

GraphMix: Improved Training of GNNs for Semi-Supervised Learning

Vikas Verma^{1,2}, Meng Qu¹, Kenji Kawaguchi³, Alex Lamb¹, Yoshua Bengio¹, Juho Kannala², Jian Tang¹

¹ Mila - Québec Artificial Intelligence Institute, Montréal, Canada, ² Aalto University, Finland, ³ Massachusetts Institute of Technology (MIT), USA

Abstract

We present *GraphMix*, a regularization method for Graph Neural Network based semi-supervised object classification, whereby we propose to train a fully-connected network jointly with the graph neural network via parameter sharing and interpolation-based regularization. Further, we provide a theoretical analysis of how GraphMix improves the generalization bounds of the underlying graph neural network, without making any assumptions about the “aggregation” layer or the depth of the graph neural networks. We experimentally validate this analysis by applying GraphMix to various architectures such as Graph Convolutional Networks, Graph Attention Networks and Graph-U-Net. Despite its simplicity, we demonstrate that GraphMix can consistently improve or closely match state-of-the-art performance using even simpler architectures such as Graph Convolutional Networks, across three established graph benchmarks: Cora, Citeseer and Pubmed citation network datasets, as well as three newly proposed datasets: Cora-Full, Co-author-CS and Co-author-Physics.

1 Introduction

Due to the presence of graph-structured data across a wide variety of domains, such as biological networks, citation networks and social networks, there have been several attempts to design neural networks, known as graph neural networks (GNN), that can process arbitrarily structured graphs. Early work includes (Gori, Monfardini, and Scarselli 2005; Scarselli et al. 2009) which propose a neural network that can directly process most types of graphs e.g., acyclic, cyclic, directed, and undirected graphs. More recent approaches include (Bruna et al. 2013; Henaff, Bruna, and LeCun 2015; Defferrard, Bresson, and Vandergheynst 2016; Kipf and Welling 2016; Gilmer et al. 2017; Hamilton, Ying, and Leskovec 2017; Veličković et al. 2018, 2019; Qu, Bengio, and Tang 2019; Gao and Ji 2019; Ma et al. 2019), among others. Many of these approaches are designed for addressing the problem of semi-supervised learning over graph-structured data (Zhou et al. 2018). Much of these research efforts have been dedicated to developing novel architectures.

Here we instead propose an architecture-agnostic method for regularized training of GNNs for semi-supervised

node classification. Recently, regularization based on data-augmentation has been shown to be very effective in other types of neural networks but how to apply these techniques in GNNs is still under-explored. Our proposed method GraphMix¹ is a unified framework that draws inspiration from interpolation based data augmentation (Zhang et al. 2018; Verma et al. 2019a) and self-training based data-augmentation (Laine and Aila 2016; Tarvainen and Valpola 2017; Verma et al. 2019b; Berthelot et al. 2019). We show that with our proposed method, we can achieve state-of-the-art performance even when using simpler GNN architectures such as Graph Convolutional Networks (Kipf and Welling 2017), with no additional memory cost and with minimal additional computation cost. Further, we conduct a theoretical analysis to demonstrate the effectiveness of the proposed method over the underlying GNNs.

2 Problem Definition and Preliminaries

Problem Setup: We are interested in the problem of semi-supervised node and edge classification using graph-structured data. We can formally define such graph-structured data as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{A}, \mathcal{X})$, where $\mathcal{V} = \mathcal{V}_l \cup \mathcal{V}_u$ is the union of labeled (\mathcal{V}_l) and unlabeled (\mathcal{V}_u) nodes in the graph with cardinalities n_l and n_u , and \mathcal{A} is the adjacency matrix representing the edges between the nodes of \mathcal{V} , $\mathcal{X} \in \mathbb{R}^{(n_l+n_u) \times d}$ is the input node features. Each node v belongs to one out of C classes and can be labeled with a C -dimensional one-hot vector $y_v \in \mathbb{R}^C$. Given the labels $Y_l \in \mathbb{R}^{n_l \times C}$ of the labeled nodes \mathcal{V}_l , the task is to predict the labels $Y_u \in \mathbb{R}^{n_u \times C}$ of the unlabeled nodes \mathcal{V}_u .

Graph Neural Networks: Graph Neural Networks (GNN) learn the l_{th} layer representations of a sample i by leveraging the representations of the samples $NB(i)$ in the neighbourhood of i . This is done by using an aggregation function that takes as an input the representations of all the samples along with the graph structure and outputs the aggregated representation. The aggregation function can be defined using the Graph Convolution layer (Kipf and Welling 2017),

¹Code available at <https://github.com/vikasverma1077/GraphMix>
Appendix available at <https://arxiv.org/abs/1909.11715>
Correspondence: vikasverma.iitm@gmail.com

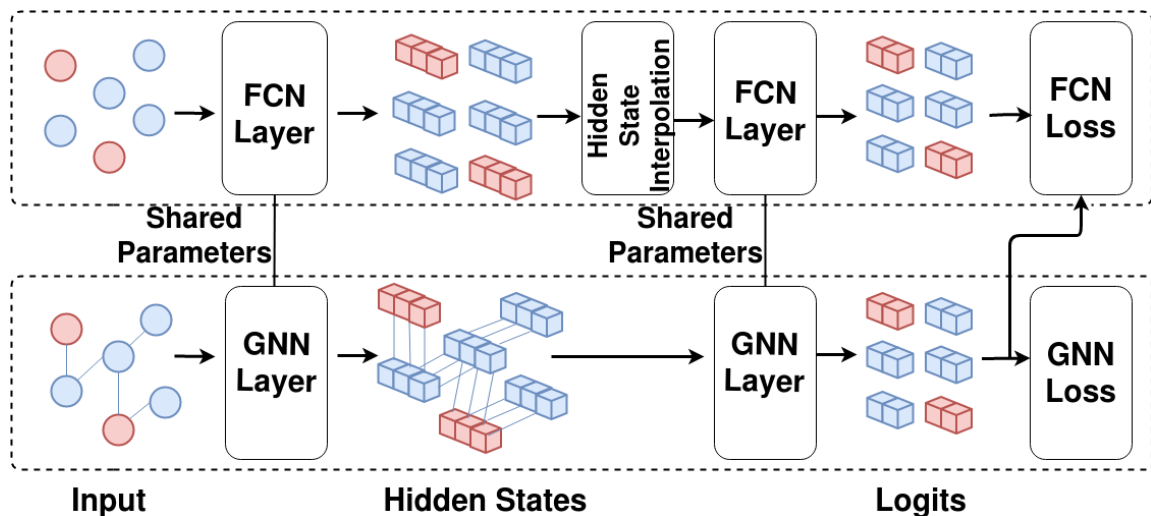


Figure 1: The procedure for training with GraphMix . The labeled and unlabeled nodes are shown with different colors in the graph. GraphMix augments the training of a baseline Graph Neural Network (GNN) with a Fully-Connected Network (FCN). The FCN is trained by interpolating the hidden states and the corresponding labels. This leads to better features which are transferred to the GNN via sharing the linear transformation parameters W (in Equation 1) of the GNN and FCN layers. Furthermore, the predictions made by the GNN for unlabeled data are used to augment the input data for the FCN. The FCN and the GNN losses are minimized jointly by alternate minimization.

