

AVERAGES OF UNLABELED NETWORKS: GEOMETRIC CHARACTERIZATION AND ASYMPTOTIC BEHAVIOR

BY ERIC D. KOLACZYK^{1,*}, LIZHEN LIN², STEVEN ROSENBERG^{1,**},
JACKSON WALTERS^{1,†} AND JIE XU^{1,‡}

¹Department of Mathematics and Statistics, Boston University, *kolaczyk@bu.edu; **sr@bu.edu; †jackwalt@bu.edu; ‡xujie@bu.edu

²Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, lizhen.lin@nd.edu

It is becoming increasingly common to see large collections of network data objects, that is, data sets in which a network is viewed as a fundamental unit of observation. As a result, there is a pressing need to develop network-based analogues of even many of the most basic tools already standard for scalar and vector data. In this paper, our focus is on averages of unlabeled, undirected networks with edge weights. Specifically, we (i) characterize a certain notion of the space of all such networks, (ii) describe key topological and geometric properties of this space relevant to doing probability and statistics thereupon, and (iii) use these properties to establish the asymptotic behavior of a generalized notion of an empirical mean under sampling from a distribution supported on this space. Our results rely on a combination of tools from geometry, probability theory and statistical shape analysis. In particular, the lack of vertex labeling necessitates working with a quotient space modding out permutations of labels. This results in a nontrivial geometry for the space of unlabeled networks, which in turn is found to have important implications on the types of probabilistic and statistical results that may be obtained and the techniques needed to obtain them.

1. Introduction. Over the past 15–20 years, as the field of network science has exploded with activity, the majority of attention has been focused upon the analysis of (usually large) *individual* networks; see [23, 27, 35], for example. While it is unlikely that the analysis of individual networks will become any less important in the near future, it is likely that in the context of the modern era of “big data” there will soon be an equal need for the analysis of (possibly large) *collections* of (sub)networks, that is, collections of network data objects.

We are already seeing evidence of this emerging trend. For example, the analysis of massive online social networks like Facebook can be facilitated by local analyses, such as through extraction of ego-networks (e.g., [22]). Similarly, the 1000 Functional Connectomes Project, launched a few years ago in imitation of the data-sharing model long-common in computational biology, makes available a large number of fMRI functional connectivity networks for use and study in the context of computational neuroscience (e.g., [11]). It would seem, therefore, that in the near future networks of small to moderate size will themselves become standard, high-level data objects.

Faced with databases in which networks are the fundamental unit of data, it will be necessary to have in place a network-based analogue of the “Statistics 101” tool box, extending standard tools for scalar and vector data to network data objects. The extension of such classical tools to network-based datasets, however, is not immediate, since networks are not inherently Euclidean objects. Rather, formally they are combinatorial objects, defined simply through two sets, of vertices and edges, respectively, possibly with an additional corresponding set of weights. Nevertheless, initial work in this area demonstrates that networks can be

Received October 2018.

MSC2010 subject classifications. Primary 62E20, 62G20; secondary 53C20.

Key words and phrases. Fundamental domain, Fréchet mean, graphs.

associated with certain natural Euclidean subsets and, furthermore, demonstrates that through a combination of tools from geometry, probability theory and statistical shape analysis it should be possible to develop a comprehensive, mathematically rigorous and computationally feasible framework for producing the desired analogues of classical tools.

For example, in our recent work [21] we have characterized the geometry of the space of all labeled, undirected networks with edge weights, that is, consisting of graphs $G = (V, E, W)$, for weights $w_{ij} = w_{ji} \geq 0$, where equality with zero holds if and only if $\{i, j\} \notin E$. This characterization allowed us in turn to establish a central limit theorem for an appropriate notion of a network empirical mean, as well as analogues of classical one- and two-sample hypothesis testing procedures. Other results of this type include additional work on asymptotics for network empirical means [36] and regression modeling with a network response variable, where for the latter there have been both frequentist [14] and Bayesian [19] proposals. Work in this area continues at a quick pace; see, for example, [2], which proposes a classification model based on network-valued inputs and [20], which proposes a nonparametric Bayes model for distributions on populations of networks. Earlier efforts in this space have focused on the specific case of trees. Contributions of this nature include work on central limit theorems in the space of phylogenetic trees [4, 10] and work by Marron and colleagues [3, 39] in the context of so-called object-oriented data analysis with trees.

To the best of our knowledge, nearly all such work to date pertains to the case of *labeled* networks: that is, to networks in which the vertices V have unique labels, for example, $V = \{1, \dots, d\}$. In fact, *unlabeled* networks have received decidedly less attention in the network science literature as a whole but nevertheless arise in various important settings. For example, surveys of so-called “personal networks” are common in social network analysis, wherein individuals (“egos”) are surveyed for a list of d others (“alters”) with whom they share a certain relationship (e.g., friendship, colleague, etc.) and only common patterns across networks in the structure of the relationships among these others within each network are of interest. This leads to analyses that either ignore vertex labels or for which vertex labels are simply not available (e.g., through de-identification); see [33], for example. Other variations on this idea include, for example, the study of neighborhood subgraphs in the context of online social networks (e.g., [32, 37]), where now the number of vertices d in these neighborhoods typically varies. More generally, the study of ego-networks (whether inclusive or exclusive of the ego vertex and its edges—confusingly, the term is used for both cases) under privacy constraints can increasingly be expected to be a natural source of collections of unlabeled networks.

In this paper, our focus is on averages of unlabeled, undirected networks with edge weights. Adopting a perspective similar to that in our previous work [21], we (i) characterize a certain notion of the space of all such networks, (ii) describe key topological and geometric properties of this space relevant to doing probability and statistics thereupon, and (iii) use these properties to establish the asymptotic behavior of a generalized notion of an empirical mean under sampling from a distribution supported on this space. In particular, adopting the notion of a Fréchet mean, we establish a corresponding strong law of large numbers and a central limit theorem. In contrast to [21], where the corresponding space of networks was found to form a smooth manifold, here the lack of vertex labeling necessitates working with a quotient space modding out permutations of labels. As a result, we have only an orbifold—a more general geometric structure—which in turn is found to have important implications on the types of probabilistic and statistical results that may be obtained and the techniques needed to obtain them.

The nature of our work is in the spirit of statistics on manifolds and statistical shape analysis, which employs the geometry of manifolds or shape spaces for defining Fréchet means and developing large sample theory of their sample counterparts for inference. See [6] for a

rather comprehensive treatment on the subject. Our approach to studying the entire family of networks subject to an equivalence relation under a group action, via forming the associated quotient or moduli space, is a common theme in modern geometry, including gauge theory [17], symplectic topology [34] and algebraic geometry [15, 38]. The appearance of orbifolds, often much more complicated than in our case, is quite common.

Finally, there is a large literature on graph limits, for which substantial work has been done on analysis of appropriate spaces of networks (e.g., [30]). While there are high-level connections between graphons as a space of equivalence classes (as studied in the literature) and unlabeled networks as another space of equivalence classes (as studied here), a meaningful comparison between the two is not immediate. For example, while one natural source of unlabeled networks is as finite-dimensional restrictions of graphons, our framework encompasses both these and more general cases (e.g., *dependent* generation of (non)edges, conditional on a graphon). On the other hand, the distance defined in our framework for weighted graphs is the Euclidean norm taken over the equivalence classes of weighted matrices under the group action of the permutation group, while the cut distance commonly used for graphons is defined over the equivalent classes of bivariate functions under the group action of all measure-preserving transformations. The permutation group is a finite group, as compared to the infinite dimensional group of all measure-preserving transformations. Ultimately, however, we note too that, whereas the focus in the graphon literature typically is on the case of a single network asymptotically increasing in size (or finite-dimensional restrictions thereof), here the focus is on asymptotics in many networks, with the dimension fixed. These and related issues suggest that a simple comparison is unlikely to be forthcoming.

The organization of this paper is as follows. In Section 2, we present our characterization of the space of unlabeled networks. Results from our investigation of the asymptotic behavior of the Fréchet empirical mean are then provided in Section 3. While a strong law of large numbers is found to emerge under quite general conditions, establishing just when conditions dictated by the current state of the art for central limit theorems on manifolds hold turns out to be a decidedly more subtle exercise. This latter is the focus of Section 4. Some additional discussion of open problems may be found in Section 5. The Appendices discuss implementation issues for the main theoretical results in the paper, and the Supplementary Material [28] contain several results related to the latter.

After we submitted this paper, we became aware of related work of B. Jain. In [24], Theorem 4.2, a result is stated that is equivalent to our main result on the uniqueness of the Fréchet mean (Theorem 4.5). However, no proof is given in [24], although the companion paper [25] contains much of the material needed for a proof. Here, in Section 4, we give a complete proof of this result. In addition, we also give two extensions in Supplement C. Proofs of both the main result and the extensions are nontrivial. On the computational side, our work gives an algorithm for the computation of fundamental domains (see Appendix A), while [25], Section 3.3, gives an algorithm for iteratively approximating the sample mean for the Fréchet function of a distribution on the space of unlabeled networks. Both algorithms are addressing NP-hard problems.

2. The space of unlabeled networks. Our ultimate focus in this paper is on a certain well-defined notion of an “average” on elements drawn randomly from a “space” of unlabeled networks and on the statistical behavior of such averages. Accordingly, we need to establish and understand the relevant topology and geometry of this space. We do so by associating labeled networks with vectors and mapping those to unlabeled networks through the use of equivalence classes in an appropriate quotient space. In this section, we provide relevant definitions, characterization and illustrations of this space of unlabeled networks.

2.1. *The topological space of unlabeled networks.* Let $G = (V, E)$ be a labeled, undirected graph/network with weighted edges and with d vertices/nodes. We always think of E as having $D := \binom{d}{2}$ elements, where some of the edge weights can be zero. We think of the edge weight between vertices i and j as the strength of some unspecified relationship between i and j .

Let Σ_d be the group of permutations of $\{1, 2, \dots, d\}$. A permutation $\sigma \in \Sigma_d$ of the d vertex labels technically produces a new graph σG , but with no new information. To define σG precisely, note that the weight function $w_G : E \rightarrow \mathbb{R}_{\geq 0}$ can be thought of as a symmetric function $w_G : V \times V \rightarrow \mathbb{R}_{\geq 0}$, with $w_G(i, j)$ the weight of the edge joining vertex i and vertex j in G . Therefore, the action of Σ_d on w_G is given by

$$(\sigma \cdot w_G)(i, j) = w_G(\sigma^{-1}(i), \sigma^{-1}(j)).$$

(The inverse guarantees that $(\sigma\tau) \cdot w_G = \sigma \cdot (\tau \cdot w_G)$.) Note that for general G , not all permutations of the entries of w_G are of the form $\sigma \cdot w_G$, as w_G may have $[d(d-1)/2]!$ distinct permutations and Σ_d has $d!$ elements

In summary, σG is defined to be the graph on d vertices with weight function $\sigma \cdot w_G : E \rightarrow \mathbb{R}_{\geq 0}$. Let $\mathcal{G} = \mathcal{G}_d$ be the set of all labeled graphs with d vertices. Then the quotient space

$$\mathcal{U}_d = \mathcal{G}_d / \Sigma_d$$

is the space of unlabeled graphs, the object we want to study. This means that an unlabeled network $[G] \in \mathcal{U}_d$ is an equivalence class

$$[G] := \{\sigma \cdot G : \sigma \in \Sigma_d\}.$$

As we now explain, $\mathcal{G}_d$ looks like an explicit subset of $\mathbb{R}^d$, and so is easy to picture. In contrast, the quotient space $\mathcal{U}_d$ is difficult to picture. Nevertheless, as we describe in the following paragraphs, the topology of $\mathcal{U}_d$ may be characterized through standard point-set topology techniques, with the conclusion that everything works as well as possible. Readers who wish can safely skip to the examples in Section 2.2.

