Geometric averages of partitioned datasets*

Tom Needham† and Thomas Weighill‡

Abstract. We introduce a method for jointly registering ensembles of partitioned datasets in a way which is both geometrically coherent and partition-aware. Once such a registration has been defined, one can group partition blocks across datasets in order to extract summary statistics, generalizing the commonly used order statistics for scalar-valued data. By modeling a partitioned dataset as an unordered k-tuple of points in a Wasserstein space, we are able to draw from techniques in optimal transport. More generally, our method is developed using the formalism of local Fréchet means in symmetric products of metric spaces. We establish basic theory in this general setting, including Alexandrov curvature bounds and a verifiable characterization of local means. Our method is demonstrated on ensembles of political redistricting plans to extract and visualize basic properties of the space of plans for a particular state, using North Carolina as our main example.

Key words. Wasserstein space, barycenter, symmetric product, clustering, redistricting

AMS subject classifications. 62R20, 51F99

arXiv:2107.03460v1 [cs.CG] 7 Jul 2021

1. Introduction.

Clustering of data is a fundamental task in unsupervised machine learning which searches for a partitioning of a given dataset which is optimal with respect to a given objective. This paper introduces statistical methods for the study of ensembles of datasets, such that each dataset comes with a predefined partition into k subsets, say, via some clustering algorithm. In particular, we consider the problem of computing the mean (or *barycenter*) of such an ensemble—this yields a mean of the underlying datasets overlaid with a mean partitioning. This framework has general applications to comparison of clusterings, for instance allowing one to quantify the stability of a clustering algorithm with respect to perturbations of an underlying dataset or to fuse distributed data which has been pre-clustered; see [53] for more examples. Our primary motivation is an application to *political redistricting*, where each dataset in the ensemble consists of a *districting plan*, or a partitioning of a given geographical region into districts. Once the mean of an ensemble of partitioned datasets has been computed, the partition blocks of each dataset can be assigned labels $\{1, 2, \dots, k\}$ by registering to the mean, allowing for partition-aware statistical analysis of the ensemble. This generalizes the classical idea of order statistics of an ensemble of scalar-valued datasets (see Example 13 for details). In our redistricting application, this registration allows for direct comparison of demographic and political statistics of districts across different plans.

A k-partitioned dataset can be modeled as an unordered k-tuple of distributions on the data space; in other words, an unordered k-tuple of points in the associated Wasserstein space. With a view toward a general theory, we develop our approach in the context of unordered k-tuples of points in an arbitrary metric space X. This space of k-tuples is referred to as the symmetric product $\mathsf{SP}^k X$ (also called the *sample space* in [26]), and is simply the quotient

* **Funding:** The second author was supported by NSF grant OIA-1937095.

†Florida State University (tneedham@fsu.edu).

‡The University of North Carolina at Greensboro (t_weighill@uncg.edu).

of the space of ordered k-tuples X^k by the order-permuting action of the symmetric group S_k. The main examples we are interested in are when X is $\mathbb{R}^n$, Wasserstein space $W_2(\mathbb{R}^n)$, a manifold M, or $\mathsf{SP}^m Y$ where Y is one of these spaces (that is, we consider $\mathsf{SP}^k\mathsf{SP}^m Y$). Our goal is then to study theoretical and computational aspects of the computation of means or barycenters of subsets $S \subset \mathsf{SP}^k X$. We show that under mild assumptions, $\mathsf{SP}^k X$ has curvature unbounded from above (Theorem 7), so that general existence and uniqueness results for barycenters do not directly apply. We can nonetheless characterize local barycenters of subsets $S \subseteq \mathsf{SP}^k X$ (Theorem 11) and we prove that, for many spaces of interest, the labeling of the points in S given by a best matching to a local barycenter is unique (Corollary 29). This allows us to implement an algorithm (Algorithm 3.1) for computing or approximating local barycenters in $\mathsf{SP}^k X$.

As was mentioned above, our target application in this paper is *political redistricting*: the process of dividing up a territory into pieces for the purpose of electing representatives. For example, in the United States every state is divided up into a number of Congressional districts roughly proportional to its population, with one member of the U.S. House of Representatives being elected from each of these districts. Applying our theory and Algorithm 3.1 produces a new method for visualizing and analyzing large ensembles of computer-generated redistricting plans. The analysis of large ensembles of redistricting plans has recently become prominent in research and litigation surrounding redistricting and gerrymandering. Using our method, we are able to label the districts in thousands of computer-generated redistricting plans in a coherent way and then examine the political and geographic features of the districts assigned to a given label. We demonstrate the value in this approach by comparing Congressional district-level election outcomes for two elections in North Carolina. We also analyze enacted and proposed plans within our framework in a way that complements recent work on quantifying gerrymandering in North Carolina and which answers a clear need for "local analysis" [35] of proposed maps.

The outline of the paper is as follows. In Section 2 we define the symmetric product space and establish some of its geometric properties. In Section 3 we develop the theory of p-barycenters in symmetric product spaces. This theory is used to formulate an algorithm for computing local p-barycenters. Section 4 gives some simple examples of this algorithm in practice. Section 5 contains the application to redistricting ensembles. We conclude this introductory section with a survey of related work.

Optimal transport. Optimal transport (OT) problems consist of finding the best way to transport a source distribution to a target distribution within a metric space. This problem was first posed by Monge in the eighteenth century [39] and was reformulated in the 1940s by Kantorovich [30] as a linear program, leading to significant progress and interest. Today, OT-based methods are applied in fields including statistics, machine learning, computer graphics and economics—see general references [44, 52] for details of theoretical and computational aspects of OT. Optimal transport connects with this paper in two ways. Firstly, Wasserstein space (the space of distributions endowed with an optimal transport distance between them) provides a key example of a space X for which we want to study $\mathsf{SP}^k X$. Secondly, $\mathsf{SP}^k X$ itself can be identified with a subset of Wasserstein space over X (Proposition 4), so our work can be considered as finding barycenters in a subset of Wasserstein space over a complicated underlying space. Barycenters in Wasserstein space have already been applied in areas such as

texture analysis [46], shape interpolation [48] and color transfer [23]. Wasserstein barycenters were introduced and studied for Euclidean spaces in [2], sparking a surge of interest in their theory, as well as methods for computing or approximating them [13, 17, 18, 45, 54]. Most relevant for the present paper is the theory developed for the manifold setting [32] and the Euclidean discrete case [3], as well as the exact and regularized algorithms in [18]. As mentioned above, the work in this paper can be formulated as finding barycenters in some subset of Wasserstein space, but we require more complicated underlying spaces (such as another Wasserstein space) and the use of local barycenters, which forces us to develop new theory specific to these spaces that is not currently found in the literature. A related thread is the theory and computation of barycenters of sets of persistence diagrams [15, 51], which treats barycenters in a metric space with curvature unbounded from above.

Redistricting ensembles. A key question in redistricting research is to determine whether a proposed or enacted redistricting plan is a gerrymander or not—that is, was some agenda other than the basic redistricting requirements of the state driving the line-drawing? A prominent approach in both research and litigation is to compare the plan to a large ensemble of alternatives generated by an algorithm that takes into account some or all of the redistricting criteria for the particular state and level of government [14, 29, 27, 21, 19, 4, 12, 11]. These ensembles are designed to represent the intractably large set of possible alternative plans, and are necessarily very geographically diverse. Typically, these ensembles are analyzed (and compared against the plan being evaluated) at the level of summary statistics – for example, the number of Republican seats won under historical vote data. Two recent papers have also analyzed spatial characteristics of redistricting ensembles: via graph optimal transport [1] and topological data analysis [40]. Our method allows us to combine a geometric perspective in line with these two papers with a classical summary statistics approach.

Geometry of symmetric product spaces. The main theoretical object of study in this paper is the k-fold symmetric product $\mathsf{SP}^k X$. The recent paper [26] studies the geometry of this space with the W_p metrics defined in the next section. In particular, they show that $\mathsf{SP}^k X$ is a stratified space and that a Fréchet mean (with respect to the metric on X) of a subset $S \subseteq X$ of size k is a projection of the corresponding point of $\mathsf{SP}^k X$ onto its lowest dimensional stratum. Our work, by contrast, studies barycenters of subsets of $\mathsf{SP}^k X$ relative to the W_p metric on $\mathsf{SP}^k X$. We should also mention the notion of unordered configuration space (see, e.g., [25]), which is the proper subspace of $\mathsf{SP}^k X$ consisting of points with k distinct entries, and which sometimes appears in applications to robotics.

2. Symmetric products and the W_p distance.

In this section, we formally define the symmetric product metric and establish some of its basic properties.

2.1. The W_p metric.

We use the notation $\langle k \rangle = \{1, \ldots, k\}$ and denote the i^{th} coordinate of an (ordered) tuple $\mathbf{v}$ by $\mathbf{v}_i$. For context, we recall the definition of Wasserstein distance.

Definition 1. *Let X be a metric space on which every finite Borel measure is a Radon measure. Given two Borel probability measures α and β, the* p-Wasserstein distance *is defined by*

$$W_p(\alpha, \beta) = \left(\inf_{\gamma \in \mathcal{U}(\alpha,\beta)} \int d(x,y)^p d\gamma(x,y) \right)^{1/p},$$

where $\mathcal{U}(\alpha, \beta)$ is the set of Borel measures on X^2 with marginals α and β.

If $\alpha = \sum_{i=1}^{m} \mathbf{a}_i \delta_{x_i}$ and $\beta = \sum_{i=1}^{n} \mathbf{b}_i \delta_{y_i}$ are discrete, then we can equivalently write

$$W_p(\alpha, \beta) = \left(\min_{P \in U(\mathbf{a}, \mathbf{b})} \sum_{i,j} d(x_i, y_j)^p P_{ij} \right)^{1/p},$$

where P ranges over the set $U(\mathbf{a}, \mathbf{b})$ of $n \times m$ matrices such that $P\mathbb{1} = \mathbf{a}$ and $P^T\mathbb{1} = \mathbf{b}$, with $\mathbb{1}$ denoting the column vector of the appropriate size with all entries equal to one. Throughout this paper, we will denote by $W_p(X)$ the *p-Wasserstein space over X* – that is, the set of Borel probability measures on X with finite p^{th} moment endowed with the p-Wasserstein metric.

We now introduce the main theoretical context for studying unordered data.

Definition 2. *Let X be a set. The symmetric group S_k of bijections $\phi : \langle k \rangle \to \langle k \rangle$ acts on the product set X^k by permuting entries of ordered k-tuples. We denote the action of a bijection $\pi \in S_k$ on $\mathbf{x} \in X^k$ by $\pi\mathbf{x}$; this action is given explicitly by the formula $(\pi\mathbf{x})_i = \mathbf{x}_{\pi(i)}$. The* k-fold symmetric product of X *is the quotient of X^k by the action of the symmetric group S_k, denoted*

$$\mathsf{SP}^k X := X^k / S_k.$$

We denote equivalence under this S^k-action by $\mathbf{x} \sim \mathbf{x}'$ and we denote the equivalence class of $\mathbf{x} \in X^k$ by $[\mathbf{x}]$.

If (X, d) is a metric space then $\mathsf{SP}^k X$ comes with a natural family of metrics.

Definition 3. *Let (X, d) be a metric space and let $p \geq 1$. Define the* p-Wasserstein distance *on $\mathsf{SP}^k X$ as follows. Let $\mathbf{x} = (x_1, \ldots, x_k)$ and $\mathbf{y} = (y_1, \ldots, y_k)$. Then*

$$W_p([\mathbf{x}], [\mathbf{y}])^p := \min_{\pi} \frac{1}{k} \sum_{i=1}^{k} d(x_i, y_{\pi(i)})^p \tag{2.1}$$

where the minimum ranges over all bijections $\pi : \langle k \rangle \to \langle k \rangle$. We call a bijection realizing this minimum an optimal matching from $\mathbf{x}$ to $\mathbf{y}$.

Let d_{ℓ^p} denote the *ℓ^p-metric* on X^k, given by

$$d_{\ell^p}(\mathbf{x}, \mathbf{y})^p = \sum_{i=1}^{k} d(\mathbf{x}_i, \mathbf{y}_i)^p.$$

Using the fact that S_k acts by isometries on d_{ℓ^p}, we have

$$W_p([\mathbf{x}], [\mathbf{y}]) = \min_{\mathbf{x}' \sim \mathbf{x},\ \mathbf{y}' \sim \mathbf{y}} \frac{1}{k} d_{\ell^p}(\mathbf{x}', \mathbf{y}') = \min_{\mathbf{y}' \sim \mathbf{y}} \frac{1}{k} d_{\ell^p}(\mathbf{x}, \mathbf{y}').$$

This relation makes it easy to check that W_p is indeed a metric, and that it induces the quotient topology on $\mathsf{SP}^k X$ when X^k is endowed with the ℓ^p metric.

We next give a precise relationship between the Wasserstein p-metric on $\mathsf{SP}^k X$ and the classical Wasserstein metric on $W_p(X)$. The result follows easily from the Birkhoff-von Neumann theorem; see [26, Lemma 4.7] for details.

Proposition 4. *The map*

$$\iota : \mathsf{SP}^k X \to W_p(X), \quad [\mathbf{x}] \mapsto \frac{1}{k}\sum_{i=1}^{k} \delta_{\mathbf{x}_i} \tag{2.2}$$

is an isometric embedding of $(\mathsf{SP}^k X, W_p)$ *into Wasserstein space* $W_p(X)$.

2.2. Geodesics and curvature in $\mathsf{SP}^k X$. In this section we derive some results about the basic geometry of the metric space $(\mathsf{SP}^k X, W_p)$. Along the way, we recall basic notions of metric geometry, following [7, 8].

Let (X, d) be a metric space. A *constant speed geodesic* in X is a map $[a, b] \to X$ of some interval which is an isometric embedding up to a multiplicative constant. The space X is a *geodesic space* if any two points $x, y \in X$ can be joined by a constant-speed geodesic path $\gamma : [0, 1] \to X$ of length $d(x, y)$—that is, a path satisfying

$$d(\gamma(s), \gamma(t)) = d(x, y) \cdot |s - t|, \qquad \gamma(0) = x,\ \gamma(1) = y.$$

Such a path is called a *minimal geodesic joining* x *to* y. A geodesic $\gamma : [a, b] \to X$ *branches at* $t_0 \in (a, b)$ if there exists another geodesic $\bar{\gamma} : [a, b] \to X$ such that $\gamma|_{[a,t_0]} = \bar{\gamma}|_{[a,t_0]}$ but $\gamma \neq \bar{\gamma}$ on some interval $(t_0, t_0 + \varepsilon)$. If X has no branching geodesics, we say that X is *non-branching*.