Graph Attention Layer (Veličković et al. 2018), or any general message passing layer (Gilmer et al. 2017). Formally, let $\mathbf{h}^{(l)} \in \mathbb{R}^{n \times k}$ be a matrix containing the k -dimensional representation of n nodes in the l_{th} layer, then:

$$\mathbf{h}^{(l+1)} = \sigma(\text{AGGREGATE}(\mathbf{h}^{(l)} \mathbf{W}, \mathcal{A})) \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{k \times k'}$ is a linear transformation matrix, k' is the dimension of $(l+1)_{th}$ layer, AGGREGATE is the aggregation function that utilizes the graph adjacency matrix \mathcal{A} to aggregate the hidden representations of neighbouring nodes and σ is a non-linear activation function, e.g. ReLU.

Interpolation Based Regularization Techniques: Recently, interpolation-based techniques have been proposed for regularizing neural networks. We briefly describe some of these techniques here. Mixup (Zhang et al. 2018) trains a neural network on the convex combination of input and targets, whereas Manifold Mixup (Verma et al. 2019a) trains a neural network on the convex combination of the hidden states of a randomly chosen hidden layer and the targets. While Mixup regularizes a neural network by enforcing that the model output should change linearly in between the examples in the input space, Manifold Mixup regularizes the neural network by learning better (more discriminative) hidden states.

Formally, suppose $T_\theta(\mathbf{x}) = (f \circ g)_\theta(\mathbf{x})$ is a neural network parametrized with θ such that $g : \mathbf{x} \rightarrow \mathbf{h}$ is a function that maps input sample to hidden states, $f : \mathbf{h} \rightarrow \hat{\mathbf{y}}$ is a function that maps hidden states to predicted output, λ is a random variable drawn from $\text{Beta}(\alpha, \alpha)$ distribution, $\text{Mix}_\lambda(\mathbf{a}, \mathbf{b}) = \lambda * \mathbf{a} + (1 - \lambda) * \mathbf{b}$ is an interpolation function, \mathcal{D} is the data distribution, (\mathbf{x}, \mathbf{y}) and $(\mathbf{x}', \mathbf{y}')$ is a pair of labeled examples sampled from distribution \mathcal{D} and ℓ be a loss

function such as cross-entropy loss, then the Manifold Mixup Loss is defined as:

$$\mathcal{L}_{\text{MM}}(\mathcal{D}, T_\theta, \alpha) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \mathbb{E}_{(\mathbf{x}', \mathbf{y}') \sim \mathcal{D}} \mathbb{E}_{\lambda \sim \text{Beta}(\alpha, \alpha)} \ell(f(\text{Mix}_\lambda(g(\mathbf{x}), g(\mathbf{x}'))), \text{Mix}_\lambda(\mathbf{y}, \mathbf{y}')). \quad (2)$$

We use above Manifold Mixup loss for training an auxiliary Fully-connected-network as described in Section 3.

3 GraphMix

3.1 Motivation

Data Augmentation is one of the simplest and most efficient technique for regularizing a neural network. In the domains of computer vision, speech and natural language, there exist efficient data augmentation techniques, for example, random cropping, translation or Cutout (Devries and Taylor 2017) for computer vision, (Ko et al. 2015) and (Park et al. 2019) for speech and (Xie et al. 2017) for natural language. However, data augmentation for the graph-structured data remains under-explored. There exists some recent work along these lines but the prohibitive computation cost (see Section 5) introduced by these methods make them impractical for real-world large graph datasets. Based on these limitations, our main objective is to propose an efficient data augmentation technique for graph datasets.

Recent work based on interpolation-based data augmentation (Zhang et al. 2018; Verma et al. 2019a) has seen sizable improvements in regularization performance across a number of tasks. However, these techniques are not directly applicable to graphs for an important reason: *Although we can create additional nodes by interpolating the features and corresponding labels, it remains unclear how these new nodes*

must be connected to the original nodes via synthetic edges such that the structure of the whole graph is preserved. To alleviate this issue, we propose to train an auxiliary Fully-Connected Network (FCN) using Manifold Mixup as discussed in Section 3.2. Note that the FCN only uses the node features (not the graph structure), thus the Manifold mixup loss in Eq. 2 can be directly used for training the FCN.

Interpolation based data-augmentation techniques have an added advantage for training GNNs: A vanilla GNN learns the representation of each node by iteratively aggregating information from the neighbors of that node (Equation 1). However, this induces the problem of *oversmoothing* (Li, Han, and Wu 2018; Xu et al. 2018) while training GNNs with many layers. Due to this limitation, GNNs are trained only with a few layers, and thus they can only leverage the local neighbourhood of each node for learning its representations, without leveraging the representations of the nodes which are multiple hops away in the graph. This limitation can be addressed using the interpolation-based method such as Manifold Mixup: in Manifold Mixup, the representations of a randomly chosen pair of nodes is used to facilitate better representation learning; it is possible that the randomly chosen pair of nodes will not be in the local neighbourhood of each other. Based on these challenges and motivations we present our proposed approach GraphMix for training GNNs in the following Section.

3.2 Method

We first describe GraphMix at a high-level and then give a more formal description. GraphMix augments the vanilla GNN with a Fully-Connected Network (FCN). The FCN loss is computed using *Manifold Mixup* as discussed below and the GNN loss is computed normally. *Manifold Mixup* training of FCN facilitates learning more discriminative node representations (Verma et al. 2019a). An important question is how these more discriminative node representations can be transferred to the GNN? One potential approach could involve maximizing the mutual information between the hidden states of the FCN and the GNN using formulations similar to those proposed by (Hjelm et al. 2019; Sun et al. 2020). However, this requires optimizing additional network parameters. Instead, we propose parameter sharing between FCN and GNN to facilitate the transfer of discriminative node representations from the FCN to the GNN. It is a viable option because as mentioned in Eq 1, a GNN layer typically performs an additional operation (*AGGREGATE*) on the linear transformations of node representations (which are essentially pre-activation representations of the FCN layer). Using the more discriminative representations of the nodes from FCN, as well as the graph structure, the GNN loss is computed in the usual way to further refine the node representations. In this way we can exploit the improved representations from *Manifold Mixup* for training GNNs.

In Section 3.3, without making any assumption about the aggregation function and the depth of the graph neural network, we show that GraphMix improves the generalization of the underlying graph neural network. This makes GraphMix applicable to various kind of architectures having different aggregation functions, such as weighed averaging in GCN

(Kipf and Welling 2016), attention based aggregation in GAT (Veličković et al. 2018) and graph-pooling/unpooling operations in Graph U-Nets (Gao and Ji 2019). In the aforementioned sense, GraphMix procedure is highly flexible: it can be applied to any underlying GNN as long as the underlying GNN applies parametric transformations to the node features.

