

The Gerrymandering Jumble: Map Projections Permute Districts' Compactness Scores

Assaf Bar-Natan^a, Lorenzo Najt^b, and Zachary Schutzman^c

^aUniversity of Toronto, Department of Mathematics, Toronto, ON, Canada;

^bUniversity of Wisconsin, Department of Mathematics, Madison, WI, USA;

^cUniversity of Pennsylvania, Department of Computer and Information Science, Philadelphia, PA, USA

ABSTRACT

In political redistricting, the *compactness* of a district is used as a quantitative proxy for its fairness. Several well-established, yet competing, notions of geographic compactness are commonly used to evaluate the shapes of regions, including the *Polsby-Popper score*, the *convex hull score*, and the *Reock score*, and these scores are used to compare two or more districts or plans. In this paper, we prove mathematically that any *map projection* from the sphere to the plane reverses the ordering of the scores of some pair of regions for all three of these scores. We evaluate these results empirically on United States congressional districts and demonstrate that this order-reversal does occur in practice with respect to commonly-used projections. Furthermore, the Reock score ordering in particular appears to be quite sensitive to the choice of map projection.

1. Introduction

Striving for the *geographic compactness* of electoral districts is a traditional principle of redistricting (Altman, 1998), and, to that end, many jurisdictions have included the criterion of compactness in their legal code for drawing districts. Some of these include Iowa's measuring the perimeter of districts (Iowa Code §42.4(4)), Maine's minimizing travel time within a district (Maine Statute §1206-A), and Idaho's avoiding drawing districts which are 'oddly shaped' (Idaho Statute 72-1506(4)). Such measures can vary widely in their precision, both mathematical and otherwise. Computing the perimeter of districts is a very clear definition, minimizing travel time is less so, and what makes a district oddly shaped or not seems rather challenging to consider from a rigorous standpoint.

While a strict definition of when a district is or is not 'compact' is quite elusive, the purpose of such a criterion is much easier to articulate. Simply put, a district which is bizarrely shaped, such as one with small tendrils grabbing many distant

This is an Accepted Manuscript of an article published by Taylor & Francis in *Cartography and Geographic Information Systems* on May 13, 2020, available online: <http://www.tandfonline.com/10.1080/15230406.2020.1737575>. Please cite that version as the Version of Record.

This work was partially completed while the authors were at the Voting Rights Data Institute in the summer of 2018, which was generously supported by the Amar G. Bose Grant. L. Najt was also partially supported by a grant from the NSF GEAR Network (DMS-1107452, DMS-1107263, DMS-1107367) and an NSF RTG award (DMS-1502553).

Corresponding Author: Zachary Schutzman. Email: ianzach@seas.upenn.edu

chunks of territory, probably wasn't drawn like that by accident. Such a shape need not be drawn for nefarious purposes, but its unusual nature should trigger closer scrutiny. Measures to compute the geographic compactness of districts are intended to formalize this quality of 'bizarreness' mathematically. We briefly note here that the term *compactness* is somewhat overloaded, and that we exclusively use the term to refer to the shape of geographic regions and not to the topological definition of the word.

People have formally studied geographic compactness for nearly two hundred years, and, over that period, scientists and legal scholars have developed many formulas to assign a numerical measure of 'compactness' to a region such as an electoral district (Young, 1988). Three of the most commonly discussed formulations are the *Polsby-Popper score*, which measures the normalized ratio of a district's area to the square of its perimeter, the *convex hull score*, which measures the ratio of the area of a district to the smallest convex region containing it, and the *Reock score*, which measures the ratio of the area of a district to the area of the smallest circular disc containing it. Each of these measures is appealing at an intuitive level, since they assign to a district a single scalar value between zero and one, which presents a simple method to compare the relative compactness of two or more districts. Additionally, the mathematics underpinning each is widely understandable by the relevant parties, including lawmakers, judges, advocacy groups, and the general public.

However, none of these measures truly discerns which districts are 'compact' and which are not. For each score, we can construct a mathematical counterexample for which a human's intuition and the score's evaluation of a shape's compactness differ. A region which is roughly circular but has a jagged boundary may appear compact to a human's eye, but such a shape has a very poor Polsby-Popper score. Similarly, a very long, thin rectangle appears non-compact to a person, but has a perfect convex hull score. Additionally, these scores often do not agree. The long, thin rectangle has a very poor Polsby-Popper score, and the ragged circle has an excellent convex hull score. These issues are well-studied by political scientists and mathematicians alike (Barnes and Solomon, 2020; Frolov, 1975; Maceachren, 1985; Polsby and Popper, 1991).

In this paper, we propose a further critique of these measures, namely *sensitivity under the choice of map projection*. Each of the compactness scores named above is defined as a tool to evaluate geometric shapes in the plane, but in reality we are interested in analyzing shapes which sit on the surface of the planet Earth, which is (roughly) spherical. When we analyze the geometric properties of a geographic region, we work with a *projection* of the Earth onto a flat plane, such as a piece of paper or the screen of a computer. Therefore, when a shape is assigned a compactness score, it is implicitly done with respect to some choice of map projection. We prove that this may have serious consequences for the comparison of districts by these scores. Because there is no projection from the sphere to the plane which preserves 'too many' metric properties and most compactness scores synthesize several of these properties, it is unreasonable to expect any projection to preserve the numerical values of these scores for all regions. However, since there are projections which preserve *some* geometric properties, such as those which preserve the area of all regions or conformal projections which preserve the angle of intersection of all line segments, we might ask a weaker question and consider whether there is a projection which can preserve the *induced ordering* of a compactness score over all regions.

In particular, we consider the Polsby-Popper, convex hull, and Reock scores on the sphere, and demonstrate that for any choice of map projection, there are two regions, A and B , such that A is more compact than B on the sphere but B is more compact

than A when projected to the plane. We prove our results in a theoretical context before evaluating the extent of this phenomenon empirically. We find that with real-world examples of Congressional districts, the effect of the commonly-used *Plate carée* projection, which treats latitude-longitude coordinates as Cartesian coordinate pairs, on the convex hull and Polsby-Popper scores is relatively minor, but the impact on Reock scores is quite dramatic, which may have serious implications for the use of this measure as a tool to evaluate geographic compactness.

1.1. Organization

For each of the compactness scores we analyze, our proof that no map projection can preserve their order follows a similar recipe. We first use the fact that any map projection which preserves an ordering must preserve the *maximizers* in that ordering. In other words, if there is some shape which a score says is “the most compact” on the sphere but the projection sends this to a shape in the plane which is “not the most compact”, then whatever shape *does* get sent to the most compact shape in the plane leapfrogs the first shape in the induced ordering. For all three of the scores we study, such a maximizer exists.

Using this observation, we can restrict our attention to those map projections which preserve the maximizers in the induced ordering, then argue that any projection in this restricted set must permute the order of scores of some pair of regions.

Preliminaries We first introduce some definitions and results which we will use to prove our three main theorems. Since spherical geometry differs from the more familiar planar geometry, we carefully describe a few properties of spherical lines and triangles to build some intuition in this domain.