Proposition 5. *Let* (X, d) *be a geodesic metric space and endow* $\mathsf{SP}^k X$ *with the* W_p *metric.*

1. *The map*

$$\rho : X \to \mathsf{SP}^k X, \qquad x \mapsto (x, \ldots, x) \tag{2.3}$$

   is an isometric embedding of (X, d) *into* $(\mathsf{SP}^k X, W_p)$.
2. *Let* π *be an optimal matching of* $\mathbf{x}, \mathbf{y} \in \mathsf{SP}^k X$ *and let* $\gamma_i : [0, 1] \to X$ *be a minimal geodesic between* x_i *and* $y_{\pi(i)}$ *in* X. *Define*

$$\gamma : [0, 1] \to \mathsf{SP}^k X, \qquad t \mapsto [\gamma_1(t), \ldots, \gamma_k(t)].$$

   Then γ *is a minimal geodesic in* $(\mathsf{SP}^k X, W_p)$ *joining* $[\mathbf{x}]$ *to* $[\mathbf{y}]$ *for any* $p \in [1, \infty]$. *In particular,* $\mathsf{SP}^k X$ *is a geodesic space.*

Remark 6. *The map* $X \to W_2(X)$ *taking* x *to the Dirac measure* δ_x *is well known to be an isometric embedding [50, Proposition 2.10]. This map factors as the composition* $\iota \circ \rho$, *where* $\rho : X \to \mathsf{SP}^k X$ *is the isometric embedding defined above in* (2.3) *and* $\iota : \mathsf{SP}^k X \to W_2(X)$ *is the isometric embedding defined in* (2.2). *The image of* ρ *is the lowest dimensional stratum (referred to as the 1-skeleton) in the stratified space structure of* $\mathsf{SP}^k X$ *described in [26].*

Proof. Point 1 follows by a simple computation: for $x, x' \in X$, we have

$$W_p(\rho(x), \rho(x'))^p = \min_{\pi} \frac{1}{k}\sum_{i=1}^{k} d\left(\rho(x)_i, \rho(x')_{\pi(i)}\right)^p = \frac{1}{k}\sum_{i=1}^{k} d(x, x')^p = d(x, x')^p.$$

To prove point 2, let $D_i := d(x_i, y_{\pi(i)})$ and $D := W_p([\mathbf{x}], [\mathbf{y}])$. By standard arguments, it suffices to show that for all $0 \leq s \leq t \leq 1$,

$$W_p(\gamma(s), \gamma(t)) \leq (t - s)D$$

(see, e.g., [16, Lemma 1.3]). This is straightforward: for $p \in [1, \infty)$ we have

$$W_p(\gamma(s), \gamma(t))^p = \min_{\tilde{\pi}} \sum_{i=1}^{k} d(\gamma_i(s), \gamma_{\tilde{\pi}(i)}(t))^p$$
$$\leq \sum_{i=1}^{k} d(\gamma_i(s), \gamma_i(t))^p = \sum_{i=1}^{k} (t-s)^p D_i^p = (t-s)^p D^p,$$

and the $p = \infty$ case is similar. ∎

The rest of this section deals with curvature of metric spaces, in the sense of Alexandrov. For $\kappa \in \mathbb{R}$, let M_κ denote the 2-dimensional space form of constant curvature κ, with metric denoted d_κ and metric diameter denoted D_κ. For three points x, y, z in a metric space (X, d), one can always find points $\bar{x}, \bar{y}, \bar{z} \in M_\kappa$ with $d(x, y) = d_\kappa(\bar{x}, \bar{y})$, $d(y, z) = d_\kappa(\bar{y}, \bar{z})$ and $d(z, x) = d_\kappa(\bar{z}, \bar{x})$. We say that $\bar{x}, \bar{y}, \bar{z}$ give a *comparison triangle* for x, y, z. A geodesic metric space (X, d) is said to have *curvature bounded below* (respectively, *above*) by κ if for any three points x, y, z such that $d(x, y) + d(y, z) + d(z, x) \leq 2D_\kappa$ and for any minimal geodesic $\gamma : [0, 1] \to X$ joining x to y, it holds that $d(z, \gamma(t)) \geq d_\kappa(\bar{z}, \bar{\gamma}(t))$ (respectively, $d(z, \gamma(t)) \leq d_\kappa(\bar{z}, \bar{\gamma}(t))$) for all $t \in [0, 1]$, where $\bar{x}, \bar{y}, \bar{z}$ give a comparison triangle for x, y, z in M_κ and $\bar{\gamma} : [0, 1] \to M_\kappa$ is a minimizing geodesic joining $\bar{x}$ to $\bar{y}$.

A geodesic metric space with curvature bounded below by zero is called an *Alexandrov space with nonnegative curvature*. In this case, the relevant inequality can be expressed as

$$d(z, \gamma(t))^2 \geq (1-t)d(z, x)^2 + td(z, y)^2 - t(1-t)d(x, y)^2 \tag{2.4}$$

for all x, y, z, where γ is a minimal geodesic joining x to y [41, Section 2.1]. A geodesic metric space with curvature bounded above by κ is called a CAT(κ) *space*. We deal below with Alexandrov spaces with nonnegative curvature. This category includes most spaces of interest from our data analysis perspective, including Euclidean spaces, complete Riemannian manifolds whose sectional curvature is not everywhere negative and Wasserstein spaces $W_2(X)$, where X is itself an Alexandrov space of nonnegative curvature [41, Section 2.1].

We now state the main result of this section, which describes Alexandrov curvature bounds for symmetric products. We restrict our attention to the W_2 metric in this setting—this is sensible, since even the standard ℓ^p space $(\mathbb{R}^n, d_{\ell^p})$ is an Alexandrov space with nonnegative curvature if and only if $p = 2$ (in an ℓ^p-space, one can show that the Alexandrov inequality (2.4) implies the Parallelogram Law). By similar reasoning, $(\mathbb{R}^k, d_{\ell^p})$ is CAT(0) if $p = 2$ and is otherwise not CAT(κ) for any κ—see also [7, Proposition II.1.14].

Theorem 7. *Let (X, d) be a geodesic metric space and endow $\mathsf{SP}^k X$ with the W_2 metric.*

1. *X is an Alexandrov space with nonnegative curvature if and only if $\mathsf{SP}^k X$ is.*
2. *If X is not a one point space, a 1-manifold or a 1-manifold with boundary then $\mathsf{SP}^k X$ is not* CAT(κ) *for any κ.*

Theoretical results and computational tools regarding barycenters in CAT(κ) spaces exist in the literature [55]; for example, barycenters in CAT(0) spaces are unique [49]. This theorem indicates that these methods cannot be generally applied to the symmetric product spaces of interest, motivating the new theory developed in Section 3. Results which are similar to point

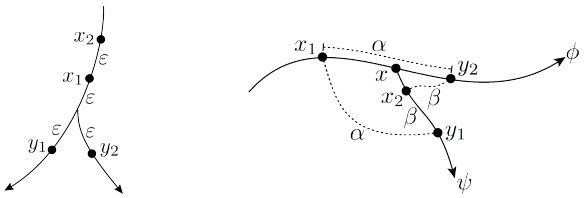

Figure 1: Schematic figures for the proof of Theorem 7.

1 are established for Wasserstein spaces $W_2(X)$ in [50, Proposition 2.10] and [33, Theorem A.8]. It is well known that Wasserstein spaces do not inherit upper curvature bounds—see, e.g. [5, Remark 2.10], which shows that if X is CAT(0) then $W_2(X)$ is not, unless X is a isometric to an interval. A result in a similar spirit to point 2 is proved for the space of persistence diagrams (in the context of topological data analysis) in [51, Proposition 2.4].

Proof of Theorem 7. If $(\mathsf{SP}^k X, W_2)$ has nonnegative curvature then so does (X, d), since X embeds isometrically in $\mathsf{SP}^k X$ via the map ρ defined in (2.3). Now suppose that X has nonnegative curvature. Then so does (X, d_{ℓ^2}) [9, Proposition 4.1]. The quotient $\mathsf{SP}^k X = X^k/S_k$ of the nonnegatively curved space (X^k, d_{ℓ^2}) by the isometric action of the finite group S_k is therefore nonnegatively curved by [9, Corollary 4.6]. This completes the proof of 1.

To prove point 2, it suffices to show that $\mathsf{SP}^k X$ contains arbitrarily close points joined by distinct geodesics [7, Proposition II.1.4]. Moreover, it suffices to prove the claim for $k = 2$.

First suppose that X contains a branching geodesic. Consider the point configurations $\mathbf{x} = (x_1, x_2)$ and $\mathbf{y} = (y_1, y_2)$ lying on the branching geodesic near the branch point, as shown in the lefthand side of Figure 1, where $\varepsilon > 0$ is arbitrarily small. Then both of the possible matchings of $\mathbf{x}$ and $\mathbf{y}$ are optimal with respect to W_2. Let γ_i (respectively, γ_i') be the geodesic from x_i to y_i (respectively, x_i to $y_{\pi(i)}$, where π is the non-identity element of S_2) whose image is contained in the branching geodesic. Then $\gamma := [\gamma_1, \gamma_2]$ and $\gamma' := [\gamma_1', \gamma_2']$ are both minimizing geodesics in $\mathsf{SP}^2 X$, by Proposition 5, point 2. These geodesics are distinct—for example, $\gamma(2\varepsilon/3) \neq \gamma'(2\varepsilon/3)$. Since $\varepsilon > 0$ was arbitrary, $\mathsf{SP}^2 X$ is not CAT(κ) for any κ.

Next, suppose that (X, d) is non-branching. We may also assume that X is CAT(κ) for some κ—otherwise the claim follows immediately, since X embeds isometrically into $\mathsf{SP}^k X$ by Proposition 5. We construct a configuration of points in X which yields nonunique geodesics between arbitrarily close points in $\mathsf{SP}^2 X$ as follows. We claim that there exists a point $x \in X$ which lies in the relative interior of a geodesic such that the image of the geodesic is not surjective onto any neighborhood of x. Indeed, let $a \in X$ be a point with no 1-manifold or 1-manifold with boundary chart. Choose a minimizing geodesic α from a to a point b in a small neighborhood of a. This geodesic is not onto any neighborhood of a, so we may choose a geodesic β from the midpoint m of α to another point c not in the image of α. If there exists a neighborhood U of m such that $\mathrm{Im}(\alpha) \cap \mathrm{Im}(\beta) \cap U = \{m\}$, then we set $x = m$. Otherwise, the nonbranching and CAT(κ) conditions imply that a lies in the relative interior of β and we set $x = a$. Now we fix an $x \in X$ in the relative interior of a geodesic path γ which is not surjective onto any neighborhood of x. Without loss of generality, we parameterize $\gamma : [-1, 1] \to X$ with $\gamma(0) = x$. Choose another geodesic $\psi : [0, 1] \to X$ with $\psi(0) = x$ and $\gamma(1)$ not in the image

of γ. We have that $\mathrm{Im}(\gamma) \cap \mathrm{Im}(\psi) \cap U = \{x\}$ for some small neighborhood U of x, by the assumption that X is non-branching and CAT(κ). Set $x_1 := \gamma(t_1)$ for some arbitrarily small $t_1 < 0$, $y_1 := \psi(s_1)$ for some arbitrarily small $s_1 > 0$ and $\alpha := d(x_1, y_1) > 0$. Next, let $s_2 := \alpha - c \cdot |t_1|$, where c is the constant speed of γ, and $y_2 := \gamma(s_2)$, so that $d(x_1, y_2) = \alpha$. Consider the path from y_2 to y_1, through x, following the images of the geodesic paths γ and ψ. By continuity, there is a point x_2 on this path such that $d(x_2, y_1) = d(x_2, y_2) =: \beta$. Note that $x_2 \neq x_1$. A schematic of this configuration is shown in the righthand side of Figure 1.

The configuration constructed above has the property that either matching of $\mathbf{x} := (x_1, x_2)$ and $\mathbf{y} := (y_1, y_2)$ is optimal with respect to W_2. Let γ_i (respectively, γ_i') be the geodesic from x_i to y_i (respectively, to $y_{\pi(i)}$), where π is the non-identity element of S_2), so that $[\gamma_1, \gamma_2]$ and $[\gamma_1', \gamma_2']$ are both minimizing geodesics in $\mathsf{SP}^2 X$ from $[\mathbf{x}]$ to $[\mathbf{y}]$. We claim that these minimizing geodesics are distinct. For sufficiently small $\eta > 0$, we have $\gamma_1(t) \neq \gamma_1'(t)$ for all $t \in (0, \eta)$. This is because X is non-branching, so γ_1 and γ_2 can't coincide on any interval (as their endpoints are distinct), while the CAT(κ) condition on X implies that the distinct geodesics γ_1 and γ_2 can't intersect arbitrarily close to x_1. On the other hand $\gamma_1(0) = x_1 \neq x_2 = \gamma_2'(0)$, so $\gamma_1(t) \neq \gamma_2'(t)$ for all t in some small interval $[0, \eta)$, by continuity. In particular, there exists t such that $\gamma_1(t) \notin \{\gamma_1'(t), \gamma_2'(t)\}$, which implies $[\gamma_1, \gamma_2] \neq [\gamma_1', \gamma_2']$. Since this construction can be done in an arbitrarily small neighborhood of x, this completes the proof. ∎

Example 8. *The symmetric product space $\mathsf{SP}^k\mathbb{R} = \mathbb{R}^k/S_k$, endowed with the W_2 metric, is isometric to the space $Y := \{\mathbf{x} \in \mathbb{R}^k \mid \mathbf{x}_1 \leq \mathbf{x}_2 \leq \cdots \leq \mathbf{x}_k\}$, endowed with the ℓ^2 metric, via the map taking $[\mathbf{x}]$ to its sorted representation—this follows by standard results on one-dimensional optimal transport [47, Section 3.1]. Thus $\mathsf{SP}^k\mathbb{R}$ is* CAT(0)*, and a similar argument works for $\mathsf{SP}^k X$ when X is isometric to an interval.*

3. Barycenters in Symmetric Products.

This section introduces our main object of study: barycenters of subsets of $\mathsf{SP}^k X$ with respect to the W_p metric.

3.1. Local p-barycenters in $\mathsf{SP}^k X$.

In this section we characterize local p-barycenters in $\mathsf{SP}^k X$ in terms of local barycenters in X and optimal matchings.

Definition 9. *Let X be a metric space and $S \subseteq X$ be a finite subset. Then a* p-barycenter *of S is a minimizer of the* p-Fréchet functional *associated to S given by:*

$$f(x) = f_{S,p}(x) := \sum_{s \in S} d(s, x)^p \tag{3.1}$$

A local p-barycenter *is a local minimum of the functional f.*

We will use the following technical lemma.

