

# MEDIAN ACTIVATION FUNCTIONS FOR GRAPH NEURAL NETWORKS

Luana Ruiz<sup>†</sup>, Fernando Gama<sup>†</sup>, Antonio G. Marques<sup>\*</sup> and Alejandro Ribeiro<sup>†</sup>

<sup>†</sup> Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, USA

<sup>\*</sup> Department of Signal Theory and Communications, King Juan Carlos University, Madrid, Spain

## ABSTRACT

Graph neural networks (GNNs) have been shown to replicate convolutional neural networks' (CNNs) superior performance in many problems involving graphs. By replacing regular convolutions with linear shift-invariant graph filters (LSI-GFs), GNNs take into account the (irregular) structure of the graph and provide meaningful representations of network data. However, LSI-GFs fail to encode local *nonlinear* graph signal behavior, and so do regular activation functions, which are nonlinear but pointwise. To address this issue, we propose median activation functions with support on graph neighborhoods instead of individual nodes. A GNN architecture with a trainable multiresolution version of this activation function is then tested on synthetic and real-world datasets, where we show that median activation functions can improve GNN capacity with marginal increase in complexity.

**Index Terms**— graph neural networks, big data, activation functions, graph signal processing, median filters

## 1. INTRODUCTION

Modern problems of social and economic interest such as identifying the origin of epidemics on transportation networks and deciding which products to market to users of a social network have motivated the development of an extensive framework to model and analyze big data on networks [1]. In particular, many of these tasks require information processing architectures that are adapted to the inherently irregular structure of graphs. Graph convolutional neural networks (GNNs) are one such example. GNNs are the extension of convolutional neural networks (CNNs) to the graph domain and, like them, have demonstrated remarkable performance in multiple classification and regression tasks [2–6]. They are good at generating graph data representations because their convolutional layers are based off of linear-shift invariant graph filters (LSI-GFs) that encode the irregular structure of the underlying graph into the filtered signals [7].

Naturally, GNNs give good approximations of more than just linear functions, and the reason for that is the fact that each GNN layer composes the LSI-GFs with a nonlinear operator. In NNs, this operator is called the activation function, and common choices for it are the sigmoid, the tanh and the rectified linear unit (ReLU) [8]. Because these are all pointwise operations, in the

GNN they can be computed at each node individually, without looking at the rest of the graph. While separating the local mixing from the non-linearities may render the analysis and implementation of the NN easier, architectures that allow for non-scalar nonlinear activation functions (i.e., activation functions whose input is a vector collecting the features in the neighborhood of a given node) are more general, either because they are better proxies of the actual underlying data model or because they achieve the same representation capabilities with less parameters.

Building upon this observation, we introduce *median activation functions* as activation functions whose output at each node depends on the values of the input at the neighboring nodes. These functions will provide smooth and analytically tractable nonlinear transformations which, in some cases, can be related to median graph filters [9, 10]. The main focus of this paper is on the case where the activation function is a weighted sum of medians taken at different graph resolutions, whose linear weights can be trained together with the other linear parameters of the GNN. While parametrized activation function training [11, 12] and activation functions that adapt from graph training data [13] have been explored in the existing literature, they are all pointwise and, as such, unable to learn local graph topologies in the way that trainable median activations do. This ability is demonstrated in experiments performed on both synthetic and real-world datasets, where GNN architectures with median and ReLU activations are compared. Our results show that median activation functions can further improve GNN performance by improving capacity while only marginally increasing the number of trainable parameters.

## 2. GRAPH NEURAL NETWORKS

Consider a training set  $\mathcal{T} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{T}|}$  with  $|\mathcal{T}|$  samples where  $\mathbf{x} \in \mathcal{X}$  are the input data and  $\mathbf{y} \in \mathcal{Y}$  are known target representations that we intend to learn from. By learning we mean that we want to obtain suitable representations  $\hat{\mathbf{y}} \in \mathcal{Y}$  of unseen inputs  $\mathbf{x}' \notin \mathcal{T}$ ; to accomplish this, we train a model  $f : \mathcal{X} \rightarrow \mathcal{Y}$  such that  $\hat{\mathbf{y}} = f(\mathbf{x})$  and  $f$  minimizes the value of a loss function  $\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y}_i)$  on  $\mathcal{T}$ . We make the assumption that the unseen inputs  $\mathbf{x}' \notin \mathcal{T}$  are drawn from the same or from a similar distribution as the  $\mathbf{x} \in \mathcal{T}$ , in which case it is reasonable to expect  $f$  to output accurate representations of  $\mathbf{x}'$ .

