

Analysis of Graph Convolutional Networks using Neural Tangent Kernels

Mahalakshmi Sabanayagam¹[0000-0002-3476-8435], Pascal Mattia
Esser¹[0000-0002-1235-7918], and Debarghya
Ghoshdastidar^{1,2}[0000-0003-0202-7007]

¹ Technical University of Munich
{sabanaya, esser, ghoshdas}@in.tum.de
² Munich Data Science Institute

Abstract. Graph Convolutional Networks (GCNs) have emerged as powerful tools for learning on network structured data. Although empirically successful, GCNs exhibit certain behaviour that has no rigorous explanation—for instance, the performance of GCNs significantly degrades with increasing network depth, whereas it improves marginally with depth using skip connections. This paper focuses on semi-supervised learning on graphs, and explores the above observations through the lens of Neural Tangent Kernels (NTKs). To analyse the influence of depth, we derive NTKs corresponding to infinitely wide GCNs with and without skip connections and allowing non-linear output layer. While the constancy property of NTK is lost with the non-linear output layer, we show empirically that the approximation is similar to linear output layer. Using the newly derived NTK we analyze the influence of depth in GCNs and provide a comparison of different skip connections.

Keywords: Graph Convolutional Networks · Neural Tangent Kernels · Semi-Supervised Learning.

1 Introduction

Graph structured data are ubiquitous in various domains, including social network analysis, bioinformatics, communications engineering among others. In recent years, graph neural networks have become an indisputable choice for various learning problems on graphs, and have been employed in a wide range of applications across domains. Several variants of graph neural networks have been proposed, including graph convolutional network [Kipf and Welling, 2017], graph recurrent network [Scarselli et al., 2008, Li et al., 2016], graph attention network [Velickovic et al., 2018], to name a few. The popularity of graph neural networks can be attributed to their ability to tackle two conceptually different learning problems on graphs. In *supervised learning on graphs*, each data instance is a graph and the goal is to predict a label for each graph (for example, a protein structure). In contrast, *semi-supervised learning on graphs* (also called *node classification* or *graph transduction*) refers to the problem of predicting the labels of

nodes in a single graph. For instance, given the memberships of a few individuals in a social network, the goal is to predict affiliations of others.

This work focuses on the latter problem of semi-supervised learning. GCNs, along with its variants that locally aggregate information in the neighbourhood of each node, have proved to be superior methods in practice [Defferrard et al., 2016, Kipf and Welling, 2017, Chen et al., 2018a, Wu et al., 2019, Chen et al., 2020], outperforming classical, and well-studied, graph embedding based approaches. Among the different variants of GCNs, we focus on the methods based on approximations of spectral graph convolutions [Defferrard et al., 2016, Kipf and Welling, 2017], rather than spatial graph convolutions [Hamilton et al., 2017, Xu et al., 2019]. Surprisingly, these papers suggest shallow networks for the best performance, and unlike the standard neural networks that gain advantage with depth, the performance of GCN has been reported to decrease for deeper nets. This appears to be due to the over smoothing effect of applying many convolutions, that is, with repeated application of the graph diffusion operator in each layer, the feature information gets averaged out to a degree where it becomes uninformative. As a solution to this, Chen et al. [2020] and Kipf and Welling [2017] proposed different formulations of skip connections in GCNs that overcome the smoothing effect and thus outperform the vanilla GCN empirically. These networks achieve state-of-the-art results by directly operating on graphs which enables effective capturing of the complex structural information as well as the features associated with the entities. However, similar to standard neural networks, tuning the hyper-parameters is particularly hard due to the highly non-convex objective and the over-parameterised setup making it computationally intense. As a result, there is no theoretical framework that supports rigorous analysis of graph neural networks. Moreover, the graph convolutions increase the difficulty of analysis. Motivated by this, we are interested in a more formal approach to analyze GCNs and, specifically, to understand the influence of depth.

Explaining the empirical evidence of deep neural networks through mathematical rigour is an active area of research. In contrast, theoretical analysis of graph neural networks has been limited in the literature. From the perspective of learning theory, generalisation error bounds have been derived for graph neural networks using complexity measures like VC Dimension and Rademacher complexity [Esser et al., 2021, Scarselli et al., 2018, Garg et al., 2020]. However, it is often debated whether generalisation error bounds can explain the performance of deep neural networks [Neyshabur et al., 2017]. Another line of research relies on the connection between graph convolutions and belief propagation [Dai et al., 2016] to analyse the behaviour of graph neural networks in both supervised and semi-supervised settings using cavity methods and mean field approaches [Zhou et al., 2020b, Kawamoto et al., 2019, Chen et al., 2018b]. However, the above lines of research do not completely explain the empirical trends observed in GCNs, especially with regards to the aspects analysed in our work.

In this paper, we explain the empirically observed trends of GCNs using the recently introduced *Neural Tangent Kernel* (NTK) [Jacot et al., 2018]. NTK was proposed to describe the behaviour and generalisation properties of randomly

initialised fully connected neural networks during training by gradient descent with infinitesimally small learning rate. Jacot et al. [2018] also showed that, as the network width increases, the change in the kernel during training decreases and hence, asymptotically, one may replace an infinitely wide neural network by a deterministic kernel machine, where the kernel (NTK) is defined by the gradient of the network with respect to its parameters as

$$\Theta(x, x') = \mathbb{E}_{W \sim \mathcal{N}} \left[\left\langle \frac{\partial F(W, x)}{\partial W}, \frac{\partial F(W, x')}{\partial W} \right\rangle \right]. \quad (1)$$

Here $F(W, x)$ represents the output of the network at data point x and the expectation is with respect to W , that is, all the parameters of the network randomly sampled from Gaussian distribution \mathcal{N} . There has been criticism of the ‘infinite width’ assumption being too strong to model real (finite width) neural networks, and empirical results show that NTK often performs worse than the practical networks [Arora et al., 2019, Lee et al., 2019]. Nevertheless, theoretical insights on neural network training gained from NTK have proved to be valuable, particularly in showing how gradient descent can achieve good generalisation properties [Du et al., 2019a]. Subsequent works have derived NTK to analyse different neural network architectures in infinite width limit, including convolutional networks, recurrent networks among others [Arora et al., 2019, Du et al., 2018, 2019a, Alemohammad et al., 2021]. The most relevant work in the context of our discussion is the work of Du et al. [2019b] that derived NTK for graph neural networks in the supervised setting (each graph is a data instance to be classified) and empirically showed that graph NTK outperforms most graph neural networks as well as other graph kernels for the problem of graph classification.

Focus of this paper and contributions. The focus of the present paper differs from existing work on graph NTK [Du et al., 2019b] in two key aspects—we derive NTK for semi-supervised node classification with non-linear output layer and, more importantly, we use the derived NTKs to rigorously analyse corresponding GCN architectures and demonstrate the cause for surprising trends observed empirically in GCNs, as opposed to standard deep neural networks. More precisely, we make the following contributions:

1. In Section 2, we derive the NTKs for GCNs used in semi-supervised node classification [Kipf and Welling, 2017, Wu et al., 2019] in infinite width limit. In contrast to simplifying assumptions in most NTKs derivations, we allow a non-linear (sigmoid) pooling in the last layer—a natural choice in practical networks for binary classification (can also be extended to the multi class setting through *one vs. all* strategy). Although non-linear output layer results in non-constant NTK, our experiments demonstrate that the consequence of linear approximation is similar for both sigmoid and linear output layers. In addition, we show that the NTK captures the general trend in performance of GCN.

2. Due to the observation that NTK is a hyper-parameter free alternative which approximates the behaviour of GCNs, we analyse the impact of skip connections in GCNs [Chen et al., 2020, Kipf and Welling, 2017] by deriving the

corresponding NTKs in Section 3. Our empirical studies suggest the choice of skip connection for improved performance and the application of NTK to assess relative importance of structure and feature information in graph datasets.

3. In Section 4, we explain an empirical finding—unlike vanilla GCNs, the performance of NTK for certain skip connections converge with network depth. This is because the NTKs for skip connections converge with network depth, whereas this is less prominent in the case of NTK for vanilla GCNs.

We conclude in Section 5, and provide the NTK derivations and further experimental details in the appendix.

Notation. We represent the matrix Hadamard (entry-wise) product by \odot and the scalar product by $\langle \cdot, \cdot \rangle$. We use $M^{\odot k}$ to denote Hadamard product of matrix M with itself repeated k times. Let $\mathcal{N}(\mu, \Sigma)$ be Gaussian distribution with mean μ and co-variance Σ . For a function $\sigma(\cdot)$, $\dot{\sigma}(\cdot)$ represents its derivative. We use $\mathbf{1}_{n \times n}$ for the $n \times n$ matrix of ones, I_n for identity matrix of size $n \times n$, $\mathbb{E}[\cdot]$ for expectation, $\|\cdot\|_F, \|\cdot\|_2, \|\cdot\|_\infty$ denote Frobenius, Euclidean and Spectral norms, respectively, and $[d] = \{1, 2, \dots, d\}$.

