

# Reducing Oversmoothing in Graph Neural Networks by Changing the Activation Function

Dimitrios Kelesis<sup>a,b,\*</sup>, Dimitrios Vogiatzis<sup>b,c</sup>, Georgios Katsimpras<sup>b</sup>, Dimitris Fotakis<sup>a</sup> and Georgios Paliouras<sup>b</sup>

<sup>a</sup>School of Electrical Engineering and Computer Science, National Technical University of Athens, Greece

<sup>b</sup>National Center for Scientific Research "Demokritos", Greece

<sup>c</sup>American College of Greece, Deree

**Abstract.** The performance of Graph Neural Networks (GNNs) deteriorates as the depth of the network increases. That performance drop is mainly attributed to oversmoothing, which leads to similar node representations through repeated graph convolutions. We show that in deep GNNs the activation function plays a crucial role in oversmoothing. We explain theoretically why this is the case and propose a simple modification to the slope of ReLU to reduce oversmoothing. The proposed approach enables deep networks without the need to change the network architecture or to add residual connections. We verify the theoretical results experimentally and further show that deep networks, which do not suffer from oversmoothing, are beneficial in the presence of the "cold start" problem, i.e. when there is no feature information about unlabeled nodes.

## 1 Introduction

Graph Neural Networks (GNNs) utilize message passing or neighborhood aggregation schemes to extract representations for nodes and their neighborhoods. GNNs have achieved good results on a variety of graph analytics tasks, such as node classification [25, 7], link prediction [15, 24] and graph classification [13]. As a result, they play a key role in graph representation learning. One of the most prominent GNN models is the Graph Convolutional Network (GCN) [10], which creates node representations, by averaging the representations (embeddings) of the node's immediate neighbors. Several studies have shown that, the performance of GCNs deteriorates, when their architecture becomes deeper [14].

The success of deep CNNs on many tasks, like image classification, naturally led to several attempts towards building deep GNNs for node classification [10, 14, 23]. Increasing the model's depth (and the number of parameters it has), would allow more accurate representational learning to occur. Most of the existing approaches have failed to develop a sufficiently deep architecture that achieves good performance. Therefore, there is a need to design new models which can efficiently scale to a large number of layers. The aim of this paper is to investigate the contributing factors, that compromise the performance of deep GNNs and develop a method to address them.

The performance drop of deep GNNs, is associated with several factors, including vanishing gradients, overfitting, as well as the phenomenon called oversmoothing [14, 23, 12, 22]. Oversmoothing has

been shown to be associated with graph convolution, a type of Laplacian operator. Li et al. [14] proved that applying that operator repeatedly, makes node representations converge to a stationary point. At that point, all of the initial information (i.e. node features' information) is lost through the Laplacian smoothing. Consequently, oversmoothing hurts the performance by making node features indistinguishable across different classes.

Several attempts have been made to reduce oversmoothing, utilizing a variety of techniques. Chen et al. [4] have introduced residual connections in order to 'short circuit' initial node features to subsequent layers. Using residuals to enhance the influence of the close neighborhood of a node, may partially address oversmoothing, but can also hurt performance in large graphs, as shown in the results of [5]. Rong et al. [20] and Huang et al. [8] attempted to reduce oversmoothing by altering graph topology. The latter methods aimed to reduce oversmoothing, while avoiding to inject initial node information and without forcing the model to focus on the local neighborhood of each node.

In this work, we address the oversmoothing problem in deep GNNs, without using residual connections. We provide what is to the best of our knowledge the first study regarding the role of the activation function and the learning rate per layer of the model to oversmoothing and propose a new method to address the problem. We confirm our hypotheses both experimentally and theoretically. The new method is shown to prevent node embeddings from converging to the same point, thus leading to better node representations. We summarize our main contributions as follows.

•**Role of Activation Function in Oversmoothing:** We prove theoretically the connection between the activation function and oversmoothing. In fact, we show the relation between the slope of ReLU and the singular values of weight matrices, which are known to be associated with oversmoothing [18, 3]. We have also verified our theoretical results experimentally.

•**Role of Learning Rate in Oversmoothing:** Our analysis on the effect of the slope of ReLU to oversmoothing has a direct extension to the learning rates used per layer of the network. We conducted further experiments to study the effect of tuning the learning rates, showing that this approach could also reduce oversmoothing, but it is less practical.

•**The power of Deep GNNs:** We have performed extensive experiments using up to 64-layer networks, tackling oversmoothing with

\* Corresponding Author. Email: dkelesis@iit.demokritos.gr.

the proposed method. We further show the benefits that such deep GNNs can provide in the presence of reduced information, such as in a “cold start” situation, where node features are available only for the labeled nodes in a node classification setting.

## 2 Notations and Preliminaries

### 2.1 Notations

In order to illustrate the problem of oversmoothing, we consider the task of semi-supervised node classification on a graph. The graph to be analysed is  $G(V, E, X)$ , with  $|V| = N$  nodes  $u_i \in V$ , edges  $(u_i, u_j) \in E$  and  $X = [x_1, \dots, x_N]^T \in R^{N \times C}$  denotes the initial node features (every node has a feature vector of dimensionality  $C$ ). The edges form an adjacency matrix  $A \in R^{N \times N}$  where edge  $(u_i, u_j)$  is associated with element  $A_{i,j}$ .  $A_{i,j}$  can take arbitrary real values indicating the weight (strength) of edge  $(u_i, u_j)$ . Node degrees are represented through a diagonal matrix  $D \in R^{N \times N}$ , where each element  $d_i$  represents the sum of edge weights connected to node  $i$ . During training, only the labels of a subset  $V_l \in V$  are available. The task is to learn a node classifier, that predicts the label of each node using the graph topology and the given feature vectors.

