

A Riemannian Corollary of Helly's Theorem

Alexander Rusciano

*Department of Mathematics, University of California, Berkeley, U.S.A.
rusciano@math.berkeley.edu*

Received: July 4, 2018

Accepted: February 7, 2020

We introduce a notion of halfspace for Hadamard manifolds that is natural in the context of convex optimization. For this notion of halfspace, we generalize a classic result of Grünbaum, which itself is a corollary of Helly's theorem. Namely, given a probability distribution on the manifold, there is a point for which all halfspaces based at this point have at least $\frac{1}{n+1}$ of the mass. As an application, the subgradient oracle complexity of convex optimization is polynomial in the size of the parameters defining the problem.

Keywords: Helly's theorem, geodesic convexity, convex optimization.

2010 Mathematics Subject Classification: 52A01.

1. Overview

1.1. Introduction

The extrema of functions are of fundamental importance in mathematics and its applications. Much of numerical optimization studies this topic. Most of the theory focuses on convex functions, as it has proven hard to find other classes that are both useful and tractable. The motivation for this paper comes from the desire to expand the boundaries of this class of tractable functions.

Rigorous study of convergence rates for optimizing convex functions on Hadamard manifolds was initiated in [15] for first-order methods. That is, their work considered gradient descent methods for simply connected manifolds of non-positive sectional curvature. Such manifolds are diffeomorphic to $\mathbb{R}^n$ and exhibit natural convex functions. In a sense, they give new classes of functions for which optimization is tractable.

Still, as far as the author is aware, all known algorithms for general convex optimization on Riemannian manifolds have iteration complexity depending polynomially on ϵ^{-1} . To achieve better convergence rates, further conditions are added such as strong convexity, dominated gradients, or recently robust second-order [15], [14], [1]. One major unresolved question for Hadamard manifolds like SL_n/SO_n is, does convexity enable algorithms whose time complexity depends polynomially on $\log(\epsilon^{-1})$?

For Euclidean optimization, cutting plane methods are the standard, general approach to get $\log(\epsilon^{-1})$ complexity. It is well known that the minimum of a convex

function lies in the halfspace opposite the subgradient direction. Cutting plane methods use this fact to reduce the size of the feasible set. One feature of $\mathbb{R}^n$ that enables this approach to succeed is the existence of what are commonly termed centerpoints. A precise definition of centerpoint is given in Definition 1.8. Roughly speaking, if c is to be a centerpoint for a set S , then no halfspace based at c should contain too large (or small) a fraction of the volume of S . Ellipsoid methods explicitly maintain a radially symmetric set, so the center of the current ellipse provides a perfect centerpoint. Thus a subgradient at the center of an ellipse allows one to eliminate half of the ellipse. For more general subsets $S \subset \mathbb{R}^n$, Grünbaum's result in [7] shows the existence of a centerpoint c for which any halfspace based at c contains at most a $\frac{n}{n+1}$ -fraction of the mass of S .

As generalizing this result of Grünbaum is the main goal of this work and we fundamentally build from it in the proof, we summarize his result as Theorem 1.1. Note the statement is for a general probability measure, a fact we will make use of by applying it to the Riemannian volume measure.

Theorem 1.1. *A $\frac{1}{n+1}$ -centerpoint c exists for any probability measure on $\mathbb{R}^n$, endowed with the usual Borel σ -algebra. Here c being a $\frac{1}{n+1}$ -centerpoint means any halfspace based at c contains at least a $\frac{1}{n+1}$ -fraction of the mass of the probability distribution.*

We replicate this result in the more general setting of Hadamard manifolds in the hope that others find the result encouraging, useful, or intrinsically interesting. Though we present the results here to provide motivation, the definitions in Section 1.2 make the statements precise. The main result is from Section 2,

Theorem 1.2. *Suppose μ is a probability distribution on a Hadamard manifold M of dimension n , and μ is absolutely continuous with respect to the Riemannian volume measure vol_g . Then there exists a $\frac{1}{n+1}$ -centerpoint c for the measure μ . If we assume the support of μ is contained in a (geodesic) convex set S , then $c \in S$. Moreover, even for the uniform measure on convex and compact S , this value $\frac{1}{n+1}$ cannot in general be improved.*

We noted the sharpness of $\frac{1}{n+1}$ in the above theorem because this contrasts with the guaranteed existence of $\frac{1}{e}$ -centerpoints for the uniform measure on convex subsets of Euclidean space [7].

The theorem leads to a bound on the number of subgradient oracle calls needed to optimize a function.

Theorem 1.3. *Suppose the subset S of Hadamard manifold M of dimension n is (geodesic) convex, and that $f : S \rightarrow \mathbb{R}$ is a (geodesic) convex L -Lipschitz function. Additionally assume the minimum of f , denote by x_* , is in the ϵ -interior of S , meaning that the open ball centered at x_* of Riemannian radius ϵ is contained within S . Then it is possible to find a point $x \in S$ such that $f(x) - f(x_*) \leq \epsilon$ using $O(n^2 \log(n L \text{vol}_g(S) \epsilon^{-1}))$ subgradient oracle calls, where $\text{vol}_g(S)$ denotes the Riemannian volume of S .*

1.2. Definitions and notation

In this section we give the definitions needed to frame the problem and results. Only basic notions of Riemannian geometry are needed in this paper; these are surveyed in Appendix A, with the present section mainly providing non-standard or less common definitions.

For the remainder of this paper we study triples (M, g, μ) , where M is an n -dimensional, simply-connected manifold equipped with a complete Riemannian metric g of non-positive sectional curvature. Such Riemannian manifolds (M, g) are called Hadamard manifolds. For each point $x \in M$ we denote by $\langle \cdot, \cdot \rangle_x$ the inner product, defined by g , on the tangent space $T_x M$. The metric g also induces the Riemannian volume measure, denoted by vol_g . In addition to this, we consider a probability measure μ , which we assume to be absolutely continuous with respect to vol_g . For the motivating application μ is taken to be the Riemannian volume measure restricted to a subset $S \subset M$, i.e. $\mu = (\mathbb{1}_S / \text{vol}_g(S)) \cdot \text{vol}_g$. The metric g further induces a notion of distance between points on the manifold, allowing us to define geodesics, which are locally length minimizing paths. For $x \in M$, we denote by $\exp_x : T_x M \rightarrow M$ the exponential map based at x , which maps tangent vectors to geodesics passing through x . As explained in the Riemannian geometry overview in Appendix A, the exponential map is a diffeomorphism when (M, g) is a Hadamard manifold.