NNs are the architecture of our choice in coming up with  $f$ . They are built from the concatenation of a series of computational layers which, in turn, perform two basic operations: a linear transform and a pointwise nonlinearity that is also referred to

Supported by USA NSF CCF 1717120 and ARO W911NF1710438, and Spanish MINECO TEC2016-75361-R.

as activation function. Layer  $\ell$  will output  $\mathbf{x}_\ell$  given by

$$\mathbf{x}_\ell = \sigma_\ell(T_{\mathbf{A}_\ell}(\mathbf{x}_{\ell-1})) = \sigma_\ell(\mathbf{A}_\ell \mathbf{x}_{\ell-1}), \ell = 1, \dots, L \quad (1)$$

where  $\mathbf{x}_\ell \in \mathcal{X}_\ell$ ,  $T_{\mathbf{A}_\ell} : \mathcal{X}_{\ell-1} \rightarrow \mathcal{X}_\ell$  is the linear transform and  $\sigma_\ell : \mathcal{X}_\ell \rightarrow \mathcal{X}_\ell$  is the pointwise nonlinearity. We take the input  $\mathbf{x}_0$  and the output  $\mathbf{x}_L$  to be  $\mathbf{x}$  and  $\hat{\mathbf{y}}$  respectively.

In this paper, we are interested in the case where  $\mathbf{x}$  is a graph signal. A graph problem of this sort would be, for instance, the task of predicting a power outage in a given city based on information collected at the region's meteorological stations, modeled as the vertices of a geodesic graph [15]. Graph signals are represented as vectors  $\mathbf{x} \in \mathbb{R}^N$ ; the individual components  $[\mathbf{x}]_i$  are the values of the signal at each node  $i$  of the graph, which we denote as  $\mathcal{G}$ . For our purposes,  $\mathcal{G}$  is fully determined by the triplet  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ , where  $\mathcal{V} = \{1, \dots, N\}$  is the set of nodes,  $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$  is the set of edges and  $\mathcal{W} : \mathcal{E} \rightarrow \mathbb{R}$  is the weight function of the graph,  $\mathcal{W}(i, j) = w_{ij}$ . Additionally,  $\mathcal{G}$  can be represented in matrix form: the graph adjacency matrix  $\mathbf{W} \in \mathbb{R}^{N \times N}$ , where  $[\mathbf{W}]_{ij} := w_{ji}$ , and the unnormalized graph Laplacian  $\mathbf{L} := \mathbf{D} - \mathbf{W}$ , with  $\mathbf{D} = \text{diag}(\mathbf{W}\mathbf{1})$  the degree matrix of  $\mathcal{G}$ , are the two most common representations.

Both  $\mathbf{W}$  and  $\mathbf{L}$  are examples of the so-called *graph shift operator* (GSO) of the graph. A GSO  $\mathbf{S} \in \mathbb{R}^{N \times N}$  of  $\mathcal{G}$  is defined as any matrix satisfying  $[\mathbf{S}]_{ij} \neq 0$  if and only if  $(i, j) \in \mathcal{E}$  or  $i = j$ . In this paper, the GSO  $\mathbf{M} \in \{0, 1\}^{N \times N}$  such that  $[\mathbf{M}]_{ij} = 1$  for every  $(i, j) \in \mathcal{E}$  is a particularly useful one. Its powers  $\mathbf{M}^r$ ,  $r \geq 0$ , allow us to define the  $r$ -hop neighborhood of a node  $i$  as  $\mathcal{N}_i^r$ , that is, the set of nodes that are exactly at  $r$  hops from  $i$ .

