

Non-backtracking Graph Neural Networks

Anonymous authors

Paper under double-blind review

Abstract

The celebrated message-passing updates for graph neural networks allow representing large-scale graphs with local and computationally tractable updates. However, the updates suffer from backtracking, i.e., a message flowing through the same edge twice and revisiting the previously visited node. Since the number of message flows increases exponentially with the number of updates, the redundancy in local updates prevents the graph neural network from accurately recognizing a particular message flow relevant for downstream tasks. In this work, we propose to resolve such redundancy issue via the non-backtracking graph neural network (NBA-GNN) that updates a message without incorporating the message from the previously visited node. We theoretically investigate how NBA-GNN alleviates the over-squashing of GNNs, and establish a connection between NBA-GNN and the impressive performance of non-backtracking updates for stochastic block model recovery. Furthermore, we empirically verify the effectiveness of our NBA-GNN on long-range graph benchmark and transductive node classification problems.

1 Introduction

Recently, graph neural networks (GNNs) (Kipf & Welling, 2017; Hamilton et al., 2017; Xu et al., 2019) have shown great success in various applications, such as molecular property prediction (Gilmer et al., 2017) and community detection (Bruna & Li, 2017). Such success can be largely attributed to the message-passing structure of GNNs, which provides a computationally tractable way of incorporating the overall graph through iterative updates based on local neighborhoods. However, the message-passing structure also brings challenges due to the parallel updates and memoryless behavior of messages passed along the graph.

In particular, the message flow in a GNN is prone to backtracking, where the message from vertex i to vertex j is reincorporated in the subsequent message from j to i , e.g., the left image of step 2 in Figure 1. Since the message-passing iteratively aggregates the information, the GNN inevitably encounters an exponential surge in the number of message flows, proportionate to the vertex degrees. This issue is compounded by backtracking, which accelerates the growth of message flows with redundant information.

The effectiveness of non-backtracking updates has been extensively explored in non-GNN message-passing algorithms or random walks (Fitzner & van der Hofstad, 2013; Rappaport et al., 2017) (Figure 1). For example, given a pair of vertices i, j , the belief propagation algorithm (Pearl, 1982) forbids an $i \rightarrow j$ message from incorporating the $j \rightarrow i$ message. Researchers have similarly considered non-Markovian walks (Alon et al., 2007), i.e., walks that do not consecutively traverse the same edge twice. Such classic algorithms have demonstrated great success in applications like probabilistic graphical model inference and stochastic block




Figure 1: Comparison between message flows of simple (left) and non-backtracking (right) updates.

models (Bordenave et al., 2015). In particular, the spectrum of the non-backtracking operator contains more useful information than that of the adjacency matrix in revealing the hidden structure of a graph model (Bordenave et al., 2015).

However, despite their promise, non-backtracking updates have received limited attention in the context of GNNs. While Chen et al. (2018) already considered combining non-backtracking updates with the typical GNN updates, the updates are not designed to prevent backtracking updates, and the analysis is limited to the optimization landscape of linear GNNs. Hence, a thorough investigation of the benefits of using non-backtracking updates in GNNs is warranted.

Contribution. In this paper, we propose to use a non-backtracking graph neural network (NBA-GNN) to resolve the redundant messages, the over-squashing phenomenon. We employ non-backtracking updates on the messages, i.e., forbid a message from vertex i to vertex j from being incorporated in the message from vertex j to i . To this end, we associate hidden features with transitions between a pair of vertices, e.g., $h_{j \rightarrow i}$, and update them from features associated with non-backtracking transitions, e.g., $h_{k \rightarrow j}$ for $k \neq i$.

To motivate our work, we formulate “message flows” as the sensitivity of a GNN concerning walks in the graph. Then we explain how the message flows are redundant; the GNN’s sensitivity of a walk with backtracking transitions can be covered by that of other non-backtracking walks. We explain how the redundancy harms the GNN since the number of walks increases exponentially as the number of layers grows and the GNN becomes insensitive to a particular walk information. Hence, reducing the redundancy by simply considering non-backtracking walks would benefit the message-passing updates to better recognize each walk’s information. As a motivating example, we provide a connection between our sensitivity analysis to the over-squashing phenomenon for GNN (Topping et al., 2022; Black et al., 2023; Di Giovanni et al., 2023) in terms of access time.

Furthermore, we analyze our NBA-GNNs from the perspective of over-squashing and their expressive capability to recover sparse stochastic block models (SBMs). To this end, we prove that NBA-GNN improves the Jacobian-based measure of over-squashing (Topping et al., 2022) compared to its original GNN counterpart.

Next, we investigate NBA-GNN’s proficiency in node classification within SBMs and its ability to distinguish between graphs originating from the Erdős–Rényi model or the SBM, from the results of (Stephan & Massoulié, 2022; Bordenave et al., 2015). Unlike traditional GNNs that operate on adjacency matrices and necessitate an average degree of at least $\Omega(\log n)$, NBA-GNN demonstrates the ability to perform node classification with a substantially lower average degree bound of $\omega(1)$ and $n^{o(1)}$. Furthermore, the algorithm can accurately classify graphs even when the average degree remains constant.

Finally, we empirically evaluate our NBA-GNN on the long-range graph benchmark (Dwivedi et al., 2022) and transductive node classification problems (Sen et al., 2008; Pei et al., 2019). We observe that our NBA-GNN demonstrates competitive performance and even achieves state-of-the-art performance on the long-range graph benchmark. For the node classification tasks, we demonstrate that NBA-GNN consistently improves over its conventional GNN counterpart.

To summarize, our contributions are as follows:

- We propose NBA-GNN as a solution for the message flow redundancy problem in GNNs.
- We analyze how NBA-GNN alleviates over-squashing and is expressive enough to recover sparse stochastic block models with an average degree of $o(\log n)$.
- We empirically verify our NBA-GNNs to show state-of-the-art performance on the long-range graph benchmark and consistently improve over the conventional GNNs for the transductive node classification tasks.

2 Related works

Non-backtracking Algorithms. Non-backtracking updates have been considered by many classical algorithms (Newman, 2013; Kempton, 2016). Belief propagation (Pearl, 1982) infers the marginal distribution on probabilistic graphical models, and has demonstrated success for tree graphs (Kim & Pearl, 1983) and

graphs with large girth (Murphy et al., 1999). Moreover, Mahé et al. (2004) and Aziz et al. (2013) suggest graph kernels between labeled graphs utilizing non-backtracking walks, and Krzakala et al. (2013) has first used it for node classification. Furthermore, the non-backtracking has been shown to yield better spectral separation properties, and its eigenspace contains information about the hidden structure of a graph model (Bordenave et al., 2015; Stephan & Massoulié, 2022).

Non-backtracking and GNNs. We also note that there have been similar approaches of applying non-backtracking to GNN. Chen et al. (2018) first used the non-backtracking operator in GNNs, though they do not prevent backtracking and only target community detection tasks. Also, Chen et al. (2022) has computed non-redundant tree for every node to eliminate redundancy, but inevitably suffering high complexity. We emphasize the distinction between prior works and our NBA-GNN in Appendix A from the lens of (i) computational complexity and (ii) empirical performance.

Over-squashing of GNNs. When a node receives information from a k -hop neighbor node, an exponential number of messages passes through node representations with fixed-sized vectors. This leads to the loss of information denoted as *over-squashing* (Alon & Yahav, 2021), and has been formalized in terms of sensitivity (Topping et al., 2022; Di Giovanni et al., 2023). Hence, sensitivity is defined as the Jacobian of a final node feature at a GNN layer with respect to the initial node representation, and can be upper bounded via the graph topology. Stemming from this, graph rewiring methods alleviate over-squashing by adding or removing edges to compute an optimal graph (Topping et al., 2022; Black et al., 2023; Di Giovanni et al., 2023). Moreover, transformers have been applied to consider global aspects (Ying et al., 2021; Kreuzer et al., 2021; Rampásek et al., 2022; Shirzad et al., 2023).

Expressivity of GNNs for the SBM. Certain studies focus on analyzing the expressive power of GNNs using variations of the SBM (Holland et al., 1983). Fountoulakis et al. (2023) established conditions for the existence of graph attention networks (GATs) that can precisely classify nodes in the contextual stochastic block model (CSBM) with high probability. Similarly, Baranwal et al. (2023) investigated the effects of graph convolutions within a network on the XOR-CSBM. The preceding works primarily focused on the probability distribution of node features, such as the distance between the means of feature vectors. On the other hand, Kanatsoulis & Ribeiro (2023) conducted analysis on the expressivity utilizing linear algebraic tools and eigenvalue decomposition of graph operators.

3 Non-backtracking Graph Neural Network

3.1 Motivation from Sensitivity Analysis

We first explain how the conventional message-passing updates are prone to backtracking. To be specific, consider a simple, undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{N}(i)$ denotes the neighbor of the node i . Each node $i \in \mathcal{V}$ is associated with a feature x_i . Then the conventional graph neural networks (GNNs), i.e., message-passing neural networks (MPNNs) (Gilmer et al., 2017), iteratively update the t -th layer node-wise hidden feature $h_i^{(t)}$ as follows:

$$h_i^{(t+1)} = \phi^{(t)}\left(h_i^{(t)}, \left\{\psi^{(t)}\left(h_i^{(t)}, h_j^{(t)}\right) : j \in \mathcal{N}(i)\right\}\right), \quad (1)$$

where $\phi^{(t)}$ and $\psi^{(t)}$ are architecture-specific non-linear update and permutation invariant aggregation functions, respectively. Our key observation is that the message from a node feature $h_i^{(t)}$ to the node feature $h_j^{(t+1)}$ is reincorporated in the node feature $h_i^{(t+2)}$, e.g., Figure 3a shows the computation graph of conventional GNNs with redundant messages.

Sensitivity Analysis. To concretely describe the consequences of backtracking in message-passing updates, we formulate the sensitivity of the final node feature $h_i^{(T)}$ with respect to the input as follows:

$$\sum_{j \in \mathcal{V}} \frac{\partial h_i^{(T)}}{\partial h_j^{(0)}} = \sum_{s \in \mathcal{W}(i)} \prod_{t=1}^T \frac{\partial h_{s(t)}^{(t)}}{\partial h_{s(t-1)}^{(t-1)}}, \quad (2)$$

where $h_i^{(0)} = x_i$, $\mathcal{W}(i)$ denotes the set of T -step walks ending at node i , and $s(t)$ denotes the t -th node in the walk $s \in \mathcal{W}(i)$. Intuitively, this equation shows that a GNN with T layer recognizes the graph via aggregation of random walks with length T . Our key observation from Equation (2) is on how the feature $h_i^{(T)}$ is insensitive to the information to an initial node feature $h_j^{(0)}$, due to the information being “squashed” by the aggregation over the exponential number of walks $\mathcal{W}(i)$. A similar analysis has been conducted on how a node feature $h_i^{(T)}$ is insensitive to the far-away initial node feature $h_j^{(0)} = x_j$, i.e., the over-squashing phenomenon of GNNs (Topping et al., 2022).

Redundancy of walks with backtracking.

In particular, a walk s randomly sampled from $\mathcal{W}(i)$ is likely to contain a transition that backtracks, i.e., $s(t) = s(t+2)$ for some $t < T$. Then the walk s would be redundant since the information is contained in two other walks in $\mathcal{W}(i)$: $s(0), \dots, s(t+1)$ and $s(0), \dots, s(t)s(t+1), s(t+2) = s(t), s(t+3), \dots, s(T)$, as illustrated in Figure 2. This idea leads to the conclusion that non-backtracking walks, i.e., walks that do not contain backtracking transitions, are sufficient to express the information in the walks $\mathcal{W}(i)$. Since the exponential number of walks in $\mathcal{W}(i)$ causes the GNN to be insensitive to a particular walk information, it makes sense to design a non-backtracking GNN that is sensitive to the constrained set of non-backtracking walks.




Figure 2: Two non-backtracking walks (right) are sufficient to express information contained in a walk with backtracking transition.

Relation to Over-squashing. Finally, we point out an intriguing motivation for our work in terms of over-squashing. In particular, we note that Di Giovanni et al. (2023) analyzed the lower bound for the Jacobian obstruction that measures the degree of over-squashing in terms of access time with respect to a simple random walk. They conclude that the degree of over-squashing, i.e., the size of Jacobian obstruction, is higher for a pair of nodes with longer access time. Hence, to design a GNN architecture robust to over-squashing, one could (i) propose a random walk that has shorter access time for a pair of nodes in the graph, and (ii) design a GNN that aligns with the random walk. Non-backtracking random walks have been empirically shown and believed to generally yield faster access time than simple random walks (Lin & Zhang, 2019; Fasino et al., 2023). Therefore, one could aim to design a GNN architecture that aligns with the non-backtracking random walks.

Access Time of Random Walks. However, to the best of our knowledge, there is no formal proof of scenarios where non-backtracking random walks yield a shorter access time. As a motivating example, we provide a theoretical result comparing the access time of non-backtracking and simple random walks for tree graphs. Since non-backtracking random walks do not guarantee a walk of a certain length, we make use of begrudgingly backtracking random walks (Rappaport et al., 2017), which modifies non-backtracking random walks to remove “dead ends” for tree graphs. For the full proof, please refer to Appendix B.

Proposition 1. *Given a tree $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a pair of nodes $i, j \in \mathcal{V}$, the access time of begrudgingly backtracking random walk is equal or smaller than that of a simple random walk. The equality holds if and only if the walk length is 1.*




Figure 3: Computation graph of typical GNN and NBA-GNN predicting node “0”. (a) Redundant messages increase the size of the computation graph, proportional to the number of layers. (b) NBA-GNN assigns a pair of features for each edge and update them via non-backtracking message passing. By reducing redundant messages, it results a simplified computation graph compared to typical GNNs.

3.2 Method Description

In this section, we present the **Non-BA**cktracking GNNs (NBA-GNN) with the motivation described in Section 3.1. Given an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, our NBA-GNN associates a pair of hidden features $h_{i \rightarrow j}, h_{j \rightarrow i}$ for each edge $\{i, j\}$. Then the non-backtracking message passing update for a hidden feature $h_{j \rightarrow i}^{(t)}$ is defined as follows:

$$h_{j \rightarrow i}^{(t+1)} = \phi^{(t)} \left(h_{j \rightarrow i}^{(t)}, \left\{ \psi^{(t)} \left(h_{k \rightarrow j}^{(t)}, h_{j \rightarrow i}^{(t)} \right) : k \in \mathcal{N}(j) \setminus \{i\} \right\} \right), \quad (3)$$

where $\phi^{(t)}$ and $\psi^{(t)}$ are backbone-specific non-linear update and permutation-invariant aggregation functions at the t -th layer, respectively. For example, $\psi^{(t)}$ and $\phi^{(t)}$ are multi-layer perceptron and summation over a set for the graph isomorphism network (Xu et al., 2019), respectively. Given the update in Equation (3), one can observe that the message $h_{i \rightarrow j}^{(t)}$ is never incorporated in the message $h_{j \rightarrow i}^{(t+1)}$, and hence the update is free from backtracking. Note that line graph neural networks (Chen et al., 2018, LGNN) also applied non-backtracking operators in GNNs. However, it does not address the issue of redundant messages, as it continues to use the adjacency matrix. Our key contribution is the use of the non-backtracking operator specifically to tackle the message redundancy issue. For a detailed comparison, please refer to Appendix A.

Initialization and Node-wise Aggregation of Messages. The message at the 0-th layer $h_{i \rightarrow j}^{(0)}$ is initialized by encoding the node features x_i, x_j , and the edge feature e_{ij} using a non-linear function ϕ . After updating hidden features for each edge based on Equation (3), we apply a permutation-invariant pooling over all the messages for graph-wise predictions. Since we use hidden features for each edge, we construct the node-wise predictions at the final layer as follows:

$$h_i = \sigma \left(\rho \left\{ h_{j \rightarrow i}^{(T)} : j \in \mathcal{N}(i) \right\}, \rho \left\{ h_{i \rightarrow j}^{(T)} : j \in \mathcal{N}(i) \right\} \right), \quad (4)$$

where σ is a non-linear aggregation function with different weights for incoming edges $j \rightarrow i$ and outgoing edges $i \rightarrow j$, ρ is a non-linear aggregation function invariant to the permutation of nodes in $\mathcal{N}(i)$. We provide a computation graph of NBA-GNNs in Figure 3b to summarize our algorithm.

