
Hyperbolic Graph Embedding with Enhanced Semi-Implicit Variational Inference

Ali Lotfi Rezaabad

Rahi Kalantari

Sriram Vishwanath

Mingyuan Zhou

Jonathan I Tamir

The University of Texas at Austin

corresponding authors: {alotfi, jtamir}@utexas.edu

Abstract

Efficient modeling of relational data arising in physical, social, and information sciences is challenging due to complicated dependencies within the data. In this work we build off of semi-implicit graph variational auto-encoders to capture higher order statistics in a low-dimensional graph latent representation. We incorporate hyperbolic geometry in the latent space through a Poincaré embedding to efficiently represent graphs exhibiting hierarchical structure. To address the naive posterior latent distribution assumptions in classical variational inference, we use semi-implicit hierarchical variational Bayes to implicitly capture posteriors of given graph data, which may exhibit heavy tails, multiple modes, skewness, and highly correlated latent structures. We show that the existing semi-implicit variational inference objective provably reduces information in the observed graph. Based on this observation, we estimate and add an additional mutual information term to the semi-implicit variational inference learning objective to capture rich correlations arising between the input and latent spaces. We show that the inclusion of this regularization term in conjunction with the Poincaré embedding boosts the quality of learned high-level representations and enables more flexible and faithful graphical modeling. We experimentally demonstrate that our approach outperforms exist-

ing graph variational auto-encoders both in Euclidean and in hyperbolic spaces for edge link prediction and node classification.

1 Introduction and Related Work

Rich relational data represented as undirected graphs arise in a variety of applications spanning physical, social, and information sciences. These include knowledge graphs, gene expression networks, social graphs in social media graphs, citation graphs, recommendation systems, transportation networks, and cellular networks, to name a few. Downstream analyses of these graph data including edge link prediction, clustering, classification, filtering, and denoising are of particular interest and may reveal insight into the underlying processes governing the data. As the number of nodes and edges grows, it is both expensive and challenging to elucidate this information directly from the observations due to computation and complicated structure.

Low-dimensional graph representations aim to address these difficulties by learning a small number of summary statistics of the graph’s features, e.g., through unsupervised learning, and using this representation learning for downstream tasks (Bengio et al. [2013]). Initial work in this space focused on deterministic graph embedding in the latent space; see Cai et al. [2018] for a survey on these methods. More recently, representations have been extended to model uncertainty in the graph by assigning each node to a random variable in latent space (Bojchevski and Günnemann [2018], He et al. [2015], Dos Santos et al. [2016], Kipf and Welling [2016], Davidson et al. [2018], Hasanzadeh et al. [2019]). A central idea in many of these works is the use of a tractable approach to estimating the latent posterior distribution through a variational inference framework introduced by Kingma and Welling

[2013]. To ease restrictions on the expressibility of the posterior distribution, semi-implicit variational inference (SIVI, Yin and Zhou [2018]) and its extension to graph representation (Hasanzadeh et al. [2019]) have been developed to model highly correlated latent structures, thus lifting the need for an explicit probability density function required by classical variational inference techniques.

There remain two additional challenges in succinctly and accurately representing relational data, which we address in this work. First, while SIVI can characterize data exhibiting heavy tails, multiple modes, skewness, and highly correlated latent structures, we show that the existing SIVI objective provably reduces information in the observed graph. This is because the log-likelihood maximization has the undesirable effect of reducing the mutual information between the input observations and latent representation. Thus, while the reconstruction error may be low, the quality of the representation may also suffer. Based on this observation, we estimate and add an additional mutual information term to the SIVI objective to capture the correlation between the input and latent spaces. By optimizing over this term, we are able to capture a more faithful latent representation of the graph.

A second limitation is the use of a Euclidean embedding, which is often not suitable for graphical data exhibiting tree-like hierarchical structure (Adcock et al. [2013], Ravasz and Barabási [2003]), spherical structure (Fisher et al. [1987]), or other complicated dependencies between nodes (Steyvers and Tenenbaum [2005]). Recognizing this limitation, the authors in Davidson et al. [2018] replace the Gaussian prior with a Mises-Fisher distribution, which imposes a Gaussian distribution on the hypersphere, and thus allows for a hyperspherical embedding. However, as the approach still relies on graph variational auto-encoders (GVAE) consisting of an analytical prior, neither higher-order statistics (e.g. skewness, multi-modality) nor correlations between input and latent space are captured. A number of works have recently explored the utility of hyperbolic latent embeddings for graphs. Nickel and Kiela [2017] show that a deterministic Poincaré embedding improved link prediction compared to a Euclidean embedding. Similarly, Mathieu et al. [2019] extended VAEs to Poincaré embeddings and showed promising preliminary results on link prediction compared to the Euclidean counterpart. The work by Ovinnikov [2019] formulated an alternative objective function based on Wasserstein loss in hyperbolic space. Both Nickel and Kiela [2017] and Ovinnikov [2019] considered a broad range of hierarchical data, but did not fully explore downstream applications on graphs. In this work we carry this direction forward and endow our

latent space with a hyperbolic Poincaré embedding, while also including our enhanced-SIVI framework.

Combining these ideas, we arrive at our *enhanced semi-implicit hyperbolic graph embedding* (ESI-HGE). The main contributions of our work are the following:

1. We show that the SIVI objective introduced in Yin and Zhou [2018] provably reduces the mutual information between the input and latent representation.
2. We propose to correct the objective function with an additional mutual information term and provide a computationally tractable approach for estimating it based on Monte Carlo sampling.
3. We demonstrate empirically that the inclusion of the enhanced SIVI objective in conjunction with a hyperbolic embedding in latent space is able to capture higher-level feature representation and improve quantitative metrics on downstream tasks.

To validate our approach, we compare ESI-HGE to other leading methods for edge link prediction and unsupervised node classification on three popular citation network datasets (Sen et al. [2008]). We show that in both tasks, our proposed method outperforms other state-of-the-art methods. We empirically demonstrate that our approach captures more mutual information between input and latent spaces. We also explore the representation quality on synthetically generated hierarchical image-graph data and show that the proposed embedding naturally represents the graph structure. Reference code to reproduce our results is available at https://github.com/utcsilab/ESI_HGE.