Additionally, we propose to use the predicted targets from the GNN to augment the training set of the FCN. In this way, both the FCN and the GNN facilitate each other’s learning process. Both the FCN loss and the GNN loss are optimized in an alternating fashion during training. At inference time, predictions are made using only the GNN.

A diagram illustrating GraphMix is presented in Figure 1 and the full algorithm is presented in Appendix A.3. Further, we draw similarities and difference of GraphMix w.r.t. Co-training framework in the Appendix A.2.

So far we have presented the general design of GraphMix, now we present GraphMix more formally. Given a graph \mathcal{G} , let $(\mathbf{X}_l, \mathbf{Y}_l)$ be the input features and the labels of the labeled nodes \mathcal{V}_l and let (\mathbf{X}_u) be the input features of the unlabeled nodes \mathcal{V}_u . Let F_θ and G_θ be a FCN and a GNN respectively, which share the parameters θ . The FCN loss from the labeled data is computed using Eq. 2 as follows:

$$\mathcal{L}^{\text{supervised}} = \mathcal{L}_{\text{MM}}((\mathbf{X}_l, \mathbf{Y}_l), F_\theta, \alpha) \quad (3)$$

For unlabeled nodes \mathcal{V}_u , we compute the prediction $\hat{\mathbf{Y}}_u$ using the GNN:

$$\hat{\mathbf{Y}}_u = G_\theta(\mathbf{X}_u) \quad (4)$$

We note that recent state-of-the-art semi-supervised learning methods use a *teacher* model to accurately predict targets for the unlabeled data. The teacher model can be realized as a temporal ensemble of the *student* model (the model being trained) (Laine and Aila 2016) or by using an Exponential Moving Average (EMA) of the parameters of the student model (Tarvainen and Valpola 2017; Verma et al. 2019b). Different from these approaches, we use the GNN as a teacher model for predicting the targets for the FCN. This is due to the fact that GNNs leverage graph structure, which in practice, allows them to make more accurate predictions than the temporal ensemble or EMA ensemble of FCN (although there is no theoretical guarantee for this).

Moreover, to improve the accuracy of the predicted targets in Eq 4, we applied the average of the model prediction on K random perturbations of an input sample along with sharpening as discussed in Appendix A.1.

Using the predicted targets for unlabeled nodes, we create a new training set $(\mathbf{X}_u, \hat{\mathbf{Y}}_u)$. The loss from the unlabeled data for the FCN is computed as:

$$\mathcal{L}^{\text{unsupervised}} = \mathcal{L}_{\text{MM}}((\mathbf{X}_u, \hat{\mathbf{Y}}_u), F_\theta, \alpha) \quad (5)$$

Total loss for training the FCN is given as the weighted sum of above two loss terms.

$$\mathcal{L}_{\text{FCN}} = \mathcal{L}^{\text{supervised}} + w(t) * \mathcal{L}^{\text{unsupervised}} \quad (6)$$

where $w(t)$ is a sigmoid ramp-up function (Tarvainen and Valpola 2017) which increases from zero to its max value γ during the course of training.

Now let us assume that the loss for an underlying GNN is $\mathcal{L}_{\text{GNN}} = \ell(G_\theta(\mathbf{X}_l), \mathbf{Y}_l)$; the overall GraphMix loss for the joint training of the FCN and GNN can be defined as the weighted sum of the GNN loss and the FCN loss:

$$\mathcal{L}_{\text{GraphMix}} = \mathcal{L}_{\text{GNN}} + \lambda * \mathcal{L}_{\text{FCN}} \quad (7)$$

However, throughout our experiments, optimizing FCN loss and GNN loss alternatively at each training epoch achieved better test accuracy (discussed in Appendix A.12). This has an added benefit that it removes the need to tune weighing hyper-parameter λ .

For *Manifold Mixup* training of FCN, we apply *mixup* only in the hidden layer. Note that in (Verma et al. 2019a), the authors recommended applying mixing in a randomly chosen layer (which also includes the input layer) at each training update. However, we observed under-fitting when applying *mixup* randomly at the input layer or the hidden layer. Applying *mixup* only in the input layer also resulted in underfitting and did not improve test accuracy.

Memory and Computational Requirements: One of the major limitations of current GNNs, which prohibits their application to real-world large datasets, is their memory complexity. For example, the fastest implementation of GCN, which stores the entire adjacency matrix \mathcal{A} in the memory, has the memory complexity $O(|\mathcal{V}|^2)$. Implementations with lower memory requirements are possible but they incur higher latency cost due to repeatedly loading parts of adjacency matrix in the memory. Due to these reasons, methods which have additional memory requirement in comparison to the baseline GNNs, are less appealing in practice. In GraphMix, since the parameters of the FCN and GNN are shared, there is no additional memory requirement. Furthermore, GraphMix does not add any *significant computation cost* over the underlying GNN, because the underlying GNN is trained in the standard way and the FCN training requires trivial additional computation cost for computing the predicted-targets (Appendix A.1) and the interpolation function ($\text{Mix}_\lambda(\mathbf{a}, \mathbf{b})$ in Section 2).

3.3 Analysis

In this subsection, we study how GraphMix impacts the generalization bound of a underlying GNN. Our analysis, which is based on Rademacher complexity (Bartlett and Mendelson 2002), provides a new generalization bound for GraphMix, which shows how adding regularization to training the FCN with pseudolabels improves overall generalization. We rely on the effect of changing one sample in Manifold Mixup and the fact that the weights are shared by a GNN and the corresponding FCN in GraphMix.

Let G be a fixed graph with n total nodes. Define $z_i = (x_i, y_i)$ to be the pair of the feature and the true label of the i -th node. Without loss of generality, let $S = (z_1, \dots, z_m)$ be the training set with the labeled nodes where $m = n_l$ and data points z_1, \dots, z_m are sampled according to an unknown distribution \mathcal{D} to form the labeled node set S . In this subsection, we follow the previous paper on graph neural networks by assuming that all samples are i.i.d. (including replacement sample) (Verma and Zhang 2019).

Let Γ be a finite set of the hyperparameters. For every hyperparameter $\gamma \in \Gamma$, define \mathcal{F}_γ to be a distribution-dependent hypothesis space corresponding the hyperparameter γ . That is, $\mathcal{F}_\gamma = \{f_\gamma : (\exists S \in \mathcal{S})[f_\gamma = \mathcal{A}_\gamma(S)]\}$ where \mathcal{A}_γ is an algorithm that outputs the hypothesis f_γ given a dataset S , and \mathcal{S} is the set of training datasets such that the probability of $S \in \mathcal{S}$ according to \mathcal{D} is one. For each $f_\gamma \in \mathcal{F}_\gamma$, $f_\gamma(\cdot; G)$ represents GNN with the graph G and $f_\gamma(\cdot; G_0)$ represents FCN where G_0 is the null graph version of G (i.e., G without edges). Let $\mathcal{R}_m^\ell(\mathcal{F}_\gamma)$ be the Rademacher complexity (Bartlett and Mendelson 2002) of the set $\{(x, y) \mapsto \ell(f_\gamma(x; G), y) : f_\gamma \in \mathcal{F}_\gamma\}$.