2 Derivation of NTK for non-linear Vanilla GCN

Before presenting the NTK for GCN, we discuss the theoretical implications of the NTK for networks with non-linear output layer. While non-linear output layers are standard in practice, Liu et al. [2020] showed that the NTKs for such networks do not remain constant during training. The constancy of NTK during training is of great significance as this allows for deterministic characterisation of the network using the NTK in infinite width limit as detailed below: Consider the Taylor expansion of network $F(w, x)$ modeled by parameters w and data x around the weight initialisation w_0 :

$$F(w, x) \approx F(w_0, x) + \nabla_w F(w_0, x)^T (w - w_0) + \frac{1}{2} (w - w_0)^T \nabla_w^2 F(w_0, x) (w - w_0).$$

For NTK to be constant, $F(w, x)$ should be linear which implies the second order term should vanish as the width of the network m goes to infinity. Note that $(w - w_0)^T \nabla_w^2 F(w_0, x) (w - w_0) \leq \|w - w_0\|_2 \|\nabla_w^2 F(w_0, x)\|_\infty$.

Theoretical implications of non-linear last layer. Liu et al. [2020] showed that the second order term of the Taylor-expansion does not vanish for network with non-linear output layer: We note that $\|(w_0 - w)\|_2 = O(1)$ and for non-linear output layer $\|\nabla_w^2 F(w_0, x)\|_\infty = O(1)$. On the other hand, for linear output layer $\|\nabla_w^2 F(w_0, x)\|_\infty = O(1/\sqrt{m})$ and therefore NTK is constant only for linear output layer.

Empirical observation. Contrasting the linear output layer assumption, most networks used in practice rely heavily on the use of non-linear output layer. Therefore we investigate empirically how well the graph NTKs approximate networks with non-linear output layer as compared to linear output networks as theoretically suggested by Liu et al. [2020]. For this purpose, we consider node classification setup and measure prediction difference between graph NTKs and




Fig. 1. (left) Comparison of GCN and NTK prediction. **(middle/right)** Performance of NTK vs GCN in linear and non-linear architectures. The performance trend of NTK matches its corresponding GCN in both the architectures. Experiments performed on Cora.

trained GCNs. Formally, for the unlabeled nodes u , we compute the prediction difference $\frac{1}{|u|} \sum_{i=1}^u 1\{y_i^{NTK} \neq y_i^{GCN}\}$ where y_i^{NTK} and y_i^{GCN} are predictions from NTK and GCN respectively. This measure gives an estimate of how well the graph NTK approximate its corresponding trained GCN. The results for different architectures of GCNs (formally described later) with linear and sigmoid outputs for depths $\{1, 2, 4, 8\}$ are illustrated in the left plot of Figure 1. In contrast to the theory, we observe empirically that the NTKs approximate GCNs with both linear and non-linear outputs similarly, supporting the hypothesis that NTKs are good representative of the networks despite the non-constancy of the kernel in non-linear output case. This observation suggests theoretical gap at least in the graph setting and we leave the investigation as a possible future work.

2.1 Formal Setup of GCN

We consider the problem of node classification in graphs in a semi-supervised setting,³ where the labels are observed only for a subset of the nodes. Formally, given a graph with n nodes and a set of node features $\{x_i\}_{i=1}^n \subset \mathbb{R}^f$, we may assume without loss of generality that the set of observed labels $\{y_i\}_{i=1}^m$ correspond to first m nodes. We consider a binary classification problem such that $y_i \in \{\pm 1\}$ for simplicity of the derivation and then extend it to a K -class setting in the experimental section. The goal is to predict the $n - m$ unknown labels $\{y_i\}_{i=m+1}^n$. We represent the observed labels of m nodes as $Y \in \{\pm 1\}^{m \times 1}$, and the node features as $X \in \mathbb{R}^{n \times f}$ with the assumption that entire X is available during training. We define S to be the graph diffusion operator. The analysis holds for any diffusion S , but for simulations, we consider the symmetric degree normalized diffusion $S := (D + I_n)^{-\frac{1}{2}}(A + I_n)(D + I_n)^{-\frac{1}{2}}$ where A is the adjacency matrix and D is the degree matrix. We adapt the GCN proposed by Kipf

³ More precisely, transductive setting as we assume all features are available during training at the same time.

and Welling [2017] as follows,

$$F_W(X, S) := \Phi \left(\sqrt{\frac{c_\sigma}{h_d}} S \dots \sigma \left(\sqrt{\frac{c_\sigma}{h_1}} S \sigma(SXW_1) W_2 \right) \dots W_{d+1} \right) \quad (2)$$

where d is the network depth and $W := \{W_i \in \mathbb{R}^{h_{i-1} \times h_i}\}_{i=1}^{d+1}$ is the set of learnable weight matrices with $h_0 = f$ and $h_{d+1} = 1$, and $\Phi : \mathbb{R} \rightarrow (-1, +1)$ is re-scaled sigmoid since we consider binary node classification with labels in $\{\pm 1\}$, h_i is the size of layer $i \in [d]$ and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is the point-wise activation function. We initialise all the weights to be i.i.d $\mathcal{N}(0, 1)$ and optimise it using stochastic gradient descent. We study the limiting behavior of this network with respect to the width, that is, $h_1, \dots, h_d \rightarrow \infty$.

Remark 1 (c_σ). While this setup is similar to Kipf and Welling [2017], it is important to note that we additionally consider the normalisation $\sqrt{c_\sigma/h_i}$ for layer i to ensure that the input norm is approximately preserved. Here, c_σ is a scaling factor to normalize the input in the initialization phase and $c_\sigma = (\mathbb{E}_{u \sim \mathcal{N}(0,1)} [(\sigma(u))^2])^{-1}$ from Du et al. [2019a].

2.2 NTK for Vanilla GCN

We derive the NTK for vanilla GCN by first rewriting $F_W(X, S)$ as defined in (2) using the following recursive definitions with $g_1 := SX$:

$$g_i := \sqrt{\frac{c_\sigma}{h_{i-1}}} S \sigma(f_{i-1}) \quad \forall i \in \{2, \dots, d+1\}, \quad f_i := g_i W_i \quad \forall i \in [d+1]$$

$$\text{Output: } F_W(X, S) := \Phi(f_{d+1}), \quad \text{where } \Phi(x) := 2(1 + \exp(-x))^{-1} - 1 \quad (3)$$

Using the definitions in (3), the gradient with respect to W_i can be written as

$$\frac{\partial F_W(X, S)}{\partial W_i} := g_i^T b_i \quad \text{with} \quad b_i := \sqrt{\frac{c_\sigma}{h_i}} S^T b_{i+1} W_{i+1}^T \odot \dot{\sigma}(f_i) \quad (4)$$

and $b_{d+1} := \dot{\Phi}(f_{d+1})$. We derive the NTK, as defined in (1), using the recursive definition of $F_W(X, S)$ in (3) and its derivative in (4). The following theorem defines the NTK between every pair of nodes, and the $n \times n$ NTK matrix can be computed at once, as shown below (proof in appendix).

Theorem 1 (NTK for Vanilla GCN). *For the vanilla GCN defined in (2), the NTK Θ is given by*

$$\Theta = \left[\sum_{i=1}^{d+1} \Sigma_i \odot (SS^T)^{\odot(d+1-i)} \odot \left(\bigodot_{j=i}^{d+1-i} \dot{E}_j \right) \right] \odot_{f \sim \mathcal{N}(0, \Sigma_d)} \mathbb{E} \left[\dot{\Phi}(f) \dot{\Phi}(f)^T \right] \quad (5)$$

Here $\Sigma_i \in \mathbb{R}^{n \times n}$ is the co-variance between nodes of the layer f_i , and is given by $\Sigma_1 := SXX^T S^T$, $\Sigma_i := SE_{i-1} S^T$ with $E_i := c_\sigma \mathbb{E}_{f \sim \mathcal{N}(0, \Sigma_i)} [\sigma(f) \sigma(f)^T]$ and $\dot{E}_i := c_\sigma \mathbb{E}_{f \sim \mathcal{N}(0, \Sigma_i)} [\dot{\sigma}(f) \dot{\sigma}(f)^T]$.

Each entry of the expected matrix in (5) can be approximately computed as follows. For $\Delta \in \mathbb{R}^{2 \times 2}$,

$$\mathbb{E}_{(p,q) \sim \mathcal{N}(0,\Delta)} \left[\dot{\Phi}(p) \dot{\Phi}(q) \right] = \frac{1}{4} - \frac{\Delta_{00} + \Delta_{11}}{16} + \frac{\Delta_{00}\Delta_{11} + 2\Delta_{01}^2}{64} + \frac{\Delta_{00}^2 + \Delta_{11}^2}{32} + \frac{\epsilon^3}{16}$$

for $|\epsilon| \leq \max\{\Delta_{00}, \Delta_{11}\}$.

Inference using NTK. The NTK matrix $\Theta \in \mathbb{R}^{n \times n}$ defines the pairwise kernel among all labeled and unlabeled nodes, where each entry Θ_{pq} represents the kernel between nodes (or features) x_p and x_q . For inference, consider the sub-matrix $\Theta_l \in \mathbb{R}^{m \times m}$ that consists of the kernel computed between all pairs of labeled nodes, and $\Theta_u \in \mathbb{R}^{(n-m) \times m}$ that consists of the kernel computed between all pairs of unlabeled and labeled nodes. In the case of squared loss minimisation by stochastic gradient descent with infinitesimally small learning rate $\eta \rightarrow 0$, the training dynamics resemble kernel regression [Arora et al., 2019]. Hence, the labels for unlabeled nodes Y_u can be inferred as

$$Y_u = \Theta_u \Theta_l^{-1} Y \in \mathbb{R}^{n-m} \quad (6)$$

which, when thresholded entry-wise at 0, yields the class prediction for unlabeled nodes.