Since we are interested in graph signals, the generic architecture described by (1) has to be adapted to process this type of data. In GNNs, this is done by imposing that  $T_{\mathbf{A}}(\mathbf{x})$  be an LSI-GF, where  $\mathbf{A}$  in (1) takes the form of  $\mathbf{H} : \mathbb{R}^N \rightarrow \mathbb{R}^N$  such that

$$\mathbf{H} := \sum_{k=0}^{K-1} h_k \mathbf{S}^k, \quad (2)$$

with  $h_0, \dots, h_{K-1}$  the filter coefficients and  $K$  the number of filter taps.

Another differentiating characteristic of GNNs is the number of filters that they have per layer. Like CNNs designed for image processing applications, GNNs also rely on different filters to extract distinct local features at each layer. Each of these filters outputs a linear feature that could either be combined with others through a summarizing operation or be considered individually, at the price of increasing the subsequent layer's dimensionality.

Extending the activation functions  $\sigma$  to GNNs is immediate because these operations are inherently pointwise. Applying  $\sigma$  to an arbitrary graph signal  $\mathbf{y}$  yields  $[\sigma(\mathbf{y})]_i = \sigma([\mathbf{y}]_i)$ ; this means that the outcome of the activation function at node  $i$  depends solely on the value of the input signal at this node. In this paper, we argue that, since the overall goal is to merge information across nodes in a nonlinear fashion, considering only scalar activation functions can be too restrictive. This idea — that it takes nonlinear operators to capture nonlinear relationships in data — is not new and has for instance motivated the authors of [16] to come up with a dynamic-nonlinear method for online identifica-

tion of graph topologies that does better than traditional linear structural vector autoregressive models.

In this context, we put forward *median activation functions* as local operators that provide a nonlinear yet relatively smooth measure of central tendency in variously sized regions of a graph around each node. The main reasons for choosing this family of nonlinearities are the fact that they preserve tractability (both computationally and analytically) and are effective in many inference and signal processing tasks. With median activations, we expect to provide an alternative, more efficient way of increasing GNNs' overall capacity [8] to wider or deeper layered architectures, that risk overfitting when parameters are added in excess.

### 3. MEDIAN ACTIVATION FUNCTIONS

The median activation functions introduced in this paper differ from conventional activation functions in two ways. The first and most important difference is in their support, which is now a graph neighborhood instead of a single node. The second is the choice of nonlinear operation to perform — the median graph filter. Indeed, the activation function discussed in the next paragraph is simply a  $r$ -hop median graph filter where, for every node  $i$ , the filter window is defined as the node's extended  $r$ -hop neighborhood  $\bar{\mathcal{N}}_i^r$ ,

$$\bar{\mathcal{N}}_i^r := \bigcup_{k=0}^r \mathcal{N}_i^k. \quad (3)$$

To make the distinction between the two activation functions introduced in this paper, we will refer to this first one as the *static median activation*. We introduce it as a first step to later define a more interesting operator, namely the *dynamic median activation*. The static median activation function is defined as follows.

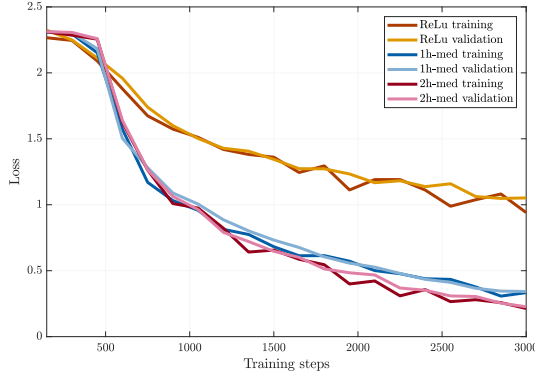
**Definition 1.** Consider a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ , a graph signal  $\mathbf{x} \in \mathbb{R}^N$  on  $\mathcal{G}$  and the number of hops  $r \in \mathbb{N}$ . We define the  $r$ -hop median graph filter as  $r\text{-med}_{\mathcal{G}}(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}^N$  such that

$$[r\text{-med}_{\mathcal{G}}(\mathbf{x})]_i := \text{med}\{x_{j_1}, x_{j_2}, \dots, x_{j_{N_i^r}}\} \text{ with } N_i^r = |\bar{\mathcal{N}}_i^r| \text{ and } j_1, j_2, \dots, j_{N_i^r} \in \bar{\mathcal{N}}_i^r. \quad (4)$$

For  $r > 0$  and  $\mathcal{G}$  without unitary components (isolated nodes),  $r\text{-med}_{\mathcal{G}}(\mathbf{x})$  defines the **static median activation function** for GNN architectures.

