

Generalizing Graph Convolutional Neural Networks with Edge-Variant Recursions on Graphs

Elvin Isufi, Fernando Gama, and Alejandro Ribeiro

Abstract—This paper reviews graph convolutional neural networks (GCNNs) through the lens of edge-variant graph filters. The edge-variant graph filter is a finite order, linear, and local recursion that allows each node, in each iteration, to weigh differently the information of its neighbors. By exploiting this recursion, we put forth a general framework for GCNNs which considers state-of-the-art solutions as particular cases. This framework results useful to *i)* understand the tradeoff between local detail and the number of parameters of each solution and *ii)* provide guidelines for developing a myriad of novel approaches that can be implemented locally in the vertex domain. One of such approaches is presented here showing superior performance w.r.t. current alternatives in graph signal classification problems.

Index Terms—Graph convolutional neural networks; graph signal processing; graph filters; edge-variant.

I. INTRODUCTION

Graph convolutional neural networks (GCNNs) are gaining momentum in addressing a variety of classification and regression tasks for data that live in irregular spaces, such as network data [1]. With the aim to replicate the success of traditional CNNs, GCNNs play a fundamental role in semi-supervised learning on graphs and classifying different graph signals (i.e., values indexed by the nodes of the graph).

A central aspect of GCNNs is the extension of the convolution operation to graph signals. The seminal work [2] defined convolution as the point-wise multiplication in the graph spectral domain between the projected graph signal in this domain and the transfer function of a learnable filter. Such an approach finds solid grounds on graph signal processing (GSP) theory [3], where the graph spectrum plays the role of the Fourier basis for graph signals. Subsequently, several works exploited this connection and proposed computationally lighter GCNN models. In particular, [4]–[7] used the so-called polynomial graph filters [3], while [8], [9] relied on graph filters having a rational transfer function [10]. Differently, [11] designed architectures by using the node-variant (NV) graph filters [12], while [13] introduced MIMO approaches to learn from multiple features.

While the above works introduce techniques that extend CNNs to graphs, they are mostly based on (well-motivated) analogies with classical neural networks. Such a strategy presents, however, its own limitations towards extending these methods to more involved ones that better exploit the graph.

The main aim of this paper is to provide a general framework that unifies state-of-the-art GCNN architectures facilitating comparison, showing the limitations, and highlighting

their potential. For such a goal, it explores the so-called edge-variant (EV) graph filters [14], [15]. The EV graph filter is a local, linear, and finite-order recursion in the node domain where each node weighs differently, in each iteration, the information in its neighborhood. Therefore, it presents the most general linear and local operation that a node can do –gather information from all neighbors and weight each of them differently. This local property puts the EV as a computationally efficient candidate (only local information is exchanged) for capturing detail at the node connection level.

Nevertheless, in the general form, the learnable parameters of the EV graph filter depend on the number of graph edges. To tackle the latter, we first show how state-of-the-art solutions fall under the EV recursion and how they impose parsimonious models on the learnable parameters, rendering their number independent of the graph dimensions. Then, we explore such insights to provide guidelines for designing a variety of novel architectures in the spirit of the EV recursion whose number of parameters is independent of the graph dimensions.

In a nutshell, the contributions of this paper are: *i)* To formulate a general framework for GCNNs through EV graph filters. *ii)* To show how the state-of-the-art approaches are specific parameterizations of this EV recursion. *iii)* To present rigorous design guidelines for GCNNs based on the EV recursion that preserve locality and whose number of parameters is independent of the graph dimension. *iv)* To introduce one new such an architecture and show its superior performance for graph signal classification tasks.

II. BACKGROUND

A. Graphs and graph filters

Let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ be a weighted graph with vertex set $\mathcal{V} = \{1, \dots, N\}$ of cardinality $|\mathcal{V}| = N$ and edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ composed of $|\mathcal{E}| = M$ ordered pairs $(i, j) \in \mathcal{E}$ iff there is an edge between nodes i and j . For each node i , we define the neighborhood set $\mathcal{N}_i = \{j : (j, i) \in \mathcal{E}\}$ as the set of nodes connected to i . The sparsity of the edge set of \mathcal{G} is represented by an $N \times N$ matrix \mathbf{S} , named the graph shift operator matrix, where $[\mathbf{S}]_{i,j} \neq 0$ if $(j, i) \in \mathcal{E}$ or $i = j$. Candidates for \mathbf{S} are the graph adjacency \mathbf{W} , the graph Laplacian matrix (undirected graphs) or any of their normalizations. For generality, in the sequel, we will focus on directed graphs.

