

Shape recognition via Wasserstein distance

(To appear in Quarterly of Applied Mathematics)

Wilfrid Gangbo*

School of Mathematics, Georgia Institute of Technology
Atlanta, GA 30332, USA gangbo@math.gatech.edu

Robert J. McCann†

Department of Mathematics, Brown University
Providence, RI 02912, USA mccann@math.brown.edu

February 23, 1999

Abstract

The Kantorovich-Rubinstein-Wasserstein metric defines the distance between two probability measures μ and ν on \mathbf{R}^{d+1} by computing the cheapest way to transport the mass of μ onto ν , where the cost per unit mass transported is a given function $c(\mathbf{x}, \mathbf{y})$ on \mathbf{R}^{2d+2} . Motivated by applications to shape recognition, we analyze this transportation problem with the cost $c(\mathbf{x}, \mathbf{y}) = |\mathbf{x} - \mathbf{y}|^2$ and measures supported on two curves in the plane, or more generally on the boundaries of two domains $\Omega, \Lambda \subset \mathbf{R}^{d+1}$. Unlike the theory for measures which are absolutely continuous with respect to Lebesgue, it turns out not to be the case that μ -a.e. $\mathbf{x} \in \partial\Omega$ is transported to a single image $\mathbf{y} \in \partial\Lambda$; however, we show the images of \mathbf{x} are almost surely collinear and parallel the normal to $\partial\Omega$ at \mathbf{x} . If either domain is strictly convex, we deduce that the solution to the optimization problem is unique. When both domains are uniformly convex, we prove a regularity result showing the images of $\mathbf{x} \in \partial\Omega$ are always collinear, and both images depend on \mathbf{x} in a continuous and (continuously) invertible way. This produces some unusual extremal doubly stochastic measures.

©1999 Brown University.

*WG gratefully acknowledges the support of National Science Foundation grant DMS 9622734.

†RJM gratefully acknowledges the support provided by an American Mathematical Society Centennial Fellowship and National Science Foundation grant DMS 9622997. His current address is: Department of Mathematics, University of Toronto, Toronto, Ontario, Canada M5S 3G3. E-mail: mccann@math.toronto.edu

Contents

Introduction	2
1 Background Theory, General Curves and Surfaces	9
2 Unique Solutions for Strictly Convex Shapes	13
3 Pairs of Maps Between Strictly Convex Curves	18
4 Continuous Maps Between Convex Hypersurfaces	27
A Appendix on Convex Domains	31
References	32

Introduction

In his 1993 thesis, D. Fry proposed an algorithm which he showed met with some success in enabling a computer to identify leaves of New England from their outlines [9]. Like many other approaches to shape recognition, this algorithm involved comparing the unknown leaf to a catalog of standard leaves, and deciding which known leaf afforded the best fit. The novelty of Fry's approach lay in his criterion for measuring goodness of fit between the unknown sample and the catalog leaves: his proposal (which should be compared with alternatives discussed in Mumford [18]) was to distribute unit mass uniformly along each leaf boundary, and then calculate the total cost of transporting the mass from the boundary of the sample leaf to the specified distribution on the catalog leaf. The cost per-unit-mass-transported was expressed as a sum

$$\left(\frac{\text{distance}}{\text{transported}} \right)^2 + \beta \left[1 - \cos \left(\frac{\text{angle}}{\text{rotated}} \right) \right] \approx \left(\frac{\text{distance}}{\text{transported}} \right)^2 + \frac{\beta}{2} \left(\frac{\text{angle}}{\text{rotated}} \right)^2,$$

weighted by a parameter β measuring the significance of local reorientations relative to translations. Parameterizing the two leaf boundaries by piecewise Lipschitz curves σ and $\tau : [0, 1] \rightarrow \mathbf{R}^2$ (each having constant speed) in the plane, Fry's distance between them is defined by

$$d_\beta^2(\sigma, \tau) := \inf_{\tilde{\gamma}} \int_0^1 \int_0^1 \left\{ |\sigma(s) - \tau(t)|^2 + \frac{\beta}{2} \left| \frac{\dot{\sigma}(s)}{|\dot{\sigma}|} - \frac{\dot{\tau}(t)}{|\dot{\tau}|} \right|^2 \right\} d\tilde{\gamma}(s, t), \quad (1)$$

the infimum being taken over all *doubly stochastic* measures $\tilde{\gamma}$ on the unit square: i.e. $\tilde{\gamma}[B \times [0, 1]] = \mathcal{H}^1[B] = \tilde{\gamma}[[0, 1] \times B]$ for each Borel set $B \subset [0, 1]$, where \mathcal{H}^d denotes d -dimensional Hausdorff measure (so \mathcal{H}^1 is Lebesgue measure). Thus each comparison involves computing the solution to a Monge-Kantorovich transportation problem (1),

c.f. Rachev [19]. When $\beta = 0$, this amounts to metrizing the distance between leaf boundaries using the Kantorovich-Rubinstein-Wasserstein L^2 metric [14][24], or the Wasserstein distance as it has come to be known in information theory [12].

Motivated by the desire to understand what features of the curves this distance is sensitive to, we analyze the $\beta = 0$ problem by examining the measures $\tilde{\gamma}$ which minimize (1). Such $\tilde{\gamma}$ are called *optimal*. Since the minimization problem amounts to an infinite-dimensional linear program, at least one optimal joint measure is known to exist (e.g. Kellerer [15]), and to be an extreme point in the convex set of doubly stochastic measures. The geometrical properties of this measure can be studied through its *support*, which is denoted $\text{spt } \tilde{\gamma}$ and refers to the smallest closed subset $K \subset [0, 1]^2$ of the square carrying full mass $\tilde{\gamma}[K] = 1$. Having $(s_o, t_o) \in \text{spt } \tilde{\gamma}$ means it is efficient to transport mass from the point $\mathbf{x} = \sigma(s_o)$ on the first curve to $\mathbf{y} = \tau(t_o)$ on the second one. Our main results establish topological and geometrical properties of this support, and show strict convexity of either curve is enough to ensure uniqueness of the optimal measure. We give examples demonstrating some ways in which our hypotheses are necessary and our conclusions sharp. Of course, $\beta \neq 0$ proved important in Fry's study of object recognition, but unfortunately remains beyond the scope of our techniques to analyze.

Fixing $\beta = 0$, it is convenient to dispense with the parameterizations σ and τ and rephrase the transportation problem in terms of the normalized arclength measures μ and ν of the two curves in the plane. More generally, given two Borel probability measures μ and ν on \mathbf{R}^{d+1} , the Wasserstein distance between them is defined by

$$d^2(\mu, \nu) := \inf_{\gamma \in \Gamma(\mu, \nu)} \int_{\mathbf{R}^{d+1} \times \mathbf{R}^{d+1}} |\mathbf{x} - \mathbf{y}|^2 d\gamma(\mathbf{x}, \mathbf{y}). \quad (2)$$

Here $\Gamma(\mu, \nu)$ denotes the set of all Borel measures on $\mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ having μ and ν as marginals: $\mu[B] = \gamma[B \times \mathbf{R}^{d+1}]$ and $\gamma[\mathbf{R}^{d+1} \times B] = \nu[B]$ for each Borel set $B \subset \mathbf{R}^{d+1}$. When $\mu = \sigma_{\#} \mathcal{H}^1$ and $\nu = \tau_{\#} \mathcal{H}^1$ measure arclength along two curves (defined e.g. by using (28) to *push forward* Lebesgue measure from the interval to the plane), then $d(\mu, \nu) = d_0(\sigma, \tau)$ is easily verified. Moreover, each optimal measure $\tilde{\gamma}$ for (1) corresponds to an optimal measure $\gamma \in \Gamma(\mu, \nu)$ for (2) obtained by pushing $\tilde{\gamma}$ forward through the map $(s, t) \rightarrow (\sigma(s), \tau(t))$ from the square $[0, 1]^2$ to $\mathbf{R}^2 \times \mathbf{R}^2$.

In the more familiar form (2) to which we now restrict ourselves, the Monge-Kantorovich problem has begun to be quite well-understood after much scrutiny through the past decade; see Brenier [4] or Gangbo & McCann [11] for references. One of the initial insights of Brenier [3], Smith & Knott [23], and Rüschemdorf & Rachev [21], was that minimizers to (2) could be characterized by the existence of a convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R} \cup \{+\infty\}$ whose *subdifferential* $\partial\psi \subset \mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ contains the support of every optimal measure $\gamma \in \Gamma(\mu, \nu)$. *Support* again refers to the smallest closed set $\text{spt } \gamma \subset \mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ of full mass, while $\partial\psi$ denotes the set of all $(\mathbf{x}, \mathbf{y}) \in \mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ satisfying

$$\psi(\mathbf{z}) - \psi(\mathbf{x}) \geq \mathbf{y} \cdot (\mathbf{z} - \mathbf{x}) \quad (\forall \mathbf{z} \in \mathbf{R}^{d+1}); \quad (3)$$

each $(\mathbf{x}, \mathbf{y}) \in \partial\psi$ corresponds to a hyperplane which touches but does not cross the graph of ψ . The convex function ψ arises as a Lagrange multiplier to the constraints μ and ν ; its existence was originally deduced from the duality theory of Kantorovich [13]. For alternative approaches consult Gangbo [10] or McCann [16].

As Brenier also realized, when μ is absolutely continuous with respect to Lebesgue, so ψ is differentiable on a set $\text{dom } \nabla\psi \subset \mathbf{R}^{d+1}$ of full μ -measure, the optimizer γ in (2) will be unique: its full mass lies on the graph $\{(\mathbf{x}, \nabla\psi(\mathbf{x})) \mid \mathbf{x} \in \text{dom } \nabla\psi\}$ of the gradient of ψ . This means μ -a.e. point \mathbf{x} must be mapped to the unique destination $\mathbf{y} = \nabla\psi(\mathbf{x})$ for transportation to be efficient; the Borel map $\mathbf{t} := \nabla\psi$ pushes μ forward to ν , and it does so optimally. The same conclusions were subsequently extended to the case where μ vanishes merely on all Lipschitz *hypersurfaces* (i.e. surfaces of codimension one) in \mathbf{R}^{d+1} [16][11]. And when μ and ν are given by bounded densities on two domains Ω and $\Lambda \subset \mathbf{R}^{d+1}$, Caffarelli [5][6] exploited connections with the Monge-Ampère equation to show that smoothness of the map $\mathbf{t} = \nabla\psi$ follows from convexity of Λ .

Nevertheless, for the present application we are interested precisely in the case in which both measures concentrate on hypersurfaces: namely curves in \mathbf{R}^2 . In this case, Fry's numerical evidence suggested that the support of the optimal measure γ might fail to concentrate on the graph of any map [9, Fig. 3.5]. This evidence is reproduced in Fig. 1(a), which provides a 'movie' illustrating the optimal matching between a pair of pentagons, depicted as an evolution from the convex to the non-convex shape. This evolution is obtained by projecting $\text{spt } \gamma \subset \mathbf{R}^2 \times \mathbf{R}^2$ onto the plane through a sequence of nine maps $\pi_\lambda(\mathbf{x}, \mathbf{y}) = (1 - \lambda)\mathbf{x} + \lambda\mathbf{y}$ which interpolate between $\pi_0(\mathbf{x}) = \mathbf{x}$ and $\pi_1(\mathbf{y}) = \mathbf{y}$. Note how certain stretches of boundary seem to disintegrate into two pieces just after time $\lambda = 0$, the two pieces being rearranged into the final curve when $\lambda = 1$. Fry claims this disintegration became finer and finer when the grid size was reduced. Although this picture may seem pathological, he noted that the presence of such phenomena varied with the initial and final shapes Ω and $\Lambda \subset \mathbf{R}^2$; for comparison see Fig. 1(b).

In trying to reconcile the theory with these numerical experiments, we have been led to natural examples in which the optimal measure γ can be proved to be unique, though its support fails to concentrate on the graph of a single map. Even when both curves are smooth and strictly convex perturbations of the same isosceles triangle an optimal map may not exist; see Fig. 2 and the development in Examples 3.12–3.13. What does remain true is that the images of μ -a.e. \mathbf{x} will be collinear: they lie on a line perpendicular to the first curve's tangent at \mathbf{x} , as one observes in Figs. 1(a) and 2. This follows from the tangential differentiability of the Kantorovich potential ψ almost everywhere along the curve: in those sections which disintegrate into multiple images it is the normal differentiability of ψ which fails. Finally, when both curves are strictly convex, the tangential derivative of ψ will be shown not only to exist but to be continuous everywhere (assuming the curves are differentiable). Since each line intersects a strictly convex curve twice at most, this means each point $\mathbf{x} \in \text{spt } \mu$ can have only two images: denoted $\mathbf{t}^+(\mathbf{x})$ and $\mathbf{t}^-(\mathbf{x}) \in \text{spt } \nu$, they correspond to the two

limits of $\nabla\psi(\mathbf{x}_k)$ obtained as $\mathbf{x}_k \rightarrow \mathbf{x}$ from outside or inside the curve, respectively. The *outer trace* $\mathbf{t}^+ : \text{spt } \mu \rightarrow \text{spt } \nu$ is shown to give a global homeomorphism between the curves, while the *inner trace* \mathbf{t}^- is continuous and continuously invertible on the closure of the set $S_2 := \{\mathbf{x} \in \text{spt } \mu \mid \mathbf{t}^+(\mathbf{x}) \neq \mathbf{t}^-(\mathbf{x})\}$ where it differs from \mathbf{t}^+ . Together, the graphs of these two maps cover the support $\text{spt } \gamma$ of the optimal measure. It is interesting that the geometry of the traces determines unambiguously how each bit of the first curve must divide itself among its potential destinations when the curve at \mathbf{x} splits into two segments of masses $d\mu_1(\mathbf{x})$ and $d\mu(\mathbf{x}) - d\mu_1(\mathbf{x})$ whose positions $(1-\lambda)\mathbf{x} + \lambda\mathbf{t}^+(\mathbf{x})$ and $(1-\lambda)\mathbf{x} + \lambda\mathbf{t}^-(\mathbf{x})$ evolve toward the second curve with $\lambda \in [0, 1]$; compare with Seethoff and Shiflett's theorem [22]. In our case, the rule determining μ_1 is that if $d\mu(\mathbf{x}) \neq d\mu_1(\mathbf{x})$ (as e.g. for $\mathbf{x} \in S_2$ in Fig. 2), then no point but \mathbf{x} of the first curve arrives at the point $\mathbf{t}^+(\mathbf{x})$ on the second curve (Lemma 2.5). This rule implies the graph of \mathbf{t}^+ is covered by $\text{spt } \gamma$. Uniqueness of the optimal measure is a more important consequence, implying in turn that γ is extremal in the convex set $\Gamma(\mu, \nu)$.

Our results extend to hypersurfaces in all dimensions, where they may also find application and — except for continuous differentiability of ψ — are not much harder to prove. Thus the general framework we choose involves measures μ and ν , each supported on the boundary of a domain $\Omega \subset \mathbf{R}^{d+1}$ which will be assumed to be bounded, Lipschitz, convex, etc. as required. Typical examples are the *surface measures* $\mu = z_1 \mathcal{H}^d|_{\partial\Omega}$ and $\nu = z_2 \mathcal{H}^d|_{\partial\Lambda}$, obtained by restricting d -dimensional Hausdorff measure to the boundaries of two domains Ω and $\Lambda \subset \mathbf{R}^{d+1}$ and normalizing with constants $z_1 := 1/\mathcal{H}^d(\partial\Omega)$ and $z_2 := 1/\mathcal{H}^d(\partial\Lambda)$. However, any Borel probability measures satisfying $\text{spt } \mu \subset \partial\Omega$ and $\text{spt } \nu \subset \partial\Lambda$ may also be considered. In this case one interprets the support of μ as a set covered to various depths using a unit volume of paint, and the Wasserstein distance $d(\mu, \nu)$ as measuring the minimum work (or elastic energy) required to rearrange the paint so as to cover $\text{spt } \nu$ to a depth specified by ν . This metric takes into account not only the geometry of the supports but also the weighting of points, thus combining useful aspects of the Hausdorff distance and the area (or 'template') metric [18][9].

The first section of the paper is devoted mainly to recovering relevant background theory concerning the Monge-Kantorovich problem (2) and convex functions [1][25], before using it to deduce that the images of μ -a.e. \mathbf{x} are collinear provided μ concentrates no mass on any submanifold of codimension *two* in \mathbf{R}^{d+1} (Corollary 1.5). For the images to be perpendicular to the tangent space at $\mathbf{x} \in \partial\Omega$, we assume μ is absolutely continuous with respect to surface measure on the *Lipschitz domain* $\Omega \subset \mathbf{R}^{d+1}$, Lipschitz meaning that $\partial\Omega$ can be parameterized as the graph of a Lipschitz function in suitable coordinates near any point, and therefore has a well-defined tangent space at \mathcal{H}^d -a.e. \mathbf{x} . This condition is verified explicitly for convex domains in (68). We impose the further geometrical restriction on the target that *any straight line intersects $\partial\Lambda$ at most k times* when we want to conclude that almost every $\mathbf{x} \in \partial\Omega$ has finitely many images in Corollary 1.7.

Section 2 establishes uniqueness of the minimizer γ for (2) under the assumption

that $\Omega \subset \mathbf{R}^{d+1}$ is strictly convex, and μ absolutely continuous with respect to its surface measure $\mathcal{H}^d|_{\partial\Omega}$. The target ν can be any Borel probability measure on \mathbf{R}^{d+1} . A simple example with μ and ν measuring arclength along two rectangles shows uniqueness can fail without strict convexity: Fig. 3 and Example 2.1.

The third section marks the beginning of a regularity theory: it shows tangential differentiability along $\partial\Omega$ of the convex function ψ whose subdifferential $\partial\psi \supset \text{spt } \gamma$ contains the support of the optimal measure. The proof is carried out for boundary measures on strictly convex domains in the plane, leaving the extension to higher dimensions for Section 4. We assume that μ (and ν) are *free from atoms* — meaning $\mu[\{\mathbf{x}\}] = 0$ for every $\mathbf{x} \in \mathbf{R}^{d+1}$ — while $\partial\Lambda = \text{spt } \nu$ and $\partial\Omega = \text{spt } \mu$. The key idea is *monotonicity* of $\partial\psi$, which follows from (3) and the definition:

Definition 0.1 (Monotonicity) *The set $M \subset \mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ is monotone if all $(\mathbf{x}_1, \mathbf{y}_1) \in M$ and $(\mathbf{x}_2, \mathbf{y}_2) \in M$ satisfy*

$$(\mathbf{y}_1 - \mathbf{y}_2) \cdot (\mathbf{x}_1 - \mathbf{x}_2) \geq 0. \quad (4)$$

A monotone relation is generally viewed as a multivalued mapping on \mathbf{R}^{d+1} which rotates pairs of points by no more than 90° ; in keeping with this tradition we use the notation $M(\mathbf{x}) := \{\mathbf{y} \mid (\mathbf{x}, \mathbf{y}) \in M\}$ to denote the images of \mathbf{x} . We exploit the fact that the subdifferential $\partial\psi \subset \mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ is also closed (because our convex functions are continuous) to derive continuity of the maps $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ and $\mathbf{t}^-|_{\overline{S_2}}$ whose graphs cover $\text{spt } \gamma$. Both maps are said to be *monotone*, meaning their graphs are monotone subsets of $\mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$. They are continuously invertible because of symmetry in $\mu \leftrightarrow \nu$. The results of this section and the next are summarized by Theorem 3.8, which asserts that the optimal measure $\gamma \in \Gamma(\mu, \nu)$ satisfies

$$\{(\mathbf{x}, \mathbf{t}^+(\mathbf{x}))\}_{\mathbf{x} \in \text{spt } \mu} \subset \text{spt } \gamma \subset \{(\mathbf{x}, \mathbf{t}^+(\mathbf{x}))\}_{\mathbf{x} \in \partial\Omega} \cup \{(\mathbf{x}, \mathbf{t}^-(\mathbf{x}))\}_{\mathbf{x} \in S_2}, \quad (5)$$

where \mathbf{t}^+ and \mathbf{t}^- are the traces of $\nabla\psi$ on the boundary of Ω , and $S_2 := \partial\Omega \setminus \text{dom } \nabla\psi$. As we know from the examples studied at the end of Section 3, the result (5) is optimal in the sense that the first containment cannot be replaced by an equality: the full mass of γ is not contained by any graph. On the other hand, we cannot say whether the second containment generally reduces to equality.