What this definition conveys is that the static median activation function takes a graph signal as input, fetches each node's extended  $r$ -hop neighborhood, and computes  $N$  medians of the signal in these neighborhoods to output another signal on the same graph. If  $N_i^r$  in equation (4) is even, we stipulate  $[r\text{-med}_{\mathcal{G}}(\mathbf{x})]_i = x_{[N_i^r/2+1]}$ , that is, the median is given by the  $(N_i^r/2 + 1)$ th order statistics of  $x_{j_1}, \dots, x_{j_{N_i^r}}$ . We point out that, as long as  $r > 0$ , the operations in (4) yield nonlinear activation functions.

While the graph filter from Definition 1 has the advantage of enlarging the activation function support beyond a single point or node, the  $r$ -hop static median activation is only able to capture nonlinear relationships in the  $r$ -hop neighborhood of each node. However, the behavior of the graph signal at a given node might



**Fig. 1:** Training and validation losses in one round of simulation of the source localization problem on the Twitter graph, for the ReLu, 1-hop dynamic median and 2-hop dynamic median architectures .

be more complex than what this function is able to seize; for instance, it could be as related to the behavior of the signal on closer nodes (within the  $r$ -hop neighborhood) as to the signal dynamics on more distant but highly influential nodes (e.g. nodes with high centrality). The dynamic median activation function presented in Definition 2 addresses this issue by taking into account the medians in node neighborhoods of increasing size (up to  $R$  hops, where  $R$  is the median filter’s reach) and weighing them linearly. We define it as follows.

**Definition 2.** Consider a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ , a graph signal  $\mathbf{x} \in \mathbb{R}^N$  defined on  $\mathcal{G}$ , the filter reach  $R \in \mathbb{N}$  and the linear weight vector  $\omega \in \mathbb{R}^{R+1}$ . We define the  $R$ -hop weighted graph filter  $R\text{-med}_{\mathcal{G}}^{\omega}(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}^N$  as

$$\begin{aligned} [R\text{-med}_{\mathcal{G}}^{\omega}(\mathbf{x})]_i &:= \omega^{\top} \mathbf{z}_i, \\ \mathbf{z}_i &\in \mathbb{R}^{R+1} \text{ and } [\mathbf{z}_i]_r = [(r)\text{-med}_{\mathcal{G}}(\mathbf{x})]_i \end{aligned} \quad (5)$$

with  $r = 0, \dots, R$  and  $i = 1, \dots, N$ . For  $R > 0$  and  $\mathcal{G}$  without unitary components (isolated nodes),  $R\text{-med}_{\mathcal{G}}^{\omega}(\mathbf{x})$  defines the **dynamic median activation function** for GNN architectures.

The introduction of linear weight variables in the above definition is desirable for two reasons. The first one is that they allow assigning different levels of importance to the median in each neighborhood. The second is the fact that, because these variables are linear, their gradients are easy to compute and so they can be updated via backpropagation. In other words, this activation function can be trained, which is why we call it *dynamic*. It is worth pointing out that even though the dynamic median activation function is linear in the weights  $\omega$ , it is still nonlinear in the input graph signal  $\mathbf{x}$  as long as  $R > 0$ .

## 4. NUMERICAL EXPERIMENTS

In this section, we present simulation results obtained with GNN architectures using *dynamic median activation functions* in two applications. In the first, we try to predict the source of a synthetic diffusion process on social networks; in the second, we

attribute authorship of text excerpts to 19th century writers. Our architecture’s performance is compared with that of a GNN that shares the same design parameters in the convolutional layer, except that the activation function is a (pointwise) ReLu.