Lemma 10. *Let $\mathbf{x}, \mathbf{y} \in X^k$. Then there is an $\varepsilon > 0$ such that for any $\mathbf{x}'$ with $d_{\ell^p}(\mathbf{x}', \mathbf{x}) < \varepsilon$, every optimal matching of $\mathbf{y}$ to $\mathbf{x}'$ is an optimal matching of $\mathbf{y}$ to $\mathbf{x}$.*

Proof. If all matchings are optimal from $\mathbf{y}$ to $\mathbf{x}$, then there is nothing to prove. Otherwise, let ψ be a matching from $\mathbf{y}$ to $\mathbf{x}$ which minimizes $\sum d(y_i, x_{\psi(i)})^p$ among all non-optimal matchings from $\mathbf{y}$ to $\mathbf{x}$, and set

$$\delta := \left(\sum d(y_i, x_{\psi(i)})^p\right)^{1/p} - k^{1/p} \cdot W_p([\mathbf{y}], [\mathbf{x}]).$$

By definition, $\delta > 0$. Suppose that for some $\mathbf{x}'$, there is an optimal matching ϕ' from $\mathbf{y}$ to $\mathbf{x}'$ which is not optimal from $\mathbf{y}$ to $\mathbf{x}$. Pick any optimal matching ϕ from $\mathbf{y}$ to $\mathbf{x}$. We have that

$$R := d_{\ell^p}(\mathbf{y}, \phi\mathbf{x}) \geq d_{\ell^p}(\mathbf{y}, \phi\mathbf{x}') - d_{\ell^p}(\phi\mathbf{x}', \phi\mathbf{x}) = d_{\ell^p}(\mathbf{y}, \phi\mathbf{x}') - d_{\ell^p}(\mathbf{x}', \mathbf{x}).$$

Since ϕ' is an optimal matching from $\mathbf{y}$ to $\mathbf{x}'$ and ϕ is not, we have

$$\begin{aligned} d_{\ell^p}(\mathbf{y}, \phi\mathbf{x}') - d_{\ell^p}(\mathbf{x}', \mathbf{x}) &> d_{\ell^p}(\mathbf{y}, \phi'\mathbf{x}') - d_{\ell^p}(\mathbf{x}', \mathbf{x}) \\ &\geq d_{\ell^p}(\mathbf{y}, \phi'\mathbf{x}) - d_{\ell^p}(\phi'\mathbf{x}, \phi'\mathbf{x}') - d_{\ell^p}(\mathbf{x}', \mathbf{x}) \\ &= d_{\ell^p}(\mathbf{y}, \phi'\mathbf{x}) - 2d_{\ell^p}(\mathbf{x}', \mathbf{x}) \geq R + \delta - 2d_{\ell^p}(\mathbf{x}', \mathbf{x}). \end{aligned}$$

Putting everything together gives $R > R + \delta - 2d_{\ell^p}(\mathbf{x}', \mathbf{x})$, which implies that $d_{\ell^p}(\mathbf{x}', \mathbf{x}) > \frac{\delta}{2}$. Therefore it is only possible to find an optimal matching of $\mathbf{y}$ to $\mathbf{x}'$ which is not an optimal matching of $\mathbf{y}$ to $\mathbf{x}$ if $d_{\ell^p}(\mathbf{x}', \mathbf{x}) > \delta/2$, so the claim follows for any $\varepsilon < \delta/2$. ∎

Definition 11. *Let $\mathbf{x} \in X^k$ and let $S \subseteq X^k$ be a finite subset. We say that $\mathbf{x}$ is* stationary *with respect to S if the following holds: for every choice of $\Phi = (\phi_{\mathbf{s}})_{\mathbf{s}\in S}$, where $\phi_{\mathbf{s}}$ is an optimal matching from $\mathbf{x}$ to $\mathbf{s}$, and for every $1 \leq i \leq k$, $\mathbf{x}_i$ is a local p-barycenter of the set*

$$S_i^{\Phi} := \{\mathbf{s}_{\phi_{\mathbf{s}}(i)} \mid \mathbf{s} \in S\}.$$

Theorem 12. *Let $\mathbf{x} \in X^k$ and let $S \subseteq X^k$ be a finite subset. Then $[\mathbf{x}]$ is a local p-barycenter in $\mathsf{SP}^k X$ of $[S] = \{[\mathbf{s}] \mid \mathbf{s} \in S\}$ if and only if $\mathbf{x}$ is* stationary *with respect to S.*

Proof. First suppose that $[\mathbf{x}]$ is a local p-barycenter of $[S]$ but $\mathbf{x}$ is not stationary. That means there exist optimal matchings $\Phi = (\phi_{\mathbf{s}})_{\mathbf{s}\in S}$ and an i such that $\mathbf{x}_i$ is not a local p-barycenter of S_i^{Φ}. Thus for every $\varepsilon > 0$, there is an $\mathbf{x}_i'$ within ε of $\mathbf{x}_i$ such that $f(x_i') < f(x_i)$ where f is the p-Fréchet functional associated to S_i^{ϕ}. Let $\mathbf{x}'$ be $\mathbf{x}$ with the i^{th} coordinate replaced by $\mathbf{x}_i'$. We have

$$f(\mathbf{x}') = \sum_{\mathbf{s}\in S} d([\mathbf{x}'], [\mathbf{s}])^p = \sum_i \sum_{\mathbf{s}\in S} d(\mathbf{x}_i', \mathbf{s}_{\phi_{\mathbf{s}}'(i)})^p,$$

where $\phi_{\mathbf{s}}'$ is an optimal matching from $\mathbf{x}'$ to $\mathbf{s}$. Since these matchings are optimal,

$$\begin{aligned} f(\mathbf{x}') = \sum_{\mathbf{s}\in S} d([\mathbf{x}'], [\mathbf{s}])^p &\leq \sum_i \sum_{\mathbf{s}\in S} d(\mathbf{x}_i', \mathbf{s}_{\phi_{\mathbf{s}}(i)})^p \\ &< \sum_i \sum_{\mathbf{s}\in S} d(\mathbf{x}_i, \mathbf{s}_{\phi_{\mathbf{s}}(i)})^p = \sum_{\mathbf{s}\in S} d([\mathbf{x}], [\mathbf{s}])^p = f(\mathbf{x}). \end{aligned}$$

But since $\mathbf{x}_i'$ is within ε of $\mathbf{x}_i$, we have that $d([\mathbf{x}], [\mathbf{x}']) \leq \varepsilon$. Thus $[\mathbf{x}]$ cannot be a local p-barycenter of $[S]$.

To prove the converse, suppose that $\mathbf{x}$ is stationary. For any choice $\Phi = (\phi_s)_{s\in S}$ of optimal matchings as in Definition 11, $\mathbf{x}$ is a local minimum of the function

$$f_{\Phi}(\mathbf{x}) = \sum_i \sum_{\mathbf{s}\in S} d(\mathbf{x}_i, \mathbf{s}_{\phi_{\mathbf{s}}(i)})^p$$

Algorithm 3.1 Local p-barycenter in $\mathsf{SP}^k X$

```
(stationary, x̄, n) ← (false, x_0, |S|)
while stationary = false do
  for j = 1 to n do
    R_j ← {s' ~ s^j | W_p([s^j], [x̄])^p = (1/k) d_{ℓ^p}(s', x̄)^p}
  end for
  stationary ← true
  for (t^1, ..., t^n) ∈ ∏_j R_j do
    for i = 1 to k do
      x̂_i ← LB({t_i^j | j ∈ [n]}, x̄_i)
    end for
    if x̄ ≠ x̂ then
      (stationary, x̄) ← (false, x̂)
    end if
  end for
end while
return [x̄]
```

because each $\mathbf{x}_i$ is a local p-barycenter of the set $S_i^\Phi = \{\mathbf{s}_{\phi_\mathbf{s}(i)} \mid \mathbf{s} \in S\}$. Since there are only finitely many choices for Φ, we can fix $\varepsilon > 0$ such that $\mathbf{x}$ is a mininum of f_Φ on $B(\mathbf{x}, \varepsilon)$ for any choice of Φ. By Lemma 10 and the fact that there are finitely many elements in S, there is a $\varepsilon > \varepsilon' > 0$ such that if $d(\mathbf{x}, \mathbf{x}') < \varepsilon'$ then Φ can be chosen such that it is also a set of optimal matchings from $\mathbf{x}'$ to each element of S.

It follows that if $d([\mathbf{x}], [\mathbf{x}']) < \varepsilon'$, then there is a $\mathbf{x}'' \sim \mathbf{x}'$ with $d_{\ell^p}(\mathbf{x}, \mathbf{x}'') < \varepsilon'$ such that Φ can be chosen to be a set of optimal matchings from $\mathbf{x}''$ to elements of S. We have

$$\sum_{s\in S} W_p([\mathbf{x}'], [\mathbf{s}])^p = \sum_{s\in S} W_p([\mathbf{x}''], [\mathbf{s}])^p = \frac{1}{k}\sum_i \sum_{\mathbf{s}\in S} d(\mathbf{x}_i'', \mathbf{s}_{\phi_\mathbf{s}(i)})^p$$
$$\geq \frac{1}{k}\sum_i \sum_{\mathbf{s}\in S} d(\mathbf{x}_i, \mathbf{s}_{\phi_\mathbf{s}(i)})^p = \sum_{s\in S} W_p([\mathbf{x}], [\mathbf{s}])^p$$

We conclude that $[\mathbf{x}]$ is a local p-barycenter for $[S]$. ∎

Example 13. *Consider a subset $S \subseteq \mathsf{SP}^k\mathbb{R}$. Then a 2-barycenter of S consists of points $(a_1, \ldots, a_k)$ where for each i, a_i is the average of all the i^{th} largest entries of the points in S. Indeed, an optimal W_2 matching between two vectors in $\mathbb{R}^k$ is given by pairing up the i^{th} largest values for each i, and Theorem 12 dictates that each point in the barycenter is a 2-barycenter (in this case, mean) of the points matched to it. We can therefore view our barycenter method for labeling data as a generalization of order statistics to more complicated spaces.*

3.2. Computing local barycenters. We will now describe an iterative algorithm aimed at finding local p-barycenters in $\mathsf{SP}^k X$, inspired by the algorithm for persistence diagrams in

[51]. We state the algorithm in a very general context where convergence to a solution in finite time is not guaranteed, before giving some cases where a finite number of iterations produces a local p-barycenter.

Definition 14. *Let X be a metric space and $F(X)$ the set of all finite subsets of X. By a* p-descent operator *we mean a function $\mathcal{LB} : F(X) \times X \to X$ such that for any $S \in F(X)$ and $x \in X$ one of the following is true:*

(a) either $\sum_{s\in S} d(\mathcal{LB}(S,x),s)^p < \sum_{s\in S} d(x,s)^p$, or

(b) $\mathcal{LB}(S,x) = x$ and x is a local p-barycenter of S,

In particular, if X admits unique p-barycenters and $\mathcal{LB}(S,x)$ is defined to be the p-barycenter of S for any x, then $\mathcal{LB}$ is a p-descent operator. More generally, a p-descent operator might be one or more steps of a gradient descent method. Our method for finding local p-barycenters in $\mathsf{SP}^k X$ given a subset S, a p-descent operator $\mathcal{LB}$ on X and an initial seed $\mathbf{x}_0 \in X^k$ is described in Algorithm 3.1.

Remark 15. *Algorithm 2 in [18] describes a method for finding approximate 2-barycenters in $P_k(\mathbb{R}^d, \Theta)$, the subset of $W_2(\mathbb{R}^d)$ consisting of those measures with support of size at most k and weights in some chosen set of k-tuples Θ. When Θ contains only the k-tuple $(1/k, \dots, 1/k)$, Algorithm 2 in [18] (with the parameter choice $\theta = 1$) is equivalent to Algorithm 1 in this section with $X = \mathbb{R}^d$ and $\mathcal{LB}(S,x) = \frac{1}{|S|}\sum_{\mathbf{s}\in S} \mathbf{s}$.*

If the while loop terminates in Algorithm 3.1, then we have the exact conditions for $\bar{\mathbf{x}}$ to be stationary with respect to S (Definition 11), so by Theorem 12, we have found a local p-barycenter. As mentioned above, however, whether or not the algorithm terminates might depend on the definition of $\mathcal{LB}$. In all cases, we can observe the quantity

$$D = \sum_j d_{\ell^p}(\bar{\mathbf{x}}, \mathbf{t}^j)^p = \sum_{i,j} d_X(\bar{\mathbf{x}}_i, \mathbf{t}_i^j)^p,$$

where $\mathbf{t}^j \in R_j$ for each j (note that D does not depend on the choice of $\mathbf{t}^j$), and note that D is strictly reduced during all but the last iteration of the while loop. Indeed, redefining R_j in Line 6 cannot increase D, while the definition of the local p-barycenter operator $\mathcal{LB}$ guarantees that when Line 14 is executed, one of the sums $\sum_i d_X(\bar{\mathbf{x}}_i, \mathbf{t}_i^j)^p$ is reduced.

Proposition 16. *Suppose Algorithm 3.1 always terminates in a finite number of steps for some X and $\mathcal{LB}$. For $S \subseteq X$ and $x \in X$, define $\mathcal{A}(S,x)$ to be the output of Algorithm 3.1 applied to the subset S with initial value $\mathbf{x}_0 = x$. Then $\mathcal{A}$ defines a p-descent operator.*

Proof. Note that $D = \sum_{\mathbf{s}^i\in S} W_p(\bar{\mathbf{x}}_0, \mathbf{s}^i)^p$ at Line 8 on the first iteration of the `while` loop, after which $\sum_{\mathbf{s}^i\in S} W_p(\bar{\mathbf{x}}_0, \mathbf{s}^i)^p \le D$. Since D is strictly reduced during all but the last iteration of the `while` loop, $\mathcal{A}$ satisfies Definition 14(a). The only way the `while` loop terminates after a single iteration is if the output $[\bar{\mathbf{x}}]$ equals the initial value $\mathbf{x}_0$, which shows condition (b). ∎

Since D is strictly reduced at each iteration, if we can show that there are only finitely many possible values D can take, then we are done. If $\mathcal{LB}(S,\cdot)$ admits finitely many values for each S, then D can take only finitely many values. Indeed, in this case the sets R_j constrain $\bar{x}$, and hence D, to finitely many values. Moreover, there are only finitely many possibilities for the sets R_j since each R_j is a subset of $[\mathbf{s}^j]$. Thus we get the following cases where Algorithm 3.1 is guaranteed to terminate.

Proposition 17. *Let X be one of the following:*
(a) a space such that every finite subset has a non-empty, finite set of local p-barycenters,
(b) $\mathsf{SP}^k Y$ where Y is a space satisfying (a) above.
Suppose $\mathcal{LB}$ is chosen so that $\mathcal{LB}(S, x)$ is always a local p-barycenter of S. Then Algorithm 3.1 terminates after finitely many steps.

Proof. Given the above discussion, we have only to show that $\mathcal{LB}(S, \cdot)$ admits finitely many values for fixed S. For spaces of type (a) this is obvious. For spaces of type (b) we note that by Theorem 12, local barycenters in $\mathsf{SP}^k Y$ are completely determined by the sets S_i^ϕ in Definition 11 and a choice of local p-barycenter for each S_i^ϕ. Since there are only finitely many possibilities for the S_i^ϕ, we get the required result. ∎

Examples of spaces satisfying condition (a) in Proposition 17 for $p = 2$ include $\mathbb{R}^n$ and more generally any Hadamard manifold [10]. An example of a space which does not satisfy (a) is the 2-sphere: any point along the equator is a 2-barycenter of the set consisting of the north and south poles. In the applications to redistricting in Section 5 we will work with $\mathsf{SP}^k X$ where $X = \mathsf{SP}^M \mathbb{R}^2$, and the p-descent operator will be Algorithm 3.1, but applied to X; Corollary 17 thus applies doubly to this setting (see Section 5 for details).

