
Neighborhood Growth Determines Geometric Priors for Relational Representation Learning

Melanie Weber
Princeton University
mw25@math.princeton.edu

Abstract

The problem of identifying geometric structure in heterogeneous, high-dimensional data is a cornerstone of representation learning. While there exists a large body of literature on the embeddability of canonical graphs, such as lattices or trees, the heterogeneity of the relational data typically encountered in practice limits the applicability of these classical methods. In this paper, we propose a combinatorial approach to evaluating embeddability, i.e., to decide whether a data set is best represented in Euclidean, Hyperbolic or Spherical space. Our method analyzes nearest-neighbor structures and local neighborhood growth rates to identify the geometric priors of suitable embedding spaces. For canonical graphs, the algorithm’s prediction provably matches classical results. As for large, heterogeneous graphs, we introduce an efficiently computable statistic that approximates the algorithm’s decision rule. We validate our method over a range of benchmark data sets and compare with recently published optimization-based embeddability methods.

1 Introduction

A key challenge in data science is the identification of geometric structure in high-dimensional data. Such structural understanding is of great value for designing efficient algorithms for optimization and learning tasks. Classically, the structure of data has been studied under an Euclidean assumption. The simplicity of vector spaces and the wide range of well-studied tools

and algorithms that assume such structure make this a natural approach. However, lately, it has been recognized that Euclidean spaces do not necessarily allow for the most ‘natural’ representation, at least not in the low-dimensional regime. Recently, the representation of data in hyperbolic space has gained significant interest (Nickel and Kiela, 2017; Chamberlain et al., 2017; Sala et al., 2018; Tifrea et al., 2019). The intrinsic hierarchical structure of data sets ranging from social networks to wordnets has been related to “tree-likeness” and in turn to hyperbolic embeddability, since trees embed with low distortion into hyperbolic space (Sarkar, 2012). On the other hand, there is a long tradition for spherical embeddings in computer vision and shape analysis, where volumetric data is efficiently represented in spherical space.

In this work, we study the question of embeddability in the context of relational representation learning: For a given set of pairwise similarity relations, we want to determine the geometric priors of an embedding space that reflects the intrinsic structure of the data. Optimization-based embeddability methods (Sala et al., 2018; Gu et al., 2019) rely on *Multi-dimensional Scaling (MDS)*, that require performing large-scale minimization tasks to determine suitable embedding parameters. Furthermore, such methods require an a priori fixed embedding dimension. Here, we introduce a purely combinatorial approach, that efficiently determines suitable priors through a direct analysis of the data’s discrete geometry without a priori dimensionality assumptions.

The paper is structured as follows: We will first analyze the relation between embeddability and neighborhood growth rates. Expanding neighborhoods (*exponential growth*) exhibit tree-like properties and contribute to the hyperbolicity of the data set. On the other hand, cycles have a contracting effect, as they slow down local growth. Therefore, slowly expanding neighborhoods are an indicator of good embeddability into Euclidean (*linear growth*) or Spherical space (*sublinear growth*). To extend this framework from canonical graphs to heterogeneous relational data, we introduce a regular-

ization that ensures uniform node degrees and therefore allows for a direct comparison of growth in diverse graph neighborhoods. We then introduce a statistic (*3-regular score*) that aggregates local growth information across the graph. Based on the 3-regular score, we determine the geometric priors of the most suitable embedding space. For canonical graphs (N -cycles, $(\sqrt{N} \times \sqrt{N})$ -lattices and b -ary trees) we give a proof that the approach matches classical embeddability results. Furthermore, we establish a relation between the 3-regular score and discrete Ricci curvature (Forman, 2003; Ollivier, 2010), a concept from Discrete Geometry that has been linked to embeddability (Weber et al., 2017b).

The introduced method is purely combinatorial with a computational complexity linear in the average neighborhood size multiplied by the number of nodes. Moreover, as a local analysis, it can be efficiently parallelized.

1.1 Related Work

Embeddings for Representation Learning. The theoretical foundation of Euclidean embeddability has been laid out by (Bourgain, 1985) and (Johnson and Lindenstrauss, 1984). Here, the relation of *intrinsic* and *metric dimension* is of special interest: One can show with a volume argument that data X with intrinsic dimension $\dim(X)$ can be embedded with metric dimension $\Theta\left(\frac{\dim(X)}{\log(\alpha)}\right)$ and distortion α into Euclidean space. (Abraham et al., 2008) study this relation algorithmically and derive distortion bounds. Recently, optimization-based hyperbolic embeddings have gained a surge of interest in the representation learning community. (Nickel and Kiela, 2017, 2018; Chamberlain et al., 2017; Tifrea et al., 2019) proposed optimization-based frameworks for embedding similarity data into hyperbolic space. (Sala et al., 2018) analyze representation trade-offs in hyperbolic embeddings with varying geometric priors. (Gu et al., 2019) introduce mixed-curvature embeddings by studying embeddability onto product manifolds. A related approach by (Weber and Nickel, 2018) explores connections between graph motifs and hyperbolic vs. spherical embeddability.

Spectral Approaches. (Wilson et al., 2014) propose a spectral approach that determines embedding parameters by minimizing the magnitude of the first (spherical case) or second (hyperbolic case) eigenvalue. In this approach, the sign of the curvature is a hyperparameter of the objective and has to be prior known. However, the framework is only valid for isometrically embeddable data, leading to inaccuracies in heterogeneous data. In addition, spectral methods have limited scalability on large-scale data.

