than one unit in any discretization index. The set of *neighbors* of \mathbf{x} is denoted $N(\mathbf{x})$ (defined in Equation (3.9), below). Because regions of μ -measure zero need not be transported to any particular \mathbf{y}_i , boxes of positive weight that are adjacent to such regions are always retained. We refer to such a box as an *edge box*. Thus, the set of edge boxes is

$$\text{edg}(A^r) := \{\mathbf{x} \in A^r \mid \mu(\mathbf{x}) > 0 \text{ and } \exists \mathbf{x}_n \in N(\mathbf{x}) \text{ such that } \mu(\mathbf{x}_n) = 0\}. \quad (2.8)$$

Because A contains the support of μ , every box of positive mass that is adjacent to the boundary of A is an edge box.

A box whose neighbors and itself all have positive measure is referred to as an *internal box*. The set of internal boxes is

$$\text{int}(A^r) := \{\mathbf{x} \in A^r \mid \mu(\mathbf{x}) > 0 \text{ and } \mu(\mathbf{x}_n) > 0 \text{ for all } \mathbf{x}_n \in N(\mathbf{x})\}. \quad (2.9)$$

Boxes of μ -measure zero are not part of $\text{edg}(A^r)$ or $\text{int}(A^r)$ and they are discarded when the optimal transport problem is solved. We need not be concerned about losing a region A_i due to this discard process, since this would imply $\mu(A_i) = 0$ (and hence $\nu(\mathbf{y}_i) = 0$, which contradicts the conditions in Section 1.2).

Region interiors are identified by comparing the destination of each $\mathbf{x} \in \text{int}(A^r)$ to the destinations of its neighbors. Edge boxes are never considered part of a region interior, so they are passed directly to B^r .

In order to remove identified region interiors, we also maintain a running total of the untransported mass, given by *partial measure* $\tilde{\nu}$. To preserve the balance of the transport problem, each time a region \mathbf{x}^r is transported from A to \mathbf{y}_i , the remaining amount that can be transported to \mathbf{y}_i , $\tilde{\nu}(\mathbf{y}_i)$, must be reduced by $\mu(\mathbf{x}^r)$.

We can approximate the Wasserstein distance P^* by generating a running total over region interiors: \tilde{P} . This \tilde{P} is an increasing function of r , and for all r , $P^* \geq \tilde{P}$. The Wasserstein distance over any remaining boundary region is evaluated at completion.

Remark 2. *Further approximations may be required for a truly general algorithm. Depending on μ , it may be necessary to approximate the mass of each box, $\mu(\mathbf{x}^r)$. Depending on μ and c , the Wasserstein distance over each box, given by $\int_{\mathbf{x}^r} c(\mathbf{z}, \mathbf{y}_i) d\mu(\mathbf{z})$, may also require approximation. However, in this work we assume that the integrals can be computed exactly. In practice, this is not a significant limitation. Most numerical applications focus on the exactly-computable cases where μ is uniform and c is the Euclidean or squared-Euclidean distance. Furthermore, as we show in Section 4.1, the set of exactly-computable options is quite large.*

Boundary method algorithm

(0) Set $\tilde{P} = 0$, $\tilde{\nu} = \nu$, and $r = 1$. Create $A^r = A^1$ from A .

(1) Solve the discretized transport solution.

(2) For each $\mathbf{x} \in \text{int}(A^r)$:

Are the neighbors of \mathbf{x} all transported to the same \mathbf{y}_i ?

• If so, then \mathbf{x}^r is in the interior of A_i :

- [optional] Add $\int_{\mathbf{x}^r} c(\mathbf{z}, \mathbf{y}_i) d\mu(\mathbf{z})$ to \tilde{P} .
- Reduce the value of $\tilde{\nu}(\mathbf{y}_i)$ by $\mu(\mathbf{x}^r)$.
- Remove \mathbf{x} from $\text{int}(A^r)$.

The sets $\text{edge}(A^r)$ and the reduced set $\text{int}(A^r)$ combine to form B^r .

(3) Is the desired refinement reached?

• If not:

- Refine B^r to create A^{r+1} , increment r , and go to Step (1).

Optionally, once the desired refinement level is reached:

(4) Use B^r to identify $(n - 1)$ appropriate shift differences $\{a_{ij}\}$ and solve for the shifts $\{a_i\}_{i=1}^n$.

(5) Use \tilde{P} and B^r to approximate $W_1(\mu, \nu)$.

To illustrate the iterative portion of the boundary method, Steps (1) and (2), we present the following example.

Example 2.1. Let $X = Y = [0, 1]^2$. Assume μ is the uniform probability density, so for all measurable sets $S \subseteq A$, $\mu(S) = |S|$, and that ν has uniform discrete probability density, so $\nu(y_i) = 1/n$ for $1 \leq i \leq n$. Take $n = 5$, with the five points where ν has nonzero density distributed as shown in Figure 1.

Let c be the squared Euclidean norm, $\|\mathbf{y} - \mathbf{x}\|_2^2$. Suppose a discretization with width 2^{-5} is sufficient to provide the desired accuracy and that we apply the boundary method with initial width 2^{-4} .

Figure 1 shows the state of the boundary method during the first iteration. In Figure 1(a), we have just completed Step (1): the discrete transport map has been computed, but we have not identified interior points or added anything to the partial transport cost \tilde{P} . Figure 1(b) shows the state of the algorithm after Step (2): the interior regions have been identified (shown in gray), and the partial transport cost has been computed for those regions, giving us $\tilde{P} = 0.01387$.

Figure 2 shows the state of the boundary method algorithm during the second iteration. In Figure 2(a), Step (1) has just been completed. As can be seen by comparing Figure 1(b) to Figure 2(a), the boundary and interior regions are the same ones that we had at the end of the first iteration, but refining the boundary set to width $w_2 = 2^{-5}$ allows us to compute a more refined transport map. Because Step (1) does not add to the identified interior regions, the partial Wasserstein distance \tilde{P} is also unchanged from Figure 1(b).

After Step (2) of the second iteration, shown in Figure 2(b), more of the interiors have been identified. The partial transport cost shows a corresponding increase: we now have $\tilde{P} = 0.02898$. Because we have achieved our desired refinement, a width of 2^{-5} , we move on to Step (4), ending the iterative process.

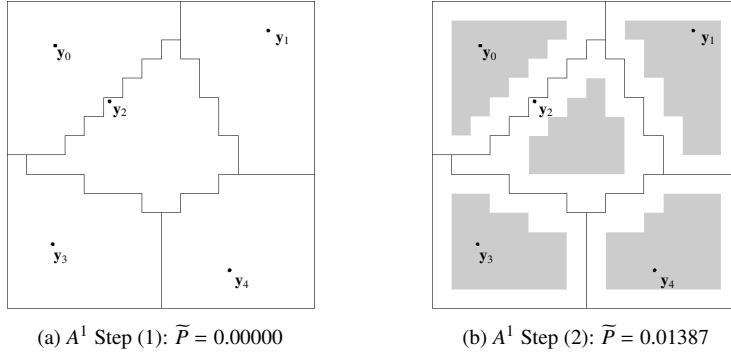


Figure 1: Iteration A^1 of Example 2.1 illustrated: $w_1 = 2^{-4}$




Figure 2: Iteration A^2 of Example 2.1 illustrated: $w_2 = 2^{-5}$

2.2.1. Step (1): solving the discrete optimal transport problem

The proofs in Section 3 assume the discrete solver is exact, but in practice we achieve good results using solvers whose error satisfies reasonable bounds. Thus, the ideal discrete algorithm should be *fast*, have controlled error, and possess reasonable scaling properties. To satisfy these requirements, and to bypass the shortcomings of standard discrete approaches, we have turned to the distributed relaxation methods known as auction algorithms; see [39] and [40]. (As it turns out, there are natural connections between auction algorithms and the Oliker-Prussner algorithm for semi-discrete transport; see [41] for details).

We chose to apply a new auction algorithm, the *general auction*, which we developed and presented in [42]. The general auction is so named because it is based directly on the (more general) real-valued transport problem, rather than the integer-valued assignment problem which forms the foundation of other auction algorithms. As described in [42], it offers significant performance advantages over other auction algorithms. Public domain C++ software implementing the general auction can be found on the Internet at [43].

2.2.2. Step (4): computing the shifts

Once we have reached a desired level of refinement for the boundary, we can use the set B' to identify $(n - 1)$ shift differences a_{ij} . Finding the shift differences is not necessary once we have the boundary (which is why Step (4) is optional), but the shift differences allow one to

reconstruct the entire transport map.

By completing Step (4), one can reduce the transport map in \mathbb{R}^d to a set of n real numbers a_i , greatly reducing storage requirements. Also, building the reconstructed transport map, and comparing the value of each $\mu(A_i)$ to its corresponding $\nu(\mathbf{y}_i)$, effectively evaluates the actual (vs. worst case) error associated with the boundary method's solution.

It is also worth considering that the exact shifts $\{a_i\}_{i=1}^n$ correspond to a transport map giving the exact optimal solution of our semi-discrete problem. The approximated shifts $\{\tilde{a}_i\}_{i=1}^n$, unless generating the same shift differences, correspond to a transport map giving the exact optimal solution to a *different* semi-discrete problem, one whose measure ν at each $\mathbf{y}_i, i \in \mathbb{N}_n$, corresponds to the value of $\mu(\tilde{A}_i)$. Hence, $|\mu(\tilde{A}_i) - \nu(\mathbf{y}_i)|$ is the error in measure when approximating A_i by \tilde{A}_i .

2.2.3. Step (5): approximating the Wasserstein distance

Because some applications focus on determining the transport map, rather than the Wasserstein distance, Step (5) is optional. One could also skip the computation of \tilde{P} in Step (2), since the Wasserstein distance can be computed in full using only the transport map defined by the boundary set. However, we find it convenient to compute as much of the distance as possible within the boundary method algorithm, establishing \tilde{P} one box at a time during Step (2). By the time we reach Step (5), the partial Wasserstein distance \tilde{P} includes the exact cost of all the identified interior regions, and all that remains is to determine the cost of the regions associated with B^r .

3. Mathematical support

In this section, we provide mathematical support for the boundary method, assuming that all computations are solved exactly: both the discrete optimal transport problems handled by the general auction and the determinations of mass and Wasserstein distance for individual boxes (see Remark 2). We present three types of results: on the shift characterization, on our system of equations, and, finally, on the boundary method itself.

3.1. Semi-discrete optimal transport and the shift characterization

Here we examine the features of the shift characterization, defined in Section 1.3, and consider what they can tell us about the semi-discrete optimal transport problem itself.

First, in Lemmas 3.1 and 3.2, we develop theoretical support for the boundary method.

Lemma 3.1. *Fix $i \in \mathbb{N}_n$. If $\mathbf{x} \in A_i$ and $j \in \mathbb{N}_n, j \neq i$, then the following hold:*

$$g_{ij}(\mathbf{x}) \leq a_i - a_j, \quad (3.1)$$

$$g_{ij}(\mathbf{x}) = a_i - a_j \iff \mathbf{x} \in A_{ij}, \text{ and} \quad (3.2)$$

$$g_{ij}(\mathbf{x}) < a_i - a_j \iff \mathbf{x} \in A_i \setminus A_j. \quad (3.3)$$

PROOF. Let us show Equation (3.1). By the definitions of A_i and F ,

$$a_i - c(\mathbf{x}, \mathbf{y}_i) = F(\mathbf{x}) \geq a_j - c(\mathbf{x}, \mathbf{y}_j).$$

Rearranging terms gives

$$c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{x}, \mathbf{y}_j) \leq a_i - a_j.$$

To show Equation (3.2), first note that Section 2.1 already explains how $\mathbf{x} \in A_{ij}$ implies $g_{ij}(\mathbf{x}) = a_i - a_j$. Consider the converse: Assume that $g_{ij}(\mathbf{x}) = a_i - a_j$. Rewriting, we find that $a_j - c(\mathbf{x}, \mathbf{y}_j) = a_i - c(\mathbf{x}, \mathbf{y}_i) = F(\mathbf{x})$. This implies $\mathbf{x} \in A_j$, and since $\mathbf{x} \in A_i$, therefore $\mathbf{x} \in A_{ij}$. Equation (3.3) is a consequence of Equations (3.1) and (3.2).

Lemma 3.2. For all $i, j \in \mathbb{N}_n, i \neq j$,

- (a) If $c(\mathbf{y}_i, \mathbf{y}_j) = a_i - a_j$, then $A_j \subseteq A_{ij}$.
- (b) If $c(\mathbf{y}_i, \mathbf{y}_j) < a_i - a_j$, then $A_j = \emptyset$.