Observation 18. *A local p-barycenter in any Wasserstein space $W_p(X)$ is necessarily a global p-barycenter. Indeed, if μ is a local p-barycenter of $\mu_1, \ldots, \mu_k$ and there exists a μ' with $\sum_i W_p(\mu_i, \mu)^p > \sum_i W_p(\mu_i, \mu')^p$ then for any small $\varepsilon > 0$ let μ_ε be the convex combination $\mu_\varepsilon = \varepsilon\mu' + (1-\varepsilon)\mu$. Taking appropriate convex combinations of optimal plans between μ_i and μ and μ' respectively, one can show $\sum W_p(\mu_i, \mu_\varepsilon)^p < \sum W_p(\mu_i, \mu)^p$. This is a contradiction since μ is a local p-barycenter and the distance between μ and μ_ε can be made arbitrarily small in W_p distance. This all remains true if we restrict ourselves to the subset of $W_p(X)$ consisting of measures with finite support.*

Remark 19. *One can also consider local p-barycenters of a set S* on a subset $\Omega \subseteq X$, *that is, local minimizers in the subspace Ω of the functional $f_{S,p}$ from Equation 3.1. For example, in [18] the authors consider minimizing $f_{S,p}$ on the subset of $W_2(X)$ given by discrete measures with support of size at most m. Theorem 12 holds if we replace "local p-barycenter" by "local p-barycenter on Ω" in the definition of stationary (Definition 11) and replace "local p-barycenter in $\mathsf{SP}^k X$" with "local p-barycenter in $\mathsf{SP}^k X$ on $\mathsf{SP}^k \Omega$". The proof is the same. If we replace "local p-barycenter" in the definition of p-descent operator (Definition 14) by "local p-barycenter on Ω", then Algorithm 3.1 becomes an algorithm for finding local p-barycenters on $\mathsf{SP}^k \Omega$. This more general setup will become relevant when we apply Algorithm 3.1 to unbalanced partitions of point clouds in Section 4.2.*

3.3. Indexing by the barycenter. Let S be a set of points in X^k, and let $[\mathbf{x}]$ be a (local) barycenter of the set $[S] \subseteq \mathsf{SP}^k X$. A choice of representative $\hat{\mathbf{x}} \in [\mathbf{x}]$ gives rise naturally to a way of choosing a representative $\hat{\mathbf{s}} \in [\mathbf{s}]$ for each $[\mathbf{s}] \in S$. Indeed, for $[\mathbf{s}] \in S$, choose an optimal matching π from $\hat{\mathbf{x}}$ to $\mathbf{s}$, and define $\hat{\mathbf{s}}_i = \mathbf{s}_{\pi(i)}$. In other words, label the local barycenter first, and then reorder the points in S so that they are each labeled by a best matching to the local barycenter. In many applications, this reordering of the elements of S is at least as important as the local barycenter itself. In general, however, there may be multiple choices of optimal matching π, so an additional condition is required for the choice of representative

$\hat{\mathbf{s}}_i$ to be unique. In this section we outline some cases for which $\hat{\mathbf{s}}_i$ is defined, focussing on 2-barycenters and the 2-Wasserstein distance.

Definition 20. *Let X be a metric space. We say that X has the* p-barycenter one-point-change (p-OPC) property *if for any finite set of points $x_1, \dots, x_k$ with local p-barycenter $\hat{x}$, the point $\hat{x}$ is not a local p-barycenter of $x_1', \ x_2 \ \dots, x_k$ if $x_1 \neq x_1'$.*

In other words, the OPC property means that local barycenters change whenever exactly one point changes. For the case $p = 2$ and $X = \mathbb{R}^n$ with the Euclidean metric, (local) barycenters are given by the coordinate means, and so it is clear that Euclidean spaces have the 2-OPC property. Note the OPC property does not require that the local barycenter change if another point is removed or added. Indeed, in $\mathbb{R}^n$ the barycenter of a subset S is the same as the barycenter of the subset $S \cup \{\bar{s}\}$ where $\bar{s}$ is the barycenter of S.

Proposition 21. *If X has the p-OPC property, then so does $\mathsf{SP}^k X$ with the W_p metric.*

Proof. If the subset $[S] \subseteq \mathsf{SP}^k X$ has local p-barycenter $[\mathbf{x}]$ then by Theorem 11, for every $1 \leq i \leq k$, $\mathbf{x}_i$ is a local p-barycenter of the set $S_i^\Phi = \{\mathbf{s}_{\phi_{\mathbf{s}}(i)} \mid \mathbf{s} \in S\}$, where $\Phi = (\phi_{\mathbf{s}})_{\mathbf{s}\in S}$ is a set of optimal matchings from $\mathbf{x}$ to each $\mathbf{s} \in S$. If one of the elements of $[S]$ changes then at least one of the S_i^Φ changes by exactly one element. Thus since X has the p-OPC property, $[\mathbf{x}]$ is no longer a local p-barycenter. ∎

Proposition 22. *Let M be a connected, compact Riemannian manifold of dimension n. Then M has the 2-OPC property.*

Proof. Consider a finite set of points $x_1, \dots, x_k$ with local 2-barycenter (in X) $\bar{x}$. For each i, let $\gamma_i : [0,1] \to M$ be a minimal geodesic from $\bar{x}$ to x_i. We claim that for any $\varepsilon \in [0,1]$, $\bar{x}$ is a 2-barycenter of $x_1', \dots, x_k'$ where $x_i' = \gamma_i(\varepsilon)$. Suppose that $\sum d(y, x_i')^2 < \sum d(\bar{x}, x_i')^2$ for some $y \in M$. Then

$$\sum_i d(y, x_i)^2 \leq \sum_i \left(d(y, x_i') + d(x_i', x_i)\right)^2 = \sum_i \left(d(y, x_i')^2 + d(x_i', x_i)^2 + 2d(y, x_i')d(x_i', x_i)\right).$$

Applying the Cauchy-Schwarz Inequality and the facts that $d(\bar{x}, x_i') = \varepsilon d(\bar{x}, x_i)$ and $d(x_i', x_i) = (1-\varepsilon)d(\bar{x}, x_i)$, we see that the latter quantity is upper bounded by

$$\begin{aligned}
&\sum_i d(y, x_i')^2 + \sum_i d(x_i', x_i)^2 + 2\left[\sum_i d(y, x_i')^2 \cdot \sum_i d(x_i', x_i)^2\right]^{1/2} \\
&\quad < \sum_i d(\bar{x}, x_i')^2 + \sum_i d(x_i', x_i)^2 + 2\left[\sum_i d(\bar{x}, x_i')^2 \cdot \sum_i d(x_i', x_i)^2\right]^{1/2} \\
&\quad = \sum_i d(\bar{x}, x_i')^2 + \sum_i d(x_i', x_i)^2 + 2\varepsilon(1-\varepsilon)\sum_i d(\bar{x}, x_i)^2 = \sum_i d(\bar{x}, x_i)^2.
\end{aligned}$$

Since $\bar{x}$ is a local 2-barycenter of the x_i, this implies that y cannot be too near $\bar{x}$, so we have that $\bar{x}$ is indeed a local 2-barycenter of the x_i' as well. Thus, by replacing x_i with x_i' for sufficiently small ε, we can assume from here on that the x_i are contained in a neighborhood of $\bar{x}$ for which the exponential map $\exp_{\bar{x}} : \mathbb{R}^n \to M$ is a diffeomorphism. In particular, $\log_{\bar{x}} x_i$ is defined for

each i. We now recall that in this situation the gradient of the function $p \mapsto d(x_i, p)^2$ at $\bar{x}$ is given by $-2\log_{\bar{x}} x_i$. Thus if $\bar{x}$ is a local 2-barycenter, we have $\sum_i \log_{\bar{x}} x_i = 0$, so that x_1 is completely determined by $\bar{x}$ and the other x_i, as required. ∎

When X has the OPC property, optimal representative choices are always unique:

Proposition 23. *Let X be a space with the p-OPC property. Let S be a set of points in X^k, and let $[\mathbf{x}]$ be a (local) barycenter of the set $[S] \subseteq \mathsf{SP}^k X$. Then for every $\mathbf{s} \in S$, there is a unique representative $\hat{\mathbf{s}} \in [\mathbf{s}]$ minimizing*

$$\sum_{i=1}^{k} d(\hat{\boldsymbol{x}}_i, \hat{\mathbf{s}}_i)^p, \tag{3.2}$$

Proof. For each $\mathbf{s}$, choose a representative $\hat{\mathbf{s}}$ as above and suppose that $\tilde{\mathbf{s}}$ was another possible choice of representative for $\mathbf{s}$ with $\tilde{\mathbf{s}}_m \neq \hat{\mathbf{s}}_m$. By Theorem 12, $\mathbf{x}_m$ is a local p-barycenter of both $Y := \{\hat{\mathbf{s}}_m \mid \mathbf{s} \in S\}$ and $(Y \setminus \hat{\mathbf{s}}_m) \cup \tilde{\mathbf{s}}_m$. But this is a contradiction since X has the p-OPC property. ∎

We now prove that certain 2-Wasserstein spaces have the 2-OPC property. For the absolutely continuous case (Proposition 26), we will need the following two results.

Theorem 24 ([6] for $\mathbb{R}^n$, [36] for M). *Let M be a connected, compact Riemannian manifold or $\mathbb{R}^n$. Let $\mu_1, \mu_2 \in W_2(M)$ be two probability measures which are absolutely continuous with respect to volume. Then the Wasserstein distance $W_2(\mu_1, \mu_2)$ is realized by a unique measure γ. Moreover, this γ is concentrated on the graph of a measureable mapping T over μ_1.*

Theorem 25 ([2, 32]). *Let M be a connected, compact Riemannian manifold or $\mathbb{R}^n$ and let $W_2(M)$ be the 2-Wasserstein space of probability measures on M. Let $\mu_1, \ldots, \mu_k$ in $W_2(M)$ be a finite set of elements of $W_2(M)$ which are all absolutely continuous with respect to volume. Then:*

- *the $\mu_1, \ldots, \mu_k$ admit a unique barycenter $\bar{\mu}$ which is absolutely continuous with respect to volume, and*
- *if we denote by T_k the optimal map from $\bar{\mu}$ to μ_k as guaranteed by Theorem 24 then for $\bar{\mu}$-almost every z, z is the unique barycenter of the points $T_1(z), T_2(z), \ldots, T_k(z)$.*

Proof. For the case where M is a connected, compact Riemannian manifold, the statements are special cases of Theorem 5.1 and Lemma 4.3 in [32] respectively. For the case $M = \mathbb{R}^n$, the first statement was first proven by Agueh and Carlier in [2], and the proof of Lemma 4.3 in [32] works for the second part without modifications. ∎

Proposition 26. *Let M be a connected, compact Riemannian manifold or $\mathbb{R}^n$ and let $W_2(M)$ be the 2-Wasserstein space of probability measures on M. Let $\mu_1, \ldots, \mu_k$ in $W_2(M)$ be a finite set of elements of $W_2(M)$ which are absolutely continuous with respect to volume. Then if $\bar{\mu}$ is a (local) 2-barycenter of $\{\mu_1, \mu_2, \ldots, \mu_k\}$ in $W_2(M)$, then $\bar{\mu}$ is not a (local) 2-barycenter of $\mu_1', \mu_2 \ldots, \mu_k$ when $\mu_1 \neq \mu_1'$.*

Proof. Assume for contradiction that $\bar{\mu}$ is a 2-barycenter of $\mu_1', \mu_2 \ldots, \mu_k$ when $\mu_1 \neq \mu_1'$. By Theorem 25, for $\bar{\mu}$-almost every z, z is the unique barycenter of the points $T_1(z)$, $T_2(z)$,…,$T_k(z)$, where T_k is defined as in Theorem 25. In addition, for $\bar{\mu}$- almost every z, z is

the unique barycenter of the points $T_1'(z), T_2(z), \ldots, T_k(z)$, where T_1' is the optimal map from $\bar{\mu}$ to μ_1'. Since M has the p-OPC property, this means that $T_1'(z) = T_1(z)$ for $\bar{\mu}$-almost every z. It follows that for any measureable subset A, $\mu_1(A) = \bar{\mu}(T_1^{-1}(A)) = \bar{\mu}(T_1'^{-1}(A)) = \mu_1'(A)$, a contradiction since $\mu_1 \neq \mu_1'$. ∎

We now consider the situation where $\mu_1, \ldots, \mu_k$ are finitely supported distributions over $\mathbb{R}^n$, since these are often used in applications to data, or to approximate continuous distributions. In this situation Anderes, Borgwadt and Miller [3] showed the existence of a finite subset S such that any 2-barycenter of $\mu_1, \ldots, \mu_k$ has support contained in S. Let $\mathcal{S} = \{s_0, s_1, \ldots, s_K\}$ denote the union of the above set S and all the supports of $\mu_1, \ldots, \mu_k$, and let $\bar{\mu}$ be any 2-barycenter of $\mu_1, \ldots, \mu_k$. For each i, choose an optimal coupling between each μ and μ_i, encoded as a list of values $y_i^{j\ell}$ where $y_i^{j\ell}$ denotes the coupling between point s_j in μ_i and point s_ℓ in μ. In particular, we have the following identities

$$\forall_i \forall_\ell \sum_j y_i^{j\ell} = \mu(\{s_\ell\}), \quad \forall_i \forall_j \sum_\ell y_i^{j\ell} = \mu_i(\{s_j\}).$$

Lemma 27 (Lemma 1 in [3]). *With the notation above, for any $s_j \in \mathsf{supp}(\bar{\mu})$, there is for each i a unique point $s_{j_i} \in \mathcal{S}$ for which $y_i^{j_i j} > 0$. Moreover, we have $s_j = \sum_i s_{j_i}/k$, so that in particular s_j is the 2-barycenter of the s_{j_i}.*

Intuitively this means that the discrete situation in $\mathbb{R}^n$ is very similar to the absolutely continuous case originally considered in [2]: the support of the Wasserstein barycenter consists of the (metric) barycenters of points it is coupled to. It is therefore unsuprising that the p-OPC property holds in this setting.