Convex Hull For the convex hull score, we first show that any projection which preserves the maximizers of the convex hull score ordering must maintain certain geometric properties of shapes and line segments between the sphere and the plane. Using this, we demonstrate that no map projection from the sphere to the plane can preserve these properties, and therefore no such convex hull score order preserving projection exists.

Reock For the Reock score, we follow a similar tack, first showing that any order-preserving map projection must also preserve some geometric properties and then demonstrating that such a map projection cannot exist.

Polsby-Popper To demonstrate that there is no projection which maintains the score ordering induced by the Polsby-Popper score, we leverage the difference between the *isoperimetric inequalities* on the sphere and in the plane, in that the inequality for the plane is scale invariant in that setting but not on the sphere, in order to find a pair of regions in the sphere, one more compact than the other, such that the less compact one is sent to a circle under the map projection.

Empirical Results We finally examine the impact of the Cartesian latitude-longitude map projection on the convex hull, Reock, and Polsby-Popper scores and the ordering of regions under these scores. While the impacts of the projection on the con-

vex hull and Polsby-Popper scores and their orderings are not severe, the Reock score and the Reock score ordering both change dramatically under the map projection.

2. Preliminaries

We begin by introducing some necessary observations, definitions, and terminology which will be of use later.

2.1. Spherical Geometry

In this section, we present some basic results about spherical geometry with the goal of proving Girard's Theorem, which states that the area of a triangle on the unit sphere is the sum of its interior angles minus π . Readers familiar with this result should feel free to skip ahead.

We use \mathbb{R}^2 to denote the Euclidean plane with the usual way of measuring distances,

$$d(x, y) = \sqrt{(x - y)^2};$$

similarly, \mathbb{R}^3 denotes Euclidean 3-space. We use \mathbb{S}^2 to denote the *unit 2-sphere*, which can be thought of as the set of points in \mathbb{R}^3 at Euclidean distance one from the origin.

In this paper, we only consider the sphere and the plane, and leave the consideration of other surfaces, measures, and metrics to future work.

Definition 2.1. On the sphere, a **great circle** is the intersection of the sphere with a plane passing through the origin. These are the circles on the sphere with radius equal to that of the sphere. See Figure 1 for an illustration.

Definition 2.2. Lines in the plane and great circles on the sphere are called **geodesics**. A **geodesic segment** is a line segment in the plane and an arc of a great circle on the sphere.

The idea of geodesics generalizes the notion of 'straight lines' in the plane to other settings. One critical difference is that in the plane, there is a unique line passing through any two distinct points and a unique line segment joining them. On the sphere, there will typically be a unique great circle and *two* geodesic segments through a pair of points, with the exception of one case.

Definition 2.3. A **triangle** in the plane or the sphere is defined by three distinct points and the shortest geodesics connecting each pair of points.

Observation 1. Given any two points p and q on the sphere which are not antipodal, meaning that our points aren't of the form $p = (x, y, z)$ and $q = (-x, -y, -z)$, there is a unique great circle through p and q and therefore two geodesic segments joining them.

If p and q are antipodal, then any great circle containing one must contain the other as well, so there are infinitely many such great circles. For any two non-antipodal points on the sphere, one of the geodesic segments will be shorter than the other. This shorter geodesic segment is the shortest path between the points and its length is the metric distance between p and q .




Figure 1. A great circle on the sphere with its identifying plane.




Figure 2. Two great circles meet at antipodal points.

We now have enough terminology to show a very important fact about spherical geometry. This observation is one of the salient features which distinguishes it from the more familiar planar geometry.

Claim 2.4. *Any pair of distinct great circles on the sphere intersect exactly twice, and the points of intersection are antipodes.*

Why is this weird? In the plane, it is always the case that any pair of distinct lines intersects exactly once or never, in which case we call them *parallel*. Since distinct great circles on the sphere intersect exactly twice, there is no such thing as ‘parallel lines’ on the sphere, and we have to be careful about discussing ‘the’ intersection of two great circles since they do not meet at a unique point. Furthermore, it is not the case that there is a unique segment of a great circle connecting any two points; there are two, but unless our two points are antipodes, one of the two segments will be shorter.

Another difference between spherical and planar geometry appears when computing the angles of triangles. In the planar setting, the sum of the interior angles of a triangle is always π , regardless of its area. However, in the spherical case we can construct a triangle with three right angles. The north pole and two points on the equator, one a quarter of the way around the sphere from the other, form such a triangle. Its area is one eighth of the whole sphere, or $\frac{\pi}{2}$, which is, not coincidentally, equal to $\frac{\pi}{2} + \frac{\pi}{2} + \frac{\pi}{2} - \pi$. Girard’s theorem, which we will prove below, connects the total angle to the area of a spherical triangle.

In order to show Girard’s Theorem, we need some way to translate between *angles* and *area*. To do that, we’ll use a shape which doesn’t even exist in the plane: the *diangle* or *lune*. We know that two great circles intersect at two antipodal points, and we can also see that they cut the surface of the sphere into four regions. Consider

one of these regions. Its boundary is a pair of great circle segments which connect antipodal points and meet at some angle $\theta \leq \pi$ at both of these points.




Figure 3. A lune corresponding to an angle θ .

Using that the surface area of a unit sphere is 4π , computing the area of a lune with angle θ is straightforward.

Claim 2.5. *Consider a lune whose boundary segments meet at angle θ . Then the area of this lune is 2θ .*

Now that we have a tool that lets us relate angles and areas, we can prove Girard's Theorem.

Lemma 2.6. (*Girard's Theorem*)

The sum of the interior angles of a spherical triangle is strictly greater than π . More specifically, the sum of the interior angles is equal to π plus the area of the triangle.

Proof. Consider a triangle T on the sphere with angles θ_1 , θ_2 , and θ_3 . Let $\text{area}(T)$ denote the area of this triangle. If we extend the sides of the triangle to their entire great circles, each pair intersects at the vertices of T as well as the three points antipodal to the vertices of T , and at the same angles at antipodal points. This second triangle is congruent to T , so its area is also $\text{area}(T)$. Each pair of great circles cuts the sphere into four lunes, one which contains T , one which contains the antipodal triangle, and two which do not contain either triangle. We are interested in the three pairs of lunes which do contain the triangles. We will label these lunes by their angles, so we have a lune $L(\theta_1)$ and its antipodal lune $L'(\theta_1)$, and we can similarly define $L(\theta_2)$, $L'(\theta_2)$, $L(\theta_3)$, and $L'(\theta_3)$.

We have six lunes. In total, they cover the sphere, but share some overlap. If we remove T from two of the three which contain it and the antipodal triangle from two of the three which contain it, then we have six non-overlapping regions which cover the sphere, so the area of the sphere must be equal to the sum of the areas of these six regions.

By the earlier claim, we know that the areas of the lunes are twice their angles, so we can write this as

$$4\pi = 2\theta_1 + 2\theta_1 + (2\theta_2 - \text{area}(T)) + (2\theta_2 - \text{area}(T)) \\ + (2\theta_3 - \text{area}(T)) + (2\theta_3 - \text{area}(T))$$

and rearrange to get

$$\theta_1 + \theta_2 + \theta_3 = \pi + \text{area}(T),$$

which is exactly the statement we wanted to show. \square




Figure 4. A spherical triangle and the antipodal triangle define six lunes.

