

Asymmetric Learning for Spectral Graph Neural Networks

Fangbing Liu, Qing Wang

Graph Research Lab, School of Computing, Australian National University
fangbing.liu@anu.edu.au, qing.wang@anu.edu.au

Abstract

Optimizing spectral graph neural networks (GNNs) remains a critical challenge in the field, yet the underlying processes are not well understood. In this paper, we investigate the inherent differences between graph convolution parameters and feature transformation parameters in spectral GNNs and their impact on the optimization landscape. Our analysis reveals that these differences contribute to a poorly conditioned problem, resulting in suboptimal performance. To address this issue, we introduce the concept of the block condition number of the Hessian matrix, which characterizes the difficulty of poorly conditioned problems in spectral GNN optimization. We then propose an asymmetric learning approach, dynamically preconditioning gradients during training to alleviate poorly conditioned problems. Theoretically, we demonstrate that asymmetric learning can reduce block condition numbers, facilitating easier optimization. Extensive experiments on eighteen benchmark datasets show that asymmetric learning consistently improves the performance of spectral GNNs for both heterophilic and homophilic graphs. This improvement is especially notable for heterophilic graphs, where the optimization process is generally more complex than for homophilic graphs.

Code — <https://github.com/Mia-321/asym-opt.git>

Extended version — <https://arxiv.org/abs/2412.11739>

Introduction

Graph Neural Networks (GNNs) have been extensively studied in recent years (Wu et al. 2019; Loukas 2020; Dwivedi et al. 2023; Velickovic et al. 2018; Xu et al. 2019). Inspired by the seminal work of (Kipf and Welling 2017; Defferrard, Bresson, and Vandergheynst 2016), research has delved into spectral GNNs, which extract information from graph-structured data by performing graph convolutions in the spectral domain (Wang and Zhang 2022a; Balciilar et al. 2021; Yang et al. 2022). Generally, a *spectral GNN* is defined as

$$\hat{Y} = \sigma(g_{\Theta}(M)f_W(X)). \quad (1)$$

Here, $M \in \mathbb{R}^{n \times n}$ is a graph matrix (e.g., normalized Laplacian), $X \in \mathbb{R}^{n \times d}$ is a feature matrix, g_{Θ} is a graph convolu-

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

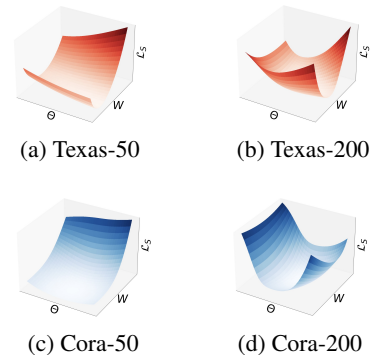


Figure 1: The landscapes of ChebNet on Texas and Cora at iterations 50 and 200. The x, y, z -axes represent the directions of Θ, W , and the empirical loss \mathcal{L}_S , respectively. The landscapes change during training, and the stretched direction of the parameters Θ and W also shifts.

tion function parameterized by Θ , f_W is a feature transformation function parameterized by W , and σ is a non-linear function (e.g., softmax). For computational efficiency and local vertex analysis (Hammond, Vandergheynst, and Grignonval 2011), the graph convolution function is typically approximated by polynomial filters in the form $g_{\Theta}(M) = \sum_{k=0}^K \theta_k T_k(M)$, where each $T_k(\bullet)$ denotes the k -th polynomial basis. The exact form of T_k is determined by the selected polynomial basis, e.g., the Chebyshev basis (Defferrard, Bresson, and Vandergheynst 2016; He, Wei, and Wen 2022) or the Bernstein basis (He et al. 2021).

Despite the success of spectral GNNs in graph learning tasks (Zhu et al. 2020; Lim et al. 2021; Pei et al. 2020), their training process remains inadequately understood. Notably, the parameters Θ and W play distinct roles in modeling and may significantly differ in dimensions. Specifically, the dimension of Θ is fewer than 20 while W has thousands of dimensions. During the training of spectral GNNs, different learning rates for parameters Θ and W are often adopted to achieve good performance. Interestingly, we observe that the contours of the loss function undergo stretching or compression along specific directions during training. In stretched directions, small alterations in parameters lead to significant changes in the loss function. Figure 1 shows the loss

landscapes of a spectral GNN model, ChebNet (Defferrard, Bresson, and Vandergheynst 2016). There always exists a stretched direction for Θ or W during training, indicating that parameters Θ and W exhibit different sensitivities to learning rates and gradient updates.

The distorted landscapes of spectral GNNs can be further explained from the view of Hessian matrix. According to (Le-Cun, Simard, and Pearlmutter 1992; Granzio, Zohren, and Roberts 2020), neural networks exhibits non-convex characteristics and the optimal learning rate of a parameter is inversely proportional to the largest eigenvalue of the Hessian matrix. The Hessian matrix \mathbf{H} of a spectral GNN contains second partial derivatives of Θ and W , which can be divided into two diagonal block matrices $\mathbf{H}_{\Theta, \Theta}$ and $\mathbf{H}_{W, W}$, and two off-diagonal block matrices. In scenarios where Θ and W exhibit minimal interference, the optimal learning rates for Θ and W correspond inversely to the largest eigenvalues of their respective diagonal block Hessian matrices. The largest eigenvalue of a Hessian matrix indicates the maximum rate of change of the gradient with respect to a small change of its corresponding parameter. In regions where the largest eigenvalue of a Hessian matrix is large, the gradient norm tends to be large as well, indicating a large loss change with respect to a small parameter change (Bonnans et al. 2006). The larger the gap between two diagonal block matrices of Hessian with respect to Θ and W respectively, the more distorted of the local landscape. Figure 2 shows the ratio of the largest eigenvalues of two diagonal block Hessian matrices (largest eigenvalue ratio) when a spectral GNN is applied on different graphs. There is a notable disparity in the largest eigenvalue ratio across different graphs. When a graph is more homophilic, the largest eigenvalue ratio is closer to one, indicating a less distorted landscape. For example, the landscapes of Texas are more distorted than those of Cora because Texas is a heterophilic graph while Cora is a homophilic graph.

To characterize the aforementioned phenomenon, particularly the relationship between the parameters Θ and W in spectral GNNs and their impact on optimization, we introduce the concept of the *block condition number* of the Hessian matrix. This metric serves as an indicator for assessing the difficulty of an optimization problem. Specifically, utilizing the block condition number, we show that poorly conditioned spectral GNNs are characterized by a large block condition number of the Hessian matrix.




Figure 2: The largest eigenvalue ratio α of ChebNet when applied to graphs of varying edge homophily ratios H_{edge} (Zhu et al. 2020) for node classification. \mathbf{H} represents the Hessian matrix and variances are depicted as grey shades.