The final section contains the key estimate used to prove Theorem 3.8 in dimensions $d > 1$. To establish the estimate (and theorem), we assume that the surface densities of μ and ν are uniformly bounded:

$$\mu < \frac{1}{\epsilon} \mathcal{H}^d|_{\partial\Omega} \quad \text{and} \quad \nu > \epsilon \mathcal{H}^d|_{\partial\Lambda} \quad (6)$$

while

$$\mu > \epsilon \mathcal{H}^d|_{\partial\Omega} \quad \text{and} \quad \nu < \frac{1}{\epsilon} \mathcal{H}^d|_{\partial\Lambda} \quad (7)$$

for some $\epsilon > 0$, where the domains Ω and $\Lambda \subset \mathbf{R}^{d+1}$ both satisfy the definition 0.2 below of *uniform convexity*. Actually, continuity of the maps \mathbf{t}^+ and $\mathbf{t}^-|_{\overline{S_2}}$ in (5)

follows just from (6), uniform convexity of Ω , and strict convexity of $\partial\Lambda \supset \text{spt } \nu$; it is only in proving $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ and $\mathbf{t}^- : \overline{S}_2 \subset \partial\Omega \rightarrow \overline{T}_2 \subset \partial\Lambda$ are continuously bijective that (7) and uniform convexity of Λ are employed. The bounds (6–7) are stronger than we required in the plane: the difference between \mathbf{R}^2 and \mathbf{R}^{d+1} in the proof is that any connected subset of $\partial\Omega \subset \mathbf{R}^2$ having more than one point contains a relatively open arc, and therefore has positive ν -measure just from $\partial\Lambda = \text{spt } \nu$ without using area estimates to appeal to (6).

Before concluding this discussion, we introduce some terminology pertaining to convex sets $\Omega \subset \mathbf{R}^{d+1}$. We say that $\mathbf{n} \in \mathbf{R}^{d+1}$ is a *generalized outward normal* to $\partial\Omega$ at $\mathbf{x} \in \partial\Omega$ if $0 \geq \mathbf{n} \cdot (\mathbf{z} - \mathbf{x})$ for all $\mathbf{z} \in \overline{\Omega}$. Thus the generalized outward normal coincides with the classical normal at points where $\partial\Omega$ is differentiable. Abusing terminology, we drop the prefix “generalized”. Observe that convexity of Ω yields a non-zero normal at every point of $\partial\Omega$. The set of *unit outward normals* at $\mathbf{x} \in \partial\Omega$ is therefore non-empty, and denoted $N_\Omega(\mathbf{x})$. Then $N_\Omega(\mathbf{x})$ consists of a single element, denoted $\mathbf{n}_\Omega(\mathbf{x})$, if and only if \mathbf{x} is a point of differentiability for $\partial\Omega$.

Definition 0.2 (Strict vs. Uniform Convexity) *A convex domain $\Omega \subset \mathbf{R}^{d+1}$ is strictly convex if its boundary $\partial\Omega$ contains no line segments. It is uniformly convex if for some $R > 0$, all $\mathbf{x} \in \partial\Omega$ and all unit normals $\hat{\mathbf{n}} \in N_\Omega(\mathbf{x})$, the entire domain $\Omega \subset \mathbf{B}(\mathbf{x} - R\hat{\mathbf{n}}; R)$ lies inside the ball of radius R having outward normal $\hat{\mathbf{n}}$ at \mathbf{x} .*

Observe that uniformly convex domains will necessarily be bounded and strictly convex: any boundary segment would lie outside of a sphere through its midpoint. Another characterization of strict convexity is that each supporting hyperplane touches $\partial\Omega$ at a single point: thus $\mathbf{x} = \mathbf{x}_o$ whenever $\mathbf{x} \in \overline{\Omega}$, $\mathbf{x}_o \in \partial\Omega$ and $\mathbf{n}_o \in N_\Omega(\mathbf{x}_o)$ satisfy

$$\mathbf{n}_o \cdot (\mathbf{x} - \mathbf{x}_o) \geq 0. \quad (8)$$

Further properties of convex sets are summarized in Appendix A.

A final feature of Fry’s distance (1) worth mentioning is its behaviour with respect to symmetries, which may prove useful in applications where one wants to disregard e.g. translations and/or dilations of the known prototype. The identity

$$d_\beta^2(\sigma, \tau + \mathbf{y}) = d_\beta^2(\sigma, \tau) + |\mathbf{y}|^2 - 2\mathbf{y} \cdot \int_0^1 [\sigma(s) - \tau(s)] ds \quad (9)$$

shows $d_\beta^2(\sigma, \tau + \mathbf{y})$ to be minimized when the center of masses of the two curves coincide with, say, the origin. For $\beta = 0$ the best dilation $\lambda \geq 0$ of the standard leaf is easily deduced by computing $d_0(\sigma, \tau)$ and using

$$d_\beta^2(\sigma, \lambda\tau) = (1 - \lambda) \left[\int_0^1 |\sigma(s)|^2 ds - \lambda \int_0^1 |\tau(t)|^2 dt \right] + \lambda d_{\beta/\lambda}^2(\sigma, \tau). \quad (10)$$

It is our pleasure to conclude by thanking David Mumford for bringing this problem to our attention, and for his ongoing inspiration and advice. We are also very grateful to Dev Sinha for providing fruitful conversations, Joan Mann and Karin Smith for their help with the figures, and David Fry for his permission to reproduce [9, Figs. 3.4 and 3.5].

Notation

For the convenience of the reader we collect together some of the notation introduced throughout the text.

- \mathcal{H}^n denotes n -dimensional Hausdorff measure on the Borel σ -algebra of sets.
- The closed segment joining \mathbf{x} to $\mathbf{y} \in \mathbf{R}^{d+1}$ is denoted by

$$[\mathbf{x}, \mathbf{y}] := \{(1 - \lambda)\mathbf{x} + \lambda\mathbf{y} \mid \lambda \in [0, 1]\}. \quad (11)$$

- If $\Omega \subset \mathbf{R}^{d+1}$ then $\overline{\Omega}$ denotes the closure, $\Omega^c := \mathbf{R}^{d+1} \setminus \Omega$ the complement, and $\text{conv}(\Omega)$ the convex hull of Ω , meaning the smallest convex set containing Ω .

- If Ω is a convex set and $\mathbf{x} \in \partial\Omega$, then $N_\Omega(\mathbf{x})$ stands for the set of all outward unit normals to $\partial\Omega$ at \mathbf{x} — a closed, geodesically convex subset of the unit sphere. When $N_\Omega(\mathbf{x})$ contains only one element, we denote that unit vector by $\mathbf{n}_\Omega(\mathbf{x})$.

- If $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R} \cup \{+\infty\}$ is not identically $+\infty$, the Legendre-Fenchel transform of ψ is the convex, lower semicontinuous function $\psi^* : \mathbf{R}^{d+1} \rightarrow \mathbf{R} \cup \{+\infty\}$ defined by

$$\psi^*(\mathbf{y}) := \sup_{\mathbf{x} \in \mathbf{R}^{d+1}} \{\mathbf{x} \cdot \mathbf{y} - \psi(\mathbf{x})\}. \quad (12)$$

Hence ψ^{**} is the greatest lower semicontinuous convex function dominated by ψ .

- The set where ψ is differentiable is denoted by $\text{dom } \nabla\psi \subset \mathbf{R}^{d+1}$.
- The subdifferential of a convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R} \cup \{+\infty\}$ is the set $\partial\psi \subset \mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ consisting of all (\mathbf{x}, \mathbf{y}) satisfying

$$\psi(\mathbf{z}) - \psi(\mathbf{x}) \geq \mathbf{y} \cdot (\mathbf{z} - \mathbf{x}), \quad (\forall \mathbf{z} \in \mathbf{R}^{d+1}).$$

If $(\mathbf{x}, \mathbf{y}) \in \partial\psi$ we may also write $\mathbf{y} \in \partial\psi(\mathbf{x})$. Recall $\mathbf{x} \in \partial\psi^*(\mathbf{y})$ whenever $\mathbf{y} \in \partial\psi(\mathbf{x})$, while the converse also holds true if ψ is convex lower semicontinuous. In that case $\partial\psi$ is a closed set. Note that $\mathbf{x} \in \text{dom } \nabla\psi$ precisely when $\partial\psi(\mathbf{x})$ consists of a single element, namely $\nabla\psi(\mathbf{x})$. In general, the set $\partial\psi(\mathbf{x}) \subset \mathbf{R}^{d+1}$ is closed and convex.

- If μ is a Borel measure on \mathbf{R}^{d+1} , we denote by $\text{spt } \mu$ the support of μ , which refers to the smallest closed set K such that $\mu[\mathbf{R}^{d+1} \setminus K] = 0$. For $S \subset \mathbf{R}^{d+1}$ Borel we denote by $\mu|_S$ the restriction of μ to S , defined by $\mu|_S[B] := \mu[B \cap S]$ for $B \subset \mathbf{R}^{d+1}$.

- Similarly, $\psi|_S$ denotes the restriction of the function ψ to $S \subset \mathbf{R}^{d+1}$.

- If μ is a Borel measure on \mathbf{R}^{d+1} and $\mathbf{t} : \mathbf{R}^{d+1} \rightarrow \mathbf{R}^n$ is a Borel map we define $\mathbf{t}_\# \mu$ to be the Borel measure on \mathbf{R}^n given by $\mathbf{t}_\# \mu[B] := \mu[\mathbf{t}^{-1}(B)]$ for $B \subset \mathbf{R}^n$. We call $\mathbf{t}_\# \mu$ the push-forward of μ by \mathbf{t} .

- We denote the identity map $\mathbf{id}(\mathbf{x}) = \mathbf{x}$ by \mathbf{id} .

- If μ and ν are two Borel measures on \mathbf{R}^{d+1} , $\Gamma(\mu, \nu)$ stands for the set of all Borel measures on $\mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ having μ and ν as their marginals: $\mu[B] = \gamma[B \times \mathbf{R}^{d+1}]$ and $\gamma[\mathbf{R}^{d+1} \times B] = \nu[B]$ for all Borel sets $B \subset \mathbf{R}^{d+1}$.

1 Background Theory, General Curves and Surfaces

The first goal of the present section is to recall a central result from the theory developed for the special form (2) of the Monge-Kantorovich problem by Brenier [3][4], Smith and Knott [23], Rüschemdorf and Rachev [21], and others. It characterizes the optimal measures $\gamma \in \Gamma(\mu, \nu)$ via the existence of a convex function ψ whose subdifferential contains their support. It is the starting point for our study. Although this function will not be unique, we can (and do) normalize ψ so that $\nabla\psi(\mathbf{x})$ belongs to $\text{spt } \nu \subset \mathbf{R}^{d+1}$. Focusing on compactly supported measures thus ensures our convex functions satisfy Lipschitz estimates throughout \mathbf{R}^{d+1} . At the end of the section we shall exploit the theorem to conclude that if μ is absolutely continuous with respect to surface measure $\mathcal{H}^d|_{\partial\Omega}$ on a Lipschitz domain $\Omega \subset \mathbf{R}^{d+1}$, then the images $\partial\psi(\mathbf{x}) \cap \text{spt } \nu$ almost surely lie on a line parallel to the normal $\mathbf{n}_\Omega(\mathbf{x})$. Although the first theorem is well-known and can be recovered from e.g. [11, Corollaries 2.4 and 2.8], we conveniently sketch a duality based proof following the argument of Brenier.

Theorem 1.1 (Optimality Criterion) *Fix Borel probability measures μ and ν of bounded support in \mathbf{R}^{d+1} . Then there is a convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}$ whose subdifferential $\partial\psi$ includes the support $\text{spt } \gamma \subset \mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$ of all minimizers $\gamma \in \Gamma(\mu, \nu)$ for (2). Moreover, at each point \mathbf{x} where ψ is differentiable, $\nabla\psi(\mathbf{x}) \in \text{spt } \nu$.*

Proof: 1. The factorization $|\mathbf{x} - \mathbf{y}|^2 = |\mathbf{x}|^2 - 2\mathbf{x} \cdot \mathbf{y} + |\mathbf{y}|^2$ implies that computing our Wasserstein distance (2) is equivalent to maximizing correlations:

$$\sup_{\gamma \in \Gamma(\mu, \nu)} \int_{\mathbf{R}^{d+1} \times \mathbf{R}^{d+1}} \mathbf{x} \cdot \mathbf{y} \, d\gamma(\mathbf{x}, \mathbf{y}). \quad (13)$$

Indeed γ_o is a maximizer for (13) if and only if it minimizes (2). Now the Kantorovich duality principle asserts the infinite dimensional linear program (13) is dual to the minimization

$$\inf_{(\psi, \phi) \in \mathcal{A}} \int_{\mathbf{R}^{d+1}} \psi(\mathbf{x}) \, d\mu(\mathbf{x}) + \int_{\mathbf{R}^{d+1}} \phi(\mathbf{y}) \, d\nu(\mathbf{y}), \quad (14)$$

in the sense that the supremum in (13) and infimum in (14) coincide; see e.g. [15, 19, 21, 4]. Here \mathcal{A} denotes the set of all pairs of functions (ψ, ϕ) from \mathbf{R}^{d+1} into $\mathbf{R} \cup \{+\infty\}$ which are lower semicontinuous and satisfy

$$\psi(\mathbf{x}) + \phi(\mathbf{y}) \geq \mathbf{x} \cdot \mathbf{y} \quad (15)$$

for all $\mathbf{x}, \mathbf{y} \in \mathbf{R}^{d+1}$. Another way of stating the duality is to say that $(\psi_o, \phi_o) \in \mathcal{A}$ and $\gamma_o \in \Gamma(\mu, \nu)$ solve

$$\int_{\mathbf{R}^{d+1} \times \mathbf{R}^{d+1}} (\psi_o(\mathbf{x}) + \phi_o(\mathbf{y}) - \mathbf{x} \cdot \mathbf{y}) \, d\gamma_o(\mathbf{x}, \mathbf{y}) = 0 \quad (16)$$

if and only if (ψ_o, ϕ_o) and γ_o optimize (14) and (2) respectively. Note that the integrand of (16) is non-negative whenever $(\psi_o, \phi_o) \in \mathcal{A}$.

2. Since the measures μ and ν have bounded support, well-known continuity and compactness principles show that (13) admits a maximizer γ_o and (14) admits a minimizer $(\psi_o, \phi_o) \in \mathcal{A}$; see e.g. Kellerer [15]. Without losing generality we may assume $\phi_o(\mathbf{y}) = +\infty$ outside $\text{spt } \nu$, and define its Legendre-Fenchel transform ϕ_o^* by

$$\phi_o^*(\mathbf{x}) := \sup_{\mathbf{y} \in \mathbf{R}^{d+1}} \{\mathbf{x} \cdot \mathbf{y} - \phi_o(\mathbf{y})\}. \quad (17)$$

Setting $\psi := \phi_o^*$, the pair $(\psi, \phi_o) \in \mathcal{A}$ continues to be optimal since $\phi_o^* \leq \psi_o$ follows from (15) and (17). Since $\mathbf{y} \notin \text{spt } \mu$ does not contribute to (17), ψ is the convex function given by

$$\psi(\mathbf{x}) = \sup_{\mathbf{y} \in \text{spt } \nu} \{\mathbf{x} \cdot \mathbf{y} - \phi_o(\mathbf{y})\}. \quad (18)$$

3. *Claim:* The compactness of $\text{spt } \nu$ and lower semicontinuity of ϕ_o in (18) imply $\nabla\psi(\mathbf{x}_o) \in \text{spt } \nu$ at each point \mathbf{x}_o where ψ is differentiable.

Proof of Claim: Fix any $\mathbf{x}_o \in \mathbf{R}^{d+1}$ where ψ happens to be differentiable. Since the maximum value (18) is attained at some $\mathbf{y}_o \in \text{spt } \nu$, the non-negative function

$$\psi(\mathbf{x}) + \phi_o(\mathbf{y}_o) - \mathbf{x} \cdot \mathbf{y}_o \geq 0 \quad (19)$$

vanishes at $\mathbf{x} = \mathbf{x}_o$. It is therefore minimized, so its derivative $\nabla\psi(\mathbf{x}_o) - \mathbf{y}_o$ (with respect to $\mathbf{x} \in \mathbf{R}^{d+1}$) must vanish. This proves the claim: $\nabla\psi(\mathbf{x}_o) = \mathbf{y}_o \in \text{spt } \nu$.

4. Recall that (16) is satisfied by our functions (ψ, ϕ_o) and any maximizer γ_o for (13). The lower semicontinuous integrand is positive outside of a closed set where it vanishes (and which therefore contains $\text{spt } \gamma_o$). Now let $(\mathbf{x}_o, \mathbf{y}_o) \in \text{spt } \gamma_o$, and observe that (19) is minimized at $\mathbf{x} = \mathbf{x}_o$. Thus $\mathbf{0} \in \partial\psi(\mathbf{x}_o) - \mathbf{y}_o$ and hence $\mathbf{y}_o \in \partial\psi(\mathbf{x}_o)$ which concludes the proof that $\text{spt } \gamma_o \subset \partial\psi$. \square

Remark 1.2 (Symmetry) *Similarly, there is a convex function $\phi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}$ for which $\mathbf{x} \in \partial\phi(\mathbf{y})$ whenever (\mathbf{x}, \mathbf{y}) belongs to the support of a minimizer γ for (2). Indeed, we may take*

$$\phi(\mathbf{y}) := \sup_{\mathbf{x} \in \text{spt } \mu} \mathbf{x} \cdot \mathbf{y} - \psi(\mathbf{x}), \quad (20)$$

in which case $\nabla\phi \in \text{spt } \mu$ holds at each point of differentiability, while $\partial\phi(\mathbf{y}) \subset \partial\psi^(\mathbf{y})$ and $\partial\psi(\mathbf{x}) \subset \partial\phi^*(\mathbf{x})$ hold for $\mathbf{y} \in \text{spt } \nu$ and $\mathbf{x} \in \text{spt } \mu$.*

Remark 1.3 (Converse) *Conversely, if $\gamma \in \Gamma(\mu, \nu)$ is supported in the subdifferential $\partial\psi$ of any convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}$, then γ minimizes (2).*

Proof of Remarks 1.2 and 1.3: We shall prove the second remark first:

1. Suppose $\gamma_o \in \Gamma(\mu, \nu)$ is supported on the subdifferential of any convex function ψ . Then for fixed $(\mathbf{x}_o, \mathbf{y}_o) \in \text{spt } \gamma_o$, the concave function $\mathbf{x} \cdot \mathbf{y}_o - \psi(\mathbf{x})$ must attain its maximum value at $\mathbf{x} = \mathbf{x}_o$ since $\mathbf{0} \in \mathbf{y}_o - \partial\psi(\mathbf{x}_o)$. This maximum value $\mathbf{x}_o \cdot \mathbf{y}_o - \psi(\mathbf{x}_o)$ is more commonly denoted by $\psi^*(\mathbf{y}_o)$ as in (12). Setting $(\psi_o, \phi_o) = (\psi, \psi^*)$ we have

just shown that the integrand in (16) vanishes on $\text{spt } \gamma$. Thus γ_o maximizes (13) — establishing Remark 1.3 — while $(\psi, \psi^*) \in \mathcal{A}$ minimize the dual problem (14).

2. To prove the remaining remark, take ψ from Theorem 1.1 and let γ_o minimize (2). Then $\text{spt } \gamma_o \subset \partial\psi$ so $(\psi_o, \phi_o) := (\psi, \psi^*) \in \mathcal{A}$ minimizes (14) as we just verified. At this point we can repeat steps 2–4 of the proof of the theorem exchanging $\psi_o \leftrightarrow \phi_o$, $\mathbf{x} \leftrightarrow \mathbf{y}$ and $\mu \leftrightarrow \nu$ to conclude that $\mathbf{x}_o \in \partial\phi(\mathbf{y}_o)$ whenever $(\mathbf{x}_o, \mathbf{y}_o) \in \text{spt } \gamma$. Note in particular that (18) becomes (20).

