

EdgeNets

Edge Varying Graph Neural Networks

Isufi, Elvin; Gama, Fernando ; Ribeiro, Alejandro

DOI

[10.1109/TPAMI.2021.3111054](https://doi.org/10.1109/TPAMI.2021.3111054)

Publication date

2022

Document Version

Final published version

Published in

IEEE Transactions on Pattern Analysis and Machine Intelligence

Citation (APA)

Isufi, E., Gama, F., & Ribeiro, A. (2022). EdgeNets: Edge Varying Graph Neural Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11), 7457-7473. Article 9536420. <https://doi.org/10.1109/TPAMI.2021.3111054>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

EdgeNets: Edge Varying Graph Neural Networks

Elvin Isufi , Fernando Gama , and Alejandro Ribeiro 

Abstract—Driven by the outstanding performance of neural networks in the structured euclidean domain, recent years have seen a surge of interest in developing neural networks for graphs and data supported on graphs. The graph is leveraged at each layer of the neural network as a parameterization to capture detail at the node level with a reduced number of parameters and computational complexity. Following this rationale, this paper puts forth a general framework that unifies state-of-the-art graph neural networks (GNNs) through the concept of EdgeNet. An EdgeNet is a GNN architecture that allows different nodes to use different parameters to weigh the information of different neighbors. By extrapolating this strategy to more iterations between neighboring nodes, the EdgeNet learns edge- and neighbor-dependent weights to capture local detail. This is a general linear and local operation that a node can perform and encompasses under one formulation all existing graph convolutional neural networks (GCNNs) as well as graph attention networks (GATs). In writing different GNN architectures with a common language, EdgeNets highlight specific architecture advantages and limitations, while providing guidelines to improve their capacity without compromising their local implementation. For instance, we show that GCNNs have a parameter sharing structure that induces permutation equivariance. This can be an advantage or a limitation, depending on the application. In cases where it is a limitation, we propose hybrid approaches and provide insights to develop several other solutions that promote parameter sharing without enforcing permutation equivariance. Another interesting conclusion is the unification of GCNNs and GATs—approaches that have been so far perceived as separate. In particular, we show that GATs are GCNNs on a graph that is learned from the features. This particularization opens the doors to develop alternative attention mechanisms for improving discriminatory power.

Index Terms—Edge varying, graph neural networks, graph signal processing, graph filters, learning on graphs

1 INTRODUCTION

DATA generated by networks is increasingly common. Examples include user preferences in recommendation systems, writer proclivities in blog networks [2], or properties of assembled molecular compounds [3]. Different from data encountered in the structured temporal or spatial domains, network data lives in high-dimensional irregular spaces. This fact makes difficult to extend tools that exploit the regularity of time and space, leading to a rising interest in novel techniques for dealing with network data [4]. Since graphs are the prominent mathematical tool to model individual node properties—product ratings, writer bias, or molecule properties—along with node dependencies—user similarities, blog hyperlinks, or molecular bonds—the interest in network data has translated into a concomitant increase in the interest in tools for processing graphs and data supported on graphs [5].

Several recent works have proposed graph neural networks (GNNs) as a means of translating to graphs the success convolutional and recurrent neural networks have attained at learning on time and space [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22].

- Elvin Isufi is with the Intelligent Systems Department, Delft University of Technology, 2628 CD, Delft, The Netherlands. E-mail: e.isufi-1@tudelft.nl.
- Fernando Gama is with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005 USA. E-mail: fgama@rice.edu.
- Alejandro Ribeiro is with the Department of Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA. E-mail: aribeiro@seas.upenn.edu.

Manuscript received 12 Mar. 2020; revised 16 June 2021; accepted 2 Sept. 2021.

Date of publication 13 Sept. 2021; date of current version 3 Oct. 2022.

(Corresponding author: Elvin Isufi.)

Recommended for acceptance by H. Kjellstrom.

Digital Object Identifier no. 10.1109/TPAMI.2021.3111054

GNNs are first concretized in [6], [7] by means of recursive neighboring label aggregations combined with pointwise nonlinearities. The convolutional GNN counterpart appears in [8] where graph convolutions are defined as pointwise operators in the Laplacian's spectrum. To avoid the cost and numerical instability of spectral decompositions, [9] approximates this spectral convolution with a Chebyshev polynomial on the Laplacian matrix. Parallel to these efforts, the field of graph signal processing has developed notions of graph convolutional filters as polynomials on a matrix representation of a graph [23], [24], [25], [26], [27], [28], [29]. This has led to GNNs described as architectures that simply replace time convolutions with graph convolutions [10], [11]. A third approach to define GNNs is to focus on the locality of convolutions by replacing the adjacency of points in time with the adjacency of neighbors in a graph; something that can be accomplished by mixing nodes' features with their neighbor's features [12], [13].

Despite their different motivations, spectral GNNs [8], [9], polynomial GNNs [10], [11], and local GNNs [12], [13] can all be seen to be equivalent to each other (Section 4). In particular, they all share the reuse of parameters across all neighborhoods of a graph as well as indifference towards the values of different neighbors—see also [30]. This is an important limitation that is tackled, e.g., by the graph attention networks (GAT) of [18], [19], [20], [21], [22] through the use of attention mechanisms [31], [32]. In this paper, we leverage edge varying graph filter [33] to provide a generic framework for the design of GNNs that can afford flexibility to use different parameters at different nodes as well as different weighing to different neighbors of a node (Section 3). Edge varying filters are linear finite order recursions that

allow individual nodes to introduce weights that are specific to the node, specific to each neighbor, and specific to the recursion index. In this way, the edge varying recursion represents a general linear operation that a node can implement locally. I.e., a general operation that relies on information exchanges only with neighbor nodes (Section 2).

In alternative to the EdgeNet, graph network [34] is a popular framework for generalizing GNNs. Graph network consist of general update and aggregation functions over nodal, edge, and entire graph features. This unification is slightly more general than the message passing neural network [35] and considers updates to be principally affected by information exchange only with the one-hop neighbors. While providing relevant insights on the local detail of order-one filter GNNs such as [12], [18], this strategy does not put emphasis on the role of the filter within the GNN or how the parameters of such filter are allocated to the different multi-hop neighbours [9], [10]. Instead, the EdgeNet framework focuses only on nodal feature aggregations to highlight the role of multi-hop exchanges within a layer and to put emphasis on how different solutions operate from a node perspective. At the same time, the EdgeNet framework allows for a filter spectral analysis [36], which provides a better understanding of the type of filters that conform the learned filter bank. This spectral perspective of GNNs will shed light, for instance, of the advantages of ARMA filters [29] in learning sharper responses with less parameters than finite impulse response graph filters [24], [25], [28].

In its most general form, the edge varying GNNs allocate different parameters to the different edges, which is of the order of the number of nodes and edges of the graph. While allocating different parameters over the edges can help exploiting the graph structure better for the learning task at hand, edge-specific parameters sacrifice the inductive capabilities across different graphs. To reduce the complexity of this parameterization, we can regularize EdgeNets in different ways by imposing restrictions on the freedom to choose different parameters at different nodes. We explain that existing GNN architectures are particular cases of EdgeNets associated with different parameter restrictions. In turn, this shows how these solutions sacrifice the degrees of freedom to gain in parameter sharing and inductive capabilities. We further utilize the insight of edge varying recursions to propose novel GNN architectures. In consequence, the novel contributions of this paper are:

- (i) We define EdgeNets, which parameterize the linear operation of neural networks through a bank of edge varying recursions. EdgeNets are a generic framework to design GNN architectures (Section 3).
- (ii) We show the approaches in [8], [9], [10], [11], [12], [13], [14], [15], [16], [17] (among others) are EdgeNets where all nodes share the parameters. We extend the representing power of these networks by adding some level of variability in weighing different nodes and different edges of a node (Section 4).
- (iii) Replacing finite length polynomials by rational functions provides an alternative parameterization of convolutional GNNs in terms of autoregressive moving

average (ARMA) convolutional graph filters [29]. These ARMA GNNs generalize rational functions based on Cayley polynomials [14] (Section 4.3).

- (iv) We show that GATs can be understood as GNNs with convolutional graph filters where a graph is learned ad hoc in each layer to represent the required abstraction between nodes. The weights of this graph choose neighbors whose values should most influence the computations at a particular node. This reinterpretation allows for the proposal of more generic GATs with higher expressive power (Section 5).

The this paper has the following structure. Section 2 reviews edge varying recursions on graphs and Section 3 introduces edge varying GNNs. To ease exposition, the GNN are grouped into two categories: Convolutional in Section 4 and attentional in Section 5. Within each category, we follow the same rationale. First, we discuss the state-of-the-art solutions as a particular case of the EdgeNet framework. Then, we discuss their architectural advantages and limitations. Finally, we leverage the EdgeNet viewpoint to propose new solutions that address some of these limitations and highlight the corresponding tradeoffs. Section 6 evaluates these solutions with numerical results and Section 7 concludes the paper.

2 EDGE VARYING LINEAR GRAPH FILTERS

Consider a weighted graph \mathcal{G} with vertex set $\mathcal{V} = \{1, \dots, N\}$, edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ composed of $|\mathcal{E}| = M$ ordered pairs (i, j) , and weight function $\mathcal{W} : \mathcal{E} \rightarrow \mathbb{R}$. For each node i , define the neighborhood $\mathcal{N}_i = \{j : (j, i) \in \mathcal{E}\}$ as the set of nodes connected to i and let $N_i := |\mathcal{N}_i|$ denote the number of elements (neighbors) in this set. Associated with \mathcal{G} is a graph shift operator matrix $\mathbf{S} \in \mathbb{R}^{N \times N}$ whose sparsity pattern matches that of the edge set, i.e., entry $S_{ij} \neq 0$ when $(j, i) \in \mathcal{E}$ or when $i = j$. Supported on the vertex set are graph signals $\mathbf{x} = [x_1, \dots, x_N]^T \in \mathbb{R}^N$ in which component x_i is associated with node $i \in \mathcal{V}$.

The adjacency of points in time signals or the adjacency of points in images codifies a sparse and local relationship between signal components. This sparsity and locality are leveraged by time or space filters. Similarly, \mathbf{S} captures the sparsity and locality of the relationship between components of a signal \mathbf{x} supported on \mathcal{G} . It is then natural to take the shift operator as the basis for defining filters for graph signals. In this spirit, let $\Phi^{(0)}$ be an $N \times N$ diagonal matrix and $\Phi^{(1)}, \dots, \Phi^{(K)}$ be a collection of K matrices sharing the sparsity pattern of $\mathbf{I}_N + \mathbf{S}$. Consider then the sequence of signals $\mathbf{z}^{(k)}$ as

$$\mathbf{z}^{(k)} = \prod_{k'=0}^k \Phi^{(k')} \mathbf{x} = \Phi^{(k:0)} \mathbf{x}, \quad \text{for } k = 0, \dots, K, \quad (1)$$

where the product matrix $\Phi^{(k:0)} := \prod_{k'=0}^k \Phi^{(k')} = \Phi^{(k)} \dots \Phi^{(0)}$ is defined for future reference. Signal $\mathbf{z}^{(k)}$ can be computed using the recursion

$$\mathbf{z}^{(k)} = \Phi^{(k)} \mathbf{z}^{(k-1)}, \quad \text{for } k = 0, \dots, K, \quad (2)$$

with initialization $\mathbf{z}^{(-1)} = \mathbf{x}$. This recursive expression implies signal $\mathbf{z}^{(k)}$ is produced from $\mathbf{z}^{(k-1)}$ using operations that are

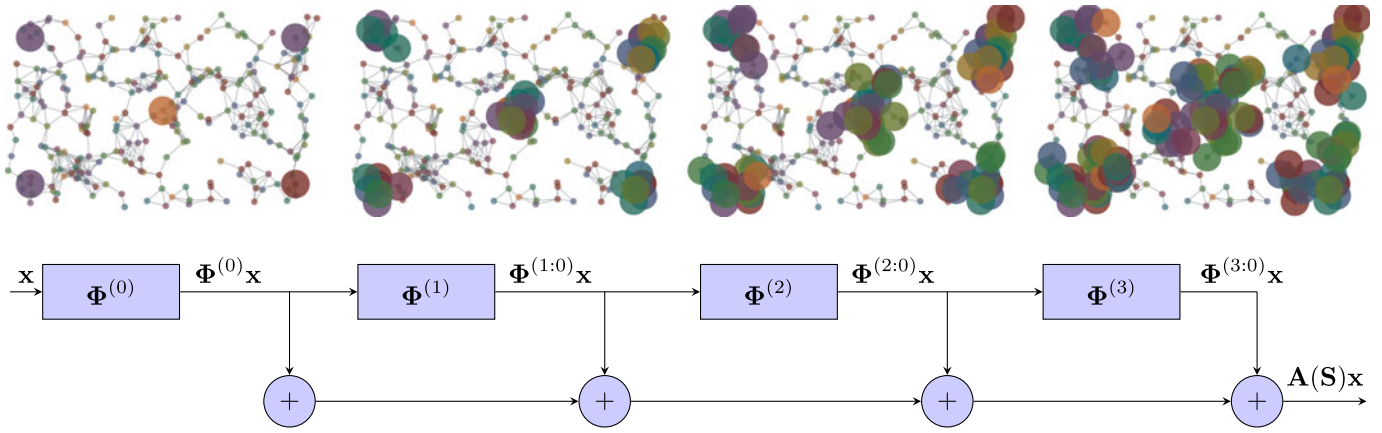


Fig. 1. Edge varying graph filters. Each edge varying matrix $\Phi^{(k)}$ acts as a different shift operator that locally combines the graph signal. (Top-left) The colored discs are centered at five reference nodes and their coverage shows the amount of local information needed to compute $\mathbf{z}^{(1)} = \Phi^{(1:0)}\mathbf{x}$ at these nodes. The coverage of the discs in the other graphs shows the signal information needed by the reference nodes to produce the successive outputs. (Bottom) Schematic illustration of the edge varying filter output of order $K = 3$.

local in the graph. Indeed, since $\Phi^{(k)}$ shares the sparsity pattern of \mathbf{S} , node i computes its component $z_i^{(k)}$ as

$$z_i^{(k)} = \sum_{j \in \mathcal{N}_i \cup i} \Phi_{ij}^{(k)} z_j^{(k-1)}. \quad (3)$$

Particularizing (3) to $k = 0$, it follows each node i builds the i th entry of $\mathbf{z}^{(0)}$ as a scaled version of its signal \mathbf{x} by the diagonal matrix $\Phi^{(0)}$, i.e., $z_i^{(0)} = \Phi_{ii}^{(0)} x_i$. Particularizing to $k = 1$, (3) yields the components of $\mathbf{z}^{(1)}$ depend on the values of the signal \mathbf{x} at most at neighboring nodes. Particularizing to $k = 2$, (3) shows the components of $\mathbf{z}^{(2)}$ depend only on the values of signal $\mathbf{z}^{(1)}$ at neighboring nodes which, in turn, depend only on the values of \mathbf{x} at their neighbors. Thus, the components of $\mathbf{z}^{(2)}$ are a function of the values of \mathbf{x} at most at the respective two-hop neighbors. Repeating this argument iteratively, $z_i^{(k)}$ represents an aggregation of information at node i coming from its k -hop neighborhood —see Fig. 1.