PROOF. For Part (a), because c satisfies the triangle inequality, for all $\mathbf{x} \in A$,

$$\begin{aligned} c(\mathbf{x}, \mathbf{y}_i) &\leq c(\mathbf{x}, \mathbf{y}_j) + c(\mathbf{y}_j, \mathbf{y}_i) \\ c(\mathbf{x}, \mathbf{y}_i) &\leq c(\mathbf{x}, \mathbf{y}_j) + a_i - a_j \\ a_j - c(\mathbf{x}, \mathbf{y}_j) &\leq a_i - c(\mathbf{x}, \mathbf{y}_i). \end{aligned} \tag{3.4}$$

Suppose $\mathbf{x} \in A_j$. Then $a_i - c(\mathbf{x}, \mathbf{y}_i) \geq a_j - c(\mathbf{x}, \mathbf{y}_j) = F(\mathbf{x})$. Because F is defined as the maximum such difference, this implies $a_i - c(\mathbf{x}, \mathbf{y}_i) = F(\mathbf{x})$, and so $\mathbf{x} \in A_i$. Further, since \mathbf{x} is an element of A_i and A_j , $\mathbf{x} \in A_{ij}$. Therefore, $A_j \subseteq A_{ij}$.

To show (b), note that (3.4) now gives $a_j - c(\mathbf{x}, \mathbf{y}_j) < a_i - c(\mathbf{x}, \mathbf{y}_i)$. Hence, for all $\mathbf{x} \in A$, $F(\mathbf{x}) \geq a_i - c(\mathbf{x}, \mathbf{y}_i) > a_j - c(\mathbf{x}, \mathbf{y}_j)$. Therefore, $A_j = \emptyset$.

Lemma 3.3. $F(\mathbf{x})$ is a continuous function of \mathbf{x} .

PROOF. The ground cost c is defined as a continuous function in $X \times Y$. Thus, for all $i \in \mathbb{N}_n$, $a_i - c(x, y_i)$ is a continuous function of \mathbf{x} . Since F is the maximum of a finite set of continuous functions, F is itself a continuous function of \mathbf{x} .

Definition 3.4 (F μ -partitions A). Let F be as defined in Equation (1.14), and the sets A_i as defined in Equation (1.15) for $i \in \mathbb{N}_n$. Then one says F μ -partitions the set A , or F is called a μ -partition, if

1. $\mu(A) < \infty$,
2. for all $i, j \in \mathbb{N}_n, i \neq j, \mu(A_{ij}) = 0$,
3. $\sum_{i=1}^n \mu(A_i) = \mu(A)$, and
4. for all $i \in \mathbb{N}_n, \mu(A_i) = v(\mathbf{y}_i) > 0$.

Lemma 3.5. Suppose one has a semi-discrete transport problem, as described in Section 1.2. Let F be as defined in Equation (1.14), and the sets A_i as defined in Equation (1.15) for $i \in \mathbb{N}_n$. Then F μ -partitions A if and only if $\mu(B) = 0$.

PROOF. If F μ -partitions A , by Definition 3.4, $\mu(B) = 0$. For the converse, assume F and the sets A_i are defined as given. Because μ is a probability density function, $\mu(A) = 1 < \infty$. Because μ is a non-negative measure, $\mu(B) = 0$ implies that, for all $i, j \in \mathbb{N}_n, i \neq j, \mu(A_{ij}) = 0$.

For any μ -measurable set $S \subseteq X, S = S_1 \cup S_2$,

$$\mu(S_1) + \mu(S_2) = \mu(S) + \mu(S_1 \cap S_2),$$

and since $\mu(X) < \infty$,

$$\mu(S) = \mu(S_1) + \mu(S_2) - \mu(S_1 \cap S_2).$$

Proceeding inductively, it follows that

$$S = \bigcup_{i=1}^n S_i, \text{ all } \mu\text{-measurable} \quad \implies \quad \mu(S) = \sum_{i=1}^n \mu(S_i) - \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n \mu(S_i \cap S_j).$$

Thus,

$$1 = \mu(A) = \sum_{i=1}^n \mu(A_i) - \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n \mu(A_{ij}) = \sum_{i=1}^n \mu(A_i).$$

For all $i, j \in \mathbb{N}_n, i \neq j, \mu(A_i \cap A_j) = 0$, and therefore $\mu(A_i) = \nu(\mathbf{y}_i)$.

Given our definition of the semi-discrete problem in Section 1.2, Corollary 4 of [5] provides a sufficient condition for the existence of a Monge solution that is unique μ -a.e. For convenience, we restate their conclusion here, as the following:

Theorem 3.6. *Suppose that the support of ν is finite, c is continuous, μ is tight, and*

$$\mu(\{\mathbf{x} \in A \mid g_{ij}(\mathbf{x}) = k\}) = 0 \quad \forall i, j \in \mathbb{N}_n, i \neq j, \quad \forall k \in \mathbb{R}. \quad (3.5)$$

Then there exists an optimal transport map, $T : X \rightarrow Y$, that solves the Monge problem, and T is μ -a.e. unique.

This condition leads directly to the following theorem.

Theorem 3.7. *A semi-discrete transport problem, as described in Section 1.2, has an associated transport map T , a function F , as described in Equation (1.14), and sets $\{A_i\}_{i=1}^n$, as described in Equation (1.15), such that for all $\mathbf{x} \in A$,*

$$\mathbf{x} \in \mathring{A}_i \text{ for some } i \in \mathbb{N}_n \quad \implies \quad T(\mathbf{x}) = \mathbf{y}_i. \quad (3.6)$$

In other words, F μ -partitions A and T agrees with F on $A \setminus B$. Furthermore, T is unique μ -a.e.

PROOF. Consider the requirements given in Section 1.2. Condition (2) ensures that ν is finite, and Condition (3) implies c is continuous. We know that $A \subseteq \mathbb{R}^d$, so A is a Polish space, and Condition (1)(c) assures us that A is compact. Therefore, as a consequence of Prokhorov's theorem, the probability measure μ must be tight. Because Condition (3) requires that the ground cost is equal to a p -norm with $p \in (1, \infty)$,

$$|\{\mathbf{x} \in A \mid g_{ij}(\mathbf{x}) = k\}| = 0 \quad \forall i, j \in \mathbb{N}_n, i \neq j, \quad \forall k \in \mathbb{R}.$$

A set with d -dimensional Lebesgue measure zero must have dimension lower than d , so Condition (1)(b) implies that

$$\mu(\{\mathbf{x} \in A \mid g_{ij}(\mathbf{x}) = k\}) = 0 \quad \forall i, j \in \mathbb{N}_n, i \neq j, \quad \forall k \in \mathbb{R},$$

as required by Equation (3.5) (see [44] for another argument). Therefore, the conditions of Theorem 3.6 are satisfied.

Let the function F and sets $\{A_i\}_{i=1}^n$ be as described in Definition 1.4. For any $i, j \in \mathbb{N}_n, i \neq j, A_{ij} \subseteq \{\mathbf{x} \in A \mid g_{ij}(\mathbf{x}) = k\}$ for some fixed $k \in \mathbb{R}$. Hence, it follows that $\mu(B) = 0$, and thus, by Lemma 3.5, F μ -partitions A . Therefore, we can construct a transport plan T that satisfies the semi-discrete problem and agrees with F on $A \setminus B$. Furthermore, by Theorem 3.6, T is unique μ -a.e.

3.2. Existence of linearly independent boundary equations

To prove the existence of $(n - 1)$ linearly independent equations of the form shown in Equation (2.6), we will investigate the structure of the boundary set using a connected graph G defined as follows:

Definition 3.8. *Let G be a graph with n vertices v_1, \dots, v_n . The edge (v_i, v_j) is contained in the edge set of G if and only if A_{ij} is non-empty. We refer to G as the adjacency graph of our transport problem.*

Lemma 3.9. *G is a connected graph.*

PROOF. Assume to the contrary that G is not a connected graph. Then we can write G as the union of two disjoint nonempty subgraphs, $G = G_1 \cup G_2$, such that no vertex v_1 in G_1 has a path connecting it to any vertex v_2 in G_2 .

Define

$$\tilde{A}_1 := \bigcup_{v_i \in G_1} A_i \quad \text{and} \quad \tilde{A}_2 := \bigcup_{v_j \in G_2} A_j.$$

Since $G_1 \neq \emptyset$ and $G_2 \neq \emptyset$, $\tilde{A}_1 \neq \emptyset$ and $\tilde{A}_2 \neq \emptyset$. Because G_1 and G_2 are disjoint, and no paths connect them, it follows that $\tilde{A}_1 \cap \tilde{A}_2 = \emptyset$. Since the union of G_1 and G_2 is G , $\tilde{A}_1 \cup \tilde{A}_2 = A$.

Suppose $A_i \subseteq \tilde{A}_1$, $A_j \subseteq \tilde{A}_2$. Then $A_{ij} = \emptyset$. A is a closed and bounded set, and the definition given in Equation (1.15) implies that A_i and A_j must also be closed and bounded. Thus, A_i and A_j are disjoint compact sets in the Hausdorff space \mathbb{R}^d . This implies A_i and A_j are separated by some positive distance ϵ_{ij} . Because this is true for all $A_i \subseteq \tilde{A}_1$ and $A_j \subseteq \tilde{A}_2$, there exists $\epsilon > 0$, the minimum over all such ϵ_{ij} .

Let $\mathbf{x}_1 \in \tilde{A}_1$, $\mathbf{x}_2 \in \tilde{A}_2$, and for all $t \in [0, 1]$, define

$$\mathbf{x}_t = (1 - t)\mathbf{x}_1 + t\mathbf{x}_2.$$

Because $\epsilon > 0$, there exists $(t_0, t_1) \subseteq [0, 1]$, $|t_1 - t_0| \geq \epsilon$, such that $t \in (t_0, t_1)$ implies $\mathbf{x}_t \notin \tilde{A}_1 \cup \tilde{A}_2 = A$. This contradicts the convexity of A . Hence, G is connected.

Corollary 3.10. *If $i \in \mathbb{N}_n$, there exists $j \in \mathbb{N}_n$, such that $j \neq i$ and $A_{ij} \neq \emptyset$.*

PROOF. Assume the contrary for some i . Since $n \geq 2$, G includes at least two vertices, and v_i is disconnected from the rest of G , which contradicts Lemma 3.9.

Corollary 3.11. *The boundary set B is nonempty, and for each $\mathbf{x} \in B$, there exist $i, j \in \mathbb{N}_n$ such that $i \neq j$ and $\mathbf{x} \in A_{ij}$.*

PROOF. This follows from Corollary 3.10 and the definition of B in Equation (2.2).

Lemma 3.12. *Let G be the adjacency graph of the transport problem, and let H be a subgraph of G that includes all n vertices. Define the system of equations*

$$S := \{a_i - a_j = a_{ij} \mid (v_i, v_j) \in \text{the edge set of } H\}. \quad (3.7)$$

The system of equations S is linearly independent with respect to the shifts $\{a_i\}_{i=1}^n$ if and only if H contains no cycles.

PROOF. (\implies) Suppose H contains the cycle $(v_{i_1}, v_{i_2}, \dots, v_{i_k}, v_{i_1})$. Then S contains the linear system

$$M \begin{bmatrix} a_{i_1} \\ \vdots \\ a_{i_{k-1}} \\ a_{i_k} \end{bmatrix} = \begin{bmatrix} a_{i_1 i_2} \\ \vdots \\ a_{i_{k-1} i_k} \\ a_{i_k i_1} \end{bmatrix}, \quad \text{where } M = \begin{bmatrix} 1 & -1 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & 1 & -1 \\ -1 & & & & & 1 \end{bmatrix}.$$

Because $\det(M) = 0$, we know S is linearly dependent.

(\impliedby) Suppose instead that S is linearly dependent. Given the form of the equations in S , we can assume without loss of generality that S contains the equations $a_{ij_{j+1}} = a_{i_j} - a_{i_{j+1}}$, $\forall j \in \mathbb{N}_{k-1}$, and that $a_{i_1 i_k} = a_{i_1} - a_{i_k}$ is also in S . By the definition of S , these equations imply that the edges $(v_1, v_2), (v_2, v_3), \dots, (v_{k-1}, v_k)$, and (v_k, v_1) are contained in H . Together, these edges generate the cycle $(v_{i_1}, v_{i_2}, \dots, v_{i_k}, v_{i_1})$, so H contains at least one cycle.

Theorem 3.13. *There exists at least one system of exactly $(n - 1)$ equations of the form $a_i - a_j = a_{ij}$ that is linearly independent with respect to the set of shifts $\{a_i\}_{i=1}^n$. No system of n independent equations exists.*

PROOF. Because G is a connected graph, we can always create a spanning tree H that is a subgraph of G . Let S be the corresponding set of linear equations, defined as described in (3.7). As a spanning tree, H contains $(n - 1)$ edges and H has no cycles, so by Lemma 3.12, we know S contains exactly $(n - 1)$ linearly independent equations.

Suppose a set S of n linearly independent equations exists, all of the form $a_i - a_j = a_{ij}$. Because there are n unknowns in the set of shifts, there is exactly one solution set $\{a_i\}_{i=1}^n$. Fix $\sigma \neq 0$ and for all $i \in \mathbb{N}_n$, define $\tilde{a}_i = a_i + \sigma$. For each equation in S , $\tilde{a}_i - \tilde{a}_j = a_i - a_j = a_{ij}$. Thus, $\{\tilde{a}_i\}_{i=1}^n$ is also a solution to S . This contradicts the uniqueness of $\{a_i\}_{i=1}^n$, and therefore no such set of n linearly independent equations exists.

