

Efficient and Stable Graph Scattering Transforms via Pruning

Ioannidis, Vassilis N.; Chen, Siheng; Giannakis, Georgios B.

TR2021-018 March 20, 2021

Abstract

Graph convolutional networks (GCNs) have well-documented performance in various graph learning tasks, but their analysis is still at its infancy. Graph scattering transforms (GSTs) offer training-free deep GCN models that extract features from graph data, and are amenable to generalization and stability analyses. The price paid by GSTs is exponential complexity in space and time that increases with the number of layers. This discourages deployment of GSTs when a deep architecture is needed. The present work addresses the complexity limitation of GSTs by introducing an efficient so-termed pruned (p)GST approach. The resultant pruning algorithm is guided by a graph-spectrum-inspired criterion, and retains informative scattering features on-the-fly while bypassing the exponential complexity associated with GSTs. Stability of the novel pGSTs is also established when the input graph data or the network structure are perturbed. Furthermore, the sensitivity of pGST to random and localized signal perturbations is investigated analytically and experimentally. Numerical tests showcase that pGST performs comparably to the baseline GST at considerable computational savings. Furthermore, pGST achieves comparable performance to state-of-the-art GCNs in graph and 3D point cloud classification tasks. Upon analyzing the pGST pruning patterns, it is shown that graph data in different domains call for different network architectures, and that the pruning algorithm may be employed to guide the design choices for contemporary GCNs.

IEEE Transactions on Pattern Analysis and Machine Intelligence 2021

Efficient and Stable Graph Scattering Transforms via Pruning

Vassilis N. Ioannidis, *Student Member, IEEE*, Siheng Chen, *Member, IEEE*,
and Georgios B. Giannakis, *Fellow, IEEE*,



Abstract—Graph convolutional networks (GCNs) have well-documented performance in various graph learning tasks, but their analysis is still at its infancy. Graph scattering transforms (GSTs) offer training-free deep GCN models that extract features from graph data, and are amenable to generalization and stability analyses. The price paid by GSTs is exponential complexity in space and time that increases with the number of layers. This discourages deployment of GSTs when a deep architecture is needed. The present work addresses the complexity limitation of GSTs by introducing an efficient so-termed pruned (p)GST approach. The resultant pruning algorithm is guided by a graph-spectrum-inspired criterion, and retains informative scattering features on-the-fly while bypassing the exponential complexity associated with GSTs. Stability of the novel pGSTs is also established when the input graph data or the network structure are perturbed. Furthermore, the sensitivity of pGST to random and localized signal perturbations is investigated analytically and experimentally. Numerical tests showcase that pGST performs comparably to the baseline GST at considerable computational savings. Furthermore, pGST achieves comparable performance to state-of-the-art GCNs in graph and 3D point cloud classification tasks. Upon analyzing the pGST pruning patterns, it is shown that graph data in different domains call for different network architectures, and that the pruning algorithm may be employed to guide the design choices for contemporary GCNs.

1 INTRODUCTION

The abundance of graph-abiding data calls for advanced learning techniques that complement nicely standard machine learning tools when the latter cannot be directly employed, e.g. due to irregular data inter-dependencies. Permeating the benefits of deep learning to graph data, graph convolutional networks (GCNs) offer a versatile and powerful framework to learn from complex graph data [1]. GCNs and variants thereof have been remarkably successful in social network analysis, 3D point cloud processing, recommender systems and action recognition. However, researchers have recently reported less consistent perspectives on the desirable GCN designs. For example, experiments in social network analysis have argued that deeper GCNs marginally increase the learning performance [2], whereas a method for 3D point cloud segmentation achieves state-of-the-art performance with a 56-layer GCN network [3]. These ‘controversial’ empirical findings motivate theoretical analysis to

understand the fundamental performance-defining factors, and the resultant design choices for high-performance GCNs.

Aiming to bestow GCNs with theoretical guarantees, one promising path is to study graph scattering transforms (GSTs) – an analysis framework that has been advocated to assess stability and explain the success of deep neural networks (DNNs) [4], [5]. GSTs are non-trainable GCNs comprising a cascade of graph filter banks followed by nonlinear activation functions. The graph filter banks are designed analytically to scatter an input graph signal into multiple channels. GSTs extract stable features of graph data that can be utilized for downstream graph learning tasks [6], with competitive performance especially when the number of training examples is small. Under certain conditions on the graph filter banks, GSTs are endowed with energy conservation properties [7], as well as stability that amounts to robustness to graph topology deformations [8]. Inherited from scattering transforms, GSTs however are known to incur exponential complexity in space and time that increases with the number of layers (GST depth) [4], [5]. Furthermore, stability should not come at odds with sensitivity. A filter’s output should be sensitive to and able to cope with perturbations of large magnitude. Current GST efforts have not addressed transform sensitivity to input noise. Lastly, graph data in different domains (e.g., social networks versus 3D point clouds) have distinct properties, which prompts domain-adaptive GST designs.

1.1 Contributions

The present paper develops a data-adaptive pruning approach to systematically retain informative GST features, which justifies the term pruned graph scattering transform (pGST). The pruning decisions are guided by a criterion promoting alignment (matching) of the input graph spectrum with that of the graph filters. The optimal pruning decisions are provided on-the-fly, and alleviate the exponential complexity of GSTs. We prove that the pGST is stable to perturbations of the input graph data and those of the network structure. Under certain conditions on the perturbation energy, the resultant pruning patterns before and after the network and input perturbations are identical, and the overall pGST is stable. Further, the sensitivity of pGST to random and localized noise is theoretically and experimentally investigated. It turns out that pGST is more sensitive to noise that is localized in the graph spectrum relative to noise that is uniformly spread over the spectrum. This is appealing because pGST can detect transient changes in the graph spectral domain, while ignoring small random

- *V. N. Ioannidis and G. B. Giannakis are with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN. Emails: {ioann006,georgios}@umn.edu. S. Chen is with Mitsubishi Electric Research Laboratories, Cambridge, MA. Email: schen@merl.com. The work was carried out during V. N. Ioannidis’ internship at Mitsubishi Electric Research Laboratories in the summer of 2019.*

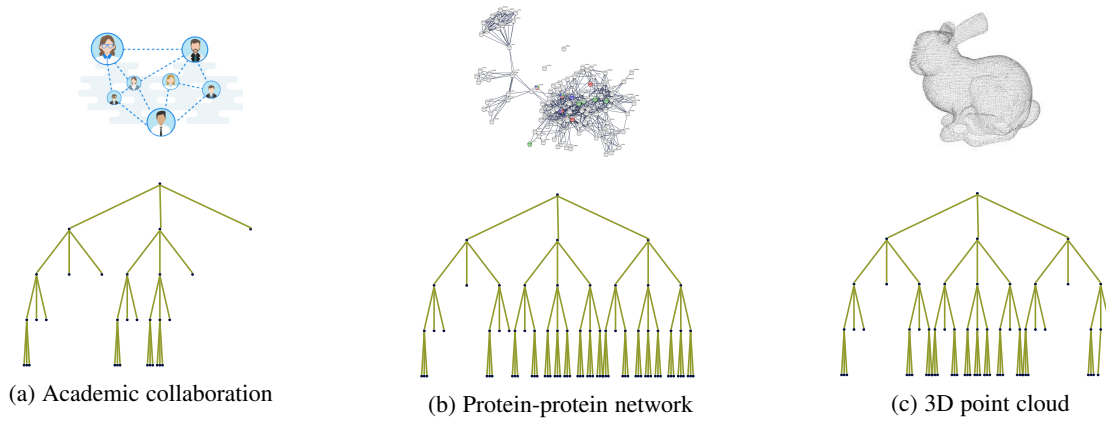


Fig. 1: Illustration of the same pGST applied to different graph datasets. For the social network (a) most GST branches are pruned, suggesting that most information is captured by local interactions.

perturbations. With extensive experiments we showcase that the proposed pGSTs perform similar to, and in certain cases better than, the baseline GSTs that use all scattering features, while achieving significant computational savings. Furthermore, the *efficient* and *stable* features extracted can be utilized towards graph classification and 3D point cloud recognition. Even without any training on the feature extraction step, the performance is comparable to state-of-the-art deep supervised learning approaches, particularly when training data are scarce. By analyzing the pruning patterns of the pGST, we deduce that graph data in different domains call for distinct network architectures; see Fig. 1. Finally, we establish that the number of pGST layers after pruning can be utilized to guide the design parameters of contemporary GCNs.

A preliminary version of this work was presented in [9]. Relative to this, the novelty here is fourfold. First, novel theoretical results are derived that establish stability of the pGST to structural perturbations. An additional interpretation of the pruning criterion as a rate-distortion tradeoff is presented. Furthermore, sensitivity analysis of pGST reveals that the transform is more sensitive to noise that is localized in the graph spectrum relative to noise that is uniformly spread over the spectrum, which is appealing because pGST can detect transient changes in the graph spectral domain. Finally, new experimental results showcase: i) the effect of localized and random noise to the pruning algorithm; ii) the competitive performance of pGST for semi-supervised learning tasks; and, iii) the strong link between design choices in GCNs, and the pruning patterns of pGST.

2 RELATED WORK

2.1 Graph convolutional networks

GCNs rely on a layered processing architecture comprising trainable graph convolutional operations to linearly combine features per graph neighborhood, followed by pointwise nonlinear functions applied to the linearly transformed features [1]. Complex GCNs and their variants have shown remarkable success in graph semi-supervised learning [10]–[12] and graph classification tasks [13]. GCNs as simple as a single-layer linear module can offer high performance in certain social network learning applications [2]. On the other hand, a 56-layer GCN has been employed to achieve state-of-the-art performance in 3D point cloud segmentation [3]. Whether simple or complex, designing GCNs guided by properties of the graph data at hand is of paramount importance.

2.2 Graph scattering transforms

To understand the success of GCNs analytically, recent works study the stability properties of GSTs with respect to metric deformations of the domain [7], [8], [14]. GSTs generalize scattering transforms [4], [5] to non-Euclidean domains. GSTs are a cascade of graph filter banks and nonlinear operations that is organized in a tree-structured architecture. The number of extracted GST features grows exponentially with the number of layers. Theoretical guarantees for GSTs are obtained after fixing the graph filter banks to implement a set of graph wavelets. The work in [7] establishes energy conservation properties for GSTs given that certain energy-preserving graph wavelets are employed, and also prove that GSTs are stable to graph structure perturbations; see also [14] that focuses on diffusion wavelets. On the other hand, [8] proves stability to relative metric deformations for a wide class of graph wavelet families. These contemporary works shed light into the stability and generalization capabilities of GCNs. However, stable transforms are not necessarily informative, and albeit highly desirable, a principled approach to selecting informative GST features remains still an uncharted venue.

2.3 Neural network based compression

This work advocates pruning of GSTs to avoid the exponential growth of features with the network depth, which is naturally related to deep neural network (DNN) based compression [15], [16]. State-of-the-art DNNs typically entail a large number of network layers and corresponding trainable parameters, which introduce excessive computational and memory costs [17], [18]. Confronting these challenges, neural network compression aims at reducing the network parameters and pruning the architecture to facilitate practical deployment of DNN-based solutions. Typical compression techniques prune redundant parameters during training, while minimizing the effect of the pruned parameters to the learned features [15], [16]. Recent approaches perform network pruning at initialization before training the DNN [19].

Although DNN-based compression is a fruitful direction, pruning GCNs has not received commensurate attention. Typical GCNs are applied to relatively small data sets and require only a small number of layers and training parameters to attain state-of-the-art learning performance [2], [10], [11], which explains why GCN-based compression is yet to be explored. Deep GCNs are emerging however, to deal with web-scale graphs [3], [20].

