
Return of ChebNet: Understanding and Improving an Overlooked GNN on Long-Range Tasks

Ali Hariri^{1,*,\dagger} Álvaro Arroyo^{2,*} Alessio Gravina^{3,*} Moshe Eliasof⁴
Carola-Bibiane Schönlieb⁴ Davide Bacciu³ Kamyar Azizzadenesheli⁵
Xiaowen Dong² Pierre Vandergheynst¹

Abstract

ChebNet, one of the earliest spectral GNNs, has largely been overshadowed by Message Passing Neural Networks (MPNNs), which gained popularity for their simplicity and effectiveness in capturing local graph structure. Despite their success, MPNNs are limited in their ability to capture long-range dependencies between nodes. This has led researchers to adapt MPNNs through *rewiring* or make use of *Graph Transformers*, which compromises the computational efficiency that characterized early spatial message-passing architectures, and typically disregards the graph structure. Almost a decade after its original introduction, we revisit ChebNet to shed light on its ability to model distant node interactions. We find that out-of-box, ChebNet already shows competitive advantages relative to classical MPNNs and GTs on long-range benchmarks, while maintaining good scalability properties for high-order polynomials. However, we uncover that this polynomial expansion leads ChebNet to an unstable regime during training. To address this limitation, we cast ChebNet as a stable and non-dissipative dynamical system, which we coin *Stable-ChebNet*. Our *Stable-ChebNet* model allows for stable information propagation, and has controllable dynamics which do not require the use of eigendecompositions, positional encodings, or graph rewiring. Across several benchmarks, *Stable-ChebNet* achieves near state-of-the-art performance.

1 Introduction

Graph Neural Networks (GNNs) [81, 43, 74, 67, 14, 25, 44] have emerged as a prevalent framework for handling data defined on graphs. Graph convolutional networks have their roots in spectral approaches that extend convolutional filters to non-Euclidean domains. The first practical instantiation of a GNN was proposed by [14], which leveraged the eigenbasis of the graph Laplacian to perform spectral filtering, as an attempt to generalize image convolutions to non-euclidean structures. However, this formulation required costly eigen-decompositions at each layer. Defferrard *et al.* [25] addressed this inefficiency by approximating spectral filters with truncated Chebyshev polynomials, giving rise to *ChebNet*, the first tractable and localized spectral GNN. By parameterizing filters as K -order polynomials of the Laplacian, ChebNet could aggregate information from K -hop neighborhoods without repeated eigendecompositions, thereby enabling scalable spectral convolution on large graphs.

¹École Polytechnique Fédérale de Lausanne (EPFL)

²University of Oxford

³University of Pisa

⁴University of Cambridge

⁵NVIDIA Research

*Equal contribution

[†]Correspondence to: Ali Hariri (ali.hariri@epfl.ch)

In 2017, Kipf and Welling distilled ChebNet into a simpler, first-order approximation now known as the *Graph Convolutional Network (GCN)* by (i) restricting the polynomial order to one and (ii) tying filter coefficients across hops [57]. This yielded an architecture that was both lightweight and effective: a GCN with few layers would achieve strong node-classification performance on standard homophilic benchmarks, and its $\mathcal{O}(|E|)$ complexity made it practical for large-scale graphs. GCN’s efficiency and strong locality bias quickly made it the default baseline, and subsequent *message-passing neural networks (MPNNs)* adopted a similar paradigm of iterative neighborhood aggregation [41].

Despite their popularity, MPNNs exhibit pronounced shortcomings when made deeper and when capturing long-range dependencies [33]. Repeated neighborhood aggregations tend to cause *representational collapse*, where node features become indistinguishable (often referred to as “oversmoothing”) [15, 69], and information from distant nodes is “squashed” through narrow bottlenecks, limiting the ability to model global context [1]. In recent years, researchers have proposed various strategies to overcome these limitations. To mitigate oversmoothing, several models have drawn

on principles from physics [28, 12] to preserve feature diversity across layers. At the same time, efforts to capture long-range dependencies have led to graph-rewiring techniques [82, 49, 7] that add or reweight edges to shorten information pathways, as well as the emergence of graph transformers [30], which replace purely local aggregation with global self-attention mechanisms. Although these advances can alleviate depth-related pathologies, they often trade off numerous benefits that made MPNNs appealing: scalability, parameter efficiency, and to only process information along the graph’s edges. In this context, ChebNet and other spectral GNNs are usually relegated to a footnote-mentioned only as a predecessor to GCN, which is typically not revisited as a competitive baseline.

In this work, we revisit ChebNet from first principles. We demonstrate that the original ChebNet (without any rewiring or attention mechanisms) already delivers state-of-the-art performance or is close on long-range graph tasks while scaling gracefully to large graphs. By deriving and analyzing ChebNet’s linearized dynamics, we prove that enlarging its receptive field introduces signal-propagation instabilities. To overcome this, we propose *Stable-ChebNet*, a minimal set of architectural modifications that restore stable propagation for arbitrarily large receptive fields, supported by both theoretical guarantees and empirical validation. To build intuition for our *Stable-ChebNet* framework, Figure 1 illustrates how classical ChebNet filters (top) can exhibit unbounded dynamics, whereas our antisymmetric, forward-Euler discretization yields smooth, stable propagation (bottom). An intuitive heat-transfer analogy helps explain the difference: if we inject ‘heat’ at seed nodes, the high-order filters of vanilla ChebNet diffuse and dissipate this heat so that distant nodes cool rapidly. In contrast, the non-dissipative dynamics induced by the antisymmetric, forward-Euler step in *Stable-ChebNet* preserve energy, keeping temperatures higher at nodes many hops away.

Across a suite of challenging long-range node- and graph-level benchmarks, *Stable-ChebNet* matches or outperforms state-of-the-art message-passing neural networks and graph transformers, while retaining the ChebNet backbone. We hope this work will reignite interest in spectral GNNs as a scalable, theoretically grounded alternative for long-range graph modeling.

Contributions and Outline

- In Section 3.1, we empirically demonstrate that vanilla ChebNet can achieve very strong performance on long-range benchmarks without incurring prohibitive computational cost.
- In Section 3.2, we analyze ChebNet’s signal-propagation dynamics, providing exact sensitivity analysis, and theoretically and empirically prove the emergence of instability for large filter order.




Figure 1: **Top:** Vanilla ChebNet. **Bottom:** Stable-ChebNet. While the original ChebNet’s high-order Chebyshev filters induce unstable dynamics resulting in dissipative behavior, our Stable-ChebNet yields bounded propagation through layers.

- In Section 3.3, we introduce Stable-ChebNet, which enforces layer-wise stability.
- In Section 4, we empirically validate that Stable-ChebNet consistently outperforms MPNNs, rewiring methods, and graph transformers across a number of tasks.

2 Background

2.1 Background on Spectral Graph Neural Networks

Spectral GNNs extend the notion of convolution to graphs by leveraging the eigen-decomposition of the graph Laplacian. Given an undirected graph $G = (V, E)$ with normalized Laplacian $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$, any graph signal $\mathbf{X} \in \mathbb{R}^n$ can be filtered in the spectral domain via $\mathbf{Y} = \mathbf{U}g_{\theta}(\mathbf{\Lambda})\mathbf{U}^{\top}\mathbf{X}$, where $\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{\top}$ diagonalizes the Laplacian, $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$ its eigenvalues, and g_{θ} is a learnable spectral response. Early methods directly parameterize $g_{\theta}(\mathbf{\Lambda})$, but require an expensive eigendecomposition of the Laplacian, which can be computationally and memory intensive [14]. ChebNet alleviates the cost of an explicit eigendecomposition by approximating g_{θ} using the recurrence relation for a K -th order Chebyshev polynomial in \mathbf{L} [25]. The latter defines g_{θ} as $g_{\theta}(\mathbf{\Lambda}) \approx \sum_{k=0}^K \Theta_k T_k(\tilde{\mathbf{L}})$, where $T_k(\tilde{\mathbf{L}})$ is the k -th polynomial of $\tilde{\mathbf{L}}$ with $\tilde{\mathbf{L}} = \frac{2\mathbf{L}}{\lambda_{max}} - \mathbf{I}_n$. The spectral convolution can then be written without any eigendecomposition as the truncated expansion:

$$\mathbf{Y} = \sum_{k=0}^K \Theta_k T_k(\tilde{\mathbf{L}}) \mathbf{X} \quad (1)$$

where $\tilde{\mathbf{L}} = \frac{2\mathbf{L}}{\lambda_{max}} - \mathbf{I}_n$ enabling efficient, localized filtering in $\mathcal{O}(K|E|)$ time.

2.2 MPNNs and their Limitations

Message-Passing Neural Networks (MPNNs) define a general framework in which node features are iteratively updated by exchanging “messages” along edges. At each layer l , every node v aggregates information from its neighbors $u \in \mathcal{N}(v)$ and combines it with its own representation. This formulation unifies many graph models, including graph convolutional networks (GCNs) [57] and graph attention networks (GATs) [83]. While this local neighborhood aggregation captures structural information effectively, it has a limited capacity to model long-range interactions within the graph. This is due to the phenomenon of *over-squashing*, an information bottleneck that impedes effective information flow among distant nodes [1, 82, 27]. Numerous techniques have emerged to address this limitation such as graph rewiring [49, 7], Graph Transformers [71, 80, 79] in addition to some enhanced spatial methods that tackle over-squashing through combined local and global information [75, 40], or through non-dissipativity achieved by antisymmetric weight parameterization [45, 46] or port-Hamiltonian systems [54].

However, some of the aforementioned methods suffer from substantial overhead due to denser graph shift operators or the use of all-pairs interactions. Specifically, [79] and other graph transformers increase computational complexity through dense attention-maps; [40] relies on costly full eigendecomposition operations; and graph rewiring techniques heavily pre-process the graph topology, incurring $\mathcal{O}(n^3)$ time in the case of [49]. For a detailed discussion on the relevant literature, we point the reader to the Appendix A.

2.3 The Connection between ChebNet and GCN

Although often interpreted as disparate models, GCN [57] can be derived as a special case of ChebNet [25] by truncating the Chebyshev expansion to $K = 1$ and making additional simplifications such as approximating $\tilde{\mathbf{L}} \approx \mathbf{I} - \mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$ and adding self-loops to improve numerical stability. These choices introduced a strong locality bias, which aligned with widely used homophilic benchmarks and significantly reduced computational costs. Due to some of these strengths, GCN has become the de facto backbone for many modern GNNs, and has led to ChebNet being somewhat forgotten, under the assumption that it does not perform well and will not scale well in mid- to large-size graphs due to its spectral nature. For instance, until recently, ChebNet was almost never included in popular GNN benchmarks, such as [31], or in those designed to evaluate long-range dependencies, such as those in the Long Range Graph Benchmark

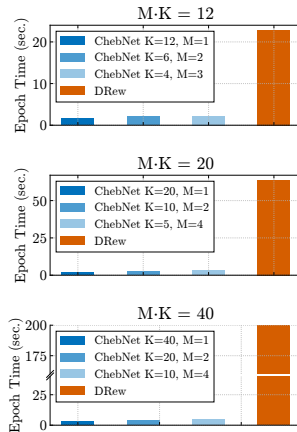


Figure 2: Epoch times for DReW and ChebNet for different receptive fields on the peptides-func task. M is the number of layers, and K is the number of filters.