We will need one more fact about spherical triangles before we conclude this section. It follows immediately from the Spherical Law of Cosines.

Fact 1. An equilateral triangle is equiangular, and vice versa, where *equilateral* means that the three sides have equal length and *equiangular* means that the three angles all have the same measure.

An astute reader may notice that this result is also true of planar triangles, and the planar version follows from Propositions I.6 and I.8 in Euclid's *Elements* (Byrne, 1847; Crowell, 2016). Since Euclid's proof doesn't rely on the existence of parallel lines, this fact can alternatively be shown using his argument.

2.2. Some Definitions

Now that we have the necessary tools of spherical geometry, we will wrap up this section with a battery of definitions. We carefully lay these out so as to align with an intuitive understanding of the concepts and to appease the astute reader who may be concerned with edge cases, geometric weirdness, and nonmeasurability. Throughout, we implicitly consider all figures on the sphere to be strictly contained in a hemisphere.

Definition 2.7. A **region** is a non-empty, open subset Ω of \mathbb{S}^2 or \mathbb{R}^2 such that Ω is bounded and its boundary is piecewise smooth.

We choose this definition to ensure that the *area* and *perimeter* of the region are well-defined concepts. This eliminates pathological examples of open sets whose boundaries have non-zero area or edge cases like considering the whole plane a 'region'.

Definition 2.8. A **compactness score function** \mathcal{C} is a function from the set of all regions to the non-negative real numbers or infinity. We can compare the scores of

any two regions, and we adopt the convention that *more compact* regions have *higher* scores. That is, region A is at least as compact as region B if and only if $\mathcal{C}(A) \geq \mathcal{C}(B)$.

The final major definition we need is that of a *map projection*. In reality, the regions we are interested in comparing sit on the surface of the Earth (i.e. a sphere), but these regions are often examined after being projected onto a flat sheet of paper or computer screen, and so have been subject to such a projection.

Definition 2.9. A **map projection** φ is a diffeomorphism from a region on the sphere to a region of the plane.

We choose this definition, and particularly the term *diffeomorphism*, to ensure that φ is smooth, its inverse φ^{-1} exists and is smooth, and both φ and φ^{-1} send regions in their domain to regions in their codomain. Throughout, we use φ to denote such a function from a region of the sphere to a region of the plane and φ^{-1} , to denote the inverse which is a function from a region of the plane back to a region of the sphere.

Since the image of a region under a map projection φ is also a region, we can examine the compactness score of that region both before and after applying φ , and this is the heart of the problem we address in this paper. We demonstrate, for several examples of compactness scores \mathcal{C} , that the order induced by \mathcal{C} is different than the order induced by $\mathcal{C} \circ \varphi$ for *any* choice of map projection φ .

Definition 2.10. We say that a map projection φ **preserves the compactness score ordering** of a score \mathcal{C} if for any regions Ω, Ω' in the domain of φ , $\mathcal{C}(\Omega) \geq \mathcal{C}(\Omega')$ if and only if $\mathcal{C}(\varphi(\Omega)) \geq \mathcal{C}(\varphi(\Omega'))$ in the plane.

This is a weaker condition than simply preserving the raw compactness scores. If there is some map projection which results in adding .1 to the score of each region, the raw scores are certainly not preserved, but the ordering of regions by their scores is. Additionally, φ preserves a compactness score ordering if and only if φ^{-1} does.

Definition 2.11. A **cap** on the sphere \mathbb{S}^2 is a region on the sphere which can be described as all of the points on the sphere to one side of some plane in \mathbb{R}^3 . A cap has a *height*, which is the largest distance between this cutting plane and the cap, and a *radius*, which is the radius of the circle formed by the intersection of the plane and the sphere. See Figure 5 for an illustration.




Figure 5. The height h and radius r of a spherical cap.

3. Convex Hull

We first consider the *convex hull score*. We briefly recall the definition of a convex set and then define this score function.

Definition 3.1. A set in \mathbb{R}^2 or \mathbb{S}^2 is **convex** if every shortest geodesic segment between any two points in the set is entirely contained within that set.

Definition 3.2. Let $\text{conv}(\Omega)$ denote the *convex hull* of a region Ω in either the sphere or the plane, which is the smallest convex region containing Ω . Then we define the *convex hull score* of Ω as

$$\text{CH}(\Omega) = \frac{\text{area}(\Omega)}{\text{area}(\text{conv}(\Omega))}.$$

Since the intersection of convex sets is a convex set, there is a unique smallest (by containment) convex hull for any region Ω .




Figure 6. A region Ω and its convex hull.

Suppose that our map projection φ does preserve the ordering of regions induced by the convex hull score. We begin by observing that such a projection must preserve certain geometric properties of regions within its domain.

Lemma 3.3. *Let φ be a map projection from some region of the sphere to a region of the plane. If φ preserves the convex hull compactness score ordering, then the following must hold:*

- (1) φ and φ^{-1} send convex regions in their domains to convex regions in their codomains.
- (2) φ sends every segment of a great circle in its domain to a line segment in its codomain. That is, it preserves geodesics.¹
- (3) There exists a region U in the domain of φ such that for any regions $A, B \subset U$, if A and B have equal area on the sphere, then $\varphi(A)$ and $\varphi(B)$ have equal area in the plane. The same holds for φ^{-1} for all pairs of regions inside of $\varphi(U)$.

Proof. The proof of (1) follows from the idea that any projection which preserves the convex hull score ordering of regions must preserve the maximizers in that ordering. Here, the maximizers are convex sets.

¹Such a projection is sometimes called a *geodesic map*.




Figure 7. Two equal area regions A and B removed from U to form the regions X and Y .

To show (2) we suppose, for the sake of contradiction, that there is some geodesic segment s in U such that $\varphi(s)$ is not a line segment. Construct two convex spherical polygons L and M inside of U which both have s as a side.




Figure 8. If $\varphi(s)$ is not a line segment, then one of $\varphi(M)$ or $\varphi(L)$ is not convex.

By (1), φ must send both of these polygons to convex regions in the plane, but this is not the case. All of the points along $\varphi(s)$ belong to both $\varphi(L)$ and $\varphi(M)$, but since $\varphi(s)$ is not a line segment, we can find two points along it which are joined by some line segment which contains points which only belong to $\varphi(L)$ or $\varphi(M)$, which means that at least one of these convex spherical polygons is sent to something non-convex in the plane, which contradicts our assumption. See Figure 8 for an illustration.

That φ^{-1} sends line segments in the plane to great circle segments on the sphere is shown analogously. This completes the proof of (2).

To show (3), let U be some convex region in the domain of φ . Take A, B to be regions of equal area such that A and B are properly contained in the interior of U , as in Figure 7. Define two new regions $X = U \setminus A$ and $Y = U \setminus B$, i.e. these regions are equal to U with A or B deleted, respectively.

The cap U is itself the convex hull of both X and Y , and since A and B have equal area, we have that $\text{CH}(X) = \text{CH}(Y)$. Since U is a cap, it is convex, so by (1), $\varphi(U)$ is also convex. Since φ preserves the ordering of convex hull scores and X and Y had equal scores on the sphere, φ must send X and Y to regions in the plane which also have the same convex hull score as each other. Furthermore, the convex hulls of $\varphi(X)$ and $\varphi(Y)$ are $\varphi(U)$.

