

GRAPH-BASED NEURAL NETWORK MODELS WITH MULTIPLE SELF-SUPERVISED AUXILIARY TASKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Self-supervised learning is currently gaining a lot of attention, as it allows neural networks to learn robust representations from large quantities of unlabeled data. Additionally, multi-task learning can further improve representation learning by training networks simultaneously on related tasks, leading to significant performance improvements. In this paper, we propose a general framework to improve graph-based neural network models by combining self-supervised auxiliary learning tasks in a multi-task fashion. Since Graph Convolutional Networks are among the most promising approaches for capturing relationships among structured data points, we use them as a building block to achieve competitive results on standard semi-supervised graph classification tasks.

1 INTRODUCTION

In the last decade, neural networks approaches that can deal with with structured data have been gaining a lot of traction (Scarselli et al., 2009; Bruna et al., 2013; Defferrard et al., 2016; Kipf & Welling, 2017; Manessi et al., 2020). Due to the prevalence of data structured in the form of graphs, the capability to explicitly exploit structural relationships among data points is particularly useful in improving the performance for a variety of tasks. Graph Convolutional Networks (GCNs, Kipf & Welling (2017)) stand out as a particularly successful iteration of such networks, especially for semi-supervised problems. GCNs act to encode graph structures, while being trained on a supervised target loss for all the nodes with labels. This technique is able to share the gradient information from the supervised loss through the graph adjacency matrix and to learn representations exploiting both labeled and unlabeled nodes. Although GCNs can stack multiple graph convolutional layers in order to capture high-order relations, these architectures suffer from “over-smoothing” when the number of layers increases (Li et al., 2018), thus making difficult to choose an appropriate number of layers.

If we have a dataset with enough labels, supervised learning can usually achieve good results. Unfortunately, to label a large amount of data is an expensive task. In general, the amount of unlabelled data is substantially more than the data that has been human curated and labelled. It is therefore valuable to find ways to make use of this unlabelled data. A potential solution to this problem comes if we can get labels from unlabelled data and train unsupervised dataset in a supervised manner. Self-supervision achieves this by automatically generating additional labelled signals from the available unlabelled data, using them to learn representations. A possible approach in deep learning involves taking a complex signal, hiding part of it from the network, and then asking the network to fill in the missing information (Doersch & Zisserman, 2017).

Additionally, it is found that joint learning of different tasks can improve performance over learning them individually, given that at least a subset of these tasks are related to each other (Caruana, 1997). This observation is at the core of multi-task learning. Precisely, given T tasks $\{\mathcal{T}_i\}_{i=1}^T$ where a subset of them are related, multi-task learning aims to help improve the learning of a model for $\{\mathcal{T}_i\}_{i=1}^T$ by using the knowledge contained in all or some of the T tasks (Zhang & Yang (2017a)).

The idea presented in this paper is to learn self-supervised auxiliary tasks in a multi-task framework, to improve the performance achieved by a neural network-based graph architecture. The proposed framework is general but, considering the promising results of the GCN, we decided to experiment it in semi-supervised classification problems on graphs employing GCN as a base building block. This new framework allows us to achieve competitive results on standard datasets and helps to reduce the aforementioned “over-smoothing” limitation of deep GCNs.

We investigate three different auxiliary tasks:

autoencoding: with which we aim at extracting node representations robust enough to allow both semi-supervised classification as well as vertex features reconstruction;

corrupted features reconstruction: with which we try to extract node representations that allows to reconstruct some of the vertex input features, starting from an embedding built from a corrupted version of them. This auxiliary task can be seen as the graph equivalent of reconstructing one of the color channels of a RGB image using the other channels in computer vision self-supervised learning;

corrupted embeddings reconstruction: with which we try to extract node representations robust to embedding corruption. This is similar to the aforementioned auxiliary task, with the difference that the reconstruction is performed on the node embeddings instead of the vertex features.

These three tasks are intrinsically self-supervised, since the labels are directly extracted from the input graph and its vertex features. To the best of our knowledge, this is the first attempt to combine multiple self-supervised auxiliary task with Graph Convolutional Networks using a multi-task framework.

The paper is organized as follows: in Section 2 the related works are summarized; in Section 3 we describe our methods; in Section 4 a detailed comparison against GCN on a standard public datasets is presented; Section 5 reports conclusions and future works.

2 RELATED WORKS

In recent years, graph representation learning have gained a lot of attention. These techniques can be divided in three main categories: (i) random walk-based; (ii) factorization-based; (iii) neural network-based. In the first group, *node2vec* (Grover & Leskovec, 2016) and *Deepwalk* (Perozzi et al., 2014) are worth mentioning. The former is an efficient and scalable algorithm for feature learning that optimizes a novel network-aware, neighborhood preserving objective function, using stochastic gradient descent. The latter uses truncated random walks to efficiently learn representations for vertices in graphs. These latent representations, which encode graph relations in a vector space, can be easily exploited by standard statistical models to produce state-of-the-art results.

Among the factorization based methods, Xu et al. (2013) presents a semi-supervised factor graph model that can exploit the relationships among the nodes. In this approach, each vertex is modeled as a variable node and the various relationships are modeled as factor nodes.