Begrudgingly Backtracking Update. While the messages from our update are resistant to backtracking, a message $h_{j \rightarrow i}^{(t)}$ may get trapped in node i for the special case when $\mathcal{N}(i) = \{j\}$. To resolve this issue, we introduce a simple trick coined begrudgingly backtracking update (Rappaport et al., 2017) that updates $h_{i \rightarrow j}^{(t+1)}$ using $h_{j \rightarrow i}^{(t)}$ only when $\mathcal{N}(i) = \{j\}$. We empirically verify the effectiveness of begrudgingly backtracking updates in Section 5.3.

Implementation. To better understand our NBA-GNN, we provide an example of non-backtracking message-passing updates with a GCN backbone (Kipf & Welling, 2017), coined NBA-GCN. The message

update at the t -th layer of NBA-GCN can be written as follows:

$$h_{j \rightarrow i}^{(t+1)} = h_{j \rightarrow i}^{(t)} + \sigma_{\text{GCN}} \left(\frac{1}{|\mathcal{N}(j)| - 1} \mathbf{W}^{(t)} \sum_{k \in \mathcal{N}(j) \setminus \{i\}} h_{k \rightarrow j}^{(t)} \right), \quad (5)$$

where σ_{GCN} is an element-wise nonlinear function, e.g., rectified linear unit (Agarap, 2018), $\mathbf{W}^{(t)}$ is the weight matrix, and messages are normalized by their population $|\mathcal{N}(j)| - 1$.

4 Theoretical Analysis

In this section, we analyze the proposed NBA-GNN framework. To be specific, we show that (a) our NBA-GNN improves the upper bound for sensitivity-based measures of GNN over-squashing and (b) NBA-GNNs can detect the underlying structure of SBMs even for very sparse graphs.

4.1 Sensitivity Analysis on Over-squashing

While Chen et al. (2018) initially introduced the non-backtracking operator in GNNs, they did not explore its theoretical implications. Hence, we first analyze how NBA-GNNs alleviate the over-squashing issue. To this end, following Topping et al. (2022), we use the Jacobian of node-wise output with respect to another node initial feature to assess over-squashing effect. We first derive the Jacobian quantification of over-squashing for the NBA-GNNs, then demonstrate that the upper bound for Jacobian is larger than that of the conventional GNNs derived by Topping et al. (2022). To be specific, we interpret the *sensitivity bound* in terms of the *random walk* that a GNN aligns with. We then prove that the non-backtracking random walk yields a higher sensitivity than simple random walks. To the best of our knowledge, this is the first attempt to compare the degree of over-squashing between GNNs aligning with random walks.

To derive our analysis, we bring the notation of non-backtracking matrix $B \in \{0, 1\}^{2|\mathcal{E}| \times 2|\mathcal{E}|}$ and the incidence matrix $C \in \mathbb{R}^{2|\mathcal{E}| \times |\mathcal{V}|}$ following Chen et al. (2018) to describe the message-passing of NBA-GNN and node-wise aggregation via linear operation, respectively. To be specific, the backtracking matrix B and the incidence matrix C are defined as follows:

$$B_{(\ell \rightarrow k), (j \rightarrow i)} = \begin{cases} 1 & \text{if } k = j, \ell \neq i \\ 0 & \text{otherwise} \end{cases}, \quad C_{(k \rightarrow j), i} = \begin{cases} 1 & \text{if } j = i \text{ or } k = i \\ 0 & \text{otherwise} \end{cases}.$$

We also define D_{out} and D_{in} as the matrices representing out-degree and in-degree of NBA-GNN, respectively, capturing the count of outgoing and incoming edges for each edge. These are diagonal matrices with $(D_{\text{out}})_{(j \rightarrow i), (j \rightarrow i)} = \sum_{\ell \rightarrow k} B_{(j \rightarrow i), (\ell \rightarrow k)}$ and $(D_{\text{in}})_{(j \rightarrow i), (j \rightarrow i)} = \sum_{\ell \rightarrow k} B_{(\ell \rightarrow k), (j \rightarrow i)}$. Next, we introduce \hat{B} as the normalized non-backtracking matrix augmented with self-loops: $\hat{B} = (D_{\text{out}} + I)^{-\frac{1}{2}} (B + I) (D_{\text{in}} + I)^{-\frac{1}{2}}$. Finally, we let \tilde{C} as the matrix where $\tilde{C}_{(k \rightarrow j), i} = C_{(k \rightarrow j), i} + C_{(j \rightarrow k), i}$. Then, one obtains the following sensitivity bound of NBA-GNN.

Lemma 1. *Consider two nodes $i, j \in \mathcal{V}$ with a random walk distance T given a $(T - 1)$ -layer NBA-GNN as described in Equation (3) and Equation (4). Suppose $|\nabla \phi^{(t)}|, |\nabla \sigma| \leq \alpha$, $|\nabla \psi^{(t)}|, |\nabla \rho| \leq \beta$, and $|\nabla \phi| \leq \gamma$ for $0 \leq t < T$. Then the following holds:*

$$\left\| \frac{\partial h_j}{\partial x_i} \right\| \leq (\alpha\beta)^T \gamma (\tilde{C}^\top \hat{B}^{(T-1)} \tilde{C})_{j,i}.$$

We provide the proof in Appendix C. Lemma 1 states how the over-squashing effect is controlled by the power of \hat{B} . Consequently, one can infer that increasing the upper bound likely leads to alleviation of the GNN over-squashing effect (Topping et al., 2022).

From this motivation, we provide an analysis to support our claim on how the NBA-GNN suffers less from over-squashing effect due to its larger sensitivity bound. The key point is that the *sensitivity bounds aligns with a random walk*, and the non-backtracking random walks results a larger sensitivity bound.

Table 1: Comparison of conventional MPNNs and GNNs in the long-range graph benchmark, with and without Laplacian positional encoding (LapPE). We also denote the relative improvement by Impr.

Model	Peptides-func		Peptides-struct		PascalVOC-SP	
	AP \uparrow	Impr.	MAE \downarrow	Impr.	F1 \uparrow	Impr.
GCN	0.5930 \pm 0.0023		0.3496 \pm 0.0013		0.1268 \pm 0.0060	
+ NBA	0.6951 \pm 0.0024	+17%	0.2656 \pm 0.0009	+22%	0.2537 \pm 0.0054	+100%
+ NBA+LapPE	0.7206 \pm 0.0028	+22%	0.2472 \pm 0.0008	+29%	0.3005 \pm 0.0010	+137%
GIN	0.5498 \pm 0.0079		0.3547 \pm 0.0045		0.1265 \pm 0.0076	
+ NBA	0.6961 \pm 0.0045	+27%	0.2534 \pm 0.0025	+29%	0.3040 \pm 0.0119	+140%
+ NBA+LapPE	0.7071 \pm 0.0067	+29%	0.2424 \pm 0.0010	+32%	0.3223 \pm 0.0010	+155%
GatedGCN	0.5864 \pm 0.0077		0.3420 \pm 0.0013		0.2873 \pm 0.0219	
+ NBA	0.6429 \pm 0.0062	+10%	0.2539 \pm 0.0011	+26%	0.3910 \pm 0.0010	+36%
+ NBA+LapPE	0.6982 \pm 0.0014	+19%	0.2466 \pm 0.0012	+28%	0.3969 \pm 0.0027	+38%

Theorem 1. Let \hat{A} denote the degree-normalized matrix. Then, for any pair of nodes $i, j \in \mathcal{V}$ with distance T , the sensitivity bound of NBA-GNN is larger than that of conventional GNNs (Topping et al., 2022), i.e., $(\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i} \geq (\hat{A}^T)_{j,i}$. For d -regular graphs, $(\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i}$ decays slower by $O(d^{-T})$, while $(\hat{A}^T)_{j,i}$ decays with $O((d+1)^{-T})$.

Note that $(\alpha\beta)^T(\hat{A}^T)_{j,i}$ provides an upper bound for the sensitivity term in conventional GNNs (Topping et al., 2022). We provide the full proof in Appendix C, based on comparing the degree-normalized number of non-backtracking and simple walks from node i to node j .

4.2 Expressive Power of NBA-GNN on SBMs

In the literature on the expressive capabilities of GNNs, comparisons with the well-known k -WL test are common. However, even when certain models surpass the k -WL test in performance, evaluating their relative merits remains a nontrivial task. Furthermore, due to the substantial performance gap between the 1-WL and 3-WL tests, many algorithms fall into a range between these two tests, making it more difficult to compare them with each other. It is also worth noting that comparing GNNs with the WL test does not always accurately reflect their performance on real-world datasets.

To address these issues, several studies have turned to spectral analysis of GNNs. From a spectral viewpoint, GNNs can be seen as functions of the eigenvectors and eigenvalues of the given graph. NT & Maehara (2019) showed that GNNs operate as low-pass filters on the graph spectrum, and Balcilar et al. (2020) analyzed the use of various GNNs as filters to extract the relevant graph spectrum and measure their expressive power. Moreover, Oono & Suzuki (2020) argue that the expressive power of GNNs is influenced by the topological information contained in the graph spectrum.

The eigenvalues and the corresponding adjacency matrix eigenvectors play a pivotal role in establishing the fundamental limits of community detection in SBM, as evidenced by Yun & Proutière (2019). The adjacency matrix exhibits a spectral separation property, and an eigenvector containing information about the assignments of the vertex community becomes apparent (Lei & Rinaldo, 2015). Furthermore, by analyzing the eigenvalues of the adjacency matrix, it is feasible to determine whether a graph originates from the Erdős-Rényi (ER) model or the SBM (Erdős et al., 2013; Avrachenkov et al., 2015). However, these spectral properties are particularly salient when the average degree of the graph satisfies $\Omega(\log n)$. For graphs with average degrees $o(\log n)$, vertices with higher degrees predominate, affecting eigenvalues and complicating the discovery of the underlying structure of the graph (Benaych-Georges et al., 2019).

In contrast, the non-backtracking matrix exhibits several advantageous properties even for constant-degree cases. In (Stephan & Massoulié, 2022), the non-backtracking matrix demonstrates a spectral separation

property and establishes the presence of an eigenvector containing information about vertex community assignments, when the average degree only satisfies $\omega(1)$ and $n^{o(1)}$. Furthermore, Bordenave et al. (2015) have demonstrated that by inspecting the eigenvalues of the non-backtracking matrix, it is possible to discern whether a graph originates from the ER model or the SBM, even when the graph’s average degree remains constant. This capability sets NBA-GNNs apart and enhances their performance in both node and graph classification tasks, especially in sparse settings. These lines of reasoning lead to the formulation of the following theorem.

Theorem 2. (Informal) *Assume the average degree in the stochastic block model satisfies the conditions of being at least $\omega(1)$ and $n^{o(1)}$. In such a scenario, the NBA-GNN can map from graph \mathcal{G} to node labels.*

Theorem 3. (Informal) *Suppose we have a pair of graphs with a constant average degree, one generated from the stochastic block model and the other from the Erdős–Rényi model. In this scenario, the NBA-GNN is capable of distinguishing between them.*

This theorem integrates prior work on the non-backtracking matrix into GNNs, contributing to a deeper understanding of the expressive power within GNN structures. Such an approach has the potential to advance the field by introducing new perspectives and methodologies within the GNN framework. For an in-depth exploration of this argument, please refer to Appendix D. The rationale behind these valuable properties of the non-backtracking matrix B in sparse scenarios lies in the fact that the matrix B^k exhibits similarity to the k -hop adjacency matrix, while A^k is mainly influenced by high-degree vertices. For these reasons, NBA-GNNs would outperform traditional GNNs in both node and graph classification tasks, particularly in sparse graph environments.

5 Experiment

In this section, we assess the effectiveness of NBA-GNNs across multiple benchmarks on graph classification, graph regression, and node classification tasks. Additionally, we conduct a detailed comparison, including complexity analysis, with existing methods aiming to reduce redundancy (see Appendix A). Detailed experimental information is also provided in Appendix E.

5.1 Long-range Graph Benchmark

The long-range graph benchmark (LRGB) (Dwivedi et al., 2022) is a set of tasks that require learning long-range interactions. We validate our method using three datasets from the LRGB benchmark: graph classification (Peptides-func), graph regression (Peptides-struct), and node classification (PascalVOC-SP). We use performance scores from (Dwivedi et al., 2022) and from each baseline papers: subgraph based GNNs (Abu-El-Haija et al., 2019; Michel et al., 2023; Giusti et al., 2023), graph Transformers (Kreuzer et al., 2021; Rampásek et al., 2022; Shirzad et al., 2023; He et al., 2023), and graph rewiring methods (Gasteiger et al., 2019; Gutteridge et al., 2023). For NBA-GNNs and NBA-GNNs with begrudgingly backtracking, we report the one with better performance. Furthermore, LapPE, i.e., Laplacian positional encoding (Dwivedi et al., 2023), is applied, as it enhances the performance of NBA-GNNs in common cases.

As one can see in Table 1, NBA-GNNs show improvement regardless of the combined backbone GNNs, i.e., GCN (Kipf & Welling, 2017), GIN (Xu et al., 2019), and GatedGCN (Bresson & Laurent, 2018). Furthermore, even against a variety of baselines in Table 2, at least one NBA-GNNs shows competitive performance with the best baseline in LRGB. Also, it is noteworthy that the improvement of NBA-GNNs is higher in dense graphs, PascalVOC-SP having an average degree of 8 while Peptides-func, Peptides-struct has an average degree of 2.

5.2 Transductive Node Classification Tasks

We conduct experiments on three citation networks (Cora, CiteSeer, and Pubmed) (Sen et al., 2008), and three heterophilic datasets (Texas, Wisconsin, and Cornell) (Pei et al., 2019) for transductive node classification. In our experiments, we employed various GNN architectures (Kipf & Welling, 2017; Hamilton et al., 2017; Veličković et al., 2018), along with their corresponding NBA versions, as illustrated in Table 3.

Table 2: Evaluation of NBA-GNN on the LRGB benchmark. We color the **first-**, **second-** and **third-** best results. Performance within a standard deviation of one another is considered equal. Non-reported values are denoted by -.

Method	Model	Peptides-func AP \uparrow	Peptides-struct MAE \downarrow	PascalVOC-SP F1 \uparrow
GNNs	GCN	0.5930 \pm 0.0023	0.3496 \pm 0.0013	0.1268 \pm 0.0060
	GCN+LapPE	0.6213 \pm 0.0060	0.3250 \pm 0.0180	0.1370 \pm 0.0077
	GIN	0.5498 \pm 0.0079	0.3547 \pm 0.0045	0.1265 \pm 0.0076
	GIN+LapPE	0.5877 \pm 0.0044	0.3369 \pm 0.0026	0.1302 \pm 0.0105
	GatedGCN	0.5864 \pm 0.0077	0.3420 \pm 0.0013	0.2873 \pm 0.0219
	GatedGCN+LapPE	0.6069 \pm 0.0035	0.3357 \pm 0.0006	0.2860 \pm 0.0085
Subgraph GNNs	MixHop-GCN	0.6592 \pm 0.0036	0.2921 \pm 0.0023	0.2506 \pm 0.0133
	MixHop-GCN+LapPE	0.6843 \pm 0.0049	0.2614 \pm 0.0023	0.2218 \pm 0.0174
	PathNN	0.6816 \pm 0.0026	0.2545 \pm 0.0032	-
	CIN++	0.6569 \pm 0.0117	0.2523 \pm 0.0013	-
Transformers	Transformer+LapPE	0.6326 \pm 0.0126	0.2529 \pm 0.0016	0.2694 \pm 0.0098
	GraphGPS+LapPE	0.6535 \pm 0.0041	0.2500 \pm 0.0005	0.3748 \pm 0.0109
	SAN+LapPE	0.6384 \pm 0.0121	0.2683 \pm 0.0043	0.3230 \pm 0.0039
	Expformer	0.6527 \pm 0.0043	0.2481 \pm 0.0007	0.3966 \pm 0.0027
	Graph MLP-Mixer/ViT	0.6970 \pm 0.0080	0.2449 \pm 0.0016	-
Graph Rewiring	DIGL+MPNN	0.6469 \pm 0.0019	0.3173 \pm 0.0007	0.2824 \pm 0.0039
	DIGL+MPNN+LapPE	0.6830 \pm 0.0026	0.2616 \pm 0.0018	0.2921 \pm 0.0038
	DRew-GCN+LapPE	0.7150 \pm 0.0044	0.2536 \pm 0.0015	0.1851 \pm 0.0092
	DRew-GIN+LapPE	0.7126 \pm 0.0045	0.2606 \pm 0.0014	0.2692 \pm 0.0059
	DRew-GatedGCN+LapPE	0.6977 \pm 0.0026	0.2539 \pm 0.0007	0.3314 \pm 0.0024
NBA-GNNs (Ours)	NBA-GCN	0.6951 \pm 0.0024	0.2656 \pm 0.0009	0.2537 \pm 0.0054
	NBA-GCN+LapPE	0.7207 \pm 0.0028	0.2472 \pm 0.0008	0.3005 \pm 0.0010
	NBA-GIN	0.6961 \pm 0.0045	0.2775 \pm 0.0057	0.3040 \pm 0.0119
	NBA-GIN+LapPE	0.7071 \pm 0.0067	0.2424 \pm 0.0010	0.3223 \pm 0.0063
	NBA-GatedGCN	0.6429 \pm 0.0062	0.2539 \pm 0.0011	0.3910 \pm 0.0010
	NBA-GatedGCN+LapPE	0.6982 \pm 0.0014	0.2466 \pm 0.0012	0.3969 \pm 0.0027

Table 3: Comparison of conventional GNNs and their NBA-GNN counterpart on transductive node classification tasks for citation networks and heterophilic datasets, with and without Laplacian positional encoding(LapPE). We mark the best numbers in bold.

