
Exact Representation of Sparse Networks with Symmetric Nonnegative Embeddings

Sudhanshu Chanpuriya¹, Ryan A. Rossi², Anup Rao², Tung Mai²,
Nedim Lipka², Zhao Song², and Cameron Musco³

¹University of Illinois Urbana-Champaign, schariya@illinois.edu

²Adobe Research, {ryrossi, anuprao, tumai, lipka, zsong}@adobe.com

³University of Massachusetts Amherst, cmusco@cs.umass.edu

Abstract

Graph models based on factorization of the adjacency matrix often fail to capture network structures related to links between dissimilar nodes (heterophily). We introduce a novel graph factorization model that leverages two nonnegative vectors per node to interpretably account for links between both similar and dissimilar nodes. We prove that our model can exactly represent any graph with low *arboricity*, a property that many real-world networks satisfy; our proof also applies to related models but has much greater scope than the closest prior bound, which is based on low *max degree*. Our factorization also has compelling properties besides expressiveness: due to its symmetric structure and nonnegativity, fitting the model inherently finds node communities, and the model’s link predictions can be interpreted in terms of these communities. In experiments on real-world networks, we demonstrate our factorization’s effectiveness on a variety of tasks, including community detection and link prediction.

1 Introduction

Graphs data naturally arises in a variety of fields including sociology (Mason & Verwoerd, 2007), biology (Scott, 1988), and computer networking (Bonato, 2004). A key task in machine learning for graph data is forming models of graphs that can predict edges between nodes, form useful representations of nodes, and reveal interpretable structure in the graph, such as detecting clusters of nodes. Many graph models fall under the framework of edge-independent graph generative models, which output the probabilities of edges existing between any pair of nodes. The parameters of such models can be trained iteratively on the network, or some fraction of the network which is known, in the link prediction task, i.e., by minimizing a predictive loss. To choose among these models, one must consider two criteria: 1) whether the model can express structures of interest in the graph, 2) whether the model expresses these structures in an interpretable way.

Expressiveness of low-dimensional embeddings As real-world graphs are high-dimensional objects, graph models generally compress information about the graph. For example, dot product models associate each node with a real-valued “embedding” vector; the predicted probability of a link between two nodes increases with the similarity of their embeddings. These models can alternatively be seen as factorizing the graph’s adjacency matrix to approximate it with a low-rank matrix. Recent work of Seshadhri et al. (2020) has shown that dot product models are limited in their ability to model common structures in real-world graphs, such as triangles incident only on low-degree nodes. In response, Chanpuriya et al. (2020) showed that with the logistic principal components analysis (LPCA) model, which has two embeddings per node (i.e., using the dot product of the ‘left’ embedding of one node and the ‘right’ embedding of another), not only can such structures be represented, but further, any graph can be exactly represented with embedding vectors whose lengths are linear in the

max degree of the graph. There are two keys to this result. First is the presence of a nonlinear linking function in LPCA; since adjacency matrices are generally not low-rank, exact low-rank factorization is impossible without a linking function. Second is that having two embeddings rather than one allows for expression of non-positive semidefinite (PSD) matrices. As discussed in [Peysakhovich & Bottou \(2021\)](#), that the single-embedding models can only represent PSD matrices precludes representation of ‘heterophilous’ structures in graphs; heterophilous structures are those wherein dissimilar nodes are linked, in contrast to more intuitive ‘homophilous’ linking between similar nodes.

Interpretability and node clustering Beyond being able to capture a given network accurately, it is often desirable for a graph model to form interpretable representations of nodes and to produce edge probabilities in an interpretable fashion. Dot product models can achieve this by restricting the node embeddings to be nonnegative. Nonnegative factorization has long been used to decompose data into parts ([Donoho & Stodden, 2003](#)). In the context of graphs, this entails decomposing the set of nodes of the network into clusters or communities. In particular, each entry of the nonnegative embedding vector of a node represents the intensity with which the node participates in a community. This allows the edge probabilities output by dot product models to be interpretable in terms of coparticipation in communities. Depending on the model, these vectors may have restrictions such as a sum-to-one requirement, meaning the node is assigned a categorical distribution over communities. The least restrictive and most expressive case is that of soft assignments to overlapping communities, where the entries can vary totally independently. In such models, which include the BIGCLAM model of [Yang & Leskovec \(2013\)](#), the output of the dot product may be mapped through a nonlinear link function (as in LPCA) to produce a probability for each edge, i.e., to ensure the values lie in $[0, 1]$.

Summary of contributions The key contributions of this work are as follows:

- We prove that LPCA admits exact low-rank factorizations of graphs with bounded *arboricity*, which is the minimum number of forests into which a graph’s edges can be partitioned. By the Nash-Williams theorem, arboricity is a measure of a graph’s density in that, letting S denote an induced subgraph and n_S and m_S denote the number of nodes and edges in S , arboricity is the maximum over all S of $\lceil \frac{m_S}{n_S-1} \rceil$. Our result is more applicable to real-world graphs than the prior result of [Chanpuriya et al. \(2020\)](#) for graphs with bounded max degree, since sparsity is a common feature of real networks, whereas low max degree is not.
- We introduce a graph model that extends LPCA and is both highly expressive and interpretable. Our model incorporates two embeddings per node and a nonlinear linking function, and hence is able to express both heterophily and overlapping communities. At the same time, our model is based on symmetric nonnegative matrix factorization, so it outputs link probabilities that are interpretable in terms of the communities it detects. While prior graph factorizations incorporate some aspects of nonnegativity, heterophily, and/or nonlinearity, our proposed model lies at the intersection of all three.
- We show how any graph with a low-rank factorization in the LPCA model also admits a low-rank factorization in our community-based model. This means that the guarantees on low-rank representation for bounded max degree and arboricity also apply to our model.
- In experiments, we show that our method is competitive with and often outperforms other comparable models on real-world graphs in terms of representing the network, doing interpretable link prediction, and detecting communities that align with ground-truth.

2 Motivating Example

To demonstrate how heterophily can manifest in networks, as well as how models which assume homophily can fail to represent such networks, we provide a simple synthetic example involving a graph of matches between users of a dating app. Suppose that users of this app are generally seeking partners of a different gender (to simplify this example, we assume that each user of the app is either a man or a woman), and suppose that each user comes from one of ten cities. Members from the same city are likely to match with each other; this typifies homophily, wherein links occur between similar nodes. Furthermore, users having the same gender are unlikely to match with each other; this typifies heterophily. [Figure 1](#) shows an instantiation of such an adjacency matrix with 1000 nodes, which are randomly assigned to man or woman and to one of the ten cities. We recreate this network with our proposed embedding model and with BIGCLAM, which explicitly assumes

homophily. We also compare with the SVD of the adjacency matrix, which outputs the best (lowest Frobenius error) low-rank approximation that is possible without a nonlinear linking function. Since SVD lacks nonnegativity constraints on the factors, we do not expect interpretability. In Figure 2, we show how BIGCLAM captures only the ten communities based on city, i.e., only the homophilous structure, and fails to capture the heterophilous distinction between men and women. We also plot the error of the reconstructions as the embedding length increases. There are $10 \cdot 2 = 20$ different kinds of nodes, meaning the expected adjacency matrix is rank-20, and our model maintains the lowest error up to this embedding length; by contrast, BIGCLAM is unable to decrease error after capturing city information with length-10 embeddings. In Figure 3, we visualize the features generated by the three methods, i.e., the factors returned by each factorization. Our model’s factors capture the relevant latent structure in an interpretable way. By contrast, SVD’s factors are harder to interpret, and BIGCLAM does not represent the heterophilous structure.

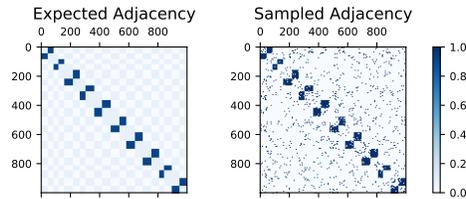


Figure 1: The motivating synthetic graph. The expected adjacency matrix (left) and the sampled matrix (right); the latter, which is passed to the training algorithms, is produced by treating the entries of the former as parameters of Bernoulli distributions and sampling. The network is approximately a union of ten bipartite graphs, each of which correspond to men and women in one of the ten cities.