2 Background

2.1 Graph VAE

Introduced by Kipf and Welling [2016], the graph variational autoencoder (GVAE) is a framework on top of VAEs (Kingma and Welling [2013]) that is capable of embedding attributes of nodes in an undirected graph on to a lower dimensional space. Formally, given an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with total number of nodes $N = |\mathcal{V}|$, we define \mathbf{A} as its associated adjacency matrix and $\mathbf{X} \in \mathbb{R}^{N \times M}$ as its feature matrix (M features per node). We introduce a stochastic latent variable (of lower dimension) F associated with each node, summarized by the matrix $\mathbf{Z} \in \mathbb{R}^{N \times F}$. Since optimizing the true posterior, $p_\theta(\mathbf{Z}|\mathbf{X}, \mathbf{A})$, is fairly challenging, GVAEs instead employ a variational posterior $q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{A})$, which is parameterized by a graph convolutional network (GCN) with parameters ϕ :

$$\begin{aligned}
 q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{A}) &= \prod_{i=1}^N q_\phi(\mathbf{z}_i|\mathbf{X}, \mathbf{A}), \\
 q_\phi(\mathbf{z}_i|\mathbf{X}, \mathbf{A}) &= \mathcal{N}(\mathbf{z}_i|\boldsymbol{\mu}_i, \text{diag}(\boldsymbol{\sigma}_i^2)), \\
 \boldsymbol{\mu}_i &= \text{GCN}_\mu(\mathbf{X}, \mathbf{A}), \quad \log \boldsymbol{\sigma}_i = \text{GCN}_\sigma(\mathbf{X}, \mathbf{A}).
 \end{aligned} \tag{1}$$

Here, $\boldsymbol{\mu}_i$ and $\boldsymbol{\sigma}_i$ denote the mean and standard deviation associated with the i -th node, respectively. The generative model is given by an inner product of the corresponding higher level representations as follows:

$$\begin{aligned}
 p(\mathbf{A}|\mathbf{Z}) &= \prod_{i=1}^N \prod_{j=1}^N p(A_{i,j}|\mathbf{z}_i, \mathbf{z}_j), \\
 p(A_{i,j} = 1|\mathbf{z}_i, \mathbf{z}_j) &= \text{sigmoid}(\mathbf{z}_i^\top \mathbf{z}_j).
 \end{aligned} \tag{2}$$

The optimal values of the GCN parameters ϕ can be obtained by maximizing the evidence lower bound (ELBO, [Wainwright and Jordan \[2008\]](#)):

$$\max_{\phi} \mathbb{E}_{q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{A})} [\log p(\mathbf{A}|\mathbf{Z})] - \text{KL}(q_\phi(\mathbf{Z}|\mathbf{X}, \mathbf{A})||p(\mathbf{Z})), \tag{3}$$

where KL denotes the Kullback-Leibler (KL) divergence. Furthermore, $p(\mathbf{Z})$ is defined to be a Gaussian prior, i.e., $p(\mathbf{Z}) = \prod_i \mathcal{N}(\mathbf{z}_i|0, \mathbf{I})$.

2.2 Semi-Implicit Variational Inference

As demonstrated by [Yin and Zhou \[2018\]](#), semi-implicit variational inference (SIVI) provides a computationally simple yet powerful approach to capturing multi-modality and skewness of posterior distributions, while vanilla variational inference is often insufficient. The authors empirically show that SIVI can closely match the accuracy of MCMC, however, with far lower computational cost. The latent posterior distribution expresses how a node in the graph is represented in the latent space given its features and adjacency information. In many works (e.g., [Bojchevski and Günnemann \[2018\]](#), [Davidson et al. \[2018\]](#), [Dos Santos et al. \[2016\]](#), [He et al. \[2015\]](#), [Kipf and Welling \[2016\]](#), [Mathieu et al. \[2019\]](#)), the latent representation assumes a uni-modal distribution with uncorrelated latent variables. To enrich this simplistic representation, SIVI develops a posterior representation which captures multi-modality, correlation and skewness in the latent representation of each node. As this representation is more faithful to the data, it improves downstream tasks including edge link prediction and classification using the low dimensional graph representation ([Hasanzadeh et al. \[2019\]](#)).

The semi-implicit VAE enriches the posterior distribution by using a hierarchical generative model where the distribution $h_\phi(\mathbf{Z}|\mathbf{X})$ is represented by a mixture model consisting of an infinite number of mixtures with weights given by an implicit distribution $q_\phi(\boldsymbol{\Psi}|\mathbf{X})$. As a result, the latent variable distribution ($h_\phi(\mathbf{Z})$) can be written as $h_\phi(\mathbf{Z}|\mathbf{X}) = \int q(\mathbf{Z}|\boldsymbol{\Psi}, \mathbf{X})q_\theta(\boldsymbol{\Psi}|\mathbf{X})d\boldsymbol{\Psi}$. In the semi-implicit VAE, $\boldsymbol{\Psi}$ is drawn from an implicit distribution ($q_\phi(\boldsymbol{\Psi})$) and used as parameter for the exponential family distribution $q_\theta(\mathbf{Z}|\boldsymbol{\Psi}, \mathbf{X})$. To sample from $q_\theta(\boldsymbol{\Psi}|\mathbf{X})$, one generates samples from a simple distribution $q(\boldsymbol{\epsilon})$ and maps \mathbf{X} and $\boldsymbol{\epsilon}$ to $\boldsymbol{\Psi}$ using $\boldsymbol{\Psi} = g_\phi(\mathbf{X}, \boldsymbol{\epsilon})$. Since $h_\phi(\mathbf{Z}|\mathbf{X})$ is an implicit distribution, it is not possible to explicitly evaluate the ELBO for $h_\phi(\mathbf{Z}|\mathbf{X})$. Instead, a lower bound ($\underline{\mathcal{L}}_K$) is evaluated, which has been shown to asymptotically increase to the real ELBO, $\underline{\mathcal{L}}$:

$$\underline{\mathcal{L}} = \mathbb{E}_{q(\mathbf{X})} \mathbb{E}_{q_\phi(\boldsymbol{\Psi}|\mathbf{X})} \mathbb{E}_{q_\theta(\mathbf{Z}|\boldsymbol{\Psi}, \mathbf{X})} \log \left(\frac{p(\mathbf{X}, \mathbf{Z})}{q_\phi(\mathbf{Z}|\boldsymbol{\Psi}, \mathbf{X})} \right). \tag{4}$$