The collection of signals $\mathbf{z}^{(k)}$ behaves like a sequence of scaled shift operations except that instead of shifting the signal in time, the signal is diffused through the graph (the signal values are *shifted* between neighboring nodes). Leveraging this interpretation, the graph filter output \mathbf{u} is defined as the sum

$$\mathbf{u} = \sum_{k=0}^K \mathbf{z}^{(k)} = \sum_{k=0}^K \Phi^{(k:0)} \mathbf{x}. \quad (4)$$

A filter output in time is a sum of scaled and shifted copies of the input signal. That (4) behaves as a filter follows from interpreting $\Phi^{(k:0)}$ as a scaled shift, which holds because of its locality. Each shift $\Phi^{(k:0)}$ is a recursive composition of individual shifts $\Phi^{(k)}$. These individual shifts represent different operators that respect the structure of \mathcal{G} while reweighing individual edges differently when needed.

For future reference, define the filter matrix $\mathbf{A}(\mathbf{S})$ so (4) rewrites as $\mathbf{u} = \mathbf{A}(\mathbf{S})\mathbf{x}$. For this to hold, the filter matrix must be

$$\mathbf{A}(\mathbf{S}) = \sum_{k=0}^K \Phi^{(k:0)} = \sum_{k=0}^K \left(\prod_{k'=0}^k \Phi^{(k')} \right). \quad (5)$$

Following [33], $\mathbf{A}(\mathbf{S})$ is a K th order edge varying graph filter. Each matrix $\Phi^{(k)}$ contains at most $M + N$ nonzero elements corresponding to the nonzero entries of $\mathbf{I}_N + \mathbf{S}$; thus, the total number of parameters defining filter $\mathbf{A}(\mathbf{S})$ in (5) is $K(M + N) + N$. For short filters, this is smaller than the N^2 components of an arbitrary linear transform. Likewise, in computing $\mathbf{z}^{(k)} = \Phi^{(k)}\mathbf{z}^{(k-1)}$ as per (2) incurs a computational complexity of order $\mathcal{O}(M + N)$. This further results in an overall computational complexity of order $\mathcal{O}(K(M + N))$ for obtaining the filter output \mathbf{u} in (4). This reduced number of parameters and computational cost is leveraged next to define graph neural network (GNN) architectures with a controlled number of parameters and computational complexity matched to the graph sparsity.

Remark 1. The presence of the edge (j, i) in graph \mathcal{G} is interpreted here as signal components x_j and x_i being related by the given structure in the data. The shift operator entry S_{ij} is a measure of the expected similarity. Larger entries indicate linked signal components are more related to each other. Therefore, the definition of the shift operator \mathbf{S} makes it a valid stand-in for any graph representation matrix. Forthcoming discussions are valid whether \mathbf{S} is an adjacency or a Laplacian matrix in any of their various normalized and unnormalized forms. We use \mathbf{S} to keep discussions generic.

3 EDGE VARYING GRAPH NEURAL NETWORKS

Edge varying graph filters are the basis for defining GNN architectures through composition with pointwise nonlinear functions. Formally, consider a set of L layers indexed by $l = 1, \dots, L$ and let $\mathbf{A}_l(\mathbf{S}) = \sum_{k=0}^K \Phi_l^{(k:0)}$ be the graph filter used at layer l . A GNN is defined by the recursive expression

$$\mathbf{x}_l = \sigma\left(\mathbf{A}_l(\mathbf{S})\mathbf{x}_{l-1}\right) = \sigma\left(\sum_{k=0}^K \Phi_l^{(k:0)} \mathbf{x}_{l-1}\right), \quad (6)$$

where we convene that $\mathbf{x}_0 = \mathbf{x}$ is the input to the GNN and \mathbf{x}_L is its output. To augment the representation power of GNNs, it is customary to add multiple node features per layer. We do this by defining matrices $\mathbf{X}_l = [\mathbf{x}_l^1, \dots, \mathbf{x}_l^{F_l}] \in \mathbb{R}^{N \times F_l}$ in which each column \mathbf{x}_l^f represents a different graph signal at layer l . These so-called features are cascaded through layers where they are processed with edge varying graph filters and composed with pointwise nonlinearities according to

$$\mathbf{X}_l = \sigma \left(\sum_{k=0}^K \Phi_l^{(k;0)} \mathbf{X}_{l-1} \mathbf{A}_{lk} \right), \quad (7)$$

where $\mathbf{A}_{lk} \in \mathbb{R}^{F_{l-1} \times F_l}$ is a parameter matrix that affords flexibility to process different features with different filter parameters. It is ready to see that (7) represents a bank of edge varying graph filters $\mathbf{A}_l^{fg}(\mathbf{S})$ applied to a set of F_{l-1} input features \mathbf{x}_{l-1}^g to produce a set of F_l output features \mathbf{x}_l^f . Indeed, if we let $a_{lk}^{fg} = [\mathbf{A}_{lk}]_{fg}$ denote the (f, g) th entry of \mathbf{A}_{lk} , (7) produces a total of $F_{l-1} F_l$ intermediate features of the form¹

$$\mathbf{u}_l^{fg} = \mathbf{A}_l^{fg}(\mathbf{S}) \mathbf{x}_{l-1}^g = \sum_{k=0}^K a_{lk}^{fg} \Phi_l^{fg, (k;0)} \mathbf{x}_{l-1}^g, \quad (8)$$

for $g = 1, \dots, F_{l-1}$ and $f = 1, \dots, F_l$. The features \mathbf{u}_l^{fg} are then aggregated across all g and passed through a pointwise nonlinearity to produce the output features of layer l as

$$\mathbf{x}_l^f = \sigma \left(\sum_{g=1}^{F_{l-1}} \mathbf{u}_l^{fg} \right). \quad (9)$$

At layer $l = 1$ the input feature is a graph signal $\mathbf{x}_0^1 = \mathbf{x}$. This feature is passed through F_1 filters to produce F_1 higher-level features as per (8). The latter are then processed by a pointwise nonlinearity [cf. (9)] to produce F_1 output features \mathbf{x}_1^1 . The subsequent layers $l > 1$ start with F_{l-1} input features \mathbf{x}_{l-1}^g that are passed through the filter bank $\mathbf{A}_l^{fg}(\mathbf{S})$ [cf. (8)] to produce the higher-level features \mathbf{u}_l^{fg} . These are aggregated across all $g = 1, \dots, F_{l-1}$ and passed through a nonlinearity to produce the layer's output features \mathbf{x}_l^f [cf. (9)]. In the last layer $l = L$, we consider without loss of generality the number of output features is $F_L = 1$. This single feature $\mathbf{x}_L^1 = \mathbf{x}_L$ is the output of the edge varying GNN or, for short, EdgeNet. Remark the EdgeNet aggregates at each layer information from neighbors that are up to K hops away [cf. (7)]. This increases its flexibility to process intermediate features and generalizes the masking aggregation rule in [37], which can be seen as an EdgeNet of order $K = 1$.

The EdgeNet output is a function of the input signal \mathbf{x} and the collection of filter banks \mathbf{A}_l^{fg} [cf. (5)]. Group the filters in the filter tensor $\mathcal{A}(\mathbf{S}) = \{\mathbf{A}_l^{fg}(\mathbf{S})\}_{lfg}$ so that to define the GNN output as the mapping

1. Throughout the paper, we will denote *any* graph filter by $\mathbf{A}_l^{fg}(\mathbf{S})$ to indicate that it is a matrix depending on the graph shift operator \mathbf{S} . When this filter contains additional parameters rather than those on the edges (i.e., $\Phi_l^{fg, (k;0)}$), we will indicate them with scalars a_{lk}^{fg} [cf. (8)]. For consistency, when expressing the bank of filters in a single recursion [cf. (7)], we will group parameters a_{lk}^{fg} into the matrix \mathbf{A}_{lk} .

$$\Psi(\mathbf{x}; \mathcal{A}(\mathbf{S})) := \mathbf{x}_L \quad \text{with } \mathcal{A}(\mathbf{S}) = \left\{ \mathbf{A}_l^{fg}(\mathbf{S}) \right\}_{lfg}. \quad (10)$$

The filter parameters are trained to minimize a loss over a training set of input-output pairs $\mathcal{T} = \{(\mathbf{x}, \mathbf{y})\}$. This loss measures the difference between the EdgeNet output \mathbf{x}_L and the true value \mathbf{y} averaged over the examples $(\mathbf{x}, \mathbf{y}) \in \mathcal{T}$.

As it follows from (5), the number of parameters in each filter is $K(M + N) + N$. This gets scaled by the number of filters per layer $F_{l-1} F_l$ and the number of layers L . To provide an order bound on the number of parameters defining the EdgeNet set the maximum feature number $F = \max_l F_l$ and observe the number of parameters per layer is of order $(K(M + N) + N)F^2 + F^2$. Likewise, the computational complexity at each layer is of order $\mathcal{O}(K(M + N)F^2)$. This number of parameters and computational complexity are expected to be smaller than the corresponding numbers of a fully connected neural network. This is a consequence of exploiting the sparse nature of edge varying filters [cf. (4) and (5)]. A GNN can be then considered as an architecture that exploits the graph structure to reduce the number of parameters of a fully connected neural network. The implicit hypothesis is those signal components associated with different nodes are processed together in accordance with the nodes' proximity in the graph.

We will show different existing GNN architectures are particular cases of (8)–(9) using different subclasses of edge varying graph filters (Section 4) and the same is true for graph attention networks (Section 5). Establishing these relationships allows the proposal of natural architectural generalizations that increase the descriptive power of GNNs while still retaining manageable complexity.

Remark 2. The key property of the EdgeNet is to allocate trainable parameters for each edge in each shift. While this formulation improves the expressive power of a GNN, it affects its inductive capability over graphs (i.e., the ability to generalize to new unseen graphs) [38]. In the full form (7), the EdgeNet does not have inductive capabilities for graphs but only for graph signals. That is, it cannot be applied to test cases where new unseen graphs are present but can only be applied to test cases where new unseen graph signals are present for a fixed graph support. We shall see in the next section that graph convolutional neural networks sacrifice instead the degrees of freedom to gain inductive capabilities also for graphs.

Remark 3. In the proposed EdgeNet, we considered graphs with single edge features, i.e., each edge is described by a single scalar. However, even when the graph has multiple edge features, say E , the EdgeNet extends readily to this scenario. This can be obtained by seeing the multi-edge featured graph as the union of E graphs $\mathcal{G}_e = (\mathcal{V}, \mathcal{E}_e)$ with identical node set \mathcal{V} and respective shift operator matrix \mathbf{S}_e . For $\{\Phi^{e(k)}\}$ being the collection of the edge varying parameter matrices [cf. (1)] relative to the shift operator \mathbf{S}_e , the l th layer output \mathbf{X}_l [cf. (7)] becomes

$$\mathbf{X}_l = \sigma \left(\sum_{e=1}^E \sum_{k=0}^K \Phi_l^{e(k;0)} \mathbf{X}_{l-1} \mathbf{A}_{lk}^e \right). \quad (11)$$

I.e., the outputs of each filter are aggregated also over the edge-feature dimension. The number of parameters and computational complexity get scaled by E . The GNN architectures discussed in the remainder of this manuscript, as a special case of the EdgeNet, are readily extendable to the multi-edge feature scenario by replacing (7) with (11). The approach in [39] is the particular case for (11) with $K=1$ and the parameter matrix reduced to a scalar.

4 GRAPH CONVOLUTIONAL NEURAL NETWORKS

Graph convolutional neural networks (GCNNs) have shown great success to learning representations for graph data with prominent variants introduced in [9], [10], [11], [12]. All these variants be written as GNN architectures in which the edge varying component in (7) is fixed and given by powers of the shift operator matrix $\Phi_l^{(k;0)} = \mathbf{S}^k$

$$\mathbf{X}_l = \sigma \left(\sum_{k=0}^K \mathbf{S}^k \mathbf{X}_{l-1} \mathbf{A}_{lk} \right). \tag{12}$$

By comparing (8) with (12), it follows this particular restriction yields a tensor $\mathcal{A}(\mathbf{S})$ with filters of the form

$$\mathbf{A}_l^{fg}(\mathbf{S}) = \sum_{k=0}^K a_{lk}^{fg} \mathbf{S}^k, \tag{13}$$

for some order K and scalar parameters $a_{l0}^{fg}, \dots, a_{lK}^{fg}$. Our focus in this section is to discuss variations on (13). To simplify the discussion, we omit the layer and feature indices and for the remainder of this section write

$$\mathbf{A}(\mathbf{S}) = \sum_{k=0}^K a_k \mathbf{S}^k. \tag{14}$$

The filters in (14) are of the form in (5) with $\Phi^{(0)} = a_0 \mathbf{I}_N$ and $\Phi^{(k;0)} = a_k \mathbf{S}^k$ for $k \geq 1$. By particularizing \mathcal{G} to the line graph, (14) represents a linear time-invariant filter described by a regular convolution. This justifies using the qualifier convolutional for an architecture with filters of the form (14).