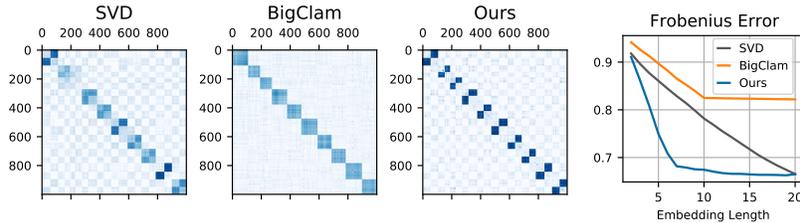


Figure 2: Left: Reconstructions of the motivating synthetic graph of Figure 1 with SVD, BIGCLAM, and our model, using 12 communities or singular vectors. Note the lack of the small diagonal structure in BIGCLAM’s reconstruction; this corresponds to its inability to capture the heterophilous interaction between men and women. Right: Frobenius error when reconstructing the motivating synthetic graph of Figure 1 with SVD, BIGCLAM, and our model, as the embedding length is varied. The error is normalized by the sum of the true adjacency matrix (i.e., twice the number of edges).

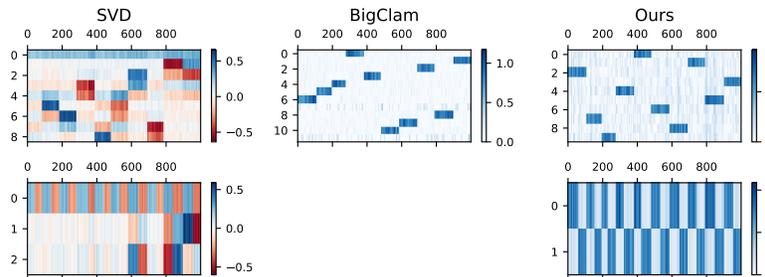


Figure 3: Factors resulting from decomposition of the motivating synthetic graph of Figure 1 with the three models, using 12 communities or singular vectors. The top/bottom rows represent the positive/negative eigenvalues corresponding to homophilous/heterophilous communities (note that BIGCLAM does not include the latter). The homophilous factors from BIGCLAM and our model reflect the 10 cities, and the heterophilous factor from our model reflect men and women. The factors from SVD are harder to interpret. Note that the order of the communities in the factors is arbitrary.

3 Community-Based Graph Factorization

Consider the set of undirected, unweighted graphs on n nodes, i.e., the set of graphs with symmetric adjacency matrices in $\{0, 1\}^{n \times n}$. We propose an edge-independent generative model for such graphs. Given nonnegative parameter matrices $\mathbf{B} \in \mathbb{R}_+^{n \times k_B}$ and $\mathbf{C} \in \mathbb{R}_+^{n \times k_C}$, we set the probability of an edge existing between nodes i and j to be the (i, j) -th entry of matrix $\tilde{\mathbf{A}}$:

$$\tilde{\mathbf{A}} := \sigma(\mathbf{B}\mathbf{B}^\top - \mathbf{C}\mathbf{C}^\top), \quad (1)$$

where σ is the logistic function. Here k_B, k_C are the number of homophilous/heterophilous clusters. Intuitively, if $\mathbf{b}_i \in \mathbb{R}_+^{k_B}$ is the i -th row of matrix \mathbf{B} , then \mathbf{b}_i is the affinity of node i to each of the k_B homophilous communities. Similarly, $\mathbf{c}_i \in \mathbb{R}_+^{k_C}$ is the affinity of node i to the k_C heterophilous communities. As an equivalent statement, for each pair of nodes i and j , $\tilde{\mathbf{A}}_{i,j} := \sigma(\mathbf{b}_i \mathbf{b}_j^\top - \mathbf{c}_i \mathbf{c}_j^\top)$. We will soon discuss the precise interpretation of this model, but the idea is roughly similar to the attract-repel framework of Peysakhovich & Bottou (2021). When nodes i and j have similar ‘attractive’ \mathbf{b} embeddings, i.e., when $\mathbf{b}_i \mathbf{b}_j^\top$ is high, the likelihood of an edge between them increases, hence why the \mathbf{B} factor is homophilous. By contrast, the \mathbf{C} factor is ‘repulsive’/heterophilous since, when $\mathbf{c}_i \mathbf{c}_j^\top$ is high, the likelihood of an edge between i and j decreases.

Alternate expression We note that the model above can also be expressed in a form that normalizes cluster assignments and is more compact, in that it combines the homophilous and heterophilous cluster assignments. Instead of \mathbf{B} and \mathbf{C} , this form uses a matrix $\mathbf{V} \in [0, 1]^{n \times k}$ and a diagonal matrix $\mathbf{W} \in \mathbb{R}^{k \times k}$, where $k = k_B + k_C$ is the total number of clusters. In particular, let \mathbf{m}_B and \mathbf{m}_C be the vectors containing the maximums of each column of \mathbf{B} and \mathbf{C} . By setting

$$\begin{aligned} \mathbf{V} &= (\mathbf{B} \cdot \text{diag}(\mathbf{m}_B^{-1}); \mathbf{C} \cdot \text{diag}(\mathbf{m}_C^{-1})) \\ \mathbf{W} &= \text{diag}((+\mathbf{m}_B^2; -\mathbf{m}_C^2)), \end{aligned} \quad (2)$$

the constraint on \mathbf{V} is satisfied. Further, $\mathbf{V}\mathbf{W}\mathbf{V}^\top = \mathbf{B}\mathbf{B}^\top - \mathbf{C}\mathbf{C}^\top$, so

$$\tilde{\mathbf{A}} := \sigma(\mathbf{B}\mathbf{B}^\top - \mathbf{C}\mathbf{C}^\top) = \sigma(\mathbf{V}\mathbf{W}\mathbf{V}^\top). \quad (3)$$

Here, if $\mathbf{v}_i \in [0, 1]^k$ is the i -th row of matrix \mathbf{V} , then \mathbf{v}_i is the soft (normalized) assignment of node i to the k communities. The diagonal entries of \mathbf{W} represent the strength of the homophily (if positive) or heterophily (if negative) of the communities. For each entry, $\tilde{\mathbf{A}}_{i,j} = \sigma(\mathbf{v}_i \mathbf{W} \mathbf{v}_j^\top)$. We use these two forms interchangeably throughout this work.

Interpretation The edge probabilities output by this model have an intuitive interpretation. Recall that there are bijections between probability $p \in [0, 1]$, odds $o = \frac{p}{1-p} \in [0, \infty)$, and logit $\ell = \log(o) \in (-\infty, +\infty)$. The logit of the link probability between nodes i and j is $\mathbf{v}_i^\top \mathbf{W} \mathbf{v}_j$, which is a summation of terms $\mathbf{v}_{ic} \mathbf{v}_{jc} \mathbf{W}_{cc}$ over all communities $c \in [k]$. If the nodes both fully participate in community c , that is, $\mathbf{v}_{ic} = \mathbf{v}_{jc} = 1$, then the edge logit is changed by \mathbf{W}_{cc} starting from a baseline of 0, or equivalently, the odds of an edge is multiplied by $\exp(\mathbf{W}_{cc})$ starting from a baseline odds of 1; if either of the nodes participates only partially in community c , then the change in logit and odds is accordingly prorated. Homophily and heterophily also have a clear interpretation in this model: homophilous communities, which are expressed in \mathbf{B} , are those with $\mathbf{W}_{cc} > 0$, where two nodes both participating in the community increases the odds of a link, whereas communities with $\mathbf{W}_{cc} < 0$, which are expressed in \mathbf{C} , are heterophilous, and coparticipation decreases the odds of a link.

