
Geometric Network Comparisons

Dena Marie Asta

Department of Engineering & Public Policy
Department of Statistics
Carnegie Mellon University
Pittsburgh, PA 15213
dasta@andrew.cmu.edu

Cosma Rohilla Shalizi

Department of Statistics
Carnegie Mellon University
Pittsburgh, PA 15213
cshalizi@stat.cmu.edu

Abstract

Network analysis needs tools to compare networks and assess the significance of differences between networks. We propose a principled statistical approach to network comparison that approximates networks as probability distributions on negatively curved manifolds. We outline the theory, as well as implement the approach on simulated networks, where its accuracy can be confirmed.

1 INTRODUCTION

Many scientific questions about networks amount to problems of *network comparison*: one wants to know whether networks observed at different times, or in different locations, or under different environmental or experimental conditions, actually differ in their structure. Such problems arise in neuroscience (e.g., comparing subjects with disease conditions to healthy controls, or the same subject before and after learning), in biology (e.g., comparing gene- or protein- interaction networks across species, developmental stages or cell types), and in social science (e.g., comparing different social relations within the same group, or comparing social groups which differ in some outcome). That the graphs being compared are not identical or even isomorphic is usually true, but scientifically unhelpful. What we need is a way to say if the difference between the graphs exceeds what we should expect from mere population variability or stochastic fluctuations. Network comparison, then, is a kind of two-sample testing, where we want to know whether the two samples could have come from the same source distribution. It is made challenging by the fact that the samples being compared are very structured, high-dimensional objects (networks), and more challenging because we often have only *one* graph in each sample.

We introduce a method for network comparison. The crucial idea is to approximate networks by continuous geo-

metric objects, namely probability densities, and then do two-sample bootstrap tests on those densities. Specifically, we draw on recent work showing how many real-world networks are naturally embedded in hyperbolic (negatively curved) manifolds. Graphs then correspond to clouds of points in hyperbolic space, and can be viewed as being generated by sampling from an underlying density on that space. We estimate a separate density for each of the two networks being compared, calculate the distance between those densities, and compare it to the distance expected under sampling from a pooled density estimate.

Our method, while conceptually fairly straightforward, is admittedly more complicated than the current practice in the scientific literature, which is to compare networks by taking differences in *ad hoc* descriptive statistics (e.g., average shortest path lengths, or degree distributions). It is very hard to assess the statistical significance of these differences, and counter-examples are known where the usual summary statistics fail to distinguish graphs which are qualitatively radically different (e.g., grid-like graphs from highly clustered tree-like ones). Similarly, whole-graph metrics and similarity measures are of little statistical use, without probability models to gauge their fluctuations. Below, we show through simulations that our method let us do network comparisons where (i) we can assess significance, (ii) power is high for qualitative differences, and (iii) when we detect differences, we also get some idea *how* the networks differ.

2 MOTIVATION AND BACKGROUND

A fundamental issue with network comparison, mentioned in the introduction, is that we often have *only* two networks to compare, and nonetheless need to make some assessment of statistical significance. This can obviously only be done by regarding the networks as being drawn from (one or more) probability models, and restricting the form of the model so that an observation of a single graph is informative about the underlying distribution. That is, we must restrict ourselves to network models which obey some sort

of law of large numbers or ergodic theorem within a single graph, or else we always have $n = 1$. As in any other testing problem, the better the alignment between the model's restrictions and actual properties of the graphs, the more efficiently the test will use the available information.

Salient properties of actual networks Over the last two decades, it has become clear that many networks encountered in the real world, whether natural or human-made, possess a number of mathematically striking properties (Newman, 2010). They have highly right-skewed degree distributions, they show the “small-world effect” of short average path lengths (growing only logarithmically with the number of nodes) but non-trivial transitivity of links, and high clusterability, often with a hierarchical arrangement of clusters. This is all a far cry from what is expected of conventional random graphs. While a large literature of parametric stochastic models has developed to try to account for these phenomena (Newman, 2010), there are few situations where a data analyst can *confidently* assert that one of these models is even approximately well-specified.

Current approaches to network comparison The typical approach in the literature is *ad hoc* comparison of common descriptive statistics on graphs (path lengths, clustering coefficients, etc.). These statistics are often misapplied, as in the numerous incorrect claims to have found “power law” or “scale-free” networks (Clauset et al., 2009), but that is not the fundamental issue. Even the recent authoritative review of, and advocacy for, the “connectomics” approach to neuroscience by Sporns (2010) takes this approach. Disturbingly, Henderson and Robinson (2011) show that, with commonly used choices of statistics and criteria, this approach cannot distinguish between complex, hierarchically-structured networks, and simple two-dimensional grids (such as a grid over the surface of the cortex).