(LRGB) [33] and the numerous studies that leveraged this dataset to assess the effectiveness of graph rewiring, positional encodings, or Transformer-based architectures. ChebNet was similarly disregarded from large-scale graph evaluations such as the OGB benchmarks [55], presumably due to prevailing assumptions about its scalability.

3 Analyzing and Improving ChebNet from First Principles

3.1 The Effectiveness and Scalability of Vanilla ChebNet

In this subsection, we perform two high-level empirical tests to challenge the commonplace assumption that spectral GNNs inherently suffer from poor performance and limited scalability. We do so by firstly testing ChebNet [25] on a long-range test on the Ring Transfer dataset from [27], which has become a de facto benchmark for state-of-the-art methods seeking to model long-distance dependencies on graphs. Furthermore, to test scalability, we compare epoch training times on the `peptides-func` dataset from [33] with respect to state-of-the-art rewired MPNNs [49] based on a GCN backbone. We compare ChebNet with different numbers of filters K and layers M_{Cheb} with an M_{MPNN} -layer MPNN. To ensure a fair comparison, we ensure that $M_{\text{Cheb}}K = M_{\text{MPNN}}$ so that the two methods would have the same receptive field. The results are shown in Figures 2 and 3.

As seen from Figure 3, ChebNet is capable of performing long-range retrieval on the ring transfer dataset for rings of up to 50 nodes, which is a substantial improvement over a regular GCN. While this finding aligns with intuition, as we will formalize in the following sections, it may




Figure 3: Test accuracy on RingTransfer.

nevertheless surprise practitioners, since ChebNet is not typically included as a benchmark in long-range studies such as that in [30]. On the other hand, as seen in Figure 2, ChebNet scales gracefully on standard tasks such as the `peptides-func` dataset. When compared to DRew, a state-of-the-art method that leverages both dynamic rewiring and delay to propagate information through a GCN (or other MPNN) backbone, we observe that DRew’s epoch times are almost up to two orders of magnitude larger than ChebNet’s. We note further that this inefficiency is not unique to this particular baseline: Graph Transformers incur quadratic scaling in the number of nodes and, by effectively discarding the underlying graph structure, trade-off inductive bias for increased compute, while many rewiring approaches depend on cubic-time algorithms (e.g., Floyd–Warshall or eigendecompositions). Together, these examples underscore how numerous contemporary techniques intended to overcome traditional message-passing limitations actually erode computational benefits, whereas ChebNet, a more natural spectral baseline that generalizes GCN, delivers both training speed and strong performance out-of-the-box.

3.2 Signal Propagation Analysis of ChebNet

In this subsection, we conduct a sensitivity analysis of ChebNet via the spectral norm of the Jacobian of node features, providing an exact characterization of ChebNet’s information flow through different layers and between pairs of nodes, in the spirit of [3]. Specifically, we begin by analyzing the layer-wise Jacobian for Spectral GNNs that use polynomial filters. In this setting, we demonstrate that the layer-wise Jacobian becomes unstable as the polynomial order K increases. Lastly, we investigate the sensitivity of node pairs when using ChebNet. We provide the proofs for the statements in Appendix B.

Lemma 3.1 (Layer-Wise Jacobian for a Spectral GNN). *Consider a linear spectral GNN whose layer-wise update is performed through the following polynomial filter $f(\mathbf{X}) = \sum_{k=1}^K T_k(\mathbf{L}) \mathbf{X} \Theta_k$, where $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the node feature matrix, $T_k(\mathbf{L}) \in \mathbb{R}^{n \times n}$ is the k -th polynomial of the Laplacian $\mathbf{L} \in \mathbb{R}^{n \times n}$, and $\Theta_k \in \mathbb{R}^{d \times d'}$ are learnable weight matrices. Then, the vectorized Jacobian $\mathbf{J} = \partial \text{vec}(f(\mathbf{X})) / \partial \text{vec}(\mathbf{X})$ is*

$$\mathbf{J} = \sum_{k=1}^K \Theta_k^\top \otimes T_k(\mathbf{L}). \tag{2}$$

Given the layer-wise Jacobian from Lemma 3.1, we proceed to analyze the dynamics of a Spectral GNN in Theorem 1 below. Specifically, we focus on the case where the polynomial filter can be approximated by powers of the Laplacian, i.e., $T_k(\mathbf{L}) = \mathbf{L}^k$.




Figure 4: Singular-value spectra of the graph-wise Jacobian in the complex plane for vanilla ChebNet with increasing polynomial order K .

Theorem 1 (Layer-Wise Jacobian singular-value distribution). *Assume the setting of Lemma 3.1, with $T_k(\mathbf{L}) = \mathbf{L}^k$ and \mathbf{L} the symmetric normalized Laplacian, and let all $\Theta_k \in \mathbb{R}^{d \times d}$ be initialized with i.i.d. $\mathcal{N}(0, \sigma^2)$ entries. Denote the eigenvalues of \mathbf{L} as $\{\lambda_1, \dots, \lambda_n\}$, the squared singular values of $\Theta_k \Theta_k^T$ as $\{\mu_{1,k}, \dots, \mu_{d,k}\}$, and the squared singular values of the Jacobian by $\gamma_{i,j}$. Then, for sufficiently large d the empirical eigenvalue distribution of $\Theta_k \Theta_k^T$ converges to the Marchenko-Pastur distribution $\forall k$. Then, the mean and variance of each $\gamma_{i,j}$ are*

$$\mathbb{E}[\gamma_{i,j}] = \sigma^2 \sum_{k=1}^K \lambda_i^{2k}, \quad (3)$$

$$\text{Var}[\gamma_{i,j}] = \sigma^4 \left(\sum_{k=1}^K \lambda_i^{2k} \right)^2. \quad (4)$$

Theorem 1 shows that the singular values spectrum of the layer-wise Jacobian depends on the sum over the powers of the normalized Laplacian’s eigenvalues. This indicates that larger polynomial orders K push the singular value-spectrum towards unstable dynamics, as empirically demonstrated in Figure 4. Stacking several layers of large filter orders will therefore severely hinder the trainability of a Spectral GNN of this form.

Beyond analyzing the information propagation dynamics between different layers, we are interested in quantifying the communication ability between distant nodes in the graph. Recent literature has proposed to measure information flow in the graph by evaluating the sensitivity of a node embedding after l layers (i.e., hops of propagation) with respect to the input of another node using the node-wise Jacobian [82, 27, 3], i.e., $\partial \mathbf{x}_u^{(l)} / \partial \mathbf{x}_v^{(0)}$. Following this approach, we measure how sensitive a node embedding of ChebNet at an arbitrary layer l with respect to the initial features of another node, to better illustrate the long-range propagation capabilities of spectral GNNs.

Theorem 2 (ChebNet Sensitivity). *Consider a Chebyshev-based Graph Neural Network (ChebNet) defined as:*

$$\mathbf{X}^{(l+1)} = \sum_{k=0}^K T_k(\mathbf{L}) \mathbf{X}^{(l)} \mathbf{W}_k^{(l)}, \quad (5)$$

where $\mathbf{L} \in \mathbb{R}^{n \times n}$ is the graph Laplacian, $T_k(\mathbf{L})$ is the k -th Chebyshev polynomial of the Laplacian, and $\mathbf{W}_k^{(l)}$ are learnable weight matrices. Assume activation function σ is identity and let $\mathbf{X}^{(0)}$ be the input features.

Then, the sensitivity of node v with respect to node u after l layers is given by:

$$\frac{\partial \mathbf{x}_v^{(l)}}{\partial \mathbf{x}_u^{(0)}} = \left(\prod_{l=0}^{l-1} \left(\sum_{k=0}^K T_k(\mathbf{L}) \mathbf{W}_k^{(l)} \right) \right)_{v,u}. \quad (6)$$

This result indicates that the sensitivity of ChebNet is closely tied to the polynomial order K . We note that, in the case of $K = 1$, the sensitivity aligns with that of standard MPNNs¹ (such as GCN), exhibiting reduced long-range communication. However, for $K > 1$, the higher polynomial orders significantly enhance the ChebNet’s sensitivity, enabling more effective long-range propagation and improving the overall capacity to capture distant dependencies in the graph. Therefore, we conclude that a large filter order K is needed to enable long-range communication between nodes in the graph, but this will result in unstable training dynamics. This serves to explain the decay in performance in Figure 3. In the next subsection, we will propose a remedy to this issue.

¹The sensitivity upper bound of standard MPNNs is further discussed in Appendix B.5.

3.3 Stable-ChebNet: Stability with Antisymmetric Parameterization

As discussed in Section 3.2, although ChebNet demonstrates strong long-range propagation capabilities, increasing the polynomial order can introduce significant instability into the model dynamics. To address this challenge, we propose Stable-ChebNet, a simple yet effective modification of classical ChebNet aimed at improving its stability. Recent literature has demonstrated that the effectiveness of neural architectures can be significantly improved by framing them as stable, non-dissipative dynamical systems [50, 18, 45, 46, 54]. The core idea behind these approaches is to carefully regulate the spectrum of the Jacobian matrix to ensure the network operates within a stable regime. Specifically, this behavior can be achieved by constraining the eigenvalues of the Jacobian to be purely imaginary. Under this constraint, the input graph information is effectively propagated through the successive transformations into the final nodes’ representation. Motivated by this line of work, we begin by reformulating ChebNet as a continuous-time differential equation. Specifically, we consider the following ordinary differential equation (ODE):

$$\frac{d\mathbf{X}(t)}{dt} = \sum_{k=0}^K T_k(\mathbf{L})\mathbf{X}(t)\mathbf{W}_k \quad (7)$$

for time $t \in [0, T]$ and subject to the initial condition (i.e., the input features) $\mathbf{X}(0) = \mathbf{X}^{(0)}$. In other words, the dynamics of the system (i.e., the continuous flow of information over the graph) is now described as the ChebNet update rule. To ensure the Jacobian of this system has purely imaginary eigenvalues, a straightforward approach is to use antisymmetric weight matrices² and the symmetrically normalized Laplacian. This choice, as formalized in the following theorem, directly enforces the desired spectral property, leading to inherently stable dynamics.

Theorem 3 (Purely Imaginary Eigenvalues). *Let \mathbf{L} be the symmetric normalized Laplacian and $-\mathbf{W}_k = \mathbf{W}_k^\top$ $\forall k = 0, \dots, K$, then the graph-wise Jacobian of the ODE in Equation (7) has purely imaginary eigenvalues, i.e.,*

$$Re(\lambda_i(\mathbf{J})) = 0, \forall i. \quad (8)$$

Even in this section, we provide the proofs for the statements in Appendix B. Theorem 3 shows that by enforcing antisymmetry in the weight matrices and leveraging the symmetric structure of the Laplacian, we guarantee that the Jacobian has purely imaginary eigenvalues, ensuring that node representations remain sensitive to input features of far away nodes without suffering from the instability of standard ChebNet.

