

---

# Graphite: Iterative Generative Modeling of Graphs

---

Aditya Grover<sup>1</sup> Aaron Zweig<sup>1</sup> Stefano Ermon<sup>1</sup>

## Abstract

Graphs are a fundamental abstraction for modeling relational data. However, graphs are discrete and combinatorial in nature, and learning representations suitable for machine learning tasks poses statistical and computational challenges. In this work, we propose *Graphite*, an algorithmic framework for unsupervised learning of representations over nodes in large graphs using deep latent variable generative models. Our model parameterizes variational autoencoders (VAE) with graph neural networks, and uses a novel iterative graph refinement strategy inspired by low-rank approximations for decoding. On a wide variety of synthetic and benchmark datasets, Graphite outperforms competing approaches for the tasks of density estimation, link prediction, and node classification. Finally, we derive a theoretical connection between message passing in graph neural networks and mean-field variational inference.

## 1. Introduction

Latent variable generative modeling is an effective approach for unsupervised representation learning of high-dimensional data (Loehlin, 1998). In recent years, representations learned by latent variable models parameterized by deep neural networks have shown impressive performance on many tasks such as semi-supervised learning and structured prediction (Kingma et al., 2014; Sohn et al., 2015). However, these successes have been largely restricted to specific data modalities such as images and speech. In particular, it is challenging to apply current deep generative models for large scale graph-structured data which arise in a wide variety of domains in physical sciences, information sciences, and social sciences.

To effectively model the relational structure of large graphs for deep learning, prior works have proposed to use *graph*

*neural networks* (Gori et al., 2005; Scarselli et al., 2009; Bruna et al., 2013). A graph neural network learns node-level representations by parameterizing an iterative message passing procedure between nodes and their neighbors. The tasks which have benefited from graph neural networks, including semi-supervised learning (Kipf & Welling, 2017) and few shot learning (Garcia & Bruna, 2018), involve *encoding* an input graph to a final output representation (such as the labels associated with the nodes). The inverse problem of learning to *decode* a hidden representation into a graph, as in the case of a latent variable generative model, is a pressing challenge that we address in this work.

We propose *Graphite*, a latent variable generative model for graphs based on variational autoencoding (Kingma & Welling, 2014). Specifically, we learn a directed model expressing a joint distribution over the entries of adjacency matrix of graphs and latent feature vectors for every node. Our framework uses graph neural networks for inference (encoding) and generation (decoding). While the encoding is straightforward and can use any existing graph neural network, the decoding of these latent features to reconstruct the original graph is done using a multi-layer iterative procedure. The procedure starts with an initial reconstruction based on the inferred latent features, and iteratively refines the reconstructed graph via a message passing operation. The iterative refinement can be efficiently implemented using graph neural networks. In addition to the Graphite model, we also contribute to the theoretical understanding of graph neural networks by deriving equivalences between message passing in graph neural networks with mean-field inference in latent variable models via kernel embeddings (Smola et al., 2007; Dai et al., 2016), formalizing what has thus far has been largely speculated empirically to the best of our knowledge (Yoon et al., 2018).

In contrast to recent works focussing on generation of small graphs *e.g.*, molecules (You et al., 2018; Li et al., 2018), the Graphite framework is particularly suited for representation learning on large graphs. Such representations are useful for several downstream tasks. In particular, we demonstrate that representations learned via Graphite outperform competing approaches for graph representation learning empirically for the tasks of density estimation (over entire graphs), link prediction, and semi-supervised node classification on synthetic and benchmark datasets.

---

<sup>1</sup>Department of Computer Science, Stanford University, USA. Correspondence to: Aditya Grover <adityag@cs.stanford.edu>.

## 2. Preliminaries

Throughout this work, we assume that all probability distributions admit absolutely continuous densities on a suitable reference measure. Consider a weighted undirected graph  $G = (V, E)$  where  $V$  and  $E$  denote index sets of nodes and edges respectively. Additionally, we denote the (optional) feature matrix associated with the graph as  $\mathbf{X} \in \mathbb{R}^{n \times m}$  for an  $m$ -dimensional signal associated with each node, for *e.g.*, these could refer to the user attributes in a social network. We represent the graph structure using a symmetric adjacency matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$  where  $n = |V|$  and the entries  $A_{ij}$  denote the weight of the edge between node  $i$  and  $j$ .

### 2.1. Weisfeiler-Lehman algorithm

The Weisfeiler-Lehman (WL) algorithm (Weisfeiler & Lehman, 1968; Douglas, 2011) is a heuristic test of graph isomorphism between any two graphs  $G$  and  $G'$ . The algorithm proceeds in iterations. Before the first iteration, we label every node in  $G$  and  $G'$  with a scalar *isomorphism invariant* initialization (*e.g.*, node degrees). That is, if  $G$  and  $G'$  are assumed to be isomorphic, then an isomorphism invariant initialization is one where the matching nodes establishing the isomorphism in  $G$  and  $G'$  have the same labels (a.k.a. messages). Let  $\mathbf{H}^{(0)} = [h_1^{(0)}, h_2^{(0)}, \dots, h_n^{(0)}]^T$  denote the vector of initializations for the nodes in the graph at iteration  $l \in \mathbb{N} \cup 0$ . At every iteration  $l > 0$ , we perform a relabelling of nodes in  $G$  and  $G'$  based on a message passing update rule:

$$\mathbf{H}^{(l)} \leftarrow \text{hash} \left( \mathbf{A} \mathbf{H}^{(l-1)} \right) \quad (1)$$

where  $\mathbf{A}$  denotes the adjacency matrix of the corresponding graph and  $\text{hash} : \mathbb{R}^n \rightarrow \mathbb{R}^n$  is any suitable hash function *e.g.*, a non-linear activation. Hence, the message for every node is computed as a hashed sum of the messages from the neighboring nodes (since  $A_{ij} \neq 0$  only if  $i$  and  $j$  are neighbors). We repeat the process for a specified number of iterations, or until convergence. If the label sets for the nodes in  $G$  and  $G'$  are equal (which can be checked using sorting in  $O(n \log n)$  time), then the algorithm declares the two graphs  $G$  and  $G'$  to be isomorphic.