Let $L_{\text{GNN}}(S, f_\gamma)$ be the \mathcal{L}_{GNN} with the GNN $f_\gamma(\cdot; G)$ and labeled data points S as $L_{\text{GNN}}(S, f_\gamma) = \frac{1}{m} \sum_{i=1}^m \ell(f_\gamma(x_i; G), y_i)$. Let $L_{\text{FCN}}(S, f_\gamma)$ be \mathcal{L}_{FCN} with the FCN $f_\gamma(\cdot; G_0)$ and labeled data points S as $L_{\text{FCN}}(S, f_\gamma) = \mathcal{L}_{\text{MM}}(S, f_\gamma(\cdot; G_0), \alpha) + \frac{n_u}{m} \mathcal{L}_{\text{MM}}(\mathbf{X}_u^S, \hat{\mathbf{Y}}_u^S, f_\gamma(\cdot; G_0), \alpha)$ where $(\mathbf{X}_u^S, \hat{\mathbf{Y}}_u^S)$ is the unlabeled nodes $(\mathbf{X}_u, \hat{\mathbf{Y}}_u)$ that corresponds to the labeled node set S . Let $L_{\text{GraphMix}}(S, f_\gamma) = L_{\text{GNN}}(S, f_\gamma) + \lambda L_{\text{FCN}}(S, f_\gamma)$. Let c be the upper bound on per-sample loss as $c \geq \ell(f_\gamma(x_i; G), y_i)$ and $c \geq \ell(f_\gamma(x_i; G_0), y_i)$. For example, $c = 1$ for training and test error (or 0-1 loss).

Theorem 1 provides a generalization bound for GraphMix, which shows that GraphMix can improve the generalization bound of the underlying GNN under the condition of $V > 0$ as discussed below.

Theorem 1. *For any $\delta > 0$, with probability at least $1 - \delta$, the following holds: for all $\gamma \in \Gamma$ and all $f_\gamma \in \mathcal{F}_\gamma$, we have that $\mathbb{E}_{x, y \sim \mathcal{D}}[\ell(f_\gamma(x; G), y)] - L_{\text{GraphMix}}(S, f_\gamma) \leq \mathcal{R}_m^\ell(\mathcal{F}_\gamma) + c\sqrt{\frac{\ln(|\Gamma|/\delta)}{2m}} - \lambda V$, where $V = \mathbb{E}_{S' \sim \mathcal{D}^m}[\inf_{f_\gamma \in \mathcal{F}_\gamma} L_{\text{FCN}}(S', f_\gamma)] - 4c\sqrt{\frac{\ln(|\Gamma|/\delta)}{2m}}$.*

The proof of Theorem 1 is given in the following. The generalization bound in Theorem 1 becomes a generalization bound for GNN without GraphMix when $\lambda = 0$ as desired. By comparing the generalization bound with $\lambda = 0$ (no GraphMix) and $\lambda > 0$ (with GraphMix), we can see that GraphMix improves the generalization bound of underlying GNN when $V = \mathbb{E}_{S' \sim \mathcal{D}^m}[\inf_{f_\gamma \in \mathcal{F}_\gamma} L_{\text{FCN}}(S', f_\gamma)] - 4c\sqrt{\frac{\ln(|\Gamma|/\delta)}{2m}} > 0$. Here, the first term $\mathbb{E}_{S' \sim \mathcal{D}^m}[\inf_{f_\gamma \in \mathcal{F}_\gamma} L_{\text{FCN}}(S', f_\gamma)]$ increases as the hypothesis space \mathcal{F}_γ gets “smaller” or has less complexity. Thus, GraphMix improves the generalization bound of an underlying GNN when the hyperparameter search over $\gamma \in \Gamma$ results in \mathcal{F}_γ of moderate complexity such that the first term is greater than $4c\sqrt{\frac{\ln(|\Gamma|/\delta)}{2m}}$. The first term contains the manifold mixup loss $L_{\text{FCN}}(S', f_\gamma)$ over random training datasets $S' (\neq S)$, which is expected to be larger when compared with that of the standard loss without manifold mixup.

For each fixed \mathcal{F}_γ , the generalization bound in Theorem 1 goes to zero as $m \rightarrow \infty$ since $\mathcal{R}_m^\ell(\mathcal{F}_\gamma) \rightarrow 0$ and $V \rightarrow V_0 \geq 0$ as $m \rightarrow \infty$. The training loss is also minimized when the trained model $f_\gamma \in \mathcal{F}_\gamma$ fits well to the particular given training data set S . Therefore, given a particular training dataset S , the expected loss is minimized if

we conduct a hyperparameter search over $\gamma \in \Gamma$ such that $f_\gamma \in \mathcal{F}_\gamma$ minimize the training loss for the given S but \mathcal{F}_γ has moderate complexity to not being able to minimize the manifold mixup losses for other datasets $S' (\neq S)$ drawn from \mathcal{D} . Unlike the standard data points, the data points generated during manifold mixup in GraphMix are typically not memorizable or interpolated by FCN.

3.4 Proof of Theorem 1

In the proof, we analyse the effect of changing one sample in Manifold Mixup and GNN-generated targets by utilizing the fact that the weights are shared by a GNN and the corresponding FCN. The the fact of sharing weights also results in the generalization bound that relates the generalization in FCN via Manifold Mixup to the generalization of the GNN.

Proof. Let $\gamma \in \Gamma$ be fixed. Define $\varphi(S) = \sup_{f_\gamma \in \mathcal{F}_\gamma} \mathbb{E}_{x, y \sim \mathcal{D}}[\ell(f_\gamma(x; G), y)] - L_{\text{GraphMix}}(S, f_\gamma)$. We first provide an upper bound on $\varphi(S)$ by using McDiarmid's inequality. To apply McDiarmid's inequality to $\varphi(\mathcal{D})$, we compute an upper bound on $|\varphi(S) - \varphi(S')|$ where S and S' be two training datasets differing by exactly one point of an arbitrary index i_0 ; i.e., $S_i = S'_i$ for all $i \neq i_0$ and $S_{i_0} \neq S'_{i_0}$.

$$\begin{aligned} \varphi(S') - \varphi(S) &\leq \sup_{f_\gamma \in \mathcal{F}_\gamma} L_{\text{GraphMix}}(S, f_\gamma) - L_{\text{GraphMix}}(S', f_\gamma) \\ &= \sup_{f_\gamma \in \mathcal{F}_\gamma} (L_{\text{GNN}}(S, f_\gamma) - L_{\text{GNN}}(S', f_\gamma)) \\ &\quad + \lambda(L_{\text{FCN}}(S, f_\gamma) - L_{\text{FCN}}(S', f_\gamma)) \end{aligned}$$

where the first line follows the property of the supremum, $\sup(a) - \sup(b) \leq \sup(a - b)$, and the second line follows the definition of L_{GraphMix} .