The simulated GNN architectures all consist of a single convolutional layer followed by a fully connected softmax layer, which we use to compute the cross-entropy loss. The convolutional layer contains 32 filters with 5 filter taps each; each filter uses the adjacency matrix normalized by its largest eigenvalue as the GSO and outputs a single feature. The GNNs were trained without dropout using the ADAM algorithm, a variation of stochastic gradient descent that keeps an exponentially decaying average of past gradients, with decaying factors  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$  [17].

### 4.1. Source Localization on Facebook and Twitter

In this experiment, we consider a graph  $\mathcal{G}$  with  $N$  nodes and generate graph diffusion processes  $\mathbf{x}(t)$  [18] with origin at an arbitrary node  $c$ . Suppose that  $\mathcal{G}$  has adjacency matrix  $\mathbf{W}$  and that  $\mathbf{x}(0) = \mathbf{x}_0$ , where  $\mathbf{x}_0$  is such that  $[\mathbf{x}_0]_i = 1$  if and only if  $i = c$  and 0 everywhere else. The diffused signals  $\mathbf{x}(t)$ ,  $t = 1, 2, \dots$ , are generated as  $\mathbf{x}(t) = \mathbf{W}^t \mathbf{x}_0$ . By picking random values of  $t$  and placing the corresponding  $\mathbf{x}(t)$  in tuples  $(\mathbf{x}(t), c)$ , we can train a GNN that when presented with a new, unobserved process  $\mathbf{x}'(t)$ , is able to accurately predict its source  $c'$ .

This experiment is performed in two graphs: a 234-node undirected graph formed by a single connected component of the Facebook friend network; and a directed ego network [19] from Twitter with 237 nodes and a single weakly connected component. These networks were built under the Stanford Network Analysis Project and can be found in [20]. Although the diffusion processes used to train the GNNs in this paper are synthetic, they are a good representation of how rumors spread on social networks, for instance [21].

In both the Facebook and Twitter experiments, 3 GNNs were trained on the same randomly generated diffusion datasets at a time: one with ReLu activations, one with 1-hop dynamic median activations and one with 2-hop dynamic median activations. In total, ten rounds of simulations corresponding to ten different datasets were conducted for each problem. We used 10,000 training and 200 test samples in every round, and training was done in batches of 100 samples with learning rate 0.001. In the Facebook problem, we restricted the number of classes to 5, corresponding to the 5 nodes with highest degree. For Twitter, we allowed for the 10 nodes with highest degree to originate the diffusion process. The Facebook GNNs were trained in 40 epochs, while the Twitter ones only needed 30 training epochs in total.

Figure 1 shows the training and validation losses for every analyzed architecture in one of the 10 rounds of simulations on the Twitter graph. In this realization, it is clear that architectures with median activation functions incur a dramatic increase in capacity and sensibly reduce the generalization error while only adding up to 3 extra parameters to the GNN (cf. 1). In particular, the best results are obtained with the median activation function with largest reach ( $R = 2$ ). Average test accuracy results for both the Facebook and Twitter problems are reported in Table 1. While in

Archit.	Average accuracy (%)		Parameters
	Facebook	Twitter	
ReLU	70.30 $\pm$ 9.08	62.10 $\pm$ 11.91	160
1h-med	<b>79.25 <math>\pm</math> 1.72</b>	97.60 $\pm$ 4.44	162
2h-med	79.80 $\pm$ 2.39	<b>99.2 <math>\pm</math> 2.24</b>	163

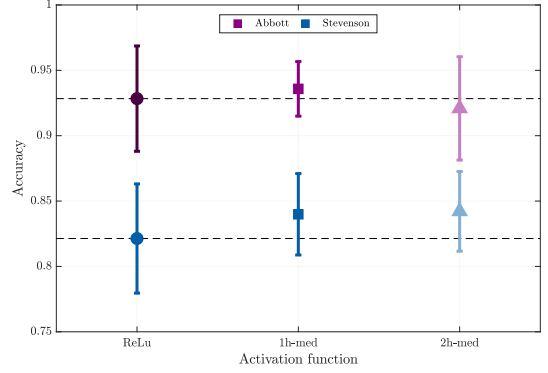
**Table 1:** Test accuracy for ReLu, 1-hop median and 2-hop median architectures in the source localization problem. Average and standard deviation for 200 test samples over 10 rounds of simulations on both the Facebook and Twitter graphs. The number of parameters in the convolutional layer of each architecture is reported in the rightmost column.

