

On the Transferability of Spectral Graph Filters

Levie, Ron; Isufi, Elvin; Kutyniok, Gitta

DOI

[10.1109/SampTA45681.2019.9030932](https://doi.org/10.1109/SampTA45681.2019.9030932)

Publication date

2019

Document Version

Final published version

Published in

2019 13th International Conference on Sampling Theory and Applications, SampTA 2019

Citation (APA)

Levie, R., Isufi, E., & Kutyniok, G. (2019). On the Transferability of Spectral Graph Filters. In A. Amini (Ed.), *2019 13th International Conference on Sampling Theory and Applications, SampTA 2019: Bordeaux, France* [9030932] (2019 13th International Conference on Sampling Theory and Applications, SampTA 2019). SampTA. <https://doi.org/10.1109/SampTA45681.2019.9030932>

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.

On the Transferability of Spectral Graph Filters

Ron Levie*, Elvin Isufi[†], and Gitta Kutyniok*

*Technische Universität Berlin, Institut für Mathematik, Berlin, Germany

[†]Delft University of Technology, Department of Microelectronics, Delft, Netherlands.

Emails: levie@math.tu-berlin.de, E.Isufi-1@tudelft.nl, kutyniok@math.tu-berlin.de

Abstract—This paper focuses on spectral filters on graphs, namely filters defined as elementwise multiplication in the frequency domain of a graph. In many graph signal processing settings, it is important to transfer a filter from one graph to another. One example is in graph convolutional neural networks (ConvNets), where the dataset consists of signals defined on many different graphs, and the learned filters should generalize to signals on new graphs, not present in the training set. A necessary condition for transferability (the ability to transfer filters) is stability. Namely, given a graph filter, if we add a small perturbation to the graph, then the filter on the perturbed graph is a small perturbation of the original filter. It is a common misconception that spectral filters are not stable, and this paper aims at debunking this mistake. We introduce a space of filters, called the Cayley smoothness space, that contains the filters of state-of-the-art spectral filtering methods, and whose filters can approximate any generic spectral filter. For filters in this space, the perturbation in the filter is bounded by a constant times the perturbation in the graph, and filters in the Cayley smoothness space are thus termed linearly stable. By combining stability with the known property of equivariance, we prove that graph spectral filters are transferable.

I. INTRODUCTION

The success of convolutional neural networks (ConvNets) on Euclidean domains ignited an interest in recent years in extending these methods to graph structured data. A graph ConvNet is a mapping, that receives a signal defined over the vertices of a graph, and returns a value in some output space. A graph ConvNet consists of many layers of computations, where each layer computes a set of filters of the output of the previous layer, followed by a pointwise nonlinearity, and optionally a pooling step and non-convolution layers. In a machine learning setting, the general architecture of the ConvNet is fixed, but the specific filters to use in each layer are free parameters. In training, the filter coefficients are optimized to minimize some loss function. In some situations, both the graph and the signal defined on the graph are variables in the input space of the ConvNet. In these situations, if two graphs represent the same underlying phenomenon, and the two signals given on the two graphs are similar in some sense, the output of the ConvNet on both signals should be similar as well. This property is typically termed transferability, and is an essential requirement if we wish the ConvNet to generalize well on the test set. Analyzing and proving transferability is the subject of this paper.

A necessary condition of any reasonable definition of transferability is stability. Namely, given a filter, if the topology of a graph is perturbed, then the filter on the perturbed

graph is close to the filter on the un-perturbed graph. Without stability it is not even possible to transfer a filter from a graph to another very close graph, and thus stability is necessary for transferability. Previous work studied the behavior of graph filters with respect variations in the graph. [1] provided numerical results on the robustness of polynomial graph filters to additive Gaussian perturbations of the eigenvectors of the graph Laplacian. Since the eigendecomposition is not stable to perturbations in the topology of the graph, this result does not prove robustness to such perturbations. [2] showed that the expected graph filter under random edge losses is equal to the accurate output. However, [2] did not bound the error in the output in terms of the error in the graph topology. In this paper we show the linear stability of graph filters to general perturbations in the topology.