4 Related Work

Community detection via interpretable factorizations There is extensive prior work on community detection and node clustering (Schaeffer, 2007; Aggarwal & Wang, 2010; Nascimento & De Carvalho, 2011), perhaps the most well-known being the normalized cuts algorithm of Shi & Malik (2000), which produces a clustering based on the entrywise signs of an eigenvector of the graph Laplacian matrix. However, the clustering algorithms which are most relevant to our work are those based on non-negative matrix factorization (NMF) (Lee & Seung, 1999; Berry et al., 2007; Wang & Zhang, 2012; Gillis, 2020), many of which can be seen as integrating nonnegativity constraints into

the broader, well-studied *random dot product graph* (RDPG) model (Young & Scheinerman, 2007; Scheinerman & Tucker, 2010; Athreya et al., 2017). One such algorithm is that of Yu et al. (2005), which approximately factors a graph’s adjacency matrix $A \in \{0, 1\}^{n \times n}$ into two positive matrices H and Λ , where $H \in \mathbb{R}_+^{n \times k}$ is left-stochastic and $\Lambda \in \mathbb{R}_+^{k \times k}$ is diagonal, such that $H\Lambda H^\top \approx A$. Here H represents a soft clustering of the n nodes into k clusters, while the diagonal entries of Λ represent the prevalence of edges within clusters. Note the similarity of the factorization to our model, save for the lack of a nonlinearity. Other NMF approaches include those of Ding et al. (2008), Yang et al. (2012), Kuang et al. (2012), and Kuang et al. (2015) (SYMNMF).

Modeling heterophily Much of the existing work on graph models has an underlying assumption of network homophily (Johnson et al., 2010; Noldus & Van Mieghem, 2015). There has been significant recent interest in the limitations of graph neural network (GNN) models (Duvenaud et al., 2015; Kipf & Welling, 2017; Hamilton et al., 2017) at addressing network heterophily (NT & Maehara, 2019; Zhu et al., 2020; Zheng et al., 2022), as well as proposed solutions (Pei et al., 2020; Yan et al., 2021), but relatively less work for models applicable to clustering. Some existing NMF approaches to clustering do naturally model heterophilous structure in networks. The model of Nourbakhsh et al. (2014), for example, is similar to that of Yu et al. (2005), but allows the cluster affinity matrix Λ to be non-diagonal; this allows for inter-cluster edge affinity to exceed intra-cluster edge affinity, so heterophily can arise in this model, though it is not a focus of their work. Another example is the model in Miller et al. (2009), which is similar to ours, though it restricts the cluster assignment matrix V to be binary; additionally, their training algorithm is not based on gradient descent as ours is, and it does not scale to larger networks. More recently, Rubin-Delanchy et al. (2017) and Peysakhovich & Bottou (2021) propose simple decompositions which allow for representation of non-PSD adjacency matrices. The model in the latter work is a factorization of the form $A \approx D + BB^\top - CC^\top$, where $D \in \mathbb{R}^{n \times n}$ is diagonal and $B, C \in \mathbb{R}^{n \times k}$ are low-rank; excluding the diagonal D term, the model in the former work is algebraically identical. The authors discuss how, interestingly, this model separates the homophilous and heterophilous structure into different factors, namely B and C , corresponding to positive and negative eigenvalues, respectively. Note that these decompositions do not include a nonlinear linking function, which is crucial to our exact factorization results, and the respective works do not investigate constraining the factors to be nonnegative.

Overlapping communities and exact embeddings Many models discussed above focus on the single-label clustering task and thus involve highly-constrained factorizations (e.g., sum-to-one conditions). We are interested in the closely related but distinct task of multi-label clustering, also known as overlapping community detection (Xie et al., 2013; Javed et al., 2018), which involves less constrained, more expressive factorizations. The BIGCLAM algorithm of Yang & Leskovec (2013) uses the following generative model for this task: the probability of a link between two nodes i and j is given by $1 - \exp(-\mathbf{f}_i \cdot \mathbf{f}_j)$, where $\mathbf{f}_i, \mathbf{f}_j \in \mathbb{R}_+^k$ represent the intensities with which the nodes participate in each of the k communities. Note that BIGCLAM assumes strict homophily of the communities: two nodes participating in the same community always increases the probability of a link. However, this model allows for expression of very dense intersections of communities, which the authors observe is generally a characteristic of real-world networks. To ensure that output entries are probabilities, BIGCLAM’s factorization includes a nonlinear linking function (namely, $f(x) = 1 - e^{-x}$), like our model and LPCA. Recent work outside clustering and community detection on graph generative models (Rendsburg et al., 2020; Chanpuriya et al., 2020) suggests that incorporating a nonlinear linking function can greatly increase the expressiveness of factorization-based graph models, to the point of being able to exactly represent a graph. This adds to a growing body of literature on expressiveness guarantees for embeddings on relational data (Sala et al., 2018; Bhattacharjee & Dasgupta, 2020; Boratko et al., 2021). Most relevant to our work, as previously discussed, Chanpuriya et al. (2020) provide a guarantee for exact low-rank representation of graphs with bounded max degree when using the LPCA factorization model. In this work, we provide a new such guarantee, except for bounded arboricity, which is more applicable to real-world networks, and extend these guarantees to our community-based factorization.

5 Theoretical Results

We first restate the main result from Chanpuriya et al. (2020) on exact representation of graphs with bounded max degree using the logistic principal components analysis (LPCA) model, which

reconstructs a graph $\mathbf{A} \in \{0, 1\}^{n \times n}$ using logit factors $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times k}$ via

$$\mathbf{A} \approx \sigma(\mathbf{X}\mathbf{Y}^\top). \quad (4)$$

Note that unlike our community-based factorization, the factors of the LPCA model are not nonnegative, and the factorization does not reflect the symmetry of the undirected graph's adjacency matrix. Regardless of the model's interpretability, the following theorem provides a significant guarantee on its expressiveness. We use the following notation: given a matrix \mathbf{M} , let $H(\mathbf{M})$ denote the matrix resulting from entrywise application of the Heaviside step function to \mathbf{M} , that is, setting all positive entries to 1, negative entries to 0, and zero entries to $1/2$.

Theorem 5.1 (Exact LPCA Factorization for Bounded-Degree Graphs). *Let $\mathbf{A} \in \{0, 1\}^{n \times n}$ be the adjacency matrix of a graph G with maximum degree c . Then there exist matrices $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times (2c+1)}$ such that $\mathbf{A} = H(\mathbf{X}\mathbf{Y}^\top)$.*

This corresponds to arbitrarily small approximation error in the LPCA model (Equation 4) because, provided such factors \mathbf{X}, \mathbf{Y} for some graph \mathbf{A} , we have that $\lim_{s \rightarrow \infty} \sigma(s\mathbf{X}\mathbf{Y}^\top) = H(\mathbf{X}\mathbf{Y}^\top) = \mathbf{A}$. That is, we can scale the factors larger to reduce the error to an arbitrary extent.

We expand on this result in two ways. First, give a new bound for exact embedding in terms of arboricity, rather than max degree. This increases the applicability to real-world networks, which often are sparse (i.e., low arboricity) and have right-skewed degree distributions (i.e., high max degree). Second, we show that any rank- k LPCA factorization can be converted to our model's symmetric nonnegative factorization with $O(k)$ communities. This extends the guarantees on the LPCA model's power for exact representation of graphs, both the prior guarantee in terms of max degree and our new one in terms of arboricity, to our community-based model as well. In Appendix A.1, we also introduce a natural family of graphs - Community Overlap Threshold (COT) graphs - for which our model's community-based factorization not only exactly represents the graph, but also must capture some latent structure to do so with sufficiently low embedding dimensionality.