The NTK derived in (5) holds for vanilla GCN with arbitrary activation function in (2). Since the focus of this work is explaining the empirical performance trends of GCNs, we focus on specific activation functions that fix the network architecture allowing the NTK to be evaluated exactly. We first consider a linear activation, that results in the SGC network [Wu et al., 2019], and derive the NTK as follows.

Corollary 1 (Linear GCN). Consider $\sigma(x) := x$ in $F_W(X, S)$, then $E_i = c_\sigma \Sigma_i$ and $\dot{E}_i = c_\sigma \mathbf{1}_{n \times n}$ in Theorem 1, resulting in the following NTK

$$\Theta = c_\sigma^d \left[\sum_{i=1}^{d+1} \left(S^i X X^T (S^T)^i \right) \odot (SS^T)^{\odot(d+1-i)} \right] \odot_{f \sim \mathcal{N}(0, \Sigma_d)} \mathbb{E} \left[\dot{\Phi}(f) \dot{\Phi}(f)^T \right].$$

where the last expectation is approximated as in Theorem 1. The natural choice of normalisation constant c_σ is $c_\sigma = 1$ based on Remark 1.

Considering a non-linear network with ReLU activation, the NTK can be computed as shown below.

Corollary 2 (ReLU GCN). Consider $\sigma(x) := \text{ReLU}(x)$ in $F_W(X, S)$. The NTK kernel is computed as in (5), where given Σ_i at each layer, one can evaluate the entries of E_i and \dot{E}_i using a result from Bietti and Mairal [2019] as

$$\begin{aligned} (E_i)_{pq} &= \frac{c_\sigma}{2} \sqrt{(\Sigma_i)_{pp} (\Sigma_i)_{qq}} \kappa_1 \left(\frac{(\Sigma_i)_{pq}}{\sqrt{(\Sigma_i)_{pp} (\Sigma_i)_{qq}}} \right) \quad \text{and} \quad (7) \\ (\dot{E}_i)_{pq} &= \frac{c_\sigma}{2} \kappa_0 \left(\frac{(\Sigma_i)_{pq}}{\sqrt{(\Sigma_i)_{pp} (\Sigma_i)_{qq}}} \right), \end{aligned}$$

where $\kappa_0(x) := \frac{1}{\pi}(\pi - \arccos(x))$ and $\kappa_1(x) := \frac{1}{\pi}(x(\pi - \arccos(x)) + \sqrt{1-x^2})$. Based on Remark 1, the natural choice for normalisation constant c_σ is $c_\sigma = 2$.

2.3 Empirical Analysis of Depth

Many studies have shown that the performance of vanilla GCN drastically drops with depth due to the over smoothing effect of convolutional layers [Li et al., 2018, Kipf and Welling, 2017, Chen et al., 2020]. To validate it, we empirically study the performances of GCN and its NTK counterpart. We evaluate the performances of linear and ReLU GCNs as stated in Corollary 1 and 2, respectively, and their corresponding NTKs for different depths $d = \{1, 2, 4, 8\}$ by fixing the size of hidden layers $h_i = 2000$, same across all layers to reduce the number of hyper-parameters and set the learning rate η to 0.01.

NTK captures the performance trend of GCN. The right plots of Figure 1 show the performance of both the GCN architectures with its NTK counterpart evaluated on 7 class dataset *Cora* [McCallum et al., 2000]. The performance of GCN decreases with depth in both linear and non-linear architectures, as observed in other papers. This trend is also confirmed in NTK and thus making it a suitable method to analyse finite width GCN, despite the fact that the actual performance of the NTK is usually worse than the corresponding GCN. While there is a drop in performance in both the GCNs and the corresponding NTKs, the drop is not as drastic as it has been reported in other papers. This is due to two factors: first, increasing network size h_i in combination with appropriate learning rate η can reduce the performance drop with depth. Second, using correct normalisation c_σ (see Remark 1) stabilises training, thereby enabling the network to learn faster and achieve best results. In addition, we observe that the performances of NTKs with linear and non-linear output layer are very close.

3 NTK Analysis of Skip Connections

Skip connections [Chen et al., 2020, Kipf and Welling, 2017] are one way to overcome the performance degradation with depth in GCNs, but little is known about the effectiveness of different forms of available skip connections. Inspired by the observation of the previous section that the NTK is a hyper-parameter free model that captures the trends of GCNs, we use NTK to investigate different skip connections for GCN. We consider two formulations of skip connections with two variants each that are described in subsequent sections. To facilitate skip connections, we need to enforce constant layer size, that is, $h_i = h_{i-1}$. Therefore, we transform the input layer to H_0 of size $n \times h$ where h is the hidden layer size. This transformation is necessary as otherwise we would have to assume $h_i = f \forall i \in [d]$ and $h_i \rightarrow \infty$ would not be possible. For this work, we do not consider this transformation as a learnable parameter in the network. As we consider constant layer size, the NTKs are derived considering $h \rightarrow \infty$. We first define a skip connection related to the one in Kipf and Welling [2017], where

the skip connection is added to the features before convolution (we refer to it as pre-convolution or Skip-PC).

Definition 1 (Skip-PC). *In a Skip-PC (pre-convolution) network, the transformed input H_0 is added to the hidden layers before applying the diffusion, leading to the changes in the recursive definition of (3) with $g_1 := \sqrt{\frac{c_\sigma}{h}} S \sigma_s(H_0)$, $f_i := g_i W_i \forall i \in [d+1]$ and*

$$g_i := \sqrt{\frac{c_\sigma}{h}} S (\sigma(f_{i-1}) + \sigma_s(H_0)) \quad \forall i \in \{2, \dots, d+1\}, \quad (8)$$

where $\sigma_s(\cdot)$ can be linear or ReLU accounting for two different skip connections.

We refer to the network with linear $\sigma_s(\cdot)$ and ReLU $\sigma_s(\cdot)$ as Linear Skip-PC and ReLU Skip-PC, respectively. The above definition deviates from Kipf and Welling [2017] in the fact that we skip to the input layer instead of the previous layer. This particular change helps in evaluating the importance of graph information in a dataset which we discuss in the following section. We also consider a skip connection similar to the one described in Chen et al. [2020].

Definition 2 (Skip- α). *Given an interpolation coefficient $\alpha \in (0, 1)$ and a function $\sigma_s(\cdot)$, a Skip- α network is defined such that the transformed input H_0 and the hidden layer are interpolated linearly, which changes the recursive definition in (3) as $g_1 := \sqrt{\frac{c_\sigma}{h}} ((1-\alpha)S\sigma_s(H_0) + \alpha\sigma_s(H_0))$, $f_i := g_i W_i \forall i \in [d+1]$ and*

$$g_i := \sqrt{\frac{c_\sigma}{h}} ((1-\alpha)S\sigma(f_{i-1}) + \alpha\sigma_s(H_0)) \quad \forall i \in \{2, \dots, d+1\}, \quad (9)$$

Similar to Skip-PC, $\sigma_s(\cdot)$ can be linear or ReLU accounting for two different skip connections. We refer to the network with linear $\sigma_s(\cdot)$ and ReLU $\sigma_s(\cdot)$ as Linear Skip- α and ReLU Skip- α , respectively. Chen et al. [2020] recommends the choices for α as 0.1 or 0.2.

Remark 2 (Change of the normalization factor c_σ due to Skip connections). Note that the normalisation constant c_σ for GCN with skip connections is not the same as defined in Remark 1 of vanilla GCN, since we add the transformed input to the hidden layers and derive it for a GCN with $\sigma(x) := \text{ReLU}(x)$, to be $\simeq 0.67$.

3.1 NTK for GCN with Skip Connections

We derive NTKs for the skip connections – Skip-PC and Skip- α . Both the NTKs maintain the form presented in Theorem 1 with the following changes to the covariance matrices. Let $\tilde{E}_0 = \mathbb{E}_{f \sim \mathcal{N}(0, \Sigma_0)} [\sigma_s(f) \sigma_s(f)^T]$.

Corollary 3 (NTK for Skip-PC). *The NTK for an infinitely wide Skip-PC network is as presented in Theorem 1 where E_i is defined as in the theorem, but Σ_i is defined as*

$$\Sigma_0 := X X^T, \quad \Sigma_1 := S \tilde{E}_0 S^T \quad \text{and} \quad \Sigma_i := S E_{i-1} S^T + \Sigma_1. \quad (10)$$




Fig. 2. (left/middle) Performance of NTKs corresponding to the different skip connections where Skip- α is plotted for $\alpha = 0.2$. **(right)** Impact of α in Skip- α evaluated on Cora and WebKB datasets.

Corollary 4 (NTK for Skip- α). *The NTK for an infinitely wide Skip- α network is as presented in Theorem 1 where E_i is defined as in the theorem, but Σ_i is defined with $\Sigma_0 := XX^T$,*

$$\begin{aligned}\Sigma_1 &:= (1 - \alpha)^2 SE_0 S^T + \alpha(1 - \alpha)(SE_0 + E_0 S^T) + \alpha^2 E_0 \\ \Sigma_i &:= (1 - \alpha)^2 SE_{i-1} S^T + \alpha^2 \tilde{E}_0.\end{aligned}\tag{11}$$

Both Corollary 1 and 2 for linear and ReLU activations, respectively, hold for the derived NTKs corresponding to Skip-PC and Skip- α .