Along with \mathcal{G} , consider a set of signal values (features) $\mathbf{x} = [x_1, \dots, x_N]^T \in \mathbb{R}^N$ in which component x_i resides on node i . By exploiting the coupling between the graph \mathcal{G} and the graph signal \mathbf{x} , it is possible to compute a graph harmonic analysis for \mathbf{x} similarly to the one performed for temporal and image signals. Specifically, given the eigendecomposition $\mathbf{S} = \mathbf{U}\mathbf{A}\mathbf{U}^{-1}$, the graph Fourier transform (GFT) of \mathbf{x} is

The authors are with the Department of Electrical & Systems Engineering, University of Pennsylvania, 19104 Philadelphia, United States. This work is supported by NSF CCF 1717120, ARO W911NF1710438, ARL DCIST CRA W911NF-17-2-0181, ISTC-WAS and Intel DevCloud. E-mails: {eisufi, fgama, aribeiro}@seas.upenn.edu.

$\hat{\mathbf{x}} = \mathbf{U}^{-1}\mathbf{x}$. Likewise, the inverse transform is $\mathbf{x} = \mathbf{U}\hat{\mathbf{x}}$. Here, \mathbf{U} contains along the columns the oscillating modes of the graph and $\hat{\mathbf{x}}$ are the respective Fourier coefficients. The eigenvalues Λ represent the spectral support for $\hat{\mathbf{x}}$ and are commonly referred to as the graph frequencies [3].

Given the Fourier expansion, we can now filter \mathbf{x} directly in the spectral domain. That is, for $h : \mathbb{C} \rightarrow \mathbb{R}$ being the filter spectral response (transfer function), the filter output $\hat{\mathbf{y}} = h(\Lambda)\hat{\mathbf{x}}$ is computed as the convolution (pointwise multiplication in the spectral domain) between the filter transfer function $h(\Lambda)$ and the GFT of \mathbf{x} . By means of the inverse GFT, the vertex domain output becomes

$$\mathbf{y} = \mathbf{H}(\mathbf{S})\mathbf{x} \quad \text{and} \quad \mathbf{H}(\mathbf{S}) = \mathbf{U}h(\Lambda)\mathbf{U}^{-1}. \quad (1)$$

However, (1) is not local, since in computing the output y_i , node i needs access to the graph signal of non-neighboring nodes. To account for the locality, we can define the filtering operation directly in the vertex domain as the aggregation of neighboring information. The node i output for an order one local filter is

$$y_i = \sum_{j \in \mathcal{N}_i \cup i} \phi_{i,j} x_j \quad (2)$$

where the scalar parameter $\phi_{i,j}$ weighs the information of the neighboring node j . Nevertheless, this direct vertex domain definition does not enjoy a spectral behavior analysis limiting the connection with the convolution operation.

One way to link the spectral and the vertex domain filtering is to consider the polynomial graph filters [3] with output

$$\begin{aligned} \mathbf{y} &= \phi_0 \mathbf{I}_N \mathbf{x} + \phi_1 \mathbf{S} \mathbf{x} + \dots + \phi_K \mathbf{S}^K \mathbf{x} \\ &\triangleq \phi_0 \mathbf{w}^{(0)} + \phi_1 \mathbf{w}^{(1)} + \dots + \phi_K \mathbf{w}^{(K)}. \end{aligned} \quad (3)$$

Due to the locality of \mathbf{S} , \mathbf{y} can be obtained in the vertex domain through local information exchange. The state $\mathbf{w}^{(0)}$ is simply the graph signal, while $\mathbf{w}^{(1)} = \mathbf{S}\mathbf{x}$ consists of one-hop information exchange between adjacent nodes. The higher order states $\mathbf{w}^{(k)} = \mathbf{S}^k \mathbf{x}$ are computed recursively as $\mathbf{w}^{(k)} = \mathbf{S}\mathbf{w}^{(k-1)}$, i.e., by exchanging with the neighbors the previous intermediate state $\mathbf{w}^{(k-1)}$. Such an implementation amounts for a complexity of order $\mathcal{O}(MK)$. By means of the GFT, the filtering operation in (3) has the transfer function

$$h(\Lambda) = \sum_{k=0}^K \phi_k \Lambda^k. \quad (4)$$

Therefore, we conclude that the output \mathbf{y} in (3) consists of the convolution between a graph filter with a polynomial transfer function (4) and \mathbf{x} . This filter enjoys a local implementation and captures detail in a neighborhood of radius K from the node.

B. Edge-variant graph filters

The edge-variant graph filter is a finite order recursion implemented in the vertex domain in the form similar to (2) [15]. Let $\{\Phi^{(k)}\}_{k=1}^K \in \mathbb{R}^{N \times N}$ be a collection of K matrices that share the sparsity pattern of $\mathbf{I}_N + \mathbf{S}$. The intermediate states of the EV filter are computed recursively as $\mathbf{w}^{(1)} = \Phi^{(1)}\mathbf{x}$; $\mathbf{w}^{(2)} = \Phi^{(2)}\mathbf{w}^{(1)} = \Phi^{(2)}\Phi^{(1)}\mathbf{x}$; and

$$\mathbf{w}^{(k)} = \Phi^{(k)}\mathbf{w}^{(k-1)} = \Phi^{(k)}\Phi^{(k-1)} \dots \Phi^{(2)}\Phi^{(1)}\mathbf{x}. \quad (5)$$