Arboricity bound for exact representation We will use the following well-known fact: the rank of the entrywise product of two matrices is at most the product of their individual ranks, that is,

$$\text{rank}(\mathbf{X} \circ \mathbf{Y}) \leq \text{rank}(\mathbf{X}) \cdot \text{rank}(\mathbf{Y}).$$

Theorem 5.2 (Exact LPCA Factorization for Bounded-Arboricity Graphs). *Let $\mathbf{A} \in \{0, 1\}^{n \times n}$ be the adjacency matrix of an undirected graph G with arboricity α . Then there exist embeddings $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times (4\alpha^2+1)}$ such that $\mathbf{A} = H(\mathbf{X}\mathbf{Y}^\top)$.*

Proof. Let the undirected graph \mathbf{A} have arboricity α , i.e., the edges can be partitioned into α forests. We produce a directed graph \mathbf{B} from \mathbf{A} by orienting the edges in these forests so that each node's edges point towards its children. Now $\mathbf{A} = \mathbf{B} + \mathbf{B}^\top$, and every node in \mathbf{B} has in-degree at most α .

Let $\mathbf{V} \in \mathbb{R}^{n \times 2\alpha}$ be the Vandermonde matrix with $V_{i,j} = t^{j-1}$. For any $\mathbf{c} \in \mathbb{R}^{2\alpha}$, $[\mathbf{V}\mathbf{c}](t) = \sum_{j=1}^{2\alpha} c_j \cdot t^{j-1}$, that is, $\mathbf{V}\mathbf{c} \in \mathbb{R}^n$ is a degree- (2α) polynomial with coefficients \mathbf{c} evaluated at the integers $t \in [n] = \{1, \dots, n\}$. Let \mathbf{b}_i be the i^{th} column of \mathbf{B} . We seek to construct a polynomial such that for t with $\mathbf{b}_i(t) = 1$, $[\mathbf{V}\mathbf{c}_i](t) = 0$, and $[\mathbf{V}\mathbf{c}_i](t) < 0$ elsewhere; that is, when inputting an index $t \in [n]$ such that the t^{th} node is an in-neighbor of the i^{th} node, we want the polynomial to output 0, and for all other indices in $[n]$, we want it to have a negative output. Letting $N(i)$ denote the in-neighbors of the i^{th} node, a simple instantiation of such a polynomial in t is $-1 \cdot \prod_{j \in N(i)} (t - j)^2$. Note that since all nodes have in-degree at most α , this polynomial's degree is at most 2α , and hence there exists a coefficient vector $\mathbf{c}_i \in \mathbb{R}^{2\alpha}$ encoding this polynomial.

Let $\mathbf{C} \in \mathbb{R}^{n \times 2\alpha}$ be the matrix resulting from stacking such coefficient vectors for each of the n nodes. Consider $\mathbf{P} = \mathbf{V}\mathbf{C} \in \mathbb{R}^{n \times n}$. $\mathbf{P}_{i,j}$ is 0 if $\mathbf{B}_{i,j} = 1$ and negative otherwise. Then $(\mathbf{P} \circ \mathbf{P}^\top)_{i,j}$ is 0 when either $\mathbf{B}_{i,j} = 1$ or $(\mathbf{B}^\top)_{i,j} = 1$ and positive otherwise; equivalently, since $\mathbf{A} = \mathbf{B} + \mathbf{B}^\top$, $(\mathbf{P} \circ \mathbf{P}^\top)_{i,j} = 0$ iff $\mathbf{A}_{i,j} = 1$. Take any positive ϵ less than the smallest positive entry of $\mathbf{P} \circ \mathbf{P}^\top$. Letting \mathbf{J} be an all-ones matrix, define $\mathbf{M} = \epsilon\mathbf{J} - (\mathbf{P} \circ \mathbf{P}^\top)$. Note that $\mathbf{M}_{i,j} > 0$ if $\mathbf{A}_{i,j} = 1$ and $\mathbf{M}_{i,j} < 0$ if $\mathbf{A}_{i,j} = 0$, that is, $\mathbf{M} = H(\mathbf{A})$ as desired. Since $\text{rank}(\mathbf{J}) = 1$ and $\text{rank}(\mathbf{P}) \leq 2\alpha$, by the bound on the rank of entrywise products of matrices, the rank of \mathbf{M} is at most $(2\alpha)^2 + 1$. ■

Exact representation with community factorization LPCA factors $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times k}$ can be processed into nonnegative factors $\mathbf{B} \in \mathbb{R}_+^{n \times k_B}$ and $\mathbf{C} \in \mathbb{R}_+^{n \times k_C}$ such that $k_B + k_C = 6k$ and

$$\mathbf{B}\mathbf{B}^\top - \mathbf{C}\mathbf{C}^\top = \frac{1}{2}(\mathbf{X}\mathbf{Y}^\top + \mathbf{Y}\mathbf{X}^\top). \quad (5)$$

As a rough outline of the argument that follows, we will need $6k$ columns in the new factors \mathbf{B}, \mathbf{C} , up from the k columns in \mathbf{X}, \mathbf{Y} , because accounting for the possible asymmetry in $\mathbf{X}\mathbf{Y}^\top$ will double the required columns, and accounting for the nonnegativity of \mathbf{B}, \mathbf{C} will triple the required columns. Observe that the left-hand side can only represent symmetric matrices, but $\mathbf{X}\mathbf{Y}^\top$ is not necessarily symmetric even if $H(\mathbf{X}\mathbf{Y}^\top) = \mathbf{A}$ for a symmetric \mathbf{A} . For this reason, we use a symmetrization: let $\mathbf{L} = \frac{1}{2}(\mathbf{X}\mathbf{Y}^\top + \mathbf{Y}\mathbf{X}^\top)$. Note that $H(\mathbf{L}) = H(\mathbf{X}\mathbf{Y}^\top)$, so if $\mathbf{X}\mathbf{Y}^\top$ constitutes an exact representation of \mathbf{A} in that $H(\mathbf{X}\mathbf{Y}^\top) = \mathbf{A}$, so too do both expressions for \mathbf{L} in Equation 5. Pseudocode for the procedure of constructing \mathbf{B}, \mathbf{C} given \mathbf{X}, \mathbf{Y} is given in Algorithm 1. The concept of this algorithm is to first separate the logit matrix \mathbf{L} into a sum and difference of rank-1 components via eigendecomposition. Each of these components can be written as $+\mathbf{v}\mathbf{v}^\top$ or $-\mathbf{v}\mathbf{v}^\top$ with $\mathbf{v} \in \mathbb{R}^n$, where the sign depends on the sign of the eigenvalue. Each component is then separated into a sum and difference of three outer products of nonnegative vectors, via the following Lemma 5.3.

Lemma 5.3. *Let $\phi : \mathbb{R} \rightarrow \mathbb{R}$ denote the ReLU function, i.e., $\phi(z) = \max\{z, 0\}$. For any vector \mathbf{v} ,*

$$\mathbf{v}\mathbf{v}^\top = 2\phi(\mathbf{v})\phi(\mathbf{v})^\top + 2\phi(-\mathbf{v})\phi(-\mathbf{v})^\top - |\mathbf{v}||\mathbf{v}|^\top.$$

Proof. Take any $\mathbf{v} \in \mathbb{R}^k$. Then

$$\begin{aligned} \mathbf{v}\mathbf{v}^\top &= (\phi(\mathbf{v}) - \phi(-\mathbf{v})) \cdot (\phi(\mathbf{v})^\top - \phi(-\mathbf{v})^\top) \\ &= +\phi(\mathbf{v})\phi(\mathbf{v})^\top + \phi(-\mathbf{v})\phi(-\mathbf{v})^\top - \phi(\mathbf{v})\phi(-\mathbf{v})^\top - \phi(-\mathbf{v})\phi(\mathbf{v})^\top \\ &= +2\phi(\mathbf{v})\phi(\mathbf{v})^\top + 2\phi(-\mathbf{v})\phi(-\mathbf{v})^\top - (\phi(\mathbf{v}) + \phi(-\mathbf{v})) \cdot (\phi(\mathbf{v}) + \phi(-\mathbf{v}))^\top \\ &= +2\phi(\mathbf{v})\phi(\mathbf{v})^\top + 2\phi(-\mathbf{v})\phi(-\mathbf{v})^\top - |\mathbf{v}||\mathbf{v}|^\top, \end{aligned}$$

where the first step follows from $\mathbf{v} = \phi(\mathbf{v}) - \phi(-\mathbf{v})$, and the last step from $|\mathbf{v}| = \phi(\mathbf{v}) + \phi(-\mathbf{v})$. ■