**GCN**, originally proposed by [10], utilizes a feed forward propagation as:

$$H^{(l+1)} = \sigma(\hat{A}H^{(l)}W^{(l)}) \quad (1)$$

where  $H^{(l)} = [h_1^{(l)}, \dots, h_N^{(l)}]$  are node representations (or hidden vectors or embeddings) at the  $l$ -th layer, with  $h_i^{(l)}$  standing for the hidden representation of node  $i$ ;  $\hat{A} = \hat{D}^{-1/2}(A + I)\hat{D}^{-1/2}$  denotes the augmented adjacency matrix after self-loop addition, where  $\hat{D}$  corresponds to the degree matrix;  $\sigma(\cdot)$  is a nonlinear element-wise function, i.e. the activation function, ReLU; and  $W^{(l)}$  is the trainable weight matrix of the  $l$ -th layer. A generic ReLU function with slope  $\alpha$  is defined as  $\text{ReLU}(x) = \max(\alpha \cdot x, 0)$ .

### 2.2 Understanding Oversmoothing

GNNs achieve state-of-the-art performance in a variety of graph-based tasks. Despite their success, models like GCN [10] and GAT [21] experience a performance drop, when stacking multiple layers. To a large extent, this is attributed to oversmoothing due to repeated graph convolutions. In their analysis, [14] showed that graph convolution is a special form of Laplacian smoothing. In fact, they proved that the new representation of each node is formed by a weighted average of its own representation and that of its neighbors. This mechanism allows the node representations within each (graph) cluster, i.e. highly connected group of nodes, to become more similar and improves the performance on semi-supervised tasks on graphs. When stacking multiple layers, the smoothing operation is repeated multiple times leading to oversmoothing of node representations, i.e., the hidden representations of all nodes become similar, resulting in information loss.

### 2.3 Deep GNN limitations

Therefore, oversmoothing leads node representations to converge to a fixed point as the network’s depth increases [14]. At that point, node representations contain information relevant to the graph topology and disregard the input features. Oono & Suzuki [18] have generalized the idea in [14] by taking into consideration that the ReLU activation function maps to a positive cone. They explain oversmoothing as the convergence to a subspace, instead of the convergence to a

fixed point. A similar approach is presented in [3], offering a different perspective to the oversmoothing problem using Dirichlet energy. We now look closer to the proposal of [18], which we will use as a basis for our analysis. The main result in [18] is the definition of the convergence speed towards a subspace  $M$ , where the distance between node representations tends to zero. We denote as  $d_M(X)$  the distance between the feature vector  $X$  and the subspace, where the oversmoothing is prevalent. When that distance approaches to zero it indicates that node representations have been over-smoothed. For this distance, [18] prove an interesting property.

**Theorem 1 (Oono & Suzuki [18])** *Let the largest singular value of the weight matrix  $W_{lh}$  be  $s_{lh}$  and  $s_l = \prod_{h=1}^{H_l} s_{lh}$ , where  $W_{lh}$  is the weight matrix of layer  $h$  and  $H_l$  is the network’s depth, following the notation of the original paper. Then it holds that  $d_M(f_l(X)) \leq s_l \lambda d_M(X)$  for any  $X \in R^{N \times C}$ , where  $f(\cdot)$  is the forward pass of a GNN layer (i.e.  $\sigma(AW_{lh})$ ).*

Theorem 1 indicates, that the deeper the network the smaller the distance of node representations from the subspace  $M$ . If the maximum singular values are small then node representations asymptotically approach  $M$ , for any initial values of node features. Extending the above theorem, the authors conclude to the following estimate about the speed of convergence to the oversmoothing subspace.

**Corollary 2 (Oono & Suzuki [18])** *Let  $s = \sup_{l \in N_+} s_l$ . then  $d_M(X^{(l)}) = O((s\lambda)^l)$ , where  $l$  is the layer number and if  $s\lambda < 1$  the distance from oversmoothing subspace exponentially approaches zero. Where  $\lambda$  is the smallest non-zero eigenvalue of  $I - \hat{A}$ .*

According to the authors, sufficiently deep GCN will certainly suffer from oversmoothing under some conditions (details can be found in [18]). We build upon this result, aiming to develop a consistent approach that reduces oversmoothing and enables deep architectures.

### 2.4 Normalization in deep neural networks

Of relevance to oversmoothing are also methods that perform normalization for deep neural networks [9, 1]. In such methods, the output of each neuron is normalized, in order to keep a portion of the initial feature variance. Related work proposed Self Normalized Networks (SNN), that use a different activation function (SeLU) [11]. These models perform self normalization inside each neuron, in order to keep the variance between consecutive layers stable. They have managed to enable deep Fully Connected models and achieve good performance.

A recent work regarding normalization in GNNs is the Paimorm method [26], which aims to keep constant the total pairwise distance between node representations. Compared to SNNs, Paimorm performs normalization needing a constant value as hyper-parameter determined per dataset, while SNNs normalize the output utilizing their activation function, which can be used in a family of neural networks. The SeLU activation function has a saturating region to reduce data variance and a slope slightly greater than 1 in order to increase data variance, when needed.