An additional regularizer is used to avoid degeneracy:

$$B_K = \mathbb{E}_{q(\mathbf{X})} \mathbb{E}_{\boldsymbol{\Psi}, \boldsymbol{\Psi}^{(1)}, \dots, \boldsymbol{\Psi}^{(K)} \sim q_\phi(\boldsymbol{\Psi}|\mathbf{X})} \left[\text{KL}(q_\phi(\mathbf{Z}|\boldsymbol{\Psi}, \mathbf{X})||h_K(\mathbf{Z}|\mathbf{X})) \right], \tag{5}$$

where

$$h_K(\mathbf{Z}|\mathbf{X}) = \frac{1}{K+1} \left[q_\phi(\mathbf{Z}|\boldsymbol{\Psi}, \mathbf{X}) + \sum_{k=1}^K q_\phi(\mathbf{Z}|\boldsymbol{\Psi}^{(k)}, \mathbf{X}) \right].$$

While the semi-implicit VAE (using the asymptotic lower bound) can capture the correlation and multi-modality among latent dimensions of \mathbf{Z} , we illustrate in Section 3 that it may compromise the correlation with the input and the corresponding latent representation. We empirically show that we can improve the results by incorporating the mutual information between the input and latent representations during training. This helps the objective function to both preserve the correlation among latent dimensions of the \mathbf{Z} , as well as to encourage a high level of correlation between the input and its latent representation.

2.3 Hyperbolic (Poincaré) Space

Within the Riemannian geometry framework, hyperbolic spaces are manifolds with constant negative curvature and are of particular interest for embedding hierarchical structures. There are multiple models for an n -dimensional hyperbolic space; among these are the hyperboloid H_{nC} (also known as the Lorentz model) ([Nickel and Kiela \[2018\]](#)), and the Poincaré ball P_{nC} ([Nickel and Kiela \[2017\]](#)). Specifically, in hyperbolic

space, distances grow exponentially as a point moves away from the origin, and to describe the shortest path between two distant points, we should pass through a common parent (i.e., the origin) which gives rise to a hierarchical or tree-like structure. In this section, we briefly review the Poincaré model for hyperbolic geometry.

A manifold \mathcal{M} is a set of points \mathbf{z} , which are locally similar to a linear space. For Each point \mathbf{z} on \mathcal{M} we can define its n -dimensional *tangent space* $T_{\mathbf{z}}\mathcal{M}$, which is a first order approximation of the manifold around the point \mathbf{z} . Based on this, we can use an *exponential mapping* to transform a vector in the tangent space at \mathbf{v} to the manifold based on the following formulation:

$$\exp_{\mathbf{z}}^c(\mathbf{v}) = \mathbf{z} \oplus_c \left(\tanh \left(\sqrt{c} \frac{\lambda_{\mathbf{z}}^c \|\mathbf{v}\|}{2} \right) \frac{\mathbf{v}}{\sqrt{c} \|\mathbf{v}\|} \right), \quad (6)$$

where c is the constant negative curvature, and \oplus_c is the *Möbius addition* of two vectors on the manifold and defined as

$$\mathbf{z} \oplus_c \mathbf{y} = \frac{(1 + 2c\langle \mathbf{z}, \mathbf{y} \rangle + c\|\mathbf{y}\|^2) \mathbf{z} + (1 - c\|\mathbf{z}\|^2) \mathbf{y}}{1 + 2c\langle \mathbf{z}, \mathbf{y} \rangle + c^2\|\mathbf{z}\|^2\|\mathbf{y}\|^2}.$$

Conversely, one can map a vector on the manifold to the tangent space using the *logarithmic mapping* defined as

$$\log_{\mathbf{z}}^c(\mathbf{y}) = \frac{2}{\sqrt{c}\lambda_{\mathbf{z}}^c} \tanh^{-1} \left(\sqrt{c} \|\mathbf{z} \oplus_c \mathbf{y}\| \right) \frac{-\mathbf{z} \oplus_c \mathbf{y}}{\|\mathbf{z} \oplus_c \mathbf{y}\|}, \quad (7)$$

where $\lambda_{\mathbf{z}}^c = \frac{2}{1 - c\|\mathbf{z}\|^2}$.

Based on the aforementioned mapping, we can map a Gaussian distribution defined in the vector space (tangent space) to the Poincaré manifold using the exponential mapping in (6). After passing the Gaussian distribution $\eta \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma})$ through the exponential mapping we get the *wrapped normal*, defined as

$$\mathcal{N}_p^W(\mathbf{z} | \boldsymbol{\mu}, \boldsymbol{\sigma}) = \mathcal{N}(\lambda_{\boldsymbol{\mu}}^c \log_{\boldsymbol{\mu}}(\mathbf{z}) | \mathbf{0}, \boldsymbol{\sigma}) \left(\frac{\sqrt{c} d_p^c(\boldsymbol{\mu}, \mathbf{z})}{\sinh(\sqrt{c} d_p^c(\boldsymbol{\mu}, \mathbf{z}))} \right)^{d-1}, \quad (8)$$

where $d_p^c(\boldsymbol{\mu}, \mathbf{z})$ is the Poincaré distance between $\boldsymbol{\mu}$ and \mathbf{z} , defined as

$$d_p^c(\mathbf{z}, \boldsymbol{\mu}) = \frac{1}{\sqrt{c}} \cosh^{-1} \left(1 + 2c \frac{\|\mathbf{z} - \boldsymbol{\mu}\|^2}{(1 - c\|\mathbf{z}\|^2)(1 - c\|\boldsymbol{\mu}\|^2)} \right). \quad (9)$$

3 Enhanced Semi-Implicit Hyperbolic Graph Embedding

In this section, we propose a generative model to create a richer latent representation of the graph nodes in hyperbolic space. We model the posterior distribution using the SIVI framework; in addition, we show why semi-implicit VAE's representation can compromise the correlation between input \mathbf{X} and latent code \mathbf{Z} and propose a solution to address this problem. Since we use a semi-implicit wrapped normal representation for our latent model, we start with the semi-implicit VAE objective function given by $\underline{\mathcal{L}}_K = \underline{\mathcal{L}} + B_K$. Each of these components can be expressed by the following Lemma.

