



# HHS Public Access

Author manuscript

*Stat Sci.* Author manuscript; available in PMC 2020 November 23.

Published in final edited form as:

*Stat Sci.* 2019 August ; 34(3): 428–453. doi:10.1214/19-sts702.

## The Geometry of Continuous Latent Space Models for Network Data

**Anna L. Smith [Assistant Professor],**

Department of Statistics, University of Kentucky, 725 Rose Street, Lexington, Kentucky 40536, USA

**Dena M. Asta [Assistant Professor],**

Department of Statistics, The Ohio State University, 1958 Neil Avenue, Columbus, Ohio 43210, USA

**Catherine A. Calder [Professor]**

Department of Statistics and Data Sciences, The University of Texas at Austin, 2317 Speedway, Stop D9800, Austin, Texas 78712, USA

### Abstract

We review the class of continuous latent space (statistical) models for network data, paying particular attention to the role of the geometry of the latent space. In these models, the presence/absence of network dyadic ties are assumed to be conditionally independent given the dyads' unobserved positions in a latent space. In this way, these models provide a probabilistic framework for embedding network nodes in a continuous space equipped with a geometry that facilitates the description of dependence between random dyadic ties. Specifically, these models naturally capture homophilous tendencies and triadic clustering, among other common properties of observed networks. In addition to reviewing the literature on continuous latent space models from a geometric perspective, we highlight the important role the geometry of the latent space plays on properties of networks arising from these models via intuition and simulation. Finally, we discuss results from spectral graph theory that allow us to explore the role of the geometry of the latent space, independent of network size. We conclude with conjectures about how these results might be used to infer the appropriate latent space geometry from observed networks.

### Keywords

Geometric curvature; graph Laplacian; latent variable; network model

## 1. UNDERSTANDING NETWORK DATA THROUGH GEOMETRIC EMBEDDINGS

An important consideration in the development of an analytic strategy for understanding network data is the expected complexity of the types or patterns of connections in the

network. In a statistical framework, we can formalize this notion through a particular dependence structure among the potential ties in the network (Kolaczyk and Csárdi, 2014; Airol di et al., 2008a). For example, a simple statistical model might treat each pair of nodes, or dyad, independently (Erdős and Rényi, 1960; Gilbert, 1959). Of course, this is typically not a reasonable assumption. For example, in a social network, it is (at least heuristically) reasonable to assume that “the friend of my friend is also a friend of mine” (Rapoport, 1953). Indeed, this transitivity effect has strong theoretical support in the social network setting (Simmel, 1950; Holland and Leinhardt, 1970; James, 1990) and has been observed in a variety of empirically observed networks as well (Davis, 1970; Goodreau, Kitts and Morris, 2009; Watts, 1999a). This expectation for complex nondyadic dependence among nodes is not unique to social networks or the social sciences, but instead permeates network applications from various disciplines. However, note that such complex interactions directly contradict an assumption of dyad independence. If there is a connection between the  $i$ th and  $j$ th node and a connection between the  $j$ th and  $k$ th node, this tells us something about the probability of a connection between the  $i$ th and  $k$ th node.

This of course raises a natural question: What type of dependence is appropriate and how might this dependence affect an analysis of network data? One popular solution comes from the class of exponential random graph models, or ERGMs (Frank and Strauss, 1986; Wasserman and Pattison, 1996; Snijders, 2002; Robins et al., 2007). These models assume a particular dependence regime for the full set of potential ties in the network. In Frank and Strauss’s (1986) original formulation, Markov dependence is assumed (i.e., where all dyads that share a node are dependent), but more recent specifications have explored alternative dependence regimes (e.g., local conditional Markov dependence in Snijders et al., 2006).

Alternatively, we might consider representing this complex network data in an unobserved, latent lower-dimensional space and specifying a dependence regime there. This approach draws on ideas developed by Lazarsfeld, Henry and Anderson (1968) which proposed using latent variables to simplify complicated dependence structures and develop simpler model specifications: in these analyses, an unobserved latent variable is constructed such that the distribution of the data given the latent variables has a simple form. Of course, this idea has been used in a variety of settings, from factor analysis to item response theory to mixture models (Spearman, 1904; van der Linden and Hambleton, 2013; Bartholomew, Knott and Moustaki, 2011), and so it is no surprise that such a technique might be helpful in the network data setting. Additionally, we might consider choosing this lower-dimensional space to be a metric space, in which case we are embedding the nodes of the network within a particular geometry (e.g., Hoff, Raftery and Handcock, 2002; McCormick and Zheng, 2015; Airol di et al., 2008b; Krioukov et al., 2010). Although the details of the embedding may vary across applications or analytical methods, nodes’ positions in the latent space will correspond to features in the observed network with nodes that are connected in the network being closer together (or more generally, being more similar) in the latent space. That is, dependence in the network can be induced through distances, or similarities, between nodes in this lower-dimensional latent space. In practical settings, note that this framework provides a highly intuitive interpretation of the network data. For example, in the social network setting, we can interpret the embedding as a map of nodes’ relative positions in a

latent “social space,” capturing the variety of unobserved social forces that may have influenced the observed pattern of social ties.

Traditionally, most latent space approaches have suggested embedding the nodes of the network in a Euclidean latent space (Hoff, Raftery and Handcock, 2002). Although Euclidean space is certainly familiar and relatively easy to work with, it is not clear that it is the most appropriate geometry for network data or that Euclidean space is best equipped to reveal the variety of potentially interesting network features (Krioukov et al., 2010). More specifically, we argue that the geometry of the latent space plays an important role in determining both the fitness of the analysis for the particular observed network data as well as the types of questions which the analysis is equipped to answer. In order to tease out this relationship, we will consider relating observed network features to latent space geometry through two different lenses: first, through common network summary statistics (e.g., transitivity and average degree; Kolaczyk and Csárdi, 2014) and second, through graph Laplacians which describe interesting features of networks and whose spectrum, under suitable assumptions of the latent space, is not influenced by network size.

We begin by summarizing existing approaches to the analysis of network data which take advantage of a lower-dimensional embedding of the nodes of the network in an unobserved latent space. In Section 2, we discuss commonly used similarity measures, focusing on methodology that is most amenable to the traditional statistical modeling framework. In Section 3.1, we review the variety of different types of latent spaces—from the familiar Euclidean space to more abstract spaces, like ultrametric space—in the continuous latent space model literature, in order to point out the impact the choice of geometry for the latent space can have on an analysis of network data. We argue that the geometry of the latent space can play an important role in understanding network data, highlighting the advantages of a newly proposed latent space for network data, hyperbolic space, in Section 3.2. In Section 4, we describe typical features of networks embedded in spaces with different curvatures using simulated network data. In this way, we connect the geometry of the latent space to features of networks that we typically care about, such as centralization measures and properties of the degree distribution. Finally in Section 5, we examine graph Laplacians to elucidate the relationship between the types of questions or discoveries that are possible in an analysis of network data which uses a latent space of a particular geometry.

## 2. CONTINUOUS LATENT SPACE MODELS FOR NETWORK DATA

A fully parametric model for network data provides a natural way to think about utilizing a lower-dimensional latent space for the nodes in the network. Letting the nodes in the network be indexed by  $i, j = 1, \dots, n$ , then a generic form for such a model is as given below:

$$\begin{aligned} Y_{ij} &\stackrel{\text{ind}}{\sim} \text{Bernoulli}(p_{ij}), \quad i \neq j \\ \text{logit}(p_{ij}) &= \alpha + s(z_i, z_j), \\ z_i &\in \mathcal{Z}^t, z_i \stackrel{\text{ind}}{\sim} f_t(z | \psi), \end{aligned} \tag{2.1}$$

where  $Y_{ij}$  is a binary indicator of a tie between nodes  $i$  and  $j$ ,  $z_i$  is the  $i$ th node’s  $t$ -dimensional unobserved latent vector in the  $t$ -dimensional space  $\mathcal{Z}^t$ , and  $\psi$  parameterizes

the distribution of the latent vectors in  $\mathcal{Z}^t$ . The generic latent space model assumes that network ties are conditionally independent, given the dyad's latent vectors. The probability of a tie consists of a baseline tendency for tie formation,  $\alpha$ , and depends on some similarity measure,  $s(\cdot, \cdot)$ , of the latent vectors so that the probability of a tie is higher for nodes with more similar latent vectors and lower for nodes with latent vectors that are dissimilar. In this sense, the latent vectors can represent the observed pattern of ties in the network as well as any unobserved attributes of the nodes that may have contributed to the pattern of observed ties. For example, in a social network, the latent vectors are typically interpreted as individuals' positions in a latent "social space," representing the variety of unobserved social forces present in the observed network. In most cases, the similarity measure,  $s(\cdot, \cdot)$ , is chosen so that transitivity is a natural consequence of the model specification; if nodes  $i$  and  $j$  are tied (i.e.,  $s(z_i, z_j)$  is large) and nodes  $j$  and  $k$  are tied (i.e.,  $s(z_j, z_k)$  is large), then  $i$  and  $k$  are likely tied (i.e.,  $s(z_i, z_k)$  is likely large). This model can easily be generalized to incorporate homophily on observed attributes by simply adding a  $\beta'x_{ij}$  term, where  $x_{ij}$  is a vector of (dyad-level) observed covariates for the dyad or pair formed by the  $i$ th and  $j$ th nodes. These latent space models typically require extra structure (e.g., hierarchical priors on the components of  $\psi$ ) in order to model community structure and degree heterogeneity.

Here, we focus on methods that use a *continuous* latent space (CLS), although the use of a discrete space is also possible (e.g., the stochastic block models proposed by Snijders and Nowicki, 1997, described briefly in Section 2.1). In practice, the dimension of the latent space,  $t$ , must be specified by the researcher, though traditional model selection techniques can be implemented to inform this choice. Without loss of generality, we will assume  $t = 2$ , which allows for easy visualization of the latent vectors and is a popular choice. Note that this model can be easily extended to networks with nonbinary (i.e., weighted) ties using ideas from generalized linear models. Further, note that model (2.1) is written for directed ties, and subsumes the setting of networks with undirected ties by simply replacing the condition  $i \sim j$  with  $i < j$  (though this may have implications for the choice of  $s(\cdot, \cdot)$ ). It is also worth pointing out that many existing network analysis methods which take advantage of a latent space representation do not require a fully parametric model, as outlined above, but instead fall under the umbrella of semi-parametric or nonparametric techniques. We will highlight this distinction in our discussion of these methods below.

In the rest of this section, we describe existing latent space models for network data in terms of their similarity measure,  $s(z_i, z_j)$ . In our discussion, we will group existing methods according to whether the proposed similarity measure has a dot product form, uses a distance metric, or is a graphon. However, these categories are not strictly distinct. For example, it may certainly be possible to rewrite a dot product model as a distance metric or perhaps to re-specify a distance model as a graphon. Further understanding of the connections between these models provides an interesting avenue for future research (we will revisit this idea briefly in Section 2.3). Finally, we should point out that one of the most popularly cited papers regarding latent space models for networks, Hoff, Raftery and Handcock (2002), actually proposes two distinct versions of a model that uses a Euclidean latent space. First, they propose a version which uses a distance metric (which they refer to as the "distance model" and which we will discuss later in Section 3.1) and then secondly,

they propose a version which utilizes a dot product similarity measure (which they refer to as the “projection model” and which we discuss first, in the Section 2.1).

## 2.1 Dot Product Models

Many latent space methods for network data use a simple dot product similarity measure,

$$s(z_i, z_j) = z_i' z_j.$$

In Euclidean space, the dot product measures the positioning of the vectors in the space. For vectors in  $\mathbb{R}^t$ , the dot product is related to both the angular distance,  $\phi$ , and the Euclidean distance,  $d(z_i, z_j)$ , in the following ways:

$$\begin{aligned} z_i' z_j &= |z_i| |z_j| \cos \phi \\ &= \frac{1}{2} \{ |z_i|^2 + |z_j|^2 - [d(z_i, z_j)]^2 \}. \end{aligned}$$

The dot product in  $\mathbb{R}^t$  can be interpreted as the scaled cosine similarity (scaled by the magnitude of  $z_i$  and  $z_j$ ) or equivalently as a function of the squared distance between unit vectors in the direction of  $z_i$  and  $z_j$ . As can be seen in Figure 1, while the dot product in Euclidean space is certainly related to Euclidean distance, they are not equivalent measures. Although Euclidean distance is certainly much easier to visualize, the dot product is simpler to calculate and thus has computational benefits.

Most notably, the second version of the latent space model introduced in Hoff, Raftery and Handcock (2002), the so-called “projection model,” fits within this framework (the first version of their proposed model uses a distance metric and will be discussed in detail in Section 3.1). Hoff, Raftery and Handcock’s (2002) projection model is specifically designed to handle directed network ties and is specified as follows:

$$\begin{aligned} s(z_i, z_j) &= -a_i z_i' z_j, \\ z_i &\in t - \text{dimensional hypersphere,} \end{aligned} \tag{2.2}$$

where  $a_i > 0$  is a sociality parameter for the  $i$ th individual. Note that  $s(z_i, z_j)$  can be interpreted as the signed scalar projection of  $v_i$  onto  $v_j$ , where  $v_i = a_i z_i$  and  $v_j = a_j z_j$  are the scaled latent positions for the  $i$ th and  $j$ th node, scaled by their sociality parameters,  $a_i$  and  $a_j$  respectively (see Figure 2, which we discuss in more detail in Section 3.1.2; note that the notation used here differs from that used in Hoff, Raftery and Handcock, 2002). The sociality parameters model a node’s tendency to send ties (i.e., to be “social”). Of course, the authors point out that this model could be easily adjusted to account for variability in nodes’ tendency to receive ties, by replacing  $a_j$  with  $a_j$ .

In practice, the authors suggest reparameterizing this model with alternative latent vectors,  $v_j = a_j z_j$ , so that

$$s(v_i, v_j) = \frac{v_i' v_j}{|v_j|},$$

$$v_i \in \mathbb{R}^t.$$

This specification of the model is more parsimonious since it does not require a separate model for the sociality parameters. Further, since the  $v_i$  are now arbitrary positions in  $\mathbb{R}^t$ , we can simply choose  $f_\lambda(v|\psi)$  to be a standard Normal distribution, as suggested for both versions of the latent space model proposed in Hoff, Raftery and Handcock (2002).

Other examples of dot product latent space models include McCormick and Zheng's (2015) model for aggregated relational data (ARD) on the hypersphere (discussed in more detail in Section 3.1) and Hoff's (2005) bilinear model, which incorporates additive node-level sociality effects to allow for degree heterogeneity,

$$s(z_i, z_j) = a_i + b_j + z_i' z_j$$

$$a_i \stackrel{\text{iid}}{\sim} \text{Normal}(0, \sigma_a^2),$$

$$b_j \stackrel{\text{iid}}{\sim} \text{Normal}(0, \sigma_b^2)$$

$$z_i \stackrel{\text{iid}}{\sim} \text{Multivariate Normal}(0, \sigma^2 I_t),$$

where  $a_i$  and  $b_j$  are sociality effects for the  $i$ th and  $j$ th node, respectively. Hoff (2005) motivates the interpretation of  $\sigma_a^2$ ,  $\sigma_b^2$  and  $\sigma^2$  as random effects variances, highlighting similarities between the model and a traditional ANOVA model. In fact, these sociality effects mimic the similarity measures used in later models (e.g., Krivitsky et al., 2009, see Section 3.1 for more discussion), but the idea of modeling sender and receiver effects has a long history (Holland and Leinhardt, 1981; Fienberg and Wasserman, 1981; van Duijn, Snijders and Zijlstra, 2004).

Dot product models in Euclidean space have been further developed by Young and Scheinerman (2007), Nickel (2009) and Pao, Coppersmith and Priebe (2011) as the class of random dot product graphs, where  $\mathcal{X}^t$  is the  $[0, 1]$  box in  $\mathbb{R}^t$  and  $s(z_i, z_j)$  is the dot product, or some function of it.

We can also further generalize the type of similarity measure used in dot product models to a quadratic-like form,

$$s(z_i, z_j) = z_i' A z_j, \tag{2.3}$$

where  $A$  is a  $t \times t$ -dimensional matrix. This specification subsumes the models discussed previously in this section, by simply choosing  $A$  to be the identity matrix. Examples include Hoff's (2008) eigenmodel for undirected networks, Hoff's (2009) SVD-inspired extension to directed networks, and Minhas, Hoff and Ward's (2016) class of additive and multiplicative effects (AME) models.