## 3 Understanding and dealing with oversmoothing

Based on the mathematical definition of oversmoothing by [18], in this section we establish a connection with the training process of

GNNs. In particular, we analyse the role of the activation function and the learning rate and we propose modified versions of GNNs to address the issue.

### 3.1 Theoretical Analysis

We start by establishing the connection between oversmoothing and variance reduction of node representations. Consider oversmoothing as the convergence to a subspace of the feature space, where node representations are almost the same. In that particular subspace, the initial variance of feature vectors has been massively reduced. Using ReLU, this unwanted oversmoothing effect is ‘harsh’, due to the mapping of negative input to zero. We use the term ‘harsh’ to indicate the smoothing case, where representations are mapped to a single point (i.e. zero), instead of converging to a subspace.

We will now show how the slope of the ReLU activation function affects oversmoothing, starting with two important assumptions: (a) the non-exploding gradients, and (b) independence of ReLU’s probability not to output zero. We use these assumptions to extract bounds on the weights of the network and then use these bounds to determine the relationship between the largest singular value and the elements of each weight matrix. Given that relationship, we connect our analysis with existing literature regarding oversmoothing through Theorem 1.

**Assumption 3** For each layer  $l$  there exists a number  $G_l$  which is the upper bound to the gradients of the output of the subsequent layer  $(l+1)$  with respect to the weight elements of  $W^{(l)}$ , i.e.

$$\frac{dO^{(l+1)}}{dw_{old_{i,j}}^{(l)}} \leq G_l, \quad \forall w_{old_{i,j}}^{(l)} \quad (2)$$

where  $O^{(l+1)}$  is the output of  $(l+1)$ -th layer.

This assumption needs to hold in all cases, in order to avoid the exploding gradient case, that would not allow the learning process to converge.

**Assumption 4** The probability, that the output of the ReLU function does not equal zero at layer  $l$  is independent of the outputs of the previous layers, in a feed-forward ReLU neural network, i.e.

$$P(\text{ReLU}(h^{(l)}) \neq \vec{0} | \forall h^{(j)}, j < l) \leq p, \quad p \in [0, 1] \quad (3)$$

where  $h^{(l)}$  is the node representation at layer  $l$ , ReLU is applied piece-wise on that representation and  $\vec{0}$  is the all zero vector.

Following the results of [16] we get that  $P(\text{ReLU}(h^{(l)}) \neq \vec{0} | \forall h^{(j)}, j < l) = 1/2$  for feed-forward networks (FFs), on the condition that the output of the previous layers is positive. That condition does not necessary hold in GNNs, due to the role of the adjacency matrix in the aggregation scheme. In FFs, when the output of a layer is zero, so are the outputs of all subsequent layers. In GNNs there might be cases, where the output of layer  $l$  (a node representation) is zero but the aggregation using the adjacency matrix produces a non-zero representation in layer  $(l+1)$ . Apart from this difference, GNNs can be considered a special case of FFs, with the proper wiring due to the adjacency matrix aggregation. Therefore, the above assumption holds also for GNNs.

**Lemma 5** For a network of depth  $dep$  the total gradient reaching the  $l$ -th layer (i.e.,  $\frac{dJ}{dw_{old_{i,j}}^{(l)}}$ ) in order to update  $W^{(l)}$  is bounded by:

$$\frac{dJ}{dw_{old_{i,j}}^{(l)}} \leq \alpha^{(dep-l)} \cdot G_l \quad (4)$$

Where  $J$  is the model’s loss function (i.e. Cross Entropy),  $\alpha$  stands for the ReLU slope and  $G_l$  is the upper bound of gradients of the output of the subsequent layer  $(l+1)$  with respect to the weight elements of  $W^{(l)}$ .

The proof of Lemma 5 is shown in Appendix A and is based on the chain rule for backpropagation. Given the derivative of the ReLU function and the upper bound  $G_l$  of Assumption 3, we can derive the bound of Equation 4. To compute the gradient with respect to weight element  $w_{i,j}^{(l)}$  of layer  $l$ , we repeatedly differentiate nested ReLU functions, leading to a product of their slopes (or zero). In the final differentiation step, we get the gradient of the output of  $(l+1)$ -th layer with respect to  $w_{i,j}^{(l)}$ , which is bounded by  $G_l$ .

**Lemma 6** While model’s loss, through gradients, flows backwards, some weight elements do not receive updates, because we have dying ReLUs (i.e. ReLUs that output zero) [16]. The total number of weight elements getting updated at layer  $l$  is bounded by:

$$\#\{w_{i,j}^{(l)}\} \leq p^{(dep-l)} \cdot d^2 \quad (5)$$

where  $d$  is the largest of the two dimensions of  $W^{(l)}$ , i.e. there are at most  $d^2$  elements in  $W^{(l)}$ , if it is a square matrix.

The proof of Lemma 6 appears in Appendix B. Given Lemma 5, the gradient flowing backwards, with respect to a weight element  $w_{i,j}^{(l)}$ , contains a product of ReLU derivatives. In order for  $w_{i,j}^{(l)}$  to get an update, all ReLU derivatives need to be non-zero, otherwise the gradient will be zeroed and  $w_{i,j}^{(l)}$  will not be updated. Additionally, due to Assumption 4, weight elements located at the lower weight matrices tend to receive fewer updates. This is due to the fact, that the further the gradient flows backwards the more ReLU derivative factors appear in it. Based on Assumption 4, the probability of all of them to be non-zero decreases.