There are generally two approaches to defining convolution on graphs, both generalizing the standard convolution on Euclidean domains [3], [4]. Spatial approaches generalize the idea of a sliding window to graphs. Here, the main challenge is to define a way to translate a filter kernel along the vertices of the graph. Some popular examples of spatial methods are [5], [6], [7]. Spectral methods are inspired by the convolution theorem in Euclidean domains, that states that convolution in the spatial domain is equivalent to pointwise multiplication in the frequency domain. The challenge here is to define the frequency domain and the Fourier transform of graphs. The basic idea is to define the graph Laplacian, or some other graph operator that we interpreted as a shift operator, and to use its eigenvalues as frequencies and its eigenvectors as the corresponding pure harmonics [8]. Decomposing a graph signal to its pure harmonic coefficients is by definition the graph Fourier transform, and filters are defined by multiplying the different frequency components by different values. For some examples of spectral methods see, e.g., [9], [10], [11], [12]. Additional references for both methods can be found in [4].

The great majority of researchers from the graph ConvNet community currently focus on developing spatial methods. One typical motivation for favoring spatial methods is the claim that spectral methods are not transferable, and thus do not generalize well on graphs unseen in the training set. The goal in this paper is to debunk this misconception, and to show that state-of-the-art spectral graph filtering methods are transferable. This paper does not argue against spatial methods, but shows the potential of spectral approaches to cope with datasets having varying graphs. We would like to

encourage researches to reconsider spectral methods in such situations. Interestingly, [13] obtained state-of-the-art results using spectral graph filters on variable graphs, without any modification to compensate for the “non-transferability”.

II. PRELIMINARIES

A. Spectral graph filters

Consider an undirected weighted graph $\mathcal{G} = \{E, V, \mathbf{W}\}$, with vertices $V = \{1, \dots, N\}$, edges E , and adjacency matrix \mathbf{W} . The adjacency matrix \mathbf{W} is symmetric and represents the weights of the edges, where $w_{n,m}$ is nonzero only if vertex n is connected to vertex m . Consider the degree matrix \mathbf{D} , defined as the diagonal matrix with entries $d_{n,n} = \sum_{m=1}^N w_{n,m}$.

The frequency domain of a graph is determined by choosing a shift operator, namely a self-adjoint operator Δ that respects the connectivity of the graph. As a prototypical example, we consider the unnormalized Laplacian $\Delta = \mathbf{D} - \mathbf{W}$, which depends (affine) linearly on \mathbf{W} . Other examples of common shift operators are the normalized Laplacian $\Delta_n = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2}$, and the adjacency matrix itself. In this paper we call a generic self-adjoint shift operator *Laplacian*, and denote it by Δ . Denote the eigenvalues of Δ by $\{\lambda_n\}_{n=1}^N$, and the eigenvectors by $\{\phi_n : V \rightarrow \mathbb{C}\}_{n=1}^N$. The Fourier transform of a graph signal $f : V \rightarrow \mathbb{C}$ is given by the vector of frequency intensities

$$\mathcal{F}f = (\langle f, \phi_n \rangle)_{n=1}^N,$$

where $\langle u, v \rangle$ is an inner product in \mathbb{C}^N , e.g., the standard dot product. The inverse Fourier transform of the vector $(v_n)_{n=1}^N$ is given by

$$\mathcal{F}^*(v_n)_{n=1}^N = \sum_{n=1}^N v_n \phi_n.$$

Since $\{\phi_n\}_{n=1}^N$ is an orthonormal basis, \mathcal{F}^* is the inverse of \mathcal{F} . A spectral graph filter \mathbf{G} based on the coefficients $(g_n)_{n=1}^N$ is defined by

$$\mathbf{G}f = \sum_{n=1}^N g_n \langle f, \phi_n \rangle \phi_n. \quad (1)$$

Any spectral filter defined by (1) is permutation *equivariant*, namely, does not depend on the indexing of the vertices. Re-indexing the vertices in the input, results in the same re-indexing of vertices in the output.