3.2 Empirical Analysis

Despite studies [Chen et al., 2020, Kipf and Welling, 2017] showing that having skip-connections gives a significant performance advantage, there is no clear way to choose one formulation of the skip connection over others due to the empirical fluctuation in GCN training. This practical problem can again be seen in the NTK setting as the derived NTKs have similar structure except the covariance between the nodes, thus making it difficult to compare analytically. Therefore, we empirically study the performance of different NTKs in order to determine the preferred formulation, thereby avoiding computational intensive hyper-parameter tuning. In addition, we show that the NTK corresponding to Skip- α can be used for assessing the relevance of structure and feature information of graph in a dataset.

Experimental setup. We evaluate the performance of NTKs corresponding to GCNs with skip connections for depth one to ten using non-linear activation $\sigma_s(x) := \text{ReLU}(x)$ (see (2)) for the GCNs. The linear transformation of the input X is done by $H_0 = XT$ where T is a $f \times h$ matrix and each entry is sampled from $\mathcal{N}(0, 1)$. The interpolation coefficient α in Skip- α is chosen to be $\{0.1, 0.2, 0.5\}$. NTKs for all the formulations of skip connections discussed in the previous section are evaluated on different multi-class datasets, namely *Cora*,

Citeseer [Giles et al., 1998] and *WebKB* [Craven et al., 1998]. Figure 2 shows the empirical observations.

We validate the expected performance advantage of GCN with skip connections over vanilla GCN, through NTK counterparts, and observe the following main findings.

Shallow net for GCNs with skip connection. Empirical analysis reveals a distinct behavior of skip connections with $\sigma_s(\cdot)$ being linear and ReLU, which is illustrated in the left plot of Figure 2. We observe that there is no performance gain with depth for both Skip-PC and Skip- α , and hence we restrict our focus to shallow depths. With regards to linear and ReLU $\sigma_s(\cdot)$, the performances are close and there is not a significant difference between the architectures. In addition, this experiment also validates the general practice of using shallow nets for GCNs. Hence, we propose skip connections with ReLU $\sigma_s(\cdot)$ and using shallow nets to achieve best performance in practice.

NTK as a model to assess relevance of structure and feature information of graphs. In the left plot of Figure 2, we notice that the performance of Skip- α on WebKB improved significantly as compared to Skip-PC and moreover, its performance continued to improve with depth, which is in contrast to other datasets. We further investigate this by analysing the interpolation coefficient α , and the corresponding results on Cora and WebKB datasets are shown in the right plot of Figure 2. Large value of α in Skip- α implies that more importance is given to feature information than the structural information of the graph. Therefore, from the figure, we infer that the structural information is not as important as the feature information for WebKB which is in contrast to Cora. Besides, NTK is a ready-to-use model without the need for hyper-parameter tuning. As a result, we propose NTK corresponding to Skip- α as a stand-alone model to determine the relative importance of structure and feature information in tasks where GCNs are employed.

4 Convergence of NTK with depth

In Figure 2, we observe that the performance of NTKs corresponding to GCNs with skip connections does not change significantly beyond a certain depth. We investigate this behaviour of the NTK further by measuring the amount of change between NTKs of different depths. To this end, we consider the alignment between the NTKs in the eigenspace following Fowlkes et al. [2004, Section 4.2]. Figure 3 shows the alignment of the NTKs for the discussed non-linear ReLU architectures ($\sigma(\cdot) := \text{ReLU}$ in (2)), evaluated on *Cora* dataset.

The learning happens in shallow depth. The different alignment plots illustrate the general influence of depth in GCN. We observe significant changes in the alignment between NTKs of shallow depths indicating that this is where learning happens. Since the NTKs for both vanilla GCN and GCN with skip connections converge with depth, it is clear that deep GCNs have no advantage or in other words, no new information is learned at deeper depths.

Influence of Skip connection. In addition, we observe that the NTKs reach almost perfect alignment with depth for GCNs with skip connection, suggesting that the networks reached saturation in learning as well. We can further distinguish the presented skip connections: overall Skip-PC has slow convergence most likely because the skip connection facilitates learning; Skip- α converges fast.




Fig. 3. Convergence of NTK with depth for all the discussed ReLU architectures, evaluated on Cora dataset. The plots show perfect alignment of NTKs for higher depths in GCNs with skip connections.

5 Conclusion

In this work, we derive NTKs for semi-supervised GCNs, including different formulations of skip connections. The deterministic hyper-parameter free nature of NTK makes it preferable choice for analysing its neural network counterpart since it captures the behaviour of the networks very well, as demonstrated in our experiments. While Liu et al. [2020] shows that the NTK does not remain constant under non-linear output layer theoretically, we observe empirically that this more practically used setup can still be analyzed using NTK. Under consideration of those findings, the primary goal of our work is to use NTK to advance our understanding of GCN, particularly on the impact of depth and network architectures. From our analysis, we suggest the NTK corresponding to the skip connection Skip- α to determine the relative importance of structure and feature information in graphs, which we believe to be of great practical value. There is a possibility of expanding the application of NTK to analyse robustness or explainability of GCNs, or other contexts that involve repeated training of networks. Another direction of research is to incorporate practical considerations of network architecture in the NTK derivation. The present paper allows sigmoid functions in the output layer, which is included through a Taylor expansion. It would be interesting to derive NTKs considering max-pooling or dropout, and use NTKs to analyse the impact of these techniques on network performance.

Bibliography

- Sina Alemohammad, Zichao Wang, Randall Balestriero, and Richard Baraniuk. The recurrent neural tangent kernel. In *International Conference on Learning Representations*, 2021.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In *Annual Conference on Neural Information Processing Systems, NeurIPS*, 2019.
- Alberto Bietti and Julien Mairal. On the inductive bias of neural tangent kernels. In *NeurIPS 2019-Thirty-third Conference on Neural Information Processing Systems*, volume 32, pages 12873–12884, 2019.
- Jie Chen, Tengfei Ma, and Cao Xiao. Fastgcn: Fast learning with graph convolutional networks via importance sampling. In *International Conference on Learning Representations*, 2018a.
- Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. Simple and deep graph convolutional networks. In *International Conference on Machine Learning*, pages 1725–1735. PMLR, 2020.
- Zhengdao Chen, Lisha Li, and Joan Bruna. Supervised community detection with line graph neural networks. In *International Conference on Learning Representations*, 2018b.
- Mark Craven, Andrew McCallum, Dan PiPasquo, Tom Mitchell, and Dayne Freitag. Learning to extract symbolic knowledge from the world wide web. Technical report, Carnegie-mellon univ pittsburgh pa school of computer Science, 1998.
- Hanjun Dai, Bo Dai, and Le Song. Discriminative embeddings of latent variable models for structured data. In *International conference on machine learning*, pages 2702–2711. PMLR, 2016.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *NIPS*, 2016.
- Pedro Domingos. Every model learned by gradient descent is approximately a kernel machine. *arXiv preprint arXiv:2012.00152*, 2020.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685. PMLR, 2019a.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2018.
- Simon S Du, Kangcheng Hou, Barnabás Póczos, Ruslan Salakhutdinov, Ruosong Wang, and Keyulu Xu. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019b.
- Pascal Mattia Esser, Leena C. Vankadara, and Debarghya Ghoshdastidar. Learning theory can (sometimes) explain generalisation in graph neural networks.

- In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2021.
- Charless Fowlkes, Serge Belongie, Fan Chung, and Jitendra Malik. Spectral grouping using the nystrom method. *IEEE transactions on pattern analysis and machine intelligence*, 26(2):214–225, 2004.
- Vikas Garg, Stefanie Jegelka, and Tommi Jaakkola. Generalization and representational limits of graph neural networks. In *International Conference on Machine Learning*, pages 3419–3430. PMLR, 2020.
- C Lee Giles, Kurt D Bollacker, and Steve Lawrence. Citeseer: An automatic citation indexing system. In *Proceedings of the third ACM conference on Digital libraries*, pages 89–98, 1998.
- William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1025–1035, 2017.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: convergence and generalization in neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 8580–8589, 2018.
- Tatsuro Kawamoto, Masashi Tsubaki, and Tomoyuki Obuchi. Mean-field theory of graph neural networks in graph partitioning. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124007, dec 2019. <https://doi.org/10.1088/1742-5468/ab3456>.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018.
- Jaehoon Lee, Lechao Xiao, Samuel S Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In *NeurIPS*, 2019.
- Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- Chaoyue Liu, Libin Zhu, and Mikhail Belkin. On the linearity of large non-linear models: When and why the tangent kernel is constant. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3(2):127–163, 2000.
- Behnam Neyshabur, Srinadh Bhojanapalli, David Mcallester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.
- Franco Scarselli, Ah Chung Tsoi, and Markus Hagenbuchner. The vapnik–chervonenkis dimension of graph and recursive neural networks. *Neural Networks*, 108:248 – 259, 2018.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *stat*, 1050:4, 2018.
- Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR, 2019.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020a.
- Pengfei Zhou, Tianyi Li, and Pan Zhang. Phase transitions and optimal algorithms for semisupervised classifications on graphs: From belief propagation to graph convolution network. *Physical Review Research*, 2(3):033325, 2020b.