More formally, Pao et al. (2011) study the power of tests based on such summaries to detect departures from the null hypothesis of completely independent and homogeneous edges (Erdos-Renyi graphs) in the direction of independent but heterogeneous edges. Their results were inconclusive, and neither their null nor the alternative models are plausible for real-world networks. Apart from this, essentially nothing is known about either the significance of such comparisons or their power, how to combine comparisons of different descriptive statistics, which statistics to use, or if significant differences are found, how to infer changes in structure from them. The issue of statistical significance also afflicts graph metrics and similarity measures, even those with plausible rationales in graph theory (e.g., that of Koutra et al. 2013).

Hunter et al. (2008) show one way to check goodness-of-fit for a model of a single network, using simulations to

check whether the observed values of various graph statistics are plausible under the model's sampling distribution. But they are unable to combine checks with different statistics, cannot find the power of such tests, and do not touch on differences across networks.

More relevantly to comparisons, Middendorf et al. (2005) use machine-learning techniques to classify networks as coming from one or another of various generative models, taking features of the network (such as the counts of small sub-graphs, or “motifs”) as the inputs to the classifier. They demonstrate good operating characteristics in simulations, but rely on having a good set of generative models to start with.

The approach to network comparison most similar to ours is Tang et al. (2014), which, like our proposed methods, models the nodes as drawn from densities on a latent space and attaches edges based on the geometric relationship between node coordinates. The primary difference between both approaches is the choice of latent space. Tang et al. (2014) use a Euclidean inner product space, allowing for an algebraic method of network inference. Our choice is motivated by the desire to pick a latent space that matches geometric properties of the real-world networks we aim to study.

A final related approach to network comparison is Rosvall and Bergstrom (2010), which like our proposed methods, uses bootstrap resampling from models fit to the original networks to assess significance of changes. The goal there however is not to detect global changes in the network structure, but local changes in which nodes are most closely tied to one another.

Hyperbolic geometry of networks While waiting for scientifically-grounded parametric models, we seek a class of non-parametric models which can accommodate the stylized facts of complex networks. Here we draw on the more recent observation that for many real-world networks, if we view them as metric spaces with distance given by shortest path lengths, the resulting geometry is *hyperbolic* (Albert et al., 2014, Kennedy et al., 2013, Krioukov et al., 2010), rather than Euclidean. Said another way, many real-world networks can be naturally embedded into negatively-curved continuous spaces. Indeed, Krioukov et al. (2010) show that if one draws points representing nodes according to a “quasi-uniform” distribution on the hyperbolic plane (see (2) below), and then connects nodes with a probability that decays according to the hyperbolic distance between the representative points, one naturally obtains graphs showing right-skewed degree distributions, short average path lengths, and high, hierarchical clusterability.

Continuous latent space models The model of (Krioukov et al., 2010) is an example of a *continuous latent*

space model, characterized by a metric space (M, ρ) , a link probability function W , and a probability density f on M , the *node density*. Points representing nodes are drawn iid from f , and edges form independently between nodes at x and y with probability $W(x, y) = W(\rho(x, y))$ decreasing in the distance. As a hierarchical model,

$$\begin{aligned} Z_i &\sim_{iid} f \\ A_{ij}|Z_1, \dots, Z_n &\sim_{ind} W(\rho(Z_i, Z_j)) \end{aligned} \quad (1)$$

where A_{ij} is the indicator variable for an edge between nodes i and j . Holding M, ρ, W fixed, but allowing f to vary, we obtain different distributions over graphs. Two densities f, g on M determine the same distribution over graphs if f is the image of g under some isometry of (M, ρ) . Note that node *densities* can be compared regardless of the number of nodes in the observed graphs.

The best-known continuous latent space model for social networks is that of Hoff et al. (2002), where the metric space is taken to be Euclidean and the density f is assumed to be Gaussian. Our general methodology for network comparison could certainly be used with such models. However, the striking properties of large real-world graphs, such as their highly-skewed degree distributions, lead us to favor the sort of hyperbolic model used by Krioukov et al. (2010), but without their restrictive assumptions on f . Rather, we will show how to non-parametrically estimate the node density from a single observed graph, and then reduce network comparison to a comparison of these probability densities.

Continuous latent space models are themselves special cases of models called *graphons*, lifting the restriction that M be a metric space, and requiring of the edge probability function $W(x, y)$ only that it be measurable and symmetric in its arguments¹. Any distribution over infinite graphs which is invariant under permuting the order of the nodes turns out to be a mixture of such graphons (Kallenberg, 2005, ch. 7). Moreover, as one considers larger and larger graphs, the properties of the observed graph uniquely identify the generating graphon (Diaconis and Janson, 2008); what almost comes to the same thing, the limit of a sequence of growing graphs is a graphon (Borgs et al., 2006, Lovász, 2012, Borgs et al., 2014). One might, then, try to use our approach to compare graphons with estimated f and W . While graphon estimation is known to be possible in principle (Bickel et al., 2011, Choi and Wolfe, 2014), there are no published, computationally feasible methods to do it. Moreover, we expect to gain power by tailoring our models to enforce salient network properties, as described above. Accordingly, we turn to some of the important as-

¹Graphons are often *defined* to have $M = [0, 1]$ and f Lebesgue measure. One can show that any graphon over another measure space or with another node density is equivalent to one of this form, i.e., generates the same distribution over infinite graphs (Kallenberg, 2005, ch. 7).