Regarding the singular values (denoted by  $s_l(\cdot)$ ) of a weight matrix  $W^{(l)}$  at the  $l$ -th layer, the largest of them is given by:

$\max(s_l(W^{(l)})) = \|W^{(l)}\|_2 \leq \|W^{(l)}\|_F = \sqrt{\sum |w_{i,j}^{(l)}|^2}$ , where  $\|\cdot\|_F$  is the Frobenius norm and  $\|\cdot\|_2$  is the spectral norm. Using the general weight updating rule (i.e.,  $w_{new_{i,j}}^{(l)} = w_{old_{i,j}}^{(l)} + \eta \cdot \frac{dJ}{dw_{old_{i,j}}^{(l)}}$ ), where  $J$  is the model’s loss and  $\eta$  is the learning rate, we arrive at the main theorem of this work.

**Theorem 7** The upper bound of the largest singular value of the weight matrix  $W^{(l)}$  at layer  $l$  for a GNN model, utilizing a ReLU activation function, depends on the slope of the function. That bound is given per layer and shows the effect of each iteration of updates on the weight matrix. We denote with  $W_{old}$  and  $W_{new}$  the weight matrices before and after the update during an iteration of the training process respectively.

$$\max(s_l(W_{new}^{(l)})) \leq \|W_{old}^{(l)}\|_F + \sqrt{3} \cdot p^{\left(\frac{dep-l}{2}\right)} \cdot d \cdot \alpha^{(dep-l)} \cdot B_l \quad (6)$$

where  $B_l = \eta \cdot G_l$ ,  $dep$  is the network’s depth,  $d$  is the largest dimension of the  $W^{(l)}$  matrix,  $p$  is the upper bound of the probability of ReLU not to output zero and  $\alpha$  is ReLU’s slope.

The proof of Theorem 7 appears in Appendix C and uses the upper bound of the largest singular value by the Frobenius norm, expanded according to the weight update rule. Separating weight elements into

two sets (updated and not updated) we identify the gradient values responsible for the weight updates. These values are bounded (Lemma 5), as is the number of updated elements (Lemma 6), leading to Equation 6. The resulting upper bound shows the connection of the slope of ReLU and the upper bound of the probability of ReLU not to output zero with the largest singular value of the weight matrix. That largest singular value is connected with the oversmoothing problem and as we mentioned above could act as a resource to reduce it.

**Theorem 8** *The lower bound of the largest singular value of the weight matrix  $W^{(l)}$  at layer  $l$  for a GNN model, utilizing the ReLU activation function, depends on the slope of the function. That bound is defined per layer and shows the effect of each iteration of updates on the weight matrix. We denote with  $W_{old}$  and  $W_{new}$  the weight matrices before and after the update that happens during an iteration of the training process respectively.*

$$\max(s_l(W_{new}^{(l)})) \geq \sqrt{2\sqrt{3} \cdot \max(w_{old_{i,j}}^{(l)}) \cdot B_l^l \cdot \alpha^{\frac{dep-l}{2}}} \quad (7)$$

where  $B_l^l = \eta \cdot L_l$ ,  $L_l$  is the lower bound to the gradient with respect to the output of the subsequent layer  $\left(\frac{dO^{(l+1)}}{w_{old_{i,j}}^{(l)}} \geq L_l\right)$ ,  $dep$  is the network's depth and  $\alpha$  is ReLU's slope.

The proof of Theorem 8 appears in Appendix C and is analogous to the one used in Theorem 7. Theorem 8 indicates the effect of ReLU's slope to the lower bound of the largest singular value of the weight matrix. From Equation 7 we observe that, increasing the slope of ReLU results in an increment of the lower bound of the maximum singular value of the weight matrix. Therefore, a larger ReLU slope yields greater largest singular value, which in turn reduces the oversmoothing effect.

### 3.2 Alleviating Oversmoothing

We transfer the idea of SeLU to GNNs, focusing on the part that increases the variance, because the repetition of the Laplacian operator acts as a variance reducer. In order to avoid oversmoothing in deep GNNs using this approach, we need to identify a 'sweet' spot, where variance reduction from graph convolution is counteracted as needed by the slope of the activation function. Typically in GNNs, ReLU is used with a slope value  $\alpha = 1$ . As a result, the second exponential factor in Equation 6 can be ignored (it is always equal to 1). This pushes the bound for the largest singular value of weight matrices towards a fixed low value, when the architecture of the network gets deeper ( $dep$  increases), and especially at the lower layers (small  $l$ ) of the network. This is due to the first exponential factor of Equation (6) converging to zero for large values of  $dep$ .

The restriction of the largest singular values to low values, increases the speed of convergence to the oversmoothing subspace, as stated in Corollary 2 from [18]. According to the corollary, the speed depends on  $\lambda$  and  $s$ . The former parameter ( $\lambda$ ) is a property of the data (largest eigenvalue of the adjacency matrix), while the latter ( $s$ ) is the product of the largest singular values of the weight matrices.

Realizing the importance of the slope of ReLU in the training process, we move on to propose a simple modification of the activation function that reduces oversmoothing. Following [16], we proceed our analysis with  $p = 1/2$  (see Assumption 4), which simplifies calculations. In particular, we observe that a slope of 2 makes the second exponential factor prevail over the first one, leading to a new combined factor of  $2^{\left(\frac{dep-l}{2}\right)}$ . This in turn increases the upper bound for

the largest singular value, restricting the influence of the layer index ( $l$ ) and depth ( $dep$ ) in Equation (6). It is worth noting, that the upper bound should not be constant across all layers (which could be achieved by setting a corresponding value for the slope), because different layers of the network have different roles and the goal of a GNN is to bring intra-class representations close, while keeping inter-class representations apart. Moreover, an increased value of the slope pushes the lower bound of the largest singular value, as shown in Equation 7.