A Proofs of NTKs for GCN and GCN with Skip Connections

We provide proofs of Theorem 1 and all corollaries with additional empirical results in this section.

A.1 Proof of NTK for Vanilla GCN (Theorem 1)

Co-variance between Nodes.

We will first derive the co-variance matrix of size $n \times n$ for each layer comprising of co-variance between any two nodes p and q . The co-variance between p and q in f_1 and f_i are derived below. We denote p -th row of matrix M as M_p . throughout our proofs.

$$\begin{aligned}
\mathbb{E} \left[(f_1)_{pk} (f_1)_{qk'} \right] &= \mathbb{E} \left[(g_1 W_1)_{pk} (g_1 W_1)_{qk'} \right] \\
&= \mathbb{E} \left[\sum_{r=1}^{h_0} (g_1)_{pr} (W_1)_{rk} \sum_{s=1}^{h_0} (g_1)_{qs} (W_1)_{sk'} \right] \\
&\stackrel{(W_1)_{xy} \sim \mathcal{N}(0,1)}{=} 0 \quad ; \text{ if } r \neq s \text{ or } k \neq k' \\
\mathbb{E} \left[(f_1)_{pk} (f_1)_{qk} \right] &\stackrel{r=s}{\stackrel{k=k'}{=}} \mathbb{E} \left[\sum_{r=1}^{h_0} (g_1)_{pr} (g_1)_{qr} (W_1)_{rk}^2 \right] \\
&\stackrel{(W_1)_{xy} \sim \mathcal{N}(0,1)}{=} \sum_{r=1}^{h_0} (g_1)_{pr} (g_1)_{qr} \\
&= \left\langle (g_1)_p, (g_1)_q \right\rangle \tag{12}
\end{aligned}$$

$$\begin{aligned}
\mathbb{E} \left[(f_i)_{pk} (f_i)_{qk} \right] &\stackrel{r=s}{\stackrel{k=k'}{=}} \mathbb{E} \left[\sum_{r=1}^{h_{i-1}} (g_i)_{pr} (g_i)_{qr} (W_i)_{rk}^2 \right] \\
&\stackrel{(W_i)_{xy} \sim \mathcal{N}(0,1)}{=} \sum_{r=1}^{h_{i-1}} (g_i)_{pr} (g_i)_{qr} \\
&= \left\langle (g_i)_p, (g_i)_q \right\rangle \tag{13}
\end{aligned}$$

$$\begin{aligned}
(12) : \quad \left\langle (g_1)_p, (g_1)_q \right\rangle &= \left\langle (SX)_p, (SX)_q \right\rangle \\
&= S_p X X^T S_q^T \\
&= (\Sigma_1)_{pq} \tag{14}
\end{aligned}$$

$$\begin{aligned}
 (13) : \quad \langle (g_i)_p, (g_i)_q \rangle &= \frac{c_\sigma}{h_{i-1}} \langle (S\sigma(f_{i-1}))_p, (S\sigma(f_{i-1}))_q \rangle \\
 &= \frac{c_\sigma}{h_{i-1}} \sum_{k=1}^{h_{i-1}} (S\sigma(f_{i-1}))_{pk} (S\sigma(f_{i-1}))_{qk} \\
 &\stackrel{h_{i-1} \rightarrow \infty}{=} c_\sigma \mathbb{E} \left[(S\sigma(f_{i-1}))_{pk} (S\sigma(f_{i-1}))_{qk} \right] \quad ; \text{ law of large numbers} \\
 &= c_\sigma \mathbb{E} \left[\left(\sum_{r=1}^n S_{pr} \sigma(f_{i-1})_{rk} \right) \left(\sum_{s=1}^n S_{qs} \sigma(f_{i-1})_{sk} \right) \right] \\
 &= c_\sigma \mathbb{E} \left[\sum_{r=1}^n \sum_{s=1}^n S_{pr} S_{qs} \sigma(f_{i-1})_{rk} \sigma(f_{i-1})_{sk} \right] \\
 &\stackrel{(a)}{=} \sum_{r=1}^n \sum_{s=1}^n S_{pr} (E_{i-1})_{rs} S_{sq}^T \\
 &= S_p E_{i-1} S_q^T \\
 &= (\Sigma_i)_{pq} \tag{15}
 \end{aligned}$$

(a): using $\mathbb{E}[(f_{i-1})_{rk} (f_{i-1})_{sk}] = (\Sigma_{i-1})_{rs}$ and the definition of E_{i-1} in Theorem 1.

NTK for Vanilla GCN.

Let us first evaluate the tangent kernel component from W_i respective to nodes p and q . The following two results are needed to derive it.

Result 1 (Inner Product of Matrices).

Let a and b be vectors of size $d_1 \times 1$ and $d_2 \times 1$, then

$$\begin{aligned}
 \langle ab^T, ab^T \rangle &= \text{tr} \left(ab^T (ab^T)^T \right) \\
 &= \text{tr} (ab^T ba^T) \\
 &= \text{tr} (a^T ab^T b) \\
 &= (a^T a) \odot (b^T b) \\
 &= \langle a, a \rangle \odot \langle b, b \rangle \tag{16}
 \end{aligned}$$

Result 2 $\langle (b_r)_p, (b_r)_q \rangle$.

We evaluate $\langle (b_r)_p, (b_r)_q \rangle = (b_r b_r^T)_{pq}$ which appears in the gradient.

$$\begin{aligned}
(b_r b_r^T)_{pq} &= \frac{c_\sigma}{h_r} \sum_{k=1}^{h_r} (S^T b_{r+1} W_{r+1}^T)_{pk} \dot{\sigma}(f_r)_{pk} (S^T b_{r+1} W_{r+1}^T)_{qk} \dot{\sigma}(f_r)_{qk} \\
&= \frac{c_\sigma}{h_r} \sum_{k=1}^{h_r} \sum_{i,j}^{n, h_{r+1}} S_{ip} (b_{r+1})_{ij} (W_{r+1})_{kj} \dot{\sigma}(f_r)_{pk} \dot{\sigma}(f_r)_{qk} \sum_{i',j'}^{n, h_{r+1}} S_{i'q} (b_{r+1})_{i'j'} (W_{r+1})_{kj'} \\
&= \frac{c_\sigma}{h_r} \sum_{i,j}^{n, h_{r+1}} \sum_{i',j'}^{n, h_{r+1}} (b_{r+1})_{ij} (b_{r+1})_{i'j'} S_{ip} S_{i'q} \sum_{k=1}^{h_r} (W_{r+1})_{kj} \dot{\sigma}(f_r)_{pk} \dot{\sigma}(f_r)_{qk} (W_{r+1})_{kj'} \\
&= \sum_{j,j'}^{h_{r+1}, h_{r+1}} (S^T b_{r+1})_{pj} (S^T b_{r+1})_{qj'} \frac{c_\sigma}{h_r} \sum_{k=1}^{h_r} (W_{r+1})_{kj} \dot{\sigma}(f_r)_{pk} \dot{\sigma}(f_r)_{qk} (W_{r+1})_{kj'} \\
&\stackrel{h_r \rightarrow \infty}{=} \sum_j^{h_{r+1}} (S^T b_{r+1})_{pj} (S^T b_{r+1})_{qj} c_\sigma \mathbb{E} \left[(W_{r+1}^2)_{kj} \dot{\sigma}(f_r)_{pk} \dot{\sigma}(f_r)_{qk} \right] \quad ; 0 \text{ for } j \neq j' \\
&\stackrel{(b)}{=} \left\langle (S^T b_{r+1})_p, (S^T b_{r+1})_q \right\rangle c_\sigma \mathbb{E} [\dot{\sigma}(f_r)_{pk} \dot{\sigma}(f_r)_{qk}] \\
&\stackrel{(16)}{=} (SS^T)_{pq} \langle b_{r+1}, b_{r+1} \rangle_{pq} c_\sigma \mathbb{E} [\dot{\sigma}(f_r)_{pk} \dot{\sigma}(f_r)_{qk}] \\
&= (SS^T)_{pq} \langle b_{r+1}, b_{r+1} \rangle_{pq} \left(\dot{E}_r \right)_{pq} \tag{17}
\end{aligned}$$

(b): $(W_{r+1})_{kj}$ is independent and $\mathbb{E} \left[(W_{r+1}^2)_{kj} = 1 \right]$.

Now, lets derive $\left\langle \left(\frac{\partial F}{\partial W_i} \right)_p, \left(\frac{\partial F}{\partial W_i} \right)_q \right\rangle$ and $\left\langle \left(\frac{\partial F}{\partial W_1} \right)_p, \left(\frac{\partial F}{\partial W_1} \right)_q \right\rangle$ using the above results.

$$\begin{aligned}
\left\langle \left(\frac{\partial F}{\partial W_i} \right)_p, \left(\frac{\partial F}{\partial W_i} \right)_q \right\rangle &= \left\langle (g_i)_p^T (b_i)_p, (g_i)_q^T (b_i)_q \right\rangle \\
&\stackrel{(16)}{=} \left\langle (g_i)_p, (g_i)_q \right\rangle \odot \left\langle (b_i)_p, (b_i)_q \right\rangle \\
&\stackrel{(15), (17)}{=} (\Sigma_i)_{pq} (SS^T)_{pq} \langle b_{r+1}, b_{r+1} \rangle_{pq} \left(\dot{E}_r \right)_{pq} \\
&\stackrel{(c)}{=} (\Sigma_i)_{pq} \left((SS^T)_{pq} \right)^{d+1-i} \left(\prod_{j=i}^{d+1-i} \left(\dot{E}_j \right)_{pq} \right) \langle b_{d+1}, b_{d+1} \rangle_{pq} \\
&\stackrel{(d)}{=} (\Sigma_i)_{pq} \left((SS^T)_{pq} \right)^{d+1-i} \left(\prod_{j=i}^{d+1-i} \left(\dot{E}_j \right)_{pq} \right) \left(\dot{\Phi}(f_{d+1}) \dot{\Phi}(f_{d+1})^T \right)_{pq} \tag{18}
\end{aligned}$$

(c): repeated application of (17).