3.3. Discretization for the boundary method

In the first two subsections below, we give some results on how the grid-points interact with the underlying space. In sections 3.3.3 and 3.3.4 we present error bounds. In section 3.3.5 we consider issues of volume and containment: here we ensure that one can have $B \subseteq \bar{B}^r$ for all r , and show that $|\bar{B}^r| \rightarrow 0$ as $r \rightarrow \infty$. Finally, Section 3.3.6 puts bounds on the error for the Wasserstein distance approximation.

3.3.1. Discretization definitions

As described in Section 2.2, we discretize the region A using a regular Cartesian grid, and refine the grid over multiple iterations, with the aim of refining only the grid region containing the boundary set.

Let $r \in \mathbb{N}$ be the current discretization level, and $w = w_r$ be the *width* of the discretization at level r . Let A^r be the r -th *point set*, the set of points \mathbf{x} included in the r -th discretization of A . Since we discard boxes of μ -measure zero during the transport step, assume without loss of generality that $\mu(\mathbf{x}^r) > 0$ for all $\mathbf{x} \in A^r$. Let $A_i^r = A_i \cap A^r$ for $i \in \mathbb{N}_n$.

Let \mathcal{V} be the set of adjacency vectors for all discretizations of A . We choose \mathcal{V} to be the linear combinations of the standard unit vectors, e_1, \dots, e_d , with coefficients ± 1 . We specifically exclude the zero vector from the set, so $|\mathcal{V}| = 3^d - 1$. If $d = 2$, \mathcal{V} equals

$$\mathcal{V} := \{(-1, -1), (0, -1), (-1, 0), (-1, 1), (1, -1), (1, 0), (0, 1), (1, 1)\}. \quad (3.8)$$

For all r and all $\mathbf{x} \in A^r$, the points in A^r that are adjacent to \mathbf{x} constitute a subset of the *neighbors* of \mathbf{x} ,

$$N(\mathbf{x}) := \{\mathbf{x} + w_r \mathbf{v} \mid \mathbf{v} \in \mathcal{V}\}. \quad (3.9)$$

Lemma 3.14. *If $\mathbf{x} \in N(\mathbf{x}_0)$, then $\mathbf{x}_0 \in N(\mathbf{x})$.*

PROOF. This follows from Equation (3.9) and our definition of the adjacency vectors: for all $k \in \mathbb{N}_d$, $e_k \in \mathcal{V} \iff -e_k \in \mathcal{V}$.

We now formalize our idea of the r -th interior and boundary point sets used in our discretization. For all $i \in \mathbb{N}_n$, define the r -th iteration *interior point set* associated with A_i as

$$\mathring{A}_i^r := \{\mathbf{x} \in A_i^r \mid \forall \mathbf{v} \in \mathcal{V}, \mathbf{x} + w_r \mathbf{v} \in A_j^r \implies j = i\}. \quad (3.10)$$

Define the r -th *boundary point set* as

$$B^r := A^r \setminus \bigcup_{i=1}^n \mathring{A}_i^r, \quad (3.11)$$

and let $B_i^r := B^r \cap A_i$ for all $i \in \mathbb{N}_n$. The r -th *evaluation region*, the subset of A enclosed by the discretization A^r , is defined as

$$\bar{A}^r := \{\mathbf{x}^r \mid \mathbf{x} \in A^r\}, \quad (3.12)$$

and the r -th *boundary region*, the subset of A enclosed by the boundary point set B^r , is given by

$$\bar{B}^r := \{\mathbf{x}^r \mid \mathbf{x} \in B^r\}. \quad (3.13)$$

3.3.2. Distance bounds

Though the discretization is fully defined, it still needs to be related back to the sets A_{ij} and the boundary set B . To do this, we first bound the distance separating B^r and A_{ij} .

Lemma 3.15. *Suppose $B_i^r \neq \emptyset$. For each $\mathbf{x}_i \in B_i^r$, either $\mathbf{x}_i \in \text{edg}(A^r)$ or there exists a point $\mathbf{x}_j = \mathbf{x}_i + w_r \mathbf{v}$, with $\mathbf{v} \in \mathcal{V}$, such that $\mathbf{x}_j \in B_j^r$ for some $j \neq i$. Thus, if $\mathbf{x}_i \notin \text{edg}(A^r)$, the distance from \mathbf{x}_i to the set A_{ij} , as measured with respect to the ground cost c , is bounded above by $c(\mathbf{x}_i, \mathbf{x}_j)$.*

PROOF. Assume $\mathbf{x}_i \in B_i^r \setminus \text{edg}(A^r)$. By the definition of B^r , there exists $\mathbf{x}_j = \mathbf{x}_i + w_r \mathbf{v} \in A_j^r \cup N(\mathbf{x}_0)$ for some $j \neq i$. By Lemma 3.14, $\mathbf{x}_i \in N(\mathbf{x}_j)$, and since $\mathbf{x}_i \in A_i^r$, we have $\mathbf{x}_j \in B_j^r$. Thus, $\mathbf{x}_i \in A$ and $\mathbf{x}_j \in A$, and because A is convex, this implies

$$\{t\mathbf{x}_i + (1-t)\mathbf{x}_j \mid t \in [0, 1]\} \subseteq A.$$

By Lemma 3.3, F is continuous on A . Therefore, because $\mathbf{x}_i \in A_i$ and $\mathbf{x}_j \in A_j$, there exists $t_* \in [0, 1]$ such that $\mathbf{b} = t_* \mathbf{x}_i + (1-t_*) \mathbf{x}_j \in A_{ij}$. Then $\mathbf{b} = \mathbf{x}_i + (1-t_*)w_r \mathbf{v}$, so by applying the ground cost we have

$$\|\mathbf{b} - \mathbf{x}_i\|_p = \|(1-t_*)w_r \mathbf{v}\|_p = (1-t_*)\|w_r \mathbf{v}\|_p \leq \|w_r \mathbf{v}\|_p = \|\mathbf{x}_j - \mathbf{x}_i\|_p.$$

Therefore, $c(\mathbf{x}_i, \mathbf{b}) \leq c(\mathbf{x}_i, \mathbf{x}_j)$.

Because we can bound the ground cost between the points in $B^r \setminus \text{edg}(A^r)$ and the set A_{ij} in terms of the ground cost between neighboring points, it is worth identifying a bound on that ground cost between neighbors.

Lemma 3.16. *For the r -th iteration of the boundary method, there exists a maximum M_r such that, for all $\mathbf{x}_i \in B_i^r$ and $\mathbf{x}_j \in B_j^r$, where $i, j \in \mathbb{N}_n$ and $i \neq j$, if $\mathbf{x}_j \in N(\mathbf{x}_i)$, then $c(\mathbf{x}_i, \mathbf{x}_j) \leq M_r \leq w_r d^{1/p}$.*

PROOF. Let \mathbf{x}_i and \mathbf{x}_j be defined as above. By applying the definition given in Equation (3.11), $\mathbf{x}_j = \mathbf{x}_i + w_r \mathbf{v}$ for some $\mathbf{v} \in \mathcal{V}$. For our Cartesian grid \mathcal{V} , $\|\mathbf{v}\|_p$ achieves its maximum when $\mathbf{v} = \mathbb{1}_d = (1, \dots, 1) \in \mathbb{R}^d$, so

$$c(\mathbf{x}_i, \mathbf{x}_j) = \|w_r \mathbf{v}\|_p = w_r \|\mathbf{v}\|_p \leq w_r \|\mathbb{1}_d\|_p = w_r d^{1/p}.$$

Therefore, there exists maximum M_r such that, for all $\mathbf{x}_i \in B_i^r$ and $\mathbf{x}_j \in B_j^r$, $c(\mathbf{x}_i, \mathbf{x}_j) \leq M_r \leq w_r d^{1/p}$.

3.3.3. Error bounds for shift differences

In order to bound the error on the Wasserstein distance, we merely require a finite bound on the errors for the individual shift differences, a_{ij} . However, accurately computing the shift differences themselves is also important, and for that reason, we also present theorems that more finely bound the error on a_{ij} for important ground cost functions. Because estimates are generated using one or more computations of $g_{ij}(\mathbf{x})$, the magnitude of these errors is dependent on the point(s) chosen.

Lemma 3.17. *Let $\mathbf{x} \in A$ and $i, j \in \mathbb{N}_n$ such that $i \neq j$. The error resulting from approximating a_{ij} at \mathbf{x} is bounded above by $|\alpha_{ij}(\mathbf{x})| \leq 2c(\mathbf{x}, \mathbf{b})$, where*

$$\alpha_{ij}(\mathbf{x}) := [c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{b}, \mathbf{y}_i)] + [c(\mathbf{b}, \mathbf{y}_j) - c(\mathbf{x}, \mathbf{y}_j)], \quad (3.14)$$

and \mathbf{b} is the point in A_{ij} nearest to \mathbf{x} with respect to the ground cost.

PROOF. Assume $\mathbf{b} \in A_{ij}$ is the closest point in A_{ij} to \mathbf{x} . Then

$$c(\mathbf{b}, \mathbf{y}_i) - c(\mathbf{b}, \mathbf{y}_j) = g_{ij}(\mathbf{b}) = a_{ij}.$$

For every $\mathbf{x} \in A$, there exists some $\alpha_{ij}(\mathbf{x}) \in \mathbb{R}$ such that

$$c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{x}, \mathbf{y}_j) = a_{ij} + \alpha_{ij}.$$

By rearrangement and substitution, we have

$$\begin{aligned} \alpha_{ij}(\mathbf{x}) &= -a_{ij} + c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{x}, \mathbf{y}_j) \\ &= -[c(\mathbf{b}, \mathbf{y}_i) - c(\mathbf{b}, \mathbf{y}_j)] + c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{x}, \mathbf{y}_j) \\ &= [c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{b}, \mathbf{y}_i)] + [c(\mathbf{b}, \mathbf{y}_j) - c(\mathbf{x}, \mathbf{y}_j)]. \end{aligned}$$

Since c satisfies the triangle inequality,

$$c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{b}, \mathbf{y}_i) \leq c(\mathbf{x}, \mathbf{b}) + c(\mathbf{b}, \mathbf{y}_i) - c(\mathbf{b}, \mathbf{y}_i) = c(\mathbf{x}, \mathbf{b}).$$

Thus,

$$|c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{b}, \mathbf{y}_i)| \leq |c(\mathbf{x}, \mathbf{b})| = c(\mathbf{x}, \mathbf{b}),$$

and, by a similar line of reasoning, $|c(\mathbf{b}, \mathbf{y}_j) - c(\mathbf{x}, \mathbf{y}_j)| \leq c(\mathbf{x}, \mathbf{b})$. Therefore,

$$|\alpha_{ij}(\mathbf{x})| \leq |c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{b}, \mathbf{y}_i)| + |c(\mathbf{b}, \mathbf{y}_j) - c(\mathbf{x}, \mathbf{y}_j)| \leq 2c(\mathbf{x}, \mathbf{b}).$$

In addition to bounding the error for individual points \mathbf{x} , we can also establish meaningful global bounds.

Lemma 3.18. *Let α_{\max} be the maximum value of $|\alpha_{ij}(\mathbf{x})|$ over all $\mathbf{x} \in \bar{B}^r$ and $i, j \in \mathbb{N}_n$, such that: (1) $i \neq j$, (2) $\mathbf{x} \in \mathbf{x}_i^r$ for some $\mathbf{x}_i \in B_i^r$, and (3) $B_j^r \cap N(\mathbf{x}_i) \neq \emptyset$. Then $\alpha_{\max} \leq 4w_r d^{1/p}$ and for all $\mathbf{x} \in \bar{B}^r$, $\|\mathbf{x} - B\|_p \leq 2w_r d^{1/p}$.*

PROOF. Suppose $\mathbf{x} \in \bar{B}^r$. By the definition of our grid, \mathbf{x} is contained in some $G = \text{Conv}(S)$, where S is a finite set of neighboring grid points. For each $\mathbf{x}_a, \mathbf{x}_b \in S$, $\mathbf{x}_b \in N(\mathbf{x}_a)$, and hence $\mathbf{x}_b = \mathbf{x}_a + w_r \mathbf{v}$ for some $\mathbf{v} \in \mathcal{V}$. Since $\|\mathbf{v}\|_p \leq d^{1/p}$, $c(\mathbf{x}_a, \mathbf{x}_b) \leq w_r d^{1/p}$. Because \mathbf{x}_a and \mathbf{x}_b were arbitrarily chosen, this is true of every pair of vertices of G . By the definition of G , \mathbf{x} can be written as a convex combination of the points in S . Therefore, for any fixed $\mathbf{x}_0 \in S$, $c(\mathbf{x}, \mathbf{x}_0) \leq w_r d^{1/p}$.