As for standard differential-equation-inspired neural architectures, a numerical discretization method is needed to solve Equation (7). We solve the equation with a simple finite difference scheme, i.e., forward Euler’s method, yielding the following node update equation

$$\mathbf{X}^{(l+1)} = \mathbf{X}^{(l)} + \epsilon \left(\sum_{k=0}^K T_k(\mathbf{L})\mathbf{X}^{(l)}(\mathbf{W}_k - \mathbf{W}_k^\top - \gamma\mathbf{I}) \right) \quad (9)$$

where \mathbf{I} is the identity matrix, $\gamma \in \mathbb{R}$ is a hyper-parameter that maintains the stability of the forward Euler method, and $\epsilon \in \mathbb{R}_+$ is the discretization step.

We refer to the ODE in Equation (9) as Stable-ChebNet, and in the following, we show that this new formulation achieves second-order stability. This is a critical improvement, as a naive Euler discretization of the original ChebNet (without imposing constraints on the Jacobian eigenvalues) would result in only first-order stability, which is considerably more prone to numerical instability. Specifically, without these constraints, the model can exhibit exponential growth or decay in the gradients, significantly limiting its ability to capture long-range dependencies [3]. In contrast, second-order stable systems, like Stable-ChebNet, maintain controlled gradient dynamics over longer timescales, allowing for effective long-range information propagation.

²A matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ is antisymmetric (i.e., skew-symmetric) if $-\mathbf{A} = \mathbf{A}^\top$

Table 1: Mean test set $\log_{10}(\text{MSE})$ and standard deviation averaged over 4 random weight initializations for each configuration on the Graph Property Prediction dataset. The lower the better.

Model	Diameter	SSSP	Eccentricity
GCN	0.7424 ± 0.0466	0.9499 ± 0.0001	0.8468 ± 0.0028
GAT	0.8221 ± 0.0752	0.6951 ± 0.1499	0.7909 ± 0.0222
GraphSAGE	0.8645 ± 0.0401	0.2863 ± 0.1843	0.7863 ± 0.0207
GIN	0.6131 ± 0.0990	-0.5408 ± 0.4193	0.9504 ± 0.0007
GCNII	0.5287 ± 0.0570	-1.1329 ± 0.0135	0.7640 ± 0.0355
DGC	0.6028 ± 0.0050	-0.1483 ± 0.0231	0.8261 ± 0.0032
GRAND	0.6715 ± 0.0490	-0.0942 ± 0.3897	0.6602 ± 0.1393
A-DGN w/ GCN backbone	0.2271 ± 0.0804	-1.8288 ± 0.0607	0.7177 ± 0.0345
ChebNet	-0.1517 ± 0.0343	-1.8519 ± 0.0539	-1.2151 ± 0.0852
Stable-ChebNet (ours)	-0.2477 ± 0.0526	-2.2111 ± 0.0160	-2.1043 ± 0.0766

Theorem 4 (Non-exponential Information Growth or Decay with Antisymmetric Weights). *Consider the Stable-type Chebyshev Graph Neural Network (Stable-ChebNet) defined by:*

$$\mathbf{X}^{(l+1)} = \mathbf{X}^{(l)} + \epsilon \sum_{k=0}^K T_k(\mathbf{L}) \mathbf{X}^{(l)} \mathbf{W}_k^{(l)}, \quad (10)$$

with small step size $\epsilon > 0$ and antisymmetric weight matrices:

$$(\mathbf{W}_k^{(l)})^\top = -\mathbf{W}_k^{(l)}, \quad \forall k, l. \quad (11)$$

Then the Jacobian $J^{(l)}$ of the layer does not lead to exponential growth or decay across layers. Specifically, we have:

$$\|\mathbf{J}^{(l)}\|_2 = 1 + O(\epsilon^2). \quad (12)$$

Conversely, for general weights without the antisymmetric property, exponential growth or decay of the Jacobian norm typically occurs.

4 Experiments

We evaluate ChebNet and its stable formulation (**Stable-ChebNet**) across a variety of settings to thoroughly assess long-range capabilities. In this section, we report and discuss the performance of both models on these benchmarks. We report additional experiments on heterophilic node classification tasks from [70] in Appendix F. We run our experiments on a single A100 GPU and provide the full details on the hyperparameter search for all datasets in Appendix D, and baseline and datasets details in Appendix C.

Graph Property Prediction Dataset. We evaluate our model’s ability to predict long-range graph properties using a synthetic dataset developed by Corso et.al in [24] under the experimental setup of [45]. The dataset consists of undirected graphs drawn from a diverse set of random and structured families (Erdős–Rényi, Barabási–Albert, caterpillar, etc), ensuring a broad coverage of topological properties. Each graph contains between 25 and 35 nodes as per the setup of Gravina et.al in [45] (in contrast to the 15–25 node range originally used by [24]), thus increasing task complexity and raising the need for long-range information propagation. Finally, each node is assigned a single scalar feature sampled uniformly at random from the interval $[0, 1]$. We provide a detailed comparison in Table 1.

Results. Compared to classical ChebNet, Stable-ChebNet yields consistent and significant gains across all three tasks. On the *Diameter* task, classical ChebNet achieves a $\log_{10}(\text{MSE})$ of -0.15 , whereas Stable-ChebNet improves it to -0.25 . On the *Single Source Shortest Path (SSSP)* task, the gain is larger and goes from -1.85 using ChebNet to -2.21 with our Stable-ChebNet form. Finally, on *Eccentricity*, where the difficulty is highest due to the necessity to propagate information about the most distant nodes individually, classical ChebNet achieves a $\log_{10}(\text{MSE})$ of -1.22 while Stable-ChebNet reaches -2.10 , reducing the average prediction error by more than an order of magnitude relative to the baseline. Relative to other baselines, Stable-ChebNet dominates most models based on standard message-passing or modified diffusion mechanisms. Methods like GCN and GCNII are badly over-squashed, barely reaching negative $\log_{10}(\text{MSE})$ values.

Over-Squashing Analysis on Barbell Graphs. To further investigate the robustness of our model to oversquashing, we use as a benchmark the barbell regression tasks introduced in [5]. In this task, a model that fails to transfer any information across the single bridge edge will produce an essentially random constant and obtain a mean-squared error (MSE) close to 1; an error in the 0.4–0.6 band indicates that only a partial amount of information has overcome the bottleneck. Errors around ≈ 0.25 and below suggest that the oversquashing has been effectively overcome. In this work, we compare Stable-ChebNet’s performance on barbell graphs of varying sizes ($N = 10, 25, 50, 100$) against numerous baselines, mainly an MLP and variants of MPNNs such as GCN [57], GAT [83], and SAGE [51]. Further description of the task is found in Appendix C.2.

Results. We observe in Figure 5 that both a classical ChebNet and Stable-ChebNet successfully learn the small $N = 10$ case with negligible error. However, for moderate graph sizes with $N = 50$, a classical ChebNet with fixed $K = 8$ already sits in the "partial collapse" regime as its MSE increases to around 0.90 and slides towards the random-guess area as N keeps growing (Table 2). For the same range of hops K , replacing the standard update with our stable Euler-based formulation keeps the error almost two orders of magnitude smaller with an MSE below 0.20, confirming that the non-dissipative time-stepping effectively prevents the over-squashing phenomenon.




Figure 5: Mean Squared Error (MSE) comparison of various MPNN baselines at different node counts N of Barbell graphs.

Table 2: Mean squared error (MSE) of ChebNet and Stable-ChebNet on the over-squashing experiment for barbell graphs. Left: sizes $N = 50, 70$ for $K = 9$ and 10 . Right: size $N = 100$ for $K = 20$.

Method	K	50	70	Method	K	100
ChebNet	$K = 9$	0.32 ± 0.39	1.08 ± 0.05	ChebNet	$K = 20$	0.87 ± 0.05
	$K = 10$	0.05 ± 0.00	1.08 ± 0.01		Stable-ChebNet (ours)	$K = 20$
Stable-ChebNet	$K = 9$	0.17 ± 0.11	0.47 ± 0.49			
	$K = 10$	0.05 ± 0.00	0.06 ± 0.03			

Open-Graph Benchmark. To evaluate real-world applicability on large-scale graphs, we run experiments for node-level tasks on two large-scale graph datasets from the Open-Graph Benchmark (OGB) [55]. `ogbn-arxiv` is a citation network in which each node corresponds to an academic paper, and the task is node classification by predicting the subject area of unseen papers. The other dataset we use, `ogbn-proteins`, is a protein–protein interaction network aimed at inferring protein functions. To ensure a fair comparison and emphasize efficiency, we limit the number of parameters in our models to be within the same range as those used in existing and most recent OGB benchmarks.

Results. Table 3 reports the performance on the `ogbn-arxiv` citation network, where ChebNet achieves around 73% test accuracy, while Stable-ChebNet boosts the performance further to 75.7%, outperforming all other methods, including a variety of MPNNs and Graph Transformers such as GraphGPS [71] and Exphormer [80]. Similarly, on the `ogbn-proteins` interaction network, ChebNet attains about 77.6% accuracy compared to Stable-ChebNet’s 79.5% (Table 4). Competing approaches achieve nearly 72% for MPNN-based methods, while Transformer-based models’ performances range from 77.4% for NodeFormer [88] to 79.5% for SGFormer [89]. Hence, Stable-ChebNet remarkably competes and often outperforms state-of-the-art models on this benchmark, demonstrating that the Euler formulation consistently narrows the gap with and in some cases overtakes Transformer baselines such as SGFormer [89] and Spexphormer [79]. Together, these results demonstrate that augmenting ChebNet with an Euler step not only addresses the classical ChebNet’s shortcomings on long-range information propagation but also performs effectively well on graphs with hundreds of thousands of nodes, in contrast to regular-sized graphs seen in previous experiments.

Table 3: Accuracy on ogbn-arxiv.

Model	ogbn-arxiv
GCN	71.74 \pm 0.29
ChebNet	73.27 \pm 0.23
ChebNetII	72.32 \pm 0.23
GraphSAGE	71.49 \pm 0.27
GAT	72.01 \pm 0.20
NodeFormer	59.90 \pm 0.42
GraphGPS	70.92 \pm 0.04
GOAT	72.41 \pm 0.40
EXPHORMER+GCN	72.44 \pm 0.28
SPEXPHORMER	70.82 \pm 0.24
Stable-ChebNet (ours)	75.73 \pm 0.51

Table 4: Accuracy on ogbn-proteins.

Model	ogbn-proteins
MLP	72.04 \pm 0.48
GCN	72.51 \pm 0.35
ChebNet	77.55 \pm 0.43
SGC	70.31 \pm 0.23
GCN-NSAMPLER	73.51 \pm 1.31
GAT-NSAMPLER	74.63 \pm 1.24
SIGN	71.24 \pm 0.46
NodeFormer	77.45 \pm 1.15
SGFormer	79.53 \pm 0.38
SPEXPHORMER	80.65 \pm 0.07
Stable-ChebNet (ours)	79.55 \pm 0.34

Long-Range Graph Benchmark (LRGB). LRGB [33] is a collection of GNN benchmarks that evaluate models on tasks involving long-range interactions. We use two of its molecular-property datasets: Peptides-func for graph classification and Peptides-struct for graph regression.