It is also worth pointing out that the traditional stochastic block model (SBM) can be viewed as a special case of these quadratic-like latent space models. Generally, the stochastic block model assumes that the nodes can be separated into  $t$  unobserved latent classes where the membership vectors,  $z_i$ , are unknown. From model (2.1), assume  $\alpha = 0$  and let

$$\begin{aligned} s(z_i, z_j) &= p_{\pi(z_i), \pi(z_j)}, \\ z_i &\sim \text{Multinomial}_t(1, \psi), \end{aligned}$$

where  $\pi$  maps nodes' latent vectors to  $t$  classes,  $p_{l_1 l_2}$  parameterizes ties between nodes in the  $l_1$ th and  $l_2$ th classes, and  $\psi$  is a  $t$ -dimensional vector of membership probabilities. The model treats nodes within the same class as stochastically equivalent and parameterizes the probability of ties between nodes belonging to different classes by the  $p_{l_1 l_2}$ 's. Note that we can rewrite the similarity function in the quadratic-like form of equation (2.3) by letting  $A$  be a  $t \times t$  matrix of the  $\{p_{l_1 l_2}\}$  parameters. In the original versions of this model, the number of classes  $t$  was assumed to be known, though a variety of methods now exist to help understand the implication of different choices of  $t$  (Saldaña, Yu and Feng, 2017; Chen and Lei, 2018).

Airoldi et al. (2008b) further extend this model by allowing nodes to have simultaneous partial memberships in each of the  $t$  classes in the mixed membership stochastic block model (MMSBM) with

$$\begin{aligned} s(z_i, z_j) &= z_i' A z_j, \\ z_i &\sim \text{Dirichlet}_t(\psi), \end{aligned}$$

where again  $A$  is the block diagonal matrix of the between-class tie probability vector and  $\psi$  is a  $t$ -dimensional vector of membership probabilities. In this setting, the  $z_i$ 's are still interpreted as unobserved latent membership vectors and further, can be viewed geometrically as positions on a  $(t - 1)$ -dimensional simplex. In fact, in the original specification of the stochastic block model, the  $z_i$  can be viewed geometrically as the *vertices* of a  $(t - 1)$ -dimensional simplex. In the MMSBM, the membership vectors have a geometric interpretation on the simplex: distance from a node's position on the simplex to a particular vertex (community) is an indication of how strongly that node's behavior in the network matches the behavior of the community as a whole.

## 2.2 Distance Models

A natural choice for  $s(\cdot, \cdot)$  is a (negative) distance metric, in which case we can interpret the  $z_i$ 's as latent positions in a  $t$ -dimensional latent metric space,  $\mathcal{X}^t$ . Again letting the nodes in the network be indexed by  $i, j = 1, \dots, n$ , we will refer to these special cases as latent distance models:

$$\begin{aligned} Y_{ij} &\stackrel{\text{ind}}{\sim} \text{Bernoulli}(p_{ij}), \quad i \neq j \\ \text{logit}(p_{ij}) &= \alpha - d(z_i, z_j), \\ z_i &\in \mathcal{X}^t, z_i \stackrel{\text{ind}}{\sim} f_t(z | \psi), \end{aligned} \tag{2.4}$$

where  $d$  is the natural metric on  $\mathcal{X}^t$  and  $f_t$  describes the distribution of latent positions in  $\mathcal{X}^t$ . In a latent distance model, transitivity effects are guaranteed by the triangle inequality. Note also that these models are most appropriate for undirected ties, since distances are inherently symmetric. These models also allow us to visualize the nodes' latent positions by simply plotting estimates,  $\hat{z}_i$ , in  $\mathcal{X}^t$  which can help illuminate previously unnoticed patterns in the data. For example, one could plot the latent positions from a null model and color the positions according to some covariate that could later be included in the model. Although this procedure is available to all versions of latent space models (whether distance is the chosen similarity measure or not), interpretation of relationships among nodes' latent positions in these plots is certainly more intuitive in the latent distance models.

Further, note that the concept of using distance to model dependence is not unique to the class of latent distance models for network data. In fact, this idea has had a long tradition of success in the statistical modeling literature. For example, in methods for continuously-indexed spatial data, distance is obviously a natural part of the modeling framework. Even in cases where spatial data is observed at an aggregate level (i.e., in some type of districts or areal units) the spatial weights matrix used in spatial regression models, such as CAR and SAR models, contains information about distance through a nearest neighbors matrix (where the  $ij$ th entry is one if the  $i$ th district neighbors the  $j$ th district and zero otherwise). In fact, in many spatial applications, non-Euclidean distances can be used to better represent the underlying dependence structure of the particular data generating process. For example, a researcher might use stream or flow distance when considering chemical concentrations in a watershed area or perhaps travel time or travel cost for a group of city commuters. Further, models for phylogenetic data often make use of Hamming distance<sup>1</sup> to represent similarity among species when creating a phylogenetic tree. In short, using latent distances to model dependence among nodes in a network allows for an intuitive representation of an otherwise complex dependence structure and draws on the well-developed practices of using distance to model dependent data in other fields of statistical research.

Examples of latent space approaches which use distance include Hoff, Raftery and Handcock (2002), Handcock, Raftery and Tantrum (2007), Krivitsky et al. (2009) and will be discussed in more detail in Section 3.1.

### 2.3 Graphons

Graphons provide a different perspective than the models discussed thus far, but in many cases, the graphon approach is essentially equivalent to the CLS approach. To be precise, a *graphon* is a measurable function  $W$  which maps  $[0,1]^2$  to  $[0, 1]$ , and is symmetric in its coordinates (Lovász, 2012). The graphon approach can be regarded as the CLS model given in (2.1) by taking the  $z_i$ 's to be uniformly and independently drawn points in the unit interval and setting  $p_{ij} = W(z_i, z_j)$ . On the other hand, a CLS model can be regarded as a graphon by setting  $W(z_i, z_j) = p_{ij}$ , as long as (and is often the case that) the unit interval can be identified with the CLS model's latent space,  $\mathcal{X}^t$ , along a (likely discontinuous) isomorphism mod 0 of probability measure spaces. In this sense, we can think of the graphon approach as a CLS

<sup>1</sup>The Hamming distance between two vectors counts the number of positions at which the corresponding vectors differ.



model with a simple latent space—the unit interval—but with a very general class of (often complicated) functions,  $W$ , that specify the connection probabilities. Graphons are commonly used to study asymptotic behavior and networks that grow in size. For example, mixtures of graphons are exactly the distributions of infinite graphs invariant under permutations of the nodes (Lovász, 2012). Conversely, it is possible to formalize a sense in which a limit of suitable sequences of growing graphs uniquely identifies a graphon up to graphon equivalence (i.e., up to relabeling; Diaconis and Janson, 2008).

While the specification of graphon and CLS models are similar, the standard inferential strategies for the two types of models vary. Inference in CLS models focuses on the estimation of unknown parameters characterizing the distribution of the positions in latent space, in addition to parameters capturing standard fixed and random effects components of the model. On the other hand, inference for graphon models focuses on estimation of the  $W$  function; in CLS models, the functional form of the connection probabilities is treated as fixed. In many settings, constraints on allowable graphon models are used to increase the efficiency of the estimation process. For example, a graphon can be practically estimated from samples of a stochastic blockmodel (Abbe, Bandeira and Hall, 2016; Amini et al., 2013; Bickel and Chen, 2009), efficiently under an assumption of Hölder continuity on the  $W$ -function (Gao, Lu and Zhou, 2015; Wolfe and Olhede, 2013). However, it is difficult to impose a simple condition on a  $W$ -function, in order to constrain salient network characteristics of interest. In contrast, choice of latent space and constraints on the parameters of the distribution of latent positions in a CLS model can impose natural constraints on small world properties, degree distributions and clusterability (cf. Krioukov et al., 2010).

Graphons and CLS models also differ in how they are used to compare random networks. Two graphons are often compared by a *cut-metric*, an infima of  $L_p$ -distances between the  $W$ -functions *up to graphon equivalence* (Lovász, 2012). Two CLS distance models with a manifold as a fixed latent space and connection probability are often compared by an infima of  $L_p$ -distances between the node densities *up to a distance-preserving bijection* of latent space. The latter distance implicitly takes into account the geometry of the latent space; the former distance, only dependent up to measure-preserving bijections, cannot possibly take into account any interesting features of the latent space.

### 3. LATENT SPACE GEOMETRIES

In this section, we continue our description of continuous latent space models for network data, paying particular attention to the geometric properties of the latent space assumed under each model. In Section 3.1, we review existing approaches in Euclidean space (Section 3.1.1); elliptic space, such as the surface of a three-dimensional sphere (Section 3.1.2); and other spaces (Section 3.1.3). In Section 3.2, we highlight properties of hyperbolic space and detail its promising potential application in the class of continuous latent space models for network data, inspired by recent network analysis methods which take advantage of a hyperbolic latent space (Krioukov et al., 2010; Asta and Shalizi, 2014; Aldecoa, Orsini and Krioukov, 2015). We will discuss various interesting aspects of the different choices for the geometry of the latent space in each subsection, but we also have provided a summary of

some of the most important differences between Euclidean, elliptic and hyperbolic space in Table 1.

### 3.1 Literature Review

**3.1.1 Euclidean space.**—Perhaps the most intuitive specification of the latent space model is the Euclidean latent distance model, where  $\mathcal{X}^t$  is the familiar  $t$ -dimensional Euclidean space and  $d$  is Euclidean distance,

$$\begin{aligned} d(z_i, z_j) &= \|z_i - z_j\| \\ &= \sqrt{(z_i - z_j)'(z_i - z_j)}. \end{aligned}$$

In fact, this version of the model was introduced by Hoff, Raftery and Handcock (2002) as the “distance model” and is certainly the most widely used latent space model for network data. The authors specify  $f_i(z|\psi)$  as a standard normal distribution.

Handcock, Raftery and Tantrum (2007) extend this model to capture community structure in a network by letting  $f(z|\psi)$  be a mixture of  $K$  multivariate normal distributions:

$$z_i \sim \sum_{k=1}^K \lambda_k \text{MVN}_t(\mu_k, \sigma_k^2 I_t),$$

where  $\lambda_k$  is the probability that a node belongs to the  $k$ th group ( $\sum_{k=1}^K \lambda_k = 1$ ). Under this model, the latent positions are spatially clustered in  $\mathcal{X}^t$ . Thus, this model can also be viewed as a generalization of the stochastic block model (see Section 3.1.3 for more details), with the  $k$ th mixture component representing the  $k$ th class, which also allows for transitivity within blocks and homophily on attributes. On the other hand, Handcock, Raftery and Tantrum’s (2007) model can be thought of as an extension of Hoff, Raftery and Handcock’s (2002) version of the Euclidean latent distance model which introduces more structure in the latent positions (e.g., spatial clustering) to better represent potential community structure in the observed network.

Although developed for directed ties, Handcock, Raftery and Tantrum’s (2007) model does not explicitly model degree heterogeneity. This extension is provided by Krivitsky et al. (2009) where the addition of additive individual random effects act as sociality or popularity effects to accommodate the differing tendency of nodes to send or receive ties, with

$$\begin{aligned} s(z_i, z_j) &= a_i + b_j - d(z_i, z_j), \\ a_i &\stackrel{\text{iid}}{\sim} \text{Normal}(0, \sigma_a^2), \\ b_j &\stackrel{\text{iid}}{\sim} \text{Normal}(0, \sigma_b^2), \end{aligned}$$

where  $a_j$  is the sender effect for the  $j$ th node and  $b_j$  is the receiver effect for the  $j$ th node. The authors’ extension of Handcock, Raftery and Tantrum’s (2007) model borrows from the form of Hoff’s (2005) bilinear model, mentioned in Section 2.1. In both cases, the sender

and receiver effects mimic the specification of an ANOVA model and can be interpreted similarly.

The latent Euclidean distance model has also been extended to the case of dynamic networks (Sarkar and Moore, 2006; Sewell and Chen, 2015).

While a Euclidean latent space certainly allows for easy visualization and interpretation of estimated latent positions and may be desirable given the familiarity of Euclidean geometry, it is not clear whether this geometry is best suited to represent the complex dependencies we typically encounter in network data. We will discuss this point further in Section 3.2.

**3.1.2 The unit hypersphere.**—As discussed in Section 2.1, Hoff, Raftery and Handcock (2002) introduce two versions of a latent space model: the distance model in Euclidean space (see Section 3.1.1) and the so-called projection model, which uses a dot product similarity measure and embeds nodes on the unit circle (i.e., a one-dimensional unit hypersphere).

Note that this projection model is *not* an example of a latent distance model, since the natural distance metric on the hypersphere is great-circle distance (arc length). However, on the unit hypersphere, the great-circle distance between two latent positions,  $z_i$  and  $z_j$  is equivalent to their angular distance,  $\phi$ , and can be calculated as the inverse cosine of the dot product of  $z_i$  and  $z_j$ . Thus, the similarity measure in Hoff, Raftery and Handcock's (2002) projection model (see Equation (2.2)) is simply the *cosine similarity* between the latent position vectors, scaled by a sociality parameter,  $a_i$  (see Figure 2 for some examples). When  $z_i$  and  $z_j$  are close (and the angle between them is acute, as in the left panel of Figure 2),  $z_i'z_j$  will be positive and when they are far apart (and the angle between them is obtuse, as in the right panel of Figure 2),  $z_i'z_j$  will be negative. Thus, this choice of  $s(\cdot, \cdot)$  preserves the relationship between the probability of a tie and some measure of “closeness.” That is, just as in the latent distance models, nodes that are “closer” together are more likely to be tied. Further, if this model were used for networks with undirected ties, the sociality parameters (the  $a_i$ 's) would disappear, leaving  $s(\cdot, \cdot)$  as the cosine similarity, which is a monotonic function of the natural distance metric,  $\phi$ , on the unit hypersphere. However, although cosine similarity is not a proper distance metric itself (since, e.g., it is not nonnegative), it is a popular similarity measure. Compared to a true latent distance model on the unit hypersphere, this choice of  $s(\cdot, \cdot)$  changes the interpretation of  $\alpha$  (and any parameters for covariates, if included) in model (2.1) and also re-weights the distances, since  $\cos\phi$  is a nonlinear function of  $\phi$ . For example, given a fixed network, we would not expect an estimate of  $\alpha$  under the distance model to match an estimate under the projection model. Generically,  $\alpha$  represents a baseline tendency for tie formation, *relative to the average expected similarity* under the chosen model. For example, consider a two-dimensional latent space and assume that each dimension of the latent positions are independently and identically distributed according to a standard Normal, it is relatively easy to show that under the distance model the average expected similarity is  $\sqrt{\pi}$  while under the projection model this same quantity is zero.

McCormick and Zheng (2015) extend the projection model to the setting of aggregated relational data (ARD), an example of incomplete network data. In this setting, surveyed individuals answer “how many  $X$ ’s do you know?,” where  $X$  represents particular, typically hard-to-reach, groups of interest such as drug users, terrorists, or homeless individuals. Thus, rather than observe the entire network, we only observe connections between surveyed individuals and the groups of interest; we do not observe ties between surveyed individuals and actual members of the group nor ties between members of the group. To accommodate this type of data, the authors first propose a model for the complete case (if the network had been entirely observed) and then derive a model for the aggregated observations. The proposed model for the complete case is a slight extension of Hoff, Raftery and Handcock’s (2002) projection model, where cosine similarity is again incorporated:

$$s(z_i, z_j) = a_i + a_j + \zeta z_i^t z_j^t,$$

where  $a_i$  is a gregariousness or sociality effect for the  $i$ th node and  $\zeta$  scales the overall influence of the latent component. Compared to Hoff, Raftery and Handcock’s (2002) projection model, McCormick and Zheng (2015) specify a sociality effect that is additive rather than multiplicative and that accounts for differences between the sender and the receiver. They also incorporate an additional scale parameter,  $\zeta$ , which modifies the influence of the latent vectors. Latent positions for the survey participants are modeled as uniformly distributed across the hypersphere.

**3.1.3 Other spaces.**—Although both the SBM and MMSBM allow for geometric interpretations of the latent vectors (i.e., as positions on the simplex), neither is an example of a latent distance model.<sup>2</sup> Although some metrics have been developed in the context of compositional data, they generally are only intended for points on the open simplex and exclude points on the boundary. For example, of the popular Aitchison distance metric defined for the open simplex, Aitchison et al. (2000) note that when one of the coordinates of a point in the simplex tends to zero, the distance between it and other points in the simplex will tend toward infinity. The authors explain, “There is nothing surprising about this feature; it is merely recognizing that a composition with one of the parts absent may be chemically, physically, or biologically completely different from compositions with all components positive. Doveton’s (1998) perfect martini with its (gin, dry martini, sweet martini) composition is completely different from a cocktail with no gin but only dry and sweet martini present.” For compositional data, this may be a natural assumption but it is not clear whether it is necessary in the setting of latent membership vectors.