In the last group, we find all the works that have revisited the problem of generalizing neural networks to work on structured graphs, some of them achieving promising results in domains that have been previously dominated by other techniques. Gori et al. (2005) and Scarselli et al. (2009) formalize a novel neural network model, the Graph Neural Network. This model maps a graph and its nodes into a D -dimensional Euclidean space in order to learn a final classification/regression model. Bruna et al. (2013) approach the graph structured data by proposing two generalizations of Convolutional Neural Networks (CNNs): one based on a hierarchical clustering of the domain and another based on the spectrum of the graph (computed using the Laplacian matrix). Defferrard et al. (2016) extend the spectral graph theory approach of the previous work by providing efficient numerical schemes to design fast localized convolutional filters on graphs, achieving the same computational complexity of classical CNNs. Kipf & Welling (2017) build on this idea by introducing GCNs. They exploit a localized first-order approximation of the spectral graph convolutions framework (Hammond et al., 2011). Recently, Veličković et al. (2018) have applied attention mechanism to graph neural networks to improve model performance.

However, the majority of the methods belonging the three aforementioned categories require a large amount of labelled data, which can limit their applicability. On the other hand, unsupervised algorithms, such as Hamilton et al. (2017); Grover et al. (2019); Veličković et al. (2019) do not require any external labels, but their performances usually suffer when compared to supervised techniques.

Self-supervised learning can be considered a branch of unsupervised learning, where virtually unlimited supervised signals are generated from the available data and used to learn representations. This learning framework finds many applications, ranging from language modeling (Wu et al., 2019;

Mikolov et al., 2013; Radford et al., 2018), to robotics (Jang et al., 2018), and computer vision (Zhang et al., 2016; Ledig et al., 2017; Pathak et al., 2016; Zhang et al., 2017; Noroozi & Favaro, 2016; Doersch et al., 2015). Applied to graph representation learning, Sun et al. (2020) proposed a multi-stage self-supervised framework, called M3S, showing some empirical success.

Multi-task learning approaches can be divided in five many categories: (i) feature learning; (ii) low-rank approaches; (iii) task clustering; (iv) task relation learning; (v) decomposition (Zhang & Yang, 2017b). In the feature learning approach, it is assumed that different tasks share a common feature representation based on the original features. In the Multi-Task Feature Learning method, task specific hidden representations within a shallow network are obtained by learning the feature covariance for all the tasks, in turn allowing to decouple the learning of the different tasks (Argyriou et al., 2007; 2008). A common approach applied in the deep learning setting is to have the different tasks share the first several hidden network layers, including task-specific parameters only in the subsequent layers (Zhang et al., 2014; Mrkšić et al., 2015; Li et al., 2014). A more complex approach in deep learning is the cross-stitch network, proposed by Misra et al. (2016), in which each task has its own independent hidden layers that operate on learned linear combinations of the activation maps of the previous layers.

The low-rank approaches assume that the model parameters of different tasks share a low-rank subspace (Ando & Zhang, 2005). Pong et al. (2010) propose to regularize the model parameters by means of the trace norm regularizer, in order to exploit the property of the trace norm to induce low rank matrices. The same idea has been applied in deep learning by Yang & Hospedales (2016).

Another approach is to assume that different tasks form several clusters, each of which consists of similar tasks. This can be thought of as clustering algorithms on the task level, while the conventional clustering algorithms operate on the data level. Thrun & O’Sullivan (1996) introduced the first implementation of this idea for binary classification tasks that are defined over the same input space. Bakker & Heskes (2003) followed the idea to recast the neural networks used in the feature learning approach in a Bayesian settings, where the weights of the task specific final layers are assumed to have a Gaussian mixture as a prior. Similarly, Xue et al. (2007) build on the previous idea by changing the prior to a Dirichlet process.

In the task relation based approaches, the task relatedness (e.g. task correlation or task covariance) is used to drive the joint training of multiple tasks. In the early works, these relations are assumed to be known in advance. They are used to design regularizers to guide the learning of multiple tasks, so that the more similar two tasks are, the closer the corresponding model parameters are expected to be (Evgeniou et al., 2005; Kato et al., 2008). However, in most applications, task relations are not available and need to be automatically estimated from data. Bonilla et al. (2008) go into this direction by exploiting Gaussian processes and defining a multivariate normal prior on the functional output of all the task outputs, whose covariance is trained from data and represents the relation between the different tasks.

Finally, in the decomposition approach, it is assumed that the matrix whose row vectors are the weights of each of the tasks can be decomposed as a linear combination of two or more sub-matrices, where each sub-matrix is suitably regularized (Jalali et al., 2010; Chen et al., 2012; Zhong & Kwok, 2012).

3 METHODS

In this section, we introduce the formalization of a multi-task self-supervised GCN for semi-supervised classification. We will first give some preliminary definitions, including of a Graph Convolutional (GC) layer and multi-task target loss. We then proceed by showing the auxiliary tasks that can be learned jointly with the semi-supervised classification loss. Finally, we introduce the overall architecture we used in our experiments.

3.1 PRELIMINARIES

Let $Y_{i,j}$ be the i -th row, j -th column element of the matrix \mathbf{Y} . \mathbf{I}_d is the identity matrix in \mathbb{R}^d ; softmax and ReLU are the *soft-maximum* and the *rectified linear unit* activation functions (Goodfellow et al., 2016). Note that all the activation functions act element-wise when applied to a matrix.