Using the results of [18] we can further connect the choice of slope for the activation function to the speed of convergence to the oversmoothing subspace. The proposed method increases the bound for each largest singular value, which in turn decreases the speed of convergence to the oversmoothing subspace, according to Corollary 2.

### 3.3 Modifying the slope of ReLU: Limitations

Oono & Suzuki [18] have proved that any deep enough GCN-like GNN will eventually reach the oversmoothing subspace. In theory, this also holds for our method, when the GNN gets very deep. Could this be avoided by increasing the slope of ReLU further? Unfortunately not, as we cannot make the slope too large. It is known, that by increasing the slope of ReLU too much, one may face the problem of exploding gradients [2], that impedes the learning procedure. In fact, if we increase the slope of ReLU too much, the lower bound in Equation 7 over increases, the upper bound in Equation 6 becomes too loose and the performance degrades. Our proof suggests, that a slope equal to 2 avoids oversmoothing and our experiments have shown that it also avoids the exploding gradients problem. In fact, the choice of 2 is not exact, but that particular value seems to work well as depth increases. Our experiments have shown, that the proposed method is not sensitive on the exact slope and that values around 2 also manage to reduce oversmoothing.

### 3.4 Modifying the learning rate, instead of the slope of ReLU

Instead of changing the slope of ReLU, we could opt to modify the Learning Rate (LR) per layer in order to increase the lower bound in Equation 7 and the upper bound in Equation 6. In particular, a different learning rate ( $\eta$ ) per layer leads to a different  $B_l$  ( $= \eta \cdot G_l$ ) per layer, counteracting the problematic first exponential factor of Equation 6.

However, determining the right LR per layer is a non-trivial task. Intuitively, LR should be larger in the lower layers and smaller in the upper ones. Keeping the slope equal to 1, the first problematic exponential term is smaller (note that  $p < 1$ ) in the lower layers, reducing the value of the upper bound. To avoid this, one needs to tune LR carefully. Additionally, LR affects heavily the learning process and large LR values might lead to oscillations and poor learning performance. Modifying the slope and the learning rate combined led to negligible improvements. Preliminary results with modified LR values are shown in Appendix D.

### 3.5 Why do we need deep GNNs?

Oversmoothing is only a problem for deep GNNs, leading to the question of whether and when deep GNNs are really needed. Most of the existing benchmark datasets for various graph analytics tasks do not seem to justify the need for deep networks. Due to their homophilic nature, useful information for each node resides in its close

neighbors (usually 2 or 3 hops away). A task where deeper architectures could be needed is the “cold start” problem; namely the situation where many node features are missing. The problem of missing features is called “cold start”, because it resembles the situation of a new product/user arriving at a recommender system. The system has no prior knowledge about the new arrival, yet has to make some recommendations. In this scenario, the hope is that deeper GNNs could recover features from more distant nodes, in order to create informative representations.

## 4 Experiments

### 4.1 Experimental Setup

**Datasets:** Aligned to most of the literature, we use four well-known benchmark datasets: *Cora*, *CiteSeer*, *Pubmed* and the less homophilic dataset *Texas*. The statistics of the datasets are reported in the Appendix. For *Cora*, *CiteSeer* and *Pubmed* we use the same data splits as in [10], where all nodes except the ones used for training and validation are used for testing. For *Texas* we use the same splits as in [19].

**Models:** We utilize two different GNNs as our base models for the proposed methodology; namely GCN [10] and GAT [21]. We compare the results of these models with and without the modification of the slope of ReLU for varying number of layers. We do not compare against methods utilizing residual connections, because the aim of this work is to show that oversmoothing could be avoided without “short-circuiting” initial information to latter layers. We use residual GNNs as a baseline in the Extended Experiments section, Appendix E.

**Hyperparameters:** We set the number of hidden units (of each layer) of GCN and GAT to 128 for both models across all datasets. The  $L_2$  regularization is set with penalty  $5 \cdot 10^{-4}$  for both models and learning rate of  $10^{-3}$ . We vary the depth between 2 and 64 layers. The number of attention heads for GAT is set to 1.

**Configuration:** Each experiment is run 10 times and we report the average performance over these runs. We train all models for 200 epochs using Cross Entropy as a loss function. For our experiments we used an RTX 3080ti GPU.

### 4.2 Experimental Results

#### Reducing oversmoothing:

Table 1 presents the classification performance of the two base models (GCN and GAT) on all four datasets, with and without the modified slope of ReLU (called Slope2GNN) for varying number of layers. Additionally, it presents results using SeLU, instead of ReLU, since it also modifies the slope of the function, while additionally normalizing node representations, as explained in section 2.4. To the best of our knowledge, this is the first time SeLU is used with GNNs. Table 1 shows that methods with a modified activation function reduce oversmoothing significantly. As expected, baselines maintain a high performance for shallow architectures in homophilic datasets, where oversmoothing is not a problem. The modified activation functions (Slope2GNN and SeLU) consistently improve the test accuracy as the number of layers ( $\#L$ ) increases. Note that even the proposed method suffers performance degradation for very deep GNNs, due to the unavoidable nature of oversmoothing [18]. GCN and GAT average over information in the close neighborhood of the node, assuming homophily, which is not the case in the *Texas* dataset. Interestingly, low homophily in this dataset makes oversmoothing appear at a larger depth, compared to other datasets. This is because connected

nodes do not have similar representations and the model needs more Laplacian operations to oversmooth them.