Model	Cora	CiteSeer	PubMed	Texas	Wisconsin	Cornell
GCN	0.8658 \pm 0.0060	0.7532 \pm 0.0134	0.8825 \pm 0.0042	0.6162 \pm 0.0634	0.6059 \pm 0.0438	0.5946 \pm 0.0662
+ LapPE	0.8592 \pm 0.0083	0.7572 \pm 0.0132	0.8817 \pm 0.0040	0.6216 \pm 0.0584	0.6000 \pm 0.0600	0.5703 \pm 0.0547
+ NBA	0.8722\pm0.0095	0.7585 \pm 0.0175	0.8826 \pm 0.0044	0.7108\pm0.0796	0.7471\pm0.0386	0.6108 \pm 0.0614
+ NBA+LapPE	0.8720 \pm 0.0129	0.7609\pm0.0186	0.8827\pm0.0048	0.6811 \pm 0.0595	0.7471\pm0.0466	0.6378\pm0.0317
GraphSAGE	0.8632 \pm 0.0158	0.7559 \pm 0.0161	0.8864 \pm 0.0030	0.7108 \pm 0.0556	0.7706 \pm 0.0403	0.6027 \pm 0.0625
+ LapPE	0.8700 \pm 0.0117	0.7608 \pm 0.0144	0.8895\pm0.0047	0.7162 \pm 0.0653	0.7647 \pm 0.0453	0.6189 \pm 0.0484
+ NBA	0.8702\pm0.0083	0.7586 \pm 0.0213	0.8871 \pm 0.0044	0.7270 \pm 0.0905	0.7765\pm0.0508	0.6459\pm0.0691
+ NBA+LapPE	0.8650 \pm 0.0120	0.7621\pm0.0172	0.8870 \pm 0.0037	0.7486\pm0.0612	0.7647 \pm 0.0531	0.6378 \pm 0.0544
GAT	0.8694 \pm 0.0119	0.7463 \pm 0.0159	0.8787 \pm 0.0046	0.6054 \pm 0.0386	0.6000 \pm 0.0491	0.4757 \pm 0.0614
+ LapPE	0.8686 \pm 0.0152	0.7512 \pm 0.0154	0.8775 \pm 0.0045	0.6135 \pm 0.0404	0.6294 \pm 0.0448	0.5108 \pm 0.0769
+ NBA	0.8722\pm0.0120	0.7549 \pm 0.0171	0.8829\pm0.0043	0.6622 \pm 0.0514	0.7059 \pm 0.0562	0.5838\pm0.0558
+ NBA+LapPE	0.8692 \pm 0.0098	0.7561\pm0.0175	0.8822 \pm 0.0047	0.6730\pm0.0348	0.7314\pm0.0531	0.5784 \pm 0.0640

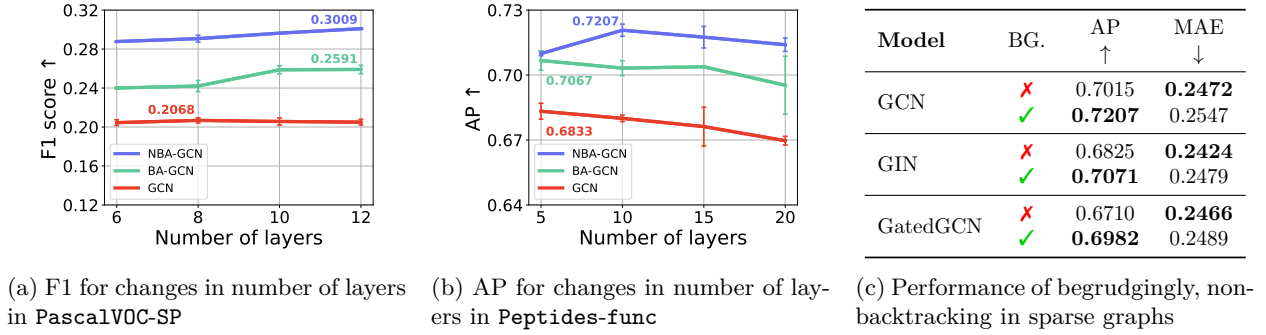


Figure 4: Ablation studies on the components of NBA-GNN.

The results indicate that the non-backtracking improves the performance of all GNN variants. Notably, it demonstrates significant enhancements in heterophilic datasets. Given that related nodes in heterophilic graphs are often widely separated (Zheng et al., 2022), the ability of NBA-GNN to alleviate over-squashing plays a vital role in classifying nodes in such scenarios.

5.3 Ablation Studies

In this section, we conduct ablation studies to empirically verify our framework. For simplicity, we use BA for backtracking GNNs and BG for begrudgingly backtracking GNNs. All experiments are averaged over 3 seeds, and hyper-parameters for GCN from Tönshoff et al. (2023) are used for a fair comparison.

Non-backtracking vs. Backtracking. We first verify whether the performance improvements indeed stem from the non-backtracking updates. To this end, we compare our NBA-GNN with a backtracking variant, coined BA-GNN. To be clear, BA-GNN allows backtracking update prohibited in NBA-GNN, i.e., using $h_{i \rightarrow j}^{(\ell)}$ to update $h_{j \rightarrow i}^{(\ell+1)}$. From Figures 4a and 4b, one can see how NBA-GNN outperforms the BA-GNN consistently irrelevant to the numbers of layers. Intriguingly, one can also observe that BA-GNN outperforms the naïve backbone, i.e., GCN, consistently.

Begrudgingly Backtracking Updates in Sparse Graphs. Additionally, we investigate the effect of begrudgingly backtracking in sparse graphs, i.e., Peptides-func, and Peptides-struct, in Figure 4c. One can see that begrudgingly backtracking is effective in Peptides-func, and shows similar performance in Peptides-struct (Note that the PascalVOC-SP does not have a vertex with degree one).

6 Conclusion

We have introduced a message passing framework applicable to any GNN architectures, to alleviated over-squashing in GNNs. As theoretically shown, NBA-GNNs mitigate over-squashing in terms of sensitivity and enhance its expressive power for both node and graph classification tasks on SBMs. Additionally, we demonstrated that NBA-GNNs achieve competitive performance on the LRGB benchmark, and outperform conventional GNNs across transductive node classification tasks, even in heterophilic datasets.

References

- Emmanuel Abbe. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research*, 18(1):6446–6531, 2017.
- Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *international conference on machine learning*, pp. 21–29. PMLR, 2019.
- Abien Fred Agarap. Deep learning using rectified linear units (relu). *arXiv preprint arXiv:1803.08375*, 2018.

- Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9(04):585–603, 2007.
- Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. In *International Conference on Learning Representations*, 2021.
- Konstantin Avrachenkov, Laura Cottatellucci, and Arun Kadavankandy. Spectral properties of random matrices for stochastic block model. In *2015 13th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, pp. 537–544. IEEE, 2015.
- Furqan Aziz, Richard C Wilson, and Edwin R Hancock. Backtrackless walks on a graph. *IEEE transactions on neural networks and learning systems*, 24(6):977–989, 2013.
- Muhammet Balcilar, Guillaume Renton, Pierre Héroux, Benoit Gaüzère, Sébastien Adam, and Paul Honeine. Analyzing the expressive power of graph neural networks in a spectral perspective. In *International Conference on Learning Representations*, 2020.
- Aseem Baranwal, Kimon Fountoulakis, and Aukosh Jagannath. Effects of graph convolutions in multi-layer networks. In *The Eleventh International Conference on Learning Representations*, 2023.
- Florent Benaych-Georges, Charles Bordenave, and Antti Knowles. Largest eigenvalues of sparse inhomogeneous erdős-rényi graphs. *The Annals of Probability*, 47(3):1653–1676, 2019.
- Mitchell Black, Zhengchao Wan, Amir Nayyeri, and Yusu Wang. Understanding oversquashing in gnns through the lens of effective resistance. In *International Conference on Machine Learning*, pp. 2528–2547. PMLR, 2023.
- Charles Bordenave, Marc Lelarge, and Laurent Massoulié. Non-backtracking spectrum of random graphs: community detection and non-regular ramanujan graphs. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pp. 1347–1357. IEEE, 2015.
- Xavier Bresson and Thomas Laurent. Residual gated graph convnets. In *International Conference on Learning Representations*, 2018.
- Joan Bruna and X Li. Community detection with graph neural networks. *stat*, 1050:27, 2017.
- Rongqin Chen, Shenghui Zhang, Ye Li, et al. Redundancy-free message passing for graph neural networks. *Advances in Neural Information Processing Systems*, 35:4316–4327, 2022.
- Zhengdao Chen, Lisha Li, and Joan Bruna. Supervised community detection with line graph neural networks. In *International Conference on Learning Representations*, 2018.
- Francesco Di Giovanni, Lorenzo Giusti, Federico Barbero, Giulia Luise, Pietro Lio, and Michael M Bronstein. On over-squashing in message passing neural networks: The impact of width, depth, and topology. In *International Conference on Machine Learning*, pp. 7865–7885. PMLR, 2023.
- Vijay Prakash Dwivedi, Ladislav Rampášek, Michael Galkin, Ali Parviz, Guy Wolf, Anh Tuan Luu, and Dominique Beaini. Long range graph benchmark. *Advances in Neural Information Processing Systems*, 35:22326–22340, 2022.
- Vijay Prakash Dwivedi, Chaitanya K Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24(43):1–48, 2023.
- László Erdős, Antti Knowles, Horng-Tzer Yau, and Jun Yin. Spectral statistics of erdős-rényi graphs i: Local semicircle law. *The Annals of Probability*, pp. 2279–2375, 2013.
- Dario Fasino, Arianna Tonetto, and Francesco Tudisco. Hitting times for second-order random walks. *European Journal of Applied Mathematics*, 34(4):642–666, 2023.
- Robert Fitzner and Remco van der Hofstad. Non-backtracking random walk. *Journal of Statistical Physics*, 150:264–284, 2013.

- Kimon Fountoulakis, Amit Levi, Shenghao Yang, Aseem Baranwal, and Aukosh Jagannath. Graph attention retrospective. *Journal of Machine Learning Research*, 24(246):1–52, 2023.
- Johannes Gasteiger, Stefan Weißenberger, and Stephan Günnemann. Diffusion improves graph learning. *Advances in neural information processing systems*, 32, 2019.
- Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pp. 1263–1272. PMLR, 2017.
- Lorenzo Giusti, Teodora Reu, Francesco Ceccarelli, Cristian Bodnar, and Pietro Liò. Cin++: Enhancing topological message passing. *arXiv preprint arXiv:2306.03561*, 2023.
- Benjamin Gutteridge, Xiaowen Dong, Michael M Bronstein, and Francesco Di Giovanni. Drew: Dynamically rewired message passing with delay. In *International Conference on Machine Learning*, pp. 12252–12267. PMLR, 2023.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *NeurIPS*, 30, 2017.
- Xiaoxin He, Bryan Hooi, Thomas Laurent, Adam Perold, Yann LeCun, and Xavier Bresson. A generalization of vit/mlp-mixer to graphs. In *International Conference on Machine Learning*, pp. 12724–12745. PMLR, 2023.
- Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.
- Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020.
- Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-lsc: A large-scale challenge for machine learning on graphs. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Charilaos I. Kanatsoulis and Alejandro Ribeiro. Representation power of graph neural networks: Improved expressivity via algebraic analysis, 2023.
- Mark Condie Kempton. Non-backtracking random walks and a weighted ihara’s theorem. *Open Journal of Discrete Mathematics*, 2016.
- JinHyung Kim and Judea Pearl. A computational model for causal and diagnostic reasoning in inference systems. In *International Joint Conference on Artificial Intelligence*, pp. 0–0, 1983.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- Devin Kreuzer, Dominique Beaini, Will Hamilton, Vincent Létourneau, and Prudencio Tossou. Rethinking graph transformers with spectral attention. *Advances in Neural Information Processing Systems*, 34: 21618–21629, 2021.
- Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proceedings of the National Academy of Sciences*, 110(52):20935–20940, 2013.
- Jing Lei and Alessandro Rinaldo. Consistency of spectral clustering in stochastic block models. *The Annals of Statistics*, pp. 215–237, 2015.
- Yuan Lin and Zhongzhi Zhang. Non-backtracking centrality based random walk on networks. *The Computer Journal*, 62(1):63–80, 2019.

- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- Pierre Mahé, Nobuhisa Ueda, Tatsuya Akutsu, Jean-Luc Perret, and Jean-Philippe Vert. Extensions of marginalized graph kernels. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 70, 2004.
- Gaspard Michel, Giannis Nikolentzos, Johannes F Lutzeyer, and Michalis Vazirgiannis. Path neural networks: Expressive and accurate graph neural networks. In *International Conference on Machine Learning*, pp. 24737–24755. PMLR, 2023.
- Kevin P Murphy, Yair Weiss, and Michael I Jordan. Loopy belief propagation for approximate inference: an empirical study. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pp. 467–475, 1999.
- MEJ Newman. Spectral community detection in sparse networks. *arXiv preprint arXiv:1308.6494*, 2013.
- Hoang NT and Takanori Maehara. Revisiting graph neural networks: All we have is low-pass filters. *arXiv preprint arXiv:1905.09550*, 2019.
- Kenta Oono and Taiji Suzuki. Graph neural networks exponentially lose expressive power for node classification. In *International Conference on Learning Representations*, 2020.
- Judea Pearl. Reverend bayes on inference engines: A distributed hierarchical approach. In *Proceedings of the AAAI conference on artificial intelligence*, pp. 133–136, 1982.
- Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. In *International Conference on Learning Representations*, 2019.
- Ladislav Rampášek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. Recipe for a general, powerful, scalable graph transformer. *Advances in Neural Information Processing Systems*, 35:14501–14515, 2022.
- Brian Rappaport, Anuththari Gamage, and Shuchin Aeron. Faster clustering via non-backtracking random walks. *stat*, 1050:26, 2017.
- Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- Hamed Shirzad, Ameya Velingker, Balaji Venkatachalam, Danica J Sutherland, and Ali Kemal Sinop. Exphormer: Sparse transformers for graphs. In *International Conference on Machine Learning*, 2023.
- Ludovic Stephan and Laurent Massoulié. Non-backtracking spectra of weighted inhomogeneous random graphs. *Mathematical Statistics and Learning*, 5(3):201–271, 2022.
- Jan Tönshoff, Martin Ritzert, Eran Rosenbluth, and Martin Grohe. Where did the gap go? reassessing the long-range graph benchmark. In *The Second Learning on Graphs Conference*, 2023.
- Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature. In *International Conference on Learning Representations*, 2022.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*, pp. 40–48. PMLR, 2016.

- Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? *Advances in Neural Information Processing Systems*, 34:28877–28888, 2021.
- Se-Young Yun and Alexandre Proutière. Optimal cluster recovery in the labeled stochastic block model. *Advances in Neural Information Processing Systems*, 29, 2016.
- Se-Young Yun and Alexandre Proutière. Optimal sampling and clustering in the stochastic block model. *Advances in Neural Information Processing Systems*, 32, 2019.
- Xin Zheng, Yixin Liu, Shirui Pan, Miao Zhang, Di Jin, and Philip S Yu. Graph neural networks for graphs with heterophily: A survey. *arXiv preprint arXiv:2202.07082*, 2022.