Lemma 1: *The semi-implicit VAE objective components can be shown to have the following representation:*

$$\begin{aligned} \underline{\mathcal{L}}_K &= \underline{\mathcal{L}} + B_K, \text{ where:} \\ \underline{\mathcal{L}} &= \mathbb{E}_{q(\mathbf{X})q_{\phi}(\boldsymbol{\Psi}|\mathbf{X})q(\mathbf{Z}|\boldsymbol{\Psi}, \mathbf{X})} [\log p(\mathbf{X}|\mathbf{Z})] \\ &\quad - I(\mathbf{Z}; \mathbf{X} | \boldsymbol{\Psi}) - I(\mathbf{Z}; \boldsymbol{\Psi}) - \text{KL}(q(\mathbf{Z})||p(\mathbf{Z})), \\ &\text{and } B_K \leq I(\mathbf{Z}; \boldsymbol{\Psi}|\mathbf{X}). \end{aligned} \quad (10)$$

Proof is provided in the Supplementary Material.

In the proof, it has been shown explicitly how each component of the semi-implicit VAE asymptotic lower bound ($\underline{\mathcal{L}}_K = \underline{\mathcal{L}} + B_K$) is related to the expressions in (10). Next we discuss how the components of the objective function play a role to achieve our goal of a richer latent representation. As shown in Lemma 1, it can be inferred that the maximization of $\underline{\mathcal{L}}_K$ will lead to the maximization of the log-likelihood of reconstruction which is a desirable outcome. Furthermore, the optimization in (10) involves the minimization of the mutual information between the observed input and latent representations. As a result, the objective function may reduce the input-latent-representation correlation, specifically when we use very expressive modeling tools such as deep neural networks [Rezaabad and Vishwanath, 2019]. In addition, it can be observed that maximization of objective will results in minimization of $I(\mathbf{Z}; \boldsymbol{\Psi})$. However, we note that this term has been compensated by adding the $\lim_{K \rightarrow \infty} B_K$. Indeed, one can show that

$$\begin{aligned} \lim_{K \rightarrow \infty} B_K &\leq I(\mathbf{Z}; \boldsymbol{\Psi}) - I(\mathbf{X}; \mathbf{Z}; \boldsymbol{\Psi}), \text{ where} \\ I(\mathbf{X}; \mathbf{Z}; \boldsymbol{\Psi}) &= \int q_{\phi}(\mathbf{X}, \mathbf{Z}, \boldsymbol{\Psi}) \left[\log \frac{q_{\phi}(\mathbf{X}, \mathbf{Z})q_{\phi}(\mathbf{Z}, \boldsymbol{\Psi})q_{\phi}(\mathbf{X}, \boldsymbol{\Psi})}{q_{\phi}(\mathbf{X}, \mathbf{Z}, \boldsymbol{\Psi})q(\mathbf{X})q_{\phi}(\mathbf{Z})q_{\phi}(\boldsymbol{\Psi})} \right] d\mathbf{X}d\mathbf{Z}d\boldsymbol{\Psi}. \end{aligned} \quad (11)$$

Adding $\lim_{K \rightarrow \infty} B_K$ in the semi-implicit objective function (10), will not address all the concerns which

we have discussed earlier. Also, there are certain theoretical and practical drawbacks of adding that term to objective function which we will discuss then in following. 1- Based on (11), while the additional term may compensate for the minimization of mutual information between \mathbf{Z} and Ψ ($I(\mathbf{Z}; \Psi)$), it will not address the concern with regard to minimization of $I(\mathbf{Z}; \mathbf{X}|\Psi)$; 2- in addition, it will minimize the mutual information of the tuple $(\mathbf{X}, \mathbf{Z}, \Psi)$; and 3- it is not possible to have the K grow without bound in practice. Specifically, for graph convolution neural networks in which optimization of parameters is conducted on the graph, the semi-implicit VAE lower bound optimization gets prohibitively expensive as K grows. For small K s where the computation is practical on large graph data, B_K will not be able to fully compensate for the $I(\mathbf{Z}; \Psi)$ minimization.

Finally, $\text{KL}(q(\mathbf{Z})||p(\mathbf{Z}))$ is independent of \mathbf{X} and will not help in modeling the correlation among input and latent representations. In the next section, we discuss how we address the minimization of the mutual information between the latent representations \mathbf{Z} and observed features \mathbf{X} .

3.1 Semi-Implicit Objective Enhancement

As we discussed, the correlation between (input) and the latent variables may be compromised due to the nature of the semi-implicit VAE objective function. Here we propose a practical and effective enhancement of the semi-implicit VAE by adding a mutual information term to the objective function. It is noteworthy to mention, the added mutual information term should not only address the minimization of mutual information in 10, but also, it should be evaluated effectively and efficiently. As mentioned previously, B_K will compensate for the minimization of mutual information between Ψ and \mathbf{Z} when K grows without bound, but the compensation is limited for finite and small values of K . Furthermore, as shown in (10), the mutual information between \mathbf{X} and \mathbf{Z} is minimized due to the maximization of $\underline{\mathcal{L}}_K$. Therefore, to compensate for minimization of $I(\mathbf{Z}; \Psi)$ and $I(\mathbf{X}; \mathbf{Z}|\Psi)$, we add $I((\mathbf{X}, \Psi); \mathbf{Z})$ to the objective of the semi-implicit VAE. The objective of our *Enhanced Semi-Implicit VAE (ESI-VAE)* is defined as follows:

$$\mathcal{L}_{\text{ESI-VAE}}(\phi) = \underline{\mathcal{L}} + B_K + \gamma I(\mathbf{X}, \Psi; \mathbf{Z}), \quad (12)$$

where $\gamma > 0$ is a regularization parameter for the mutual information. In practice, γ is treated as a hyperparameter and optimized through cross-validation. The added term $I(\mathbf{X}, \Psi; \mathbf{Z})$ can be expanded as $I(\mathbf{Z}; \Psi) + I(\mathbf{X}; \mathbf{Z}|\Psi)$ which shows how it explicitly compensates for mutual information terms minimized in (10). The computation of $I(\mathbf{X}, \Psi; \mathbf{Z})$ imposes a

challenge which has been addressed by different works (Esmaili et al. [2019], Gabri  et al. [2018]). The challenge is to approximate the marginal distribution $q_\phi(\mathbf{Z})$ which is not computationally tractable. To approximate $q_\phi(\mathbf{Z})$ by the samples, one can use Monte-Carlo sampling (Esmaili et al. [2019]) over a large amount of data, which could be expensive and cause optimization instability. An alternative solution is to leverage the dual representation of the mutual information (Belghazi et al. [2018]) which is tractable to compute and stable to optimize. We use f -divergence dual estimation to approximate the mutual information lower bound by solving the optimization problem on the space of functions with finite expectations in our latent domain. These methods for estimation offer a stable and computationally efficient technique to approximate the mutual information which can be effectively implemented by drawing samples from the joint and marginals of the posterior.