By definition, we have

$$\text{CH}(X) = \text{CH}(Y)$$

and by the construction of X and Y , we have

$$\frac{\text{area}(\varphi(U)) - \text{area}(\varphi(A))}{\text{area}(\varphi(U))} = \frac{\text{area}(\varphi(U)) - \text{area}(\varphi(B))}{\text{area}(\varphi(U))}$$

$$\text{area}(\varphi(A)) = \text{area}(\varphi(B))$$

which is what we wanted to show. The proof that φ^{-1} also has this property is analogous. □

We can now show that no map projection can preserve the convex hull score ordering of regions by demonstrating that there is no projection from a patch on the sphere to the plane which has all three of the properties described in Lemma 3.3.

Theorem 3.4. *There does not exist a map projection with the three properties in Lemma 3.3*

Proof. Assume that such a map, φ , exists, and restrict it to U as above. Let $T \subset U$ be a small equilateral spherical triangle centered at the center of U . Let T_1 and T_2 be two congruent triangles meeting at a point and each sharing a face with T , as in Figure 9.




Figure 9. The spherical regions T, T_1, T_2 .

We first argue that the images of $T \cup T_1$ and $T \cup T_2$ are parallelograms.

Without loss of generality, consider $T \cup T_1$. By construction, it is a convex spherical quadrilateral. By symmetry, its geodesic diagonals on the sphere divide it into four triangles of equal area. To see this, consider the geodesic segment which passes through the vertex of T opposite the side shared with T_1 which divides T into two smaller triangles of equal area. Since T is equilateral, this segment meets the shared side at a right angle at the midpoint, and the same is true for the area bisector of T_1 . Since both of these bisectors meet the shared side at a right angle and at the same point, together they form a single geodesic segment, the diagonal of the quadrilateral. Since the diagonal cuts each of T and T_1 in half, and T and T_1 have the same area, the four triangles formed in this construction have the same area.

Since φ sends spherical geodesics to line segments in the plane, it must send $T \cup T_1$ to a Euclidean quadrilateral Q whose diagonals are the images of the diagonals of the spherical quadrilateral $T \cup T_1$.

Since φ sends equal area regions to equal area regions, it follows that the diagonals of Q split it into four equal area triangles.

We now argue that this implies that Q is a Euclidean parallelogram by showing that its diagonals bisect each other. Since the four triangles formed by the diagonals of Q are all the same area, we can pick two of these triangles which share a side and consider the larger triangle formed by their union. One side of this triangle is a diagonal d_1 of Q and its area is bisected by the other diagonal d_2 , which passes through d_1 and its opposite vertex. The area bisector from a vertex, called the *median*, passes through the midpoint of the side d_1 , meaning that the diagonal d_2 bisects the diagonal d_1 . Since this holds for any choice of two adjacent triangles in Q , the diagonals must bisect each other, so Q is a parallelogram.




Figure 10. The image under φ of T, T_1, T_2 which form the quadrilateral in the plane.

Since $T \cup T_1$ and $T \cup T_2$ are both spherical quadrilaterals which overlap on the spherical triangle T , the images of $T \cup T_1$ and $T \cup T_2$ are Euclidean parallelograms of equal area which overlap on a shared triangle $\varphi(T)$. See Figure 10 for an illustration.

Because the segment ℓ is parallel to m_1 and m_2 , m_1 and m_2 are parallel to each other, and because they meet at the point shared by all three triangles, m_1 and m_2 together form a single segment parallel to ℓ . Therefore, the image of the three triangles forms a quadrilateral in the plane. Therefore, the image of $T \cup T_1 \cup T_2$ has a boundary consisting of four line segments.

To find the contradiction, consider the point on the sphere shared by T , T_1 , and T_2 . Since these triangles are all equilateral spherical triangles, the three angles at this point are each strictly greater than $\frac{\pi}{3}$ radians, because the sum of interior angles on a triangle is strictly greater than π . so, the total measure of the three angles at this point is greater than π . Therefore, the two geodesic segments which are part of the boundaries of T_1 and T_2 meet at this point at an angle of measure strictly greater than π . Therefore, together they do not form a single geodesic. On the sphere, the region $T \cup T_1 \cup T_2$ has a boundary consisting of five geodesic segments whereas its image has a boundary consisting of four, which contradicts the assumption that φ and φ^{-1} preserve geodesics. \square

This implies that no map projection can preserve the ordering of regions by their convex hull scores, which is what we aimed to show.

4. Reock

Let $\text{circ}(\Omega)$ denote the *smallest bounding circle* (smallest bounding cap on the sphere) of a region Ω . Then the *Reock score* of Ω is

$$\text{Reock}(\Omega) = \frac{\text{area}(\Omega)}{\text{area}(\text{circ}(\Omega))}.$$

We again consider what properties a map projection φ must have in order to preserve the ordering of regions by their Reock scores.

Lemma 4.1. *If φ preserves the ordering of regions induced by their Reock scores, then the following must hold:*

- (1) φ sends spherical caps in its domain to Euclidean circles in the plane, and φ^{-1} does the opposite.
- (2) There exists a region U in the domain of φ such that for any regions $A, B \subset U$, if A and B have equal area on the sphere, then $\varphi(A)$ and $\varphi(B)$ have equal area in the plane. The same holds for φ^{-1} for all pairs of regions inside of $\varphi(U)$.

Proof. Similarly to the convex hull setting, the proof of (1) follows from the requirement that φ preserves the maximizers in the compactness score ordering. In the case of the Reock score, the maximizers are caps in the sphere and circles in the plane.

To show (2), let κ be a cap in the domain of φ , and let $A, B \subset \kappa$ be two regions of equal area properly contained in the interior of κ . Then, define two new regions $X = \kappa \setminus A$ and $Y = \kappa \setminus B$, which can be thought of as κ with A and B deleted, respectively.

Since κ is the smallest bounding cap of X and Y and since A and B have equal areas, $\text{Reock}(X) = \text{Reock}(Y)$. Furthermore, by (1), φ must send κ to some circle in the plane, which is the smallest bounding circle of $\varphi(X)$ and $\varphi(Y)$. Since φ preserves the ordering of Reock scores, it must be that $\varphi(X)$ and $\varphi(Y)$ have identical Reock scores in the plane.

By definition, we can write

$$\begin{aligned} \text{Reock}(X) &= \text{Reock}(Y) \\ \frac{\text{area}(\varphi(X))}{\text{area}(\varphi(\kappa))} &= \frac{\text{area}(\varphi(Y))}{\text{area}(\varphi(\kappa))} \end{aligned}$$

and by the construction of X and Y , we have

$$\begin{aligned} \frac{\text{area}(\varphi(\kappa)) - \text{area}(\varphi(A))}{\text{area}(\varphi(\kappa))} &= \frac{\text{area}(\varphi(\kappa)) - \text{area}(\varphi(B))}{\text{area}(\varphi(\kappa))} \\ \text{area}(\varphi(A)) &= \text{area}(\varphi(B)), \end{aligned}$$

meaning that $\text{area}(\varphi(A)) = \text{area}(\varphi(B))$. Thus, for all pairs of regions of the same area inside of κ , the images under φ of those regions will have the same area as well.