An *undirected* graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is defined by its set of the nodes (or vertices), \mathcal{V} , and set of the edges, \mathcal{E} . For each vertex $v_i \in \mathcal{V}$ let $\mathbf{v}_i \in \mathbb{R}^d$ be the corresponding feature vector. Moreover, let \mathbf{A} be the adjacency matrix of the graph \mathcal{G} ; namely, $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ where $A_{i,j} = A_{j,i} = w_{ij}$ if and only if there is an edge between the i -th and j -th vertices and the edge has weight w_{ij} . In the case of an unweighted graph, $w_{ij} = 1$. The symbol \mathbf{X} will denote instead the vertex-features matrix $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$, i.e. the matrix whose row vectors are the \mathbf{x}_i .

The mathematics of the GC layer (Kipf & Welling, 2017) is here briefly recalled, since it is a basic building block of the following network architectures. Given a graph with adjacency matrix $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ and vertex-feature matrix $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$, the GC layer with M output nodes (also called channels) and $\mathbf{B} \in \mathbb{R}^{d \times M}$ weight matrix is defined as the function GC_M from $\mathbb{R}^{|\mathcal{V}| \times d}$ to $\mathbb{R}^{|\mathcal{V}| \times M}$ such as $\text{GC}_M(\mathbf{X}) := \hat{\mathbf{A}}\mathbf{X}\mathbf{B}$, where $\hat{\mathbf{A}}$ is the re-normalized adjacency matrix, i.e. $\hat{\mathbf{A}} := \tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2}$ with $\tilde{\mathbf{A}} := \mathbf{A} + \mathbf{I}_{|\mathcal{V}|}$ and $\tilde{D}_{kk} := \sum_l \tilde{A}_{kl}$. Note that the GC layer can be seen as localized first-order approximation of spectral graph convolution (Defferrard et al., 2016), with the additional *renormalization trick* in order to improve numerical stability (Kipf & Welling, 2017).

Consider now a multi-task problem, made of T tasks, indexed by $t = 1, \dots, T$. All the tasks share the input space \mathcal{X} and have the task-specific output spaces \mathcal{Y}_t . We suppose that each task t is associated to a parametric hypothesis class f_t (e.g. a neural network architecture) such that $f_t(\mathbf{x}; \vartheta_{\text{sh}}, \vartheta_t) = \hat{y}^t$, where $\mathbf{x} \in \mathcal{X}$, $y^t \in \mathcal{Y}_t$, ϑ_{sh} is a parameter vector shared among the hypothesis classes of different tasks, and ϑ_t is task-specific. The joint training of each of the f_t is achieved by means of empirical risk minimization:

$$\operatorname{argmin}_{\vartheta_{\text{sh}}, \vartheta_1, \dots, \vartheta_T} \sum_{t=1}^T w_t \mathcal{R}_t(\vartheta_{\text{sh}}, \vartheta_t), \quad (1)$$

where $w_t \in \mathbb{R}^+$ and $\mathcal{R}_t(\vartheta_{\text{sh}}, \vartheta_t)$ are the task-specific empirical risks. Precisely, $\mathcal{R}_t(\vartheta_{\text{sh}}, \vartheta_t) := \frac{1}{N} \sum_i \mathcal{L}_t(f_t(\mathbf{x}_i; \vartheta_{\text{sh}}, \vartheta_t), y_i^t)$ with \mathcal{L}_t the task-specific loss function, \mathbf{x}_i the feature vectors of the i -th training sample, y_i^t the target variable of the i -th training sample corresponding to the t -th task, and N the total number of training samples. Roughly speaking, the multi-task objective of Equation (1) is the conic combination with weights w_t of the empirical risk of each task. A basic justification for taking the weighted combination is due to the fact that it is not possible to define global optimality in the multi-task setting. Indeed, consider two sets of solutions $(\vartheta_{\text{sh}}, \vartheta_1, \vartheta_2)$ and $(\bar{\vartheta}_{\text{sh}}, \bar{\vartheta}_1, \bar{\vartheta}_2)$ such that $\mathcal{R}_1(\vartheta_{\text{sh}}, \vartheta_1) < \mathcal{R}_1(\bar{\vartheta}_{\text{sh}}, \bar{\vartheta}_1)$ and $\mathcal{R}_2(\vartheta_{\text{sh}}, \vartheta_2) > \mathcal{R}_2(\bar{\vartheta}_{\text{sh}}, \bar{\vartheta}_2)$, i.e. $(\vartheta_{\text{sh}}, \vartheta_1, \vartheta_2)$ is the best solution for the first task, while $(\bar{\vartheta}_{\text{sh}}, \bar{\vartheta}_1, \bar{\vartheta}_2)$ reaches optimality in the second task. It is not possible to compare these two solutions without a pairwise measure. A possible way to put them on the same footing is by mean of the conic combination of Equation (1).

The weights w_t will be considered as static hyper-parameters of the training procedure in the remaining of the paper. It is worth mentioning that also other approaches exist in which the weights are dynamically computed or obtained through an heuristic (Chen et al., 2018; Kendall et al., 2018).

It is worth noting that the framework we are considering is usually called *hard parameter sharing*, i.e. there are some parameters ϑ_{sh} that are shared among all the tasks. On the other hand, in *soft parameter sharing*, all parameters are task-specific but they are jointly constrained by means of regularization.

3.2 THE TASKS

This section is organized as follows: in 3.2.1 the main task is defined; in 3.2.2, 3.2.3, 3.2.4 the auxiliary tasks are formalized.

3.2.1 THE MAIN TASK

As mentioned before, we will consider the semi-supervised classification of graph nodes as our main task. However, what follows can easily be extended to other main tasks as well.