Discrete Curvature. In addition to spectral ap-

proaches, discrete curvature has recently gained interest as a structural graph characteristic. Gromov’s δ -hyperbolicity (Gromov, 1987), a discrete notion of sectional curvature, has been used to study the hyperbolicity of relational data (Chen et al., 2013; Gu et al., 2019). Discrete notions of Ricci curvature were studied as graph characteristics (Ollivier, 2010; Weber et al., 2017a,b), providing insight into the local geometry of the underlying relational data.

1.2 Contributions

In the present paper, we connect the structure of nearest neighbor relations to the intrinsic geometry of relational data. We argue that the growth rates of graph neighborhoods serve as a proxy for the geometric priors of a suitable embedding space. To account for the heterogeneity of complex data sets, we perform a regularization that allows for a low-cost embedding of any graph into a 3-regular graph quasi-isometrically. In the regularized setting, where node degrees are uniform, we directly compare local neighborhood growth rates and deduce the curvature of the embedding space.

We show that our classification scheme matches theoretical results for canonical graphs. Furthermore, we establish a relation to discrete Ricci curvature. For analyzing complex, heterogeneous data sets as typically encountered in ML applications, we introduce a statistic (*3-regular score*) that aggregates local growth information and approximates the algorithm’s decision rule. A series of validation experiments on classic benchmark graphs and real-world data sets demonstrates the applicability of the proposed approach. Finally, we compare our method to recently published embeddability benchmarks, validating that the 3-regular score predicts the lowest-distortion embedding.

2 Background and Notation

2.1 Model Spaces

We consider canonical Riemannian manifolds with constant curvature as embedding spaces, which can be characterized through the following set of *model spaces* $\{M_\kappa^d\}$:

1. $M_0^n = \mathbb{R}^n$ denotes the canonical **Euclidean space** with the inner product $\langle u, v \rangle_E = \sum_{i=1}^n u_i v_i$ that gives rise to the Euclidean norm $\|v\|_E = \sqrt{v_1^2 + \dots + v_n^2}$ and the metric $d_E(u, v) = \|u - v\|_E$.
2. The n -**sphere** $M_{\kappa>0}^n = \mathbb{S}^n = \{v \in \mathbb{R}^{n+1} : \|v\|_E = \sqrt{\kappa}\}$ is an embedded submanifold of the \mathbb{R}^{n+1} with constant positive curvature. A canonical metric is given by $d_S(u, v) = \arccos(\langle u, v \rangle_E)$.

3. The **hyperboloid** $M_{\kappa < 0}^n = \mathbb{H}^n = \{v \in \mathbb{R}^{n+1} : \|v\|_H = \sqrt{-\kappa}, v_1 > 0\}$ is a manifold with constant negative curvature. It is defined with respect to the *Minkowski inner product*

$$\begin{aligned} \langle u, v \rangle_H &= u^T \text{diag}(-1, 1, \dots, 1)v \\ &= -u_1 v_1 + u_2 v_2 + \dots + u_n v_n, \end{aligned}$$

which gives rise to the hyperbolic metric $d_H(u, v) = \text{acosh}(-\langle u, v \rangle_H)$ and norm $\|u\|_H = \sqrt{\langle u, v \rangle_H}$.

Here, we focus on the canonical model spaces $M_{0, \pm 1}$. Table 1 summarizes important geometric properties that will be used in the following sections. Note that $M_{\pm 1}$ can be easily generalized to arbitrary curvatures ($|\kappa| \neq 1$) by multiplying the respective distance functions by $\frac{1}{\sqrt{|\kappa|}}$. In the following, we will drop the subscript E when referring to the Euclidean notions. For a more comprehensive overview on model spaces, see, e.g. (Bridson and Haefliger, 1999).

2.2 Graph Motifs and Local Topology

The present paper focuses on relational data, i.e., we assume access to a measure of similarity between any two elements. Natural representations of such data are graphs $G = \{V, E\}$, where V denotes the set of vertices or nodes (representing data points) and E the set of edges (representing relations). Additional features may be given through weights on the edges which we encode in the weight functions $\omega_E(e) : E \rightarrow \mathbb{R}$.

The importance of graph motifs for understanding the higher-order structure of graphs has long been recognized and intensely studied (Chen et al., 2013; Maugis et al., 2017; Verbeek and Suri, 2016). Motifs are commonly defined as characteristic local connectivity patterns that occur in varying sizes and frequencies. While there is no canonical classification, trees and cycles have emerged as prevalent motifs in the study of network topology, due to having the greatest topological stability (i.e., the highest Euler characteristic) (Maugis et al., 2017). A random walk initiated at the root of a tree will never return to its point of origin, but expand into space. On the other hand, a random walk within a cycle is guaranteed (or, in a circle with outgoing connections, likely) to return to its origin, introducing a local contraction. This naturally relates to local growth rates in graph neighborhoods: While trees intrinsically encode exponential growth, cycles introduce a local contraction, resulting in sublinear growth rates. We will connect these ideas with the problem of embeddability.