The same construction works in reverse, which demonstrates that φ^{-1} also sends regions of equal area in some circle in the plane to regions of equal area in the sphere. \square

We can now show that no such φ exists. Rather than constructing a figure on the sphere and examining its image under φ , it will be more convenient to construct a figure in the plane and reason about φ^{-1} .

Theorem 4.2. *There does not exist a map projection with the two properties in Lemma 4.1.*

Proof. Assume that such a φ does exist and restrict its domain to a cap κ as above. This corresponds to a restriction of the domain of φ^{-1} to a circle in the plane. Inside of this circle, draw seven smaller circles of equal area tangent to each other as in Figure 11.




Figure 11. Seven circles arranged as in the construction for Theorem 4.2.

Under φ^{-1} , they must be sent to an similar configuration of equal-area caps on the sphere .

However, the radius of a of a spherical cap is determined by its area, so since the areas of these caps are all the same, their radii must be as well. Thus, the midpoints of these caps form six equilateral triangles on the sphere which meet at a point. However, this is impossible, as the three angles of an equilateral triangle on the sphere must all be greater than $\frac{\pi}{3}$, but the total measure of all the angles at a point must be equal to 2π , which contradicts the assumption that such a φ exists. \square

This shows that no map projection exists which preserves the ordering of regions by their Reock scores.

5. Polsby-Popper

The final compactness score we analyze is the *Polsby-Popper score*, which takes the form of an *isoperimetric quotient*, meaning it measures how much area a region's perimeter encloses, relative to all other regions with the same perimeter.

Definition 5.1. The Polsby-Popper score of a region Ω is defined to be

$$\text{PP}(\Omega) = \frac{4\pi \cdot \text{area}(\Omega)}{\text{perim}(\Omega)^2}$$

in either the sphere or the plane, and area and perim are the area and perimeter of Ω , respectively.

The ancient Greeks were first to observe that if Ω is a region in the plane, then $4\pi \cdot \text{area}(\Omega) \leq \text{perim}(\Omega)^2$, with equality if and only if Ω is a circle. This became

known as the *isoperimetric inequality* in the plane. This means that, in the plane, $0 \leq \text{PP}(\Omega) \leq 1$, where the Pólya-Popper score is equal to 1 only in the case of a circle. We can observe that the Pólya-Popper score is scale-invariant in the plane.

An isoperimetric inequality for the sphere exists, and we state it as the following lemma. For a more detailed treatment of isoperimetry in general, see Osserman (1979), and for a proof of this inequality for the sphere, see Rado (1935).

Lemma 5.2. *If Ω is a region on the sphere with area A and perimeter P , then $P^2 \geq 4\pi A - A^2$ with equality if and only if Ω is a spherical cap.*

A consequence of this is that among all regions on the sphere with a fixed area A , a spherical cap with area A has the shortest perimeter. However, the key difference between the Pólya-Popper score in the plane and on the sphere is that on the sphere, there is no scale invariance; two spherical caps of different sizes will have different scores.

Lemma 5.3. *Let S be the unit sphere, and let $\kappa(h)$ be a cap of height h . Then $\text{PP}(\kappa(h))$ is a monotonically increasing function of h .*

Proof. Let $r(h)$ be the radius of the circle bounding $\kappa(h)$. We compute:

$$\begin{aligned} 1 &= r(h)^2 + (1 - h)^2, \text{ by right triangle trigonometry} \\ &= r(h)^2 + 1 - 2h + h^2 \end{aligned}$$

Rearranging, we get that $r(h)^2 = 2h - h^2$, which we can plug in to the standard formula for perimeter:

$$\text{perim}_S(\kappa(h)) = 2\pi r(h) = 2\pi\sqrt{2h - h^2}$$

We can now use the Archimedian equal-area projection defined by $(x, y, z) \rightarrow \left(\frac{x}{\sqrt{x^2+y^2}}, \frac{y}{\sqrt{x^2+y^2}}, z\right)$ to compute $\text{area}_S(\kappa(h)) = 2\pi h$ and plug it in to get:

$$\text{PP}_S(\kappa(h)) = \frac{4\pi(2\pi h)}{4\pi^2(2h - h^2)} = \frac{2}{2 - h}$$

Which is a monotonically increasing function of h . □

Corollary 5.4. *On the sphere, Pólya-Popper scores of caps are monotonically increasing with area,*

Using this, we can show the main theorem of this section, that no map projection from a region on the sphere to the plane can preserve the ordering of Pólya-Popper scores for all regions.

Theorem 5.5. *If $\varphi : U \rightarrow V$ is a map projection from the sphere to the plane, then there exist two regions $A, B \subset U$ such that the Pólya-Popper score of B is greater than that of A in the sphere, but the Pólya-Popper score of $\varphi(A)$ is greater than that of $\varphi(B)$ in the plane.*




Figure 12. The construction of regions A and B in the proof of Theorem 5.5.

Proof. Let φ be a map projection, and let $\kappa \subset U$ be some cap. We will take our regions A and B to lie in κ . Set B to be a cap contained in κ . Let Σ be a circle in the plane such that $\Sigma \subsetneq \varphi(B)$ and let $A = \varphi^{-1}(\Sigma)$. See Figure 12 for an illustration.

We now use the isoperimetric inequality for the sphere and Corollary 5.4 to claim that A does not maximize the Poldby-Popper score in the sphere.

To see this, take \hat{A} to be a cap in the sphere with area equal to that of A . Note that since the area of \hat{A} is less than the area of the cap B , it follows that we can choose $\hat{A} \subset B$.

By the isoperimetric inequality of the sphere, $\text{PP}_S(\hat{A}) \geq \text{PP}_S(A)$. Since map projections preserve containment, $\Sigma \subsetneq \varphi(B)$ implies that $A \subsetneq B$, meaning that $\text{area}(\hat{A}) = \text{area}(A) \leq \text{area}(B)$. By Corollary 5.4, we know that $\text{PP}_S(\hat{A}) < \text{PP}_S(B)$, and combining this with the earlier inequality, we get

$$\text{PP}_S(A) \leq \text{PP}_S(\hat{A}) < \text{PP}_S(B)$$

Since $\Sigma = \varphi(A)$ maximizes the Poldby-Popper score in the plane, but A does not do so in the sphere, we have shown that φ does not preserve the maximal elements in the score ordering, and therefore it cannot preserve the ordering itself. \square

The reason why every map projection fails to preserve the ordering of Poldby-Popper scores is because the score itself is constructed from the *planar* notion of isoperimetry, and there is no reason to expect this formula to move nicely back and forth between the sphere and the plane. This proof crucially exploits a scale invariance present in the plane but not the sphere. If we consider any circle in the plane, its Poldby-Popper score is equal to one, but that is not true of every cap in the sphere.