Corollary 28. *Let $W_2(\mathbb{R}^n)$ be the 2-Wasserstein space of probability measures on $\mathbb{R}^n$. Let $\mu_1, \ldots, \mu_k$ in $W_2(M)$ be a finite set of discrete distributions with finite support. Then if $\bar{\mu}$ is a (local) 2-barycenter of $\{\mu_1, \mu_2, \ldots, \mu_k\}$ in $W_2(\mathbb{R}^n)$, then $\bar{\mu}$ is not a (local) 2-barycenter of $\mu_1', \mu_2 \ldots, \mu_k$ when $\mu_1 \neq \mu_1'$.*

Proof. Define the couplings $y_i^{j\ell}$ as above, and define $y_1'^{j\ell}$ to be an optimal coupling between μ and μ_1'. Applying Lemma 27 to the set of couplings $y_i^{j\ell}$ as well as this set with $y_1^{j\ell}$ replaced by $y_1'^{j\ell}$, we see that since $\mathbb{R}^n$ has the 2-OPC property, we have $y_1'^{j\ell} > 0 \Leftrightarrow y_1^{j\ell} > 0$. This, together with $\sum_j y_1'^{j\ell} = \mu(\{s_\ell\})$ gives us $y_1'^{j\ell} = y_1^{j\ell}$ for all j, ℓ, hence $\mu_1 = \mu_1'$. ∎

Summarizing, we get the following sufficient conditions for unique optimal representatives.

Corollary 29. *Suppose X is one of the following spaces:*

- *a connected, compact manifold M;*
- *the subset of $W_2(M)$ consisting of absolutely continuous measures with respect to volume, where M is a connected, compact manifold;*
- *the subset of $W_2(\mathbb{R}^n)$ consisting of finitely supported measures;*
- *$\mathsf{SP}^k Y$ for Y either a compact connected manifold or $\mathbb{R}^n$.*

Let S be a set of points in X^k, and let $[\mathbf{x}]$ be a (local) barycenter of the set $[S] \subseteq \mathsf{SP}^k X$ where $\mathsf{SP}^k X$ is endowed with the 2-Wasserstein metric. Then for every $\mathbf{s} \in S$, there is a unique representative $\hat{\mathbf{s}} \in [\mathbf{s}]$ minimizing $\sum_{i=1}^k d(\hat{\boldsymbol{x}}_i, \hat{\mathbf{s}}_i)^2$.

The next example gives a class of spaces which do not have the 2-OPC property.

Example 30. *Let X be any space that admits a branching geodesic $\phi : [a, b] \to X$. Without loss of generality, suppose $\phi(b) \neq \bar{\phi}(b)$ and that the geodesic branches at $t_0 = (a + b)/2$. It is easy to check that $\phi(t_0)$ is the 2-barycenter of the set $\{\phi(a), \phi(b)\}$ as well as of the set $\{\bar{\phi}(a), \bar{\phi}(b)\}$. Since $\phi(a) = \bar{\phi}(a)$, X does not have the 2-OPC property. For a concrete example of a space X with the branching property, consider the space X obtained by gluing three copies of the ray $[0, \infty)$ at zero with d the shortest path metric.*

4. Examples.

To illustrate the flexibility and intuitive outputs of Algorithm 3.1, we provide in this section some experimental results on simple datasets.

4.1. A simple non-Euclidean example.

In this section we demonstrate Algorithm 3.1 in a case where the descent operator does not always pick out a local barycenter, but where the algorithm nonetheless converges to a local barycenter in finite time. In this example X will also not be a Euclidean space, but rather the circle S^1 with the usual geodesic metric. We will use Algorithm 3.1 to compute the local 2-barycenter of a set of points in $\mathsf{SP}^2 S^1$. For two points x, y, we denote by $\angle xy$ the signed clockwise angle from x to y measured in $[-\pi, \pi]$, addressing the choice between $-\pi$ and π as a special case when necessary. The distance between two points is therefore equal to the absolute value of the angle between them.

For $S \subseteq S^1$ and $x \in S^1$, we define $\mathcal{LB}(S, x)$ by the formula $\angle x\mathcal{LB}(S, x) = \frac{1}{|S|}\sum_{s\in S} \angle xs$. If S contains antipodes of x, we define $\angle xs$ for all of the antipodes to be whichever of $\{-\pi, \pi\}$ makes $D' = \sum_{s\in S}(\angle xs - \frac{1}{|S|}\sum_{s\in S} \angle xs)^2$ smaller. If D' is the same for $\pm\pi$, we choose π.

Proposition 31. *With $\mathcal{LB}$ defined as above, $\mathcal{LB}$ is a 2-descent operator. Moreover, for fixed S, $B(S, -)$ takes only finitely many values.*

Proof. It is easy to check that if $\mathcal{LB}(S, x) \neq x$ then $\sum_{s\in S} d(s, \mathcal{LB}(S, x))^2 < \sum_{s\in S} d(s, x)^2$. Suppose then that $\mathcal{LB}(S, x) = x$. We first observe that none of the points in S are diametrically opposite x. Indeed, if some were, then $\frac{1}{|S|}\sum_{s\in S} \angle xs$ must change based on the choice between $-\pi$ and π. Denote one set of chosen angles by $\angle_0 sx$ and the other by $\angle_1 sx$, and suppose without loss of generality that $\frac{1}{|S|}\sum_{s\in S} \angle_0 xs = 0$. Then we get

$$D' = \sum_{s\in S}(\angle_1 xs - \frac{1}{|S|}\sum_{s\in S} \angle_1 xs)^2 < \sum_{s\in S}(\angle_1 xs)^2 = \sum_{s\in S}(\angle_0 xs)^2$$

so that $\angle_1 sx$ is the correct choice, contradicting our assumption that $\mathcal{LB}(S, x) = x$. Since S does not contain an antipode of x, it is easy to check that x is a local barycenter of S. To show the second part of the statement, observe that $\mathcal{LB}(S, x) = \mathcal{LB}(S, x')$ if there are no points of S on one of the arcs between the antipodes of x and x'. It follows that $\mathcal{LB}(S, x)$ is constant on at most $|S|$ arcs whose complement is at most $|S|$ points, so $B(S, x)$ has only finite many possible values for a fixed subset S. ∎

We therefore conclude based on the discussion in Section 3.2 that Algorithm 3.1 terminates after a finite number of iterations when we define $\mathcal{LB}$ as above. To demonstrate this experimentally, we generate 10 random unordered pairs of points on the circle, shown as the leftmost plot in Figure 2. We then run Algorithm 3.1 with two different seeds and plot the barycenter locations as well as the points matched to each component of the barycenter. These are shown in the rightmost two columns of Figure 2. The two different seeds produce different local 2-barycenters of the input set.

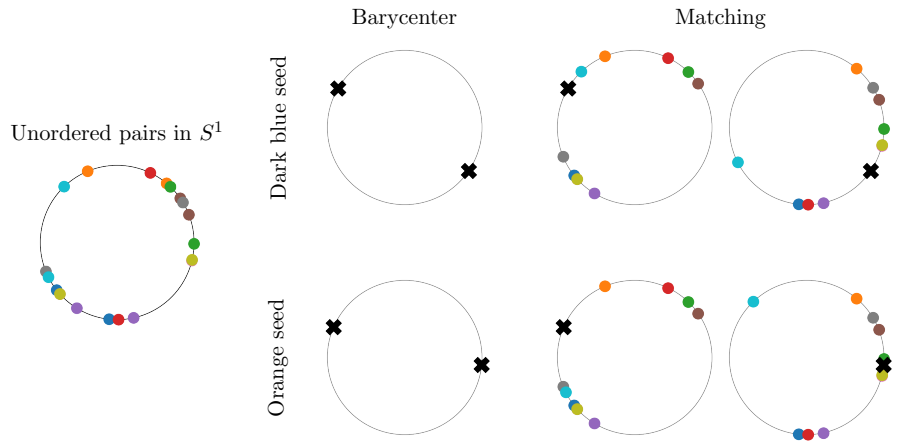

Figure 2: Demonstrating Algorithm 3.1 for a set of points in $\mathsf{SP}^2 S^1$ using the 2-descent operator defined in Section 4.1. For two different initial seeds, the barycenter and the points matched to each point in the barycenter via a unique optimal matching are shown.

4.2. Clustering algorithm consistency. We can use our method to compare the consistency of different clustering methods on point cloud data. To do this, we generate point clouds $P \subseteq \mathbb{R}^2$ consisting of 5,000 points each based on the test datasets provided by `scikit-learn` [42], labeled A–E in Figure 3. For each point cloud, we randomly subsample 10 sets $\{P_1, \ldots, P_{10}\}$ of 500 points each and partition these subsets using a clustering method. If we identify a point cloud with its empirical distribution, we can view a partitioned point cloud as an element of $\mathsf{SP}^k W_2(\mathbb{R}^2)$, where k is the number of clusters. We can therefore compute a local 2-barycenter on $\mathsf{SP}^k \Omega \subseteq \mathsf{SP}^k W_2(\mathbb{R}^2)$ of the ten partitioned point clouds, where Ω is a suitable subspace of $\mathsf{SP}^k W_2(\mathbb{R}^2)$ (see Remark 19). For our experiments, we chose Ω to be the set of atomic distributions on 100 atoms, which is isometric to $\mathsf{SP}^{100}\mathbb{R}^2$ by Proposition 5. The descent operator we used for this application is the restricted version of Algorithm 2 in [18] implemented in [24]. For the initial seed $\mathbf{x}_0$, we under- and oversample the clusters in the first partition as needed to obtain k sets of 100 points each.

For this experiment we test four different clustering methods from `scikit-learn`: k-means, spectral clustering, Ward agglomerative clustering and single linkage clustering. We set $k = 3$ for datasets A, D and E and $k = 2$ for the remaining two. For each point cloud and clustering method, we compute a barycenter B as described above and display them in the middle four columns of Figure 3. We also compute the distance $W_2(B, P_i)$ from the barycenter to each partitioned point cloud and show the distribution of these distances as boxplots.

The $W_2(B, P_i)$ distances shown in the boxplots can be read as a constistency measure in that they indicate how similar the partitioned point clouds were. Note that these values are not a measure of the quality of the clusterings, only their consistency across sub-datasets. Indeed, k-means was very consistent on dataset D despite the barycenter not resembling the intuitively "correct" clustering. Single linkage clustering was the least consistent on datasets A, C and E, and the most consistent on dataset B. The components of the single linkage

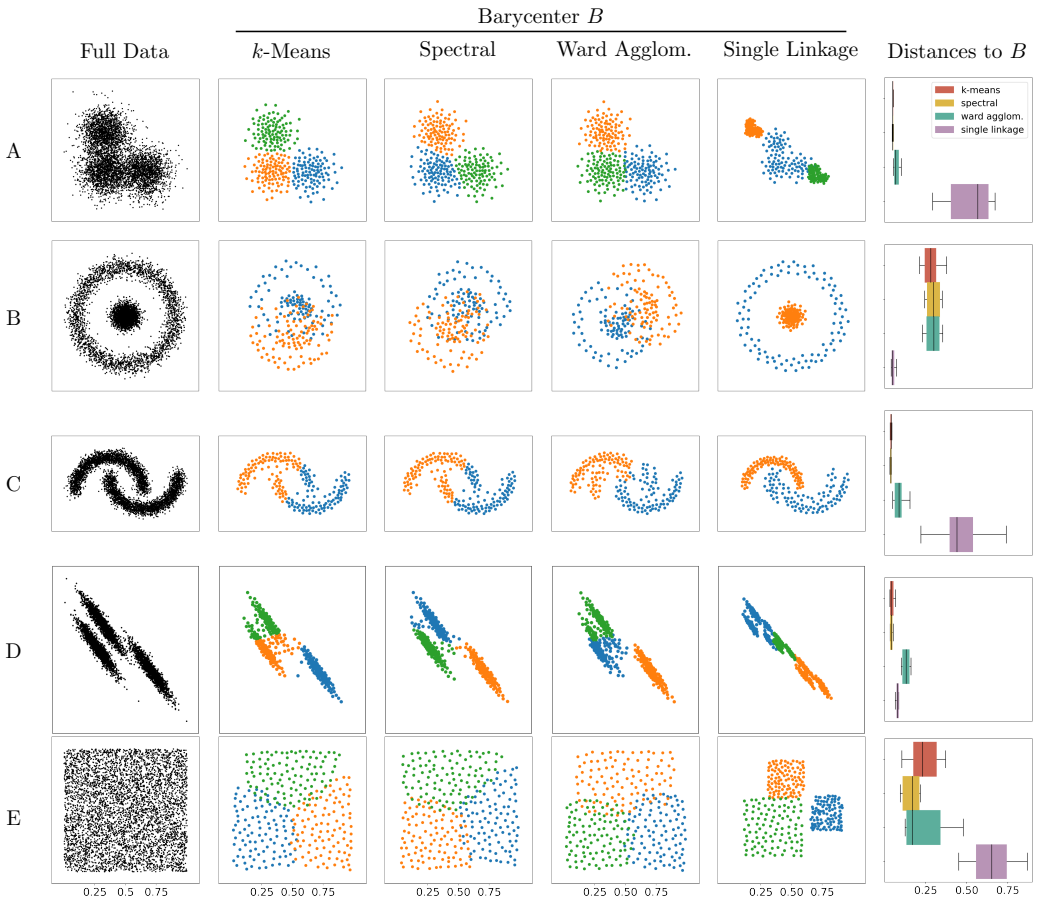

Figure 3: Comparing the consistency of clustering algorithms by subsampling. For each full data set (shown on the left), we generate 10 randomly subsampled datasets and cluster each of them into 2 or 3 clusters using a particular clustering method. We then compute the barycenter of the resulting partitioned point clouds (the middle plots). The distribution of distances from each partition to the barycenter are shown as boxplots (right).

barycenter for A and E resemble affine transformations of the full data set; this is because all of the partitions contained one large component and other very small clusters.

5. Applications to redistricting. In this section, we demonstrate applications of our method to political redistricting. Given a state X, which can be viewed as a polygonal planar region once an appropriate projection is chosen, we view a district $D \subseteq X$ as a probability distribution P_D on X whose support approximates that district. In practice, districts are built out of smaller territorial units such as Census blocks or voting precincts, so to discretize the problem, we assume that the measure of each territorial unit under P_D is concentrated at its centroid. The distribution P_D is therefore a sum of Dirac distributions with support all the centroids of territorial units assigned to D. In this paper, we will consider two different

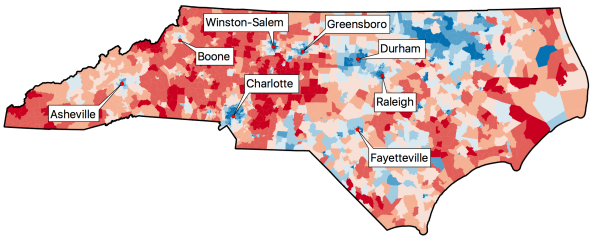

Figure 4: North Carolina's precincts showing two-way vote shares for the 2016 Presidential election. Red indicates more Republican and blue indicates more Democratic.

ways to define P_D: the *area-weighted representation* in which the measure of a territorial unit p under P_D is proportional to its area, and a *population-weighted representation* in which the measure of a precinct p under P_D is proportional to its population.