The appeal of the graph convolutional filters in (14) is that they reduce the number of parameters from the $K(M + N) + N$ of the edge varying filters in (5) to just $K + 1$; yielding also a computational complexity of order $\mathcal{O}(KM)$. While we can reduce the number of parameters in several ways, the formulation in (14) is of note because it endows the resulting GNN with equivariance to permutations of the labels of the graph. We state this property formally in the following proposition.

Proposition 1. *Let \mathbf{x} be a graph signal defined on the vertices of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with shift operator \mathbf{S} . Consider also the output of a GCNN $\Psi(\mathbf{x}; \mathcal{A}(\mathbf{S}))$ [cf. (10)] with input \mathbf{x} and tensor $\mathcal{A}(\mathbf{S}) = \{\mathbf{A}(\mathbf{S})\}$ composed of filters of the form in (14). Then, for a permutation matrix \mathbf{P} , it holds that*

$$\mathbf{P}^T \Psi(\mathbf{x}; \mathcal{A}(\mathbf{S})) = \Psi(\mathbf{P}^T \mathbf{x}; \mathcal{A}(\mathbf{P}^T \mathbf{S} \mathbf{P})).$$

That is, the GCNN output operating on the graph \mathcal{G} with input \mathbf{x} is a permuted version of the GCNN output operating on the permuted graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ with permuted shift operator $\mathbf{S}' = \mathbf{P}^T \mathbf{S} \mathbf{P}$ and permuted input signal $\mathbf{x}' = \mathbf{P}^T \mathbf{x}$.

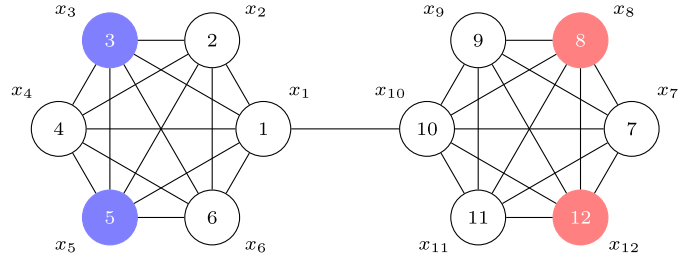


Fig. 2. Permutation equivariance of machine learning on graphs. Many tasks in machine learning on graphs are equivariant to permutations (cf. Proposition 1) but not all are. E.g., we expect agents 3, 5, 8, and 12 to be interchangeable from the perspective of predicting product ratings from the ratings of other nodes. But from the perspective of community classification we expect 3 and 5 or 8 and 12 to be interchangeable, but 3 and 5 are not interchangeable with 8 and 12.

Proof. See Appendix A,² available in the online supplemental material. \square

Proposition 1 establishes the output of a GCNN is independent of node labeling. This is important not just because graph signals are independent of labeling —therefore, so should be their processing— but because it explains how GCNNs exploit the internal signal symmetries. If two parts of the graph are topologically identical and the nodes support identical signal values, a GCNN yields identical outputs [40], [41], [42].

It must be emphasized that permutation equivariance is of use only inasmuch as this is a desirable property of the considered task. Permutation equivariance holds in, e.g., recommendation systems but does not hold in, e.g., community classification. In the graph in Fig. 2, we expect agents 3, 5, 8, and 12 to be interchangeable from the perspective of predicting product ratings from the ratings of other nodes. But from the perspective of community classification, we expect 3 and 5 or 8 and 12 to be interchangeable, but 3 and 5 are not interchangeable with 8 and 12.

When equivariance is not a property of the task, GCNNs are not expected to do well. GCNNs will suffer in any problem in which local detail around a node is important. This is because the filter in (14) forces all nodes to weigh the information of all k -hop neighbors with the same parameter a_k irrespectively of the relative importance of different nodes and different edges. To avoid this limitation, we can use a GNN that relies on the edge varying filters (5) in which each node i learns a different parameter $\Phi_{ij}^{(k)}$ for each neighbor j . These two cases are analogous to CNNs processing time signals with conventional convolutional filters as opposed to a neural network that operates with arbitrarily time varying filters (i.e., filters whose coefficients change in time). The appealing intermediate solution is to use filters with controlled edge variability to mix the advantage of a permutation equivariant parameterization (Proposition 1) with the processing of local detail. We introduce architectures that construct different versions of filters with controlled edge variability in Sections 4.1, 4.2, and 4.3.

Remark 4. Along with the above-referred works, also the works in [15], [16], [17] and [13] use versions of the convolutional filter in (14). In specific, [15], [16], [17] consider

² Which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2021.3111054>.

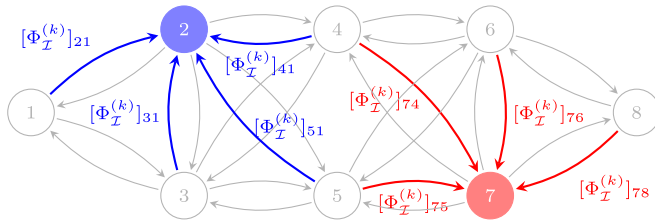


Fig. 3. Hybrid edge varying filter [cf. (17)]. The nodes in set $\mathcal{I} = \{2, 7\}$ are highlighted. Nodes 2 and 7 have edge varying parameters associated with their incident edges. All nodes, including 2 and 7, also use the global parameter a_k as in a regular convolutional graph filter.

single shifts on the graph with shift operator a learnable weight matrix, a Gaussian kernel, and a random-walk, respectively. The work in [13] adopts multi-layer perceptrons along the feature dimension at each node, before exchanging information with their neighbors. This is equivalent to (12) with the first layers having order $K = 0$ (depending on the depth of the MLP), followed by a final layer of order $K = 1$.

4.1 GNNs With Controlled Edge Variability

To build a GNN that fits between a permutation equivariant GCNN [cf. (14)] and a full edge varying GNN [cf. (5)], we use different filter parameters in different parts of the graph. Formally, let $\mathcal{B} = \{\mathcal{B}_1, \dots, \mathcal{B}_B\}$ be a partition of the node set into B blocks with block \mathcal{B}_i having B_i nodes. Define the tall matrix $\mathbf{C}_B \in \{0, 1\}^{N \times B}$ such that $[\mathbf{C}_B]_{ij} = 1$ if node i belongs to block \mathcal{B}_j and 0 otherwise. Let also $\mathbf{a}_B^{(k)} \in \mathbb{R}^B$ be a vector of block parameters of filter order k . *Block varying* graph filters are then defined as

$$\mathbf{A}(\mathbf{S}) = \sum_{k=0}^K \text{diag}(\mathbf{C}_B \mathbf{a}_B^{(k)}) \mathbf{S}^k. \quad (15)$$

Filters in (15) use parameters $[\mathbf{a}_B^{(k)}]_i$ for all nodes $i \in \mathcal{B}_i$.

Block varying filters belong to the family of node varying graph filters [28] and are of the form in (5) with

$$\Phi^{(k,0)} = \text{diag}(\mathbf{C}_B \mathbf{a}_B^{(k)}) \mathbf{S}^k. \quad (16)$$

Substituting (16) into (5) generates block varying GNNs [43]. Block varying GNNs have $B(K+1)F^2$ parameters per layer and a computational complexity of order $\mathcal{O}(KF^2M)$.

Alternatively, we can consider what we call *hybrid* filters that are defined as linear combinations of convolutional filters and edge varying filters that operate in a subset of nodes —see Fig. 3. Formally, let $\mathcal{I} \subset \mathcal{V}$ denote an important subset of $I = |\mathcal{I}|$ nodes and define the shift matrices $\Phi_{\mathcal{I}}^{(k)}$ such that the diagonal matrix $\Phi_{\mathcal{I}}^{(0)}$ has entries $[\Phi_{\mathcal{I}}^{(0)}]_{ii} \neq 0$ for all $i \in \mathcal{I}$ and $[\Phi_{\mathcal{I}}^{(k)}]_{ij} = 0$ for all $i \notin \mathcal{I}$ or $(i, j) \notin \mathcal{E}$ and $k \geq 1$. That is, the parameter matrices $\Phi_{\mathcal{I}}^{(k)}$ may contain nonzero elements only at rows i that belong to set \mathcal{I} and with the node j being a neighbor of i . We define hybrid filters as those of the form

$$\mathbf{A}(\mathbf{S}) = \sum_{k=0}^K \left(\prod_{k'=0}^k \Phi_{\mathcal{I}}^{(k')} + a_k \mathbf{S}^k \right). \quad (17)$$

Substituting (17) in (5) generates hybrid GNNs. In essence, nodes $i \in \mathcal{I}$ learn edge dependent parameters which may

also be different at different nodes, while nodes $i \notin \mathcal{I}$ learn global parameters.

Hybrid filters are defined by a number of parameters that depends on the total neighbors of all nodes in the importance set \mathcal{I} . Define then $M_{\mathcal{I}} = \sum_{i \in \mathcal{I}} N_i$ and observe $\Phi_{\mathcal{I}}^{(0)}$ has I nonzero entries since it is a diagonal matrix, while $\Phi_{\mathcal{I}}^{(k)}$ for $k \geq 1$ have respectively $M_{\mathcal{I}}$ nonzero values. We then have $KM_{\mathcal{I}} + I$ parameters in the edge varying filters and $K+1$ parameters in the convolutional filters. We therefore have a total of $(I + KM_{\mathcal{I}} + K + 1)F^2$ parameters per layer in a hybrid GNN. The implementation cost of a hybrid GNN layer is of order $\mathcal{O}(KF^2(M + N))$ since both terms in (17) respect the graph sparsity.

Block GNNs depend on the choice of blocks \mathcal{B} and hybrid GNNs on the choice of the importance set \mathcal{I} . We explore the use of different heuristics based on centrality and clustering measures in Section 6 where we will see that the choice of \mathcal{B} and \mathcal{I} is in general problem specific.

4.2 Spectral Graph Convolutional Neural Networks

The convolutional operation of the graph filter in (14) can be represented in the spectral domain. To do so, consider the input-output relationship $\mathbf{u} = \mathbf{A}(\mathbf{S})\mathbf{x}$ along with the eigenvector decomposition of the shift operator $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$. Projecting the input and output signals in the eigenvector space of \mathbf{S} creates the so-called graph Fourier transforms $\tilde{\mathbf{x}} := \mathbf{V}^{-1}\mathbf{x}$ and $\tilde{\mathbf{u}} := \mathbf{V}^{-1}\mathbf{u}$ [44] which allow us to write

$$\tilde{\mathbf{u}} := \left(\sum_{k=0}^K a_k \mathbf{\Lambda}^k \right) \tilde{\mathbf{x}}. \quad (18)$$

Eq. (18) reveals convolutional graph filters are pointwise in the spectral domain, due to the diagonal nature of the eigenvalue matrix $\mathbf{\Lambda}$. We can, therefore, define the filter's spectral response $a : \mathbb{R} \rightarrow \mathbb{R}$ as the function

$$a(\lambda) = \sum_{k=0}^K a_k \lambda^k, \quad (19)$$

which is a single-variable polynomial characterizing the graph filter $\mathbf{A}(\mathbf{S})$. If we allow for filters of order $K = N - 1$, there is always a set of parameters a_k such that $a(\lambda_i) = \tilde{a}_i$ for any set of spectral response \tilde{a}_i [25]. Thus, training over the set of spectral parameters $a(\lambda_1), \dots, a(\lambda_N)$ is equivalent to training over the space of (nodal) parameters a_0, \dots, a_{N-1} . GCNNs were first introduced in [8] using the spectral representation of graph filters in (19).

By using edge varying graph filters [cf. (5)], we can propose an alternative parameterization of the space of filters of order N which we will see may have some advantages. To explain this better let \mathcal{J} be the index set defining the zero entries of $\mathbf{S} + \mathbf{I}_N$ and let $\mathbf{C}_{\mathcal{J}} \in \{0, 1\}^{|\mathcal{J}| \times N^2}$ be a binary selection matrix whose rows are those of \mathbf{I}_{N^2} indexed by \mathcal{J} . Let also \mathbf{B} be a basis matrix that spans the null space of

$$\mathbf{C}_{\mathcal{J}} \text{vec}(\mathbf{V}^{-1} * \mathbf{V}), \quad (20)$$

where $\text{vec}(\cdot)$ is the column-wise vectorization operator and $"*" is the Khatri-Rao product [45]. Then, the following proposition from [33] quantifies the spectral response of a particular class of the edge varying graph filter in (5).$

Proposition 2. Consider the subclass of the edge varying graph filters in (5) where the parameter matrices $[\Phi^{(0)} + \Phi^{(1)}]$ and $\Phi^{(k)}$ for all $k = 2, \dots, K$ are restricted to the ones that share the eigenvectors with \mathbf{S} , i.e., $[\Phi^{(0)} + \Phi^{(1)}] = \mathbf{V}\mathbf{\Lambda}^{(1)}\mathbf{V}^{-1}$ and $\Phi^{(k)} = \mathbf{V}\mathbf{\Lambda}^{(k)}\mathbf{V}^{-1}$ for all $k = 2, \dots, K$. The spectral response of this subclass of edge varying filter has the form

$$a(\mathbf{\Lambda}) = \sum_{k=1}^K \left(\prod_{k'=1}^k \mathbf{\Lambda}^{(k')} \right) = \sum_{k=1}^K \prod_{k'=1}^k \text{diag}(\mathbf{B}\boldsymbol{\mu}^{(k')}), \quad (21)$$

where \mathbf{B} is an $N \times b$ basis kernel matrix that spans the null space of (20) and $\boldsymbol{\mu}^{(k)}$ is a $b \times 1$ vector containing the expansion parameters of $\mathbf{\Lambda}^{(k)}$ into \mathbf{B} .