For the first term,

$$\begin{aligned} L_{\text{GNN}}(S, f_\gamma) - L_{\text{GNN}}(S', f_\gamma) &= \\ \frac{1}{m} (\ell(f_\gamma(x_{i_0}; G), y_{i_0}) - \ell(f_\gamma(x'_{i_0}; G), y'_{i_0})) &\leq \frac{c}{m}, \end{aligned}$$

where we used the fact that given a fixed G and a fixed f_γ , $\ell(f_\gamma(x_i; G), y_i) = \ell(f_\gamma(x'_i; G), y'_i)$ for $i \neq i_0$. This holds since $f_\gamma(\cdot; G)$ does not depend on S given a G and a f_γ , even though $f_\gamma(x_i; G)$ contains the aggregation functions over the graph G .

For the second term,

$$\begin{aligned} \mathcal{L}_{\text{MM}}(S, f_\gamma(\cdot; G_0), \alpha) - \mathcal{L}_{\text{MM}}(S', f_\gamma(\cdot; G_0), \alpha) \\ \leq \frac{c(2m-1)}{m^2} \leq \frac{2c}{m}, \end{aligned}$$

where we use the fact that $\mathcal{L}_{\text{MM}}(S, f_\gamma(\cdot; G_0), \alpha)$ has m^2 terms and $2m - 1$ terms differ for S and S' , each of which is bounded by the constant c . Similarly, $\mathcal{L}_{\text{MM}}((\mathbf{X}_u^S, \hat{\mathbf{Y}}_u^S), f_\gamma(\cdot; G_0), \alpha) - \mathcal{L}_{\text{MM}}((\mathbf{X}_u^{S'}, \hat{\mathbf{Y}}_u^{S'}), f_\gamma(\cdot; G_0), \alpha) \leq \frac{2c}{n_u}$, since the labels $\hat{\mathbf{Y}}_u^S$ and $\hat{\mathbf{Y}}_u^{S'}$ are determined by $f_\gamma(x_i; G)$, and $f_\gamma(x_i; G) = f_\gamma(x'_i; G)$ for $i \neq i_0$, given a fixed G and a fixed f_γ . Therefore,

$$L_{\text{FCN}}(S, f_\gamma) - L_{\text{FCN}}(S', f_\gamma) \leq \frac{4c}{m}.$$

Using these upper bounds,

$$\varphi(S') - \varphi(S) \leq \frac{c(1+4\lambda)}{m}.$$

Similarly, $\varphi(S) - \varphi(S') \leq \frac{c(1+4\lambda)}{m}$. Thus, by McDiarmid's inequality, for any $\delta > 0$, with probability at least $1 - \delta/|\Gamma|$,

$$\varphi(S) \leq \mathbb{E}_{\bar{S}}[\varphi(\bar{S})] + c(1+4\lambda)\sqrt{\frac{\ln(|\Gamma|/\delta)}{2m}}.$$

Moreover,

$$\begin{aligned} &\mathbb{E}_{\bar{S}}[\varphi(\bar{S})] + \lambda \mathbb{E}_{\bar{S}} \left[\inf_{f_\gamma \in \mathcal{F}_\gamma} L_{\text{FCN}}(\bar{S}, f_\gamma) \right] \\ &\leq \mathbb{E}_{\bar{S}} \left[\sup_{f_\gamma \in \mathcal{F}_\gamma} \mathbb{E}_{\bar{x}', \bar{y}' \sim \mathcal{D}}[\ell(f_\gamma(\bar{x}'; G), \bar{y}')] - L_{\text{GNN}}(\bar{S}, f_\gamma) \right] \\ &\leq \mathbb{E}_{\bar{S}, \bar{S}'} \left[\sup_{f_\gamma \in \mathcal{F}_\gamma} \frac{1}{m} \sum_{i=1}^m (\ell(f_\gamma(\bar{x}'_i; G), \bar{y}'_i) - \ell(f_\gamma(\bar{x}_i; G), \bar{y}_i)) \right] \\ &\leq \mathbb{E}_{\xi, \bar{\mathcal{D}}, \bar{\mathcal{D}'}} \left[\sup_{f_\gamma \in \mathcal{F}_\gamma} \frac{1}{m} \sum_{i=1}^m \xi_i (\ell(f_\gamma(\bar{x}'_i; G), \bar{y}'_i) - \ell(f_\gamma(\bar{x}_i; G), \bar{y}_i)) \right] \\ &\leq 2\mathcal{R}_n(\Theta) \end{aligned}$$

where the second line and the last line use the subadditivity of supremum, the third line uses the Jensen's inequality and the convexity of the supremum, the fourth line follows that for each $\xi_i \in \{-1, +1\}$, the distribution of each term $\xi_i (\ell(f_\gamma(\bar{x}'_i; G), \bar{y}'_i) - \ell(f_\gamma(\bar{x}_i; G), \bar{y}_i))$ is the distribution of $(\ell(f_\gamma(\bar{x}'_i; G), \bar{y}'_i) - \ell(f_\gamma(\bar{x}_i; G), \bar{y}_i))$ since \bar{S} and \bar{S}' are drawn iid with the same distribution.

Therefore, for any $\delta > 0$, with probability at least $1 - \delta/|\Gamma|$, all $f_\gamma \in \mathcal{F}_\gamma$,

$$\begin{aligned} &\mathbb{E}_{x, y \sim \mathcal{D}}[\ell(f_\gamma(x; G), y)] - L_{\text{GraphMix}}(S, f_\gamma) \\ &\leq \mathcal{R}_m^\ell(\mathcal{F}_\gamma) + c\sqrt{\frac{\ln(|\Gamma|/\delta)}{2m}} - \lambda V. \end{aligned}$$

by taking union bounds over all $\gamma \in \Gamma$, we obtain the statement of this theorem. \square

4 Experiments

We present results for GraphMix algorithm using standard benchmark datasets and the standard architecture in Section 4.1 and 4.3. We also conduct an ablation study on GraphMix in Section A.5 to understand the contribution of various components to its performance. Refer to Appendix A.4 for dataset details and A.8 for implementation and hyperparameter details.

4.1 Semi-supervised Node Classification

For baselines, we choose GCN (Kipf and Welling 2017), and the recent state-of-the-art graph neural networks GAT (Veličković et al. 2018), GMNN (Qu, Bengio, and Tang 2019) and Graph U-Net (Gao and Ji 2019), as well as the recently proposed regularizers for graph neural networks. GraphMix(GCN), GraphMix(GAT) and GraphMix(Graph

Algorithm	Cora	Citeseer	Pubmed
GCN	81.30±0.66	70.61±0.22	79.86±0.34
GAT	82.70±0.21	70.40±0.35	79.05±0.64
Graph U-Net	81.74±0.54	67.69±1.10	77.73 ±0.98
BVAT*	83.6±0.5	74.0±0.6	79.9±0.4
DropEdge*	82.8	72.3	79.6
GraphSGAN*	83.0±1.3	73.1±1.8	-
GraphAT*	82.5	73.4	-
GraphVAT*	82.6	73.7	-
GraphMix (GCN)	83.94±0.57	74.72±0.59	80.98±0.55
GraphMix (GAT)	83.32±0.18	73.08±0.23	81.10±0.78
GraphMix (Graph U-Net)	82.47±0.76	69.31±1.52	78.91±1.25

Table 1: Results of node classification (% test accuracy) on the standard split of datasets. [*] means the results are taken from the corresponding papers. We conduct 100 trials and report mean and standard deviation over the trials (refer to Table 8 in the Appendix for comparison with other methods on standard Train/Validation/Test split).