FIGURE 1: **Models of \mathbb{H}_2** A connected component of the hyperboloid $x_3^2 = 1 + x_1^2 + x_2^2$ (left), with the metric given by the shortest possible Minkowski length of a path between points along the surface, is isometric to the Poincaré half-plane (right) under a suitable non-Euclidean metric. The half-plane is tiled into regions of equal area with respect to the metric. (Images from Rocchini (2007), under a Creative Commons license.)

pects of hyperbolic geometry.

2.1 HYPERBOLIC SPACES

Hyperbolic spaces are metric spaces which are negatively curved — the angles in a triangle of geodesics sum to less than 180 degrees. The oldest example of such a space is the surface of (one sheet of) the hyperboloid, the surface of points $(x_1, x_2, x_3) \in \mathbb{R}^3$ such that

$$x_1^2 + x_2^2 - x_3^2 = -1,$$

with the distance between points taken to be the smallest possible Minkowski length of a path between them along the surface. Another, and perhaps even more basic, example of a hyperbolic space is a tree, again with the shortest-path metric. Our starting data will be observed networks, which are typically at least locally tree-like, and so also possess a hyperbolic geometry (Jonckheere et al., 2008).

As explained above, we aim to represent this discrete hyperbolic geometry with a density over a continuous hyperbolic space. For concreteness, we will focus on the hyperbolic plane \mathbb{H}_2 , whose most basic geometric model is just the surface of the hyperboloid. It will be more convenient to work with another model of \mathbb{H}_2 : the *Poincaré half-plane* of \mathbb{C} ,

$$\mathbb{H}_2 = \{x + iy \mid x \in \mathbb{R}, y \in (0, \infty)\}$$

equipped with the metric $d\rho^2 = (dx^2 + dy^2)/y^2$.

As mentioned above, (Krioukov et al., 2010) showed that if the density of nodes on the Poincaré half-plane is one of the *quasi-uniform* densities,

$$q_{\delta, R}(re^{i\theta}) = \frac{\delta \sinh \delta r}{2\pi(\sinh r) \cosh(\delta R - 1)}, \quad \delta > 0 \quad (2)$$

one obtains graphs which reproduce the stylized facts of right-skewed degree distributions, clusterability, etc., for




FIGURE 2: **Densities on \mathbb{H}_2** 1000 points drawn iidly from quasi-uniform densities, Eq. 2 (top; $\delta = 1, 10, 30$ from left to right, $R = 1$ throughout), and from hyperbolic Gaussian densities, Eq. 8 (bottom, $\sigma = 0.05, 0.1, 0.3$ from left to right).

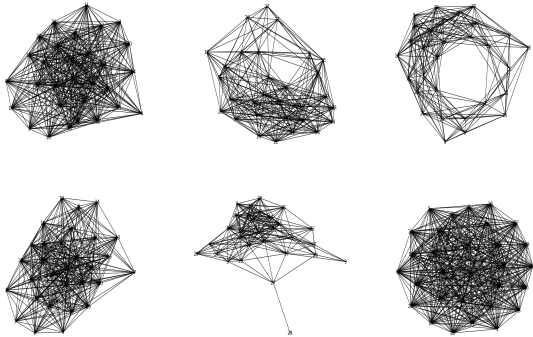


FIGURE 3: **Hyperbolic latent-space graphs** Graphs formed by drawing 30 node locations as in Fig. 2, and applying the link probability function $W(x, y) = \Theta(\rho(x, y) - 1.5)$. Note how the graphs in the bottom row become more clustered as the δ parameter increases from left to right.

a wide range of link probability functions W , including Heaviside step functions $\Theta(\rho - c)$. Note that the mode of q is always at $0 + i$, with the parameter $\delta > 0$ controlling the dispersion around the mode, and $R > 0$ being an over-all scale factor. As δ grows, the resulting graphs become more clustered.

We will introduce another family of densities on \mathbb{H}_2 , the hyperbolic *Gaussians*, in the next section.

Fig. 2 shows samples from quasi-uniform distributions on \mathbb{H}_2 , and Fig. 3 the resulting graphs. While we will use such networks as test cases, we emphasize that we will go beyond (2) to a fully nonparametric estimation of the node density.

3 METHOD

Our goal is to compare networks by comparing node densities. Our procedure for estimating node densities has in




FIGURE 4: *Schematic of network inference*

turn two steps (Figure 4): we embed the nodes of an observed network into \mathbb{H}_2 (§3.1), and then estimate a density from the embedded points (§3.2). We may then compare the observed difference between estimated node densities from two graphs to what would be expected if we observed two graphs drawn from a common node density (§3.3).