3. Comparing (12) with (20) yields $\psi^* \geq \phi$; we claim equality holds on $\text{spt } \nu$. Indeed, since our measures are compactly supported, for each $\mathbf{y}_o \in \text{spt } \nu$ there is some $\mathbf{x}_o \in \text{spt } \mu$ such that $(\mathbf{x}_o, \mathbf{y}_o) \in \text{spt } \gamma_o \subset \partial\psi$. This implies $\psi^*(\mathbf{y}_o) = \mathbf{x}_o \cdot \mathbf{y}_o - \psi(\mathbf{x}_o)$ as in 1 above, after which (20) yields $\psi^*(\mathbf{y}_o) \leq \phi(\mathbf{y}_o)$ to establish the equality claim. Keeping $\mathbf{y}_o \in \text{spt } \nu$, from 3 we deduce that any $\mathbf{x} \in \partial\phi(\mathbf{y}_o)$ also belongs to $\partial\psi^*(\mathbf{y}_o)$ because

$$\psi^*(\mathbf{y}) - \psi^*(\mathbf{y}_o) \geq \phi(\mathbf{y}) - \phi(\mathbf{y}_o) \geq \mathbf{x} \cdot (\mathbf{y} - \mathbf{y}_o)$$

holds for all $\mathbf{y} \in \mathbf{R}^d$. Thus $\partial\phi(\mathbf{y}_o) \subset \partial\psi^*(\mathbf{y}_o)$. For $\mathbf{x}_o \in \text{spt } \mu$, the same computation yields $\partial\psi(\mathbf{x}_o) \subset \partial\phi^*(\mathbf{x}_o)$ after verifying:

4. Claim: $\phi^* \geq \psi$ holds on \mathbf{R}^{d+1} , with equality on $\text{spt } \mu$.

Proof of Claim: Legendre transforming inequality 3 yields $\phi^* \geq \psi^{**} = \psi$, where the identity follows from convexity and continuity of ψ in Theorem 1.1. It remains to check equality on $\text{spt } \mu$. From (20) one has $\phi := \psi_\infty^*$, where

$$\psi_\infty(\mathbf{x}) := \begin{cases} \psi(\mathbf{x}) & \text{if } \mathbf{x} \in \text{spt } \mu \\ +\infty & \text{otherwise.} \end{cases} \quad (21)$$

Thus $\psi_\infty \geq \psi_\infty^{**} = \phi^* \geq \psi$, since Legendre transforming twice yields the greatest lower semi-continuous convex function dominated by ψ_∞ . The definition (21) of ψ_∞ forces equalities throughout $\text{spt } \mu$. \square

Corollary 1.4 (Extremal Images) *For the convex function ψ of the theorem and all $\mathbf{x} \in \mathbf{R}^{d+1}$, the extreme points of $\partial\psi(\mathbf{x})$ will lie on $\text{spt } \nu$.*

Proof: Rockafellar [20, §25.6] asserts that $\partial\psi(\mathbf{x})$ is the closed convex hull $\overline{\text{conv}S(\mathbf{x})}$ of the set $S(\mathbf{x})$ of all limits of sequences of the form $\nabla\psi(\mathbf{x}_1), \nabla\psi(\mathbf{x}_2), \dots$, such that $\mathbf{x}_k \in \text{dom } \nabla\psi$ converges to \mathbf{x} . Theorem 1.1 asserts that $S(\mathbf{x}) \subset \text{spt } \nu$ and is bounded. Thus $\partial\psi(\mathbf{x}) = \overline{\text{conv}S(\mathbf{x})}$ and its extreme points lie in $\overline{S(\mathbf{x})} \subset \text{spt } \nu$ as desired. \square

The next corollary states that if μ vanishes on all submanifolds of codimension two, the optimal images of μ -a.e. \mathbf{x} will be collinear. Though not needed here, should μ merely vanish on the Lipschitz submanifolds of codimension $k+1$ in \mathbf{R}^{d+1} , the same argument yields the images of \mathbf{x} almost surely lie in a k -dimensional affine subspace.

Corollary 1.5 (Collinearity) *If the measure μ of the theorem vanishes on each Lipschitz submanifold of dimension $d - 1$ in \mathbf{R}^{d+1} , the images $\partial\psi(\mathbf{x}) \cap \text{spt } \nu$ of \mathbf{x} will be collinear for μ -a.e. $\mathbf{x} \in \mathbf{R}^{d+1}$.*

Proof: Take the convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}$ of the theorem, and let $Z \subset \mathbf{R}^{d+1}$ denote the set of all \mathbf{x} such that the convex set $\partial\psi(\mathbf{x})$ has dimension greater than or equal to two. Then Zajícěk [25] (and also Alberti [1]) show the Borel set Z can be covered by countably many Lipschitz submanifolds of dimension $d - 1$. Thus $\mu[Z] = 0$. For any $\mathbf{x} \in \mathbf{R}^{d+1} \setminus Z$, which is to say μ -almost every \mathbf{x} , all of the points in $\partial\psi(\mathbf{x})$ must lie on a single line: if even three points failed to be collinear their convex hull would form a two-dimensional triangle inside $\partial\psi(\mathbf{x})$. Since ψ is locally Lipschitz, $\partial\psi(\mathbf{x})$ is a bounded convex set in one dimension: at most a line segment. \square

Since we are primarily interested in boundary measures on a domain $\Omega \subset \mathbf{R}^{d+1}$, it is useful to refine this corollary by relating the orientation of the segment to the geometry of $\partial\Omega$. This is accomplished in the following lemma, where we use $[\mathbf{y}, \mathbf{z}]$ to denote the line segment $[\mathbf{y}, \mathbf{z}] := \{\lambda\mathbf{y} + (1 - \lambda)\mathbf{z} \mid \lambda \in [0, 1]\}$ in \mathbf{R}^{d+1} .

Lemma 1.6 (Tangential Differentiability a.e.) *Fix a Lipschitz domain $\Omega \subset \mathbf{R}^{d+1}$ and convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}$. At \mathcal{H}^d -a.e. boundary point \mathbf{x} , the surface $\partial\Omega$ is differentiable and ψ is tangentially differentiable; i.e., $\partial\psi(\mathbf{x}) = [\mathbf{t}^+(\mathbf{x}), \mathbf{t}^-(\mathbf{x})]$ parallels the outward normal: $\mathbf{t}^+(\mathbf{x}) - \mathbf{t}^-(\mathbf{x}) = \lambda(\mathbf{x})\mathbf{n}_\Omega(\mathbf{x})$ with $\lambda(\mathbf{x}) \geq 0$.*

Proof: Being convex, ψ is locally Lipschitz throughout \mathbf{R}^{d+1} . Thus its restriction $\psi|_{\partial\Omega}$ is locally Lipschitz with respect to the induced metric on the Lipschitz submanifold $\partial\Omega \subset \mathbf{R}^{d+1}$. Rademacher's theorem assures differentiability of $\psi|_{\partial\Omega}$ and $\partial\Omega$ except on a Borel set $Z \subset \partial\Omega$ of zero area $\mathcal{H}^d[Z] = 0$. At $\mathbf{x} \in \partial\Omega \setminus Z$ the (Lipschitz) function $\psi|_{\mathbf{R}^{d+1}}$ is therefore tangentially differentiable along $\partial\Omega$; we claim nothing about normal derivatives. Now the derivative of ψ in direction $\mathbf{p} \in \mathbf{R}^{d+1}$ is given [20, §23.4] by

$$\nabla_{\mathbf{p}}\psi(\mathbf{x}) := \lim_{\lambda \rightarrow 0^+} \lambda^{-1}[\psi(\mathbf{x} + \lambda\mathbf{p}) - \psi(\mathbf{x})] = \sup_{\mathbf{y} \in \partial\psi(\mathbf{x})} \mathbf{y} \cdot \mathbf{p}. \quad (22)$$

Tangential differentiability means $\nabla_{\mathbf{p}}\psi(\mathbf{x}) = -\nabla_{-\mathbf{p}}\psi(\mathbf{x})$ provided $\mathbf{p} \cdot \mathbf{n}_\Omega(\mathbf{x}) = 0$. For each such \mathbf{p} , (22) yields $\partial\psi(\mathbf{x}) \subset H_{\mathbf{p}} := \{\mathbf{y} \in \mathbf{R}^{d+1} \mid \mathbf{y} \cdot \mathbf{p} = \nabla_{\mathbf{p}}\psi(\mathbf{x})\}$. The intersection of these affine subspaces $H_{\mathbf{p}} \supset \partial\psi(\mathbf{x})$ is a line parallel to $\mathbf{n}_\Omega(\mathbf{x})$. Now $\partial\psi(\mathbf{x})$ is a closed convex subset of this line, and bounded by the Lipschitz constant of ψ near \mathbf{x} . Either it must degenerate to a single point $\partial\psi(\mathbf{x}) = \{\nabla\psi(\mathbf{x})\} = \{\mathbf{t}^+(\mathbf{x})\} = \{\mathbf{t}^-(\mathbf{x})\}$, or else to a closed segment with endpoints denoted by $\mathbf{t}^+(\mathbf{x})$ and $\mathbf{t}^-(\mathbf{x})$. Here the superscripts \pm are assigned so that $\mathbf{t}^-(\mathbf{x}) - \mathbf{t}^+(\mathbf{x})$ is anti-parallel to $\mathbf{n}_\Omega(\mathbf{x})$. Setting $\lambda(\mathbf{x}) := |\mathbf{t}^+(\mathbf{x}) - \mathbf{t}^-(\mathbf{x})|$ concludes the proof of the lemma. \square

The map $\mathbf{t}^+ : \partial\Omega \rightarrow \mathbf{R}^{d+1}$ of the lemma corresponds to the gradient of ψ as computed from outside the domain Ω , while $\mathbf{t}^- : \partial\Omega \rightarrow \mathbf{R}^{d+1}$ gives the gradient of ψ as computed from inside the domain Ω . Where ψ is tangentially differentiable

along $\partial\Omega$, these two gradients are well-defined and coincide in all but their normal components. Should the normal components also agree, $\partial\psi(\mathbf{x}) = \{\nabla\psi(\mathbf{x})\}$ and $\nabla\psi(\mathbf{x}) = \mathbf{t}^+(\mathbf{x}) = \mathbf{t}^-(\mathbf{x})$. This point of view makes it easy to show that \mathbf{t}^+ and \mathbf{t}^- are Borel maps on $\partial\Omega$ as claimed in the following section. They shall be referred to as the *outer trace* (\mathbf{t}^+) and *inner trace* (\mathbf{t}^-) of $\nabla\psi$ on $\partial\Omega$.

Now suppose μ and ν are boundary measures on a pair of domains Ω and Λ . Since we are interested in obtaining maps $\mathbf{t} : \partial\Omega \rightarrow \partial\Lambda$, the collinearity property motivates the following condition on Λ :

$$\text{any straight line intersects } \partial\Lambda \text{ in at most } k \text{ points.} \tag{23}$$

Strictly convex domains $\Lambda \subset \mathbf{R}^{d+1}$ satisfy (23) with $k = 2$.

Corollary 1.7 (Images a.s. Parallel Normal) *Let Ω and $\Lambda \subset \mathbf{R}^{d+1}$ be bounded domains, with probability measures μ on $\partial\Omega$ and ν on $\partial\Lambda$. Assume Ω is Lipschitz, μ is absolutely continuous with respect to surface measure $\mathcal{H}^d|_{\partial\Omega}$, and Λ satisfies (23). There exist k maps $\mathbf{t}_1, \dots, \mathbf{t}_k : \partial\Omega \rightarrow \partial\Lambda$, whose graphs $\bigcup_1^k \{(\mathbf{x}, \mathbf{t}_i(\mathbf{x})) \mid \mathbf{x} \in \partial\Omega\} \subset \partial\psi$ carry the full mass of all minimizers γ to (2). At μ -a.e. \mathbf{x} , $\partial\Omega$ is differentiable and has $\mathbf{t}_i(\mathbf{x}) - \mathbf{t}_{i+1}(\mathbf{x})$ as a (possibly degenerate) outward normal, $i = 1, \dots, k - 1$.*

Proof: Let ψ denote the convex function from Theorem 1.1. Lemma 1.6 yields a Borel set $Z \subset \partial\Omega$ of \mathcal{H}^d and hence μ measure zero, such that $\partial\Omega$ is differentiable at each $\mathbf{x} \notin Z$ and $\partial\psi(\mathbf{x}) = [\mathbf{t}^+(\mathbf{x}), \mathbf{t}^-(\mathbf{x})]$ is a segment parallel to $\mathbf{n}_\Omega(\mathbf{x})$. This segment intersects $\partial\Lambda$ in at most k places by assumption (23), and at least 1 by Corollary 1.4. We label the intersections $\mathbf{t}_1(\mathbf{x}), \dots, \mathbf{t}_k(\mathbf{x})$, ordered so that $[\mathbf{t}_i(\mathbf{x}) - \mathbf{t}_{i+1}(\mathbf{x})] \cdot \mathbf{n}_\Omega(\mathbf{x}) \geq 0$. Thus $\mathbf{t}_i(\mathbf{x}) - \mathbf{t}_{i+1}(\mathbf{x})$ parallels $\mathbf{n}_\Omega(\mathbf{x})$ and is a (possibly degenerate) outward normal at $\mathbf{x} \in \partial\Omega$.

Finally, suppose $\gamma \in \Gamma(\mu, \nu)$ minimizes (2) so $\text{spt } \gamma \subset \partial\psi$ from Theorem 1.1. Also $\text{spt } \gamma \subset \partial\Omega \times \partial\Lambda$. The preceding paragraph shows that $\partial\psi \cap [\partial\Omega \times \partial\Lambda] \subset \bigcup_1^k \{(\mathbf{x}, \mathbf{t}_i(\mathbf{x})) \mid \mathbf{x} \in \partial\Omega \setminus Z\} \cup [Z \times \partial\Lambda]$. Since $\gamma[Z \times \partial\Lambda] = \mu[Z] = 0$, we conclude that the full mass of γ is carried by the graphs of the \mathbf{t}_i . \square

2 Unique Solutions for Strictly Convex Shapes

When μ and ν are supported on orthogonal subspaces of \mathbf{R}^{d+1} , it is known that the minimizing solution in (2) is not unique. On the contrary, every feasible measure $\gamma \in \Gamma(\mu, \nu)$ has the same cost: each is shown to minimize (2) by constructing a convex function ψ whose subdifferential contains $\text{spt } \mu \times \text{spt } \nu$; as in [16, Remark 15], one takes $\psi := 0$ on $\text{conv}(\text{spt } \mu)$ and $\psi := +\infty$ elsewhere. An elaboration this example shows uniqueness continues to fail when μ and ν measure arclength along a rectangle and its rotation by 90° : one can only hope for uniqueness with strictly convex shapes.

Example 2.1 (Uniqueness Fails for Rectangles) *Let $\Omega \subset \mathbf{R}^2$ be the 2×6 rectangle $HCDG$ and Λ its rotation $A'D'E'F'$ by 90° . Taking $\mu := \mathcal{H}^1|_{\partial\Omega}$ and $\nu := \mathcal{H}^1|_{\partial\Lambda}$ to measure arclength around their respective boundaries, the optimal measures γ in (2) can be computed explicitly (by checking their supports lie in the subdifferential $\partial\psi$ of the convex potential $\psi(x_1, x_2) = f(x_1) + g(x_2)$ given by*

$$\begin{aligned} f(s) &= g(s) := (s^2 + 5)/2 && \text{where } |s| \leq 1, \\ f(s) &:= |3s|, \quad g(s) := |s| + 2 && \text{where } |s| \geq 1, \end{aligned} \tag{24}$$

and invoking Remark 1.3); they turn out not to be unique. Indeed, for efficient transportation the line segments \overline{AB} , \overline{CD} , \overline{EF} and \overline{GH} must be mapped to $\overline{A'B'}$, $\overline{C'D'}$, $\overline{E'F'}$ and $\overline{G'H'}$ respectively in the orientation preserving way. The segment \overline{AH} must also be transported to $\overline{A'H'}$, but this time the map can reverse orientations, preserve them, or do anything in between. In fact, it need not be a map at all, because every joint measure with marginals $\mathcal{H}^1|_{\overline{AH}}$ and $\mathcal{H}^1|_{\overline{A'H'}}$ has the very same cost. Similarly one has the freedom to transport each of the segments \overline{BC} , \overline{DE} and \overline{FG} to its respective image $\overline{B'C'}$, $\overline{D'E'}$ or $\overline{F'G'}$ in an arbitrary way. Thus all three intermediate states depicted in Fig. 3 are possible, and the optimization (2) yields no preference among them. Note that this degeneracy is not removed by incorporating Fry's penalty $\beta \neq 0$ for local rotations.

The remainder of this section is devoted to proving that as long as one of the two shapes — say $\Omega \subset \mathbf{R}^{d+1}$ — is strictly convex, and μ absolutely continuous with respect to its surface measure $\mathcal{H}^d|_{\partial\Omega}$, then the optimal measure $\gamma \in \Gamma(\mu, \nu)$ is unique. In the context of Lemma 1.6 and its corollary, each $\mathbf{x} \in \partial\Omega$ can have multiple images $\mathbf{t}_1(\mathbf{x}), \mathbf{t}_2(\mathbf{x}), \dots, \mathbf{t}_k(\mathbf{x}) \in [\mathbf{t}^+(\mathbf{x}), \mathbf{t}^-(\mathbf{x})]$. The difficulty is to decide what fraction of the mass of μ near \mathbf{x} is transported to each of the $\mathbf{t}_i(\mathbf{x})$. This issue is resolved by Lemma 2.5, which specifies that when $\mathbf{t}_i(\mathbf{x}) \neq \mathbf{t}^-(\mathbf{x})$ no point but \mathbf{x} on the first surface $\partial\Omega$ supplies $\mathbf{t}_i(\mathbf{x}) \in \partial\Lambda$. In principle, the Jacobian of $\mathbf{t}_i : \partial\Omega \rightarrow \partial\Lambda$ at \mathbf{x} and the density $d\nu/d\mathcal{H}^d|_{\partial\Lambda}$ at $\mathbf{y} = \mathbf{t}_i(\mathbf{x})$ then determine what fraction of $d\mu/d\mathcal{H}^d|_{\partial\Omega}$ must move from \mathbf{x} to \mathbf{y} . And after the images $\mathbf{t}_i \neq \mathbf{t}^-$ are saturated, any remaining mass at \mathbf{x} is moved to \mathbf{t}^- . Unfortunately, the proof as outlined would be complicated by the requirement (among other things) that \mathbf{t}_i be differentiable μ -a.e. Such technicalities can be avoided in practice by first pushing whatever mass of ν is unambiguously destined for \mathbf{x} backwards through the map $\mathbf{y} \rightarrow \nabla\phi(\mathbf{y})$ (from Remark 1.2) to obtain a measure $\mu_1 \leq \mu$, and then pushing the balance $\mu - \mu_1$ of the mass of μ forward through the map $\mathbf{x} \rightarrow \mathbf{t}^-(\mathbf{x})$ to obtain a complete picture of γ . This strategy is used to prove Theorem 2.6. We first need to verify that the inner trace $\mathbf{t}^-(\mathbf{x})$ is Borel on $\partial\Omega$ and recall a few facts about *pushing measures forward* through Borel maps: Definition 2.3 and Lemma 2.4.

Lemma 2.2 (Traces are Borel) *Fix a Lipschitz domain $\Omega \subset \mathbf{R}^{d+1}$ and convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}$. The outer and inner traces $\mathbf{t}^+ : \partial\Omega \rightarrow \mathbf{R}^{d+1}$ and $\mathbf{t}^- : \partial\Omega \rightarrow \mathbf{R}^{d+1}$ of $\nabla\psi$ on $\partial\Omega$ are Borel maps (defined \mathcal{H}^d -a.e. by Lemma 1.6).*

Proof: 1. Recall that the gradient of a Lipschitz function is defined as the limit of continuous approximants (finite differences), hence Borel. By Rademacher's theorem, these limits converge outside a Borel set of measure zero. From this one can readily deduce the well-known fact that the Gauss map $\mathbf{n}_\Omega : \partial\Omega \rightarrow \partial\mathbf{B}(\mathbf{0}, 1) \subset \mathbf{R}^{d+1}$ on a Lipschitz hypersurface is a Borel map (defined \mathcal{H}^d -a.e.).