We proceed with definitions related to convexity,

Definition 1.4. $S \subset M$ is *convex* if points $x, y \in S$ are joined by a unique length-minimizing geodesic contained in S . $\square$

Definition 1.5. A function $f : M \rightarrow \mathbb{R}$ is *convex* on its domain if its restrictions to geodesics are convex in t . That is, $f(\exp_x(tv)) : \mathbb{R} \rightarrow \mathbb{R}$ is a convex function in t . For such a convex function f , a tangent vector $w \in T_x M$ is said to be a *subgradient* at x if for any $v \in T_x M$,

$$f(\exp_x(tv)) \geq f(x) + t\langle w, v \rangle_x.$$

The set of subgradients at x is known as the *subdifferential* at x , and is denoted by ∂f_x . $\square$

We note Theorem 4.5 in [13] proves that convex functions have a non-empty subdifferential at all points. Accordingly, when the convex function f is discussed, we will assume a subgradient oracle that for any x outputs some $w \in \partial f_x$. As explained in Appendix A, the gradient of a differentiable convex function is a subgradient. Therefore the subgradient oracle for such a function can be explicit.

Let us present two examples of convex functions for motivation. The first is general, the second specific.

- The distance to a convex subset (Definition 1.4) of a Hadamard manifold is convex [2]. Thus finding the point minimizing the mean distance or mean squared-distance to a set of points is a convex optimization problem.

- Identify SL_n/SO_n with the set of positive-definite matrices of determinant 1 [2]. Then for arbitrary $B_i \in GL_n$, the function

$$\log \det \left(\sum_{i=1}^m B_i^T X B_i \right)$$

defined on SL_n/SO_n is convex in X [12]. Minimizing such a function can be used to find the optimal Brascamp-Lieb constant [4]. Up to scaling, all symmetric spaces of non-compact type (examples of Hadamard manifolds) embed as totally geodesic submanifolds of these spaces. Therefore restrictions of this function to such submanifolds give many more examples. Minimizing this function has figured prominently in recent theoretical computer science research, notably in [1]. Their work succeeds in developing an optimization procedure depending polynomially on $\log(\epsilon^{-1})$ for such functions, but the approach relies on special properties of this family of functions.

With these examples in mind, let us return for two more important definitions,

Definition 1.6. An open *halfspace* based at $x \in M$ is formed by applying $\exp_x(\cdot)$ to a halfspace of $T_x M$. For a given $v \in T_x M$ we denote halfspaces by

$$H_x(v) := \{\exp_x(w) \mid w \in T_x M, \langle w, v \rangle_x < 0\}. \quad \square$$

Although such halfspaces are not convex sets in the general setting of Hadamard manifolds, they are naturally produced by cutting planes for convex functions. This notion of cutting plane is justified by the following lemma,

Lemma 1.7. Consider a convex function $f : S \rightarrow \mathbb{R}$, where S is a convex subset of Hadamard manifold M . Then for any $x \in S$ and any subgradient $v \in \partial f_x$, the minimum of f within S is either attained at x or lies within $H_x(v) \cap S$. Moreover, if $y \in S \setminus H_x(v)$, then $f(y) \geq f(x)$.

Proof. If $y \notin H_x(v)$, the corresponding $v' = \exp_x^{-1}(y)$ satisfies $\langle v', v \rangle_x \geq 0$, and we have $f(y) \geq f(x) + \langle v, v' \rangle_x \geq f(x)$. $\square$

Cutting plane methods need to find a point for which no halfspace based at that point has too much of the feasible set's volume. This can be captured through the notion of a centerpoint. Our definition of centerpoint technically could be applied to any probability measure on any space with a notion of halfspace. However, in proving Theorem 1.2, we further restrict to the notion of halfspace we defined for Hadamard manifolds, and require the probability measure to be absolutely continuous with respect to the Riemannian volume measure vol_g .

Definition 1.8. A β -centerpoint of the probability measure μ on a Hadamard manifold M is a point c such that $\mu(H_c(v)) \leq 1 - \beta$ for all $v \in T_c M$. $\square$

Theorem 1.2 claimed that even for the uniform measure on convex subsets of M , a $\frac{1}{n+1}$ -centerpoint is the best we can guarantee. We included this comment both to contrast with $\mathbb{R}^n$, as well as to include a concrete illustration. It is not difficult to present such an example through studying $\mathbb{H}^n$, the model space of constant -1 sectional curvature.

Towards this end, let us briefly state the important features of the Klein model of $\mathbb{H}^n$ that we require. We identify $\mathbb{H}^n$ with the open Euclidean unit ball $B(1) \subset \mathbb{R}^n$, which we take to be centered at the origin. This set is equipped with the metric $g = (dx_1^2 + \cdots + dx_n^2)/(1 - x_1^2 - \cdots - x_n^2)$. This leads to a volume form of

$$\text{vol}_g = \frac{1}{(1 - x_1^2 - \cdots - x_n^2)^{\frac{n+1}{2}}} dx^1 \wedge \cdots \wedge dx^n.$$

Critically for our exposition, this model of hyperbolic space has its geodesics appear as Euclidean lines; thus Riemannian halfspaces appear as halfspaces intersected with $B(1)$. Other work such as [3] has found this useful in studying convex objects in $\mathbb{H}^n$.