Let's consider a K -class semi-supervised classification problem; thus the output space of the main task can be written as $\mathcal{Y}_{\text{main}} := \{\mathbf{y} \in \mathbb{R}^K \mid y_k \in \{0, 1\}, \sum_k y_k = 1\}$, i.e. the space of one-hot encoded K -class vectors. By denoting with $\mathcal{V}_1 \subseteq \mathcal{V}$ the subset of the labeled nodes of the graph \mathcal{G} ,

the empirical risk $\mathcal{R}_{\text{main}}$ of the main task, corresponding to a cross-entropy loss, can be written as:

$$\mathcal{R}_{\text{main}} := \frac{1}{|\mathcal{V}_1|} \sum_{i \in \mathcal{V}_1} \sum_{k=1}^K y_k \log f_{\text{main}}(\mathbf{x}_i; \vartheta_{\text{sh}}, \vartheta_{\text{main}}), \quad \text{with } 0 \times \log 0 = 0.$$

We make the assumption that $f_{\text{main}} := g_{\text{sh}} \circ g_{\text{main}}$, with $\partial g_{\text{sh}} / \partial \vartheta_{\text{main}} = \partial g_{\text{main}} / \partial \vartheta_{\text{sh}} = 0$, namely, f_{main} can be seen as the function composition of a vertex feature embedding function g_{sh} parameterized only by ϑ_{sh} , followed by a task specific classification head g_{main} parameterized by ϑ_{main} only. As we will see later, g_{sh} is shared with the auxiliary tasks. Finally, g_{sh} , g_{main} , and all the functions we will discuss further ahead are considered differentiable almost everywhere.

3.2.2 AUTOENCODING

The objective in the autoencoding task (also called AE) is to reconstruct the graph vertex features from an encoding thereof. Using the mean squared error reconstruction loss, the corresponding empirical risk \mathcal{R}_{AE} can be written as:

$$\mathcal{R}_{\text{AE}} := \frac{1}{|\mathcal{V}_{\text{AE}}|} \sum_{\substack{i \in \mathcal{V}_{\text{AE}} \\ \mathcal{V}_{\text{AE}} \subseteq \mathcal{V}}} \|\mathbf{x}_i - f_{\text{AE}}(\mathbf{x}_i; \vartheta_{\text{sh}}, \vartheta_{\text{AE}})\|_2^2, \quad \text{with } f_{\text{AE}} := g_{\text{sh}} \circ g_{\text{AE}}, \quad \frac{\partial g_{\text{AE}}}{\partial \vartheta_{\text{sh}}} = 0.$$

Namely, the autoencoder is made of the encoder function g_{sh} also present in the main task, and a decoder component specified by g_{AE} that depends on the task specific parameters ϑ_{AE} only.

3.2.3 CORRUPTED FEATURES RECONSTRUCTION

The aim of this task (also called FR) is to reconstruct the graph vertex features from an encoding of a corrupted version of them. Namely, the goal is to train an autoencoder that it is able to restore vertex features starting from a vertex-feature matrix \mathbf{X} that has some columns zeroed out, i.e. corrupted.

We distinguish two methods: (i) *partial reconstruction*, where we aim at outputting the restored features only; (ii) *full reconstruction*, where we aim at outputting also the non-corrupted ones.

Let \mathcal{M} be the subset $\mathcal{M} \subset \{1, \dots, d\}$, and $\mathbf{P}_{\mathcal{M}} \in \mathbb{R}^{d \times d}$ the diagonal matrix such as its i -th diagonal elements are 1 for all $i \notin \mathcal{M}$, and 0 otherwise, i.e. $\mathbf{P}_{\mathcal{M}}$ is the identity matrix with some elements equal to zero. When applied to a column vector $\mathbf{v} \in \mathbb{R}^d$, such a matrix has the property of zero-ing out all the vector elements corresponding to the indexes belonging to \mathcal{M} .

Thanks to $\mathbf{P}_{\mathcal{M}}$, and considering the mean squared error reconstruction loss, the empirical risk $\mathcal{R}_{\text{FR}}^f$ corresponding to the corrupted *full* features reconstruction can be written as:

$$\mathcal{R}_{\text{FR}}^f := \frac{1}{|\mathcal{V}_{\text{FR}}|} \sum_{\substack{i \in \mathcal{V}_{\text{FR}} \\ \mathcal{V}_{\text{FR}} \subseteq \mathcal{V}}} \|\mathbf{x}_i - f_{\text{FR}}^f(\mathbf{P}_{\mathcal{M}}\mathbf{x}_i; \vartheta_{\text{sh}}, \vartheta_{\text{FR}}^f)\|_2^2, \quad \text{with } f_{\text{FR}}^f := g_{\text{sh}} \circ g_{\text{FR}}^f, \quad \frac{\partial g_{\text{FR}}^f}{\partial \vartheta_{\text{sh}}} = 0. \quad (2)$$

Namely, f_{FR}^f acts as a *denoising* autoencoder, with the input corrupted by the matrix $\mathbf{P}_{\mathcal{M}}$ (for some arbitrary chosen \mathcal{M}), and as encoder the function g_{sh} also present in the main task. The decoder component is specified by g_{FR}^f , that depends on task specific parameters ϑ_{FR}^f only.