Since $\Phi^{(1)}$ shares the support with $\mathbf{S} + \mathbf{I}_N$, the state $\mathbf{w}^{(1)}$ accounts also for the scaling of \mathbf{x} through the diagonal elements (i.e., each node i scales its own signal with a different parameter $\phi_{i,i}^{(1)}$). The higher order states $\mathbf{w}^{(k)}$ for $k > 1$ are again obtained recursively since the parameter matrices $\Phi^{(k)}$ respect the graph structure. Put differently, each $\Phi^{(k)}$ considers a different parameter $\phi_{i,j}^{(k)}$ for each edge $(j, i) \in \mathcal{E}$ and adds potential self-loops through $\phi_{i,i}^{(k)}$. Node i computes then the k th order state as

$$w_i^{(k)} = \sum_{j \in \mathcal{N}_i \cup i} \phi_{i,j}^{(k-1)} w_j^{(k-1)}. \quad (6)$$

By defining $\Phi^{(k:1)} = \prod_{\kappa=1}^k \Phi^{(\kappa)}$ and putting back together all terms in (5), the output of an order K EV graph filter is

$$\mathbf{y} = \sum_{k=0}^K \Phi^{(k:1)} \mathbf{x} = \sum_{k=1}^K \overbrace{\left(\prod_{\kappa=1}^k \Phi^{(\kappa)} \right)}^{\mathbf{H}(\mathbf{S})} \mathbf{x}. \quad (7)$$

The total number of parameters of the edge-variant graph filter is $K(M + N)$, which is in general smaller than the N^2 parameters of an arbitrary linear transform $\mathbf{y} = \mathbf{H}\mathbf{x}$. In computing the output \mathbf{y} in (7) the EV graph filter incurs in an overall computational complexity of order $\mathcal{O}(K(M + N))$ which is similar to that of (3) since in general $M \approx N$ or we can consider an EV recursion without self-loops.

The ability to capture local detail at the edge level and the reduced implementation complexity is leveraged next to define graph neural networks (GNN) with a controlled number of parameters and computational complexity matched to the sparsity pattern of the graph.

III. EDGE-VARIANT GRAPH NEURAL NETWORKS

Consider a training set $\mathcal{T} = \{\mathbf{x}_i, \mathbf{z}_i\}_{i=1}^{|\mathcal{T}|}$ composed of $|\mathcal{T}|$ examples of inputs $\mathbf{x}_i \in \mathcal{X}$ and output representations $\mathbf{z}_i \in \mathcal{Z}$. A GNN leverages the underlying graph representation of the data and learns a model $f(\cdot; \mathbf{S}) : \mathcal{X} \rightarrow \mathcal{Z}$ such that $\tilde{\mathbf{z}}_i = f(\mathbf{x}_i; \mathbf{S})$ minimizes some loss function $\mathcal{L}(\tilde{\mathbf{z}}_i, \mathbf{z}_i)$ for $\mathbf{z}_i \in \mathcal{T}$ and generalizes well for $\mathbf{z}_i \notin \mathcal{T}$.

To capture several levels of detail, model $f(\cdot, \mathbf{S})$ is layered into the cascade of L functions $f(\cdot, \mathbf{S}) = f^L(\cdot, \mathbf{S}) \circ f^{L-1}(\cdot, \mathbf{S}) \circ \dots \circ f^1(\cdot, \mathbf{S})$ each consisting of a succession of linear transforms and nonlinearities. Layer l produces as output a collection of F_l higher level signal features obtained through processing the F_{l-1} features $\mathbf{z}_1^{l-1}, \dots, \mathbf{z}_{F_{l-1}}^{l-1}$ computed at the previous layer. The f th higher level feature is computed as

$$\mathbf{z}_f^l = \sigma^l \left(\sum_{g=1}^{F_{l-1}} \mathbf{H}_{f,g}^l(\mathbf{S}) \mathbf{z}_g^{l-1} \right) \quad (8)$$

where $\sigma^l(\cdot)$ represents the nonlinearity that might be pointwise (e.g., ReLU) or graph dependent [16] and $\mathbf{H}_{f,g}^l(\mathbf{S})$ leverages the graph structure to relate the g th input feature \mathbf{z}_g^{l-1} to the f th output feature \mathbf{z}_f^l .

The graph basically serves as a parameterization to reduce both the computational complexity and the number of parameters. GCNNs, in particular, consider $\mathbf{H}_{f,g}^l(\mathbf{S})$ to be a graph

filter that has a spectral interpretation such as (1) or (4). In the sequel, we consider $\mathbf{H}_{f,g}^l(\mathbf{S})$ to be an EV graph filter and show that current approaches represent different parameterizations to induce the spectral convolution into the EV GNN.