#### The effect of deeper networks on different activation functions:

Focusing on the use of SeLU, although it seems to make the models resistant to oversmoothing, the effect seems to be lost for GCNs with 64 layers. This may be due to the sensitivity of the function on the choice of slope value, which is lower than 2. Moreover, the saturating region of SeLU, used to reduce data variance, might act in favor of oversmoothing. Finally, SNNs aim to keep the variance of the model stable, ignoring the effect of the Laplacian smoothing in GNNs.

To further examine the behavior of each method, as the number of layers increases, in Figure 1, we present a more detailed progression of the accuracy of GCNs, as the networks get deeper on the *Cora* dataset. The effect of oversmoothing and its reduction by the two methods that modify ReLU is very apparent here. Furthermore, the simpler Slope2GNN approach seems to be more resistant to the increase in the depth of the network.

#### The need for deep networks:

Having shown the benefit of using the modified activation functions in deep GNNs, we move to a set of experiments that aims to highlight also the value of using such deep architectures. In particular, we report our experimental results in the “cold start” scenario, as described in section 3.5. The cold-start datasets that we use in these experiments are generated by removing feature vectors from unlabeled nodes and replacing them with all-zero vectors. For each combination of GNN and activation function, we present in Table 2 the best performance achieved and the depth at which the model attains that performance. The main observation in these results is the improvement in terms of accuracy with the use of deeper GNNs. This improvement is only attainable with the modified activation functions that reduce the effect of oversmoothing. The vanilla versions of GCN and GAT cannot go further than the performance of their shallow versions. Worth-mentioning is also the fact that the simple Slope2GNN approach manages to benefit from the deeper architectures even in the hard *Texas* dataset. This is not the case for SeLU, which achieves its best scores with very shallow networks. This may indicate the need to tune the slope of the activation function, as discussed above.

#### Slope sensitivity:

Finally, we performed an experiment to understand the effect of different ReLU slope values on oversmoothing. Figure 2 presents results with GCN on the *Cora* dataset, showing that the proposed method is not sensitive to the exact slope value. Slope values around 2 seem also capable of reducing oversmoothing. Experiments on other datasets are included in Appendix E.

## 5 Related Work

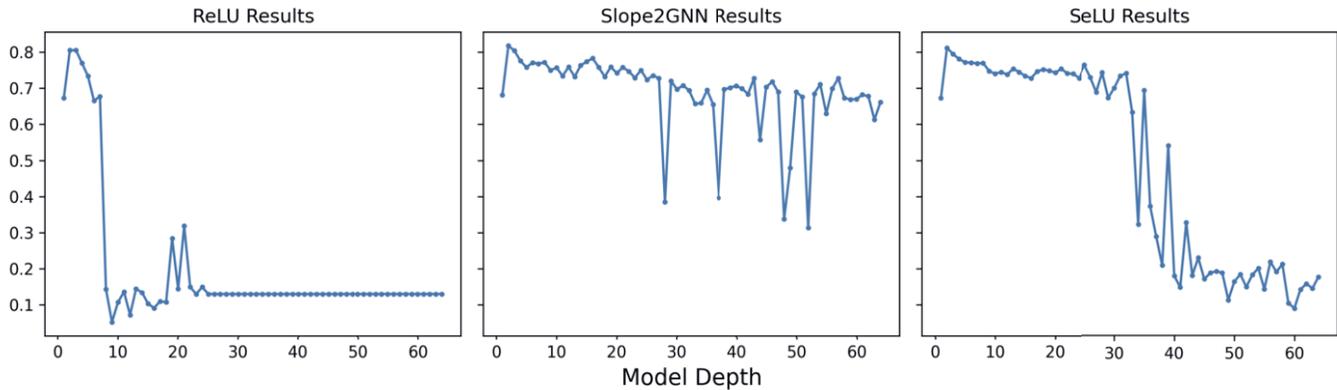
Li et al. [14] were the first to introduce the problem of oversmoothing in GNNs. An initial approach of activation function alternatives to avoid reduction of the rank of the feature space is the one provided by [17], where tanh is used instead of ReLU. Subsequent work, [18, 3] further analyzed and theoretically proved the existence of the problem and showed that it is unavoidable as the depth increases. Xu et al. [23] proposed Jumping Knowledge Networks (JK-Networks) as an attempt to address the problem, by injecting skip connections to the GNNs. The model kept information from lower layers and combined them directly with the final layer, before labeling the node. Following a similar approach, [4] proposed the use of residual connections between layers to enable deep architectures and alleviate the problem. Both of these approaches explicitly inject part of the initial

**Table 1.** Performance comparison of vanilla GCN and GAT, against SeLU and Slope2GNN enhanced versions of the same models, in *Cora*, *CiteSeer*, *Pubmed*, *Texas*. Average test node classification accuracy (%) for networks of different depth. With **bold** is the best performing model for each depth and each dataset.