The “ $k$ -dim” WL algorithm extends the 1-dim algorithm above by simultaneously passing messages of length  $k$  (each initialized with some isomorphism invariant scheme). A positive test for isomorphism requires equality in all  $k$  dimensions for nodes in  $G$  and  $G'$  after the termination of message passing. This algorithmic test is a heuristic which guarantees no false negatives but can give false positives wherein two non-isomorphic graphs can be falsely declared isomorphic. Empirically, the test has been shown to fail on some regular graphs but gives excellent performance on real-world graphs (Shervashidze et al., 2011).

### 2.2. Graph neural networks

Intuitively, the WL algorithm encodes the structure of the graph in the form of messages at every node. Graph neural networks (GNN) build on this observation and parameterize an unfolding of the iterative message passing procedure which we describe next.

A GNN consists of many layers, indexed by  $l \in \mathbb{N}$  with each layer associated with an activation  $\eta_l$  and a dimensionality  $d_l$ . In addition to the input graph  $\mathbf{A}$ , every layer  $l \in \mathbb{N}$  of the GNN takes as input the activations from the previous layer  $\mathbf{H}^{(l-1)} \in \mathbb{R}^{n \times d_{l-1}}$ , a family of linear transformations  $\mathcal{F}_l : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ , and a matrix of learnable weight parameters  $\mathbf{W}_l \in \mathbb{R}^{d_{l-1} \times d_l}$  and optional bias parameters  $\mathbf{B}_l \in \mathbb{R}^{n \times d_l}$ . Recursively, the layer wise propagation rule in a GNN is given by:

$$\mathbf{H}^{(l)} \leftarrow \eta_l \left( \mathbf{B}_l + \sum_{f \in \mathcal{F}_l} f(\mathbf{A}) \mathbf{H}^{(l-1)} \mathbf{W}_l \right) \quad (2)$$

with the base cases  $\mathbf{H}^{(0)} = \mathbf{X}$  and  $d_0 = m$ . Here,  $m$  is the feature dimensionality. If there are no explicit node features, we set  $\mathbf{H}^{(0)} = \mathbf{I}_n$  (identity) and  $d_0 = n$ . Several variants of graph neural networks have been proposed in prior work. For instance, graph convolutional networks (GCN) (Kipf & Welling, 2017) instantiate graph neural networks with a permutation equivariant propagation rule:

$$\mathbf{H}^{(l)} \leftarrow \eta_l \left( \mathbf{B}_l + \tilde{\mathbf{A}} \mathbf{H}^{(l-1)} \mathbf{W}_l \right) \quad (3)$$

where  $\tilde{\mathbf{A}} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$  is the symmetric diagonalization of  $\mathbf{A}$  given the diagonal degree matrix  $\mathbf{D}$  (*i.e.*,  $D_{ii} = \sum_{(i,j) \in E} A_{ij}$ ), and same base cases as before. Comparing the above with the WL update rule in Eq. (1), we can see that the activations for every layer in a GCN are computed via parameterized, scaled activations (messages) of the previous layer being propagated over the graph, with the hash function implicitly specified using an activation function  $\eta_l$ .

Our framework is agnostic to instantiations of message passing rule of a graph neural network in Eq. (2), and we use graph convolutional networks for experimental validation due to the permutation equivariance property. For brevity, we denote the output  $\mathbf{H}$  for the final layer of a multi-layer graph neural network with input adjacency matrix  $\mathbf{A}$ , node feature matrix  $\mathbf{X}$ , and parameters  $\langle \mathbf{W}, \mathbf{B} \rangle$  as  $\mathbf{H} = \text{GNN}_{\langle \mathbf{W}, \mathbf{B} \rangle}(\mathbf{A}, \mathbf{X})$ , with appropriate activation functions and linear transformations applied at each hidden layer of the network.

## 3. Generative Modeling via Graphite

For generative modeling of graphs, we are interested in learning a parameterized distribution over adjacency matri-




Figure 1. Latent variable model for Graphite. Observed evidence variables in gray.

ces  $\mathbf{A}$ . In this work, we restrict ourselves to modeling graph *structure* only, and any additional information in the form of node features  $\mathbf{X}$  is incorporated as conditioning evidence.

In Graphite, we adopt a latent variable approach for modeling the generative process. That is, we introduce latent variable vectors  $\mathbf{Z}_i \in \mathbb{R}^k$  and evidence feature vectors  $\mathbf{X}_i \in \mathbb{R}^m$  for each node  $i \in \{1, 2, \dots, n\}$  along with an observed variable for each pair of nodes  $A_{ij} \in \mathbb{R}$ . Unless necessary, we use a succinct representation  $\mathbf{Z} \in \mathbb{R}^{n \times k}$ ,  $\mathbf{X} \in \mathbb{R}^{n \times m}$ , and  $\mathbf{A} \in \mathbb{R}^{n \times n}$  for the variables henceforth. The conditional independencies between the variables can be summarized in the directed graphical model (using plate notation) in Figure 1. We can learn the model parameters  $\theta$  by maximizing the marginal likelihood of the observed adjacency matrix conditioned on  $\mathbf{X}$ :

$$\max_{\theta} \log p_{\theta}(\mathbf{A}|\mathbf{X}) = \log \int_{\mathbf{Z}} p_{\theta}(\mathbf{A}, \mathbf{Z}|\mathbf{X}) d\mathbf{Z} \quad (4)$$

Here,  $p(\mathbf{Z}|\mathbf{X})$  is a fixed prior distribution over the latent features of every node *e.g.*, isotropic Gaussian. If we have multiple graphs in our dataset, we maximize the expected log-likelihoods over all the corresponding adjacency matrices. We can obtain a tractable, stochastic evidence lower bound (ELBO) to the above objective by introducing a variational posterior  $q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X})$  with parameters  $\phi$ :

$$\log p_{\theta}(\mathbf{A}|\mathbf{X}) \geq \mathbb{E}_{q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X})} \left[ \log \frac{p_{\theta}(\mathbf{A}, \mathbf{Z}|\mathbf{X})}{q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X})} \right] \quad (5)$$

The lower bound is tight when the variational posterior  $q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X})$  matches the true posterior  $p_{\theta}(\mathbf{Z}|\mathbf{A}, \mathbf{X})$  and hence maximizing the above objective optimizes for the parameters that define the best approximation to the true posterior within the variational family (Blei et al., 2017). We now discuss parameterizations for specifying  $q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X})$  (*i.e.*, encoder) and  $p_{\theta}(\mathbf{A}|\mathbf{Z}, \mathbf{X})$  (*i.e.*, decoder).