A. Properties of the edge-variant neural network layer

From (7) and (8), we build the EV GNN layer as

$$\mathbf{H}_{f,g}^l(\mathbf{S}) = \sum_{k=1}^K \Phi_{f,g}^{l,(k:1)} \quad (9)$$

characterized by the following properties.

First, it does not require the knowledge of \mathbf{S} , but only of its support. This is because, differently from current solutions, it will learn from the training data \mathcal{T} a collection of K parameter matrices $\Phi^{(k)}$, where each of them acts as a different graph shift operator. Therefore, the EV GNN represents a robust learning strategy for data residing on graphs whose edge weights are known up to some uncertainty, known only partially, or not known at all, such as biological networks [17].

Second, the computational complexity of each layer is linear in the graph parameters. By setting, $F = \max_i F_i$, the overall complexity of the EV layer is of order $\mathcal{O}(KF^2(M+N))$ matching that of current state-of-the-art GCNN approaches.

Third, the number of parameters per layer is, at most, $KF^2(M+N)$. The latter, although allowing the EV to have the maximum degrees of freedom given a topology, may often be a limitation for large graphs or when $|\mathcal{T}|$ is small. Our goal in the next section is, therefore, to show that GCNN layers proposed in the literature are particular cases of (9). Establishing these relationships allows the proposal of novel solutions that increase the descriptive power while preserving an efficient implementation complexity.

B. Parametrizations

Polynomial GCNN. Several variants of GCNNs introduced in the literature use at each layer graph filters of the form

$$\mathbf{H}(\mathbf{S}) = \sum_{k=0}^K \phi_k \mathbf{S}^k \quad (10)$$

where in the filter definition $\mathbf{H}(\mathbf{S}) = \mathbf{H}_{f,g}^l(\mathbf{S})$ we omitted the layer and feature indices to simplify notation. This is the case of [6], [7] which use general polynomials and [4], [5] that consider Chebyshev polynomials.

These filters can be expressed in the form (9) by restricting the parameter matrices to $\Phi^{(k:1)} = \phi_k \mathbf{S}^k$ for $k > 1$ and $\Phi^{(1)} = \phi_0 \mathbf{I}_N + \phi_1 \mathbf{S}$. In other words, the EV and the polynomial recursions represent two extremes to implement graph filters locally. The EV recursion allows each node i to learn for each iteration k a different parameter $\phi_{i,j}^{(k)}$ that reflects the importance of j 's node features to node i . The polynomial implementation instead forces all nodes to weigh the information of all neighbors with the same parameter ϕ_k within the k th iteration. However, this restriction makes the number of parameters $(K+1)F^2$ independent from M and N .

This way of parameterizing the EV recursion creates new opportunities for proposing a myriad of intermediate solutions that extend (10) towards the edge-variant implementation (9).

One such an approach may be a recursion that in addition to (10) considers parameters also for the most critical edges (e.g., edges without which the graph becomes disconnected).

Remark 1: Along with the above works, also [18]–[20] and [8], [9] fall under the lines of the polynomial filtering (10). In specific, [18] considers single shifts on graphs using as graph shift operator a learnable weight matrix, [19] considers a Gaussian kernel to mix the neighboring node information, while [20] uses random walks. The works in [8], [9], although aiming to build a GCNN layer by using graph filters with a rational transfer function, approximate the inherited matrix inverse in the vertex domain by a finite iterative algorithm. This finite approximation implicitly transforms these techniques into polynomial recursions whose order depend on the number of iterations (see also [10] for more detail). \square

Spectral GCNN. We here establish a link between the edge-variant recursion (9) and the spectral GCNN [2] to provide more insights on its convolutional behavior. The spectral GCNN exploits (1) and learns directly the filter a transfer function $h(\Lambda) = \text{diag}(h(\lambda_0), \dots, h(\lambda_{N-1}))$. To keep the number of parameters independent from N , $h(\Lambda)$ is parameterized as

$$h(\Lambda) = \text{diag}(\mathbf{B}\mathbf{b}) \quad (11)$$

where $\mathbf{B} \in \mathbb{R}^{N \times b}$ is a prefixed kernel matrix and $\mathbf{b} \in \mathbb{R}^b$ are the $b \ll N$ learnable parameters. Therefore, the number of parameters for each layer is at most bF^2 , while the computational complexity is of order $\mathcal{O}(N^2)$ required to compute the GFT of the features. Additionally, such an approach requires the eigendecomposition of \mathbf{S} (order $\mathcal{O}(N^3)$ to be computed once) and the learned filter $\mathbf{H}(\mathbf{S}) = \mathbf{U} \text{diag}(\mathbf{B}\mathbf{b}) \mathbf{U}^{-1}$ does not capture the local detail around each vertex.