Fix an ordering of the vertices $1, \dots, d$, and take the lexicographic ordering $\{(i, j) : 1 \leq i < j \leq d\}$ on the set of edges. (Thus $(i, j) < (k, \ell)$ if $i < k$ or $i = k$ and $j < \ell$.) Given this ordering, we get an injection

$$\alpha : \mathcal{G}_d \rightarrow \mathbb{R}^D, \quad \alpha(G) = (w_1(G), \dots, w_D(G)),$$

where $w_i(G)$ is the weight of the i th edge of G . The image of α is the first “octant” $\mathcal{O}_D = \{\vec{x} = (x^1, \dots, x^D) : x^i \geq 0\}$. For simplicity, we choose the standard Euclidean metric on $\mathcal{O}_D$. This pulls back via α to a metric on $\mathcal{G}_d$ with the desirable property that two networks are close iff their edge weights are close. Similarly, the standard topology on $\mathcal{O}_D$ (an open ball in $\mathcal{O}_D$ is the intersection of an open $\mathbb{R}^D$ -Euclidean ball with $\mathcal{O}_D$) pulls back to a topology on $\mathcal{G}_d$. (This just means that $A \subset \mathcal{G}_d$ is open iff $\alpha(U)$ is open in $\mathcal{O}_D$. This makes α a homeomorphism.) Just as in $\mathbb{R}^D$, the metric and topology are compatible: a sequence of graphs/weight vectors $\vec{x}_i$ in $\mathcal{O}_D$ converges to a graph/weight vector $\vec{x}$ in the topology of $\mathcal{O}_D$ iff the distance from $\vec{x}_i$ to $\vec{x}$ goes to zero.

Via the bijection α , the action of Σ_d on $\mathcal{G}_d$ transfers to an action on $\mathcal{O}_D$. First, $\sigma \in \Sigma_d$ acts on $\{1, \dots, D\}$ by $\sigma \cdot i = j$ if i corresponds to the edge (i_1, i_2) and j corresponds to the edge $(\sigma(i_1), \sigma(i_2))$. Then σ acts on $\mathcal{O}_D$ by $\sigma \cdot \vec{x} = (x^{\sigma^{-1}(1)}, \dots, x^{\sigma^{-1}(D)})$. Since we have arranged the actions to be compatible with $\alpha : \mathcal{G}_d \rightarrow \mathcal{O}_D$, we get a well-defined bijection $\bar{\alpha}$:

$$\mathcal{U}_d = \mathcal{G}_d / \Sigma_d \xrightarrow{\bar{\alpha}} \mathcal{O}_D / \Sigma_d, \quad \bar{\alpha}[G] = [\alpha(G)].$$

From now on, we just denote $\bar{\alpha}$ by α .

To complete the topological discussion, we note that $\alpha : \mathcal{U}_d \rightarrow \mathcal{O}_D/\Sigma_d$ is a homeomorphism if we give both sides the quotient topology: for the map $q : \mathcal{G}_d \rightarrow \mathcal{U}_d$ taking a graph to its equivalence class, a set $U \subset \mathcal{U}_d$ is open iff $q^{-1}(U)$ is open in $\mathcal{O}_D$. The quotient topology on $\mathcal{O}_D/\Sigma_d$ is defined similarly.

2.2. *Examples of quotient spaces.* As a warmup, we first give a simple example of a quotient space resulting from the action of a finite group on a Euclidean space. This particular example is important in providing a relevant nonnetwork analogy to our network-based results. We will revisit it frequently throughout the paper.

EXAMPLE 2.1. The group $\mathbb{Z}_4 = \{0, 1, 2, 3\}$ acts on the plane $\mathbb{R}^2$ by rotation counter-clockwise by 90 degrees: specifically, for $k \in \mathbb{Z}_4$ and $z \in \mathbb{R}^2 = \mathbb{C}$,

$$k \cdot z = e^{ik\pi/2} \cdot z.$$

Thus $0 \cdot z = z$, $1 \cdot z = e^{i\pi/2}z$, etc. A point in the quotient space $\mathbb{C}/\mathbb{Z}_4$ is the set $[z_0] = \{e^{ik\pi/2}z_0 : k \in \mathbb{Z}_4\}$. The set $[z_0]$ is called the orbit of z_0 under $\mathbb{Z}_4$. Note that every orbit is a four element set except for the exceptional orbit $[\vec{0}] = \{\vec{0}\}$.

The closed first quadrant $F = \{(x, y) : x \geq 0, y \geq 0\}$ is a *fundamental domain* for this action, that is, each orbit $[z_0]$ has a unique representative/element in F , except possibly for the orbits of points on the boundary $\partial F = \{(x, y) : x = 0 \text{ or } y = 0\}$ of F . Orbits of boundary points like $[(5, 0)]$ have two representatives $(5, 0), (0, 5)$ in F , while the origin of course has only one representative.

Here is a precise definition of a fundamental domain for the action of a group G on a set S .

DEFINITION 2.2. $F \subset S$ is a fundamental domain for the action of G if: (i) S is the union of the orbits of F ($S = \bigcup_{k \in G} k \cdot F$);

(ii) orbits can intersect only at boundary points ($k_1 \cdot F \cap k_2 \cdot F = \emptyset$ or $k_1 \cdot F \cap k_2 \cdot F \subset \partial(k_1 \cdot F) \cap \partial(k_2 \cdot F)$).

In this example, $G = \mathbb{Z}_4$ and $S = \mathbb{C}$. It follows that the quotient map $q : F \rightarrow \mathbb{C}/\mathbb{Z}_4$ is surjective, a homeomorphism on the interior of F (where $\mathbb{C}/\mathbb{Z}_4$ has the quotient topology), and finite-to-one on the boundary of F .

If we want to picture a set that is bijective to $\mathbb{C}/\mathbb{Z}_4$, we could take, for example, F' to be F minus the positive y -axis. This is not so helpful topologically or geometrically, as the points $[(5, 0)]$ and $[(0, 5.01)]$ have close representatives $(5, 0), (5.01, 0)$ in F , while their representatives $(5, 0)$ and $(0, 5.01)$ are not close in F' . In particular, the sequence $(10^{-k}, 5)$ does not converge in F' , but the orbits $[10^{-k}, 5]$ converge to $[0, 5] = [5, 0]$ in $\mathbb{C}/\mathbb{Z}_4$. Thus F' does not give us a good picture of $\mathbb{C}/\mathbb{Z}_4$ topologically.

In summary, it is much better to keep both positive axes in F , and to consider $\mathbb{C}/\mathbb{Z}_4$ as (in bijection with) F with the boundary points $(a, 0)$ and $(0, a)$ “glued together.” More precisely, we have a bijection

$$\beta : \mathcal{F} := \frac{F}{(a, 0) \sim (0, a)} \rightarrow \mathbb{C}/\mathbb{Z}_4,$$

where the denominator indicates that the two point set $\{(a, 0), (0, a)\}$ ($a \neq 0$) is one point of $\mathcal{F}$, while all other points of F correspond to a single point in $\mathcal{F}$. At the price of this gluing, we now have that β is a homeomorphism: in particular, $\lim_{i \rightarrow \infty} x_i = x$ in $\mathcal{F}$ iff $\lim_{i \rightarrow \infty} \beta(x_i) = \beta(x)$ in $\mathbb{C}/\mathbb{Z}_4$. (Technical remark: $\mathcal{F}$ gets the quotient topology from the standard topology on F and the obvious surjection $q : F \rightarrow \mathcal{F}$.)




FIG. 1. In the figure on the left, F is a fundamental domain F for the action of $\mathbb{Z}_4$ on $\mathbb{C}$. The four point orbit of $(5, 0)$ and the one point orbit of $\vec{0}$ are shown. In the figure on the right, the quotient space $\mathbb{C}/\mathbb{Z}_4$ is drawn as a hollow cone given by taking F and gluing $(x, 0)$ to $(0, x)$.

Although this seems a little involved, it is quite easy to perform the gluing in $\mathcal{F}$ in rubber sheet topology: stretching the interior of F to allow the gluing of the two axes shows that $\mathcal{F}$, and hence $\mathbb{C}/\mathbb{Z}_4$ is a hollow cone; see Figure 1.¹

EXAMPLE 2.3. We discuss the case of a network with three vertices. This is a deceptively easy case, as $3 = d = D$ implies that every permutation of the D edge weights comes from a permutation in Σ_d . In higher dimensions, the details are more complicated.

In Supplement A, we describe the quotient space $\mathcal{U}_3$ of unlabeled graphs directly. However, it is easier to picture $\mathcal{U}_3$ by finding a fundamental domain F inside $\mathcal{O}_3$ for the action of Σ_3 . As in the previous example, F is a closed set such that the quotient map $q|_F : F \rightarrow \mathcal{U}_3$ is a continuous surjection, a homeomorphism from the interior of F to its image, and a finite-to-one map on the boundary ∂F of F . Thus F represents $\mathcal{U}_3$ bijectively except for some gluings on the boundary. This is illustrated in Figure 2, where $F = \{(x, y, z) : x \geq y \geq z \geq 0\}$. Again, the case $d = 3$ is deceptively easy, as F is a bijection even on ∂F .

The situation is more complicated for graphs with 4 (or more) vertices. For $d = 4$, if we label the edges as $(1, 2), \dots, (3, 4)$, then the weight vectors $(1, 1, 1, 0, 0, 0)$ and $(1, 1, 0, 1, 0, 0)$ have the same distributions of ones and zeros, but correspond to binary graphs which are not in the same orbit of Σ_4 . In particular, the region $\{(x_1, \dots, x_6) : x_1 \geq x_2 \geq \dots \geq x_6\}$ is not a fundamental domain for the action of Σ_4 .

While a fundamental domain is harder to find in high dimensions (see Section 4), the overall structure of $\mathcal{U}_d$ for general d is similar to the $d = 3$ case, with just increased notation.

THEOREM 2.4. The space of unlabeled graphs $\mathcal{U}_d = \mathcal{G}/\Sigma_d = \mathcal{O}_D/\Sigma_d$ is a stratified space.

The proof is sketched in Supplement A. By [29], $\mathbb{R}^D/\Sigma_d$ is PL or Lipschitz homeomorphic to $\mathbb{R}^{D-1} \times \mathbb{R}_{\geq 0}$, but the proof does not give a cell decomposition of $\mathbb{R}^D/\Sigma_d$, much less of $\mathcal{U}_d$. We do have information about the topology of $\mathcal{U}_d$: in Supplement A, we prove that $\mathcal{U}_d$ is contractible, and more surprisingly, that the natural slice of $\mathcal{U}_d$ given by the hyperplane $\sum_{i=1}^D x_i = 1$ is contractible. The practical implication of these results is that the usual topological invariants of $\mathcal{U}_d$ and its slice (the fundamental group, the homology/cohomology groups) provide no information.

¹Color versions of all figures are in Supplement E. In figures below, regions called red appear as white in black-and-white reproductions, and regions called blue appear as gray.




FIG. 2. As explained in Section 4.1, the infinite solid cone, which is the region $\{x \geq y \geq z \geq 0\}$, is a fundamental domain F for unlabeled networks with three nodes. With the convention that the bottom side of the triangle has weight x , the left-hand side has weight y , and the right-hand side has weight z , the network with edge weights 1, 2, 3 corresponds to the point P in the interior of the cone. Other networks shown are color coded to correspond to points on faces or edges of the cone.

3. Network averages and their asymptotic behavior. In this section, we define the mean of a distribution Q on the space of networks and investigate the asymptotic behavior of the empirical (or sample) mean network based on an i.i.d. sample of networks from Q . Statistical inference can be carried out based on the asymptotic distribution of the empirical mean. We illustrate with an example from hypothesis testing. The results of the previous section, characterizing the topology and geometry of the space of unlabeled networks, are essential for achieving our goals in this section.