**Encoding using forward message passing.** Typically we use the mean field approximation for defining the variational family and hence:

$$q_{\phi}(\mathbf{Z}|\mathbf{A}, \mathbf{X}) \approx \prod_{i=1}^n q_{\phi_i}(\mathbf{Z}_i|\mathbf{A}, \mathbf{X}) \quad (6)$$

Additionally, we would like to make distributional assumptions on each variational marginal density  $q_{\phi_i}(\mathbf{Z}_i|\mathbf{A}, \mathbf{X})$  such that it is reparameterizable and easy-to-sample, such that the gradients w.r.t.  $\phi_i$  have low variance (Kingma & Welling, 2014). In Graphite, we assume isotropic Gaussian variational marginals with diagonal covariance. The parameters for the variational marginals  $q_{\phi_i}(\mathbf{Z}|\mathbf{A}, \mathbf{X})$  are specified using a graph neural network:

$$\boldsymbol{\mu}, \boldsymbol{\sigma} = \text{GNN}_{\phi}(\mathbf{A}, \mathbf{X}) \quad (7)$$

where  $\boldsymbol{\mu}$  and  $\boldsymbol{\sigma}$  denote the vector of means and standard deviations for the variational marginals  $\{q_{\phi_i}(\mathbf{Z}_i|\mathbf{A}, \mathbf{X})\}_{i=1}^n$  and  $\phi = \{\phi_i\}_{i=1}^n$  are the full set of variational parameters.

**Decoding using reverse message passing.** For specifying the observation model  $p_{\theta}(\mathbf{A}|\mathbf{Z}, \mathbf{X})$ , we cannot directly use a graph neural network since we do not have an input graph for message passing. To sidestep this issue, we propose an iterative two-step approach that alternates between defining an intermediate graph and then gradually refining this graph through message passing. Formally, given a latent matrix  $\mathbf{Z}$  and an input feature matrix  $\mathbf{X}$ , we iterate over the following sequence of operations:

$$\widehat{\mathbf{A}} = \frac{\mathbf{Z}\mathbf{Z}^T}{\|\mathbf{Z}\|^2} + \mathbf{1}\mathbf{1}^T, \quad (8)$$

$$\mathbf{Z}^* = \text{GNN}_{\theta}(\widehat{\mathbf{A}}, [\mathbf{Z}|\mathbf{X}]) \quad (9)$$

where the second argument to the GNN is a concatenation of  $\mathbf{Z}$  and  $\mathbf{X}$ . The first step constructs an intermediate weighted graph  $\widehat{\mathbf{A}} \in \mathbb{R}^{n \times n}$  by applying an inner product of  $\mathbf{Z}$  with itself and adding an additional constant of 1 to ensure entries are non-negative. And the second step performs a pass through a parameterized graph neural network. We can repeat the above sequence to gradually refine the feature matrix  $\mathbf{Z}^*$ . The final distribution over graph parameters is obtained using an inner product step on  $\mathbf{Z}^* \in \mathbb{R}^{n \times k^*}$  akin to Eq. (8), where  $k^* \in \mathbb{N}$  is determined via the network architecture. For efficient sampling, we assume the observation model factorizes:

$$p_{\theta}(\mathbf{A}|\mathbf{Z}, \mathbf{X}) = \prod_{i=1}^n \prod_{j=1}^n p_{\theta}^{(i,j)}(A_{ij}|\mathbf{Z}^*). \quad (10)$$

The distribution over the individual edges can be expressed as a Bernoulli or Gaussian distribution for unweighted and real-valued edges respectively. *E.g.*, the edge probabilities for an unweighted graph are given as  $\text{sigmoid}(\mathbf{Z}^* \mathbf{Z}^{*T})$ .

Table 1. Mean reconstruction errors and negative log-likelihood estimates (in nats) for autoencoders and variational autoencoders respectively on test instances from six different generative families. Lower is better.

	Erdoes-Renyi	Ego	Regular	Geometric	Power Law	Barabasi-Albert
GAE	221.79 ± 7.58	197.3 ± 1.99	198.5 ± 4.78	514.26 ± 41.58	519.44 ± 36.30	236.29 ± 15.13
Graphite-AE	<b>195.56</b> ± 1.49	<b>182.79</b> ± 1.45	<b>191.41</b> ± 1.99	<b>181.14</b> ± 4.48	<b>201.22</b> ± 2.42	<b>192.38</b> ± 1.61
VGAE	273.82 ± 0.07	273.76 ± 0.06	275.29 ± 0.08	274.09 ± 0.06	278.86 ± 0.12	274.4 ± 0.08
Graphite-VAE	<b>270.22</b> ± 0.15	<b>270.70</b> ± 0.32	<b>266.54</b> ± 0.12	<b>269.71</b> ± 0.08	<b>263.92</b> ± 0.14	<b>268.73</b> ± 0.09

### 3.1. Scalable learning & inference in Graphite

For representation learning of large graphs, we require the encoding and decoding steps in Graphite to be computationally efficient. On the surface, the decoding step involves inner products of potentially dense matrices  $\mathbf{Z}$ , which is an  $O(n^2k)$  operation. Here,  $k$  is the dimension of the per-node latent vectors  $\mathbf{Z}_i$  used to define  $\hat{\mathbf{A}}$ .

For any intermediate decoding step as in Eq. (8), we propose to offset this expensive computation by using the associativity property of matrix multiplications for the message passing step in Eq. (9). For notational brevity, consider the simplified graph propagation rule for a GNN:

$$\mathbf{H}^{(l)} \leftarrow \eta_l \left( \hat{\mathbf{A}} \mathbf{H}^{(l-1)} \right)$$

where  $\hat{\mathbf{A}}$  is defined in Eq. (8).