In Schweinberger and Snijders (2003), the authors specify a latent distance model in an ultrametric latent space. Unlike other versions of the latent space models considered here, this model does not specify  $\mathcal{X}^t$  as a traditional metric space. Recall that a metric space

<sup>2</sup>Of course, utilizing  $s(z_i, z_j) = -d(z_i, z_j)$  in the original stochastic block model is meaningless, since all vertices of the simplex are, by design, equidistant. In the mixed-membership stochastic block model, this choice would significantly change the interpretation of the model: from a node’s tie probabilities depending on a weighted sum of the between-class tie probabilities, weighted by the node’s membership vector, to a node’s tie probabilities depending on distances in the space of all membership vectors.

consists of a set,  $M$ , and distance (metric),  $d$ , which satisfies the following properties for all points  $x, y$  and  $z \in M$ :

1. Nonnegativity:  $d(x, y) \geq 0$ .
2. Identity of Indiscernibles:  $d(x, y) = 0 \Leftrightarrow x = y$ .
3. Symmetry:  $d(x, y) = d(y, x)$ .
4. Triangle Inequality:  $d(x, z) \leq d(x, y) + d(y, z)$ .

An ultrametric space, or non-Archimedean space, is a metric space where  $d$  additionally satisfies the strong triangle inequality

$$d(x, z) \leq \max\{d(x, y), d(y, z)\},$$

which, among other things, means that all triangles are isosceles with the unequal side being shortest and that every point in a disc is a center of that disc, so that two discs intersect only if one disc completely contains the other. Not surprisingly, visualizing an ultrametric space and understanding the behavior of distances within it can be difficult. However, ultrametric space can be visualized via dendrograms or trees. For example,  $\mathbb{Q}_p$ , the field of  $p$ -adic numbers, is a well-known ultrametric space which can be visualized as the tree in the right panel of Figure 3 (see Holly, 2001, for a thorough motivation of this figure). For simplicity, consider a tree of the familiar animal species classification system in the left panel of Figure 3. Naturally, the distance between species is the height of the smallest (inverted) tree between them so that, for example, species within the same genus are closer ( $d = 1$ ) than those who only share a common order ( $d = 3$ ). Returning to the intersection of discs in ultrametric space, we illustrate one example of a disc visualized in the ultrametric tree for the animal species classification system in the left panel of Figure 3. For species  $g$ , the disc centered at this species with radius  $\gamma = 3$  is the set of all leaves in the subtree descending from the node at distance level  $\gamma$ . If we consider any arbitrary point (or species) within this disc, say  $x$ , then the disc centered at  $x$  is the set of leaves in the subtree descending from the (unique) node above  $x$  at distance level  $\gamma$ ; this yields the exact same disc.

Schweinberger and Snijders (2003) motivate this choice for  $\mathcal{Z}^f$  by noting that social scientists' understanding of latent settings<sup>3</sup> structures mimic properties guaranteed in ultrametric spaces. More specifically, the latent settings should be nonoverlapping, the interaction within settings should be stronger than the interaction between settings, and hierarchical nesting of the settings is expected. The authors fit the model to a fraternity interaction network and provide a topological map (e.g., Figure 4) of the estimated latent positions where each level of the mapped social "mountain" can be interpreted as a social setting. An example of the style of resulting settings groups is provided in Figure 4.

---

<sup>3</sup>Here, the authors refer to the settings terminology emphasized by Pattison and Robins (2002) where a setting is a close-knit cluster of actors that are strongly tied.

### 3.2 Hyperbolic Space

Recent work in latent space models for network data has resulted in a novel proposal: the use of negatively curved  $t$ -dimensional hyperbolic space,  $\mathbb{H}^t$ . In the following section, we review some recent work in this area and provide some intuition for why hyperbolic space may be particularly amenable to modeling important network characteristics. We refer the reader to Appendix A for more details about the general geometric properties of hyperbolic space. First suggested by Krioukov et al. (2010), modeling and analyzing networks in latent hyperbolic spaces is rapidly being investigated (Aldecoa, Orsini and Krioukov, 2015; Asta and Shalizi, 2014; Smith, 2017, among others). A generic latent distance model in two-dimensional hyperbolic space can be written as follows:

$$\begin{aligned} Y_{ij} &\stackrel{\text{ind}}{\sim} \text{Bernoulli}(p_{ij}), \\ \text{logit}(p_{ij}) &= \alpha - d(z_i, z_j), \\ z_i &\in \mathbb{H}^2, \quad z_i \stackrel{\text{ind}}{\sim} f_2(z | \psi), \end{aligned} \quad (3.1)$$

where

$$\begin{aligned} d(z_i, z_j) &= d((r_i, \phi_i), (r_j, \phi_j)) \\ &= \text{acosh}\{\cosh(r_i)\cosh(r_j) \\ &\quad - \sinh(r_i)\sinh(r_j)\cos(\Delta\phi)\}, \end{aligned} \quad (3.2)$$

and

$$\Delta\phi = \pi - |\pi - |\phi_i - \phi_j||,$$

so that  $d(z_i, z_j)$  is the appropriate distance metric in hyperbolic space (using polar coordinates for convenience) and  $f_2(z | \psi)$  is a distribution for positions in 2-dimensional hyperbolic latent space (extending this specification to  $t$ -dimensional hyperbolic space is trivial). Of course, the exact form of this model and the choice of  $f_\lambda(z | \psi)$  can be altered—perhaps to respect particular network features or to accommodate different network modeling assumptions.

Krioukov et al. (2010) provides a thorough motivation for the use of hyperbolic space in the class of latent space models for network data, stemming from the observation that many real world networks tend to be tree-like and that the geometry of a tree behaves much like the geometry of hyperbolic space. Informally, trees can be thought of as discrete hyperbolic spaces since  $\mathbb{H}^2$  can be constructed as metrically equivalent to  $b$ -ary trees, trees with a branching factor of  $b$ . Further, the number of nodes at distance (path length)  $r$  in a  $b$ -ary tree grows exponentially. Thus, trees need an exponential amount of space for branching and only hyperbolic space is able to accommodate this since hyperbolic space is “bigger” than Euclidean space. For example, consider linearly increasing the radius of a sphere; in Euclidean space, the circumference of the sphere also grows linearly, but in hyperbolic space the circumference grows exponentially (see Table 1).

In this sense, hyperbolic space can be thought as having “more space” than Euclidean space and this fact has also been used to help visualize large networks (Lamping, Rao and Pirolli,

1995; Munzner, 1997). For example, visualizing a network in the Poincaré disk (see Appendix A and Figure 12 for a more detailed description of the Poincaré disk) is like viewing the network through a fish-eye lens, with nodes at the center of the disk in focus and displayed with the most detail while nodes that are further away are pushed towards the boundary of the disk and are diminished in size.<sup>4</sup> In fact, because hyperbolic space is “bigger” than Euclidean space, it cannot be exactly embedded within it. Thus, unlike elliptic space where we can, for example, easily picture the surface of a hypersphere in  $\mathbb{R}^3$ , there exist many models for visualizing hyperbolic space, each of which highlights different features of hyperbolic geometry (see Appendix A for more details). One popular frame of reference for 2-dimensional hyperbolic space is the hyperboloid model where points on the hyperbolic plane are modeled as points on the surface of the upper sheet of a hyperboloid in  $\mathbb{R}^3$ . This model can help us to visualize why hyperbolic space may be useful for understanding network data: Consider any two individuals in a social network and draw the shortest path between them. This path will most likely pass through some subset of individuals who are more centrally located in the network (e.g., individuals that have many connections, particularly to other highly connected individuals). Similarly, consider any two points on a hyperboloid; the shortest path between them must pass through the “more central” points located in the bottom or bowl of the hyperboloid, those points nearest the origin (we develop this idea further in Figure 5).

Krioukov et al. (2010) also provide a higher-level argument based on the observed hierarchical structure of many real world networks (Clauset, Moore and Newman, 2008). More specifically, note that networks, by definition, connect heterogeneous nodes. Consider the familiar example of a social network, where the nodes represent individuals who may have different genders, ages, family backgrounds, etc. This heterogeneity implies that some sort of classification or grouping exists. We can imagine grouping the nodes of a social network first by gender, then by age group within gender, etc., creating a hidden hierarchical structure. Finally, the relationships among these groups could be represented as a tree-like structure. Thus, networks can be thought of as (even approximate) trees, which we have already described as living in hyperbolic space. And in fact, many networks exhibit tree-like characteristics (see, e.g., Abu-Ata and Dragan, 2016). Note that this line of thought mirrors the justification for the use of ultrametric space provided by Schweinberger and Snijders (2003). In fact, ultrametric spaces are very closely related to hyperbolic space (see, e.g., Ibragimov, 2014). However, in generic hyperbolic space, the natural distance metric does not satisfy the strong triangle inequality.

The key aspect of hyperbolic space which allows for a tree’s exponential branching and which can accommodate complex network structure is its negative curvature. This directly influences the behavior of familiar geometric properties in hyperbolic space. For example, triangles in hyperbolic space appear skinny or thin, with the sum of their angles being less than  $180^\circ$  (see Table 1). Similarly, distances in hyperbolic space behave differently than

---

<sup>4</sup>Recently, some video games have taken advantage of this feature of hyperbolic space as well, including the maze-like game developed by Madore and the role playing game developed by Rogue and Rogue (2011) and described here Celinska and Kocczynski (2017); both games have online demonstrations and can help provide some intuition for movement and distances within a hyperbolic space.

those in Euclidean space. For example, in Figure 5 we simulate 500 points in 2-dimensional hyperbolic space (using the hyperboloid model of  $\mathbb{H}^2$  for visualization and simulating points uniformly<sup>5</sup> within a disk of radius  $R$ ), and connect all points within  $\gamma = R$  units of each other. This is equivalent to replacing the logit link function in model (3.1) with a Heaviside step function,  $p_{ij} = \Theta\{\gamma - d(z_i, z_j)\}$ , where the Heaviside step function is defined by

$$\Theta(x) = \begin{cases} 1 & x \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Of course, rather than deterministically connecting points within a certain distance, we could easily use the original model outlined in (3.1). However, although the probabilistic link function allows for greater variation in simulated networks, the Heaviside step function instead emphasizes the role of the distance metric in the latent space network model. In the left panel of Figure 5, we have also drawn circles of radius  $\gamma = R$  about two particular simulated points. First, it should be clear that distances behave differently in hyperbolic space, since these circles look very different than what we would expect in Euclidean space. Second, note that most of the lines or connections are radially oriented. In fact, if we imagine this simple simulation as a toy latent space network model, we see that this tendency towards radially oriented connections might create networks that have heterogeneous degree distributions and are more centralized. Points or nodes near the base of the hyperbola are the more central nodes in the network, with lots of connections to nodes further from the origin, near the periphery. These peripheral nodes have fewer connections to the less central nodes along the periphery (fewer connections to other nodes that are also far from the origin) and are most likely connected to the central nodes near the origin. This type of centralization is often observed in real world networks and, for the latent space model, is simply an artifact of the negative curvature of the latent hyperbolic space.

Similarly, we can mimic this construction in positively-curved space, embedding nodes on the surface of a 3-dimensional hypersphere as suggested by Hoff, Raftery and Handcock (2002) and McCormick and Zheng (2015). In the right panel of Figure 5, we simulate 500 points on the surface of the upper half of a 2-dimensional hypersphere (naturally visualized in 3-dimensional Euclidean space), again simulating points uniformly across this space. Again, we use the toy version of the latent distance model which incorporates the Heaviside step function, so that all points within  $\gamma = R$  units of each other are connected. Note that the hyperbolic space's striking pattern of radially oriented ties does not appear in positively-curved space, indicating that we will not see much centralization or heterogeneity of the degree distribution. Instead, the network simulated on the surface of the hypersphere appears to have more community structure, with sets of nodes that are highly interconnected but have fewer connections to other nodes outside of their set. To our knowledge, there has been no formal development of theoretical or practical reasons as to why positively-curved space might be particularly amenable to latent space models for network data. However, in McCormick and Zheng's (2015) application of this model to aggregated relational data, the

---

<sup>5</sup>See equation (3.3) and the following discussion for a more detailed description of the uniform distribution for discs in hyperbolic space.



authors mention a few appealing characteristics of working on the surface of a hypersphere. First, calculations are simplified by the fact that the space is finite so that there is no need to restrict latent positions to some subset of the space (i.e., a (finite) disk of radius  $R$  in (infinite) hyperbolic space). Second, because the space is finite, this provides a natural nonzero lower bound for the probability of a tie in the network. This assumption is more amenable to many practical settings, especially for aggregated relational data since the probability that an individual is tied to at least one member of a relatively large group of individuals is likely nonzero in practice.

Krioukov et al. (2010) derive a hyperbolic latent space model which directly connects model parameters to important features of network structure, but departs from our specification of a generic latent space model in (2.1). The authors outline the following model in  $\mathbb{H}^2$ :

$$\begin{aligned} Y_{ij} &\stackrel{\text{iid}}{\sim} \text{Bernoulli}(p_{ij}), \\ \text{logit}(p_{ij}) &= \frac{1}{2T} \{R(n, \bar{k}, \gamma, T) - d(z_i, z_j)\}, \\ r_i &\stackrel{\text{iid}}{\sim} p(r \mid R, \alpha(\gamma, T)), \\ \phi_i &\stackrel{\text{iid}}{\sim} \text{Uniform}(0, 2\pi), \end{aligned} \tag{3.3}$$

where the latent positions are specified in polar coordinates,  $z_i = (r_i, \phi_i)$  and distance between the  $i$ th and  $j$ th latent position is given by (3.2). The latent positions are assumed to be quasi-uniformly distributed within a disk of radius  $R$  centered at the origin,

$$\begin{aligned} p(r \mid R, \alpha) &= \alpha \frac{\sinh(\alpha r)}{\cosh(\alpha R) - 1}, \\ \alpha &= \begin{cases} \frac{1}{2}(\gamma - 1), & T \leq 1, \\ \frac{1}{2T}(\gamma - 1), & T > 1, \end{cases} \end{aligned}$$

which is roughly equivalent to an Exponential( $\alpha$ ) distribution. Note that  $p(r \mid R, \alpha)$  distributes points exactly uniformly within a disk of radius  $R$  when  $\alpha = 1$  (see Figure 6 for examples of other values of  $\alpha$ ). In the model, the radius of this disk is a function of expected network characteristics,  $R \equiv R(n, \bar{k}, \gamma, T)$  where  $\bar{k}$  is the expected average degree,  $\gamma$  is the exponent of the power-law degree distribution, and  $T$  controls the amount of clustering, and satisfies the following equation:

$$\begin{aligned} \bar{k} &= \frac{N}{\pi} \int_0^R p(r_1 \mid R, \alpha) \int_0^R p(r_2 \mid R, \alpha) \\ &\times \int_0^\pi \text{logit}(p_{ij}) d\phi_1 dr_1 dr_2. \end{aligned}$$

In practice, numerical integration is used to calculate  $R$  which complicates potential inference for this particular specification of the model. Note that this formulation is similar to the generic model outlined in (2.1), with the exception of the baseline tendency for tie formation,  $\alpha$ . Here, model (2.1)'s  $\alpha$  is given a very specific form and in fact depends on an important feature of the latent space—its radius or bound,  $R$ —which itself is a function of

parameters which correspond to important features of network data. In this sense, model (2.1) 's parameter for a baseline tendency for tie formation has been replaced by a term that represents the complicated interplay between features of the latent space and expected network properties.

## 4. IMPLICATIONS OF LATENT SPACE GEOMETRY ON NETWORK SUMMARY STATISTICS

Of course, Krioukov et al.'s (2010) line of reasoning—that the negative curvature of hyperbolic geometry accommodates complex network structure—leads to a rather natural question: what role does the curvature of the latent space play in latent space models for network data? By considering networks simulated from latent distance models in spaces of varying curvature, we will explore the function of geometric curvature in the latent space model. Further, we show that, for the purposes of building a hierarchical network model, changing the geometry of the latent space can parsimoniously grow the complexity of the latent space network model while maintaining the intuitive appeal of using distance to model dependence.