In order to further reduce the computational cost of the algorithm, we approximate each P_D by sampling M points from it. We choose $M = 40$ in this section based on the stability analysis in Section A. We can therefore identify P_D with an element of $\mathsf{SP}^M\mathbb{R}^2$, and a k-district redistricting plan with an element of $\mathcal{P} = \mathsf{SP}^k\mathsf{SP}^M\mathbb{R}^2$. To compute a local 2-barycenter in $\mathcal{P}$, we use Algorithm 3.1 with a descent operator $\mathcal{LB}_1(S, x)$ given by applying Algorithm 3.1 to $S \subseteq \mathsf{SP}^M\mathbb{R}^2$ using seed x. The descent operator for this inner algorithm is given by $\mathcal{LB}_2(S, x) = \frac{1}{|S|}\sum_{\mathbf{s}\in S}\mathbf{s}$. Theorem 12 and Proposition 17 guarantee that Algorithm 3.1 returns a local 2-barycenter in finite time and Corollary 29 guarantees that there are unique labelings for each redistricting plan that arise from an optimal matching to that barycenter. Note that while enacted redistricting plans often come pre-equipped with identifiers such as "1st Congressional District", "2nd Congressional district", and so on, plans generated by a computer do not, hence the need for the geometry-aware labeling provided by our method.

We demonstrate our method on Congressional redistricting in North Carolina, an area which has seen a great deal of litigation in the last decade, and which has already been extensively studied using ensemble methods in the literature [28, 34, 40, 43]. For some context on North Carolina's political geography, we show the 2016 Presidential two-way vote shares by precinct in Figure 4. An examination of the stability of our method with respect to choice of initial seed and M can be found in Section A; in particular, we give evidence for a high degree of stability for both population- and area-weighted representations. For the sake of brevity, we will only show the population-weighted versions of plots in the main text and relegate the area-weighted version to Appendix (we will however mention any notable differences between the two). More details about the computational aspects of the study are available in Section B as well as a link to our code base.

5.1. Visualizing ensembles. We generate 1,000 Congressional plans for North Carolina using the ReCom Markov chain algorithm [22], each with $k = 13$ districts.[1] We then represent each plan in $\mathsf{SP}^k\mathsf{SP}^M\mathbb{R}^2$ as described above and compute a local 2-barycenter. Figure 5 shows the barycenter as an element of $\mathsf{SP}^k\mathsf{SP}^M\mathbb{R}^2$ for the population-weighted representation; each set of $M = 40$ points of a given color gives one component of the barycenter which is then matched to one district in each plan in the ensemble. The accompanying heat maps show

[1]Note that in the next redistricting cycle, North Carolina will have $k = 14$ Congressional districts, but we keep with the numbers for the 2010-2020 cycle since we want to compare plans enacted during that era to the ensemble.

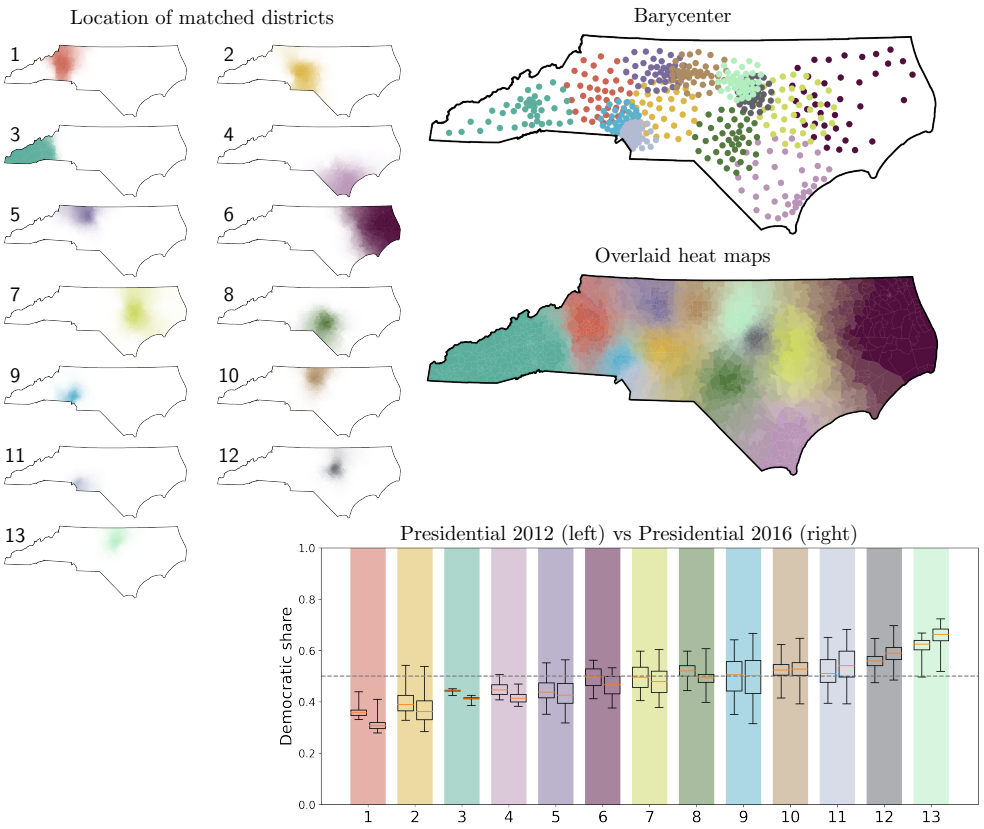

Figure 5: Barycenter for one ensemble of Congressional plans for North Carolina, with population-weighted representations used for the districts, approximated by 40-point samples. The heat maps show the location of districts matched to each component of the barycenter. The boxplots show vote shares for the ensemble the 2012 and 2016 Presidential elections.

where the districts matched to each component lie. The heat maps each appear concentrated and have low overlap with one another, demonstrating that our method is able to successfully label districts based on their geography. Figure 10 shows the area-weighted analysis; the heat maps are almost identical, while the form of the barycenters is very different: the points in the population-weighted barycenter concentrate on population-dense areas such as cities, while the points in each component of the area-weighted barycenter are more evenly spread.

Figures 13 and 12 show barycenters for ReCom ensembles on a selection of 19 states whose precinct data is available from [38]. To determine the number of districts for each state, we use Congressional apportionments from the 2020 Census, which as of the writing of this paper had not gone into effect.[2] In particular, Figures 13 and 12 show North Carolina with 14 districts instead of the 13 used in the analysis in this section.

[2]Population balance is still based on 2010 Census data as that is what was available from [38] at time of writing

5.2. Comparing elections. In this section, we compare two-way Democratic vote shares for our ensemble based on data from the 2012 and 2016 Presidential elections. Figure 5 shows these vote shares for each district label[3]. For ease of visualization, we sort the district labels by mean Democratic vote share under the 2016 data. The statewide two-way vote shares for these two elections differ by less than 1 percent: Obama received 49% of the two-way vote in 2012 and Clinton received 48% in 2016. A common hypothesis used to model shifts in voting patterns in the political science literature is the *uniform swing hypothesis* (see e.g. [31]); in this scenario this hypothesis would dictate that Democratic vote shares in every district or even every precinct would shift by 1% between these two elections. Our method allows us to probe the validity of this hypothesis at the district level for a typical plan. In this case the shift is noticeably non-uniform. Indeed, between 2012 and 2016, the more Republican Districts 1–8, mostly larger, more rural districts, tended to become even more Republican. The Democratic-leaning districts 11–13 near Charlotte and Raleigh, however, became more Democratic. This finding is consistent with the examination of (non)-uniformity of vote shifts in North Carolina using topological data analysis in [40].

5.3. Comparing enacted plans in NC. Since the districts in a computer-generated plan are unsorted, a common approach found in redistricting research and litigation is to sort them by Democratic vote share (under some choice of vote data) and do the same to the plan being evaluated [28, 34, 20, 22]. In other words, these methods are based on the *order statistics* of the vote shares in an ensemble. Figure 6 demonstrates this technique for our ensemble using 2016 Presidential vote data, along with heat maps showing where the districts for each rank lie. Clearly, the districts at a particular rank need not be geographically similar. The result is that if a district from a proposed plan is an outlier compared to districts from the ensemble of the same rank, it is impossible to say whether that district is itself unusually drawn or if it is merely placed in an unusually low or high rank as a result of unusual vote shares in other districts (which may be in completely different part of the state).

A method based on comparing districts which are geographically similar has already been proposed by Mattingly in a blog post [35] and was also presented in [34]. Both of these methods rely on grouping districts who share a particular geographic unit; our method takes a more geometric approach by looking at the geometric distance between districts to group them into geographic clusters. Our method can be applied to any ensemble, and answers the call by Mattingly in [35] for a "more geographically localized analysis" than order statistics.

We compare the ensemble against the Congressional plans used in 2012, 2016 and 2020 respectively, as well as a plan proposed by a bipartisan panel of judges (which we call the 'Judges' plan). The 2012, 2016 and Judges plan were analyzed in [35] and the 2012 and 2016 plans were found to be gerrymanders based on a order-statistics comparison with a neutral ensemble (generated by a different algorithm). We represent these plans in the same way as those in the ensemble and match them to the barycenter to get a labeling (which need not agree with the legal names of these districts). Figure 7 shows the results.

We will make the somewhat arbitrary choice to call a district an *outlier* if it has Democratic vote share which is more than one percentage point (0.01 on the plots) outside the 1st–99th

[3]Following the convention in [21], the boxes show the 25th–75th percentile range, while the whiskers indicate the 1st and 99th percentiles.

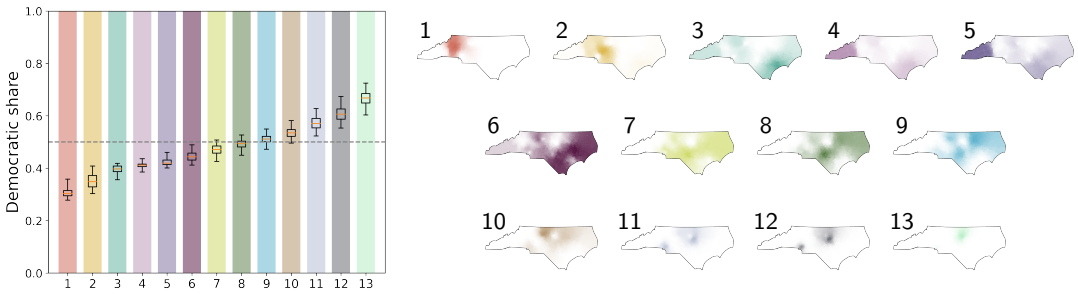

Figure 6: Another way to label districts found in redistricting literature is to order them by Democratic vote share. Applying this with 2016 Presidential vote data to the neutral ensemble used in this section results in the boxplots above showing the vote shares for each label. The heat maps indicate the location of the districts assigned to each label, which are geographically very diffuse.

percentile range of the ensemble. In the 2012 plan, we note that there are five outliers: Districts 2, 3, 7, 12 and 13. District 2 (legally named the 12^{th} Congressional District) in the 2012 plan has a highly unusual shape and is therefore not just a vote share outlier but a geometric outlier too, making it hard to compare with the ensemble. However, since 2012's District 2 has a higher Democratic vote share than any district in the entire ensemble, it does not really matter where it is placed, it will still be an outlier. In the 2016 plan we find three outliers: Districts 3, 7 and 9. Note the unusually low Democratic vote shares in District 3 in both the 2012 and 2016 plan despite the very narrow range in ensemble values. This phenomenon was observed in [40] using other methods. The Judges plan has no outliers, while the 2020 plan has one: District 9. In the area-weighted analysis (Figure 11 in the Appendix), some labelings change, but not the number of outliers in each plan.

6. Conclusion and Future Work. We have introduced a method for labeling unordered k-tuples in a geometrically coherent way using local barycenters in symmetric product spaces. The algorithm (Algorithm 3.1) for computing these local barycenters is very general and depends on a inner local barycenter operation in a modular way. We have demonstrated how this method enables a new analysis technique for redistricting ensembles, effectively summarizing and organizing large sample sets from the non-linear and extremely diverse set of possible redistricting plans for a state. Beyond redistricting, we expect this approach to have applications to problems in machine learning involving successive applications of k-means or other classifiers to partition multiple datasets into unlabeled clusters.

This work suggests many directions of future research. For example, it would be interesting to study theoretical properties of local p-descent operators on non-Euclidean data, as in Section 4.1 for the circle. The development of more advanced statistical machine learning methods for symmetric product spaces would be useful for more in-depth analysis of clustering algorithms, as in Section 4.2, or for studying spaces of districting plans in further detail.

Acknowledgements. The authors would like to thank Justin Solomon for discussions early on in the project, and Olivia Walch for the color scheme used for districts in the paper.

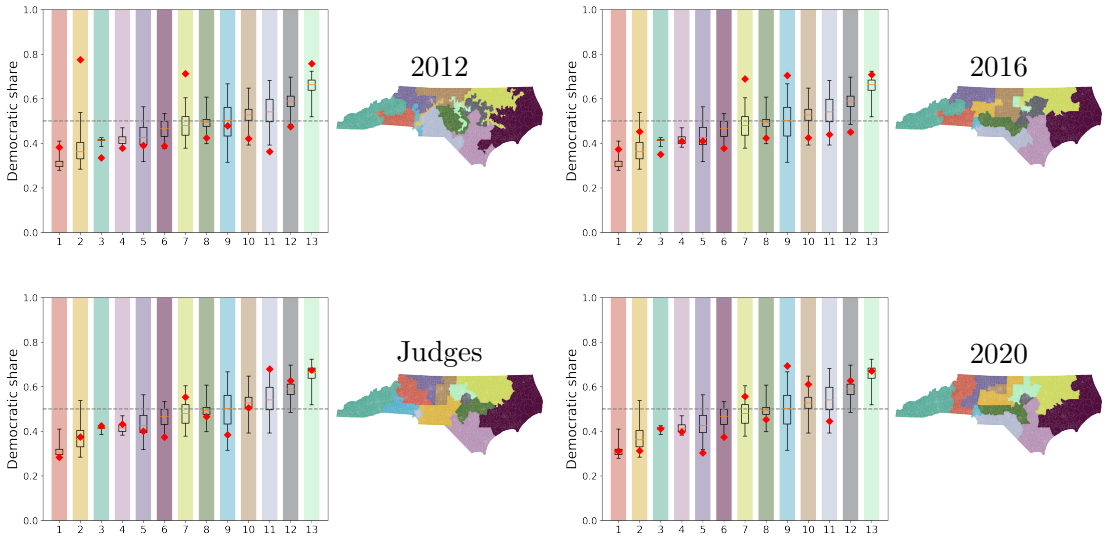

Figure 7: Comparison of four enacted or proposed Congressional plans for North Carolina using votes shares from the Presidential 2016 race. Population-weighting was used. The red diamonds in each boxplot indicate the vote shares of the plan being evaluated for the district matched to that component of the barycenter. The maps on the right show the plan in question colored by a best matching to the barycenter.