Such overparameterized GCNs motivate the development of novel compression techniques. Our work can be seen as a stepping stone towards pruning GCNs; meanwhile, pruning GSTs versus GCNs/DNNs has differences in two aspects. First, GSTs are nontrainable feature extractors, whereas GCNs/DNNs introduce trainable parameters that should be considered by the compression algorithm. Second, GST pruning is inherently an one-shot process and pruning is performed in an online and even adaptive fashion, whereas pruning GCNs or DNNs should be performed offline.

3 BACKGROUND

Consider an undirected graph $\mathcal{G} := \{\mathcal{V}, \mathcal{E}\}$ with node set $\mathcal{V} := \{v_i\}_{i=1}^N$, and edge set $\mathcal{E} := \{e_i\}_{i=1}^E$. Its connectivity is described by the *graph shift matrix* $\mathbf{S} \in \mathbb{R}^{N \times N}$, whose (n, n') th entry $S_{nn'}$ is nonzero if $(n, n') \in \mathcal{E}$ or if $n = n'$. A typical choice for \mathbf{S} is the symmetric adjacency or the Laplacian matrix. Further, each node can be also associated with a few attributes. Collect attributes across all nodes in the matrix $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_F] \in \mathbb{R}^{N \times F}$, where each column $\mathbf{x}_f \in \mathbb{R}^N$ can be regarded as a *graph signal*.

3.1 Graph Fourier transform

A Fourier transform corresponds to the expansion of a signal over bases that are invariant to filtering; here, this graph frequency basis is the eigenbasis of the shift matrix \mathbf{S} . Henceforth, \mathbf{S} is assumed normal with $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$, where $\mathbf{V} \in \mathbb{R}^{N \times N}$ forms the graph Fourier basis, and $\mathbf{\Lambda} \in \mathbb{R}^{N \times N}$ is the diagonal matrix of corresponding eigenvalues $\lambda_0, \dots, \lambda_{N-1}$, that can be thought of as graph frequencies. The graph Fourier transform (GFT) of $\mathbf{x} \in \mathbb{R}^N$ is $\hat{\mathbf{x}} = \mathbf{V}^\top \mathbf{x} \in \mathbb{R}^N$, while the inverse transform is $\mathbf{x} = \mathbf{V}\hat{\mathbf{x}}$. The vector $\hat{\mathbf{x}}$ represents the signal's expansion in the eigenvector basis and describes the graph spectrum of \mathbf{x} . The inverse GFT reconstructs the graph signal from its graph spectrum by combining graph frequency components weighted by the coefficients of the signal's graph Fourier transform. GFT is a theoretical tool that has been popular for analyzing graph data in the graph spectral domain.

3.2 Graph convolutional networks

GCNs permeate the benefits of CNNs from processing Euclidean data to modeling graph structured data. GCNs model graph data through a succession of layers, each consisting of a graph convolutional operation (a.k.a. graph filter), a pointwise nonlinear function $\sigma(\cdot)$, and oftentimes also a pooling operation. Given $\mathbf{x} \in \mathbb{R}^N$, the graph convolution operation diffuses each node's information to its neighbors to obtain $\mathbf{S}\mathbf{x}$ with n th entry $[\mathbf{S}\mathbf{x}]_n = \sum_{n' \in \mathcal{N}_n} S_{nn'} x_{n'}$ being a weighted average of the one-hop neighboring features. Successive application of \mathbf{S} will reach multi-hop neighbors, spreading the information across the network. Summing up, a K th-order graph convolutional operation (graph filtering) is

$$h(\mathbf{S})\mathbf{x} := \sum_{k=0}^K w_k \mathbf{S}^k \mathbf{x} = \mathbf{V} \hat{h}(\mathbf{\Lambda}) \hat{\mathbf{x}} \quad (1)$$

where the graph filter $h(\cdot)$ is parameterized by the learnable weights $\{w_k\}_{k=0}^K$, and in the graph spectral domain $\hat{h}(\mathbf{\Lambda}) = \sum_{k=0}^K w_k \mathbf{\Lambda}^k$. In the graph vertex domain, the learnable weights capture the influences from various orders of neighbors; and in the graph spectral domain, those weights adaptively adjust the focus and emphasize certain graph frequency bands. GCNs employ various graph filter banks per layer, and learn the parameters that minimize a predefined learning objective, such as classification, or regression.

3.3 Graph scattering transforms

GSTs are the training-free counterparts of GCNs, where the parameters of graph convolutions are fixed based on a design criterion. Per GST layer, the input is graph filtered using a filter bank $\{h_j(\mathbf{S})\}_{j=1}^J$, an elementwise nonlinear function $\sigma(\cdot)$, and a pooling operator U . At the first layer, the input $\mathbf{x} \in \mathbb{R}^N$ constitutes the first scattering feature vector $\mathbf{z}_{(0)} := \mathbf{x}$. Next, $\mathbf{z}_{(0)}$ is processed through $\{h_j(\cdot)\}_{j=1}^J$ and $\sigma(\cdot)$ to generate $\{\mathbf{z}_{(j)}\}_{j=1}^J$ with $\mathbf{z}_{(j)} := \sigma(h_j(\mathbf{S})\mathbf{z}_{(0)})$. At the second layer, the same operation is repeated per j . This yields a tree structure with J branches stemming out from non-leaf node; see also Fig. 2. The ℓ th layer of the tree includes J^ℓ nodes. Each node of layer ℓ is indexed by the path $p^{(\ell)}$ of the sequence of ℓ graph convolutions applied to the input \mathbf{x} , i.e. $p^{(\ell)} := (j^{(1)}, j^{(2)}, \dots, j^{(\ell)})$.¹

The scattering feature vector at the tree node indexed by $(p^{(\ell)}, j)$ at layer $\ell + 1$, is

$$\mathbf{z}_{(p^{(\ell)}, j)} = \sigma(h_j(\mathbf{S})\mathbf{z}_{(p^{(\ell)})}) \quad (2)$$

where $p^{(\ell)}$ holds the list of indices of the parent nodes ordered by ancestry, and all path $p^{(\ell)}$ in the tree with length ℓ are included in the path set $\mathcal{P}^{(\ell)}$ with $|\mathcal{P}^{(\ell)}| = 2^\ell$. The memoryless nonlinearity $\sigma(\cdot)$ disperses the graph frequency representation through the spectrum, and endows the GST with increased discriminating power [8]. By exploiting graph sparsity, the computational complexity of (2) is $\mathcal{O}(KE)$, where $E = |\mathcal{E}|$ is the number of edges in \mathcal{G} . Each feature vector $\mathbf{z}_{(p^{(\ell)})}$ is summarized by an aggregation operator $U(\cdot)$ defining a scalar scattering coefficient $\phi_{(p^{(\ell)})} := U(\mathbf{z}_{(p^{(\ell)})})$, where $U(\cdot)$ is typically an averaging operator that reduces dimensionality of the extracted features. The scattering coefficient per tree node reflects the activation level at a certain graph frequency band.

These scattering coefficients are collected across all tree nodes to form a scattering feature map

$$\Phi(\mathbf{x}) := \left\{ \left\{ \phi_{(p^{(\ell)})} \right\}_{p^{(\ell)} \in \mathcal{P}^{(\ell)}} \right\}_{\ell=0}^L \quad (3)$$

where $|\Phi(\mathbf{x})| = \sum_{\ell=0}^L J^\ell$. The GST operation resembles a forward pass of a trained GCN. This is why several works study GST stability under perturbations of \mathbf{S} in order to understand the working mechanism of GCNs [7], [8], [14].

4 PRUNED GRAPH SCATTERING TRANSFORMS

While the representation power of GST increases with the number of layers, the computational and storage complexity of the transform also increase exponentially with the number of layers due to its scattering nature. Hence, even if informative features become available at deeper layers, the associated exponential complexity of extracting such features is prohibitive. On the other hand, various input data (e.g., social networks and 3D point clouds) may have distinct properties, leading to different GST feature maps. In some cases, just a few nodes of deep layers are informative; and in other cases, nodes of shallow layers convey most of the information; see Fig. 1. This motivates the pursuit of a tuned GST to adaptively capture informative nodes.

Aspiring to improve GST models, we introduce a *pruned* (p)GST to systematically retain informative nodes without extra

1. A tree node is fully specified by its corresponding path.

2. Any analytical function $h(\mathbf{S})$ can be written as a polynomial of \mathbf{S} with maximum degree $N - 1$ [21].




Fig. 2: Scattering pattern of a pGST with $J = 3$ and $L = 3$. Dashed lines represent pruned branches. An example \mathbf{x} and GFTs of filter banks are depicted too. The third filter $j = 3$ at $\ell = 1$ is pruned because it generates no output ($\mathbf{z}_{(3)} = \mathbf{0}$).

complexity. Our novel pGST reduces the exponential complexity and adapts GST to different graph data. Furthermore, pGST offers a practical mechanism to understand the architecture of GCNs. Based on the pruning patterns, the proposed pGST suggests when a deeper GCN is desirable, or, when a shallow one will suffice. Pruning the wavelet packets has been traditionally employed for compression in image processing applications [22], where the pruning is guided by a rate-distortion optimality criterion.

In this work, we consider a graph spectrum inspired criterion. Intuitively, each GST node will be associated with a unique subband of the graph spectrum. When the subband of a node does not have sufficient overlap with the graph signal spectrum, this node cannot capture the underlying graph signal, and should be pruned. Consider for specificity a smooth \mathbf{x} , where connected nodes have similar signal values. This implies sparse (low-rank) representation in the graph spectral domain; that is, $\hat{\mathbf{x}} := \mathbf{V}^\top \mathbf{x} \in \mathbb{R}^N$ and $[\hat{\mathbf{x}}]_n = 0$ for $n \geq b$. The graph spectrum of the j th output is then

$$\mathbf{V}^\top h_j(\mathbf{S})\mathbf{x} = \text{diag}(\hat{h}_j(\boldsymbol{\lambda}))\hat{\mathbf{x}} = [\hat{h}_j(\lambda_1)\hat{x}_1, \hat{h}_j(\lambda_2)\hat{x}_2, \dots, \hat{h}_j(\lambda_N)\hat{x}_N]^\top \quad (4)$$

where λ_n is the n th eigenvalue of \mathbf{S} and each frequency \hat{x}_n is weighted by the corresponding transformed eigenvalue $\hat{h}_j(\lambda_n)$. Hence, if the support of the spectrum $\{\hat{h}_j(\lambda_n)\}_n$ is not included in the support of $[\hat{\mathbf{x}}]_n$, the j th graph filter output will not capture any information; that is, $h_j(\mathbf{S})\mathbf{x} = \mathbf{0}_N$; see Fig. 2. Thus, identifying such graph filters and pruning the corresponding tree nodes will result in a parsimonious and thus computationally efficient GST.

4.1 Pruning criterion

Motivated by this last observation, we introduce a pruning criterion to select the scattering branches per node by maximizing the alignment between the graph spectrum of the filters and the scattering features. Per node p , the optimization problem is

$$\begin{aligned} \max_{\{f_{(p,j)}\}_{j=1}^J} & \sum_{j=1}^J \left(\sum_{n=1}^N (\hat{h}_j(\lambda_n)^2 - \tau) [\hat{\mathbf{z}}_{(p)}]_n^2 \right) f_{(p,j)} \quad (5) \\ \text{s. t.} & f_{(p,j)} \in \{0, 1\}, \quad j = 1, \dots, J \end{aligned}$$

where $\hat{\mathbf{z}}_{(p)} := \mathbf{V}\mathbf{z}_{(p)}$ is the graph spectrum of the scattering feature vector $\mathbf{z}_{(p)}$; τ is a user-specific threshold; and, $f_{(p,j)}$ stands for the pruning assignment variable indicating whether node (p, j) is active ($f_{(p,j)} = 1$), or, it should be pruned ($f_{(p,j)} = 0$). The objective in (5) promotes retaining tree nodes that maximize the alignment

of the graph spectrum of $\hat{\mathbf{z}}_{(p)}$ with that of $\hat{h}_j(\boldsymbol{\lambda})$. The threshold τ introduces a minimum spectral value to locate those nodes whose corresponding graph spectral response is small, meaning $\hat{h}_j(\lambda_n)^2 \ll \tau$. Note that criterion (5) is evaluated per tree node p , thus allowing for a flexible and scalable design.