First, note that the curvature of a space is determined by the extrinsic radius,  $r$ ; here, we are examining a hyperbolic space with curvature  $r = -1$ , a Euclidean space with curvature  $r = 0$ , and an elliptic space with curvature  $r = 1$  (see Figure 7). For consistency, we will define  $\mathcal{Z}$  to be a disk with intrinsic radius,  $R$ , centered at the origin (or its analogue<sup>6</sup>) in each of the three latent spaces and will assume that the latent positions are uniformly distributed within this disk. For any geometric space,  $X$ , a disk with intrinsic radius  $R$  is the set of all points in  $X$  having distance less than or equal to  $R$  from a specific fixed point. As a result, all three latent spaces,  $\mathcal{Z}$ , have the same diameter (maximum distance between pairs of points). This corresponds to the latent spaces plotted in Figure 7, each of which has the same intrinsic radius but different curvatures. To better visualize these spaces, we have drawn Figure 7 with intrinsic radius  $R = \pi/2$ , but in the some of the simulations that follow we will use  $R = \pi$ , resulting in an elliptic latent space that is the entire sphere. Note that  $R$  controls how much of each latent space,  $X$ , is being used in the model as  $\mathcal{Z}$ ; latent positions are sampled only within disks of intrinsic radius  $R$ .

Specifying latent spaces that are comparable across the three geometries is essential to ensuring fair comparisons of any resulting network features. However, doing so has required that we deviate from some of the standard assumptions of popular CLS models; for example, in the Hoff, Raftery and Handcock (2002) model, latent positions are drawn from a bivariate standard normal distribution. The relationship between the *distribution* of the latent positions in a given latent space (as opposed to the *geometry* of the latent space) and emergent network features is not the focus of this paper and remains an interesting topic for further research. We will rely on the comparable latent spaces defined above and depicted in Figure 7 (with some modification to the value of  $R$ ) in the simulations described in this section (and

---

<sup>6</sup>To be concrete, in Euclidean space the disk is centered about  $(0, 0)$ . Using the models of elliptic and hyperbolic space in  $\mathbb{R}^3$  (see Figure 7), the disks in elliptic and hyperbolic space are each centered about  $(0, 0, 1)$ .

depicted in Figures 8, 9 and 10) as well as in Section 5 (and Figure 11 included there as well as Figure 13 in the Appendix).

Although we might consider directly examining simulated networks under each of these models, binary data can be particularly difficult to study and binary relational data (i.e., networks) are no exception. For example, in the latent space model, even if  $p_{ij}$  smoothly varies across potential ties in the network, this smoothness cannot be manifested in the observed network, since observed ties are by definition either present or absent. Thus, in our simulation study, we use a Heaviside step link function as was used in Figure 5, again to emphasize the roles of the different distance metrics in each latent geometry. Here, we consider smoothly varying  $\gamma$ , the Heaviside step function cut point, and examining differences in network structure across three latent geometries: Euclidean space, hyperbolic space and elliptic space. Letting  $i, j = 1, \dots, n$  index the nodes in the network, the model used in this simulation exercise is thus as follows:

$$\begin{aligned} Y_{ij} &= \Theta\{\gamma - d(z_i, z_j)\}, \quad i < j \\ z_i &\in \mathcal{X}^t, z_i \stackrel{\text{iid}}{\sim} \text{Uniform}(R) \end{aligned} \quad (4.1)$$

which is equivalent to replacing the logit link function in model (3.1) with the Heaviside step function. In this case, the tie variables,  $Y_{ij}$ , can still technically be written as Bernoulli variables but are actually deterministically, rather than probabilistically, specified given the set of latent positions. Of course, while these choices greatly simplify the model originally outlined by (3.1), they allow us to more directly consider the role of geometry in the latent space model for network data.

First, it is worth noting that the typical distribution of distances varies greatly across these geometries. For example, consider distances from a set of points simulated under each model<sup>7</sup> in Figure 8. The distribution of distances for latent positions on the hyperboloid looks very different than for those in the other spaces. In hyperbolic space, there are few points that are very close together and the distribution of distances is left skewed. This distribution likely better represents our intuition about the distances between latent positions for nodes in real world social networks—any randomly selected pair of individuals is not likely to have a tie (i.e., any pair of positions is likely far apart). In both Euclidean<sup>8</sup> and elliptic space, the distribution of distances is much more symmetric. Recall that we have chosen the same distribution for the latent positions (uniform within a disk of radius  $R$ ) across all three geometries, so that observed differences in the distributions of the distances in Figure 8 is due to the differing curvature of the spaces. However, to better illustrate the differences between these geometric spaces, in the simulations described in this section we have increased the intrinsic radius to  $R = 7.5$  in hyperbolic space (leaving  $R = \pi/2$  in Euclidean and elliptic space<sup>9</sup>), since many interesting features of this space are more striking further away from the origin.

<sup>7</sup>Variation across simulations does not appear to impact the conclusions drawn here.

<sup>8</sup>Although not depicted here, using a standard normal distribution (or even altering the standard deviation so that the maximum expected distance between latent positions mimics that of elliptic space) in Euclidean space (as in Hoff, Raftery and Handcock, 2002) results in a distribution of pairwise distances that is right-skewed and a degree distribution that implies that most individuals are very social, but there may be a few actors with very few ties.

Naturally, one can imagine how these differences in the distribution of distances might affect a network's degree distribution (see the center row in Figure 8). In the social network setting, a hyperbolic space implies that most individuals are relatively asocial with only a few very popular actors; while a Euclidean or elliptic space implies that there are some loners (some actors with very few ties) and some popular actors, but most actors display roughly similar levels of sociality. Of course, our description of the intuitive degree distribution in a hyperbolic latent space is a natural way of describing degree heterogeneity, a long observed trait of many real world networks.

The results of our simulation study where we vary the cut point in the Heaviside step function,  $\gamma$ , and examine changes in network features is displayed in Figures 9 and 10. Even after specifying comparable latent spaces,<sup>10</sup> we still observe some variation in the range of pairwise distances across the latent geometries (consider the  $x$ -axes of the density plots for Euclidean and elliptic space, where  $R = \pi/2$  in each case, in Figure 8). In our simulation exercise, we accommodate for this difference in scaling by dividing the distances by the maximum simulated distance in each simulation, so that all distances are contained within  $[0, 1]$ . As a result of this scaling, note that as  $\gamma$  increases, the density of each network will increase according to a smooth linear trend which is uniform across all three spaces. This is a result of the fact that the latent positions in each space are uniformly distributed and ties in the network depend directly on the natural distance metric in each space. In this sense, the scaling adjusts for the slight variation in the range of simulated pairwise distances and also results in networks that are comparable in terms of network density, for any particular  $\gamma$ . We consider networks of various size ( $n = 20, 50, 100$ ) embedded in Euclidean, hyperbolic, and elliptic latent geometries. For each geometry and network size considered, we simulate 5000 networks according to model (4.1), and slowly increase  $\gamma$  from 0 to 1 by increments of 0.20. We plot the average of common network summary measures<sup>11</sup>—clustering, average path length, degree centrality, betweenness centrality, closeness centrality, and modularity<sup>12</sup>—of the resulting simulated networks as a function of  $\gamma$  in Figures 9 and 10. Bands on these plots are pointwise minima and maxima across the simulations—the jaggedness of these bands is a consequence of finite size effects for the simulated networks.

In Figure 9, we compare networks in hyperbolic space to those in Euclidean space. Note that networks in hyperbolic space can achieve much higher levels of degree centrality, betweenness and closeness, and that this effect is more prominent for larger networks. This coincides with our previous intuition for the behavior of the degree distribution of networks in hyperbolic space. Note also, that this improved ability to model degree heterogeneity does not come at the expense of any dramatic losses in clustering or average path length. In

<sup>9</sup>Note that in elliptic space,  $R$  is bounded above by  $\pi$ , so increasing  $R$  and keeping  $R$  equal across all three geometries is not possible

<sup>10</sup>As for Figure 8, we again use an increased  $R = 7.5$  in hyperbolic space in order to fully capture the features of this space.

<sup>11</sup>All network statistics were computed using the `igraph` (Csardi and Nepusz, 2006) and `network` (Butts, 2008) packages in R (R Core Team, 2016).

<sup>12</sup>To provide a quick sense of any observed community structure, we use the `cluster_edge_betweenness` (Newman and Girvan, 2004) and `modularity` (Clauset, Newman and Moore, 2004) function provided in the `igraph` (Csardi and Nepusz, 2006) package for R (R Core Team, 2016). This modularity statistic is included simply to provide an exploratory look at the extent to which community structure may be observed across the latent geometries, particularly for elliptic space. In practice, any reliable confirmatory conclusions would require more detailed consideration of the community detection model, including the expected number and size of the communities, and modularity measure for any particular set of network data.

Figure 10, we see that networks in elliptic space do not differ greatly from those in Euclidean space in terms of these network statistics.<sup>13</sup>

## 5. USING GRAPH LAPLACIANS TO IDENTIFY AN APPROPRIATE GEOMETRY

As mentioned above, many common descriptive summary measures of network structure are sensitive to the size of the network. Thus, while the behavior of these measures across latent spaces of different geometries can provide us with some indication of the role of geometry in these models, it is difficult to formally relate these statistics to properties of the latent space. This issue is rooted in the fact that it is difficult to find a common language describing both networks and geometric spaces; such a common language is a prerequisite for comparing properties between the two sorts of structures. However, as we will examine in this section, both the geometry of a space and the combinatorics of a graph or network can be described by linear operators whose eigenvalues are comparable.

For geometric spaces, the Laplace-Beltrami operator,  $\nabla_M^2$ , is an operator on smooth functions which map general Riemannian manifolds,  $M$ , to  $\mathbb{R}$ . For a function,  $f = f(x_1, x_2, \dots, x_n)$ , when  $M = \mathbb{R}^n$ , the Laplace-Beltrami operator can be expressed in terms of partial derivatives:

$$\nabla_{\mathbb{R}^n}^2 f = \sum_{i=1}^n \frac{\partial^2 f}{\partial x_i^2}, \quad (5.1)$$

and hence when  $M = \mathbb{R}$ , the Laplace-Beltrami operator is simply the second derivative,  $\nabla_{\mathbb{R}}^2 f = f''$ . Even for general Riemannian manifolds,  $M$ , the Laplace-Beltrami operator,  $\nabla_M^2$ , is essentially defined by (5.1) with respect to local coordinates  $x_1, x_2, \dots, x_n$ . The eigenvalues of the Laplace-Beltrami operator describe many (though not all) geometric properties of the space  $M$ . The first positive eigenvalue of the Laplace-Beltrami operator is positively related to the curvatures of certain compact manifolds (Lichnerowicz, 1958). More generally, a conjectural generalization of Weyl's Law (Weyl, 1911) describes a relationship between the volumes of compact domains  $D$  in any Riemannian manifold  $M$  and the eigenvalues of the Laplace-Beltrami operator  $\nabla_M^2$ , restricted to those functions with support  $D$  (Ivrii, 1980). Thus, in particular settings, we can describe a direct relationship between the eigenvalues of the Laplace-Beltrami operator and much of the geometry of a space.

The *graph Laplacian*  $\nabla_G^2$  of a finite graph  $G$  is the operator on functions  $V_G \rightarrow \mathbb{R}$ , from nodes  $V_G$  defined by the rule

<sup>13</sup>Although not displayed here, there are some slight differences between Euclidean and elliptic space when latent positions are normally distributed within Euclidean space, as in the Hoff, Raftery and Handcock (2002) model. In this case, we observe some slight increases in the degree heterogeneity in elliptic space, but this difference is not quite as exaggerated as in hyperbolic space (particularly for  $n = 100$ ). Recall that in the simulated network on the surface of the hypersphere in Figure 5, we observed some indications of a community structure, with distinct groups of nodes having more connections within the group and fewer connections across groups. And in fact, we see some difference in the modularity statistics across the two spaces.

$$(\nabla_G^2 \phi)(v) = \deg(v)\phi(v) - \sum_{(v,w) \in E_G} \phi(w).$$

After ordering vertices from 1 to  $n$ ,  $\nabla_G^2$  can be expressed by

$$\nabla_G^2 = D_G - A_G,$$

where  $A_G$  is the adjacency matrix of  $G$  and  $D_G$  is the diagonal matrix whose entries are the degrees of vertices of  $G$ . We note that there are other versions of the graph Laplacian, including the symmetric normalized graph Laplacian,  $I - D_G^{-1/2} A_G D_G^{-1/2}$ , with which the reader may be familiar.

The graph Laplacian is a discrete analogue of the Laplace–Beltrami operator in the following sense. Just as the Laplace–Beltrami generalizes the second derivative, the graph Laplacian generalizes a *difference of differences*. For example, let  $I$  be the discrete, graph-theoretic approximation

$$v_1 \text{ --- } v_2 \text{ --- } \dots \text{ --- } v_n,$$

of the open interval  $(0, 1)$ , where horizontal lines represent connections between the vertices of  $I$ ,  $v_i$  for  $i = 1, \dots, n$ . Just as  $\nabla_{(0,1)}^2 f = f''$ ,  $\nabla_I^2 \phi$  is defined on the inner  $(1 < i < n)$  vertices  $v_i$  by

$$(\nabla_I^2 \phi)(v_i) = (\phi(v_i) - \phi(v_{i-1})) - (\phi(v_{i+1}) - \phi(v_i)).$$

The eigenvalues of  $\nabla_G^2$  describe many (though not all) combinatorial features of a graph  $G$ . For example, the algebraic multiplicity of the eigenvalue 0 reveals how many connected components  $G$  has. The second smallest eigenvalue of  $\nabla_G^2$  quantifies how interconnected the nodes in  $G$  are.

At least intuitively, the Laplace–Beltrami operator  $\nabla_M^2$  is some sort of limiting approximation of graph Laplacians  $\nabla_G^2$  of graphs  $G$  sampled from a latent distance model whose latent space is the Riemannian manifold  $M$ . Unfortunately, the closest sort of theoretical results of this nature (cf. Belkin and Niyogi, 2005; Ting, Huang and Jordan, 2011; Hein, Audibert and von Luxburg, 2007) use a modification of  $\nabla_G^2$  when the edges of  $G$  are weighted by distances (e.g., Euclidean point-cloud data.)

We test this intuition with some simulations. Using the comparable latent spaces defined in Section 4, we consider models where nodes are uniformly distributed within a disk of radius  $R = \pi$  in Euclidean, hyperbolic and elliptic latent space. For each real world social network

we examine, we generate 100 simulated networks from model (2.4) (with  $\alpha = 0$ ), each with the same number of nodes as the real world social network, for each CLS model.<sup>14</sup>

In Figure 11 moving from top to bottom and left to right, we consider the following real world network examples: the famous Zachary's Karate Club network which records social interactions of  $n = 34$  members of a karate club in the 1970s (Zachary, 1977); the canonical Florentine marriage network in which the  $n = 16$  nodes represent powerful families in Renaissance Florence and ties represent marital alliances (Padgett, 1994; Butts, 2008); Coleman's high school network in which the  $n = 73$  nodes represent boys in a small high school in Illinois during the fall of 1957 and ties are self-reported friendships (Coleman, 1975; Butts, 2016); the neural network of the *Caenorhabditis elegans* (*C. elegans*) worm where the  $n = 297$  nodes represent neurons which are tied if at least one synapse or gap junction exists between them (Watts and Strogatz, 1998; Opsahl, 2009); the road transportation network for Chicago where the  $n = 1467$  nodes represent traffic zones and ties represent links by freeways or other roads (Eash et al., 1979; Kunegis, 2013); and the power grid of the western U.S. where the  $n = 4941$  nodes represent either a generator, a transformer or a substation and ties represent a power supply line (Watts and Strogatz, 1998; Kunegis, 2013). For each of these networks, plots of the eigenvalues in decreasing order suggest that the real world networks exhibit geometric features strongly reminiscent of hyperbolic space. Further, for the nonsocial networks (see the networks in the right panel of Figure 11), note that the eigenvalues for these networks are pushed even further towards the bottom left side of the plots. This indicates that perhaps these networks require latent spaces that are more negatively curved or take better advantage of the negatively curved space (i.e., have a larger  $R$ <sup>15</sup>). See Appendix B for a discussion of how hyperbolic features manifest in networks simulated from common, intuitive models for network data.