Algorithm 1 follows from Lemma 5.3 and constitutes a constructive proof of the following theorem:

Theorem 5.4 (Exact Community Factorization from Exact LPCA Factorization). *Given a symmetric matrix $\mathbf{A} \in \{0, 1\}$ and $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times k}$ such that $\mathbf{A} = H(\mathbf{X}\mathbf{Y}^\top)$, there exist nonnegative matrices $\mathbf{B} \in \mathbb{R}_+^{n \times k_B}$ and $\mathbf{C} \in \mathbb{R}_+^{n \times k_C}$ such that $k_B + k_C = 6k$ and $\mathbf{A} = H(\mathbf{B}\mathbf{B}^\top - \mathbf{C}\mathbf{C}^\top)$.*

Algorithm 1 Converting LPCA Factors to Community Factors

Input logit factors $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times k}$

Output $\mathbf{B} \in \mathbb{R}_+^{n \times k_B}, \mathbf{C} \in \mathbb{R}_+^{n \times k_C}$ such that $k_B + k_C = 6k$
and $\mathbf{B}\mathbf{B}^\top - \mathbf{C}\mathbf{C}^\top = \frac{1}{2}(\mathbf{X}\mathbf{Y}^\top + \mathbf{Y}\mathbf{X}^\top)$

- 1: Set $\mathbf{Q} \in \mathbb{R}^{n \times 2k}$ and $\boldsymbol{\lambda} \in \mathbb{R}^{2k}$ by truncated eigendecomposition such that $\mathbf{Q} \times \text{diag}(\boldsymbol{\lambda}) \times \mathbf{Q}^\top = \frac{1}{2}(\mathbf{X}\mathbf{Y}^\top + \mathbf{Y}\mathbf{X}^\top)$
 - 2: $\mathbf{B}^* \leftarrow \mathbf{Q}^+ \times \text{diag}(\sqrt{+\boldsymbol{\lambda}^+})$, where $\boldsymbol{\lambda}^+, \mathbf{Q}^+$ are the positive eigenvalues/vectors
 - 3: $\mathbf{C}^* \leftarrow \mathbf{Q}^- \times \text{diag}(\sqrt{-\boldsymbol{\lambda}^-})$, where $\boldsymbol{\lambda}^-, \mathbf{Q}^-$ are the negative eigenvalues/vectors
 - 4: $\mathbf{B} \leftarrow (\sqrt{2}\phi(\mathbf{B}^*); \sqrt{2}\phi(-\mathbf{B}^*); |\mathbf{C}^*|) \triangleright \phi$ and $|\cdot|$ are entrywise ReLU and absolute value
 - 5: $\mathbf{C} \leftarrow (\sqrt{2}\phi(\mathbf{C}^*); \sqrt{2}\phi(-\mathbf{C}^*); |\mathbf{B}^*|)$
 - 6: **return** \mathbf{B}, \mathbf{C}
-

As stated in the introduction to this section, Theorem 5.4 extends any upper bound on the exact factorization dimensionality from the LPCA model to our community-based model. That is, up to a constant factor, the bound in terms of max degree from Theorem 5.1 and the bound in terms of arboricity from Theorem 5.2 also apply to our model; for brevity, we state just the latter here.

Corollary 5.5 (Exact Community Factorization for Bounded-Arboricity Graphs). *Let $\mathbf{A} \in \{0, 1\}^{n \times n}$ be the adjacency matrix of an undirected graph G with arboricity α . Then there exist nonnegative embeddings $\mathbf{B} \in \mathbb{R}_+^{n \times k_B}$ and $\mathbf{C} \in \mathbb{R}_+^{n \times k_C}$ such that $k_B + k_C = 6(4\alpha^2 + 1)$ and $\mathbf{A} = H(\mathbf{B}\mathbf{B}^\top - \mathbf{C}\mathbf{C}^\top)$.*

Note that Corollary 5.5 is purely a statement about the capacity of our model; Theorem 5.2 stems from a constructive proof based on polynomial interpolation, and therefore so too does this corollary. We do not expect this factorization to be informative about the graph’s latent structure. In the following Section 6, we will fit the model with an entirely different algorithm for downstream applications.

6 Experiments

We now present a training algorithm to fit our model, then evaluate our method on a benchmark of five real-world networks.

6.1 Training Algorithm

Given an input graph $A \in \{0, 1\}^{n \times n}$, we find low-rank nonnegative matrices B and C such that the model produces $\tilde{A} = \sigma(BB^\top - CC^\top) \in (0, 1)^{n \times n}$ as in Equation 1 which approximately matches A . In particular, we train the model to minimize the sum of binary cross-entropies of the link predictions over all pairs of nodes:

$$R = - \sum_{ij} \left(A_{ij} \log(\tilde{A}_{ij}) + (1 - A_{ij}) \log(1 - \tilde{A}_{ij}) \right). \quad (6)$$

We fit the parameters by gradient descent over this loss, as well as L_2 regularization of the factors B and C , subject to the nonnegativity of B and C . This algorithm is fairly straightforward; pseudocode is given in Algorithm 2. This is quite similar to the training algorithm of Chanpuriya et al. (2020), but in contrast to that work, which only targets an exact fit, we explore the expression of graph structure in the factors and their utility in downstream tasks. Regularization of the factors is implemented to this end to avoid overfitting. Though in the main paper we outline and evaluate a non-stochastic version of the training algorithm, it can generalize straightforwardly to a more scalable stochastic version, e.g., by sampling links and non-links for the loss function and using projected SGD. In Appendix A.2, we discuss an industry application of our model to tabular dataset completion, for which we employ such stochastic training. Here, we use the simpler non-stochastic training to isolate the impact of model capacity, which is the focus of this work, as opposed to optimization.

Algorithm 2 Fitting the Constrained Model

Input adjacency matrix $A \in \{0, 1\}^{n \times n}$, regularization weight $\lambda \geq 0$, # of iterations I ,
of homo/heterophilous communities k_B/k_C

Output fitted factors $B \in \mathbb{R}_+^{n \times k_B}$, $C \in \mathbb{R}_+^{n \times k_C}$ such that $\sigma(BB^\top - CC^\top) \approx A$

- 1: Initialize B, C by setting entries to independent samples of $\text{Unif}(0, 1/\sqrt{k_B}), \text{Unif}(0, 1/\sqrt{k_C})$
 - 2: **for** $i \leftarrow 1$ to I **do**
 - 3: $\tilde{A} \leftarrow \sigma(BB^\top - CC^\top)$
 - 4: $R \leftarrow - \sum_{ij} \left(A_{ij} \log(\tilde{A}_{ij}) + (1 - A_{ij}) \log(1 - \tilde{A}_{ij}) \right)$
 - 5: $R \leftarrow R + \lambda (\|B\|_F^2 + \|C\|_F^2)$
 - 6: Calculate $\partial_{B,C} R$, the gradient of R w.r.t. B, C , via differentiation through Steps 2 to 4
 - 7: Update B, C to minimize R using $\partial_{B,C} R$, subject to $B, C \geq 0$
 - 8: **end for**
 - 9: **return** B, C
-

Implementation details Our implementation uses PyTorch (Paszke et al., 2019) for automatic differentiation and minimizes loss using the SciPy (Jones et al., 2001) implementation of the L-BFGS (Liu & Nocedal, 1989; Zhu et al., 1997) algorithm with default hyperparameters and up to a max of 200 iterations of optimization. We set regularization weight $\lambda = 10$ as in Yang & Leskovec (2013). We include code in the form of a Jupyter notebook (Pérez & Granger, 2007) demo.