2.3 Embeddability

An *embedding* between metric spaces (X_1, d_1) and (X_2, d_2) is described as a map $\phi : X_1 \rightarrow X_2$. Here, we consider embeddings of relational data into canonical model spaces, i.e. we want to embed the graph metric (V, d_G) of a (weighted or unweighted) graph $G = \{V, E\}$ using a map $\phi : V \rightarrow \mathcal{M}_{0, \pm 1}^d$, where d_G denotes the usual path distance metric. The goodness of an embedding is measured in terms of distortion. We denote the *additive* distortion $c_A \geq 0$ of the map ϕ as

$$|d_G(u, v) - d_{\mathcal{M}}(\phi(u), \phi(v))| \leq c_A \quad \forall u, v \in V,$$

and the *multiplicative distortion* $c_M \geq 0$ as

$$\begin{aligned} d_{\mathcal{M}}(\phi(u), \phi(v)) &\leq d_G(u, v) \\ &\leq c_M d_{\mathcal{M}}(\phi(u), \phi(v)) \quad \forall u, v \in V. \end{aligned}$$

Note that for an isometric map $c_A = 0$ and $c_M = 1$.

While little is known about the embeddability of large, heterogeneous graphs, there exists a large body of literature on the embeddability of canonical graphs. The, to our knowledge, best known results for multiplicative distortion are summarized in Table 2. In the following we develop a computational method that applies not only to canonical graphs, but to any relational data set.

3 Methods

We determine the geometric priors of a suitable embedding space with a two-step method, (1) by performing a *regularization* that enforces uniform node degrees while preserving structural information and (2) by analyzing local neighborhood growth rates to determine the dominating geometry (*β -regular score*).

3.1 Regularization

Relational data as typically encountered in data science applications is very heterogeneous, making it difficult to draw a conclusion on the global geometry from local analysis. Our first step is therefore the regularization of the graph's connectivity structure that will allow for a more efficient comparison of local neighborhood growth rates and, in turn, the local geometry. We will use throughout the paper the following (conventional) notation: When analyzing the neighborhood of a vertex $v \in V$, we say that v is the *root* or *center* of the neighborhood. Neighborhood directionality is always assumed from the root outwards, the start of an outward facing edge is called *parent*, the end *child*. The root has no parent. Furthermore, $\text{diam}(G)$ denotes the diameter of the graph.

	Euclidean \mathbb{R}^d	Spherical \mathbb{S}^d	Hyperboloid \mathbb{H}^d
Space	\mathbb{R}^n	$\{x \in \mathbb{R}^{n+1} : \langle x, x \rangle = 1\}$	$\{x \in \mathbb{R}^{n+1} : \langle x, x \rangle_H = -1, x_0 > 0\}$
$\langle u, v \rangle$	$\sum_{i=1}^n u_i v_i$	$\sum_{i=1}^n u_i v_i$	$-u_0 v_0 + \sum_{i=1}^n u_i v_i$
$d(u, v)$	$\sqrt{\langle u - v, u - v \rangle}$	$\arccos(\langle u, v \rangle)$	$\operatorname{arcosh}(-\langle u, v \rangle_H)$
Curvature	$\kappa = 0$	$\kappa = 1$	$\kappa = -1$
Canonical graph	$(\sqrt{N} \times \sqrt{N})$ -lattice	N -cycle	regular N -tree

Table 1: Geometric properties of model spaces.

	Euclidean \mathbb{R}^d	Spherical \mathbb{S}^d	Hyperboloid \mathbb{H}^d
$(\sqrt{N} \times \sqrt{N})$ - lattice	$c_M \leq O(1)^a$	-	$c_M \geq O(\sqrt{N}/\log(N))^a$
N -cycle	$c_M \leq O(1)^a$	$c_M \leq O(1)^a$	$c_M \geq O(N/\log(N))^a$
b -regular tree (size N)	$c_M \leq O(N^{\frac{1}{d-1}})^c$	-	$c_M \leq O(1 + \epsilon)^b$

 Table 2: Known results on embeddability of canonical graphs. For ^{a)} see (Verbeek and Suri, 2016), for ^{b)} see (Sarkar, 2012) and ^{c)} (Gupta, 1999). Here, N denotes the size of the graph, d the embedding dimension and $\epsilon > 0$ a small real number.

We utilize a quasi-isometric embedding (Bermudo et al., 2013) that allows for embedding any (connected) graph into a three-regular graph, i.e. a graph with uniform node degrees ($\deg(v) = 3$ for all $v \in V$). The regularization algorithm is shown schematically in Fig. 1, for more details see Appendix A. One can show the following bound on the distortion induced by this transformation:

Theorem 3.1 (Bermudo et al. (2013)). $G \hookrightarrow^\phi G_3$ is a $(\epsilon + 1, \epsilon)$ -quasi-isometric embedding, i.e. $c_M = O(1 + \epsilon)$ and $c_A = O(\epsilon)$.

3.2 Estimating local neighborhood growth rates

In order to decide the geometry of a suitable embedding space, we want to analyze *neighborhood growth rates*. Consider first a continuous, metric space $(\mathcal{X}, d_{\mathcal{X}})$. The δ -neighborhood of a point $x \in \mathcal{X}$ is defined as the set of points within a distance δ , i.e.

$$\mathcal{B}_\delta(x) = \{y \in \mathcal{X} : d_{\mathcal{X}}(x, y) \leq \delta\}. \quad (1)$$

In Euclidean space, the volume of $\mathcal{B}_\delta(x)$ is growing at a polynomial rate in δ . However, in hyperbolic space, the volume growth is exponential. Therefore, the local volume growth of neighborhoods serves as a proxy for the space’s global geometry.