The optimization problem in (5) is nonconvex since $f_{(p,j)}$ is a discrete variable. Furthermore, recovering $\hat{\mathbf{z}}_{(p)}$ requires an eigendecomposition of the Laplacian matrix that incurs complexity $\mathcal{O}(N^3)$. Nevertheless, by exploiting the structure in (5), we will develop an efficient pruning algorithm attaining the maximum of (5), as asserted by the following theorem.

Theorem 1. *The optimal $\{f_{(p,j)}^*\}_j$ are given by*

$$f_{(p,j)}^* = \begin{cases} 1 & \text{if } \frac{\|\mathbf{z}_{(p,j)}\|^2}{\|\mathbf{z}_{(p)}\|^2} > \tau, \\ 0 & \text{otherwise.} \end{cases}, \quad j = 1, \dots, J \quad (6)$$

The optimal variables $f_{(p,j)}^*$ are given by comparing the energy of the input $\mathbf{z}_{(p)}$ to that of the output $\mathbf{z}_{(p,j)}$ per graph filter j that can be evaluated with linear complexity of order $\mathcal{O}(N)$. Our pruning criterion leads to a principled and scalable means of selecting the GST nodes to be pruned. The pruning objective is evaluated per node p , and pruning decisions are made on-the-fly. Hence, when $f_{(p)}^* = 1$, node p is active and the graph filter bank will be applied to $\mathbf{z}_{(p)}$, expanding the tree to the next layer; otherwise, the GST will not be expanded further at node p , which can effect exponential savings in computations. An example of such a pruned tree is depicted in Fig. 2. Evidently, the hyperparameter τ controls the input-to-output energy ratio. A large τ corresponds to an aggressively pruned scattering tree, while a small τ amounts to a minimally pruned scattering tree. The stable and efficient representation extracted by the pGST is then defined as

$$\Psi(\mathbf{x}) := \{\phi_{(p)}\}_{p \in \mathcal{T}}$$

where \mathcal{T} is the set of active tree nodes $\mathcal{T} := \{p \in \mathcal{P} \mid f_{(p)}^* = 1\}$.

Our pruning approach provides a concise version of GSTs and effects savings in computations as well as memory. Although the worst-case complexity of pGST is still exponential, a desirable complexity can be realized by properly selecting τ . As a byproduct, the scattering patterns of pGSTs reveal the appropriate depths and widths of the GSTs for different graph data; see also Fig. 1. The pruning approach so far is an unsupervised one, since the labels are not assumed available. Note that the GFT is used throughout

for analytical purposes, and neither the pruning algorithm nor the pGST requires explicit calculation of the GFT.

4.2 Rate-distortion tradeoff

Our pruning objective is closely related to the optimal rate-distortion objective in signal processing [23], [24]. By judiciously selecting τ , the proposed pruning criterion in (6) finds also the optimal variables to the following problem

$$\begin{aligned} \max_{\{f_{(p,j)}\}_{j=1}^J} & \sum_{j=1}^J \frac{\|\mathbf{z}_{(p,j)}\|^2}{\|\mathbf{z}_{(p)}\|^2} f_{(p,j)} \\ \text{s. t.} & f_{(p,j)} \in \{0, 1\}, \quad j = 1, \dots, J \\ & \sum_{j=1}^J f_{(p,j)} \leq K \end{aligned} \quad (7)$$

where K denotes the maximum number of retained scattering features, that can be dictated by computational complexity constraints. The optimal solution to (7) is to set $f_{(p,j)} = 1$ for the K scattering features with larger energy ratio $\|\mathbf{z}_{(p,j)}\|^2 / \|\mathbf{z}_{(p)}\|^2$ and the rest to zero $f_{(p,j)} = 0$. Indeed, by properly choosing τ , the criterion in (6) can be employed to find the K channels with maximum energy, where τ would be the K th larger energy ratio. Such energy preservation objective has been employed in signal processing in the context of PCA wavelets [23]. Indeed, the hyperparameter τ reflects the sweet spot in the rate-distortion tradeoff [24].

5 STABILITY AND SENSITIVITY OF PGST

In this section, we study the stability and sensitivity of pGST when the input graph data and the network topology are perturbed. We will establish that pGST is stable in the presence of feature or network perturbations with bounded power. We will further analyze pGST sensitivity to input perturbations. To establish these results, we consider graph wavelets that form a frame with frame bounds A and B [25], meaning for $\mathbf{x} \in \mathbb{R}^N$, it holds that, $A^2 \|\mathbf{x}\|^2 \leq \sum_{j=1}^J \|h_j(\mathbf{S})\mathbf{x}\|^2 \leq B^2 \|\mathbf{x}\|^2$. In the graph vertex domain, A and B characterize the numerical stability of recovering \mathbf{x} from $\{h_j(\mathbf{S})\mathbf{x}\}_j$. In the graph spectral domain, they reflect the ability of a graph filter bank to amplify \mathbf{x} along each graph frequency. Tight frame bounds, satisfying $A^2 = B^2$, are of particular interest because such wavelets lead to enhanced numerical stability and faster computations [26]. The frame property of a graph wavelet plays an instrumental role in proving GST stability to perturbations of the underlying graph structure [7], [8], [14].

5.1 Stability to graph data perturbations

Perturbations present in \mathbf{x} may be attributed to modeling errors, or adversarial intervention aiming to poison learning. Consider the perturbed graph signal

$$\tilde{\mathbf{x}} := \mathbf{x} + \boldsymbol{\delta} \in \mathbb{R}^N \quad (8)$$

where $\boldsymbol{\delta} \in \mathbb{R}^N$ is the perturbation vector. We wish to study how and under what conditions our pGST is affected by such perturbations. A stable transformation should have a similar output under small input perturbations.

Before establishing that our pGST is stable, we first show that GST is stable to small perturbations in \mathbf{x} . Prior art deals mainly with GST stability to structure perturbations [7], [8], [14].

Lemma 1. Consider the GST $\Phi(\cdot)$ with L layers and J graph filters; and suppose that the graph filter bank forms a frame with bound B , while \mathbf{x} and $\tilde{\mathbf{x}}$ are related via (8). It then holds that

$$\frac{\|\Phi(\mathbf{x}) - \Phi(\tilde{\mathbf{x}})\|}{\sqrt{|\Phi(\mathbf{x})|}} \leq \sqrt{\frac{\sum_{\ell=0}^L (B^2 J)^\ell}{\sum_{\ell=0}^L J^\ell}} \|\boldsymbol{\delta}\|. \quad (9)$$

The squared difference of the GSTs is normalized by the number of scattering features in $\Phi(\cdot)$, that is $|\Phi(\mathbf{x})| = \sum_{\ell=1}^L J^\ell$. The bound in (9) relates to the frame bound of the wavelet filter bank. Clearly, for tight frames with $B = 1$, the normalized stability bound (9) is tight. Let $\tilde{\mathcal{T}}$ be the structure of the pruned tree for $\Psi(\tilde{\mathbf{x}})$. The following lemma asserts that the pGST offers the same pruned tree for the original and the perturbed inputs.

Lemma 2. Let $\Psi(\cdot)$ denote the pGST with L layers and J graph filters; $\tilde{\mathbf{z}}_p$ the perturbed scattering feature at node p ; and, $\boldsymbol{\delta}_p := \mathbf{z}_p - \tilde{\mathbf{z}}_p$. If for all $p \in \mathcal{P}$ and $j = 1, \dots, J$, we have

$$\left| \|h_j(\mathbf{S})\mathbf{z}_p\|^2 - \tau \|\mathbf{z}_p\|^2 \right| > \|h_j(\mathbf{S})\boldsymbol{\delta}_p\|^2 + \tau \|\mathbf{z}_p\|^2 - \|\tilde{\mathbf{z}}_p\|^2. \quad (10)$$

it then follows that

- i) The pruned scattering transform will output the same tree for $\Psi(\mathbf{x})$ and $\Psi(\tilde{\mathbf{x}})$; that is, $\mathcal{T} = \tilde{\mathcal{T}}$; and,
- ii) With $g(\mathbf{x}) := \|h_j(\mathbf{S})\mathbf{x}\|^2 - \tau \|\mathbf{x}\|^2$, a necessary condition for (10) is

$$|g(\mathbf{z}_p)| > g(\boldsymbol{\delta}_p). \quad (11)$$

According to (11), Lemma 2 can be interpreted as a signal-to-noise-ratio (SNR) condition because under $g(\boldsymbol{\delta}_p) > 0$, it is possible to write (11) as $|g(\mathbf{z}_p)|/g(\boldsymbol{\delta}_p) > 1$. Lemma 2 provides a per-layer and branch condition for pGST to output the same pruned scattering tree for the original or the perturbed signal. This condition is also experimentally validated in Sec. 6, where the structure in \mathcal{T} remains the same for Fig. 9b and Fig. 9a.

By combining Lemmas 4 and 2, we arrive at the following stability result for the pGST network.

Theorem 2. Consider the pGST transform $\Psi(\cdot)$ with L layers and J graph filters; and suppose that the graph filter bank forms a frame with bound B , while \mathbf{x} and $\tilde{\mathbf{x}}$ are related via (8). The pGST is stable to bounded perturbations $\boldsymbol{\delta}$, in the sense that

$$\frac{\|\Psi(\mathbf{x}) - \Psi(\tilde{\mathbf{x}})\|}{\sqrt{|\Psi(\mathbf{x})|}} \leq \sqrt{\frac{\sum_{\ell=0}^L F_\ell B^{2\ell}}{\sum_{\ell=0}^L F_\ell}} \|\boldsymbol{\delta}\| \quad (12)$$

where $F_\ell := |\mathcal{P}^{(\ell)} \cup \mathcal{T}|$ is the number of active scattering features at layer ℓ , and $|\Psi(\mathbf{x})| = \sum_{\ell=0}^L F_\ell$ the number of retained scattering features.

The bound in (12) is linear in the perturbation power $\|\boldsymbol{\delta}\|$, and hence pGST is stable to perturbations in the input.

5.2 Stability to structural perturbations

We next study the effect of network perturbations to the proposed pGST. The goal here is to establish that pGST is also stable under structural perturbations that can be due to noise or adversarial attacks [27], [28]. Vanilla scattering transforms are invariant to translations and stable to perturbations that resemble translations [4]. Likewise, GSTs are invariant to permutations and stable to perturbations that are close to permutations [7], [8]. In the

graph context, permutations are regarded as rearrangements of the vertex indices.