## 6. DISCUSSION

As we have demonstrated in this simulation exercise, using a negatively curved latent space allows us to smoothly grow the complexity of the latent space model (i.e., accommodate more complicated network structure, as was expected from the intuitive arguments provided by Krioukov et al., 2010 and as was demonstrated in Figure 9) without adding any additional structure to the model, such as explicit functions of network statistics as is done in ERGMs (Frank and Strauss, 1986; Snijders et al., 2006), or node-level random effects (Hoff, 2005; Krivitsky et al., 2009), or mixture distributions to spatially cluster the latent positions (Handcock, Raftery and Tantrum, 2007). Furthermore, models in hyperbolic latent space (such as (3.3) and (4.1)) preserve the use of distance to model dependence, a common feature of many popular statistical models, and is able to draw upon the long success of latent space models for network data. We have provided an illustration of how a hyperbolic

<sup>14</sup>In these simulations, we purposefully utilize a relatively simplistic CLS model in order to highlight differences due solely to the geometry of the latent spaces; better model fit could most likely be obtained by using models that can incorporate more features specific to the real world network being considered, such as the Hoff, Raftery and Handcock (2002) model or model (3.3) (Krioukov et al., 2010), for example.

<sup>15</sup>As mentioned in Section 4, many interesting features of hyperbolic space are more striking further away from the origin. Recall that in the simulations for Figure 11, we set  $R = \pi$  across all spaces to make the latent spaces as comparable as possible; in elliptic space,  $R$  is bounded above by  $\pi$ . Increasing  $R$ , perhaps according to model (3.1) or the approximate formula provided in Krioukov et al. (2010), could result in networks that better match the observed eigenvalues for these nonsocial networks.

latent space can produce generated networks which align more closely to real-world networks in local and global structure (e.g., Figure 11), as opposed to the popular choice of a Euclidean latent space. A more thorough investigation to fully understand in which particular cases (i.e., under what kinds of assumptions about the network data generating mechanism) this holds is a promising avenue for future research.

The study of the types of graphs generated by latent space network models across different geometries, which we have provided here, should prove very helpful to practitioners in many fields. Note that the viewpoint we have adopted here is commonly used in Bayesian analysis, for example in the assessment of posterior predictive distributions, where we are interested in understanding the behavior of a model through the data that it generates. This view of network analysis can help improve practitioners' understanding of what kind of behavior to expect from different network models and, thus, which model may be most appropriate in a particular practical setting or application.

Of course, many open questions about the relationship between geometric curvature and network properties remain unanswered. For example, what is the precise role of the distribution for the latent positions in a geometric space of a particular curvature? Future work might also consider continuing this investigation by considering a parametric family of spaces parametrized by their curvature,  $\delta$ : from spaces of negative curvature ( $\delta < 0$ ) to Euclidean space ( $\delta = 0$ ) to spaces with positive curvature ( $\delta > 0$ ). In this setting, we might imagine modeling observed networks by averaging over features of networks from both hyperbolic and elliptic space, for example. We might also imagine performing inference for  $\delta$  and using estimates of  $\delta$  to describe properties of the observed network, or perhaps, to make comparisons across a set of observed networks.

Another exciting line of future research might consider a flipped perspective on the discussion provided here. More specifically, given a particular set of network data can we use summary measures of observed network structure and/or graph Laplacians to infer the appropriate latent space geometry. This type of approach might prove useful in model selection.

## ACKNOWLEDGMENTS

Smith and Calder were partially supported by grants from the National Science Foundation (NSF DMS-1209161), the National Institutes of Health (NIH R01 HD088545) and The Ohio State University Institute for Population Research (NIH P2CHD058484).

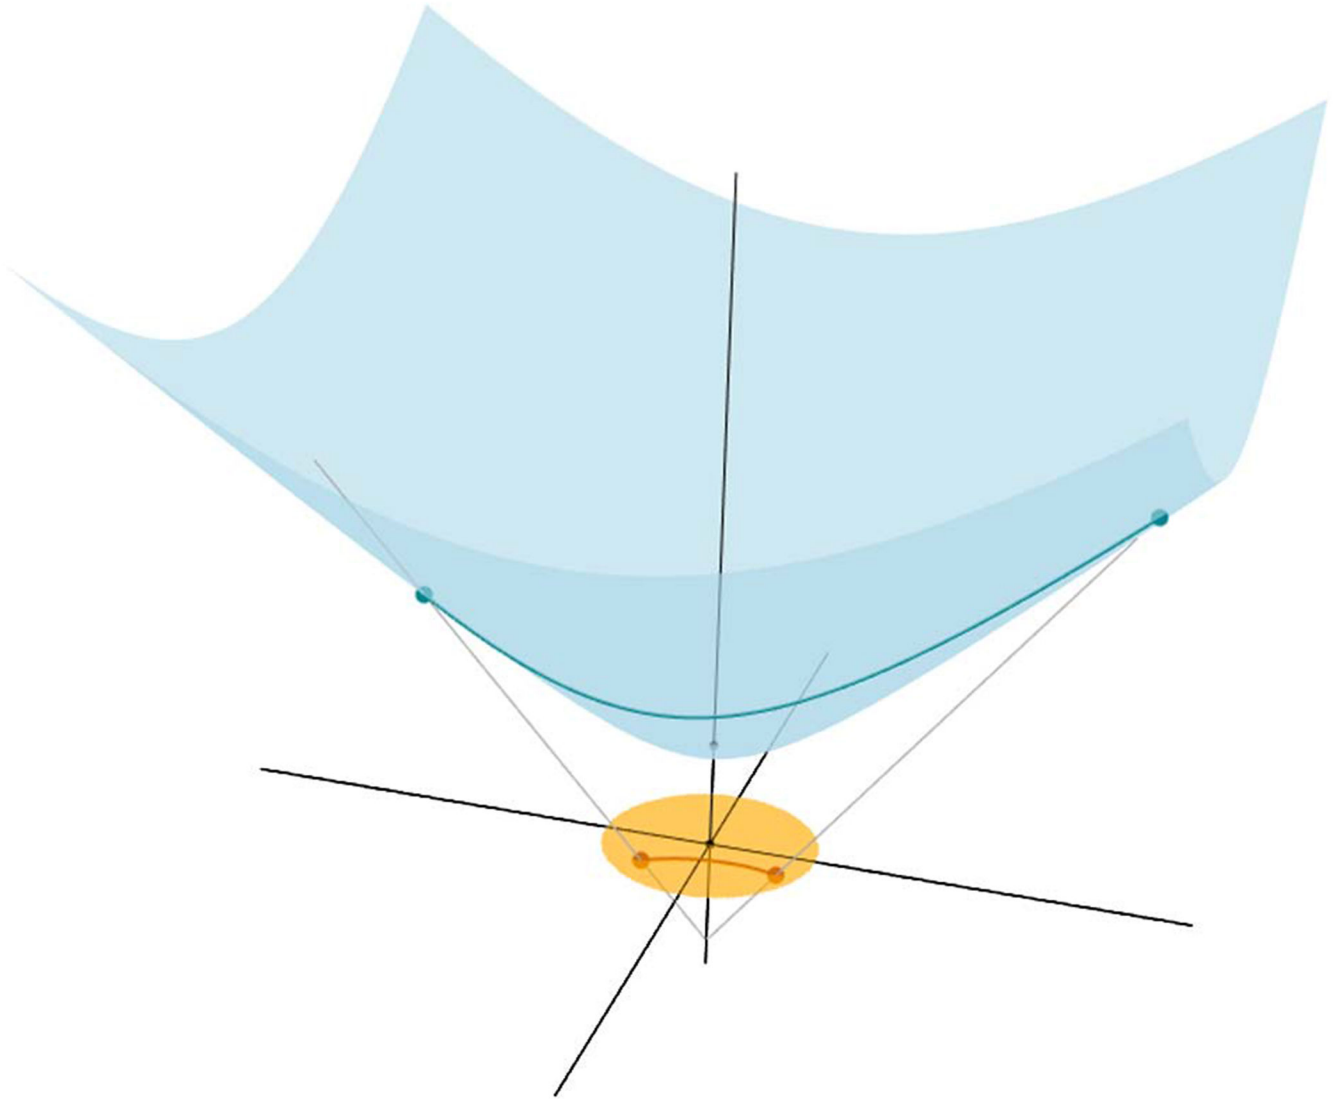
## APPENDIX A:: MODELS FOR VISUALIZING HYPERBOLIC SPACE

First, note that hyperbolic space is an example of a non-Euclidean space. Recall that Euclidean space is described by Euclid's five postulates:

1. Each pair of points can be joined by one and only one straight line segment.
2. Any straight line segment can be indefinitely extended in either direction.
3. There is exactly one circle of any given radius with any given center.
4. All right angles are congruent to one another.



5. Given any straight line and a point not on it, there exists one and only one straight line which passes through that point and never intersects the first line, no matter how far they are extended.

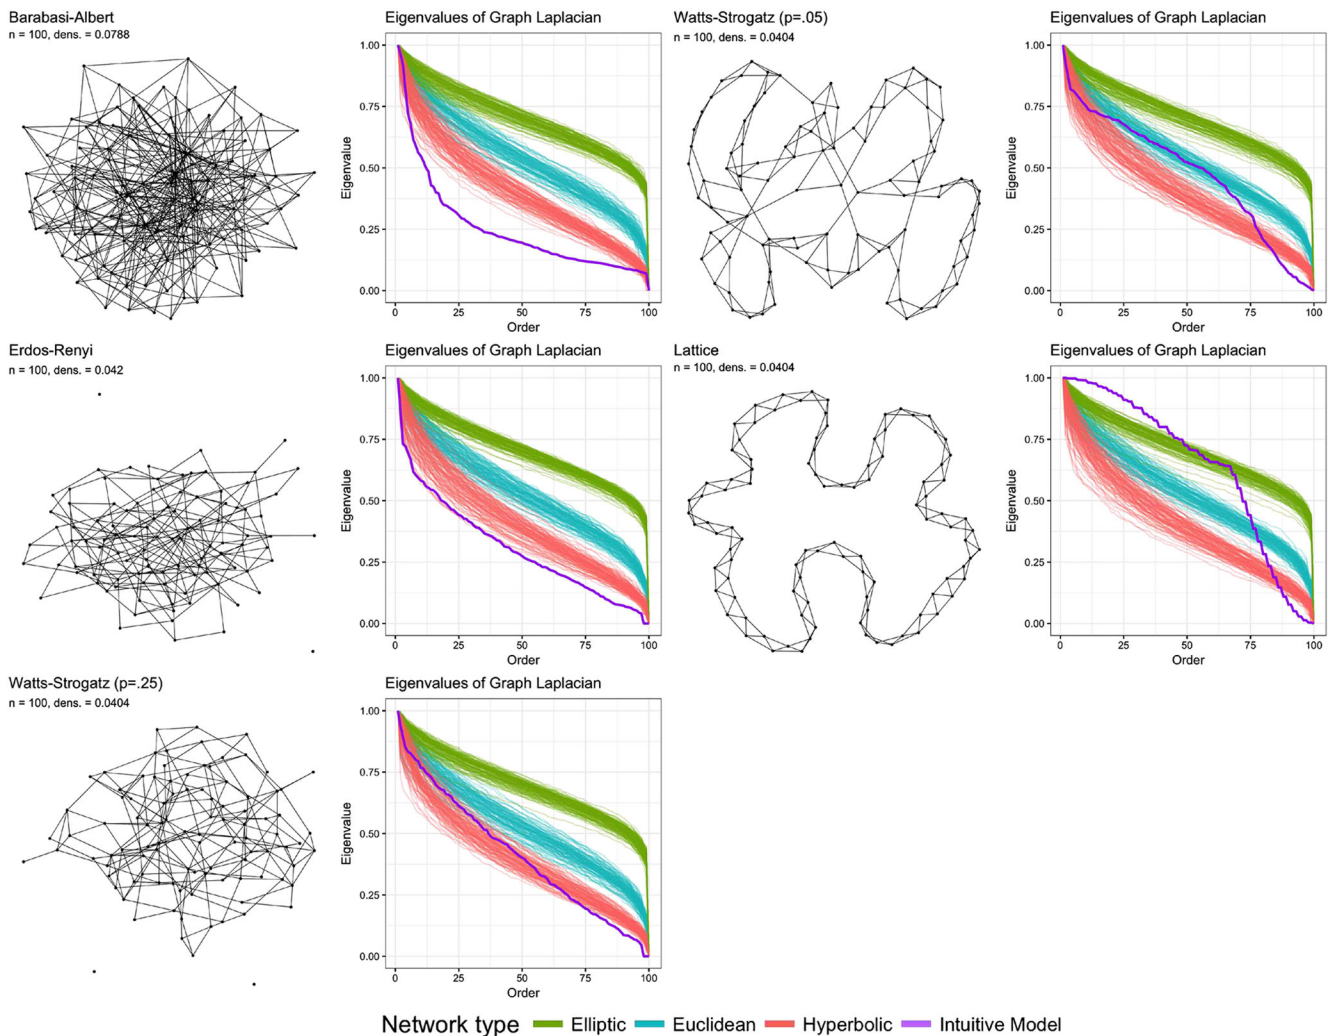


**Fig. 12.** Visualizing hyperbolic space. Turquoise points are plotted on the hyperboloid model (light blue), where the shortest distance between two points is a line across the surface of the hyperboloid. Orange points are plotted on the Poincaré model, where the shortest distance between two points is the arc of the circle which intersects the two points and is perpendicular to the boundary. Points in the Poincaré model are stereographic projections of points in the hyperboloid model—imagine sitting at  $z = -1$  and looking through a disk of radius one, centered at  $z = 0$ , up at the points on the hyperboloid (grey lines).

where the last postulate is referred to as the “parallel postulate.” Non-Euclidean spaces are formed by replacing the parallel postulate with some alternate behavior. In hyperbolic space, it is replaced with the following:

- 5\*. Given any straight line and a point not on it, there are at least two distinct lines passing through that point which do not intersect with the line.

Thus, in hyperbolic space, there exists infinitely many parallel lines (see Table 1). In order to visualize the geometric properties of hyperbolic space, there are four popular analytic models: the Beltrami–Klein model, the Poincaré disk model, the Poincaré half-plane model and the hyperboloid model (although other models certainly exist). Each of these models define a hyperbolic plane which satisfies the axioms of a hyperbolic geometry (1, 2, 3, 4 and 5\*). Moving forward, we will



**Fig. 13.** *Intuitive models of network structure. In these plots, networks simulated from intuitive models of network structure—the Barabási–Albert small world model, the Erdős–Rényi model, the Watts–Strogatz model and a perfect lattice—are compared to networks simulated*

from our general latent distance model (Equation (2.4)) across different geometries. Each row includes a network plot and a plot of the eigenvalue curves, where the color corresponds to the geometry of the latent space. In the eigenvalue plots, each network is represented by a curve which plots the network's eigenvalues against their order. Parameters for the intuitive network models are specified such that each network has roughly the same density.

primarily rely on the hyperboloid model to visualize positions in hyperbolic space but will sometimes refer to the Poincaré disk model as well.

## Hyperboloid model.

Although specified for  $n$ -dimensional hyperbolic space, we will typically focus on a 2-dimensional hyperbolic space,  $\mathbb{H}^2$ , and describe this version of the model here. Points on the hyperbolic plane are modeled as points on the surface of the upper (or lower) sheet of a hyperboloid in 3-dimensional Euclidean space,  $\mathbb{R}^3$  (consider the blue points in Figure 12). That is, all points  $(x, y, z)$  in the plane satisfy the following:

$$z^2 - y^2 - x^2 = 1,$$

with  $z > 0$ . Distance between points in the hyperboloid model is simply distance along the surface of the hyperboloid; this mimics visualization of distance in elliptic space, as the distance (arc length) along the surface of the hypersphere.

## Poincaré disk model.

Unlike the hyperboloid model, the Poincaré disk model represents (infinite) hyperbolic space within a finite space. This model is formed by the stereographic projection of the hyperboloid in  $\mathbb{R}^3$  onto the unit disc at  $z = 0$ . As points in this model move closer to the boundary of the disk, they approach infinity; in the hyperboloid model, this is equivalent to moving further up the surface of the hyperboloid (consider the orange points in Figure 12). Distances in this model are semicircles perpendicular to the boundary of the Poincaré disk.