3.1 GRAPH EMBEDDING

An *embedding* of a graph G is a mapping of its nodes V_G to points into a continuous metric space (M, ρ) which preserves the structure of the graph, or tries to. Specifically, the distances between the representative points should match the shortest-path distances between the nodes, as nearly as possible. This is a multidimensional scaling problem, where typically one seeks the embedding $\phi : V_G \mapsto M$ minimizing

$$\sum_{(v,w) \in V_G^2} (\rho_G(v, w) - \rho(\phi(v), \phi(w)))^2, \quad (3)$$

where ρ_G is the shortest-path-length metric on V_G . Classically, when $M = \mathbb{R}^n$ and ρ is the Euclidean metric, the arg-min of (3) can be found by spectral decomposition of the matrix of $\rho_G(v, w)$ values (Hand et al., 2001, ch. 3).

Spectral decomposition does not however give the arg-min of (3) when $M = \mathbb{H}_2$ with the appropriate non-Euclidean metric. While the solution could be approximated by gradient descent (Cvetkovski and Crovella, 2011), we follow Begelfor and Werman (2005) in changing the problem slightly. They propose minimizing

$$\sum_{(v,w) \in V_G^2} (\cosh \rho_G(v, w) - \cosh \rho(\phi(v), \phi(w)))^2 \quad (4)$$

which can be done exactly via a spectral decomposition. Specifically, let $R_{ij} = \cosh \rho_G(i, j)$, whose leading eigenvector is u_1 and whose trailing eigenvectors are u_2 and u_3 . Then the i^{th} row of the matrix $(u_1 u_2 u_3)$ gives the \mathbb{H}_2 coordinates for node i . If R has one positive eigenvalue, exactly

2 negative eigenvalues, and all remaining eigenvalues vanish, this defines an exact isometric embedding (Begelfor and Werman, 2005).

We have not found a way of estimating the node density which avoids the initial step of embedding. Our method is, however, fairly indifferent as to *how* the nodes are embedded, so long as this is done well, and in a way which does not pre-judge the form of the node density.




FIGURE 5: **Re-embedded Generated Graphs** Results of embedding simulated graphs, formed as in Fig. 3, back into \mathbb{H}_2 . Comparison with Fig. 2 illustrates the fidelity of the embedding process.

3.2 DENSITY ESTIMATION

Having embedded the graph into \mathbb{H}_2 , we estimate the node density. Our procedure for doing so is more easily grasped by first reviewing the connections between kernel density estimation, convolution, and Fourier transforms in Euclidean space.

Kernel density estimation in Euclidean space as convolution In Euclidean space, kernel density estimation smooths out the empirical distribution by adding a little bit of noise around each observation. Given observations $z_1, z_2, \dots, z_n \in \mathbb{R}^p$, and a normalized kernel function K_h , the ordinary kernel density estimator $\hat{f}^{n,h}$ at a point $z \in \mathbb{R}^p$ is

$$\begin{aligned} \hat{f}^{n,h}(z) &= \frac{1}{n} \sum_{i=1}^n K_h(z - z_i) \\ &= \int_{\mathbb{R}^p} K_h(z - z') \left(\frac{1}{n} \sum_{i=1}^n \delta(z' - z_i) \right) dz' \\ &= \int_{\mathbb{R}^p} K_h(z - z') \hat{P}_n(dz') \\ &= (K_h * \hat{P}_n)(z) \end{aligned}$$

where the third line defines the empirical measure \hat{P}_n , and $*$ denotes convolution. In words, the kernel density estimate is the convolution of the empirical measure with the

kernel. Here the role of the kernel K_h is not so much to be a distribution over the Euclidean space, as a distribution over *translations* of the space: $K_h(z - z_i)$ is really the density at the translation mapping the data point z_i into the operating point z . As it happens, the group of translations of \mathbb{R}^p is also \mathbb{R}^p , but when we adapt to non-Euclidean spaces, this simplifying coincidence goes away.

Since, in Euclidean space, the Fourier transform \mathcal{F} converts convolutions into products (Stein and Weiss, 1971),

$$\mathcal{F}[\hat{f}^{n,h}](s) = \mathcal{F}[K_h](s) \mathcal{F}[\hat{P}_n](s)$$

This relation often greatly simplifies computing $\hat{f}^{n,h}$. It also lets us define the bandwidth h , through the relation $\mathcal{F}[K_h](s) = \mathcal{F}[K](hs)$.

It is well known that kernel density estimators on \mathbb{R}^p , with $h \rightarrow 0$ at the appropriate rate in n , are minimax-optimal in their L_2 risk (van der Vaart, 1998). With suitable modifications, this still holds for compact manifolds (Pelletier, 2005), but the hyperbolic plane \mathbb{H}_2 is not compact.