Appendix

Table of Contents

A Comparison with Related Works	16
A.1 Analysis on Space and Time Complexity	16
A.2 Line Graph Neural Networks	17
A.3 Redundancy-free Graph Neural Network	17
A.4 Stochastic Block Models	18
B Proofs for Section 3.1	19
B.1 Preliminaries	19
B.2 Access Time of Simple Random Walks	19
B.3 Access Time of Begrudgingly Backtracking Random Walks	21
B.4 Proof of Proposition 1	26
C Proofs for Section 4.1	27
C.1 Preliminaries	27
C.2 Proof of Lemma 1	27
C.3 Proof of Proposition 1	29
D Proofs for Section 4.2	30
D.1 Proof of Theorem 2	30
D.2 Proof of Theorem 3	31
D.3 Proof of Lemma 8	31
D.4 Proof of Lemma 9	31
E Experiment Details	35
E.1 Implementation	35
E.2 Long-range Graph Benchmark	36
E.3 Transductive Node Classification	37

A Comparison with Related Works

This section delves into the detailed exploration of key related works, highlighting essential distinctions from our NBA-GNN. In addition, we present simple experiments to prove the superiority of our model.

A.1 Analysis on Space and Time Complexity

A.1.1 Space Complexity

NBA-GNNs construct message for every edge considering directions, and connects them in a non-backtracking matter. This requires $2|\mathcal{E}|$ number of messages, and $(d_{avg} - 1)|\mathcal{E}|$ connections among messages, where d_{avg} is the average degree of the graph. This has not been a bottleneck at practice, and can be reduced by adjusting the batch size.

A.1.2 Time Complexity

As we connect the messages in a non-backtracking matter, we must calculate the relation between edges in the pre-processing process. This non-backtracking edge adjacency can be computed in $\mathcal{O}(|\mathcal{E}|^2)$, and $\mathcal{O}(d_{avg}|\mathcal{E}|)$ if the data is provided in the form of adjacency list. However, this is only required once per data, can be pre-computed, independent with the number of layers, and re-used for runs.

Table 4: Average training and test time (s) per epoch of models for **Peptides-func**, **Peptides-struct**, **PascalVOC-SP** dataset. The shortest time for each dataset has been highlighted in bold.

Split	Dataset	PathNN-SP	PathNN-SP+	PathNN-AP	GraphGPS+LapPE	DRew-GCN+LapPE	NBA-GCN+LapPE
Train	Peptides-func	12.94 (64)	18.15 (64)	N/A	N/A	8.64 (128)	6.66 (200)
	Peptides-struct	11.59 (64)	8.75 (64)	N/A	N/A	12.63 (128)	7.01 (128)
	PascalVOC-SP	N/A	N/A	N/A	N/A	62.17 (20)	69.50 (30)
Test	Peptides-func	3.349 (64)	5.156 (64)	5.977 (64)	0.639 (128)	0.933 (128)	0.541 (200)
	Peptides-struct	2.120 (64)	1.417 (64)	1.372 (64)	0.926 (128)	1.053 (128)	0.532 (128)
	PascalVOC-SP	N/A	N/A	N/A	1.307 (32)	6.468 (20)	5.866 (30)

Furthermore, we have measured the average training and test time per epoch for the **Peptides-func** dataset. Again, every experiment has been tested on a single RTX 3090, with the largest batch size that fits the GPU. In Table 4, NBA-GNN shows competitive time compared to other baselines, though it does suffer in graphs with a high average degree. Parenthesis refers to the batch size.

Table 5: Training time and memory usage on the large datasets for the transductive node classification task.

Datasets	Model	Time (s)	# Param. (K)	Memory (MiB)	Model	Time (s)	# Param. (K)	Memory (MiB)
Cora	GCN	12.22	2317	24.01	GCN+NBA	38.95	3051	158.86
	SAGE	10.74	3104	27.01	SAGE+NBA	47.07	3837	161.86
	GAT	17.72	3130	27.21	GAT+NBA	58.20	3864	162.06
CiteSeer	GCN	15.74	3479	60.68	GCN+NBA	52.44	5375	373.36
	SAGE	14.95	4265	63.68	SAGE+NBA	59.94	6161	376.36
	GAT	19.68	4292	63.87	GAT+NBA	66.92	6188	376.56
PubMed	GCN	42.66	1838	47.9	GCN+NBA	242.66	2094	436.9
	SAGE	56.10	2624	50.9	SAGE+NBA	320.92	2880	439.9
	GAT	64.80	2650	51.0	GAT+NBA	416.85	2906	440.0

We also report the training time per 1,000 epochs and memory usage on the large datasets for the transductive node classification task in Table 5. Note that DRew and PathNN fail to train on these datasets using a single GPU due to memory overflow.

A.1.3 Time Complexity for Preprocessing

We have also compared the theoretical & practical time required for preprocessing of PathNN-SP (Michel et al., 2023), DRew (Gutteridge et al., 2023), RFGNN (Chen et al., 2022), and our NBA-GNN for the

Peptides-func dataset. For preprocessing, PathNN requires the single shortest path between nodes, DRew requires the k -hop neighbors information, RFGNN requires k -depth non-redundant tree, and NBA-GNN requires the non-backtracking edge adjacency. We denote $|V|$ as the number of nodes, $|E|$ as the number of edges, b as the branching factor, k as the number of layers, d_{avg} as the average degree of nodes in a graph. Once again, one can easily see that previous works are dependent on the number of layers, but NBA-GNN remains **irrelevant to the number of layers**. Computation time of preprocessing has been measured on a single RTX 3090 (RFGNN codes were not reported by the authors). One can see that from Table 6 that NBA-GNN shows superiority in both theoretical time complexity and actual computation time.

Table 6: Comparison on theoretical and practical time complexity for preprocessing **Peptides-func** dataset.

Model	PathNN-SP	DRew	RFGNN	NBA-GNN
Theoretical time complexity	$ V * b^k$	$ E \sim k * E $	$ V ! / V - k - 1 !$	$d_{avg} * E $
Practical time complexity (s)	666	123	N/A	68

A.2 Line Graph Neural Networks

Line Graph Neural Networks (LGNN) (Chen et al., 2018) were the pioneers in applying the non-backtracking operator to GNNs. However, it does not resolve redundancy issues, as it employs both an adjacency matrix and non-backtracking matrix in a layer. In other words, our main contribution over LGNN is in consideration of the non-backtracking operator specifically for the message redundancy problem, in an end-to-end manner. Notably, the computational complexity of LGNN is acknowledged to be close to $|V| \log(|V|)$ by its authors.

In our experiments on **Peptides-func**, we compare the performance and complexity of LGNN and NBA-GCN+LapPE, as detailed in Table 7. Given that LGNN was initially proposed using only node degrees for node features, we incorporated an additional node feature encoder, similar to our NBA-GNN setup. Specifically, we used a batch size of 64, hidden dimension of 24, and 4 layers for LGNN. Refer to Appendix E.2 for the hyperparameters of NBA-GCN. The drawbacks of LGNN in terms of time and space complexity become evident when compared to NBA-GNN, while requiring space and time for computing both simple random walks and non-backtracking walks, but result in poor performance.

Table 7: Comparison of LGNN and NBA-GCN on **Peptides-func**.

Model	Training AP \uparrow	Test AP \uparrow	Test time per epoch (s) \downarrow	GPU usage (%) \downarrow
LGNN	0.4202	0.3778	87.761	97.10
NBA-GCN+LapPE	0.9724	0.7207	0.541	51.18

A.3 Redundancy-free Graph Neural Network

Redundancy-Free Graph Neural Network (RFGNN) (Chen et al., 2022) shares the common motivation with NBA-GNN, aiming to reduce redundant messages in the computation graph. RFGNN achieves this by constructing a tree for every node, coined Truncated ePath Tree (TPT). TPT ensures that there are no repeated nodes along the simple path from the root to the leaf, except for the root node which can appear twice in the path. This approach differs significantly from NBA-GNN, which eliminates redundancy by identifying non-backtracking edge adjacency.

In terms of complexity, RFGNN is known to have a space and time complexity of $\mathcal{O}\left(\frac{|V|!}{|V-t-1|!}\right)$, where $|V|$ is the number of nodes and t is the number of GNN layers. In contrast, NBA-GNN achieves superior complexity with space complexity $\mathcal{O}(2|E|)$ and time complexity $\mathcal{O}(d_{avg}|E|)$, which remains **irrelevant to the number of layers**. The preprocessing process time complexity, involved in finding non-backtracking edge adjacency, is $\mathcal{O}(|E|^2)$. This scalability allows NBA-GNN to handle larger graph dataset compared to RFGNN. Furthermore, in theoretical aspects, our work distinguishes itself by providing the sensitivity upper bound of non-backtracking updates, rather than comparing the relative inference of paths.

A.4 Stochastic Block Models

We also conduct a comparative analysis for our method on two sparse Stochastic Block Models (SBMs) with distinct parameters for Belief propagation (BP), and LGNN: (a) a binary assortative SBM ($n = 400$, $C = 2$, $p = 20/n$, $q = 10/n$), and (b) a 5-community dissociative SBM ($n = 400$, $C = 5$, $p = 0$, $q = 18/n$). To be specific, n denotes the number of vertices, C denotes the number of classes, and p and q denote edge probabilities within and across communities, respectively. Table 8 illustrates the average test accuracy across 100 graphs. Notably, BP, known for achieving the information-theoretic threshold, exhibited the best performance, consistent with expectations. Additionally, NBA-GNN outperforms LGNN in both scenarios.

Model	Parameters (n, C, p, q)	LGNN	NBA-GNN	BP
(a) Binary assortative SBM	(400, 2, 20/ n , 10/ n)	0.4885	0.4900	0.5303
(b) 5-community dissociative SBM	(400, 5, 0, 18/ n)	0.1821	0.1888	0.2869

Table 8: Average test accuracy of LGNN, NBA-GNN and BP on two sparse SBMs, where n represents the number of vertices, C is the number of classes, and p and q denote edge probabilities within and across communities, respectively.

B Proofs for Section 3.1

In this section, we present the findings discussed in Section 3.1. Similar to Di Giovanni et al. (2023), we concentrate on the relationship between over-squashing and access time of non-backtracking random walks. Our study establishes that the access time using a begrudgingly backtracking random walk (BBRW) is smaller than that of a simple random walk (SRW) between two nodes, in tree graphs. Also, it is noteworthy that the gap between these two access times increases as the length of the walk grows.

We have derived formulas to compare the access times between BBRW and SRW. First, we show that on the tree graph, the access time equals the sum of access times between neighboring nodes. We note that the access time between neighboring nodes, which is the cut-point, can be represented in terms of return time. The formulas for return time in Fasino et al. (2023) and Lemma 6 allow us to derive the formulas for access time between neighboring nodes. Finally, we derive and compare the access time of BBRW and SRW.

B.1 Preliminaries

B.1.1 Simple and Begrudgingly Random Walks

A random walk on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a sequence of \mathcal{V} -valued random variable $x_0, x_1, x_2 \dots$ where x_{n+1} is chosen randomly from neighborhood of x_n . Different types of random walks have different probabilities for selecting neighboring nodes.

Simple random walk (SRW) choose a next node j uniformly from the neighbors of current node i :

$$P(x_{n+1} = j | x_n = i) = \begin{cases} \frac{1}{d_i} & \text{if } (i, j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases}.$$

On the other hand, begrudgingly backtracking random walk (BBRW) tries to avoid the previous node when there is another option:

$$P(x_{n+2} = k | x_{n+1} = j, x_n = i) = \begin{cases} \frac{1}{d_j - 1} & \text{when } (j, k) \in \mathcal{E}, k \neq i, \text{ and } |\mathcal{N}(j) \setminus \{i\}| \geq 1 \\ 1 & \text{when } (j, k) \in \mathcal{E} \text{ and } |\mathcal{N}(j) \setminus \{i\}| = 0 \\ 0 & \text{otherwise} \end{cases}.$$

B.1.2 Access Time

Consider a SRW starting at a node $i \in \mathcal{V}$. Let T_i denote the time when a SRW first arrived at node i , $T_i := \min\{n \geq 0 | x_n = i\}$, and \bar{T}_i be the time when a random walk first arrived at node i after the first step, $\bar{T}_i := \min\{n > 0 | x_n = i\}$. With slight abuse of notation, we define access time $\mathfrak{t}(i, j)$ from i to j , access time $\mathfrak{t}(i \rightarrow j, k)$ from i to k where $x_1 = j$, and return time $\mathfrak{t}(i)$ in a graph \mathcal{G} as follows:

$$\begin{aligned} \mathfrak{t}(i, j) &:= \mathbb{E}[T_j | x_0 = i] \\ \mathfrak{t}(i \rightarrow j, k) &:= \mathbb{E}[T_k | x_1 = j, x_0 = i] \\ \mathfrak{t}(i; \mathcal{G}) &:= \mathbb{E}[\bar{T}_i | x_0 = i] \end{aligned}$$

We similarly denote access time and return time of BBRW as $(\tilde{\mathfrak{t}}(i, j), \tilde{\mathfrak{t}}(i \rightarrow j, k))$ and $\tilde{\mathfrak{t}}(i; \mathcal{G})$, respectively. Note that the first step of BBRW shows the same behavior as the SRW since there is no previous node on the first step.

Finally, for a tree graph \mathcal{G} and pair of nodes i, j in Proposition 1, we denote the unique paths between i and j as (v_0, \dots, v_N) with $i = v_0, j = v_N$. We also let N denote the distance between the nodes i and j .

B.2 Access Time of Simple Random Walks

In this Section, we derive formulas for the access time of simple random walks on tree graphs, Proposition 2. We show it by the following process:

1. We decompose the access time for two nodes i and j into a summation of access time of neighboring nodes in the path, i.e., $\mathfrak{t}(v_n, v_{n+1})$ for $n \in 0, \dots, N-1$. (Lemma 2)
2. We evaluate the access time of neighboring nodes in the path, e.g., $\mathfrak{t}(v_n, v_{n+1})$, in using the number of edges in a graph. (Lemma 3)

B.2.1 Decomposition of Access Time

First, we show that the access time is equal to the sum of access times between neighboring nodes. When a random walker travels from v_0 to v_N , it must pass through all nodes v_n on the paths. We can consider the entire time taken as the summation of the time intervals between when the walker arrived at v_n and when it arrived at v_{n+1} . It is an intuitive way of understanding Lemma 2.

Lemma 2. *Given a tree \mathcal{G} and path (v_0, \dots, v_N) ,*

$$\mathfrak{t}(v_0, v_N) = \sum_{n=0}^{N-1} \mathfrak{t}(v_n, v_{n+1}).$$

Proof. The proof outline is as follows. Given a unique path (v_0, \dots, v_N) between v_0 and v_N , any walk between (v_0, v_N) can be decomposed into a series of $N-1$ walks between $(v_0, v_1), (v_1, v_2), \dots, (v_{N-1}, v_N)$ such that the walk between (v_n, v_{n+1}) does not contain a node v_{n+1} except at the end point. The expected length of each walk is $t(v_n, v_{n+1})$.

To be specific, consider the following decomposition.

$$\mathfrak{t}(v_0, v_N) = \mathbb{E}[T_{v_N} | \mathbf{x}_0 = v_0] = \sum_{n=0}^{N-1} \mathbb{E}[T_{v_{n+1}} - T_{v_n} | \mathbf{x}_0 = v_0] \quad (6)$$

From the Markov property of SRW:

$$\mathbb{E}[T_{v_{n+1}} - T_{v_n} | \mathbf{x}_0 = v_0] = \mathfrak{t}(v_n, v_{n+1})$$

Thus,

$$\mathfrak{t}(v_0, v_N) = \sum_{n=0}^{N-1} \mathbb{E}[T_{v_{n+1}} - T_{v_n} | \mathbf{x}_0 = v_0] = \sum_{n=0}^{N-1} \mathfrak{t}(v_n, v_{n+1})$$

□

B.2.2 Access Time between Neighbors

Now, we derive the formula for the access time between neighboring nodes.

Lemma 3. *Given a tree graph \mathcal{G} and adjacent nodes i, j , the associated access time $\mathfrak{t}(i, j)$ is defined as follows:*

$$\mathfrak{t}(i, j) = 1 + 2|\mathcal{E}(\mathcal{G}_i)|,$$

where $\mathcal{E}(\mathcal{G})$ is edge set of graph \mathcal{G} and \mathcal{G}_i is the subtree produced by deleting edge (i, j) and choosing connected component of i .

Proof. Given a random walk from i to j that only contains j once, every time the walk lands at the node i , its future events can be categorized into two scenarios:

1. The walk transitions from node i to j based on transitioning with respect to the edge $(i, j) \in \mathcal{E}$ with probability $\frac{1}{d_i}$. The walk terminates.
2. The walk fails to reach j and continues the walk in the subtree \mathcal{G}_i until arriving at the node i again. Note that the walk cannot arrive at node j without arriving at the node i in prior. In other words, the walk continues for the return time of i with respect to the graph \mathcal{G}_i .