Because $\mathbf{x} \in \bar{B}^r$, $\text{Conv}(S) \cap B^r$ must be nonempty. Assume without loss of generality that $\mathbf{x}_0 = \mathbf{x}_i \in B_i^r$ for some $i \in \mathbb{N}_n$. By Lemma 3.17, there must exist a point $\mathbf{x}_j \in B_j^r$, a neighbor of \mathbf{x}_i , with $j \neq i$, and a point $\mathbf{b} \in A_{ij}$ such that $c(\mathbf{x}_i, \mathbf{b}) \leq c(\mathbf{x}_i, \mathbf{x}_j) \leq w_r d^{1/p}$. Applying the triangle inequality, we find that

$$c(\mathbf{x}, \mathbf{b}) \leq c(\mathbf{x}, \mathbf{x}_i) + c(\mathbf{x}_i, \mathbf{b}) \leq 2w_r d^{1/p}.$$

Therefore, $\|\mathbf{x} - B\|_p \leq 2w_r d^{1/p}$ and $\alpha_{\max} \leq 4w_r d^{1/p}$.

3.3.4. Error bound for ground costs

In preparation for bounding the Wasserstein distance error, we now bound the error on the ground cost c with respect to individual points in \bar{B}^r .

Lemma 3.19. *Let $\tilde{\pi}^*$ be an approximated transport plan with associated transport map \tilde{T} . Suppose π^* is an optimal transport plan with associated map T , and let \mathbf{x} in A such that $T(\mathbf{x}) = \mathbf{y}_i$, but $\tilde{T}(\mathbf{x}) = \mathbf{y}_j$. Then the error in the ground cost at the point \mathbf{x} is equal to $|g_{ij}(\mathbf{x})|$. Furthermore, there exists γ_{\max} such that, for all such $\mathbf{x} \in A$ with $T(\mathbf{x}) = \mathbf{y}_i$ and $\tilde{T}(\mathbf{x}) = \mathbf{y}_j$ for some $i \neq j$,*

$$|g_{ij}(\mathbf{x})| \leq \gamma_{\max} \leq \max_{\substack{1 \leq i < n \\ i < j \leq n}} |a_{ij}| + 4w_r d^{1/p} < \infty. \quad (3.15)$$

PROOF. Let $\mathbf{x} \in A$ such that $T(\mathbf{x}) = \mathbf{y}_i$, but $\tilde{T}(\mathbf{x}) = \mathbf{y}_j$. Then the error in the ground cost at \mathbf{x} equals

$$|c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{x}, \mathbf{y}_j)| = |g_{ij}(\mathbf{x})|.$$

As a consequence of Lemma 3.18:

$$|g_{ij}(\mathbf{x})| = |c(\mathbf{x}, \mathbf{y}_i) - c(\mathbf{x}, \mathbf{y}_j)| \leq |a_{ij}| + |\alpha_{ij}(\mathbf{x})| \leq \max_{\substack{1 \leq i < n \\ i < j \leq n}} |a_{ij}| + 4w_r d^{1/p} < \infty.$$

The result is independent of \mathbf{x} , i , and j , and therefore there must exist some $\gamma_{\max} \leq \max_{\substack{1 \leq i < n \\ i < j \leq n}} |a_{ij}| + 4w_r d^{1/p} < \infty$.

3.3.5. Volume and containment for the boundary region

As shown in Section 3.3.4, the ground cost error for individual points is finitely bounded over a wide range of admissible ground cost functions. By definition, the measure μ is bounded. We propose to identify the largest possible region in which the ground cost error can be non-zero, and to show that the area of that region goes to zero as r goes to infinity. With this, we will show that the boundary method converges with respect to the Wasserstein distance.

In Equation (3.13), we defined a region \bar{B}^r based on the point set B^r . For this, we need to know that we can choose an initial width w_1 such that, for all iterations r , $B \subset \bar{B}^r$. Theorem 3.20 presents evidence that such a width exists, and gives a sense of the relevant features driving the choice of w_1 . For details about the numerical considerations involved, see Remark 3.

Theorem 3.20. *There exists an initial width w_1 such that, for all w_r such that $w_r \leq w_1$, $\mathbf{x} \in \mathring{A}_i^r$ implies $\mathbf{x}^r \subseteq \mathring{A}_i$, the strict interior of A_i , as defined by Equation (2.3).*

PROOF. Let $\mathcal{B}(\mathbf{x}, s)$ indicate the open ball of radius s (with respect to the p -norm c) centered at \mathbf{x} and $C(\mathbf{x}, s)$ indicate the d -dimensional cube with side length s (with respect to the Euclidean distance) centered at \mathbf{x} . Because c is a p -norm, for each $i \in \mathbb{N}_n$, $\mathbf{y}_i \in \mathring{A}_i$, and therefore there exists $\delta_i > 0$ such that $\mathcal{B}(\mathbf{y}_i, \delta_i) \subseteq \mathring{A}_i$. Thus, there exist $\varepsilon > 0$ and $\delta \geq \varepsilon$, such that, for any $i \in \mathbb{N}_n$, $c(\mathbf{x}, \mathbf{y}_i) < \delta$ implies $C(\mathbf{x}, 4s) \subseteq \mathring{A}_i$ for all $s \leq \varepsilon$.

Let

$$S := A \setminus \bigcup_{i \in \mathbb{N}_n} \mathcal{B}(\mathbf{y}_i, \delta).$$

Because S is a closed set minus a finite number of open sets, S is closed.

If all A_{ij} are hyperplanes on S , the claim is self-evident for all $w_1 \leq \varepsilon$, so assume instead that at least one A_{ij} is not a hyperplane on S . Let

$$G_1 := \sup_{\substack{\mathbf{x} \in S \\ i, j \in \mathbb{N}_n \\ i \neq j}} |\nabla g_{ij}(\mathbf{x})| \quad \text{and} \quad G_2 := \sup_{\substack{\mathbf{x} \in S \\ i, j \in \mathbb{N}_n \\ i \neq j}} |\nabla^2 g_{ij}(\mathbf{x})|.$$

There exists a maximum directional magnitude with respect to the Euclidean distance,

$$M = \sup_{\substack{\mathbf{x} \in S \\ \mathbf{u} \in \mathbb{R}^d \\ i \in \mathbb{N}_n}} |x_{\mathbf{u}} - y_{\mathbf{u}}^i| \leq l \sqrt{d} < \infty,$$

where $|x_{\mathbf{u}} - y_{\mathbf{u}}^i|$ is the magnitude of the vector $\mathbf{x} - \mathbf{y}_i$ projected parallel to the direction of \mathbf{u} . Because $c \in C^2(S)$, G_1 and G_2 are well-defined. For any $\mathbf{x} \in S$ and any unit direction vector $\mathbf{u} \in \mathbb{R}^d$,

$$\begin{aligned} |\nabla_{\mathbf{u}} g_{ij}(\mathbf{x})| &= \left| \frac{(x_{\mathbf{u}} - y_{\mathbf{u}}^i) |x_{\mathbf{u}} - y_{\mathbf{u}}^i|^{p-2}}{(c(\mathbf{x}, \mathbf{y}_i))^{p-1}} - \frac{(x_{\mathbf{u}} - y_{\mathbf{u}}^j) |x_{\mathbf{u}} - y_{\mathbf{u}}^j|^{p-2}}{(c(\mathbf{x}, \mathbf{y}_j))^{p-1}} \right| \\ &\leq 2 \frac{M^{p-1}}{\delta^{p-1}} < \infty. \end{aligned}$$

and

$$|\nabla_{\mathbf{u}}^2 g_{ij}(\mathbf{x})| = (p-1) \left| \frac{|x_{\mathbf{u}} - y_{\mathbf{u}}^i|^{p-2}}{(c(\mathbf{x}, \mathbf{y}_i))^{p-1}} - \frac{|x_{\mathbf{u}} - y_{\mathbf{u}}^j|^{p-2}}{(c(\mathbf{x}, \mathbf{y}_j))^{p-1}} \right|$$

$$\begin{aligned} & \left| -\frac{|x_{\mathbf{u}} - y_{\mathbf{u}}^j|^{p-2}}{(c(\mathbf{x}, \mathbf{y}_j))^{p-1}} + \frac{|x_{\mathbf{u}} - y_{\mathbf{u}}^j|^{2p-2}}{(c(\mathbf{x}, \mathbf{y}_j))^{2p-1}} \right| \\ & \leq 2(p-1) \frac{M^{p-2}}{\delta^{p-1}} \left[1 + \frac{M^p}{\delta^p} \right] < \infty. \end{aligned}$$

Hence, $G_1 < \infty$ and $G_2 < \infty$.

Assume the Gaussian curvature of the set A_{ij} at a point $\mathbf{x} \in A_{ij}$ is given by the function $K_{ij}(\mathbf{x})$, and when $K_{ij}(\mathbf{x}) \neq 0$ the radius of curvature is given by $R_{ij}(\mathbf{x}) = |K_{ij}(\mathbf{x})|^{-1}$. Because $K_{ij}(\mathbf{x})$ is defined as a product of first and second directional derivatives of g_{ij} , and those derivatives are bounded, there exists a maximum absolute Gaussian curvature for B on S , given by

$$K := \sup_{\substack{i, j \in \mathbb{N}_n \\ i \neq j \\ \mathbf{x} \in A_{ij} \cap S}} |K_{ij}(\mathbf{x})| < \infty.$$

Because at least one A_{ij} is not a hyperplane, $K > 0$. Because $K < \infty$, for any $i, j \in \mathbb{N}_n, i \neq j$ and any $\mathbf{x} \in A_{ij} \cap S$, the radius of curvature is bounded below: $R_{ij}(\mathbf{x}) \geq K^{-1} > 0$.

Let $\tilde{\varepsilon} = \frac{2}{K\sqrt{d}}$. Suppose $s \leq \min\{\varepsilon, \tilde{\varepsilon}\}$, $\mathbf{x}_0 \in \mathring{A}_i^r$ for some $i \in \mathbb{N}_n$, and that $\mu(A_j \cap C(\mathbf{x}_0, s)) > 0$ for some $j \in \mathbb{N}_n, j \neq i$.

The set $C(\mathbf{x}_0, 2s)$ is the cube surrounding \mathbf{x}_0 and its neighbors. Because $\mathbf{x}_0 \in \mathring{A}_i^r$, A_{ij} cannot be a hyperplane in $C(\mathbf{x}_0, 2s)$, and so R_{ij} is well-defined on $C(\mathbf{x}_0, 2s)$. If there exist $k \in \mathbb{N}_n$ and $\mathbf{x} \in C(\mathbf{x}_0, 2s)$ such that $c(\mathbf{x}, \mathbf{y}_k) < \delta$, then $C(\mathbf{x}_0, 2s) \subseteq C(\mathbf{x}, 4\varepsilon) \subseteq \mathring{A}_k$, and since $\mathbf{x}_0 \in \mathring{A}_i^r$, this implies $k = i$ and $C(\mathbf{x}_0, 2s) \subseteq \mathring{A}_i$. This implies $C(\mathbf{x}_0, s) \cap A_j = \emptyset$, which contradicts the claim that $\mu(A_j \cap C(\mathbf{x}_0, s)) > 0$. Therefore, $c(\mathbf{x}, \mathbf{y}_k) \geq \delta$ for all $k \in \mathbb{N}_n$ and $\mathbf{x} \in C(\mathbf{x}_0, 2s)$. This implies $C(\mathbf{x}_0, 2s) \subseteq S$. Hence, the intersection of the boundary A_{ij} with the cube $C(\mathbf{x}_0, 2s)$ must have a point with minimum radius of curvature,

$$\mathbf{x}_m := \arg \inf_{\mathbf{x} \in A_{ij} \cap C(\mathbf{x}_0, 2s)} R_{ij}(\mathbf{x}),$$

and since $\mathbf{x}_m \in S$, it must be the case that $R_{ij}(\mathbf{x}_m) \geq 1/\kappa$.

Because $\mu(A_j \cap C(\mathbf{x}_0, s)) > 0$, but $\mathbf{x}_0 \notin A_j$, there must exist $\mathbf{x}_c \in A_{ij} \cap C(\mathbf{x}_0, s)$. Hence, within the cube $C(\mathbf{x}_0, 2s)$, there must be a d -dimensional sphere (or partial sphere) of radius $R_{ij}(\mathbf{x}_m)$, not in A_i , whose boundary intersects \mathbf{x}_c ("partial" because the sphere may be cut off by one or more of the planes bounding the cube). Call this (partial) sphere \tilde{S} .

Since $\mathbf{x}_0 \in \mathring{A}_i^r$, it must be the case that $\tilde{S} \cap \{N(\mathbf{x}_0) \cup \{\mathbf{x}_0\}\} = \emptyset$. Because $\mathbf{x}_c \in C(\mathbf{x}_0, s)$, and the maximum distance between grid points in $C(\mathbf{x}_0, 2s)$ is $s\sqrt{d}$, this requires $R_{ij}(\mathbf{x}_m) < s\sqrt{d}/2$. Hence, there exists $\mathbf{x}_m \in A_{ij} \cap S$ such that

$$R_{ij}(\mathbf{x}_m) < \frac{s\sqrt{d}}{2} \leq \frac{1}{\kappa}.$$

This contradicts $R_{ij}(\mathbf{x}_m) \geq 1/\kappa$. Thus, it must be the case that for all $j \in \mathbb{N}_n, j \neq i$ implies $\mu(C(\mathbf{x}, s) \cap A_j) = 0$, and therefore $C(\mathbf{x}, s) \subseteq \mathring{A}_i$.