For the construction demonstrating that $\frac{1}{n+1}$ cannot in general be improved in Theorem 1.2, the idea is simply that convex polyhedra in $\mathbb{H}^n$ have their volume concentrated towards the vertices. Let $T(1)$ be the closed, regular n -simplex inscribed in $B(1)$. It can be checked that $T(1)$ has finite volume; such objects are called ideal polyhedra and have been studied extensively. By symmetry, we will see that the origin $\vec{0}$ is the optimal centerpoint for $T(1)$. However, note that a hyperplane through $\vec{0}$ parallel to any of the faces will contain exactly 1 of the $n+1$ vertices. An application of the Gauss-Bonnet theorem can be used to check that in the case of $\mathbb{H}^2$, the area of the halfspace containing a single vertex is $\frac{\pi}{3}$. The opposite halfspace, which contains the other two vertices, is of area $\frac{2\pi}{3}$. The author conjectures the result to hold as well for higher dimensions.

For our purpose, it is more direct to modify the example slightly,

Proposition 1.9. *Let $T(1+\delta)$ be a closed, regular n -simplex inscribed in $B(1+\delta)$. We take $\delta > 0$ small enough so that $B(1)$ is not inscribed in $T(1+\delta)$. Also define $S_\epsilon = T(1+\delta) \cap B(1-\epsilon)$, which we observe to be convex and compact. Then the optimal centerpoint of S_ϵ approaches being a $\frac{1}{n+1}$ -centerpoint as $\epsilon \rightarrow 0$.*

Proof. It is clear that $S = T(1+\delta) \cap B(1)$ has unbound Riemannian volume, because each of the $n+1$ vertices of $T(1+\delta)$ lies outside of $B(1)$. However, $S_\epsilon = T(1+\delta) \cap B(1-\epsilon)$ has finite volume for any $\epsilon > 0$, and is convex and compact. Define the set of probability measures consisting of the uniform measure restricted to S_ϵ , i.e. $\mu_\epsilon = (\mathbb{1}_{S_\epsilon}/\text{vol}_g(S_\epsilon)) \cdot \text{vol}_g$.

By symmetry, the origin $\vec{0}$ is the optimal centerpoint for each μ_ϵ . In more detail, the optimal centerpoint for S_ϵ is the solution to minimizing the function $G(y) := \sup_{\hat{v} \in S^{n-1}} \mu_\epsilon(H_y(\hat{v}))$. Viewing $G(y)$ as a function on $S_\epsilon \subset \mathbb{R}^n$, Lemma 2.2 establishes that it is quasi-convex. Again, this is possible because in this model of $\mathbb{H}^n$, geodesics appear as Euclidean straight lines. One characterization of quasi-convex functions f is $f(tx + (1-t)y) \leq \max(f(x), f(y))$ for all $0 \leq t \leq 1$. Thus, suppose $x_* \neq \vec{0}$ is the optimal centerpoint. Because S_ϵ exhibits tetrahedral symmetry, we see there are $n+1$ optimal centerpoints, and their convex hull includes $\vec{0}$. Using quasi-convexity, any points in the convex hull of these $n+1$ optimal centerpoints must also be optimal. This proves $\vec{0}$ is the optimal centerpoint, using the symmetry of the set S_ϵ and quasi-convexity of the centerpoint function $G(y)$.

Now consider a hyperplane through $\vec{0}$ parallel to one of the faces of S_ϵ , and denote by H^+ the resulting halfspace containing only one of the vertices of $T(1 + \delta)$. Because the volumes close to the vertices of $T(1 + \delta)$ diverge at equal rates, it follows that as $\epsilon \rightarrow 0$,

$$\mu_\epsilon(H^+) = \text{vol}_g(H^+ \cap S_\epsilon) / \text{vol}_g(S_\epsilon) \rightarrow \frac{1}{n+1}. \quad \square$$

This concludes our proof of the sharpness of Theorem 1.2.

1.3. Overview and conclusion

The remainder of this paper is organized as follows:

- Section 2 analyzes the existence of centerpoints on Hadamard manifolds.
- Section 3 presents the brief application of the above to upper bound subgradient oracle complexity.
- Appendix A recalls the relevant notions of Riemannian geometry and provides references.

To be clear, the problem of developing an efficient optimization procedure is far from resolved. However, our results show that there is not an information theoretic obstacle to developing cutting plane methods for Hadamard manifolds.

We hope our main result is of interest and encourages others to study centerpoints in the manifold setting. Targeting optimization procedures, we believe focusing on the spaces SL_n/SO_n would be of greatest interest, both for theory and applications. Computing a centerpoint from a discrete point set would be a notable advancement. It would also be useful to be able to sample from the Riemannian volume restricted to a convex subset.

2. Existence of centerpoints

One might wonder if the centroid of a convex set of a manifold is an adequate centerpoint. Here centroid refers to the center of mass, the point minimizing the average squared-distance. After all, the centroid of a convex subset of $\mathbb{R}^n$ is an approximately optimal centerpoint [7]. However, this is tied closely to the fact that cross-sectional areas of a convex set in $\mathbb{R}^n$ follow a log-concave probability distribution - a consequence of the Brunn-Minkowski inequality. On the other hand, for a manifold with negative sectional curvature, the distribution of cross-sectional areas is not necessarily even unimodal. This reflects the fact that manifold versions of the Brunn-Minkowski inequality use curvature lower bounds as parameters, and are qualitatively different in negative curvature compared to $\mathbb{R}^n$ [6]. Helly's Theorem is somewhat the opposite, as it holds in situations in which the distance function is convex. Moreover, as cited in Appendix A, Hadamard manifolds have convex distance functions. One can find in [9] and [8] proofs that amount to,

Theorem 2.1. *Let M be an n -dimensional Riemannian manifold of non-positive sectional curvature. Suppose we are given a convex compact set C and a family $\{C_\alpha\} \subset C$ of closed convex sets. Then if for an arbitrary selection of $n+1$ sets $C_{\alpha_1} \cap \dots \cap C_{\alpha_{n+1}} \neq \emptyset$, it follows that $\cap_\alpha C_\alpha \neq \emptyset$.*

The paper [8] actually proves this result for $\text{Cat}(0)$ geodesic spaces.