The two scenarios implies that, every time the walk arrives at node i , the walk terminates with probability $\frac{1}{d_i}$. Then the number of trials follow the geometric distribution and consequently, the average number of trial is d_i . In other words, the walk falls into the scenario of type 2, for $d_i - 1$ times in average. We have to traverse at least one edge (i, j) . The expected total penalty is the product of average failure penalty and the average number of failure. Thus,

$$\mathfrak{t}(i, j) = 1 + (d_i - 1)\mathfrak{t}(i; \mathcal{G}_i), \quad (7)$$

where $\mathfrak{t}(i; \mathcal{G}_i)$ is the return time of i with respect to the subgraph \mathcal{G}_i . Since the return time is the ratio of the sum of the degree to the degree of the node, as stated by [Fasino et al. \(2023\)](#), the following equation holds:

$$\mathfrak{t}(i; \mathcal{G}_i) = \frac{2|\mathcal{E}(\mathcal{G}_i)|}{d_i - 1},$$

which directly implies our conclusion of $\mathfrak{t}(i, j) = 1 + 2|\mathcal{E}(\mathcal{G}_i)|$. \square

B.2.3 Main Result

Using [Lemma 2](#) and [Lemma 3](#), we arrive at our main result for SRW on trees.

Proposition 2. *Given a tree \mathcal{G} and pair of nodes $i, j \in \mathcal{V}$, the following equations holds for the access time of SRW between u and v .*

$$\mathfrak{t}(i, j) = \sum_{n=0}^{N-1} (1 + 2|\mathcal{E}(\mathcal{G}_n)|),$$

where $\mathcal{E}(\mathcal{G})$ is edge set of graph \mathcal{G} and \mathcal{G}_n is the subtree produced by deleting edge (v_n, v_{n+1}) and choosing connected component of v_n .

Proof. The proof is a straightforward combination of [Lemma 2](#) and [Lemma 3](#).

$$\mathfrak{t}(i, j) = \sum_{n=0}^{N-1} \mathfrak{t}(v_n, v_{n+1}) = \sum_{n=0}^{N-1} (1 + 2|\mathcal{E}(\mathcal{G}_n)|).$$

\square

B.3 Access Time of Begrudgingly Backtracking Random Walks

In this section, we derive formulas for the access time of begrudgingly backtracking random walks on tree graphs. As mentioned in the preliminaries, we use $\tilde{\mathfrak{t}}$ to denote the access time of begrudgingly backtracking. At a high level, this section consists of the following proofs.

1. We decompose the access time for two nodes v_0 and v_N into summations of access time of neighboring nodes in the path. ([Lemma 4](#))
2. The decomposed term in 1. can be formulated using (i) the return time and (ii) the access time between neighboring nodes. ([Lemma 5](#))
3. The return time of a node can be formulated in terms of the number of edges and the degree of a node. ([Lemma 6](#))
4. The access time between neighboring nodes can be formulated using the number of edges and the degree of the node. ([Lemma 7](#))

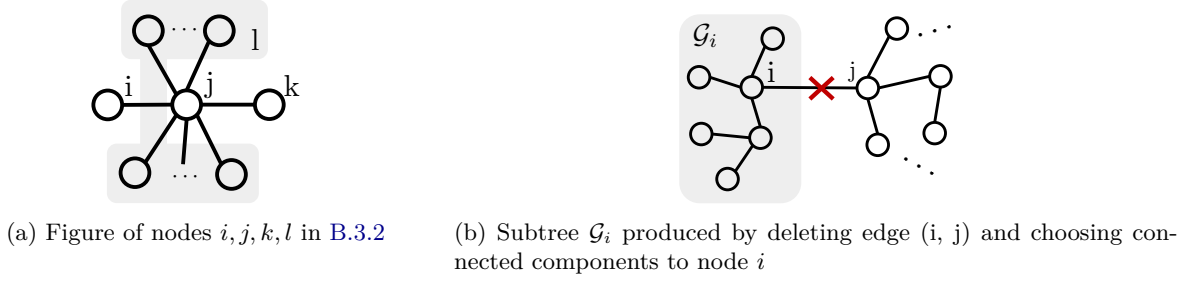


Figure 5: References for section B.3.2

B.3.1 Decomposition of Access Time

We start by decomposing the access time $\tilde{\mathbf{t}}(i, j)$ similar to the one in Lemma 2. However, a key difference exists in the BBRW random walk. When the walk first arrives at node v_n , it previously passed the node v_{n-1} . Therefore, we cannot return to v_{n-1} upon the first failure to reach v_{n+1} .

Lemma 4. *Consider a tree \mathcal{G} and path (v_0, \dots, v_N) . Then the access time of BBRW between node v_0 and v_N can be derived as follows:*

$$\tilde{\mathbf{t}}(v_0, v_N) = \tilde{\mathbf{t}}(v_0, v_0 \rightarrow v_1) + \sum_{l=1}^{N-1} \left(\tilde{\mathbf{t}}(v_{n-1} \rightarrow v_n, v_{n+1}) - 1 \right)$$

Proof. The proof is similar to Lemma 2 such that we decompose the walk into a series of $N-1$ walks between $(v_0, v_1), (v_1, v_2), \dots, (v_{N-1}, v_N)$ such that the walk between v_n, v_{n+1} does not contain a node v_{n+1} except at the endpoint. The expected length of each walk is $\tilde{\mathbf{t}}(v_{n-1} \rightarrow v_n, v_{n+1}) - 1$. Note that $\tilde{\mathbf{t}}(v_{n-1} \rightarrow v_n, v_{n+1})$ counts the length of walk from v_{n-1} to v_{n+1} , hence should be subtracted to represent the length of walk from v_n to v_{n+1} .

To be specific, the proof of Lemma 2, we mentioned that Equation (6) holds for general random walks.

$$\tilde{\mathbf{t}}(v_0, v_N) = \sum_{n=0}^{N-1} \mathbb{E}[T_{v_{n+1}} - T_{v_n} | \mathbf{x}_0 = v_0]$$

When the BBRW random walk reaches v_n at T_{v_n} , it previously passed the node v_{n-1} . (i.e., $\mathbf{x}_{T_{v_n}-1} = v_{n-1}$). Therefore, we can think the random walk after the time $t = T_{v_n} - 1$ as random walk starting at v_{n-1} with the second node $\mathbf{x}_1 = v_n$. Thus,

$$\mathbb{E}[T_{v_{n+1}} - T_{v_n} | \mathbf{x}_0 = v_0] = \tilde{\mathbf{t}}(v_{n-1} \rightarrow v_n, v_{n+1}) - 1 \quad \text{for } l \geq 1.$$

□

B.3.2 Expressing Transition-conditioned Access Time Using Subgraph-return Time

We further prove the following formula for computing the transition-conditioned access time $\tilde{\mathbf{t}}(v_{n-1} \rightarrow v_n, v_{n+1})$.

Lemma 5. *Given a tree $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and three nodes i, j , and k such that $(i, j), (j, k) \in \mathcal{E}$,*

$$\tilde{\mathbf{t}}(i \rightarrow j, k) - 1 = \tilde{\mathbf{t}}(j, k) - \frac{d_i - 1}{d_j} \tilde{\mathbf{t}}(i; \mathcal{G}_i) - \frac{2}{d_j}, \quad (8)$$

where $\tilde{\mathbf{t}}(i; \mathcal{H})$ is return time of node i in a subgraph \mathcal{H} , and \mathcal{G}_i is the subtree produced by deleting edge (i, j) and choosing connected components of i .

Proof. We start with establishing basic equalities for access time of BBRW.

The first equality describes how the access time of a random walk from $i \rightarrow j$ to k can be expressed by its subset $j \rightarrow l, k$ where $l \neq i$ due to the non-backtracking property.

$$\tilde{\mathbf{t}}(i \rightarrow j, k) - 1 = \frac{1}{d_j - 1} + \sum_{l \in \mathcal{N}(j) \setminus \{i, k\}} \frac{1}{d_j - 1} \tilde{\mathbf{t}}(j \rightarrow l, k) \quad (9)$$

The next equality describes how the access time from j to k can be decomposed by considering subsequent transitions to some l in the neighborhood of j .

$$\tilde{\mathbf{t}}(j, k) = \frac{1}{d_j} \sum_{l \in \mathcal{N}(j)} \tilde{\mathbf{t}}(j \rightarrow l, k), \quad (10)$$

When plugging in Equation (10) into Equation (9), one can see that we need to consider $l = i, k$ for $\tilde{\mathbf{t}}(j \rightarrow l, k)$. In case of k , it is trivially 1. For the case of i , i.e., $\tilde{\mathbf{t}}(j \rightarrow i, k)$, it can be defined as follow:

$$\tilde{\mathbf{t}}(j \rightarrow i, k) = 1 + (d_i - 1)\tilde{\mathbf{t}}(i; \mathcal{G}_i) + \tilde{\mathbf{t}}(i \rightarrow j, k). \quad (11)$$

This is similar to Equation (7), where \mathcal{G}_i describes how the walk from $j \rightarrow i$ to k can be divided into two scenarios: (i) continues the walk in \mathcal{G}_i or (ii) i transitions into j with probability $1/(d_i - 1)$.

Starting from Equation (9), one can derive the following relationship:

$$\begin{aligned} \tilde{\mathbf{t}}(i \rightarrow j, k) - 1 &= \frac{1}{d_j - 1} \cdot 1 + \sum_{l \in \mathcal{N}(j) \setminus \{i, k\}} \frac{1}{d_j - 1} \tilde{\mathbf{t}}(j \rightarrow l, k) \\ &\stackrel{(a)}{=} \frac{1}{d_j - 1} + \frac{1}{d_j - 1} \left[\sum_{l \in \mathcal{N}(j)} \tilde{\mathbf{t}}(j \rightarrow l, k) - \tilde{\mathbf{t}}(j \rightarrow i, k) - \tilde{\mathbf{t}}(j \rightarrow k, k) \right] \\ &\stackrel{(b)}{=} \frac{1}{d_j - 1} + \frac{1}{d_j - 1} \left[d_j \tilde{\mathbf{t}}(j, k) - \tilde{\mathbf{t}}(j \rightarrow i, k) - 1 \right] \\ &\stackrel{(c)}{=} \frac{1}{d_j - 1} \left[d_j \tilde{\mathbf{t}}(j, k) - (d_i - 1)\tilde{\mathbf{t}}(i; \mathcal{G}_i) - \tilde{\mathbf{t}}(i \rightarrow j, k) - 1 \right], \end{aligned}$$

where (a) is from Equation (9), (b) is from Equation (10) and $\tilde{\mathbf{t}}(j \rightarrow k, k) = 1$, (c) is from Equation (11).

Now, by multiplying $d_j - 1$ on both sides, we get

$$\begin{aligned} (d_j - 1) \left(\tilde{\mathbf{t}}(i \rightarrow j, k) - 1 \right) &= d_j \tilde{\mathbf{t}}(j, k) - (d_i - 1)\tilde{\mathbf{t}}(i; \mathcal{G}_i) - \tilde{\mathbf{t}}(i \rightarrow j, k) - 1 \\ &= d_j \tilde{\mathbf{t}}(j, k) - (d_i - 1)\tilde{\mathbf{t}}(i; \mathcal{G}_i) - \left(\tilde{\mathbf{t}}(i \rightarrow j, k) - 1 \right) - 2. \end{aligned}$$

Thus, we can conclude for $\tilde{\mathbf{t}}(i \rightarrow j, k) - 1$:

$$\tilde{\mathbf{t}}(i \rightarrow j, k) - 1 = \tilde{\mathbf{t}}(j, k) - \frac{d_i - 1}{d_j} \tilde{\mathbf{t}}(i; \mathcal{G}_i) - \frac{2}{d_j}.$$

□

B.3.3 Return Time with respect to a Subgraph

Now we need the formula for the return time $\tilde{\mathbf{t}}(i; \mathcal{G})$ with respect to a tree-graph \mathcal{G} . For the return time, we prove that the return time of BBRW is less than or equal to that of SRW.

Lemma 6. *Given a tree $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a node i , the return time of i is,*

$$\tilde{\mathfrak{t}}(i; \mathcal{G}) = \frac{2|\mathcal{E}|}{d_i}$$

Proof. We can think the tree as rooted tree where root is i . We use mathematical induction with tree height of i .

Consider the base case where the height of tree is 1. Then, whatever we choose as the next node x_1 from $x_0 = i$, we return to i at second transition (i.e., $x_2 = i$) since all the neighbors of i is leaf node. Since $d_i = |\mathcal{E}|$ for tree with height 1, $\tilde{\mathfrak{t}}(i; \mathcal{G}) = 2 = \frac{2|\mathcal{E}|}{d_i}$.

Now, assume that the lemma holds for the tree with a height less than $k \geq 1$. It suffices to show that the lemma also holds for a tree with height $k + 1$ and its root i .

From the same perspective in the proofs of the [Lemma 3](#), we can view the random walk returning to i as follows:

1. We choose a node $x_1 = l$ from $\mathcal{N}(i)$ uniformly. Then, from node l , we return to l to reach i . (i.e., we have to pay penalty amounts to return time of l)
2. In node l , we try to reach i by selecting edge (l, i) with probability $\frac{1}{d_l - 1}$. Thus, the average number of failure of failure is $d_l - 2$.
3. If we fail to reach i from l , we should return to l for a next chance to reach i . i.e., a average failure penalty amounts to a return time of l in subtree with root l .

Hence,

$$\tilde{\mathfrak{t}}(i; \mathcal{G}) = 1 + \sum_{l \in \mathcal{N}(i)} \frac{1}{d_i} \left\{ 1 + \tilde{\mathfrak{t}}(l; \mathcal{G}_l) \cdot \left((d_l - 2) + 1 \right) \right\} = 2 + \sum_{l \in \mathcal{N}(i)} \frac{1}{d_i} \cdot \tilde{\mathfrak{t}}(l; \mathcal{G}_l) \cdot (d_l - 1),$$

where $\tilde{\mathfrak{t}}(l; \mathcal{G})$ is return time of l with respect to \mathcal{G} and \mathcal{G}_l is the subtree with root l .

Since the subtree with root l has height less than k , $\tilde{\mathfrak{t}}(l; \mathcal{G}_l) = \frac{2|\mathcal{E}(\mathcal{G}_l)|}{d_l - 1}$. Therefore,

$$\begin{aligned} \tilde{\mathfrak{t}}(i; \mathcal{G}) &= 2 + \frac{1}{d_i} \sum_{l \in \mathcal{N}(i)} \tilde{\mathfrak{t}}(l; \mathcal{G}_l) \cdot (d_l - 1) \\ &= 2 + \frac{1}{d_i} \sum_{l \in \mathcal{N}(i)} \frac{2|\mathcal{E}(\mathcal{G}_l)|}{d_l - 1} \cdot (d_l - 1) \\ &= 2 + \frac{1}{d_i} \sum_{l \in \mathcal{N}(i)} 2|\mathcal{E}(\mathcal{G}_l)| \\ &= \frac{2 \left(d_i + \sum_{l \in \mathcal{N}(i)} |\mathcal{E}(\mathcal{G}_l)| \right)}{d_i} \\ &= \frac{2|\mathcal{E}|}{d_i} \end{aligned}$$

□

B.3.4 Access Time between Neighbors

For access time between neighbors, we get the similar result to [Lemma 3](#).

Lemma 7. *Given a tree \mathcal{G} and adjacent nodes i, j ,*

$$\tilde{\mathfrak{t}}(i, j) = 1 + 2|\mathcal{E}(\mathcal{G}_i)| \cdot \frac{d_i - 1}{d_i},$$

where $\mathcal{E}(\mathcal{G})$ is edge set of graph \mathcal{G} , and \mathcal{G}_i is the subtree produced by deleting edge (i, j) and choosing connected component of i .

Proof. From the same perspective in the proofs of the [Lemma 3](#), we analyze success probability for each trial.

For the first trial, success probability is $\frac{1}{d_i}$ since there is no previous node. After the first trial, success probability is fixed to $\frac{1}{d_i-1}$. On the each trial, the average failure penalty amounts to $\tilde{\mathfrak{t}}'(i)$, which is the return time of i on subtree \mathcal{G}_i .