Instead of directly taking an inner product of  $\mathbf{Z}$  with itself, we note that the subsequent operation involves another matrix multiplication and hence, we can perform right multiplication instead. If  $d_l$  and  $d_{l-1}$  denote the size of the layers  $\mathbf{H}^{(l)}$  and  $\mathbf{H}^{(l-1)}$  respectively, then the time complexity of propagation based on right multiplication is given by  $O(nkd_{l-1} + nd_{l-1}d_l)$ .

The above trick sidesteps the quadratic  $n^2$  complexity for decoding in the intermediate layers without any loss in statistical accuracy. The final layer however still involves an inner product with respect to  $\mathbf{Z}^*$  between potentially dense matrices. However, since the edges are generated independently, we can approximate the loss objective by performing a Monte Carlo evaluation of the reconstructed adjacency matrix parameters in Eq. (10). By adaptively choosing the number of entries for Monte Carlo approximation, we can trade-off statistical accuracy for computational budget.

## 4. Experimental Evaluation

We evaluate Graphite on tasks involving entire graphs, nodes, and edges. We consider two variants of our proposed framework: the *Graphite-VAE*, which corresponds to a directed latent variable model as described in Section 3 and *Graphite-AE*, which corresponds to an autoencoder trained

to minimize the error in reconstructing an input adjacency matrix. For unweighted graphs (*i.e.*,  $\mathbf{A} \in \{0, 1\}^{n \times n}$ ), the reconstruction terms in the objectives for both Graphite-VAE and Graphite-AE minimize the negative cross entropy between the input and reconstructed adjacency matrices. For weighted graphs, we use the mean squared error. Additional hyperparameter details are described in Appendix B.

### 4.1. Reconstruction & density estimation

In the first set of tasks, we evaluate learning in Graphite based on held-out reconstruction losses and log-likelihoods estimated by the learned Graphite-VAE and Graphite-AE models respectively on a collection of graphs with varying sizes. In direct contrast to modalities such as images, graphs cannot be straightforwardly reduced to a fixed number of vertices for input to a graph convolutional network. One simplifying modification taken by [Bojchevski et al. \(2018\)](#) is to consider only the largest connected component for evaluating and optimizing the objective, which we appeal to as well. Thus by setting the dimensions of  $\mathbf{Z}^*$  to a maximum number of vertices, Graphite can be used for inference tasks over entire graphs with potentially smaller sizes by considering only the largest connected component.

We create datasets from six graph families with fixed, known generative processes: the Erdoes-Renyi, ego-nets, random regular graphs, random geometric graphs, random Power Law Tree and Barabasi-Albert. For each family, 300 graph instances were sampled with each instance having 10 – 20 nodes and evenly split into train/validation/test instances. As a benchmark comparison, we compare against the Graph Autoencoder/Variational Graph Autoencoder (GAE/VGAE) ([Kipf & Welling, 2016](#)). The GAE/VGAE models consist of an encoding procedure similar to Graphite. However, the decoder has no learnable parameters and reconstruction is done solely through an inner product operation (such as the one in Eq. (8)).

The mean reconstruction errors and the negative log-likelihood results on a test set of instances are shown in Table 1. Both Graphite-AE and Graphite-VAE outperform AE and VGAE significantly on these tasks, indicating the usefulness of learned decoders in Graphite.

Table 2. Citation network statistics

	Nodes	Edges	Node Features	Labels
Cora	2708	5429	1433	7
Citeseer	3327	4732	3703	6
Pubmed	19717	44338	500	3

#### 4.2. Link prediction


The task of link prediction is to predict whether an edge exists between a pair of nodes (Loehlin, 1998). Even though Graphite learns a distribution over graphs, it can be used for predictive tasks within a *single* graph. In order to do so, we learn a model for a random, connected training subgraph of the true graph. For validation and testing, we add a balanced set of positive and negative (false) edges to the original graph and evaluate the model performance based on the reconstruction probabilities assigned to the validation and test edges (similar to *denoising* of the input graph). In our experiments, we held out a set of 5% edges for validation, 10% edges for testing, and train all models on the remaining subgraph. Additionally, the validation and testing sets also each contain an equal number of non-edges.

**Datasets.** We compared across standard benchmark citation network datasets: Cora, Citeseer, and Pubmed with papers as nodes and citations as edges (Sen et al., 2008). The node-level features correspond to the text attributes in the papers. The dataset statistics are summarized in Table 2.


**Baselines and evaluation metrics.** We evaluate performance based on the Area Under the ROC Curve (AUC) and Average Precision (AP) metrics. We evaluated Graphite-VAE and Graphite-AE against the following baselines: Spectral Clustering (SC) (Tang & Liu, 2011), DeepWalk (Perozzi et al., 2014), node2vec (Grover & Leskovec, 2016), and GAE/VGAE (Kipf & Welling, 2016). SC, DeepWalk, and node2vec do not provide the ability to incorporate node features while learning embeddings, and hence we evaluate them only on the featureless datasets.

**Results.** The AUC and AP results (along with standard errors) are shown in Table 3 and Table 4 respectively averaged over 50 random train/validation/test splits. On both metrics, Graphite-VAE gives the best performance overall. Graphite-AE also gives good results, generally outperforming its closest competitor GAE.

**Qualitative evaluation.** We visualize the embeddings learned by Graphite and given by a 2D t-SNE projection (Maaten & Hinton, 2008) of the latent feature vectors (given as rows for  $\mathbf{Z}$  with  $\lambda = 0.5$ ) on the Cora dataset in Figure 2. Even without any access to label information for



(a) Graphite-AE



(b) Graphite-VAE

Figure 2. t-SNE embeddings of the latent feature vectors for the Cora dataset. Colors denote labels.

the nodes during training, the name models are able to cluster the nodes (papers) as per their labels (paper categories).

#### 4.3. Semi-supervised node classification

Given labels for a subset of nodes in an underlying graph, the goal of this task is to predict the labels for the remaining nodes. We consider a *transductive* setting, where we have access to the test nodes (without their labels) during training.

Closest approach to Graphite for this task is a supervised graph convolutional network (GCN) trained end-to-end. We consider an extension of this baseline, wherein we augment the GCN objective with the Graphite objective and a hyperparameter to control the relative importance of the two terms in the combined objective. The parameters  $\phi$  for the encoder are shared across these two objectives, with an additional GCN layer for mapping the encoder output to softmax probabilities over the requisite number of classes. All parameters are learned jointly.