Consider a perturbed topology given by the $N \times N$ perturbed (and permuted) shift matrix $\tilde{\mathbf{S}}$. The set of permutations that make $\tilde{\mathbf{S}}$ close to \mathbf{S} are given by

$$\mathcal{P}_0 := \arg \min_{\mathbf{P} \in \mathcal{P}} \|\mathbf{P}^\top \tilde{\mathbf{S}} \mathbf{P} - \mathbf{S}\| \quad (13)$$

where \mathbf{P} is an $N \times N$ permutation matrix having entries 0 – 1. Further, consider the following set of perturbation matrices [8]

$$\mathcal{D} := \left\{ \mathbf{\Delta} : \mathbf{P}^\top \tilde{\mathbf{S}} \mathbf{P} = \mathbf{S} + \mathbf{\Delta}^\top \mathbf{S} + \mathbf{S} \mathbf{\Delta}, \mathbf{P} \in \mathcal{P}_0 \right\} \quad (14)$$

where $\mathbf{\Delta}$ denotes the $N \times N$ perturbation matrix. The error is captured by the term $\mathbf{\Delta}^\top \mathbf{S} + \mathbf{S} \mathbf{\Delta}$. Given the set of perturbation matrices in (14), we consider the distance between $\tilde{\mathbf{S}}$ and \mathbf{S} as

$$d(\mathbf{S}, \tilde{\mathbf{S}}) = \min_{\mathbf{\Delta} \in \mathcal{D}} \|\mathbf{\Delta}\|. \quad (15)$$

Without loss of generality, let $\mathbf{P} = \mathbf{I}$, otherwise fix a $\mathbf{P}_0 \in \mathcal{P}$ and define $\tilde{\mathbf{S}}$ to equal $\mathbf{P}_0^\top \tilde{\mathbf{S}} \mathbf{P}_0$. Hence, the perturbed topology can be written as³

$$\tilde{\mathbf{S}} = \mathbf{S} + \mathbf{\Delta}^\top \mathbf{S} + \mathbf{S} \mathbf{\Delta}. \quad (16)$$

Let $\tilde{\Psi}(\cdot)$ denote the pruned scattering transform that is based on the perturbed topology $\tilde{\mathbf{S}}$, and $\tilde{\mathcal{T}}$ be the structure of the pruned tree for $\tilde{\Psi}(\mathbf{x})$. The following lemma asserts that pGST outputs the same pruned tree for both the original and the perturbed graph.

Lemma 3. *Suppose \mathbf{S} and $\tilde{\mathbf{S}}$ satisfy $d(\mathbf{S}, \tilde{\mathbf{S}}) \leq \mathcal{E}/2$; and that for $\mathbf{\Delta} \in \mathcal{D}$ with eigendecomposition $\mathbf{\Delta} = \mathbf{U} \text{diag}(\mathbf{d}) \mathbf{U}^\top$ it holds that $\|\mathbf{\Delta}/d_{\max} - \mathbf{I}\| \leq \mathcal{E}$, where d_{\max} is the eigenvalue of $\mathbf{\Delta}$ with maximum absolute value. Suppose the graph filter bank forms a frame with bound B , and $h(\lambda)$ satisfies the integral Lipschitz constraint $|\lambda h'(\lambda)| \leq C_0$. Let $\tilde{\mathbf{z}}_p$ denote the perturbed scattering feature at the tree node p and $\delta_p := \mathbf{z}_p - \tilde{\mathbf{z}}_p$. If for all nodes $p \in \mathcal{P}$ of layer ℓ and $j = 1, \dots, J$, it holds that*

$$\left| \|h_j(\mathbf{S})\mathbf{z}_p\|^2 - \tau \|\mathbf{z}_p\|^2 \right| > (\ell \mathcal{E} C_0 B^{\ell-1} \|\mathbf{x}\|)^2 + \tau \|\delta_p\|^2 \quad (17)$$

the pruned scattering transform will then output the same tree for $\Psi(\mathbf{x})$ and $\tilde{\Psi}(\mathbf{x})$; that is, $\mathcal{T} = \tilde{\mathcal{T}}$.

Lemma 3 establishes conditions under which the pGST based on the original and perturbed graphs will output the same scattering tree. The assumption on the eigenvalues of the perturbation, that is $\|\mathbf{\Delta}/d_{\max} - \mathbf{I}\| \leq \mathcal{E}$, limits the structural changes in the graph such as edge insertions or deletions; see also [8]. The integral Lipschitz constraint $|\lambda h'(\lambda)| \leq C_0$ is an additional stability requirement on the wavelet $h(\cdot)$. Note that these conditions are not required to establish stability to input noise in Theorem 2. By combining Lemma 3 with Proposition 3 in [8], we can show the stability of pGST to structural perturbations.

Theorem 3. *Under the conditions of Lemma 3 it holds that*

$$\frac{\|\Psi(\mathbf{x}) - \tilde{\Psi}(\mathbf{x})\|}{\sqrt{|\Psi(\mathbf{x})|}} \leq \mathcal{E} C_0 \sqrt{\frac{\sum_{\ell=0}^L F_\ell \ell^2 B^{2\ell}}{\sum_{\ell=0}^L F_\ell}} \|\mathbf{x}\| \quad (18)$$

where $F_\ell := |\mathcal{P}^{(\ell)} \cup \mathcal{T}|$ is the number of active scattering features at layer ℓ , and $|\Psi(\mathbf{x})| = \sum_{\ell=0}^L F_\ell$ the number of retained scattering features.

3. Unless it is stated otherwise the ℓ_2 norms are employed.

The bound in (18) is linear in the perturbation \mathcal{E} , which establishes the stability of pGST. As the bound depends also linearly in the Lipschitz constant C_0 and the number of layers L , these terms also appear in the stability results for traditional GST [8]. For tight frames with $B = 1$, the bound in (18) can be further simplified. By contrasting the bounds in (12) and (18), we deduce that (12) is tighter than (18) by a factor ℓ . This follows because the topology perturbation in (18) implies a perturbed wavelet that is repeatedly applied to the scattering features. Hence, the error in (18) is accumulated layer by layer, giving rise to the ℓ factor in (18). On the other hand, for (12) the perturbation is introduced only at the first layer $\ell = 0$. Theorems 2 and 3 can be combined to guarantee stability under *joint* perturbations of both the input and the topology.

5.3 Sensitivity to graph data perturbations

Next, the sensitivity of pGST is analyzed for different types of noises in input signals⁴. Although the stability of a transform is unquestionably important, pGST should be sensitive to signal perturbation that have some specific pattern. For example, signal perturbations may be distributed over all the nodes but are localized in the graph spectrum.

Specifically, two classes of input signal perturbations are considered, each having distinct graph spectral properties.

Highly localized noise. Here the energy of the perturbation signal is localized in the graph spectrum; that is,

$$\mathbf{E}_L^{(\epsilon)} = \{\delta_L \in \mathbb{R}^N \mid \hat{\delta}_L = \mathbf{V} \delta_L, \|\hat{\delta}_L\| \leq \epsilon, \|\hat{\delta}_L\|_0 = 1\}. \quad (19)$$

Random noise. Here the energy of a perturbation signal is uniformly spread over the graph spectrum; that is,

$$\mathbf{E}_R^{(\epsilon)} = \{\delta_R \in \mathbb{R}^N \mid \hat{\delta}_R = \mathbf{V} \delta_R, [\hat{\delta}_R]_n^2 \leq \epsilon/N\}. \quad (20)$$

The following corollary establishes the sensitivity bound for these classes of perturbation signals.

Corollary 1. *Let $\Psi(\cdot)$ be a pGST with L layers and $J = N$ filter. The j th filter has a transform function $h_j(\mathbf{S})$ with $\hat{h}_j(\lambda_n) = 1$ for $n = j$, and 0 for $n \neq j$. Consider two types of the perturbed inputs: $\tilde{\mathbf{x}}_L := \mathbf{x} + \delta_L$, where $\delta_L \in \mathbf{E}_L$, and $\tilde{\mathbf{x}}_R := \mathbf{x} + \delta_R$, where $\delta_R \in \mathbf{E}_R$. It then holds that*

$$\begin{aligned} \|\Psi(\mathbf{x}) - \Psi(\tilde{\mathbf{x}}_R)\| &\leq \sqrt{\left(\sum_{\ell=1}^L \frac{F_\ell B^{2\ell}}{N} + 1 \right)} \epsilon \\ \|\Psi(\mathbf{x}) - \Psi(\tilde{\mathbf{x}}_L)\| &\leq \sqrt{\left(\sum_{\ell=2}^L F'_\ell B^{2\ell} + 2 \right)} \epsilon \end{aligned}$$

where $F_\ell := |\mathcal{P}^{(\ell)} \cup \mathcal{T}_R|$ and $F'_\ell := |\mathcal{P}^{(\ell)} \cup \mathcal{T}_L|$ are the retained features in the cases of random and localized noise.

Corollary 1 suggests that pGST is more sensitive to δ_L than to δ_R . This is appealing because pGST can detect transient changes in the graph spectral domain while ignoring small random perturbations. These results will be numerically corroborated in Sec. 6. Specifically, Fig. 9 shows that relative to the original \mathcal{T} , the

4. The sensitivity analysis for structural perturbations does not provide interesting results; see the experimental validation in Section 6.3.2. Highly localized signal perturbation would affect only a few scattering filters; however, arbitrary structural perturbation would affect all scattering filters $\{h_j(\mathbf{S})\}_j$. Furthermore, different from data perturbations, the perturbed graph introduces errors at each GST layer. This makes the analysis non-trivial.

pruned tree changes considerably when the noise is localized in the input spectrum, while the tree remains the same for random noise.

6 EXPERIMENTS

This section evaluates the performance of our pGST in various graph classification tasks. Graph classification amounts to predicting a label y_i given \mathbf{x}_i and \mathbf{S}_i for the i th graph. Our pGST extracts the efficient and stable representation $\Psi(\mathbf{x}_i)$, which is utilized as a feature vector for predicting y_i . During training, the structure of the pGST \mathcal{T} is determined, which is kept fixed during validation and testing. The parameter τ is selected via cross-validation. Our goal is to provide tangible answers to the following research questions.

- RQ1** How does the proposed pGST compare to GST?
- RQ2** How does pGST compare to state-of-the-art GCNs in graph-based classification tasks?
- RQ3** Given graph data, what are pruned scattering patterns?
- RQ4** What is the impact of signal perturbations in the scattering patterns?
- RQ5** Can the effective GCN depth linked with the pGST depth?

6.1 Pruned GST compared to traditional GST

To address RQ1, we reproduce the experiments of two tasks in [8]: authorship attribution and source localization. For the scattering transforms, we consider three implementations of graph filter banks: the diffusion wavelets (DS) in [14], the monic cubic wavelets (MCS) in [25] and the tight Hann wavelets (THS) in [26].⁵ The scattering transforms use $J = 5$ filters, $L = 5$ layers, and $\tau = 0.01$. The extracted features from GSTs are subsequently utilized by a linear support vector machine (SVM) classifier.

Authorship attribution amounts to determining if a certain text was written by a specific author. Each text is represented by a graph with $N = 244$, where words (nodes) are connected based on their relative positions in the text, and \mathbf{x} is a bag-of-words representation of the text; see also [14]. Fig. 3 (a) reports the classification accuracy as the number of training samples (texts) increases. GSTs utilize $\sum_{\ell=1}^5 5^\ell = 781$ scattering coefficients, while pGSTs rely only on $|\mathcal{T}| = 61$ for PDS, $|\mathcal{T}| = 30$ for PMCS, and $|\mathcal{T}| = 80$ for PTHS. Evidently, pGST achieves comparable performance as the baseline GST, whereas pGST uses only a subset of features (12.8%, 3.8% and 10.2%, respectively). The SVM classifier provides a coefficient that weighs each scattering scalar. The magnitude of each coefficient shows the importance of the corresponding scattering feature in the classification. Fig. 3 (c) depicts the percentage of features retained after pruning in the top-2 $|\mathcal{T}|$ most important GST features given by the SVM classifier. It is observed, that although pGST does not take into account the labels, the retained features are indeed informative for classification.

Source localization amounts to recovering the source of a rumor given a diffused signal over a Facebook subnetwork with $N = 234$; see the detailed settings in [14]. Fig. 3 (b) shows the classification accuracy of the scattering transforms as the SNR (in dB) increases. In accordance to Lemma 1 and Theorem 2, both pGST and GST are stable over a wide range of SNR. Furthermore, the performance of pGST matches that of GST, while the pGST uses only a subset of features. Finally, Fig. 4 depicts the runtime of the different scattering approaches, where the computational advantage of the pruned methods is evident.