That the halfspace notion of Definition 1.6 is not typically convex limits the applicability of this generalization of Helly's Theorem. The remainder of this section proves a result that could be considered a Riemannian variant of the well-known corollary of Helly's theorem cited as Theorem 1.1. To generalize that result, we rely on a few simple regularity properties of sets of Euclidean centerpoints, which we now collect. In the following lemma, the halfspaces are Euclidean halfspaces, and D is the Hausdorff distance. That is,

$$D(A, B) := \max\left\{\sup_{b \in B} \inf_{a \in A} |a - b|, \sup_{a \in A} \inf_{b \in B} |a - b|\right\},$$

where $|\cdot|$ denotes the Euclidean norm. In other words, $D(A, B)$ is the minimal value ϵ so that A is contained in the ϵ -fattened version of B and vice versa. Also recall the total variation distance between two probability distributions is

$$\sup_{A \in \mathcal{F}} |\mu_1(A) - \mu_2(A)|,$$

where $|\cdot|$ denotes the absolute value, as there will be no confusion. Here A can be any measurable set, the collection of which is labeled $\mathcal{F}$.

Lemma 2.2. *Let $\{\mu_x(\cdot)\}$ be a family of probability measures on $\mathbb{R}^n$ that share a compact support Y . Assume the measures are indexed by members x of a compact metric space X with metric d , and the measures $\mu_x(\cdot)$ vary continuously with respect to total variation distance. Define the Euclidean centrality function*

$$G: X \times Y \rightarrow \mathbb{R} \quad \text{by} \quad G(x, y) := \sup_{\hat{v} \in S^{n-1}} \mu_x(H_y(\hat{v})),$$

in order to measure how good of a centerpoint y is for distribution μ_x . Then G is continuous under the product topology and $G(x, \cdot)$ is a quasi-convex function for a fixed x . Fixing an arbitrary $\alpha > 0$, also define the sets

$$U_x := \left\{ y \in \mathbb{R}^n \mid G(x, y) \in \left(0, 1 - \frac{1}{n+1} + \alpha \right] \right\},$$

in order to explicitly propose the set of centerpoints for distribution μ_x . For any x , these sets have a non-empty interior. Moreover fix $x \in X$ and suppose $\text{supp}(\mu_x)$ is a connected set. Then as $x_i \rightarrow x$, necessarily $D(U_{x_i}, U_x) \rightarrow 0$.

Proof. Each $\{y \mid \mu_x(H_y(\hat{v})) < a\}$ is a halfspace. Indeed, there is a maximal halfspace with normal $\hat{v}$ of mass a , and the previous set is precisely the points contained in this halfspace. Therefore the intersection over all $\hat{v}$ is a convex set. This shows that preimages under $G(x, \cdot)$ of sets $(-\infty, a)$ are convex, which is the definition of quasi-convex.

As we are using the product topology, the domain of G , which we denote by $K = X \times Y$, is compact. Because $g(x, y, \hat{v}) : (x, y, \hat{v}) \mapsto \mu(H_x(\hat{v}))$ is continuous and K is compact, g is uniformly continuous on $K \times S^{n-1}$.

Thus given $\epsilon > 0$, one can choose δ so that if $d(x, x') < \delta$ and $|y - y'| < \delta$, then

$$|g(x, y, \hat{v}_0) - g(x', y', \hat{v}_0)| < \epsilon$$

holds for any $\hat{v}_0$. Because of the compactness in the last argument, we may let $G(x, y) = g(x, y, \hat{v}_{x,y})$. Therefore

$$\begin{aligned} G(x, y) - G(x', y') &= g(x, y, \hat{v}_{x,y}) - g(x', y', \hat{v}_{x',y'}) \\ &> g(x, y, \hat{v}_{x,y}) - (g(x, y, \hat{v}_{x',y'}) + \epsilon) \geq -\epsilon \end{aligned}$$

holds. Switching roles gives the reverse inequality, $G(x, y) - G(x', y') < \epsilon$, which proves continuity.

We move on to the final claim. Recall the Hausdorff distance is the maximum distance it might require to travel from a point in one of the sets to the other set. We argue by contradiction that $U_{(\cdot)}$ converges to U_x in the Hausdorff distance metric. If we assume to the contrary, then either (i) there exists a sequence of points $y_{n_i} \in U_{x_{n_i}}$ such that y_{n_i} are bounded away from U_x or (ii) there exists a sequence of points $y_{n_i} \in U_x$ bounded away from $U_{x_{n_i}}$.

In situation (i), compactness implies an accumulation point y for the sequence y_{n_i} . However, continuity of G requires $y \in U_x$, because $x_{n_i} \rightarrow x$ and each $G(x_{n_i}, y_{n_i})$ is in $(0, 1 - \frac{1}{n+1} + \alpha]$. This contradicts the premise that y_{n_i} are bound away from U_x . In particular, this shows that the maximum distance from a point in U_{x_i} to the set U_x is going to 0.

In situation (ii), again by compactness there is an accumulation point $y \in U_x$ that is bounded away from infinitely many of the $U_{x_{n_i}}$. Because we have assumed $\alpha > 0$, Theorem 1.1 and the continuity of $G(x, \cdot)$ imply U_x has an interior. Note that in proving continuity of $G(x, \cdot)$, we used the absolute continuity of μ_x with respect to Lebesgue measure.