6.2 Datasets

We use five fairly common mid-size datasets ranging from around 1K to 10K nodes. The selection of these five datasets is partly based on the presence of ground-truth multi-labels, which allows for evaluating the overlapping clustering methods. Statistics for these datasets, including the degeneracy

Table 1: Statistics of datasets used in our experiments. As in Sun et al. (2019), for YOUTUBE and AMAZON, we take only nodes that participate in at least one of the largest 5 ground-truth communities.

Name	Reference	Nodes	Edges	Labels	Max Degree	Degeneracy
BLOG	Tang & Liu (2009)	10,312	333,983	39	3,992	114
YOUTUBE	Yang & Leskovec (2015)	5,346	24,121	5	628	19
POS	Qiu et al. (2018)	4,777	92,406	40	3,644	49
PPI	Breitkreutz et al. (2007)	3,852	76,546	50	593	29
AMAZON	Yang & Leskovec (2015)	794	2,109	5	29	6

of each network, are given in Table 1. Note that degeneracy is an upper bound on arboricity. We note that the mid-sized networks used in our empirical work actually underemphasize the significance of our theoretical arboricity bound: see, e.g., the real-world networks in Pashanasangi & Seshadhri (2021), which have up to tens of millions of nodes but still have degeneracies at most in the hundreds.

BLOG is a social network of relationships between online bloggers; the node labels represent interests of the bloggers. Similarly, **YOUTUBE** is a social network of YouTube users, and the labels represent groups that the users joined.

POS is a word co-occurrence network: nodes represent words, and there are edges between words which are frequently adjacent in a section of the Wikipedia corpus. Each node label represents the part-of-speech of the word. **PPI** is a subgraph of the protein-protein interaction network for Homo Sapiens. Labels represent biological states. Finally, **AMAZON** is a co-purchasing network: nodes represent products, and there are edges between products which are frequently purchased together. Labels represent categories of products.

While social networks like the former two in this list are generally dominated by homophily (McPherson et al., 2001), the latter three should exhibit significant heterophily. For co-purchasing networks like AMAZON, depending on the product, two of the same kind of product are generally not co-purchased, e.g., Pepsi and Coke, as discussed in Peysakhovich & Bottou (2021). Though less intuitively accessible, there is also prior discussion of disassortativity in word adjacencies (Foster et al., 2010; Zweig, 2016), as well as in PPI networks (Newman, 2002; Hase et al., 2010).

6.3 Results

Expressiveness First, we investigate the expressiveness of our generative model, that is, the fidelity with which it can reproduce an input network. In Section 1, we used a simple synthetic network to show that our model is more expressive than others due to its ability to represent heterophilous structures in addition to homophilous ones. We now evaluate the expressiveness of our model on real-world networks. As with the synthetic graph, we fix the number of communities or singular vectors, fit the model, then evaluate the reconstruction error. In Figure 4, we compare the results of our model with those of SVD, BIGCLAM (which is discussed in detail in Section 4), and SYMNMF (Kuang et al., 2015). SYMNMF simply factors the adjacency matrix as $A \approx HH^T$, where $H \in \mathbb{R}_+^{n \times k}$; note that, like SVD, SYMNMF does not necessarily output a matrix whose entries are probabilities (i.e., bounded in $[0, 1]$), and hence it is not a natural graph generative model like ours and BIGCLAM.

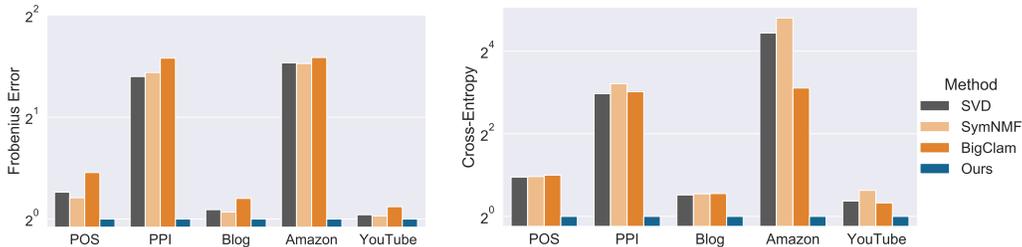


Figure 4: Reconstruction error on real-world networks, relative to our model’s error.

For each method, we fix the number of communities or singular vectors at the ground-truth number. For this experiment only, we are not concerned with learning the latent structure of the graph; the only goal is accurate representation of the network with limited parameters. So, for a fair comparison with SVD, we do not regularize the training of the other methods. Our method consistently has the lowest reconstruction error, both in terms of Frobenius error and entrywise cross-entropy (Equation 6). We particularly highlight the improvement over BIGCLAM; the salient difference between these models, both of which are factorization-based, include a nonlinear link function, and assign node communities using nonnegative factors, is the presence of the heterophilous $-CC^T$ term in our model. Thus, the improvement directly reflects the value of incorporating heterophily into the model. Interestingly, we find the most significant improvement exactly on the three datasets which have been noted to exhibit significant heterophily: POS, PPI, and AMAZON.

Similarity to ground-truth communities To assess the interpretability of clusters generated by our method, we evaluate the similarity of these clusters to ground-truth communities (i.e., class labels), and we compare other methods for overlapping clustering. We additionally compare to another recent but non-generative approach, the vGRAPH method of Sun et al. (2019), which is based on link clustering; the authors found their method to generally achieve state-of-the-art results in this task. For all methods, we set the number of communities to be detected as the number of ground-truth communities. We report F1-Score as computed in Yang & Leskovec (2013). See Figure 5 (left): the performance of our method is competitive with SYMNMF, BIGCLAM, and vGraph.

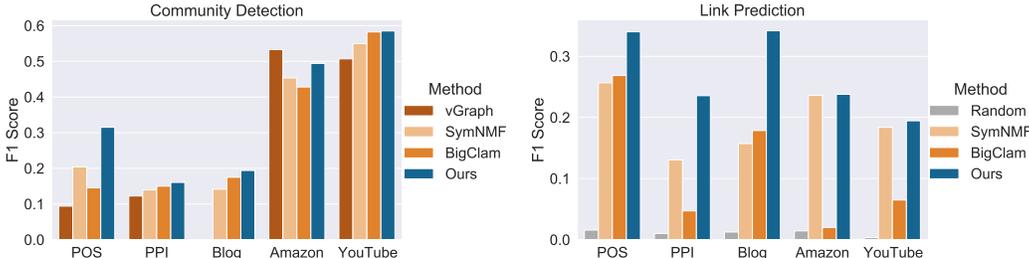


Figure 5: Left: Similarity of recovered communities to ground-truth labels of real-world datasets. (Note: vGRAPH is omitted on BLOG due to memory limitations.) Right: Accuracy of link prediction.

Interpretable link prediction We assess the predictive power of our generative model on the link prediction task. As discussed in Section 3, the link probabilities output by our model are interpretable in terms of a clustering of nodes that it generates; we compare results with our method to those with other models which permit similar interpretation, namely BIGCLAM and SYMNMF. We randomly select 10% of node pairs to hold out, fit the models on the remaining 90%, then use the trained models to predict links between node pairs in the held out 10%. As a baseline, we also show results for randomly predicting link or no link with equal probability. See Figure 5 (right). The performance of our method is competitive with or exceeds those of the other methods in terms of F1 Score.

7 Conclusion

We introduce a community-based graph generative model based on symmetric nonnegative matrix factorization which can represent both homophily and heterophily. We add to a prior guarantee of exact representation for bounded degree graphs with a broader guarantee for bounded *arboricity* graphs, and we show that both of these guarantees apply to our more interpretable graph model. We illustrate our model’s capabilities with experiments on a synthetic motivating example. Experiments on real-world networks show its effectiveness on several key tasks. Broadly, our results suggest that incorporating heterophily into methods for networks can improve both theoretical grounding and empirical performance, while maintaining interpretability. Future directions include deeper understanding of the expressiveness of low-rank logit models and convergence of training algorithms.

Acknowledgments and Disclosure of Funding

This project was partially supported by an Adobe Research grant, along with NSF Grants 2046235 and 1763618.