Now, we will consider the partial features reconstruction. Let $\mathbf{I}_{\mathcal{M}} \in \mathbb{R}^{|\mathcal{M}| \times d}$ be a rectangular matrix whose i -th row is a zero vector, with only a 1 at the j -th position, with $j \in \mathcal{M}$. When applied to a column vector $\mathbf{v} \in \mathbb{R}^d$, such a matrix has the property of selecting the vector elements corresponding to the indexes belonging to \mathcal{M} . Leveraging $\mathbf{I}_{\mathcal{M}}$, the empirical risk $\mathcal{R}_{\text{FR}}^p$ corresponding to the corrupted *partial* features reconstruction can be written similarly as:

$$\mathcal{R}_{\text{FR}}^p := \frac{1}{|\mathcal{V}_{\text{FR}}|} \sum_{\substack{i \in \mathcal{V}_{\text{FR}} \\ \mathcal{V}_{\text{FR}} \subseteq \mathcal{V}}} \|\mathbf{I}_{\mathcal{M}}\mathbf{x}_i - f_{\text{FR}}^p(\mathbf{P}_{\mathcal{M}}\mathbf{x}_i; \vartheta_{\text{sh}}, \vartheta_{\text{FR}}^p)\|_2^2, \quad \text{with } f_{\text{FR}}^p := g_{\text{sh}} \circ g_{\text{FR}}^p, \quad \frac{\partial g_{\text{FR}}^p}{\partial \vartheta_{\text{sh}}} = 0. \quad (3)$$

3.2.4 CORRUPTED EMBEDDINGS RECONSTRUCTION

Similarly to the previous task, the aim is to reconstruct ‘‘something’’ from a corrupted version of it. In this case (also called ER), the goal is to reconstruct the embeddings produced by some encoder, in order to make the embeddings resilient to noise. Also in this case, the corruption is achieved by zero-ing out some entries, distinguishing two methods: (i) *partial reconstruction*, where we aim at

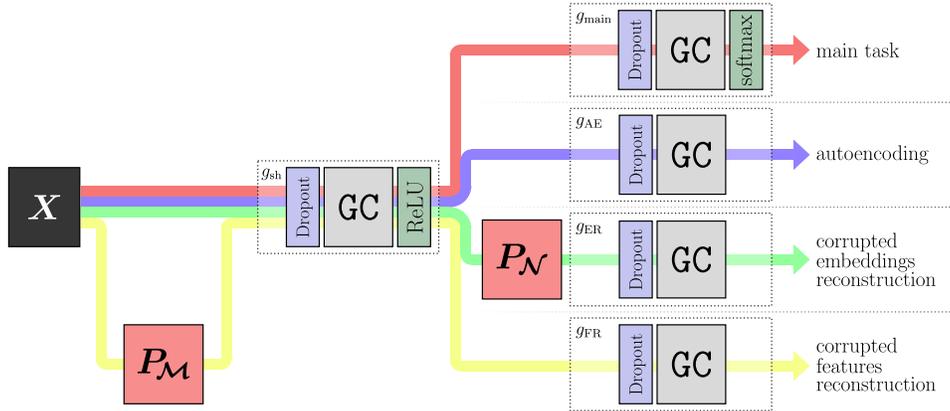


Figure 1: The drawing shows the architecture of the network described in Section 3.3, which is the one used in our experiments. The network is made of a shared encoder g_{sh} , followed by four heads, one devoted to the main task g_{main} , and the other three to each of the auxiliary tasks g_{AE} , g_{FR} , g_{ER} .

outputting the restored embeddings only; (ii) *full reconstruction*, where we aim at outputting the restored embeddings as well as the non-corrupted ones.

Considering the full reconstruction case, the mean squared error loss, and \mathcal{N} as the set containing the corrupted embedding index, we can write the empirical risk \mathcal{R}_{ER}^f corresponding to the corrupted *full* embeddings reconstruction can be written as:

$$\mathcal{R}_{ER}^f := \frac{1}{|\mathcal{V}_{ER}|} \sum_{\substack{i \in \mathcal{V}_{ER} \\ \mathcal{V}_{ER} \subseteq \mathcal{V}}} \|g_{sh}(x_i) - f_{ER}^f(x_i; \vartheta_{sh}, \vartheta_{ER}^f)\|_2^2, \quad \text{with } f_{ER}^f := g_{sh} \circ P_{\mathcal{N}} \circ g_{ER}^f, \quad \frac{\partial g_{ER}^f}{\partial \vartheta_{sh}} = 0. \quad (4)$$

Namely, we use the function g_{sh} also present in the main task as our encoder producing vertex embeddings, we corrupt them by zero-ing out some of them with $P_{\mathcal{N}}$, and we try to reconstruct them with the decoder defined by g_{ER}^f .

The *partial* reconstruction version can be obtained by leveraging $I_{\mathcal{N}}$ as made in Section 3.2.3.

3.3 THE FINAL NETWORK

The formalism introduced in the previous section does not make any assumptions about the mathematical form of the various encoder and decoder functions. This means that the functions g_{sh} , g_{main} , g_{AE} , g_{FR} , g_{ER} can be modeled by an arbitrary neural network, with any kind of layers, number of units, and activation functions. In the rest of the paper, we will restrict our analysis by using as a foundation block the GC layer (Kipf & Welling, 2017).

Our overall network, is composed of a shared encoder g_{sh} , and four output heads g_{main} , g_{AE} , g_{FR}^f , g_{ER}^f , one per task. Note that we are going to present explicitly only the *full reconstruction* variant of the network, since the *partial reconstruction* version can be easily derived.