As a first subcase of (ii), we assume y is in the interior of U_x . We show by contradiction that $G(x, y) < 1 - \frac{1}{n+1} + \alpha$. Assuming to the contrary, there would be $\hat{v}$ such that $G(x, y) = 1 - \frac{1}{n+1} + \alpha = \mu_x(H_y(\hat{v}))$, and we may select $p \in H_y(-\hat{v}) \cap U_x$ because y is in the interior of U_x . As $p \in U_x$ and $\mu_x(H_y(\hat{v})) = 1 - \frac{1}{n+1}$, it must be the case that $\mu_x(H_y(-\hat{v}) \cap H_p(\hat{v})) = 0$, hence

$$H_y(-\hat{v}) \cap H_p(\hat{v}) \cap \text{supp}(\mu_x) = \emptyset,$$

and this implies the boundary of $H_{\frac{y+p}{2}}(\hat{v})$ is disjoint from $\text{supp}(\mu_x)$. Then

$$\text{supp}(\mu_x) \cap H_{\frac{y+p}{2}}(-\hat{v}) \quad \text{and} \quad \text{supp}(\mu_x) \cap H_{\frac{y+p}{2}}(\hat{v})$$

are non-empty sets whose union is $\text{supp}(\mu_x)$. As these sets are open in the induced topology on $\text{supp}(\mu_x)$, this contradicts the assumption that $\text{supp}(\mu_x)$ is connected. Therefore we have shown by contradiction that $G(x, y) < 1 - \frac{1}{n+1} + \alpha$. We immediately conclude from the continuity of G that $y \in U_{x_i}$ for large enough i .

In the event that y is not in the interior of U_x , we can still select an interior point $y' \in U_x$ that is arbitrarily close to y , because the set is open and convex. For any

such y' , the prior argument establishes that $y' \in U_{x_i}$ for large enough i . Therefore, we conclude that the distance between y and U_{x_i} is going to 0, contradicting our assumption. This completes the proof showing $D(U_{x_i}, U_x) \rightarrow 0$. $\square$

As a comment on the proof, the assumption that $\text{supp}(\mu_x)$ is connected was essential in the final conclusion of the proof. This assumption, along with a few others like the use of α , are bootstrapped out of the eventual theorem we prove. It would be nice to eliminate the absolute continuity assumption on μ_x by using a more general convergence tool like Wasserstein distance. We necessarily lose G 's continuity, but it remains lower semi-continuous. However, these topological properties alone were insufficient for proving U_x had an interior point or analogous “deep” point, which were necessary in proving Hausdorff convergence for U_x .

We are now ready for the key step in proving the main result.

Proposition 2.3. *Let μ be a probability measure whose support lies within a compact subset S of a Hadamard manifold M . Further assume μ is absolutely continuous with respect to the Riemannian volume measure vol_g and has connected support. Then there exists a $\frac{1}{n+1}$ -centerpoint for μ .*

Before going into the proof details, let us consider a conceptual overview of the proof. We will define a continuous function F from S to itself, and an application of Brouwer's theorem will show there is a fixed point. We design F so that the fixed point is a $(\frac{1}{n+1} - \alpha)$ -centerpoint. The α is inherited from the definition of U_x in Lemma 2.2, and is removed at the end of the proof. In designing F , we adopt normal coordinates at x and pull back the measure μ from M (i.e. the measure of $U \subset \mathbb{R}^n$ is $\mu(\exp_x(U))$). In these coordinates, there is a Euclidean-convex set of Euclidean centerpoints U_x provided by the previous lemma, for the pulled-back measure. We select the closest of these centerpoints to x and denote this point by u_x . Finally, $F(x)$ is then defined by projecting u_x onto S . As stated precisely in the appendix, it is the Hadamard assumption that implies a strictly convex distance function, making this projection possible.

The technical part of the proof mostly involves showing continuity of $F(x)$, as it is not hard to show that fixed points are $(\frac{1}{n+1} - \alpha)$ -centerpoints. The main obstacle is to show that u_x varies continuously. To establish this, we note that the pulled back measures vary continuously with respect to total variation. Then Lemma 2.2 shows that the Euclidean centerpoint sets $U_x, U_{x'}$ are close in Hausdorff distance, provided x, x' are close. Combining this with convexity of the centerpoint sets, we are able to make $|u_x - u_{x'}|$ small.

We now provide the details.

Proof. We may WLOG assume S is a closed Riemannian ball of radius R . By parallel transport we may fix a smooth orthonormal frame $V = (\vec{e}_1, \dots, \vec{e}_n)$ on S , thereby determining normal coordinate charts at each $x \in S$ defined by

$$\psi_x : y \mapsto \exp_x(y^i \vec{e}_i(x)).$$

Note that $\psi_x(y)$ varies smoothly both in x and y . We may pull back the measure μ by ψ_x to give the measures $\mu_x(y)dy$. The absolute continuity of these pull

back measures with respect to Lebesgue measure is due to μ being absolutely continuous with respect to the Riemannian volume measure. Smoothness of parallel transport ensures that the coordinate charts vary smoothly and therefore the $\mu_x(y)$ vary continuously with respect to total variation distance. Finally, since ψ_x are diffeomorphisms, we see that for each $x \in S$ the measure $\mu_x(y)$ has connected support. The set S is of radius R . Therefore in applying Lemma 2.2, we may choose Y to be the closed ball of radius $2R$.

First fix $\alpha > 0$. For all $x \in S$, Lemma 2.2 then establishes the existence of non-empty compact convex sets $U_x \subset \mathbb{R}^n$ of Euclidean $(\frac{1}{n+1} - \alpha)$ -centerpoints. There is a unique point $u_x \in U_x$ that is closest to x . However, it is not necessarily the case that u_x is inside $\psi_x^{-1}(S)$, because $\psi_x^{-1}(S)$ is not convex with respect to the Euclidean metric. To work around this, project u_x onto S . That is,

$$F(x) := \pi(u_x) := \arg \min_{s \in S} d(s, \psi_x(u_x)),$$

where by $d(\cdot, \cdot)$ we mean the Riemannian distance. The projection is well-defined and continuous by [2, Corollary 5.6]. Therefore F is well-defined. In the following we show

- If $F(x) = x$, then x is a $(\frac{1}{n+1} - \alpha)$ -centerpoint contained in S .
- $F(x)$ is continuous.

Then since S is a closed ball, an application of Brouwer's fixed point theorem yields the desired result.