3.1. Network averages through Fréchet means. Let Q be some distribution on a general metric space (M, ρ) . One can define the Fréchet function $f(p)$ on M as

$$(3.1) \quad f(p) = \int_M \rho^2(p, z) Q(dz) \quad (p \in M).$$

If f is finite on M and has a unique minimizer,

$$(3.2) \quad \mu = \operatorname{argmin}_p f(p),$$

then μ is called the *Fréchet mean* of Q (with respect to the metric ρ). Otherwise, the minimizers of the Fréchet function form a *Fréchet mean set* C_Q . Given an i.i.d. sample

$X_1, \dots, X_n \sim Q$ on M , the empirical Fréchet mean can be defined by replacing Q with the empirical distribution $Q_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\cdot)$, that is,

$$(3.3) \quad \mu_n = \operatorname{argmin}_p \frac{1}{n} \sum_{i=1}^n \rho^2(p, X_i).$$

When M is a manifold, one can equip M with a metric space structure through an embedding into some Euclidean space or employing a Riemannian structure of M . Respectively, ρ can be taken to be the Euclidean distance after embedding (extrinsic distance) or the geodesic distance (intrinsic distance), giving rise to extrinsic and intrinsic means. Asymptotic theory for extrinsic and intrinsic analysis has been developed in [6, 8, 9], and applied to many manifolds of interest (see, e.g., [5, 18]).

Now take $M = \mathcal{U}_d$, the space of unlabeled networks with d nodes, our space of interest, and let Q be a distribution on $\mathcal{U}_d$. Given an i.i.d. sample $X_1, \dots, X_n$ from Q , in order to define the Fréchet mean μ of Q and empirical Fréchet mean μ_n of Q_n , one needs an appropriate choice of distance on $\mathcal{U}_d$. Given the quotient space structure characterized in the previous section, that is, $\mathcal{U}_d = \mathcal{G}_d / \Sigma_d$, a natural choice for the distance ρ is the *Procrustean distance* d_P , where

$$(3.4) \quad d_P([\vec{x}], [\vec{y}]) := \min_{\sigma_1, \sigma_2 \in \Sigma_d} d_E(\sigma_1 \cdot \vec{x}, \sigma_2 \cdot \vec{y}),$$

for unlabeled networks $[\vec{x}], [\vec{y}] \in \mathcal{U}_d$, with $\vec{x}$ denoting the vectorized representation of a representative network x . We recall that $\mathcal{G}_d$ is the set of all labeled graphs with d vertices and Σ_d is the group of permutations of $\{1, 2, \dots, d\}$.

In order to carry out statistical inference based on μ_n , defined with respect to the distance (3.4), some natural and fundamental questions related to μ and μ_n need to be addressed, which we aim to do in the following subsections. Here are some of the most crucial ones:

1. (Consistency.) What are the consistency properties of the network empirical mean μ_n , that is, μ_n a consistent estimator of the population Fréchet mean μ ? Can we establish some notion of a law of large numbers for μ_n ?
2. (Uniqueness of Fréchet mean.) This question is concerned with establishing general conditions on Q for uniqueness of the Fréchet mean μ . In general, this a challenging task—indeed, the lack of general uniqueness conditions for Fréchet means is still one of the main hurdles for carrying out intrinsic analysis on manifolds [26]. To date, the most general results in the literature for generic manifolds [1] force the support of Q to be a small geodesic ball to guarantee uniqueness of the intrinsic Fréchet mean. We address this question for the space of unlabeled networks in Section 4.
3. (CLT.) Once conditions for uniqueness of μ are provided, the next key question is whether one can derive the limiting distribution for μ_n for purposes of statistical inference, for example, proving a central limit type of theorem for μ_n , which in turn might be used for hypothesis testing.

We first illustrate the difficult nature of these problems (in particular for question 2 above) through the example $C = \mathbb{R}^2 / \mathbb{Z}_4$ in Section 2, by explicitly constructing a distribution on C that has nonunique Fréchet means.

EXAMPLE 3.1 (Example $C = \mathbb{R}^2 / \mathbb{Z}_4$ continued). As in Example 2.1 in Section 2, the quotient $C = \mathbb{R}^2 / \mathbb{Z}_4$ is a cone. Working in polar coordinates and taking $F_0 = [0, \infty) \times [-\pi/4, \pi/4]$ to be a fundamental domain, we consider probability distributions of the form $\nu(r, \theta) = \frac{1}{Z} R(r) \chi(\theta)$, where $Z = \int_{F_0} R(r) \chi(\theta) dr d\theta$.




FIG. 3. Plot of the probability distribution $v(r, \theta)$ in (3.5), with $\alpha = 15$. The distribution peaks in the red region where v is large, and is small in the blue region. The left-hand side shows the plot on $\mathbb{R}^2$, and the right-hand side shows that plot on the cone.

We can explicitly compute the Fréchet function $f(x)$ with respect to v . For $x = (r, \theta) \in F_0$, $F_\theta = [0, \infty) \times [-\pi/4 + \theta, \pi/4 + \theta]$ is a fundamental domain. For $y \in F_\theta$, $d_P([x], [y]) = d_E(x, y)$. Then

$$f(x) = \int_{F_\theta} \|x - y\|^2 v(y) dy = \frac{1}{Z} (c_1 \chi_1 r^2 - 2c_2 \chi_2(\theta) r + c_3 \chi_1),$$

where $c_k = \int_0^\infty r^k R(r) dr$ for $k = 1, 2, 3$, $\chi_1(\theta) = \int_{\theta-\pi/4}^{\theta+\pi/4} \chi(t) dt$ and $\chi_2(\theta) = \int_{\theta-\pi/4}^{\theta+\pi/4} \chi(t) \cos(\theta - t) dt$.

The Fréchet mean occurs when $\partial f / \partial r = \partial f / \partial \theta = 0$, which is difficult to compute in general. Consider the special case

$$(3.5) \quad v(r, \theta) = \frac{1}{Z} \exp(-(r - \alpha)^2),$$

where α is a fixed constant, $\chi(\theta) \equiv 1$, $Z = (\pi^{3/2}/4)(1 + \text{erf}(\alpha))$; this distribution for $\alpha = 15$ is plotted in Figure 3. The minimum for this f occurs at

$$r_0 = \frac{\sqrt{2}(2\alpha + \sqrt{\pi}e^{\alpha^2}(2\alpha^2 + 1)(\text{erf}(\alpha) + 1))}{\pi^{3/2}e^{\alpha^2}\alpha(\text{erf}(\alpha) + 1) + \pi},$$

with θ arbitrary. When α is large, $r_0 \approx \frac{2\sqrt{2}}{\pi}\alpha$. For $\alpha = 15$, $r_0 \approx 13.5348$. This shows that v has a circle's worth of Fréchet means; the θ -independence of v implies θ -independence of the Fréchet means. One can see this in Figure 4 where the Fréchet function $f(x)$ is minimal and most blue on a circle of radius approximately 13.53 (corresponding to the red circle on the cone in Figure 3).

3.2. A strong law of large numbers. Before establishing the limiting distribution for μ_n , a natural first step is to explore the consistency properties of μ_n . Drawing on Theorem 3.3 in [6] for general metric spaces, we have the following result.

THEOREM 3.2. *Let Q be a distribution on $\mathcal{U}_d$, let C_Q be the set of means of Q with respect to the Procrustean distance d_P and let C_{Q_n} be the set of empirical means with respect to a sample of unlabeled networks $X_1, \dots, X_n$. Assume that the Fréchet function is everywhere finite. Then the following holds: (a) the Fréchet mean set C_Q is nonempty and*




FIG. 4. Plot of the Fréchet function $f(x)$ with $\alpha = 15$ on $\mathbb{R}^2$ and on the cone. The Fréchet mean occurs in the “most blue” region, on a circle of radius approximately 13.53.

compact; (b) for any $\epsilon > 0$, there exists a positive integer-valued random variable $N \equiv N(\epsilon)$ and a Q -null set $\Omega(\epsilon)$ such that

$$(3.6) \quad C_{Q_n} \subseteq C_Q^\epsilon \doteq \{p \in \mathcal{U}_d : d_P(p, C_Q) < \epsilon\} \quad \forall n \geq N$$

outside of $\Omega(\epsilon)$; (c) if C_Q is a singleton, that is, the Fréchet mean μ is unique, then μ_n converges to μ almost surely.

PROOF. We first prove that every closed and bounded subset of $M = \mathcal{U}_d$ is compact.

Let F be a fundamental domain for the action of Σ_d on $\mathcal{G}_d$, as defined in Definition 2.2, with the associated projection $q : F \rightarrow \mathcal{U}_d$. This map is continuous and a diffeomorphism on the interior of F . Take a closed and bounded set S in $\mathcal{U}_d$. Because q is continuous, $q^{-1}(S)$ is closed. We now show that $q^{-1}(S)$ is also bounded. S is contained in a ball centered at some $[\vec{z}] \in \mathcal{U}_d$ with radius r , so $d_P([\vec{z}], [\vec{x}]) < r$ for $[\vec{x}] \in S$. Now say that the largest entry in $[\vec{z}]$ (in any ordering of the entries of $\vec{z}$) is C . If the largest entry in $\vec{x}$ (in any ordering) is greater than $C + r$, then $d_P([\vec{z}], [\vec{x}]) > (C + r) - C$, a contradiction. (This holds because under any permutation σ of the entries of $\vec{x}$, one entry in $d_E(\vec{z}, \sigma \cdot \vec{x})$ is greater than $(C + r) - C$.) Thus for any choice of $\vec{x} \in [\vec{x}] \in S$, $\|\vec{x}\| \leq \sqrt{D}(C + r)$. Thus $q^{-1}(S)$ is contained in the ball of radius $\sqrt{D}(C + r)$ centered at the origin, and thus is bounded.

Since F is a closed subset of $\mathbb{R}^D$, the closed and bounded set $q^{-1}(S)$ is compact. Since q is continuous, S is compact.

Then by Theorem 3.3 in [6], (a) and (b) follow.

Part (c) follows from [40] under the uniqueness of Fréchet mean. $\square$

REMARK 3.3. For every model,

$$(3.7) \quad Q \in \mathcal{Q} = \{Q : C_Q \text{ is a singleton with a finite Fréchet function}\},$$

the sample Fréchet mean μ_n is a strongly consistent estimate of μ with respect to this model.

3.3. A central limit theorem. The goal of this section is to derive a central limit theorem for the empirical Fréchet mean, as an important precursor for statistical inference. One of the key challenges is to establish geometric conditions on distributions on $\mathcal{U}_d$ which ensure the uniqueness of the population Fréchet mean. We discuss and address the uniqueness issue in detail in Section 4. Here, our central limit theorem assumes that the uniqueness conditions of Section 4 are met.

Let $q : \mathcal{O}_D \rightarrow \mathcal{U}_d$ be the projection from the space of labeled networks to the space of unlabeled networks.

THEOREM 3.4. Assume Q' has support on a compact set $K' \subset \mathcal{O}_D$ defined in Theorem 4.5, so that the pushdown measure $Q = q_* Q'$ supported on $K = q(K')$ has a unique Fréchet mean μ . Let μ_n be the empirical Fréchet mean of an i.i.d. sample $X_1, \dots, X_n \sim Q$ with respect to the distance (3.4). Let $\phi = q^{-1}$. Then we have

$$(3.8) \quad \sqrt{n}(\phi(\mu_n) - \phi(\mu)) \xrightarrow{L} N(0, \Sigma),$$

where $\Sigma = \Lambda^{-1} C \Lambda^{-1}$ with the Hessian matrix

$$\Lambda = (\mathbb{E}[D_{r,s} \|\phi(\mu) - \phi(X_1)\|])_{r,s=1,\dots,D},$$

and C is the covariance matrix of $\{D_r \|\phi(\mu) - \phi(X_1)\|\}_{r=1,\dots,D}$.