A widely adopted strategy for addressing poorly conditioned optimization problems is *preconditioning* (Saad 2003). Advanced adaptive learning rate optimizers tackle poorly conditioned problems by assigning different learning rates to different parameters, calculated based on a gradient history (Kingma and Ba 2014; Duchi, Hazan, and Singer 2011). However, these optimizers overlook the inherent parameter magnitude differences exhibited by the parameters Θ and W during the optimization of spectral GNNs. To accelerate the learning process and potentially achieve superior solutions, we propose *asymmetric learning* for training spectral GNNs. This new method involves scaling the gradient of the empirical loss function asymmetrically during each iteration. The scaling factors within the asymmetric preconditioner are determined by the *gradient-parameter norm ratio*, which is closely associated with the largest eigenvalues of the block Hessian matrices $\mathbf{H}_{\Theta, \Theta}$ and $\mathbf{H}_{W, W}$. Notably, this asymmetric preconditioning of the gradient corresponds to an effective preconditioner for the Hessian matrix \mathbf{H} . Under reasonable assumptions, we demonstrate that asymmetric preconditioning can effectively reduce the block condition number of the Hessian matrix \mathbf{H} , thereby facilitating more efficient optimization of spectral GNNs.

Contributions. In summary, our contributions are: (1) We reveal that the inherent difference between the graph convolution parameter Θ and the feature transformation parameter W in spectral GNNs leads to a poorly conditioned optimization problem. (2) We introduce the concept of block condition number to quantify optimization difficulty. (3) We introduce an asymmetric learning approach to improve the training and performance of spectral GNNs. (4) Our extensive experiments on eighteen benchmark datasets demonstrate that asymmetric learning consistently enhances spectral GNN performance across various graphs.

Problem Analysis

Let $G = \{\mathcal{V}, \mathcal{E}\}$ be an undirected graph with a node set \mathcal{V} , an edge set \mathcal{E} , and $n = |\mathcal{V}|$. Each node v_i is associated with a feature vector x_i and a label y_i . We use X to denote the node feature matrix, where the i -th row corresponds to x_i . Let $S = (X, \{y_i\}_{i=1}^m)$ be the training set containing m labeled nodes randomly selected from all nodes of the graph. Similarly, we obtain the validation set $\mathcal{D}_{val} = (X, \{y_i\}_{i=m+1}^{m+q})$ containing q labeled nodes. A loss function $\ell(\hat{y}_i, y_i; \Theta, W)$ estimates the distance between the truth node label y_i and the prediction \hat{y}_i from a spectral GNN parameterized by Θ, W in Eq. (1). The empirical loss is $\mathcal{L}_S(\Theta, W) = \frac{1}{m} \sum_{i=0}^m \ell(\hat{y}_i, y_i; \Theta, W)$, and the validation loss is $\mathcal{L}_{\mathcal{D}_{val}}(\Theta, W) = \frac{1}{q} \sum_{i=m+1}^{m+q} \ell(\hat{y}_i, y_i; \Theta, W)$. We use $\|\bullet\|_2$ to denote the ℓ_2 norm of \bullet , $vec(\bullet)$ to vectorize a matrix by concatenating its rows, and $d_\Theta = |\Theta|$ and $d_W = |vec(W)|$ to denote the total parameter number of Θ and W , respectively.

Poorly Conditioned Spectral GNNs. We begin by presenting the classic definition of a poorly conditioned problem

and then introduce a new concept called the *block condition number*. Using this new concept, we analyze the poorly conditioned nature of spectral GNNs.

Definition 1 (Poorly Conditioned Problem (Golub and Van Loan 2013)). Let $\mathcal{L}_S(\Psi)$ be a convex objective function, parameterized by Ψ , and \mathbf{H} be its associated Hessian matrix. If \mathbf{H} is ill-conditioned, the optimization of $\mathcal{L}_S(\Psi)$ is poorly conditioned. The extent to which \mathbf{H} is ill-conditioned is quantified by the *condition number* of \mathbf{H} , defined as

$$\kappa(\mathbf{H}) = \frac{\lambda_{\max}(\mathbf{H})}{\lambda_{\min}(\mathbf{H})}. \quad (2)$$

Here, $\kappa(\mathbf{H})$ is the ratio of the absolute values of the maximum eigenvalue $\lambda_{\max}(\mathbf{H})$ to the minimum eigenvalue $\lambda_{\min}(\mathbf{H})$.

Remark 2. Eigenvectors corresponding to the maximum eigenvalue $\lambda_{\max}(\mathbf{H})$ and the minimum eigenvalue $\lambda_{\min}(\mathbf{H})$ represent the most and least sensitive directions of the loss function \mathcal{L}_S , respectively. A significant gap between $\lambda_{\max}(\mathbf{H})$ and $\lambda_{\min}(\mathbf{H})$ induces numerical instability during optimization and slow down the convergence rate (Bonnans et al. 2006).

In spectral GNNs, the Hessian matrix \mathbf{H} can be partitioned into two diagonal blocks $\mathbf{H}_{\Theta, \Theta}$ and $\mathbf{H}_{W, W}$, and two non-diagonal blocks $\mathbf{H}_{\Theta, W}$ and $\mathbf{H}_{W, \Theta}$. The ij -th element of $\mathbf{H}_{\Theta, \Theta}$ is $\frac{\partial^2 \mathcal{L}_S(\Theta, W)}{\partial \Theta_i \partial \Theta_j}$, where $i, j \in [1, d_{\Theta}]$. The ij -th element of $\mathbf{H}_{W, W}$ is $\frac{\partial^2 \mathcal{L}_S(\Theta, W)}{\partial W_i \partial W_j}$, where $i, j \in [1, d_W]$. The ij -th element of $\mathbf{H}_{\Theta, W}$ (equivalently $\mathbf{H}_{W, \Theta}^T$) is $\frac{\partial^2 \mathcal{L}_S(\Theta, W)}{\partial \Theta_i \partial W_j}$, with $i \in [1, d_{\Theta}]$ and $j \in [1, d_W]$.

To understand how the parameters Θ and W influence the optimization process, we propose the notion of block condition number for spectral GNNs.