3.2 Dual form

To promote the correlation between input (\mathbf{X}) and latent representation (\mathbf{Z}), the following term is added to the semi-implicit VAE objective function, $I(\mathbf{X}, \Psi; \mathbf{Z}) = \text{KL}(q(\mathbf{Z}, \mathbf{X}, \Psi)||q(\mathbf{Z})q(\mathbf{X}, \Psi))$. We elaborate in the following Lemma:

Lemma 2: *Substituting the KL-divergence with the variational f -divergence, a lower bound for the objective function in (12) is given by:*

$$\begin{aligned} \mathcal{L}_{\text{ESI-VAE}}(\phi) \geq & \max_T \underline{\mathcal{L}}_K + \gamma \left[\mathbb{E}_{q_\phi(\mathbf{X}, \Psi, \mathbf{Z})} [T(\mathbf{X}, \Psi, \mathbf{Z})], \right. \\ & \left. - \mathbb{E}_{q_\phi(\mathbf{X}, \Psi)q_\phi(\mathbf{Z})} [f^*(T(\mathbf{X}, \Psi, \mathbf{Z}))] \right] \end{aligned}$$

where $f(T) = T \log T$, f^* is the convex conjugate function of f , and T represents all possible functions such that their expectations are finite. See Supplementary Material for the proof.

To evaluate $\mathbb{E}_{q_\phi(\mathbf{X}, \Psi, \mathbf{Z})}[\cdot]$ and $\mathbb{E}_{q(\mathbf{X}, \Psi)q_\phi(\mathbf{Z})}[\cdot]$ in a tractable way, we take the following approach. It is possible to draw samples from $(\mathbf{X}^{(i)}, \Psi^{(i)}, \mathbf{Z}^{(i)}) \sim q_\phi(\mathbf{X}, \Psi, \mathbf{Z}) = q_\phi(\mathbf{Z}|\mathbf{X}, \Psi)q_\phi(\Psi|\mathbf{X})q(\mathbf{X})$ by first sampling ϵ from $q(\epsilon)$. Having ϵ and \mathbf{X} (observed graph), we can evaluate Ψ from $\Psi = g_\phi(\mathbf{X}, \epsilon)$, and then generate \mathbf{Z} from $q_\phi(\mathbf{Z}|\mathbf{X}, \Psi)$. To create the samples of the marginalized distributions $q_\phi(\mathbf{Z})$ and $q_\phi(\mathbf{X}, \Psi)$, we randomly choose a pair $(\mathbf{X}, \Psi)^{(j)}$ then sample from $\mathbf{Z} \sim q_\phi(\mathbf{Z}|\mathbf{X}, \Psi)^{(i)}$, $i \neq j$. In practice, however, we can effectively get samples from the observed attributes and then permute the representations \mathbf{Z} . This trick is first used in Arcones and Gine [1992], and proved to be sufficiently accurate so long as the training set size is large enough. To be specific, we take

correlated tuples of $(\mathbf{X}, \Psi, \mathbf{Z})$ generated previously, fix the data on the (\mathbf{X}, Ψ) axis and permute the \mathbf{Z} samples to decorrelate (\mathbf{X}, Ψ) and \mathbf{Z} to reflect how the samples of marginalized distributions are generated. It is worth mentioning that the consistency of the estimation we have used relies on the size of the family T and number of samples. Based on the universal approximation theorem for the neural networks (Hornik et al. [1989]), T in Lemma 2 can be closely approximated by neural networks. Moreover, for graph embedding, one need to feed the whole dataset to the networks which leads to near accurate estimation of the two expectations in Lemma 2.

3.3 Training

The full training algorithm is described in Algorithm 1 and the details are discussed in the following subsections.

Objective: The final objective can be readily extended to graph convolution neural networks defined in Poincaré space first by defining \mathcal{L}^g and B_K^g :

$$\mathcal{L}^g = \mathbb{E}_{q(\mathbf{X}, \mathbf{A})} \mathbb{E}_{q_\phi(\Psi | \mathbf{X}, \mathbf{A})} \mathbb{E}_{q_\phi(\mathbf{Z} | \Psi, \mathbf{X}, \mathbf{A})} \left[\log \left(\frac{p(\mathbf{A}, \mathbf{X}, \mathbf{Z})}{q_\phi(\mathbf{Z} | \Psi, \mathbf{X}, \mathbf{A})} \right) \right], \quad (13)$$

$$B_K^g = \mathbb{E}_{q(\mathbf{X}, \mathbf{A})} \mathbb{E}_{\Psi, \Psi^{(1)}, \dots, \Psi^{(K)} \sim q_\phi(\Psi | \mathbf{X}, \mathbf{A})} \left[\text{KL}(q_\phi(\mathbf{Z} | \Psi, \mathbf{X}, \mathbf{A})) || h_K^g(\mathbf{Z} | \mathbf{X}, \mathbf{A}) \right],$$

where,

$$h_K^g(\mathbf{Z} | \mathbf{X}, \mathbf{A}) = \frac{1}{K+1} [q_\phi(\mathbf{Z} | \Psi, \mathbf{X}, \mathbf{A}) + \sum_{k=1}^K q_\phi(\mathbf{Z} | \Psi^{(k)}, \mathbf{X}, \mathbf{A})].$$

Finally, we can formulate our proposed ESI-HGE objective as follows:

$$\mathcal{L}_{\text{ESI-HGE}} = \mathcal{L}^g + B_K^g + \gamma \max_T \left[\mathbb{E}_{q_\phi(\mathbf{X}, \Psi, \mathbf{Z})} [T(\mathbf{X}, \mathbf{Z}, \Psi)] - \mathbb{E}_{q_\phi(\mathbf{X}, \Psi) q_\phi(\mathbf{Z})} [f^*(T(\mathbf{X}, \Psi, \mathbf{Z}))] \right], \quad (14)$$

where we seek to estimate ϕ as well as T (parameters of the networks) that maximize the objective.