Table 3. Area Under the ROC Curve (AUC) for link prediction (\* denotes dataset with features). Higher is better.

	Cora	Citeseer	Pubmed	Cora*	Citeseer*	Pubmed*
SC	89.9 ± 0.20	91.5 ± 0.17	<b>94.9</b> ± 0.04	-	-	-
DeepWalk	85.0 ± 0.17	88.6 ± 0.15	91.5 ± 0.04	-	-	-
node2vec	85.6 ± 0.15	89.4 ± 0.14	91.9 ± 0.04	-	-	-
GAE	90.2 ± 0.16	92.0 ± 0.14	92.5 ± 0.06	93.9 ± 0.11	94.9 ± 0.13	96.8 ± 0.04
VGAE	90.1 ± 0.15	92.0 ± 0.17	92.3 ± 0.06	94.1 ± 0.11	96.7 ± 0.08	95.5 ± 0.13
Graphite-AE	91.0 ± 0.15	92.6 ± 0.16	94.5 ± 0.05	94.2 ± 0.13	96.2 ± 0.10	<b>97.8</b> ± 0.03
Graphite-VAE	<b>91.5</b> ± 0.15	<b>93.5</b> ± 0.13	94.6 ± 0.04	<b>94.7</b> ± 0.11	<b>97.3</b> ± 0.06	97.4 ± 0.04

Table 4. Average Precision (AP) scores for link prediction (\* denotes dataset with features). Higher is better.

	Cora	Citeseer	Pubmed	Cora*	Citeseer*	Pubmed*
SC	92.8 ± 0.12	94.4 ± 0.11	<b>96.0</b> ± 0.03	-	-	-
DeepWalk	86.6 ± 0.17	90.3 ± 0.12	91.9 ± 0.05	-	-	-
node2vec	87.5 ± 0.14	91.3 ± 0.13	92.3 ± 0.05	-	-	-
GAE	92.4 ± 0.12	94.0 ± 0.12	94.3 ± 0.5	94.3 ± 0.12	94.8 ± 0.15	96.8 ± 0.04
VGAE	92.3 ± 0.12	94.2 ± 0.12	94.2 ± 0.04	94.6 ± 0.11	97.0 ± 0.08	95.5 ± 0.12
Graphite-AE	92.8 ± 0.13	94.1 ± 0.14	95.7 ± 0.06	94.5 ± 0.14	96.1 ± 0.12	<b>97.7</b> ± 0.03
Graphite-VAE	<b>93.2</b> ± 0.13	<b>95.0</b> ± 0.10	<b>96.0</b> ± 0.03	<b>94.9</b> ± 0.13	<b>97.4</b> ± 0.06	97.4 ± 0.04

Table 5. Classification accuracies (\* denotes dataset with features). Baseline numbers from Kipf &amp; Welling (2017).

	Cora*	Citeseer*	Pubmed*
SemiEmb	59.0	59.6	71.1
DeepWalk	67.2	43.2	65.3
ICA	75.1	69.1	73.9
Planetoid	75.7	64.7	77.2
GCN	81.5	70.3	79.0
Graphite	<b>82.1</b> ± 0.06	<b>71.0</b> ± 0.07	<b>79.3</b> ± 0.03

**Results.** The classification accuracy of the semi-supervised models is given in Table 5. We find that Graphite-hybrid outperforms the competing models on all datasets and in particular the GCN approach which is the closest baseline. Recent work in Graph Attention Networks shows that extending GCN by incorporating attention can boost performance on this task (Veličković et al., 2018). Using GATs in place of GCNs for parameterizing Graphite could yield similar performance boost in future work.

## 5. Theoretical Analysis

In this section, we derive a theoretical connection between message passing in graph neural networks and approximate inference in related undirected graphical models.

### 5.1. Kernel embeddings

We first provide a brief background on kernel embeddings. A kernel defines a notion of similarity between pairs of objects (Schölkopf & Smola, 2002; Shawe-Taylor & Cristianini, 2004). Let  $K : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$  be the kernel function defined over a space of objects, say  $\mathcal{Z}$ . With every kernel function  $K$ , we have an associated feature map  $\psi : \mathcal{Z} \rightarrow \mathcal{H}$  where  $\mathcal{H}$  is a potentially infinite dimensional feature space.

Kernel methods can be used to specify embeddings of *distributions* of arbitrary objects (Smola et al., 2007; Gretton et al., 2007). Formally, we denote these functional mappings as  $T_\psi : \mathcal{P} \rightarrow \mathcal{H}$  where  $\mathcal{P}$  specifies the space of all distributions on  $\mathcal{Z}$ . These mappings, referred to as kernel embeddings of distributions, are defined as:

$$T_\psi(p) := \mathbb{E}_{Z \sim p}[\psi(Z)] \quad (11)$$

for any  $p \in \mathcal{P}$ . We are particularly interested in kernels with feature maps  $\psi$  that define injective embeddings, *i.e.*, for any pair of distributions  $p_1$  and  $p_2$ , we have  $T_\psi(p_1) \neq T_\psi(p_2)$  if  $p_1 \neq p_2$ . For injective embeddings, we can compute functionals of any distribution by directly applying a corresponding function on its embedding. Formally, for every function  $\mathcal{O} : \mathcal{P} \rightarrow \mathbb{R}^d$ ,  $d \in \mathbb{N}$  and injective embedding  $T_\psi$ , there exists a function  $\tilde{\mathcal{O}}_\psi : \mathcal{H} \rightarrow \mathbb{R}^d$  such that:

$$\mathcal{O}(p) = \tilde{\mathcal{O}}_\psi(T_\psi(p)) \quad \forall p \in \mathcal{P}. \quad (12)$$

Informally, we can see that the operator  $\tilde{\mathcal{O}}_\psi$  can be defined as the composition of  $\mathcal{O}$  with the inverse of  $T_\psi$ .