5. PDS, PMCS, PTHS denote the pruned versions of these transforms.

Dataset	Graphs	Features F	Max N per graph
Collab	5000	1	492
D&D	1178	89	5748
Enzymes	600	3	126
Proteins	1113	3	620

TABLE 1: Dataset characteristics

	Method	Data Set			
		ENZYMES	D&D	COLLAB	PROTEINS
Kernel	SHORTEST-PATH	42.32	78.86	59.10	76.43
	WL-OA	60.13	79.04	80.74	75.26
	PATCHYSAN	–	76.27	72.60	75.00
	GRAPHSAGE	54.25	75.42	68.25	70.48
GNNs	ECC	53.50	74.10	67.79	72.65
	SET2SET	60.15	78.12	71.75	74.29
	SORTPOOL	57.12	79.37	73.76	75.54
	DIFFPOOL-DET	58.33	75.47	82.13	75.62
Scattering	DIFFPOOL-NO LP	62.67	79.98	75.63	77.42
	DIFFPOOL	64.23	81.15	75.50	78.10
	GSC	53.88	76.57	76.88	74.03
	GST	59.84	79.28	77.32	76.23
	PGST (Ours)	60.25	81.27	78.40	78.57

TABLE 2: Graph classification accuracy.

6.1.1 Ablation study

Fig. 5 reports how the pGST is affected by varying the threshold τ in the task of source localization, with $J = 6$ and $L = 5$. Fig. 5 (a) shows the classification accuracy that generally decreases as τ increases since the number of active features $|\mathcal{T}|$ decreases; cf. Fig. 5 (b). Fig. 5 (c) reports the runtime in seconds. Fig. 6, 7 showcase the classification performance of pGST with $\tau = 0.01$ for varying L with $J = 3$, and for varying J with $L = 3$, respectively. It is observed, that the classification performance generally improves with L and J .

6.2 Pruned GST for graph-based classification tasks

In response to RQ2, we consider three graph-based classification tasks: graph classification, graph-based 3D point cloud classification, and semi-supervised node classification.

6.2.1 Graph classification

We compare the proposed pGST with the following state-of-the-art methods⁶. The kernel methods shortest-path [29], and Weisfeiler-Lehman optimal assignment (WL-OA) [30]; the deep learning approaches PatchySan [31], GraphSage [32], edge-conditioned filters in CCNs (ECC) [33], Set2Set [34], SortPool [35], and DiffPool [13]; and the geometric scattering classifier (GSC) [6]. Results are presented on the protein data sets D&D, Enzymes and Proteins, and the scientific collaboration data set Collab. The parameters of these datasets are listed in Table 1. Since the Collab dataset did not have nodal features, \mathbf{x} was selected as the vector that holds the node degrees. We performed 10-fold cross validation and report the classification accuracy averaged over the 10 folds. The gradient boosting classifier is employed for pGST and GST with parameters chosen based on the performance on the validation set.

6. For the competing approaches we report the 10-fold cross-validation numbers reported by the original authors; see also [13].




Fig. 3: Classification accuracy against number of samples in the authorship attribution (a) and SNR in dB for source localization (b). The percentage of features after pruning retained in the top-2| \mathcal{T} | most important GST features given by the SVM classifier (c).




Fig. 4: Runtime comparison of the scattering transforms for Fig. 3.

The graph scattering transforms use the MC wavelet with $L = 5$, $J = 5$, and $\tau = 0.01$. Table 2 lists the classification accuracy of the proposed and competing approaches. Even without any training on the feature extraction step, the performance of pGST is comparable to the state-of-the-art deep supervised learning approaches across all datasets. GST and pGST outperform also GSC, since the latter uses a linear SVM to classify the scattering features.

6.2.2 Point cloud classification

We further test pGST in classifying 3D point clouds. Given a point cloud, a graph can be created by connecting points (nodes) to their nearest neighbors based on their Euclidian distance. Each node is also associated with 6 scalars denoting its x-y-z coordinates and RGB colors. For this experiment, GSTs are compared against PointNet++ [36], [37], 3dShapeNets [38] and VoxNet [39], that are state-of-the-art deep learning approaches. Fig. 8 reports the classification accuracy for the ModelNet40 dataset [38] for increasing L . In Fig. 8 (a) 9,843 clouds are used for training and 2,468 for testing using the gradient boosting classifier; whereas, in Fig. 8 (b) only 615 clouds are used for training and the rest for testing using a fully connected neural network classifier with 3 layers. The scattering transforms use an MC wavelet with $J = 5$ for Fig. 8 (a), and $J = 9$ for Fig. 8 (b). Fig. 8 showcases that scattering transforms are competitive to state-of-the-art approaches, while pGST outperforms GST. This may be attributed to overfitting effects, since a large number of GST features is not informative. Furthermore, the exponential complexity of GSTs prevents their application with $L = 6$. Fig. 8 (b) shows that when the training data are scarce, GST and pGST outperform the PointNet++, which requires a large number of training data to optimize over the network parameters.

6.2.3 Semi-supervised node classification

A task of major importance at the intersection of machine learning and network science is semi-supervised learning (SSL) over graphs.

Method	Cora
ManiReg [40]	59.5
SemiEmb [41]	59.0
LP [42]	68.0
Planetoid [43]	75.7
GCN [10]	81.5
GAT [11]	83.0
GST [7]	81.9
pGST (ours)	81.9

TABLE 3: SSL classification accuracy.

Given the topology \mathbf{S} , the features in the $N \times F$ matrix \mathbf{X} , and labels only at a subset \mathcal{L} of nodes $\{y_n\}_{n \in \mathcal{L}}$ with $\mathcal{L} \subset \mathcal{V}$, the goal of is to predict the labels $\{y_n\}_{n \in \mathcal{U}}$ of the unlabeled set of nodes. For this task, we utilize the pGST algorithm to extract the feature vectors $\{\mathbf{z}_{(p)}\}_{p \in \mathcal{T}}$ from \mathbf{X} , which are then utilized by a fully connected neural network with 3 layers to predict the missing labels. To facilitate comparison, we reproduce the experimental setup of [10], namely the same split of the data for training, validation, and testing sets, and compare the various methods in the Cora dataset with $N = 2,708$ nodes, $C = 7$ classes and $F = 1,433$ features. The pGST algorithm employs the TH wavelet with $L = 5$, $J = 3$, and $\tau = 0.01$. We also compare to [7], where the GST is employed, the full scattering coefficients are extracted and then dimensionality reduction is effected using PCA to handle the large number of features. The so obtained features are processed by a fully connected network. Table 3 reports the classification accuracy of several state-of-the-art methods. Our method that performs unsupervised feature extraction is highly competitive and outperforms all except the GAT approach. It is worth noting that pGST performs comparably to the GST approach in [7], which suggests that our spectrum-inspired criterion has a ‘PCA-like’ effect. However, PCA needs all the scattering features to process them offline, while pGST prunes the scattering features on-the-fly, and allows for increased scalability.

6.3 Scattering patterns analysis

To address RQ3, we depict the scattering structures of pGSTs, with an MC wavelet, $J = 3$, and $L = 5$, for the Collab, Proteins, and ModelNet40 datasets in Fig. 1. Evidently, graph data from various domains require an adaptive scattering architecture. Specifically, most tree nodes for the academic collaboration dataset are pruned, and hence most informative features reside at the shallow layers. This is consistent with the study in [2], which experimentally shows

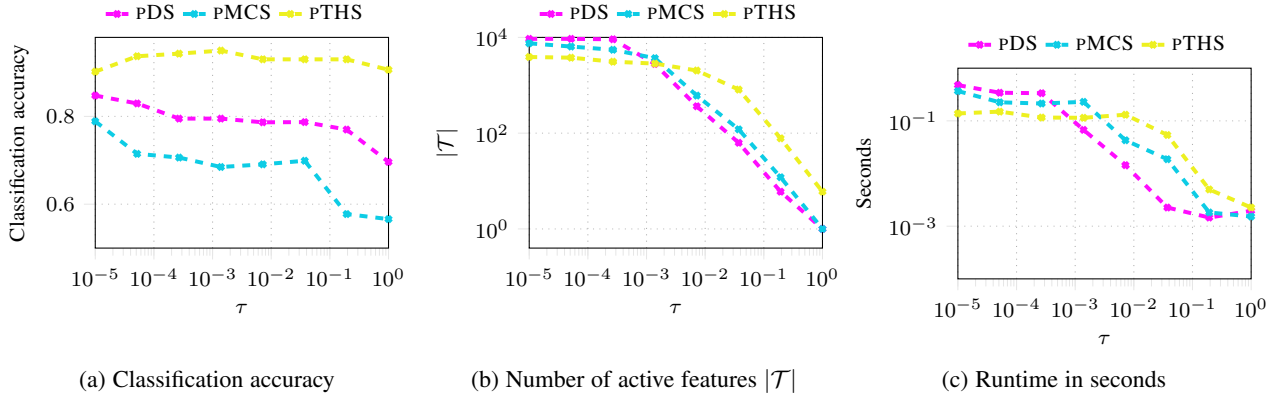


Fig. 5: Performance of pGSTs for varying τ .




Fig. 6: Classification accuracy over L




Fig. 7: Classification accuracy over J

that deeper GCNs do not contribute as much for social network data. These findings are further supported by the small-world phenomenon in social networks, which suggests that the diameter of social networks is small [44]. On the other hand, the tree nodes for a 3D point cloud are minimally pruned, which is in line with the work in [3] that showcases the advantage of deep GCNs in 3D point clouds classification.

6.3.1 Sensitivity of the pGST tree to signal perturbations

Next, we evaluate noise effects in the pruning algorithm to provide an answer to RQ4. Corollary 1 suggests that localized noise will have a more prominent effect on the pGST than random noise. This is appealing because pGST can detect transient changes in the graph spectral domain, while ignoring small random perturbations. Here, we attempt to establish the changes to the *structure* in the scattering patterns of pGST under localized and random noise.

For this experiment, the Cora dataset is tested with $N = 2,708$ nodes, $C = 7$ classes and $F = 1,433$ features. We draw random and localized noise from (20) and (19) respectively, which is added to the nodal features. The noise is selected so that the signal to noise (SNR) energy ratio is -20dB . We observe that for smaller noise power, the scattering patterns are invariant, which corroborates the theoretical result in Lemma 2.

Fig. 9 plots the scattering patterns of the pGST using the TH wavelet with $L = 3$, $J = 5$, and $\tau = 0.1$. It is observed that the set of pruned features given the original features without noise in Fig. 9a, is almost the same as the set of pruned features given the randomly perturbed features in Fig 9b. On the other hand, the set of pruned features when the localized noise is employed, is quite different; see Fig. 9c.

6.3.2 Sensitivity of the pGST tree to structural perturbations

For the same setting as Fig. 9 the effect of the structural perturbations in the structure of pGST is also tested. The adjacency matrix was perturbed by adding noise in the eigenvalues such that the SNR is -20dB . Random noise in this scenario means that the noise is distributed over the eigenvalues, whereas localized noise affects only a subset of the eigenvalues significantly. Fig. 10 shows that the localized and random noise do not alter the scattering patterns significantly. This is in accordance with the stability results in Theorem 3.

6.3.3 Scattering patterns insights for GCN design

Finally, we empirically explore the connections between PGST and GCN design, in response to RQ5. We validate the scattering patterns in Fig. 1 for the Protein datasets. The validation employs DiffPool that is a state-of-the-art GCN; see [13].