Definition 3 (Block Condition Number). For a spectral GNN parameterized by Θ and W , the *block condition number* of its Hessian matrix \mathbf{H} with respect to the empirical loss is defined as:

$$\kappa'(\mathbf{H}) = \frac{\max(\lambda_{\max}(\mathbf{H}_{\Theta, \Theta}), \lambda_{\max}(\mathbf{H}_{W, W}))}{\min(\lambda_{\max}(\mathbf{H}_{\Theta, \Theta}), \lambda_{\max}(\mathbf{H}_{W, W}))}. \quad (3)$$

Remark 4. The block condition number $\kappa'(\mathbf{H})$ serves as a metric to quantify the disparity between the maximum eigenvalues of the block Hessian matrices $\mathbf{H}_{\Theta, \Theta}$ and $\mathbf{H}_{W, W}$. A spectral GNN is considered *poorly conditioned* if $\kappa'(\mathbf{H})$ is large. According to the interlacing theorem (Horn and Johnson 2012), when a symmetric matrix is divided into blocks according to parameters, the eigenvalues of the submatrices interlace with the eigenvalues of the full matrix. In the case where the Hessian matrix \mathbf{H} is positive semi-definite, we have $\max(\lambda_{\max}(\mathbf{H}_{\Theta, \Theta}), \lambda_{\max}(\mathbf{H}_{W, W})) \leq \lambda_{\max}(\mathbf{H})$ and $\lambda_{\min}(\mathbf{H}) \leq \min(\lambda_{\max}(\mathbf{H}_{\Theta, \Theta}), \lambda_{\max}(\mathbf{H}_{W, W}))$. Therefore, $\kappa'(\mathbf{H}) \leq \kappa(\mathbf{H})$. A problem identified as poorly conditioned using κ' will also be poorly conditioned when assessed using κ . However, the converse does not hold true. Thus, the block condition number offers a more stringent criterion for determining whether a problem is poorly conditioned, as it directly models the impact of the differences between the parameters Θ and W on the optimization process.

In spectral GNNs, the block condition number for heterophilic graphs tends to be larger than that for homophilic graphs. When $\lambda_{\max}(\mathbf{H}_{\Theta, \Theta}) > \lambda_{\max}(\mathbf{H}_{W, W})$, the block condition number $\kappa'(\mathbf{H})$ is reduced to the largest eigenvalue ratio shown in Figure 2, i.e., $\kappa'(\mathbf{H}) = \alpha$ when applying ChbNet on those four datasets. Thus, the optimization on heterophilic graphs is generally more difficult than the optimization on homophilic graphs.

The relation between the optimal learning rate η^* and the Hessian matrix \mathbf{H} is that $\eta^* = \frac{1}{\lambda_{\max}(\mathbf{H})}$ (Granziol, Zohren, and Roberts 2020). For spectral GNNs, we observe that $\lambda_{\max}(\mathbf{H}_{\Theta, \Theta}) \leq \lambda_{\max}(\mathbf{H})$ and $\lambda_{\max}(\mathbf{H}_{W, W}) \leq \lambda_{\max}(\mathbf{H})$. Consequently, the learning rates must satisfy $\eta \leq \eta_{\Theta}$ and $\eta \leq \eta_W$, where η , η_{Θ} , and η_W denote the learning rates for all parameters, Θ , and W , respectively. Given that $\lambda_{\max}(\mathbf{H}_{\Theta, \Theta}) \neq \lambda_{\max}(\mathbf{H}_{W, W})$ in most cases, employing different learning rates for Θ and W is a common practice in the training of spectral GNNs, as evidenced in recent studies (He et al. 2021; He, Wei, and Wen 2022).

Asymmetric Learning

In this section, we first introduce the *gradient-parameter norm ratio* (GPNR), which is computed based on the magnitudes of gradients and parameter values. We then discuss the relationship between GPNR and the maximum eigenvalue of the Hessian matrix. After that, we propose *asymmetric preconditioning* on the gradients of spectral GNNs and examine its impact on the Hessian matrix. We theoretically demonstrate that our asymmetric preconditioning can mitigate the poorly conditioned optimization of spectral GNNs under mild assumptions.

Gradient-Parameter Norm Ratio The update rule for gradient descent in spectral GNNs is given by

$$\begin{aligned} \Theta^t &= \Theta^{t-1} - \eta \nabla_{\Theta} \mathcal{L}_S(\Theta^t, W^t); \\ W^t &= W^{t-1} - \eta \nabla_W \mathcal{L}_S(\Theta^t, W^t), \end{aligned} \quad (4)$$

where η denotes the learning rate.

We introduce the notion of the gradient-parameter norm ratio to estimate the update speed of parameters.




Figure 3: The Gradient-Parameter Norm Ratio (GPNR) of ρ_{Θ} is larger than ρ_W during iterations. It means the updating speed of Θ is faster than that of W .

Definition 5 (Gradient-Parameter Norm Ratio). For a differentiable function \mathcal{L}_S parameterized by Ψ , the gradient-parameter norm ratio (GPNR) of Ψ is defined as

$$\rho_\Psi = \frac{\|\nabla \mathcal{L}_S(\Psi)\|_2}{\|\Psi\|_2}, \quad (5)$$

where $\nabla \mathcal{L}_S(\Psi)$ is the gradient of \mathcal{L}_S with respect to Ψ .

Remark 6. The GPNRs of parameters Θ^t and W^t are $\rho_\Theta^t = \frac{\|\nabla_\Theta \mathcal{L}_S(\Theta^t, W^t)\|_2}{\|\Theta^t\|_2}$ and $\rho_W^t = \frac{\|\nabla_W \mathcal{L}_S(\Theta^t, W^t)\|_2}{\|W^t\|_2}$, respectively, at each iteration t . These ratios are effective metrics quantifying the update speeds of Θ and W . Generally, parameters should be updated at similar speeds (You, Gitman, and Ginsburg 2017), meaning that parameters with larger magnitudes undergo larger updates. However, our analysis of ρ_Θ^t and ρ_W^t , as shown in Figure 3, reveals that ρ_Θ^t is typically larger than ρ_W^t , indicating that Θ is updated more rapidly than W .

We further demonstrate the close relationship between the GPNR and the Hessian matrix.

Assumption 7 (Point Proximity). Let Ψ be a point near a critical point Ψ^* . If $\|\Psi - \Psi^*\|_2$ is close to zero, we assume $\|\Psi - \Psi^*\|_2 \leq \|\Psi\|_2$.

Remark 8. In spectral GNNs, $\Psi = [\Theta; W]$. As the dimension of W is usually in thousands, there is a high probability that $\|\Psi\|_2$ is greater than zero. Thus, when $\|\Psi - \Psi^*\|_2$ is a very small number approximating zero, it is natural to assume that $\|\Psi - \Psi^*\|_2 \leq \|\Psi\|_2$.

Proposition 9 (GPNR and Maximum Eigenvalue of Hessian). For a neural network parameterized by Ψ and the Hessian matrix \mathbf{H}_Ψ of empirical loss at point Ψ , when Ψ is near a critical point Ψ^* under Assumption 7, we have $\rho_\Psi \leq \lambda_{max}(\mathbf{H}_\Psi)$.

Remark 10. The Hessian matrix \mathbf{H}_Ψ provides insights into how gradients change in the vicinity of the point Ψ . The empirical loss \mathcal{L}_S changes rapidly along the director of the eigenvector corresponding to $\lambda_{max}(\mathbf{H}_\Psi)$. Consequently, a small perturbation around Ψ can result in a significant increase in the gradient $\nabla_\Psi \mathcal{L}_S(\Psi)$, leading to a large $\rho_\Psi = \frac{\|\nabla_\Psi \mathcal{L}_S(\Psi)\|_2}{\|\Psi\|_2}$. When the parameter is near a critical point, $\lambda_{max}(\mathbf{H}_\Psi)$ is greater than or equal to the GPNR. Hence, if the GPNR is large, then $\lambda_{max}(\mathbf{H}_\Psi)$ is also large.