Algorithm	Cora	Citeseer	Pubmed
GCN	77.84±1.45	72.56±2.46	78.74±0.99
GAT	77.74±1.86	70.41±1.81	78.48±0.96
Graph U-Net	77.59±1.60	67.55±0.69	76.79±2.45
GraphMix (GCN)	82.07±1.17	76.45±1.57	80.72±1.08
GraphMix (GAT)	80.63±1.31	74.08±1.26	80.14±1.51
GraphMix (Graph-U-Net)	80.18±1.62	72.85±1.71	78.47±0.64

Table 2: Results of node classification (% test accuracy) using 10 random Train/Validation/Test split of datasets. We conduct 100 trials and report mean and standard deviation over the trials.

U-Net) refer to the methods where underlying GNNs are GCN, GAT and Graph U-Net respectively. Refer to Appendix A.8 for implementation and hyperparameter details. (Shchur et al. 2018) demonstrated that the performance of the current state-of-the-art Graph Neural Networks on the standard train/validation/test split of the popular benchmark datasets (such as Cora, Citeseer, Pubmed, etc) is significantly different from their performance on the random splits. For fair evaluation, they recommend using multiple random partitions of the datasets. Along these lines, we created 10 random splits of the Cora, Citeseer and Pubmed with the same train/ validation/test number of samples as in the standard split. We also provide the results for the standard train/validation/test split. The results are presented in Table 2. We observe that GraphMix always improves the accuracy of the underlying GNNs such as GCN, GAT and Graph-U-Net across all the dataset, with GraphMix(GCN) achieving the best results. We further present results with fewer labeled samples in Appendix A.7. We observe that the relative increase in test accuracy using GraphMix over the baseline GNN is more pronounced when the labeled samples are fewer. This makes GraphMix particularly appealing for very few labeled data problems.

Algorithm	Cora-Full	Coauthor-CS	Coauthor-Physics
GCN*	62.2±0.6	91.1±0.5	92.8±1.0
GAT*	51.9±1.5	90.5±0.6	92.5±0.9
MoNet*	59.8±0.8	90.8±0.6	92.5±0.9
GS-Mean*	58.6±1.6	91.3±2.8	93.0±0.8
GCN	60.13±0.57	91.27±0.56	92.90±0.92
Graph-U-Net	59.82±0.39	90.89±0.43	92.57±0.81
GraphMix (GCN)	61.80±0.54	91.83±0.51	94.49±0.84
GraphMix (Graph-U-Net)	60.92 ± 0.51	91.44 ± 0.46	93.78 ± 0.79

Table 3: Comparison of GraphMix with other methods (% test accuracy), for Cora-Full, Coauthor-CS, Coauthor-Physics. * refers to the results reported in (Shchur et al. 2018).

Algorithm	Bit OTC	Bit Alpha
DeepWalk	63.20	62.71
GMNN	66.93	65.86
GCN	65.72±0.38	64.00±0.19
GraphMix (GCN)	66.35±0.41	65.34±0.19

Table 4: Results on Link Classification (%F1 score). * means the results are taken from the corresponding papers.

4.2 Results on Larger Datasets

We also provide results on three recently proposed datasets which are relatively larger than standard benchmark datasets (Cora/Citeseer/Pubmed). We use Cora-Full dataset proposed in (Bojchevski and Günnemann 2018) and Coauthor-CS and Coauthor-Physics datasets proposed in (Shchur et al. 2018). The results are presented in Table 3². We observe that GraphMix(GCN) improves the results over GCN for all the three datasets with a significant margin. We note that we did minimal hyperparameter search for GraphMix(GCN) as mentioned in Appendix A.8. The details of the datasets is given in Appendix A.4.

4.3 Semi-supervised Link Classification

In the semi-supervised link classification problem, the task is to predict the labels of the remaining links, given a graph and labels of a few links. Following (Taskar et al. 2004), we can formulate the link classification problem as a node classification problem, i.e., given an original graph G , we construct a dual Graph G' , the node set V' of the dual graph corresponds to the link set E' of the original graph. The nodes in the dual graph G' are connected if their corresponding links in the graph G share a node. The attributes of a node in the dual graph are defined as the index of the nodes of the corresponding link in the original graph. Using this formulation, we present results on link classification on Bit OTC and Bit Alpha benchmark datasets in the Table 4. As the numbers of the positive and negative edges are strongly imbalanced, we report the F1 score. Our results show that GraphMix(GCN)

²We do not provide results for GAT based experiments in Table 3 and Table 4 because we ran out of GPU memory required to run these experiments with larger (higher number of nodes) datasets.

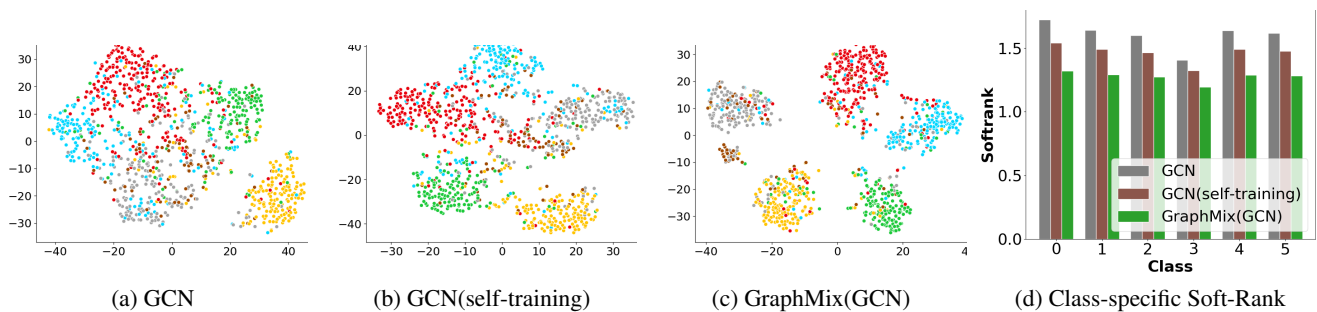


Figure 2: Two-dimensional representation of the hidden states of Citeseer dataset using (a) GCN, (b) GCN(self-training), (c) GraphMix(GCN), and Soft-Rank (d). GraphMix(GCN) learns better separated representations.

improves the performance over the baseline GCN method and is comparable with the recently proposed state-of-the-art method GMNN (Qu, Bengio, and Tang 2019) for both the datasets.