Figure 3. Interpreting message passing in Graph Neural Networks via Kernel Embeddings and Mean-field inference

## 5.2. Connections with mean-field inference

Locality preference for representational learning is a key inductive bias for graphs. We formulate this using an (undirected) graphical model  $\mathcal{G}$  over  $\mathbf{X}$ ,  $\mathbf{A}$ , and  $\{\mathbf{Z}_1, \dots, \mathbf{Z}_n\}$ . As in a GNN, we assume that  $\mathbf{X}$  and  $\mathbf{A}$  are observed and specify conditional independence structure in a conditional distribution over the latent variables, denoted as  $r(\mathbf{Z}_1, \dots, \mathbf{Z}_n | \mathbf{A}, \mathbf{X})$ . We are particularly interested in models that satisfy the following property.

**Property 1.** *The edge set  $E$  defined by the adjacency matrix  $\mathbf{A}$  is an undirected 1-map for the distribution  $r(\mathbf{Z}_1, \dots, \mathbf{Z}_n | \mathbf{A}, \mathbf{X})$ .*

In words, the above property implies that according to the conditional distribution over  $\mathbf{Z}$ , any individual  $\mathbf{Z}_i$  is independent of all other  $\mathbf{Z}_j$  when conditioned on  $\mathbf{A}$ ,  $\mathbf{X}$ , and the neighboring latent variables of node  $i$  as determined by the edge set  $E$ . See Figure 3 for an illustration.

A mean-field (MF) approximation for  $\mathcal{G}$  approximates the conditional distribution  $r(\mathbf{Z}_1, \dots, \mathbf{Z}_n | \mathbf{A}, \mathbf{X})$  as:

$$r(\mathbf{Z}_1, \dots, \mathbf{Z}_n | \mathbf{A}, \mathbf{X}) \approx \prod_{i=1}^n q_{\phi_i}(\mathbf{Z}_i | \mathbf{A}, \mathbf{X}) \quad (13)$$

where  $\phi_i$  denotes the set of parameters for the  $i$ -th variational marginal. These parameters are optimized by minimizing the KL-divergence between the variational and the true conditional distributions:

$$\min_{\phi_1, \dots, \phi_n} \text{KL} \left( \prod_{i=1}^n q_{\phi_i}(\mathbf{Z}_i | \mathbf{A}, \mathbf{X}) \parallel r(\mathbf{Z}_1, \dots, \mathbf{Z}_n | \mathbf{A}, \mathbf{X}) \right) \quad (14)$$

Using standard variational arguments (Wainwright et al., 2008), we know that the optimal variational marginals assume the following functional form:

$$q_{\phi_i}(\mathbf{Z}_i | \mathbf{A}, \mathbf{X}) = \mathcal{O}_{\mathcal{G}}^{MF}(\mathbf{Z}_i, \{q_{\phi_j}\}_{j \in \mathcal{N}(i)}) \quad (15)$$

where  $\mathcal{N}(i)$  denotes the neighbors of  $\mathbf{Z}_i$  in  $\mathcal{G}$  and  $\mathcal{O}$  is a function determined by the fixed point equations that depends on the potentials associated with  $\mathcal{G}$ . Importantly, the above functional form suggests that the optimal marginals in mean field inference are locally consistent that they are only a function of the neighboring marginals. An iterative algorithm for mean-field inference is to perform message passing over the underlying graph until convergence. With an appropriate initialization at  $l = 0$ , the updated marginals at iteration  $l \in \mathbb{N}$  are given as:

$$q_{\phi_i}^{(l)}(\mathbf{Z}_i | \mathbf{A}, \mathbf{X}) = \mathcal{O}_{\mathcal{G}}^{MF}(\mathbf{Z}_i, \{q_{\phi_j}^{(l-1)}\}_{j \in \mathcal{N}(i)}) \quad (16)$$

We will sidestep deriving  $\mathcal{O}$ , and instead use the kernel embeddings of the variational marginals to directly reason in the embedding space. That is, we assume we have an injective embedding for each marginal  $q_{\phi_i}$  given by  $\boldsymbol{\mu}_i = \mathbb{E}_{\mathbf{Z}_i \sim q_{\phi_i}}[\psi(\mathbf{Z}_i)]$  for some feature map  $\psi: \mathbb{R}^k \rightarrow \mathbb{R}^{k'}$  and directly use the equivalence established in Eq. (12) iteratively. For mean-field inference via message passing as in Eq. (16), this gives us the following recursive expression for iteratively updating the embeddings at iteration  $l \in \mathbb{N}$ :

$$\boldsymbol{\mu}_i^{(l)} = \tilde{\mathcal{O}}_{\psi, \mathcal{G}}^{MF}(\{\boldsymbol{\mu}_j^{(l-1)}\}_{j \in \mathcal{N}(i)}) \quad (17)$$

with an appropriate base case for  $\boldsymbol{\mu}_i^{(0)}$ . We then have the following result:

**Theorem 2.** *Let  $\mathcal{G}$  be any undirected latent variable model such that the conditional distribution  $r(\mathbf{Z}_1, \dots, \mathbf{Z}_n | \mathbf{A}, \mathbf{X})$  expressed by the model satisfies Property 1.*

*Then there exists a choice of  $\eta_l$ ,  $\mathcal{F}_l$ ,  $\mathbf{W}_l$ , and  $\mathbf{B}_l$  such that for all  $\{\boldsymbol{\mu}_i^{(l-1)}\}_{i=1}^n$ , the GNN propagation rule in Eq. (2) is computationally equivalent to updating  $\{\boldsymbol{\mu}_i^{(l-1)}\}_{i=1}^n$  via a first order approximation of Eq. (17).*

*Proof.* See Appendix A. □

While  $\eta_l$  and  $\mathcal{F}_l$  are typically fixed beforehand, the parameters  $\mathbf{W}_l$ , and  $\mathbf{B}_l$  are directly learned from data in practice. Hence we have shown that a GNN is a good model for computation with respect to latent variable models that attempt to capture inductive biases relevant to graphs, *i.e.*, ones where the latent feature vector for every node is conditionally independent from everything else given the feature vectors of its neighbors (and  $\mathbf{A}$ ,  $\mathbf{X}$ ). Note that such a graphical model would satisfy Property 1 but is in general different

from the posterior specified by the one in Figure 1. However if the true (but unknown) posterior on the latent variables for the model proposed in Figure 1 could be expressed as an equivalent model satisfying the desired property, then Theorem 2 indeed suggests the use of GNNs for parameterizing variational posteriors, as we do so in the case of Graphite.