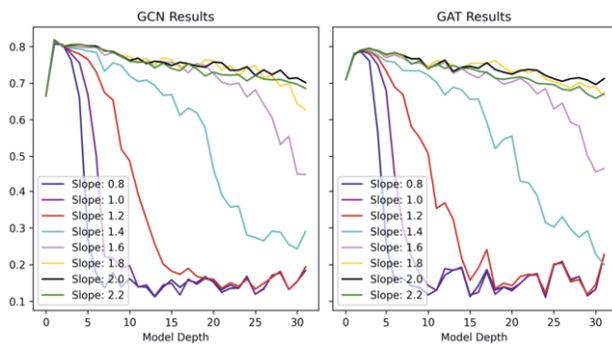
# Layers	Method	Accuracy (%)							
		Cora		CiteSeer		Pubmed		Texas	
		GCN	GAT	GCN	GAT	GCN	GAT	GCN	GAT
2	Original	81.38	77.79	70.52	<b>69.04</b>	77.65	<b>77.33</b>	<b>59.01</b>	<b>59.54</b>
	Slope2GNN	<b>81.84</b>	<b>77.81</b>	<b>70.55</b>	68.24	<b>78.34</b>	77.02	57.93	56.67
	SeLU	81.74	76.59	69.91	66.81	77.80	76.83	59.01	57.21
4	Original	78.09	78.06	65.21	63.98	<b>76.75</b>	<b>76.34</b>	<b>59.55</b>	58.28
	Slope2GNN	<b>80.39</b>	<b>79.54</b>	<b>67.57</b>	<b>67.27</b>	76.46	75.59	57.66	57.39
	SeLU	79.85	79.33	67.53	67.24	74.09	72.62	58.02	<b>58.65</b>
8	Original	23.56	33.11	32.54	31.33	49.08	55.86	<b>57.48</b>	<b>57.65</b>
	Slope2GNN	76.54	<b>78.33</b>	65.80	65.55	<b>75.00</b>	<b>74.20</b>	57.18	56.85
	SeLU	<b>79.41</b>	78.19	<b>67.27</b>	<b>67.58</b>	73.43	71.93	56.85	57.47
16	Original	14.92	15.53	17.73	17.80	28.95	29.88	39.91	41.71
	Slope2GNN	<b>76.81</b>	<b>75.00</b>	57.10	56.13	<b>76.35</b>	73.81	<b>57.21</b>	<b>57.53</b>
	SeLU	76.38	74.92	<b>61.99</b>	<b>62.23</b>	76.11	<b>74.27</b>	52.88	57.44
32	Original	14.02	14.67	16.69	18.67	38.49	29.34	25.50	18.64
	Slope2GNN	<b>73.21</b>	<b>69.62</b>	47.20	49.16	<b>75.78</b>	<b>74.63</b>	<b>55.79</b>	54.14
	SeLU	68.77	69.36	<b>47.90</b>	<b>52.97</b>	71.80	73.60	55.77	<b>57.75</b>
64	Original	12.48	15.18	17.65	13.35	31.73	28.62	27.39	06.76
	Slope2GNN	<b>61.59</b>	24.17	<b>46.46</b>	<b>26.24</b>	<b>71.68</b>	38.01	46.31	41.44
	SeLU	26.85	<b>30.55</b>	27.10	20.42	40.90	<b>41.00</b>	<b>58.38</b>	<b>56.30</b>

**Table 2.** Comparison of different models and activation functions on the “cold start” problem. We show accuracy (%) on the test set, using GCN and GAT as backbone GNN models. Only the features of the nodes in the training set are available to the model. We also show at what depth (i.e. # Layers) each model achieves its best accuracy.

Model		GCN		GAT	
Dataset	Method	Accuracy (%)	#L	Accuracy (%)	#L
Cora	Original	64.85	4	59.74	3
	Slope2GNN	<b>73.47</b>	19	<b>72.78</b>	20
	SeLU	73.00	22	72.65	23
CiteSeer	Original	41.94	4	38.50	4
	Slope2GNN	<b>49.88</b>	21	<b>49.63</b>	26
	SeLU	49.30	19	49.38	20
Pubmed	Original	60.16	4	50.45	4
	Slope2GNN	<b>72.61</b>	32	71.14	32
	SeLU	72.50	23	<b>71.41</b>	26
Texas	Original	32.40	4	30.10	2
	Slope2GNN	<b>33.33</b>	6	<b>31.00</b>	4
	SeLU	31.62	3	30.81	2



**Figure 1.** Comparison between three different activation functions for a GCN on *Cora* dataset. Y-axis shows model accuracy on test nodes, while varying the model’s depth (shown on the x-axis).



**Figure 2.** Varying ReLU slope results on *Cora* dataset.

information into higher layers of the network, implicitly downgrading the importance of intermediate layers. On the other hand, our method enables deep architectures, without the need to inject initial information, and allows the model to learn how to compress the initial information and what amount of it to maintain. We do not compare against these methods, because the aim of this work is to show that oversmoothing could be avoided without “short-circuiting” initial information to latter layers.

DropEdge [20] and DropNode [6] are two alternative methods that address oversmoothing by altering the graph topology. In particular, they remove either edges or nodes at random, in order to slow down message passing. In contrast to such approaches, our method does not change the graph topology, which may have unpredictable effects, but rather relies on a small adjustment of the slope of the activation function. An additional advantage of the approach proposed in this paper is that it does not have any hyper-parameters to be tuned, in contrast to all of the alternatives mentioned above. Hyperparameter tuning is also required by the Pairnorm method, which was mentioned in section 2.4.