Encoder and decoder: Like conventional encoders in GVAEs, the ESI-HGE encoder returns the mean and the variance and then uses the exponential mapping defined in (6) to get parameters of

$q_\phi(\mathbf{Z}^i | \Psi^i, \mathbf{X}, \mathbf{A})$, which is defined as a wrapped normal distribution. We use the technique proposed by Mathieu et al. [2019] to sample from the wrapped normal distribution. For the decoder, our generative model is given by an inner product in the vector space.

Mutual information estimation: The function T is defined to be a neural network where its inputs are $(\mathbf{X}, \Psi, \mathbf{Z})$. As Ψ and \mathbf{Z} are in the Poincaré space, we first pass them through a *gyroplane* layer to map from the Poincaré space to Euclidean space (Ganea et al. [2018]).

4 Experiments and Results

We evaluated the utility of our proposed graph embedding method by conducting experiments on different analytical tasks: edge link prediction, node classification, latent space interpretability, and mutual information stored in the latent codes.

4.1 Link Prediction

We investigate link prediction accuracy of our embedding on three popular citation networks (CORA, Citeseer, and Pubmed) to facilitate comparisons to other state-of-the-art methods. We report the graph data statistics for these networks in the Supplementary Material. We split each dataset into training, validation, and test sets containing, 85%, 5%, and 10% of the total network links, respectively. To avoid performance differences due to random sampling we used identical training/validation/test splits for the comparative methods when possible.

Setup: We optimized the network for 1000 epochs, where we used edge link prediction accuracy on the validation set to tune the network hyper-parameters including early stopping. Unless stated differently, in all experiments we set our latent variable dimension to $F = 16$, and we use a generative network with one hidden layer of size 32. We compare our method with and without the enhanced SIVI objective (Eq. 12) to separately evaluate the effect of the additional mutual information term. To aid in fair comparison, we used code from the authors of Kipf and Welling [2016], Bojchevski and Günnemann [2018], and Hasanzadeh et al. [2019] to implement GAE, GVAE, G2G, and SIG-VAE¹. We directly used the results for DW and SEAL published by Perozzi et al. [2014] and Zhang et al. [2018], respectively. Due to the prohibitively

¹We note that implementation details in comparative methods did not always include specific parameter settings necessary for exact replication of their published results. Specifically, we were not able to reproduce the results in Hasanzadeh et al. [2019] using the authors' available code.

Algorithm 1 Enhanced Semi-Implicit Hyperbolic Graph Embedding**Require:** \mathbf{X} , \mathbf{A} , latent variable dimension \mathbf{Z}_{dim} , optimizers: G_ϕ , G_T

- 1: get the adjacency matrix \mathbf{A} and nodes attributes \mathbf{X} ; create train/test/validation
- 2: initialize ϕ , T , G
- 3: **repeat**
- 4: get samples of $\mathbf{z} \sim q_\phi(\mathbf{Z}|\mathbf{X}, \Psi, \mathbf{A})$, $\Psi \sim q_\phi(\Psi|\mathbf{X}, \mathbf{A})$ and create the tuple $\{(\mathbf{X}, \Psi, \mathbf{Z})\}_{i=1}^N$,
- 5: *permute* the tuple along the \mathbf{Z} -column to get $\{(\mathbf{X}, \Psi, \hat{\mathbf{Z}})\}_{i=1}^N$
- 6: $\phi \leftarrow G \left[\nabla_\phi \left(\mathcal{L}^g + B_K^g + \frac{\gamma}{N} \left(\sum_{i=1}^N [T(\mathbf{X}^{(i)}, \Psi^{(i)}, \mathbf{Z}^{(i)})] - [f^*(T(\mathbf{X}^{(i)}, \Psi^{(i)}, \hat{\mathbf{Z}}^{(i)}))] \right) \right) \right]$
- 7: $t \leftarrow G_t \left[\nabla_t \left(\frac{1}{N} \sum_{i=1}^b T(\mathbf{X}^{(i)}, \Psi^{(i)}, \mathbf{Z}^{(i)}) - \frac{1}{N} \sum_{i=1}^N f^*(T(\mathbf{X}^{(i)}, \Psi^{(i)}, \mathbf{Z}^{(i)})) \right) \right]$
- 8: **until** convergence

Table 1: Link prediction performance in networks with node attributes.

Method	CORA		Citeseer		Pubmed	
	AUC	AP	AUC	AP	AUC	AP
DW [Perozzi et al., 2014]	83.1±0.01	85.0±0.00	80.5±0.02	83.6±0.01	84.4±0.00	84.1±0.00
SEAL [Zhang et al., 2018]	90.1±0.10	83.0±0.30	83.6±0.20	77.6±0.20	96.7±0.10	90.1±0.10
GAE [Kipf and Welling, 2016]	91.0±0.02	92.0±0.03	89.5±0.04	89.9±0.05	96.4±0.00	96.5±0.00
GVAE [Kipf and Welling, 2016]	91.4±0.01	92.6±0.01	90.8±0.02	92.0±0.02	94.4±0.02	94.7±0.02
G2G [Bojchevski and Günnemann, 2018]	92.1±0.90	92.6±0.80	94.5±0.40	93.2±0.70	94.3±0.30	93.4±0.50
SIG-VAE [Hasanzadeh et al., 2019]	91.2±0.05	92.0±0.10	93.5±0.02	94.3±0.10	96.5±0.30	96.0±0.50
EVAE	91.6±0.02	92.3±0.02	94.7±0.10	94.5±0.15	-	-
HGE	91.8±0.10	92.7±0.15	93.1±0.40	94.2±0.70	-	-
SI-HGE	92.1±0.01	91.2±0.10	94.2±0.10	94.9±0.10	93.7 ±0.45	94.5 ±0.05
ESI-HGE	92.4±0.20	93.2±0.02	95.9±0.10	96.2±0.2	96.6±0.15	96.8±0.35

large size of the Pubmed dataset, we report results directly from the relevant work for this dataset. We include full implementation details in the Supplementary Material.