## 6. Discussion & Related Work

Our framework effectively marries *probabilistic modeling* and *representation learning* on graphs. We review some of the dominant prior works in these fields below.

**Probabilistic modeling of graphs.** The earliest *probabilistic models of graphs* proposed to generate graphs by creating an edge between any pair of nodes with a constant probability (Erdős & Rényi, 1959). Several alternatives have been proposed since; *e.g.*, the small-world model generates graphs that exhibit local clustering (Watts & Strogatz, 1998), the Barabasi-Albert models preferential attachment wherein high-degree nodes are likely to form edges with newly added nodes (Barabasi & Albert, 1999), the stochastic block model is based on inter and intra community linkages (Holland et al., 1983) etc. We direct the interested reader to prominent surveys on this topic (Newman, 2003; Mitzenmacher, 2004; Chakrabarti & Faloutsos, 2006).

**Representation learning on graphs.** For *representation learning on graphs*, there are broadly three kinds of approaches: matrix factorization, random walk based approaches, and graph neural networks. We include a brief discussion on the first two kinds in Appendix C and refer the reader to Hamilton et al. (2017b) for a recent survey.

Graph neural networks, a collective term for networks that operate over graphs using message passing, have shown success on several downstream applications, *e.g.*, (Duvenaud et al., 2015; Li et al., 2016; Kearnes et al., 2016; Kipf & Welling, 2017; Hamilton et al., 2017a) and the references therein. Gilmer et al. (2017) provides a comprehensive characterization of these networks in the message passing setup. We used Graph Convolution Networks, partly to provide a direct comparison with GAE/VGAE and leave the exploration of other GNN variants for future work.

**Latent variable models for graphs.** Hierarchical Bayesian models parameterized by deep neural networks have been recently proposed for graphs (Hu et al., 2017; Wang et al., 2017). Besides being restricted to single graphs, these models are limited since inference requires running expensive Markov chains (Hu et al., 2017) or are task-specific (Wang et al., 2017). Johnson (2017) and Kipf et al. (2018) generate graphs as latent representations learned directly from data. In contrast, we are interested

in modeling observed (and not latent) relational structure. Finally, there has been a fair share of recent work for generation of special kinds of graphs, such as parsed trees of source code (Maddison & Tarlow, 2014) and SMILES representations for molecules (Olivecrona et al., 2017).

Several deep generative models for graphs have recently been proposed. Amongst adversarial generation approaches, Wang et al. (2018) and Bojchevski et al. (2018) model local graph neighborhoods and random walks on graphs respectively. Li et al. (2018) and You et al. (2018) model graphs as sequences and generate graphs via autoregressive procedures. Adversarial and autoregressive approaches are successful at generating graphs, but do not directly allow for inferring latent variables via encoders. Latent variable generative models have also been proposed for generating small molecular graphs (Jin et al., 2018; Samanta et al., 2018; Simonovsky & Komodakis, 2018). These methods involve an expensive decoding procedure that limits scaling to large graphs. Finally, closest to our framework is the GAE/VGAE approach (Kipf & Welling, 2016) discussed in Section 4. Pan et al. (2018) extends this approach with an adversarial regularization framework but retain the inner product decoder. Our work proposes a novel multi-step decoding mechanism based on graph refinement.

## 7. Conclusion & Future Work

We proposed Graphite, a scalable deep generative model for graphs based on variational autoencoding. The encoders and decoders in Graphite are parameterized by graph neural networks that propagate information locally on a graph. Our proposed decoder performs a multi-layer iterative decoding comprising of alternate inner product operations and message passing on the intermediate graph.

Current generative models for graphs are not permutation-invariant and are learned by feeding graphs with a fixed or heuristic ordering of nodes. This is an exciting challenge for future work, which could potentially be resolved by incorporate graph representations robust to permutation invariances (Verma & Zhang, 2017) or modeling distributions over permutations of node orderings via recent approaches such as NeuralSort (Grover et al., 2019). Extending Graphite for modeling richer graphical structure such as heterogeneous and time-varying graphs, as well as integrating domain knowledge within Graphite decoders for applications in generative design and synthesis *e.g.*, molecules, programs, and parse trees is another interesting future direction.

Finally, our theoretical results in Section 5 suggest that a principled design of layerwise propagation rules in graph neural networks inspired by additional message passing inference schemes (Dai et al., 2016; Gilmer et al., 2017) is another avenue for future research.



## Acknowledgements

This research has been supported by Siemens, a Future of Life Institute grant, NSF grants (#1651565, #1522054, #1733686), ONR (N00014-19-1-2145), AFOSR (FA9550-19-1-0024), and an Amazon AWS Machine Learning Grant. AG is supported by a Microsoft Research Ph.D. fellowship and a Stanford Data Science Scholarship. We would like to thank Daniel Levy for helpful comments on early drafts.