REFERENCES

[1] T. Abrishami, N. Guillen, P. Rule, Z. Schutzman, J. Solomon, T. Weighill, and S. Wu, *Geometry of graph partitions via optimal transport*, SIAM Journal on Scientific Computing, 42 (2020), pp. A3340–A3366.

[2] M. Agueh and G. Carlier, *Barycenters in the wasserstein space*, SIAM Journal on Mathematical Analysis, 43 (2011), pp. 904–924.

[3] E. Anderes, S. Borgwardt, and J. Miller, *Discrete wasserstein barycenters: optimal transport for discrete data*, Mathematical Methods of Operations Research, 84 (2016), pp. 389–409.

[4] S. Bangia, C. V. Graves, G. Herschlag, H. S. Kang, J. Luo, J. C. Mattingly, and R. Ravier, *Redistricting: Drawing the line*, arXiv:1704.03360, (2017).

[5] J. Bertrand and B. Kloeckner, *A geometric study of wasserstein spaces: Hadamard spaces*, Journal of Topology and Analysis, 4 (2012), pp. 515–542.

[6] Y. Brenier, *Décomposition polaire et réarrangement monotone des champs de vecteurs*, CR Acad. Sci. Paris Sér. I Math., 305 (1987), pp. 805–808.

[7] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, vol. 319, Springer Science & Business Media, 2013.

[8] D. Burago, I. D. Burago, Y. Burago, S. Ivanov, S. V. Ivanov, and S. A. Ivanov, *A course in metric geometry*, vol. 33, American Mathematical Soc., 2001.

[9] Y. Burago, M. Gromov, and G. Perel'man, *Ad alexandrov spaces with curvature bounded below*, Russian mathematical surveys, 47 (1992), pp. 1–58.

[10] É. Cartan, *La géométrie des espaces de Riemann*, Gauthier-Villars, 1928.

[11] J. Chen and J. Rodden, *Cutting through the thicket: Redistricting simulations and the detection of partisan gerrymanders*, Election Law Journal, 14 (2015), pp. 331–345.

[12] J. Chen, J. Rodden, et al., *Unintentional gerrymandering: Political geography and electoral bias in*

legislatures, Quarterly Journal of Political Science, 8 (2013), pp. 239–269.
[13] S. Chewi, T. Maunu, P. Rigollet, and A. J. Stromme, *Gradient descent algorithms for bures-wasserstein barycenters*, in Conference on Learning Theory, PMLR, 2020, pp. 1276–1304.
[14] M. Chikina, A. Frieze, and W. Pegden, *Assessing significance in a Markov chain without mixing*, Proceedings of the National Academy of Sciences, 114 (2017), pp. 2860–2864.
[15] S. Chowdhury, *Geodesics in persistence diagram space*, arXiv preprint arXiv:1905.10820, (2019).
[16] S. Chowdhury and F. Mémoli, *Explicit geodesics in gromov-hausdorff space*, Electronic Research Announcements, 25 (2018), p. 48.
[17] S. Claici, E. Chien, and J. Solomon, *Stochastic wasserstein barycenters*, in International Conference on Machine Learning, PMLR, 2018, pp. 999–1008.
[18] M. Cuturi and A. Doucet, *Fast computation of wasserstein barycenters*, in International conference on machine learning, PMLR, 2014, pp. 685–693.
[19] D. DeFord and M. Duchin, *Redistricting reform in Virginia: Districting criteria in context*, Virginia Policy Review, (2019).
[20] D. DeFord, M. Duchin, and J. Solomon, *Comparison of districting plans for the virginia house of delegates*, tech. report, MGGG, 2018. https://mggg.org/VA-report.pdf.
[21] D. DeFord, M. Duchin, and J. Solomon, *Recombination: A family of Markov chains for redistricting*, Submitted, (2019).
[22] D. DeFord, M. Duchin, and J. Solomon, *Recombination: A family of markov chains for redistricting*, Harvard Data Science Review, (2021), https://doi.org/10.1162/99608f92.eb30390f, https://hdsr.mitpress.mit.edu/pub/1ds8ptxu. https://hdsr.mitpress.mit.edu/pub/1ds8ptxu.
[23] S. Ferradans, N. Papadakis, G. Peyré, and J.-F. Aujol, *Regularized discrete optimal transport*, SIAM Journal on Imaging Sciences, 7 (2014), pp. 1853–1882.
[24] R. Flamary, N. Courty, A. Gramfort, M. Z. Alaya, A. Boisbunon, S. Chambon, L. Chapel, A. Corenflos, K. Fatras, N. Fournier, et al., *Pot: Python optimal transport*, Journal of Machine Learning Research, 22 (2021), pp. 1–8.
[25] R. Ghrist, *Configuration spaces, braids, and robotics*, in Braids: Introductory Lectures on Braids, Configurations and Their Applications, World Scientific, 2010, pp. 263–304.
[26] P. Harms, P. W. Michor, X. Pennec, and S. Sommer, *Geometry of sample spaces*, arXiv preprint arXiv:2010.08039, (2020).
[27] G. Herschlag, H. S. Kang, J. Luo, C. V. Graves, S. Bangia, R. Ravier, and J. C. Mattingly, *Quantifying gerrymandering in North Carolina*, arXiv:1801.03783, (2018).
[28] G. Herschlag, H. S. Kang, J. Luo, C. V. Graves, S. Bangia, R. Ravier, and J. C. Mattingly, *Quantifying Gerrymandering in North Carolina*, arXiv:1801.03783 [physics, stat], (2018), http://arxiv.org/abs/1801.03783 (accessed 2018-11-14). arXiv: 1801.03783.
[29] G. Herschlag, R. Ravier, and J. C. Mattingly, *Evaluating partisan gerrymandering in Wisconsin*, arXiv:1709.01596, (2017).
[30] L. V. Kantorovich, *On the translocation of masses*, in Dokl. Akad. Nauk. USSR (NS), vol. 37, 1942, pp. 199–201.
[31] J. N. Katz, G. King, and E. Rosenblatt, *Theoretical foundations and empirical evaluations of partisan fairness in district-based democracies*, American Political Science Review, 114 (2020), pp. 164–178.
[32] Y.-H. Kim and B. Pass, *Wasserstein barycenters over riemannian manifolds*, Advances in Mathematics, 307 (2017), pp. 640–683.
[33] J. Lott and C. Villani, *Ricci curvature for metric-measure spaces via optimal transport*, Annals of Mathematics, (2009), pp. 903–991.
[34] R. v. C. C. Mathematicians' Amicus Brief, *Amicus brief of mathematicians, law professors, and students in support of appelleees and affirmance*. Amicus Brief, Supreme Court of the United States, Rucho et al. v. Common Cause et al., March 2018.
[35] J. Mattingly, *Localized view of quantifying gerrymandering*, March 2019, https://sites.duke.edu/quantifyinggerrymandering/2019/03/04/localized-view-of-quantifying-gerrymandering/ (accessed 2021-07-07).
[36] R. J. McCann, *Polar factorization of maps on riemannian manifolds*, Geometric & Functional Analysis GAFA, 11 (2001), pp. 589–608.
[37] Metric Geometry and Gerrymandering Group, *mggg/gerrychain: v0.2.12*, July 2019, https://

github.com/mggg/gerrychain.

[38] Metric Geometry and Gerrymandering Group and R. Buck, *mggg-states*, September 2019, https://github.com/mggg-states.

[39] G. Monge, *Mémoire sur la théorie des déblais et des remblais*, Histoire de l'Académie Royale des Sciences de Paris, (1781).

[40] T. W. Moon Duchin, Tom Needham, *The (homological) persistence of gerrymandering*, Foundations of Data Science, (2021).

[41] S.-i. Ohta, *Barycenters in alexandrov spaces of curvature bounded below*, Advances in geometry, 12 (2012), pp. 571–587.

[42] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Scikit-learn: Machine learning in Python*, Journal of Machine Learning Research, 12 (2011), pp. 2825–2830.

[43] W. Pegden, J. Rodden, and S. Wang, *Brief of amici curiae professors wesley pegden, jonathan rodden, and samuel wang in support of appellees.* Supreme Court of the United States, March 2018.

[44] G. Peyré and M. Cuturi, *Computational optimal transport*, Foundations and Trends in Machine Learning, 11 (2019), pp. 355–607.

[45] G. Puccetti, L. Rüschendorf, and S. Vanduffel, *On the computation of wasserstein barycenters*, Journal of Multivariate Analysis, 176 (2020), p. 104581.

[46] J. Rabin, G. Peyré, J. Delon, and M. Bernot, *Wasserstein barycenter and its application to texture mixing*, in International Conference on Scale Space and Variational Methods in Computer Vision, Springer, 2011, pp. 435–446.

[47] S. T. Rachev and L. Rüschendorf, *Mass Transportation Problems: Volume I: Theory*, vol. 1, Springer Science & Business Media, 1998.

[48] J. Solomon, F. De Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas, *Convolutional wasserstein distances: Efficient optimal transportation on geometric domains*, ACM Transactions on Graphics (TOG), 34 (2015), pp. 1–11.

[49] K. T. Sturm, *Probability measures on metric spaces of nonpositive curvature, heat kernels and analysis on manifolds, graphs, and metric spaces (paris, 2002)*, Contemp. Math., 338 (2003), pp. 357–390.

[50] K.-T. Sturm et al., *On the geometry of metric measure spaces*, Acta mathematica, 196 (2006), pp. 65–131.

[51] K. Turner, Y. Mileyko, S. Mukherjee, and J. Harer, *Fréchet means for distributions of persistence diagrams*, Discrete & Computational Geometry, 52 (2014), pp. 44–70.

[52] C. Villani, *Topics in Optimal Transportation*, no. 58 in Graduate Studies in Mathematics, American Mathematical Soc., 2003.

[53] S. Wagner and D. Wagner, *Comparing clusterings: an overview*, vol. Technical Report 2006-04, Universität Karlsruhe, Fakultät für Informatik Karlsruhe, 2007.

[54] L. Yang, J. Li, D. Sun, and K.-C. Toh, *A fast globally linearly convergent algorithm for the computation of wasserstein barycenters.*, J. Mach. Learn. Res., 22 (2021), pp. 1–37.

[55] T. Yokota, *Convex functions and barycenter on CAT (1)-spaces of small radii*, Journal of the Mathematical Society of Japan, 68 (2016), pp. 1297–1323.

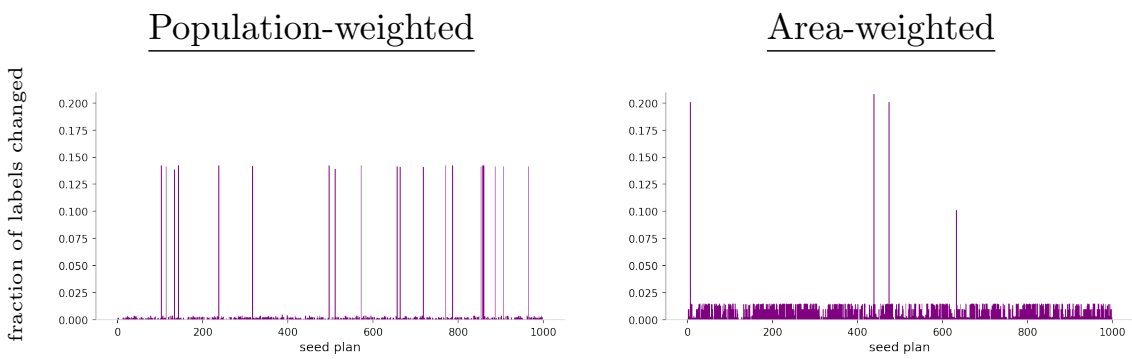

Figure 8: Measuring the dependence of the method on the choice of seed plan $\mathbf{x}_0$. These plots show the discrepancy between the district labels when using the first plan (seed 0) and the i^{th} plan as the seed.

Appendix A. Choice of seed and M for redistricting application. In this section we test the dependence of the ensemble barycenter on the choice of seed for Algorithm 3.1 and also motivate the choice of $M = 40$ sample points per district. Since we mainly interested in the labeling of districts induced by a given barycenter, we measure the discrepancy between seeds by the number of relabelings required, assuming good labelings for the resulting barycenters. To be precise, for a barycenter B with a chosen ordering, let B_i be all the districts labeled District i by matching to B. We define the *discrepancy* $D(B, B')$ to be the fraction

$$\min_{\phi:\langle k\rangle\to\langle k\rangle} \sum_i \frac{|B_i \setminus B'_{\phi(i)}|}{|B_i|}$$

where ϕ ranges over all bijections $\langle k\rangle \to \langle k\rangle$.

We run Algorithm 3.1 1000 times on the neutral ensemble, each time using a different plan as a seed x_0. For each pair $0 \le j < 1000$, we compute the discrepancy between the barycenters coming from seed 0 (the one used in the previous sections) and seed j and display these values in Figure 8. Comparing the population-weighted and area-weighted representations, we see that the population-weighted version has 20 seeds with $> 5\%$ change, while the area-weighted version has 4. On the other hand, for other seeds the discrepancy for the population-weighted representation was generally lower than the area-weighted version. Overall, for both versions, at least 98% of seeds had less than 2% difference from seed 0.

In order to find the right number of points to sample from each district, we first sample $M = 40$ points from each district. We then run Algorithm 3.1 with a fixed seed plan, using only the first i sample points for each district for each $i \in \{1, \dots, 40\}$ to produce a series of 40 barycenters and labelings. For each $t \in \{1, \dots, 39\}$, we compute the discrepancy between the barycenter with t sample points and the barycenter with $t + 1$ sample points. Figure 9 shows the results. We see that for both the population-weighted and area-weighted representations, the discrepancy between successive values of t drops to around 1% at around $t = 20$ and remains low thereafter.

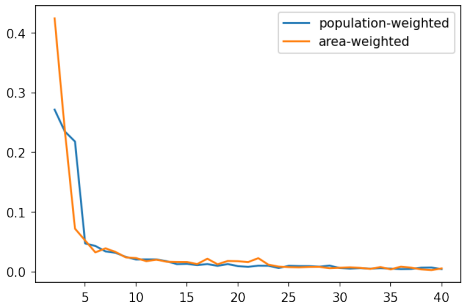

Figure 9: Discrepancy between using the first t points and using the first $t+1$ points of the 40-point samples to compute the barycenter for the neutral ensemble.

Appendix B. Computational details. For the experiments in Section 5, we implemented Algorithm 3.1 in Python.[4] The (outer) 2-descent operator $\mathcal{LB}_1$ is implemented using the Python Optimal Transport library [24] (the particular function used implements the restricted version of Algorithm 2 in [18] discussed in Remark 15). Cleaned population and vote data was obtained from the `mggg-states` repository [38]. For the neutral ensemble, we sampled every 50th plan from an ensemble of 50,000 plans generated using the ReCom algorithm [21] implemented in the Python library GerryChain [37]. The chain was constrained to generate only contiguous districts and population deviation less than 2% from the ideal district population. To give some idea of the computational cost of the method, computing the (population-weighted) barycenter for the neutral ensemble of 1000 Congressional plans for North Carolina (represented in Figure 5) was performed on a single core of a high performance cluster and took about 3,200 seconds (25 iterations) to complete.