4.4 Visualization of the Learned Features

In Figure 2, we present an analysis of the features learned by GraphMix for Cora dataset using the t-SNE (van der Maaten and Hinton 2008) based 2D visualization of the hidden states. We observe that GraphMix learns hidden states which are better separated and condensed than GCN and GCN(self-training). Here, GCN(self-training) refers to training a GCN in a normal way but with additional self-prediction based targets for the unlabeled samples. We further evaluate the Soft-rank (refer to Appendix A.10) of the class-specific hidden states to demonstrate that GraphMix(GCN) makes the class-specific hidden states more concentrated than GCN and GCN(self-training), as shown in Figure 2d. Refer to Appendix A.11 for 2D representation of other datasets.

5 Related Work

- Semi-supervised Learning over Graph Data:** There exists a long line of work for semi-supervised learning over graph data (Zhu and Ghahramani 2002; Zhu, Ghahramani, and Lafferty 2003). Contrary to previous methods, the recent GNN based approaches define the convolutional operators using the neighbourhood information of the nodes (Kipf and Welling 2017; Veličković et al. 2018). These convolution operator based method exhibit state-of-the-results for semi-supervised learning over graph data, hence much of the recent attention is dedicated to proposing architectural changes to these methods (Qu, Bengio, and Tang 2019; Gao and Ji 2019; Ma et al. 2019). Unlike these methods, we propose a regularization technique that can be applied to any of these GNNs which uses a parameterized transformation on the node features.
- Data Augmentation:** It is well known that the generalization of a learning algorithm can be improved by enlarging the training data size. Since labeling more samples is labour-intensive and costly, data-augmentation has become a popular technique for enlarging the training data size. Mixing based algorithms are a particular class of data-augmentation methods in which additional training

samples are generated by interpolating the samples (either in the input or hidden space) and/or their corresponding targets. Mixup (Zhang et al. 2018), Manifold Mixup (Verma et al. 2019a), AMR (Beckham et al. 2019), ISD (Jeong et al. 2021) are notable examples of this class of algorithms. Unlike these methods which have been proposed for the fixed topology datasets such as images, we propose interpolation based data-augmentation for graph-structured data.

- Regularizing Graph Neural Networks:** Regularizing Graph Neural Networks has drawn some attention recently. GraphSGAN (Ding, Tang, and Zhang 2018) first uses an embedding method such as DeepWalk (Perozzi, Al-Rfou, and Skiena 2014) and then trains generator-classifier networks in the adversarial learning setting to generate fake samples in the low-density region between sub-graphs. BVAT (Deng, Dong, and Zhu 2019) and (Feng et al. 2019) generate adversarial perturbations to the features of the nodes while taking graph structure into account. These methods require significant additional computation: GraphScan requires computing node embedding as a preprocessing step, BVAT and (Feng et al. 2019) require additional gradient computation for computing adversarial perturbations. Unlike these methods, GraphMix does not introduce any significant additional computation since it is based on interpolation-based techniques and self-generated targets.

6 Discussion

GraphMix is a simple and efficient regularizer for semi-supervised node classification using graph neural networks. Our extensive experiments demonstrate state-of-the-art performance using GraphMix on benchmark datasets. Our theoretical analysis compares generalization bounds of GraphMix vs the underlying GNNs. Further, we conduct a systematic ablation study to understand the effect of different components in the performance of GraphMix. The strong empirical results of GraphMix suggest that in parallel to designing new architectures, exploring better regularization for graph-structured data is a promising avenue for research. A future research direction is to jointly model the node features and edges of the graph such that they can be further used for generating the synthetic interpolated nodes and their corresponding connectivity to the other nodes in the graph.

References

- Bartlett, P. L.; and Mendelson, S. 2002. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* 3(Nov): 463–482.
- Beckham, C.; Honari, S.; Verma, V.; Lamb, A.; Ghadiri, F.; Devon Hjelm, R.; Bengio, Y.; and Pal, C. 2019. On Adversarial Mixup Resynthesis. *arXiv e-prints* arXiv:1903.02709.
- Belkin, M.; Niyogi, P.; and Sindhvani, V. 2006. Manifold Regularization: A Geometric Framework for Learning from Labeled and Unlabeled Examples. *J. Mach. Learn. Res.* 7: 2399–2434. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1248547.1248632>.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. 2019. MixMatch: A Holistic Approach to Semi-Supervised Learning. *arXiv e-prints* arXiv:1905.02249.
- Blum, A.; and Mitchell, T. 1998. Combining Labeled and Unlabeled Data with Co-training. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory, COLT' 98*, 92–100. New York, NY, USA: ACM. ISBN 1-58113-057-0. doi:10.1145/279943.279962. URL <http://doi.acm.org/10.1145/279943.279962>.
- Bojchevski, A.; and Günnemann, S. 2018. Deep Gaussian Embedding of Graphs: Unsupervised Inductive Learning via Ranking. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=r1ZdKJ-0W>.
- Bruna, J.; Zaremba, W.; Szlam, A.; and LeCun, Y. 2013. Spectral Networks and Locally Connected Networks on Graphs. *CoRR* abs/1312.6203.
- Chapelle, O.; Schölkopf, B.; and Zien, A. 2010. *Semi-Supervised Learning*. The MIT Press, 1st edition. ISBN 0262514125, 9780262514125.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In Lee, D. D.; Sugiyama, M.; Luxburg, U. V.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 29*, 3844–3852.
- Deng, Z.; Dong, Y.; and Zhu, J. 2019. Batch Virtual Adversarial Training for Graph Convolutional Networks. *CoRR* abs/1902.09192. URL <http://arxiv.org/abs/1902.09192>.
- Devries, T.; and Taylor, G. W. 2017. Improved Regularization of Convolutional Neural Networks with Cutout. *CoRR* abs/1708.04552. URL <http://arxiv.org/abs/1708.04552>.
- Ding, M.; Tang, J.; and Zhang, J. 2018. Semi-supervised Learning on Graphs with Generative Adversarial Nets. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18*, 913–922. New York, NY, USA: ACM. ISBN 978-1-4503-6014-2. doi:10.1145/3269206.3271768. URL <http://doi.acm.org/10.1145/3269206.3271768>.
- Feng, F.; He, X.; Tang, J.; and Chua, T. 2019. Graph Adversarial Training: Dynamically Regularizing Based on Graph Structure. *CoRR* abs/1902.08226. URL <http://arxiv.org/abs/1902.08226>.
- Gao, H.; and Ji, S. 2019. Graph U-Nets. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 2083–2092. Long Beach, California, USA: PMLR. URL <http://proceedings.mlr.press/v97/gao19a.html>.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural Message Passing for Quantum Chemistry. In *ICML*.
- Gori, M.; Monfardini, G.; and Scarselli, F. 2005. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, 729–734. IEEE.
- Grandvalet, Y.; and Bengio, Y. 2005. Semi-supervised Learning by Entropy Minimization. In Saul, L. K.; Weiss, Y.; and Bottou, L., eds., *Advances in Neural Information Processing Systems 17*, 529–536.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. In *NIPS*.
- Henaff, M.; Bruna, J.; and LeCun, Y. 2015. Deep Convolutional Networks on Graph-Structured Data. *ArXiv* abs/1506.05163.
- Hjelm, R. D.; Fedorov, A.; Lavoie-Marchildon, S.; Grewal, K.; Bachman, P.; Trischler, A.; and Bengio, Y. 2019. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=Bklr3j0cKX>.
- Jeong, J.; Verma, V.; Hyun, M.; Kannala, J.; and Kwak, N. 2021. Interpolation-based semi-supervised learning for object detection. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Ko, T.; Peddinti, V.; Povey, D.; and Khudanpur, S. 2015. Audio augmentation for speech recognition. In *INTERSPEECH*.
- Kumar, S.; Hooi, B.; Makhija, D.; Kumar, M.; Faloutsos, C.; and Subrahmanian, V. 2018. Rev2: Fraudulent user prediction in rating platforms. In *WSDM*.
- Kumar, S.; Spezzano, F.; Subrahmanian, V.; and Faloutsos, C. 2016. Edge weight prediction in weighted signed networks. In *ICDM*.
- Laine, S.; and Aila, T. 2016. Temporal Ensembling for Semi-Supervised Learning. *CoRR* abs/1610.02242. URL <http://arxiv.org/abs/1610.02242>.
- Lee, D.-H. 2013. Pseudo-Label : The Simple and Efficient Semi-Supervised Learning Method for Deep Neural Networks URL <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.664.3543>. Accessed on 25 October 2019.
- Li, Q.; Han, Z.; and Wu, X.-M. 2018. Deeper Insights into Graph Convolutional Networks for Semi-Supervised Learning. In *AAAI*.