We first show that fixed points are centerpoints. We argue by contradiction and assume x is a fixed point of F which is not a $(\frac{1}{n+1} - \alpha)$ -centerpoint. Observe that $u_x \neq \vec{0}$, because this would imply x is a $(\frac{1}{n+1} - \alpha)$ -centerpoint. Further, we see that $\psi_x(H_0(-u_x)) \cap S \neq \emptyset$, since $H_{u_x}(-u_x) \subset H_0(-u_x)$ and

$$\mu(\psi_x(H_{u_x}(-u_x))) > \frac{1}{n+1} - \alpha > 0$$

by the centerpoint property. Next choose $s \in \psi_x(H_{\vec{0}}(-u_x)) \cap S$, and consider the geodesic between x, s . This geodesic is contained in S by the assumption that S is convex. Triangle inequalities in the form of Toponogov's Theorem [5] (or convexity of the distance function as a simple alternative) show that, initially, moving from x to s along the geodesic decreases the distance to u_x . This means it is not the case that $\pi(\psi_x(u_x)) = x$.

Next we consider the continuity claim. Once we show $u_x \in \mathbb{R}^n$ varies continuously with respect to $x \in S$, then the continuity of $F(x)$ follows because, as noted in Appendix A, the projection is also continuous. As a first step, we remark that the pull-back probability densities $\mu_x(y)$ vary continuously with respect x , because they are defined by smoothing varying diffeomorphisms ψ_x . Moreover, as S is compact and μ is supported on S , we may assume the y are taken from a compact set. Then we may apply uniform continuity to show there is δ so that $d(x, x') < \delta$ implies $|\mu_x(y) - \mu_{x'}(y)| < \epsilon$ for any y . This establishes continuity for the family of measures $\mu_x(y)dy$, with respect to total variation distance. We can now make use of the regularity properties provided by Lemma 2.2.

From the lemma's last part, by requiring $d(x, x') < \delta$ for small enough δ , one can ensure $D(U_x, U_{x'}) < \epsilon$. Let $h_x \in U_x$ be the point closest to $u_{x'}$; this ensures $|u_{x'} - h_x| < \epsilon$. It is also not difficult to see that $|u_{x'}| - |u_x| < \epsilon$. Therefore $|h_x| - |u_x| < 2\epsilon$. Critically, the Euclidean distance to the origin is strongly convex and u_x minimizes it on the Euclidean convex set U_x , which also includes h_x . Therefore, qualitatively, since $|h_x|$ and $|u_x|$ are similar in value, we know that $|h_x - u_x|$ is small. Making this quantitative through the Euclidean law of cosines,

$$|u_x - h_x|^2 \leq |h_x|^2 - |u_x|^2 = (|h_x| - |u_x|)(|h_x| + |u_x|) < \epsilon R$$

where a sufficiently large R can be taken to be twice the diameter of S .

We conclude $|u_x - u_{x'}| \leq |u_x - h_x| + |h_x - u_{x'}| < \sqrt{\epsilon R} + \epsilon$,

which establishes continuity for $F(x)$. This essentially completes the proof, but recall that we have used a small α parameter to define U_x , and this resulted in our proving only the existence of $(\frac{1}{n+1} - \alpha)$ -centerpoints for $\alpha > 0$. However, the continuity of the centerpoint function $\sup_{\hat{v} \in S^{n-1}} H_x(\hat{v})$ on S follows from the same argument used to prove continuity of G in Lemma 2.2. From this and compactness of S , may conclude the existence of $\frac{1}{n+1}$ -centerpoints. $\square$

This nearly proves the main part of Theorem 1.2. The main difference is the absence of a few simplifying assumptions, namely compactness and connected support. The fact that $x \in S$ provided S is convex was also postponed. We complete the proof here.

Proof for Theorem 1.2. We first remove the connected support assumption. For measures μ supported on compact S , we may WLOG assume S is a ball and therefore connected. Then we may define the probability measures

$$(1 - \epsilon)\mu + \epsilon \frac{\mathbb{1}_S}{\text{vol}_g(S)} \cdot \text{vol}_g.$$

Proposition 2.3 applies to these measures. Thus it is clear that we may construct $(\frac{1}{n+1} - \epsilon)$ -centerpoints for μ . Again using the continuity of the centerpoint function and compactness of S as at the end of Proposition 2.3's proof, it follows that a $\frac{1}{n+1}$ -centerpoint exists for μ .

Next we extend the result by removing the assumption that S is compact. Fixing some point $x \in M$, we may define the family of compact sets $S_i = S \cap \bar{B}_g(x, i)$, where i ranges over the positive integers and $\bar{B}_g(x, i)$ denotes the closed ball of radius i around x . These sets satisfy $\lim_{i \rightarrow \infty} \mu(S_i) = 1$. Applying Proposition 2.3 to these S_i , we get points s_i that are at least $(\frac{1}{n+1} - \mu(S_i^C))$ -centerpoints for μ . Moreover, these s_i must all lie in some compact set $C \subset S$. Indeed, by the analog of the separating hyperplane theorem proven in Lemma A.2, any point $p \notin B_g(x, i)$ will have a halfspace $H_p(\hat{v}) \cap S_i = \emptyset$, and thus be at most a $\mu(S_i^C)$ -centerpoint.

As in Lemma 2.2, the function $G: C \rightarrow \mathbb{R}$ defined by $G(y) := \sup_{\hat{v} \in S^{n-1}} \mu(H_y(\hat{v}))$ is continuous. Because $\lim_{i \rightarrow \infty} G(s_i) \leq 1 - \frac{1}{n+1}$, C is compact, and G is continuous,

it follows that there is some $c \in C \subset S$ such that $G(c) \leq 1 - \frac{1}{n+1}$. Therefore this c is a $\frac{1}{n+1}$ -centerpoint.

If we additionally assume S is convex, then the separating hyperplane theorem again gives that the centerpoint satisfies $c \in S$.

Finally, we remind the reader that Proposition 1.9 established the second part of the theorem, concerning sharpness. $\square$

3. Upper bound on needed subgradient calls

We require one final lemma for the application to convex optimization. In this lemma, as in the past, vol_g will denote the Riemannian volume measure and $B_g(x, r)$ the open ball of radius r around x .