Results. A detailed leaderboard is shown in Table 5. It can be seen that Stable-ChebNet improves upon its vanilla counterpart. Together with S2GCN, it reaches an average precision (AP) above 70 on Peptide-func and a Mean Absolute Error (MAE) below 0.26 on the regression task, bearing in mind that S2GCN requires a more expensive full Laplacian eigendecomposition. Overall, our model achieves competitive performance on peptide structures with results competing with and often outperforming some well-known graph-based models including graph transformers such as Exphormer [80] and GraphViT [53], state space models including Graph Mamba and GMN [84], and rewiring methods like DRew [49]. It is worth noting that the gain in AP for DRew comes at the cost of computing positional encodings (Laplacian eigenvectors) for every graph before training, while Stable-ChebNet does not use any positional encodings.

Heterophilic benchmarks We further assess Stable-ChebNet on node-classification tasks explicitly designed to stress performance under heterophily, following the standardized protocol of Platonov et al. (“Roman-empire”, “Amazon-ratings”, “Minesweeper” and “Tolokers”) [70]. We keep the exact data processing, splits, and metrics recommended therein. Concretely, we report accuracy on Roman-empire and Amazon-ratings, and ROC-AUC on Minesweeper and Tolokers averaging over four random initializations as in the protocol.

Link to oversmoothing, heterophily, and long-rangeness. Our heterophilic results (Table 9) should not be over-interpreted as “evidence of long-range propagation” or as a direct antidote to oversmoothing. The recent position paper by Arnaiz-Rodríguez & Errica [2] argues that several widespread assumptions in the literature are often conflated: (i) that heterophily is inherently detrimental while homophily is beneficial, (ii) that long-range propagation is best evaluated on heterophilic graphs, and (iii) that performance degradation mainly arises from oversmoothing. They show that heterophily, long-range interactions, and oversmoothing are orthogonal factors: a graph may be heterophilic yet dominated by local dependencies, or homophilic yet require long-range reasoning. Hence, evaluations should focus on the nature of the learning task, not merely on global homophily ratios. In this light, our Stable-ChebNet scores on Roman-empire, Minesweeper, and Tolokers demonstrate that a stable spectral propagator has competitive performance on standardized heterophily benchmarks, but they do not by themselves certify long-rangeness. Those conclusions are better supported by our dedicated long-range tests and stability analysis.

5 Conclusion

In this work, we have re-examined ChebNet, one of the earliest spectral GNNs, from first principles, uncovering its innate ability to capture long-range dependencies via higher-order polynomial filters, but also its susceptibility to unstable propagation dynamics as the polynomial order grows. By casting ChebNet as a continuous-time ODE and imposing antisymmetric weight constraints, we introduced

Table 5: Long-range benchmark results. AP is the target metric on *peptides-func* (higher is better), and MAE is the target metric on *peptides-struct* (lower is better).

Model Type	Model	peptides-func (AP \uparrow)	peptides-struct (MAE \downarrow)
<i>Transformer</i>	SAN+LapPE	63.84 \pm 1.21	0.2683 \pm 0.0043
	TIGT	66.79 \pm 0.74	0.2485 \pm 0.0015
	Specformer	66.86 \pm 0.64	0.2550 \pm 0.0014
	Expformer	65.27 \pm 0.43	0.2481 \pm 0.0007
	G.MLPMixer	69.21 \pm 0.54	0.2475 \pm 0.0015
	Graph ViT	69.42 \pm 0.75	0.2449 \pm 0.0016
	GRIT	69.88 \pm 0.82	0.2460 \pm 0.0012
<i>Rewiring</i>	LASER	64.40 \pm 0.10	0.3043 \pm 0.0019
	DRew-GCN	69.96 \pm 0.76	0.2781 \pm 0.0028
	+PE	71.50 \pm 0.44	0.2536 \pm 0.0015
<i>State Space</i>	Graph Mamba	67.39 \pm 0.87	0.2478 \pm 0.0016
	GMN	70.71 \pm 0.83	0.2473 \pm 0.0025
	MP-SSM	69.93 \pm 0.52	0.2458 \pm 0.0017
<i>GNN</i>	A-DGN	59.75 \pm 0.44	0.2874 \pm 0.0021
	ChebNet	69.61 \pm 0.33	0.2627 \pm 0.0033
	ChebNetII	68.19 \pm 0.27	0.2618 \pm 0.0058
	GCN	68.60 \pm 0.50	0.2460 \pm 0.0007
	GRAMA	70.93 \pm 0.78	0.2436 \pm 0.0022
	GRAND	57.89 \pm 0.62	0.3418 \pm 0.0015
	GraphCON	60.22 \pm 0.68	0.2778 \pm 0.0018
	PH-DGN	70.12 \pm 0.45	0.2465 \pm 0.0020
	SWAN	67.51 \pm 0.39	0.2485 \pm 0.0009
	PathNN	68.16 \pm 0.26	0.2545 \pm 0.0032
	CIN++	65.69 \pm 1.17	0.2523 \pm 0.0013
	S2GCN	72.75 \pm 0.66	0.2467 \pm 0.0019
	+PE	73.11 \pm 0.66	0.2447 \pm 0.0032
	Stable-ChebNet (ours)	70.32 \pm 0.26	0.2542 \pm 0.0030

Stable-ChebNet, whose forward Euler discretization yields non-dissipative, second-order–stable information flow without resorting to costly eigendecompositions, positional encodings, or graph rewiring. We provide a theoretical analysis of Stable-ChebNet showing purely imaginary Jacobian spectra and bounded layerwise sensitivity. We support the analysis with extensive experiments on a variety of synthetic and other long-range graph benchmarks. Stable-ChebNet consistently matches or outperforms state-of-the-art message-passing, rewiring, state-space, and transformer-based models, while retaining the fundamental properties of Chebyshev filters. Future work can focus on generalizing this ODE-based framework to broader spectral GNN families, whereby the stability analysis can be extended to other polynomial types and orthogonal bases.

Impact Statement. This work aims to advance the field of machine learning on graph-structured data which are abundant in the real world. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

Code Availability All code and datasets will be made publicly available at <https://github.com/ahariri13/Stable-ChebNet> to support reproducibility and future work.

Acknowledgments and Disclosure of Funding

AG and DB acknowledge funding from EU-EIC EMERGE (Grant No. 101070918). XD acknowledges support from the Oxford-Man Institute of Quantitative Finance and EPSRC No. EP/T023333/1.

References

- [1] Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. In *International Conference on Learning Representations*, 2021.
- [2] Adrian Arnaiz-Rodriguez and Federico Errica. Oversmoothing, "oversquashing", heterophily, long-range, and more: Demystifying common beliefs in graph machine learning. *arXiv preprint arXiv:2505.15547*, 2025.
- [3] Álvaro Arroyo, Alessio Gravina, Benjamin Gutteridge, Federico Barbero, Claudio Gallicchio, Xiaowen Dong, Michael Bronstein, and Pierre Vandergheynst. On vanishing gradients, over-smoothing, and over-squashing in gnn: Bridging recurrent and graph learning. *arXiv preprint arXiv:2502.10818*, 2025.
- [4] Hugo Attali, Davide Buscaldi, and Nathalie Pernelle. Rewiring techniques to mitigate over-squashing and oversmoothing in gnn: A survey. *arXiv preprint arXiv:2411.17429*, 2024.
- [5] Jacob Bamberger, Federico Barbero, Xiaowen Dong, and Michael Bronstein. Bundle neural networks for message diffusion on graphs. *arXiv preprint arXiv:2405.15540*, 2024.
- [6] Jacob Bamberger, Benjamin Gutteridge, Scott le Roux, Michael M Bronstein, and Xiaowen Dong. On measuring long-range interactions in graph neural networks. *arXiv preprint arXiv:2506.05971*, 2025.
- [7] Federico Barbero, Ameya Velingker, Amin Saberi, Michael M. Bronstein, and Francesco Di Giovanni. Locality-aware graph rewiring in GNNs. In *The Twelfth International Conference on Learning Representations*, 2024.
- [8] Ali Behrouz and Farnoosh Hashemi. Graph mamba: Towards learning on graphs with state space models. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 119–130, 2024.
- [9] Mitchell Black, Zhengchao Wan, Amir Nayyeri, and Yusu Wang. Understanding oversquashing in gnn through the lens of effective resistance. In *International Conference on Machine Learning*, pages 2528–2547. PMLR, 2023.
- [10] Deyu Bo, Chuan Shi, Lele Wang, and Renjie Liao. Specformer: Spectral graph neural networks meet transformers. *arXiv preprint arXiv:2303.01028*, 2023.
- [11] Deyu Bo, Xiao Wang, Chuan Shi, and Huawei Shen. Beyond low-frequency information in graph convolutional networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(5):3950–3957, May 2021.
- [12] Cristian Bodnar, Francesco Di Giovanni, Benjamin P. Chamberlain, Pietro Liò, and Michael M. Bronstein. Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in GNNs, 2022.
- [13] Xavier Bresson and Thomas Laurent. Residual Gated Graph ConvNets. *arXiv preprint arXiv:1711.07553*, 2018.
- [14] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- [15] Chen Cai and Yusu Wang. A note on over-smoothing for graph neural networks. *arXiv preprint arXiv:2006.13318*, 2020.
- [16] Andrea Ceni, Alessio Gravina, Claudio Gallicchio, Davide Bacciu, Carola-Bibiane Schonlieb, and Moshe Eliasof. Message-Passing State-Space Models: Improving Graph Learning with Modern Sequence Modeling. *arXiv preprint arXiv:2505.18728*, 2025.
- [17] Benjamin Paul Chamberlain, James Rowbottom, Maria Gorinova, Stefan Webb, Emanuele Rossi, and Michael M Bronstein. GRAND: Graph neural diffusion. In *International Conference on Machine Learning (ICML)*, pages 1407–1418. PMLR, 2021.