In discrete space, instead of analyzing volume growth, we characterize the local growth of neighborhoods. We denote the r -neighborhood of a vertex $v \in V(G)$ as

$$\mathcal{N}_r(v) = \{u \in V(G) : d_G(u, v) \leq r \leq \operatorname{diam}(G)\}. \quad (2)$$

We say that $\mathcal{N}_r(v)$ is *exponentially expanding*, if it grows exponentially in r and *linearly expanding*, if it grows at least linearly in r . Otherwise, we call $\mathcal{N}_r(v)$ *sublinearly*

expanding. Thanks to the regularized structure of our graphs, we can quantify precisely the corresponding neighborhood growth laws:

1. *exponentially expanding*: $\gamma_{EE}(v, R) = |v| + \sum_{r=1}^R 3 \cdot 2^{r-1}$;
2. *linearly expanding*: $\gamma_{LE}(v, R) = |v| + \sum_{r=1}^R 3r$.

Here, $|v|$ denotes the size of the root structure, i.e., $|v| = 1$, if $\deg(v) \leq 3$ and $|v| = \deg(v)$ otherwise, due to the transformation of star-nodes into three-regular rings (see Fig. 1 and Appendix A). Then we say that the R -neighborhood of a vertex $u \in V(G)$ is exponentially expanding, if $|\mathcal{N}_R(u)| \geq \gamma_{EE}(R)$, linearly expanding, if $|\mathcal{N}_R(u)| \geq \gamma_{LE}(R)$ and sublinearly expanding otherwise. For canonical graphs, we get the following neighborhood growth rates (a proof can be found in Appendix B):

Theorem 3.2 (Neighborhood growth in canonical graphs). For $1 \leq R \leq \operatorname{diam}(G)$, every R -neighborhood in (i) a b -regular tree is exponentially expanding, (ii) an $(\sqrt{N} \times \sqrt{N})$ -lattice is linearly expanding and (iii) an N -cycle is sublinearly expanding.

Utilizing the link between neighborhood growth and global geometry that we discussed above, we introduce the following decision rule for the geometric prior of the embedding space ($\operatorname{sign}(\kappa)$):

- If $\mathcal{N}_R(v)$ is *exponentially expanding* $\forall v \in V$ and $1 \leq R \ll N$, assume $\operatorname{sign}(\kappa) < 0$, i.e., embed into \mathbb{H}^d ;
- If $\mathcal{N}_R(v)$ is *linearly expanding* $\forall v \in V$ and $1 \leq R \ll N$, assume $\operatorname{sign}(\kappa) = 0$, i.e., embed into \mathbb{R}^d ;

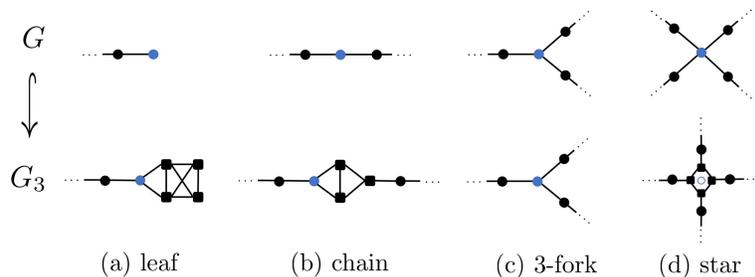


Figure 1: Regularization $G \leftrightarrow G_3$ following (Bermudo et al., 2013). A vertex $v \in V$ is called *leaf*, if it has degree 1, *chain* if it has degree 2, *3-fork*, if it has degree 3 and *star* otherwise. The transformation maps each vertex and the edges connecting to its direct neighbors onto the respective three-regular structure as predicted by the vertex' degree. Circles represent original vertices, squares auxiliary vertices that have been added to enforce uniform node degrees.

- If $\mathcal{N}_R(v)$ is *sublinearly expanding* $\forall v \in V$ and $1 \leq R \ll N$, assume $\text{sign}(\kappa) > 0$, i.e., embed into \mathbb{S}^d .

Note that these result match known embeddability results for canonical graphs (see Table 2). In the following section we introduce a statistic that allows for applying this decision rule to heterogeneous relational data also.

3.3 3-Regular Score

The heterogeneity commonly encountered in relational data makes it impossible to generalize global growth rates from a local analysis as in Thm. 3.2. We will typically find a mixture of local growth rates, that is not covered by the decision rule above. Instead, we focus on determining the globally *dominating* geometry: We analyze growth rates locally and then compute an average across the graph, weighted by the size of the respective R -neighborhood. The resulting statistic, to which we refer as the *3-regular score*, can be computed as follows:

$$A = \sum_{v \in V} \sigma(v) |\mathcal{N}_R(v)| ;$$

$$\sigma(v) = \begin{cases} 1, & \mathcal{N}_R(v) \text{ sublinearly expanding} \\ -1, & \mathcal{N}_R(v) \text{ exponentially expanding} \\ 0, & \text{otherwise} . \end{cases}$$

To determine the geometric priors, we apply the following decision rule:

- if $A > 0$, assume $\text{sign}(\kappa) > 0$, i.e., embed into \mathbb{S}^d ;
- if $A < 0$, assume $\text{sign}(\kappa) < 0$, i.e., embed into \mathbb{H}^d ;
- and if $A \approx 0$, assume $\text{sign}(\kappa) = 0$, i.e., embed into \mathbb{R}^d .

For weighted networks, we perform the same regularization and computation of the 3-regular score, but replace in the growth rate estimations 3 with the weighted node degree of the center. When determining $|\mathcal{N}_R(v)|$, i.e., the set of all neighbors up to distance R from v , the metric is the weighted path distance \tilde{d}_G . Consequently, we count all neighbors v' with $\tilde{d}_G(v, v') \leq R$.