Here, D_r denotes the partial derivative with respect to the r th direction, $D_{r,s}$ denotes second partial derivatives and $\xrightarrow{L}$ means convergence in law or distribution.

PROOF. Since μ_n converges to μ almost surely under our support condition for Q' , one can find a small open neighborhood U of μ inside K such as that $P(\mu_n \in U) \rightarrow 1$. Let $S = q^{-1}(U)$, which is an open subset of $\mathbb{R}^D$. Note that $\phi : U \rightarrow S$ is a homeomorphism. By the proof of Theorem 4.5, the projection map q is a Euclidean isometry. Therefore, for any vectorized network $\vec{x} \in S$ and $[\vec{z}] \in U$, one has

$$d_P^2([\vec{x}], [\vec{z}]) = d_P^2(\phi^{-1}(\vec{x}), [\vec{z}]) = \|\vec{x} - \vec{z}\|^2,$$

where $\vec{z} = \phi([\vec{z}])$. Thus $d_P^2(\phi^{-1}(\vec{x}), [\vec{z}])$ is twice differentiable in $\vec{x}$ for any $[\vec{z}] \in U$. Tracing through the definition of the smooth structure on U induced from the standard structure on $\mathbb{R}^D$, we see that $d_P^2([\vec{x}], [\vec{z}])$ is twice differentiable in $[\vec{x}]$. One can also verify the conditions (A5) and (A6) on the Hessian matrices of Theorem 2.2 in [7], and our theorem follows. $\square$

As an immediate consequence of this central limit theorem, we can define natural analogues of classical hypothesis tests. For example, consider the construction of a statistical test for two or more independent samples using the same framework. Assume that we have k independent sets of networks on d vertices, and consider the problem of testing whether or not these sets have in fact been drawn from the same population. Formally, we have independent samples $X_{ij} \sim Q_j$, for $i = 1, \dots, n_j$ and $j = 1, \dots, k$. Each of these k populations has an unknown mean, denoted $\mu^{(j)}$. Then, as a direct corollary to Theorem 3.4, we have the following asymptotic result.

COROLLARY 3.5. Assume that the distributions $Q_1, \dots, Q_k$ satisfy the conditions of Theorem 3.4. Moreover, also assume that $n_j/n \rightarrow p_j$ for every sample, with $n := \sum_j n_j$, and $0 < p_j < 1$. Then, under $H_0 : \phi(\mu^{(1)}) = \dots = \phi(\mu^{(k)})$, we have

$$T_k := \sum_{j=1}^k n_j (\phi(\mu_{j,n_j}) - \phi(\mu_n))' \widehat{\Xi}^{-1} (\phi(\mu_{j,n_j}) - \phi(\mu_n)) \longrightarrow \chi_{(k-1)D}^2,$$

where μ_{j,n_j} denotes the empirical mean of the j th sample, μ_n represents the grand empirical mean of the full sample, and $\widehat{\Xi} := \sum_{j=1}^k \widehat{\Xi}_j / n_j$ is a pooled estimate of covariance, with the $\widehat{\Xi}_j$'s denoting the individual covariance matrices estimates of each subsample.

As previously noted, this central limit theorem and such corollaries hold only if the population Fréchet mean(s) is unique. This depends crucially on the nontrivial geometry of the space of unlabeled networks. The following section deals exclusively with this issue.

4. Geometric requirements for uniqueness of the Fréchet mean. Underlying the central limit theorem in Theorem 3.4, the basic question is: which compact subsets K of $\mathcal{U}_d$ have a unique Fréchet mean? We have seen in Section 2 that $\mathcal{U}_d$ may be difficult to work with, while a fundamental domain $F \subset \mathcal{O}_D$ for the action of Σ_d on the space of labeled networks $\mathcal{O}_D$ seems more tractable. Indeed, finding the Fréchet mean for a distribution supported in F is a standard center of mass calculation in Euclidean space. However, it is not clear that this Fréchet mean in F projects to the Fréchet mean in the quotient space $\mathcal{U}_d = \mathcal{O}_D / \Sigma_d$, because the metric used to compute Fréchet means in $\mathcal{U}_d$ is the Procrustean distance, which may or may not equal the Euclidean distance.

In Section 4.1, we find a fundamental domain F by a standard procedure (Lemma 4.1), and find compact subsets $K' \subset F$ for which the Fréchet mean in $\mathcal{O}_d$ is guaranteed to project to the Fréchet mean of $K \subset \mathcal{U}_d$, where $K = q(K')$ is the projection of K' under the quotient/projection map $q : \mathcal{O}_D \rightarrow \mathcal{U}_d = \mathcal{O}_D / \Sigma_d$. This is the content of the main result in this section (Theorem 4.5). We also show that this result in our particular setting is an improvement of the best result for general Riemannian manifolds due to Afsari [1] (see Figure 7). In Section 4.2, we generalize Theorem 4.5 in two different directions: Theorem 4.7 allows distributions with small tails outside the compact set K above, and Theorem 4.8 shows that a unique Fréchet mean can be found for any compact set K after first embedding K isometrically into a larger Banach space.

4.1. The main result on the uniqueness of the Fréchet mean. We now discuss the construction of a fundamental domain F . By Definition 2.2, a fundamental domain is characterized by: (i) every weight vector x can be permuted by some $\sigma \in \Sigma_d$ to a network $\sigma \cdot x$ in F , (ii) if $\vec{w} \in \mathcal{O}_D$ has $\vec{w} \in F \cap \sigma F$ for $\sigma \in \Sigma_d$, $\sigma \neq \text{Id}$, then $\vec{w} \in \partial F$. (As a technical note, we always consider ∂F with respect to the induced topology on $\mathcal{O}_D$ from the standard topology on $\mathbb{R}^D$.) Once F has been constructed, we are guaranteed that the projection map $q : \mathcal{O}_D \rightarrow \mathcal{U}_d$, $q(\vec{w}) = [\vec{w}]$ restricts to a surjective map $q : F \rightarrow \mathcal{U}_d$ which is a homeomorphism from the interior F° of F to $q(F^\circ)$. (In fact, q is a diffeomorphism on this region by definition of the smooth structure on $q(F^\circ)$.)

It is convenient to center a choice of fundamental domain on a weight vector with trivial stabilizer.

DEFINITION 4.1. A vector $\vec{w} = (w_1, \dots, w_D) \in \mathcal{O}_D$ is *distinct* if it has trivial stabilizer for the action of Σ_d : that is, if $\sigma \cdot \vec{w} \neq \vec{w}$ for all $\sigma \in \Sigma_d$, $\sigma \neq \text{id}$.

A vector with trivial stabilizer is also called a vector with trivial automorphism group. Weight vectors with all different entries are distinct, which implies that the distinct vectors are dense in $\mathcal{O}_D$.

For example, consider the two graphs G_1 and G_2 in Figure 5. Both share the same connectivity pattern (i.e., are isomorphic), but have different weight vectors. The weight vector $\vec{w}^1 = (20, 20, 7, \dots)$ of G_1 satisfies $\sigma \cdot \vec{w}^1 = \vec{w}^1$ for $\sigma = (23) \in \Sigma_7$. In contrast, for all $\sigma \in \Sigma_7 \setminus \{\text{id}\}$, $\sigma \cdot \vec{w}^2 \neq \vec{w}^2$, where $\vec{w}^2$ is the weight vector of G_2 , because the two 20's belong to nodes with different valences. Thus even though $\vec{w}^1, \vec{w}^2$ have the same set of weights, $\vec{w}^2$ is distinct, while $\vec{w}^1$ is not.

We now explain a standard procedure to construct a fundamental domain as one region in the Voronoi diagram of the orbit of a distinct vector. (Geometers call this a Dirichlet domain.) Let d_P be the Procrustean distance on $\mathcal{O}_D / \Sigma_d$. From now on, we just write $\vec{w}$ instead of $[\vec{w}]$ for elements of $\mathcal{U}_d$.




FIG. 5. Two networks of identical connectivity with (G_1) and without (G_2) distinct weight vectors.

LEMMA 4.1. Fix a distinct vector $\vec{w} \in \mathcal{O}_D$. Set

$$\begin{aligned} F = F_{\vec{w}} &= \{\vec{w}' \in \mathcal{O}_D : d_E(\vec{w}, \vec{w}') \leq d_E(\vec{w}, \sigma \cdot \vec{w}'), \forall \sigma \in \Sigma_d\} \\ &= \{\vec{w}' \in \mathcal{O}_D : d_E(\vec{w}, \vec{w}') = d_P(\vec{w}, \vec{w}')\}. \end{aligned}$$

Then:

- (i) F is a fundamental domain for the action of Σ_d on $\mathcal{G}_d$.
- (ii) F is a solid cone with polyhedral cross section.

In the proof, we use the fact that $\vec{w}$ is distinct just below (4.1).

PROOF. (i) First, for fixed $\vec{w}_1$, a minimum of $d_E(\vec{w}, \sigma \cdot \vec{w}_1)$ is attained, since Σ_d is finite. Thus every network (characterized by its weight vector $\vec{w}_1$) has a permutation in F .

Second, we can rewrite F as

$$F = \{\vec{w}_1 \in \mathcal{O}_D : d_E(\vec{w}, \vec{w}_1) \leq d_E(\sigma \cdot \vec{w}, \vec{w}_1), \forall \sigma \in \Sigma_d\}.$$

Let $\vec{w}_1 \in F \cap \sigma F$. Then $\sigma^{-1} \vec{w}_1 \in F$, and

$$\begin{aligned} d_E(\vec{w}, \vec{w}_1) &= \min_{\tau} d_E(\vec{w}, \tau \vec{w}_1), d_E(\vec{w}, \sigma^{-1} \vec{w}_1) = \min_{\tau} d(\vec{w}, \tau \sigma^{-1} \vec{w}_1) \\ &= \min_{\tau} d(\vec{w}, \tau \vec{w}_1). \end{aligned}$$

Thus

$$(4.1) \quad d_E(\vec{w}, \vec{w}_1) = d_E(\vec{w}, \sigma^{-1} \vec{w}_1) = d_E(\sigma \vec{w}, \vec{w}_1),$$

so $\vec{w}_1$ is equidistant to $\vec{w}$ and $\sigma \vec{w}$. Since $\vec{w}$ is distinct, $\sigma \vec{w} \neq \vec{w}$ for $\sigma \neq \text{id}$. Thus $\vec{w}_1$ lies on a hyperplane in $\mathcal{O}_D$ defined by (4.1), and any ball around $\vec{w}'$ contains points that are closer to $\vec{w}$ than to $\sigma \vec{w}$, and points that are farther from $\vec{w}$ than from $\sigma \vec{w}$. Therefore, $\vec{w}_1 \in \partial F$.

(ii) We can construct F' for the action of Σ_d on all of $\mathbb{R}^D$ by taking the set of hyperplanes H_σ of points equidistant from $\vec{w}$ and $\sigma \vec{w}$ for $\sigma \in \Sigma_d$, and taking the connected component of $\mathcal{O}_D \setminus \bigcup_{\sigma} H_\sigma$ containing $\vec{w}$. Since these hyperplanes all pass through the origin, this component is a solid cone on the origin. The boundary is given by a union of hyperplanes, so the cross section is a polyhedron. (Not all hyperplanes contribute edges to the cross sectional polygon; see the next example.) Moreover, Σ_d preserves $\mathcal{O}_D$, so a fundamental domain in $\mathcal{O}_D$ is given by $F = F' \cap \mathcal{O}_D$. The boundary planes of F are either boundary planes of F' or (subspaces of) the boundary of $\mathcal{O}_D$. $\square$

This Dirichlet/Voronoi fundamental domain depends on a choice of $\vec{w}$. In particular, for a fixed distinct $\vec{w}_0$, we can guarantee that $\vec{w}_0$ is in the interior of F by setting $\vec{w} = \vec{w}_0$ in the lemma.