Thus, average number of trial until first success is,

$$(\text{Average number of trial}) = \frac{1}{d_i} \cdot 1 + \frac{d_i - 1}{d_i} \cdot \left(1 + (d_i - 1)\right)$$

The average number of failure is as follows:

$$\begin{aligned} (\text{Average number of failure}) &= (\text{Average number of trial}) - 1 \\ &= \frac{1}{d_i} \cdot 1 + \frac{d_i - 1}{d_i} \cdot (1 + d_i - 1) - 1 \\ &= \frac{(d_i - 1)^2}{d_i} \end{aligned}$$

Thus, expected total penalty is as follows:

$$(\text{Expected total penalty}) = \tilde{\mathfrak{t}}(i; \mathcal{G}_i) \cdot \frac{(d_i - 1)^2}{d_i}$$

Since $\tilde{\mathfrak{t}}(i; \mathcal{G}_i) = \frac{2|\mathcal{E}(\mathcal{G}_i)|}{d_i-1}$ by [Lemma 6](#), $\tilde{\mathfrak{t}}(u, v) = 1 + 2|\mathcal{E}(\mathcal{G}_i)| \frac{d_i-1}{d_i}$. □

B.3.5 Main Result

Combining the results provides the following theorem.

Proposition 3. *Given a tree \mathcal{G} and pair of nodes i, j , the following equations holds for the access time of BBRW between i and j .*

$$\tilde{\mathfrak{t}}(i, j) = \sum_{n=0}^{N-1} \left(1 + 2|\mathcal{E}(\mathcal{G}_n)| \cdot \frac{d_{v_n} - 1}{d_{v_n}}\right) - \sum_{l=1}^{N-1} \frac{2|\mathcal{E}(\mathcal{G}_{n-1})|}{d_{v_n}} - \sum_{l=1}^{N-1} \frac{2}{d_{v_n}},$$

where $\mathcal{E}(\mathcal{G})$ is the edge set of \mathcal{G} and \mathcal{G}_n is the subtree produced by deleting edge (v_n, v_{n+1}) and choosing connected component of v_n .

Proof. By Lemma 4 and Equation (8),

$$\begin{aligned}
\tilde{\mathfrak{t}}(v_0, v_N) &= \tilde{\mathfrak{t}}(v_0, v_1) + \sum_{l=1}^{N-1} \left\{ \tilde{\mathfrak{t}}(v_{n-1} \rightarrow v_n, v_{n+1}) - 1 \right\} \\
&= \tilde{\mathfrak{t}}(v_0, v_1) + \sum_{l=1}^{N-1} \left\{ \tilde{\mathfrak{t}}(v_n, v_{n+1}) - \frac{d_{v_{n-1}} - 1}{d_{v_n}} \tilde{\mathfrak{t}}'(v_{n-1}) - \frac{2}{d_{v_n}} \right\} \\
&= \sum_{n=0}^{N-1} \tilde{\mathfrak{t}}(v_n, v_{n+1}) - \sum_{l=1}^{N-1} \frac{d_{v_{n-1}} - 1}{d_{v_n}} \tilde{\mathfrak{t}}'(v_{n-1}) - \sum_{l=1}^{N-1} \frac{2}{d_{v_n}} \\
&= \sum_{n=0}^{N-1} \left(1 + 2|\mathcal{E}(\mathcal{G}_l)| \cdot \frac{d_{v_n} - 1}{d_{v_n}} \right) - \sum_{l=1}^{N-1} \frac{2|\mathcal{E}(\mathcal{G}_{l-1})|}{d_{v_n}} - \sum_{l=1}^{N-1} \frac{2}{d_{v_n}}
\end{aligned}$$

□

B.4 Proof of Proposition 1

Finally, we compare the access time of two random walks in a tree. Recall Proposition 1.

Proposition 1. *Given a tree $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a pair of nodes $i, j \in \mathcal{V}$, the access time of begrudgingly backtracking random walk is equal or smaller than that of a simple random walk. The equality holds if and only if the walk length is 1.*

Proof. Let $\mathcal{E}(\mathcal{G})$ be the edge set of \mathcal{G} and \mathcal{G}_n be the subtree produced by deleting edge (v_n, v_{n+1}) and choosing connected component of v_n . Then,

$$\begin{aligned}
\tilde{\mathfrak{t}}(v_0, v_N) - \mathfrak{t}(v_0, v_N) &= \sum_{n=0}^{N-1} \left(1 + 2|\mathcal{E}(\mathcal{G}_n)| \cdot \frac{d_{v_n} - 1}{d_{v_n}} \right) - \sum_{l=1}^{N-1} \frac{2|\mathcal{E}(\mathcal{G}_{n-1})|}{d_{v_n}} \\
&\quad - \sum_{l=1}^{N-1} \frac{2}{d_{v_n}} - \sum_{n=0}^{N-1} (1 + 2|\mathcal{E}(\mathcal{G}_n)|) \\
&= - \sum_{n=0}^{N-1} \frac{2|\mathcal{E}(\mathcal{G}_n)|}{d_{v_n}} - \sum_{l=1}^{N-1} \frac{2|\mathcal{E}(\mathcal{G}_{n-1})|}{d_{v_n}} - \sum_{l=1}^{N-1} \frac{2}{d_{v_n}} \leq 0
\end{aligned}$$

Therefore the access time of two nodes are always less or equal in begrudgingly backtracking random walk than simple random walks, where equality holds for random walks with length 1.

□

C Proofs for Section 4.1

C.1 Preliminaries

Let \mathcal{G} be a graph with a set of n vertices \mathcal{V} , a set of m edges $\mathcal{E} \in \mathcal{V}^2$. We use x_i to denote the node-wise feature for $i \in \mathcal{V}$, and d_i for the degree of node $i \in \mathcal{V}$. The adjacency matrix $A \in \mathbb{R}^{n \times n}$ encodes the connectivity of graph \mathcal{G} . For node $i \in \mathcal{V}$, we define the set of incident outgoing edges of i as $N_e^+(i)$, the set of incident incoming edges of i as $N_e^-(i)$, and let $N_e(i) = N_e^+(i) \cup N_e^-(i)$ be the set of all incident edges of node i . Also, recall the non-backtracking matrix $B \in \{0, 1\}^{2|\mathcal{E}| \times 2|\mathcal{E}|}$ and the incidence matrix $C \in \mathbb{R}^{2|\mathcal{E}| \times |\mathcal{V}|}$:

$$B_{(\ell \rightarrow k), (j \rightarrow i)} = \begin{cases} 1 & \text{if } k = j, \ell \neq i \\ 0 & \text{otherwise} \end{cases}, \quad C_{(k \rightarrow j), i} = \begin{cases} 1 & \text{if } j = i \text{ or } k = i \\ 0 & \text{otherwise} \end{cases}.$$

We also define D_{out} and D_{in} as the out-degree and in-degree matrices of NBA-GNN, respectively, counting the number of outgoing and incoming edges for each edge. These are diagonal matrices with $(D_{out})_{(j \rightarrow i), (j \rightarrow i)} = \sum_{\ell \rightarrow k} B_{(j \rightarrow i), (\ell \rightarrow k)}$ and $(D_{in})_{(j \rightarrow i), (j \rightarrow i)} = \sum_{\ell \rightarrow k} B_{(\ell \rightarrow k), (j \rightarrow i)}$, for each index $j \rightarrow i$. Next, we introduce \hat{B} as the normalized non-backtracking matrix augmented with self-loops: $\hat{B} = (D_{out} + I)^{-\frac{1}{2}}(B + I)(D_{in} + I)^{-\frac{1}{2}}$. Then one obtains the following sensitivity bound of NBA-GNN. We also let \tilde{C} denote a matrix where $\tilde{C}_{(k \rightarrow j), i} = C_{(k \rightarrow j), i} + C_{(j \rightarrow k), i}$.

Consequently, one can express our NBA-GNN updates:

$$h_{j \rightarrow i}^{(t+1)} = \phi^{(t)} \left(h_{j \rightarrow i}^{(t)}, \sum_{(k, \ell) \in \mathcal{E}} \hat{B}_{(\ell \rightarrow k), (j \rightarrow i)} \psi^{(t)} \left(h_{\ell \rightarrow k}^{(t)}, h_{j \rightarrow i}^{(t)} \right) \right), \quad (12)$$

where $\phi^{(t)}$ and $\psi^{(t)}$ corresponds to the update and the aggregation as described in Equation (3). Next, the node-wise aggregation step can be described as follows:

$$h_i = \sigma \left(\sum_{(j, k) \in \mathcal{E}} C_{(k \rightarrow j), i} \rho \left(h_{k \rightarrow j}^{(T)} \right), \sum_{(j, k) \in \mathcal{E}} C_{(j \rightarrow k), i} \rho \left(h_{j \rightarrow k}^{(T)} \right) \right).$$

C.2 Proof of Lemma 1

The original sensitivity bound from Topping et al. (2022) for a hidden representation of node j and an initial feature of node i is defined as the following.

Proposition 4. Sensitivity bound. (Topping et al., 2022) Assume an MPNN defined in Equation (12). Let two nodes $i, j \in \mathcal{V}$ with distance T . If $|\nabla \phi^{(t)}| \leq \alpha$ and $|\nabla \psi^{(t)}| \leq \beta$ for $0 \leq t < T$, then

$$\left\| \frac{\partial h_j^{(T)}}{\partial x_i} \right\| \leq (\alpha\beta)^T (\hat{A}^T)_{j, i}.$$

Now, we show that the sensitivity bound for non-backtracking GNNs can be defined as following. Following Topping et al. (2022), we assume the node features and hidden representations are scalar for better understanding.

Lemma 1. Consider two nodes $i, j \in \mathcal{V}$ with a random walk distance T given a $(T-1)$ -layer NBA-GNN as described in Equation (3) and Equation (4). Suppose $|\nabla \phi^{(t)}|, |\nabla \sigma| \leq \alpha$, $|\nabla \psi^{(t)}|, |\nabla \rho| \leq \beta$, and $|\nabla \phi| \leq \gamma$ for $0 \leq t < T$. Then the following holds:

$$\left\| \frac{\partial h_j}{\partial x_i} \right\| \leq (\alpha\beta)^T \gamma (\tilde{C}^\top \hat{B}^{(T-1)} \tilde{C})_{j, i}.$$

Proof. Just a straight calculation using the chain rule is enough to derive the above upper bound.

$$\begin{aligned}
\left\| \frac{\partial h_j}{\partial x_i} \right\| &= \left\| \partial_1 \sigma \partial_2 \rho \left(\sum_{k \rightarrow \ell \in N_e^-(j)} \frac{\partial h_{k \rightarrow \ell}^{(T-1)}}{\partial x_i} \right) + \partial_1 \sigma \partial_3 \rho \left(\sum_{k \rightarrow \ell \in N_e^+(j)} \frac{\partial h_{k \rightarrow \ell}^{(T-1)}}{\partial x_i} \right) \right\| \\
&= \left\| \partial_1 \sigma \partial_2 \rho \left(\sum_{k \rightarrow \ell \in N_e^-(j)} \frac{\partial h_{k \rightarrow \ell}^{(T-1)}}{\partial x_i} \right) \right\| \\
&\leq \alpha \beta \left(\sum_{k \rightarrow \ell \in N_e^-(j)} \left\| \frac{\partial h_{k \rightarrow \ell}^{(T-1)}}{\partial x_i} \right\| \right), \tag{13}
\end{aligned}$$

where the bound for the derivatives were $|\nabla \sigma| \leq \alpha, |\nabla \rho| \leq \beta$.

Thus considering the message update from Equation (12),

$$\begin{aligned}
\frac{\partial h_{k \rightarrow \ell}^{(T-1)}}{\partial x_i} &= \partial_1 \phi^{(T-2)} \frac{\partial h_{k \rightarrow \ell}^{(T-2)}}{\partial x_i} + \partial_2 \phi^{(T-2)} \left(\sum_{k' \rightarrow \ell' \in N_e^-(k)} \hat{B}_{k' \rightarrow \ell', k \rightarrow \ell} \frac{\partial h_{k \rightarrow \ell}^{(T-2)}}{\partial x_i} \right) \\
&= \partial_2 \phi^{(T-2)} \left(\sum_{k' \rightarrow \ell' \in N_e^-(k)} \hat{B}_{k' \rightarrow \ell', k \rightarrow \ell} \frac{\partial h_{k \rightarrow \ell}^{(T-2)}}{\partial x_i} \right),
\end{aligned}$$

since the distance between node i and node j is at least T , therefore $\frac{\partial h_{k \rightarrow \ell}^{(T-2)}}{\partial x_i} = 0$.

Now, let the path from node i to node j as $s(i, j)$, where s_t denotes the t -th node in the walk s , i.e., $s_0 = i, s_T = j$. Using the bound of the derivative of functions, we can iterate like the following.

$$\begin{aligned}
&\left\| \frac{\partial h_{k \rightarrow \ell}^{(T-1)}}{\partial x_i} \right\| \\
&\leq \alpha \beta \left(\sum_{k' \rightarrow \ell' \in N_e^-(k)} \hat{B}_{k' \rightarrow \ell', k \rightarrow \ell} \left\| \frac{\partial h_{k \rightarrow \ell}^{(T-2)}}{\partial x_i} \right\| \right) \\
&\leq (\alpha \beta)^{T-1} \left(\sum_{(s_0, \dots, s_T) \in s(i, j)} \hat{B}_{s_0 \rightarrow s_1, s_1 \rightarrow s_2} \cdots \hat{B}_{s_{T-2} \rightarrow s_{T-1}, s_{T-1} \rightarrow s_T} \cdot \left\| \frac{\partial h_{s_0 \rightarrow s_1}^{(0)}}{\partial x_i} \right\| \right) \\
&= (\alpha \beta)^{T-1} \left(\sum_{(s_0, \dots, s_T) \in s(i, j)} \prod_{t=1}^{T-1} \hat{B}_{s_{t-1} \rightarrow s_t, s_t \rightarrow s_{t+1}} \left\| \frac{\partial h_{s_0 \rightarrow s_1}^{(0)}}{\partial x_i} \right\| \right) \\
&\leq (\alpha \beta)^{T-1} \gamma \left(\sum_{(s_0, \dots, s_T) \in s(i, j)} \left(\prod_{t=1}^{T-1} \hat{B}_{s_{t-1} \rightarrow s_t, s_t \rightarrow s_{t+1}} \right) \right) \tag{14}
\end{aligned}$$

Substitute the inequality Equation (14) into Equation (13) to get the final result. Then, we can get

$$\begin{aligned}
\left\| \frac{\partial h_j}{\partial x_i} \right\| &\leq (\alpha \beta)^T \gamma \left(\sum_{(s_0, \dots, s_T) \in s(i, j)} \left(\prod_{t=1}^{T-1} \hat{B}_{s_{t-1} \rightarrow s_t, s_t \rightarrow s_{t+1}} \right) \right) \\
&= (\alpha \beta)^T \gamma (\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j, i}
\end{aligned}$$

□

since paths can be expressed using the power of adjacency-type matrix.

C.3 Proof of Proposition 1

We have defined the sensitivity bound for NBA-GNNs. Now, we show that the sensitivity bound of NBA-GNNs are bigger than the sensitivity bound of GNNs.

Theorem 1. *Let \hat{A} denote the degree-normalized matrix. Then, for any pair of nodes $i, j \in \mathcal{V}$ with distance T , the sensitivity bound of NBA-GNN is larger than that of conventional GNNs (Topping et al., 2022), i.e., $(\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i} \geq (\hat{A}^T)_{j,i}$. For d -regular graphs, $(\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i}$ decays slower by $O(d^{-T})$, while $(\hat{A}^T)_{j,i}$ decays with $O((d+1)^{-T})$.*

Proof. Identical to the notations above, we denote the list of nodes from node i to node j as path $s(i, j)$ where $s(i, j) = (s_0 = i, s_1, \dots, s_{T-1}, s_T = j)$.

$$\begin{aligned}
(\hat{A}^T)_{j,i} &= \sum_{(s_0, \dots, s_T) \in s(i,j)} \hat{A}_{s_0, s_1} \cdots \hat{A}_{s_{T-1}, s_T} \\
&= \sum_{(s_0, \dots, s_T) \in s(i,j)} (d_i + 1)^{-\frac{1}{2}} \cdot (d_j + 1)^{-\frac{1}{2}} \cdot \prod_{t=1}^{T-1} (d_{s_t} + 1)^{-1} \\
(\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i} &= \sum_{(s_0, \dots, s_T) \in s(i,j)} d_{s_0}^{-\frac{1}{2}} \cdot \hat{B}_{s_0 \rightarrow s_1, s_1 \rightarrow s_2} \cdots \hat{B}_{s_{T-2} \rightarrow s_{T-1}, s_{T-1} \rightarrow s_T} \cdot d_{s_T}^{-\frac{1}{2}} \\
&= \sum_{(s_0, \dots, s_T) \in s(i,j)} d_{s_0}^{-\frac{1}{2}} \cdot d_{s_1}^{-1} \cdots d_{s_{T-1}}^{-1} \cdot d_{s_T}^{-\frac{1}{2}} \\
&= \sum_{(s_0, \dots, s_T) \in s(i,j)} d_{s_0}^{-\frac{1}{2}} \cdot d_{s_T}^{-\frac{1}{2}} \cdot \prod_{t=1}^{T-1} d_{s_t}^{-1}
\end{aligned}$$

Now, for a path $(s_0 = i, \dots, s_T = j)$,

$$(d_i + 1)^{-\frac{1}{2}} (d_j + 1)^{-\frac{1}{2}} \prod_{t=1}^{T-1} (d_{s_t} + 1)^{-1} \leq d_i^{-\frac{1}{2}} d_j^{-\frac{1}{2}} \prod_{t=1}^{T-1} d_{s_t}^{-1}$$

Each term in $(\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i}$ is larger than $(\hat{A}^T)_{j,i}$, thus $(\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i} \geq (\hat{A}^T)_{j,i}$ is trivial.