- Lu, Q.; and Getoor, L. 2003. Link-based Classification. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning, ICML'03*, 496–503. AAAI Press. ISBN 1-57735-189-4. URL <http://dl.acm.org/citation.cfm?id=3041838.3041901>.
- Ma, J.; Cui, P.; Kuang, K.; Wang, X.; and Zhu, W. 2019. Disentangled Graph Convolutional Networks. In *ICML*.
- Miyato, T.; Ichi Maeda, S.; Koyama, M.; and Ishii, S. 2018. Virtual Adversarial Training: a Regularization Method for Supervised and Semi-supervised Learning. *IEEE transactions on pattern analysis and machine intelligence*.
- Monti, F.; Boscaini, D.; Masci, J.; Rodola, E.; Svoboda, J.; and Bronstein, M. M. 2016. Geometric deep learning on graphs and manifolds using mixture model CNNs. *CoRR* abs/1611.08402. URL <http://arxiv.org/abs/1611.08402>.
- Park, D. S.; Chan, W.; Zhang, Y.; Chiu, C.-C.; Zoph, B.; Cubuk, E. D.; and Le, Q. V. 2019. SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition. *arXiv e-prints* arXiv:1904.08779.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *KDD*.
- Qu, M.; Bengio, Y.; and Tang, J. 2019. GMNN: Graph Markov Neural Networks. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 5241–5250. Long Beach, California, USA: PMLR.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2009. The Graph Neural Network Model. *Trans. Neur. Netw.* 20(1): 61–80. ISSN 1045-9227. doi: 10.1109/TNN.2008.2005605. URL <http://dx.doi.org/10.1109/TNN.2008.2005605>.
- Shchur, O.; Mumme, M.; Bojchevski, A.; and Günnemann, S. 2018. Pitfalls of Graph Neural Network Evaluation. *CoRR* abs/1811.05868. URL <http://arxiv.org/abs/1811.05868>.
- Sun, F.-Y.; Hoffman, J.; Verma, V.; and Tang, J. 2020. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=r1lff2NYvH>.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in Neural Information Processing Systems 30*, 1195–1204.
- Taskar, B.; Wong, M.-F.; Abbeel, P.; and Koller, D. 2004. Link prediction in relational data. In *NIPS*.
- van der Maaten, L.; and Hinton, G. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9: 2579–2605. URL <http://www.jmlr.org/papers/v9/vandermaaten08a.html>.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *ICLR*.
- Veličković, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep graph infomax. In *ICLR*.
- Verma, S.; and Zhang, Z.-L. 2019. Stability and generalization of graph convolutional neural networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1539–1548.
- Verma, V.; Lamb, A.; Beckham, C.; Najafi, A.; Mitliagkas, I.; Lopez-Paz, D.; and Bengio, Y. 2019a. Manifold Mixup: Better Representations by Interpolating Hidden States. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 6438–6447. Long Beach, California, USA: PMLR. URL <http://proceedings.mlr.press/v97/verma19a.html>.
- Verma, V.; Lamb, A.; Juho, K.; Bengio, Y.; and Lopez-Paz, D. 2019b. Interpolation Consistency Training for Semi-supervised Learning. In Kraus, S., ed., *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*. ijcai.org. doi:10.24963/ijcai.2019. URL <https://doi.org/10.24963/ijcai.2019>.
- Weston, J.; Ratle, F.; Mobahi, H.; and Collobert, R. 2012. Deep Learning via Semi-Supervised Embedding. In Montavon, G.; Orr, G.; and Müller, K. R., eds., *In Neural Networks: Tricks of the Trade*. Springer, second edition.
- Xie, Z.; Wang, S. I.; Li, J.; Lévy, D.; Nie, A.; Jurafsky, D.; and Ng, A. Y. 2017. Data Noising as Smoothing in Neural Network Language Models. *ArXiv* abs/1703.02573.
- Xu, K.; Li, C.; Tian, Y.; Sonobe, T.; Kawarabayashi, K.-i.; and Jegelka, S. 2018. Representation Learning on Graphs with Jumping Knowledge Networks. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 5453–5462. Stockholm, Sweden: PMLR. URL <http://proceedings.mlr.press/v80/xu18c.html>.
- Yang, Z.; Cohen, W.; and Salakhutdinov, R. 2016. Revisiting Semi-Supervised Learning with Graph Embeddings. In *ICML*.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. *International Conference on Learning Representations* URL <https://openreview.net/forum?id=r1Ddp1-Rb>.
- Zhou, J.; Cui, G.; Zhang, Z.; Yang, C.; Liu, Z.; and Sun, M. 2018. Graph Neural Networks: A Review of Methods and Applications. *CoRR* abs/1812.08434. URL <http://arxiv.org/abs/1812.08434>.
- Zhu, X.; and Ghahramani, Z. 2002. Learning from Labeled and Unlabeled Data with Label Propagation. URL <http://www.scholar.google.com/url?sa=U&q=http://www.cs.cmu.edu/~zhuxj/pub/propagate.ps.gz>. Accessed on 25 October 2019.
- Zhu, X.; Ghahramani, Z.; and Lafferty, J. D. 2003. Semi-supervised learning using gaussian fields and harmonic functions. In *ICML*.