## 6 Conclusion

In this paper we have shown the important role that the slope of the ReLU activation function plays in the oversmoothing problem

in GNNs. We have further proposed a simple modification that reduces drastically the problem. We have illustrated the benefits of the approach in a set of experiments with different datasets and GNNs of different depth. Additionally, we showed the improvement in accuracy one can achieve through deep architectures that do not suffer from oversmoothing. This was evident in a set of experiments that simulated the “cold start” problem of missing node features in GNNs. The simple link between the activation function and oversmoothing unveiled in this paper, opens up a range of interesting hypotheses to be investigated in the future. In particular, we would like to study the effect of changing the activation function in various GNN architectures and learning methods. Additionally, we would like to study alternative activation functions that could provide further benefits. Finally, we aim to study in more detail the approach of varying the learning rate per layer.

## Acknowledgements

The research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) under the 4th Call for HFRI PhD Fellowships (Fellowship Number: 10860).

## References

- [1] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton, ‘Layer normalization’, *CoRR*, **abs/1607.06450**, (2016).
- [2] Yoshua Bengio, Patrice Y. Simard, and Paolo Frasconi, ‘Learning long-term dependencies with gradient descent is difficult’, *IEEE Trans. Neural Networks*, **5**(2), 157–166, (1994).
- [3] Chen Cai and Yusu Wang, ‘A note on over-smoothing for graph neural networks’, *CoRR*, **abs/2006.13318**, (2020).
- [4] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li, ‘Simple and deep graph convolutional networks’, in *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1725–1735. PMLR, (2020).
- [5] Weilin Cong, Morteza Ramezani, and Mehrdad Mahdavi, ‘On provable benefits of depth in training graph convolutional networks’, in *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, eds., Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, pp. 9936–9949, (2021).

- [6] Tien Huu Do, Duc Minh Nguyen, Giannis Bekoulis, Adrian Munteanu, and Nikos Deligiannis, ‘Graph convolutional neural networks with node transition probability-based message passing and dropout regularization’, *Expert Syst. Appl.*, **174**, 114711, (2021).
- [7] William L. Hamilton, Zitao Ying, and Jure Leskovec, ‘Inductive representation learning on large graphs’, in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, eds., Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, pp. 1024–1034, (2017).
- [8] Wei Huang, Yayong Li, Weitao Du, Richard Y. D. Xu, Jie Yin, Ling Chen, and Miao Zhang, ‘Towards deepening graph neural networks: A gntk-based optimization perspective’, in *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, (2022).
- [9] Sergey Ioffe and Christian Szegedy, ‘Batch normalization: Accelerating deep network training by reducing internal covariate shift’, in *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, eds., Francis R. Bach and David M. Blei, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 448–456. JMLR.org, (2015).
- [10] Thomas N. Kipf and Max Welling, ‘Semi-supervised classification with graph convolutional networks’, in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, (2017).
- [11] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter, ‘Self-normalizing neural networks’, *CoRR*, **abs/1706.02515**, (2017).
- [12] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann, ‘Personalized embedding propagation: Combining neural networks on graphs with personalized pagerank’, *CoRR*, **abs/1810.05997**, (2018).
- [13] Johannes Klicpera, Janek Groß, and Stephan Günnemann, ‘Directional message passing for molecular graphs’, in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, (2020).
- [14] Qimai Li, Zhichao Han, and Xiao-Ming Wu, ‘Deeper insights into graph convolutional networks for semi-supervised learning’, in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, eds., Sheila A. McIlraith and Kilian Q. Weinberger, pp. 3538–3545. AAAI Press, (2018).
- [15] David Liben-Nowell and Jon M. Kleinberg, ‘The link prediction problem for social networks’, in *Proceedings of the 2003 ACM CIKM International Conference on Information and Knowledge Management, New Orleans, Louisiana, USA, November 2-8, 2003*, pp. 556–559. ACM, (2003).
- [16] Lu Lu, Yeonjong Shin, Yanhui Su, and George E. Karniadakis, ‘Dying relu and initialization: Theory and numerical examples’, *CoRR*, **abs/1903.06733**, (2019).
- [17] Sitao Luan, Mingde Zhao, Xiao-Wen Chang, and Doina Precup, ‘Break the ceiling: Stronger multi-scale deep graph convolutional networks’, in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, eds., Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, pp. 10943–10953, (2019).
- [18] Kenta Oono and Taiji Suzuki, ‘Graph neural networks exponentially lose expressive power for node classification’, in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, (2020).
- [19] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang, ‘Geom-gcn: Geometric graph convolutional networks’, in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, (2020).
- [20] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang, ‘Dropedge: Towards deep graph convolutional networks on node classification’, in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, (2020).
- [21] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio, ‘Graph attention networks’, in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, (2018).
- [22] Felix Wu, Amauri H. Souza Jr., Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Q. Weinberger, ‘Simplifying graph convolutional networks’, in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, eds., Kamalika Chaudhuri and Ruslan Salakhutdinov, volume 97 of *Proceedings of Machine Learning Research*, pp. 6861–6871. PMLR, (2019).
- [23] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka, ‘Representation learning on graphs with jumping knowledge networks’, in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, eds., Jennifer G. Dy and Andreas Krause, volume 80 of *Proceedings of Machine Learning Research*, pp. 5449–5458. PMLR, (2018).
- [24] Muhan Zhang and Yixin Chen, ‘Link prediction based on graph neural networks’, in *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, eds., Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, pp. 5171–5181, (2018).
- [25] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen, ‘An end-to-end deep learning architecture for graph classification’, in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, eds., Sheila A. McIlraith and Kilian Q. Weinberger, pp. 4438–4445. AAAI Press, (2018).
- [26] Lingxiao Zhao and Leman Akoglu, ‘Pairnorm: Tackling oversmoothing in gnns’, in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, (2020).