2. Claim: If $\partial\Omega$ is differentiable and ψ is tangentially differentiable at a point $\mathbf{x} \in \partial\Omega$, then the i -th component of $\mathbf{t}^+(\mathbf{x}) = (t_1^+(\mathbf{x}), \dots, t_{d+1}^+(\mathbf{x}))$ is given by

$$t_i^+(\mathbf{x}) = \lim_{\lambda \rightarrow 0^+} \lambda^{-1} [\psi(\mathbf{x} + \lambda \mathbf{e}_i) - \psi(\mathbf{x})] \quad (25)$$

if $\mathbf{e}_i \cdot \mathbf{n}_\Omega(\mathbf{x}) \geq 0$, and otherwise by the same limit but with $\lambda \rightarrow 0^-$ from below.

Proof of Claim: From (22) recall

$$\lim_{\lambda \rightarrow 0^+} \lambda^{-1} [\psi(\mathbf{x} + \lambda \mathbf{e}_i) - \psi(\mathbf{x})] = \sup_{\mathbf{y} \in \partial\psi(\mathbf{x})} \mathbf{y} \cdot \mathbf{e}_i \quad (26)$$

$$\text{and } \lim_{\lambda \rightarrow 0^-} \lambda^{-1} [\psi(\mathbf{x} + \lambda \mathbf{e}_i) - \psi(\mathbf{x})] = \inf_{\mathbf{y} \in \partial\psi(\mathbf{x})} \mathbf{y} \cdot \mathbf{e}_i. \quad (27)$$

Since Lemma 1.6 forces $\partial\psi(\mathbf{x}) = [\mathbf{t}^+(\mathbf{x}), \mathbf{t}^-(\mathbf{x})]$ to be a line segment whose orientation $\mathbf{t}^+(\mathbf{x}) - \mathbf{t}^-(\mathbf{x})$ parallels $\mathbf{n}_\Omega(\mathbf{x})$, we may evaluate the supremum and infimum explicitly. Indeed, the supremum (26) is given by $\mathbf{t}^+(\mathbf{x}) \cdot \mathbf{e}_i$ when $\mathbf{e}_i \cdot \mathbf{n}_\Omega(\mathbf{x}) \geq 0$, while the infimum (27) is given by $\mathbf{t}^-(\mathbf{x}) \cdot \mathbf{e}_i$ if $\mathbf{e}_i \cdot \mathbf{n}_\Omega(\mathbf{x}) \leq 0$. This establishes Claim 2.

3. Now let $S \subset \partial\Omega$ be the Borel set on which $\partial\Omega$ is differentiable and ψ is tangentially differentiable. The subset $S_i = \{\mathbf{x} \in S \mid \mathbf{n}_\Omega(\mathbf{x}) \cdot \mathbf{e}_i \geq 0\}$ and its complement $S_i^c := S \setminus S_i$ are Borel since the Gauss map $\mathbf{n}_\Omega(\cdot)$ was Borel in step 1. On S_i , $t_i^+(\mathbf{x})$ is given by the restriction to S_i of a limit (25) of continuous functions, hence $t_i^+ : S_i \rightarrow \mathbf{R}$ is Borel. Similarly, the second part of Claim 2 gives $t_i^+ : S_i^c \rightarrow \mathbf{R}$ as a measurable limit of continuous functions when $\lambda \rightarrow 0^-$. Thus we conclude that each component $t_i^+ : S \rightarrow \mathbf{R}$ is Borel measurable and hence $\mathbf{t}^+ : S \rightarrow \mathbf{R}^{d+1}$ is Borel.

4. Interchanging Ω with $\mathbf{R}^{d+1} \setminus \Omega$ corresponds to reversing the directions of the outer unit normal $\mathbf{n}_\Omega(\mathbf{x})$ along $\partial\Omega$, hence to interchanging \mathbf{t}^+ and \mathbf{t}^- . Thus we conclude from step 3 that $\mathbf{t}^- : S \rightarrow \mathbf{R}^{d+1}$ is also Borel. Since the complement of $S \subset \partial\Omega$ has zero area $\mathcal{H}^d[\partial\Omega \setminus S] = 0$, the proof of the lemma is complete. \square

Definition 2.3 (Push-forward) A measure ω on $(\mathbf{X}, \mathcal{X})$ and measurable map $\mathbf{t} : \mathbf{X} \rightarrow \mathbf{Y}$ induce a measure on $(\mathbf{Y}, \mathcal{Y})$ called the push-forward of ω through \mathbf{t} , denoted $\mathbf{t}_{\#}\omega$ and defined by

$$(\mathbf{t}_{\#}\omega)[B] := \omega[\mathbf{t}^{-1}(B)] \quad (28)$$

for each set $B \in \mathcal{Y}$. More generally, the map \mathbf{t} need only be defined on a subset $S \in \mathcal{X}$ of full measure $\mu[\mathbf{X} \setminus S] = 0$.

Note that for measurable functions $f : \mathbf{Y} \rightarrow \mathbf{R}$, it follows from Definition 2.3 that the change of variables formula is satisfied:

$$\int_{\mathbf{Y}} f(\mathbf{y}) d\mathbf{t}_{\#}\omega(\mathbf{y}) = \int_{S \subset \mathbf{X}} f(\mathbf{t}(\mathbf{x})) d\omega(\mathbf{x}). \quad (29)$$

The following lemma shows that any joint measure γ supported on the graph of a function $\{(\mathbf{x}, \mathbf{t}(\mathbf{x})) \mid \mathbf{x} \in S\}$ can be recovered from \mathbf{t} and its left marginal μ , by pushing μ forward through the map $\mathbf{id} \times \mathbf{t} : \mathbf{x} \rightarrow (\mathbf{x}, \mathbf{t}(\mathbf{x}))$. Here $\mathbf{id}(\mathbf{x}) = \mathbf{x}$ denotes the identity map. In alternative (Dirac) notation the conclusion is often expressed by writing $d\gamma(\mathbf{x}, \mathbf{y}) = \delta(\mathbf{y} - \mathbf{t}(\mathbf{x}))d\mu(\mathbf{x})$.

Lemma 2.4 (Measures on Graphs are Push-Forwards) *Let (\mathbf{X}, d) and (\mathbf{Y}, ρ) be metric spaces with a Borel measure μ on \mathbf{X} and Borel map $\mathbf{t} : S \rightarrow \mathbf{Y}$ defined on a (Borel) subset $S \subset \mathbf{X}$ of full measure $\mu[\mathbf{X} \setminus S] = 0$. If a non-negative Borel measure γ on the product space $\mathbf{X} \times \mathbf{Y}$ has left marginal μ and satisfies*

$$\int_{\mathbf{X} \times \mathbf{Y}} \rho(\mathbf{t}(\mathbf{x}), \mathbf{y}) d\gamma(\mathbf{x}, \mathbf{y}) = 0, \quad (30)$$

then $\gamma = (\mathbf{id} \times \mathbf{t})_{\#}\mu$.

Proof: The argument generalizes the proof of [16, Proposition 10]. It suffices to show that the measure $(\mathbf{id} \times \mathbf{t})_{\#}\mu$ coincides with γ on products $U \times V$ of Borel sets $U \subset \mathbf{X}$ and $V \subset \mathbf{Y}$; the semi-algebra of such products generates the Borel sets in $\mathbf{X} \times \mathbf{Y}$. Observe that the graph $G := \{(\mathbf{x}, \mathbf{t}(\mathbf{x})) \mid \mathbf{x} \in S\}$ coincides with the zeros in $S \times \mathbf{Y}$ of the measurable function $\rho(\mathbf{t}(\mathbf{x}), \mathbf{y})$. This graph must therefore be a Borel set of full γ measure by (30). Since

$$(U \times V) \cap G = ((U \cap \mathbf{t}^{-1}(V)) \times \mathbf{Y}) \cap G$$

it follows that

$$\begin{aligned} \gamma[U \times V] &= \gamma[(U \cap \mathbf{t}^{-1}(V)) \times \mathbf{Y}] \\ &= \mu[U \cap \mathbf{t}^{-1}(V)] \\ &= ((\mathbf{id} \times \mathbf{t})_{\#}\mu)[U \times V], \end{aligned}$$

which proves that $\gamma = (\mathbf{id} \times \mathbf{t})_{\#}\mu$. □

For boundary measures on convex domains, the next proposition implies that $\mathbf{x} \in \partial\Omega$ alone supplies destinations in the half-open segment $[\mathbf{t}^+(\mathbf{x}), \mathbf{t}^-(\mathbf{x})) \subset \partial\psi(\mathbf{x})$. On a non-convex domain, a similar lemma would hold at exposed points of $\text{conv}(\Omega)$.

Lemma 2.5 (Sole Supplier) *Fix a strictly convex domain $\Omega \subset \mathbf{R}^{d+1}$ and convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}$. At $\mathbf{x} \in \partial\Omega$, let $\partial\Omega$ be differentiable and ψ tangentially differentiable. Then if $(\mathbf{x}, \mathbf{y}) \in \partial\psi$, either (i) $\overline{\Omega} \cap \partial\psi^*(\mathbf{y}) = \{\mathbf{x}\}$ or else (ii) $\mathbf{y} = \mathbf{t}^-(\mathbf{x})$, where \mathbf{t}^- denotes the inner trace of $\nabla\psi$ on $\partial\Omega$ defined in Lemma 1.6.*

Proof: Recall that $\partial\psi(\mathbf{x}) = [\mathbf{t}^+(\mathbf{x}), \mathbf{t}^-(\mathbf{x})]$ with $\mathbf{t}^+(\mathbf{x}) - \mathbf{t}^-(\mathbf{x})$ parallel to $\mathbf{n}_\Omega(\mathbf{x})$ by Lemma 1.6. Thus when (ii) fails, $\mathbf{t}^-(\mathbf{x}) - \mathbf{y} \neq \mathbf{0}$ is antiparallel to $\mathbf{n}_\Omega(\mathbf{x})$. Assuming (ii) fails, take $\mathbf{z} \in \overline{\Omega} \cap \partial\psi^*(\mathbf{y})$. Monotonicity of $\partial\psi$ then yields $(\mathbf{z} - \mathbf{x}) \cdot (\mathbf{y} - \mathbf{t}^-(\mathbf{x})) \geq 0$ so $(\mathbf{z} - \mathbf{x}) \cdot \mathbf{n}_\Omega(\mathbf{x}) \geq 0$. Strict convexity of Ω implies $\mathbf{z} = \mathbf{x}$ as in (8). Thus (i) holds unless $\overline{\Omega} \cap \partial\psi^*(\mathbf{y})$ is empty, which cannot happen since $(\mathbf{y}, \mathbf{x}) \in \partial\psi^*$ was a hypothesis. □

Theorem 2.6 (Uniqueness of Correlation with a Strictly Convex Shape)

Fix Borel probability measures μ and ν of bounded support in \mathbf{R}^{d+1} . If μ is absolutely continuous with respect to the surface measure $\mathcal{H}^d|_{\partial\Omega}$ of a strictly convex domain $\Omega \subset \mathbf{R}^{d+1}$, then the infimum (2) is uniquely attained. Indeed the unique optimal measure in $\Gamma(\mu, \nu)$ is given by $\gamma = \gamma_1 + \gamma_2$, where

$$\gamma_1 = (\nabla\phi \times \mathbf{id})_{\#}\nu_1 \quad \text{and} \quad \gamma_2 = (\mathbf{id} \times \mathbf{t}^-)_{\#}\mu_2, \quad (31)$$

$\nu_1 := \nu|_{\text{dom } \nabla\phi}$ is the restriction of ν to the subset of \mathbf{R}^{d+1} where ϕ is differentiable, $\mu_1 := (\nabla\phi)_{\#}\nu_1$ and $\mu_2 := \mu - \mu_1$. Here ϕ and ψ are the convex functions of (20) and Theorem 1.1, while $\mathbf{t}^- : \partial\Omega \rightarrow \mathbf{R}^{d+1}$ is the inner trace of $\nabla\psi$ on $\partial\Omega$.

Proof: At least one optimal measure exists according to, e.g. Givens & Shortt [12].

1. Suppose the infimum (2) is attained by $\gamma \in \Gamma(\mu, \nu)$. Denoting the Borel set where ϕ is differentiable by $T := \text{dom } \nabla\phi \subset \mathbf{R}^{d+1}$, let us define $\gamma_1 := \gamma|_{\mathbf{R}^{d+1} \times T}$ and $\gamma_2 := \gamma - \gamma_1$ and try to derive (31). When this derivation is complete, uniqueness will have been established since (31) expresses γ completely in terms of μ , ν , ψ and ϕ (which depend only on μ and ν in Theorem 1.1 and not on γ).

2. The reflection $\gamma_1^* := R_{\#}\gamma_1$ of γ_1 under $R(\mathbf{x}, \mathbf{y}) := (\mathbf{y}, \mathbf{x})$ assigns full mass to the Borel set

$$(T \times \mathbf{R}^{d+1}) \cap \partial\phi = \{(\mathbf{x}, \nabla\phi(\mathbf{x})) \mid \mathbf{x} \in T\}$$

since $\text{spt } \gamma_1^* \subset \partial\phi$ by Remark 1.2. Thus

$$\begin{aligned} \int_{\mathbf{R}^{d+1} \times \mathbf{R}^{d+1}} |\nabla\phi(\mathbf{x}) - \mathbf{y}| d\gamma_1^*(\mathbf{x}, \mathbf{y}) &= \int_{(T \times \mathbf{R}^{d+1}) \cap \partial\phi} |\nabla\phi(\mathbf{x}) - \nabla\phi(\mathbf{x})| d\gamma_1^*(\mathbf{x}, \mathbf{y}) \\ &= 0. \end{aligned}$$

Since the left marginal of γ_1^* agrees with $\nu_1 := \nu|_{\text{dom } \nabla\phi}$, applying Lemma 2.4 yields $\gamma_1^* = (\mathbf{id} \times \nabla\phi)_{\#}\nu_1$. Reflecting proves the first equality in (31).

3. The last conclusion shows the left marginal of γ_1 to be $\mu_1 := (\nabla\phi)_{\#}\nu_1$. The left marginal of $\gamma_2 := \gamma - \gamma_1$ must therefore coincide with $\mu_2 := \mu - \mu_1$. Let $S \subset \partial\Omega$ denote the set where $\partial\Omega$ is differentiable and ψ is tangentially differentiable. It is a Borel set of full μ_2 measure on which the inner trace $\mathbf{t}^- : S \rightarrow \mathbf{R}^{d+1}$ of $\nabla\psi$ is Borel and defined by Lemmas 1.6 and 2.2. Setting $T^c = \text{spt } \nu \setminus T$, we claim:

4. *Claim:* If $(\mathbf{x}, \mathbf{y}) \in (S \times T^c) \cap \partial\psi$ then $\mathbf{y} = \mathbf{t}^-(\mathbf{x})$.

Proof of Claim: Let $(\mathbf{x}, \mathbf{y}) \in (S \times T^c) \cap \partial\psi$. According to Lemma 2.5, either (i) $\overline{\Omega} \cap \partial\psi^*(\mathbf{y}) = \{\mathbf{x}\}$ or (ii) $\mathbf{y} = \mathbf{t}^-(\mathbf{x})$. We shall show that (i) implies differentiability of ϕ at \mathbf{y} thus contradicting $\mathbf{y} \in T^c$. Recall Remark 1.2 yields $\partial\phi(\mathbf{y}) \subset \partial\psi^*(\mathbf{y})$. Also, the bounded convex set $\partial\phi(\mathbf{y})$ has its extreme points on $\text{spt } \mu \subset \overline{\Omega}$ by Corollary 1.4, hence is contained in the convex set $\overline{\Omega}$. Thus (i) implies $\partial\phi(\mathbf{y}) \subset \{\mathbf{x}\}$, which in turn yields $\mathbf{x} = \nabla\phi(\mathbf{y})$, contradicting $\mathbf{y} \in T^c$. The only alternative (ii) establishes the claim.

5. Theorem 1.1 asserts $\text{spt } \gamma_2 \subset \partial\psi$. Using this gives

$$\int_{\mathbf{R}^{d+1} \times T^c} |\mathbf{t}^-(\mathbf{x}) - \mathbf{y}| d\gamma_2(\mathbf{x}, \mathbf{y}) = \int_{(S \times T^c) \cap \partial\psi} |\mathbf{t}^-(\mathbf{x}) - \mathbf{y}| d\gamma_2(\mathbf{x}, \mathbf{y}) \quad (32)$$

$$= 0, \quad (33)$$

where the second equality follows from claim 4. On the other hand, the definition of $\gamma_2 = \gamma|_{\mathbf{R}^{d+1} \times T^c}$ yields

$$\int_{\mathbf{R}^{d+1} \times T} |\mathbf{t}^-(\mathbf{x}) - \mathbf{y}| d\gamma_2(\mathbf{x}, \mathbf{y}) = 0. \quad (34)$$

Summing (33) and (34) allows us to conclude $\gamma_2 = (\mathbf{id} \times \mathbf{t}^-)_\# \mu_2$ from Lemma 2.4. This establishes (31) and completes the proof of the theorem. \square

3 Pairs of Maps Between Strictly Convex Curves

The main goal of this section is to develop a (partial) regularity theory for the convex function ψ from Theorem 1.1 along the boundary of Ω , when μ and ν are a pair of measures satisfying the following hypotheses:

Definition 3.1 (Suitable Boundary Measures) *Let Ω and $\Lambda \subset \mathbf{R}^{d+1}$ be bounded, strictly convex domains with Borel probability measures μ on $\partial\Omega$ and ν on $\partial\Lambda$. The pair of boundary measures (μ, ν) is said to be suitable if (i) μ has no atoms and (ii) $\text{spt } \nu = \partial\Lambda$. For $d > 1$ we furthermore require (iii) the surface density bounds*

$$\mu < \epsilon^{-1} \mathcal{H}^d|_{\partial\Omega} \quad \text{and} \quad \nu > \epsilon \mathcal{H}^d|_{\partial\Lambda} \quad (35)$$

hold for some $\epsilon > 0$, and (iv) the uniform convexity of $\Omega \subset \mathbf{R}^{d+1}$.

If in addition, the above hypotheses remain satisfied when the roles of $\mu \leftrightarrow \nu$ and $\Omega \leftrightarrow \Lambda$ are interchanged, we say the pair (μ, ν) is symmetrically suitable.

Our first task is to establish tangential differentiability of ψ not only almost everywhere, but everywhere along $\partial\Omega \subset \mathbf{R}^{d+1}$. In the plane this follows from the next proposition (through Corollary 4.4). The more complicated argument in higher dimensions $d > 1$ is postponed until Section 4. Since we know this result will be established, we proceed to state Theorem 3.8 — the culmination of the present section — in all dimensions. It asserts that the outer and inner traces of $\nabla\psi$ on $\partial\Omega$ yield a pair of continuous maps \mathbf{t}^+ and $\mathbf{t}^-|_{\{\mathbf{x} \in \partial\Omega | \mathbf{t}^- \neq \mathbf{t}^+\}}$ covering the support of the optimal measure. Two concluding examples demonstrate the necessity of both maps. In these examples, non-tangential differentiability of ψ must fail.

Proposition 3.2 (Images Surely Parallel Normal) *Let $\Omega, \Lambda \subset \mathbf{R}^2$ be bounded, strictly convex domains with suitable measures μ on $\partial\Omega$ and ν on $\partial\Lambda$. Take the convex function ψ from Theorem 1.1. If \mathbf{n}_o is any normal to Ω at $\mathbf{x}_o \in \partial\Omega$ and both $\mathbf{y}_o, \mathbf{y} \in \partial\psi(\mathbf{x}_o)$, then $\mathbf{y}_o - \mathbf{y}$ is parallel to \mathbf{n}_o .*

Proof: Recall that $\partial\psi(\mathbf{x}_o)$ is a closed, non-empty, convex set whose extreme points lie on $\partial\Lambda = \text{spt } \nu$ by Corollary 1.4. Thus $\partial\psi(\mathbf{x}_o) \subset \bar{\Lambda}$.