Lemma 3.1. *Suppose f is convex and L -Lipschitz on its convex domain $S \subset M$, where M is a Hadamard manifold. Additionally assume the minimum x_* is in the ϵ -interior of S , meaning $B_g(x_*, \epsilon) \subset S$. Now suppose we are given a sequence of cutting planes $H_{c_i}(v_i)$, $i = 1, 2, \dots, N$ with $c_i \in S$ and $v_i \in \partial f_{c_i}$ such that the remaining feasible set $S' := S \cap_i H_{c_i}(v_i)$ satisfies the volume bound*

$$\text{vol}_g(S' := S \cap_i H_{c_i}(v_i)) < \frac{(\epsilon/L)^n}{n^n}.$$

Then one of the c_i satisfies $f(c_i) - f(x_) \leq \epsilon$.*

Proof. The main fact to be established is that $\text{vol}_g(B_g(x_*, \frac{\epsilon}{L})) > \text{vol}_g(S')$, as this implies that one of the complements of the halfspaces $H_{c_i}(v_i)$ must have intersected $B_g(x_*, \frac{\epsilon}{L})$. By volume comparison methods (see [5]), the volume of this geodesic ball is greater or equal to the volume of a Euclidean ball of equal radius. A reference justifying the exact version required is included in Appendix A as Theorem A.1. Hence we obtain

$$\text{vol}_g(B_g(x_*, \frac{\epsilon}{L})) > \frac{(\epsilon/L)^n}{n^n} > \text{vol}_g(S'),$$

by using $\frac{\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2} + 1)} \frac{\epsilon^n}{L^n} > \frac{1}{n^n} \frac{\epsilon^n}{L^n}$ in the first inequality.

It follows that there exists a point $x' \in B(x_*, \frac{\epsilon}{L})$ that lies in the complement of one of the halfspaces $H_{c_i}(v_i)$. From Lemma 1.7, $f(c_i) \leq f(x')$. The Lipschitz bound on f then gives

$$f(c_i) - f(x_*) \leq f(x') - f(x_*) \leq L \cdot d(x', x_*) \leq \epsilon. \quad \square$$

The proof of Theorem 1.3 is now a rather straightforward consequence.

Proof for Theorem 1.3. Lemma 3.1 shows that one of the origins of the cuts is ϵ from optimal for the function f as soon as the remaining set, denoted by S' , has volume $O(\frac{\epsilon^n}{n^n L^n})$.

We must only bound the number of halfspaces needed to reduce the volume of S' to this amount. Proceeding iteratively, apply Theorem 1.2 with μ being the

Riemannian volume measure restricted to S' (i.e. $\mu = (\mathbb{1}_{S'}/\text{vol}_g(S')) \cdot \text{vol}_g$). As the support of μ is contained in the convex set S , Theorem 1.2 shows that we may choose the cut centers to be $\frac{1}{n+1}$ -centerpoints $c_i \in S$ for the remaining set $S' \subset S$, so that the volume is reduced by a factor $(1 - \frac{1}{n+1})$ each cut. This means the number of iterations needed is $O(n^2 \log(nL \text{vol}_g(S)\epsilon^{-1}))$. $\square$

Acknowledgments. Many people have listened patiently and offered advice which has been helpful for this work. In particular, many thanks to Alexander Appleton, Richard Bamler, Andrew Hanlon, and Nikhil Srivastava.

Appendix

A. Riemannian overview

We will be working in the setting of Riemannian geometry, but will not use much machinery. We provide an informal overview. The definitions we introduce here are generally standard and formalized in introductory texts, one such being [10].

An n -dimensional (smooth) manifold M can be understood as a space that is locally diffeomorphic to $\mathbb{R}^n$, so we identify such subsets of M with coordinates $(x_1, \dots, x_n)$. This allows us to define smooth curves $\gamma: \mathbb{R} \rightarrow M$, by requiring that their coordinate representations $(x_1(t), \dots, x_n(t))$ are smooth. We may define velocities $\gamma'(t)$ by associating them with $(x'_1(t), \dots, x'_n(t))$, leading to the notion of the tangent spaces $T_x M \cong \mathbb{R}^n$.

Riemannian manifolds additionally specify a metric for measuring the size of these velocities, by defining an inner product $\langle \cdot, \cdot \rangle_x$ on the tangent space of each $x \in M$. This immediately enables the definition of curve length, as $\int |\gamma'(t)|_{\gamma(t)} dt$. It also gives a method of measuring volume; if g_{ij} is the bilinear form for the metric in a local coordinate choice, then $\sqrt{|g|} dx^1 \wedge \dots \wedge dx^n$ is the Riemannian volume form.

It also turns out to be helpful to compute directional derivatives for vector fields (or acceleration along curves). Requiring a few natural conditions leads to a unique Riemannian connection $\nabla: T_x M \times T_x M \rightarrow T_x M$ determined by the metric. It is known as the Levi-Civita connection. In the coordinates of a local frame $E = (\vec{e}_1, \dots, \vec{e}_n)$, which provides a basis for the tangent spaces of a neighborhood, the Riemannian connection is given by

$$\nabla_{\vec{e}_i} \vec{e}_j = \Gamma_{ji}^k \vec{e}_k,$$

for the Christoffel symbols Γ_{ij}^k . When the acceleration of a curve is 0, i.e. we have $\nabla_{\gamma'(t)} \gamma'(t) \equiv 0$, we say that the curve is a geodesic. This is a second order non-linear ODE system for $\gamma(t) = (x_1(t), \dots, x_n(t))$: $\ddot{x}^k(t) + \dot{x}^i \dot{x}^j \Gamma_{ij}^k(x(t)) = 0$.

A unique solution will exist locally, provided we specify the initial position and velocity. The exponential map is defined by $\exp_p(v) = \gamma(1)$. For Hadamard manifolds, the exponential map is well-defined for any values of p and v .