Fig. 11 depicts the pruned scattering patterns for the three protein datasets (Enzyme, Protein and DD). Most connections after $l = 5$ are pruned. This suggests that for protein data $L = 5$ graph convolution layers capture most information.

Fig. 12 shows the performance of DiffPool for these datasets as the number of GCN layers increases. The performance of DiffPool does not improve significantly for more than 5 GCN layers, which corroborates the insights obtained from the pGST patterns.

7 CONCLUSIONS

This paper developed a novel approach to pruning the graph scattering transform. The proposed pGST relies on a graph-spectrum-based data-adaptive criterion to prune less informative features

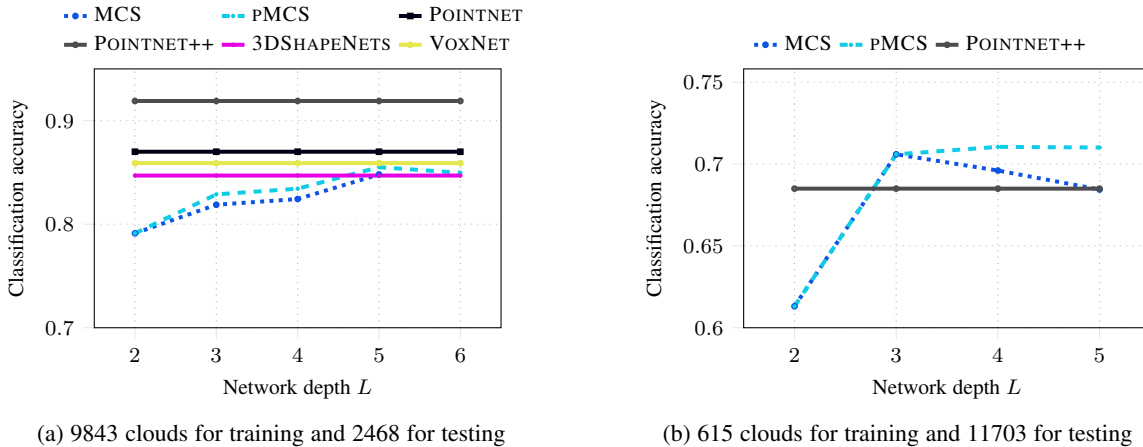


Fig. 8: 3D point cloud classification.




Fig. 9: The pGST applied to the Cora dataset with and without signal perturbation, corroborating the result of Corollary 1.





Fig. 10: The pGST applied to the Cora dataset with and without structural perturbations.

Fig. 11: Illustration of the pGST applied to three protein datasets. Observe that after $l = 5$ most features are pruned.

on-the-fly, and effectively reduce the computational complexity of GSTs. Stability of pGST is established in the presence of perturbations in the input or the network structure. Sensitivity analysis of pGST reveals that the transform is more sensitive to noise that is localized in the graph spectrum relative to noise that is uniformly spread over the spectrum, which is appealing since pGST can detect transient changes in the graph spectral domain. The novel pGST extracts efficient and stable features of graph data. Experiments demonstrated i) the performance gains of pGSTs relative to GSTs; ii) that pGST is competitive in a variety of graph classification tasks; and (iii) graph data from different domains exhibit unique pruned scattering patterns, which calls for adaptive network architectures.

Acknowledgments. This work was fully supported by Mitsubishi Electric Research Laboratories.

8 PROOFS

8.1 Proof of Theorem 1

By its definition, the objective in (5) can be rewritten as

$$\begin{aligned} & \sum_{j=1}^J \left(\sum_{n=1}^N (\hat{h}_j(\lambda_n)^2 - \tau) [\hat{\mathbf{z}}_{(p)}]_n^2 \right) f_j \\ &= \sum_{j=1}^J \hat{\mathbf{z}}_{(p)}^\top \left(\text{diag}(\hat{h}_j(\boldsymbol{\lambda}))^2 - \tau \mathbf{I} \right) \hat{\mathbf{z}}_{(p)} f_j. \end{aligned}$$




Fig. 12: Performance of DiffPool for different number of GCN layers.

By introducing the scalars $\alpha_j := \widehat{\mathbf{z}}_{(p)}^\top (\text{diag}(\widehat{h}_j(\boldsymbol{\lambda}))^2 - \tau \mathbf{I}) \widehat{\mathbf{z}}_{(p)}$ for $j = 1, \dots, J$, (5) can be rewritten as

$$\begin{aligned} \max_{f_j} \quad & \sum_{j=1}^J \alpha_j f_j \\ \text{s. t.} \quad & f_j \in \{0, 1\}, \quad j = 1, \dots, J. \end{aligned} \quad (21)$$

The optimization problem in (21) is nonconvex since f_j is a discrete variable. However, maximizing the sum in (21) amounts to setting $f_j = 1$ for the positive α_j over j . Such an approach leads to the optimal pruning assignment variables

$$f_j^* = \begin{cases} 1 & \text{if } \alpha_j > 0, \\ 0 & \text{if } \alpha_j < 0. \end{cases}, \quad j = 1, \dots, J \quad (22)$$

The rest of the proof focuses on rewriting α_j as

$$\alpha_j = \widehat{\mathbf{z}}_{(p)}^\top (\text{diag}(\widehat{h}_j(\boldsymbol{\lambda}))^2 - \tau \mathbf{I}) \widehat{\mathbf{z}}_{(p)} \quad (23)$$

$$= \|\text{diag}(\widehat{h}_j(\boldsymbol{\lambda})) \widehat{\mathbf{z}}_{(p)}\|^2 - \tau \|\widehat{\mathbf{z}}_{(p)}\|^2 \quad (24)$$

Furthermore, since matrix \mathbf{V} is orthogonal, it holds that $\|\widehat{\mathbf{z}}_{(p)}\|^2 = \|\mathbf{V}^\top \mathbf{z}_{(p)}\|^2 = \|\mathbf{z}_{(p)}\|^2$ from which it follows that

$$\|\text{diag}(\widehat{h}_j(\boldsymbol{\lambda})) \widehat{\mathbf{z}}_{(p)}\|^2 = \|h_j(\mathbf{S}) \mathbf{z}_{(p)}\|^2 \quad (25)$$

$$\begin{aligned} &= \|\sigma(h_j(\mathbf{S}) \mathbf{z}_{(p)})\|^2 \\ &= \|\mathbf{z}_{(p,j)}\|^2 \end{aligned} \quad (26)$$

where the second line follows because $\sigma(\cdot)$ is applied elementwise, and does not change the norm.

8.2 Proof of Lemma 1

By the definition in (3), it holds that

$$\|\Phi(\mathbf{x}) - \Phi(\tilde{\mathbf{x}})\|^2 = \sum_{\ell=0}^L \sum_{p^{(\ell)} \in \mathcal{P}^{(\ell)}} |\phi_{(p^{(\ell)})} - \tilde{\phi}_{(p^{(\ell)})}|^2 \quad (27)$$

which suggests bounding each summand in (27) as

$$|\phi_{(p^{(\ell)})} - \tilde{\phi}_{(p^{(\ell)})}| = |U(\mathbf{z}_{(p^{(\ell)})}) - U(\tilde{\mathbf{z}}_{(p^{(\ell)})})| \quad (28)$$

$$\leq \|U\| \|\mathbf{z}_{(p^{(\ell)})} - \tilde{\mathbf{z}}_{(p^{(\ell)})}\| \quad (29)$$

where (29) follows since the norm is a sub-multiplicative operator. Next, we will show the recursive bound

$$\begin{aligned} & \|\mathbf{z}_{(p^{(\ell)})} - \tilde{\mathbf{z}}_{(p^{(\ell)})}\| \\ &= \|\sigma(h_{j^{(\ell)}}(\mathbf{S}) \mathbf{z}_{(p^{(\ell-1)})}) - \sigma(h_{j^{(\ell)}}(\mathbf{S}) \tilde{\mathbf{z}}_{(p^{(\ell-1)})})\| \end{aligned} \quad (30)$$

$$\leq \|\sigma(\cdot)\| \|h_{j^{(\ell)}}(\mathbf{S}) \mathbf{z}_{(p^{(\ell-1)})} - h_{j^{(\ell)}}(\mathbf{S}) \tilde{\mathbf{z}}_{(p^{(\ell-1)})}\| \quad (31)$$

$$\leq \|h_{j^{(\ell)}}(\mathbf{S}) \mathbf{z}_{(p^{(\ell-1)})} - h_{j^{(\ell)}}(\mathbf{S}) \tilde{\mathbf{z}}_{(p^{(\ell-1)})}\| \quad (32)$$

$$\leq \|h_{j^{(\ell)}}(\mathbf{S})\| \|\mathbf{z}_{(p^{(\ell-1)})} - \tilde{\mathbf{z}}_{(p^{(\ell-1)})}\| \quad (33)$$

where (31), (33) hold because the norm is a sub-multiplicative operator, and (32) follows since the nonlinearity is nonexpansive, i.e. $\|\sigma(\cdot)\| < 1$. Hence, by applying (33) $\ell - 1$ times, the following condition holds

$$\|\mathbf{z}_{(p^{(\ell)})} - \tilde{\mathbf{z}}_{(p^{(\ell)})}\| \leq \|h_{j^{(\ell)}}(\mathbf{S})\| \|h_{j^{(\ell-1)}}(\mathbf{S})\| \cdots \|h_{j^{(1)}}(\mathbf{S})\| \|\mathbf{x} - \tilde{\mathbf{x}}\| \quad (34)$$

and by further applying the frame bound and (8), we deduce that

$$\|\mathbf{z}_{(p^{(\ell)})} - \tilde{\mathbf{z}}_{(p^{(\ell)})}\| \leq B^\ell \|\delta\| \quad (35)$$

Combining (29), (35) and the average operator property $\|U\| = 1$ it holds that

$$|\phi_{(p^{(\ell)})} - \tilde{\phi}_{(p^{(\ell)})}| \leq B^\ell \|\delta\| \quad (36)$$

By applying the bound (36) for all entries in the right hand side of (27) it follows that

$$\|\Phi(\mathbf{x}) - \Phi(\tilde{\mathbf{x}})\|^2 \leq \sum_{\ell=0}^L \sum_{p^{(\ell)} \in \mathcal{P}^{(\ell)}} B^{2\ell} \|\delta\|^2 \quad (37)$$

By factoring out $\|\delta\|$ and observing that the sum in the right side of (37) does not depend on the path index p it follows that

$$\|\Phi(\mathbf{x}) - \Phi(\tilde{\mathbf{x}})\|^2 \leq \left(\sum_{\ell=0}^L |\mathcal{P}^{(\ell)}| B^{2\ell} \right) \|\delta\|^2 \quad (38)$$

Finally, since the cardinality of the paths at ℓ is $|\mathcal{P}^{(\ell)}| = J^\ell$ and $\sum_{\ell=0}^L (B^2 J)^\ell = ((B^2 J)^L) / (B^2 J - 1)$ it holds

$$\|\Phi(\mathbf{x}) - \Phi(\tilde{\mathbf{x}})\| \leq \sqrt{\frac{(B^2 J)^L}{B^2 J - 1}} \|\delta\| \quad (39)$$

8.3 Proof of Lemma 2

We will prove the case for $\ell = 0$, where $\mathbf{z}_{p^{(0)}} = \mathbf{x}$, since the same proof holds for any ℓ . First, we adapt (10) to the following

$$|\|h_j(\mathbf{S})\mathbf{x}\|^2 - \tau\|\mathbf{x}\|^2| > \|h_j(\mathbf{S})\boldsymbol{\delta}\|^2 + \tau\|\mathbf{x}\|^2 - \|\tilde{\mathbf{x}}\|^2. \quad (40)$$