F is a solid cone cut out by at most $d! - 1 + D$ hyperplanes, where $d! - 1$ is the order of $\Sigma_d \setminus \{\text{Id}\}$ and D is the number of coordinate hyperplanes. Thus this construction of F is not very practical except in low dimensions. Looking back at Figure 2, the infinite solid cone is the fundamental domain for the distinct vector $(3, 2, 1)$. See Appendix A for an algorithm that computes the fundamental domain and examples for $d = 3, 4$.

We now give more information about F . The following result, although interesting, is not used below.

PROPOSITION 4.2. *Let F be the fundamental domain associated to a distinct vector $\vec{w} \in \mathcal{O}_D$. All distinct vectors in F have a representative in the interior F° of F .*

PROOF. See Supplement B. $\square$

REMARK 4.3. Any nondistinct vector $\vec{w}'$ has an arbitrarily close distinct vector $\vec{w}$. Then $d_P(\vec{w}, \vec{w}') = d_E(\vec{w}, \vec{w}')$, and this remains true for vectors close to $\vec{w}$. Therefore, $\vec{w}'$ is in the interior F° of $F = F_{\vec{w}}$. Thus for any nondistinct vector $\vec{w}'$, we can find a fundamental domain that contains $\vec{w}'$ in its interior.

The next lemma gives a sense of the minimal size of F . Let F_c denote the solid cone with vertex at the origin, axis $\vec{w}$ and cone angle c . Let $a = a_{\vec{w}}$ be the smallest angle between $\vec{w}$ and $\sigma \cdot \vec{w}$ for $\sigma \in \Sigma_d$, for $\sigma \neq \text{id}$. Of course, $a/2$ is the smallest angle between $\vec{w}$ and a hyperplane boundary of F .

LEMMA 4.2. *F contains the solid cone $F_{a/2}$.*

PROOF. We claim that for vectors $\vec{u}, \vec{\ell}, \vec{v}$, the angles formed by them satisfy

$$\angle(\vec{u}, \vec{\ell}) + \angle(\vec{\ell}, \vec{v}) \geq \angle(\vec{u}, \vec{v}).$$

To prove this, we may assume the $\vec{u}, \vec{v}, \vec{\ell}$ are unit vectors. Let $\vec{\ell}'$ be the projection of $\vec{\ell}$ into the plane of $\vec{u}$ and $\vec{v}$. Assume that $\angle(\vec{u}, \vec{\ell}') + \angle(\vec{\ell}', \vec{v}) = \angle(\vec{u}, \vec{v})$. Then $\cos(\angle(\vec{u}, \vec{\ell}')) = \vec{u} \cdot \vec{\ell}' / |\vec{\ell}'| \geq \vec{u} \cdot \vec{\ell}'$, and $\cos(\angle(\vec{u}, \vec{\ell})) = \vec{u} \cdot \vec{\ell} = \vec{u} \cdot \vec{\ell}' + \vec{u} \cdot (\vec{\ell} - \vec{\ell}') = \vec{u} \cdot \vec{\ell}'$, since $\vec{\ell} - \vec{\ell}'$ is normal to the $(\vec{u}, \vec{v})$ -plane. Thus $\angle(\vec{u}, \vec{\ell}') \leq \angle(\vec{u}, \vec{\ell})$. Similarly, $\angle(\vec{\ell}', \vec{v}) \leq \angle(\vec{\ell}, \vec{v})$. Thus $\angle(\vec{u}, \vec{\ell}) + \angle(\vec{\ell}, \vec{v}) \geq \angle(\vec{u}, \vec{\ell}') + \angle(\vec{\ell}', \vec{v}) = \angle(\vec{u}, \vec{v})$. The other possibility is that $\angle(\vec{u}, \vec{\ell}') - \angle(\vec{\ell}', \vec{v}) = \angle(\vec{u}, \vec{v})$ or the same with $\vec{u}$ and $\vec{v}$ switched. In this case, $\cos(\angle(\vec{u}, \vec{\ell})) = \vec{u} \cdot \vec{\ell}' \leq \cos(\angle(\vec{u}, \vec{\ell}'))$ implies $\angle(\vec{u}, \vec{\ell}) \geq \angle(\vec{u}, \vec{\ell}')$.

Thus for a fixed permutation σ , we have $\angle(\vec{w}, \vec{u}) + \angle(\vec{u}, \sigma \cdot \vec{w}) \geq \angle(\vec{w}, \sigma \cdot \vec{w}) \geq a$. Therefore,

$$(4.2) \quad \vec{u} \in F_{a/2} \implies \angle(\vec{u}, \sigma \cdot \vec{w}) \geq a/2.$$

We may assume that $|\vec{w}| = |\vec{u}|$. Because $\vec{u}$ and the $\sigma \cdot \vec{w}$ all lie on the sphere of radius $|\vec{w}|$, the distances between $\vec{u}, \vec{w}, \sigma \cdot \vec{w}$ are proportional to the angles they form with the origin. Thus (4.2) implies that $d^E(\vec{u}, \sigma \cdot \vec{w}) \geq d^E(\vec{u}, \vec{w})$, so $\vec{u} \in F$. $\square$

Although the topology of the interior F° and its image $q(F^\circ)$ are the same, their geometries are very different; this underlies the difference in general between the Fréchet means in F° and $q(F^\circ)$.

EXAMPLE 4.4 (Example $C = \mathbb{R}^2 / \mathbb{Z}_4$ continued). F° is the open first quadrant. Let $\vec{v}, \vec{\ell} \in F^\circ$ cut out angles α, β with the positive x -axis, respectively. By the law of cosines, the Euclidean distance between $\vec{v}, \vec{\ell}$ is less than the distance between $\vec{v}$ and $R_{\pi/2}\vec{\ell}$, the rotation

of $\vec{\ell}$ by $\pi/2$ radians counterclockwise, iff $|\alpha - \beta| < \pi/4$. Thus for $|\alpha - \beta| > \pi/4$, $d_E(\vec{v}, \vec{w}) < d_P([\vec{v}], [\vec{\ell}])$. Thus distances in F^o and C are different.

This affects the Fréchet means. Let $\nu = (4/3\pi)r dr d\theta$ be the uniform probability measure on F supported on $\{(r, \theta) \in [1, 2] \times [0, \pi/2]\}$. The Fréchet mean on F is the center of mass $(3/2, 3/2)$. The cone C has a circle action which rotates points equidistant from the vertex, and the Procrustean distance is clearly invariant under this action. This implies that if $[(r_0, \theta_0)]$ is a Fréchet mean on C , so is $[(r_0, \theta)]$ for all θ . Therefore, we can compute the Fréchet mean at $\theta = \pi/4$. Since $|(\pi/4) - \beta| \leq \pi/4$ in F , $d_E((r_0, \pi/4), (r, \theta)) = d_P((r_0, \pi/4), (r, \theta))$ for all $(r, \theta) \in F$. The previous computation gives $r_0 = 3/2$. We conclude that the Fréchet mean on C is the entire circle $r = 3/2$.

We now find a subcone of F such that the Fréchet mean of a compact convex set K inside this subcone projects to the unique Fréchet mean of the associated quotient space $K = q(K')$ inside $\mathcal{U}_d$. As explained at the beginning of this section, this allows us to derive a central limit theorem on K .

THEOREM 4.5. *Let $K \subset \mathcal{U}_d$ be such that there exists a compact convex set $K' \subset F_{a/4}$ with $q : K' \rightarrow K$ a homeomorphism. Then the Fréchet mean μ_K of K is unique and satisfies $q(\mu_{K'}) = \mu_K$.*

See Figure 6 for a schematic picture. This result is discussed without a full proof in [24, 25]. Note that $\mu(K)$ can be computed by finding the center of mass $\mu_{K'}$ of K' (with respect to the pullback of a distribution Q supported in K) by standard integrals, and then projecting to $\mathcal{U}_d$.

PROOF. Take $\vec{u}, \vec{v} \in F_{a/4}$ with $|\vec{u}| = |\vec{v}| = 1$. As in the previous lemma, $\angle(\vec{u}, \vec{w}) + \angle(\vec{w}, \vec{v}) \geq \angle(\vec{u}, \vec{v})$, with equality iff $\vec{u}, \vec{v}, \vec{w}$ are coplanar. Thus

$$\angle(\vec{u}, \vec{v}) \leq 2(a/4).$$

On the other hand, for $\sigma \neq \text{id}$, we have as above

$$\angle(\vec{u}, \sigma \cdot \vec{v}) + \angle(\sigma \cdot \vec{v}, \sigma \cdot \vec{w}) \geq \angle(\vec{u}, \sigma \cdot \vec{w}) \Rightarrow \angle(\vec{u}, \sigma \cdot \vec{v}) \geq \angle(\vec{u}, \sigma \cdot \vec{w}) - \angle(\vec{v}, \vec{w}).$$

Let the plane containing $\vec{0}, \vec{u}, \sigma \cdot \vec{w}$ intersect ∂F at a line containing the unit vector $\vec{z}$. Then

$$\angle(\vec{u}, \sigma \cdot \vec{w}) = \angle(\vec{u}, \vec{z}) + \angle(\vec{z}, \sigma \cdot \vec{w}) \geq (a/2 - a/4) + a/2 = 3a/4.$$

Therefore,

$$\angle(\vec{u}, \sigma \cdot \vec{v}) \geq 3a/4 - a/4 = a/2 \geq \angle(\vec{u}, \vec{v}).$$

Since $|\vec{u}| = |\vec{v}| = |\sigma \cdot \vec{v}| = 1$, the distance between these vectors is proportional to their angles. This implies that

$$d_E(\vec{u}, \vec{v}) = d_P([\vec{u}], [\vec{v}])$$

on $F_{a/4}$.

As a compact convex subset of Euclidean space, K' has a unique Fréchet mean. It follows that K is compact convex. $q : K' \rightarrow K$ is a homeomorphism, so for $[\vec{x}] \in K$, the Fréchet functions on K' and K satisfy

$$f([\vec{x}]) = \int_K d_P^2([\vec{x}], [\vec{y}]) q_* Q(d[\vec{y}]) = \int_{K'} d_E^2(\vec{x}, \vec{y}) Q(d\vec{y}) = f(\vec{x}).$$

Thus the Fréchet mean $\mu_{K'}$ of f on K' projects to μ_K , the unique Fréchet mean of f on K . $\square$




FIG. 6. $K' \subset F_{a/4}$, the blue cone. K and K' are homeomorphic via q . The Euclidean distance between points $\vec{x}, \vec{y} \in F_{a/4}$ is the same as the Procrustean distance between their orbits $[\vec{x}], [\vec{y}]$, so K' and K are actually isometric. The Fréchet means of K and K' are related by $\mu_K = q(\mu_{K'})$. In particular, the Fréchet mean of K is unique.

From Theorem 4.5, we derive the main result Theorem 3.4.

We now prove that Theorem 4.5 is a quantitative improvement over the optimal estimate for general Riemannian manifolds due to Afsari.

THEOREM 4.6 ([1], Theorem 2.1). *Let M be a complete Riemannian manifold with sectional curvatures at most Δ and injectivity radius ι_M . Set*

$$\rho_0 = \frac{1}{2} \min \left\{ \iota_M, \frac{\pi}{\Delta} \right\},$$

with the convention that $\frac{\pi}{\Delta} = \infty$ if $\Delta \leq 0$. For any $\rho < \rho_0$, a geodesic ball of radius ρ in M has a unique Fréchet mean.