The shared encoder g_{sh} is the same used in Kipf & Welling (2017), i.e. a dropout layer (Srivastava et al., 2014) with 50% dropout rate followed by GC layer made of 16 units and a ReLU activation function: $g_{sh} = \text{Dropout}(0.5) \circ \text{GC}_{16} \circ \text{ReLU}$. The main task classification head g_{main} is made of a dropout layer followed by a GC layer and a softmax activation, where the units of the GC layer depends on the number of classification classes ($g_{main} = \text{Dropout}(0.5) \circ \text{GC} \circ \text{softmax}$).

All the auxiliary task heads g_{AE} , g_{FR}^f , g_{ER}^f are made of a dropout layer followed by a GC layer made of 16 units, a ReLU activation function, another dropout layer, and a final GC layer with no activation function: $\text{Dropout}(0.5) \circ \text{GC}_{16} \circ \text{ReLU} \circ \text{Dropout}(0.5) \circ \text{GC}$. Note that the number of nodes of the last GC layer depends on the dimension of the vector that we have to reconstruct.

The resulting 4-head network is represented in Figure 1, and it is trained by minimization of the empirical risk given by $w_{main}\mathcal{R}_{main} + w_{AE}\mathcal{R}_{AE} + w_{FR}\mathcal{R}_{FR}^f + w_{ER}\mathcal{R}_{ER}^f$. The parameters ϑ_{sh} , ϑ_{main} , ϑ_{AE} , ϑ_{FR}^f , ϑ_{ER}^f are trained by stochastic gradient descent. As in previous works (Kipf & Welling, 2017), the gradient update is performed batch-wise, using the full dataset for every training iteration.

The proposed framework inherits memory and time complexity from the underlying layers we chose to use for each of the functions g_{sh} , g_{main} , g_{AE} , g_{FR}^f , g_{ER}^f . Thus, for the architecture here presented, it means a memory complexity linear in the number of edges for a sparse representation of the adjacency matrix A , and a time complexity linear in the number of edges (Kipf & Welling, 2017).

Note that by forcing some of w_{AE} , w_{FR}^f , and w_{ER}^f to be identically equal to zero, we can achieve with the same network architecture a settings where some of the auxiliary tasks are effectively deactivated, and in the limit case where all of them are equal to zero we recover the standard GCN.

4 EXPERIMENTAL RESULTS

4.1 DATASETS AND EXPERIMENTAL SETUP

We test our models on semi-supervised classification using the standard datasets Citeseer, Cora, and Pubmed (Sen et al., 2008). These are citation networks, where graph vertexes correspond to documents and (undirected) edges to citations. The vertex features are a bag-of-words representation of the documents. Each node is associated to a class label. The Cora dataset contains 2.708 nodes, 5.429 edges, 7 classes and 1.433 features per node. The Citeseer dataset contains 3.327 nodes, 4.732 edges, 6 classes and 3.703 features per node. The Pubmed dataset contains 19.717 nodes, 44.338 edges, 3 classes and 500 features per node.

For the training phase, we used the same setting adopted by Kipf & Welling (2017), which in turns follows the experimental setup of Yang et al. (2016). We allow for only 20 nodes per class to be used for training. However, since this is a semi-supervised setting, the training procedure can use the features of all the nodes. The test and validation sets comprise 1.000 and 500 nodes respectively. All the train/test/splitting used in the experiments are the ones used by Kipf & Welling (2017).

Since as already mentioned in Section 3.3 we opted to use the GC layers as foundational building blocks, the fair baseline comparison is with GCNs. We conducted two kinds of experiments. In the first round of tests, we focused on one-hidden layer architectures with 16 units, so that we could compare these results directly with the one presented in Kipf & Welling (2017). For each of the three datasets, the networks that we compared against the baseline and against each other are the ones showed in Table 1. Each of them are instances of the multi-head architecture presented in Section 3.3, with some of the heads deactivated. We used the same amount of dropout, L2 regularization, optimizer (Kingma & Ba, 2015), and learning rate used by Kipf & Welling (2017), unless otherwise stated. The sub-tasks weights appearing in Equation (1) have been tuned for all the networks by means of grid search. For the network *main* + FR, we tuned whether it is better the full or the partial reconstruction and we found the optimal number of reconstructed features by searching through the values 100, 200, 400, 800 for the Citeseer and Cora, and through the values 50, 100, 200 for Pubmed¹. The resulting optimal values have been used also for the networks *main* + AE + FR and *main* + AE + FR + ER. Similarly, for the network *main* + FR we tuned whether it is better the full or the partial reconstruction and we found the optimal number of reconstructed embeddings by searching through the values 2, 4, 8 for all the datasets. The resulting optimal values have been used also for the networks *main* + AE + ER and *main* + AE + FR + ER as well. Finally, we tuned the sets \mathcal{V}_{AE} , \mathcal{V}_{FR} , \mathcal{V}_{ER} by comparing the two limiting cases $\mathcal{V}_{\text{AE}} = \mathcal{V}_{\text{FR}} = \mathcal{V}_{\text{ER}} = \mathcal{V}_1$ and $\mathcal{V}_{\text{AE}} = \mathcal{V}_{\text{FR}} = \mathcal{V}_{\text{ER}} = \mathcal{V}$.

With Citeseer and Cora we trained for 5.000 epochs, while with Pubmed we trained for 2.500 epochs. During the training, the learning rate was reduced by a factor 10 if the multi-task loss on the validation set did not improve for 40 epochs in a row. In all the cases, we selected the best performing epoch on the validation set to assess the final performance on the test set. Each network has been trained and tested 10 times, each time with a different random weights initialization, and randomized reconstruction features and embeddings (if applicable).