References

- Aggarwal, C. C. and Wang, H. A survey of clustering algorithms for graph data. In *Managing and Mining Graph Data*, pp. 275–301. Springer, 2010.
- Athreya, A., Fishkind, D. E., Tang, M., Priebe, C. E., Park, Y., Vogelstein, J. T., Levin, K., Lyzinski, V., and Qin, Y. Statistical inference on random dot product graphs: a survey. *The Journal of Machine Learning Research*, 18(1):8393–8484, 2017.
- Berry, M. W., Browne, M., Langville, A. N., Pauca, V. P., and Plemmons, R. J. Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics & Data Analysis*, 52(1):155–173, 2007.
- Bhattacharjee, R. and Dasgupta, S. What relations are reliably embeddable in euclidean space? In *Algorithmic Learning Theory*, pp. 174–195. PMLR, 2020.
- Bonato, A. A survey of models of the web graph. In *Workshop on Combinatorial and Algorithmic Aspects of Networking*, pp. 159–172. Springer, 2004.
- Boratko, M., Zhang, D., Monath, N., Vilnis, L., Clarkson, K. L., and McCallum, A. Capacity and bias of learned geometric embeddings for directed graphs. *Advances in Neural Information Processing Systems*, 34:16423–16436, 2021.
- Breitkreutz, B.-J., Stark, C., Reguly, T., Boucher, L., Breitkreutz, A., Livstone, M., Oughtred, R., Lackner, D. H., Bähler, J., Wood, V., et al. The biogrid interaction database: 2008 update. *Nucleic acids research*, 36(suppl_1):D637–D640, 2007.
- Chanpuriya, S., Musco, C., Sotiropoulos, K., and Tsourakakis, C. Node embeddings and exact low-rank representations of complex networks. *Advances in Neural Information Processing Systems*, 33, 2020.
- Ding, C., Li, T., and Jordan, M. I. Nonnegative matrix factorization for combinatorial optimization: Spectral clustering, graph matching, and clique finding. In *2008 Eighth IEEE International Conference on Data Mining*, pp. 183–192. IEEE, 2008.
- Donoho, D. L. and Stodden, V. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in Neural Information Processing Systems 16*, pp. 1141–1148. MIT Press, 2003.
- Duvenaud, D., Maclaurin, D., Aguilera-Iparraguirre, J., Gómez-Bombarelli, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems 28*, pp. 2224–2232, 2015.
- Foster, J. G., Foster, D. V., Grassberger, P., and Paczuski, M. Edge direction and the structure of networks. *Proceedings of the National Academy of Sciences*, 107(24):10815–10820, 2010.
- Gillis, N. *Nonnegative Matrix Factorization*. SIAM, 2020.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Hase, T., Niimura, Y., and Tanaka, H. Difference in gene duplicability may explain the difference in overall structure of protein-protein interaction networks among eukaryotes. *BMC Evolutionary Biology*, 10(1):1–15, 2010.
- Javed, M. A., Younis, M. S., Latif, S., Qadir, J., and Baig, A. Community detection in networks: A multidisciplinary review. *Journal of Network and Computer Applications*, 108:87–111, 2018.

- Johnson, S., Torres, J. J., Marro, J., and Munoz, M. A. Entropic origin of disassortativity in complex networks. *Physical Review Letters*, 104(10):108702, 2010.
- Jones, E., Oliphant, T., Peterson, P., et al. SciPy: Open source scientific tools for Python, 2001. URL <http://www.scipy.org/>.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*, 2017.
- Kuang, D., Ding, C., and Park, H. Symmetric nonnegative matrix factorization for graph clustering. In *Proceedings of the 2012 SIAM International Conference on Data Mining*, pp. 106–117. SIAM, 2012.
- Kuang, D., Yun, S., and Park, H. Symnmf: nonnegative low-rank approximation of a similarity matrix for graph clustering. *Journal of Global Optimization*, 62(3):545–574, 2015.
- Lee, D. D. and Seung, H. S. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.
- Liu, D. C. and Nocedal, J. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1-3):503–528, 1989.
- Mason, O. and Verwoerd, M. Graph theory and networks in biology. *IET Systems Biology*, 1(2): 89–119, 2007.
- McPherson, M., Smith-Lovin, L., and Cook, J. M. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.
- Miller, K. T., Griffiths, T. L., and Jordan, M. I. Nonparametric latent feature models for link prediction. In *Advances in Neural Information Processing Systems 22*, pp. 1276–1284. Curran Associates, Inc., 2009.
- Nascimento, M. C. and De Carvalho, A. C. Spectral methods for graph clustering—a survey. *European Journal of Operational Research*, 211(2):221–231, 2011.
- Newman, M. E. Assortative mixing in networks. *Physical Review Letters*, 89(20):208701, 2002.
- Noldus, R. and Van Mieghem, P. Assortativity in complex networks. *Journal of Complex Networks*, 3(4):507–542, 2015.
- Nourbakhsh, F., Buló, S. R., and Pelillo, M. A matrix factorization approach to graph compression. In *2014 22nd International Conference on Pattern Recognition*, pp. 76–81. IEEE, 2014.
- NT, H. and Maehara, T. Revisiting graph neural networks: All we have is low-pass filters. *arXiv preprint arXiv:1905.09550*, 2019.
- Pashanasangi, N. and Seshadhri, C. Faster and generalized temporal triangle counting, via degeneracy ordering. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 1319–1328, 2021.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019.
- Pei, H., Wei, B., Chang, K. C., Lei, Y., and Yang, B. Geom-gcn: Geometric graph convolutional networks. In *8th International Conference on Learning Representations, ICLR 2020*, 2020.
- Pérez, F. and Granger, B. E. IPython: a system for interactive scientific computing. *Computing in Science and Engineering*, 9(3):21–29, May 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.53. URL <https://ipython.org>.
- Peysakhovich, A. and Bottou, L. An attract-repel decomposition of undirected networks. *arXiv preprint arXiv:2106.09671*, 2021.

- Qiu, J., Dong, Y., Ma, H., Li, J., Wang, K., and Tang, J. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pp. 459–467. ACM, 2018.
- Rendsburg, L., Heidrich, H., and Von Luxburg, U. Netgan without gan: From random walks to low-rank approximations. In *International Conference on Machine Learning*, pp. 8073–8082. PMLR, 2020.
- Rubin-Delanchy, P., Cape, J., Tang, M., and Priebe, C. E. A statistical interpretation of spectral embedding: the generalised random dot product graph. *arXiv preprint arXiv:1709.05506*, 2017.
- Sala, F., De Sa, C., Gu, A., and Ré, C. Representation tradeoffs for hyperbolic embeddings. In *International Conference on Machine Learning*, pp. 4460–4469. PMLR, 2018.
- Schaeffer, S. E. Graph clustering. *Computer Science Review*, 1(1):27–64, 2007.
- Scheinerman, E. R. and Tucker, K. Modeling graphs using dot product representations. *Computational statistics*, 25(1):1–16, 2010.
- Scott, J. Social network analysis. *Sociology*, 22(1):109–127, 1988.
- Seshadhri, C., Sharma, A., Stolman, A., and Goel, A. The impossibility of low-rank representations for triangle-rich complex networks. *Proceedings of the National Academy of Sciences*, 117(11): 5631–5637, 2020.
- Shi, J. and Malik, J. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- Sun, F., Qu, M., Hoffmann, J., Huang, C., and Tang, J. vgraph: A generative model for joint community detection and node representation learning. In *Advances in Neural Information Processing Systems 32*, pp. 512–522, 2019.
- Tang, L. and Liu, H. Relational learning via latent social dimensions. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 817–826. ACM, 2009.
- Wang, Y.-X. and Zhang, Y.-J. Nonnegative matrix factorization: A comprehensive review. *IEEE Transactions on Knowledge and Data Engineering*, 25(6):1336–1353, 2012.
- Xie, J., Kelley, S., and Szymanski, B. K. Overlapping community detection in networks: The state-of-the-art and comparative study. *ACM Computing Surveys*, 45(4):43:1–43:35, 2013.
- Yan, Y., Hashemi, M., Swersky, K., Yang, Y., and Koutra, D. Two sides of the same coin: Heterophily and oversmoothing in graph convolutional neural networks. *arXiv preprint arXiv:2102.06462*, 2021.
- Yang, J. and Leskovec, J. Overlapping community detection at scale: a nonnegative matrix factorization approach. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, pp. 587–596, 2013.
- Yang, J. and Leskovec, J. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems*, 42(1):181–213, 2015.
- Yang, Z., Hao, T., Dikmen, O., Chen, X., and Oja, E. Clustering by nonnegative matrix factorization using graph random walk. In *Advances in Neural Information Processing Systems*, pp. 1079–1087, 2012.
- Young, S. J. and Scheinerman, E. R. Random dot product graph models for social networks. In *International Workshop on Algorithms and Models for the Web-Graph*, pp. 138–149. Springer, 2007.
- Yu, K., Yu, S., and Tresp, V. Soft clustering on graphs. In *Advances in Neural Information Processing Systems*, pp. 1553–1560, 2005.