## APPENDIX B:: HYPERBOLIC FEATURES IN INTUITIVE NETWORK MODELS

Perhaps non-CLS models for network data also incorporate some of the same features we observed in the networks simulated in a hyperbolic space. In order to examine this idea, we replicate the analysis used to create Figure 11 with networks simulated from common intuitive models for networks. Here, we examine the Barabási–Albert model (Barabási and Albert, 1999), the Erdős–Rényi model (Erdős and Rényi, 1960), the Watts–Strogatz model (Watts and Strogatz, 1998) and a perfect lattice.<sup>16</sup> The eigenvalue plots indicate that many of these network models capture the same kind of features we observe for networks in a

<sup>16</sup>For the Barabási–Albert model, the number of edges added in each time step is 4 (one time step is used to simulate the network); for the Erdős–Rényi model, the probability of a tie is 0.08080808; for the Watts–Strogatz model, a neighborhood of 2 vertices of the lattice are initially connected and the rewiring probability is set to 0.25 and 0.05, respectively. All networks are simulated with  $n = 100$  nodes using standard functions in the `igraph` package for R and plotted using the default settings of the `gnet` function in the `GGally` package for R.

hyperbolic space, particularly for models that create highly heterogeneous networks (like the Barabási–Albert model) or highly random networks (like the Erdős–Rényi model and the Watts–Strogatz model with larger rewiring probability). For example, the Barabási–Albert model uses a preferential attachment mechanism to generate scale-free networks that are characterized by highly right-skewed degree distributions, as we observed in hyperbolic space in Figure 8. However, the perfect lattice and some variations of the Watts–Strogatz model (which modifies a perfect lattice with random rewirings) provide nice nonexamples; these highly regular networks, which are neither very heterogeneous nor very random (see Figure 3 in Solé and Valverde, 2004, for more discussion), do not exhibit hyperbolic features.

## REFERENCES


- Abbe E, Bandeira AS and Hall G (2016). Exact recovery in the stochastic block model. *IEEE Trans. Inform. Theory* 62 471–487. [MR3447993](#)
- Abu-Ata M and Dragan FF (2016). Metric tree-like structures in real-world networks: An empirical study. *Networks* 67 49–68. [MR3439658](#)
- Airoldi EM, Blei DM, Fienberg SE, Goldenberg A, Xing EP and Zheng AX (2008a). Statistical Network Analysis: Models, Issues, and New Directions: ICML 2006 Workshop on Statistical Network Analysis, Pittsburgh, PA, USA, June 29, 2006, Revised Selected Papers 4503. Springer, Berlin.
- Airoldi EM, Blei DM, Fienberg SE and Xing EP (2008b). Mixed membership stochastic blockmodels. *J. Mach. Learn. Res* 9 1981–2014. [PubMed: 21701698]
- Aitchison J, Barceló-Vidal C, Martín-Fernández J and Pawłowsky-Glahn V (2000). Logratio analysis and compositional distance. *Mathematical Geology* 32 271–275.
- Aldecoa R, Orsini C and Krioukov D (2015). Hyperbolic graph generator. *Comput. Phys. Commun* 196 492–496.
- Amini AA, Chen A, Bickel PJ and Levina E (2013). Pseudo-likelihood methods for community detection in large sparse networks. *Ann. Statist* 41 2097–2122. [MR3127859](#)
- Asta D and Shalizi CR (2014). Geometric network comparison. Under review. Preprint. Available at [arXiv:1411.1350](#).
- Barabási A-L and Albert R (1999). Emergence of scaling in random networks. *Science* 286 509–512. [MR2091634](#) [PubMed: 10521342]
- Bartholomew D, Knott M and Moustaki I (2011). *Latent Variable Models and Factor Analysis: A Unified Approach*, 3rd ed. Wiley Series in Probability and Statistics. Wiley, Chichester [MR2849614](#)
- Belkin M and Niyogi P (2005). Towards a theoretical foundation for Laplacian-based manifold methods In *Learning Theory*. Lecture Notes in Computer Science 3559 486–500. Springer, Berlin [MR2203282](#)
- Bickel PJ and Chen A (2009). A nonparametric view of network models and Newman-Girvan and other modularities. *Proc. Natl. Acad. Sci. USA*. 106 21068–21073. [PubMed: 19934050]
- Butts CT (2008). Network: A package for managing relational data in R. *J. Stat. Softw* 24.
- Butts CT (2016). sna: Tools for Social Network Analysis. R package version 2.4. <https://CRAN.R-project.org/package=sna>.
- Celinska D and Kopczyński E (2017). Programming languages in GitHub: A visualization in hyperbolic plane. In *Proceedings of the Eleventh International AAAI Conference on Web and Social Media (ICWSM 2017)* 727–728 Association for the Advancement of Artificial Intelligence Menlo Park, CA.
- Chen K and Lei J (2018). Network cross-validation for determining the number of communities in network data. *J. Amer. Statist. Assoc* 113 241–251. [MR3803461](#)

- Clauset A, Moore C and Newman MEJ (2008). Hierarchical structure and the prediction of missing links in networks. *Nature* 453 98–101. [PubMed: 18451861]
- Clauset A, Newman ME and Moore C (2004). Finding community structure in very large networks. *Phys. Rev. E* 70 066111.
- Coleman JS (1964). *Introduction to Mathematical Sociology*. Free Press, New York.
- Csardi G and Nepusz T (2006). The igraph software package for complex network research. *Inter Journal Complex Systems* 1695.
- Davis JA (1970). Clustering and hierarchy in interpersonal relations: Testing two graph theoretical models on 742 sociomatrices. *Am. Sociol. Rev* 843–851.
- Diaconis P and Janson S (2008). Graph limits and exchangeable random graphs. *Rend. Mat. Appl* (7) 28 33–61. [MR2463439](#)
- Doveton H (1998). Beyond the perfect martini: Teaching the mathematics of petrological logs. In *Proceedings of IAMG98, The Fourth Annual Conference of the International Association for Mathematical Geology 71–75* ACM Press/Addison-Wesley Co., De Frede, Naples.
- Eash R, Chon K, Lee Y and Boyce D (1979). Equilibrium traffic assignment on an aggregated highway network for sketch planning. *Transportation Research* 13 243–257.
- Erdős P and Rényi A (1960). On the evolution of random graphs. *Magy. Tud. Akad. Mat. Kut. Intez. Kozl* 5 17–61. [MR0125031](#)
- Fienberg SE and Wasserman SS (1981). Categorical data analysis of single sociometric relations. *Sociol. Method.* 12 156–192.
- Frank O and Strauss D (1986). Markov graphs. *J. Amer. Statist. Assoc* 81 832–842. [MR0860518](#)
- Gao C, Lu Y and Zhou HH (2015). Rate-optimal graphon estimation. *Ann. Statist* 43 2624–2652. [MR3405606](#)
- Gilbert EN (1959). Random graphs. *Ann. Math. Stat* 30 1141–1144. [MR0108839](#)
- Goodreau SM, Kitts JA and Morris M (2009). Birds of a feather, or friend of a friend? Using exponential random graph models to investigate adolescent social networks. *Demography* 46 103–125. [PubMed: 19348111]
- Handcock MS, Raftery AE and Tantrum JM (2007). Model-based clustering for social networks. *J. Roy. Statist. Soc. Ser. A* 170 301–354. [MR2364300](#)
- Hein M, Audibert J-Y and von Luxburg U (2007). Graph Laplacians and their convergence on random neighborhood graphs. *J. Mach. Learn. Res* 8 1325–1368. [MR2332434](#)
- Hoff PD (2005). Bilinear mixed-effects models for dyadic data. *J. Amer. Statist. Assoc* 100 286–295. [MR2156838](#)
- Hoff P (2008). Modeling homophily and stochastic equivalence in symmetric relational data In *Advances in Neural Information Processing Systems* 657–664.
- Hoff PD (2009). Multiplicative latent factor models for description and prediction of social networks. *Computational and Mathematical Organization Theory* 15 261–272.
- Hoff PD, Raftery AE and Handcock MS (2002). Latent space approaches to social network analysis. *J. Amer. Statist. Assoc* 97 1090–1098. [MR1951262](#)
- Holland PW and Leinhardt S (1970). A method for detecting structure in sociometric data. *Amer. J. Sociol* 492–513.
- Holland PW and Leinhardt S (1981). An exponential family of probability distributions for directed graphs. *J. Amer. Statist. Assoc* 76 33–65. [MR0608176](#)
- Holly JE (2001). Pictures of ultrametric spaces, the  $p$ -adic numbers, and valued fields. *Amer. Math. Monthly* 108 721–728. [MR1865659](#)
- Ibragimov Z (2014). A hyperbolic filling for ultrametric spaces. *Comput. Methods Funct. Theory* 14 315–329. [MR3265364](#)
- Ivrii V. Ja. (1980). The second term of the spectral asymptotics for a Laplace-Beltrami operator on manifolds with boundary. *Funktsional. Anal. i Prilozhen.* 14 25–34. [MR0575202](#)
- James C (1990). *Foundations of Social Theory*. Belknap, Cambridge, MA.
- Kolaczyk ED and Csárdi G (2014). *Statistical Analysis of Network Data with R. Use R!* Springer, New York [MR3288852](#)

- Krioukov D, Papadopoulos F, Kitsak M, Vahdat A and Bogun A M (2010). Hyperbolic geometry of complex networks. *Phys. Rev. E* (3) 82 036106, 18. [MR2787998](#)
- Krivitsky PN, Handcock MS, Raftery AE and Hoff PD (2009). Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models. *Soc. Netw* 31 204–213.
- Kunegis J (2013). KONECT—The Koblenz Network Collection.
- Lamping J, Rao R and Pirolli P (1995). A focus+ context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* 401–408. ACM Press/Addison-Wesley, New York.
- Lazarsfeld PF, Henry NW and Anderson TW (1968). *Latent Structure Analysis* 109 Houghton Mifflin, Boston.
- Lichnerowicz A (1958). *Geometrie des groupes de transformations*, Dunod, Paris, 1958. Zentralblatt MATH 96.
- Lovász L (2012). *Large Networks and Graph Limits* American Mathematical Society Colloquium Publications 60 Amer. Math. Soc., Providence, RI [MR3012035](#)
- Madore D Hyperbolic maze. Available at <http://www.madore.org/~david/math/hyperbolic-maze.html>. Hyperbolic maze.
- McCormick TH and Zheng T (2015). Latent surface models for networks using aggregated relational data. *J. Amer. Statist. Assoc* 110 1684–1695. [MR3449064](#)
- Minhas S, Hoff PD and Ward MD (2016). Inferential approaches for network analyses: AMEN for latent factor models. Preprint. Available at [arXiv:1611.00460](https://arxiv.org/abs/1611.00460).
- Munzner T (1997). H3: Laying out large directed graphs in 3D hyperbolic space. In *IEEE Symposium on Information Visualization, 1997. Proceedings* 2–10 IEEE Press, New York.
- Newman ME and Girvan M (2004). Finding and evaluating community structure in networks. *Phys. Rev. E* 69 026113.
- Nickel CLM (2009). *Random dot product graphs a model for social networks*. Ph.D. dissertation Johns Hopkins Univ, Baltimore, MD.
- Opsahl T (2009). *Structure and evolution of weighted networks*. Ph.D thesis, Queen Mary, Univ. London.
- Padgett JF (1994). *Marriage and Elite Structure in Renaissance Florence, 1282–1500*. Social Science History Association.
- Pao H, Coppersmith GA and Priebe CE (2011). Statistical inference on random graphs: Comparative power analyses via Monte Carlo. *J. Comput. Graph. Statist* 20 395–416. [MR2847801](#)
- Pattison P and Robins G (2002). Neighborhood-based models for social networks. *Sociol. Method* 32 301–337.
- R CORE Team (2016). *R: A language and environment for statistical computing* R Foundation for Statistical Computing, Vienna, Austria.
- Rapoport A (1953). Spread of information through a population with socio-structural bias. I. Assumption of transitivity. *Bull. Math. Biophys* 15 523–533. [MR0058955](#)
- Robins G, Pattison P, Kalish Y and Lusher D (2007). An introduction to exponential random graph ( $p^*$ ) models for social networks. *Soc. Netw* 29 173–191.
- Rogue Z and Rogue T (2011). HyperRogue. Available at <http://www.roguetemple.com/z/hyper/>.
- Saldaña DF, Yu Y and Feng Y (2017). How many communities are there? *J. Comput. Graph. Statist* 26 171–181. [MR3610418](#)
- Sarkar P and Moore AW (2006). Dynamic social network analysis using latent space models In *Advances in Neural Information Processing Systems* 1145–1152.
- Schweinberger M and Snijders TA (2003). Settings in social networks: A measurement model. *Sociol. Method* 33 307–341.
- Sewell DK and Chen Y (2015). Latent space models for dynamic networks. *J. Amer. Statist. Assoc* 110 1646–1657. [MR3449061](#)
- Simmel G (1950). The sociology of Georg Simmel In *Individual and Society*. 3–84. Free Press, New York.

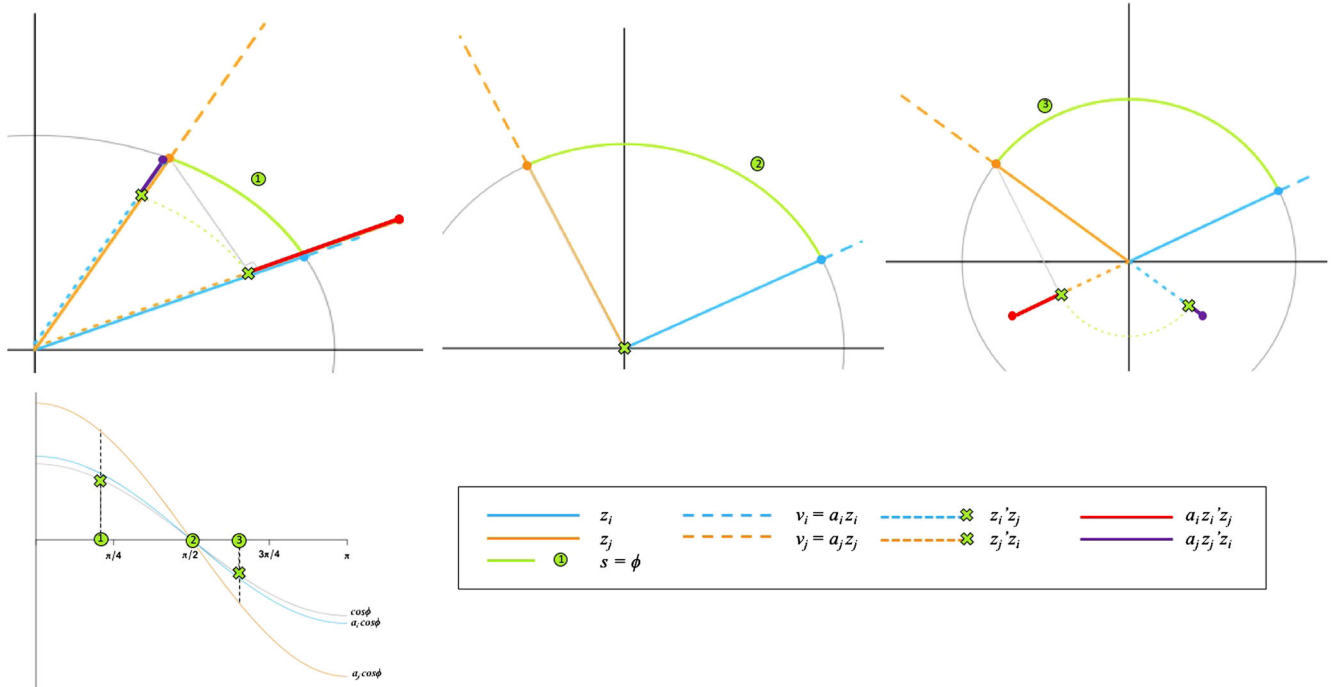
- Smith A (2017). Statistical methodology for multiple networks. PhD thesis, The Ohio State Univ., Columbus, OH.
- Snijders TA (2002). Markov chain Monte Carlo estimation of exponential random graph models. *Journal of Social Structure* 3 1–40.
- Snijders TAB and Nowicki K (1997). Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *J. Classification* 14 75–100. [MR1449742](#)
- Snijders TA, Pattison PE, Robins GL and Handcock MS (2006). New specifications for exponential random graph models. *Sociol. Method* 36 99–153.
- Solé RV and Valverde S (2004). Information theory of complex networks: on evolution and architectural constraints In *Complex Networks. Lecture Notes in Physics* 650 189–207. Springer, Berlin [MR2108978](#)
- Spearman C (1904). “General intelligence,” objectively determined and measured. *The American Journal of Psychology* 15 201–292.
- Ting D, Huang L and Jordan M (2011). An analysis of the convergence of graph Laplacians. Preprint. Available at [arXiv:1101.5435](#).
- van der Linden WJ and Hambleton RK, eds. (2013). *Handbook of Modern Item Response Theory*. Springer, New York.
- vanDuijn MAJ, Snijders TAB and Zijlstra BJH (2004).  $p_2$ : A random effects model with covariates for directed graphs. *Stat. Neerl* 58 234–254. [MR2064846](#)
- Wasserman S and Pattison P (1996). Logit models and logistic regressions for social networks. I. An introduction to Markov graphs and  $p$ . *Psychometrika* 61 401–425. [MR1424909](#)
- Watts DJ (1999a). Networks, dynamics, and the small-world phenomenon. *Amer. J. Sociol* 105 493–527.
- Watts DJ and Strogatz SH (1998). Collective dynamics of ‘small-world’ networks. *Nature* 393 440–442. [PubMed: 9623998]
- Weyl H (1911). Sur la distribution asymptotique des valeurs propres. *Nouvelles de la Societe des Sciences sur Gottingen, Mathematical-Physical Class* 110–117.
- Wolfe PJ and Olhede SC (2013). Nonparametric graphon estimation. Preprint. Available at [arXiv:1309.5936](#).
- Young SJ and Scheinerman ER (2007). Random dot product graph models for social networks In *Algorithms and Models for the Web-Graph. Lecture Notes in Computer Science* 4863 138–149. Springer, Berlin [MR2504912](#)
- Zachary WW (1977). An information flow model for conflict and fission in small groups. *Journal of Anthropological Research* 33 452–473.