- [18] Bo Chang, Minmin Chen, Eldad Haber, and Ed H. Chi. AntisymmetricRNN: A dynamical system view on recurrent neural networks. In *International Conference on Learning Representations*, 2019.
- [19] Jinsong Chen, Kaiyuan Gao, Gaichao Li, and Kun He. Nagphormer: A tokenized graph transformer for node classification in large graphs. In *Proceedings of the International Conference on Learning Representations*, 2023.
- [20] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. Simple and Deep Graph Convolutional Networks. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1725–1735. PMLR, 13–18 Jul 2020.
- [21] Eli Chien, Jianhao Peng, Pan Li, and Olgica Milenkovic. Adaptive universal generalized pagerank graph neural network. In *International Conference on Learning Representations*, 2021.
- [22] Yun Young Choi, Sun Woo Park, Minhoo Lee, and Youngho Woo. Topology-informed graph transformer. *arXiv preprint arXiv:2402.02005*, 2025.
- [23] Krzysztof Choromanski, Marcin Kuczynski, Jacek Cieszkowski, Paul L. Beletsky, Konrad M. Smith, Wojciech Gajewski, Gabriel De Masson, Tomasz Z. Broniatowski, Antonina B. Gorny, Leszek M. Kaczmarek, and Stanislaw K. Andrzejewski. Performers: A new approach to scaling transformers. *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 2020–2031, 2020.
- [24] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 13260–13271, 2020.
- [25] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems*, 29, 2016.
- [26] Chenhui Deng, Zichao Yue, and Zhiru Zhang. Polynormer: Polynomial-expressive graph transformer in linear time. In *The Twelfth International Conference on Learning Representations*, 2024.
- [27] 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*, pages 7865–7885. PMLR, 2023.
- [28] Francesco Di Giovanni, James Rowbottom, Benjamin P Chamberlain, Thomas Markovich, and Michael M Bronstein. Graph neural networks as gradient flows: understanding graph convolutions via energy. *arXiv preprint arXiv:2206.10991*, 2022.
- [29] Lun Du, Xiaozhou Shi, Qiang Fu, Xiaojun Ma, Hengyu Liu, Shi Han, and Dongmei Zhang. Gbk-gnn: Gated bi-kernel graph neural networks for modeling both homophily and heterophily. In *Proceedings of the ACM Web Conference 2022, WWW '22*, page 1550–1558, New York, NY, USA, 2022. Association for Computing Machinery.
- [30] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv preprint arXiv:2012.09699*, 2020.
- [31] Vijay Prakash Dwivedi, Chaitanya K. Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *J. Mach. Learn. Res.*, 24(1), January 2023.
- [32] Vijay Prakash Dwivedi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Graph neural networks with learnable structural and positional representations. In *International Conference on Learning Representations*, 2022.

- [33] Vijay Prakash Dwivedi, Ladislav Rampásek, 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.
- [34] Moshe Eliasof, Alessio Gravina, Andrea Ceni, Claudio Gallicchio, Davide Bacciu, and Carola-Bibiane Schönlieb. Graph Adaptive Autoregressive Moving Average Models. In *Forty-second International Conference on Machine Learning*, 2025.
- [35] Federico Errica, Henrik Christiansen, Viktor Zaverkin, Takashi Maruyama, Mathias Niepert, and Francesco Alesiani. Adaptive message passing: A general framework to mitigate oversmoothing, oversquashing, and underreaching. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 15490–15515. PMLR, 13–19 Jul 2025.
- [36] Lukas Fesser and Melanie Weber. Mitigating over-smoothing and over-squashing using augmentations of forman-ricci curvature. In Soledad Villar and Benjamin Chamberlain, editors, *Proceedings of the Second Learning on Graphs Conference*, volume 231 of *Proceedings of Machine Learning Research*, pages 19:1–19:28. PMLR, 27–30 Nov 2024.
- [37] Ben Finkelshtein, Xingyue Huang, Michael M. Bronstein, and Ismail Ilkan Ceylan. Cooperative Graph Neural Networks. In *Forty-first International Conference on Machine Learning*, 2024.
- [38] Fabrizio Frasca, Emanuele Rossi, Davide Eynard, Ben Chamberlain, Michael Bronstein, and Federico Monti. Sign: Scalable inception graph neural networks. *arXiv preprint arXiv:2004.11198*, 2020.
- [39] Johannes Gasteiger, Stefan Weiß enberger, and Stephan Günnemann. Diffusion Improves Graph Learning. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [40] Simon Markus Geisler, Arthur Kosmala, Daniel Herbst, and Stephan Günnemann. Spatio-spectral graph neural networks. *Advances in Neural Information Processing Systems*, 37:49022–49080, 2024.
- [41] 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*, pages 1263–1272. PMLR, 2017.
- [42] Lorenzo Giusti, Teodora Reu, Francesco Ceccarelli, Cristian Bodnar, and Pietro Liò. Cin++: Enhancing topological message passing. *arXiv preprint arXiv:2306.03561*, 2023.
- [43] M Gori, G Monfardini, and F Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pages 729–734. IEEE, 2005.
- [44] Alessio Gravina and Davide Bacciu. Deep Learning for Dynamic Graphs: Models and Benchmarks. *IEEE Transactions on Neural Networks and Learning Systems*, 35(9):11788–11801, 2024.
- [45] Alessio Gravina, Davide Bacciu, and Claudio Gallicchio. Anti-Symmetric DGN: a stable architecture for Deep Graph Networks. In *The Eleventh International Conference on Learning Representations*, 2023.
- [46] Alessio Gravina, Moshe Eliasof, Claudio Gallicchio, Davide Bacciu, and Carola-Bibiane Schönlieb. On oversquashing in graph neural networks through the lens of dynamical systems. In *The 39th Annual AAAI Conference on Artificial Intelligence*, 2025.
- [47] Alessio Gravina, Claudio Gallicchio, and Davide Bacciu. Non-dissipative Propagation by Randomized Anti-symmetric Deep Graph Networks. In *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, pages 25–36, Cham, 2025. Springer Nature Switzerland.

- [48] Alessio Gravina, Giulio Lovisotto, Claudio Gallicchio, Davide Bacciu, and Claas Grohnfeldt. Long Range Propagation on Continuous-Time Dynamic Graphs. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 16206–16225. PMLR, 21–27 Jul 2024.
- [49] Benjamin Gutteridge, Xiaowen Dong, Michael M Bronstein, and Francesco Di Giovanni. Drew: Dynamically rewired message passing with delay. In *International Conference on Machine Learning*, pages 12252–12267. PMLR, 2023.
- [50] E. Haber and L. Ruthotto. Stable architectures for deep neural networks. *Inverse Problems*, 34(1), 2017.
- [51] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017.
- [52] Mingguo He, Zhewei Wei, and Ji-Rong Wen. Convolutional neural networks on graphs with chebyshev approximation, revisited. *Advances in Neural Information Processing Systems*, 35:7264–7276, 2022.
- [53] 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*, pages 12724–12745. PMLR, 2023.
- [54] Simon Heilig, Alessio Gravina, Alessandro Trenta, Claudio Gallicchio, and Davide Bacciu. Port-Hamiltonian Architectural Bias for Long-Range Propagation in Deep Graph Networks. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [55] 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. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS ’20. Curran Associates Inc., 2020.
- [56] Kedar Karhadkar, Pradeep Kr. Banerjee, and Guido Montúfar. Fosr: First-order spectral rewiring for addressing oversquashing in gnns. In *The Eleventh International Conference on Learning Representations, ICLR*, 2023.
- [57] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations (ICLR)*, 2017.
- [58] Kezhi Kong, Jiu-hai Chen, John Kirchenbauer, Renkun Ni, C. Bayan Bruss, and Tom Goldstein. GOAT: A global transformer on large-scale graphs. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 17375–17390. PMLR, 23–29 Jul 2023.
- [59] 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.
- [60] Xiang Li, Renyu Zhu, Yao Cheng, Caihua Shan, Siqiang Luo, Dongsheng Li, and Weining Qian. Finding global homophily in graph neural networks when meeting heterophily. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 13242–13256. PMLR, 17–23 Jul 2022.
- [61] Huidong Liang, Haitz Sáez de Ocariz Borde, Baskaran Sripathanathan, Michael Bronstein, and Xiaowen Dong. Towards quantifying long-range interactions in graph machine learning: a large graph dataset and a measurement. *arXiv preprint arXiv:2503.09008*, 2025.
- [62] Sitao Luan, Chenqing Hua, Qincheng Lu, Liheng Ma, Lirong Wu, Xinyu Wang, Minkai Xu, Xiao-Wen Chang, Doina Precup, Rex Ying, Stan Z. Li, Jian Tang, Guy Wolf, and Stefanie Jegelka. The heterophilic graph learning handbook: Benchmarks, models, theoretical analysis, applications and challenges, 2024.

- [63] Liheng Ma, Chen Lin, Derek Lim, Adriana Romero-Soriano, Puneet K. Dokania, Mark Coates, Philip Torr, and Ser-Nam Lim. Graph inductive biases in transformers without message passing. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 23321–23337. PMLR, 23–29 Jul 2023.
- [64] Sohir Maskey, Raffaele Paolino, Aras Bacho, and Gitta Kutyniok. A fractional graph laplacian approach to oversmoothing. *Advances in Neural Information Processing Systems*, 36, 2024.
- [65] Sunil Kumar Maurya, Xin Liu, and Tsuyoshi Murata. Simplifying approach to node classification in graph neural networks. *Journal of Computational Science*, 62:101695, 2022.
- [66] Gaspard Michel, Giannis Nikolentzos, Johannes Lutzeyer, and Michalis Vazirgiannis. Path neural networks: expressive and accurate graph neural networks. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org, 2023.
- [67] Alessio Micheli. Neural Network for Graphs: A Contextual Constructive Approach. *IEEE Transactions on Neural Networks*, 20(3):498–511, 2009.
- [68] Khang Nguyen, Hieu Nong, Vinh Nguyen, Nhat Ho, Stanley Osher, and Tan Nguyen. Revisiting over-smoothing and over-squashing using ollivier-ricci curvature. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org, 2023.
- [69] Kenta Oono and Taiji Suzuki. Graph neural networks exponentially lose expressive power for node classification. In *International Conference on Learning Representations*, 2020.
- [70] Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? In *The Eleventh International Conference on Learning Representations*, 2023.
- [71] Ladislav Rampásek, 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.
- [72] T. Konstantin Rusch, Michael M. Bronstein, and Siddhartha Mishra. A Survey on Oversmoothing in Graph Neural Networks. *arXiv preprint arXiv:2303.10993*, 2023.
- [73] T Konstantin Rusch, Ben Chamberlain, James Rowbottom, Siddhartha Mishra, and Michael Bronstein. Graph-coupled oscillator networks. In *International Conference on Machine Learning*, pages 18888–18909. PMLR, 2022.
- [74] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.
- [75] Florian Sestak, Lisa Schneckenreiter, Johannes Brandstetter, Sepp Hochreiter, Andreas Mayr, and Günter Klambauer. Vn-egnn: E (3)-equivariant graph neural networks with virtual nodes enhance protein binding site identification. *arXiv preprint arXiv:2404.07194*, 2024.
- [76] Ahsan Shehzad, Feng Xia, Shagufta Abid, Ciyuan Peng, Shuo Yu, Dongyu Zhang, and Karin Verspoor. Graph transformers: A survey, 2024.
- [77] Dai Shi, Andi Han, Lequan Lin, Yi Guo, and Junbin Gao. Exposition on over-squashing problem on GNNs: Current Methods, Benchmarks and Challenges, 2023.
- [78] Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjing Wang, and Yu Sun. Masked label prediction: Unified message passing model for semi-supervised classification. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 1548–1554. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.
- [79] Hamed Shirzad, Honghao Lin, Balaji Venkatachalam, Ameya Velingker, David P Woodruff, and Danica J Sutherland. Even sparser graph transformers. *Advances in Neural Information Processing Systems*, 37:71277–71305, 2024.