- Zheng, X., Liu, Y., Pan, S., Zhang, M., Jin, D., and Yu, P. S. Graph neural networks for graphs with heterophily: A survey. *arXiv:2202.07082*, 2022.
- Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)*, 23(4):550–560, 1997.
- Zhu, J., Yan, Y., Zhao, L., Heimann, M., Akoglu, L., and Koutra, D. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in Neural Information Processing Systems*, 33, 2020.
- Zweig, K. A. Are word-adjacency networks networks? In *Towards a theoretical framework for analyzing complex linguistic networks*, pp. 153–163. Springer, 2016.

A Appendix

A.1 COT Graph Exact Representation

As a theoretical demonstration of the capability of our model to learn latent structure, we additionally show that our model can exactly represent a natural family of graphs, which exhibits both homophily and heterophily, with small k and interpretably. The family of graphs is specified below in Definition 1; roughly speaking, nodes in such graphs share an edge iff they coparticipate in some number of homophilous communities and don't coparticipate in a number of heterophilous communities. For example, the motivating graph described in Section 2 would be an instance of such a graph if an edge occurs between two users iff the two users are from the same city and have different genders.

Definition 1 (Community Overlap Threshold (COT) Graph). *An unweighted, undirected graph whose edges are determined by an overlapping clustering and a “thresholding” integer $t \in \mathbb{Z}$ as follows: for each vertex i , there are two latent binary vectors $\mathbf{b}_i \in \{0, 1\}^{k_b}$ and $\mathbf{c}_i \in \{0, 1\}^{k_c}$, and there is an edge between vertices i and j iff $\mathbf{b}_i \cdot \mathbf{b}_j - \mathbf{c}_i \cdot \mathbf{c}_j \geq t$.*

Theorem A.1 (Compact Representation of COT Graphs). *Suppose \mathbf{A} is the adjacency matrix of a COT graph on n nodes with latent vectors $\mathbf{b}_i \in \{0, 1\}^{k_b}$ and $\mathbf{c}_i \in \{0, 1\}^{k_c}$ for $i \in \{1, 2, \dots, n\}$. Let $k = k_b + k_c$. Then, for any $\epsilon > 0$, there exist $\mathbf{V} \in [0, 1]^{n \times (k+1)}$ and diagonal $\mathbf{W} \in \mathbb{R}^{(k+1) \times (k+1)}$ such that $\|\sigma(\mathbf{V}\mathbf{W}\mathbf{V}^\top) - \mathbf{A}\|_F < \epsilon$.*

Proof. Let t be the thresholding integer of the graph, and let the rows of $\mathbf{B} \in \{0, 1\}^{n \times k_b}$ and $\mathbf{C} \in \{0, 1\}^{n \times k_c}$ contain the vectors \mathbf{b} and \mathbf{c} of all nodes. Via Equation 2, we can find $\mathbf{V}^* \in [0, 1]^{n \times k}$ and diagonal $\mathbf{W}^* \in \mathbb{R}^{k \times k}$ such that $\mathbf{V}^* \mathbf{W}^* \mathbf{V}^{*\top} = \mathbf{B}\mathbf{B}^\top - \mathbf{C}\mathbf{C}^\top$. Now let

$$\mathbf{V} = (\mathbf{V}^* \quad \mathbf{1}) \quad \mathbf{W} = \begin{pmatrix} \mathbf{W}^* & 0 \\ 0 & \frac{1}{2} - t \end{pmatrix}.$$

Then $(\mathbf{V}\mathbf{W}\mathbf{V}^\top)_{ij} = \mathbf{b}_i \cdot \mathbf{b}_j - \mathbf{c}_i \cdot \mathbf{c}_j + \frac{1}{2} - t$. Hence $(\mathbf{V}\mathbf{W}\mathbf{V}^\top)_{ij} > 0$ iff $\mathbf{b}_i \cdot \mathbf{b}_j - \mathbf{c}_i \cdot \mathbf{c}_j > t - \frac{1}{2}$, which is true iff $\mathbf{A}_{ij} = 1$ by the assumption on the graph. Similarly, $(\mathbf{V}\mathbf{W}\mathbf{V}^\top)_{ij} < 0$ iff $\mathbf{A}_{ij} = 0$. It follows that

$$\lim_{s \rightarrow \infty} \sigma(\mathbf{V}(s\mathbf{W})\mathbf{V}^\top) = \lim_{s \rightarrow \infty} \sigma(s\mathbf{V}\mathbf{W}\mathbf{V}^\top) = \mathbf{A}. \quad \blacksquare$$

A.2 Application: Tabular Data Completion

We apply our link prediction algorithm to the task of data completion for categorical tabular data. We first transform such data to a graph as follows: We create a node for each row, as well as a node for each unique category for each of the columns. For each entry of the table, a link occurs between the node for the row and the node for the value in the entry, which is one of the possible categories of the column (e.g., if Bob is American, then a link occurs between the row node for ‘Bob’ and the category node for ‘American’). No links occur between two nodes for rows, or between two nodes for category values. For the negative samples, we select only edges between nodes for rows and nodes for category values; the negative samples would otherwise be dominated by pairs of nodes for rows.

We use two proprietary company datasets datasets which we call COMPANY A and COMPANY B. We use the same experimental setup as described in Section 6.3. Unlike in Section 6.3, we have a restriction on the possible links: each row node must be linked to exactly one of the value nodes for each of the categories. Hence, for each combination of row node and column, we predict a link only with the value node with the highest predicted link probability. We set the number of communities for all methods to 20 and 50 for COMPANY A and COMPANY B, respectively. We employ the stochastic version of our algorithm, based on projected SGD, for which code is also provided. We generally use a batch size of 100; we find that the optimization of BIGCLAM often diverges on COMPANY B with this batch size, so we instead use batches of size 1000 for its optimization. The results are provided in Figure 6. Notably, our proposed approach outperforms comparable (community-based factorization) methods in terms of accuracy. As a baseline, we also show the accuracy when completing the tables by simply selecting the most common categorical value for each column (“PLURALITY”). All methods outperform this baseline on both datasets.

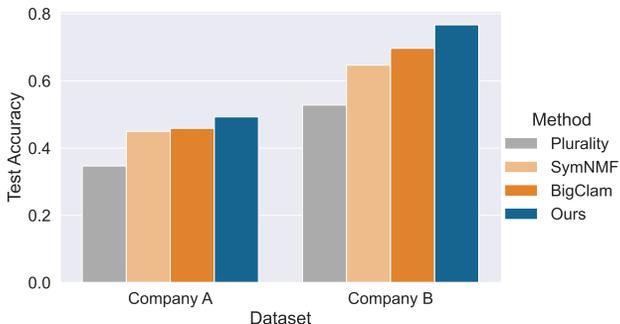


Figure 6: Test accuracy of tabular data completion on two proprietary company datasets (COMPANY A and COMPANY B).