(d): definition of b_{d+1} .

Extending (18) to all n nodes which will result in $n \times n$ matrix,

$$\begin{aligned} \left\langle \frac{\partial F}{\partial W_i}, \frac{\partial F}{\partial W_i} \right\rangle &= \Sigma_i \odot (SS^T)^{\odot d+1-i} \bigcirc_{j=i}^{d+1-i} \dot{E}_j \odot \dot{\Phi}(f_{d+1}) \dot{\Phi}(f_{d+1})^T \\ \mathbb{E}_{W_i} \left[\left\langle \frac{\partial F}{\partial W_i}, \frac{\partial F}{\partial W_i} \right\rangle \right] &= \Sigma_i \odot (SS^T)^{\odot d+1-i} \bigcirc_{j=i}^{d+1-i} \dot{E}_j \odot \mathbb{E}_{f \sim \mathcal{N}(0, \Sigma_d)} \left[\dot{\Phi}(f) \dot{\Phi}(f)^T \right] \end{aligned} \quad (19)$$

Finally, NTK Θ is,

$$\begin{aligned} \Theta &= \sum_{i=1}^{d+1} \mathbb{E}_{W_i} \left[\left\langle \frac{\partial F}{\partial W_i}, \frac{\partial F}{\partial W_i} \right\rangle \right] \\ &= \left[\sum_{i=1}^{d+1} \Sigma_i \odot (SS^T)^{\odot (d+1-i)} \odot \left(\bigcirc_{j=i}^{d+1-i} \dot{E}_j \right) \right] \odot \mathbb{E}_{f \sim \mathcal{N}(0, \Sigma_d)} \left[\dot{\Phi}(f) \dot{\Phi}(f)^T \right] \end{aligned} \quad (20)$$

We will now compute $\mathbb{E}_{f \sim \mathcal{N}(0, \Sigma_d)} \left[\dot{\Phi}(f) \dot{\Phi}(f)^T \right]$. We use Lagrange form of the remainder to approximate the Taylor's expansion for the re-scaled sigmoid function $\dot{\Phi}(\cdot)$ which gives better bound.

$$\begin{aligned} \Phi(x) &= \frac{2}{1 + \exp^{-x}} - 1 = \frac{x}{2} - \frac{x^3}{24} + \frac{x^5}{240} + \dots \\ \dot{\Phi}(x) &= \frac{1}{2} - \frac{x^2}{8} + \frac{x^4}{48} + \frac{x^6 \dot{\Phi}^6(\xi)}{6!} \quad ; \text{ last term is the Lagrange form of the remainder.} \end{aligned} \quad (21)$$

To evaluate the expectation of an entry i, j in the matrix $\dot{\Phi}(f) \dot{\Phi}(f)^T$, let us define Δ as a 2×2 co-variance matrix as follows, $\Delta = \begin{bmatrix} (\Sigma_{d+1})_{ii} & (\Sigma_{d+1})_{ij} \\ (\Sigma_{d+1})_{ji} & (\Sigma_{d+1})_{jj} \end{bmatrix}$

$$\begin{aligned} \mathbb{E}_{(x,y) \sim \Delta} \left[\dot{\Phi}(x) \dot{\Phi}(y) \right] &\stackrel{(21)}{=} \mathbb{E}_{(x,y) \sim \Delta} \left[\left(\frac{1}{2} - \frac{x^2}{8} + \frac{x^4}{48} + \frac{x^6 \dot{\Phi}^6(\xi)}{6!} \right) \left(\frac{1}{2} - \frac{y^2}{8} + \frac{y^4}{48} + \frac{y^6 \dot{\Phi}^6(\xi)}{6!} \right) \right] \\ &= \frac{1}{4} \mathbb{E}_{(x,y) \sim \Delta} \left[1 - \frac{x^2}{4} - \frac{y^2}{4} + \frac{x^4}{24} + \frac{y^4}{24} + \frac{x^2 y^2}{16} - \frac{x^4 y^2}{96} - \frac{x^2 y^4}{96} \right. \\ &\quad \left. + \frac{x^4 y^4}{576} + \frac{x^6 \dot{\Phi}^6(\xi)}{6!} \left(\frac{1}{2} - \frac{y^2}{8} + \frac{y^4}{48} + \frac{y^6 \dot{\Phi}^6(\xi)}{6!} \right) \right] \end{aligned} \quad (22)$$

Compute $\mathbb{E}_{x \sim \mathcal{N}(0, \lambda^2)} [x^k]$ and $\mathbb{E}_{(x, y) \sim \mathcal{N}(0, \Delta)} [x^i y^j]$.

$$\begin{aligned} \mathbb{E}_{x \sim \mathcal{N}(0, \lambda^2)} [x^k] &= \frac{2}{\sqrt{2\pi}\lambda} \int_0^\infty x^k \exp\left(\frac{-x^2}{2\lambda^2}\right) dx \\ &= \frac{2\lambda^k}{\sqrt{2\pi}} \int_0^\infty t^k \exp\left(\frac{-t^2}{2}\right) dt \quad ; x = \lambda t \implies dx = \lambda dt \\ &= \frac{2\lambda^k}{\sqrt{2\pi}} (k-1) \int_0^\infty t^{k-2} \exp\left(\frac{-t^2}{2}\right) dt \end{aligned}$$

$$\text{Thus, } \mathbb{E}_{x \sim \mathcal{N}(0, \lambda^2)} [x^k] = (k-1)\lambda^2 \mathbb{E}_{x \sim \mathcal{N}(0, \lambda^2)} [x^{k-2}] \quad (23)$$

$$\begin{aligned} \mathbb{E}_{(x, y) \sim \mathcal{N}(0, \Delta)} [x^i y^j] &= \mathbb{E}_{(x, y) \sim \mathcal{N}(0, \Delta)} [x^i (y \pm \alpha x)^j] \quad ; \alpha = \frac{\mathbb{E}[xy]}{\mathbb{E}[x^2]} \text{ then } x, y - \alpha x \text{ are independent} \\ &= \mathbb{E}_{(x, y) \sim \mathcal{N}(0, \Delta)} \left[x^i \left(\sum_{k=0}^j {}^j C_k (y - \alpha x)^j (\alpha x)^k \right) \right] \\ &= \mathbb{E}_{(x, y) \sim \mathcal{N}(0, \Delta)} \left[\sum_{k=0}^j {}^j C_k \alpha^k (y - \alpha x)^j x^{k+i} \right] \\ &\stackrel{(e)}{=} \sum_{k=0}^j {}^j C_k \alpha^k \mathbb{E}_{(x, y) \sim \mathcal{N}(0, \Delta)} [x^{k+i}] \mathbb{E}_{(x, y) \sim \mathcal{N}(0, \Delta)} [(y - \alpha x)^j] \end{aligned} \quad (24)$$

(e): $x, (y - \alpha x)$ are independent then $x^a, (y - \alpha x)^b$ are also independent.

Now, we evaluate (22) using (23) and (24) as follows.

$$\begin{aligned} (22) &\stackrel{(23), (24)}{=} \frac{1}{4} - \frac{1}{16} (\Sigma_{2ii} + \Sigma_{2jj}) + \frac{1}{64} (\Sigma_{2ii} \Sigma_{2jj} + 2\Sigma_{2ij}^2) \\ &\quad + \frac{1}{32} (\Sigma_{2ii}^2 + \Sigma_{2jj}^2) - \frac{1}{128} (\Sigma_{2ii}^2 \Sigma_{2jj} + \Sigma_{2ii} \Sigma_{2jj}^2 + 4\Sigma_{2ij}^2 \Sigma_{2ii} + 4\Sigma_{2ij}^2 \Sigma_{2jj}) \\ &\quad + \frac{1}{768} (3\Sigma_{2ii}^2 \Sigma_{2jj}^2 + 8\Sigma_{2ij}^4 + 24\Sigma_{2ij}^2 \Sigma_{2ii} \Sigma_{2jj}) \\ &\quad + \mathbb{E}_{(x, y) \sim \Delta} \left[\frac{x^6 \dot{\Phi}^6(\xi)}{6!} \left(\frac{1}{2} - \frac{y^2}{8} + \frac{y^4}{48} + \frac{y^6 \dot{\Phi}^6(\xi)}{6!} \right) \right] \\ &\leq \frac{1}{4} - \frac{1}{16} (\Sigma_{2ii} + \Sigma_{2jj}) + \frac{1}{64} (\Sigma_{2ii} \Sigma_{2jj} + 2\Sigma_{2ij}^2) + \frac{1}{32} (\Sigma_{2ii}^2 + \Sigma_{2jj}^2) \\ &\quad - \frac{10}{128} \epsilon^3 + \frac{35}{768} \epsilon^4 + \frac{15}{720} \epsilon^3 \quad ; |\epsilon| \leq \max\{\Delta_{00}, \Delta_{11}\}, \mathbb{E}[x^6] = 15\Delta_{00} \\ &\leq \frac{1}{4} - \frac{1}{16} (\Sigma_{2ii} + \Sigma_{2jj}) + \frac{1}{64} (\Sigma_{2ii} \Sigma_{2jj} + 2\Sigma_{2ij}^2) + \frac{1}{32} (\Sigma_{2ii}^2 + \Sigma_{2jj}^2) + \frac{1}{16} \epsilon^3 \end{aligned} \quad (25)$$

where $|\epsilon| \leq \max\{\Delta_{00}, \Delta_{11}\}$.