The proof will examine two cases and will follow by contradiction. For the first case, consider that branch j is pruned in $\Psi(\mathbf{x})$ and not pruned in $\Psi(\tilde{\mathbf{x}})$, i.e. $(j) \notin \mathcal{T}$ and $(j) \in \tilde{\mathcal{T}}$. By applying (6) for $\mathbf{z}_{(j)} = \sigma(h_j(\mathbf{S})\mathbf{x})$ there exists $C \geq 0$ such that

$$\frac{\|h_j(\mathbf{S})\mathbf{x}\|^2}{\|\mathbf{x}\|^2} \leq \tau - C \quad (41)$$

$$\|h_j(\mathbf{S})\mathbf{x}\|^2 \leq \tau\|\mathbf{x}\|^2 - C\|\mathbf{x}\|^2 \quad (42)$$

Furthermore, from (6) it holds for $\tilde{\mathbf{z}}_{(j)} = \sigma(h_j(\mathbf{S})\tilde{\mathbf{x}})$ that

$$\frac{\|h_j(\mathbf{S})\tilde{\mathbf{x}}\|^2}{\|\tilde{\mathbf{x}}\|^2} > \tau \quad (43)$$

By applying (8) to (43), and using the triangular inequality it follows that

$$\|h_j(\mathbf{S})\mathbf{x}\|^2 + \|h_j(\mathbf{S})\boldsymbol{\delta}\|^2 \geq \tau\|\tilde{\mathbf{x}}\|^2 \quad (44)$$

Next, by applying (42) it holds that

$$\tau\|\mathbf{x}\|^2 - C\|\mathbf{x}\|^2 + \|h_j(\mathbf{S})\boldsymbol{\delta}\|^2 \geq \tau\|\tilde{\mathbf{x}}\|^2 \quad (45)$$

$$\tau(\|\mathbf{x}\|^2 - \|\tilde{\mathbf{x}}\|^2) + \|h_j(\mathbf{S})\boldsymbol{\delta}\|^2 \geq C\|\mathbf{x}\|^2. \quad (46)$$

Next, by utilizing (40) and the absolute value property $|a| \geq a$ to upper-bound the left side of (46) it follows that

$$\|h_j(\mathbf{S})\mathbf{x}\|^2 - \tau\|\mathbf{x}\|^2 > C\|\mathbf{x}\|^2. \quad (47)$$

Finally, by applying (42) the following is obtained

$$0 > 2C\|\mathbf{x}\|^2 \quad (48)$$

which implies that $C < 0$. However, this contradicts (41) since $C \geq 0$. Following a similar argument we can complete the proof for the other case.

8.4 Proof of Lemma 3

The proof of Lemma 3 requires the following result.

Lemma 4. Consider the shift matrix \mathbf{S} and the perturbed matrix $\tilde{\mathbf{S}}$, such that $d(\mathbf{S}, \tilde{\mathbf{S}}) \leq \mathcal{E}/2$. Further, for $\boldsymbol{\Delta} \in \mathcal{D}$ with eigendecomposition $\boldsymbol{\Delta} = \mathbf{U} \text{diag}(\mathbf{d})\mathbf{U}^\top$ it holds that $\|\boldsymbol{\Delta}/d_{\max} - \mathbf{I}\| \leq \mathcal{E}$, where d_{\max} is the eigenvalue of $\boldsymbol{\Delta}$ with maximum absolute value. At layer ℓ consider the scattering feature indexed by $p^{(\ell)}$ of the original transform $\Psi(\mathbf{x})$ as $\mathbf{z}_{(p^{(\ell)})}$ and of the perturbed transform $\tilde{\Psi}(\mathbf{x})$ as $\tilde{\mathbf{z}}_{(p^{(\ell)})}$. Suppose also that graph filter bank forms a frame with bound B , and $h(\lambda)$ satisfies the integral Lipschitz constraint $|\lambda h'(\lambda)| \leq C_0$. It then holds that

$$\|\tilde{\mathbf{z}}_{(p^{(\ell)})} - \mathbf{z}_{(p^{(\ell)})}\| \leq \ell \mathcal{E} C_0 B^{\ell-1} \|\mathbf{x}\|. \quad (49)$$

Proof. First, we add and subtract on (49) the term $\sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\tilde{\mathbf{S}}) \cdots \sigma(h_{j^{(1)}}(\tilde{\mathbf{S}})\mathbf{x}) \cdots)$

$$\|\tilde{\mathbf{z}}_{(p^{(\ell)})} - \mathbf{z}_{(p^{(\ell)})}\| \quad (50)$$

$$\begin{aligned} &\leq \|\sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\tilde{\mathbf{S}}) \cdots \sigma(h_{j^{(1)}}(\tilde{\mathbf{S}})\mathbf{x}) \cdots) \\ &\quad - \sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\mathbf{S}) \cdots \sigma(h_{j^{(1)}}(\mathbf{S})\mathbf{x}) \cdots) \\ &\quad - \sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\tilde{\mathbf{S}}) \cdots \sigma(h_{j^{(1)}}(\tilde{\mathbf{S}})\mathbf{x}) \cdots) \\ &\quad + \sigma(h_{j^{(\ell)}}(\tilde{\mathbf{S}})\sigma(h_{j^{(\ell-1)}}(\tilde{\mathbf{S}}) \cdots \sigma(h_{j^{(1)}}(\tilde{\mathbf{S}})\mathbf{x}) \cdots)\| \end{aligned} \quad (51)$$

$$\begin{aligned} &\leq \|\sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\tilde{\mathbf{S}}) \cdots \sigma(h_{j^{(1)}}(\tilde{\mathbf{S}})\mathbf{x}) \cdots) \\ &\quad - \sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\mathbf{S}) \cdots \sigma(h_{j^{(1)}}(\mathbf{S})\mathbf{x}) \cdots)\| \\ &\quad + \|\sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\tilde{\mathbf{S}}) \cdots \sigma(h_{j^{(1)}}(\tilde{\mathbf{S}})\mathbf{x}) \cdots) \\ &\quad - \sigma(h_{j^{(\ell)}}(\tilde{\mathbf{S}})\sigma(h_{j^{(\ell-1)}}(\tilde{\mathbf{S}}) \cdots \sigma(h_{j^{(1)}}(\tilde{\mathbf{S}})\mathbf{x}) \cdots)\| \end{aligned} \quad (52)$$

By a similar argument as in (34), it holds for the first summand in (52) that

$$\begin{aligned} &\|\sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\tilde{\mathbf{S}}) \cdots \sigma(h_{j^{(1)}}(\tilde{\mathbf{S}})\mathbf{x}) \\ &\quad - \sigma(h_{j^{(\ell)}}(\mathbf{S})\sigma(h_{j^{(\ell-1)}}(\mathbf{S}) \cdots \sigma(h_{j^{(1)}}(\mathbf{S})\mathbf{x})\| \\ &\leq B\|\tilde{\mathbf{z}}_{(p^{(\ell-1)})} - \mathbf{z}_{(p^{(\ell-1)})}\| \end{aligned} \quad (53)$$

For the second summand in (52), we consider the following

$$\begin{aligned} \|\mathbf{z}_{(p^{(\ell)})}\| &\leq \|h_{j^{(\ell)}}(\mathbf{S})\| \|h_{j^{(\ell-1)}}(\mathbf{S})\| \cdots \|h_{j^{(1)}}(\mathbf{S})\| \|\mathbf{x}\| \\ &\leq B^\ell \|\mathbf{x}\| \end{aligned} \quad (54)$$

which can be derived following (34). Also by utilizing the result of Proposition 2 in [8] it can be shown that

$$\|h_j(\mathbf{S}) - h_j(\tilde{\mathbf{S}})\| \leq \mathcal{E} C_0. \quad (55)$$

where C_0 is the integral Lipschitz constant $|\lambda h'(\lambda)| \leq C_0$. By combining (53)-(55) we arrive to the following recursive condition

$$\|\tilde{\mathbf{z}}_{(p^{(\ell)})} - \mathbf{z}_{(p^{(\ell)})}\| \leq \|\tilde{\mathbf{z}}_{(p^{(\ell-1)})} - \mathbf{z}_{(p^{(\ell-1)})}\| + (\ell - 1)\mathcal{E} C_0 B^{\ell-1} \|\mathbf{x}\|$$

which can be solved to arrive at (49). \square

Now that Lemma 4 is established we proceed with the proof for Lemma 3 follows. The proof will examine two cases and will follow by contradiction. For the first case, consider that branch j is pruned in $\Psi(\mathbf{x})$ and not pruned in $\tilde{\Psi}(\mathbf{x})$, i.e. $(j) \notin \mathcal{T}$ and $(j) \in \tilde{\mathcal{T}}$. By applying (6) for $\mathbf{z}_{(p,j)} = \sigma(h_j(\mathbf{S})\mathbf{z}_{(p)})$ there exists $C \geq 0$ such that

$$\frac{\|h_j(\mathbf{S})\mathbf{z}_{(p)}\|^2}{\|\mathbf{z}_{(p)}\|^2} \leq \tau - C \quad (56)$$

$$\|h_j(\mathbf{S})\mathbf{z}_{(p)}\|^2 \leq \tau\|\mathbf{z}_{(p)}\|^2 - C\|\mathbf{z}_{(p)}\|^2 \quad (57)$$

Furthermore, from (6) it holds for $\tilde{\mathbf{z}}_{(p,j)} = \sigma(h_j(\tilde{\mathbf{S}})\tilde{\mathbf{z}}_{(p)})$ that

$$\frac{\|h_j(\tilde{\mathbf{S}})\tilde{\mathbf{z}}_{(p)}\|^2}{\|\tilde{\mathbf{z}}_{(p)}\|^2} > \tau \quad (58)$$

By using the triangle inequality to (58), we obtain

$$\|h_j(\mathbf{S})\mathbf{z}_{(p)}\|^2 + \|h_j(\tilde{\mathbf{S}})\tilde{\mathbf{z}}_{(p)} - h_j(\mathbf{S})\mathbf{z}_{(p)}\|^2 \geq \tau\|\tilde{\mathbf{z}}_{(p)}\|^2 \quad (59)$$

and upon applying (57), we arrive at

$$\tau(\|\mathbf{z}_{(p)}\|^2 - \|\tilde{\mathbf{z}}_{(p)}\|^2) + \|h_j(\tilde{\mathbf{S}})\tilde{\mathbf{z}}_{(p)} - h_j(\mathbf{S})\mathbf{z}_{(p)}\|^2 \geq C\|\mathbf{z}_{(p)}\|^2. \quad (60)$$

Next, we leverage Lemma 4 and the absolute value property $|a| \geq a$ to upper-bound the left side of (60) as

$$\tau\|\boldsymbol{\delta}_{(p)}\|^2 + (\ell \mathcal{E} C_0 B^{\ell-1} \|\mathbf{x}\|)^2 > C\|\mathbf{z}_{(p)}\|^2. \quad (61)$$

By applying (17) in Lemma 3, we deduce that

$$\|h_j(\mathbf{S})\mathbf{z}_p\|^2 - \tau\|\mathbf{z}_p\|^2 > C\|\mathbf{z}_p\|^2. \quad (62)$$

Finally, by applying (57), we find

$$0 > 2C\|\mathbf{z}_p\|^2 \quad (63)$$

which implies that $C < 0$. However, this contradicts (56) since $C \geq 0$. Following a symmetric argument, we can complete the proof for the other case too.