In the second rounds of experiments we compared two-hidden layers architectures instead. The goal was to assess if the proposed framework allows to achieve good results even when increasing the network depth, thus making less relevant to tune the number of hidden layers in GCN architectures in order to reduce the “over-smoothing” phenomenon whenever the networks become deeper. The

¹As an example, considering Citeseer, which has 3.703 dimensions, we searched through 100, 200, 400, 800 reconstructed features, meaning that we zero-ed out 3.603, 3.503, 3.303, 2.903 number of features, respectively.

Table 1: The table shows the mean classification accuracy achieved on the test sets and the standard error of the mean for both the one-hidden layer and the two-hidden layers networks on Citeseer, Cora, and Pubmed. The acronyms AE, FR, and ER refer to the three auxiliary tasks described in Section 3.2. For completeness, in the one-hidden layer case we show also the results reported by Kipf & Welling (2017), which however do not report the standard errors.

Network	Accuracy					
	One-hidden layer			Two-hidden layers		
	Citeseer	Cora	Pubmed	Citeseer	Cora	Pubmed
GCN (Kipf & Welling, 2017)	70.3%	81.5%	79.0%	—	—	—
GCN (our 10 runs)	69.84 ± 0.22%	81.13 ± 0.13%	78.63 ± 0.41%	67.51 ± 0.52%	79.74 ± 0.54%	73.43 ± 3.37%
<i>main</i> + AE	71.06 ± 0.16%	82.17 ± 0.09%	78.97 ± 0.14%	67.65 ± 0.25%	81.10 ± 0.28%	78.14 ± 0.12%
<i>main</i> + FR	70.94 ± 0.13%	82.07 ± 0.15%	78.91 ± 0.19%	67.82 ± 0.32%	80.89 ± 0.29%	78.52 ± 0.15%
<i>main</i> + ER	70.42 ± 0.20%	81.83 ± 0.12%	79.33 ± 0.07%	67.75 ± 0.28%	80.37 ± 0.20%	72.39 ± 1.49%
<i>main</i> + AE + FR	71.14 ± 0.12%	82.13 ± 0.10%	78.92 ± 0.14%	67.80 ± 0.33%	80.85 ± 0.25%	78.14 ± 0.16%
<i>main</i> + AE + ER	69.95 ± 0.14%	82.05 ± 0.14%	79.15 ± 0.13%	67.77 ± 0.29%	79.76 ± 0.23%	78.06 ± 0.18%
<i>main</i> + AE + FR + ER	70.09 ± 0.20%	81.96 ± 0.13%	79.15 ± 0.13%	67.75 ± 0.32%	79.45 ± 0.13%	78.32 ± 0.11%

experimental setup is the same as in the first round of experiments, the only difference being that each hidden GC layer has been replaced by two 16 units GC layers.

4.2 RESULTS

The left side of Table 1 shows the results for the one-hidden layer architectures. It can be seen that all the proposed self-supervised multi-task frameworks achieve better mean accuracy than a plain GCN architecture. Moreover, the best performing architecture in each dataset is always one of ours, with a corresponding improvement in performance ranging from 0.70 to 1.30 percentage points. Interestingly, the best performing architecture always shows a smaller standard error compared to the plain GCN, thus exhibiting a more stable performance at different random weights initialization.

The stable performance verified with our one-hidden layer networks are confirmed in the second rounds of experiments (see the right side of the Table 1). Also in this case, the best performing architectures are those proposed in this paper. Notice that the improved stability is not only in terms of reduced variance of the accuracy metric, but also in the reduced need of hyper-parameters tuning. Indeed, in the two-layers GCN we had to carefully tune the learning rate to achieve results comparable to the one-layer networks. Such a procedure was not needed in any of the multi-head variants, thus reducing the effort to tune their hyper-parameters. These results suggest that the proposed framework helps to alleviate the “over-smoothing” problem affecting deep GCNs.

Finally, it is worth noticing that, the best performing architectures varies depending on the dataset at hand. In particular, the full network with 4 active heads was never the best performing candidate model (but still better than the baseline).

5 CONCLUSION

We introduced a general framework to improve semi-supervised classification performance on graph structured data by jointly learn multiple self-supervised auxiliary tasks in a multi-task fashion. To achieve this goal, we exploited three different self-supervised auxiliary tasks, i.e. (i) vertex features autoencoding; (ii) corrupted vertex features reconstruction; (iii) corrupted vertex embeddings reconstruction.

The theoretical framework is agnostic to the mathematical form of the various encoder and decoder functions, in terms of kind of layers, number of units, and activation functions. We restricted our analyses by selecting the GC layer as the fundamental building block.

The experiments we performed on standard datasets showed better performance with respect to GCNs. Contrary to GCNs, in the two-hidden layers scenario the proposed architectures showed no need to tune the learning rate. These considerations suggest that the proposed framework helps to alleviate the “over-smoothing” problem affecting deep GCNs (Li et al., 2018).

The generality of the proposed framework allows for further developments in the direction of: (i) replacing the GC layer with other neural layers devoted to deal with graph structured data (e.g. Graph Attention Networks (Veličković et al., 2018)); and (ii) exploring additional auxiliary tasks.