Setting $w_1 \leq \min\{\varepsilon, \tilde{\varepsilon}\}$ completes the proof.

Remark 3. *The boundary method's strategy progressively refines the boundaries between individual regions A_i . Thus, by the method's very nature, any initial configuration must enclose*

the boundary in a way that allows it to be distinguished from the region interiors. This is the meaning behind the width w_1 considered in Theorem 3.20. In principle, w_1 may need to be quite small. In practice, the potential problems associated with an overly-large w_1 rarely occur, and they are obvious when they do. We did occasionally observe an issue when the initial w_1 was so large that a region \mathbf{x}^r could contain an entire A_i (in other words, when w_1 was significantly larger than the δ described in Theorem 3.20). In those cases, the affected region's a_i was such that $c(\mathbf{y}_i, \mathbf{y}_j) = a_i - a_j$ for some $j \neq i$, and the resulting transport plan had $\mu(A_i) = 0$. Hence, the set $\{a_i\}_{i=1}^n$ and reconstructed regions $\{A_i\}_{i=1}^n$ directly reveal when such an error has occurred.

Also, because of the nature of the iterative method, a poor choice of w_1 quickly becomes obvious in the boundary region itself. Simply put, the loss of any portion of the boundary set B destabilizes the method. Losing part of B creates a visible gap in the “wall” between two regions, and the gap increases in size with each successive iteration. This behavior seems to occur whenever some part of B is lost, no matter what the cause. For example, in our tests we observed that discarding an edge box that intersects B results in the same progressive damage to the boundary set. Not surprisingly, this also “stalls” the convergence of the Wasserstein distance in ways that are obvious during computation.

In our numerical tests, we used $w_1 \leq 1/50n$ and obtained consistently reliable results.

Next, we show that a well-chosen initial width and grid arrangement can guarantee that, for every iteration r , each point in $A^r \setminus B^r$ corresponds to a box in the interior of some region A_i .

Theorem 3.21. *Suppose w_1 is chosen as described in Theorem 3.20. Fix r , and let $w_r \leq w_1$. If $B \subseteq \bar{A}^r$, then $B \subseteq \bar{B}^r$, and hence $B \subseteq A^{r+1}$.*

PROOF. We will show the conclusions by proving that $\mathbf{x}_0 \notin \bar{B}^r$ implies $\mathbf{x}_0 \notin B$.

Suppose $\mathbf{x}_0 \notin \bar{B}^r$. If $\mathbf{x}_0 \notin \bar{A}^r$, then $\mathbf{x}_0 \notin B$, since by assumption, $B \subseteq \bar{A}^r$. Thus, we assume instead that $\mathbf{x}_0 \in \bar{A}^r \setminus \bar{B}^r$.

Because $\mathbf{x}_0 \in \bar{A}^r$, we know $\mathbf{x}_0 \in \mathbf{x}^r$ for some $\mathbf{x} \in A^r$. We have $\mathbf{x} \in A_i$ for some $i \in \mathbb{N}_n$, and so $\mathbf{x} \in A_i^r$. However, $\mathbf{x}_0 \notin \bar{B}^r$ implies $\mathbf{x}^r \not\subseteq \bar{B}^r$, so $\mathbf{x} \notin B^r$. Because, $\mathbf{x} \in A_i^r \setminus B^r = \dot{A}_i^r$, by Theorem 3.20, $\mathbf{x}^r \subseteq \dot{A}_i$. Hence, $\mathbf{x}_0 \in \dot{A}_i$. Therefore, by Equation (2.3), $\mathbf{x}_0 \notin B$.

Now that we have ensured $B \subseteq \bar{B}^r$, we aim to construct a region of controlled volume enclosing \bar{B}^r : $\bar{B}^r \subseteq \bar{B}_+^r$. Then we show that, as $r \rightarrow \infty$, the volume of \bar{B}_+^r in \mathbb{R}^d goes to zero with respect to the Lebesgue measure. This will allow us to put a convenient upper bound on the volume of \bar{B}^r in terms of the width w_r . Because \bar{B}_+^r exists solely in A , and not on the product space, we can once again rely on the Euclidean distance in \mathbb{R}^d .

Lemma 3.22. *Let the region $\bar{B}_+^r \subseteq A$ be defined as*

$$\bar{B}_+^r := \{\mathbf{x} \in A \mid \|\mathbf{x} - B\|_2 \leq 2w_r \sqrt{d}\}. \quad (3.16)$$

For all r , $\bar{B}^r \subseteq \bar{B}_+^r$.

PROOF. By definition, $\bar{B}^r \subseteq A$. Suppose $\mathbf{x} \in \bar{B}^r$. Because we are applying the Euclidean norm, Lemma 3.18 implies that $\|\mathbf{x} - B\|_2 \leq 2w_r \sqrt{d}$, and since $\mathbf{x} \in A$, $\mathbf{x} \in \bar{B}_+^r$.

Theorem 3.23. *If $\mu(B) = 0$ and $|B| = \tilde{L} < \infty$ with respect to the \mathbb{R}^{d-1} Lebesgue measure, then there exists some $L < \infty$, such that $|\bar{B}^r| \leq w_r^d L$ with respect to the \mathbb{R}^d Lebesgue measure.*

PROOF. We know $\int_{\bar{B}_+^r} d\mathbf{x} = \int_A \chi[\bar{B}_+^r](\mathbf{x}) d\mathbf{x}$. Let $\mathcal{B}(\mathbf{x}, \rho)$ be the closed ball of radius ρ centered at \mathbf{x} , and write

$$\begin{aligned} \int_A \chi[\bar{B}_+^r](\mathbf{x}) d\mathbf{x} &= \int_A \chi[\{\mathbf{x} \in A \mid \|\mathbf{x} - B\|_2 \leq 2w_r \sqrt{d}\}](\mathbf{x}) d\mathbf{x} \\ &= \int_A \chi[\{\mathbf{x} \in A \mid \mathbf{x} \in \mathcal{B}(\mathbf{z}, 2w_r \sqrt{d}) \text{ for some } \mathbf{z} \in B\}](\mathbf{x}) d\mathbf{x}, \\ &\leq \int_A \chi[B](\mathbf{z}) \left(\int_A \chi[\mathcal{B}(\mathbf{z}, 2w_r \sqrt{d})](\mathbf{x}) d\mathbf{x} \right) d\mathbf{z}. \end{aligned}$$

For all fixed \mathbf{x} ,

$$\int_A \chi[\mathcal{B}(\mathbf{x}, 2w_r \sqrt{d})](\mathbf{x}) d\mathbf{x} \leq \text{Vol}_d(2w_r \sqrt{d}),$$

where $\text{Vol}_d(\rho)$ is the volume of the d -dimensional sphere of radius ρ . By using the Γ function, this volume can be written as

$$\begin{aligned} \text{Vol}_d(2w_r \sqrt{d}) &:= \frac{\pi^{d/2}}{\frac{d}{2} \Gamma(\frac{d}{2})} (2w_r \sqrt{d})^d \\ &= \begin{cases} \frac{\pi^{d/2}}{(d/2)!} (2w_r \sqrt{d})^d & \text{if } d = 2k \text{ for some } k \in \mathbb{Z} \\ \frac{2k!(4\pi)^k}{d!} (2w_r \sqrt{d})^d & \text{if } d = 2k + 1 \text{ for some } k \in \mathbb{Z}. \end{cases} \end{aligned}$$

Because the volume is independent of the point $\mathbf{x} \in A$, we therefore have

$$\begin{aligned} \int_{\bar{B}_+^r} d\mathbf{x} &\leq \int_A \chi[B](\mathbf{z}) \int_A \chi[\mathcal{B}(\mathbf{z}, 2w_r \sqrt{d})](\mathbf{x}) d\mathbf{x} d\mathbf{z}, \\ &\leq \int_A \chi[B](\mathbf{z}) \text{Vol}_d(2w_r \sqrt{d}) d\mathbf{z} = \text{Vol}_d(2w_r \sqrt{d}) \int_B d\mathbf{z} = w_r^d L, \end{aligned}$$

where

$$L := \begin{cases} \tilde{L} \frac{\pi^{d/2}}{(d/2)!} (2\sqrt{d})^d & \text{if } d = 2k \text{ for some } k \in \mathbb{Z} \\ \tilde{L} \frac{2k!(4\pi)^k}{d!} (2\sqrt{d})^d & \text{if } d = 2k + 1 \text{ for some } k \in \mathbb{Z}. \end{cases}$$

Let $\mathbf{x} \in \bar{B}^r$. By applying Lemma 3.18 with c the Euclidean distance, we know that for all $\mathbf{x} \in \bar{B}^r$, $\|\mathbf{x} - B\|_2 \leq 2w_r \sqrt{d}$, which implies $\mathbf{x} \in \bar{B}_+^r$. Thus, $\bar{B}^r \subseteq \bar{B}_+^r$, which implies $|\bar{B}^r| \leq |\bar{B}_+^r| \leq w_r^d L$.

Remark 4. *The interplay between B , B^r , \bar{B}^r , and \bar{B}_+^r is nontrivial. Figure 3 helps to visualize it properly. In Figure 3(a), we show placement of some boundary set B^r . It is crucial that the subgrid created by B^r completely surrounds B , because that is the only way to ensure that $B \subseteq \bar{B}^r$. One can see in this image how a (very degenerate) choice of c , coupled with the right arrangement of \mathbf{y}_i 's, might allow a small and sharply curved boundary set to slip unnoticed between points.*

As Figure 3(b) illustrates, each point in B^r appears as the center of its corresponding box, and the boxes completely cover the boundary set.

The region \bar{B}_+^r is deliberately constructed to entirely cover all the boxes in \bar{B}^r . As Figure 3(c) shows, its volume can be significantly larger than that of the boxes it contains. However, the worst-case "thickness" given to \bar{B}_+^r ensures that it will always enclose both B and \bar{B}^r .




Figure 3: Detail from problem in Figure 5(a): Boundary set interactions near $A_0 \cap A_2 \cap A_3$

3.3.6. The Wasserstein distance error

Theorem 3.24. *Let P^* be the Wasserstein distance. Suppose $B \subseteq \tilde{B}^r$, and that there exists some L such that $|\tilde{B}^r| = w_r^d L < \infty$ with respect to the d -dimensional Lebesgue measure. If $\gamma_{\max} < \infty$ is the maximum error of the ground cost in the set \tilde{B}^r , and \tilde{P}^* is the Wasserstein distance approximation obtained with the boundary method, then the value of μ on A is bounded by some $M < \infty$ and*

$$|\tilde{P}^* - P^*| \leq w_r^d L M \gamma_{\max}, \quad (3.17)$$

where the bound equals the maximum possible volume of \tilde{B}^r multiplied by the maximum value of μ and the maximum error of the ground cost.

PROOF. If $\mathbf{x} \in A \setminus \tilde{B}^r$, then \mathbf{x} has been identified as being in the interior of A_i for some $i \in \mathbb{N}_n$. Thus, the cost error associated with the points outside \tilde{B}^r is zero.

Suppose instead that $\mathbf{x} \in \tilde{B}^r$. By definition, the absolute value of the difference between the correct and approximated ground costs at \mathbf{x} is less than or equal to γ_{\max} . Condition (1)(a) requires μ to be absolutely continuous, so there exists M such that, for all $\mathbf{x} \in X$, $0 \leq \mu(\mathbf{x}) \leq M < \infty$.

Therefore, the error on the Wasserstein distance is bounded above by

$$\begin{aligned} |\tilde{P}^* - P^*| &\leq \int_{\tilde{B}^r} \gamma_{\max} d\mu(\mathbf{x}) = \int_{\tilde{B}^r} \mu(\mathbf{x}) \gamma_{\max} d\mathbf{x} \\ &\leq \int_{\tilde{B}^r} (M) \gamma_{\max} d\mathbf{x} = |\tilde{B}^r| M \gamma_{\max} \leq (w_r^d L) M \gamma_{\max}. \end{aligned}$$

Remark 5. *The bounds in Theorems 3.23 and 3.24 indicate that the volume of the boundary set and the error of the computed Wasserstein distance decrease according to the dimension of the space. Thus, we should expect our numerical tests to show a quadratic (in \mathbb{R}^2) or cubic (in \mathbb{R}^3) decrease of the Wasserstein distance error. These decreases are clearly observed in practice, see Section 4.*

4. Numerical results

4.1. Test conditions

As mentioned in Remark 2, some choices of μ may make it necessary to approximate $\mu(\mathbf{x}^r)$ or $\int_{\mathbf{x}^r} c(\mathbf{z}, \mathbf{y}_i) d\mu(\mathbf{z})$. However, the majority of numerical studies we have seen restrict to simple

choices of μ (most often uniform). For this reason, we restricted our examples to cases where the cost and mass integrals can be written in a closed form.