Proof. See Appendix B, available in the online supplemental material. \square

Proposition 2 provides a subclass of the edge varying graph filters where, instead of learning $K(M+N)+N$ parameters, they learn the Kb entries $\boldsymbol{\mu}^{(1)}, \dots, \boldsymbol{\mu}^{(K)}$ in (21). These filters build the output features as a pointwise multiplication between the filter spectral response $a(\mathbf{\Lambda})$ and the input spectral transform $\tilde{\mathbf{x}} = \mathbf{V}^{-1}\mathbf{x}$, i.e., $\mathbf{u} = \mathbf{V}a(\mathbf{\Lambda})\tilde{\mathbf{x}} = \mathbf{V}a(\mathbf{\Lambda})\mathbf{V}^{-1}\mathbf{x}$. Following then the analogies with conventional signal processing, (21) represents the spectral response of a convolutional edge varying graph filter. Spectral GCNNs are a particular case of (21) with order $K = 1$ and kernel \mathbf{B} independent from the graph (e.g., a spline kernel). Besides generalizing [8], a graph-dependent kernel allows to implement (21) in the vertex domain through an edge varying filter of the form (5); hence, having a complexity of order $\mathcal{O}(K(M+N))$ in contrast to $\mathcal{O}(N^2)$ required for the graph-independent kernels. The edge varying implementation captures also local detail up to a region of radius K from a node; yet, having a spectral interpretation. Nevertheless, both the graph-dependent GNN [cf. (21)] and the graph-independent GNN [8] are more of theoretical interest since they require the eigendecomposition of the shift operator \mathbf{S} . This aspect inadvertently implies a cubic complexity in the number of nodes and an accurate learning process will suffer from numerical instabilities since it requires an order $K \approx N$; hence, high order matrix powers \mathbf{S}^k .

4.3 ARMA Graph Convolutional Neural Networks

We can increase the descriptive power of the filter in (14) by growing its order K , which allows learning filters with a more discriminative polynomial frequency response [cf. (19)]. However, this also increases the parameters and computational cost. Most importantly, it introduces numerical issues associated with high order matrix powers \mathbf{S}^k , ultimately, leading to poor interpolatory and extrapolatory performance [46]. These challenges can be overcome by considering graph filters with a rational spectral response, since rational functions have better interpolatory and extrapolatory properties than polynomials [46], [47], [48]. Rational functions can also achieve more complicated responses with lower degrees in both numerator and denominator, thus, having less learnable parameters. Autoregressive moving average (ARMA) graph filters [29] serve for such purpose and implement rational functions of the form

$$\mathbf{A}(\mathbf{S}) = \left(\mathbf{I} + \sum_{p=1}^P a_p \mathbf{S}^p \right)^{-1} \left(\sum_{q=0}^Q b_q \mathbf{S}^q \right) := \mathbf{P}^{-1}(\mathbf{S})\mathbf{Q}(\mathbf{S}), \quad (22)$$

where we have defined $\mathbf{P}(\mathbf{S}) := \mathbf{I} + \sum_{p=1}^P a_p \mathbf{S}^p$ and $\mathbf{Q}(\mathbf{S}) := \sum_{q=0}^Q b_q \mathbf{S}^q$. The ARMA filter in (22) is defined by P denominator parameters $\mathbf{a} = [a_1, \dots, a_P]^\top$ and $Q+1$ numerator parameters $\mathbf{b} = [b_0, \dots, b_Q]^\top$. The input-output relationship $\mathbf{u} = \mathbf{A}(\mathbf{S})\mathbf{x}$ of the ARMA filter can be represented in the spectral domain as [cf. (18)]

$$\tilde{\mathbf{u}} = \left(\mathbf{I} + \sum_{p=1}^P a_p \mathbf{\Lambda}^p \right)^{-1} \left(\sum_{q=0}^Q b_q \mathbf{\Lambda}^q \right) \tilde{\mathbf{x}}. \quad (23)$$

It follows that ARMA filters are also pointwise operators in the spectral domain characterized by the rational spectral response function

$$a(\lambda) = \left(\sum_{q=0}^Q b_q \lambda^q \right) / \left(1 + \sum_{p=1}^P a_p \lambda^p \right). \quad (24)$$

In particular, it follows the space of ARMA filters defined by (22) is equivalent to the space of spectral ARMA filters defined by (24) which is equivalent to the space of spectral filters in (19) and, in turn, equivalent to the graph convolutional filters in (14). That they are equivalent does not mean they have the same properties. We expect ARMA filters produce useful spectral responses with less parameters than the convolutional filters in (14) or the spectral filters in (19).

As it follows from (22), we need to compute the inverse matrix $\mathbf{P}^{-1}(\mathbf{S})$ to get the ARMA output. The latter incurs a cubic complexity, which unless the graph is of limited dimensions is computationally unaffordable. When the graph is large, we need an iterative method that exploits the sparsity of the graph to approximate the inverse with a reduced cost [29], [48]. Due to its faster convergence, we consider a parallel structure that consists of first transforming the polynomial ratio in (18) in its partial fraction decomposition form and subsequently using the Jacobi method to approximate inverse. While also other Krylov approaches are possible, the parallel Jacobi method offers a better tradeoff between computational complexity and convergence rate.

Partial Fraction Decomposition of ARMA Filters. The partial fraction decomposition of the rational function $a(\lambda)$ in (24) provides an equivalent representation of ARMA filters. Let $\boldsymbol{\gamma} = [\gamma_1, \dots, \gamma_P]^\top$ be a set of poles, $\boldsymbol{\beta} = [\beta_1, \dots, \beta_P]^\top$ a corresponding set of residuals and $\boldsymbol{\alpha} = [\alpha_0, \dots, \alpha_K]^\top$ be a set of direct terms; we can then rewrite (24) as

$$a(\lambda) = \sum_{p=1}^P \frac{\beta_p}{\lambda - \gamma_p} + \sum_{k=0}^K \alpha_k \lambda^k, \quad (25)$$

where $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ are computed from \mathbf{a} and \mathbf{b} . A graph filter whose spectral response is as in (25) is one in which the spectral variable λ is replaced by the shift operator variable \mathbf{S} . It follows that if $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$ are chosen to make (25) and (24) equivalent, the filter in (22) is, in turn, equivalent to

$$\mathbf{A}(\mathbf{S}) = \sum_{p=1}^P \beta_p \left(\mathbf{S} - \gamma_p \mathbf{I} \right)^{-1} + \sum_{k=0}^K \alpha_k \mathbf{S}^k. \quad (26)$$

The equivalence of (22) and (26) means that instead of training \mathbf{a} and \mathbf{b} in (22) we can train α , β , and γ in (26).

Jacobi Implementation of Single-Pole Filters. To circumvent the matrix inverses in (26), we first consider each single-pole filter in (26) separately and implement the input-output relationship

$$\mathbf{u}_p = \beta_p \left(\mathbf{S} - \gamma_p \mathbf{I} \right)^{-1} \mathbf{x}. \quad (27)$$

Expression (27) is equivalent to the linear equation $(\mathbf{S} - \gamma_p \mathbf{I})\mathbf{u}_p = \beta_p \mathbf{x}$, which we can solve iteratively through a Jacobi recursion. This requires us to separate $(\mathbf{S} - \gamma_p \mathbf{I})$ into diagonal and off-diagonal components. We, therefore, begin by defining the diagonal degree matrix $\mathbf{D} = \text{diag}(\mathbf{S})$ so that the shift operator can be written as

$$\mathbf{S} = \mathbf{D} + (\mathbf{S} - \mathbf{D}) := \text{diag}(\mathbf{S}) + (\mathbf{S} - \text{diag}(\mathbf{S})). \quad (28)$$

With this definition, we write $(\mathbf{S} - \gamma_p \mathbf{I}_N) = (\mathbf{D} - \gamma_p \mathbf{I}_N) + (\mathbf{S} - \mathbf{D})$, which is a decomposition on diagonal terms $(\mathbf{D} - \gamma_p \mathbf{I}_N)$ and off-diagonal terms $(\mathbf{S} - \mathbf{D})$. The Jacobi iteration k for (27) is given by the recursive expression

$$\mathbf{u}_{pk} = - \left(\mathbf{D} - \gamma_p \mathbf{I}_N \right)^{-1} \left[\beta_p \mathbf{x} - (\mathbf{S} - \mathbf{D}) \mathbf{u}_{p(k-1)} \right], \quad (29)$$

initialized with $\mathbf{u}_{p0} = \mathbf{x}$. We can unroll this iteration to write an explicit relationship between \mathbf{u}_{pk} and \mathbf{x} . To do that, we define the parameterized shift operator

$$\mathbf{R}(\gamma_p) = - \left(\mathbf{D} - \gamma_p \mathbf{I}_N \right)^{-1} (\mathbf{S} - \mathbf{D}), \quad (30)$$

and use it to write the K 'th iterate of the Jacobi recursion as

$$\mathbf{u}_{pK} = \beta_p \sum_{k=0}^{K-1} \mathbf{R}^k(\gamma_p) \mathbf{x} + \mathbf{R}^K(\gamma_p) \mathbf{x}. \quad (31)$$

For a convergent Jacobi recursion, signal \mathbf{u}_{pK} in (31) converges to the output \mathbf{u}_p of the single-pole filter in (27). Truncating (31) at a finite K yields an approximation in which single-pole filters are written as polynomials on the shift operator $\mathbf{R}(\gamma_p)$. I.e., a single-pole filter is approximated as a convolutional filter of order K [cf. (14)] in which the shift operator of the graph \mathbf{S} is replaced by the shift operator $\mathbf{R}(\gamma_p)$ defined in (30). This convolutional filter uses parameters β_p for $k = 0, \dots, K-1$ and 1 for $k = K$.

Jacobi ARMA Filters and Jacobi ARMA GNNs. Assuming we use Jacobi iterations to approximate all single-pole filters in (26) and that we truncate all of these iterations at K , we can write ARMA filters as

$$\mathbf{A}(\mathbf{S}) = \sum_{p=1}^P \mathbf{H}_K(\mathbf{R}(\gamma_p)) + \sum_{k=0}^K \alpha_k \mathbf{S}^k, \quad (32)$$

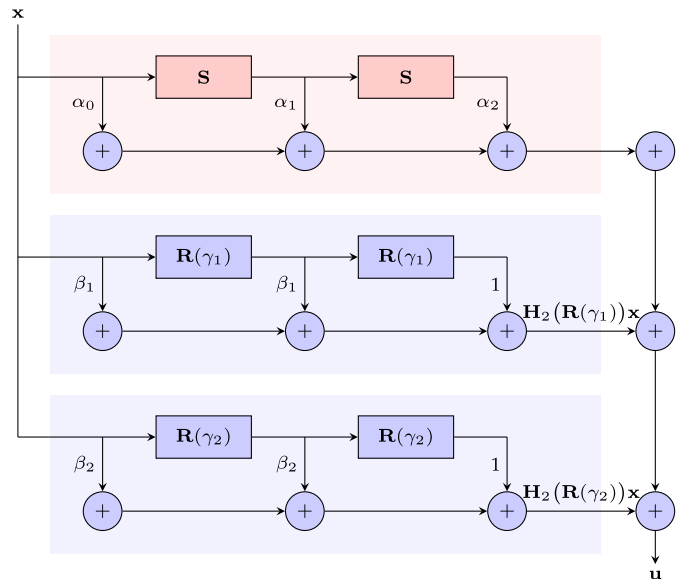


Fig. 4. Jacobi autoregressive moving average filter. The input signal x is processed by a parallel bank of filters. One of this filters is a convolutional filter of the form in (14) operating w.r.t. the shift operator \mathbf{S} (highlighted in red). The remaining filters operate w.r.t. scaled shift operators [cf. (30)] (highlighted in blue). All filter outputs are summed together to yield the overall Jacobi ARMA output.

where $\mathbf{H}_K(\mathbf{R}(\gamma_p))$ is a K order Jacobi approximation of the ARMA filter, which, as per (31) is given by

$$\mathbf{H}_K(\mathbf{R}(\gamma_p)) = \beta_p \sum_{k=0}^{K-1} \mathbf{R}^k(\gamma_p) + \mathbf{R}^K(\gamma_p). \quad (33)$$

A Jacobi ARMA filter of order (P, K) is defined by (32) and (33), see Fig. 4. The order P represents the number of poles in the filter and the order K the number of Jacobi iterations we consider appropriate to properly approximate individual single-pole filters. Notice the number of taps K in the filter $\sum_{k=0}^K \alpha_k \mathbf{S}^k$ need not be the same as the number of Jacobi iterations used in (33). But we use the same to avoid complicating notation.

For sufficiently large K (32)–(33), (26), and (22) are all equivalent expressions of ARMA filters of orders (P, Q) . We could train parameters using either of these equivalent expressions but we advocate for the use (32)–(33) as no inversions are necessary except for the elementary inversion of the diagonal matrix $(\mathbf{D} - \gamma_p \mathbf{I})$. It is interesting to note that in this latter form ARMA filters are reminiscent of the convolutional filters in (14) but the similarity is superficial. In (14), we train $K+1$ parameters a_k that multiply shift operator powers \mathbf{S}^k . In (32)–(33) we also train $K+1$ parameters of this form in the filter $\sum_{k=0}^K \alpha_k \mathbf{S}^k$ but this is in addition to the parameters β_p and γ_p of each of the single-pole filter approximations $\mathbf{H}_K(\mathbf{R}(\gamma_p))$. These single-pole filters are themselves reminiscent of the convolutional filters in (14) but the similarity is again superficial. Instead of parameters a_k that multiply shift operator powers \mathbf{S}^k , the filters in (33) train a parameters γ_p which represents a constant that is subtracted from the diagonal entries of the shift operators \mathbf{S} . The fact this is equivalent to an ARMA filter suggests (32)–(33) may help designing more discriminative filters. We corroborate in Section 6 that GNNs using (32)–(33) outperform GNNs that utilize filters (14).