## References

- Barabasi, A.-L. and Albert, R. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- Belkin, M. and Niyogi, P. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems*, 2002.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Bojchevski, A., Shchur, O., Zügner, D., and Günnemann, S. NetGAN: Generating graphs via random walks. In *International Conference on Machine Learning*, 2018.
- Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. Spectral networks and locally connected networks on graphs. In *International Conference on Learning Representations*, 2013.
- Chakrabarti, D. and Faloutsos, C. Graph mining: Laws, generators, and algorithms. *ACM Computing Surveys*, 38(1):2, 2006.
- Dai, H., Dai, B., and Song, L. Discriminative embeddings of latent variable models for structured data. In *International Conference on Machine Learning*, 2016.
- Douglas, B. L. The Weisfeiler-Lehman method and graph isomorphism testing. *arXiv preprint arXiv:1101.5211*, 2011.
- Duvenaud, D., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*, 2015.
- Erdős, P. and Rényi, A. On random graphs. *Publicationes Mathematicae (Debrecen)*, 6:290–297, 1959.
- Garcia, V. and Bruna, J. Few-shot learning with graph neural networks. In *International Conference on Learning Representations*, 2018.
- Gilmer, J., Schoenholz, S., Riley, P., Vinyals, O., and Dahl, G. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, 2017.
- Gori, M., Monfardini, G., and Scarselli, F. A new model for learning in graph domains. In *International Joint Conference on Neural Networks*, 2005.
- Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., and Smola, A. A kernel method for the two-sample-problem. In *Advances in Neural Information Processing Systems*, 2007.
- Grover, A. and Leskovec, J. node2vec: Scalable feature learning for networks. In *Knowledge Discovery and Data Mining*, 2016.
- Grover, A., Wang, E., Zweig, A., and Ermon, S. Stochastic optimization of sorting networks via continuous relaxations. In *International Conference on Learning Representations*, 2019.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, 2017a.
- Hamilton, W. L., Ying, R., and Leskovec, J. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584*, 2017b.
- Holland, P. W., Laskey, K. B., and Leinhardt, S. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.
- Hu, C., Rai, P., and Carin, L. Deep generative models for relational data with side information. In *International Conference on Machine Learning*, 2017.
- Jin, W., Barzilay, R., and Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning*, 2018.
- Johnson, D. Learning graphical state transitions. In *International Conference on Learning Representations*, 2017.
- Kearnes, S., McCloskey, K., Berndl, M., Pande, V., and Riley, P. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8):595–608, 2016.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, 2014.

- Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. Neural relational inference for interacting systems. In *International Conference on Machine Learning*, 2018.
- Kipf, T. N. and Welling, M. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- Li, Y., Tarlow, D., Brockschmidt, M., and Zemel, R. Gated graph sequence neural networks. In *International Conference on Learning Representations*, 2016.
- Li, Y., Vinyals, O., Dyer, C., Pascanu, R., and Battaglia, P. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.
- Loehlin, J. C. *Latent variable models: An introduction to factor, path, and structural analysis*. Lawrence Erlbaum Associates Publishers, 1998.
- Lu, Q. and Getoor, L. Link-based classification. In *International Conference on Machine Learning*, 2003.
- Maaten, L. v. d. and Hinton, G. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov): 2579–2605, 2008.
- Maddison, C. and Tarlow, D. Structured generative models of natural source code. In *International Conference on Machine Learning*, 2014.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, 2013.
- Mitzenmacher, M. A brief history of generative models for power law and lognormal distributions. *Internet mathematics*, 1(2):226–251, 2004.
- Newman, M. E. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.
- Olivecrona, M., Blaschke, T., Engkvist, O., and Chen, H. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9(1):48, 2017.
- Pan, S., Hu, R., Long, G., Jiang, J., Yao, L., and Zhang, C. Adversarially regularized graph autoencoder for graph embedding. In *International Joint Conference on Artificial Intelligence*, 2018.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(Oct):2825–2830, 2011.
- Perozzi, B., Al-Rfou, R., and Skiena, S. Deepwalk: On-line learning of social representations. In *Knowledge Discovery and Data Mining*, 2014.
- Samanta, B., De, A., Ganguly, N., and Gomez-Rodriguez, M. Designing random graph models using variational autoencoders with applications to chemical design. *arXiv preprint arXiv:1802.05283*, 2018.
- Saxena, A., Gupta, A., and Mukerjee, A. Non-linear dimensionality reduction by locally linear isomaps. In *Advances in Neural Information Processing Systems*, 2004.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- Schölkopf, B. and Smola, A. J. *Learning with kernels: Support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- Sen, P., Namata, G. M., Bilgic, M., Getoor, L., Gallagher, B., and Eliassi-Rad, T. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.
- Shawe-Taylor, J. and Cristianini, N. *Kernel methods for pattern analysis*. Cambridge university press, 2004.
- Shervashidze, N., Schweitzer, P., Leeuwen, E. J. v., Mehlhorn, K., and Borgwardt, K. M. Weisfeiler-Lehman graph kernels. *Journal of Machine Learning Research*, 12(Sep):2539–2561, 2011.
- Simonovsky, M. and Komodakis, N. GraphVAE: Towards generation of small graphs using variational autoencoders. *arXiv preprint arXiv:1802.03480*, 2018.
- Smola, A., Gretton, A., Song, L., and Schölkopf, B. A hilbert space embedding for distributions. In *International Conference on Algorithmic Learning Theory*, 2007.
- Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, 2015.
- Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., and Mei, Q. Line: Large-scale information network embedding. In *International World Wide Web Conference*, 2015.
- Tang, L. and Liu, H. Leveraging social media networks for classification. *Data Mining and Knowledge Discovery*, 23(3):447–478, 2011.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *International Conference on Learning Representations*, 2018.

- Verma, S. and Zhang, Z.-L. Hunt for the unique, stable, sparse and fast feature learning on graphs. In *Advances in Neural Information Processing Systems*, 2017.
- Wainwright, M. J., Jordan, M. I., et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- Wang, H., Shi, X., and Yeung, D.-Y. Relational deep learning: A deep latent variable model for link prediction. In *AAAI Conference on Artificial Intelligence*, 2017.
- Wang, H., Wang, J., Wang, J., Zhao, M., Zhang, W., Zhang, F., Xie, X., and Guo, M. GraphGAN: Graph representation learning with generative adversarial nets. In *International World Wide Web Conference*, 2018.
- Watts, D. J. and Strogatz, S. H. Collective dynamics of small-world networks. *Nature*, 393(6684):440, 1998.
- Weisfeiler, B. and Lehman, A. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsia*, 2(9): 12–16, 1968.
- Weston, J., Ratle, F., and Collobert, R. Deep learning via semi-supervised embedding. In *International Conference on Machine Learning*, 2008.
- Yang, Z., Cohen, W. W., and Salakhutdinov, R. Revisiting semi-supervised learning with graph embeddings. *arXiv preprint arXiv:1603.08861*, 2016.
- Yoon, K., Liao, R., Xiong, Y., Zhang, L., Fetaya, E., Urtasun, R., Zemel, R., and Pitkow, X. Inference in probabilistic graphical models by graph neural networks. *arXiv preprint arXiv:1803.07710*, 2018.
- You, J., Ying, R., Ren, X., Hamilton, W. L., and Leskovec, J. GraphRNN: A deep generative model for graphs. In *International Conference on Machine Learning*, 2018.