Nevertheless, this spectral interpretation is useful to understand the EV behavior. We can force the EV recursion (9) to have a spectral response by restricting all coefficient matrices to share the eigenvectors with \mathbf{S} , i.e., $\Phi^{(k)} = \mathbf{U} \Lambda^{(k)} \mathbf{U}^{-1}$ [15]. Then, the EV transfer function becomes

$$h(\Lambda) = \sum_{k=1}^K \left(\prod_{\kappa=1}^k \Lambda^{(\kappa)} \right). \quad (12)$$

Subsequently, let \mathcal{I} be the index set defining the zero entries of $\mathbf{S} + \mathbf{I}_N$. The fixed support condition for each $\Phi^{(k)}$ is

$$\mathbf{C}_{\mathcal{I}} \text{vec} \left(\Phi^{(k)} \right) = \mathbf{0}_{|\mathcal{I}|} \quad (13)$$

where $\mathbf{C}_{\mathcal{I}} \in \{0, 1\}^{|\mathcal{I}| \times N^2}$ is a selection matrix whose rows are those of \mathbf{I}_{N^2} indexed by \mathcal{I} , $\text{vec}(\cdot)$ denotes the vectorization operation, and $\mathbf{0}_{|\mathcal{I}|}$ is the $|\mathcal{I}| \times 1$ vector of all zeros. From the properties of the $\text{vec}(\cdot)$ operator, (13) becomes

$$\mathbf{C}_{\mathcal{I}} \text{vec} \left(\mathbf{U} \Lambda^{(k)} \mathbf{U}^H \right) = \mathbf{C}_{\mathcal{I}} \text{vec} \left(\mathbf{U} * \mathbf{U}^H \right) \boldsymbol{\lambda}^{(k)} = \mathbf{0}_{|\mathcal{I}|} \quad (14)$$

where “*” denotes the Khatri-Rao product and $\boldsymbol{\lambda}^{(k)} = \text{diag}(\Lambda^{(k)})$ is the $N \times 1$ vector composed by the diagonal elements of $\Lambda^{(k)}$. Put differently, (14) implies

$$\boldsymbol{\lambda}^{(k)} \in \text{null} \left(\mathbf{C}_{\mathcal{I}} \text{vec} \left(\mathbf{U} * \mathbf{U}^H \right) \right). \quad (15)$$

Finally, by considering $\mathbf{B}_{\mathbf{U}, \mathcal{I}}$ as a basis that spans the nullspace of $\mathbf{C}_{\mathcal{I}} \text{vec} \left(\mathbf{U} * \mathbf{U}^H \right)$, we can expand $\boldsymbol{\lambda}^{(k)}$ as $\boldsymbol{\lambda}^{(k)} = \mathbf{B}_{\mathbf{U}, \mathcal{I}} \boldsymbol{\mu}^{(k)}$ and write (12) as

$$h(\mathbf{\Lambda}) = \sum_{k=1}^K \left(\prod_{\kappa=1}^k \text{diag} \left(\mathbf{B}_{\mathcal{U}, \mathcal{I}} \boldsymbol{\mu}^{(\kappa)} \right) \right) \quad (16)$$

for some basis expansion coefficients $\boldsymbol{\mu}^{(k)}$.

The eigendecomposition of $\Phi^{(k)}$ implicitly reduces the total number of layer parameters from $KF^2(M+N)$ to $KF^2 \text{rank}(\mathbf{B}_{\mathcal{U}, \mathcal{I}})$ with $\text{rank}(\mathbf{B}_{\mathcal{U}, \mathcal{I}}) \ll N$. That is, there is a subclass of the EV recursion that respects the operation of convolution, but, differently from (11), it captures local detail in the vertex domain and enjoys a linear implementation complexity. This subclass has also analogies with (11), which is obtained by setting $K=1$, $\mathbf{B}_{\mathcal{U}, \mathcal{I}} = \mathbf{B}$, and $\boldsymbol{\mu}^{(1)} = \mathbf{b}$.

In general, we may conclude that the EV recursion implements a GNN layer that goes beyond convolution. Drawing analogies with linear system theory, the GCNN approaches behave as a linear time-invariant (now shift-invariant [3]) filter, while the EV graph filter behaves as linear time-varying (now shift-varying; a different shift per (k) filter that trades the convolutional interpretation with the ability to capture time-varying (now shift-varying) detail.

Node-variant GCNN. The idea to propose GNNs that extend convolution is also considered in [11], which proposed an architecture having as graph filter the recursion

$$\mathbf{H}(\mathbf{S}) = \sum_{k=0}^K \text{diag} \left(\mathbf{C}_{\mathcal{B}} \boldsymbol{\phi}_{\mathcal{B}}^{(k)} \right) \mathbf{S}^k \quad (17)$$

where $\mathcal{B} \subset \mathcal{V}$ is a set of *privileged* nodes (e.g., the $|\mathcal{B}|$ nodes with the highest degree), $\mathbf{C}_{\mathcal{B}} \in \{0, 1\}^{N \times |\mathcal{B}|}$ is a tall binary matrix, and $\boldsymbol{\phi}_{\mathcal{B}}^{(k)} \in \mathbb{R}^{|\mathcal{B}|}$ is a vector of parameters for the nodes in \mathcal{B} . In short, (17) learns for each shift, $|\mathcal{B}|$ different coefficients for the nodes in \mathcal{B} and then maps them through $\mathbf{C}_{\mathcal{B}}$ to the remaining nodes $\mathcal{V} \setminus \mathcal{B}$.