Appendix C. Supplemental and area-weighted redistricting figures.

[4]Code is available at https://github.com/thomasweighill/barymandering

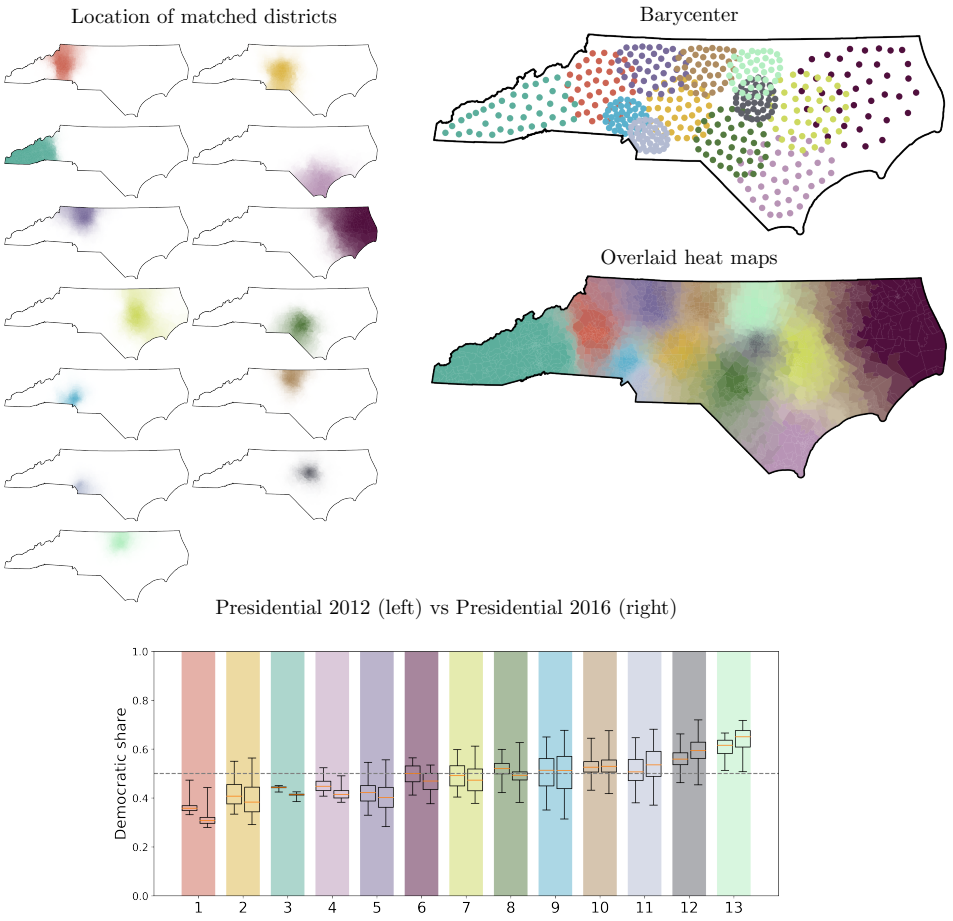

Figure 10: Barycenter for one ensemble of Congressional plans for North Carolina, with area-weighted representations used for the districts, approximated by 40-point samples. The heat maps show the location of districts matched to each component of the barycenter. The boxplots show vote shares for the ensemble using two consecutive Presidential elections: 2012 on the left and 2016 on the right.

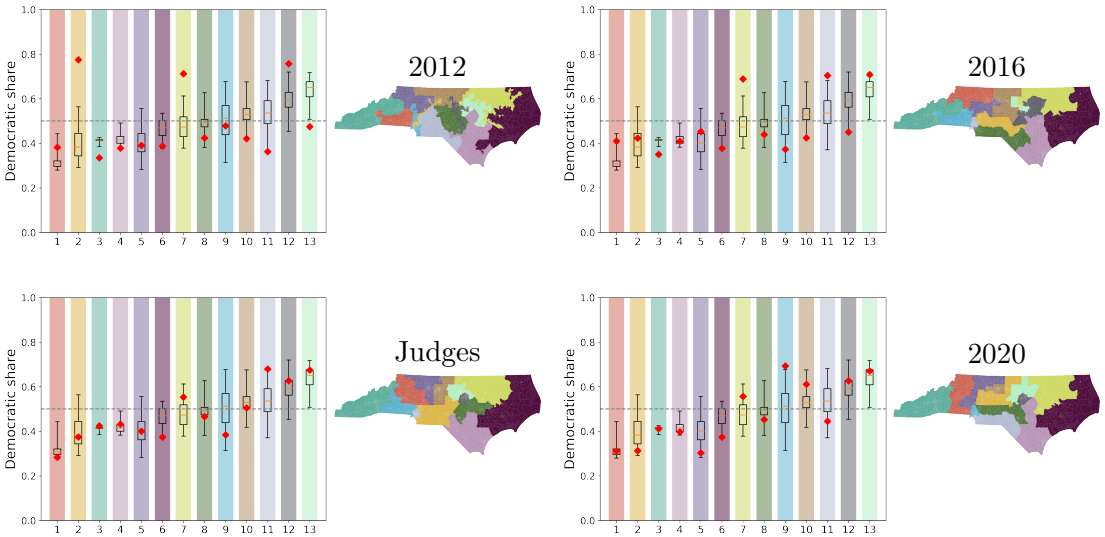

Figure 11: Comparison of four enacted or proposed Congressional plans for North Carolina using votes shares from the Presidential 2016 race. Area-weighting was used. The red diamonds in each boxplot indicate the vote shares of the plan being evaluated for the district matched to that component of the barycenter. The maps on the right show the plan in question colored by a best matching to the barycenter.

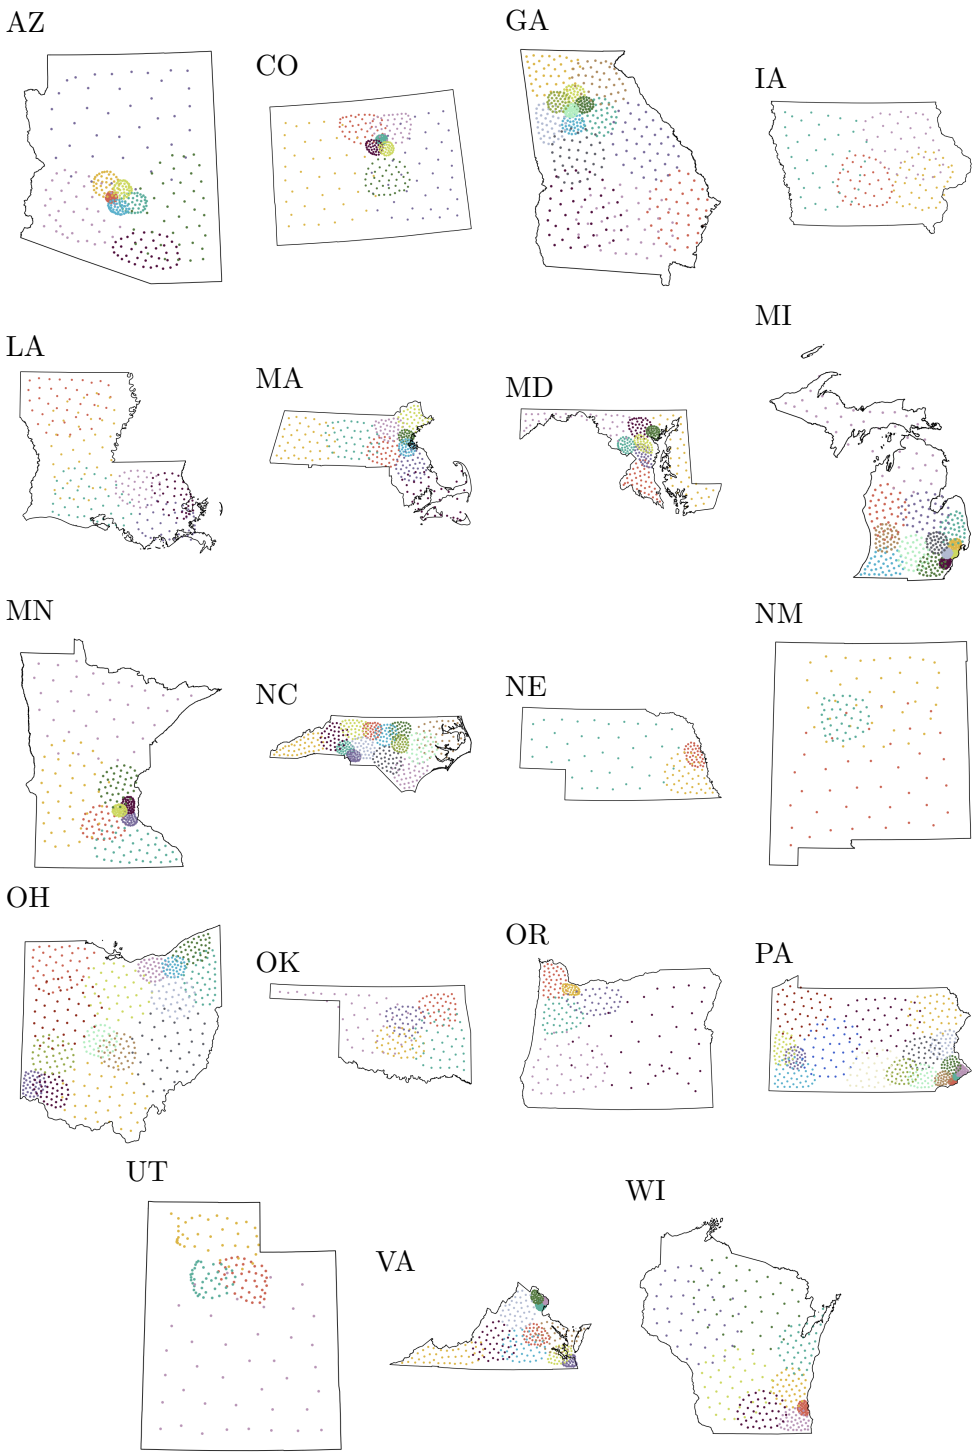

Figure 12: Area-weighted ensemble barycenters for a selection of states using precinct data from [38]. Congressional apportionment from the 2020 Census was used.

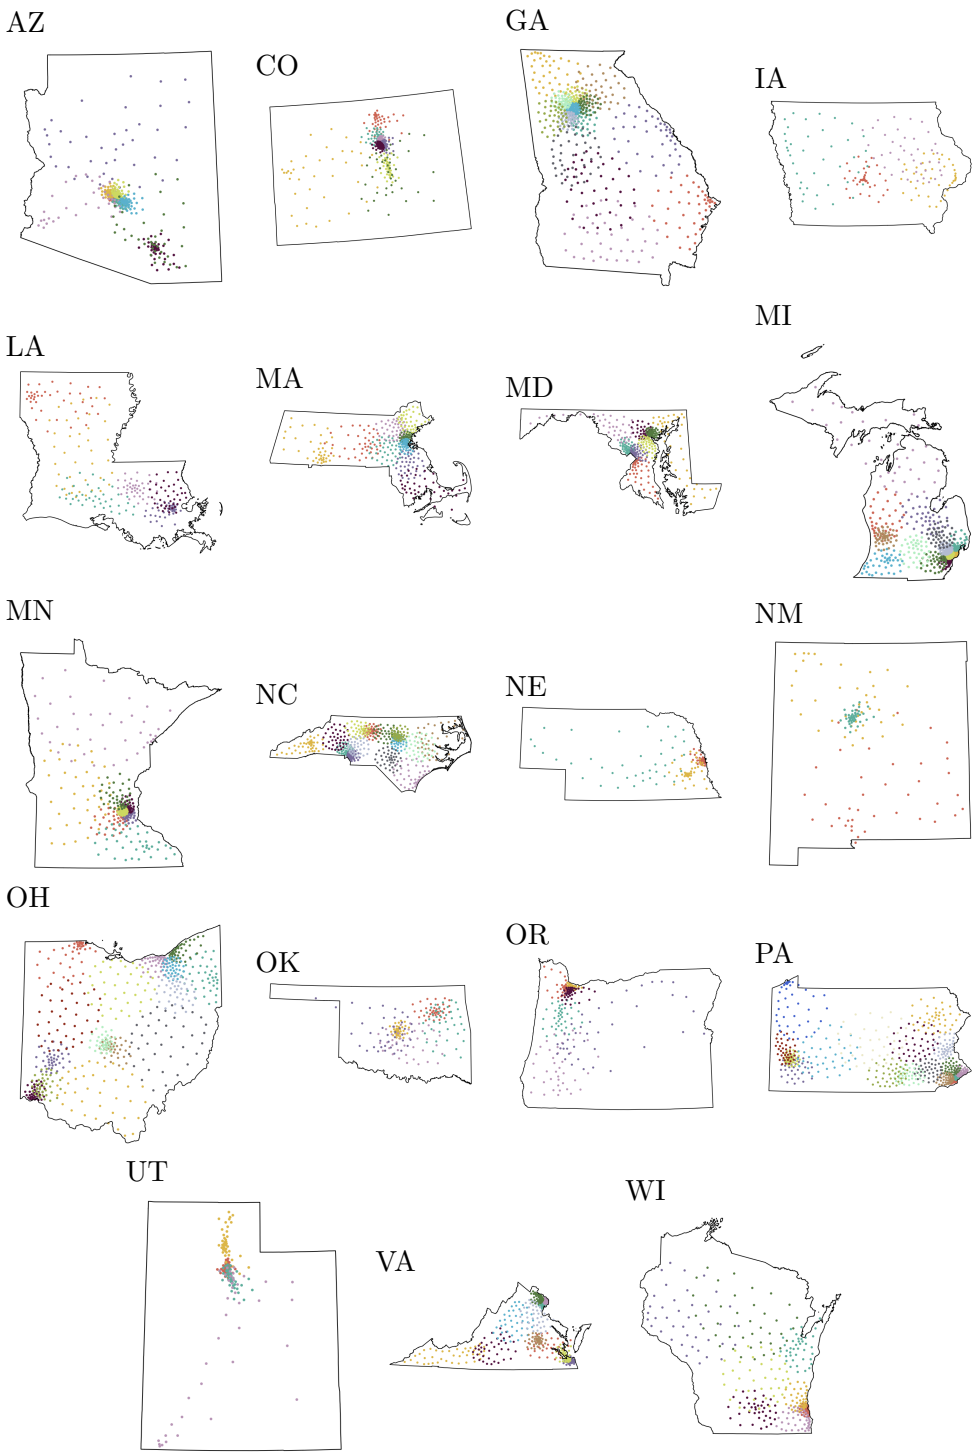

Figure 13: Population-weighted ensemble barycenters for a selection of states using precinct data from [38]. Congressional apportionment from the 2020 Census was used.