1. Claim: If $\mathbf{y} \in \Lambda \cap \partial\psi(\mathbf{x}_o)$ there exists a unique $\lambda > 0$ such that $\mathbf{y} + \lambda\mathbf{n}_o \in \partial\Lambda$. Furthermore, $\partial\phi(\mathbf{y} + \lambda\mathbf{n}_o) = \{\mathbf{x}_o\}$ where ϕ is defined in (20).

Proof of Claim: Assume $\mathbf{n}_o \neq \mathbf{0}$ (since otherwise the proposition holds trivially) and define the ray $L := \{\mathbf{y} + \lambda\mathbf{n}_o \mid \lambda \geq 0\}$. The intersection $L \cap \bar{\Lambda}$, being compact and convex, must be a segment $[\mathbf{y}, \mathbf{y} + \lambda\mathbf{n}_o]$ with one endpoint \mathbf{y} in Λ . The other endpoint $\mathbf{y} + \lambda\mathbf{n}_o$ belongs to $\partial\Lambda$, being a limit of points $\mathbf{y} + (\lambda + 1/k)\mathbf{n}_o$ outside $\bar{\Lambda}$. Convexity of Λ implies $[\mathbf{y}, \mathbf{y} + \lambda\mathbf{n}_o) \subset \Lambda$ by Rockafellar [20, Theorem 6.1], so $\mathbf{y} + \lambda\mathbf{n}_o$ is the only point where L intersects $\partial\Lambda$. This proves a unique $\lambda > 0$ exists.

Setting $\mathbf{y}_1 := \mathbf{y} + \lambda\mathbf{n}_o \in \partial\Lambda (= \text{spt } \nu)$, the second part of the claim — that \mathbf{y}_1 is supplied solely by \mathbf{x}_o — is argued like Lemma 2.5. Indeed, $\partial\phi(\mathbf{y}_1)$ is a (non-empty) subset of $\bar{\Omega} \cap \partial\psi^*(\mathbf{y}_1)$ by Remark 1.2 and Corollary 1.4, so let $\mathbf{x}_1 \in \partial\phi(\mathbf{y}_1)$. Then $\mathbf{y}_1 \in \partial\psi(\mathbf{x}_1)$ and the monotonicity of $\partial\psi$ yields

$$\lambda(\mathbf{x}_1 - \mathbf{x}_o) \cdot \mathbf{n}_o = (\mathbf{x}_1 - \mathbf{x}_o) \cdot (\mathbf{y}_1 - \mathbf{y}) \geq 0.$$

The strict convexity of $\partial\Omega$ then forces $\mathbf{x}_1 = \mathbf{x}_o$ because of (8). This concludes the proof of Claim 1.

2. We assume henceforth that $\partial\psi(\mathbf{x}_o) \subset \bar{\Omega}$ contains two distinct points $\mathbf{y}_o \neq \mathbf{y}$ since otherwise the proposition holds trivially. We shall complete the proof of the proposition by showing:

Claim: The segment $[\mathbf{y}_o, \mathbf{y}] := \{s\mathbf{y}_o + (1-s)\mathbf{y} \mid s \in [0, 1]\}$ is parallel to \mathbf{n}_o .

Proof of Claim: The entire segment $[\mathbf{y}_o, \mathbf{y}]$ lies in the convex sets $\partial\psi(\mathbf{x}_o) \subset \bar{\Lambda}$. Since $\bar{\Lambda}$ is strictly convex, a slightly shorter segment, also denoted $[\mathbf{y}_o, \mathbf{y}] \subset \Lambda$, lies in the interior of this domain. Defining $\mathbf{y}(s) = s\mathbf{y}_o + (1-s)\mathbf{y}$, for each $s \in [0, 1]$ we deduce from Claim 1 that there exists a unique $\lambda(s) > 0$ such that $\mathbf{y}(s) + \lambda(s)\mathbf{n}_o \in \partial\Lambda$ and

$$\partial\phi(\mathbf{y}(s) + \lambda(s)\mathbf{n}_o) = \{\mathbf{x}_o\}. \quad (36)$$

The uniqueness of $\lambda(s) \geq 0$ implies the map $s \rightarrow \mathbf{y}(s) + \lambda(s)\mathbf{n}_o$ is continuous, so

$$A := \{\mathbf{y}(s) + \lambda(s)\mathbf{n}_o \mid s \in [0, 1]\} \subset \partial\Omega$$

is a connected closed arc on the boundary of Ω . Either its endpoints $\mathbf{y}_o + \lambda(0)\mathbf{n}_o$ and $\mathbf{y} + \lambda(1)\mathbf{n}_o$ coincide, in which case the claim will follow, or else A contains a relatively open subset of $\partial\Lambda = \text{spt } \nu$ in which case $\nu[A] > 0$.

We preclude the latter possibility by considering any minimizer $\gamma \in \Gamma(\mu, \nu)$ for (2) (which exists by lower semicontinuity and compactness; see e.g. Givens & Shortt [12]). Since $\text{spt } \gamma \subset \partial\phi^*$ from Remark 1.2, (36) yields

$$(\mathbf{R}^2 \times A) \cap \text{spt } \gamma \subset \{\mathbf{x}_o\} \times A,$$

hence

$$\nu[A] = \int_{\mathbf{R}^2 \times A} d\gamma \leq \int_{\{\mathbf{x}_o\} \times \mathbf{R}^2} d\gamma = \mu[\{\mathbf{x}_o\}] = 0.$$

Thus A must degenerate to a single point $A = \{\mathbf{y}_o + \lambda(0)\mathbf{n}_o\} = \{\mathbf{y} + \lambda(1)\mathbf{n}_o\}$ so $\mathbf{y} - \mathbf{y}_o = (\lambda(1) - \lambda(0))\mathbf{n}_o$, to conclude the proof of the claim. \square

An elementary lemma helps to distinguish $\mathbf{t}^+(\mathbf{x})$ from $\mathbf{t}^-(\mathbf{x})$ without referring to ψ :

Lemma 3.3 *Take distinct points $\mathbf{y}_1, \mathbf{y}_2 \in \partial\Lambda$ on the boundary of a strictly convex domain $\Lambda \subset \mathbf{R}^{d+1}$. Every outer unit normal \mathbf{q}_1 to $\partial\Lambda$ at \mathbf{y}_1 satisfies $\mathbf{q}_1 \cdot (\mathbf{y}_1 - \mathbf{y}_2) > 0$. Similarly, each $\mathbf{q}_2 \in N_\Lambda(\mathbf{y}_2)$ satisfies $\mathbf{q}_2 \cdot (\mathbf{y}_1 - \mathbf{y}_2) < 0$.*

Proof: According to our definition, a (generalized) outer normal $\mathbf{q}_1 \in N_\Lambda(\mathbf{y}_1)$ satisfies $\mathbf{q}_1 \cdot (\mathbf{z} - \mathbf{y}_1) \leq 0$ for all $\mathbf{z} \in \overline{\Lambda}$. Strict convexity of Λ forces a strict inequality unless $\mathbf{z} = \mathbf{y}_1$ by (8). Setting $\mathbf{z} = \mathbf{y}_2 \neq \mathbf{y}_1$ proves the first desired inequality $\mathbf{q}_1 \cdot (\mathbf{y}_2 - \mathbf{y}_1) < 0$. The second desired inequality follows by relabelling symmetry after interchanging the indices $1 \leftrightarrow 2$. \square

For the next proposition, we recall the notation $[\mathbf{y}_1, \mathbf{y}_2]$ for line segments (11), and that extreme points of $\partial\psi(\mathbf{x})$ must lie on $\partial\Lambda = \text{spt } \nu$ in view of Corollary 1.4.

Proposition 3.4 (Trichotomy) *Let $\Omega, \Lambda \subset \mathbf{R}^{d+1}$ be bounded, strictly convex domains with suitable measures μ on $\partial\Omega$ and ν on $\partial\Lambda$. Take the convex function ψ as in Theorem 1.1. For each $\mathbf{x} \in \partial\Omega$ exactly one of the following three statements holds:*

- (o) $\partial\psi(\mathbf{x}) = \{\mathbf{y}_1\}$ with $\mathbf{n} \cdot \mathbf{q}_1 = 0$ for some pair of normals $\mathbf{n} \in N_\Omega(\mathbf{x})$ and $\mathbf{q}_1 \in N_\Lambda(\mathbf{y}_1)$;
- (i) $\partial\psi(\mathbf{x}) = \{\mathbf{y}_1\}$ with $\mathbf{n} \cdot \mathbf{q}_1 > 0$ for all pairs $\mathbf{n} \in N_\Omega(\mathbf{x})$ and $\mathbf{q}_1 \in N_\Lambda(\mathbf{y}_1)$; or
- (ii) $\partial\psi(\mathbf{x}) = [\mathbf{y}_1, \mathbf{y}_2]$, in which case $\partial\Omega$ is differentiable at \mathbf{x} and $\mathbf{n}_\Omega(\mathbf{x}) \cdot \mathbf{q}_1 > 0$ but $\mathbf{n}_\Omega(\mathbf{x}) \cdot \mathbf{q}_2 < 0$ for all $\mathbf{q}_i \in N_\Lambda(\mathbf{y}_i)$, $i = 1, 2$.

Proof: Recall that $\partial\psi(\mathbf{x})$ is a non-empty compact convex set which lies in a line parallel to $\mathbf{n}_o \in N_\Omega(\mathbf{x})$ by Proposition 3.2 (or 4.3). Thus $\partial\psi(\mathbf{x}) = [\mathbf{y}_1, \mathbf{y}_2]$ consists either of a segment parallel to \mathbf{n}_o or a single point $\mathbf{y}_1 = \mathbf{y}_2$. Either way the extreme points $\mathbf{y}_1, \mathbf{y}_2 \in \text{spt } \nu \subset \partial\Lambda$ by Corollary 1.4.

1. Assume first that $\mathbf{y}_1 \neq \mathbf{y}_2$, so neither (o) nor (i) holds. Now Proposition 3.2 implies that every outward unit normal $\mathbf{n} \in N_\Omega(\mathbf{x})$ must be parallel to $\mathbf{y}_1 - \mathbf{y}_2$, which means there is only one supporting hyperplane to Ω at \mathbf{x} . In other words, \mathbf{x} is a point of differentiability for $\partial\Omega$. Interchange \mathbf{y}_1 and \mathbf{y}_2 if necessary to ensure that $\mathbf{y}_1 - \mathbf{y}_2$ is an outward normal rather than an inward normal. Lemma 3.3 then concludes the proof that $\mathbf{n}_\Omega(\mathbf{x}) \cdot \mathbf{q}_1 > 0$ but $\mathbf{n}_\Omega(\mathbf{x}) \cdot \mathbf{q}_2 < 0$ for all $\mathbf{q}_i \in N_\Lambda(\mathbf{y}_i)$, so we are in case (ii).

2. Suppose $\mathbf{y}_1 = \mathbf{y}_2$ on the other hand. Then $\partial\psi(\mathbf{x}) = \{\mathbf{y}_1\}$ and (ii) is violated. We must show that one of the mutually exclusive claims (o) or (i) is satisfied. Choose any normal $\mathbf{n} \in N_\Omega(\mathbf{x})$ and define the line $L := \{\mathbf{y}_1 + \lambda\mathbf{n} \mid \lambda \in \mathbf{R}\}$. The convex intersection $L \cap \overline{\Lambda}$ consists of either a segment $[\mathbf{y}_0, \mathbf{y}_1]$ with endpoints on $\partial\Lambda$ or the single point $\mathbf{y}_1 \in \partial\Lambda$.

If the line intersects $\overline{\Lambda}$ only at $\mathbf{y}_1 \in \partial\Lambda$, then it can be separated from the convex interior Λ by a hyperplane: there exists a unit vector $\mathbf{q} \in \mathbf{R}^{d+1}$ and $s \in \mathbf{R}$ such that

$$\mathbf{q} \cdot \mathbf{y} \geq s \quad \text{for all } \mathbf{y} \in L, \quad (37)$$

$$\mathbf{q} \cdot \mathbf{y} < s \quad \text{for all } \mathbf{y} \in \Lambda. \quad (38)$$

Obviously $s = \mathbf{q} \cdot \mathbf{y}_1$, so we deduce $\mathbf{q} \in N_\Lambda(\mathbf{y}_1)$ from (38). Furthermore, substituting $\mathbf{y} = \mathbf{y}_1 \pm \mathbf{n}$ in (37) yields both $\mathbf{q} \cdot \mathbf{n} \geq 0$ as well as $\mathbf{q} \cdot \mathbf{n} \leq 0$, whence (o) has been verified.

The other possibility is $L \cap \bar{\Lambda} = [\mathbf{y}_0, \mathbf{y}_1]$ with $\mathbf{y}_1 \neq \mathbf{y}_0 \in \partial\Lambda$. In this case we need to prove $\mathbf{y}_1 - \mathbf{y}_0$ is parallel to \mathbf{n} rather than antiparallel. To derive a contradiction, suppose $\mathbf{y}_1 = \mathbf{y}_0 - \lambda\mathbf{n}$ with $\lambda > 0$. Since $\mathbf{y}_0 \in \text{spt } \nu$ we have $\partial\phi(\mathbf{y}_0) \subset \bar{\Omega} \cap \partial\psi^*(\mathbf{y}_0)$ non-empty, where ϕ is defined in Remark 1.2. Therefore, take $\mathbf{x}_0 \in \bar{\Omega}$ with $(\mathbf{x}_0, \mathbf{y}_0) \in \partial\psi$ and observe monotonicity yields

$$\lambda(\mathbf{x}_0 - \mathbf{x}) \cdot \mathbf{n} = (\mathbf{x}_0 - \mathbf{x}) \cdot (\mathbf{y}_0 - \mathbf{y}_1) \geq 0.$$

The strict convexity of Ω forces $\mathbf{x}_0 = \mathbf{x}$ as in (8). But then $\mathbf{y}_0 \in \partial\psi(\mathbf{x})$ contradicts our hypothesis 2: $\partial\psi(\mathbf{x}) = \{\mathbf{y}_1\}$. Thus we have proved that $\mathbf{y}_1 - \mathbf{y}_0 = \lambda\mathbf{n}$ with $\lambda > 0$. Lemma 3.3 now shows $\lambda\mathbf{q}_1 \cdot \mathbf{n} = \mathbf{q}_1 \cdot (\mathbf{y}_1 - \mathbf{y}_0) > 0$ for all $\mathbf{q}_1 \in N_\Lambda(\mathbf{y}_1)$ to complete the proof that (i) holds. \square

Observe that ψ is differentiable at \mathbf{x} in each of the first two cases; in the last case ($\mathbf{y}_1 \neq \mathbf{y}_2$) its non-tangential derivatives fail to exist since the partial derivatives $\partial\psi/\partial x_i$ from the left and the right (26–27) do not coincide unless $\mathbf{e}_i \cdot \mathbf{n}_\Omega(\mathbf{x}) = 0$.

Definition 3.5 (Twin Maps) *Given $\Omega, \Lambda, (\mu, \nu)$ and ψ as in Proposition 3.4, we decompose $\partial\Omega = S_0 \cup S_1 \cup S_2$ into three disjoint sets such that (o) holds for $\mathbf{x} \in S_0$, (i) holds for $\mathbf{x} \in S_1$ and (ii) holds for $\mathbf{x} \in S_2$. We use the extreme images $\mathbf{y}_1, \mathbf{y}_2 \in \partial\psi(\mathbf{x})$ of the proposition to define an outer map $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ and inner map $\mathbf{t}^- : S_2 \rightarrow \partial\Lambda$ by $\mathbf{t}^+(\mathbf{x}) := \mathbf{y}_1$ and $\mathbf{t}^-(\mathbf{x}) := \mathbf{y}_2$. It is convenient to extend the definition of \mathbf{t}^- to $\partial\Omega$ by setting $\mathbf{t}^-(\mathbf{x}) := \mathbf{t}^+(\mathbf{x})$ for $\mathbf{x} \in S_0 \cup S_1$.*

Definition 3.6 (Inverse Maps) *If the boundary measures (ν, μ) are symmetrically suitable, then the analogous construction starting from (20) yields a decomposition $\partial\Lambda = T_0 \cup T_1 \cup T_2$ and maps $\mathbf{s}^\pm : \partial\Lambda \rightarrow \partial\Omega$ such that $\partial\phi(\mathbf{y}) = [\mathbf{s}^+(\mathbf{y}), \mathbf{s}^-(\mathbf{y})]$ while (o) $\mathbf{y} \in T_0$ if $\mathbf{s}^+(\mathbf{y}) = \mathbf{s}^-(\mathbf{y})$ and $\mathbf{q} \cdot \mathbf{n} = 0$ for some $\mathbf{q} \in N_\Lambda(\mathbf{y})$ and $\mathbf{n} \in N_\Omega(\mathbf{s}^+(\mathbf{y}))$; (i) $\mathbf{y} \in T_1$ if $\mathbf{s}^+(\mathbf{y}) = \mathbf{s}^-(\mathbf{y})$ and $\mathbf{q} \cdot \mathbf{n} > 0$ for all $\mathbf{q} \in N_\Lambda(\mathbf{y})$ and $\mathbf{n} \in N_\Omega(\mathbf{s}^+(\mathbf{y}))$; and (ii) $\mathbf{y} \in T_2$ if $\mathbf{s}^+(\mathbf{y}) \neq \mathbf{s}^-(\mathbf{y})$, in which case $\partial\Lambda$ is differentiable at \mathbf{y} and $\mathbf{n}_\Lambda(\mathbf{y}) \cdot \mathbf{n}^+ > 0$ but $\mathbf{n}_\Lambda(\mathbf{y}) \cdot \mathbf{n}^- < 0$ for all $\mathbf{n}^\pm \in N_\Omega(\mathbf{s}^\pm(\mathbf{y}))$.*

Comparing these definitions with the discussion of Lemma 1.6, it should be clear that \mathbf{t}^+ is the trace of $\nabla\psi|_{\Omega^c}$ on $\partial\Omega$ and \mathbf{t}^- the boundary trace of $\nabla\psi|_\Omega$. Similarly, \mathbf{s}^+ and \mathbf{s}^- are the outer and inner traces respectively of $\nabla\phi$ on $\partial\Lambda$, where ϕ is defined by (20). The next proposition explores continuity properties of these four maps.

Proposition 3.7 (Twin Homeomorphisms) *Given bounded, strictly convex domains $\Omega, \Lambda \subset \mathbf{R}^{d+1}$ with suitable measures μ on $\partial\Omega$ and ν on $\partial\Lambda$ define the maps $\mathbf{t}^+, \mathbf{t}^-$ and decomposition $\partial\Omega = S_0 \cup S_1 \cup S_2$ as in Definition 3.5. Then*

- (i) *the map $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ is continuous;*

(ii) the sets S_0 and $S_0 \cup S_2 \subset \partial\Omega$ are both closed; and

(iii) the map $\mathbf{t}^- : S_0 \cup S_2 \rightarrow \partial\Lambda$ is continuous.

Furthermore, if the measures (μ, ν) are symmetrically suitable then

(iv) $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ is a homeomorphism with inverse map $\mathbf{s}^+ = (\mathbf{t}^+)^{-1}$; and

(v) $\mathbf{t}^- : S_i \rightarrow T_i$ is a homeomorphism with inverse map $\mathbf{s}^-|_{T_i}$ for $i = 2, 0$.

Here $\mathbf{s}^+, \mathbf{s}^-$ and the decomposition $\partial\Lambda = T_0 \cup T_1 \cup T_2$ are defined in Definition 3.6.