6. Empirical Evaluation

In the previous sections we showed that no projection from the sphere to the plane can preserve various compactness orderings. These theorems suggest that in general maps that distort shape cannot preserve compactness orderings. In this section we investigate empirically the consequences of calculating compactness in different map projections, demonstrating the practical relevance of our investigation and providing evidence for possible generalizations.

6.1. Commonly-Used Projections

We briefly identify four commonly-used projections in the redistricting domain, which we will use in the next section to compare the empirical effects of the choice of map projection on the compactness orderings.

Plate Carrée The *plate carrée* projection, sometimes called an *equiarectangular* projection interprets latitude-longitude coordinates on the Earth as planar $x-y$ coordinates. This map projection does not accurately reflect most geographic figures and is therefore inappropriate for most applications. The U.S. Census Bureau distributes its shapefiles in this format, trusting the user to reproject the data into a format suited for the relevant application. However, because the data is distributed in this format, redistricting analysts and stakeholders (e.g. Chen (2017); Chikina, Frieze, and Pegden (2017); League of Women Voters of Pennsylvania, *et al.* (2018)) often do not perform this reprojection step, and this has led to the *plate carrée* projection becoming a *de facto* standard in this domain.

Mercator The *Web Mercator* projection is a cylindrical projection which is popular in Web mapping applications. As a result, this is the projection used in several online redistricting software tools available to the public, including DistrictBuilder (Public Mapping Project, 2018), Dave’s Redistricting App (Bradlee, Crowley, Matheiu, and Ramesey, 2019), and Districtr (Metric Geometry and Gerrymandering Group, 2019).

Lambert The *Lambert conic* projection is a conformal projection, which means that it preserves the angles of intersection of all segments. This is colloquially interpreted as ‘preserving shape at a small scale’. This projection is used in some portions of the U.S. State Plane Coordinate System, and is therefore used in an official capacity for some states.

Albers The *Albers* projection is an *equal area* conic projection, meaning it preserves the areas of all figures. The U.S. Atlas projection for the conterminous 48 states is an Albers projection and is the default in the Maptitude for Redistricting software, which is widely used by redistricting professionals, including legislators, consultants, and advocacy groups.

6.2. Results

While we have shown mathematically that the ordering of compactness scores is necessarily permuted by any map projection, we now consider the possibility of this effect occurring in reality. If it is the case that congressional districts all have scores far enough apart that the distortion introduced by the choice of projection is not sufficient to swap the ordering of this score, then the results above are merely mathematical curiosities. Precisely stated, we ask whether reprojection affects compactness score rankings of real districts in the context of commonly-used map projections. In previous both the scientific literature and the legal landscape, this question was either unaddressed, or asserted to be answered in the negative (c.f. Chen (2017); Chikina et al. (2017); League of Women Voters of Pennsylvania, *et al.* (2018)).

In this section, we demonstrate that for the commonly-used map projections listed above and the three compactness scores we examine in the previous sections, that this

permutation effect does occur in practice, using the congressional districts from the 115th Congress.² We extract the boundaries of the districts from a U.S. Census Bureau shapefile, using the highest resolution available, drawn at a scale of 1:500,000. We then compute the convex hull, Reock, and Polsby-Popper scores with respect to common map projections³ and examine the ordering of the districts with respect to both. While this is slightly different from the mathematical framework where we compare an abstract map projection to the computation on the surface of the sphere, computing the spherical values of these scores is not a simple task, even in modern geographic information systems (GIS) software.⁴ Rather, we can observe that the numerical values of all three scores on all districts are very similar with respect to the Lambert and Albers projections. These projections preserve local shape and area, respectively, and so we can imagine the ground-truth spherical value to also be concordant with these measures.

With four different map projections and three different compactness scores, we explore several instances in which the choice of map projection distorts the compactness score ranking of districts.

We first consider the 36 congressional districts in Texas. In Section 6.2, we plot the Polsby-Popper score ordering of these districts, comparing several pairs of projections. A perfect preservation of the order would result in these points all falling on the diagonal. However, what we see in practice is that most points do lie on the diagonal but several are not, indicating a disagreement between the ordering between the two projections, although the score orders totally agree in the Mercator and Albers projections. The distortion is clearly present, although fairly mild, with the only swaps occurring being between pairs nearby in the orderings.

We observe a similarly mild, though still present, perturbation in the convex hull score orderings, shown in Figure 14. In this setting, however, the score ordering is identical between the Lambert and Albers projections. A similar observation holds at the national level, considering all 433 districts in the coterminous United States. Thus, we empirically observe that the Polsby-Popper and convex hull score orders are fairly robust to the choice of projection, although not entirely.

However, some compactness score orderings are more sensitive than others. In Figure 15, we examine the Reock score ordering for the same pairs of projections. While the permutation between the Albers and the Lambert or Mercator projections is still relatively mild, although more complex than for the Polsby-Popper score, the distortion between *Plate Carrée* and these two projections is quite dramatic. The districts at the extreme ends of the ordering are relatively undisturbed, but the districts in the middle portion get shuffled around significantly. We observe that this effect is not a result of some idiosyncrasy of Texas' districts, since a similar effect persists when we consider all of the districts in the coterminous United States, shown in Figure 16.

The Polsby-Popper score is calculated by considering a portion of the map that contains only the district itself. Since some of the projections we consider are locally very similar, and the districts themselves are very small, this gives an explanation for

²Used for the 2016 congressional elections.

³The code to compute the various compactness scores is based on Lee Hachadoorian's *compactr* project (Hachadoorian, 2018).

⁴Provided that the region is contained in an open hemisphere, and that the earth is assumed to be a perfect sphere sitting in \mathbb{R}^3 , a simple algorithm to calculate a minimum bounding cap is as follows: find the minimum bounding 3-ball of the region and intersect that ball with the Earth. Efficient algorithms exist for computing the minimum bounding ball of a collection of points (Ritter, 1990). However, since computing the minimum bounding sphere of a region is not a typical problem in GIS, the data necessary to run the algorithm is not readily available.

Polsby-Popper Score Rank for Texas Districts




Figure 13. The Polsby-Popper score rank is slightly distorted between different projections.

the robustness of the compactness orderings for that score. On the other hand, the more extreme reprojection order reversal we see in Reock scores results from the fact that its computation depends on the potentially large smallest bounding circle around the district. This circle will always be larger than the region relevant for the calculation of the convex hull score, since the convex hull of a district is always contained in any bounding circle, and all the map projections we consider distort larger shapes more severely than smaller ones. Thus, we should expect the distortion from reprojections to affect the Reock score more significantly than the convex hull score.

While the results outlined here are by no means comprehensive, they are a representative sample of the prevalence of the order-reversal phenomenon in practice. In all cases, extreme shapes remain extreme under reprojection, but the rankings of the middle-ranked districts are distorted. While the actual numerical discrepancies between the scores computed under the different projections is small, that this permutation can even occur when using ‘nice’ projections like Albers and Lambert muddies the water in discussing compactness. If value of using mathematics to describe the shape of districts is to provide a small objective frame of reference in a setting where subjective political factors play a large role, then the inconsistency even in the ordering of the districts under the scores works counter to this purpose.