An ARMA GNN has $(2P + K + 1)F^2$ parameters per layer and a computational complexity of order $\mathcal{O}(F^2P(MK + N))$. This decomposes as $\mathcal{O}(PN)$ to invert the diagonal matrices $(\mathbf{D} - \gamma_p \mathbf{I}_N)$; $\mathcal{O}(PM)$ to scale the nonzeros of $(\mathbf{S} - \mathbf{D})$ by the inverse diagonal; $\mathcal{O}(PKM)$ to obtain the outputs of Jacobi ARMA filters (32) of order K .

ARMA GNNs as EdgeNets. ARMA GNNs (ARMANets) are another subclass of the EdgeNet. To see this, consider that each shift operator $\mathbf{R}(\gamma_p)$ in (30) shares the support with $\mathbf{I}_N + \mathbf{S}$. Hence, we can express the graph filter in (32) as the union of $P + 1$ edge varying graph filters. The first P of these filters have parameters matrices of the form

$$\Phi_p^{(k:0)} = \begin{cases} \beta_p \mathbf{R}^k(\gamma_p), & k = 0, \dots, K - 1 \\ \mathbf{R}^K(\gamma_p), & k = K \end{cases},$$

while the last filter captures the direct-term with edge varying parameter matrices $\Phi^{(k:0)} = \alpha_k \mathbf{S}^K$ [cf. Section 4]. The union of these edge varying filter has the expression

$$\mathbf{A}(\mathbf{S}) = \sum_{k=0}^K \left(\sum_{p=1}^P \Phi_p^{(k:0)} + \Phi^{(k:0)} \right) \mathbf{x}, \quad (34)$$

which by grouping further the terms of the same order k leads to a single edge varying graph filter of the form in (5).

ARMANet provides an alternative parameterization of the EdgeNet that is different from that of the other polynomial convolutional filters in (14). In particular, ARMANets promote to use multiple polynomial filters of smaller order (i.e., the number of Jacobi iterations) with shared parameters between them. Each of the filters $\mathbf{H}_K(\mathbf{R}(\gamma_p))$ depends on two parameters β_p and γ_p . We believe this parameter sharing among the different orders and the different nodes is the success behind the improved performance of the ARMA GNN compared with the single polynomial filters in (14). Given also the hybrid solutions developed in Section 4.1 for the polynomial filters, a direction that may attain further improvements is that of ARMA GNN architectures with controlled edge variability.

ARMANet generalizes the architecture in [14] where instead of restricting the polynomials in (22) to Cayley polynomials, it allows the use of general polynomials. Iterative approaches to implement an ARMA graph filter [29], [49] have been recently used to develop GNNs that resemble this recursion [50], [51]. However, both [50], [51] do not implement an ARMA layer like (23) does. Instead, their propagation rule can be either seen as that used in graph recurrent neural networks [52, Eq. (5)] but with a constant input or conceptually similar to that used in iterative sparse coding [53].

5 GRAPH CONVOLUTIONAL ATTENTION NETWORKS

Graph convolutional neural networks parameterize the EdgeNet by using a fixed underlying shift operator \mathbf{S} and learning only the filter parameters. However, this shift operator may often be estimated from data disjointly from the GNN task, or its specific weights may be unknown. In these instances, we can use the graph attention mechanism

[18] to parameterize the EdgeNet in a way that we learn both the shift operator weights and the convolutional filter parameters for the task at hand. We propose the graph convolutional attention network (GCAT), which utilizes the filters as in (12) but they are convolutional in a layer-specific matrix $\Phi_l = \Phi$ that may be different from the shift operator \mathbf{S}

$$\mathbf{X}_l = \sigma \left(\sum_{k=0}^K \Phi^k \mathbf{X}_{l-1} \mathbf{A}_k \right). \quad (35)$$

Note $\mathbf{A}_k = \mathbf{A}_{lk}$ and $\Phi = \Phi_l$ are layer-dependent but we omit the layer index to simplify notation. Since matrix Φ shares the sparsity pattern of \mathbf{S} , (35) defines a GNN as per (7). Matrix Φ is learned from the features \mathbf{X}_{l-1} passed from layer $l - 1$ following the attention mechanism [18]. Specifically, we consider a trainable matrix $\mathbf{B} \in \mathbb{R}^{F_{l-1} \times F_l}$ and vector $\mathbf{e} \in \mathbb{R}^{2F_l}$, and compute the edge scores

$$\alpha_{ij} = \sigma \left(\mathbf{e}^\top \left[\begin{array}{c} [\mathbf{X}_{l-1} \mathbf{B}]_i \\ [\mathbf{X}_{l-1} \mathbf{B}]_j \end{array} \right] \right), \quad (36)$$

for all edges $(i, j) \in \mathcal{E}$. In (36), we start with the vector of features \mathbf{X}_{l-1} and mix them as per the parameters in \mathbf{B} . This produces a collection of graph signals $\mathbf{X}_{l-1} \mathbf{B}$ in which each node i has F_l features that correspond to the i th row $[\mathbf{X}_{l-1} \mathbf{B}]_i$ of the product matrix $\mathbf{X}_{l-1} \mathbf{B}$. The features at node i are concatenated with the features of node j and the resulting vector of $2F_l$ components is multiplied by vector \mathbf{e} . This product produces the score α_{ij} after passing through the nonlinearity $\sigma(\cdot)$. Note that $\mathbf{B} = \mathbf{B}_l$, $\mathbf{e} = \mathbf{e}_l$ are global parameters for all scores $\alpha_{ij} = \alpha_{lij}$ and depend on the layer index l . As is the case of \mathbf{A}_k and Φ in (35), we omitted this index for simplicity.

The score α_{ij} could be used directly as an entry for the matrix Φ but to encourage attention sparsity we pass α_{ij} through a local soft maximum operator

$$\Phi_{ij} = \exp(\alpha_{ij}) \left(\sum_{j' \in \mathcal{N}_i \cup i} \exp(\alpha_{ij'}) \right)^{-1}. \quad (37)$$

The soft maximum assigns edge weights Φ_{ij} close to 1 to the largest of the edge scores α_{ij} and weights Φ_{ij} close to 0 to the rest. See also Fig. 5a.

In Section 2, we introduced arbitrary edge varying graph filters [cf. (5)] which we leveraged in Section 3 to build edge varying GNNs [cf. (6) - (7)]. In Section 4, we pointed out that edge varying graph filters left too many degrees of freedom in the learning parametrization; a problem that we could overcome with the use of graph convolutional filters [cf. (14)]. The latter suffer from the opposite problem as they may excessively constrict the GNN. GATs provide a solution of intermediate complexity. Indeed, the filters in (35) allow us to build a GNN with convolutional graph filters where the shift operator Φ is learned ad hoc in each layer to represent the required abstraction between nodes. The edges of this shift operator try to choose neighbors whose values should most influence the computations at a particular node. This is as in any arbitrary edge varying graph filter but the novelty of GATs is to reduce the number of learnable parameters by tying edge values to matrix \mathbf{B} and vector \mathbf{e} —observe that

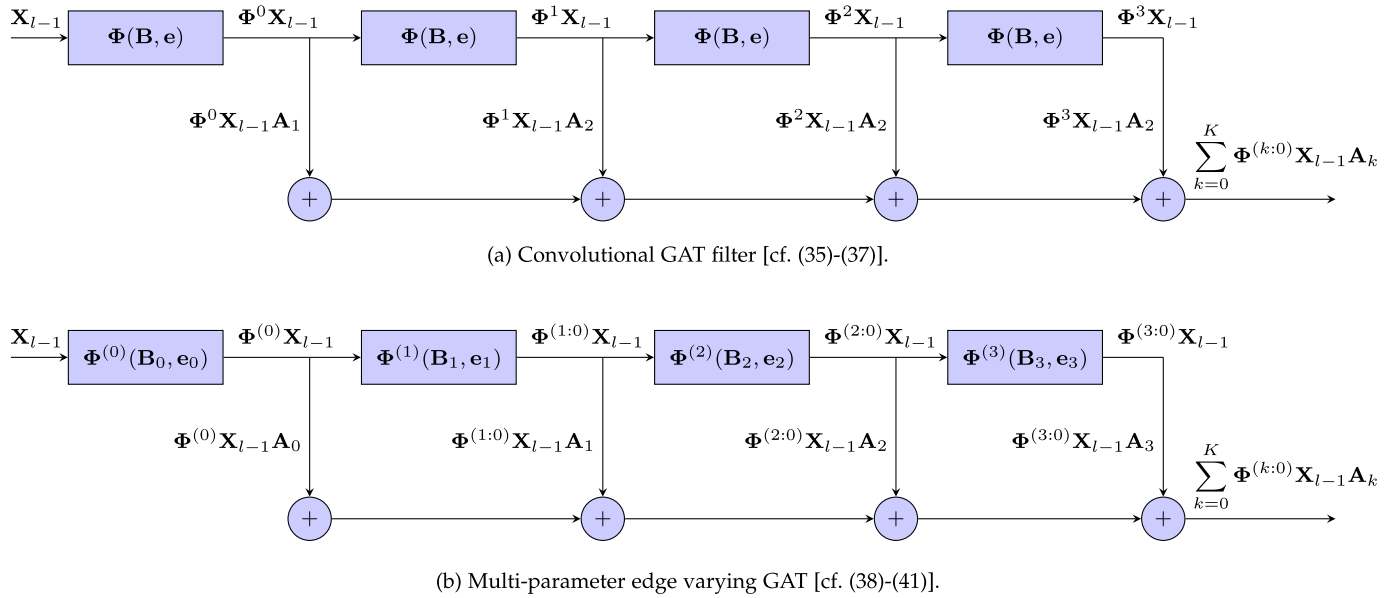


Fig. 5. Higher-order graph attention filters. (a) Graph convolutional attention filter. The input features \mathbf{X}_{l-1} are shifted by the same edge varying shift operator $\Phi(\mathbf{B}, \mathbf{e})$ and weighted by different parameter matrices \mathbf{A}_k . The edge varying parameters in all $\Phi(\mathbf{B}, \mathbf{e})$ are parameterized by the same matrix \mathbf{B} and vector \mathbf{e} following the attention mechanism. (b) Edge varying GAT filter. The input features \mathbf{X}_{l-1} are shifted by different edge varying shift operators $\Phi^{(k)}(\mathbf{B}_k, \mathbf{e}_k)$ and weighted by different parameter matrices \mathbf{A}_k . The edge varying parameters in the different $\Phi^{(k)}(\mathbf{B}_k, \mathbf{e}_k)$ are parameterized by a different matrix \mathbf{B}_k and vector \mathbf{e}_k following the attention mechanism.

in (36) \mathbf{e} is the same for all edges. Thus, the computation of scores α_{ij} depends on the $F_{l-1} \times F_l$ parameters in \mathbf{B} and the $2F_l$ parameters in \mathbf{e} . This is of order no more than F^2 if we make $F = \max_l F_l$. It follows that for the GAT in (35) the number of learnable parameters is at most $F^2 + 2F + F^2(K+1)$, which depends on design choices and is independent of the number of edges. We point out that since Φ respects the graph sparsity, the computational complexity of implementing (35) and its parameterization is of order $\mathcal{O}(F(NF + KM))$.

5.1 Edge Varying GAT Networks

The idea of using attention mechanisms to estimate entries of a shift operator Φ can be extended to estimate entries $\Phi^{(k:0)}$ of an edge varying graph filter. To be specific, we propose to implement a generic GNN as defined by recursion (7) which we repeat here for ease of reference

$$\mathbf{X}_l = \sigma \left(\sum_{k=0}^K \Phi^{(k:0)} \mathbf{X}_{l-1} \mathbf{A}_k \right). \quad (38)$$

Further recall each edge varying filter parameter matrix $\Phi^{(k:0)}$ is itself defined recursively as [cf. (5)]

$$\Phi^{(k:0)} = \Phi^{(k)} \Phi^{(k-1:0)} = \prod_{k'=0}^k \Phi^{(k')}. \quad (39)$$

We propose to generalize (36) so that we compute a different matrix $\Phi^{(k)}$ for each filter order k . Consider then matrices \mathbf{B}_k and vectors \mathbf{e}_k to compute the edge scores

$$\alpha_{ij}^{(k)} = \sigma \left(\mathbf{e}_k^T \begin{bmatrix} [\mathbf{X}_{l-1} \mathbf{B}_k]_i \\ [\mathbf{X}_{l-1} \mathbf{B}_k]_j \end{bmatrix} \right), \quad (40)$$

for all edges $(i, j) \in \mathcal{E}$. As in (36), we could use $\alpha_{ij}^{(k)}$ as edge weights in $\Phi^{(k)}$, but to promote attention sparsity we send $\alpha_{ij}^{(k)}$ through a soft maximum function to yield edge scores

$$\Phi_{ij}^{(k)} = \exp \left(\alpha_{ij}^{(k)} \right) \left(\sum_{j' \in \mathcal{N}_i \cup i} \exp \left(\alpha_{ij'}^{(k)} \right) \right)^{-1}. \quad (41)$$

Each of the edge varying matrices $\Phi^{(k)}$ for $k = 1, \dots, K$ is parameterized by the tuple of transform parameters $(\mathbf{B}_k, \mathbf{e}_k)$. Put simply, we are using a different GAT mechanism for each edge varying matrix $\Phi^{(k)}$. These learned matrices are then used to build an edge varying filter to process the features \mathbf{X}_l passed from the previous layer – see Fig. 5b. The edge varying GAT filter employs $K+1$ transform matrices \mathbf{B}_k of dimensions $F_l \times F_{l-1}$, $K+1$ vectors \mathbf{e}_k of dimensions $2F_l$, and $K+1$ matrices \mathbf{A}_k of dimensions $F_l \times F_{l-1}$. Hence, the total number of parameters for the edge varying GAT filter is at most $(K+1)(2F^2 + 2F)$. The computational complexity of the edge varying GAT is of order $\mathcal{O}(KF(NF + M))$.

Remark 5. Graph attention networks first appeared in [18].