For d -regular graphs, $d_i = d, \forall i \in \mathcal{V}$. Therefore the sensitivity bound can be written as following:

$$(\hat{A}^T)_{j,i} = (d+1)^{-T}, \quad (\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i} = d^{-T}.$$

So $(\tilde{C}^\top \hat{B}^{T-1} \tilde{C})_{j,i}$ decays slower by $O(d^{-T})$, while $(\hat{A}^T)_{j,i}$ decays with $O((d+1)^{-T})$.

□

D Proofs for Section 4.2

To assess the expressive capabilities, we initially make an assumption about the graphs, considering that it is generated from the Stochastic Block Model (SBM), which is defined as follows:

Definition 1. *Stochastic Block Model (SBM) is generated using parameters (n, K, α, P) , where n denotes the number of vertices, K is the number of communities, $\alpha = (\alpha_1, \dots, \alpha_K)$ represents the probability of each vertex being assigned to communities $\mathcal{V}_1, \dots, \mathcal{V}_K$, and P_{ij} denotes the probability of an edge $(v, w) \in \mathcal{E}$ between $v \in \mathcal{V}_i$ and $w \in \mathcal{V}_j$.*

Numerous studies have focused on the problem of achieving exact recovery of communities within the SBM. However, these investigations typically address scenarios in which the average degree is at least on the order of $\Omega(\log n)$ (Abbe, 2017). It is well-established that the information-theoretic limit in such cases can be reached through the utilization of the spectral method, as demonstrated by Yun & Proutière (2016). In contrast, when dealing with a graph characterized by the average degree much smaller, specifically $o(\log n)$, recovery using the graph spectrum becomes a more challenging endeavor. This difficulty arises due to the fact that the $n^{1-o(1)}$ largest eigenvalues and their corresponding eigenvectors are influenced by the high-degree vertices, as discussed in Benaych-Georges et al. (2019).

However, real-world benchmark datasets often fall within the category of very sparse graphs. For example, the citation network dataset discussed in Sen et al. (2008) has an average degree of less than three. In such scenarios, relying solely on an adjacency matrix may not be an efficient approach for uncovering the hidden graph structure. Fortunately, an alternative strategy is available by utilizing a non-backtracking matrix.

D.1 Proof of Theorem 2

Let's rephrase the formal version of Theorem 2 as follows:

Theorem 2. *Consider a Stochastic Block Model (SBM) with parameters $(n, 2, (\frac{1}{2}, \frac{1}{2}), (\frac{a}{n}, \frac{b}{n}))$, where $(a + b)$ satisfies the conditions of being at least $\omega(1)$ and $n^{o(1)}$. In such a scenario, the non-backtracking graph neural network can accurately map from graph \mathcal{G} to node labels, with the probability of error decreasing to 0 as $n \rightarrow \infty$.*

In Stephan & Massoulié (2022), the authors demonstrate that the non-backtracking matrix exhibits valuable properties when the average degree of vertices in the graph satisfies $\omega(1)$ and $n^{o(1)}$. When $K = 2$, we define a function $\sigma(v)$ for $v \in \mathcal{V}$, such that $\sigma(v) = 1$ if v is in the first class, and $\sigma(v) = -1$ in the second class. Then, they establish the following proposition:

Proposition 5. (Stephan & Massoulié, 2022) *Suppose we have a SBM with parameters $(n, 2, (\frac{1}{2}, \frac{1}{2}), (\frac{a}{n}, \frac{b}{n}))$. In this case, we have two eigenvalues $\mu_1 > \mu_2$ of $\text{ndiag}(\alpha)P$, and the eigenvector ϕ_2 corresponding to μ_2 , where v -th component is set to $\sigma(v)$. Then, for any n larger than an absolute constant, the eigenvalues λ_1 and λ_2 of the non-backtracking matrix B satisfies $|\lambda_i - \mu_i| = o(1)$, and all other eigenvalues of B are confined in a circle with radius $(1 + o(1))\sqrt{\frac{a+b}{2}}$. Also, there exists an eigenvector ν_2 of the non-backtracking matrix B associated with λ_2 such that*

$$\langle \nu_2, \xi_2 \rangle \geq \sqrt{1 - \frac{8}{(a+b)(a-b)^2}} + o(1) := 1 - f(a, b)$$

where $\xi_2 = \frac{T\Theta\phi_2}{\|T\Theta\phi_2\|}$, $T \in \{0, 1\}^{2m \times n}$, $T_{ei} = 1\{e_2 = i\}$ and $\Theta \in \{0, 1\}^{n \times 2}$, $\Theta_{ij} = 1$ if the vertex i is in the j -th community, and 0 otherwise.

The proposition above highlights that the non-backtracking matrix possesses a spectral separation property, even in the case of very sparse graphs. Moreover, the existence of an eigenvector ν_2 such that $\langle \nu_2, \xi_2 \rangle = 1 - o(1)$ suggests that this eigenvector contains information about the community index of vertices. The foundation for these advantageous properties of the non-backtracking matrix B in sparse scenarios can be attributed to the observation that the matrix B^k shares similarities with the k -hop adjacency matrix, while A^k is predominantly influenced by high-degree vertices. Consequently, we can establish the following theorem:

Lemma 8. *Suppose we have a SBM with parameters defined in Proposition 5. Then, there exists a function of the eigenvectors of the non-backtracking matrix that can accurately recover the original community index of vertices.*

The proof for Lemma 8 can be found in Appendix D.3 In the following, we will demonstrate that the non-backtracking GNN can compute an approximation to the top K eigenvectors mentioned in Lemma 8. To achieve this, we first require the following lemma:

Lemma 9. *Assuming that the non-backtracking matrix B has a set of orthonormal eigenvectors ν_i with corresponding eigenvalues $\lambda_1 > \dots > \lambda_K \geq \lambda_{K+1} \geq \dots \geq \lambda_{2m}$, then there exists a sequence of convolutional layers in the non-backtracking GNNs capable of computing the eigenvectors of the non-backtracking matrix.*

For the proof of Lemma 9, please refer to Appendix D.4. With this lemma in mind, we can observe that a sequence of convolutional layers, followed by a multilayer perceptron, can approximate the function outlined in Lemma 8, leveraging the universal approximation theorem. This argument leads to Theorem 2.

D.2 Proof of Theorem 3

Let’s rephrase the formal version of Theorem 3 as follows:

Theorem 3. *Consider two graphs, one generated from a SBM (\mathcal{G}) with parameters $(n, 2, (\frac{1}{2}, \frac{1}{2}), (\frac{a}{n}, \frac{b}{n}))$ and the other from an Erdős–Rényi model (\mathcal{H}) with parameters $(n, \frac{c}{n})$, for some constants $a, b, c > 1$. When $(a - b)^2 > 2(a + b)$, the non-backtracking graph neural network is capable of distinguishing between \mathcal{G} and \mathcal{H} with probability tending to 1 as $n \rightarrow \infty$.*

To establish Theorem 3, we rely on the following Proposition 6 from Bordenave et al. (2015):

Proposition 6. *Suppose we have two graphs \mathcal{G} and \mathcal{H} as defined in the formal statement of Theorem 3. Then, the eigenvalues $\lambda_i(B)$ of the non-backtracking matrix satisfy the following:*

$$\lambda_1(B_{\mathcal{G}}) = \frac{a+b}{2} + o(1), \lambda_2(B_{\mathcal{G}}) = \frac{a-b}{2} + o(1), \quad \text{and} \quad |\lambda_k(B_{\mathcal{G}})| \leq \sqrt{\frac{a+b}{2}} + o(1) \quad \text{for } k > 2$$

$$\lambda_1(B_{\mathcal{H}}) = c + o(1) \quad \text{and} \quad |\lambda_2(B_{\mathcal{H}})| \leq \sqrt{c} + o(1)$$

Proposition 6 informs us that by examining the distribution of eigenvalues, we can discern whether a graph originates from the Erdős–Rényi model or the SBM. Leveraging Lemma 9, we can obtain the top two normalized eigenvectors of the non-backtracking matrix using convolutional layers, denoted as ν_1 and ν_2 . Applying the non-backtracking convolutional layer to these vectors yields resulting vectors with ℓ_2 -norms corresponding to $\lambda_i(B)$. Consequently, we can distinguish between the two graphs, \mathcal{G} and \mathcal{H} , by examining the output of the convolutional layer in the non-backtracking GNN.

D.3 Proof of Lemma 8

Let us revisit the matrix T , defined as $T_{ei} = 1\{e_2 = i\}$, and its pseudo-inverse denoted as $T^+ = D^{-1}T^\top$, where D is a diagonal matrix containing the degrees of vertices on the diagonal. Considering the definition of ξ_2 as provided in Proposition 5, we can deduce the label of vertex v . Specifically, if $(T^+\xi_2)_v > 0$, it implies that the vertex belongs to the first class; otherwise, it belongs to the other class.

Additionally, we are aware that $|(T^+\nu_2 - T^+\xi_2)_i|_2 = O(f(a, b))$ as indicated in Proposition 5, considering the property that the sum of each row of T^+ is equal to 1. Consequently, by examining the signs of elements in the vector $T^+\nu_2$, we can classify nodes without encountering any misclassified ones.

D.4 Proof of Lemma 9

Proof. Suppose we have $f(\leq 2m)$ arbitrary vectors x_1, \dots, x_f and a matrix $X = [x_1, \dots, x_f] \in \mathbb{R}^{2m \times f}$, which has x_1, \dots, x_f as columns. Without loss of generality, we assume that $\|x_v\|_2 = 1$, and let $x_j = \sum_{i=1}^{2m} c_i^{(j)} \nu_i$

for $1 \leq j \leq f$. We will prove the lemma by showing that if we multiply X by B and repeatedly apply the Gram–Schmidt orthonormalization to the columns of the resulting matrix, the j -th column of the resulting matrix converges towards the direction of ν_j . Further, we will show that there exists a series of convolutional layers equivalent to this process.

First, we need to show $B^k x_1$ converges to the direction of ν_1 . $B^k x_1$ can be expressed as:

$$B^k x_1 = \lambda_1^k \left(c_1^{(1)} \nu_1 + c_2^{(1)} \left(\frac{\lambda_2}{\lambda_1} \right)^k \nu_2 + \dots + c_{2m}^{(1)} \left(\frac{\lambda_{2m}}{\lambda_1} \right)^k \nu_{2m} \right).$$

Let's fix $\epsilon > 0$ and take K_0 such that for all $k \geq K_0$, the inequality $\left(\frac{\lambda_2}{\lambda_1} \right)^k < \frac{\epsilon |c_1^{(1)}|}{4m}$ holds. Furthermore, we define the following for brevity:

$$e_\ell^{(k)} := \begin{cases} \frac{A^k x_1}{\|A^k x_1\|} & \text{for } \ell = 1, \\ \frac{u_\ell^{(k)}}{\|u_\ell^{(k)}\|} & \text{where } u_\ell^{(k)} = GS(B e_\ell^{(k-1)}) \text{ for } \ell \geq 2. \end{cases}$$

Then, the distance between $e_1^{(k)}$ and $\text{sign}(c_1^{(1)}) \nu_1$ is

$$\begin{aligned} \|e_1^{(k)} - \text{sign}(c_1^{(1)}) \nu_1\|_2 &= \left\| \frac{c_1^{(1)} \nu_1 + \sum_{i=2}^{2m} \left(\frac{\lambda_i}{\lambda_1} \right)^k c_i^{(1)} \nu_i}{\sqrt{(c_1^{(1)})^2 + \sum_{i=2}^{2m} \left(\frac{\lambda_i}{\lambda_1} \right)^{2k} (c_i^{(1)})^2}} - \text{sign}(c_1^{(1)}) \nu_1 \right\|_2 \\ &\stackrel{(a)}{<} \left\| \frac{c_1^{(1)} \nu_1 - \text{sign}(c_1^{(1)}) \sqrt{(c_1^{(1)})^2 + \sum_{i=2}^{2m} \left(\frac{\lambda_i}{\lambda_1} \right)^{2k} (c_i^{(1)})^2} \nu_1}{\sqrt{(c_1^{(1)})^2 + \sum_{i=2}^{2m} \left(\frac{\lambda_i}{\lambda_1} \right)^{2k} (c_i^{(1)})^2}} \right\|_2 + \frac{\epsilon}{2} \\ &= \frac{-|c_1^{(1)}| + \sqrt{(c_1^{(1)})^2 + \sum_{i=2}^{2m} \left(\frac{\lambda_i}{\lambda_1} \right)^{2k} (c_i^{(1)})^2}}{\sqrt{(c_1^{(1)})^2 + \sum_{i=2}^{2m} \left(\frac{\lambda_i}{\lambda_1} \right)^{2k} (c_i^{(1)})^2}} + \frac{\epsilon}{2} \\ &\stackrel{(b)}{\leq} \frac{\sqrt{\sum_{i=2}^{2m} \left(\frac{\lambda_i}{\lambda_1} \right)^{2k} (c_i^{(1)})^2}}{\sqrt{(c_1^{(1)})^2 + \sum_{i=2}^{2m} \left(\frac{\lambda_i}{\lambda_1} \right)^{2k} (c_i^{(1)})^2}} + \frac{\epsilon}{2} \\ &\stackrel{(c)}{<} \epsilon, \end{aligned}$$

where (a) and (c) are obtained from the assumption $\left(\frac{\lambda_2}{\lambda_1} \right)^k < \frac{\epsilon |c_1^{(1)}|}{4m}$, and for (b) we use the inequality $\sqrt{x^2 + y^2} \leq |x| + |y|$.

Moreover, we assume that for all $1 \leq \ell < L$ and any $\epsilon_\ell > 0$, there exists K_ℓ such that for all $k \geq K_\ell$, $\|e_\ell^{(k)} - \text{sign}(c_\ell^{(\ell)}) \nu_\ell\| < \epsilon_\ell$. To simplify the notations, we define $K_0 = \max_{1 \leq \ell < L} K_\ell$ and $z_\ell^{(k)} = e_\ell^{(k)} - \text{sign}(c_\ell^{(\ell)}) \nu_\ell$ for all ℓ . Then, we must show there exists K_L such that for all $k \geq K_L$, $\|z_L^{(k)}\| < \epsilon_L$. With no loss of generality, let's assume that $\text{sign}(c_\ell^{(\ell)}) = 1$ for all ℓ . Furthermore, let us fix $\epsilon_L > 0$ and $\epsilon_\ell = \min(1, \epsilon_0)$

for $1 \leq \ell < L$. Then, $u_L^{(k)}$ for $k \geq K_0$ can be written as

$$\begin{aligned} u_L^{(k+1)} &= Be_L^{(k)} - \sum_{i < L} \langle e_i^{(k+1)}, Be_L^{(k)} \rangle e_i^{(k+1)} \\ &= Be_L^{(k)} - \sum_{i < L} \left(\langle \nu_i, Be_L^{(k)} \rangle \nu_i + \langle z_i^{(k+1)}, Be_L^{(k)} \rangle \nu_i + \langle \nu_i + z_i^{(k+1)}, Be_L^{(k)} \rangle z_i^{(k+1)} \right) \\ &\stackrel{(a)}{=} \sum_{i \geq L} \langle \nu_i, Be_L^{(k)} \rangle \nu_i - y_L^{(k)}, \end{aligned}$$

where (a) comes from the definition $y_L^{(k)} := \sum_{i < L} \langle z_i^{(k+1)}, Be_L^{(k)} \rangle \nu_i + \langle \nu_i + z_i^{(k+1)}, Be_L^{(k)} \rangle z_i^{(k+1)}$.

Then, we can deduce $\|y_L^{(k)}\| \leq 3L\epsilon_0 d_{\max} \sqrt{2m}$ by the following lemma.