Proof: Let ψ and ϕ denote the convex functions of Definition 3.5, 3.6 and Remark 1.2. Strict convexity of the domains yields

$$\partial\Omega \cap \partial\psi(\mathbf{x}) = \{\mathbf{t}^+(\mathbf{x}), \mathbf{t}^-(\mathbf{x})\} \quad \text{for } \mathbf{x} \in \partial\Omega, \quad \text{while} \quad (39)$$

$$\partial\Omega \cap \partial\phi(\mathbf{y}) = \{\mathbf{s}^+(\mathbf{y}), \mathbf{s}^-(\mathbf{y})\} \quad \text{for } \mathbf{y} \in \partial\Lambda \quad (40)$$

if (μ, ν) are symmetrically suitable.

(i) *Continuity of \mathbf{t}^+ .* Take any sequence of points $\mathbf{x}_k \in \partial\Omega$ and set $\mathbf{y}_k = \mathbf{t}^+(\mathbf{x}_k)$. The definition 3.5 of \mathbf{t}^+ shows there exist a corresponding sequence of normals $(\mathbf{n}_k, \mathbf{q}_k) \in N_\Omega(\mathbf{x}_k) \times N_\Lambda(\mathbf{y}_k)$ such that $\mathbf{n}_k \cdot \mathbf{q}_k \geq 0$. Compactness of $\partial\Omega$, $\partial\Lambda$, and the unit sphere \mathbf{S}^d ensure that a subsequence, also denoted \mathbf{x}_k , can be extracted for which $(\mathbf{x}_k, \mathbf{y}_k, \mathbf{n}_k, \mathbf{q}_k)$ converge to limits denoted $(\mathbf{x}, \mathbf{y}, \mathbf{n}, \mathbf{q}) \in \partial\Omega \times \partial\Lambda \times \mathbf{S}^d \times \mathbf{S}^d$ respectively. It follows that $(\mathbf{n}, \mathbf{q}) \in N_\Omega(\mathbf{x}) \times N_\Lambda(\mathbf{y})$. Since $(\mathbf{x}_k, \mathbf{t}^+(\mathbf{x}_k))$ belong to the closed set $\partial\psi$, we have $(\mathbf{x}, \mathbf{y}) \in \partial\psi$ as well. Now $\mathbf{y} \in \partial\Omega \cap \partial\psi(\mathbf{x})$ implies $\mathbf{y} = \mathbf{t}^\pm(\mathbf{x})$ by (39). Since $\mathbf{n} \cdot \mathbf{q} \geq 0$ we know that $\mathbf{y} = \mathbf{t}^+(\mathbf{x})$. This concludes the demonstration of continuity: $\mathbf{t}^+(\mathbf{x}_k) \rightarrow \mathbf{y} = \mathbf{t}^+(\mathbf{x})$ whenever $\mathbf{x}_k \rightarrow \mathbf{x}$.

(ii) *Compactness of S_0 .* Start with a sequence $\mathbf{x}_k \in S_0 \subset \partial\Omega$ and set $\mathbf{y}_k = \mathbf{t}^+(\mathbf{x}_k)$. The definition of S_0 yields a sequence of normals $(\mathbf{n}_k, \mathbf{q}_k) \in N_\Omega(\mathbf{x}_k) \times N_\Lambda(\mathbf{y}_k)$ with $\mathbf{n}_k \cdot \mathbf{q}_k = 0$. Repeating the argument in (i) above and using continuity of \mathbf{t}^+ , we extract a convergent subsequence $(\mathbf{x}_k, \mathbf{y}_k, \mathbf{n}_k, \mathbf{q}_k) \rightarrow (\mathbf{x}, \mathbf{t}^+(\mathbf{x}), \mathbf{n}, \mathbf{q}) \in \partial\Omega \times \partial\Lambda \times N_\Omega(\mathbf{x}) \times N_\Lambda(\mathbf{t}^+(\mathbf{x}))$. This time $\mathbf{n} \cdot \mathbf{q} = \lim \mathbf{n}_k \cdot \mathbf{q}_k = 0$, which proves that $\mathbf{x} \in S_0$. Thus S_0 is compact.

(iii) *Compactness of $S_0 \cup S_2$ and continuity of \mathbf{t}^- .* Take any sequence of points $\mathbf{x}_k \in S_0 \cup S_2$ and mimic the preceding proofs, this time setting $\mathbf{y}_k = \mathbf{t}^-(\mathbf{x}_k)$. The definition of \mathbf{t}^- yields a sequence of normals $(\mathbf{n}_k, \mathbf{q}_k) \in N_\Omega(\mathbf{x}_k) \times N_\Lambda(\mathbf{y}_k)$ satisfying $\mathbf{n}_k \cdot \mathbf{q}_k \leq 0$. Extract a convergent subsequence $(\mathbf{x}_k, \mathbf{y}_k, \mathbf{n}_k, \mathbf{q}_k) \rightarrow (\mathbf{x}, \mathbf{y}, \mathbf{n}, \mathbf{q}) \in \partial\Omega \times \partial\Lambda \times N_\Omega(\mathbf{x}) \times N_\Lambda(\mathbf{y})$ as in (i) above. Noting that $(\mathbf{x}, \mathbf{y}) \in \partial\psi$ follows from $(\mathbf{x}_k, \mathbf{t}^-(\mathbf{x}_k)) \in \partial\psi$, we conclude $\mathbf{y} = \mathbf{t}^\pm(\mathbf{x})$ from (39). Finally, since $\mathbf{n}_k \cdot \mathbf{q}_k \leq 0$ either $\mathbf{n} \cdot \mathbf{q} < 0$, in which case $\mathbf{x} \in S_2$ and $\mathbf{y} = \mathbf{t}^-(\mathbf{x})$; or $\mathbf{n} \cdot \mathbf{q} = 0$, in which case $\mathbf{x} \in S_0$ and $\mathbf{y} = \mathbf{t}^-(\mathbf{x}) = \mathbf{t}^+(\mathbf{x})$. Either way, the limit point $\mathbf{x} \in S_0 \cup S_2$ demonstrates that $S_0 \cup S_2$ is compact, while $\mathbf{x}_k \rightarrow \mathbf{x}$ implies $\mathbf{t}^-(\mathbf{x}_k) \rightarrow \mathbf{y} = \mathbf{t}^-(\mathbf{x})$ so \mathbf{t}^- is continuous.

(iv) *$\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ is a homeomorphism.* We have already proved \mathbf{t}^+ is continuous, and since the hypotheses are symmetrical under interchange of $(\Omega, \mu, \psi, \mathbf{t}^\pm)$ with $(\Lambda, \nu, \phi, \mathbf{s}^\pm)$ we also have continuity of $\mathbf{s}^+ : \partial\Lambda \rightarrow \partial\Omega$. If we can prove that $\mathbf{s}^+(\mathbf{t}^+(\mathbf{x})) = \mathbf{x}$ on $\partial\Omega$ then we will have shown \mathbf{t}^+ is one-to-one and continuously invertible while \mathbf{s}^+ maps onto $\partial\Omega$. By symmetry, $\mathbf{t}^+(\mathbf{s}^+(\mathbf{y})) = \mathbf{y}$ on $\partial\Lambda$ and it follows

that $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ is a bijection with inverse \mathbf{s}^+ . Establishing (iv) is thus reduced to verifying the following claim.

Claim: $\mathbf{s}^+(\mathbf{t}^+(\mathbf{x})) = \mathbf{x}$ for all $\mathbf{x} \in \partial\Omega$.

Proof of Claim: Fix $\mathbf{x} \in \partial\Omega = \text{spt } \mu$ and set $\mathbf{y} = \mathbf{t}^+(\mathbf{x})$. From Definition 3.5, \mathbf{y} is uniquely characterized among points in $\partial\Lambda \cap \partial\psi(\mathbf{x})$ by the existence of $\mathbf{n} \in N_\Omega(\mathbf{x})$ and $\mathbf{q} \in N_\Lambda(\mathbf{y})$ such that $\mathbf{n} \cdot \mathbf{q} \geq 0$. Since $\mathbf{x} \in \text{spt } \mu$ we conclude $\mathbf{y} \in \partial\phi^*(\mathbf{x})$ from Remark 1.2. Thus $\mathbf{x} \in \partial\Omega \cap \partial\phi(\mathbf{y})$ so $\mathbf{x} = \mathbf{s}^\pm(\mathbf{y})$ by (40). Since $\mathbf{n} \cdot \mathbf{q} \geq 0$ we conclude $\mathbf{x} = \mathbf{s}^+(\mathbf{y})$ to establish the claim.

(v) $\mathbf{t}^- : S_i \rightarrow T_i$ is a homeomorphism for $i = 0, 2$. We have already shown continuity of $\mathbf{t}^- : S_i \rightarrow \partial\Lambda$ and by symmetry $\mathbf{s}^- : T_i \rightarrow \partial\Omega$. The result will be established by the next claim, which implies $\mathbf{t}^-|_{S_i}$ is one-to-one and its continuous inverse \mathbf{s}^- maps T_i onto S_i . By symmetry, it follows that \mathbf{t}^- maps S_i onto T_i and hence is a homeomorphism.

Claim: $\mathbf{t}^-(\mathbf{x}) \in T_i$ and $\mathbf{s}^-(\mathbf{t}^-(\mathbf{x})) = \mathbf{x}$ whenever $\mathbf{x} \in S_i$ for $i = 2, 0$.

Proof of Claim: Fix $\mathbf{x} \in S_i$ and set $\mathbf{y} = \mathbf{t}^-(\mathbf{x})$. For $i = 2, 0$, \mathbf{y} is characterized uniquely among points in $\partial\Lambda \cap \partial\psi(\mathbf{x})$ by the existence of $\mathbf{n} \in N_\Omega(\mathbf{x})$ and $\mathbf{q} \in N_\Lambda(\mathbf{y})$ such that $\mathbf{n} \cdot \mathbf{q} \leq 0$ according to Definition 3.5. Since $\mathbf{x} \in \partial\Omega = \text{spt } \mu$ we conclude $\mathbf{y} \in \partial\phi^*(\mathbf{x})$ from Remark 1.2. Now $\mathbf{x} \in \partial\Omega \cap \partial\phi(\mathbf{y})$ so $\mathbf{x} = \mathbf{s}^\pm(\mathbf{y})$ by (40). Since $\mathbf{n} \cdot \mathbf{q} \leq 0$ we conclude $\mathbf{x} = \mathbf{s}^-(\mathbf{y})$ and $\mathbf{y} \in T_0 \cup T_2$. If $\mathbf{x} \in S_0$ then $\mathbf{n} \cdot \mathbf{q} = 0$ and $\mathbf{y} \in T_0$, while if $\mathbf{x} \in S_2$ then $\mathbf{n} \cdot \mathbf{q} < 0$ and $\mathbf{y} \in T_2$. This concludes the proof that $\mathbf{x} = \mathbf{s}^-(\mathbf{y})$ and $\mathbf{y} = \mathbf{t}^-(\mathbf{x}) \in T_i$ for $i = 2, 0$. \square

The next theorem collects together and summarizes our results.

Theorem 3.8 (Twin Homeomorphisms Optimize) *Fix bounded, strictly convex domains $\Omega, \Lambda \subset \mathbf{R}^{d+1}$ with suitable measures μ on $\partial\Omega$ and ν on $\partial\Lambda$ (Definition 3.1). Then the graphs of a pair of continuous maps $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ and $\mathbf{t}^- : \overline{S}_2 \rightarrow \overline{T}_2$ contain the support of all minimizers $\gamma \in \Gamma(\mu, \nu)$ for (2):*

$$\{(\mathbf{x}, \mathbf{t}^+(\mathbf{x}))\}_{\mathbf{x} \in \text{spt } \mu} \subset \text{spt } \gamma \subset \{(\mathbf{x}, \mathbf{t}^+(\mathbf{x}))\}_{\mathbf{x} \in \partial\Omega} \cup \{(\mathbf{x}, \mathbf{t}^-(\mathbf{x}))\}_{\mathbf{x} \in S_2} \quad (41)$$

$$(\text{=} \partial\psi \cap (\partial\Omega \times \partial\Lambda)).$$

Here ψ is from Theorem 1.1, $S_2 = \partial\Omega \setminus \text{dom } \nabla\psi$, $T_2 \subset \partial\Lambda$, and $\mathbf{t}^+ = \mathbf{t}^-$ on $\overline{S}_2 \setminus S_2$, whereas $\mathbf{t}^+(\mathbf{x}) - \mathbf{t}^-(\mathbf{x}) \neq \mathbf{0}$ is an outward normal for $\partial\Omega$ at \mathbf{x} whenever $\mathbf{x} \in S_2$. Furthermore, \mathbf{t}^+ and $\mathbf{t}^-|_{\overline{S}_2}$ are homeomorphisms if (μ, ν) are symmetrically suitable.

Proof: Take ψ from Theorem 1.1, so that any minimizer γ of (2) is supported on the closed set $\partial\psi \cap (\partial\Omega \times \partial\Lambda)$. Define $\mathbf{t}^\pm : \partial\Omega \rightarrow \partial\Lambda$ and $S_2 \subset \partial\Omega$ using Definition 3.5. The alternatives in Proposition 3.4 make it clear that S_2 consists precisely of those points $\mathbf{x} \in \partial\Omega$ where ψ is not differentiable, and that $\mathbf{t}^+(\mathbf{x}) \neq \mathbf{t}^-(\mathbf{x})$ for $\mathbf{x} \in S_2$ while $\mathbf{t}^+ = \mathbf{t}^-$ elsewhere. From Proposition 3.2, $\mathbf{t}^+(\mathbf{x}) - \mathbf{t}^-(\mathbf{x})$ is normal to $\partial\Omega$ at $\mathbf{x} \in S_2$; it is an outward normal by Lemma 3.3. Continuity of \mathbf{t}^+ and $\mathbf{t}^-|_{\overline{S}_2}$ follow from Proposition 3.7(i)–(iii); one may take $T_2 = \mathbf{t}^-(S_2)$ in which case $\overline{T}_2 = \mathbf{t}^-(\overline{S}_2)$ follows from compactness of $\overline{S}_2 \subset \partial\Omega$. If the measures are symmetrically suitable

then Proposition 3.7(iv) shows \mathbf{t}^+ to be a homeomorphism, while (ii), (iii) and (v) combine to show $\mathbf{t}^- : S_0 \cup S_2 \rightarrow T_0 \cup T_2$ is a homeomorphism of the closed set $S_0 \cup S_2 \supset \overline{S_2}$. It remains only to verify the next two claims to establish (41) and the theorem.

1. Claim: $\partial\psi \cap (\partial\Omega \times \partial\Lambda) = \{(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \mid \mathbf{x} \in \partial\Omega\} \cup \{(\mathbf{x}, \mathbf{t}^-(\mathbf{x})) \mid \mathbf{x} \in S_2\}$.

Proof of Claim: If $\mathbf{x} \in \partial\Omega$ then $\partial\psi(\mathbf{x}) = [\mathbf{t}^+(\mathbf{x}), \mathbf{t}^-(\mathbf{x})]$ from Definition 3.5, while the extreme points $\mathbf{t}^\pm(\mathbf{x})$ lie on $\partial\Lambda$ according to Corollary 1.4. Thus $(\mathbf{x}, \mathbf{t}^\pm(\mathbf{x})) \in \partial\psi \cap (\partial\Omega \times \partial\Lambda)$, so the graphs of both \mathbf{t}^+ and \mathbf{t}^- are contained in $\partial\psi \cap (\partial\Omega \times \partial\Lambda)$. Conversely, if $(\mathbf{x}, \mathbf{y}) \in \partial\psi \cap (\partial\Omega \times \partial\Lambda)$ then strict convexity of $\overline{\Lambda} \supset \partial\psi(\mathbf{x})$ forces $\mathbf{y} \in \partial\Lambda$ to be an extreme point of $\partial\psi(\mathbf{x})$. This means $\mathbf{y} = \mathbf{t}^\pm(\mathbf{x})$. If $\mathbf{x} \in \partial\Omega \setminus S_2$ then $\mathbf{t}^-(\mathbf{x}) = \mathbf{t}^+(\mathbf{x}) = \mathbf{y}$ so $(\mathbf{x}, \mathbf{y}) \in \text{graph}(\mathbf{t}^+)$; otherwise $\mathbf{x} \in S_2$ so (\mathbf{x}, \mathbf{y}) belongs either to $\text{graph}(\mathbf{t}^+)$ or $\text{graph}(\mathbf{t}^-|_{S_2})$. This shows the union of the two graphs contains $\partial\psi \cap (\partial\Omega \times \partial\Lambda)$ thereby completing the claim.

2. Claim: If $\mathbf{x} \in \text{spt } \mu$ then $(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \in \text{spt } \gamma$ when $\gamma \in \Gamma(\mu, \nu)$ minimizes (2).

Proof of Claim: Fix $\mathbf{x} \in \text{spt } \mu$ and let γ minimize (2). Since γ is compactly supported one can find $\mathbf{y} \in \text{spt } \nu$ with $(\mathbf{x}, \mathbf{y}) \in \text{spt } \gamma \subset \partial\psi$. Then either $\mathbf{y} = \mathbf{t}^+(\mathbf{x})$ or else $\mathbf{x} \in S_2$ by Claim 1. In the first case we have $(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \in \text{spt } \gamma$ and are done, so therefore assume $\mathbf{x} \in S_2$. From Proposition 3.4 we know $\partial\Omega$ is differentiable at $\mathbf{x} \in S_2$, $(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \in \partial\psi$, and ψ is tangentially differentiable along $\partial\Omega$ (c.f. Corollary 4.4). Thus Lemma 2.5 yields (i) $\overline{\Omega} \cap \partial\psi^*(\mathbf{t}^+(\mathbf{x})) = \{\mathbf{x}\}$ since $\mathbf{t}^+(\mathbf{x}) \neq \mathbf{t}^-(\mathbf{x})$ on S_2 . On the other hand, since $\mathbf{t}^+(\mathbf{x}) \in \partial\Lambda = \text{spt } \nu$, one can find $\mathbf{z} \in \text{spt } \mu \subset \partial\Omega$ such that $(\mathbf{z}, \mathbf{t}^+(\mathbf{x})) \in \text{spt } \gamma \subset \partial\psi$. Then (i) forces $\mathbf{z} = \mathbf{x}$ to establish the claim and the theorem. \square

Remark 3.9 (Continuous Extension of \mathbf{t}^-) *If Ω has a C^1 smooth boundary, then a continuous map $\tilde{\mathbf{t}} : \partial\Omega \rightarrow \partial\Lambda$ can be defined to be the point at the opposite end of the segment $L \cap \overline{\Lambda} = [\mathbf{t}^+(\mathbf{x}), \tilde{\mathbf{t}}(\mathbf{x})]$ where the line $L := \{\mathbf{t}^+(\mathbf{x}) + \lambda \mathbf{n}_\Omega(\mathbf{x}) \mid \lambda \in \mathbf{R}\}$ through $\mathbf{t}^+(\mathbf{x})$ intersects the compact convex set $\overline{\Lambda}$. Then $\tilde{\mathbf{t}}$ coincides with \mathbf{t}^- on $S_0 \cap S_2$, though we do not expect that $\tilde{\mathbf{t}}$ will remain one-to-one throughout $\partial\Omega$.*

Remark 3.10 (Another View) *For $\mathbf{x}_o \in S_1$ and any neighbourhood $U \subset \partial\Omega \setminus \overline{S_2}$ of \mathbf{x}_o , the restriction of γ to $U \times \mathbf{R}^{d+1}$ is concentrated on the graph of $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$. Indeed, (41) yields*

$$\gamma[\{(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \mid \mathbf{x} \in U\}] = \mu[U]. \quad (42)$$

On the other hand, if the measures (μ, ν) are symmetrically suitable, we can deduce $\mathbf{t}^+(S_2) \subset T_1$ by applying Lemma 2.5 and Corollary 4.4 to each $\mathbf{x} \in S_2$ and $\mathbf{y} = \mathbf{t}^+(\mathbf{x})$. Since \mathbf{t}^+ is a homeomorphism with inverse $\mathbf{s}^+ : \partial\Lambda \rightarrow \partial\Omega$ we have

$$\{(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \mid \mathbf{x} \in \mathbf{B}(\mathbf{x}_o, r) \cap \partial\Omega\} = \{(\mathbf{s}^+(\mathbf{y}), \mathbf{y}) \mid \mathbf{y} \in \mathbf{t}^+(\mathbf{B}(\mathbf{x}_o, r)) \subset \partial\Lambda\}. \quad (43)$$

Now if $\mathbf{x}_o \in S_2$, then $\mathbf{t}^+(\mathbf{x}_o) \in T_1$ so taking $r > 0$ very small ensures $\mathbf{t}^+(\mathbf{B}(\mathbf{x}_o, r))$ is disjoint from $\overline{T_2}$. Thus there exists $r(\mathbf{x}_o) > 0$ such that

$$\gamma[\{(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \mid \mathbf{x} \in \mathbf{B}(\mathbf{x}_o, r) \cap \partial\Omega\}] = \begin{cases} \mu[\mathbf{B}(\mathbf{x}_o, r)] & \text{if } \mathbf{x}_o \in S_1 \\ \nu[\mathbf{t}^+(\mathbf{B}(\mathbf{x}_o, r))] & \text{if } \mathbf{x}_o \in S_2 \end{cases} \quad (44)$$

holds for $0 < r < r(\mathbf{x}_o)$, where the first equality follows from $U = \mathbf{B}(\mathbf{x}_o, r) \cap \partial\Omega$ in (42) and the second from $U = \mathbf{t}^+(\mathbf{B}(\mathbf{x}_o, r))$ in (42–43) and the symmetry $\mu \leftrightarrow \nu$. Note that the mass γ assigns to the graph of \mathbf{t}^+ near \mathbf{x}_o is unambiguous in either case. Moreover, a limiting version of (44) can be shown to hold at μ -a.e. $\mathbf{x}_o \in S_0$:

$$\lim_{r \rightarrow 0} \frac{\gamma[\{(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \mid \mathbf{x} \in \mathbf{B}(\mathbf{x}_o, r) \cap \partial\Omega\}]}{\mu[\mathbf{B}(\mathbf{x}_o, r)]} = 1. \quad (45)$$

This provides another approach to the uniqueness of γ in Theorem 2.6.