3.2.1 \mathbb{H}_2 -Kernel Density Estimator

Our method for density estimation on \mathbb{H}_2 is a generalization of Euclidean kernel density estimation. In \mathbb{R}^p , the kernel is a density on translations of \mathbb{R}^p . For \mathbb{H}_2 , the appropriate set of isometric transformations are not translations, but rather the class of ‘‘Möbius transformations’’ represented by the Lie group $\mathbb{S}\mathbb{L}_2$ (Terras, 1985, Huckemann et al., 2010). An \mathbb{H}_2 kernel, then, is a probability density on $\mathbb{S}\mathbb{L}_2$. We may write $K_h(z, z_i)$ to abbreviate the density the kernel K_h assigns to the Möbius transform taking z_i to z . The generalized kernel density estimator on \mathbb{H}_2 takes the form

$$\begin{aligned} \hat{f}^{n,h}(z) &= \frac{1}{n} \sum_{i=1}^n K_h(z, z_i) \\ &= (K_h * \hat{P}_n)(z) \end{aligned} \quad (5)$$

In Euclidean space, the Fourier transform analyzes functions (or generalized functions, like \hat{P}_n) into linear combinations of the eigenfunctions of the Laplacian operator. The corresponding operation for \mathbb{H}_2 is the *Helgason*, or *Helgason-Fourier*, transform \mathcal{H} (Terras, 1985). The Fourier basis functions are indexed by \mathbb{R}^p , which is the group of translations; for analogous reasons, the Helgason basis functions are indexed by $\mathbb{C} \times \mathbb{S}\mathbb{O}_2$. Many of the formal properties of the Fourier transform carry over to the Helgason transform. (See App. A.) In particular, convolution still turns into multiplication:

$$\mathcal{H}[\hat{f}^{n,h}] = \mathcal{H}[K_h] \mathcal{H}[\hat{P}_n], \quad (7)$$

where $\mathcal{H}[K_h]$ denotes the Helgason-Fourier transform of the well-defined density on \mathbb{H}_2 induced by the density K_h

on $\mathbb{S}\mathbb{L}_2$, and we define the bandwidth h through

$$\mathcal{H}[K_h](s, M) = \mathcal{H}[K](hs, M).$$

As in Euclidean density estimation, h may be set through cross-validation.

In a separate manuscript (Asta, 2014), we show that the L_2 risk of (5) goes to zero at the minimax-optimal rate, under mild assumptions on the smoothness of the true density, and of the kernel K . (This is a special case of broader results about generalized kernel density estimation on symmetric spaces.) The assumptions on the kernel are satisfied by what Huckemann et al. (2010) calls ‘‘hyperbolic Gaussians’’, densities on \mathbb{H}_2 with parameter ρ defined through their Helgason transforms,

$$\mathcal{H}[K](s, M) \propto e^{\rho s(s-1)}. \quad (8)$$

Just as the ordinary Gaussian density is the unique solution to the heat equation with a point source in Euclidean space, the hyperbolic Gaussian is the unique ($\mathbb{S}\mathbb{O}_2$ -invariant) solution to the heat equation on \mathbb{H}_2 (Terras, 1985).

3.3 NETWORK COMPARISON

Combining embedding with kernel density estimation in \mathbb{H}_2 gives us a method of estimating node densities, and so of estimating a hyperbolic latent space model for a given network. We now turn to *comparing* networks, by comparing these estimated node densities.

Our method follows the general strategy advocated in Genovese et al. (2013). Given two graphs G_1 and G_2 , we may estimate two separate network models

$$\hat{\mathcal{P}}_1 = \hat{\mathcal{P}}(G_1), \quad \hat{\mathcal{P}}_2 = \hat{\mathcal{P}}(G_2).$$

We may also pool the data from the two graphs to estimate a common model

$$\hat{\mathcal{P}}_{12} = \hat{\mathcal{P}}(G_1, G_2).$$

We calculate a distance $d^* = d(\hat{\mathcal{P}}_1, \hat{\mathcal{P}}_2)$ using any suitable divergence. We then compare d^* to the distribution of distances which may be expected under the pooled model $\hat{\mathcal{P}}_{12}$. To do so, we independently generate $G'_1, G'_2 \sim \hat{\mathcal{P}}_{12}$, and calculate

$$d(\hat{\mathcal{P}}(G'_1), \hat{\mathcal{P}}(G'_2)).$$

That is, we bootstrap two independent graphs out of the pooled model, fit a model to each bootstrapped graph, and calculate the distance between them. Repeated over many bootstrap replicates, we obtain the sampling distribution of d under the null hypothesis that G_1 and G_2 are drawn from the same source, and any differences between them are due to population variability or stochastic fluctuations.²

²This method extends easily to comparing sets of graphs, $G_{11}, G_{12}, \dots, G_{1n}$ vs. $G_{21}, G_{22}, \dots, G_{2m}$, but the notation grows cumbersome.

In our case, we have already explained how to find $\hat{\mathcal{P}}_1$ and $\hat{\mathcal{P}}_2$. Since we hold the latent space M fixed at \mathbb{H}_2 , and the link probability function W fixed, we can label our models by their node densities, $\hat{f}_1^{n,h}$ and $\hat{f}_2^{n,h}$. To obtain the pooled model $\hat{\mathcal{P}}_{12}$, we first embed G_1 and G_2 separately using generalized multidimensional scaling, and then do kernel density estimation on the union of their embedded points.