The injectivity radius is the supremum of $r > 0$ such that every geodesic ball of radius r is a topological ball. Afsari's theorem does not apply to $\mathcal{U}_d$, which is not a manifold. We also cannot apply this theorem to the more tractable interior F_w^o of F_w , which is diffeomorphic to




FIG. 7. A cross-section of the fundamental domain $\mathbf{F}$ showing the improvement of Theorem 4.5 over the Afsari result. The Afsari work gives a unique Fréchet mean of the image in $\mathcal{U}_d$ to the ball of radius $r_*/2$, where r_* is the injectivity radius at $\vec{v}$. Theorem 4.5 gives a unique Fréchet mean to the image in $\mathcal{U}_d$ of the larger compact set K' inside the $a/4$ cone. The boundary of K' can extend down to $\vec{0}$, outwards as far as the walls of the $a/4$ cone, and to any finite height.

$\mathcal{U}_d$ minus a set of measure zero, because F_w^o is not complete in the Euclidean metric, and has zero injectivity radius.

However, we can compare Afsari’s result to Theorem 4.5 on (the locally complete) geodesic balls inside F_a , as these are ordinary Euclidean balls. The Euclidean metric has zero curvature, so $\pi/\Delta = \infty$ in our convention. Therefore, we need the injectivity radius of the smooth points of cone F_a . Take $\vec{v}$ lying on the cone axis. A ball $B_{\vec{v}}$ centered at $\vec{v}$ and tangent to the cone at a point P determines a right triangle $\Delta O\vec{v}P$. This ball has radius $|\vec{v}|\sin(a/2)$, so this is the injectivity radius $\iota_{B_{\vec{v}}}$ in Theorem 4.6. Thus Afsari’s theorem applies to a ball of half this radius, denoted $B_{\vec{v}}(|\vec{v}|\sin(a/2)/2)$.

For Theorem 4.5, we can take any compact set K' inside $F_{a/4}$. To show that this theorem improves the general Afsari result, we find such a K' containing $B_{\vec{v}}(|\vec{v}|\sin(a/2)/2)$. This follows if the $F_{a/4}$ cone contains the cone containing $B_{\vec{v}}(|\vec{v}|\sin(a/2)/2)$, which has cone angle $\sin^{-1}(\sin(a/2)/2)$. Since $\sin$ is increasing for $a \in (0, \pi/2)$, it suffices to show that

(4.3)
$$\sin(a/4) \geq \sin(a/2)/2.$$

This follows from $\sin(2\theta) = 2\sin(\theta)\cos(\theta) \leq 2\sin(\theta)$.

As a result, $K' = q(K) \subset \mathcal{U}_d$ has a unique Fréchet mean, even though its radius is larger than the bound in Theorem 4.6. This just says that the Afsari bound, which is universal for all Riemannian manifolds, may have an improvement on specific manifolds. Figure 7 illustrates this improvement.

4.2. *Generalizations of Theorem 4.5.* Theorem 4.5 can be improved in two directions. The first improvement, Theorem 4.7, is a technical analytic extension which shows that the uniqueness of the Fréchet mean still holds for smooth distributions that are close to compactly supported distributions supported in the $F_{a/4}$ cone. In practical terms, this means that distributions may have small tails outside the $F_{a/4}$ cone and still have a unique Fréchet mean. The second improvement, Theorem 4.8 shows that any compact subset $E \subset \mathcal{U}_d$ has a unique Fréchet mean after isometrically embedding E into the Banach space of bounded functions

on E in the sup norm. While this is more appealing, the new Fréchet mean may not lie in the image of E and may be hard to compute.

THEOREM 4.7. *Let Q be a smooth probability distribution with support inside a compact set $E \subset F_{a/4}$. There exists $\zeta > 0$ such that for every $Q' \in B_Q^3(\zeta) \cap C_c(F)$, the Fréchet function $f_{Q'}$ has a unique minimum.*

Here, $B_Q^3(\zeta)$ denotes a ball of radius ζ around Q in a Sobolev 3-norm, and $C_c(F)$ denotes continuous functions with compact support in F . The proof is in Supplement C.

Second, we prove that any compact subset E of $\mathcal{U}_d$ admits an isometric embedding into a normed space $\mathcal{B}(E)$ such that any smooth distribution supported in E has a unique Fréchet mean in $\mathcal{B}(E)$. This Fréchet mean may not lie in the image of E .

THEOREM 4.8. *Let Q be a continuous probability distribution with support in a compact subset E of $\mathcal{U}_d$. There is an isometric embedding ι of E into $\mathcal{B}(E)$, the space of bounded functions on E in the sup norm, such that the Fréchet function f_Q extends to a function $\bar{f}_Q$ on the closed convex hull of $\iota(E)$ and such that $\bar{f}_Q$ has a unique minimum.*

The proof is also in Supplement C.

We end this section with an example of the difficulties of handling a graph with a nondistinct weight vector.

EXAMPLE 4.9. As in Figure 2, take the distinct weight vector $\vec{w} = (3, 2, 1)$. The fundamental domain is $F = F_{\vec{w}} = \{x \geq y \geq z\}$. The proof of Corollary 4.2 shows that every $\vec{w}_1 \in F_{a/2}$ has a neighborhood U on which

$$(4.4) \quad d_P([\vec{w}_2], [\vec{w}_3]) = d_E(\vec{w}_2, \vec{w}_3) \quad \text{for all } \vec{w}_2, \vec{w}_3 \in U.$$

In fact, we can take $U = F_{a/2}$ (see Example B.1 for a stronger statement).

We now show that the vector $(1, 1, 1) \in \partial F$ has no neighborhood U in F on which (4.4) holds. Note that the stabilizer group of $(1, 1, 1)$ is all of Σ_3 .

Take $\vec{w}_2 = (1 + a_1, 1 + a_2, 1 + a_3)$ with $a_1 > a_2 > a_3 > 0$, and $\vec{w}_3 = (1 + b_1, 1 + b_2, 1 + b_3)$ with $b_1 > b_2 > b_3 > 0$. For small a 's and b 's, these vectors will be arbitrarily close to $(1, 1, 1)$. Then

$$d_E^2(\vec{w}_2, \vec{w}_3) = \sum_{i=1}^3 (a_i - b_i)^2.$$

Take

$$(4.5) \quad b_1 \approx a_2, b_2 \approx a_3 \approx b_3$$

and let $\sigma = (132)$. Then

$$(4.6) \quad d_P^2([\vec{w}_2], [\vec{w}_3]) < d_E^2(\vec{w}_2, \vec{w}_3)$$

if

$$(4.7) \quad \sum_{i=1}^3 (a_i - b_i)^2 > \sum_{i=1}^3 (a_i - b_{\sigma(i)})^2.$$

The left-hand side of (4.7) is approximately $(a_1 - a_2)^2 + (a_2 - a_3)^2$, and the right-hand side is approximately $(a_1 - a_3)^2$. Thus (4.7) holds if

$$(a_1 - a_2)^2 + (a_2 - a_3)^2 > (a_1 - a_3)^2, \quad \text{i.e. if } a_1^2 + 2a_1a_3 > a_2a_1 + a_2a_3.$$

Since $a_2^2 > a_2 a_3$, we just need

$$(4.8) \quad 2a_3 > a_2.$$

So once we choose the a 's and b 's to satisfy (4.5), (4.8), we obtain (4.6), which proves the claim.

We conclude that there is no neighborhood U of $(1, 1, 1)$ inside $\mathbb{R}^3$ (not just inside F) on which Euclidean distances in U agree with Procrustean distances in $q(U) \subset \mathcal{U}_3$. This illustrates the impossibility of applying [1] to the singular point $[(1, 1, 1)]$ of $\mathcal{U}_3$.

5. Discussion. Our work pertains to the geometric and statistical foundations of unlabeled networks. Specifically, we characterize the geometry of the space of such networks, define an appropriate notion of means or averages of unlabeled networks and derive the asymptotic behavior of the empirical average network. This last result is a necessary precursor for the development of a variety of statistical inference tools in analogy to those encountered in a typical “Statistics 101” course, as we demonstrate in the context of hypothesis testing. A key technical contribution of our work is that of providing broader conditions than available from general results on manifolds for uniqueness of the Fréchet mean network of a distribution on the space of unlabeled networks.

Our work here sets the stage for a program of additional research with multiple components. First, we expect that the asymptotic theory we develop can be extended in interesting ways. For example, it is unclear if the central limit theorem Theorem 3.4 is valid for either of our extensions (Theorems 4.7, 4.8) to the main Theorem 4.5 on uniqueness of the Fréchet mean. For Theorem 4.7, the obstacles to establishing a CLT are technical differentiability issues, but there are deeper conceptual difficulties in the setting of Theorem 4.8. More broadly, when the uniqueness condition fails, there is a possibility of establishing a limit theorem based on the set distance between the sample Fréchet means and the population Fréchet means, in the spirit of analogous work done for the estimation of level sets [13, 31]. However, the Hausdorff distance between the two does not necessarily go to zero, as in the counterexample in [8], so these limit theorems will be more subtle.

Additionally, while we have adopted Euclidean distance in our work here, there are other norms (cut distance, edit distance, etc. [30]) that merit consideration, and other Riemannian metrics may be more suitable for specific classes of networks (e.g., ego-networks). Both directions will require more sophisticated techniques than the current paper: for Riemannian metrics, we expect to need Rauch-type comparison theorems as in [16], Chapter 10, while for norms like the cut distance that do not come from Riemannian metrics, we need very different metric geometry methods as in [12].

In the particular context of ego-networks—in the case where the ego node is included, and hence (uniquely among the nodes) labeled—the existence of one labeled node reduces the symmetry group of the space of networks to the stabilizer subgroup of the node. The Euclidean metric does not exploit this reduction of the symmetry group, whereas other Riemannian metrics more closely reflect this reduction. For example, in the toy example where the ego-node is connected to only one other node, the hyperbolic metric of constant negative sectional curvature singles out the ego-node. However, if more realistically the ego-node has degree greater than one, then most Riemannian metrics invariant under the stabilizer group will have both positive and negative sectional curvatures. It would be very interesting to see how our main uniqueness result (Theorem 4.5) generalizes to these metrics.

It is also natural to expand our work to treat networks of different sizes. The space $\mathcal{U}_d = \mathcal{O}_D / \Sigma_d$ of unlabeled networks on d nodes isometrically embeds into $\mathcal{U}_{d+1}$ by adding to any labeled network in $\mathcal{O}_d$ a new node with zero weights to all other nodes, and embedding Σ_d into Σ_{d+1} as the stabilizer subgroup of the new node. We can then take the direct

limit $\lim_{d \rightarrow \infty} \mathcal{U}_d = \mathcal{U}_\infty$ as the space of all unlabeled networks. It is unclear if a CLT can be expected for $\mathcal{U}_\infty$, but it should be possible to produce a CLT for the space of unlabeled networks of size at most a fixed constant.

There is also work to be done in calculating the Fréchet mean(s) for a sample of unlabeled networks and applying our results in practice (e.g., comparing the means of two subpopulations). Calculation of the Fréchet mean is NP-hard, as a simple argument shows that a brute force approach requires $O((d!)^n)$ operations for n samples of networks with d nodes. This is reinforced by the complexity of the algorithms discussed in the Appendices on calculations of fundamental domains and Fréchet integrals, due to the geometric complexity of the Procrustean distance on the space $\mathcal{U}_d$ of unlabeled networks. Jain [25] provides an iterative algorithm for calculation of the Fréchet mean with complexity $O(n(d!))$, offering a major improvement. Nevertheless, for large networks, this can still be expected to scale poorly. Ultimately, however, in order to construct confidence regions and related hypothesis testing procedures, there is the computation of the covariance to consider, which to the best of our knowledge, has yet to be addressed.

Finally, although our work here assumed weighted undirected networks, it would be important to investigate how the finite set of undirected binary networks fits into our theory. In particular, the definition and uniqueness of Fréchet means needs to be elaborated. It is crucial to understand the placement of the binary graphs inside the fundamental domains we constructed; preliminary work indicates that binary graphs are unfortunately widely scattered throughout a fundamental domain.