In this paper, (36) and (37) are proposed as an attention mechanism for the signals of neighboring nodes and GNN layers are of the form $\mathbf{X}_l = \sigma(\Phi_l \mathbf{X}_{l-1} \mathbf{A}_l)$. The latter is a particular case of either (35) or (38) in which only the term $k=1$ is not null. Our observation in this section is that this is equivalent to computing a different graph represented by the shift operator Φ . This allows for the generalization to filters of arbitrary order K [cf. (35)] and to edge varying graph filters of arbitrary order [cf. (38)]. The approaches presented in this paper can likewise be extended with the multi-head attention mechanism proposed in [18] to improve the network capacity.

TABLE 1
Properties of Different Graph Neural Network Architectures

Architecture	Expression	Order parameters* $\mathcal{O}(\cdot)$	Order of complexity** $\mathcal{O}(\cdot)$
Fully connected	n/a	$N^2 F^2$	$N^2 F^2$
Edge varying	Eq. (7)	$K(M + N)F^2$	$KF^2(M + N)$
GCNN [9], [10], [11], [12]	Eq. (12)	KF^2	$KF^2 M$
Node varying [43]	Eq. (15)	BKF^2	$KF^2 M$
Hybrid edge varying	Eq. (17)	$(\mathcal{I} + KM_{\mathcal{I}})F^2$	$KF^2(M + N)$
Spec. edge varying GNN**	Eq. (21)	KbF^2	$KF^2(M + N)$
Spec. Kernel GCNN** [8]	Eq. (21) for $K = 1$	bF^2	$N^2 F^2$
ARMANet	Eq. (22)-(33)	$(P + K)F^2$	$F^2 P(MK + N)$
GCAT	Eq. (35)	$RF^2 K$	$R(NF^2 + KFM)$
GAT [18]	Eq. (35) for $K = 1$	RF^2	$R(NF^2 + FM)$
Edge varying GAT	Eq. (38)	RKF^2	$RK(NF^2 + MF)$

The parameters and complexity are considered per layer. Architectures in bold are proposed in this work. Legend: N - number nodes; M - number of edges; F - maximum number of features; K - recursion order; b - dimension of the Kernel in (21); B - number of blocks in (15); \mathcal{I} - the set of important nodes in (15) and (17); $M_{\mathcal{I}}$ - total neighbors for the nodes in \mathcal{I} ; P - parallel J-ARMANet branches; R - parallel attention branches; *Self-loops are not considered. **The eigendecomposition cost $\mathcal{O}(N^3)$ is computed once.

5.2 Discussions

Reducing Model Complexity. As defined in (36) and (40) the attention mechanisms are separate from filtering. To reduce the number of parameters, we can equate the attention matrices \mathbf{B} or \mathbf{B}_k with the filtering matrices \mathbf{A}_k . For the GCAT in (35), the original proposal in [18] is to make $\mathbf{B} = \mathbf{A}_1$ so that (36) reduces to

$$\alpha_{ij} = \sigma \left(\mathbf{e}^T \begin{bmatrix} [\mathbf{X}_{l-1} \mathbf{A}_1]_i \\ [\mathbf{X}_{l-1} \mathbf{A}_1]_j \end{bmatrix} \right). \quad (42)$$

For the edge varying GATs in (38), it is natural to equate $\mathbf{B}_k = \mathbf{A}_k$ in which case (40) reduces to

$$\alpha_{ij}^{(k)} = \sigma \left(\mathbf{e}_k^T \begin{bmatrix} [\mathbf{X}_{l-1} \mathbf{A}_k]_i \\ [\mathbf{X}_{l-1} \mathbf{A}_k]_j \end{bmatrix} \right). \quad (43)$$

The choice in (43) removes $(K + 1)F^2$ parameters.

Accounting for Differences in Edge Weights in the Original Shift Operator. The major benefit of the GAT mechanism is to build a GNN without requiring full knowledge of \mathbf{S} . This is beneficial as it yields a GNN robust to uncertainties in the edge weights. This benefit becomes a drawback when \mathbf{S} is well estimated as it renders weights s_{ij} equivalent regardless of their relative values. One possible solution to this latter drawback is to use a weighted soft maximum operator so that the entries of $\Phi^{(k)}$ are chosen as

$$\Phi_{ij}^{(k)} = \exp \left(s_{ij} \alpha_{ij}^{(k)} \right) \left(\sum_{j' \in \mathcal{N}_i \cup i} \exp \left(s_{ij'} \alpha_{ij'}^{(k)} \right) \right)^{-1}. \quad (44)$$

Alternatively, we can resort to the use of a hybrid GAT in which we combine a GAT filter of the form in (38) with a regular convolutional filter of the form in (12)

$$\mathbf{X}_l = \sigma \left(\sum_{k=0}^K \mathbf{S}^k \mathbf{X}_{l-1} \mathbf{A}_k + \Phi^{(k:0)} \mathbf{X}_{l-1} \mathbf{A}'_k \right). \quad (45)$$

This is the GAT version of the hybrid GNN we proposed in (17). The filters in (45) account for both, the GAT learned shifts $\Phi^{(k)}$ and the original shift \mathbf{S} .

6 NUMERICAL RESULTS

This section corroborates the capacity of the different models with numerical results on synthetic and real-world graph signal classification problems. Given the different hyper-parameters for the models and the trade-offs (e.g., complexity, number of parameters, radius of local information), we aim to provide insights on which methods exploit better the graph prior for learning purposes rather than achieving the best performance. In the sequel, we provide a summary of the different architectures, while in the upcoming four section analyze in detail each experiment. Section 6.5 provides some general observations. The code used for these simulations can be found at <http://github.com/alelab-upenn/graph-neural-networks>.

Summary of the Architectures. A summary of the different architectures discussed in this paper is provided in Table 1.

- Fully connected stands in for the standard multi-layer perceptron neural network that uses no graph prior.
- The edge varying stand in for the EdgeNet in the full form in (7), which allocates different parameters per edge and shift and has the most degrees of freedom.
- The graph convolutional neural network in (12) has instead shared parameters among all nodes, edges, and shifts.
- The node varying [cf. (15)] and the hybrid edge varying [cf. (17)] are the intermediate solutions between the full-form EdgeNet and the GCNN. Both rely on selecting a set of important nodes and allocate different parameters to them. The hybrid edge varying allocates also different parameters to the edges of these important nodes. To the remaining nodes shared parameters as in the GCNN are used.
- The spectral edge varying GNN is the subclass of the full-form EdgeNet in which filters share the eigenvectors with the shift operator [cf. Proposition 2], while the spectral kernel GCNN is the equivalent representation of the GCNN in the spectral domain.
- ARMANet is the GCNN that we proposed in Section 4.3. Contrarily to the polynomial GCNN form [cf.

(12)], it has a filter with a rational spectral response rather than polynomial.

- The graph convolutional attention network and its variants are the architectures we developed in Section 5. The GCAT generalises the popular GAT [18] to a convolutional filter of higher order such that it can account for multi-hop neighbor information in each layer. The edge varying GAT [cf. (38)] is an alternative to the GCAT that uses edge varying filters together with attention instead of convolutional filters. Lastly, we also discussed two variants of the GCAT and edge varying GAT in Section 5.2 but several others developed under the EdgeNet framework.

From all these architectures, we do not evaluate the fully-connected, the spectral kernel GCNN [8], and the spectral edge varying GCNN (21) since their computational cost is higher than linear. We also leave to interested readers the extensions discussed in Section 5.2. For the remaining solutions, we trained all of them using ADAM with the standard forgetting factors [54].

6.1 Source Localization on SBM Graphs

The goal of this experiment is to identify which community in a stochastic block model (SBM) graph is the source of a diffused signal by observing different realizations in different time instants. We considered a connected and undirected graph of $N = 50$ nodes divided into five blocks each representing one community $\{c_1, \dots, c_5\}$. The intra- and inter-community edge formation probabilities are 0.8 and 0.2, respectively. The source node is one of the five nodes (i_1, \dots, i_5) with the largest degree in the respective community. The source signal $\mathbf{x}(0)$ is a Kronecker delta centered at the source node. The source signal is diffused at time $t \in [0, 50]$ as $\mathbf{x}(t) = \mathbf{S}^t \mathbf{x}(0)$, where \mathbf{S} is the graph adjacency matrix normalized by the maximum eigenvalue.

The training set is composed of 10240 tuples of the form $(\mathbf{x}(t), c_i)$ for random t and $i \in \{1, \dots, 5\}$. These tuples are used to train the EdgeNets that are subsequently used to predict the source community c_i for a testing signal $\mathbf{x}'(t)$ again for a random value of t . The validation and the test set are both composed of 2560 tuples (25% of the training set). The performance of the different algorithms is averaged over ten different graph realizations and ten data splits, for a total of 100 Monte-Carlo iterations. ADAM is run for 40 epochs with batches of 100 samples and learning rate 10^{-3} .

Architecture Parameters. For this experiment, we compared 14 different architectures. All architectures comprise the cascade of a graph filtering layer with ReLU nonlinearity and a fully connected layer with softmax nonlinearity. The architectures are: *i*) the edge varying GNN (7); *ii*) the GCNN (12); *iii*) three node varying GNNs (15), where the five important nodes are selected based on *iii - a*) maximum degree; *iii - b*) spectral proxies [55], which ranks the nodes according to their contribution to different frequencies; *iii - c*) diffusion centrality (see Appendix C, available in the online supplemental material); *iv*) three node dependent edge varying GNNs, where the five important nodes \mathcal{B} are selected similar to the node varying case; *v*) three ARMANets (32) with Jacobi iterations *v - a*) $K = 1$; *v - b*) $K = 3$; *v - c*) $K = 5$; *vii*)

the GAT network from [18]; *vii*) the GCAT network (35); and *viii*) the edge varying GAT (38) network.

Our goal is to see how the different architectures handle their degrees of freedom, while all having linear complexity. To make this comparison more insightful, we proceed with the following rationale. For the approaches in *i*) – *iv*), we analyzed filter orders in the interval $K \in \{1, \dots, 5\}$. This is the only handle we have on these filters to control the number of parameters and locality radius while keeping the same computational complexity. For the ARMANet in *v*), we set the direct term order to $K = 0$ to observe only the effect of the rational part. Subsequently, for each Jacobi iteration value K , we analyzed rational orders in the interval $P \in \{1, \dots, 5\}$ as for the former approaches. While this strategy helps us controlling the local radius, recall the ARMANet has a computational complexity slightly higher than the former four architectures. For the GAT in *vi*), we analyzed different attention heads $R \in \{1, \dots, 5\}$ such that the algorithm complexity matches those of the approaches *i*) – *iv*). The number of attention heads is the only handle in the GAT network. Finally, for the GCAT in *vii*) and the edge varying GAT in *viii*), we fixed the attention heads to $R = 3$ and analyzed different filter orders $K \in \{1, \dots, 5\}$. The latter allows comparing the impact of the local radius for the median attention head value. Recall, these architectures have again a slightly higher complexity for $K \geq 2$.

Observations. The results of this experiment are shown in Fig. 6. We make the following observations.

First, the attention-based approaches are characterized by a slow learning rate leading to a poor performance in 40 epochs. This is reflected in the higher test error of the GAT and the edge varying GAT networks. However, this is not the case for the GCAT network. We attribute the latter reduced error to the superposition of the graph convolutional to attentions that GCAT explores –all convolutional approaches learn faster. On the contrary, the error increases further for the edge varying GAT since multiple attention strategies are adopted for all $k \in \{1, \dots, K\}$ in (38). Therefore, our conclusion is the graph convolutional prior can be significantly helpful for attention mechanisms. We will see this consistent improvement in all our experiments.

Second, the edge varying GNN in (7) achieves the lowest error, although having the largest number of parameters. The convolutional approaches parameterize well the edge varying filter; hence, highlighting the benefit of the graph convolution. ARMANet is the best among the latter characterized both by a lower mean error and standard deviation. This reduced error for ARMANet is not entirely surprising since rational functions have better interpolation and extrapolation properties than polynomial ones. It is, however, remarkable that the best result is obtained for a Jacobi iteration of $K = 1$. I.e., the parameter sharing imposed by ARMANet reaches a good local optimal even with a coarse approximation of the rational function. Notice also the source localization task is not permutation equivariant, therefore, architectures that are not permutation equivariant (edge varying, node varying, hybrid edge varying, ARMA for low Jacobi orders) are expected to perform better.

Third, for the node selection strategies in architectures *iii*) and *iv*), there is no clear difference between the degree and the diffusion centrality. For the node varying GNNs,

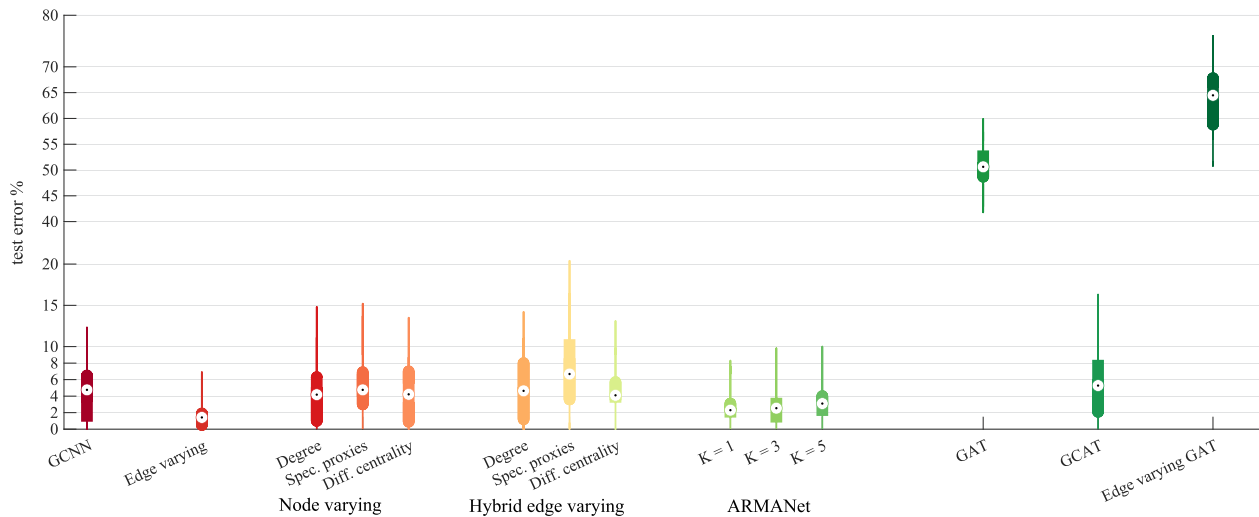


Fig. 6. Source localization test error in the stochastic block model graph. The y -axis scale is deformed to improve visibility. The thick bar interval indicates the average performance for different parameter choices (e.g., filter order, attention heads). The circle marker represents the mean value of this interval. The thin line spans an interval of one standard deviation from the average performance. The convolutional-based approaches perform better than attention-based. We attribute the poor performance of the attention techniques to the slow learning rate. Both the GAT and the edge variant GAT required more than 40 epochs to reach a local minimum. However, the graph convolutional attention network (GCAT) does not suffer from the latter issue leading to faster learning.