Remark 3.11 (Uniqueness) *From Theorem 2.6 the minimizer γ of (2) is unique whenever μ is absolutely continuous with respect to $\mathcal{H}^d|_{\partial\Omega}$. For symmetrically suitable measures, that theorem goes on to express $\gamma = \gamma_1 + \gamma_2$ explicitly by*

$$\gamma_1 = (\mathbf{id} \times \mathbf{t}^+)_{\#}\mu_1 \quad \text{and} \quad \gamma_2 = (\mathbf{id} \times \mathbf{t}^-)_{\#}\mu_2, \quad (46)$$

where $\mu_1 := (\mathbf{t}^+)^{-1}_{\#}\nu_1$, $\mu_2 := \mu - \mu_1$, and we have defined the restriction $\nu_1 := \nu|_{T_2^c}$ of ν to $T_2^c := \partial\Omega \setminus \mathbf{t}^-(S_2)$ and extended \mathbf{t}^- by setting $\mathbf{t}^-(\mathbf{x}) := \mathbf{t}^+(\mathbf{x})$ for $\mathbf{x} \in \partial\Omega \setminus S_2$. Note that μ_2 can be replaced by $\mu_2|_{S_2}$ in (46) if we also replace μ_1 by $\mu_1 + \mu_2|_{S_2^c}$.

The two maps \mathbf{t}^+ and \mathbf{t}^- provide just enough flexibility to allow the optimum (2) to be attained. Indeed, for convex sets, $\mathbf{t}^+(\mathbf{x})$ serves as the primary destination — accepting as much mass as is compatible with ν — while \mathbf{t}^- acts as a kind of auxiliary or overflow map. We interpret this second map as providing the necessary slack in a few key places to permit much of the mass of μ to be transported onto $\partial\Omega$ extremely efficiently. The following two examples show our conclusions are sharp in the sense that the support of the optimal measure cannot be contained in the graph of a single map. The second example differs from the first in that its two triangular shapes may be taken to be arbitrarily close to each other.

Example 3.12 (Reflected Triangles) *Let Ω^* be the interior of the triangle ΔIJK with vertices $I := (-a, 1)$, $J := (0, 3)$, and $K := (a, 1)$, and Λ^* its reflection $\Delta I'J'K'$ in the x -axis (Fig. 4(a)). Although our theorems apply only to strictly convex domains, these two triangles can be approximated arbitrarily well by strictly convex shapes formed by joining the same old vertices, this time with circular arcs of very small curvature $\epsilon > 0$ instead of straight lines (Fig. 4(b)). Let Ω and $\Lambda \subset \mathbf{R}^2$ denote these slightly perturbed strictly convex domains. For all \mathbf{x} on the sides $\overline{IJ} \cup \overline{JK}$ and all \mathbf{y} on the sides $\overline{I'J'} \cup \overline{J'K'}$ of the original triangles we have*

$$\mathbf{n}_{\Omega^*}(\mathbf{x}) \cdot \mathbf{n}_{\Lambda^*}(\mathbf{y}) \leq -\frac{a^2 - 4}{a^2 + 4}, \quad (47)$$

(excluding the vertices where outward unit normals are not uniquely defined). Also,

$$\mathcal{H}^1[\overline{IJ} \cup \overline{JK}] = \mathcal{H}^1[\overline{I'J'} \cup \overline{J'K'}] = 2\sqrt{a^2 + 4} \quad (48)$$

and

$$\mathcal{H}^1[\overline{TK}] = 2a. \quad (49)$$

Now fix $a > 2$ and $0 < \delta < (\sqrt{a^2 + 4} - a)/2$. Taking the perturbation $\epsilon > 0$ small enough, (47)–(49) imply that we can find an arc $A \subset \partial\Omega$ and its reflection $A' \subset \partial\Lambda$ on the boundaries of the perturbed domains such that

$$\mathbf{n}_\Omega(\mathbf{x}) \cdot \mathbf{n}_\Lambda(\mathbf{y}) < 0 \quad (50)$$

for all $\mathbf{x} \in A$ and all $\mathbf{y} \in A'$. We may assume in addition that

$$\mathcal{H}^1[A] = \mathcal{H}^1[A'] > 2\sqrt{a^2 + 4} - \delta, \quad (51)$$

and

$$\mathcal{H}^1[\partial\Omega] = \mathcal{H}^1[\partial\Lambda] < 2\sqrt{a^2 + 4} + 2a + 2\delta. \quad (52)$$

Combining (51) and (52) we have

$$\mathcal{H}^1[A] > \mathcal{H}^1[\partial\Lambda \setminus A']. \quad (53)$$

Finally, let $\mu = \mathcal{H}^1|_{\partial\Omega}$ and $\nu = \mathcal{H}^1|_{\partial\Lambda}$ measure arclength around the boundaries of Ω and Λ respectively. We claim (53) precludes the optimal measure γ in (2) from being supported on the graph of a single function. Assume on the contrary $\text{spt } \gamma$ is contained in the graph of a single map \mathbf{t} from $\partial\Omega$ into $\partial\Lambda$. From (41), the homeomorphism \mathbf{t}^+ has its graph in $\text{spt } \gamma$ so we must have $\mathbf{t} = \mathbf{t}^+$. In light of (50) and Definition 3.5 we see

$$\mathbf{t}^+[A] \subset \partial\Lambda \setminus A',$$

and so

$$\mathcal{H}^1[A] \leq \mathcal{H}^1[\partial\Lambda \setminus A'],$$

which is at a variance with (53). Consequently, in this example no optimal measure for (2) is supported on the graph of a single function.

Example 3.13 (Triangles which Differ Only Slightly) Take a pair of isosceles triangles $\triangle LMN$ and $\triangle L'M'N'$ with different side lengths but the same perimeter; e.g. the first triangle shorter and squatter than the second: $|LN| > |L'N'|$ in Fig. 5. Again our theorems apply only to strictly convex shapes, but a slight rounding of corners and edges will make the triangles smooth and strictly convex without changing their geometry significantly; i.e., the scale ϵ of the perturbation is kept very small with respect to the differences in side lengths, while reflection symmetry around a vertical axis is preserved. Let Ω and $\Lambda \subset \mathbf{R}^2$ denote these perturbed domains, rescaled if necessary to maintain the same (e.g. unit) length of perimeter. Take $\mu = \mathcal{H}^1|_{\partial\Omega}$ and $\nu = \mathcal{H}^1|_{\partial\Lambda}$ to measure arclength around their boundaries. By Theorems 2.6 and 3.8, the optimal measure γ is unique in (2), and its support contains the graph of a homeomorphism $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ while being contained in the union of that graph

with the graph of a second homeomorphism $\mathbf{t}^- : \overline{S}_2 \subset \partial\Omega \rightarrow \overline{T}_2 \subset \partial\Lambda$. The next claim demonstrates that the second homeomorphism is necessary; it implies that both graphs in (41) carry positive mass.

Claim: In this example, $\{(\mathbf{x}, \mathbf{t}^+(\mathbf{x})) \mid \mathbf{x} \in \partial\Omega\}$ is a strict subset of $\text{spt } \gamma$.

Proof of Claim: Since $\partial\Omega = \text{spt } \mu$, non-strict containment is asserted by (41). To derive a contradiction, suppose strict containment fails, i.e. $\text{spt } \gamma = \text{graph}(\mathbf{t}^+)$. Then the homeomorphism \mathbf{t}^+ must preserve arclengths because $\nu = \mathbf{t}_\#^+ \mu$ follows from $\gamma \in \Gamma(\mu, \nu)$ and Lemma 2.4; it gives an arclength reparameterization of $\partial\Lambda$ by $\partial\Omega$. Denoting the midpoints of the bases of the (perturbed) triangles by $\mathbf{x}_0 \in \partial\Omega$ and $\mathbf{y}_0 \in \partial\Lambda$ respectively, we next use symmetry to establish that $\mathbf{y}_0 = \mathbf{t}^+(\mathbf{x}_0)$. Indeed, $\mathbf{t}^+(\mathbf{x}_0)$ and $\mathbf{t}^-(\mathbf{x}_0)$ are the only two points paired with \mathbf{x}_0 in $\text{spt } \gamma$. By Lemma 3.2 they lie in a vertical line. Since γ is unique this line can only be the symmetry axis, so either $\mathbf{t}^+(\mathbf{x}_0) = \mathbf{y}_0$ or $\mathbf{t}^+(\mathbf{x}_0) = M'$. But the latter possibility is precluded since \mathbf{t}^+ was defined so $\mathbf{n}_\Omega(\mathbf{x}) \cdot \mathbf{n}_\Lambda(\mathbf{t}^+(\mathbf{x})) \geq 0$.

Having shown \mathbf{t}^+ to be arclength parameterization of $\partial\Lambda$ starting from $\mathbf{t}^+(\mathbf{x}_0) = \mathbf{y}_0$, take two points \mathbf{x}_1 and \mathbf{x}_2 on $\partial\Omega$ near N positioned so that $|\mathbf{y}_0 N'| + \epsilon < |\mathbf{x}_0 \mathbf{x}_1| < |\mathbf{x}_0 \mathbf{x}_2| < |\mathbf{x}_0 N| - \epsilon$ as in Fig. 5. The first inequality forces $\mathbf{y}_i := \mathbf{t}^+(\mathbf{x}_i)$ to lie between N' and M' on $\partial\Lambda$ since \mathbf{t}^+ is length preserving. The ϵ 's keep the \mathbf{x}_i and \mathbf{y}_i ($i = 0, 2$) away from the smoothed corners N and N' (the scale ϵ of the rounding and smoothing being too small to be detectable in Fig. 5). Apart from corrections which become negligible for ϵ very small, $\mathbf{x}_2 - \mathbf{x}_1$ points horizontally to the right while $\mathbf{y}_2 - \mathbf{y}_1$ points upwards and to the left. But this violates the monotonicity (4) of $\text{spt } \gamma$ which follows from Theorem 1.1. This contradiction establishes the claim. \square

4 Continuous Maps Between Convex Hypersurfaces

This final section contains the proof of tangential differentiability when $d > 1$ for the convex function ψ of Theorem 1.1 and suitable measures μ and ν on \mathbf{R}^{d+1} . The hypotheses (Definition 3.1) and conclusions should be compared with Caffarelli's theory which applies in a distinct though related setting [5][6]. The key result is Proposition 4.3, which shows that as long as $\gamma \in \Gamma(\mu, \nu)$ has monotone support, then the images in $\partial\Lambda = \text{spt } \nu$ of each $\mathbf{x}_o \in \text{spt } \mu$ will be collinear: they parallel the outer normals $\mathbf{n}_o \in N_\Omega(\mathbf{x}_o)$. When Theorem 3.8 and Corollary 4.4 are derived, it is here that the bounds (35) and the uniform convexity of Ω are invoked. As a preliminary lemma shows, the polar dual cone defined below is a useful construction for probing the constraints imposed by monotonicity on the area of $\partial\Lambda \cap \partial\psi(\mathbf{x}_o)$.

Definition 4.1 (Polar Dual) Given $Y \subset \mathbf{R}^{d+1}$ we define its polar dual cone Y^* by

$$Y^* := \{ \mathbf{x} \in \mathbf{R}^{d+1} \mid \mathbf{x} \cdot \mathbf{y} < 0 \text{ for all } \mathbf{y} \in Y \}.$$

In the sequel we often take Y to be the cone

$$Y(\theta) := \{ \mathbf{y} \in \mathbf{R}^{d+1} \mid \mathbf{y} \cdot \mathbf{e}_1 < -|\mathbf{y}| \cos \theta \}$$

where $\mathbf{e}_1 := (1, 0, \dots, 0) \in \mathbf{R}^{d+1}$. One then readily checks that for $\theta \in (0, \frac{\pi}{2})$,

$$Y(\theta)^* = \{ \mathbf{x} \in \mathbf{R}^{d+1} \mid 0 < |\mathbf{x}| \sin \theta \leq \mathbf{x} \cdot \mathbf{e}_1 \}. \quad (54)$$

Lemma 4.2 *Choose any $Y \subset \mathbf{R}^{d+1}$ and a monotone set $M \subset \mathbf{R}^{d+1} \times \mathbf{R}^{d+1}$. For each $\mathbf{x}_o \in \mathbf{R}^{d+1}$ the set $[\mathbf{x}_o + Y^*] \times [M(\mathbf{x}_o) + Y]$ will be disjoint from M .*

Proof: Assume on the contrary that there exist $\mathbf{x} \in Y^*$, $\mathbf{y} \in Y$ and $(\mathbf{x}_o, \mathbf{y}_o) \in M$ such that

$$(\mathbf{x}_o + \mathbf{x}, \mathbf{y}_o + \mathbf{y}) \in M.$$

Since $(\mathbf{x}_o, \mathbf{y}_o)$ also belongs to M , monotonicity (4) yields $\mathbf{x} \cdot \mathbf{y} \geq 0$. But then $\mathbf{x} \in Y^*$ and $\mathbf{y} \in Y$ contradict Definition 4.1, concluding the proof of the lemma. \square

Proposition 4.3 (Images Still Parallel Normal If $d > 1$) *Let $\Omega, \Lambda \subset \mathbf{R}^{d+1}$ be bounded, strictly convex domains with Borel probability measures μ on $\partial\Omega$ and ν on $\partial\Lambda$. Assume Ω is uniformly convex, while*

$$\mu < \epsilon^{-1} \mathcal{H}^d \llcorner_{\partial\Omega} \quad \text{and} \quad \nu > \epsilon \mathcal{H}^d \llcorner_{\partial\Lambda} \quad (55)$$

holds for some $\epsilon > 0$. Suppose a joint measure γ with marginals μ and ν has support contained in a monotone set $M \subset \mathbf{R}^{d+1} \times \bar{\Lambda}$. If $\mathbf{n}_o \in \mathbf{R}^{d+1}$ is any normal of Ω at $\mathbf{x}_o \in \partial\Omega$ and both $\mathbf{y}_o, \mathbf{y} \in M(\mathbf{x}_o)$, then $\mathbf{y}_o - \mathbf{y}$ must be parallel to \mathbf{n}_o .

Proof: It costs no generality to replace the monotone set M in the proposition by a maximal monotone set which contains it; i.e. a set $M \subset \mathbf{R}^{d+1} \times \bar{\Lambda}$ which cannot be enlarged without violating monotonicity (4). This assures convexity of $M(\mathbf{x}) := \{ \mathbf{y} \in \mathbf{R}^{d+1} \mid (\mathbf{x}, \mathbf{y}) \in M \}$ by e.g. Alberti & Ambrosio [2, Proposition 1.2].

1. Take \mathbf{n}_o normal to $\partial\Omega$ at \mathbf{x}_o . Translating and rotating coordinates if necessary, Lemma A.1 shows $\partial\Omega$ may be parameterized by a function $x_1 = k(x_2, \dots, x_{d+1}) \geq 0$ near $\mathbf{x}_o = \mathbf{0}$. In particular, $\mathbf{n} = -\mathbf{e}_1 := (-1, 0, \dots, 0)$ will be an outer normal to $\partial\Omega$ at \mathbf{x}_o (and coincides with \mathbf{n}_o unless differentiability of $\partial\Omega$ fails). Let $\theta \in (0, \frac{\pi}{2})$ and $Y := Y(\theta)$. Since $\gamma \in \Gamma(\mu, \nu)$ is supported on M , Lemma 4.2 yields

$$\begin{aligned} \nu[M(\mathbf{0}) + Y] &= \gamma[Y^{*c} \times (M(\mathbf{0}) + Y)] \\ &\leq \mu[Y^{*c}]. \end{aligned} \quad (56)$$

This combines with (55) to show the ratio of surface areas

$$\rho(\theta) := \frac{\mathcal{H}^d \llcorner_{\partial\Omega}[Y^{*c}(\theta)]}{\mathcal{H}^d \llcorner_{\partial\Lambda}[M(\mathbf{0}) + Y(\theta)]} \quad (57)$$

remains bounded away from zero: $\rho(\theta) > \epsilon^2 > 0$.

2. Let $\pi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}^d$ denote the orthogonal projection

$$\pi(x_1, x_2, \dots, x_{d+1}) = (x_2, \dots, x_{d+1})$$

and let $\mathbf{y}_o, \mathbf{y}_1 \in M(\mathbf{0}) \subset \bar{\Lambda}$. Our goal is to show that for small θ , the bound $\rho(\theta) > \epsilon^2$ forces $\pi(\mathbf{y}_o) = \pi(\mathbf{y}_1)$. This will complete the proof that $\mathbf{y}_o - \mathbf{y}_1$ is parallel to \mathbf{n} .

To derive a contradiction, suppose $\pi(\mathbf{y}_o) \neq \pi(\mathbf{y}_1)$. The convexity of $M(\mathbf{0})$ and strict convexity of Λ imply a slightly shorter segment, also denoted $[\mathbf{y}_o, \mathbf{y}_1]$, lies strictly inside Λ . Denote the distance from this segment to the boundary by

$$\delta := \inf_{\mathbf{y}, \mathbf{z}} \{|\mathbf{y} - \mathbf{z}| \mid \mathbf{y} \in [\mathbf{y}_o, \mathbf{y}_1], \mathbf{z} \in \partial\Lambda\}. \quad (58)$$

To parametrize the bottom of $\partial\Lambda$ define the convex function $h : \mathbf{R}^d \rightarrow \mathbf{R} \cup \{+\infty\}$ by

$$h(\mathbf{p}) = \inf_{\mathbf{y}} \{\mathbf{e}_1 \cdot \mathbf{y} \mid \mathbf{y} \in \bar{\Lambda}, \pi(\mathbf{y}) = \mathbf{p} \in \mathbf{R}^d\}.$$

3. *Claim:* If the distance from \mathbf{p} to the segment $\pi[\mathbf{y}_o, \mathbf{y}_1]$ is less than $\delta \sin \theta$ then $(h(\mathbf{p}), \mathbf{p}) \in M(\mathbf{0}) + Y$.