APPENDIX A: COMPUTING FUNDAMENTAL DOMAINS

In the two Appendices, we discuss the computational difficulties of implementing the theory in Section 4. In this Appendix, we show that the fundamental domain $F_{\vec{w}}$ is highly sensitive to the choice of distinct vector $\vec{w}$ as axis.

We use a Dirichlet fundamental domain for the action of Σ_d on $\mathbb{R}_{\geq 0}^D$. By Lemma 4.1, $F = F_{\vec{w}} = \bigcap_{\sigma \in \Sigma_n} \{z \in \mathbb{R}_{\geq 0}^D : d_E(\vec{w}, \vec{z}) \leq d_E(\vec{w}, \sigma \vec{z})\}$ for a distinct vector $\vec{w} \in \mathbb{R}_{\geq 0}^D$. This is the intersection of $d! + D - 1$ half-spaces, where the D coordinate half-spaces are given by the inequalities $z_j \geq 0$. The cone F is a convex, noncompact polyhedral region in $\mathbb{R}_{\geq 0}^D$. The $d! - 1$ half-space regions are given by the linear inequalities:

$$\sum_{j=1}^D (w_{\sigma(j)} - w_j) z_j \leq 0,$$

for $\sigma \neq \text{Id}$.

Sage provides efficient tools for converting an input system of linear inequalities into a minimal description of the polyhedral output region; see Supplement D.

We consider the simplest nontrivial example of graphs with four vertices: $d = 4$, $D = \binom{d}{2} = 6$.

EXAMPLE A.1. Choose the distinct vector $\vec{w} = (1, 2, 3, 4, 5, 6)$. Sage gives F_w as the intersection of 7 half-spaces or the convex hull of 7 rays [see Figure 8].

EXAMPLE A.2. Now choose $\vec{w} = (1, 2, 3, 4, 5, 6.1)$. $F_{\vec{w}}$ is now described as the convex hull of 79 rays, or the intersection of 18 half-spaces [see Figure 9].

```

[sage: load('fund_decomp.sage')]
Please specify n: 4
[sage: l=[1,2,3,4,5,6]
[sage: F=fund_domain(l)[1]; F
A 6-dimensional polyhedron in QQ^6 defined as the convex hull of 1 vertex and 7 rays
[sage: F.n_Hrepresentation()
7
[sage: F.Hrepresentation()
(An inequality (1, 0, 0, 0, 0, 0) x + 0 >= 0,
An inequality (0, 1, 0, 0, 0, 0) x + 0 >= 0,
An inequality (0, 0, 1, 0, 0, 0) x + 0 >= 0,
An inequality (0, 0, 0, 1, 0, 0) x + 0 >= 0,
An inequality (0, -1, 1, -1, 1, 0) x + 0 >= 0,
An inequality (0, -1, -1, 1, 1, 0) x + 0 >= 0,
An inequality (-1, 1, 0, 0, -1, 1) x + 0 >= 0)
[sage: F.Vrepresentation()
(A vertex at (0, 0, 0, 0, 0, 0),
A ray in the direction (0, 1, 0, 0, 1, 0),
A ray in the direction (0, 0, 1, 1, 0, 0),
A ray in the direction (1, 0, 0, 0, 0, 1),
A ray in the direction (0, 0, 0, 0, 0, 1),
A ray in the direction (0, 0, 0, 1, 1, 1),
A ray in the direction (0, 0, 0, 0, 1, 1),
A ray in the direction (0, 0, 1, 0, 1, 1))

```

FIG. 8. Sage output for computation of a fundamental domain centered at the distinct vector $\vec{w} = (1, 2, 3, 4, 5, 6)$. $F.Hrepresentation()$ lists the half-spaces whose intersection is the fundamental domain. Note that we started with $4! + \binom{4}{2} - 1 = 29$ inequalities, and have narrowed it down to 7. $F.Vrepresentation()$ lists the 7 rays whose convex hull is the fundamental domain.

APPENDIX B: COMPUTATION OF THE FRÉCHET INTEGRAL

In this Appendix, we highlight the difficulty of computing the Fréchet integral $f([\vec{x}]) = \int_{\mathcal{U}_d} d_P^2([\vec{x}], [\vec{y}]) dQ'([\vec{y}])$ for all $[\vec{x}]$. (Here, $Q' = q_*Q$ in the notation of Theorem 4.5.) Even when a fundamental domain $F_{\vec{x}}$ has been explicitly determined for a fixed $\vec{x}$, so that $f([\vec{x}]) =$

```

[sage: load('fund_domain.sage')]
Please specify n: 4
[sage: l=[1,2,3,4,5,6,1]
[sage: F=fund_domain(l)[1]; F
A 6-dimensional polyhedron in QQ^6 defined as the convex hull of 1 vertex
and 79 rays
[sage: F.n_Hrepresentation()
18
[sage: F.n_Vrepresentation()
80
[sage: F.Vrepresentation()

(A vertex at (0, 0, 0, 0, 0, 0),
A ray in the direction (1, 0, 0, 0, 0, 1),
A ray in the direction (40861997369994326590067052289, 0, 0, 0, 371472703
3635828000796950588, 40861997369994326590067052289),
A ray in the direction (8763146482339086699034463885363380263257749, 0, 0
, 0, 796649680212641841246892697107500458824308, 8763146482339087671939163
170936968054434609),
A ray in the direction (780092677063528297159301852241, 0, 78009267706353
0115920689828791, 780092677063530115920689828791, 0, 780092677063528383766
986993981),
A ray in the direction (1, 0, 1, 1, 0, 1),
A ray in the direction (180451078715427611588064303400 18045107871542771

```

FIG. 9. A partial Sage output for computation of a fundamental domain centered at the distinct vector $\vec{w} = (1, 2, 3, 4, 5, 6, 1)$. $F.Hrepresentation()$ lists the half-spaces whose intersection is the fundamental domain. Note that we started with $4! + \binom{4}{2} - 1 = 29$ inequalities, and have narrowed it down to 18. $F.Vrepresentation()$ lists the 79 rays whose convex hull is the fundamental domain.

$\int_{F_{\vec{x}}} d_E^2(\vec{x}, \vec{y}) dQ(\vec{y})$, we have seen in Appendix A that the shape of $F_{\vec{x}}$ depends delicately on $[\vec{x}]$.

We can instead divide $\mathbb{R}_{\geq 0}^D$ into $d!$ regions $F_{\sigma}(\vec{x}) = \{\vec{y} \in \mathbb{R}_{\geq 0}^D : d_P([\vec{x}], [\vec{y}]) = d_E(\vec{x}, \sigma \cdot \vec{y})\}$. Note that $F_{\sigma}(\vec{x})$ is a fundamental domain for $\sigma \cdot \vec{x}$. Since $F_{\vec{w}} = \bigcup_{\sigma \in \Sigma_d} F_{\vec{w}} \cap F_{\sigma \cdot \vec{x}}$ for a fixed distinct vector $\vec{w}$, the Fréchet integral is given by

$$f(\vec{x}) = \int_{\mathcal{U}_d} d_P^2([\vec{x}], [\vec{y}]) dQ'([\vec{y}]) = \sum_{\sigma \in \Sigma_d} \int_{F_{\vec{w}} \cap F_{\sigma \cdot \vec{x}}} d_E^2(\vec{x}, \sigma \cdot \vec{y}) dQ(\vec{y}).$$

Although on the right-hand side we now have a sum of Euclidean integrals for each $\vec{x}$, as in Appendix A it is difficult to explicitly compute $F_{\vec{w}} \cap F_{\sigma \cdot \vec{x}}$.

We illustrate this computation in the simple case $d = 3$. Already for $d = 4$, the computation becomes too lengthy for inclusion here.

EXAMPLE B.1. $d = 3$. First, we show that the fundamental domain $F_{\vec{x}}$ can be chosen to depend only on the ordering $\text{ord}(\vec{x})$ (e.g., from largest to smallest, as in Figure 2) of the components of a distinct vector $\vec{x}$. The Procrustean distance between two points is $d_P([\vec{x}], [\vec{y}]) = \min_{\sigma \in \Sigma_3} d_E(\vec{x}, \sigma \cdot \vec{y})$. To minimize the Euclidean distance, we choose σ which reorders $\vec{y}$ to match the ordering of $\vec{x}$, as any other σ cannot decrease the distance. (This uses the special fact that $\Sigma_{d=3}$ is the full permutation group of the $D = 3$ set of weight vectors.) Therefore, we can choose $F = F_{\vec{x}} = \{\vec{y} \in \mathbb{R}_{\geq 0}^3 : d_E(\vec{x}, \vec{y}) = d_P([\vec{x}], [\vec{y}])\} = \{\vec{y} \in \mathbb{R}_{\geq 0}^3 : \text{ord}(\vec{y}) = \text{ord}(\vec{x})\}$ to be independent of the distinct vector $\vec{x}$. The Fréchet integral for a compactly supported probability measure Q' on $\mathcal{U}_3$ (or equivalently for Q on F) equals

$$\begin{aligned} f([\vec{x}]) &= \int_F d_E(\vec{x}, \vec{y})^2 dQ(\vec{y}) \\ &= \|\vec{x}\|^2 \int_F dQ(\vec{y}) - 2\vec{x} \cdot \int_F \vec{y} dQ(\vec{y}) + \int_F \|\vec{y}\|^2 dQ(\vec{y}) \\ &= \|\vec{x}\|^2 - 2\vec{x} \cdot \int_F \vec{y} dQ(\vec{y}) + B \\ &= \|\vec{x}\|^2 - 2 \sum_i C_i x^i + B, \end{aligned}$$

where the second integral on the last line is the dot product of a vector and a vector valued integral, $C_i = \int_F y^i dQ(\vec{y})$, and $B = \int_F \|\vec{y}\|^2 dQ(\vec{y})$. Thus f is quadratic in $\vec{x}$ on distinct vectors. Since the distinct vectors are dense in Euclidean space, f is quadratic on all vectors. Therefore, f is strictly convex. Since we are minimizing over a convex region, f has a unique global minimum. To explicitly compute it, note that F has eight strata in varying dimensions 0, 1, 2 and 3; these are labeled (1), ..., (8) in Table 1. The restriction to each stratum is smooth away from the lower dimensional boundaries, so we can simply minimize on each open piece to find eight local minima x^* .

Here, $C' = (1/3)(C_1 + C_2 + C_3)$. The true global minimum will depend on the values of C_1 , C_2 and C_3 .

Acknowledgments. We would like to thank the referees for very helpful suggestions.

This research was supported in part by ARO Grant W911NF-15-1-0440.

The first author was supported AFOSR Grant 12RSL042.

The second author was supported by NSF Grants IIS 1663870 and DMS Career 1654579 and DARPA Grant N66001-17-1-4041.

TABLE 1
Local minima x^ on the eight strata*

#	dim	Region	$x^* =$
(1)	3	$x_3 \geq 0$ and $0 < x_1 < x_2 < x_3$	(C_1, C_2, C_3)
(2)	2	$x_1 = 0$ and $0 < x_2 < x_3$	$(0, C_2, C_3)$
(3)	2	$0 < x_1 = x_2 < x_3$	$(\frac{1}{2}(C_1 + C_2), \frac{1}{2}(C_1 + C_2), C_3)$
(4)	2	$0 < x_1 < x_2 = x_3$	$(C_1, \frac{1}{2}(C_2 + C_3), \frac{1}{2}(C_2 + C_3))$
(5)	1	$x_1 = x_2 = 0$ and $x_3 > 0$	$(0, 0, C_3)$
(6)	1	$x_1 = 0$ and $0 < x_2 = x_3$	$(0, \frac{1}{2}(C_2 + C_3), \frac{1}{2}(C_2 + C_3))$
(7)	1	$0 < x_1 = x_2 = x_3$	(C', C', C')
(8)	0	$x_1 = x_2 = x_3 = 0$	$(0, 0, 0)$

SUPPLEMENTARY MATERIAL

All the Supplementary Material (see a brief description below) can be found in <https://github.com/KolaczykResearch/Ave-Unlab-Nets>.