Performance: Table 1 reports the test-set performance measured by average precision (AP) and area under ROC curve (AUC) based on 10 runs with different random initialization for the weights of the neural network. Based on these results, our proposed method outperforms other methods across the citation network, with which we can draw couple of conclusions: 1) The inclusion of a Poincaré embedding (SI-HGE in Table 1) shows that the hyperbolic embedding improves over Euclidean SIG-VAE and provides better or comparable performance compared with other methods; and 2) ESI-HGE (obtained by enhancing the objective of SIVI and incorporating the Poincaré embedding) significantly outperforms SI-HGE and the other methods across all datasets.

Visualization: To build intuition behind the obtained results, in Figure 1 we visualize the embedded nodes of GVAE, SIG-VAE, SI-HGE, and ESI-HGE for the Citeseer dataset. The Citeseer citation graph contains five labeled classes (indicated by the colors in the figure). We note that the class labels were not used during training. As the first column of the figure

shows, we see that GVAE severely blends the latent codes of the graph for different classes, indicating mode collapsing. SIG-VAE only marginally improves this separation. By endowing the latent space with a hyperbolic geometry, SI-HGE and more significantly ESI-HGE result in more distinct embeddings for each class. In the second and third columns of Figure 1 we visualize contour plots and heatmaps of the learned posterior distributions. The GVAE embedding recovers a uni-modal distribution concentrated at the origin, while SIG-VAE is able to represent additional skewness in the distribution. Adding the Poincaré embedding provides a more expressive prior with distinct modes, but still shows blending of the different classes. In contrast, ESI-HGE clearly enriches the posterior and infers more expressive distributions with multiple modes distributed across the hyperbolic space.

4.2 Node Classification

Using the training results from the previous section, we explore classification accuracy on the citation graphs using the node embeddings as input against the labeled classes of each node. We used Cora and Citeseer as described for link prediction, and we omit Pubmed due to limited computational resources. We report the classification accuracy for competing methods from Table

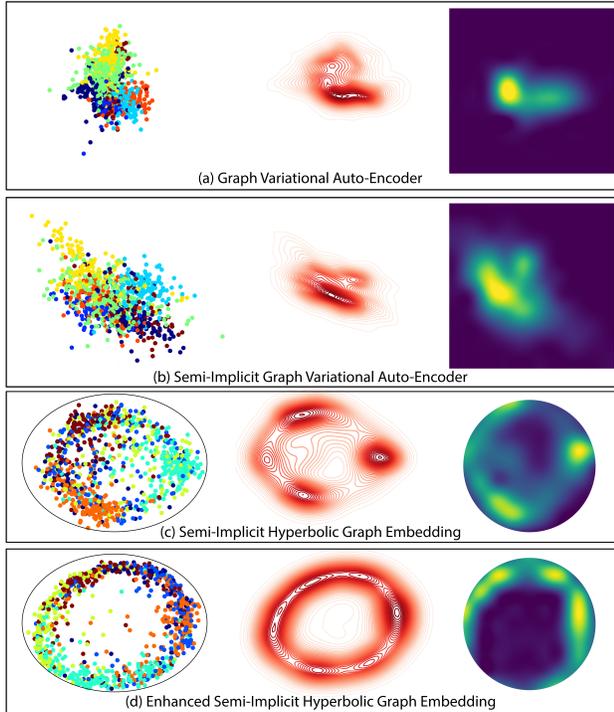


Figure 1: Visualization of Citeseer latent space for (a) GVAE, (b) SIG-VAE, (c) SI-HGE, and (d) ESI-HGE. The GVAE embedding learned a uni-modal distribution concentrated at the origin; SIG-VAE is able to capture additional skewness in the distribution; adding the Poincaré embedding provides a more expressive prior with distinct modes, but still shows blending of the different classes (please note that the Citeseer dataset has 6 distinct classes); in contrast, ESI-HGE clearly enriches the posterior representation and is able to recover more expressive distributions with multiple modes distributed across the hyperbolic space

3 in [Hasanzadeh et al., 2019].

Setup: As previously mentioned, we first trained ESI-HGE in an unsupervised manner without access to class information. After, we train a classifier where its inputs are the learned embeddings and labels are the node classes. To show robustness of the inferred latent codes, we randomly remove 10% of edges from the graphs to use for test-set performance. Additional details for the classification setup are provided in the Supplementary Material.

Performance: Classification results are summarized in Table 2. Our proposed method provides a clear improvement in classification accuracy even though it was not explicitly trained for this task. We note that our setup slightly differs from the methodology in Hasanzadeh et al. [2019], which explicitly included classification as a semi-supervised loss during VAE

training. Thus, we expect our results to further improve under this training procedure. The result suggests that the learned embeddings in ESI-HGE better capture structure in the graph compared to the other frameworks, which is supported by the richer posterior distribution visualized in Figure 1.

Table 2: Classification accuracy on citation graphs.

Method	CORA	Citeseer	Pubmed
DeepWalk	67.2	43.2	65.3
LP	68.0	45.3	63.0
ICA	75.1	69.1	73.9
Planetoid	75.7	64.7	77.2
GCN	81.5	70.3	79.0
SIG-VAE	79.7	70.4	79.3
ESI-HGE	84.4	72.8	82.1

4.3 Stored Mutual Information

Setup: Another metric that elucidates the usefulness of the learned representation is the information content stored in the embeddings. Using the same citation graphs as previously outlined, here we investigate the mutual information between the observed graph attributes \mathbf{X} and the inferred embeddings \mathbf{Z} . Similar to node classification, we first trained our model without access to class labels. Next, we used the method proposed by Belghazi et al. [2018] to estimate the mutual information. More information regarding the setup is provided in the Supplementary Material.

Performance: The results summarized in Table 3 demonstrate increased mutual information between input and latent spaces due to the enhanced SIVI objective. There is a steady increase in mutual information moving from GVAE to SIG-VAE to SI-HGE to ESI-HGE, suggesting both improved latent representation and successful modeling of the correlation.

Table 3: Mutual information between $(\mathbf{X}, \Psi), z$.