- [80] Hamed Shirzad, Ameya Velingker, Balaji Venkatachalam, Danica J Sutherland, and Ali Kemal Sinop. Exphormer: Sparse transformers for graphs. In *International Conference on Machine Learning*, pages 31613–31632. PMLR, 2023.
- [81] Alessandro Sperduti. Encoding labeled graphs by labeling raam. *Advances in Neural Information Processing Systems*, 6, 1993.
- [82] 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.
- [83] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [84] Chloe Wang, Oleksii Tsepa, Jun Ma, and Bo Wang. Graph-mamba: Towards long-range graph sequence modeling with selective state spaces. *arXiv preprint arXiv:2402.00789*, 2024.
- [85] Xiyuan Wang and Muhan Zhang. How powerful are spectral graph neural networks. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 23341–23362. PMLR, 17–23 Jul 2022.
- [86] Yifei Wang, Yisen Wang, Jiansheng Yang, and Zhouchen Lin. Dissecting the Diffusion Process in Linear Graph Convolutional Networks. In *Advances in Neural Information Processing Systems*, volume 34, pages 5758–5769. Curran Associates, Inc., 2021.
- [87] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying Graph Convolutional Networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6861–6871. PMLR, 09–15 Jun 2019.
- [88] Qitian Wu, Wentao Zhao, Zenan Li, David P Wipf, and Junchi Yan. Nodeformer: A scalable graph structure learning transformer for node classification. *Advances in Neural Information Processing Systems*, 35:27387–27401, 2022.
- [89] Qitian Wu, Wentao Zhao, Chenxiao Yang, Hengrui Zhang, Fan Nie, Haitian Jiang, Yatao Bian, and Junchi Yan. Sgformer: Simplifying and empowering transformers for large-graph representations. *Advances in Neural Information Processing Systems*, 36:64753–64773, 2023.
- [90] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations (ICLR)*, 2019.
- [91] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform bad for graph representation? In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS ’21. Curran Associates Inc., 2021.
- [92] Manzil Zaheer, Guru prasad G. H., Lihong Wang, S. V. K. N. L. Wang, Yujia Li, Jakub Konečný, Shalmali Joshi, Danqi Chen, Jennifer R. R., Zhenyu Zhang, Shalini Devaraj, and Srinivas Narayanan. Bigbird: Transformers for longer sequences. *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 12168–12178, 2020.
- [93] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931*, 2019.
- [94] Jiong Zhu, Ryan A. Rossi, Anup Rao, Tung Mai, Nedim Lipka, Nesreen K. Ahmed, and Danai Koutra. Graph neural networks with heterophily. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12):11168–11176, May 2021.
- [95] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Beyond homophily in graph neural networks: Current limitations and effective designs. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7793–7804. Curran Associates, Inc., 2020.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: All the claims found in the abstract and introduction are supported by the results and theorems discussed in the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We mention the scope of this work that is limited to ChebNet and leave the expansion to other Spectral GNN methods for future work.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The paper provides a set of theorems that justify and motivate the choice of the developed model. The corresponding proofs are also presented and can be found in the appendices.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The datasets we use are publicly available and we share the code in the supplementary materials folder. The set of swept hyperparameters to reproduce the results is also shared.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The provided repository includes all the datasets used as well as the necessary code to reproduce the experimental results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We specify that we follow the same training procedure for the datasets as their papers of reference. We also show our hyperparameter sweep in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide all our results in terms of graph performance in the format “mean \pm standard error of the mean”.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide detailed information about the used resources in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We preserve anonymity in our submission and our submission abides by the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss potential positive and negative societal impacts after the conclusion section.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The release of our data and models does not pose a direct risk of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We provide citations for all sources of code and/or data used, both in the paper and in the accompanying repository. For the datasets, we reference the original open-source work they are based on, in full accordance with their licensing terms.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release new datasets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No crowdsourcing or human subjects were involved in the experiments conducted for this paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No crowdsourcing or human subjects were involved in the experiments conducted for this paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLMs are not a core component of this paper in terms of methods.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Supplementary Related Work

Effective propagation and preservation of information on graphs remains a central challenge in deep learning on graphs, especially when long-range communication between nodes becomes fundamental for the downstream task [77]. GNNs usually rely on local neighborhood aggregation, which limits their capacity to capture interactions between distant nodes [1, 27] due to challenges such as over-smoothing [15, 69, 72] and over-squashing [1, 82, 27], which are linked to the problem of vanishing gradients [3]. Several techniques have been proposed to address this issue.

Graph rewiring techniques [39, 1, 4] modify the original topology, usually as a pre-processing step, with the aim of directly connect distant nodes and facilitate information flow. Rewiring methods can be broadly classified based on the type of information they leverage, including curvature metrics [1, 68, 36], effective resistance [9], random walks [7], spectral gap [56], and node features [49].

Similarly, Graph Transformers [76, 59, 91, 71] enable direct message passing between any pair of nodes via attention mechanisms, employing an attention mechanism on the entire graph. These models often incorporate positional encodings, such as Laplacian eigenvectors [30] or random-walk structural encodings (RWSE) [32], to encode graph structure. To reduce the quadratic complexity of the full attention mechanism, recent methods introduce methods such as sparse attention mechanisms [92, 23], Exphormer [80], and linear graph transformers [26].

Despite the success of graph rewiring methods and Graph Transformers, these approaches often introduce additional computational complexity due to the use of denser graph shift operators. An alternative strategy to enhance long-range propagation focuses on constraining the dynamics of the GNN to remain stable and non-dissipative, while maintaining the computational complexity of classical MPNNs. In this paradigm, the GNN is interpreted as a discretization of a differential equation, leveraging dynamical systems theory to maintain a constant rate of information flow between nodes. This behavior has been achieved either through antisymmetric weight parameterization [45, 46, 47, 48] or by exploiting port-Hamiltonian dynamics [54].

Recent works further explore long-range interactions in GNNs. In [6], Bamberger et al. formalize and measure interaction ranges in graph operators, while in [61] Liang et al. introduce a large-scale dataset and metric to quantify long-range dependencies. These studies complement existing approaches by providing tools to better evaluate and understand information propagation across distant nodes.

Other approaches include filtering messages in the information flow [37, 35], using a graph adaptive method based on a learnable ARMA framework [34], state space models [8, 16], or fractional power of the graph shift operator [64].

B Sensitivity Results

In this section, we provide the proofs for the statements in Section 3.2.

B.1 Proof of Lemma 3.1

Proof. Applying vectorization to $f(\mathbf{X})$ and recalling that $\text{vec}(\mathbf{AXB}) = (\mathbf{B}^\top \otimes \mathbf{A}) \text{vec}(\mathbf{X})$, we obtain

$$\text{vec}(f(\mathbf{X})) = \sum_{k=1}^K (\Theta_k^\top \otimes \mathbf{L}^k) \text{vec}(\mathbf{X}).$$

Taking derivatives:

$$\mathbf{J}_f = \frac{\partial f}{\partial \text{vec}(\mathbf{X})} = \sum_{k=1}^K (\Theta_k^\top \otimes T_k(\mathbf{L})).$$

□

B.2 Proof of Theorem 1

Proof. Define the k -hop Jacobian block $\mathbf{J}_k = \Theta_k \otimes \mathbf{L}^k$ and the full Jacobian $\mathbf{J} = \sum_{k=1}^K \mathbf{J}_k$. Because $\Theta_k \otimes \mathbf{L}^k$ has singular values $s_j(\Theta_k) |\lambda_i|^k$, the squared singular values of \mathbf{J}_k are given by $\gamma_{i,j}^{(k)} = \lambda_i^{2k} \mu_{j,k}$.

For a *square* Gaussian matrix, the empirical spectrum of $\Theta_k \Theta_k^\top$ converges to the Marchenko–Pastur (MP) law. Scaling by $\sigma^2 \lambda_i^{2k}$ yields

$$m_1^{(k)} = \sigma^2 \lambda_i^{2k} \quad (13)$$

$$m_2^{(k)} = 2(\sigma^2 \lambda_i^{2k})^2. \quad (14)$$

Independence, together with the rotational symmetry of each Θ_k implies the blocks \mathbf{J}_k are *asymptotically free*. Freeness gives additive R -transforms, hence additive *free* cumulants $\kappa_r(\sum_k \mathbf{J}_k) = \sum_k \kappa_r(\mathbf{J}_k)$. For $r = 1, 2$ the classical moments coincide with the free cumulants, so the ordinary moments of $\gamma_{i,j}$ also add:

$$m_r = \sum_{k=1}^K m_r^{(k)}, \quad r \in \{1, 2\}. \quad (15)$$

Insert Equation (13) into Equation (15):

$$m_1 = \sigma^2 \sum_{k=1}^K \lambda_i^{2k} \quad (16)$$

$$m_2 = 2\sigma^4 \sum_{k=1}^K \lambda_i^{4k}. \quad (17)$$

Finally,

$$\begin{aligned} \text{Var}[\gamma_{i,j}] &= m_2 - m_1^2 \\ &= 2\sigma^4 \sum_k \lambda_i^{4k} - \sigma^4 \left(\sum_k \lambda_i^{2k} \right)^2 \\ &= \sigma^4 \left(\sum_k \lambda_i^{2k} \right)^2. \end{aligned}$$

□

B.3 Proof of Theorem 2

Proof. Recall that the forward pass for one ChebNet layer (omitting the activation function) is:

$$\mathbf{X}^{(l+1)} = \sum_{k=0}^K T_k(\mathbf{L}) \mathbf{X}^{(l)} \mathbf{W}_k^{(l)}. \quad (18)$$

Applying vectorization, we obtain:

$$\text{vec}(\mathbf{X}^{(l+1)}) = \sum_{k=0}^K \left((\mathbf{W}_k^{(l)})^\top \otimes T_k(\mathbf{L}) \right) \text{vec}(\mathbf{X}^{(l)}). \quad (19)$$

By unrolling Equation (19) and taking the derivative with respect to $\text{vec}(\mathbf{X}^{(0)})$, we obtain the sensitivity after l layers:

$$\frac{\partial \text{vec}(\mathbf{X}^{(l)})}{\partial \text{vec}(\mathbf{X}^{(0)})} = \prod_{l=0}^{l-1} \left(\sum_{k=0}^K \left((\mathbf{W}_k^{(l)})^\top \otimes T_k(\mathbf{L}) \right) \right). \quad (20)$$

Focusing on a single feature channel (or summing across channels), we obtain the sensitivity:

$$\frac{\partial \mathbf{X}^{(l)}}{\partial \mathbf{X}^{(0)}} = \prod_{l=0}^{l-1} \left(\sum_{k=0}^K T_k(\mathbf{L}) \mathbf{W}_k^{(l)} \right). \quad (21)$$

Then, the sensitivity of node v with respect to node u after l layers is given by:

$$\frac{\partial \mathbf{x}_v^{(l)}}{\partial \mathbf{x}_u^{(0)}} = \left(\prod_{l=0}^{l-1} \left(\sum_{k=0}^K T_k(\mathbf{L}) \mathbf{W}_k^{(l)} \right) \right)_{v,u}. \quad (22)$$

□

B.4 Proof of Theorem 3

Proof. Given that the graph-wise Jacobian of Equation (7) is Equation (2), we note that each term in the Jacobian, of the form $\mathbf{W}_k^\top \otimes T_k(\mathbf{L})$, is an antisymmetric matrix. This follows from the fact that \mathbf{W}_k is antisymmetric by construction, $T_k(\mathbf{L})$ preserves the symmetry of the normalized Laplacian, and the Kronecker product of an antisymmetric matrix with a symmetric matrix is itself antisymmetric. Finally, since the sum of antisymmetric matrices remains antisymmetric, it follows that the graph-wise Jacobian of Equation (7) has purely imaginary eigenvalues. \square

B.5 Sensitivity upperbound of standard MPNNs

In the following theorem we report the sensitivity upperbound computed for standard MPNNs in [27].