This filter is another way to restrict the EV degrees of freedom, which parameterizes the coefficient matrices to $\Phi^{(1)} = \text{diag} \left(\mathbf{C}_{\mathcal{B}} \boldsymbol{\phi}_{\mathcal{B}}^{(0)} \right) + \text{diag} \left(\mathbf{C}_{\mathcal{B}} \boldsymbol{\phi}_{\mathcal{B}}^{(1)} \right) \mathbf{S}$ and $\Phi^{(k:1)} = \text{diag} \left(\mathbf{C}_{\mathcal{B}} \boldsymbol{\phi}_{\mathcal{B}}^{(k)} \right) \mathbf{S}^k$ for $k > 1$. That is, (17) is an intermediate approach between the polynomial (10) and the EV (9) recursions and allows each node $i \in \mathcal{B}$ to learn, for each k , a different parameter $\phi_i^{(k)}$ that reflects the importance of all its neighborhood to node i . The total number of parameters per layer is at most $(K+1)F^2|\mathcal{B}|$ while the computational complexity is similar to that of (10).

This different way of parameterizing the EV recursion provides alternative choices to build new intermediate architectures that lever the idea of privileged nodes while giving importance to the edge-based detail. In the sequel, we propose one such an extension that merges insights from the EV, the polynomial, and the NV architecture.

C. Hybrid edge-variant neural network layer

The hybrid edge-variant (HEV) layer considers the linear operation in (8) to be a graph filter of the form

$$\mathbf{H}(\mathbf{S}) = \sum_{k=0}^K \left(\prod_{\kappa=0}^k \Phi_{\mathcal{B}}^{(\kappa)} + \phi_k \mathbf{S}^k \right) \quad (18)$$

where $\Phi_{\mathcal{B}}^{(0)} \in \mathbb{R}^{N \times N}$ is a diagonal matrix whose i th diagonal element $\phi_{i,i}^{(0)} \neq 0$ iff node i belongs to the privileged set $\mathcal{B} \subset \mathcal{V}$; $\{\Phi_{\mathcal{B}}^{(k)}\}_{k=1}^K \in \mathbb{R}^{N \times N}$ are a collection of K matrices whose (i, j) th element $\phi_{i,j}^{(k)} \neq 0$ iff $i \in \mathcal{B}$ and $j \in \mathcal{N}_i \cup i$; and $\{\phi_k\}_{k=0}^K$ are a collection of $K+1$ scalars. Put simply, recursion (18) allows nodes in \mathcal{B} to learn node-varying parameters for $k=0$ and edge-varying parameters for $k \geq 1$, while the nodes in $\mathcal{V} \setminus \mathcal{B}$ learn global parameters similar to (10).

This approach represents yet another intermediate architecture between the full convolutional ones and the full EV. By setting $N_{\mathcal{B}, \max} = \max_{i \in \mathcal{B}} |\mathcal{N}_i|$ as the maximum number of neighbors for the nodes in \mathcal{B} , the overall number of parameters per layer is at most $F^2((K+1) + |\mathcal{B}| + K|\mathcal{B}|N_{\mathcal{B}, \max})$. Finally, the HEV implementation cost is of order $\mathcal{O}(KF^2(2M+N))$.

IV. NUMERICAL RESULTS

We compare the proposed edge-variant and hybrid edge-variant architectures with the spectral, polynomial, and the node-variant alternatives on a source localization and an author attribution problem. For both experiments, we designed all architectures (except for the spectral GCNN) to have the same computational cost. The code to generate these results can be found at <https://github.com/alelab-upenn/edgenets>.

A. Source Localization

Setup. The goal of this experiment is to find out which community in a stochastic block model (SBM) graph is the source of a diffusion process by observing different diffused signals originated at different (unknown) communities at different (unknown) time instants. \mathcal{G} is an undirected SBM graph of $N=50$ nodes divided equally into $C=5$ communities with respective intra- and inter-community edge probabilities of 0.8 and 0.2. The initial graph signal $\boldsymbol{\delta}_i \in \mathbb{R}^N$ is a Kronecker delta centered at node i and its realization at time t is $\mathbf{x}_t = \mathbf{S}^t \boldsymbol{\delta}_i$ with $\mathbf{S} = \mathbf{W} / \lambda_{\max}(\mathbf{W})$. We generated the training set $\{\mathbf{x}, i\}$ comprising 10,000 samples by selecting uniformly at random both $i \in \{1, \dots, C\}$ and $t \in \{0, \dots, N\}$. We then tested the different approaches on 200 new samples and averaged the performance over 10 different data and 10 different graph realizations for an overall of 100 Monte-Carlo runs.