Method	CORA	Citeseer	Pubmed
GVAE	1.05	1.37	3.23
SIG-VAE	2.52	2.84	4.74
SI-HGE	3.12	3.54	5.12
ESI-HGE	3.96	5.58	6.45

4.4 Interpretability and Posterior Visualization

In the Supplementary Material we outline an additional experiment on a sequence of synthetically generated images forming a graph: we construct a parent image containing a random polygon with random intensity; each successive leaf node is a copy of its parent image with the inclusion of an additional random polygon. We treat the vectorized image pixel intensities as the node features. In Figure 2 we visualize the learned latent embedding of ESI-HGE and show a natural hierarchical organization of the nodes that closely matches the true graph structure. This reveals how hyperbolic graph embedding can effectively help to preserve hierarchical representation in latent space.

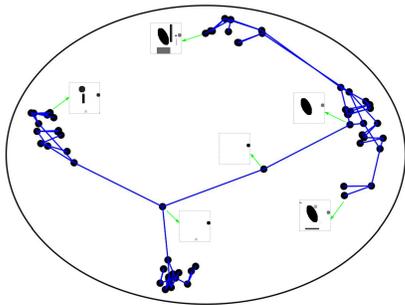


Figure 2: Visualization of the learned latent codes for the synthetic image graph data.

5 Conclusion

We have introduced an enhanced semi-implicit variational auto-encoder for relational graph data with a hyperbolic latent embedding, termed *ESI-HGE*. In addition to representing correlated structures in the latent distribution, our model additionally captures mutual information between the input and the latent space, and thus enhances the learned representations. The theoretical advantages are reflected in our experimental findings, as *ESI-HGE* outperformed other state-of-the-art methods on edge link prediction and unsupervised node classification, while also increasing interpretability in the latent space.

Acknowledgments

This work was supported by AWS Machine Learning Research Grant, the Office of Naval Research grant N00014-19-1-2590 and Army grant W911NF1910413.

References

- A. B. Adcock, B. D. Sullivan, and M. W. Mahoney. Tree-like structure in large social and information networks. In *2013 IEEE 13th International Conference on Data Mining*, pages 1–10, 2013.
- M. A. Arcones and E. Gine. On the bootstrap of u and v statistics. *Ann. Statist.*, 20(2):655–674, 06 1992. doi: 10.1214/aos/1176348650. URL <https://doi.org/10.1214/aos/1176348650>.
- M. I. Belghazi, A. Baratin, S. Rajeswar, S. Ozair, Y. Bengio, A. Courville, and R. D. Hjelm. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*, 2018.
- Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- A. Bojchevski and S. Günnemann. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. *n International Conference on Learning Representations*, 2018.
- H. Cai, V. W. Zheng, and K. C. Chang. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge and Data Engineering*, 30(9):1616–1637, 2018.
- T. R. Davidson, L. Falorsi, N. De Cao, T. Kipf, and J. M. Tomczak. Hyperspherical variational auto-encoders. *34th Conference on Uncertainty in Artificial Intelligence (UAI-18)*, 2018.
- L. Dos Santos, B. Piwowarski, and P. Gallinari. Multilabel classification on heterogeneous graphs with gaussian embeddings. In P. Frasconi, N. Landwehr, G. Manco, and J. Vreeken, editors, *Machine Learning and Knowledge Discovery in Databases*, pages 606–622, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46227-1.
- B. Esmaeili, H. Wu, S. Jain, A. Bozkurt, N. Siddharth, B. Paige, D. H. Brooks, J. Dy, and J.-W. Meent. Structured disentangled representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2525–2534. PMLR, 2019.
- N. I. Fisher, T. Lewis, and B. J. J. Embleton. *Statistical Analysis of Spherical Data*. Cambridge University Press, 1987. doi: 10.1017/CBO9780511623059.
- M. Gabrié, A. Manoel, C. Luneau, N. Macris, F. Krzakala, L. Zdeborová, et al. Entropy and mutual information in models of deep neural networks. In *Advances in Neural Information Processing Systems*, pages 1821–1831, 2018.

- O. Ganea, G. Bécigneul, and T. Hofmann. Hyperbolic neural networks. In *Advances in neural information processing systems*, pages 5345–5355, 2018.
- A. Hasanzadeh, E. Hajiramezani, K. Narayanan, N. Duffield, M. Zhou, and X. Qian. Semi-implicit graph variational auto-encoders. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 10712–10723. Curran Associates, Inc., 2019.
- S. He, K. Liu, G. Ji, and J. Zhao. Learning to represent knowledge graphs with gaussian embedding. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, CIKM '15*, page 623–632, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450337946. doi: 10.1145/2806416.2806502. URL <https://doi.org/10.1145/2806416.2806502>.
- K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- T. N. Kipf and M. Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.
- E. Mathieu, C. Le Lan, C. J. Maddison, R. Tomioka, and Y. W. Teh. Continuous hierarchical representations with poincaré variational auto-encoders. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 12565–12576. Curran Associates, Inc., 2019.
- M. Nickel and D. Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in neural information processing systems*, pages 6338–6347, 2017.
- M. Nickel and D. Kiela. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. *arXiv preprint arXiv:1806.03417*, 2018.
- S. Nowozin, B. Cseke, and R. Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems*, pages 271–279, 2016.
- I. Ovinnikov. Poincaré wasserstein autoencoder. *CoRR*, abs/1901.01427, 2019. URL <http://arxiv.org/abs/1901.01427>.
- B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710, 2014.
- E. Ravasz and A.-L. Barabási. Hierarchical organization in complex networks. *Phys. Rev. E*, 67:026112, Feb 2003. doi: 10.1103/PhysRevE.67.026112.
- A. L. Rezaabad and S. Vishwanath. Learning representations by maximizing mutual information in variational autoencoders. *arXiv preprint arXiv:1912.13361*, 2019.
- P. Sen, G. M. Namata, M. Bilgic, L. Getoor, B. Gallagher, and T. Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008. URL <http://www.cs.iit.edu/~ml/pdfs/sen-aimag08.pdf>.
- M. Steyvers and J. B. Tenenbaum. The large-scale structure of semantic networks: statistical analyses and a model of semantic growth. *Cogn Sci*, 29(1): 41–78, Jan 2005.
- M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.*, 1(1–2):1–305, Jan. 2008. ISSN 1935-8237. doi: 10.1561/2200000001. URL <https://doi.org/10.1561/2200000001>.
- M. Yin and M. Zhou. Semi-implicit variational inference. In J. Dy and A. Krause, editors, *Proceedings of Machine Learning Research*, volume 80, pages 5660–5669, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- M. Zhang et al. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*, pages 5165–5175, 2018.