Proof of Claim: To begin note that the function $f(\lambda, \tau) = \lambda/\sqrt{\lambda^2 + \tau^2}$ increases with $\lambda \in \mathbf{R}$ for every fixed τ , and that $\mathbf{z} = (z_1, \dots, z_{d+1})$ belongs to the cone $Y := Y(\theta)$ precisely when

$$f(z_1, |\pi(\mathbf{z})|) < -\cos \theta \quad \text{and} \quad \mathbf{z} \neq \mathbf{0}. \quad (59)$$

Now if $\mathbf{y} \in [\mathbf{y}_o, \mathbf{y}_1]$ and

$$|\mathbf{p} - \pi(\mathbf{y})| < \delta \sin \theta \quad (60)$$

then $\mathbf{z} := (-\delta \cos \theta, \mathbf{p} - \pi(\mathbf{y}))$ satisfies $|\mathbf{z}| \cos \theta < -\mathbf{z} \cdot \mathbf{e}_1$ and belongs to the cone Y . The definition of δ and (60) yield

$$\mathbf{y} + \mathbf{z} \in \mathbf{B}(\mathbf{y}, \delta) \subset \Lambda. \quad (61)$$

Note that $\pi(\mathbf{y} + \mathbf{z}) = \mathbf{p}$ and so $h(\mathbf{p}) < y_1 - \delta \cos \theta$. Monotonicity of $f(\cdot, \tau)$ then gives

$$f(h(\mathbf{p}) - y_1, |\mathbf{p} - \pi(\mathbf{y})|) \leq f(z_1, |\pi(\mathbf{z})|)$$

which combines with $\mathbf{z} \in Y$ and (59) to yield $(h(\mathbf{p}) - y_1, \mathbf{p} - \pi(\mathbf{y})) \in Y$. Thus $(h(\mathbf{p}), \mathbf{p}) \in \mathbf{y} + Y$ establishes the claim.

4. *Claim:* If $h(\mathbf{p}) \neq +\infty$ then $(h(\mathbf{p}), \mathbf{p}) \in \partial\Lambda$.

Proof of Claim: Let $\lambda_n \subset \mathbf{R}$ be a non-increasing sequence such that

$$(\lambda_n, \mathbf{p}) \in \bar{\Lambda} \quad \text{and} \quad h(\mathbf{p}) = \lim_{n \rightarrow +\infty} \lambda_n.$$

We have $(h(\mathbf{p}), \mathbf{p}) \in \bar{\Lambda}$ and $(h(\mathbf{p}), \mathbf{p}) \in \Lambda$ does not occur since otherwise the latter yields $(h(\mathbf{p}) - r, \mathbf{p}) \in \Lambda$, for some $r > 0$ and contradicts the definition of $h(\mathbf{p})$.

5. Combining Claims 3 and 4 we obtain

$$\{(h(\mathbf{p}), \mathbf{p}) \mid \mathbf{p} \in G\} \subset [M(\mathbf{0}) + Y] \cap \partial\Lambda, \quad (62)$$

where G is the set of all $\mathbf{p} \in \mathbf{R}^d$ satisfying $\text{dist}(\mathbf{p}, \pi[\mathbf{y}_o, \mathbf{y}_1]) < \delta \sin \theta$. By the area formula

$$\mathcal{H}^d[\{(h(\mathbf{p}), \mathbf{p}) \mid \mathbf{p} \in G\}] = \int_G \sqrt{1 + |\nabla h(\mathbf{p})|^2} d\mathcal{H}^d(\mathbf{p}) \geq \mathcal{H}^d[G]. \quad (63)$$

Noting that G is the union of two half balls of \mathbf{R}^d with radius $\delta \sin \theta$, joined by the cylinder of the same radius centered on the axis $\pi[\mathbf{y}_o, \mathbf{y}_1]$, we deduce that

$$\mathcal{H}^d[G] = \omega_{d-1} \mathcal{H}^1[\pi[\mathbf{y}_o, \mathbf{y}_1]] (\delta \sin \theta)^{d-1} + \omega_d (\delta \sin \theta)^d, \quad (64)$$

where ω_d is the volume of the unit ball in \mathbf{R}^d . Using (62)–(64) we obtain the following lower bound for the area of $[M(\mathbf{0}) + Y] \cap \partial\Lambda$:

$$\mathcal{H}^d[[M(\mathbf{0}) + Y] \cap \partial\Lambda] \geq \omega_{d-1} |\pi(\mathbf{y}_o) - \pi(\mathbf{y}_1)| (\delta \sin \theta)^{d-1}. \quad (65)$$

6. Our next goal is to bound the area of $\partial\Omega \cap Y^{*c}$ by exploiting uniform convexity. Remark A.3 yields a constant $R > 0$ such that the function $x_1 = k(x_2, \dots, x_{d+1})$ parameterizing $\partial\Omega$ near $\mathbf{x}_o = \mathbf{0}$ satisfies

$$k(\mathbf{p}) \geq |\mathbf{p}|^2/R. \quad (66)$$

Claim: For small enough $\theta > 0$ one has $\partial\Omega \cap Y(\theta)^{*c} \subset \{(k(\mathbf{p}), \mathbf{p}) \mid |\mathbf{p}| \leq R \tan \theta\}$.

Proof of Claim: Observe that $\mathbf{x} = (x_1, \pi(\mathbf{x})) \in Y^{*c} \setminus \{\mathbf{0}\}$ if and only if $x_1 < |\pi(\mathbf{x})| \tan \theta$. Thus any $\mathbf{x} \in \overline{\Omega} \cap Y^{*c}$ satisfies

$$|\mathbf{p}|^2/R \leq k(\mathbf{p}) \leq \mathbf{x} \cdot \mathbf{e}_1 \leq |\mathbf{p}| \tan \theta$$

where $\mathbf{p} = \pi(\mathbf{x})$ and Remarks A.2–A.3 have been used. One concludes $|\mathbf{p}| \leq R \tan \theta$. For small θ , say $0 < \theta < \theta_o$, this forces $|\mathbf{p}|$ and x_1 to be small enough that \mathbf{x} lies in the ball $\mathbf{B}(\mathbf{0}, r)$ of (68). Then $\mathbf{x} \in \partial\Omega$ yields $x_1 = k(\mathbf{p})$ to establish the claim.

7. Recall that $k(\mathbf{p})$ is locally Lipschitz near the origin in \mathbf{R}^d . For $0 < \theta < \theta_o$, Claim 6 and the area formula yield

$$\mathcal{H}^d[\partial\Omega \cap Y^{*c}] \leq \int_{|\mathbf{p}| \leq R \tan \theta} \sqrt{1 + |\nabla k(\mathbf{p})|^2} d\mathcal{H}^d(\mathbf{p}) \leq L(\tan \theta)^d, \quad (67)$$

where the constant L depends on R, θ_o and Ω but is independent of θ .

8. Finally, the ratio $\rho(\theta) > \epsilon^2 > 0$ of (57) can be bounded using (65) and (67):

$$\rho(\theta) \leq \frac{L \tan \theta}{\omega_{d-1} |\pi(\mathbf{y}_o) - \pi(\mathbf{y}_1)| (\delta \cos \theta)^{d-1}}.$$

Unless $\pi(\mathbf{y}_o) = \pi(\mathbf{y}_1)$, this forces $\rho(\theta)$ to zero with θ — a contradiction!

9. We conclude that $\mathbf{y}_o - \mathbf{y}_1$ must have been parallel to \mathbf{n} . If $\partial\Omega$ is differentiable at $\mathbf{x}_o = \mathbf{0}$, then $\mathbf{n}_o = \mathbf{n}$ and the proposition is complete. Otherwise, we may choose new co-ordinates obtained by rotating the $-\mathbf{e}_1$ axis slightly from \mathbf{n} toward \mathbf{n}_o . Since

the rotation is slight it is still possible to parameterize $\partial\Omega$ locally using a Lipschitz function $x_1 = \tilde{k}(x_2, \dots, x_{d+1})$ of the new co-ordinates. Moreover, $\tilde{k} \geq 0$ since both \mathbf{n} and \mathbf{n}_o are outer normals. Repeating arguments 2–8 yields $\mathbf{y}_0 - \mathbf{y}_1$ parallel to both unit normals: $-\mathbf{e}_1$ and $\mathbf{n} \neq -\mathbf{e}_1$. The only conclusion left is that $\mathbf{y}_0 - \mathbf{y}_1 = \mathbf{0}$, hence parallel to all normals. This establishes the proposition. \square

To reinforce our interpretation of this proposition, we verify tangential differentiability as a corollary. The examples discussed at the end of the previous section with $\mathbf{t}^- \neq \mathbf{t}^+$ give instances where non-tangential differentiability of ψ must fail.

Corollary 4.4 (Tangential Differentiability) *When the set $M = \partial\psi$ of Proposition 4.3 or 3.2 is the subdifferential of a convex function $\psi : \mathbf{R}^{d+1} \rightarrow \mathbf{R}$, then ψ is tangentially differentiable along $\partial\Omega$: at $\mathbf{x}_o \in \partial\Omega$ and in the coordinates of Lemma A.1, the partial derivatives $\partial\psi/\partial x_i$ exist for $i = 2, 3, \dots, d+1$. Should Ω fail to be differentiable at $\mathbf{x}_o \in \partial\Omega$, then the gradient $\nabla\psi(\mathbf{x}_o)$ exists.*

Proof: By Proposition 4.3 or 3.2 we know the convex set $\partial\psi(\mathbf{x}_o)$ is contained in a line parallel to each $\mathbf{n}_o \in N_\Omega(\mathbf{x}_o)$. Unless $N_\Omega(\mathbf{x}_o)$ consists of a single unit normal $\mathbf{n}_\Omega(\mathbf{x}_o)$, this line must degenerate to a point, namely $\nabla\psi(\mathbf{x}_o)$. Thus we may henceforth assume $\partial\Omega$ differentiable at \mathbf{x}_o , since otherwise we are done. Translating \mathbf{x}_o to the origin and choosing the coordinates of Lemma A.1, (68) and $k \geq 0$ force $\mathbf{n}_\Omega(\mathbf{x}_o) = -\mathbf{e}_1$ whence $\partial\psi(\mathbf{x}_o) \subset \{\mathbf{z} + \lambda\mathbf{e}_1 \mid \lambda \in \mathbf{R}\}$ for some $\mathbf{z} \in \mathbf{R}^{d+1}$. Choosing $i \in \{2, 3, \dots, d+1\}$ in (26–27) yields $\partial\psi/\partial x_i = \mathbf{z} \cdot \mathbf{e}_i + 0$ to complete the proof. \square

We close by speculating about Hölder continuity of the homeomorphism $\mathbf{t}^+ : \partial\Omega \rightarrow \partial\Lambda$ in Theorem 3.8. Here $S_0 \subset \partial\Omega$ refers to the closed set of Proposition 3.7.

Remark 4.5 (Hölder Continuity) *For symmetrically suitable measures on \mathbf{R}^{d+1} satisfying density bounds (6–7), we conjecture that a more delicate monotonicity argument like that of Caffarelli and McCann [7] should show the homeomorphism $\mathbf{t}^+ : \partial\Omega \setminus S_0 \rightarrow \partial\Lambda \setminus T_0$ to satisfy bi-Hölder estimates with exponent $1/d$ locally. For curves in the plane ($d = 1$), this suggests \mathbf{t}^+ is locally bi-Lipschitz outside of $S_0 \subset \partial\Omega$.*

A Appendix on Convex Domains

This appendix recalls a standard construction for parameterizing the boundary of a convex domain $\Omega \subset \mathbf{R}^{d+1}$ near any point $\mathbf{x}_o \in \partial\Omega$ in suitable Cartesian coordinates as the graph of a convex function $x_1 = k(x_2, \dots, x_{d+1})$.

Lemma A.1 (Boundary Parameterization) *Let \mathbf{x}_o lie on the boundary of a convex domain $\Omega \subset \mathbf{R}^{d+1}$. Translating $\mathbf{x}_o \in \partial\Omega$ to the origin and rotating if necessary, one can find an open ball $\mathbf{B}(\mathbf{0}, r)$ around \mathbf{x}_o and a Lipschitz function $k : \mathbf{R}^d \rightarrow [0, +\infty)$, such that*

$$\Omega \cap \mathbf{B}(\mathbf{0}, r) = \{\mathbf{x} \in \mathbf{B}(\mathbf{0}, r) \mid x_1 > k(x_2, \dots, x_{d+1})\}. \quad (68)$$

Proof: Let $f(\mathbf{x})$ denote the signed distance from $\mathbf{x} \in \mathbf{R}^{d+1}$ to the boundary of Ω :

$$f(\mathbf{x}) := \begin{cases} \text{dist}(\mathbf{x}, \partial\Omega) & \text{if } \mathbf{x} \notin \Omega \\ -\text{dist}(\mathbf{x}, \partial\Omega) & \text{if } \mathbf{x} \in \Omega. \end{cases}$$

The function f is known to be convex, e.g. [17, Lemma 4.2]. Since $f < 0$ in Ω but $f(\mathbf{x}_o) = 0$, clearly $\mathbf{0} \notin \partial f(\mathbf{x}_o)$. Let \mathbf{n} denote the point of $\partial f(\mathbf{x}_o)$ closest to the origin, and rotate co-ordinates so that \mathbf{n} lies along the negative x_1 -axis. The compact convex set $\partial f(\mathbf{x}_o)$ can be separated from the ball $\mathbf{B}^{d+1}(\mathbf{0}, |\mathbf{n}|)$ by a hyperplane which, since it passes through \mathbf{n} , can only be the plane $x_1 = -|\mathbf{n}|$. Thus $\partial f(\mathbf{x}_o)$ lies in the half-space $x_1 \leq -|\mathbf{n}| < 0$. This means that $f(x_1, x_2, \dots, x_{d+1})$ is a decreasing function of x_1 on a neighbourhood $\mathbf{B}^{d+1}(\mathbf{0}, r)$ of $\mathbf{x}_o = \mathbf{0}$. Taking $r > 0$ smaller if necessary, Clarke's non-smooth version of the implicit function theorem [8, §7.1] yields a Lipschitz function k on \mathbf{R}^d vanishing at the origin such that $f(k(\mathbf{p}), \mathbf{p}) = 0$ for $|\mathbf{p}| < r$. Since $f(\mathbf{x}) \geq -|\mathbf{n}|x_1$ it follows that $k(\mathbf{p}) \geq 0$. To verify (68) observe that on $\mathbf{B}^{d+1}(\mathbf{0}, r)$ strict monotonicity forces $f(x_1, \mathbf{p})$ to take the opposite sign from $x_1 - k(\mathbf{p})$. Thus $(x_1, \mathbf{p}) \in \Omega$ precisely when $x_1 > k(\mathbf{p})$. \square

Remark A.2 *Retaining the coordinates of the lemma, we see that the convex function*

$$k(\mathbf{p}) := \inf \{ \lambda \in \mathbf{R} \mid (\lambda, \mathbf{p}) \in \overline{\Omega} \} \quad (69)$$

also satisfies (68). Now $k : \mathbf{R}^d \rightarrow [0, +\infty]$ but remains Lipschitz near $\mathbf{0} \in \mathbf{R}^d$.

Remark A.3 (Uniform Convexity) *If $\Omega \subset \mathbf{R}^{d+1}$ is uniformly convex the function k of (68–69) satisfies $k(x_2, \dots, x_{d+1}) \geq (x_2^2 + \dots + x_{d+1}^2)/2R$ on \mathbf{R}^d , where the constant $R > 0$ is from Definition 0.2 and independent of $\mathbf{x}_o \in \partial\Omega$. Equivalence of this condition to uniform convexity is deduced from $|\mathbf{X}|^2/R \geq R - \sqrt{R^2 - |\mathbf{X}|^2} \geq |\mathbf{X}|^2/2R$.*

References

- [1] G. Alberti. On the structure of singular sets of convex functions. *Calc. Var. Partial Differential Equations* 2, 17–27 (1994).
- [2] G. Alberti and L. Ambrosio. A geometrical approach to monotone functions in \mathbf{R}^n . To appear in *Math. Z.* (1998).
- [3] Y. Brenier. Décomposition polaire et réarrangement monotone des champs de vecteurs. *C.R. Acad. Sci. Paris Sér. I Math.* 305, 805–808 (1987).
- [4] Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Comm. Pure Appl. Math.* 44, 375–417 (1991).

- [5] L.A. Caffarelli. The regularity of mappings with a convex potential. *J. Amer. Math. Soc.* 5, 99–104 (1992).
- [6] L.A. Caffarelli. Boundary regularity of maps with convex potentials — II. *Ann. of Math. (2)* 144, 453–496 (1996).
- [7] L.A. Caffarelli and R.J. McCann. Free boundary problems in optimal transport. *In preparation*.
- [8] F.H. Clarke. *Optimization and nonsmooth analysis*. Wiley, New York, 1983.
- [9] D.S. Fry. *Shape recognition using metrics on the space of shapes*. PhD thesis, Harvard University, 1993.
- [10] W. Gangbo. An elementary proof of the polar factorization of vector-valued functions. *Arch. Rational Mech. Anal.* 128, 381–399 (1994).
- [11] W. Gangbo and R.J. McCann. The geometry of optimal transportation. *Acta Math.* 177, 113–161 (1996).
- [12] C.R. Givens and R.M. Shortt. A class of Wasserstein metrics for probability distributions. *Michigan Math. J.* 31, 231–240 (1984).
- [13] L. Kantorovich. On the translocation of masses. *C.R. (Doklady) Acad. Sci. URSS (N.S.)* 37, 199–201 (1942).
- [14] L.V. Kantorovich and G.S. Rubinstein. On a functional space and certain extremum problems. *Dokl. Akad. Nauk SSSR (N.S.)* 115, 1058–1061 (1957).
- [15] H.G. Kellerer. Duality theorems for marginal problems. *Z. Wahrsch. Verw. Gebiete* 67, 399–432 (1984).
- [16] R.J. McCann. Existence and uniqueness of monotone measure-preserving maps. *Duke Math. J.* 80, 309–323 (1995).
- [17] R.J. McCann. Equilibrium shapes for planar crystals in an external field. To appear in *Comm. Math. Phys.* (1998).
- [18] D. Mumford. Mathematical theories of shape: Do they model perception. In *Geometric methods of computer vision*, vol. 1570 of the Proceedings of the Society for Photo-Optical Instrumentation Engineers, Bellingham, 1991, pp. 2–10.
- [19] S.T. Rachev. The Monge-Kantorovich mass transference problem and its stochastic applications. *Theory Probab. Appl.* 29, 647–676 (1984).
- [20] R.T. Rockafellar. *Convex analysis*. Princeton University Press, Princeton, 1972.

- [21] L. Rüschendorf and S.T. Rachev. A characterization of random variables with minimum L^2 -distance. *J. Multivariate Anal.* 32, 48–54 (1990).
- [22] T.L. Seethoff and R.C. Shiflett. Doubly stochastic measures with prescribed support. *Z. Wahrsch. Verw. Gebiete* 41, 283–288 (1978).
- [23] C. Smith and M. Knott. On the optimal transportation of distributions. *J. Optim. Theory Appl.* 52, 323–329 (1987).
- [24] L.N. Wasserstein. Markov processes over denumerable products of spaces describing large systems of automata. *Problems of Information Transmission* 5, 47–52 (1969).
- [25] L. Zajíček. On the differentiability of convex functions in finite and infinite dimensional spaces. *Czechoslovak Math. J.* 29 (104), 340–348 (1979).

List of Figures

Figure 1: Numerical simulations of $\partial\Omega$ evolving optimally onto $\partial\Lambda$ (for $\beta = 0$). Each movie starts in the upper left corner and proceeds from left to right, down the page. They were constructed by discretizing the boundaries, finding the optimal correspondence between segments of $\partial\Omega$ and $\partial\Lambda$, and interpolating the locations of those segments linearly with time λ . Courtesy of David Fry.

(a)

(b)

Figure 2: Schematic Illustration of Maps \mathbf{t}^+ and \mathbf{t}^- Splitting $\mathbf{x} \in S_2 \subset \partial\Omega$ into Two Images which Parallel the Normal $\mathbf{n}_\Omega(\mathbf{x})$.

Figure 3: Several of the Possible Intermediate States between Two Rectangles

Figure 4: No Optimal Map Links these Two Triangles (Perturbed for Strict Convexity)

(a) Reflected Triangles

(b) Strictly Convex Perturbations

Figure 5: Nor these Slight Perturbations of a Fixed Isosceles Triangle