Lemma 10. *Let $B \in \mathbb{R}^{2m \times 2m}$ be an non-backtracking matrix of the graph \mathcal{G} , and d_{\max} be the maximum degree of vertices in \mathcal{G} . Then, $\|Ax\|_2 < \epsilon_x d_{\max} \sqrt{2m}$ if the ℓ^2 -norm of the vector $x \in \mathbb{R}^{2m}$ is less than ϵ_x .*

Proof.

$$\|Bx\|_2 = \sqrt{\sum_i \left(\sum_j b_{ij} x_j \right)^2} < \sqrt{\sum_i \epsilon_x^2 \left(\sum_j b_{ij} \right)^2} \leq \epsilon_x d_{\max} \sqrt{2m}$$

□

In contrast, for any integer $p > 0$, $u_L^{(k+p)}$ is

$$\begin{aligned} u_L^{(k+p)} &= \sum_{i \geq L} \langle \nu_i, Be_L^{(k+p-1)} \rangle \nu_i - y_L^{(k+p)} \\ &= \frac{1}{\|u_L^{(k+p-1)}\|} \sum_{i \geq L} \left\langle \nu_i, \sum_{j \geq L} \langle \nu_j, Be_L^{(k+p-2)} \rangle \lambda_j \nu_j - By_L^{(k+p-2)} \right\rangle \nu_i - y_L^{(k+p)} \\ &= \frac{1}{\|u_L^{(k+p-1)}\|} \sum_{i \geq L} \left(\lambda_i \langle \nu_i, Be_L^{(k+p-2)} \rangle - \langle \nu_i, By_L^{(k+p-2)} \rangle \right) \nu_i - y_L^{(k+p)} \\ &= \frac{1}{\|u_L^{(k+p-1)}\|} \sum_{i \geq L} \langle \nu_i, Be_L^{(k+p-2)} \rangle \lambda_i \nu_i - \frac{1}{\|u_L^{(k+p-1)}\|} \sum_{i \geq L} \langle \nu_i, By_L^{(k+p-2)} \rangle \nu_i - y_L^{(k+p)} \\ &\vdots \\ &= \frac{1}{\prod_{j=1}^{p-1} \|u_L^{(k+j)}\|} \sum_{i \geq L} \langle \nu_i, Be_L^{(k)} \rangle \lambda_i^{p-1} \nu_i \\ &\quad - \sum_{j=1}^{p-1} \frac{1}{\|u_L^{(k+j)}\|} \left(\sum_{i \geq L} \langle \nu_i, By_L^{(k+j-1)} \rangle \lambda_i^{p-j-1} \nu_i \right) - y_L^{(k+p)} \\ &= \frac{1}{\prod_{j=1}^{p-1} \|u_L^{(k+j)}\|} \sum_{i \geq L} \langle \nu_i, Be_L^{(k)} \rangle \lambda_i^{p-1} \nu_i \\ &\quad - \sum_{j=1}^{p-1} \frac{1}{\|u_L^{(k+j)}\|} \left(\sum_{i \geq L} \langle \nu_i, By_L^{(k+j-1)} \rangle \lambda_i^{p-j-1} \nu_i \right) - y_L^{(k+p)} \\ &\stackrel{(a)}{=} C_0 \lambda_L^{p-1} \left(\langle \nu_L, Be_L^{(k)} \rangle \nu_L + \sum_{i > L} \langle \nu_i, Be_L^{(k)} \rangle \left(\frac{\lambda_i}{\lambda_L} \right)^{p-1} \nu_i \right) \\ &\quad - \sum_{i \geq L} \sum_{j=1}^{p-1} C_j \left(\langle \nu_i, By_L^{(k+j-1)} \rangle \lambda_i^{p-j-1} \nu_i \right) - y_L^{(k+p)}, \end{aligned}$$

where (a) stems from the definitions $C_0 = \frac{1}{\prod_{j=1}^{p-1} \|u_L^{(k+j)}\|}$ and $C_j = \frac{1}{\|u_L^{(k+j)}\|}$.

Now, define $\hat{e}_L^{(k+p)} := \frac{u_L^{(k+p)}}{C_0 \lambda_L^{p-1} \langle \nu_L, Be_L^{(k)} \rangle}$. Then,

$$\hat{e}_L^{(k+p)} = \nu_L + \sum_{i \geq L} \frac{\langle \nu_i, Ae_L^{(k)} \rangle}{\langle \nu_L, Be_L^{(k)} \rangle} \left(\frac{\lambda_i}{\lambda_L} \right)^{p-1} \nu_i - \frac{\sum_{i \geq L} \sum_{j=1}^{p-1} C_j \left(\langle \nu_i, By_L^{(k+j-1)} \rangle \lambda_i^{p-j-1} \nu_i \right) - y_L^{(k+p)}}{C_0 \lambda_L^{p-1} \langle \nu_L, Be_L^{(k)} \rangle}. \quad (15)$$

Take p_0 such that $\left(\frac{\lambda_i}{\lambda_L} \right)^{p_0-1} \leq \frac{\epsilon \langle \nu_L, Be_L^{(k)} \rangle}{4(L-N) \langle \nu_i, Be_L^{(k)} \rangle}$, then, the ℓ^2 -norm of the second term of Equation (15) is less than $\epsilon_L/4$.

Meanwhile, the ℓ^2 -norm of the third term in Equation (15) is upper-bounded by

$$\frac{3L\epsilon_0 d_{\max} \sqrt{2m}}{C_0 \lambda_L^{p-1} \langle \nu_L, Be_L^{(k)} \rangle} \left(C d_{\max} \sqrt{2m} \sum_{i \geq L} \frac{\lambda_i^{p-2} - 1}{\lambda_i - 1} + 1 \right),$$

where $C := \max_{j=1}^{p-1} C_j$.

Therefore, if we take $\epsilon_0 < \frac{\epsilon C_0 \lambda_L^{p-1} \langle \nu_L, Be_L^{(k)} \rangle}{12L d_{\max} \sqrt{2m}} \left(C d_{\max} \sqrt{2m} \sum_{i \geq L} \frac{\lambda_i^{p-2} - 1}{\lambda_i - 1} + 1 \right)^{-1}$, we can get $\|\hat{e}_L^{(k+p)} - \nu_L\| < \epsilon_L/2$ for $p > p_0$. Finally, the upper bound of $\|z_L^{(k)}\|$ is

$$\begin{aligned} \|z_L^{(k)}\| &\leq \|e_L^{(k+p)} - \hat{e}_L^{(k+p)}\| + \|\hat{e}_L^{(k+p)} - \nu_L\| \\ &\stackrel{(a)}{<} \left| \|\hat{e}_L^{(k+p)}\| - 1 \right| + \frac{\epsilon_L}{2} \\ &\stackrel{(b)}{\leq} \epsilon_L, \end{aligned}$$

where (a) follows from $e_L^{(k+p)} = \frac{\hat{e}_L^{(k+p)}}{\|\hat{e}_L^{(k+p)}\|}$, and (b) is obtained from the inequality $1 - \epsilon_L \leq \|\hat{e}_L^{(k+p)}\| - 1 \leq 1 + \epsilon_L$. \square

E Experiment Details

In this section, we provide details of NBA-GNN implementation and experiments.

E.1 Implementation

E.1.1 Message Initialization and Aggregation

For message initialization and final aggregation of messages, we have proposed functions ϕ, σ, ρ . To be specific, the message initialization can be written as following:

$$h_{i \rightarrow j}^{(0)} = \begin{cases} \phi(e_{ij}, x_i, x_j) & (\text{if } e_{ij} \text{ exists}) \\ \phi(x_i, x_j) & (\text{otherwise}) \end{cases}.$$

For our experiments, we use concatenation for ϕ , weighted sums for σ , and average for ρ .

E.1.2 Non-backtracking Updates

Here, we provide the details of using the non-backtracking operator in our NBA-GNNs. To begin, let's revisit the non-backtracking matrix $B \in \{0, 1\}^{2|\mathcal{E}| \times 2|\mathcal{E}|}$ from Section 4.1 defined as follows:

$$B_{(\ell \rightarrow k), (j \rightarrow i)} = \begin{cases} 1 & \text{if } k = j, \ell \neq i \\ 0 & \text{otherwise} \end{cases}.$$

Returning to our NBA-GNN, the non-backtracking message passing update for a hidden feature $h_{j \rightarrow i}^{(t)}$ at the $(t+1)$ -th layer, introduced in Section 3.2, is expressed as follows:

$$h_{j \rightarrow i}^{(t+1)} = \phi^{(t)} \left(h_{j \rightarrow i}^{(t)}, \left\{ \psi^{(t)} \left(h_{k \rightarrow j}^{(t)}, h_{j \rightarrow i}^{(t)} \right) : k \in \mathcal{N}(j) \setminus \{i\} \right\} \right), \quad (3)$$

where $\phi^{(t)}$ and $\psi^{(t)}$ are backbone-specific non-linear update and permutation-invariant aggregation functions at the t -th layer, respectively. Using the non-backtracking matrix B , we can rewrite Equation (3) as following:

$$H^{(t+1)} = \phi^{(t)} \left(H^{(t)}, \psi^{(t)} \left(B^\top H^{(t)}, H^{(t)} \right) \right),$$

where $H^{(t)} \in \mathbb{R}^{2|\mathcal{E}| \times d}$ are edge-wise features, i.e., each row representing $h_{j \rightarrow i}^{(t)}$.

E.1.3 Implementation Example

Now, the NBA-GNN implementation example in Section 3 using GCN (Kipf & Welling, 2017) as a backbone, can be written as the following:

$$h_{j \rightarrow i}^{(t+1)} = h_{j \rightarrow i}^{(t)} + \frac{1}{|\mathcal{N}(j)| - 1} \mathbf{W}^{(t)} \sum_{k \in \mathcal{N}(j) \setminus \{i\}} h_{k \rightarrow j}^{(t)}.$$

Recall the out-degree and in-degree matrices of NBA-GNN, denoted as D_{out} and D_{in} , where $(D_{out})_{(j \rightarrow i), (j \rightarrow i)} = \sum_{\ell \rightarrow k} B_{(j \rightarrow i), (\ell \rightarrow k)}$ and $(D_{in})_{(j \rightarrow i), (j \rightarrow i)} = \sum_{\ell \rightarrow k} B_{(\ell \rightarrow k), (j \rightarrow i)}$, for each index $j \rightarrow i$. We also introduced the normalized non-backtracking matrix augmented with self-loops as $\hat{B} = (D_{out} + I)^{-\frac{1}{2}} (B + I) (D_{in} + I)^{-\frac{1}{2}}$. In summary, Equation (5) can be expressed as follows:

$$H^{(t+1)} = \hat{B} H^{(t)} \mathbf{W}^{(t)}.$$

Hence, the message passing update in NBA-GCN is equivalent to the multiplication of non-backtracking operator and edge-wise features.

E.2 Long-range Graph Benchmark

E.2.1 Dataset Statistics

From LRGB (Dwivedi et al., 2022), we experiment for 3 tasks: graph classification (**Peptides-func**), graph regression (**Peptides-struct**), and node classification (**PascalVOC-SP**). We provide the dataset statistics in Table 9. Note that for **PascalVOC-SP**, we use SLIC compactness of 30 and edge weights are based only on super-pixels coordinates.

Table 9: Statistics of LRGB datasets

Dataset	Total Graphs	Total Nodes	Avg Nodes	Mean Deg.	Total Edges	Avg Edges	Avg Short. Path.	Avg Diameter
PascalVOC-SP	11,355	5,443,545	479.40	8.00	43,548,360	3,835.17	8.05 ± 0.18	19.40 ± 0.65
Peptides-func	15,535	2,344,859	150.94	2.04	4,773,974	307.30	20.89 ± 9.79	56.99 ± 28.72
Peptides-struct	15,535	2,344,859	150.94	2.04	4,773,974	307.30	20.89 ± 9.79	56.99 ± 28.72

E.2.2 Experiments Details

All experiment results are averaged over three runs and trained for 300 epochs, with a $\sim 500k$ parameter budget. Baseline scores were taken from each paper.

- All experiments are averaged over three seeds, 0~2.
- Baseline numbers were took from the lrgb benchmak (Dwivedi et al., 2022) and table1 of DRew(Gutteridge et al., 2023).
- We use an AdamW optimizer(Loshchilov & Hutter, 2018) with lr decay=0.1 , min lr=1e-5, momentum=0.9, and base learning rate lr=0.001 (0.0005 for **PascalVOC-SP**).
- We use cosine scheduler with reduce factor=0.5 , schedule patience=10 with 50 warm-up.
- Laplacian positional encoding were used with hidden dimension 16 and 2 layers.
- We use the "Atom Encoder", "Bond Encoder" for **Peptides-func**, **Peptides-struct** from based on OGB molecular feature (Hu et al., 2020; 2021), and the "VOCNode Encoder", "VOCEdge Encoder" for **PascalVOC-SP**.
- **PascalVOC-SP** results in Figure 4a all use the same setup described above.
- **Peptides-func** results in Figures 4b and 4c all use the same setup described above.
- GCN results in Figures 4a to 4c use the hyperparameters from (Tönshoff et al., 2023).

We searched the following range of hyperparameters, and report the best in Table 10.

- We searched layers 6 to 12 for **PascalVOC-SP** 2, layers 5 to 20 for **Peptides-func** and **Peptides-struct**.
- The hidden dimension was chosen by the maximum number in parameter budget.
- Dropout was searched from 0.0~0.8 for **PascalVOC-SP** in steps of 0.1, and 0.1~0.4 in steps of 0.1 for **Peptides-func** and **Peptides-struct**.
- We used the batch size of 30 for **PascalVOC-SP** on GPU memory, and 200 for **Peptides-func** and **Peptides-struct**.

Table 10: Best hyperparameters for each GNN, dataset in LRGB benchmark

Model	Dataset	# Param.	# Layers	hidden dim.	dropout	Batch size	# epochs
GCN	PascalVOC-SP	472k	12	180	0.7	30	200
	Peptides-func	510k	10	186	0.1	200	300
	Peptides-struct	505k	20	144	0.1	200	300
GINE	PascalVOC-SP	472k	12	180	0.7	30	200
	Peptides-func	502k	10	144	0.1	200	300
	Peptides-struct	503k	10	144	0.1	200	300
GatedGCN	PascalVOC-SP	486k	10	96	0.25	30	200
	Peptides-func	511k	10	96	0.1	200	300
	Peptides-struct	511k	8	108	0.1	200	300

E.3 Transductive Node Classification

We conducted experiments involving three citation networks (Cora, CiteSeer, and Pubmed in (Sen et al., 2008)), and three heterophilic datasets (Texas, Wisconsin, and Cornell in (Pei et al., 2019)), focusing on transductive node classification. Our reported results in Table 3 are the averages obtained from 10 different seed runs to ensure robustness and reliability.

For the citation networks, we employed the dataset splitting procedure outlined in (Yang et al., 2016). In contrast, for the heterophilic datasets, we randomly divided the nodes of each class into training (60%), validation (20%), and testing (20%) sets.

During the model training process, we utilized the AdamW optimizer (Loshchilov & Hutter, 2018) with a learning rate of $3e-5$. The training duration spanned 1,000 epochs for citation networks and 100 epochs for heterophilic datasets. Following training, we selected the best epoch based on validation accuracy for evaluation on the test dataset. The model’s hidden dimension and dropout ratio were set to 512 and 0.2, respectively, consistent across all datasets, after fine-tuning these hyperparameters on the Cora dataset. Additionally, we conducted optimization for the number of convolutional layers within the set $\{1, 2, 3, 4, 5\}$. The results revealed that the optimal number of layers is typically three for most of the models and datasets. However, there are exceptions, such as CiteSeer-GatedGCN, PubMed-{GraphSAGE, GraphSAGE+NBA+PE}, Texas-GatedGCN+NBA, Wisconsin-{GraphSAGE+NBA, GraphSAGE+NBA+PE} and Cornell-{GraphSAGE, GraphSAGE+NBA, GraphSAGE+NBA+PE, GatedGCN+NBA}, where the optimal number of layers is found to be four. Furthermore, for Cora-{GraphSAGE+NBA+PE, GAT+NBA}, CiteSeer-GraphSAGE+NBA, and Cornell-GatedGCN+NBA+PE, the optimal number of layers is determined to be five.

Table 11: Statistics of the datasets for the transductive node classification task

Dataset	Total Graphs	Num Nodes	Num Edges	Dim Features	Num Classes
Cora	1	2,708	10,556	1,433	7
CiteSeer	1	3,327	9,104	3,703	6
PubMed	1	19,717	88,648	500	3
Texas	1	183	309	1,703	5
Wisconsin	1	251	499	1,703	5
Cornell	1	183	295	1,703	5

For reference, we have provided a summary of dataset statistics in Table 11.