the Facebook graph the 1-hop and 2-hop median achieve a similar accuracy (with the 1-hop median actually performing marginally better due to the smaller standard deviation), in the Twitter graph the benefits of increasing the median reach  $R$  are significant. This could be explained not only by the fact that one graph is undirected and the other is directed, but also because while the Facebook graph has average degree 34.25, the Twitter graph has average in and out-degrees of 11.45. As such, there is much more redundancy between 1-hop and 2-hop extended neighborhoods in the Facebook graph than in the Twitter graph.

#### 4.2. Authorship Attribution of 19th Century Texts

The second application in which we assess dynamic median activation function performance is a binary classification problem based on real data. In this experiment, the graphs are author word adjacency networks (WANs) where each function word (e.g. prepositions, articles, conjunctions etc.) is a node and function words are connected through edges weighted by frequency of co-appearance in texts that are known to have been written by the author in question [22]. Because order of appearance matters, these graphs are directed. We consider WANs by two authors: Jacob Abbott and Robert Stevenson. On these graphs, graph signals are defined as the vectored frequencies of each function word in 1,000-word text excerpts. The text excerpts we consider come from a pool of twenty-one 19th century authors, including Abbott and Stevenson; then, for the GNN trained on Abbot’s WAN, for example, the objective is to classify each excerpt as having being written by Abbott or not.

Here again we trained 3 GNNs for Abbott and 3 for Stevenson, one per activation function (ReLU, 1h-med, 2h-med). The number of parameters in the convolutional layer of each GNN were, respectively, 160, 162 and 163. For each author, we conducted 10 rounds of simulations varying both the WAN and the training and test set 80-20 splits. We consider 180 function words (nodes) for Abbott and 193 for Stevenson, and the resulting WANs are connected graphs. Naturally, only the author’s texts in the training set were used to build the WAN at each simulation round. For Abbott, we used 1,080 training and 268 test samples, and, for Stevenson, 1,436 and 356 respectively. All of these datasets are evenly balanced, meaning that half of the samples correspond to texts by the author in question (labeled +1, positive) and the other half are randomly picked from the



**Fig. 2:** Test accuracy for ReLu, 1-hop median and 2-hop median architectures in the authorship attribution problem for Abbott and Stevenson. Average accuracy and standard deviation over 10 rounds of simulations.

21-author pool (labeled -1, negative). Training was done in 10 epochs in batches of 20.

Average classification accuracies for both Abbott and Stevenson are shown in Figure 2. For Stevenson, the two median activation functions incur a significant improvement in average accuracy and consistency with respect to the ReLu, with around 7 additional texts being correctly classified on average. Another thing to note for this author is that performance increases with the activation function reach. For Abbott, the average accuracy improvement is more timid and comes from the 1-hop median, while the 2-hop median degrades performance w.r.t. the ReLu. Here, this is probably less related to growing neighborhood redundancies than it is to stylistic differences such as a particular author’s average sentence length, for instance.

## 5. CONCLUSIONS

This work introduced static and dynamic median activation functions as an alternative to regular pointwise activation functions in GNNs. Their major advantage is extending activation function support to graph neighborhoods that go beyond individual nodes. This allows encoding any local nonlinear relationships existent in the graph and/or the graph signal, leading to increased GNN capacity. In particular, the dynamic median activation function takes the graph signal median in graph neighborhoods of increasing size and weighs them linearly. These weights can then be trained via backpropagation and thus yield a class of trainable activation functions. Numerical results in synthetic datasets showed that dynamic median activation functions are a better alternative to increasing GNN capacity than deepening or widening GNNs; they achieved low generalization error while keeping almost the same number of parameters as the GNN with all ReLu activations, thus decreasing the risk of overfitting. Additionally, we observed sensible classification accuracy improvements in both synthetic and real-world datasets.