We get the NTK in Theorem 1 by putting together (25) and (20).

Corollary 1 (Linear GCN). In this case, $\sigma(x) := x$ and so derivative $\dot{\sigma}(x) = 1$. Consequently, one can derive $\dot{E}_i = c_\sigma \mathbf{1}_{n \times n}$ from its definition. Therefore, we get NTK for linear GCN in Corollary 1 by substituting \dot{E}_i in general NTK equation in (20).

Corollary 2 (ReLU GCN). NTK for ReLU GCN is derived by substituting (7) in general NTK equation in (20) as discussed in the corollary.

A.2 Proof of NTK for GCN with Skip Connections (Corollary 3 and 4)

We derive the NTKs for GCNs with different skip connections, Skip-PC and Skip- α in this section. We observe that the definitions of $g_i \forall i \in [1, d+1]$ are different for GCN with skip connections from the vanilla GCN. Despite the difference, the definition of gradient with respect to W_i in (4) does not change as g_i in the gradient accounts for the change and moreover, there is no new learnable parameter since the input transformation $H_0 = XT$ where T_{ij} is sampled from $\mathcal{N}(0, 1)$ is not learnable in our setting. Given the fact that the gradient definition holds for GCN with skip connection, the NTK will retain the form from NTK for vanilla GCN as evident from the above derivation. The change in g_i will only affect the co-variance between nodes. Hence, we will derive the co-variance matrix for the discussed skip connections, Skip-PC and Skip- α in the following sections.

Skip-PC: Co-variance between nodes. The co-variance between nodes p and q in f_1 and f_i are derived below.

$$\begin{aligned}
\mathbb{E} \left[(f_1)_{pk} (f_1)_{qk} \right] &= \left\langle (g_1)_p, (g_1)_q \right\rangle \\
&= \frac{c_\sigma}{h} \left\langle (S\sigma_s(H_0))_p, (S\sigma_s(H_0))_q \right\rangle \\
&= \frac{c_\sigma}{h} \sum_{k=1}^h (S\sigma_s(H_0))_{pk} (S\sigma_s(H_0))_{qk} \\
&\stackrel{h \rightarrow \infty}{=} c_\sigma \mathbb{E} \left[(S\sigma_s(H_0))_{pk} (S\sigma_s(H_0))_{qk} \right] \quad ; \text{law of large numbers} \\
&= S_p \tilde{E}_0 S_q^T \quad ; \tilde{E}_0 = c_\sigma \mathbb{E}_{f \sim \mathcal{N}(0, XX^T)} \left[\sigma_s(f) \sigma_s(f)^T \right] \\
&= (\Sigma_1)_{pq} \tag{26}
\end{aligned}$$

$$\begin{aligned}
\mathbb{E} \left[(f_i)_{pk} (f_i)_{qk} \right] &= \left\langle (g_i)_p, (g_i)_q \right\rangle \\
&= \frac{c_\sigma}{h} \left\langle (S(\sigma(f_{i-1}) + \sigma_s(H_0)))_p, (S(\sigma(f_{i-1}) + \sigma_s(H_0)))_q \right\rangle \\
&= \frac{c_\sigma}{h} \sum_{k=1}^h (S\sigma(f_{i-1}) + S\sigma_s(H_0))_{pk} (S\sigma(f_{i-1}) + S\sigma_s(H_0))_{qk} \\
&\stackrel{h \rightarrow \infty}{=} c_\sigma \mathbb{E} \left[(S\sigma(f_{i-1}) + S\sigma_s(H_0))_{pk} (S\sigma(f_{i-1}) + S\sigma_s(H_0))_{qk} \right] \quad ; \text{law of large numbers} \\
&= c_\sigma \left[\mathbb{E} \left[(S\sigma(f_{i-1}))_{pk} (S\sigma(f_{i-1}))_{qk} \right] + \mathbb{E} \left[(S\sigma(f_{i-1}))_{pk} (S\sigma_s(H_0))_{qk} \right] \right. \\
&\quad \left. + \mathbb{E} \left[(S\sigma_s(H_0))_{pk} (S\sigma(f_{i-1}))_{qk} \right] + \mathbb{E} \left[(S\sigma_s(H_0))_{pk} (S\sigma_s(H_0))_{qk} \right] \right] \\
&= S_p E_{i-1} S_{.q}^T + c_\sigma \mathbb{E} \left[(S\sigma(f_{i-1}))_{pk} (S\sigma_s(XW_0))_{qk} \right] \\
&\quad + c_\sigma \mathbb{E} \left[(S\sigma_s(XW_0))_{pk} (S\sigma(f_{i-1}))_{qk} \right] \\
&\quad + c_\sigma \mathbb{E} \left[\sum_{r=1}^n \sum_{s=1}^n S_{pr} S_{qs} \sigma_s(XW_0)_{rk} \sigma_s(XW_0)_{sk} \right] \\
&\stackrel{(f)}{=} S_p E_{i-1} S_{.q}^T + c_\sigma S_p \mathbb{E} [\sigma_s(XW_0)_{rk} \sigma_s(XW_0)_{sk}] S_{.q}^T \\
&= S_p E_{i-1} S_{.q}^T + S_p \tilde{E}_0 S_{.q}^T \\
&= S_p E_{i-1} S_{.q}^T + (\Sigma_1)_{pq} \\
&= (\Sigma_i)_{pq} \tag{27}
\end{aligned}$$

(f): $\mathbb{E} \left[(S\sigma(f_{i-1}))_{pk} (S\sigma_s(XW_0))_{qk} \right]$ and $\mathbb{E} \left[(S\sigma_s(XW_0))_{pk} (S\sigma(f_{i-1}))_{qk} \right]$ evaluate to 0 by conditioning on W_0 first and rewriting the expectation based on this conditioning. The terms within expectation are independent when conditioned on W_0 , and hence it is $\mathbb{E}_{W_0} \left[\mathbb{E}_{\Sigma_{i-1}|W_0} \left[(S\sigma(f_{i-1}))_{pk} | W_0 \right] \mathbb{E}_{\Sigma_{i-1}|W_0} \left[(S\sigma_s(XW_0))_{qk} | W_0 \right] \right]$ by taking h in W_0 going to infinity first. Here, $\mathbb{E}_{\Sigma_{i-1}|W_0} \left[(S\sigma_s(XW_0))_{qk} | W_0 \right] = 0$.

We get the co-variance matrix for all pairs of nodes $\Sigma_1 = S\tilde{E}_0 S^T$ and $\Sigma_i = S E_{i-1} S^T + \Sigma_1$ from (26) and (27).

Skip- α : Co-variance between nodes. Let p and q be two nodes and the co-variance between p and q in f_1 and f_i are derived below.

$$\begin{aligned}
 \mathbb{E} \left[(f_1)_{pk} (f_1)_{qk} \right] &= \left\langle (g_1)_p, (g_1)_q \right\rangle \\
 &= \frac{c_\sigma}{h} \sum_{k=1}^h \left((1-\alpha)S\sigma_s(H_0) + \alpha\sigma_s(H_0) \right)_{pk} \left((1-\alpha)S\sigma_s(H_0) + \alpha\sigma_s(H_0) \right)_{qk} \\
 &\stackrel{h \rightarrow \infty}{=} c_\sigma \mathbb{E} \left[\left((1-\alpha)S\sigma_s(H_0) + \alpha\sigma_s(H_0) \right)_{pk} \left((1-\alpha)S\sigma_s(H_0) + \alpha\sigma_s(H_0) \right)_{qk} \right] \\
 &= c_\sigma \left[(1-\alpha)^2 \mathbb{E} \left[(S\sigma_s(H_0))_{pk} (S\sigma_s(H_0))_{qk} \right] \right. \\
 &\quad \left. + (1-\alpha)\alpha \left(\mathbb{E} \left[(S\sigma_s(H_0))_{pk} (\sigma_s(H_0))_{qk} \right] + \mathbb{E} \left[(S\sigma_s(H_0))_{qk} (\sigma_s(H_0))_{pk} \right] \right) \right. \\
 &\quad \left. + \alpha^2 \mathbb{E} \left[(\sigma_s(H_0))_{pk} (\sigma_s(H_0))_{qk} \right] \right] \\
 &= (1-\alpha)^2 S_p \tilde{E}_0 S_q^T + (1-\alpha)\alpha \left(S_p \cdot \left(\tilde{E}_0 \right)_{\cdot q} + \left(\tilde{E}_0 \right)_p \cdot S_q^T \right) + \alpha^2 \left(\tilde{E}_0 \right)_{pq} \\
 &= (\Sigma_1)_{pq} \tag{28}
 \end{aligned}$$