Asymmetric Preconditioner We now define asymmetric preconditioning on the gradients of spectral GNNs. Let

$$s_\Theta^t = \frac{\|\Theta^t\|_2}{\|\nabla_\Theta \mathcal{L}_S(\Theta^t, W^t)\|_2}; \quad s_W^t = \frac{\|W^t\|_2}{\|\nabla_W \mathcal{L}_S(\Theta^t, W^t)\|_2}. \quad (6)$$

Then, the *asymmetric preconditioner* for s_Θ^t and s_W^t is a diagonal preconditioner, defined as

$$R^t = \begin{pmatrix} s_\Theta^t \mathbf{I}_{d_\Theta} & \mathbf{0} \\ \mathbf{0} & s_W^t \mathbf{I}_{d_W} \end{pmatrix} \quad (7)$$

where \mathbf{I}_{d_Θ} and \mathbf{I}_{d_W} are the identity matrices of dimensions d_Θ and d_W , respectively. This yields new gradient updates

$$\begin{aligned} & [\bar{\nabla}_\Theta \mathcal{L}_S(\Theta^t, W^t); \bar{\nabla}_W \mathcal{L}_S(\Theta^t, W^t)] \\ &= R^t [\nabla_\Theta \mathcal{L}_S(\Theta^t, W^t); \nabla_W \mathcal{L}_S(\Theta^t, W^t)]. \end{aligned} \quad (8)$$

The GPNRs are also updated: $\rho_\Theta^t = \frac{\|\bar{\nabla}_\Theta \mathcal{L}_S(\Theta^t, W^t)\|_2}{\|\Theta^t\|_2} = 1$ and $\rho_W^t = \frac{\|\bar{\nabla}_W \mathcal{L}_S(\Theta^t, W^t)\|_2}{\|W^t\|_2} = 1$. Consequently, the parameters Θ and W are updated at the same rate, preventing overly aggressive or insufficient updates.

Asymmetric preconditioning affects the Hessian matrix, thereby altering the optimization landscape and the block condition number. The new Hessian matrix is given by

$$\mathbf{H}' = R^t \mathbf{H}. \quad (9)$$

We will prove later that the block condition number decreases after preconditioning, i.e., $\kappa'(\mathbf{H}') \leq \kappa'(\mathbf{H})$.

Moving Average of Parameter Norms To ensure smooth changes in the parameter norms during iterations, especially when an optimizer updates parameters aggressively, we employ the exponential moving average technique (Shalev-Shwartz and Ben-David 2014). The parameter norms π_Θ^t and π_W^t at the t -th iteration are calculated using the exponential moving average as follows:

$$\begin{aligned} \pi_\Theta^t &= \beta_{\pi_\Theta} \pi_\Theta^{t-1} + (1 - \beta_{\pi_\Theta}) \|\Theta^t\|_2; \\ \pi_W^t &= \beta_{\pi_W} \pi_W^{t-1} + (1 - \beta_{\pi_W}) \|W^t\|_2, \end{aligned} \quad (10)$$

where β_{π_Θ} and β_{π_W} are the smoothing factors in $[0, 1]$ for the parameter norms π_Θ and π_W , respectively.

Asymmetric training can be integrated with any optimizer, and the optimization procedure is elaborated in Algorithm 1. Initially, the weights Θ^0 and W^0 are randomly initialized, and the best validation loss $\mathcal{L}_{\mathcal{D}_{val}}^b$ is set to infinity. During each iteration, the raw gradient $[\nabla_\Theta \mathcal{L}_S(\Theta^t, W^t); \nabla_W \mathcal{L}_S(\Theta^t, W^t)]$ of the empirical loss is computed. Based on the specific update rules of the chosen optimizer OP , the update vectors δ_Θ^t and δ_W^t for parameters Θ and W are calculated. Subsequently, the moving average of parameter norms π_Θ^t, π_W^t and elements s_Θ^t, s_W^t in the asymmetric preconditioner R^t are calculated as follows:

$$s_\Theta^t = \frac{\|\pi_\Theta^t\|_2}{\|\delta_\Theta^t\|_2}; \quad s_W^t = \frac{\|\pi_W^t\|_2}{\|\delta_W^t\|_2}. \quad (11)$$

The minimum validation loss $\mathcal{L}_{\mathcal{D}_{val}}^b$ is then updated, along with the corresponding best parameters Θ^b, W^b . The new gradient $[\bar{\nabla}_\Theta \mathcal{L}_S(\Theta^t, W^t); \bar{\nabla}_W \mathcal{L}_S(\Theta^t, W^t)]$ is computed by preconditioning the raw gradient $[\nabla_\Theta \mathcal{L}_S(\Theta^t, W^t); \nabla_W \mathcal{L}_S(\Theta^t, W^t)]$ using the asymmetric preconditioner R^t . This new gradient serves as the input for any optimizer OP , resulting in new update vectors $\bar{\delta}_\Theta^t, \bar{\delta}_W^t$. Consequently, parameters Θ^{t+1} and W^{t+1} for the next iteration are updated accordingly.

Smaller Block Condition Number. Preconditioning the gradient $\nabla_\Psi \mathcal{L}_S(\Psi^t) = [\nabla_\Theta \mathcal{L}_S(\Theta^t, W^t); \nabla_W \mathcal{L}_S(\Theta^t, W^t)]$ with asymmetric preconditioner R^t is equivalent to preconditioning the Hessian matrix \mathbf{H}^t with R^t in the t -th iteration:

$$R^t \mathbf{H}^t \Leftrightarrow R^t [\nabla_\Theta \mathcal{L}_S(\Theta^t, W^t); \nabla_W \mathcal{L}_S(\Theta^t, W^t)]$$

where $\mathbf{H}_{ij}^t = \frac{\nabla_{\Psi_i} \mathcal{L}_S(\Psi^t)}{\nabla_{\Psi_j} \mathcal{L}_S(\Psi^t)}$, $(R^t \mathbf{H}^t)_{ij} = \frac{r_i^t \nabla_{\Psi_i} \mathcal{L}_S(\Psi^t)}{\nabla_{\Psi_j} \mathcal{L}_S(\Psi^t)}$, and r_i^t is the i -th diagonal element in R^t . To simplify the analysis,

Algorithm 1: Asymmetric Training

Input : Training set $S = (X, \{y_i\}_{i=1}^m)$, validation set $\mathcal{D}_{val} = (X, \{y_i\}_{i=m+1}^{m+q})$, loss function ℓ , learning rate η , maximum iteration number t_{max} , and optimizer $OP(\cdot)$