4.1.1. The closed-form mass $\mu(\mathbf{x}^r)$

The integral of μ over some box can be written as:

$$\mu(\mathbf{x}^r) = \int_{\mathbf{x}^r} \mu(\mathbf{z}) d\mathbf{z} = M(\mathbf{z}) \Big|_{\mathbf{z} \in \mathbf{x}^r} \quad \text{with} \quad M : X \rightarrow \mathbb{R}^{\geq 0}. \quad (4.1)$$

Since μ is a probability density function, we must have $\int_A d\mu = 1$. For convenience, let $\hat{\mu}$ denote an un-normalized version of μ , and similarly for \hat{M} .

Using the linearity of the integral, one can use linear combination of simple functions for which exact solutions are known. We can also construct more complex measures by partitioning A into disjoint subsets. In this case, however, we add an additional restriction in order to be sure that exact solutions can always be found: We μ -partition A into subsets S_1, \dots, S_σ , such that the boundaries of each S_s fall on the initial set of grid lines. Assume that for each set S_s , there exists a density function $\hat{\mu}_s$ that is exactly solvable on S_s . From these, we consider $\hat{\mu}$ (and \hat{M}) to be the piecewise functions defined on each S_s as $\hat{\mu}_s$ (and \hat{M}_s , respectively).

Most of our computations were performed in two-dimensions. For such problems, given iteration r and $\mathbf{x} = (x_1, x_2) \in A$, $\hat{\mu}(\mathbf{x}^r)$ can be written as

$$\begin{aligned} \hat{\mu}(\mathbf{x}^r) = & \hat{M}(x_1 + w_r/2, x_2 + w_r/2) - \hat{M}(x_1 + w_r/2, x_2 - w_r/2) \\ & - \hat{M}(x_1 - w_r/2, x_2 + w_r/2) + \hat{M}(x_1 - w_r/2, x_2 - w_r/2). \end{aligned} \quad (4.2)$$

The closed-form choices used in our numerical tests are shown in Table 1. As described above, we used the table entries as building blocks in the construction of more complex measures.

Table 1: Closed-form options for μ

$\hat{\mu}((x_1, x_2)) = 1$		$\hat{M}(u, v) = uv$
$\hat{\mu}((x_1, x_2)) = x_1^t x_2^t, \quad t > 0$		$\hat{M}(u, v) = (t+1)^{-2} (u^{t+1} v^{t+1})$
$\hat{\mu}((x_1, x_2)) = e^{tx_1}, \quad t \neq 0$		$\hat{M}(u, v) = t^{-1} v e^{tu}$
$\hat{\mu}((x_1, x_2)) = e^{tx_2}, \quad t \neq 0$		$\hat{M}(u, v) = t^{-1} u e^{tv}$

4.1.2. The closed-form Wasserstein distance over \mathbf{x}^r

We performed many tests where μ could be computed exactly but the Wasserstein distance could not; see Section 4 for details. In such cases, we made no attempt to approximate P^* , choosing instead to focus on the accuracy of the μ -partition generated by the approximate shift set $\{\tilde{a}_i\}_{i=1}^n$.

However, there were a number of cases in two dimensions where the choice of μ and c allowed for closed-form computations. In those cases, because the combination of c and μ gives us an exact solution, there exists $C : X \times Y \rightarrow \mathbb{R}^{\geq 0}$ such that

$$\int_{\mathbf{x}^r} c(\mathbf{z}, \mathbf{y}_i) d\mu(\mathbf{z}) = C(\mathbf{z}, \mathbf{y}_i) \Big|_{\mathbf{z} \in \mathbf{x}^r}. \quad (4.3)$$

As in Section 4.1.1, we write \hat{C} when working with $\hat{\mu}$.

Now consider $X, Y \subset \mathbb{R}^2$, $\mathbf{x} = (x_1, x_2) \in A$, and $\mathbf{y} = (y_1, y_2) \in \{\mathbf{y}_i\}_{i=1}^n$. When $\mu(\mathbf{x}^r) = 0$, the Wasserstein distance on \mathbf{x}^r is also zero. For those boxes where $\mu(\mathbf{x}^r) > 0$, we can take advantage of the uniformity to define the function \hat{C} in terms of a single variable: the component-wise distance between points given by (Δ_1, Δ_2) , where $\Delta_1 = |x_1 - y_1|$, $\Delta_2 = |x_2 - y_2|$. When the Wasserstein distance over \mathbf{x}^r can be computed and is non-zero, it takes the form

$$\int_{\mathbf{x}^r} c(\mathbf{z}, \mathbf{y}) d\hat{\mu}(\mathbf{z}) = \hat{C}(\Delta_1 + w_r/2, \Delta_2 + w_r/2) - \hat{C}(\Delta_1 + w_r/2, \Delta_2 - w_r/2) \\ - \hat{C}(\Delta_1 - w_r/2, \Delta_2 + w_r/2) + \hat{C}(\Delta_1 - w_r/2, \Delta_2 - w_r/2), \quad (4.4)$$

where $\hat{C} : \mathbb{R}^2 \rightarrow \mathbb{R}^{\geq 0}$ is an explicit function.

Table 2 gives Wasserstein distance functions \hat{C} for c the 2-norm and the p -th power of some p -norm ($p \in [1, \infty)$). By leveraging the linearity of the integral and subdividing A into disjoint sets, we can build combinations of ground costs and measures with closed form C . We used this to perform tests in \mathbb{R}^2 , with μ being either uniform or zero in relevant boxes.

Table 2: Closed-form options for C when μ is uniform or zero on A

c	$\hat{C}(u, v)$
2-norm	$\begin{cases} \frac{1}{6}u^3 \log(\sqrt{u^2 + v^2} + v) \\ \quad + \frac{1}{3}uv \sqrt{u^2 + v^2} \\ \quad + \frac{1}{6}v^3 \log(\sqrt{u^2 + v^2} + u) & \text{if}(u, v) \neq \mathbf{0} \\ 0 & \text{if}(u, v) = \mathbf{0} \end{cases}$
p -th power p -norm	$(p + 1)^{-1}(u^{p+1}v + uv^{p+1})$

4.2. Accuracy of the Wasserstein distance

4.2.1. When exact values are known.

When μ is uniform and the region boundaries are known exactly, we can use the formulas for the ground cost given in Table 2 to compute exact values. The 2-norm was used for the problems shown in Figures 4(a) and 4(b), and the 1-norm was used to generate Figure 4(c).

For two points on the Northwest-Southeast diagonal, placed as shown in Figure 4(a), the exact Wasserstein distance is equal to

$$P_{\text{NWSE}}^* := \frac{1}{96} \left[\sqrt{2} + 7\sqrt{10} + \sinh^{-1}(1) + 2\sqrt{2} \sinh^{-1}(2) + \sinh^{-1}(3) \right] \\ \approx 0.3159707808963017.$$

Table 3(a) shows the absolute error for various w_* . Note that the actual decrease in error is roughly quadratic in w_* : $|\tilde{P}_{\text{NWSE}}^* - P_{\text{NWSE}}^*| \approx 2.122(w_*)^{1.995}$.

When we have a 4×4 arrangement of boxes, with each \mathbf{y}_i in the center, as shown in Figure 4(b), the exact Wasserstein distance is equal to

$$P_{4 \times 4}^* := \frac{1}{24} \left[\sqrt{2} + \sinh^{-1}(1) \right] \approx 0.09564946455802659.$$




Figure 4: Problems where the exact Wasserstein distance and set of shifts are known

Table 3: Wasserstein errors for the NW-SE, 4×4 , and “bad” 1-norm problems

w_*	abs. error	w_*	abs. error	w_*	abs. error
2^{-9}	8.42×10^{-6}	2^{-9}	2.02×10^{-5}	2^{-9}	8.66×10^{-6}
2^{-10}	2.11×10^{-6}	2^{-10}	5.04×10^{-6}	2^{-10}	2.16×10^{-6}
2^{-11}	5.27×10^{-7}	2^{-11}	1.26×10^{-6}	2^{-11}	5.40×10^{-7}
2^{-12}	1.32×10^{-7}	2^{-12}	3.15×10^{-7}	2^{-12}	1.35×10^{-7}
2^{-13}	3.30×10^{-8}	2^{-13}	7.88×10^{-8}	2^{-13}	3.32×10^{-8}
2^{-14}	8.24×10^{-9}	2^{-14}	1.97×10^{-8}	2^{-14}	8.26×10^{-9}
2^{-15}	2.06×10^{-9}	2^{-15}	4.93×10^{-9}	2^{-15}	2.06×10^{-9}
2^{-16}	5.15×10^{-10}	2^{-16}	1.23×10^{-9}	2^{-16}	5.15×10^{-10}
2^{-17}	1.29×10^{-10}	2^{-17}	3.08×10^{-10}	2^{-17}	1.29×10^{-10}

(a) NW-SE errors

(b) 4×4 errors

(c) “Bad” 1-norm errors

Table 3(b) shows the error. Again, the observed error decrease is roughly quadratic in w_* : $|\widetilde{P}_{4 \times 4}^* - P_{4 \times 4}^*| \approx 5.254(w_*)^{1.999}$.

Recall that the theorems presented in Section 3 offer no convergence guarantee for the behavior of the 1-norm. In fact, the optimal solution may not be μ -a.e. unique, in which case the set $\{A_i\}_{i=1}^n$ may not partition A . This is exactly what happens for the problem shown in Figure 4(c). The problem is identical to that shown in Figure 4(a), except that c is the 1-norm. Because of the change in norms, the Northeast and Southwest corners of A do not have unique transport destinations. The loss of μ -a.e. uniqueness, and resulting failure to partition, is clearly visible in the figure. However, the exact Wasserstein distance can still be computed, and is equal to

$$P_{\text{bad}}^* := \frac{19}{48}.$$

Table 3(c) shows the error. Even though the theorems do not guarantee convergence, and partitioning fails, the decrease in error is nonetheless quadratic in w_* : $|\widetilde{P}_{\text{bad}}^* - P_{\text{bad}}^*| \approx 2.575(w_*)^{2.016}$.

For all three problems, we know the exact shift values: since every point in A goes to the nearest y_i , the shift differences are all zero, which means every shift should be identical. In the 4×4 and “bad” 1-norm problems, the shift values are identical for every choice of w_* . For the 4×4 problem, this is a result of computing regions that exactly correspond to the structure of

our grid. For the “bad” 1-norm problem, the exactness derives from the relative simplicity of ground cost computations. The shift values for the NWSE problem have an error whose decrease is roughly linear with respect to w_* : $|\tilde{a}_2 - \tilde{a}_1| \approx 0.339(w_*)^{1.008}$. When $w_* = 2^{-16}$, this shift error is 4.83×10^{-6} .

4.2.2. When exact values are not known.

As Section 3.3.6 shows, even if the Wasserstein distance is unknown, the Wasserstein approximation error at the end of the r -th iteration is bounded above by

$$\sum_{\mathbf{x} \in B^r} \mu(\mathbf{x}^r) \max_{\mathbf{x}_0 \in \mathbf{x}^r} g_{ij}(\mathbf{x}_0), \quad (4.5)$$

where i and j , $i \neq j$ refer to the destinations of \mathbf{x} and some neighbor. In practice, we can use continuity to refine that estimate still further, as described in Section 2.2.3.

As $w_r \rightarrow 0$, $\max_{\mathbf{x}_0 \in \mathbf{x}^r} g_{ij}(\mathbf{x}_0) \rightarrow |a_{ij}|$ for each $i \neq j$, and $\mu(\bar{B}^r) \rightarrow 0$. If the boundary method is working effectively, we can expect to see the Wasserstein distance error decreasing with respect to $\mu(\bar{B}^r)$. If μ is uniform on A , that decrease should be linear with respect to the volume $|\bar{B}^r|$.

We considered the change in the computed Wasserstein distance for Example 2.1 using three canonical ground costs: the 1-norm, the 2-norm, and the squared 2-norm. The resulting μ -partitions are shown in Figures 7(a), 5(a), and 8(a), respectively. Since μ is uniform, and $A = [0, 1]^2$, Theorem 3.24 suggests that we should see a quadratic convergence for the 2-norm. (The theorem makes no convergence claim for the 1-norm or the squared 2-norm.)

For each ground cost, we computed the Wasserstein approximation error in two ways:

1. The worst-case Wasserstein distance error bound, given by applying (4.5).
2. The rate of change with respect to a reference approximation,

$$\Delta \tilde{P}_{16}^*(w_*) := \Delta \tilde{P}_{16}^*(2^{-m}) = \frac{|\tilde{P}_m^* - \tilde{P}_{16}^*|}{|\tilde{P}_{m+1}^* - \tilde{P}_{16}^*|}, \quad \text{for } m < 15.$$

For all three ground cost functions, the worst-case Wasserstein distance error is roughly linear in w_* , and the rate of change of $\Delta \tilde{P}_{16}^*$ is roughly quadratic in w_* . Though Theorem 3.24 only guarantees quadratic convergence for the 2-norm, we also observe quadratic convergence for the 1-norm and squared 2-norm. For this example, the 1-norm generated a μ -a.e. unique partition, but comparable convergence was seen in tests where partitioning failed. Results are given in Table 4.