Models and results. We considered seven architectures each of them composed of the cascade of a graph filtering layer with ReLU nonlinearity and a fully connected layer with softmax nonlinearity. The architectures are: *a*) a spectral GCNN (11) with \mathbf{B} being a cubic spline kernel and $b=5$; *b*) a polynomial GCNN (10) of order $K=4$; *c*) two NV GNNs (17) of order $K=4$ and $|\mathcal{B}|=5$ *privileged* nodes selected by *b_i*) maximum degree and *b_{ii}*) spectral proxies [21]; *d*) an EV GNN (9) of order $K=4$; and *e*) two HEV GNNs (18) of order $K=4$ and $|\mathcal{B}|=5$ *privileged* nodes selected similarly to the NV case. We used the ADAM optimizer with a learning rate $\mu=0.001$ and decaying factors $\beta_1=0.9$ and $\beta_2=0.999$ run over 20 epochs with batches of size 100.

Table I shows the obtained results where we see that because of their increased capacity the EV and the HEV outperform the other alternatives. We observe that the hybrid approaches exploit better the *edge-varying* part when \mathcal{B} is composed of the nodes with the highest degree.

TABLE I
AVERAGE PERFORMANCE (AND STD. DEV.) FOR THE SOURCE
LOCALIZATION PROBLEM.

Model	Accuracy
Spectral	26.89 (\pm 0.87) %
Polynomial	74.55 (\pm 7.32) %
Node Variant (NV) Degree	74.77 (\pm 7.77) %
Node Variant (NV) S. Proxies	75.62 (\pm 8.19) %
Edge Variant (EV)	85.47 (\pm 10.77) %
Hybrid EV (HEV) Degree	80.53 (\pm 10.21) %
Hybrid EV (HEV) S. Proxies	75.37 (\pm 8.20) %

B. Authorship Attribution

Setup. In this experiment, we aim to classify if a text excerpt belongs to Edgar Allan Poe or to any other contemporary author. For each text excerpt, we built the graph from the word adjacency network (WAN) between function words that act as nodes. These WANs serve as stylistic signatures for the author (see [22] for full details). Fixed then WAN for Poe, we treat the frequency count of the function words as a graph signal.

In particular, we considered 846 text excerpts by Poe and randomly split the dataset into 608 training, 68 validation, and 170 testing texts. We summed the adjacency matrices of the WANs obtained from the 608 training texts to get the Poe’s signature graph. We completed the training, validation, and test sets by adding respectively other 608, 68, and 170 randomly selected texts by contemporary authors. We set \mathbf{S} to be the adjacency matrix of the Poe’s signature graph and averaged the performance over 10 different data splits.

Models and results. We analyzed the same architectures as in the previous section but set the number of output features to $F = 2$, the recursion orders to $K = 1$, $b = 2$ for the spectral GCNN, and $|\mathcal{B}| = 2$. We used the same ADAM optimizer for training now over 80 epochs with batch sizes of 100 samples.

The results in Table II show that the hybrid approaches offer the best performance highlighting the potential of solutions that consider both edge-dependent and global coefficients. In fact, the polynomial model with global coefficients suffers the most in this experiment.

V. CONCLUSION

We proposed a general framework that unifies 11 state-of-the-art GCNN architectures into one recursion, named the edge-variant recursion. This unification highlighted the different tradeoff between the number of parameters and the amount of local detail that each approach adopts. Moreover, it shows rigorous ways to choose different tradeoffs and come up with a novel and ad hoc architecture for a problem at hand that is implemented locally in the vertex domain. We here proposed one, among many, extension and showed that it outperforms current solutions for graph signal classification tasks.

REFERENCES

[1] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, “The graph neural network model,” *IEEE Trans. Neural Netw.*, vol. 20, no. 1, pp. 61–80, Jan. 2009.

[2] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, “Spectral networks and deep locally connected networks on graphs,” *arXiv:1312.6203v3 [cs.LG]*, 21 May 2014. [Online]. Available: <http://arxiv.org/abs/1213.6203>

TABLE II
AVERAGE PERFORMANCE (AND STD. DEV.) FOR THE AUTHORSHIP
ATTRIBUTION PROBLEM ON TEXT EXCERPTS BY EDGAR ALLAN POE.