Output : Trained model

```
2 Initialize weights  $\Theta^0$  and  $W^0$ , iteration number  $t = 0$ , best
  validation loss  $\mathcal{L}_{\mathcal{D}_{val}}^b = +\infty$ ;
3 while  $t \leq t_{max}$  do
4   Compute  $[\nabla_{\Theta} \mathcal{L}_S(\Theta^t, W^t); \nabla_W \mathcal{L}_S(\Theta^t, W^t)]$ ;
5    $\delta_{\Theta}^t, \delta_W^t \leftarrow OP([\nabla_{\Theta} \mathcal{L}_S(\Theta^t, W^t); \nabla_W \mathcal{L}_S(\Theta^t, W^t)])$ ;
6   Set parameter norms  $\pi_{\Theta}^t$  and  $\pi_W^t$  by Eq. (10);
7   Compute  $s_{\Theta}^t, s_W^t$  by Eq. (11);
8   Compute asymmetric preconditioner  $R^t$  by Eq. (7);
9   Compute validation loss  $\mathcal{L}_{\mathcal{D}_{val}}(\Theta^t, W^t)$ ;
10  if  $\mathcal{L}_{\mathcal{D}_{val}}(\Theta^t, W^t) \leq \mathcal{L}_{\mathcal{D}_{val}}^b$  then
11     $\Theta^b = \Theta^t, W^b = W^t$ ;
12     $\mathcal{L}_{\mathcal{D}_{val}}^b = \mathcal{L}_{\mathcal{D}_{val}}(\Theta^t, W^t)$ ;
13  end
14  Get  $[\bar{\nabla}_{\Theta} \mathcal{L}_S(\Theta^t, W^t); \bar{\nabla}_W \mathcal{L}_S(\Theta^t, W^t)]$  by Eq. (8);
15   $\bar{\delta}_{\Theta}^t, \bar{\delta}_W^t \leftarrow OP([\bar{\nabla}_{\Theta} \mathcal{L}_S(\Theta^t, W^t); \bar{\nabla}_W \mathcal{L}_S(\Theta^t, W^t)])$ ;
16   $\Theta^{t+1} = OP(\eta, \Theta^t, \bar{\delta}_{\Theta}^t)$ ;
17   $W^{t+1} = OP(\eta, W^t, \bar{\delta}_W^t)$ ;
18   $t = t + 1$ ;
19 end
20 return  $\Theta^b, W^b$ 
```

we set $\beta_{\pi_{\Theta}} = \beta_{\pi_W} = 0$ and consider the gradient descent optimizer (Shalev-Shwartz and Ben-David 2014) for training. We show that preconditioning with the asymmetric preconditioner results in a smaller block condition number of the Hessian when the following assumptions are met. The proofs are provided in detail in the Appendix.

Assumption 11 (Proportional GPNR and Maximum Eigenvalue). Let Ψ and Ψ' be two points near a critical point Ψ^* such that $\|\Psi - \Psi^*\|_2 = \epsilon_1$ and $\|\Psi' - \Psi^*\|_2 = \epsilon_2$, where $\epsilon_1 \rightarrow 0$ and $\epsilon_2 \rightarrow 0$. Denote \mathbf{H}_{Ψ} and $\mathbf{H}_{\Psi'}$ as the Hessian matrices of the empirical loss at points Ψ and Ψ' , respectively. If $\lambda_{\max}(\mathbf{H}_{\Psi}) \geq \lambda_{\max}(\mathbf{H}_{\Psi'})$, then $\rho_{\Psi} \geq \rho_{\Psi'}$.

Remark 12. The relationship between the GPNR and the largest eigenvalue of Hessian is that $\rho_{\Psi} \leq \lambda_{\max}(\mathbf{H}_{\Psi})$ when Ψ is near a critical point Ψ^* as stated in Proposition 9. Consequently, it is reasonable to assume that if the largest eigenvalues of the Hessian matrices at points Ψ and Ψ' satisfy $\lambda_{\max}(\mathbf{H}_{\Psi}) \geq \lambda_{\max}(\mathbf{H}_{\Psi'})$, then the GPNRs at these points also satisfy $\rho_{\Psi} \geq \rho_{\Psi'}$.

Assumption 13 (Mild Scaling). For the asymmetric preconditioner R^t and the Hessian matrix \mathbf{H}^t , s_{Θ}^t and s_W^t in R^t satisfy the inequality $\frac{s_{\Theta}^t}{s_W^t} \geq \frac{\lambda_{\max}(\mathbf{H}_{W,W}^t)}{\lambda_{\max}(\mathbf{H}_{\Theta,\Theta}^t)}$.

Remark 14. This assumption is also mild. The preconditioner is assumed to ensure that Proposition 15 holds. Specifically, it implies that the curvature of the loss function does not undergo significant changes after preconditioning.

Proposition 15 (Spectral Equivalent Preconditioner.). Let R^t be the asymmetric preconditioner and \mathbf{H}^t the Hessian matrix at the t -th iteration. If $\lambda_{\max}(\mathbf{H}_{\Theta,\Theta}^t) \geq \lambda_{\max}(\mathbf{H}_{W,W}^t)$ and $\frac{s_{\Theta}^t}{s_W^t} \geq \frac{\lambda_{\max}(\mathbf{H}_{W,W}^t)}{\lambda_{\max}(\mathbf{H}_{\Theta,\Theta}^t)}$, we have $\lambda_{\max}(\mathbf{H}_{\Theta,\Theta}^{t'}) \geq \lambda_{\max}(\mathbf{H}_{W,W}^{t'})$ after preconditioning $\mathbf{H}^{t'} = R^t \mathbf{H}^t$.

Remark 16. This proposition implies that the curvature of the loss function remains largely unchanged after preconditioning when the preconditioner R^t satisfies certain conditions. The sensitivity of the parameters to the loss function is maintained after preconditioning. If the loss function is more sensitive to Θ than W , then Θ remains more sensitive than W after preconditioning, and vice versa.

Theorem 17 (Smaller Block Condition Number.). Given an asymmetric preconditioner R^t , the Hessian matrix \mathbf{H}^t at the t -th iteration, and the gradient descent as the optimizer, under Assumption 7, Assumption 11 and Assumption 13, we have $\kappa'(\mathbf{H}^{t'}) \leq \kappa'(\mathbf{H}^t)$ after preconditioning $\mathbf{H}^{t'} = R^t \mathbf{H}^t$.

Remark 18. This theorem indicates that the block condition number decreases in each iteration during training when applying the asymmetric preconditioner. When the block condition number of the Hessian matrix approaches one, changes in Θ and W lead to approximately the same changes in the empirical loss. As the block condition number evaluates how poorly conditioned an optimization problem is, a small block condition number indicates an easy optimization problem.

Experiments

Experimental Setup. To empirically validate the effectiveness of our approach, we perform experiments on eighteen benchmark datasets for node classification tasks, including: six small heterophilic graphs (Texas, Wisconsin, Actor, Chameleon, Squirrel, Cornell), five large heterophilic graphs (Roman_Empire, Amazon_Ratings, Minesweeper, Tolokers, Questions) and seven homophilic graphs (Citeseer, Pubmed, Cora, Computers, Photo, Coauthor-CS, Coauthor-Physics). For each dataset, a sparse splitting is employed where nodes are randomly partitioned into train/validation/test sets with ratios of 2.5%/2.5%/95%, respectively. For Citeseer, Pubmed, and Cora datasets, the setting from (Chien et al. 2021; He, Wei, and Wen 2022) is used: 20 nodes per class are for training, 500 nodes for validation, and 1,000 nodes for testing.