Supplement A (DOI: [10.1214/19-AOS1820SUPPA](https://doi.org/10.1214/19-AOS1820SUPPA); .pdf). A description of $\mathcal{U}_3$, a sketch of Theorem 2.4, and results on the topology of $\mathcal{U}_d$.

Supplement B (DOI: [10.1214/19-AOS1820SUPPB](https://doi.org/10.1214/19-AOS1820SUPPB); .pdf). Proof of Proposition 4.2.

Supplement C (DOI: [10.1214/19-AOS1820SUPPC](https://doi.org/10.1214/19-AOS1820SUPPC); .pdf). Improvements to the main Theorem 4.5.

Supplement D (DOI: [10.1214/19-AOS1820SUPPD](https://doi.org/10.1214/19-AOS1820SUPPD); .zip). Sage code for computing fundamental domains.

Supplement E (DOI: [10.1214/19-AOS1820SUPPE](https://doi.org/10.1214/19-AOS1820SUPPE); .pdf). Color figures.

REFERENCES

[1] AFSARI, B. (2011). Riemannian L^p center of mass: Existence, uniqueness, and convexity. *Proc. Amer. Math. Soc.* **139** 655–673. [MR2736346 https://doi.org/10.1090/S0002-9939-2010-10541-5](https://doi.org/10.1090/S0002-9939-2010-10541-5)

[2] ARROYO RELIÓN, J. D., KESSLER, D., LEVINA, E. and TAYLOR, S. F. (2017). Network classification with applications to brain connectomics. Preprint. Available at [arXiv:1701.08140](https://arxiv.org/abs/1701.08140).

[3] AYDIN, B., PATAKI, G., WANG, H., BULLITT, E. and MARRON, J. S. (2009). A principal component analysis for trees. *Ann. Appl. Stat.* **3** 1597–1615. [MR2752149 https://doi.org/10.1214/09-AOAS263](https://doi.org/10.1214/09-AOAS263)

[4] BARDEN, D., LE, H. and OWEN, M. (2013). Central limit theorems for Fréchet means in the space of phylogenetic trees. *Electron. J. Probab.* **18** no. 25, 25. [MR3035753 https://doi.org/10.1214/EJP.v18-2201](https://doi.org/10.1214/EJP.v18-2201)

[5] BHATTACHARYA, A. (2008). Statistical analysis on manifolds: A nonparametric approach for inference on shape spaces. *Sankhyā* **70** 223–266. [MR2551814](https://doi.org/10.1017/CBO9781139094764)

[6] BHATTACHARYA, A. and BHATTACHARYA, R. (2012). *Nonparametric Inference on Manifolds: With Applications to Shape Spaces. Institute of Mathematical Statistics (IMS) Monographs* **2**. Cambridge Univ. Press, Cambridge. [MR2934285 https://doi.org/10.1017/CBO9781139094764](https://doi.org/10.1017/CBO9781139094764)

[7] BHATTACHARYA, R. and LIN, L. (2017). Omnibus CLTs for Fréchet means and nonparametric inference on non-Euclidean spaces. *Proc. Amer. Math. Soc.* **145** 413–428. [MR3565392 https://doi.org/10.1090/proc/13216](https://doi.org/10.1090/proc/13216)

[8] BHATTACHARYA, R. and PATRANGENARU, V. (2003). Large sample theory of intrinsic and extrinsic sample means on manifolds. I. *Ann. Statist.* **31** 1–29. [MR1962498 https://doi.org/10.1214/aos/1046294456](https://doi.org/10.1214/aos/1046294456)

[9] BHATTACHARYA, R. and PATRANGENARU, V. (2005). Large sample theory of intrinsic and extrinsic sample means on manifolds. II. *Ann. Statist.* **33** 1225–1259. [MR2195634 https://doi.org/10.1214/009053605000000093](https://doi.org/10.1214/009053605000000093)

[10] BILLERA, L. J., HOLMES, S. P. and VOGTMANN, K. (2001). Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.* **27** 733–767. [MR1867931 https://doi.org/10.1006/aama.2001.0759](https://doi.org/10.1006/aama.2001.0759)

- [11] BISWAL, B. B., MENNES, M., ZUO, X. N., GOHEL, S., KELLY, C., SMITH, S. M., BECKMANN, C. F., ADELSTEIN, J. S., BUCKNER, R. L. et al. (2010). Toward discovery science of human brain function. *Proc. Natl. Acad. Sci. USA* **107** 4734–4739.
- [12] BURAGO, D., BURAGO, Y. and IVANOV, S. (2001). *A Course in Metric Geometry*. *Graduate Studies in Mathematics* **33**. Amer. Math. Soc., Providence, RI. MR1835418 <https://doi.org/10.1090/gsm/033>
- [13] CHEN, Y.-C., GENOVESE, C. R. and WASSERMAN, L. (2017). Density level sets: Asymptotics, inference, and visualization. *J. Amer. Statist. Assoc.* **112** 1684–1696. MR3750891 <https://doi.org/10.1080/01621459.2016.1228536>
- [14] CORNEA, E., ZHU, H., KIM, P. and IBRAHIM, J. G. (2017). Regression models on Riemannian symmetric spaces. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 463–482. MR3611755 <https://doi.org/10.1111/rssb.12169>
- [15] COX, D. A. and KATZ, S. (1999). *Mirror Symmetry and Algebraic Geometry*. *Mathematical Surveys and Monographs* **68**. Amer. Math. Soc., Providence, RI. MR1677117 <https://doi.org/10.1090/surv/068>
- [16] DO CARMO, M. P. (1992). *Riemannian Geometry. Mathematics: Theory & Applications*. Birkhäuser, Inc., Boston, MA. MR1138207 <https://doi.org/10.1007/978-1-4757-2201-7>
- [17] DONALDSON, S. K. and KRONHEIMER, P. B. (1990). *The Geometry of Four-Manifolds*. *Oxford Mathematical Monographs*. Oxford University Press, New York. MR1079726
- [18] DRYDEN, I. L., KOLOYDENKO, A. and ZHOU, D. (2009). Non-Euclidean statistics for covariance matrices, with applications to diffusion tensor imaging. *Ann. Appl. Stat.* **3** 1102–1123. MR2750388 <https://doi.org/10.1214/09-AOAS249>
- [19] DURANTE, D. and DUNSON, D. B. (2014). Nonparametric Bayes dynamic modelling of relational data. *Biometrika* **101** 883–898. MR3286923 <https://doi.org/10.1093/biomet/asu040>
- [20] DURANTE, D., DUNSON, D. B. and VOGELSTEIN, J. T. (2017). Nonparametric Bayes modeling of populations of networks. *J. Amer. Statist. Assoc.* **112** 1516–1530. MR3750873 <https://doi.org/10.1080/01621459.2016.1219260>
- [21] GINESTET, C. E., LI, J., BALACHANDRAN, P., ROSENBERG, S. and KOLACZYK, E. D. (2017). Hypothesis testing for network data in functional neuroimaging. *Ann. Appl. Stat.* **11** 725–750. MR3693544 <https://doi.org/10.1214/16-AOAS1015>
- [22] GJOKA, M., KURANT, M., BUTTS, C. T. and MARKOPOULOU, A. (2010). Walking in Facebook: A case study of unbiased sampling of OSNs. In *Infocom, 2010 Proceedings IEEE* 1–9. IEEE, New York.
- [23] JACKSON, M. O. (2008). *Social and Economic Networks*. Princeton Univ. Press, Princeton, NJ. MR2435744
- [24] JAIN, B. J. (2016). Statistical graph space analysis. *Pattern Recognit.* **60** 802–812.
- [25] JAIN, B. J. (2016). On the geometry of graph spaces. *Discrete Appl. Math.* **214** 126–144. MR3546275 <https://doi.org/10.1016/j.dam.2016.06.027>
- [26] KENDALL, W. S. (1990). Probability, convexity, and harmonic maps with small image. I. Uniqueness and fine existence. *Proc. Lond. Math. Soc.* (3) **61** 371–406. MR1063050 <https://doi.org/10.1112/plms/s3-61.2.371>
- [27] KOLACZYK, E. D. (2009). *Statistical Analysis of Network Data: Methods and Models*. *Springer Series in Statistics*. Springer, New York. MR2724362 <https://doi.org/10.1007/978-0-387-88146-1>
- [28] KOLACZYK, E. D., LIN, L., ROSENBERG, S., WALTERS, J. and XU, J. (2019). Supplement to “Averages of unlabeled networks: Geometric characterization and asymptotic behavior.” <https://doi.org/10.1214/19-AOS1820SUPPA>, <https://doi.org/10.1214/19-AOS1820SUPPB>, <https://doi.org/10.1214/19-AOS1820SUPPC>, <https://doi.org/10.1214/19-AOS1820SUPPD>, <https://doi.org/10.1214/19-AOS1820SUPPE>.
- [29] LANGE, C. When is the underlying space of an orbifold a manifold with boundary? Preprint. Available at [arXiv:1509.0679](https://arxiv.org/abs/1509.0679).
- [30] LOVÁSZ, L. (2012). *Large Networks and Graph Limits*. *American Mathematical Society Colloquium Publications* **60**. Amer. Math. Soc., Providence, RI. MR3012035 <https://doi.org/10.1090/coll/060>
- [31] MASON, D. M. and POLONIK, W. (2009). Asymptotic normality of plug-in level set estimates. *Ann. Appl. Probab.* **19** 1108–1142. MR2537201 <https://doi.org/10.1214/08-AAP569>
- [32] MCAULEY, J. and LESKOVEC, J. (2014). Discovering social circles in ego networks. *ACM Trans. Knowl. Discov. Data* **8** 4.
- [33] MCCARTY, C. (2002). Structure in personal networks. *J. Soc. Struct.* **3** 20.
- [34] MCDUFF, D. and SALAMON, D. (2012). *J-Holomorphic Curves and Symplectic Topology*, 2nd ed. *American Mathematical Society Colloquium Publications* **52**. Amer. Math. Soc., Providence, RI. MR2954391
- [35] NEWMAN, M. E. J. (2010). *Networks: An Introduction*. Oxford Univ. Press, Oxford. MR2676073 <https://doi.org/10.1093/acprof:oso/9780199206650.001.0001>
- [36] TANG, R., KETCHA, M., VOGELSTEIN, J. T., PRIEBE, C. E. and SUSSMAN, D. L. (2016). Law of large graphs. Preprint. Available at [arXiv:1609.01672](https://arxiv.org/abs/1609.01672).

- [37] UGANDER, J., BACKSTROM, L. and KLEINBERG, J. (2013). Subgraph frequencies: Mapping the empirical and extremal geography of large graph collections. In *Proceedings of the 22nd International Conference on World Wide Web* 1307–1318. ACM, New York.
- [38] VAKIL, R. (2003). The moduli space of curves and its tautological ring. *Notices Amer. Math. Soc.* **50** 647–658. [MR1988577](#)
- [39] WANG, H. and MARRON, J. S. (2007). Object oriented data analysis: Sets of trees. *Ann. Statist.* **35** 1849–1873. [MR2363955](#) <https://doi.org/10.1214/009053607000000217>
- [40] ZIEZOLD, H. (1977). On expected figures and a strong law of large numbers for random elements in quasi-metric spaces. In *Transactions of the Seventh Prague Conference on Information Theory, Statistical Functions, Random Processes and of the Eighth European Meeting of Statisticians A* 591–602 (Tech. Univ. Prague, 1974). [MR0501230](#)