The decision rule is motivated by locally aggregating neighborhood growth information. Hereby σ encodes whether the neighborhood growth is locally *exponential* (indicating hyperbolic space, i.e., $\text{sign}(\kappa) = -1$ and $\sigma = -1$), *linear* (indicating Euclidean space, i.e., $\kappa = 0$ and $\sigma = 0$) or *sublinear* (indicating spherical space, i.e., $\text{sign}(\kappa) = 1$ and $\sigma = -1$). Due to the heterogeneity of the graphs, we weigh the σ s by the size of the neighborhood ($|\mathcal{N}_R(v)|$) to give large neighborhoods a larger influence on the overall score. This is motivated by the fact that the "amount" of distortion incurred is proportional to the largest subgraph of another space's canonical motif: For instance, when embedding a graph into hyperbolic space, distortion is proportional to the size of the largest cycle by a Steiner node construction (see, e.g., (Verbeek and Suri, 2016)). The resulting 3-regular score A after reweighing will then depend on the size of the graph, in particular on the number of edges in the regularized graph G_3 . Therefore, we normalize by dividing by the the number of edges in G_3 , i.e., we compare $A/\#E(G_3)$ across data sets. The dependency of A on R is explicitly given through the weights $|\mathcal{N}_R(v)|$; R is upper-bounded by the diameter of the graph.

3.4 Comparison with other discrete curvatures

The 3-regular score is conceptually related to discrete notions of curvature, such as Gromov's δ -hyperbolicity (Gromov, 1987) or discrete Ricci cur-

vature (Forman, 2003; Ollivier, 2010). While Gromov’s δ captures by construction the hyperbolicity of a graph, discrete Ricci curvature is not restricted to negative values. In this section, we analyze the relationship between the 3-regular score and discrete Ricci curvature and compare the suitability of both concepts to measuring embeddability.

In the following, we will think of data sets as graphs, where nodes represent the data points and edges the pairwise similarities between them. For simplicity, we only consider unweighted graphs. We consider both *Ollivier-Ricci curvature* (Ric_O) (Ollivier, 2010) and *Forman-Ricci curvature* (Ric_F) (Forman, 2003; Weber et al., 2017a) that have previously been analyzed in the context of large-scale data and complex networks. Although both curvatures are classically defined edge-based, we will use node-based expressions. Those can be derived by defining the Ricci curvature at a node as the aggregate curvature of its incoming and outgoing edges (see Appendix C for more details). For Ric_O , consider

$$\text{Ric}_O(u, v) = 1 - W_1(m_u, m_v),$$

where $W_1(m_u, m_v)$ is the Wasserstein-1 distance that measures the cost of transporting mass from u to v . $m_u = \frac{1}{\deg(u)}$ denotes the uniform measure on the neighborhood of u . The corresponding node-based notion is given by

$$\text{Ric}_O(v) = \frac{1}{\deg(v)} \sum_{(u,v)} \text{Ric}_O(u, v).$$

Ric_F of an edge (u, v) is defined as

$$\text{Ric}_F(u, v) = 4 - \deg(u) - \deg(v),$$

the corresponding node-based expression as

$$\text{Ric}_F(v) = \frac{1}{\deg(v)} \sum_{(u,v)} \text{Ric}_F(u, v).$$

The aggregation of local growth rates in the 3-regular score resembles the Ricci curvature’s property of "locally averaging" sectional curvature. Note that both Ricci curvatures encode only structural information of the first and second neighbors, whereas the 3-regular score measures structural information from neighbors up to distance R . Both Ric_F and the 3-regular score are very scalable due to their simple combinatorial notion. Ric_O has limited scalability on large-scale data, since the computation of Wasserstein distances requires solving a linear program for every edge.

To evaluate whether Ricci curvature can select a suitable embedding space, we consider again canonical graphs. Due to their uniform structure, the global

average curvature is equal to the local curvature at any node in the graph. We derive the following results:

Theorem 3.3. *At any node v , we have*

1. $\text{Ric}_O(v) < 0$ and $\text{Ric}_F(v) < 0$ in a b -regular tree,
2. $\text{Ric}_O \leq 0$ and $\text{Ric}_F(v) < 0$ in an $(\sqrt{N} \times \sqrt{N})$ -lattice, and
3. $\text{Ric}_O = 0$ and $\text{Ric}_F = 0$ in an N -cycle.

The proof follows from combinatorial arguments and (for Ric_O) curvature inequalities (Jost and Liu, 2014); it can be found in Appendix C. The theorem shows, that Ricci curvature, similar to Gromov’s δ , correctly detects hyperbolicity, but cannot characterize structures with non-negative curvature. We conclude, that Ricci curvature is not suitable for model selection and that the 3-regular score has broader applicability.

4 Experiments

We have shown above, that in the case of canonical graphs, our approach’s prediction matches known embeddability results. In this section, we want to experimentally validate that the *3-regular score* determines suitable embedding spaces for complex, heterogeneous data. In the following, we report “normalized” 3-regular scores, meaning that we divide by the number of edges in the regularized graph multiplied by the mean edge weight. This adjusts for differences in the average neighborhood size and therefore allows for a comparison across data sets of varying sizes.