## 6. REFERENCES

- [1] M. O. Jackson, *Social and Economic Networks*, Princeton University Press, Princeton, NJ, 2008.
- [2] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, “Convolutional neural network architectures for signals supported on graphs,” *arXiv:1805.00165v1 [eess.SP]*, 1 May 2018.
- [3] M. Defferrard, X. Bresson, and P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *Annu. Conf. Neural Inform. Process. Syst. 2016*, Barcelona, Spain, 5-10 Dec. 2016, NIPS Foundation.
- [4] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *5th Int. Conf. Learning Representations*, Toulon, France, 24-26 Apr. 2017, Assoc. Comput. Linguistics.
- [5] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, “Spectral networks and deep locally connected networks on graphs,” *arXiv:1312.6203v3 [cs.LG]*, 21 May 2014.
- [6] R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec, “Graph convolutional neural networks for web-scale recommender systems,” in *24th ACM SIGKDD Int. Conf. on Knowledge Discovery & Data Mining*, London, UK, 19-23 Aug. 2018, Assoc. for Computing Machinery.
- [7] 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.
- [8] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, The Adaptive Computation and Machine Learning Series. The MIT Press, Cambridge, MA, 2016.
- [9] S. Segarra, A. G. Marques, G. R. Arce, and A. Ribeiro, “Center-weighted median graph filters,” in *2016 IEEE Global Conf. Signal and Inform. Process.*, Washington, DC, 7-9 Dec. 2016, IEEE.
- [10] S. Segarra, A. G. Marques, G. R. Arce, and A. Ribeiro, “Design of weighted median graph filters,” in *2017 IEEE Int. Workshop Comput. Advances Multi-Sensor Adaptive Process.*, Curacao, Netherlands Antilles, 10-13 Dec. 2017, IEEE.
- [11] P. Chandra and Y. Singh, “An activation function adapting training algorithm for sigmoidal feedforward networks,” *Neurocomputing*, vol. 61, pp. 429–437, Oct. 2004.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *2015*, Santiago, Chile, 7-13 Dec. 2015, IEEE.
- [13] S. Scardapane, S. Van Vaerenbergh, D. Comminiello, and A. Uncini, “Improving graph convolutional networks with non-parametric activation functions,” *arXiv:1802.09405 [cs.NE]*, 26 Feb. 2018.
- [14] D. Owerko, F. Gama, and A. Ribeiro, “Predicting power outages using graph neural networks,” in *IEEE Global Conf. Signal and Inform. Process. 2018*, Anaheim, CA, 26-29 Nov. 2018, IEEE.
- [15] Y. Shen and G. B. Giannakis, “Online identification of directional graph topologies capturing dynamic and nonlinear dependencies,” in *2018 IEEE Data Sci. Workshop*, Lausanne, Switzerland, 4-6 June 2018, IEEE.
- [16] D. P. Kingma and J. L. Ba, “ADAM: A method for stochastic optimization,” in *3rd Int. Conf. Learning Representations*, San Diego, CA, 7-9 May 2015, Assoc. Comput. Linguistics.
- [17] S. Segarra, G. Mateos, A. G. Marques, and A. Ribeiro, “Blind identification of graph filters,” *arXiv:1604.07234 [cs.IT]*, 25 Apr. 2016.
- [18] J. Leskovec and R. Sosič, “Snap: A general-purpose network analysis and graph-mining library,” *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 8, no. 1, pp. 1, 2016.
- [19] J. Leskovec and A. Krevl, “SNAP Datasets: Stanford large network dataset collection,” <http://snap.stanford.edu/data>, June 2014.
- [20] A. Ortega, P. Frossard, J. Kovačević, J. M. F. Moura, and P. Vandergheynst, “Graph signal processing: Overview, challenges and applications,” *Proceedings of the IEEE*, vol. 106, no. 5, pp. 808–828, 2018.
- [21] 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, June 2015.