REFERENCES

- [1] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, “Geometric deep learning: going beyond euclidean data,” *IEEE Sig. Process. Mag.*, vol. 34, no. 4, pp. 18–42, 2017.
- [2] F. Wu, T. Zhang, A. H. d. Souza Jr, C. Fifty, T. Yu, and K. Q. Weinberger, “Simplifying graph convolutional networks,” in *Proc. Intl. Conf. Mach. Learn. (ICML)*, 2019, pp. 6861–6871.
- [3] G. Li, M. Müller, A. Thabet, and B. Ghanem, “DeepGCNs: Can GCNs go as deep as CNNs?” in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 9266–9275.
- [4] J. Bruna and S. Mallat, “Invariant scattering convolution networks,” *IEEE Trans. Pattern Anal. Mach. Intel.*, vol. 35, no. 8, pp. 1872–1886, 2013.
- [5] S. Mallat, “Group invariant scattering,” *Communications on Pure and Applied Mathematics*, vol. 65, no. 10, pp. 1331–1398, 2012.
- [6] F. Gao, G. Wolf, and M. Hirn, “Geometric scattering for graph data analysis,” in *Proc. Intl. Conf. Mach. Learn. (ICML)*, 2019, pp. 2122–2131.
- [7] D. Zou and G. Lerman, “Graph convolutional neural networks via scattering,” *Applied and Computational Harmonic Analysis*, vol. 49, no. 3, pp. 1046–1074, Nov. 2020.
- [8] F. Gama, J. Bruna, and A. Ribeiro, “Stability of graph scattering transforms,” in *Proc. Advances Neural Inf. Process. Syst. (NeurIPS)*, 2019, pp. 8036–8046.
- [9] V. N. Ioannidis, S. Chen, and G. B. Giannakis, “Pruned graph scattering transforms,” in *Proc. Intl. Conf. on Learn. Representations (ICLR)*, 2020.
- [10] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *Proc. Intl. Conf. on Learn. Representations (ICLR)*, Toulon, France, Apr. 2017.
- [11] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, “Graph attention networks,” in *Proc. Intl. Conf. on Learn. Representations (ICLR)*, 2018.
- [12] V. N. Ioannidis, A. G. Marques, and G. B. Giannakis, “A recurrent graph neural network for multi-relational data,” in *Proc. IEEE Int. Conf. Acoust., Speech, Sig. Process.*, London, England, May 2019, pp. 8157–8161.
- [13] Z. Ying, J. You, C. Morris, X. Ren, W. Hamilton, and J. Leskovec, “Hierarchical graph representation learning with differentiable pooling,” in *Proc. Advances Neural Inf. Process. Syst. (NeurIPS)*, 2018, pp. 4805–4815.
- [14] F. Gama, A. Ribeiro, and J. Bruna, “Diffusion scattering transforms on graphs,” in *Proc. Intl. Conf. on Learn. Representations (ICLR)*, 2019.
- [15] R. Reed, “Pruning algorithms—a survey,” *IEEE Trans. on Neural Networks*, vol. 4, no. 5, pp. 740–747, 1993.
- [16] S. Han, H. Mao, and W. J. Dally, “Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding,” in *Proc. Intl. Conf. on Learn. Representations (ICLR)*, 2016.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Proc. Advances Neural Inf. Process. Syst. (NeurIPS)*, 2012, pp. 1097–1105.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE Conf. on Comp. Vis. and Pat. Rec.*, 2016, pp. 770–778.
- [19] N. Lee, T. Ajanthan, and P. H. Torr, “Snip: Single-shot network pruning based on connection sensitivity,” in *Proc. Intl. Conf. on Learn. Representations (ICLR)*, 2019.
- [20] R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec, “Graph convolutional neural networks for web-scale recommender systems,” in *Proc. Intl. Conf. on Knowledge Disc. and Data Mining (KDD)*. ACM, 2018, pp. 974–983.
- [21] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge university press, 2012.
- [22] Z. Xiong, K. Ramchandran, and M. T. Orchard, “Space-frequency quantization for wavelet image coding,” *IEEE Trans. Image Process.*, vol. 6, no. 5, pp. 677–693, 1997.
- [23] M. K. Tsatsanis and G. B. Giannakis, “Principal component filter banks for optimal multiresolution analysis,” *IEEE Trans. Sig. Process.*, vol. 43, no. 8, pp. 1766–1777, 1995.
- [24] K. Ramchandran and M. Vetterli, “Best wavelet packet bases in a rate-distortion sense,” *IEEE Trans. Sig. Process.*, vol. 2, pp. 160–175, 1993.
- [25] D. K. Hammond, P. Vandergheynst, and R. Gribonval, “Wavelets on graphs via spectral graph theory,” *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129–150, 2011.
- [26] D. I. Shuman, C. Wiesmeyer, N. Holighaus, and P. Vandergheynst, “Spectrum-adapted tight graph wavelet and vertex-frequency frames,” *IEEE Trans. Sig. Process.*, vol. 63, no. 16, pp. 4223–4235, 2015.
- [27] D. Zügner, A. Akbarnejad, and S. Günnemann, “Adversarial attacks on kernels and applications to graph data,” in *Proc. Intl. Conf. on Knowledge Disc. and Data Mining (KDD)*, 2018, pp. 2847–2856.
- [28] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” in *Proc. Intl. Conf. on Learn. Representations (ICLR)*, 2014.
- [29] K. M. Borgwardt and H.-P. Kriegel, “Shortest-path kernels on graphs,” in *IEEE Intl. Conf. on Data Mining (ICDM)*, 2005, pp. 74–81.
- [30] N. M. Kriege, P.-L. Giscard, and R. Wilson, “On valid optimal assignment kernels and applications to graph classification,” in *Proc. Advances Neural Inf. Process. Syst. (NeurIPS)*, 2016, pp. 1615–1623.
- [31] M. Niepert, M. Ahmed, and K. Kutzkov, “Learning convolutional neural networks for graphs,” in *Proc. Intl. Conf. Mach. Learn. (ICML)*, 2016, pp. 2014–2023.
- [32] W. L. Hamilton, R. Ying, and J. Leskovec, “Representation learning on graphs: Methods and applications,” *IEEE Data Eng. Bull.*, vol. 40, no. 3, pp. 52–74, 2017.
- [33] M. Simonovsky and N. Komodakis, “Dynamic edge-conditioned filters in convolutional neural networks on graphs,” in *IEEE Conf. on Comp. Vis. and Pat. Rec.*, 2017, pp. 29–38.
- [34] O. Vinyals, S. Bengio, and M. Kudlur, “Order matters: Sequence to sequence for sets,” *Proc. Intl. Conf. on Learn. Representations (ICLR)*, 2015.
- [35] M. Zhang, Z. Cui, M. Neumann, and Y. Chen, “An end-to-end deep learning architecture for graph classification,” in *AAAI Conference on Artificial Intelligence*, 2018, pp. 4438–4445.
- [36] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation,” in *IEEE Conf. on Comp. Vis. and Pat. Rec.*, 2017, pp. 77–85.
- [37] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “Pointnet++: Deep hierarchical feature learning on point sets in a metric space,” in *Proc. Advances Neural Inf. Process. Syst. (NeurIPS)*, 2017, pp. 5099–5108.
- [38] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, “3D shapenets: A deep representation for volumetric shapes,” in *IEEE Conf. on Comp. Vis. and Pat. Rec.*, 2015, pp. 1912–1920.
- [39] D. Maturana and S. Scherer, “Voxnet: A 3d convolutional neural network for real-time object recognition,” in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 922–928.
- [40] M. Belkin, P. Niyogi, and V. Sindhwani, “Manifold regularization: A geometric framework for learning from labeled and unlabeled examples,” *J. Mach. Learn. Res.*, vol. 7, pp. 2399–2434, 2006.
- [41] J. Weston, F. Ratle, H. Mobahi, and R. Collobert, “Deep learning via semi-supervised embedding,” in *Neural Networks: Tricks of the Trade*. Springer, 2012, pp. 639–655.
- [42] X. Zhu, Z. Ghahramani, and J. D. Lafferty, “Semi-supervised learning using gaussian fields and harmonic functions,” in *Proc. Intl. Conf. Mach. Learn. (ICML)*, Washington, USA, Jun. 2003, pp. 912–919.
- [43] Z. Yang, W. W. Cohen, and R. Salakhutdinov, “Revisiting semi-supervised learning with graph embeddings,” in *Proc. Intl. Conf. Mach. Learn. (ICML)*, New York, USA, Jun. 2016, pp. 40–48.
- [44] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *Nature*, vol. 393, no. 6684, p. 440, 1998.



Vassilis N. Ioannidis (S'16) received his Diploma in Electrical and Computer Engineering from the National Technical University of Athens, Greece, in 2015, and the M.Sc. degree in Electrical Engineering from the University of Minnesota (UMN), Twin Cities, Minneapolis, MN, USA, in 2017. Currently, he is working towards the Ph.D. degree in the Department of Electrical and Computer Engineering at the University of Minnesota, Twin Cities. Vassilis received the Doctoral Dissertation Fellowship (2019-2020) from

the University of Minnesota. He also received Student Travel Awards from the IEEE Signal Processing Society (2017,2018) and from the IEEE (2018). From 2014 to 2015, he worked as a middleware consultant for Oracle in Athens, Greece, and received a Performance Excellence award. His research interests include machine learning, graph neural networks, and network science.



Siheng Chen is a research scientist at Mitsubishi Electric Research Laboratories (MERL). Before that, he was an autonomy engineer at Uber Advanced Technologies Group, working on the perception and prediction systems of self-driving cars. Before joining Uber, he was a postdoctoral research associate at Carnegie Mellon University. Chen received the doctorate in Electrical and Computer Engineering from Carnegie Mellon University in 2016, where he also received two masters degrees in Electrical and Computer Engineering and Machine Learning, respectively. He received his bachelor's degree in Electronics Engineering in 2011 from Beijing Institute of Technology, China. Chen was the recipient of the 2018 IEEE Signal Processing Society Young Author Best Paper Award. His coauthored paper received the Best Student Paper Award at IEEE GlobalSIP 2018. His research interests include graph signal processing, graph neural networks and 3D computer vision.

engineering and Machine Learning, respectively. He received his bachelor's degree in Electronics Engineering in 2011 from Beijing Institute of Technology, China. Chen was the recipient of the 2018 IEEE Signal Processing Society Young Author Best Paper Award. His coauthored paper received the Best Student Paper Award at IEEE GlobalSIP 2018. His research interests include graph signal processing, graph neural networks and 3D computer vision.



G. B. Giannakis (Fellow'97) received his Diploma in Electrical Engr. from the Ntl. Tech. Univ. of Athens, Greece, 1981. From 1982 to 1986 he was with the Univ. of Southern California (USC), where he received his MSc. in Electrical Engineering, 1983, MSc. in Mathematics, 1986, and Ph.D. in Electrical Engr., 1986. He was a faculty member with the University of Virginia from 1987 to 1998, and since 1999 he has been a professor with the Univ. of Minnesota, where he holds an ADC Endowed Chair, a University of

Minnesota McKnight Presidential Chair in ECE, and serves as director of the Digital Technology Center.

His general interests span the areas of statistical learning, communications, and networking - subjects on which he has published more than 450 journal papers, 750 conference papers, 25 book chapters, two edited books and two research monographs. Current research focuses on Data Science, and Network Science with applications to the Internet of Things, brain, and power networks with renewables. He is the (co-) inventor of 32 patents issued, and the (co-) recipient of 9 best journal paper awards from the IEEE Signal Processing (SP) and Communications Societies, including the G. Marconi Prize Paper Award in Wireless Communications. He also received the IEEE-SPS Nobert Wiener Society Award (2019); Technical Achievement Awards from the SP Society (2000) and from EURASIP (2005); the IEEE ComSoc Education Award (2019); and the IEEE Fourier Technical Field Award (inaugural recipient in 2015). He is a Fellow of the National Academy of Inventors, IEEE, and EURASIP, and has served the IEEE in a number of posts, including that of a Distinguished Lecturer for the IEEE-SPS.