Data sets We test our method on both synthetic graphs with known embeddability properties and benchmark data sets. For the former, we create data sets of similar size ($N := |V| \approx 1,000$) to allow for direct comparison. First we generate an N -CYCLE, an $(\sqrt{N} \times \sqrt{N})$ -LATTICE and a ternary TREE ($b = 3$) with N nodes. We further sample from three classic network models: The random graph model (ER) (Erdős and Rényi, 1959), the small world model (WS) (Watts and Strogatz, 1998) and the preferential-attachment model (AB) (Barabási and Albert, 1999) with different choices of hyperparameters. We sample ten networks each and report the average 3-regular score to account for structural sampling variances. Next, we analyze some classic benchmark graphs (both weighted and unweighted) which were downloaded from the Colorado Networks Index (Clauset et al., 2016). The BUNNY data was downloaded from the Stanford 3D Scanning Repository (Stanford Computer Graphics Laboratory, 2014). Finally, for validating our approach against

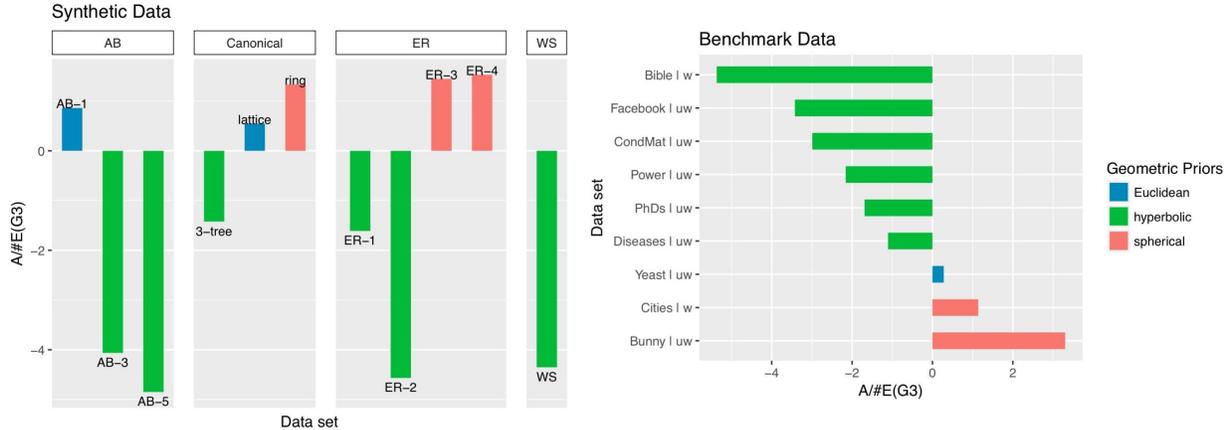


Figure 2: 3-regular scores for synthetic graphs and relational benchmark data (w: weighted, uw: unweighted).

recently published embeddability results, we analyze data sets used in (Gu et al., 2019; Nickel and Kiela, 2017), downloaded from the given original sources. We evaluate geographic distances between North American cities (CITIES, Burkardt), PhD student-adviser relationships (PHD, de Nooy et al. (2004)) and a citation networks (CONDMAT, Leskovec and Krevl (2014)). CITIES contains similarity data from which we created a nearest-neighbor graph, maintaining edges to the top 5% neighboring cities.

Results 3-regular scores for all data sets are shown in Fig. 2. For canonical graphs, the 3-regular score matches both the theoretical results of our growth rate analysis (Thm. 3.2) and embeddability results in the literature (Tab. 2). It is well known that ER undergoes phase transitions as the edge threshold increases. (Chen et al., 2013) show that ER is not hyperbolic in the low edge threshold regime, but that hyperbolicity emerges with the giant component due to its locally tree-like structure. We analyze ER shortly above (ER-3) and below (ER-4) the giant threshold ($p = \frac{1}{N}$) as well as shortly above (ER-2) and below (ER-1) the connectivity threshold ($p = \frac{\log N}{N}$). Consistent with the theoretical result of (Chen et al., 2013), we observe hyperbolicity, if there is a large giant component (ER-1) or if the graph is connected (ER-2). For AB with linear attachment ($m = 1$), the 3-regular score predicts a Euclidean embedding space to be most suitable, as opposed to hyperbolic embeddings for the case of superlinear attachment ($m > 1$). This is again consistent with theoretical results on phase transitions in the AB model. The presence of detectable network communities, as found for instance in social networks, has been repeatedly linked to a locally tree-like structure (Adcock et al., 2013; Krioukov et al., 2010). In agreement with this, the 3-regular score predicts good hyperbolic embeddability for both WS (a classic model for study-

ing community structure) and the social network data sets FACEBOOK and PHD. The wordnet BIBLE was found to embed best into hyperbolic space, in-line with the tree-likeness of such intrinsically hierarchical data. (Sala et al., 2018) observed “less hyperbolicity” in biological networks, which matches our results for DISEASES and YEAST. Finally, BUNNY, a classic benchmark for spherical embeddings, is found to embed best into spherical space.

Validation and comparison with related methods To validate our results, we compare our predicted geometric priors against recently published embeddability results by (Gu et al., 2019; Nickel and Kiela, 2017). Table 4 shows that the 3-regular score predicts the space with the smallest distortion for all benchmark data sets. Here, we follow the authors in reporting *distortion* using the following statistics: The average distortion D_{avg} , computed over all pairwise distances, and the structural distortion score MAP that measures the preservation of nearest-neighbor structures. Isometric embeddability is characterized by $D_{\text{avg}} = 0$ and $\text{MAP} = 1$. For more details, see Appendix D.