The generalized multidimensional scaling technique we use depends only on the eigendecomposition of matrices determined by shortest path lengths. Therefore the L_2 difference

$$\|\hat{f}_1^{n,h} - \hat{f}_2^{n,h}\|_2 \quad (9)$$

between two estimated node densities $\hat{f}_1^{n,h}, \hat{f}_2^{n,h}$ is 0 if and only if the original sets of vertices from the different samples are isometric and hence (9) approximates a well-defined metric d on our continuous latent space models. Moreover, since the Plancherel identity carries over to the Helgason-Fourier transform (Terras, 1985),

$$d_2(f_1, f_2) = \|\mathcal{H}[f_1] - \mathcal{H}[f_2]\|_2, \quad (10)$$

and, for our estimated node densities, $\mathcal{H}[f]$ is given by (7). Appendix B gives full details on our procedure for computing the test statistic (10).

3.4 THEORETICAL CONSIDERATIONS

Let us sum up our method, before turning to theoretical considerations. (0) We observe two graphs, G_1 and G_2 . (1) Through multi-dimensional scaling, we embed them separately in \mathbb{H}_2 (§3.1), getting two point clouds, say \mathbf{Z}_1 and \mathbf{Z}_2 . (2) From each cloud, we estimate a probability density on \mathbb{H}_2 , using hyperbolic Gaussian kernels, getting \hat{f}^{n_1, h_1} and \hat{f}^{n_2, h_2} (§3.2). We calculate $\|\hat{f}^{n_1, h_1} - \hat{f}^{n_2, h_2}\|_2$ using (10). We also form a third density estimate, $\hat{f}^{n_1+n_2, h_{12}}$, from $\mathbf{Z}_1 \cup \mathbf{Z}_2$. (3) We generate two independent graphs G_1^*, G_2^* from $\hat{f}^{n_1+n_2, h_{12}}$ according to (1), and subject these graphs to re-embedding and density estimation, obtaining \hat{f}^{n_1, h_1^*} and \hat{f}^{n_2, h_2^*} and so $\|\hat{f}^{n_1, h_1^*} - \hat{f}^{n_2, h_2^*}\|_2$. Finally, (4) repeating step (3) many times gives us the sampling distribution of the test statistic under the null hypothesis that G_1 and G_2 came from the same source, and the p -value is the quantile of $\|\hat{f}^{n_1, h_1} - \hat{f}^{n_2, h_2}\|_2$ in this distribution.

The final step of computing the p -value is a fairly unproblematic bootstrap test. The previous step of generating new graphs from the pooled model is also an unproblematic example of a model-based bootstrap. The kernel density estimates themselves are consistent, and indeed converge at the minimax rate (Asta, 2014), *given* the point clouds on the hyperbolic plane. This makes it seem that the key step is the initial embedding. Certainly, it would be convenient if the graphs G_1 and G_2 were generated by a hyperbolic




FIGURE 6: **Comparing Quasi-Uniforms** Power of our test, at size $\alpha = 0.1$, for detecting the difference between a 100-node graph generated from the quasi-uniform density $q_{1,1}$ and a 100-node graph generated from $q_{\delta,1}$, as a function of the dispersion parameter δ .

latent space model, and the embedding was a consistent estimator of the latent node locations. However, such strong conditions are not *necessary*. Suppose that if $G_1 \sim \mathcal{P}_1$ and $G_2 \sim \mathcal{P}_2 \neq \mathcal{P}_1$, then $\hat{f}_1^{n,h} \rightarrow f_1$ and $\hat{f}_2^{n,h} \rightarrow f_2$, with $\|f_1 - f_2\|_2 > 0$. Then at any nominal size (significance level) $\alpha > 0$, the power of the test will go to 1. For the nominal size of the test to match the actual size (probability of incorrectly rejecting the null hypothesis), however, will presumably require a closer alignment between the hyperbolic latent space model and the actual generating distribution.

4 SIMULATIONS

Comparison of Graphs with Quasi-Uniform Node Densities In our first set of simulation studies, we generated graphs which exactly conformed to the hyperbolic latent space model, and in fact ones where the node density was quasi-uniform (as in Fig. 3). One graph had 100 nodes, with latent locations drawn from a $q_{1,1}$ distribution; the other, also of 100 nodes, followed a $q_{\delta,1}$ distribution, with varying δ . We used 50 bootstrap replicates (pairs of resampled networks) in each test, kept the nominal size $\alpha = 0.1$, and calculated power by averaging over 25 independent graph pairs (the number of power tests). Despite the graphs having only 100 nodes, Fig. 6 shows that our test has quite respectable power.