At any $x \in M$ we may consider the image of a tangent plane σ_x spanned by $v, w \in T_x M$. Locally around x the image is a surface. The sectional curvature

of the 2-plane σ_x is defined to be the Gaussian curvature of the image surface at x . Lower and upper bounds of the sectional curvature enable generalizations of Euclidean tools like ball volume and triangle trigonometry estimates. The Bishop-Gromov volume comparison theorem is one important result along these lines. Although usually stated for its volume upper bound by assuming just a lower bound on curvature, it is understood that the proof also provides a lower volume bound [11]. Here we state a specialization of this theorem that suffices for our application,

Theorem A.1. (Bishop-Gromov) *Suppose M is a Hadamard manifold. Let vol_g denote the Riemannian volume and $B_g(x, r)$ denote the open ball of radius r around x . Then*

$$\text{vol}_g(B_g(x, r)) \leq \frac{\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2} + 1)} r^n.$$

The right side of this inequality is the volume of a Euclidean ball of radius r .

Hadamard manifolds are simply connected manifolds of non-positive sectional curvature. They have been extensively studied in mathematical literature. We collect a few commonly used facts which we made use of or provide intuition. For Hadamard manifolds,

- The exponential maps $\exp_x(\cdot)$ are diffeomorphisms from $T_x M$ to M (Cartan-Hadamard theorem).
- The distance to a point, $d(x, \cdot)$, is strictly convex. The distance to a closed, convex set is convex.
- Geodesics between points are unique and distance minimizing.
- Projection onto closed, convex sets is well defined and continuous.

All of these properties can be found in [2].

In the proof of Theorem 1.2, we made use of the separating hyperplane theorem of convex geometry. Here we briefly state and prove a version sufficient for this application.

Lemma A.2. *Suppose M is a Hadamard manifold, $S \subset M$ is a closed, convex set, and $p \notin S$. Then there is a halfspace based at p satisfying $H_p(v) \cap S = \emptyset$.*

Proof. Consider the function $f(x) = d(x, S)$. By [2] this function is convex, hence for any $p \notin S$ there exists a subgradient $v \in \partial f_p$ (see Definition 1.5). Then $H_p(v)$ is such a separating hyperplane, by Lemma 1.7. $\square$

In the introduction, we mentioned that the gradient of a differentiable convex function is a subgradient. We provide a short justification for this simple fact, as it is an important source of subgradients. In Riemannian geometry, the gradient is defined by duality using the metric; that is, ∇ satisfies $\langle \nabla f, \cdot \rangle = df(\cdot)$.

Lemma A.3. *Let $f(x): \rightarrow \mathbb{R}$ be convex along geodesics as well as differentiable. Then $f(y) \geq f(x) + \langle \nabla f(x), \exp_x^{-1}(y) \rangle_x$.*

Proof. Let $y = \exp_x(t_0 v)$. That f is convex on geodesics means $f(\exp_x(tv))$ is convex in t , so

$$f(y) \geq f(x) + t_0 \frac{d}{dt} f(\exp_x(tv))|_0.$$

But using the chain rule and that $d(\exp_x)|_0 = I$ (see [10]),

$$\frac{d}{dt} f(\exp_x(tv))|_0 = df(d\exp_x|_0(v)) = df(v) = \langle \nabla f(x), v \rangle_x. \quad \square$$

References

- [1] Z. Allen-Zhu, A. Garg, Y. Li, R. Oliveira, A. Wigderson: *Operator scaling via geodesically convex optimization, invariant theory and polynomial identity testing*, STOC 50 (2018) 172–181.
- [2] W. Ballmann: *Lectures on Spaces of Nonpositive Curvature*, Springer, Berlin (1995).
- [3] I. Benjamini, R. Eldan: *Convex hulls in hyperbolic space*, Geometriae Dedicata 160(1) (2012) 365–371.
- [4] J. Bennett, A. Carbery, M. Christ, T. Tao: *The Brascamp-Lieb inequalities: finiteness, structure and extremals*, Geom. Funct. Analysis 17 (2008) 1343–1415.
- [5] J. Cheeger, D. G. Ebin: *Comparison Theorems in Riemannian Geometry*, AMS Chelsea Publishing, Providence (2008).
- [6] D. Cordero-Erausquin, R. J. McCann, M. Schmuckenschläger: *A Riemannian interpolation inequality à la Borell, Brascamp and Lieb*, Inventiones Mathematicae 146 (2001) 219–257.
- [7] B. Grünbaum: *Partitions of mass-distributions and of convex bodies by hyperplanes*, Pacific J. Math. 10(4) (1960) 1257–1261.
- [8] S. Ivanov: *On Helly’s theorem in geodesic spaces*, Electronic Res. Announcements 21 (2014) no. 1935-9179-2014-0-109.
- [9] Y. Ledyev, J. S. Treiman, Q. J. Zhu: *Helly’s intersection theorem on manifolds of nonpositive curvature*, J. Convex Analysis 13(3-4) (2006) 785–798.
- [10] J. Lee: *Riemannian Manifolds: An Introduction to Curvature*, Springer, Berlin (1997).
- [11] P. Petersen: *Riemannian Geometry*, Springer, New York (2006).
- [12] S. Sra, R. Hosseini: *Geometric optimisation on positive definite matrices with application to elliptically contoured distributions*, in: *Proc. of the 26th International Conference on Neural Information Processing Systems. Volume 2*, NIPS’13, Lake Tahoe, Curran Associates (2013) 2562–2570.
- [13] C. Udriste: *Convex Functions and Optimization Methods on Riemannian Manifolds*, Kluwer Academic Publishers, Dordrecht (1994).
- [14] H. Zhang, S. J. Reddi, S. Sra: *Riemannian SVRG: Fast stochastic optimization on Riemannian manifolds*, in: *Advances in Neural Information Processing Systems 29*, Curran Associates (2016) 4592–4600.
- [15] H. Zhang, S. Sra: *First-order methods for geodesically convex optimization*, in: *Proceedings of the 29th Conference on Learning Theory*, PMLR (2016) 1617–1638.