Edge	Distance		Dot Product	
1 ↔ 4	0.85	(1)	-3.34	(1)
3 ↔ 5	1.17	(2)	0.22	(8)
1 ↔ 5	1.31	(3)	-0.48	(5)
1 ↔ 3	1.45	(4)	-0.49	(4)
4 ↔ 5	1.86	(5)	-0.91	(2)
2 ↔ 5	2.11	(6)	-0.79	(3)
3 ↔ 4	2.28	(7)	-0.23	(6)
1 ↔ 2	3.16	(8)	0.90	(9)
2 ↔ 3	3.22	(9)	1.95	(10)
2 ↔ 4	3.24	(10)	-0.14	(7)




**Fig. 1.** Euclidean distance and the dot product. For illustration, five points are simulated from a standard bivariate normal distribution. In the network plots, edge color indicates the size of the similarity measure (Euclidean distance for the network on the left and dot product for the network on the right) between nodes, with magenta edges representing nodes that are least similar and cyan edges for those that are closest or most similar. Note that the ranking of similarities across the two networks is not identical, for example consider the edge between nodes 3 and 5.

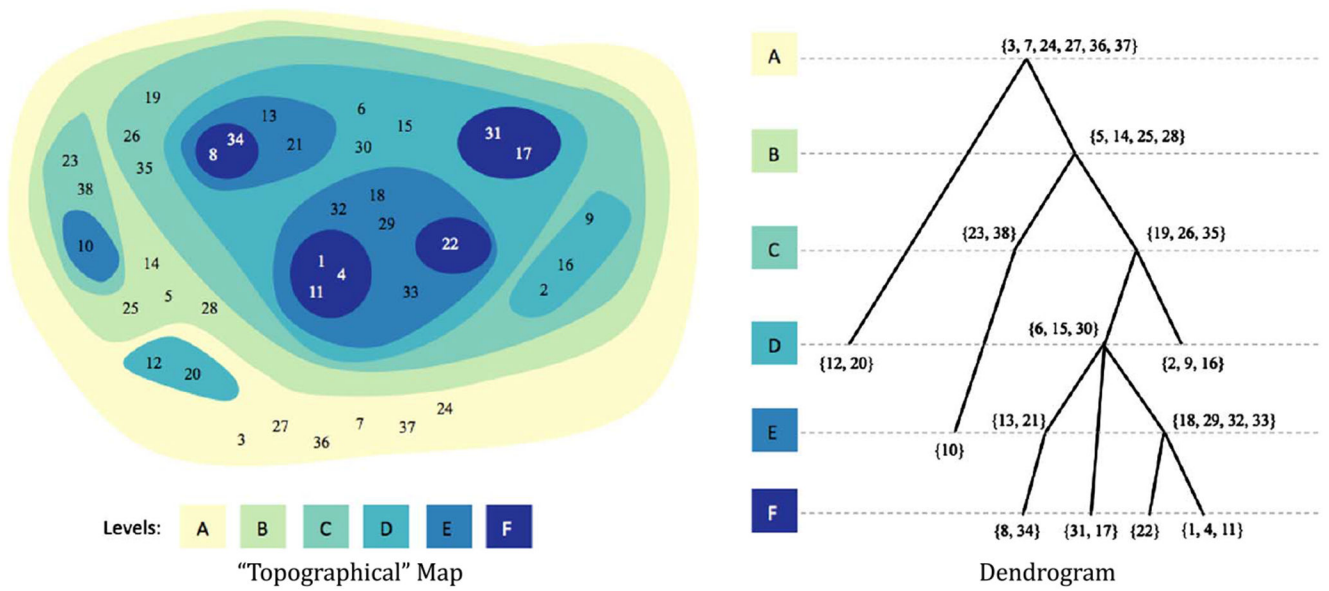




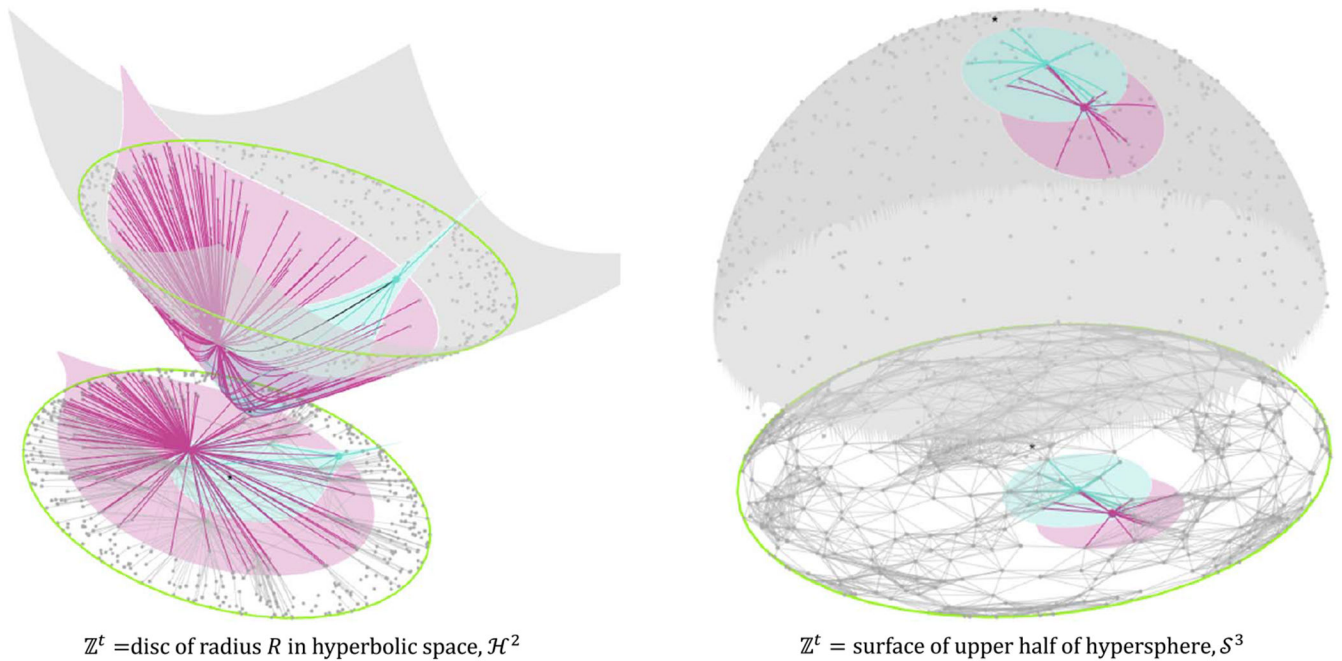
**Fig. 2.** Visual interpretation of the choice of  $s(\cdot, \cdot)$  in Hoff, Raftery and Handcock's (2002) projection model,  $t=2$ . Latent position vectors for the  $i$  th node are shown in blue and for the  $j$  th node are shown in orange. Across the top panel, we display three possible sets of latent positions, with their corresponding similarity measure in Hoff, Raftery and Handcock's (2002) model identified by green crosses. Distance in this latent space, the unit hypersphere, is great circle distance or arc length and is displayed as green arcs in the top panels and green points in the bottom left panel.




**Fig. 3.** *Examples of ultrametric space. In the left panel, we illustrate how the animal species classification system can be viewed as an ultrametric space, highlighting a disc of radius  $\gamma = 3$  centered at species  $q$ . In the right panel, we display  $\mathbb{Q}_p$ , the field of  $p$ -adic numbers, taking  $p = 3$ . Roughly speaking, two points in  $\mathbb{Q}_p$  are “close” if their difference is divisible by a large positive power of  $p$ . In the panel above, for example, 0 and 1 are closer together than are 0 and  $\frac{1}{3}$ , because  $1 - 0 = 1 = 3^0$  while  $\frac{1}{3} - 0 = 3^{-1}$ . When visualized as a dendrogram, distance in ultrametric spaces is represented by the height of the smallest (inverted) tree between two members (members are represented as points on the bottom-most line in each panel).*



**Fig. 4.** Positions in a latent ultrametric space. As a toy example, we plot “positions” in a latent ultrametric space, visualized above in the style of a topographical map, on the left, and as a dendrogram, on the right. Colors represent the levels and the nodes are indexed arbitrarily. Note that nodes’ relative positions within a particular level in the topographical map are meaningless and are simply an artifact of this particular visualization.

**Fig. 5.**

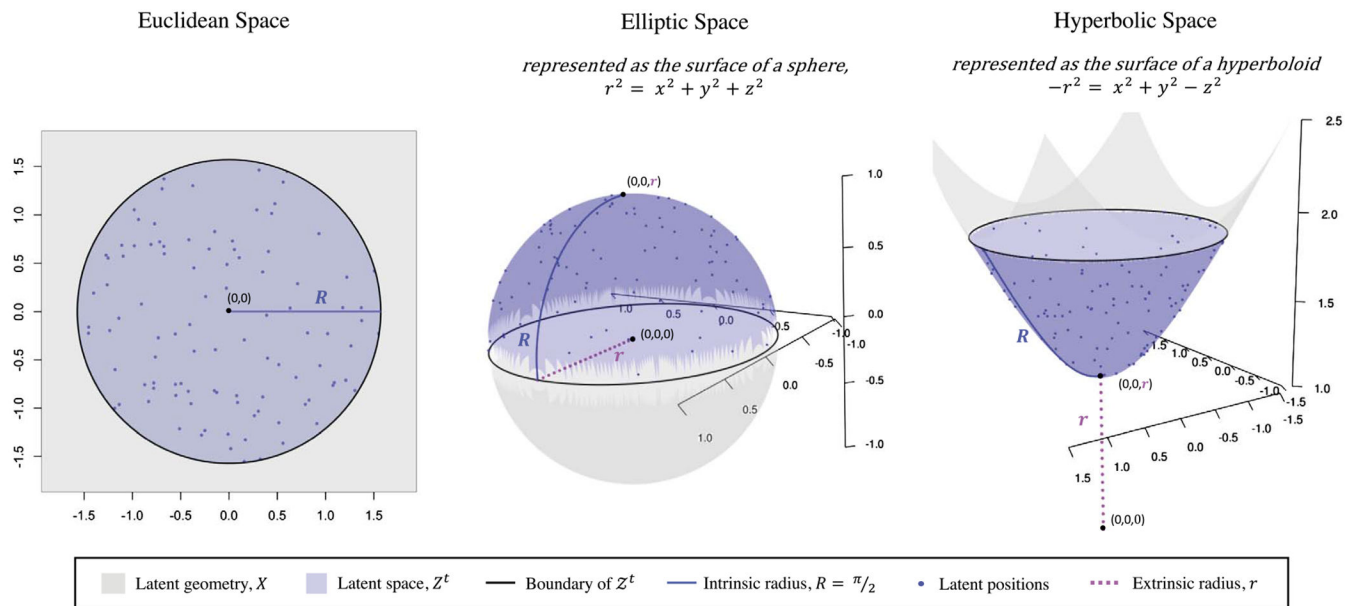
*Networks simulated in complex geometries ( $n = 500$  nodes). In each panel, the nodes are displayed both in their natural latent space (on the surface of a hyperbola—in theory, this space extends up to infinity—or the surface of a sphere, each shaded in light grey) as well as projected onto the  $xy$ -plane (the green circles below each latent space). In both models, latent positions are constrained to fall within a radius of  $R$  from the origin, the green circles identify these limiting bounds. Additionally, we highlight two nodes by drawing circles of radius  $R$  about them (in light pink and light teal) and color their connections to other nodes in the network (identified by the dark pink and dark teal lines). All other nodes are displayed as grey points with their connections identified by gray lines in the  $xy$ -plane. For hyperbolic space,  $R$  is chosen such that the average degree is eight (Krioukov et al., 2010, equation (13)) and in elliptic space,  $R = 1$ .*



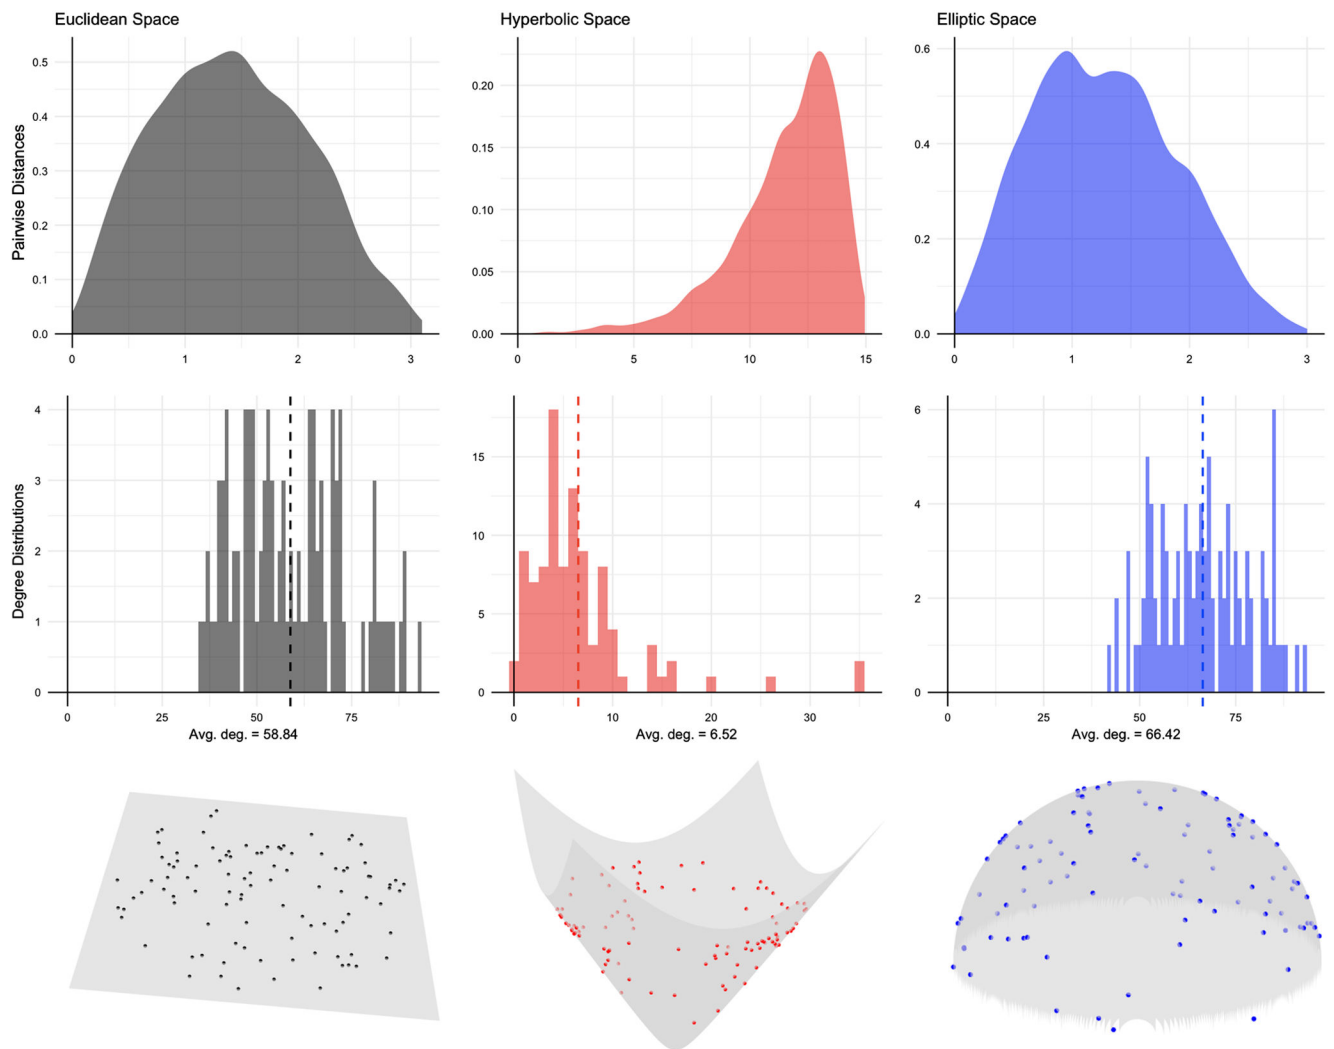
**Fig. 6.**

*Quasi-uniform distribution of latent positions in hyperbolic space. Although the distribution for the latent positions embeds nodes in hyperbolic space, here we plot only the  $x$  - and  $y$ -coordinates for these simulated positions, so that the plots above are in  $\mathbb{R}^2$ , for simplicity.*