We evaluate five prominent spectral GNN models as our baselines: ChebNet (Defferrard, Bresson, and Vandergheynst 2016), GPRGNN (Chien et al. 2021), BernNet (He et al. 2021), JacobiConv (Wang and Zhang 2022b), and ChebNetII (He, Wei, and Wen 2022), and compare their performances in the setups with and without asymmetric learning across all datasets. For each GNN model, we use GNN (S) to denote the model trained using Adam optimizer (Kingma and Ba 2014) and GNN (AS) to denote the model trained using Adam optimizer with asymmetric learning, with $\Delta \uparrow$ indicating the performance improvement.

Statistics of datasets, details about the baselines, and the setting of hyper-parameters are included in the Appendix.

Model	Texas	Wisconsin	Actor	Chameleon	Squirrel	Cornell	Roman_Empire	Amazon_Ratings
ChebNet (S)	36.76±6.76	33.17±7.83	24.94±1.33	25.65±3.56	22.58±2.56	40.52±8.15	45.32±0.65	38.87±0.26
ChebNet (AS)	50.81±8.03	33.21±5.33	25.13±0.83	38.19±1.26	27.24±1.95	40.92±7.92	48.11±0.41	39.38±0.19
$\Delta \uparrow$	+14.05	+0.04	+0.19	+12.54	+4.66	+0.40	+2.79	+0.51
ChebNetII (S)	48.44±10.87	41.33±9.25	29.29±0.8	33.48±4.59	30.8±2.65	37.69±11.1	55.06±0.32	38.33±0.42
ChebNetII (AS)	50.75±10.0	47.33±10.17	29.49±0.9	35.58±3.7	35.94±2.66	51.85±6.01	55.2±0.3	39.07±0.28
$\Delta \uparrow$	+2.31	+6.00	+0.20	+2.10	+5.14	+14.16	+0.14	+0.74
JacobiConv (S)	50.17±7.87	46.08±9.08	32.66±0.71	26.57±3.6	23.15±4.47	49.71±7.75	52.92±0.7	40.18±0.53
JacobiConv (AS)	52.66±8.9	49.5±9.75	32.78±0.71	29.95±3.37	24.71±3.83	52.83±8.73	54.16±0.38	47.17±1.37
$\Delta \uparrow$	+2.49	+3.42	+0.12	+3.38	+1.56	+3.12	+1.24	+6.99
GPRGNN (S)	48.55±7.0	40.79±3.62	30.2±0.95	30.44±4.29	24.33±2.68	46.53±7.23	55.48±1.3	39.86±0.3
GPRGNN (AS)	49.54±5.5	43.46±6.29	30.75±0.76	31.63±4.41	24.82±4.63	47.63±6.71	55.99±1.14	39.89±0.16
$\Delta \uparrow$	+0.99	+2.67	+0.55	+1.19	+0.49	+1.10	+0.51	+0.03
BernNet (S)	52.14±8.09	49.33±8.21	30.57±0.78	29.45±3.22	25.94±3.35	48.15±6.59	55.3±0.41	39.33±0.25
BernNet (AS)	53.87±9.66	52.46±6.29	30.8±0.69	30.46±5.31	27.5±3.62	49.88±8.73	55.5±0.2	39.78±0.3
$\Delta \uparrow$	+1.73	+3.13	+0.23	+1.01	+1.56	+1.73	+0.20	+0.45

Table 1: Test accuracy with/without asymmetric learning on small heterophilic datasets (Texas, Wisconsin, Actor, Chameleon, Squirrel, Cornell) and large heterophilic datasets (Roman-Empire, Amazon-Ratings). High accuracy indicates good performance.

Model	Citeseer	Pubmed	Cora	Computers	Photo	Coauthor-CS	Coauthor-Physics
ChebNet (S)	65.49±1.39	74.88±1.47	80.26±0.86	78.76±1.98	90.26±2.19	90.74±0.6	94.69±0.31
ChebNet (AS)	66.63±1.1	76.53±1.52	81.02±0.92	79.08±2.6	92.38±0.41	90.76±0.73	94.92±0.1
$\Delta \uparrow$	+1.14	+1.65	+0.76	+0.32	+2.12	+0.02	+0.23
ChebNetII (S)	70.07±1.13	78.87±1.28	82.55±0.78	83.49±0.9	92.51±0.43	92.18±0.23	94.96±0.13
ChebNetII (AS)	70.75±0.78	79.5±0.94	82.52±0.95	84.38±1.3	92.69±0.43	93.41±0.23	95.13±0.09
$\Delta \uparrow$	+0.68	+0.63	-0.03	+0.89	+0.18	+1.23	+0.17
JacobiConv (S)	67.86±1.08	76.89±1.67	77.44±1.23	84.77±0.7	92.18±0.42	93.06±0.3	95.07±0.16
JacobiConv (AS)	68.17±0.9	77.77±1.37	77.68±1.24	84.91±0.46	92.54±0.27	93.22±0.31	95.18±0.16
$\Delta \uparrow$	+0.31	+0.88	+0.24	+0.14	+0.36	+0.16	+0.11
GPRGNN (S)	70.23±1.29	79.86±1.81	83.5±0.94	83.5±0.99	92.46±0.29	92.61±0.19	95.16±0.08
GPRGNN (AS)	70.23±1.03	80.66±1.64	83.26±0.69	84.18±0.71	92.68±0.43	93.06±0.49	95.24±0.13
$\Delta \uparrow$	+0.00	+0.80	-0.24	+0.68	+0.22	+0.45	+0.08
BernNet (S)	68.64±1.54	79.58±0.87	82.37±0.96	85.1±0.97	92.65±0.22	92.3±0.27	95.29±0.15
BernNet (AS)	69.46±1.16	79.62±1.24	82.47±0.56	85.57±0.96	92.82±0.36	93.13±0.37	95.19±0.18
$\Delta \uparrow$	+0.82	+0.04	+0.10	+0.47	+0.17	+0.83	-0.10

Table 2: Test accuracy with/without asymmetric learning on homophilic graphs (Citeseer, Pubmed, Cora, Coauthor-CS, Coauthor-Physics, Amazon-Computer and Amazon-Photo). High accuracy indicates good performance.

Results and Discussion. Tables 1 and 2 present the results for heterophilic graphs and homophilic graphs, respectively. Additional results for large heterophilic graphs can be further found in Table 6 in the Appendix.