Hyperparameters The size of the local neighborhoods $\mathcal{N}_R(v)$ over which we compute the 3-regular score (determined by the neighborhood radius R) is the central hyperparameter in our analysis. Choosing R too small might leave us with too little information to properly evaluate growth rates, whereas a large R limits scalability. First, note that for $R < 3$, $\gamma_{EE} = \gamma_{LE}$ in the regularized graph G_3 . However, for $R \geq 3$ we always have $\gamma_{EE} \neq \gamma_{LE}$. Consequentially, we require $R \geq 3$. Next, we investigated experimentally if an analysis with larger neighborhood radii reveals additional geometric information by computing 3-regular scores for three data sets with different predicted geometric priors for $R \in \{3, 4, 5, 10\}$. For the $(\sqrt{N} \times \sqrt{N})$ -LATTICE, the

Data set	3-regular score	dist \mathbb{R}^{10}	dist \mathbb{S}^{10}	dist \mathbb{H}^{10}	Method
Cities	1.138 $\rightarrow \mathbb{S}^d$	$D_{avg} = 0.074$	$\mathbf{D}_{avg} = \mathbf{0.060}$	$D_{avg} = 0.093$	Gu et al.
PhD students	-1.691 $\rightarrow \mathbb{H}^d$	$D_{avg} = 0.054$ MAP = 0.869	$D_{avg} = 0.057$ MAP = 0.833	$\mathbf{D}_{avg} = \mathbf{0.050}$ $\mathbf{MAP} = \mathbf{0.931}$	Gu et al.
Power	-2.158 $\rightarrow \mathbb{H}^d$	$D_{avg} = 0.092$ $\mathbf{MAP} = \mathbf{0.886}$	$D_{avg} = 0.050$ MAP = 0.795	$\mathbf{D}_{avg} = \mathbf{0.039}$ MAP = 0.844	Gu et al.
Facebook	-3.423 $\rightarrow \mathbb{H}^d$	$D_{avg} = 0.065$ MAP = 0.580	$D_{avg} = 0.066$ MAP = 0.556	$\mathbf{D}_{avg} = \mathbf{0.060}$ $\mathbf{MAP} = \mathbf{0.782}$	Gu et al.
CondMat	-2.991 $\rightarrow \mathbb{H}^d$	MAP = 0.356	-	$\mathbf{MAP} = \mathbf{0.799}$	Nickel, Kiela

Table 3: Comparison with recently published benchmark embeddability results (Gu et al., 2019; Nickel and Kiela, 2017).

3-regular score predicts uniformly $\text{sign}(\kappa) = 0$, i.e. \mathbb{R}^d to be the most suitable embedding space. For WS we observe $\text{sign}(\kappa) = -1$ across all choices of R , predicting hyperbolic embeddability. Finally, for CITIES we observe $\text{sign}(\kappa) = 1$ across the different neighborhood radii. In consequence, the 3-regular scores reported above are all computed for $R = 3$ to maximize scalability.

5 Discussion

In this paper, we introduced a framework for determining a suitable embedding curvature for relational data. Our approach evaluates local neighborhood growth rates based on which we approximate suitable embedding curvatures. We provide theoretical guarantees for canonical graphs and introduce a statistic that efficiently aggregates local growth information, rendering the method applicable to heterogeneous, large-scale graphs (both weighted and unweighted). Moreover, we compare the 3-regular score with commonly used notions of discrete curvature in terms of their ability to measure embeddability. We find that discrete curvature is suitable for detecting hyperbolicity, but not for approximating non-negative sectional curvature. This implies that the 3-regular score is better suited for model space selection.

Contrary to related embeddability methods, our approach is purely combinatorial, circumventing the need to solve costly large-scale optimization problems. Furthermore, the method does not make any a priori assumption on the dimensionality of the embedding space as opposed to related approaches that impose dimensionality constraints or fix the dimension of the target space. Additionally, the locality of the approach confines the analysis to a small subset of the graph at any given time, allowing for a simple parallelization of the method. This increases the algorithm’s scalability

significantly.

Our results tie into the more general problem of finding data representations that reflect intrinsic geometric and topological features. This problem is three-fold: It requires us to determine (i) the sign of the curvature, which in turn determines the model space, i.e. whether to embed in hyperbolic (\mathbb{H}^d), spherical (\mathbb{S}^d) or Euclidean space (\mathbb{R}^d). Furthermore, (ii) the value of the curvature, which determines local and global geometric parameters of the embedding space, such as distance and angle relations in geodesic triangles and lastly (iii) the dimension of the embedding space. The present work mostly focuses on (i) as we restrict our analysis to canonical Riemannian manifolds with constant sectional curvature ($\kappa \in \{\pm 1, 0\}$). By combining this approach with MDS-style embedding methods (Gu et al., 2019; Sala et al., 2018; Bronstein et al., 2006) we could determine the value of the curvature (problem (ii)) also. Hereby, prior knowledge of the sign of the curvature determines the metric of the target space (with the curvature value as hyperparameter) and therefore a suitable objective function to feed into an MDS-style framework. Such a pre-analysis with combinatorial methods should significantly narrow down the search space of suitable curvature values and therefore reduce the overall computational cost. The investigation of such extensions is left for future work.

Acknowledgements

The author thanks James Gurtowski for technical help and Maximilian Nickel for helpful discussions on hyperbolic Embeddings.