REFERENCES

- Rie Kubota Ando and Tong Zhang. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6(Nov):1817–1853, 2005.
- Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Multi-task feature learning. In *Advances in neural information processing systems*, pp. 41–48, 2007.
- Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Convex multi-task feature learning. *Machine learning*, 73(3):243–272, 2008.
- Bart Bakker and Tom Heskes. Task clustering and gating for bayesian multitask learning. *Journal of Machine Learning Research*, 4(May):83–99, 2003.
- Edwin V Bonilla, Kian M Chai, and Christopher Williams. Multi-task gaussian process prediction. In *Advances in neural information processing systems*, pp. 153–160, 2008.
- Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. In *ICLR*, 2013.
- Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- Jianhui Chen, Ji Liu, and Jieping Ye. Learning incoherent sparse and low-rank patterns from multiple tasks. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(4):1–31, 2012.
- Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, pp. 794–803. PMLR, 2018.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *NIPS*, pp. 3844–3852, 2016.
- Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pp. 2070–2079. IEEE Computer Society, 2017.
- Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pp. 1422–1430, 2015.
- Theodoros Evgeniou, Charles A Micchelli, and Massimiliano Pontil. Learning multiple tasks with kernel methods. *Journal of machine learning research*, 6(Apr):615–637, 2005.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.
- Michele Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, 2: 729–734 vol. 2, 2005.
- Aditya Grover and Jure Leskovec. Node2vec: Scalable feature learning for networks. In *ACM SIGKDD*, pp. 855–864. ACM, 2016.
- Aditya Grover, Aaron Zweig, and Stefano Ermon. Graphite: Iterative generative modeling of graphs. volume 97 of *Proceedings of Machine Learning Research*, pp. 2434–2444, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v97/grover19a.html>.
- Will Hamilton, Zhitaoying, and Jure Leskovec. Inductive representation learning on large graphs. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 1024–1034. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6703-inductive-representation-learning-on-large-graphs.pdf>.

- David K. Hammond, Pierre Vandergheynst, and Rémi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Comput. Harmonic Analysis*, 30 (2):129–150, 2011.
- Ali Jalali, Sujay Sanghavi, Chao Ruan, and Pradeep K Ravikumar. A dirty model for multi-task learning. In *Advances in neural information processing systems*, pp. 964–972, 2010.
- Eric Jang, Coline Devin, Vincent Vanhoucke, and Sergey Levine. Grasp2vec: Learning object representations from self-supervised grasping. In *Conference on Robot Learning*, pp. 99–112, 2018.
- Tsuyoshi Kato, Hisashi Kashima, Masashi Sugiyama, and Kiyoshi Asai. Multi-task learning via conic programming. In *Advances in Neural Information Processing Systems*, pp. 737–744, 2008.
- Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7482–7491, 2018.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4681–4690, 2017.
- Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. *arXiv preprint arXiv:1801.07606*, 2018.
- Sijin Li, Zhi-Qiang Liu, and Antoni B Chan. Heterogeneous multi-task learning for human pose estimation with deep convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 482–489, 2014.
- Franco Manessi, Alessandro Rozza, and Mario Manzo. Dynamic graph convolutional networks. *Pattern Recognition*, 97:107000, 2020.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space. *arXiv e-prints*, art. arXiv:1301.3781, January 2013.
- Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3994–4003, 2016.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Blaise Thomson, Milica Gasic, Pei-Hao Su, David Vandyke, Tsung-Hsien Wen, and Steve Young. Multi-domain dialog state tracking using recurrent neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pp. 794–799, 2015.
- Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pp. 69–84. Springer, 2016.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2536–2544, 2016.
- Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *ACM SIGKDD*, pp. 701–710. ACM, 2014.
- Ting Kei Pong, Paul Tseng, Shuiwang Ji, and Jieping Ye. Trace norm regularization: Reformulations, algorithms, and multi-task learning. *SIAM Journal on Optimization*, 20(6):3465–3489, 2010.

- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training, 2018.
- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Trans. Neural Networks*, 20(1):61–80, 2009.
- Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Ke Sun, Zhouchen Lin, and Zhanxing Zhu. Multi-stage self-supervised learning for graph convolutional networks on graphs with few labeled nodes. In *AAAI*, pp. 5892–5899, 2020.
- Sebastian Thrun and Joseph O’Sullivan. Discovering structure in multiple learning tasks: The tc algorithm. In *ICML*, volume 96, pp. 489–497, 1996.
- Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. In *ICLR (Poster)*, 2019.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>.
- Jiawei Wu, Xin Wang, and William Yang Wang. Self-supervised dialogue learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3857–3867, 2019.
- Huan Xu, Yujiu Yang, Liangwei Wang, and Wenhua Liu. Node classification in social network via a factor graph model. In *PAKDD*, pp. 213–224, 2013.
- Ya Xue, Xuejun Liao, Lawrence Carin, and Balaji Krishnapuram. Multi-task learning for classification with dirichlet process priors. *Journal of Machine Learning Research*, 8(Jan):35–63, 2007.
- Yongxin Yang and Timothy M Hospedales. Trace norm regularised deep multi-task learning. *arXiv preprint arXiv:1606.04038*, 2016.
- Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*, pp. 40–48. PMLR, 2016.
- Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pp. 649–666. Springer, 2016.
- Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1058–1067, 2017.
- Y. Zhang and Qiang Yang. A survey on multi-task learning. *ArXiv*, abs/1707.08114, 2017a.
- Yu Zhang and Qiang Yang. A survey on multi-task learning. *arXiv preprint arXiv:1707.08114*, 2017b.
- Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Facial landmark detection by deep multi-task learning. In *European conference on computer vision*, pp. 94–108. Springer, 2014.
- Wenliang Zhong and James Kwok. Convex multitask learning with flexible task clusters. *arXiv preprint arXiv:1206.4601*, 2012.