Theorem 5 (Sensitivity upperbound of standard MPNNs, taken from [27]). *Consider a standard MPNN with l layers, where c_σ is the Lipschitz constant of the activation σ , w is the maximal entry-value over all weight matrices, and d is the embedding dimension. For $u, v \in V$ we have*

$$\left\| \frac{\partial \mathbf{h}_v^{(l)}}{\partial \mathbf{h}_u^{(0)}} \right\| \leq \underbrace{(c_\sigma w d)^l}_{\text{model}} \underbrace{(\mathbf{O}^l)_{vu}}_{\text{topology}}, \quad (23)$$

with $\mathbf{O} = c_r \mathbf{I} + c_a \mathbf{A} \in \mathbb{R}^{n \times n}$ is the message passing matrix adopted by the MPNN, with c_r and c_a are the contributions of the self-connection and aggregation term.

B.6 Proof of Theorem 4

Proof. First, recall the Jacobian explicitly:

$$\mathbf{J}^{(l)} = \mathbf{I} + \epsilon \sum_{k=0}^K \left((\mathbf{W}_k^{(l)})^\top \otimes T_k(\mathbf{L}) \right). \quad (24)$$

Define:

$$A^{(l)} = \sum_{k=0}^K \left((\mathbf{W}_k^{(l)})^\top \otimes T_k(\mathbf{L}) \right). \quad (25)$$

Antisymmetric Case: When $(\mathbf{W}_k^{(l)})^\top = -\mathbf{W}_k^{(l)}$, the matrix $A^{(l)}$ is antisymmetric, because it is a Kronecker product of antisymmetric and symmetric matrices. Its eigenvalues are purely imaginary (or symmetric about zero), meaning their real parts vanish. Thus, for the spectral radius, we have:

$$\|\mathbf{J}^{(l)}\|_2^2 = \rho \left((\mathbf{J}^{(l)})^\top \mathbf{J}^{(l)} \right) \quad (26)$$

$$= \rho \left(\mathbf{I} + \epsilon(A^{(l)} + (A^{(l)})^\top) + \epsilon^2(A^{(l)})^\top A^{(l)} \right). \quad (27)$$

Due to antisymmetry:

$$(A^{(l)})^\top + A^{(l)} = 0. \quad (28)$$

Hence:

$$\|\mathbf{J}^{(l)}\|_2^2 = \rho \left(\mathbf{I} + \epsilon^2(A^{(l)})^\top A^{(l)} \right). \quad (29)$$

The matrix $(A^{(l)})^\top A^{(l)}$ is symmetric and positive semi-definite, having nonnegative real eigenvalues. Expanding around the identity gives:

$$\|\mathbf{J}^{(l)}\|_2 = \sqrt{1 + \epsilon^2 \lambda_{\max} \left((A^{(l)})^\top A^{(l)} \right)} = 1 + O(\epsilon^2). \quad (30)$$

Thus, no exponential growth or decay occurs.

General Case of Stable-ChebNet (Without Antisymmetry): For arbitrary matrices, the linear term $(A^{(l)} + (A^{(l)})^\top)$ typically does not vanish, introducing nonzero real eigenvalues. This causes exponential growth or decay in the Jacobian norm across layers:

$$\|\mathbf{J}^{(l)}\|_2 \approx 1 \pm C\epsilon, \quad (31)$$

for some constant $C = \max_{i=0, \dots, l} C^{(i)} > 0$. Iterating over layers results in exponential instability:

$$\|\mathbf{J}^{(l)} \mathbf{J}^{(l-1)} \dots \mathbf{J}^{(0)}\|_2 \approx (1 \pm C\epsilon)^l, \quad (32)$$

which grows or decays exponentially as the depth l increases.

Thus, antisymmetric weights provide explicit protection against exponential growth or decay, while general weights typically do not.

Standard ChebNet Case: Iterating over layers, this results in exponential instability:

$$\|\mathbf{J}^{(l)} \mathbf{J}^{(l-1)} \dots \mathbf{J}^{(0)}\|_2 \approx (\pm C)^l, \quad (33)$$

which grows or decays exponentially as the depth l increases.

□

C Dataset and Baseline Description

C.1 Graph Property Prediction task description

For the graph property prediction dataset, we make use of three separate tasks:

Diameter (Graph-Level). The diameter is defined as the length of the longest shortest path between any two nodes in the graph. It requires aggregating information from distant regions of the graph.

SSSP (Node-Level). Single-Source Shortest Path requires predicting each node’s distance to a designated source node. Solving this task with GNNs places a strong emphasis on propagating information from nodes that may lie many hops away from the source.

Eccentricity (Node-Level). For each node u , the eccentricity is the length of the maximum shortest path between u and any other node. As with diameter and SSSP, accurate eccentricity estimation relies heavily on capturing long-distance relationships.

In total, the task contains 5,120 graphs for training, 640 for validation, and 1,280 for testing. The train/validation/test splits follow the same seed and setup in [45]. We train our models by optimizing mean squared error (MSE) for each task, performing a grid search over hyperparameters (e.g., learning rate, weight decay, number of layers). Each experiment is repeated across four random initializations, and we report the average performance.

Because shortest-path-related tasks inherently rely on propagating signals over large portions of a graph, they serve as a natural stress test for the capacity of GNNs to perform long-range propagation.

C.2 Barbell graph task description

A Barbell graph $B_{n,k}$ is formed by connecting two complete graphs K_n (the “bells”) with a simple path of length k (the “bridge”). Therefore, every node inside a bell has high intra-cluster connectivity ($n - 1$ neighbors), whereas the bridge nodes have degree 2 and constitute the only communication route between the two bells. In this task, the target for nodes is to output the average input feature over nodes of the opposite bell and vice-versa, as illustrated in Figure 6. **Mean squared error (MSE)** is used for node-level regression as a proxy for how severely the GNN is either *oversquashing* (failing to pass information across the narrow “bridge”) or *oversmoothing* (collapsing all node embeddings to be nearly identical). Numerical MSE outcomes are associated with each pathology.

An MSE of around 1 corresponds to oversquashing: The model is so bottlenecked by the bridge that it effectively ignores or fails to incorporate information from the other “bell.” In other words, each side of the barbell can predict its own node labels, but the information from the opposite side never gets through, leading to a characteristic level of error (≈ 1 in their chosen label distribution).

An MSE of around 30 corresponds to oversmoothing: The model passes messages so many times (or in such a way) that it “collapses” all node embeddings toward the same prediction, ignoring local distinctions within each side. Because the barbell’s node labels are diverse (randomly assigned), forcing every node toward the same value yields a much larger overall MSE (≈ 30 in their setup).

On the other hand, an MSE of around 0.5 as shown in the results simply means that the model is doing better than the severe oversquashing case (it is not completely failing to pass information across

the barbell). In other words, some amount of meaningful communication is happening between the two “bells,” and the node representations are not entirely collapsed.

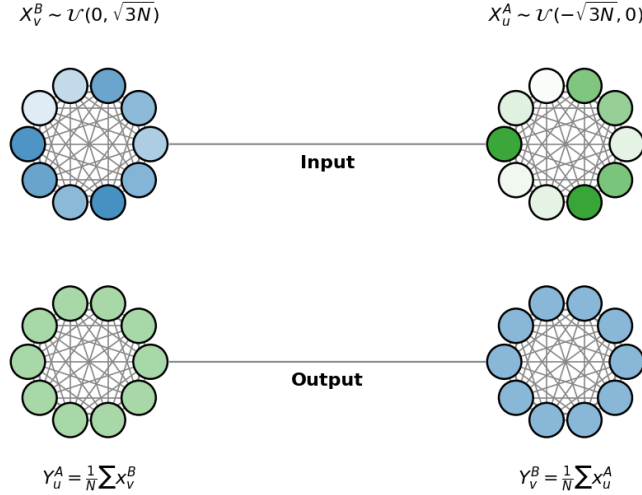


Figure 6: Illustration of input and output node features in a barbell graph setup. Adapted from Bamberger et.al [5].

C.3 Employed baselines

In our experiments, the performance of our method is compared with various state-of-the-art GNN baselines from the literature. Specifically, we consider:

- Classical GNN methods, i.e., GCN [57], GraphSAGE [51], GAT [83], GIN [90], and GCNII [20], ChebNet [25], ChebNetII [52], CIN++ [42], PathNN [66], S2GCN [40], SGC [87], SIGN [38], GRAMA [34], GatedGCN [13], CoGNN [37], H2GCN [95], CPGNN [94], FAGCN [11], GPR-GNN [21], FSGNN [65], GloGNN [60], GBK-GNN [29], JacobiConv [85];
- Differential-equation inspired GNNs (DE-GNNs), i.e., DGC [86], GRAND [17], GraphCON [73], A-DGN [45], and SWAN [46] PH-DGN [54];
- Graph Transformers, i.e., SAN [59], GraphGPS [71], GOAT [58], Exphormer [80], GRIT[63], GraphViT [53], G.MLPMixer [53] SPEXPFORMER [79], TIGT [22], SGFormer [89], NodeFormer [88], Specformer [10], GCN-nsampler and GAT-nsampler [93], GT [78], NAGphormer [19], Polynormer [26];
- Rewiring-based methods, i.e., LASER [7], and DRew [49];
- SSM-based GNN, i.e., Graph-Mamba [84], GMN [8], MP-SSM [16]

D Hyper-parameter grids

Tables 6 to 8 summarize our hyperparameter exploration: Table 6 listing the sweep ranges for Stable-ChebNet on Peptides-func, graph-property benchmarks (Diameter, SSSP, Eccentricity), and ogbn-arxiv along with ogbn-proteins experiments, respectively.

Table 6: Hyper-parameter grid for Stable-ChebNet ablation on *Peptides-func*.

Hyper-parameter	Reference	Sweep values used in ablation
Hidden dim d	140	100, 120, 140, 145, 160
Polynomial order K	10	6,8,10
Num of layers	4	3,4,5
MLP layers	2	1,2,3
Step size ε	0.5	[0.1, 1.0]
Dissipative force γ	0.05	0.001, 0.01, 0.05, 0.1
Batch size	64	32, 64, 128
Learning rate	0.001	0.0001, 0.001, 0.01
Optimizer	AdamW	AdamW
Pos-enc type	None	None, Laplacian, RW
Pos-enc dim	16	8, 16, 32

Table 7: Hyper-parameter grid and best settings for Stable-ChebNet on three synthetic graph-property benchmarks.

Hyper-parameter	Values in grid	Diam	SSSP	Ecc
Hidden dimension d	20, 30, 50	50	30	30
Number of layers	1, 2, 3, 5, 10, 20	20	5	5
Polynomial order K	3, 5, 10	4	10	10
Step size ε	0.01, 0.10, 0.20, 0.30	0.40	0.30	0.30
Dissipative force γ	0, 0.01, 0.50, 1	0.01	0.00	0.00
Activation function	<code>tanh</code> , <code>relu</code>	<code>relu</code>	<code>relu</code>	<code>relu</code>
Learning rate	0.001,0.003	0.003	0.003	0.003
Weight decay	1×10^{-6}	1×10^{-6}	1×10^{-6}	1×10^{-6}

Table 8: Hyper-parameter sweep ranges for Stable-ChebNet on *ogbn-arxiv* and *ogbn-proteins*.