References

Ittai Abraham, Yair Bartal, and Ofer Neiman. Embedding metric spaces in their intrinsic dimension. In *Proceedings of the Nineteenth Annual ACM-SIAM*

- Symposium on Discrete Algorithms*, SODA '08, pages 363–372, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- Aaron B. Adcock, Blair D. Sullivan, and Michael W. Mahoney. Tree-like structure in large social and information networks. In *2013 IEEE 13th International Conference on Data Mining*, pages 1–10, Dec 2013.
- Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. In *Science*, volume 286(5439):509–512, 1999.
- Sergio Bermudo, Jose M. Rodriguez, Jose M. Sigarreta, and Jean-Marie Vilaire. Gromov hyperbolic graphs. *Discrete Mathematics*, 313(15):1575 – 1585, 2013. ISSN 0012-365X.
- Jean Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1):46–52, Mar 1985.
- Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999. ISBN 978-3-642-08399-0 978-3-662-12494-9. URL <http://link.springer.com/10.1007/978-3-662-12494-9>.
- Alexander M. Bronstein, Michael M. Bronstein, and Ron Kimmel. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5):1168–1172, 2006.
- John Burkardt. CITIES – city distance datasets. URL <https://rdrr.io/cran/TSP/man/USCA312.html>.
- Benjamin Chamberlain, James Clough, and Marc Deisenroth. Neural embeddings of graphs in hyperbolic space. In *arXiv:1705.10359 [stat.ML]*, 2017.
- Wei Chen, Wenjie Fang, Guangda Hu, and Michael W. Mahoney. On the hyperbolicity of small-world and treelike random graphs. *Internet Mathematics*, 9(4): 434–491, 2013.
- Aaron Clauset, Ellen Tucker, and Matthias Sainz. The Colorado index of complex networks, 2016. URL <https://icon.colorado.edu>.
- Wouter de Nooy, Andrej Mrvar, and Vladimir Batagelj. Exploratory social network analysis with Pajek. Cambridge University Press, 2004.
- Paul Erdős and Alfréd Rényi. On random graphs I. In *Publicationes Mathematicae*, volume 6, page 290–297, 1959.
- Robin Forman. Bochner’s method for cell complexes and combinatorial ricci curvature. *Discrete & Computational Geometry*, 29(3):323–374, 2003.
- Mikhail Leonidovich Gromov. Hyperbolic groups. In *Essays in Group Theory*, volume 8. Mathematical Sciences Research Institute Publications. Springer, NY, 1987.
- Albert Gu, Frederic Sala, Beliz Gunel, and Christopher Ré. Learning mixed-curvature representations in product spaces. In *International Conference on Learning Representations*, 2019.
- Anupam Kumar Gupta. Embedding tree metrics into low dimensional euclidean spaces. In *STOC*, 1999.
- William Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*, volume 26 of *Contemporary Mathematics*, pages 189–206. American Mathematical Society, 1984.
- Jürgen Jost and Shiping Liu. Ollivier’s Ricci curvature, local clustering and curvature-dimension inequalities on graphs. *Discrete & Computational Geometry*, 51(2):300–322, Mar 2014.
- Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. Hyperbolic geometry of complex networks. *Phys. Rev. E*, 82:036106, Sep 2010.
- Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- Pierre-André G Maugis, Sofia C Olhede, and Patrick J Wolfe. Topology reveals universal features for network comparison. *arXiv:1705.05677*, 2017.
- Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems 30*, pages 6338–6347. 2017.
- Maximillian Nickel and Douwe Kiela. Learning continuous hierarchies in the Lorentz model of hyperbolic geometry. In *International Conference on Machine Learning*. 2018.
- Yann Ollivier. A survey of Ricci curvature for metric spaces and Markov chains. In *Probabilistic Approach to Geometry*, pages 343–381. Mathematical Society of Japan, 2010.
- Frederic Sala, Chris De Sa, Albert Gu, and Christopher Re. Representation tradeoffs for hyperbolic embeddings. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 4460–4469, 2018.
- Rik Sarkar. Low distortion Delaunay embedding of trees in hyperbolic plane. In *Graph Drawing*, pages 355–366, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-25878-7.

- Stanford Computer Graphics Laboratory. The Stanford 3D Scanning Repository, 2014. URL <http://graphics.stanford.edu/data/3Dscanrep/#bunny>.
- Alexandru Tifrea, Gary Bécigneul, and Octavian-Eugen Ganea. Poincaré glove: Hyperbolic word embeddings. *ICRL*, 2019.
- Kevin Verbeek and Subhash Suri. Metric embedding, hyperbolic space, and social networks. *Computational Geometry*, 59:1 – 12, 2016. ISSN 0925-7721.
- Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. In *Nature*, volume 393(6684):440–442, 1998.
- Melanie Weber and Maximilian Nickel. Curvature and representation learning: Identifying embedding spaces for relational data. In *NeurIPS Relational Representation Learning*, 2018.
- Melanie Weber, Emil Saucan, and Jürgen Jost. Characterizing complex networks with Forman-Ricci curvature and associated geometric flows. *Journal of Complex Networks*, 5(4):527–550, 2017a.
- Melanie Weber, Emil Saucan, and Jürgen Jost. Coarse geometry of evolving networks. *Journal of Complex Networks*, 6(5):706–732, 2017b.
- Richard C. Wilson, Edwin R. Hancock, Elzbieta Pekalska, and Robert P. W. Duin. Spherical and Hyperbolic Embeddings of Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(11):2255–2269, 2014.