Table 4: Wasserstein approximation behavior with respect to w_*

c	\tilde{P}_{16}^*	$\text{err}_{\max}(w_*)$	$\Delta \tilde{P}_{16}^*(w_*)$
1-norm	0.25702262181	$0.457(w_*)^{1.020}$	$1.186(w_*)^{2.025}$
2-norm	0.20754605961	$0.361(w_*)^{1.008}$	$4.151(w_*)^{2.023}$
squared 2-norm	0.05290682486	$0.221(w_*)^{1.008}$	$2.668(w_*)^{2.029}$

4.3. μ -Partitions in \mathbb{R}^2

4.3.1. Uniform and non-uniform measures μ and ν

We include three examples with variations of μ and ν , shown in Figure 5. All three assume c is the 2-norm.

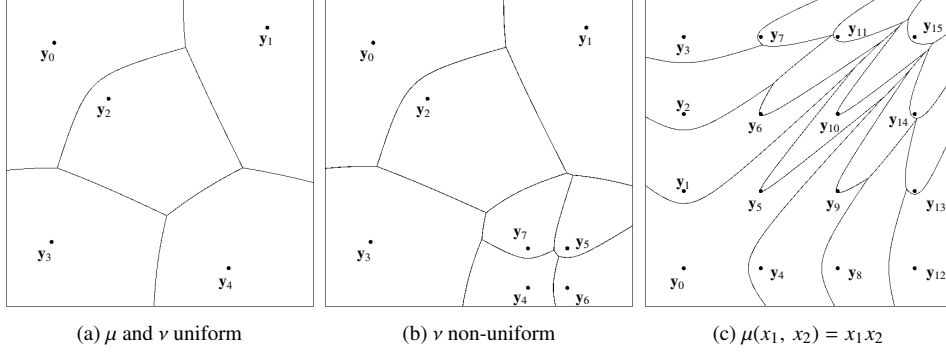


Figure 5: Partitions for uniform and non-uniform measures

In Figure 5(a), we assume μ is the uniform continuous probability distribution on A , and ν is the uniform discrete distribution with $n = 5$. The five points where $\nu = 1/5$ are placed in the positions used in [30]. Figure 5(a) shows the μ -partition obtained by the boundary method; for comparison, see [30, Figure 3 (right)].

Starting from the points shown in Figure 5(a), we next take the point y_4 and split it into four new points, each of one quarter-mass, positioned equidistantly from the point's original location. This gives us a non-uniform ν with four points of weight $1/5$ and four of weight $1/20$. We keep μ uniform. The resulting μ -partition is shown in Figure 5(b).

In Figure 5(c), we choose the nonuniform probability density $\mu(x_1, x_2) = \frac{1}{4}x_1x_2$. For ν , we choose the uniform 4×4 grid of points given in Figure 4(b). By comparing the results in Figures 4(b) and 5(c), the impact of μ 's nonuniformity becomes obvious. While the individual regions no longer have equal Lebesgue measure, each has equal μ -measure $1/16$. The larger regions in the lower-left correspond to the lower density of μ in that corner, while the smaller regions in the upper-right correspond to the higher concentration of μ -density there.

4.3.2. Discontinuous and zero-measure μ

Next, we deliberately introduce a discontinuous μ that is not strictly positive:

$$\mu(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in [0, 1/2]^2 \\ 4/3 & \text{otherwise.} \end{cases} \quad (4.6)$$

We still have $\int_A d\mu(\mathbf{x}) = 1$, so μ is a probability density function on $A = [0, 1]^2$. For ν , we use the uniform 4×4 grid shown in Figure 4(b). Figure 6 shows the results.

In Figure 6(a), we see the boundary set used to generate the solution. Points between regions are retained, as are points adjacent to regions of measure zero. No computations are done on the lower-left region, because any destination is equally valid on boxes of μ -measure zero.

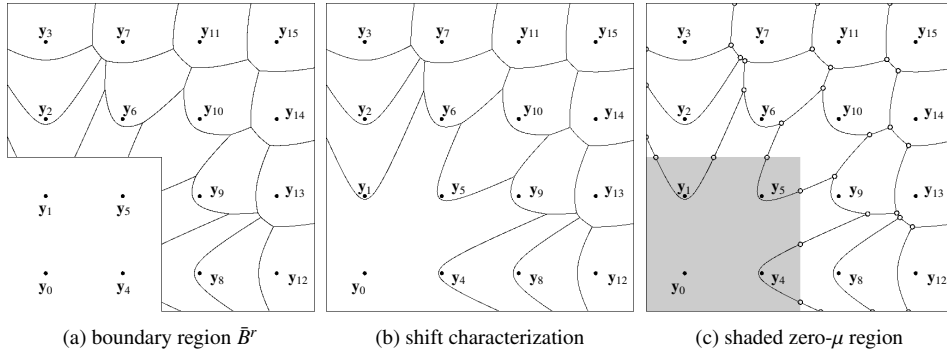


Figure 6: μ is zero in the lower-left quadrant

However, when the shift definition is applied to the semi-discrete optimal transport problem, there is only one valid shift-characterized solution over A . Figure 6(b) shows that solution. The unique shift differences force the selection of a unique boundary set B , even in the region where μ is zero.

Figure 6(c) shows the shift characterization again, but here the region of μ -zero measure is shaded, helping to confirm visually that the regions have equal μ -measure. Figure 6(c) also shows the locations of intersection points we identified using the boundary method. These intersections were used to accurately compute the set of shifts.

4.3.3. Norms as ground cost functions

The computations in Sections 4.3.1 and 4.3.2 all assume the ground cost function equals the 2-norm, but as Section 4.2.2 suggests, computation with other functions is quite possible. Using the same problem solved with the 2-norm in Figure 5(a), we generated μ -partitions for a wide range of p -norm ground costs. Results for the 1-norm, 10-norm, and ∞ -norm are shown in Figure 7. Note that the 1-norm and ∞ -norm converge to (μ -a.e. unique) solutions, even though those norms are not covered by our theorems.




Figure 7: Equal area using different ground cost norms

4.3.4. Other ground costs c

The computations above all assume that the ground cost function is a norm. However, the boundary method works equally well on much more general ground cost functions. Three examples are shown in Figure 8.




Figure 8: Equal area using non-norm ground costs

Figure 8(a) shows the result given by the squared 2-norm. Because the squared 2-norm is not itself a norm, we were only able to make the most general mathematical claims regarding its behavior. However, the boundary method has no trouble with it. In fact, as Section 4.2.2 indicates, its convergence behavior is practically identical to that of the p -norms.

When $0 < p < 1$, the p -norm formula can still be applied, even though the resulting function does not satisfy the triangle inequality. Formally, one has

$$c_p(\mathbf{x}_1, \mathbf{x}_2) := \begin{cases} \left[\sum_{k=1}^d |x_k^2 - x_k^1|^p \right]^{1/p} & p \in (0, \infty) \\ \max_{k \in \mathbb{N}_d} |x_k^2 - x_k^1| & p = \infty. \end{cases} \quad (4.7)$$

As Figure 8(b) illustrates, even these “ p -function” transport problems can be approximated. However, when $p < 1$, the regions become discontinuous and disconnected, as typified by the “spikes” on the exterior walls. (The spike on the lower right is part of the region coupled with y_3 , while the other four spikes are coupled with y_2 .) Note that the $1/2$ -function is concave. Such functions are directly applicable to transport problems involving economies of scale; e.g., see [45].

Figure 8(c) shows a ground cost function defined as a polynomial combination of p -functions with positive coefficients:

$$c(\mathbf{x}_1, \mathbf{x}_2) = 4c_2(\mathbf{x}_1, \mathbf{x}_2)^{28/5} + 61c_{1/2}(\mathbf{x}_1, \mathbf{x}_2).$$

This function, like many other “ p -polynomial” functions, is neither convex nor concave, changing behavior with distance.

4.4. μ -Partitions in \mathbb{R}^3

As we showed in Section 3, there is no theoretical obstacle to applying the boundary method to higher-dimensional problems, though visual representation becomes more complex.




Figure 9: Three-dimensional semi-discrete solution with $n = 5$

The image in Figure 9 was generated by taking c to be the 2-norm, μ the uniform continuous probability density, and ν the uniform discrete probability density with five randomly-placed non-zero points in $[0, 1]^3$. Even in this relatively simple case, it is impossible to find a single point-of-view that clearly shows all five non-zero points while clearly illustrating the boundaries of the μ -partitions. However, even though clear illustration is problematic, the computations made with the boundary method were completely successful.

4.5. Scaling behavior

One important advantage to the boundary method is its reduction of the complexity of the discretized problem, compared to traditional methods. Before considering the numerical results, it is worth developing a generalized comparison that puts this reduction in perspective:

Suppose for the sake of argument that a discretization with width 2^{-M} is required to solve a problem in \mathbb{R}^2 with N positive points in Y . Generating the full grid would create a product space $X \times Y$ of size $2^{2M}N$. Say the boundary method is used instead, with a fixed initial discretization width of 2^{-4} . Each application of Step (2) of the boundary method algorithm removes approximately half the points in A^r , so by discarding interiors the method constructs a product space of size $2^{M+4}N$.

Assume that we compute solutions for both the boundary method and the full product space using the same linear solver (e.g., the network simplex method). Using it to solve the largest boundary problem of size $2^{M+4}N$, we have $V = 2^{M+4} + N$ vertices and $E = 2^{M+4}N$ edges. Solving over the full product space gives $V = 2^{2M} + N$ vertices and $E = 2^{2M}N$ edges. Hence, even if we assumed a solver with complexity $O(V)$ (and no such solver exists), the ratio would be approximately 2^M to M . Typically, it is closer to 2^{2M} to M .

Of course, this improved complexity would be irrelevant if the constant factor was excessively large. Fortunately, that this is not the case, as our next section illustrates.

Since we focus here on the semi-discrete problem, for purposes of evaluating complexity, we assume X is discretized into W^d elements, and Y into N elements. Given W sufficiently large, the resulting network has $W^d + N \sim O(W^d)$ nodes and $W^d N$ arcs.

4.5.1. Scaling on the plane with respect to $W = 1/w_*$

Here we consider scaling on the plane with respect to $W = \frac{1}{w_*}$. We used Example 2.1, with μ and ν uniform and c the 2-norm. The locations of the 5 points where $\nu = 1/5$ were fixed as

depicted in Figure 5(a). We defined target widths $w_* = 2^{-m}$, $m \in \mathbb{N}$, and computed the time taken by the boundary method. By repeating this process for a few different location sets (and averaging them), we estimated the average scaling behavior of the boundary method with respect to W . The test results are shown in Table 5(a), and the scaling equations are on the left side of Table 6.

Table 5: Planar scaling with respect to W and N

$N = 5$			$W = 2^{10}$			$W = 2^{11}$	
W	T (sec)	S (MB)	N	T (sec)	S (MB)	T (sec)	S (MB)
2^{12}	0.855	24.540	128	6.938	17.25	22.365	33.91
2^{13}	2.005	49.100	136	12.190	18.24	36.601	35.05
2^{14}	4.497	98.210	144	10.982	17.99	29.952	36.49
2^{15}	11.025	196.400	152	13.139	18.54	36.703	41.27
2^{16}	28.093	394.400	160	11.420	18.66	34.801	40.27
2^{17}	60.577	785.800	168	15.727	20.97	44.959	40.66
2^{18}	132.397	1571.840	176	15.332	21.38	44.873	43.06
2^{19}	292.158	3151.872	184	18.243	21.38	53.689	43.20
2^{20}	640.660	6309.888	192	12.796	21.60	40.029	43.66

(a) Scaling with respect to W

(b) Scaling with respect to N

Table 6: Time and storage scaling with respect to W and N separately

$T(W) \approx 4.356 \times 10^{-5} W \ln W$	Time	$T(N) \approx 4.582 \times 10^{-2} N \ln N$
$S(W) \approx 6.015 \times 10^{-3} W$	Storage	$S(N) \approx 3.162 N^{1/2}$

4.5.2. Scaling on the plane with respect to N

To evaluate planar scaling with respect to N , we performed multiple runs in $[0, 1]^2$ where $W = 2^{11}$ was fixed and μ and ν were uniform. The $N = n$ points where $\nu = 1/n$ were placed at random locations in A . Because the resulting time data was highly dependent on point placement, it was extremely noisy. Thus, we did ten runs for each N and took the median. We started with $N = 128$, increasing by eights up to to $N = 192$, for a total of 100 tests. The results are shown in the right-hand columns of Table 5(b). See the right side of Table 6 for scaling equations with respect to N .

4.5.3. Scaling interaction of W and N on the plane

Increasing N means one must consider the scaling behavior of the boundary method with respect to N , as described in Section 4.5.2, above. However, there is another relevant limiting factor for large N : the decreasing area size $\mu(A_i) = n^{-1}$ runs up against the accuracy of the reconstruction. For the problem shown in Figure 10, the area of each region is 5.0×10^{-3} . When $w_* = 2^{-11}$, the maximum error for the area of the partition regions is 8.34×10^{-4} . This is 16.7%, about one-sixth of the size of each region.