Hyper-parameter	Sweep (arxiv)	Sweep (proteins)
Hidden dim d	128, 256, 512	256, 512, 1024
Polynomial order K	4, 5, 6, 10	5, 10, 15
Num of layers	2, 3, 4, 5	3, 5, 7
MLP layers	1, 2, 3	1, 2, 3
Step size ϵ	[0.1, 1.0]	[0.1, 1.0]
Dissipative force γ	0.01, 0.05, 0.1	0.01, 0.05, 0.1
Batch size	256, 512, 1024	512, 1024, 2048
Learning rate	0.001, 0.01, 0.05	0.0005, 0.001, 0.005
Optimizer	Adam	Adam
Pos-enc type	None, Laplacian, RW	None, Laplacian, RW
Pos-enc dim	8, 16, 32	16, 32, 64

E Eigenvalues distribution comparison

Figures 7a and 7b provide an empirical comparison of the Jacobian eigenvalue spectra for a classical ChebNet versus our Stable-ChebNet layer at $K = 5$. The top two panels illustrate ChebNet’s spectrum: the left histogram shows that both the real and imaginary parts of its eigenvalues span broadly from roughly -2 to $+2$, with pronounced peaks near the extremes signaling a large spectral radius and a susceptibility to unstable, oscillatory dynamics. The accompanying scatter plot maps these eigenvalues in the complex plane, revealing many points lying well outside the unit circle, which corroborates the theoretical prediction that high-order Chebyshev filters can push the system toward chaotic regimes.

In contrast, the bottom row portrays the spectrum of Stable-ChebNet. Its histogram (bottom left) is tightly concentrated within approximately -0.4 to $+0.4$ on both axes, indicating that eigenvalues

remain far inside the unit circle. The complex-plane scatter (bottom right) further demonstrates that all eigenvalues lie symmetrically about the imaginary axis and are bounded in magnitude, consistent with the antisymmetric weight parameterization, guaranteeing purely imaginary Jacobian eigenvalues and therefore more stable dynamics.

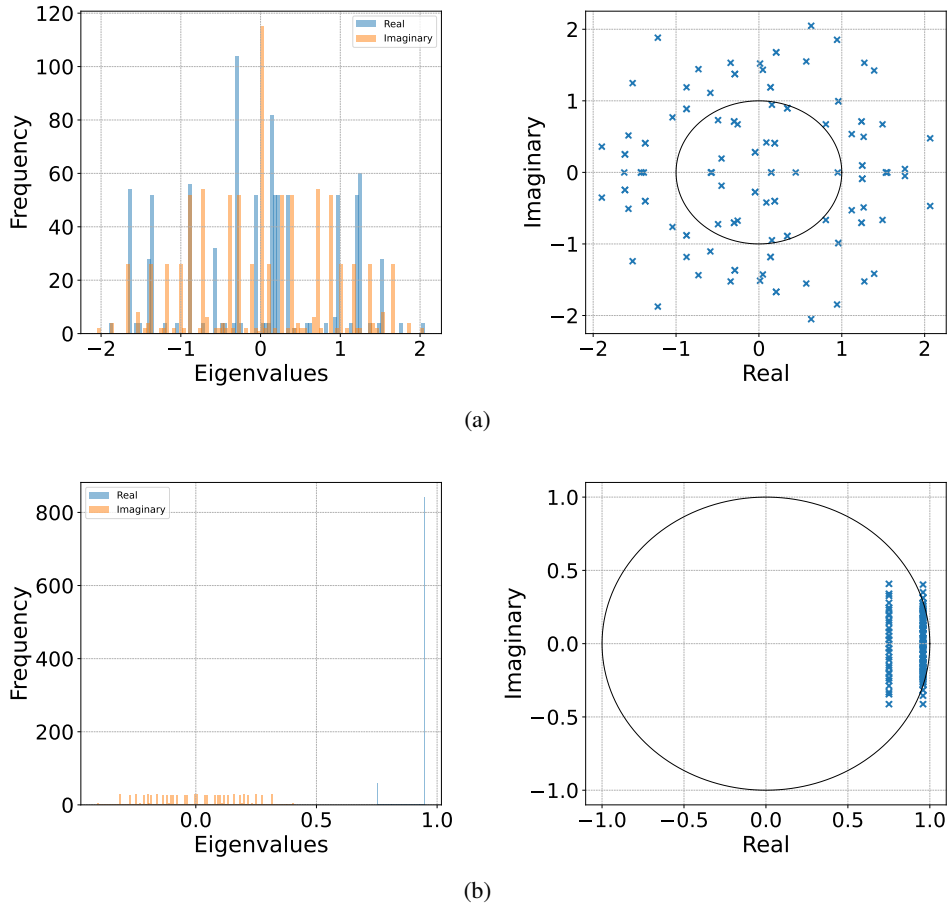


Figure 7: Comparison of eigenvalue distributions: (a) ChebNet and (b) StableChebNet.

F Additional Experiments on Heterophilic Benchmarks

To further evaluate the performance of our Stable-ChebNet, we assess its effectiveness in capturing complex relational information in heterophilic settings, where nodes belonging to same class are often connected through longer and sparser paths, we consider the five node classification tasks introduced in [70]. Specifically, we consider the “Roman-empire”, “Amazon-ratings” and “Minesweeper”, “Tolokers” datasets. We adhere to the same data and experimental setting presented in [70].

Table 9: Mean test set score and std averaged over 4 random weight initializations on heterophilic datasets. The higher, the better.

Model	Roman-empire Acc \uparrow	Amazon-ratings Acc \uparrow	Minesweeper AUC \uparrow	Tolokers AUC \uparrow
[62]				
MLP-2	66.04 \pm 0.71	49.55 \pm 0.81	50.92 \pm 1.25	74.58 \pm 0.75
SGC-1	44.60 \pm 0.52	40.69 \pm 0.42	82.04 \pm 0.77	73.80 \pm 1.35
MLP-1	64.12 \pm 0.61	38.60 \pm 0.41	50.59 \pm 0.83	71.89 \pm 0.82
MPNNs				
GAT	80.87 \pm 0.30	49.09 \pm 0.63	92.01 \pm 0.68	83.70 \pm 0.47
GAT (LapPE)	84.80 \pm 0.46	44.90 \pm 0.73	93.50 \pm 0.54	84.99 \pm 0.54
GAT (RWSE)	86.62 \pm 0.53	48.58 \pm 0.41	92.53 \pm 0.65	85.02 \pm 0.67
Gated-GCN	74.46 \pm 0.54	43.00 \pm 0.32	87.54 \pm 1.22	77.31 \pm 1.14
GCN	73.69 \pm 0.74	48.70 \pm 0.63	89.75 \pm 0.52	83.64 \pm 0.67
GCN (LapPE)	83.37 \pm 0.55	44.35 \pm 0.36	94.26 \pm 0.49	84.95 \pm 0.78
GCN (RWSE)	84.84 \pm 0.55	46.40 \pm 0.55	93.84 \pm 0.48	85.11 \pm 0.77
CO-GNN(Σ , Σ)	91.57 \pm 0.32	51.28 \pm 0.56	95.09 \pm 1.18	83.36 \pm 0.89
CO-GNN(μ , μ)	91.37 \pm 0.35	54.17 \pm 0.37	97.31 \pm 0.41	84.45 \pm 1.17
SAGE	85.74 \pm 0.67	53.63 \pm 0.39	93.51 \pm 0.57	82.43 \pm 0.44
Graph Transformers				
Expformer	89.03 \pm 0.37	53.51 \pm 0.46	90.74 \pm 0.53	83.77 \pm 0.78
NAGphormer	74.34 \pm 0.77	51.26 \pm 0.72	84.19 \pm 0.66	78.32 \pm 0.95
GOAT	71.59 \pm 1.25	44.61 \pm 0.50	81.09 \pm 1.02	83.11 \pm 1.04
GPS	82.00 \pm 0.61	53.10 \pm 0.42	90.63 \pm 0.67	83.71 \pm 0.48
GPS _{GCN} +Performer (LapPE)	83.96 \pm 0.53	48.20 \pm 0.67	93.85 \pm 0.41	84.72 \pm 0.77
GPS _{GCN} +Performer (RWSE)	84.72 \pm 0.65	48.08 \pm 0.85	92.88 \pm 0.50	84.81 \pm 0.86
GPS _{GCN} +Transformer (LapPE)	OOM	OOM	91.82 \pm 0.41	83.51 \pm 0.93
GPS _{GCN} +Transformer (RWSE)	OOM	OOM	91.17 \pm 0.51	83.53 \pm 1.06
GT	86.51 \pm 0.73	51.17 \pm 0.66	91.85 \pm 0.76	83.23 \pm 0.64
GT-sep	87.32 \pm 0.39	52.18 \pm 0.80	92.29 \pm 0.47	82.52 \pm 0.92
Polynormer	92.55 \pm 0.30	54.81 \pm 0.49	97.46 \pm 0.36	85.91 \pm 0.74
Heterophily-Designated GNNs				
CPGNN	63.96 \pm 0.62	39.79 \pm 0.77	52.03 \pm 5.46	73.36 \pm 1.01
FAGCN	65.22 \pm 0.56	44.12 \pm 0.30	88.17 \pm 0.73	77.75 \pm 1.05
FSGNN	79.92 \pm 0.56	52.74 \pm 0.83	90.08 \pm 0.70	82.76 \pm 0.61
GBK-GNN	74.57 \pm 0.47	45.98 \pm 0.71	90.85 \pm 0.58	81.01 \pm 0.67
GloGNN	59.63 \pm 0.69	36.89 \pm 0.14	51.08 \pm 1.23	73.39 \pm 1.17
GPR-GNN	64.85 \pm 0.27	44.88 \pm 0.34	86.24 \pm 0.61	72.94 \pm 0.97
H2GCN	60.11 \pm 0.52	36.47 \pm 0.23	89.71 \pm 0.31	73.35 \pm 1.01
JacobiConv	71.14 \pm 0.42	43.55 \pm 0.48	89.66 \pm 0.40	68.66 \pm 0.65
Graph SSMs				
GMN	87.69 \pm 0.50	54.07 \pm 0.31	91.01 \pm 0.23	84.52 \pm 0.21
GPS + Mamba	83.10 \pm 0.28	45.13 \pm 0.97	89.93 \pm 0.54	83.70 \pm 1.05
GRAMA _{GCN}	88.61 \pm 0.43	53.48 \pm 0.62	95.27 \pm 0.71	86.23 \pm 1.10
MP-SSM	90.91 \pm 0.48	53.65 \pm 0.71	95.33 \pm 0.72	85.26 \pm 0.93
Ours				
Stable-ChebNet	92.03 \pm 0.85	53.15 \pm 0.21	95.71 \pm 2.26	85.55 \pm 3.35