Comparison of Watts-Strogatz Graphs We have explained above, §2, why we expect hyperbolic latent space models to be reasonable ways of summarizing the structure of complex networks. However, they will also be more or less mis-specified for many networks of interest. We thus applied our methods to a class of graph distributions which do *not* follow a hyperbolic latent space model, namely Watts-Strogatz networks (Watts and Strogatz, 1998). Our

simulations used 100 node networks, with the base topology being a 1D ring with a branching factor of 40, and variable re-wiring probabilities. These graphs show the small-world property and high transitivity, but light-tailed degree distributions. Even in these cases, where the hyperbolic model is not the true generator, our comparison method had almost perfect power (Fig. 7).




FIGURE 7: **Comparing Watts-Strogatz models** Above, Watts-Strogatz graphs formed by re-wiring 1D ring lattices (85 nodes, branching factor 40) with probability p per edge; from left to right $p = 0.1, 0.2, 0.3$. Below, embeddings of the graphs into \mathbb{H}_2 . At nominal $\alpha = 0.1$, the power to detect these differences in p was 1.0 to within Monte Carlo error.

5 CONCLUSIONS

We have shown how nonparametric hyperbolic latent space models let us compare the global structures of networks. Our approach has its limits, and it may work poorly when the networks being compared are very far from hyperbolic. However, our experiments with Watts-Strogatz graphs show that it can detect differences among graph distributions from outside our model class. When we do detect a change in structure, we have a model for each network, namely their node densities, and the difference in node densities is an interpretable summary of how the networks differ. Many important directions for future work are now open. One important direction is a better handling of sparse networks, network growth, and the comparison of networks of different sizes — perhaps through some size-dependent modification of the link-probability function W , as in Krioukov et al. (2010), or the sort of scaling of graphons introduced in Borgs et al. (2014). But this should only extend our method’s scope.

Acknowledgements

Our work was supported by NSF grant DMS-1418124, NIH grant R01 NS047493, and an NSF Graduate Research Fellowship under grant DGE-1252522. We are grateful for valuable discussions with Carl Bergstrom, Elizabeth Casman, David Choi, Aaron Clauset, Steve Fienberg,

Christopher Genovese, Dmitri Krioukov, Alessandro Rinaldo, Mitch Small, Andrew Thomas, Larry Wasserman, and Chris Wiggins, and for feedback from seminar audiences at UCLA's Institute for Pure and Applied Mathematics and CMU's machine learning and social science seminar.

References

- Réka Albert, Bhaskar DasGupta, and Nasim Mobasher. Topological implications of negative curvature for biological and social networks. *Physical Review E*, 89: 032811, 2014. doi: 10.1103/PhysRevE.89.032811. URL <http://arxiv.org/abs/1403.1228>.
- Dena Asta. Kernel density estimation on symmetric spaces. arxiv:1411.4040, 2014. URL <http://arxiv.org/abs/1411.4040>.
- Evgeni Begelfor and Michael Werman. The world is not always flat, or, learning curved manifolds. Technical Report HUJI-CSE-LTR-2006-191, School of Engineering and Computer Science, Hebrew University of Jerusalem, 2005. URL <http://www.cs.huji.ac.il/~werman/Papers/cmds.pdf>.
- Peter J. Bickel, Aiyu Chen, and Elizaveta Levina. The method of moments and degree distributions for network models. *Annals of Statistics*, 39:38–59, 2011. URL <http://arxiv.org/abs/1202.5101>.
- Christian Borgs, Jennifer T. Chayes, László Lovász, Vera T. Sós, Balázs Szegedy, and Katalin Vesztegombi. Graph limits and parameter testing. In *Proceedings of the 38th Annual ACM Symposium on the Theory of Computing [STOC 2006]*, pages 261–270, New York, 2006. ACM. URL <http://research.microsoft.com/en-us/people/jchayes/Papers/TestStoc.pdf>.
- Christian Borgs, Jennifer T. Chayes, Henry Cohn, and Yufei Zhao. An L^p theory of sparse graph convergence I: Limits, sparse random graph models, and power law distributions. arxiv:1401.2906, 2014. URL <http://arxiv.org/abs/1401.2906>.
- David S. Choi and Patrick J. Wolfe. Co-clustering separately exchangeable network data. *Annals of Statistics*, 42:29–63, 2014. doi: 10.1214/13-AOS1173. URL <http://arxiv.org/abs/1212.4093>.
- Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51:661–703, 2009. URL <http://arxiv.org/abs/0706.1062>.
- Andrej Cvetkovski and Mark Crovella. Multidimensional scaling in the Poincaré disk. E-print, 2011. URL <http://arxiv.org/abs/1105.5332>.
- Persi Diaconis and Svante Janson. Graph limits and exchangeable random graphs. *Rendiconti di Matematica e delle sue Applicazioni*, 28:33–61, 2008. URL <http://arxiv.org/abs/0712.2749>.
- Christopher Genovese, Cosma Rohilla Shalizi, and Andrew C. Thomas. Network comparisons. Manuscript in preparation, 2013.
- David Hand, Heikki Mannila, and Padhraic Smyth. *Principles of Data Mining*. MIT Press, Cambridge, Massachusetts, 2001.
- J. A. Henderson and P. A. Robinson. Geometric effects on complex network structure in the cortex. *Physical Review Letters*, 107:018102, 2011. doi: 10.1103/PhysRevLett.107.018102.
- Peter D. Hoff, Adrian E. Raftery, and Mark S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97:1090–1098, 2002. URL <http://www.stat.washington.edu/research/reports/2001/tr399.pdf>.
- Stephan F. Huckemann, Peter T. Kim, Ja-Yong Koo, and Axel Munk. Möbius deconvolution on the hyperbolic plane with application to impedance density estimation. *Annals of Statistics*, 38:2465–2498, 2010. doi: 10.1214/09-AOS783. URL <http://arxiv.org/abs/1010.4202>.
- David R. Hunter, Steven M. Goodreau, and Mark S. Handcock. Goodness of fit of social network models. *Journal of the American Statistical Association*, 103:248–258, 2008. doi: 10.1198/016214507000000446. URL <http://www.csss.washington.edu/Papers/wp47.pdf>.
- Edmond Jonckheere, Poonsuk Lohsoonthorn, and Francis Bonahon. Scaled Gromov hyperbolic graphs. *Journal of Graph Theory*, pages 157–180, 2008. doi: 10.1002/jgt.20275. URL http://eudoxus2.usc.edu/jgt6396_final.pdf.
- Olav Kallenberg. *Probabilistic Symmetries and Invariance Principles*. Springer-Verlag, New York, 2005.
- W. Sean Kennedy, Onuttom Narayan, and Iraj Saniee. On the hyperbolicity of large-scale networks. arxiv:1307.0031, 2013. URL <http://arxiv.org/abs/1307.0031>.
- Danai Koutra, Joshua T. Vogelstein, and Christos Faloutsos. DELTACON: A principled massive-graph similarity function. In Joydeep Ghosh, Zoran Obradovic, Jennifer Dy, Zhi-Hua Zhou, Chandrika Kamath, and Srinivasan Parthasarathy, editors, *SIAM International Conference in Data Mining [SDM 2013]*, pages 162–170, Philadelphia, 2013. Society for Industrial and Applied Mathematics. doi: 10.1137/1.9781611972832.18. URL <http://arxiv.org/abs/1304.4657>.
- Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. Hyperbolic geometry of complex networks. *Physical Review E*, 82:

- 036106, 2010. doi: 10.1103/PhysRevE.82.036106. URL <http://arxiv.org/abs/1006.5169>.
- László Lovász. *Large Networks and Graph Limits*. American Mathematical Society, Providence, Rhode Island, 2012.
- Manuel Middendorf, Etay Ziv, and Chris Wiggins. Inferring network mechanisms: The *drosophila melanogaster* protein interaction network. *Proceedings of the National Academy of Sciences (USA)*, 102:3192–3197, 2005. URL <http://arxiv.org/abs/q-bio/0408010>.
- Mark E. J. Newman. *Networks: An Introduction*. Oxford University Press, Oxford, England, 2010.
- Henry Pao, Glen A. Coppersmith, and Carey E. Priebe. Statistical inference on random graphs: Comparative power analyses via Monte Carlo. *Journal of Computational and Graphical Statistics*, 20:395–416, 2011. doi: 10.1198/jcgs.2010.09004.
- Bruno Pelletier. Kernel density estimation on Riemannian manifolds. *Statistics and Probability Letters*, 73:297–304, 2005. doi: 10.1016/j.spl.2005.04.004.
- Claudio Rocchini. Poincare halfplane eptagonal hb, 2007. URL http://commons.wikimedia.org/wiki/File:Poincare_halfplane_eptagonal_hb.svg. Retrieved 20 October 2014.
- Martin Rosvall and Carl T. Bergstrom. Mapping change in large networks. *PLoS ONE*, 5:e8694, 2010. doi: 10.1371/journal.pone.0008694. URL <http://arxiv.org/abs/0812.1242>.
- Olaf Sporns. *Networks of the Brain*. MIT Press, Cambridge, Massachusetts, 2010.
- Elias M Stein and Guido L Weiss. *Introduction to Fourier analysis on Euclidean spaces*, volume 1. Princeton university press, 1971.
- Minh Tang, Avanti Athreya, Daniel L. Sussman, Vince Lyzinski, and Carey E. Priebe. A nonparametric two-sample hypothesis testing problem for random dot product graphs. E-print, arxiv.org, 2014. URL <http://arxiv.org/abs/1409.2344>.
- Audrey Terras. *Harmonic analysis on symmetric spaces and applications*, volume 1. Springer-Verlag, New York, 1985.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, Cambridge, England, 1998.
- Duncan J. Watts and Steven H. Strogatz. Collective dynamics of “small-world” networks. *Nature*, 393:440–442, 1998. doi: 10.1038/30918.