Q1: How does asymmetric learning perform with graphs of varying edge homophily ratios? We notice that asymmetric learning provides more performance gains on heterophilic graphs compared to homophilic ones. As shown in Table 1, on six small heterophilic datasets, there is an average accuracy improvement of 3.08%, while an accuracy increase of approximately 1.36% is achieved on Romain_Empire and Amazon_Ratings. In Table 2, the average accuracy improvement is 0.46% on homophilic graphs. These results align with our theoretical analysis, showing that asymmetric learning facilitates the optimization of poorly conditioned problems by reducing the block condition number of the Hessian matrix. Empirically, we observe that the block condition numbers of Hessian matrices are larger on heterophilic graphs compared to homophilic graphs. Consequently, asymmetric learning yields a more significant improvement in classification accuracy on heterophilic graphs than on homophilic graphs.

Q2: How does asymmetric learning perform with differ-

ent spectral GNNs? Spectral GNNs can be categorized based on their polynomial basis into orthogonal spectral GNNs (ChebNet, ChebNetII, and JacobiConv) and non-orthogonal spectral GNNs (GPRGNN and BernNet). We observe that asymmetric learning generally exhibits superior performance with orthogonal spectral GNNs compared to non-orthogonal spectral GNNs. For example, there is a 5.15% accuracy improvement with orthogonal spectral GNNs using asymmetric learning on the six small heterophilic datasets (Table 1), whereas non-orthogonal spectral GNNs show a 1.69% accuracy improvement. A plausible explanation is that the interference between different orders of orthogonal polynomial bases is significantly lower than that between non-orthogonal polynomial bases. Suppose that $\Theta = \{\theta_k\}_{k=0}^K$, where each θ_k is associated with a polynomial basis T_k of order k . Due to the interference in non-orthogonal polynomial bases, a sophisticated preconditioner is required, rather than simply scaling all θ_k uniformly by s_Θ , to achieve better performance.

Q3: How does asymmetric learning perform with graphs of different sizes? Asymmetric learning shows greater performance improvement on datasets with smaller graph sizes. In Table 1, there is an average accuracy improvement of




Figure 4: The test accuracy of ChebNet when applying asymmetric learning with varying values of β_{π_Θ} and β_{π_W} .

3.08% on smaller graphs on the left. For larger graphs on the right, improvements are smaller, with a 1.36% increase in accuracy. Specifically, spectral GNNs achieve test accuracy improvements exceeding 10% on smaller datasets such as Texas, Wisconsin, Cornell, and Chameleon, which contain fewer than 1,000 nodes. As smaller training data sizes can often lead to instability in the gradient of empirical loss during training, these findings suggest that Asymmetric learning mitigates this issue by adjusting parameter updates without risking overly aggressive updates.

Q4: How do the exponential decay parameters β_{π_Θ} and β_{π_W} affect the performance of asymmetric learning? We explore the impact of the hyperparameters β_{π_Θ} and β_{π_W} on test accuracy when using ChebNet for node classification tasks. The parameters β_{π_Θ} and β_{π_W} increase from 0 to 0.9 in steps of 0.1, with an additional value at 0.99. Figure 4 illustrates that fluctuations in performance are observed only in the Texas dataset, while marginal differences are noted in other datasets. This suggests that the performance is not sensitive to β_{π_Θ} and β_{π_W} when training samples are sufficient. Maintaining a training data splitting ratio of 2.5% can ensure an adequate amount of training data for large graphs.

Related Works

Spectral GNNs and Optimization Spectral GNNs transform graph signals from the vertex domain to the spectral domain for filtering and then applying an inverse transformation to return the filtered signals to the vertex domain (Defferrard, Bresson, and Vandergheynst 2016). Based on the design of their graph filters, spectral GNNs are categorized into polynomial and rational types. Polynomial spectral GNNs employ various polynomial basis to approximate graph filters. For example, ChebNet and ChebNetII utilize Chebyshev polynomial basis (Defferrard, Bresson, and Vandergheynst 2016; He, Wei, and Wen 2022), while JacobiConv uses Jacobi polynomial basis (Wang and Zhang 2022b). These models leverage orthogonal polynomial basis, which allow efficient approximation as approximation error decreases exponentially with increasing polynomial order. Non-orthogonal polynomial spectral GNNs are also noteworthy, such as BernNet (He et al. 2021) which uses Bernstein polynomial basis, and GPRGNN (Chien et al. 2021) which employs monomial polynomial basis. Rational spectral GNNs offer greater flexibility at the cost of increased complexity. For instance, CayleyNet applies parametric rational functions to identify

relevant frequency bands (Levie et al. 2019); ARMA incorporates moving average terms to enhance noise robustness and capture long-term dynamics on graphs (Bianchi et al. 2021). LanczosNet facilitates multi-scale analysis of graph signals (Liao et al. 2019). Polynomial spectral GNNs are computationally efficient and enable localized vertex analysis (Hammond, Vandergheynst, and Gribonval 2011). This paper focuses on polynomial spectral GNNs.

A crucial area in spectral GNNs is the optimization. Studies on optimization for GNNs are very limited. Xu et al. (2021) have investigated the effects of network depth and skip connections on the convergence behavior of linear GNNs. Izadi et al. (2020) have demonstrated the use of natural gradient methods during training to produce more effective solutions. A common practice in optimizing spectral GNNs is to apply different learning rates to graph convolution parameters and feature transformation parameters (He et al. 2021; He, Wei, and Wen 2022). Nevertheless, the rationale behind this practice remains inadequately explored.

Learning Rate Selection and Preconditioning The optimization of neural networks heavily depends on the choices of learning rates for parameter updates. Studies emphasize that the learning rate should not exceed the reciprocal of the largest eigenvalue of the Hessian matrix to prevent overshooting and convergence to inferior solutions in gradient descent algorithms (Granzol, Zohren, and Roberts 2020; Ding 2021).

The optimization of spectral GNNs is further complicated by the distinct characteristics of graph convolution parameters Θ and feature transformation parameters W , leading to a poorly conditioned problem. A popular approach to tackle poor-conditioned optimization problems is *preconditioning* (Saad 2003). First-order methods such as AdaGrad (Duchi, Hazan, and Singer 2011), RMSProp (Tieleman and Hinton 2012), and Adam (Kingma and Ba 2014) all incorporate a preconditioning component. These methods adapt learning rates based on historical gradients, assuming that all parameters and their gradients are of the same magnitude. Low-rank preconditioners are widely used in the large-scale matrix factorization and deep neural architectures (Jia et al. 2024; Gupta, Koren, and Singer 2018; Yen et al. 2024). However, there are no works specifically tailored to address the asymmetric parameters in spectral GNNs.

Conclusion and Future Work

We have demonstrated that the parameters Θ and W in spectral GNNs play distinct roles, resulting in a poorly conditioned optimization problem. To address this, we introduced the concept of the block condition number to estimate the optimization difficulty and proposed asymmetric learning, a novel preconditioning technique designed to facilitate the optimization process. Our experiments validate the efficacy of asymmetric learning, particularly in the optimization on heterophilic graphs. These findings underscore the potential of asymmetric learning to enhance the performance and training efficiency of spectral GNNs across diverse graph datasets. Future work will explore further applications of this approach to other types of GNNs and graph-based learning tasks.