Furthermore, compactness scores are used directly and indirectly in the rapidly growing area of statistical analysis of gerrymandering using *ensembles* of districting

Convex Hull Score Rank for Texas Districts




Figure 14. The convex hull score rank is slightly distorted between different projections, though not between the Lambert and Albers projections.

plans (Chen and Rodden, 2015; Chikina et al., 2017; Herschlag, Kang, Luo, Graves, Bangia, Ravier, and Mattingly, 2018; Liu, Cho, and Wang, 2015), where many possible maps are generated by a computer and used to contextualize properties of a proposed plan. In that context, compactness scores are often aggregated into a score for a districting plan, which is then used to constrain the universe of plans the algorithm generates. For example, we might insist that the average Polsby-Popper score of the generated plans not be larger than our plan of interest or assert a lower threshold for the scores of the districts individually. One underexplored question is the extent to which the dependence on the map projection affects the resulting statistical analysis of these ensembles. We emphasize that it is possible that changes to the compactness scores of the “middle of the pack” districting plans can affect the distribution from which the algorithm draws samples; for instance, under the cut-off approach the universe of allowable plans itself could change significantly if the choice of map projection shuffles which kinds of shapes have scores above and below the threshold. We refer the reader to section 6.2.2 (“the extreme outlier hypothesis”) in (Najt, DeFord, and Solomon, 2019) for more details on these questions.

Reock Score Rank for Texas Districts

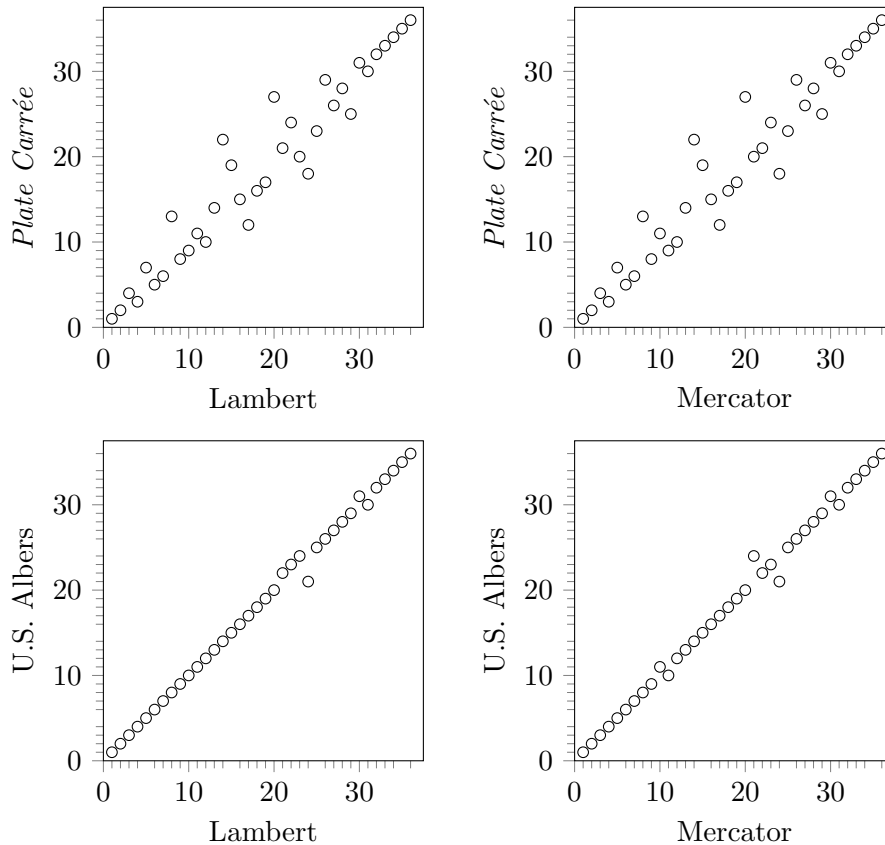


Figure 15. The Reock score ranking is distorted between several pairs of projections, with the *Plate Carrée* projection providing the most dramatic differences.

7. Discussion

We have identified a major *mathematical* weakness in the commonly discussed compactness scores in that no map projection can preserve the ordering over regions induced by these scores. This leads to several important considerations in the mathematical and popular examinations of the detection of gerrymandering.

From the mathematical perspective, rigorous definitions of compactness require more nuance than the simple score functions which assign a single real-number value to each district. *Multiscale* methods, such as those proposed by DeFord, Lavenant, Schutzman, and Solomon (2019), assign a vector of numbers or a function to a region, rather than a single number. The richer information contained in such constructions is less susceptible to perturbations of map projections. Alternatively, we can look to capturing the geometric information of a district without having to work with respect to a particular spherical or planar representation. So-called *discrete compactness* methods, such as those proposed in Duchin and Tenner (2018), extract a graph structure from the geography and are therefore unaffected by the choice of map projection, and our results suggest that this is an important advantage of these kinds of scores over traditional ones. Finally, recent work has used lab experiments to discern what qualities of a region humans use to determine whether they believe a region is compact or not

U.S. Districts Reock Score Rank: *Plate Carrée* vs. Lambert




Figure 16. The Reock score ranking is distorted by reprojection from the *Plate Carrée* projection to the Lambert projection when the entire US is considered.

(Kaufman, King, and Komisarichik, 2019). Incorporating more qualitative techniques is important, especially in this setting where the social impacts of a particular districting plan may be hard to quantify. To further complicate matters, as highlighted by Barnes and Solomon (2020), the *resolution* of the shapefile influence the computation of compactness scores, particularly the Polsby-Popper score where the detail of features like coastlines can have a massive impact on the measured perimeter of a region. For this reason, repeating the experiment in Section 6.2 for different choices of resolution results in quantitatively different (although qualitatively similar) results.

We proved our non-preservation results for three particular compactness scores which appear frequently in the context of electoral redistricting. There are countless other scores offered in legal codes and academic writing, such as definitions analogous to the Reock and convex hull scores which use different kinds of bounding regions, scores which measure the ratio of the area of the largest inscribed shape of some kind to the area of the district, and versions of these scores which replace the notion of ‘area’ with the population of that landmass. Many of these and others suffer from similar flaws as the three scores we examined in this work. It would be interesting to consider the most general version of this problem and enumerate a collection of properties such that any map projection permutes the score ordering of a pair of regions under a score with at least one of those properties.

While compactness scores are not used critically in a *legal* context, they appear frequently in the popular discourse about redistricting issues and frame the perception of the ‘fairness’ of a plan. An Internet search for a term like ‘most gerrymandered districts’ will invariably return results naming-and-shaming the districts with the most convoluted shapes rather than highlighting where more pleasant looking shapes resulted in unfair electoral outcomes.

Similarly, a sizable amount of work towards remedying such abuses focuses pri-

marily on the geometry rather than the politics of the problem. Popular press pieces (e.g. Ingraham (2014)) and academic research alike (e.g. Cohen-Addad, Klein, and Young (2017); Levin and Friedler (2019); Svec, Burden, and Dilley (2007)) describe algorithmic approaches to redistricting which use geometric methods to generate districts with appealing shapes. However, these approaches ignore all of the social and political information which are critical to the analysis of whether a districting plan treats some group of people unfairly in some way. A purely geometric approach to drawing districts implicitly supposes that the mathematics used to evaluate the geometric features of districts are unbiased and unmanipulable and therefore can provide true insight into the fairness of electoral districts. We proved here that the use of geographic compactness as a proxy for fairness is much less clear and rigid than some might expect.