TABLE 2
Source Localization Test Error in Facebook Subnetwork

Architecture	mean	std. dev.	order {1, 2, 3, 4, 5}	attention heads {1, 2, 3, 4, 5}	epochs {10, 20, 40, 100}	learning rate { 10^{-2} , 10^{-3} }
GCNN	4.0%	13.0%	3	n/a	100	10^{-3}
Edge varying	1.5%	8.4%	1	n/a	10	10^{-3}
Node varying	6.0%	15.8%	3	n/a	20	10^{-3}
Hybrid edge var.	6.6%	15.9%	2	n/a	40	10^{-3}
ARMANet	2.0%	9.7%	1	n/a	20	10^{-3}
GAT	10.9%	20.8%	n/a	1	40	10^{-3}
GCAT	8.0%	18.4%	3	1	100	10^{-3}
Edge varying GAT	7.1%	17.8%	2	3	100	10^{-3}

The goal is to grid-search the parameters to achieve a mean error of at most 2%. For the architectures that did not achieve this criterion, the minimum error is reported.

the diffusion centrality offers a lower error both in the mean and deviation. In the hybrid edge varying GNNs [cf. (17)], instead, the degree centrality achieves a lower error but pays in deviation. The spectral proxy centrality yields the worst performance.

Finally, we remark that we did not find any particular trend while changing the parameters of the different GNNs (e.g., order, attention head). A rough observation is that low order recursions are often sufficient to reach low errors.

6.2 Source Localization on Facebook Sub-Network

In the second experiment, we considered a similar source localization on a real-world network comprising a 234-used Facebook subgraph obtained as the largest connected component of the dataset in [56]. This graph has two well-defined connected communities of different size and the objective is to identify which of the two communities originated the diffusion. The performance of the different algorithms is averaged over 100 Monte-Carlo iterations. The remaining parameters for generating the synthetic data are similar as before.

Architecture Parameters. We compared the eight GNN architectures reported in the left-most column of Table 2. For the node varying and the hybrid edge varying GNNs, the important nodes are again 10% of all nodes selected based on diffusion centrality. The Jacobi number of iterations for the ARMANet is $K = 1$ while there is no direct term.

Overall, this problem is easy to solve if the GNN is hyper-tuned with enough width and depth. However, this strategy hinders the impact of the specific filter. To highlight the role of the latter, we considered minimal GNN architectures composed of one layer and two features. In turn, this allows understanding better how much the specific filter contributes to the performance. We then grid-searched all parameters in Table 2 to reach a classification error of at most 2%. For the architectures that reach this criterion, we report the smallest parameters. For the architectures that do not reach this criterion, we report the minimum achieved error and the respective parameters. Our rationale is that the minimum parameters yield a lower complexity and show better the contribution of the filter type. They also lead to faster training; the opposite holds for the learning rate.

From Table 2, we observe that only the edge varying GNN and the ARMANet reach the predefined error. Both

architectures stress our observation that low order recursions ($K = 1$) are often sufficient. Nevertheless, this is not the case for all other architectures. These observations suggest the edge varying GNN explores well its degrees of freedom and adapts well to the non-permutation equivariance of the task. The ARMANet explores the best the convolutional prior; in accordance with the former results, the Jacobi implementation does not need to run until convergence to achieve impressive results. We also conclude the convolutional prior helps to reduce the degrees of freedom of the EdgeNet but requires a deeper and/or wider network to achieve the predefined criterion. This is particularly seen in the GAT based architectures. The GCAT architecture, in here, explores the convolutional prior and reduces the error compared with the GAT. Finally, we remark for all approaches a substantially lower variance can be achieved by solely increasing the features.

6.3 Authorship Attribution

In this third experiment, we assess the performance of the different GNN architectures in an authorship attribution problem based on real data. The goal is to classify if a text excerpt belongs to a specific author or any other of the 20 contemporary authors based on word adjacency networks (WANs) [57]. A WAN is an author-specific directed graph whose nodes are function words without semantic meaning (e.g., prepositions, pronouns, conjunctions). The relative positioning of function words carries stylistic information about the author. To capture this information, we build a directed graph, where each node is a function word, and each weighted edge represents the average co-occurrence of the corresponding function words, discounted by relative distance (i.e., if the two words are next to each other, the weight is higher than if the two words are further apart). We build this graph support only once, and before the training set. The signal on top of this graph is the frequency count for the function words in text excerpts of 1,000 words. These are the graph signals that form the dataset that is used for training and testing. The WANs and the word frequency count serve as author signatures and allow learning representation patterns in their writing style. The task translates into a binary classification problem where one indicates the text excerpt is written by the author of interest and zero by any other author. A more detailed account on the creation of WANs can be found in [57] and the used dataset is available in our code.

The WANs of the respective authors have from $N = 190$ to $N = 210$ function word nodes. Following [57], we built single-author WANs for Jane Austen, Emily Brontë, and Edgar Allan Poe. For each author, we processed the texts to count the number of times each function word pair co-appears in a window of ten words. These co-appearances are imputed into an $N \times N$ matrix and normalized row-wise. The resulting matrix is used as the shift operator, which can also be interpreted as a Markov chain transition matrix. We considered a train-test split of 95% – 5% of the available texts. Around 8.7% of the training samples are used for validation. This division leads to: (i) Austen: 1346 training samples, 118 validation samples, and 78 testing samples; (ii) Brontë: 1192 training samples, 104 validation samples, 68 testing samples; (iii) Poe: 740 training samples,

TABLE 3
Authorship Attribution Test Error

Architecture	Austen	Brontë	Poe
GCNN	7.2(± 2.0)%	12.9(± 3.5)%	14.3(± 6.4)%
Edge varying	7.1(± 2.2)%	13.1(± 3.9)%	10.7(± 4.3)%
Node varying	7.4(± 2.1)%	14.6(± 4.2)%	11.7(± 4.9)%
Hybrid edge var.	6.9(± 2.6)%	14.0(± 3.7)%	11.7(± 4.8)%
ARMANet	7.9(± 2.3)%	11.6(± 5.0)%	10.9(± 3.7)%
GAT	10.9(± 4.6)%	22.1(± 7.4)%	12.6(± 5.5)%
GCAT	8.2(± 2.9)%	13.1(± 3.5)%	13.6(± 5.8)%
Edge varying GAT	14.5(± 5.9)%	23.7(± 9.0)%	18.1(± 8.4)%

The results show the average classification test error and standard deviation on 10 different training-test 95% – 5% splits.

64 validation samples, 42 testing samples. For each author, the sets are extended by a similar amount with texts from the other 20 authors shared equally between them.

Architecture Parameters. We considered again the eight GNN architectures of the former section shown in the left-most column of Table 3. Following the setup in [10], all architectures comprise a graph neural layer of $F = 32$ features with ReLU nonlinearity followed by a fully connected layer. The baseline order for all filters is $K = 4$. For the ARMANet this is also the number of denominator parameters and the order of the direct term in (32); the number of the Jacobi iterations in (33) is one. We want to show how much the rational part helps to improve the performance of the GCNN (which is the direct term in the ARMANet [cf. (32)]). The important nodes for the node varying and the hybrid edge varying are 20 ($\sim 10\%$ of N) selected with degree centrality. The GAT, GCAT, and edge varying GAT have a single attention head to highlight the role of the convolutional and edge varying recursion over it. The loss function is the cross-entropy optimized over 25 epochs with a learning rate of 0.005. The performance is averaged over ten data splits.

Table 3 shows the results of this experiment. Overall, we see again the graph convolution is a solid prior to learning meaningful representations. This is particularly highlighted in the improved performance of the GCAT for Austen and Brontë compared with the GAT even with a single attention head. These observations also suggest the GAT and the edge varying GAT architectures require multi-head approaches to achieve comparable performance. An exception is the case of Poe. In this instance, multi-head attention is also needed for the GCAT. The (approximated) rational part of the ARMANet gives a consistent improvement of the GCNN. Hence, we recommend considering the additional parameterization of the ARMANet when implementing graph convolutional neural networks, since the increased number of parameters and implementation costs are minimal. Finally, we remark the hybrid edge varying GNN improves the accuracy of the node varying counterpart.

6.4 Recommender Systems

In this last experiment, we evaluate all former architectures for movie rating prediction in a subset of the MovieLens 100K data set [58]. The full data set comprises $U = 943$ users and $I = 1,582$ movies and 100K out of $\sim 1,5M$ potential ratings. We set the missing ratings to zero. From the incomplete $U \times I$ rating matrix, we consider two scenarios: A user-based and a

TABLE 4
Average RMSE on User Graph

Archit./User-ID	405	655	13	450	276	Average
GCNN	1.09	0.72	1.18	0.82	0.66	0.89
Edge var.	1.25	0.74	1.34	0.99	0.70	1.00
Node var.	1.17	0.68	1.19	0.83	0.67	0.91
Hybrid edge var.	1.10	0.72	1.27	0.80	0.60	0.90
ARMANet	1.13	0.69	1.24	0.80	0.65	0.90
GAT	1.27	0.74	1.44	0.92	0.80	1.03
GCAT	1.09	0.71	1.12	0.77	0.65	0.87
Edge var. GAT	1.19	0.70	1.31	0.85	0.75	0.96

TABLE 5
Average RMSE on Movie Graph

Archit./Movie-ID	50	258	100	181	294	Average
GCNN	0.82	1.08	0.95	0.86	1.04	0.95
Edge var.	0.93	1.03	1.00	0.88	1.24	1.02
Node var.	0.78	1.04	1.00	0.87	1.00	0.94
Hybrid edge var.	0.75	1.02	0.98	0.82	1.08	0.93
ARMANet	0.81	1.05	1.02	0.87	1.09	0.97
GAT	0.98	1.24	1.28	1.00	1.30	1.16
GCAT	0.83	1.06	1.04	0.83	1.05	0.96
Edge var. GAT	0.81	1.04	1.01	0.86	1.07	0.96

movie-based. In a user-based scenario, we considered the 200 users that have rated the most movies as the nodes of a graph whose edges represent Pearson similarities between any two users. Each of the $I = 1,582$ movies is treated as a different graph signal whose value at a node is the rating given to that movie by a user or zero if unrated. We are interested to predict the rating of a specific user u with GNNs, which corresponds to completing the u th row of the $200 \times 1,5882$ sub-rating matrix. In a movie-based scenario, we considered the 200 movies with the largest number of ratings as nodes of a graph whose edges represent Pearson similarities between any two movies. In this instance, there are 943 graph signals: The ratings each user gives to all 200 movies is one such graph signal. We are interested to predict the rating to a specific movie i with GNNs, which corresponds to completing the i th column of the rating matrix. We remark this task is permutation equivariant, therefore, we expect architectures holding this property to perform better.

Architecture Parameters. We considered the same architectural settings as in the authorship attribution experiments to highlight consistent behaviors and differences. Following [59], we chose ten 90% – 10% splits for training and test sets and pruned the graphs to keep only the top-40 most similar connections per node. The shift operator is again the adjacency matrix normalized by the maximum eigenvalue. The ADAM learning algorithm is run over 40 epochs in batches of five and learning rate 5×10^{-3} . We trained the networks on a smooth- ℓ_1 loss and measure the accuracy through the root mean squared error (RMSE).

Tables 4 and 5 show the results for the five users and five movies with the largest number of ratings, respectively. The first thing to note is that GCAT consistently improves GAT. The latter further stresses that multi-head attentions are more needed in the GAT than in the GCAT. Second, the edge varying GNN yields the worst performance because it is not a permutation equivariant architecture. In fact, the node varying and the hybrid edge varying, which are approaches in-between permutation equivariance and local detail, work much better. This trend is observed also in the edge varying GAT results, suggesting that also the number of parameters in the edge varying is too high for this task.

6.5 General Observations

Altogether these experiments lead to four main observations. **Edge Varying GNN Useful for Non-Permutation Equivariant Tasks.** The edge varying GNN [cf. (7)] can perform

well if the task is not permutation equivariant (e.g., source localization). We have observed that minimal architectures (i.e., lower number of features F , layers L and filter order K) adapt easier to the task. This is because a minimal architecture has less degrees of freedom and can avoid overfitting. Contrarily, when the task is permutation equivariant the edge varying GNN will suffer and parameterizations matched to the task are needed.

Convolution Provides a Strong Parameterization. In permutation equivariant tasks, we have seen GCNNs are a valid parameterization. They have shown potential also in tasks that are not permutation equivariant. However, contrarily to the edge varying GNN, GCNNs require a wider architecture and with filters of higher order. We have also seen the ARMANet [cf. (22)–(33)] can improve the performance of the polynomial counterpart [cf. (12)]. This is because for the rational frequency response ARMANet implements in each layer, thus requiring less parameters. However, surprisingly to us, we have observed that even with a few Jacobi iterations –ARMA-Nets require at each layer to compute a matrix inverse, which we solve iteratively with the Jacobi method– ARMANet can often achieve a better performance than the polynomial GCNN.