Acknowledgements

We are thankful to Yun Kuen Cheung for helpful discussions. This research was supported partially by the Australian Government through the Australian Research Council's Discovery Projects funding scheme (project DP210102273).

References

- Balcilar, M.; Renton, G.; Héroux, P.; Gaüzère, B.; Adam, S.; and Honeine, P. 2021. Analyzing the Expressive Power of Graph Neural Networks in a Spectral Perspective. In *ICLR*.
- Bianchi, F. M.; Grattarola, D.; Livi, L.; and Alippi, C. 2021. Graph neural networks with convolutional arma filters. *IEEE transactions on pattern analysis and machine intelligence*, 44(7): 3496–3507.
- Bonnans, J.-F.; Gilbert, J. C.; Lemaréchal, C.; and Sagastizábal, C. A. 2006. *Numerical optimization: theoretical and practical aspects*. Springer Science & Business Media.
- Chien, E.; Peng, J.; Li, P.; and Milenkovic, O. 2021. Adaptive Universal Generalized PageRank Graph Neural Network. In *ICLR*.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *NeurIPS*.
- Ding, Y. 2021. The Impact of Learning Rate Decay and Periodical Learning Rate Restart on Artificial Neural Network. In *AIEE*.
- Duchi, J. C.; Hazan, E.; and Singer, Y. 2011. Adaptive Sub-gradient Methods for Online Learning and Stochastic Optimization. *J. Mach. Learn. Res.*, 12: 2121–2159.
- Dwivedi, V. P.; Joshi, C. K.; Luu, A. T.; Laurent, T.; Bengio, Y.; and Bresson, X. 2023. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24(43): 1–48.
- Golub, G. H.; and Van Loan, C. F. 2013. *Matrix computations*. JHU press.
- Granzol, D.; Zohren, S.; and Roberts, S. J. 2020. Learning Rates as a Function of Batch Size: A Random Matrix Theory Approach to Neural Network Training. *J. Mach. Learn. Res.*, 23: 173:1–173:65.
- Gupta, V.; Koren, T.; and Singer, Y. 2018. Shampoo: Preconditioned Stochastic Tensor Optimization. In *ICML*.
- Hammond, D. K.; Vandergheynst, P.; and Gribonval, R. 2011. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2): 129–150.
- He, M.; Wei, Z.; Huang, Z.; and Xu, H. 2021. BernNet: Learning Arbitrary Graph Spectral Filters via Bernstein Approximation. In *NeurIPS*.
- He, M.; Wei, Z.; and Wen, J.-R. 2022. Convolutional neural networks on graphs with chebyshev approximation, revisited. *NeurIPS*.
- Horn, R. A.; and Johnson, C. R. 2012. *Matrix analysis*. Cambridge university press.
- Izadi, M. R.; Fang, Y.; Stevenson, R. L.; and Lin, L. 2020. Optimization of Graph Neural Networks with Natural Gradient Descent. In *Big Data*.
- Jia, X.; Wang, H.; Peng, J.; Feng, X.; and Meng, D. 2024. Preconditioning Matters: Fast Global Convergence of Non-convex Matrix Factorization via Scaled Gradient Descent. In *NeurIPS*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *CoRR*, abs/1412.6980.
- Kipf, T.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- LeCun, Y.; Simard, P. Y.; and Pearlmutter, B. A. 1992. Automatic Learning Rate Maximization by On-Line Estimation of the Hessian's Eigenvectors. In *NeurIPS*.
- Levie, R.; Monti, F.; Bresson, X.; and Bronstein, M. M. 2019. CayleyNets: Graph Convolutional Neural Networks With Complex Rational Spectral Filters. *IEEE Transactions on Signal Processing*, 67: 97–109.
- Liao, R.; Zhao, Z.; Urtasun, R.; and Zemel, R. 2019. LanczosNet: Multi-Scale Deep Graph Convolutional Networks. In *ICLR*.
- Lim, D.; Hohne, F.; Li, X.; Huang, S.; Gupta, V.; Bhalerao, O.; and Lim, S.-N. 2021. Large Scale Learning on Non-Homophilous Graphs: New Benchmarks and Strong Simple Methods. In *NeurIPS*.
- Loukas, A. 2020. What graph neural networks cannot learn: depth vs width. In *ICLR*.
- Pei, H.; Wei, B.; Chang, K. C.-C.; Lei, Y.; and Yang, B. 2020. Geom-GCN: Geometric Graph Convolutional Networks. In *ICLR*.
- Platonov, O.; Kuznedelev, D.; Diskin, M.; Babenko, A.; and Prokhorenkova, L. 2023. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? In *ICLR*.
- Rozemberczki, B.; Allen, C.; and Sarkar, R. 2021. Multi-scale Attributed Node Embedding. *J. Complex Networks*, 9.
- Saad, Y. 2003. *Iterative Methods for Sparse Linear Systems*. SIAM.
- Shalev-Shwartz, S.; and Ben-David, S. 2014. *Understanding machine learning: From theory to algorithms*. Cambridge university press.
- Shchur, O.; Mumme, M.; Bojchevski, A.; and Günnemann, S. 2018. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*.
- Tao, T. 2012. *Topics in random matrix theory*, volume 132. American Mathematical Soc.
- Tieleman, T.; and Hinton, G. 2012. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2): 26–31.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio', P.; and Bengio, Y. 2018. Graph Attention Networks. In *ICLR*.
- Wang, X.; and Zhang, M. 2022a. How powerful are spectral graph neural networks. In *ICML*, 23341–23362.
- Wang, X.; and Zhang, M. 2022b. How powerful are spectral graph neural networks. In *ICML*.

- Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Yu, P. S. 2019. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32: 4–24.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *ICLR*.
- Xu, K.; Zhang, M.; Jegelka, S.; and Kawaguchi, K. 2021. Optimization of Graph Neural Networks: Implicit Acceleration by Skip Connections and More Depth. In *ICML*.
- Yang, M.; Shen, Y.; Li, R.; Qi, H.; Zhang, Q.; and Yin, B. 2022. A new perspective on the effects of spectrum in graph neural networks. In *ICML*.
- Yen, J.-N.; Duvvuri, S. S.; Dhillon, I.; and Hsieh, C.-J. 2024. Block Low-Rank Preconditioner with Shared Basis for Stochastic Optimization. In *NeurIPS*.
- You, Y.; Gitman, I.; and Ginsburg, B. 2017. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*.
- Zeng, H.; Zhou, H.; Srivastava, A.; Kannan, R.; and Prasanna, V. 2020. GraphSAINT: Graph Sampling Based Inductive Learning Method. In *ICML*.
- Zhu, J.; Yan, Y.; Zhao, L.; Heimann, M.; Akoglu, L.; and Koutra, D. 2020. Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs. In *NeurIPS*.