This work opens several promising avenues for further investigation. We prove strong results for the most common compactness scores, but the question remains what the most general mathematical results in this domain might be, such as giving a set of necessary and sufficient conditions for a map projection to preserve the compactness ordering with respect to a particular score, and which kinds of surfaces do and do not admit such an order-preserving diffeomorphism or describing the permutation of scores as a function of the change in curvature between the two spaces of interest. Our work demonstrates a potential issue arising from the lack of standardization in the use of map projections in redistricting applications, for instance in the statistical analysis of gerrymandering, as discussed at the end of [Section 6](#). Gaining a better understanding of these effects is crucial as these statistical methods gain both academic and legal traction. From a cartographic standpoint, understanding other redistricting-related topics beyond compactness scores where the choice of map projection might have a significant effect on the outcome is important, particularly as access to mapmaking tools and data become more widely available to the general public.

Acknowledgments

We would like to thank the participants of the Voting Rights Data Institute for many helpful discussions. Special thanks to Eduardo Chavez Heredia and Austin Eide for their help developing mathematical ideas in the early stages of this work. We would like to thank Lee Hachadoorian for inspiring the original research problem and Moon Duchin, Jeanne N. Clelland, and Anthony Pizzimenti for exceptionally helpful feedback on drafts of this work. We thank Jeanne N. Clelland, Daryl DeFord, Yael Karshon, Marshall Mueller, Anja Randecker, Caleb Stanford, and Justin Solomon for offering their wisdom through helpful conversations throughout the process and are particularly grateful to Jowei Chen for pointing us toward the relevant portion of his deposition in *Rucho* and John O’Neill for facilitating that connection.

References

- Micah Altman. Traditional districting principles: Judicial myths vs. reality. *Social Science History*, 22(2):159–200, 1998.
- Richard Barnes and Justin Solomon. Gerrymandering and compactness: Implementation flexibility and abuse. *Political Analysis (to appear)*, 2020.
- David Bradley, Terry Crowley, Mike Matheiu, and Alec Ramesey. Dave’s redistricting app, 2019. URL <http://gardow.com/davebradlee/redistricting/>.

- Oliver Byrne. *The first six books of the Elements of Euclid: in which coloured diagrams and symbols are used instead of letters for the greater ease of learners*. William Pickering, 1847.
- Jowei Chen. Common Cause, et al., vs. Robert A. Rucho, League Of Women Voters of North Carolina vs. Robert A. Rucho, deposition of Jowei Chen. *In The United States District Court For The Middle District Of North Carolina*, page 167, 2017.
- Jowei Chen and Jonathan Rodden. Cutting through the thicket: Redistricting simulations and the detection of partisan gerrymanders. *Election Law Journal*, 14(4):331–345, 2015.
- Maria Chikina, Alan Frieze, and Wesley Pegden. Assessing significance in a markov chain without mixing. *Proceedings of the National Academy of Sciences*, 114(11):2860–2864, 2017.
- Vincent Cohen-Addad, Philip N. Klein, and Neal E. Young. Balanced power diagrams for redistricting. *CoRR*, abs/1710.03358, 2017. URL <http://arxiv.org/abs/1710.03358>.
- Ben Crowell. Is an equilateral triangle the same as an equiangular triangle, in any geometry? Mathematics Stack Exchange, 2016. URL <https://math.stackexchange.com/q/95080>. URL:<https://math.stackexchange.com/q/95080> (version: 2016-11-10).
- Daryl DeFord, Hugo Lavenant, Zachary Schutzman, and Justin Solomon. Total variation isoperimetric profiles. *SIAM Applied Algebra and Geometry*, 3, 2019.
- Moon Duchin and Bridget Eileen Tenner. Discrete geometry for electoral geography. *arXiv:1808.05860*, 2018.
- Yu. S. Frolov. Measuring the shape of geographical phenomena: A history of the issue. *Soviet Geography*, 16(10):676–687, 1975.
- Lee Hachadoorian. Compactr. <https://github.com/gerrymandr/compactr>, 2018.
- Gregory Herschlag, Han Sung Kang, Justin Luo, Christy Vaughn Graves, Sachet Bangia, Robert Ravier, and Jonathan C Mattingly. Quantifying gerrymandering in north carolina. *arXiv preprint arXiv:1801.03783*, 2018.
- Idaho Statute 72-1506(4).
- Christopher Ingraham. This computer programmer solved gerrymandering in his spare time, Jun 2014. URL https://www.washingtonpost.com/news/wonk/wp/2014/06/03/this-computer-programmer-solved-gerrymandering-in-his-spare-time/?noredirect=on&utm_term=.47fccb34f63d.
- Iowa Code §42.4(4).
- Aaron Kaufman, Gary King, and Mayya Komisarchik. How to measure legislative district compactness if you only know it when you see it. *American Journal of Political Science*, 2019.
- League of Women Voters of Pennsylvania, et al. Petitioners’ brief in support of proposed remedial plans, 2 2018.
- Harry Levin and Sorelle Friedler. Automated congressional redistricting. *ACM Journal of Experimental Algorithms*, 2019.
- Yan Y Liu, Wendy K Tam Cho, and Shaowen Wang. A scalable computational approach to political redistricting optimization. In *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*, page 6. ACM, 2015.
- Alan M. Maceachren. Compactness of geographic shape: Comparison and evaluation of measures. *Geografiska Annaler. Series B, Human Geography*, 67(1):53–67, 1985. ISSN 04353684, 14680467. URL <http://www.jstor.org/stable/490799>.
- Maine Statute §1206-A.
- Metric Geometry and Gerrymandering Group. Districtr, 2019. URL <https://districtr.org/>.
- Lorenzo Najt, Daryl DeFord, and Justin Solomon. Complexity and geometry of sampling connected graph partitions. *arXiv preprint arXiv:1908.08881*, 2019.
- Robert Osserman. Bonnesen-style isoperimetric inequalities. *The American Mathematical Monthly*, 86(1):1–29, 1979.
- Daniel D Polsby and Robert D Popper. The third criterion: Compactness as a procedural safeguard against partisan gerrymandering. *Yale Law & Policy Review*, 9(2):301–353, 1991.
- Public Mapping Project. Publicmapping/districtbuilder, Dec 2018. URL <https://github.com/PublicMapping/DistrictBuilder>.

- Tibor Rado. The isoperimetric inequality on the sphere. *American Journal of Mathematics*, 57(4):765–770, 1935. ISSN 00029327, 10806377. URL <http://www.jstor.org/stable/2371011>.
- Jack Ritter. An efficient bounding sphere. *Graphics gems*, 1:301–303, 1990.
- Lukas Svec, Sam Burden, and Aaron Dilley. Applying voronoi diagrams to the redistricting problem. *The UMAP Journal*, 28(3):313–329, 2007.
- H. P. Young. Measuring the compactness of legislative districts. *Legislative Studies Quarterly*, 13(1):105–115, 1988. URL <http://www.jstor.org/stable/439947>.