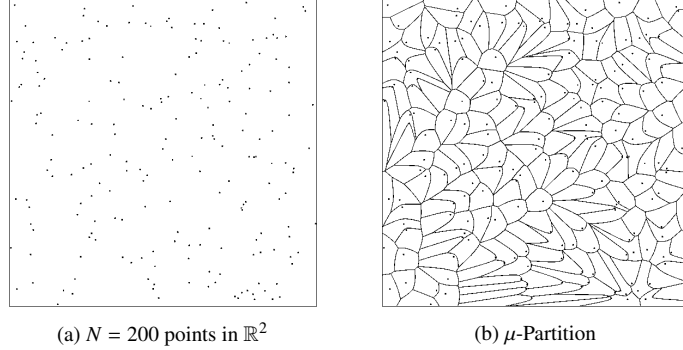


Figure 10: Partitioning with large N

If all we desire is the Wasserstein distance or the boundary set, this error need not be a concern. However, if we want an accurate set of shifts, large N requires that we increase W to match. Hence, we wanted to consider what happens as W and N increase in tandem.

As it turns out, the scaling we observe is consistent with the product of the two scaling behaviors already determined: $O(WN \log W \log N)$ with respect to time, and $O(WN^{1/2})$ with respect to storage. See Table 7 for approximate equations.

Table 7: Time and memory scaling with respect to both W and N

Time	$T(N, W) \approx 2.853 \times 10^{-6} WN \ln W \ln N$
Storage	$S(N, W) \approx 1.538 \times 10^{-3} WN^{1/2}$

4.5.4. Scaling in three dimensions and extrapolation to \mathbb{R}^d

The computations described above can be repeated in three dimensions. We scale W separately by taking a projection of Example 2.1 into the center of the cube $[0, 1]^3$. Then we consider the median of tests when $W = 2^7$ and N ranges from 8 to 80. Finally, we scale W and N together, and consider their combined behavior. Approximate scaling equations are given in Table 8.

Table 8: 3-D scaling with respect to W and N

W alone	Time	$T(W) \approx 6.878 \times 10^{-5} W^2 \ln W$
	Storage	$S(W) \approx 2.341 \times 10^{-2} W^2$
N alone	Time	$T(N) \approx 2.849 \times 10^{-1} N \ln N$
	Storage	$S(N) \approx 2.315 \times 10^2 N^{1/3}$
W and N	Time	$T(N, W) \approx 3.531 \times 10^{-6} W^2 N \ln W \ln N$
	Storage	$S(N, W) \approx 1.397 \times 10^{-1} W^2 N^{1/3}$

Taking the combined scaling equations for two and three dimensions, and extrapolating to

arbitrary dimension $d \geq 2$, we anticipate scaling of

$$T(d, N, W) \sim O(W^{d-1} N \log W \log N) \quad \text{and} \quad S(d, N, W) \sim O(W^{d-1} N^{1/d}).$$

5. Conclusions and future work

In this work, we presented the boundary method, a new technique for approximating solutions to semi-discrete optimal transportation problems. We gave an algorithmic description and mathematical justification. As we showed, by tackling only the boundary of the regions to be transported, the method has very favorable scaling properties. Under the assumption that all computations are exact, we gave sharp convergence results for p -norms with $p \in (1, \infty)$, and we presented numerical examples supporting those convergence results. We showed that the boundary method can provide accurate approximations of the partition regions and Wasserstein distance for a multitude of cost functions, including some that are not covered by our theorems: the 1-norm, the ∞ -norm, strictly convex non-norms such as the squared 2-norm, concave non-norms such as p -functions with $p \in (0, 1)$, and polynomial combinations of p -functions that are neither concave nor convex. As we also showed, even when partitioning fails, the boundary method can solve with comparable accuracy and convergence. Our future work will consider applications of the boundary method to fully continuous mass transportation problems and the impact of estimated computations on convergence.

- [1] G. Monge, Mémoire sur la théorie des déblais et des remblais, in: Histoire de l'Académie Royale des Sciences de Paris, avec les Mémoires de Mathématique et de Physique pour la même année, Académie des sciences (France), 1781, pp. 666–704, in French.
- [2] L. V. Kantorovich, On the translocation of masses, C.R. (Doklady) Acad. Sci. URSS (N.S.) 37 (1942) 199–201.
- [3] L. V. Kantorovich, On a problem of Monge, Uspekhi Mat. Nauk 3 (1948) 225–226.
- [4] C. Villani, Topics in Optimal Transportation, Vol. 58 of Graduate Studies in Mathematics, American Mathematical Society, Providence, R.I., 2003.
- [5] J. A. Cuesta-Albertos, A. Tuero-Díaz, A characterization for the solution of the Monge-Kantorovich mass transference problem, Statistics and Probability Letters 16 (2) (1993) 147–152.
- [6] A. Pratelli, On the equality between Monge's infimum and Kantorovich's minimum in optimal mass transportation, Annales de l'Institut Henri Poincaré (B): Probability and Statistics 43 (1) (2007) 1–13.
- [7] L. Rüschemdorf, Monge-Kantorovich transportation problem and optimal couplings, Jahresbericht der Deutschen Mathematiker-Vereinigung 109 (3) (2007) 113–137.
- [8] L. Rüschemdorf, L. Uckelmann, Numerical and analytical results for the transportation problem of Monge-Kantorovich, Metrika 51 (3) (2000) 245–258.
- [9] M. Muskulus, A. M. Slats, P. J. Sterk, S. Verduyn-Lunel, Fluctuations and determinism of respiratory impedance in asthma and chronic obstructive pulmonary disease, Journal of Applied Physiology 109 (2010) 1582–1591.
- [10] G. Carlier, A general existence result for the principal-agent problem with adverse selection, Journal of Mathematical Economics 35 (2001) 129–150.
- [11] S. Haker, A. Tannenbaum, Optimal mass transport and image registration, in: IEEE Workshop on Variational and Level Set Methods in Computer Vision: Proceedings: 13 July, 2001, Vancouver, Canada, IEEE Computer Society, Los Alamitos, California, 2001, pp. 29–36.
- [12] M. Cuturi, Sinkhorn distances: Lightspeed computation of optimal transport, in: Advances in Neural Information Processing Systems, Vol. 26, Curran Associates, Inc., 2013, pp. 2292–2300.
- [13] E. A. Carlen, W. Gangbo, Solution of a model Boltzmann equation via steepest descent in the 2-Wasserstein metric, Archive for Rational Mechanics and Analysis 172 (2004) 21–64.
- [14] Q. Mérigot, A multiscale approach to optimal transport, Computer Graphics Forum 30 (5) (2011) 1584–1592.
- [15] P.-A. Chiappori, R. McCann, L. Nesheim, Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness, Economic Theory 42 (2) (2010) 317–354.
- [16] P.-A. Chiappori, R. McCann, B. Pass, Multi- to one-dimensional optimal transport, to appear in Comm. Pure Appl. Math. (2016).
- [17] X. Dupuis, The semi-discrete principal agent problem, presented at “Computational Optimal Transportation” workshop, July 18–22, 2016. <http://www.crm.umontreal.ca/2016/Optimal16/pdf/dupuis.pdf>.

- [18] F. de Goes, K. Breeden, V. Ostromoukhov, M. Desbrun, Blue noise through optimal transport, *ACM Trans. Graph* 31 (6) (2012) 171:1–171:11.
- [19] F. Abedin, C. E. Gutiérrez, An iterative method for generated Jacobian equations, preprint. (2016).
- [20] T. Glimm, V. Oliker, Optical design of single reflector systems and the Monge-Kantorovich mass transfer problem, *J. Math. Sci.* 117 (3) (2003) 4096–4108.
- [21] V. Oliker, L. Prussner, On the numerical solution of the equation $\frac{\partial^2 z}{\partial x^2} \frac{\partial^2 z}{\partial y^2} - \left(\frac{\partial^2 z}{\partial x \partial y} \right)^2 = f$ and its discretizations, *I, Numerische Mathematik* 54 (1988) 271–293.
- [22] F. Aurenhammer, F. Hoffmann, B. Aronov, Minkowski-type theorems and least-squares partitioning, in: *SCG '92: Proceedings of the eighth annual symposium on Computational geometry (SOCG92 8th Annual Symposium on Computational Geometry 1992, Berlin, Germany, June 10–12, 1992)*, Association for Computing Machinery, New York, 1992, pp. 350–357.
- [23] F. Aurenhammer, F. Hoffmann, B. Aronov, Minkowski-type theorems and least-squares clustering, *Algorithmica* 20 (1) (1998) 61–76.
- [24] L. A. Caffarelli, S. A. Kochengin, V. I. Oliker, On the numerical solution of the problem of reflector design with given far-field scattering data, in: *Monge Ampère equation: applications to geometry and optimization (NSF-CBMS Conference on the Monge Ampère Equation, Applications to Geometry and Optimization, July 9–13, 1997, Florida Atlantic University)*, Vol. 226 of Contemporary Mathematics, American Mathematical Society, Providence, R.I., 1999, pp. 13–32.
- [25] B. Lévy, A numerical algorithm for L_2 semi-discrete optimal transport in 3D, *ESAIM: Mathematical Modelling and Numerical Analysis* 49 (6) (2015) 1693–1715.
- [26] J.-M. Mirebeau, Discretization of the 3D Monge-Ampère operator, between wide stencils and power diagrams, *arXiv:1503.00947* (2015).
- [27] J.-D. Benamou, B. D. Froese, A. M. Oberman, Numerical solution of the optimal transportation problem using the Monge-Ampère equation, *Journal of Computational Physics* 260 (2014) 107–126.
- [28] A. Oberman, Convergent difference schemes for degenerate elliptic and parabolic equations: Hamilton-Jacobi equations and free boundary problems, *SIAM Journal on Numerical Analysis* 44 (2006) 879–895.
- [29] B. Schmitzer, A sparse multi-scale algorithm for dense optimal transport, *arXiv:1510.05466v2* (2016).
- [30] J. W. Barrett, L. Prigozhin, A mixed formulation of the Monge-Kantorovich equations, *ESAIM: M2AN* 41 (2007) 1041–1060.
- [31] G. Bouchitté, G. Buttazzo, P. Seppecher, Shape optimization solutions via Monge-Kantorovich equation, *Comptes Rendus de l'Académie des Sciences – Series I – Mathematics* 324 (1997) 1185–1191.
- [32] J. Kitagawa, An iterative scheme for solving the optimal transportation problem, *Calculus of Variations* 51 (2014) 243–263.
- [33] X. Ma, N. S. Trudinger, X. Wang, Regularity of potential functions of the optimal transportation function, *Arch. Ration. Mech. Anal.* 177 (2) (2005) 151–183.
- [34] J. Kitagawa, Q. Mérigot, B. Thibert, A Newton algorithm for semi-discrete optimal transport, *arXiv:1603.05579* (2017).
- [35] P. Kovács, Minimum-cost flow algorithms: an experimental evaluation, *Tech. rep.*, The Egerváry Research Group (2015).
- [36] B. D. Froese, A. M. Oberman, Convergent finite difference solvers for viscosity solutions of the elliptic Monge-Ampère equation in dimensions two and higher, *Journal of Computational Physics* 260 (2014) 107–126.
- [37] R. Jordan, D. Kinderlehrer, F. Otto, The variational formulation of the Fokker-Planck equation, *SIAM Journal on Mathematical Analysis* 29 (1) (1998) 1–17.
- [38] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, G. Peyré, Iterative Bregman projections for regularized transportation problems, *SIAM Journal on Scientific Computing* 37 (2) (2015) A1111–A1138.
- [39] D. P. Bertsekas, *Network Optimization: Continuous and Discrete Models*, Athena Scientific, Belmont, Massachusetts, 1998, <http://web.mit.edu/dimitrib/www/books.htm>. Accessed: 2016-09-16.
- [40] D. P. Bertsekas, D. A. Castañón, A generic auction algorithm for the minimum cost network flow problem, *Comput. Optim. Appl.* 2 (1993) 229–260.
- [41] Q. Mérigot, A comparison of two dual methods for discrete optimal transport, in: F. Nielsen, F. Barbaresco (Eds.), *GSI 2013 — Geometric Science of Information, Aug 2013, Paris, France*, Vol. 8085 of Lecture Notes in Computer Science, Springer, 1781, pp. 389–396.
- [42] J. Walsh III, L. Dieci, General auction method for real-valued optimal transport, preprint, <http://gatech.jdwalsh03.com/index.html> (2016).
- [43] J. Walsh III, The AUCTION ALGORITHMS IN C++ project, computer software, <https://github.com/jdwalsh03/auction>.
- [44] J. Walsh III, The boundary method and general auction for optimal mass transportation and Wasserstein distance computation, Ph.D. thesis, Georgia Institute of Technology (April 2017).

[45] W. Gangbo, R. J. McCann, The geometry of optimal transportation, *Acta Mathematica* 177 (2) (1996) 113–161.