$$\begin{aligned}
 \mathbb{E} \left[(f_i)_{pk} (f_i)_{qk} \right] &= \left\langle (g_i)_p, (g_i)_q \right\rangle \\
 &= \frac{c_\sigma}{h} \sum_{k=1}^h \left((1-\alpha)S\sigma(f_{i-1}) + \alpha\sigma_s(H_0) \right)_{pk} \left((1-\alpha)S\sigma(f_{i-1}) + \alpha\sigma_s(H_0) \right)_{qk} \\
 &\stackrel{h \rightarrow \infty}{=} c_\sigma \mathbb{E} \left[\left((1-\alpha)S\sigma(f_{i-1}) + \alpha\sigma_s(H_0) \right)_{pk} \left((1-\alpha)S\sigma(f_{i-1}) + \alpha\sigma_s(H_0) \right)_{qk} \right] \\
 &= c_\sigma \left[(1-\alpha)^2 \mathbb{E} \left[(S\sigma(f_{i-1}))_{pk} (S\sigma(f_{i-1}))_{qk} \right] + \alpha^2 \mathbb{E} \left[(\sigma_s(H_0))_{pk} (\sigma_s(H_0))_{qk} \right] \right. \\
 &\quad \left. + (1-\alpha)\alpha \left(\mathbb{E} \left[(S\sigma(f_{i-1}))_{pk} (\sigma_s(H_0))_{qk} \right] + \mathbb{E} \left[(\sigma_s(H_0))_{pk} (S\sigma(f_{i-1}))_{qk} \right] \right) \right] \\
 &\stackrel{(g)}{=} (1-\alpha)^2 S_p \cdot E_{i-1} S_q^T + \alpha^2 \left(\tilde{E}_0 \right)_{pq} = (\Sigma_i)_{pq} \tag{29}
 \end{aligned}$$

(g): same argument as (f) in derivation of Σ_i in Skip-PC.

We get the co-variance matrix for all pairs of nodes $\Sigma_1 = (1-\alpha)^2 S \tilde{E}_0 S^T + \alpha(1-\alpha) \left(S \tilde{E}_0 + \tilde{E}_0 S^T \right) + \alpha^2 \tilde{E}_0$ and $\Sigma_i = (1-\alpha)^2 S E_{i-1} S^T + \alpha^2 \tilde{E}_0$ from (28) and (29).

A.3 Normalisation constant c_σ (Remark 1 and 2).

We derive the normalisation constant c_σ loosely, as the purpose of c_σ is to preserve the input norm approximately. We focus on general form of a network with skip connection (not GCN in particular), where the output vector of size h from any hidden layer l with weight matrix $W \in \mathbb{R}^{h \times h}$ and transformed input vector X_0 of size h can be written as $g_l := \sqrt{\frac{c_\sigma}{h}} (\sigma(Wg_{l-1}) + X_0) \in \mathbb{R}^{h \times 1}$. The role of

the normalisation constant c_σ is to maintain $\|g_l\|_2 \simeq \|X_0\|_2$ and is derived as follows.

$$\begin{aligned}
\|X_0\|_2^2 &= \|g_l\|_2^2 = \frac{c_\sigma}{h} \sum_{k=1}^h (\sigma(Wg_{l-1}) + X_0)_k^2 \\
\|X_0\|_2^2 &= c_\sigma \mathbb{E} \left[(\sigma(Wg_{l-1})_k)^2 + (X_0)_k^2 + 2\sigma(Wg_{l-1})_k (X_0)_k \right] \quad ; h \rightarrow \infty \\
\|X_0\|_2^2 &= c_\sigma \mathbb{E}_{u \sim \mathcal{N}(0, \|X_0\|_2^2)} \left[(\sigma(u))^2 \right] + \|X_0\|_2^2 \quad ; \mathbb{E} [\sigma(Wg_{l-1})_k (X_0)_k] = 0 \\
c_\sigma &= \left(\mathbb{E}_{u \sim \mathcal{N}(0,1)} \left[(\sigma(u))^2 \right] + 1 \right)^{-1} \quad ; \text{normalised } X_0 \tag{30}
\end{aligned}$$

We use this c_σ for GCN with skip connection in our work and it evaluates to 2/3 for $\sigma(x) := \text{ReLU}(x)$ in GCN as stated in Remark 2. The evident change for a network without skip connection is to not add X_0 in $g_l := \sqrt{\frac{c_\sigma}{h}} \sigma(Wg_{l-1})$ and consequently by following the proof, we get $c_\sigma = \left(\mathbb{E}_{u \sim \mathcal{N}(0,1)} \left[(\sigma(u))^2 \right] \right)^{-1}$ as mentioned in Remark 1.

A.4 Formal details for convergence of NTK with depth analysis

Formally, let Θ_i and Θ_j be the NTK of depth i and j , respectively, and $U_i^{(k)}$ and $U_j^{(k)}$ be the matrix of k leading eigenvectors of Θ_i and Θ_j , respectively, then the alignment between Θ_i and Θ_j is computed by $a = \frac{1}{k} \left\| U_i^{(k)T} U_j^{(k)} \right\|_F^2$, where $a \in [0, 1]$ with $a = 1$ indicating perfect alignment.

B Additional Experimental Results

B.1 Datasets for binary node classification

The additional experiments are performed by converting the datasets to have binary class by grouping the classes into two sets. Table 1 shows the label grouping for all the considered datasets *Cora*, *Citeseer* and *WebKB* and total number of nodes with the grouped labels respectively. The classes in all the datasets are well balanced and sensible to learn for binary classification problem which is proved from the performance of a simple graph neural network like linear vanilla GCN. The train-test split for each dataset is 708 and 2000 nodes for *Cora*, 312 and 2000 for *Citeseer*, and 377 and 500 for *WebKB* for all the experiments.

	Cora		Citeseer		WebKB	
	Class Groups	#nodes	Class Groups	#nodes	Class Groups	#nodes
Class 1	Neural_Networks Theory Probabilistic_Methods	1595	Agents AI ML	1435	student	415
Class 2	Case_Based Rule_Learning Reinforcement Genetic_Algorithms	1103	DB IR HCI	1877	faculty staff course project	462
Total		2708		3312		877

Table 1. Class grouping in datasets for binary node classification.

B.2 Vanilla GCN vs GCN with Skip Connections

We established ReLU GCN is preferred over linear in Section 2 and ReLU for the input transformation in Section 3. Hence, we focus on $\sigma(\cdot) := \text{ReLU}$ and $\sigma_s(\cdot) := \text{ReLU}$ with $\alpha = 0.2$ for Skip- α to validate the performance of vanilla GCN and GCN with skip connections, Skip-PC and Skip- α . We use the respective NTKs to validate the performance. Figure 4 shows that GCN with skip connection outperforms vanilla GCN even in deeper depths, and Skip- α gives better performance than Skip-PC with depth.

B.3 Convergence of NTK with depth - Cora, Citeseer, WebKB

We presented the convergence of NTK with depth for ReLU GCN with and without skip connections evaluated on multiclass Cora dataset in Figure 3. Here, we present the convergence plot for Linear GCN evaluated on binary class Cora for larger depths $d = \{1, 2, 4, 8, 16, 32, 64, 128\}$ and all discussed linear and ReLU networks evaluated on Citeseer and WebKB whose classes are grouped as per Table 1. The observation is similar to the discussion in Section 4. Figures 5, 6 and 7 show the convergence plots for linear GCN evaluated on Cora, ReLU and linear GCNs with and without skip connections for Citeseer and WebKB, respectively.




Fig. 4. Performance validation of vanilla GCN, Skip-PC and Skip- α with $\sigma(\cdot) := \text{ReLU}$, $\sigma_s(\cdot) := \text{ReLU}$ and $\alpha = 0.2$ using the respective NTKs.




Fig. 5. Convergence of NTK with depth for all the discussed linear architectures evaluated on Cora dataset.

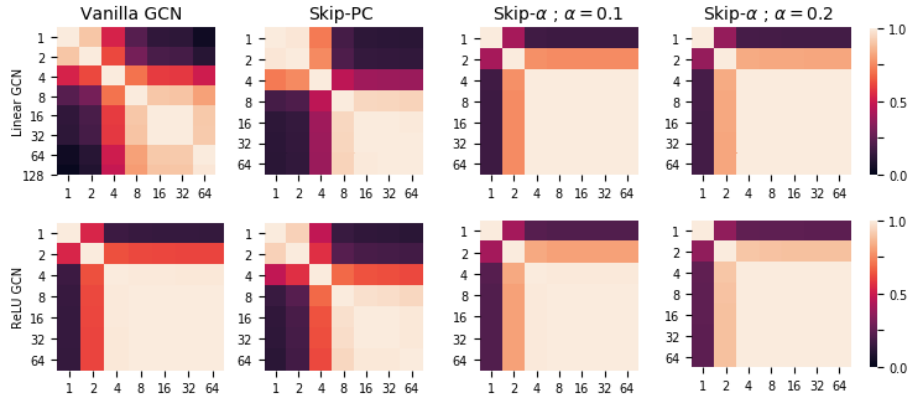


Fig. 6. Convergence of NTK with depth for all the discussed linear and ReLU architectures evaluated on Citeseer dataset.




Fig. 7. Convergence of NTK with depth for all the discussed linear and ReLU architectures evaluated on WebKB dataset.