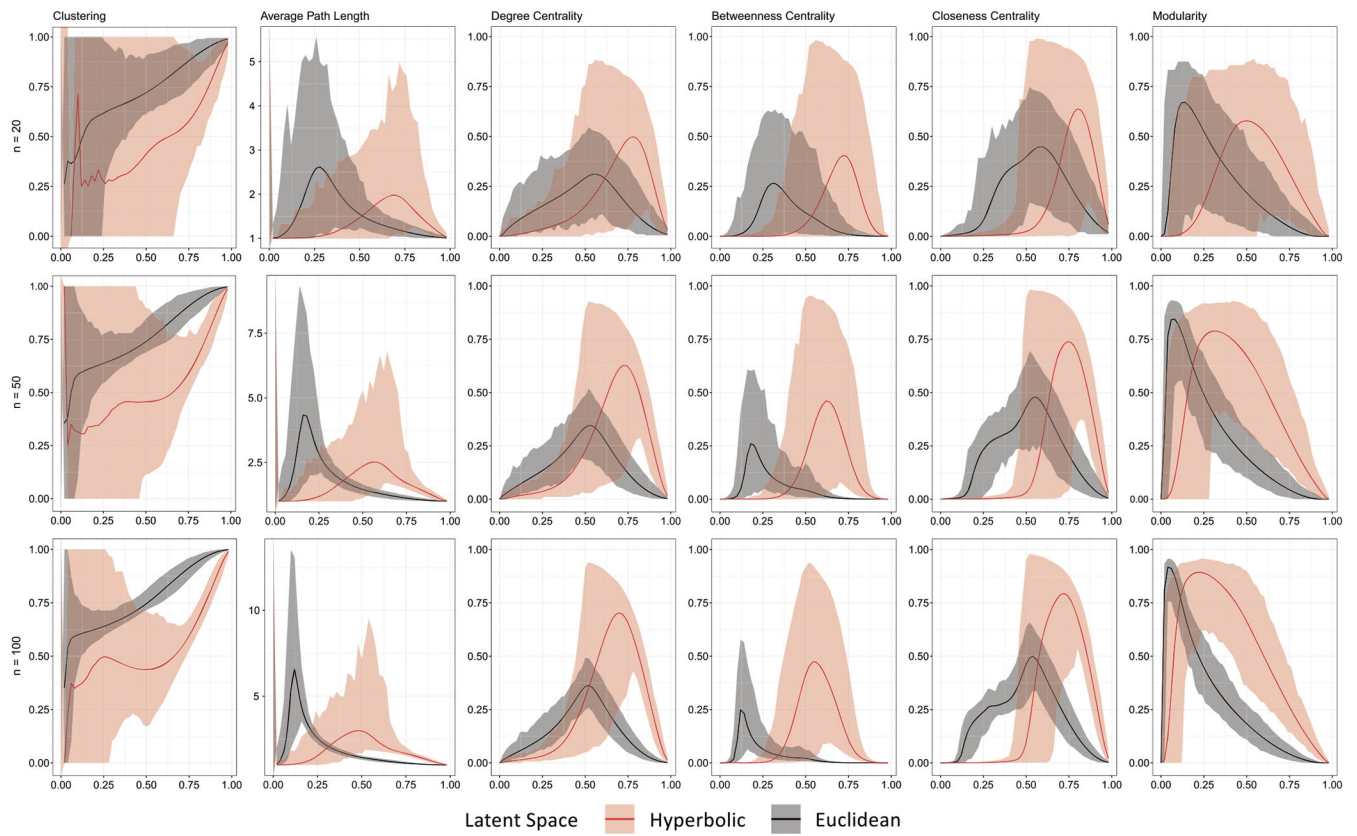
*When  $\alpha = 1$ , the points are exactly uniformly distributed and we can see that as  $\alpha$  increases, points become more clustered around the boundary.*



**Fig. 7.** *Comparable latent spaces. Latent positions are sampled uniformly from a subset (here,  $Z^t$  is defined as a disc of radius  $R = \pi/2$ ) of each latent geometry, with elliptic geometry represented as the surface of a sphere and hyperbolic geometry represented by the surface of a hyperboloid plotted in  $\mathbb{R}^3$ . The intrinsic radius,  $R$ , of the disc controls how much of each latent space,  $X$ , is being used in the model as  $Z^t$ . This should not be confused with the extrinsic radius,  $r$ , which determines the curvature of elliptic and hyperbolic space (here,  $r = 1$  and  $r = -1$ , respectively).*

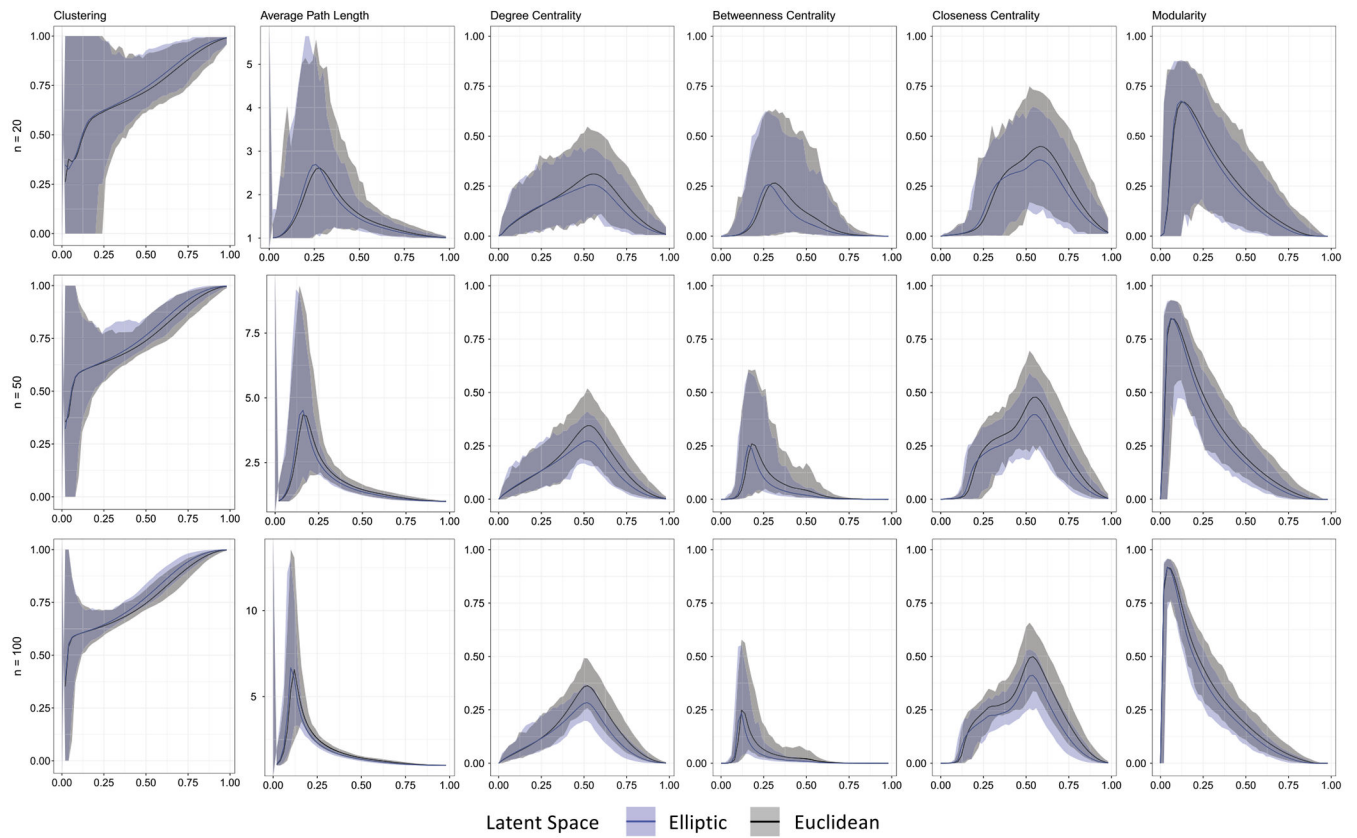


**Fig. 8.** *Distribution of distances in latent spaces with differing geometries. Panels in the bottom row plot one set of  $n = 100$  points simulated in Euclidean, hyperbolic and elliptic space, respectively, while panels in the top row plot histograms of the pairwise distances for these sets of points. The center row plots the corresponding degree distributions across the three latent geometries, where nodes in each network are connected according to model (4.1) with  $\gamma = R$ .*

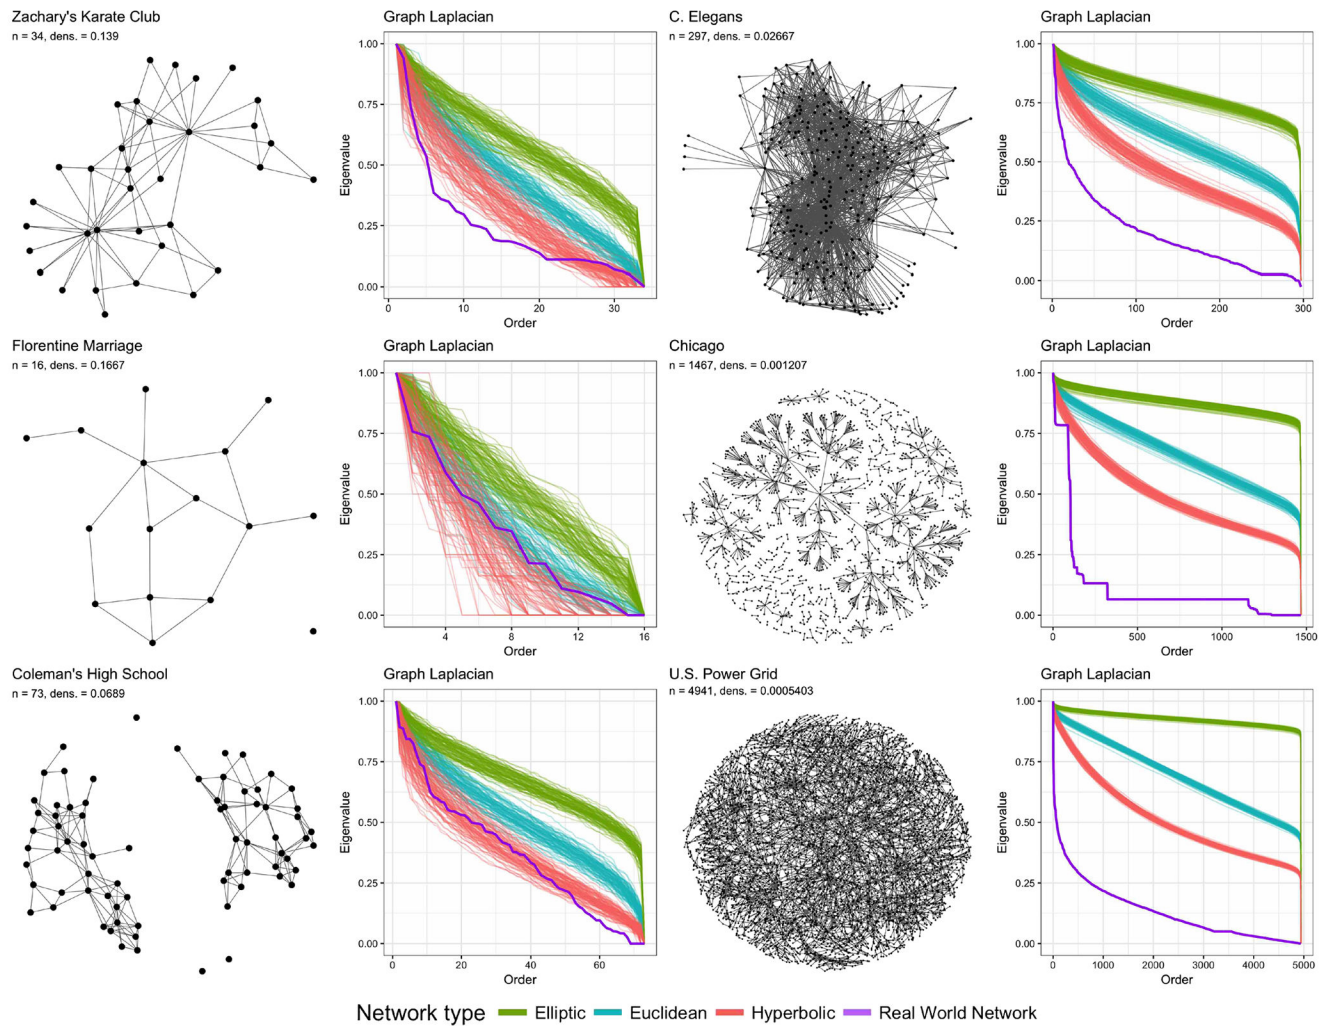


**Fig. 9.** Network statistics in hyperbolic space. The x-axis is  $\gamma$ , the Heaviside step function cut point. Distances across the geometries are made comparable by re-scaling each simulation to the  $[0, 1]$  interval. Each row of plots corresponds to simulated networks of a different size,  $n = 20, 50, 100$  and each column of plots corresponds to a different network summary measure.





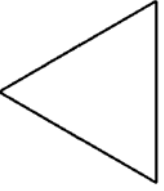
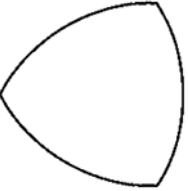
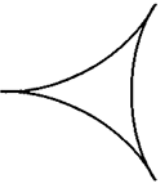
**Fig. 10.** *Network statistics in elliptic space. The x-axis is  $\gamma$ , the Heaviside step function cut point. Distances across the geometries are made comparable by re-scaling each simulation to the  $[0, 1]$  interval. Each row of plots corresponds to simulated networks of a different size,  $n = 20, 50, 100$  and each column of plots corresponds to a different network summary measure.*



**Fig. 11.** *A small collection of real world networks. In these plots, we compare a set of real world networks to networks simulated from our general latent distance model (Equation (2.4)) across different geometries. Each row includes a network plot (using the default settings of the ggnet function in the GGally package for R) and a plot of the eigenvalue curves, where the color corresponds to the geometry of the latent space. In the eigenvalue plots, each network is represented by a curve which plots the network's eigenvalues against their order. Networks on the left side of the plot are examples of social networks, that range in size from  $n = 16$  to  $n = 73$ , while networks on the right side of the plot are nonsocial networks, including a brain network, road networks and a power grid network.*

**Table 1**

*Geometric properties of curved spaces. This table is adapted from Krioukov et al. (2010)*

	Geometry		
	Euclidean	Spherical	Hyperbolic
Curvature, $K$	0	$> 0$	$< 0$
Number of parallel lines	1	0	$\infty$
Triangles look ...	normal 	thick 	thin 
PROPERTY	$\pi$	$> \pi$	$< \pi$
Sum of angles in a triangle	$2\pi r$ linear growth	$2\pi \sin(\sqrt{ K } r)$ sublinear growth	$2\pi \sinh(\sqrt{ K } r)$ exponential growth
Circumference of a circle ... as $r$ increases,	$2\pi r$ linear growth	$2\pi \sin(\sqrt{ K } r)$ sublinear growth	$2\pi \sinh(\sqrt{ K } r)$ exponential growth