Model	Accuracy
Spectral	88.88 (\pm 1.50) %
Polynomial	79.88 (\pm 15.31) %
Node Variant (NV) Degree	88.88 (\pm 2.62) %
Node Variant (NV) S. Proxies	86.12 (\pm 5.94) %
Edge Variant (EV)	89.00 (\pm 2.11) %
Hybrid EV (HEV) Degree	89.18 (\pm 1.99) %
Hybrid EV (HEV) S. Proxies	90.00 (\pm 1.21) %

- [3] A. Ortega, P. Frossard, J. Kovačević, J. M. Moura, and P. Vandergheynst, “Graph signal processing: Overview, challenges, and applications,” *Proceedings of the IEEE*, vol. 106(5), pp. 808–828, 2018.
- [4] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *5th Int. Conf. Learning Representations*. Toulon, France: Assoc. Comput. Linguistics, 24–26 Apr. 2017.
- [5] M. Defferrard, X. Bresson, and P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *Annu. Conf. Neural Inform. Process. Syst. 2016*. Barcelona, Spain: NIPS Foundation, 5–10 Dec. 2016.
- [6] J. Du, S. Zhang, G. Wu, J. M. F. Moura, and S. Kar, “Topology adaptive graph convolutional networks,” *arXiv:1710.10370v2 [cs.LG]*, 2 Nov. 2017. [Online]. Available: <http://arxiv.org/abs/1710.10370v2>
- [7] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, “Convolutional neural network architectures for signals supported on graphs,” *IEEE Trans. Signal Process.*, vol. 67, no. 4, pp. 1034–1049, Feb. 2019.
- [8] R. Levie, F. Monti, X. Bresson, and M. M. Bronstein, “CayleyNets: Graph convolutional neural networks with complex rational spectral filters,” *IEEE Trans. Signal Process.*, vol. 67(1), pp. 97–107, Jan. 2019.
- [9] F. M. Bianchi, D. Grattarola, C. Alippi, and L. Livi, “Graph neural networks with convolutional ARMA filters,” Feb. 2019. [Online]. Available: <http://arxiv.org/abs/1901.01343>
- [10] E. Isufi, A. Loukas, A. Simonetto, and G. Leus, “Autoregressive moving average graph filtering,” *IEEE Trans. Signal Process.*, vol. 65, no. 2, pp. 274–288, Jan. 2017.
- [11] F. Gama, G. Leus, A. G. Marques, and A. Ribeiro, “Convolutional neural networks via node-varying graph filters,” in *2018 IEEE Data Sci. Workshop*. Lausanne, Switzerland: IEEE, 4–6 June 2018, pp. 1–5.
- [12] S. Segarra, A. G. Marques, and A. Ribeiro, “Optimal graph-filter design and applications to distributed linear network operators,” *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4117–4131, Aug. 2017.
- [13] F. Gama, A. G. Marques, A. Ribeiro, and G. Leus, “MIMO graph filters for convolutional networks,” in *19th IEEE Int. Workshop Signal Process. Advances in Wireless Commun.* Kalamata, Greece: IEEE, June 2018.
- [14] M. Coutino, E. Isufi, and G. Leus, “Distributed edge-variant graph filters,” in *2017 IEEE Int. Workshop Comput. Advances Multi-Sensor Adaptive Process.* Curacao, Dutch Antilles: IEEE, 10–13 Dec. 2017.
- [15] M. A. C. Minguez, E. Isufi, and G. Leus, “Advances in distributed graph filtering,” *IEEE Transactions on Signal Processing*, 2019.
- [16] L. Ruiz, F. Gama, A. G. Marques, and A. Ribeiro, “Median activation functions for graph neural networks,” in *44th IEEE Int. Conf. Acoust., Speech and Signal Process.* Brighton, UK: IEEE, 12–17 May 2019.
- [17] B. Wang, A. Pourshafeie, M. Zitnik, J. Zhu, C. Bustamante, S. Batzoglou, and J. Leskovec, “Network enhancement as a general method to denoise weighted biological networks,” *Nature Communications*, vol. 9, no. 3108, pp. 1–8, Aug. 2018.
- [18] M. Simonovsky and N. Komodakis, “Dynamic edge-conditioned filters in convolutional neural networks on graphs,” in *2017 IEEE Comput. Soc. Conf. Comput. Vision and Pattern Recognition*, Honolulu, July 2017.
- [19] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. M. Bronstein, “Geometric deep learning on graphs and manifolds using mixture model CNNs,” in *2017 IEEE Comput. Soc. Conf. Comput. Vision and Pattern Recognition*. Honolulu, HI: IEEE, July 2017.
- [20] J. Atwood and D. Towsley, “Diffusion-convolutional neural networks,” in *30th Annu. Conf. Neural Inform. Process. Syst.* Barcelona, Spain: NIPS Foundation, 5–10 Dec. 2016.
- [21] A. Anis, A. Gadde, and A. Ortega, “Efficient sampling set selection for bandlimited graph signals using graph spectral proxies,” *IEEE Trans. Signal Process.*, vol. 64, no. 14, pp. 3775–3789, July 2016.
- [22] S. Segarra, M. Eisen, and A. Ribeiro, “Authorship attribution through function word adjacency networks,” *IEEE Trans. Signal Process.*, vol. 63, no. 20, pp. 5464–5478, Oct. 2015.