Hybrid Solutions Offer a Good Parameterization. While the edge varying GNN and the GCNN may outperform each other depending on the task, the node varying [cf. (15)] and the hybrid edge varying [cf. (17)] has shown a good parameterization of the EdgeNet in all tasks. We have seen the hybrid edge varying to perform overall better than the node varying due to its edge-dependent parameters. However, we did not find any consistent difference in the node selection strategies. Our rationale is that the sampling strategy of the important nodes needs to be matched with the task at hand.

Graph Convolutions Improve Attention. The graph convolutional attention network [cf. (35)] improves consistently over the GAT. This is because it generalizes the latter from a convolutional filter of order one to an arbitrary order. Working with the GCAT has shown to accelerate the learning procedure of the GAT and perform well in any task. In addition, the GCAT may result effective even with a single attention head. Instead, by generalizing the GAT to an edge varying GAT [cf. (38)] we have not seen substantial improvement than that seen in the GCAT. This indicates that the attention mechanism may not be the best strategy to learn different shift operators in each layer.

7 CONCLUSION

This paper introduced EdgeNets: GNN architectures that allow each node to collect information from its direct neighbors and apply different weights to each of them. EdgeNets preserve the state-of-the-art implementation complexity and provide a single recursion that encompasses all state-of-the-art architectures. By showcasing how each solution is a particular instance of the EdgeNet, we provided guidelines to develop more expressive GNN architectures, yet without compromising the computational complexity. This paper, in specific, proposed eight GNN architectures that can be readily extended to scenarios containing multi-edge features.

The EdgeNet link showed a tight connection between the graph convolutional and graph attention mechanism, which have been so far treated as two separate approaches. We found the graph attention network learns the weight of a graph and then performs an order one convolution over this learned graph. Following this link, we introduced the concept of graph convolutional attention networks, which is an EdgeNet that jointly learns the edge weights and the parameters of a convolutional filter.

We advocate the EdgeNet as a more formal way to build GNN solutions. However, further research is needed in three main directions. First, research should be done to explore the connection between the EdgeNets and receptive fields. This will lead to different parameterizations and architectures. Second, works needs to be done for assessing the capabilities of EdgeNets solutions to handle graph isomorphisms [60], [61]. Third, theoretical work is also needed to characterize how the stability of the EdgeNet to link perturbations [42], [62].

ACKNOWLEDGMENTS

This work was supported by NSF CCF 1717120, ARO W911NF1710438, ARL DCIST CRA W911NF-17-2-0181, ISTC-WAS and Intel DevCloud. Part of this work has been presented in [1].

REFERENCES

- [1] E. Isufi, F. Gama, and A. Ribeiro, "Generalizing graph convolutional neural networks with edge-variant recursions on graphs," in *Proc. 27th Eur. Signal Process. Conf.*, 2019, pp. 1–5.
- [2] L. Tang and H. Liu, "Relational learning via latent social dimensions," in *Proc. 15th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2009, pp. 817–826.
- [3] N. Wale, I. A. Watson, and G. Karypis, "Comparison of descriptor spaces for chemical compound retrieval and classification," *Knowl. Inf. Syst.*, vol. 14, no. 3, pp. 347–375, 2008.
- [4] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: Going beyond euclidean data," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 18–42, Jul. 2017.
- [5] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 83–98, May 2013.
- [6] F. Scarselli, S. L. Yong, M. Gori, M. Hagenbuchner, A. C. Tso, and M. Maggini, "Graph neural networks for ranking web pages," in *Proc. IEEE/WIC/ACM Int. Conf. Web Intell.*, 2005, pp. 1–7.
- [7] 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.
- [8] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, "Spectral networks and deep locally connected networks on graphs," in *Proc. 2nd Int. Conf. Learn. Representations*, 2014, pp. 1–14.
- [9] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Proc. 30th Conf. Neural Inf. Process. Syst.*, 2016, pp. 3844–3858.
- [10] 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.
- [11] J. Du, J. Shi, S. Kar, and J. M. F. Moura, "On graph convolution for graph CNNs," in *Proc. IEEE Data Sci. Workshop*, 2018, pp. 239–243.
- [12] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. 5th Int. Conf. Learn. Representations*, 2017, pp. 1–14.
- [13] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?," in *Proc. 7th Int. Conf. Learn. Representations*, 2019, pp. 1–17.
- [14] 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, no. 1, pp. 97–109, Jan. 2019.
- [15] M. Simonovsky and N. Komodakis, "Dynamic edge-conditioned filters in convolutional neural networks on graphs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 29–38.
- [16] 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 *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5425–5434.
- [17] J. Atwood and D. Towsley, "Diffusion-convolutional neural networks," in *Proc. 30th Conf. Int. Neural Inf. Process. Syst.*, 2016, pp. 2001–2009.
- [18] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *Proc. 6th Int. Conf. Learn. Representations*, 2018, pp. 1–12.
- [19] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," 2019, *arXiv:1901.00596v3*.
- [20] J. Zhou et al., "Graph neural networks: A review of methods and applications," 2019, *arXiv:1812.08434v4*.
- [21] Z. Zhang, P. Cui, and W. Zhu, "Deep learning on graphs: A survey," 2018, *arXiv:1812.04202v1*.
- [22] J. B. Lee, R. A. Rossi, S. Kim, N. K. Ahmed, and E. Koh, "Attention models in graphs: A survey," 2018, *arXiv:1807.07984v1*.
- [23] G. Taubin, "Geometric signal processing on polygonal meshes," in *Proc. 21st Annu. Conf. Eur. Assoc. Comput. Graph.*, 2000, pp. 1–11.
- [24] D. I. Shuman, P. Vandergheynst, D. Kressner, and P. Frossard, "Distributed signal processing via Chebyshev polynomial approximation," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 4, no. 4, pp. 736–751, Dec. 2018.
- [25] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, Apr. 2013.
- [26] S. K. Narang and A. Ortega, "Compact support biorthogonal wavelet filterbanks for arbitrary undirected graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 19, pp. 4673–4685, Oct. 2013.
- [27] O. Teke and P. P. Vaidyanathan, "Extending classical multirate signal processing theory to graphs—Part I: Fundamentals," *IEEE Trans. Signal Process.*, vol. 65, no. 2, pp. 409–422, Jan. 2017.
- [28] 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.
- [29] 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.
- [30] J. You, J. Leskovec, K. He, and S. Xie, "Graph structure of neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 10 881–10 891.
- [31] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proc. 3rd Int. Conf. Learn. Representations*, 2015, pp. 1–15.
- [32] A. Vaswani et al., "Attention is all you need," in *Proc. 31st Conf. Neural Inf. Process. Syst.*, 2017, pp. 1–11.
- [33] M. Coutino, E. Isufi, and G. Leus, "Advances in distributed graph filtering," *IEEE Trans. Signal Process.*, vol. 67, no. 9, pp. 2320–2333, May 2019.
- [34] P. W. Battaglia et al., "Relational inductive biases, deep learning, and graph networks," 2018, *arXiv:1806.01261*.
- [35] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1263–1272.
- [36] A. Ortega, P. Frossard, J. Kovačević, J. M. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proc. IEEE*, vol. 106, no. 5, pp. 808–828, May 2018.

- [37] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *Proc. AAAI Conf. Artif. Intell.*, 2018, pp. 7444–7452.
- [38] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 1025–1035.
- [39] V. N. Ioannidis, A. G. Marques, and G. B. Giannakis, "A recurrent graph neural network for multi-relational data," in *Proc. 44th IEEE Int. Conf. Acoust. Speech Signal Process.*, 2019, pp. 8157–8161.
- [40] H. Maron, H. Ben-Hamu, N. Shamir, and Y. Lipman, "Invariant and equivariant graph networks," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [41] R. Levie, E. Isufi, and G. Kutyniok, "On the transferability of spectral graph filters," in *Proc. IEEE 13th Int. Conf. Sampling Theory Appl.*, 2019, pp. 1–5.
- [42] F. Gama, J. Bruna, and A. Ribeiro, "Stability properties of graph neural networks," 2019, *arXiv:1905.04497v2*.
- [43] F. Gama, G. Leus, A. G. Marques, and A. Ribeiro, "Convolutional neural networks via node-varying graph filters," in *Proc. IEEE Data Sci. Workshop*, 2018, pp. 1–5.
- [44] A. Sandyhaila and J. M. F. Moura, "Discrete signal processing on graphs: Frequency analysis," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3042–3054, Jun. 2014.
- [45] S. Liu and G. Trenkler, "Hadamard, Khatri-Rao, Kronecker and other matrix products," *Int. J. Inf. Syst. Sci.*, vol. 4, no. 1, pp. 160–177, 2008.
- [46] N. A. Heckert and J. J. Filliben, "Exploratory data analysis," in *NIST/SEMATECH E-Handbook of Statistical Methods*, Gaithersburg, MD, USA: NIST, 2003, ch. 1.
- [47] L. N. Trefethen, *Approximation Theory and Approximation Practice, Extended ed.* Philadelphia, PA, USA: SIAM, 2019.
- [48] J. Liu, E. Isufi, and G. Leus, "Filter design for autoregressive moving average graph filters," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 5, no. 1, pp. 47–60, Mar. 2019.
- [49] E. Isufi, A. Loukas, and G. Leus, "Autoregressive moving average graph filtering: A stable distributed implementation," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2017, pp. 4119–4123.
- [50] F. M. Bianchi, D. Grattarola, C. Alippi, and L. Livi, "Graph neural networks with convolutional ARMA filters," 2019, *arXiv:1901.01343v5*.
- [51] W. A. S. Wijesinghe and Q. Wang, "DFNets: Spectral CNNs for graphs with feedback-looped filters," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 6007–6018.
- [52] L. Ruiz, F. Gama, and A. Ribeiro, "Gated graph recurrent neural networks," *IEEE Trans. Signal Process.*, vol. 68, pp. 6303–6318, 2020.
- [53] K. Gregor and Y. LeCun, "Learning fast approximations of sparse coding," in *Proc. 27th Int. Conf. Mach. Learn.*, 2010, pp. 399–406.
- [54] D. P. Kingma and J. L. Ba, "ADAM: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Representations*, 2015, pp. 1–15.
- [55] 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, Jul. 2016.
- [56] J. McAuley and J. Leskovec, "Learning to discover social circles in ego networks," in *Proc. 26th Conf. Neural Inf. Process. Syst.*, 2012, pp. 539–547.
- [57] 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.
- [58] F. M. Harper and J. A. Konstan, "The MovieLens datasets: History and context," *ACM Trans. Interactive Intell. Syst.*, vol. 5, no. 4, pp. 1–19, 2015.
- [59] L. Ruiz, F. Gama, A. G. Marques, and A. Ribeiro, "Invariance-preserving localized activation functions for graph neural networks," *IEEE Trans. Signal Process.*, vol. 68, no. 68, pp. 127–141, Nov. 25, 2019.
- [60] H. Maron, H. Ben-Hamu, H. Serviansky, and Y. Lipman, "Provably powerful graph networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 2156–2167.
- [61] C. Morris et al., "Weisfeiler and Leman go neural: Higher-order graph neural networks," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 4602–4609.
- [62] R. Levie, W. Huang, L. Bucci, M. M. Bronstein, and G. Kutyniok, "Transferability of spectral graph convolutional neural networks," 2019, *arXiv:1907.12972*.
- [63] A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson, "The diffusion of microfinance," *Science*, vol. 341, no. 6144, Jul. 2013, Art no. 1 236498.



Elvin Isufi received the MSc (cum laude) degree from the University of Perugia, Perugia, Italy, in 2014, and the PhD degree from the Delft University of Technology, Delft, The Netherlands, in 2019. He is currently an assistant professor with the Multimedia Computing Group, Delft University of Technology. Prior to that, he was a postdoctoral researcher with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, Pennsylvania. His research interests include intersection of signal processing, mathematical modeling, machine learning, and network theory. He has received paper awards at DSLW2021 and CAMSAP 2017.



Fernando Gama received the electronic engineer degree from the University of Buenos Aires, Argentina, in 2013, the MA degree in statistics from the Wharton School, University of Pennsylvania, in 2017, and the PhD degree in electrical and systems engineering from the University of Pennsylvania, Philadelphia, Pennsylvania, in 2020. He is currently a postdoctoral research fellow with the Electrical and Computer Engineering Department, Rice University. He has been a visiting researcher with TU Delft, The Netherlands, in 2017, a research intern with Facebook Artificial Intelligence Research, Montreal, Canada, in 2018, and a postdoctoral scholar with UC Berkeley in 2020–2021. His research interests include the field of information processing and machine learning over network data. He has been awarded a Fulbright scholarship for international students and he has received a Best Student Paper Award at EUSIPCO 2019.



Alejandro Ribeiro received the BSc degree in electrical engineering from the Universidad de la Republica Oriental del Uruguay, Montevideo, Uruguay, in 1998, and the MSc and PhD degrees in electrical engineering from the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, Minnesota, in 2005 and 2007, respectively. He joined the University of Pennsylvania (Penn), Philadelphia, in 2008 where he is currently professor of electrical and systems engineering. His research interests include the applications of statistical signal processing to the study of networks and networked phenomena. His focus is on structured representations of networked data structures, graph signal processing, network optimization, robot teams, and networked control. He received an Outstanding Research Award from Intel in 2019, the 2014 O. Hugo Schuck Best Paper Award, and paper awards at ICASSP 2020, EUSIPCO 2019, CDC 2017, 2016 SSP Workshop, 2016 SAM Workshop, 2015 Asilomar SSC Conference, ACC 2013, ICASSP 2006, and ICASSP 2005. His teaching has been recognized with the 2017 Lindback Award for distinguished teaching and the 2012 S. Reid Warren, Jr. Award presented by Penn's undergraduate student body for outstanding teaching. He is a Fulbright scholar class of 2003 and a Penn Fellow class of 2015.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.**