Spectral filters defined by (1) have two disadvantages. First, as shown in Subsection II-B, they are not transferable. Second, they entail high computational complexity. Formula (1) requires the computation of the eigendecomposition of the Laplacian Δ , which is computationally demanding and can be unstable when the number of vertices N is large. Moreover, there is no general “graph FFT” algorithm for computing the Fourier transform of a signal $f \in L^2(V)$, and (1) requires computing the frequency components $\langle f, \phi_n \rangle$ and their summation directly.

To overcome these two limitations, state-of-the-art methods, like [10], [14], [11], [12], are based on *functional calculus*. Functional calculus is the theory of applying functions $g :$

$\mathbb{C} \rightarrow \mathbb{C}$ on normal operators in Hilbert spaces. In the special case of a self-adjoint or unitary operator \mathbf{T} with a discrete spectrum, $g(\mathbf{T})$ is defined by

$$g(\mathbf{T})f = \sum_n g(\lambda_n) \langle f, \phi_n \rangle \phi_n, \quad (2)$$

for any vector f in the Hilbert space, where $\{\lambda_n, \phi_n\}$ is the eigendecomposition of the operator \mathbf{T} . The operator $g(\mathbf{T})$ is normal for general $g : \mathbb{C} \rightarrow \mathbb{C}$, self-adjoint for $g : \mathbb{C} \rightarrow \mathbb{R}$, and unitary for $g : \mathbb{C} \rightarrow e^{i\mathbb{R}}$ (where $e^{i\mathbb{R}}$ is the unit complex circle).

Definition (2) is canonical in the following sense. In the special case where

$$g(\lambda) = \frac{\sum_{l=0}^L c_l \lambda^l}{\sum_{l=0}^L d_l \lambda^l}$$

is a rational function, $g(\mathbf{T})$ can be defined in two ways. First, by (2), and second by compositions, linear combinations, and inversions, as

$$g(\mathbf{T}) = \left(\sum_{l=0}^L c_l \mathbf{T}^l \right) \left(\sum_{l=0}^L d_l \mathbf{T}^l \right)^{-1} \quad (3)$$

It can be shown that (2) and (3) are equivalent. Moreover, definition (2) is also canonical in regard to non-rational functions. Loosely speaking, if a rational function q approximates the function g , then the operator $q(\mathbf{T})$ approximates the operator $g(\mathbf{T})$.

Implementation (3) overcomes the limitation of definition (1), where now filters are defined via (2) with polynomial or rational function g . By relying on the spatial operations of compositions, linear combinations, and inversions, the computation of a spectral filter is carried out entirely in the spatial domain, without ever resorting to spectral computations. Thus, no eigendecomposition and Fourier transforms are ever computed. The inversions in $g(\mathbf{T})f$ involve solving systems of linear equations, which can be computed directly if N is small, or by some iterative approximation method for large N . Methods like [10], [15], [8], [12] use polynomial filters, and [14], [11], [13] use rational function filters. We term spectral methods based on functional calculus *functional calculus filters*.

B. The misconception of non-transferability of spectral graph filters

The non-transferability claim is formulated based on the sensitivity of the Laplacian eigendecomposition to small perturbation in \mathbf{W} , or equivalently in Δ . Namely, a small perturbation of Δ can result in a large perturbation of the eigendecomposition $\{\lambda_n, \phi_n\}_{n=1}^N$, which results in a large change in the filter defined via (1). This claim was stated in [3] only for spectral filters defined via (1), for which it is true. However, later papers misinterpreted this claim and applied it to functional calculus filters. The instability argument does not prove non-transferability, since state-of-the-art spectral methods do not explicitly use the eigenvectors, and do not parametrize the filter coefficients g_n via the index

n of the eigenvalues. Instead, state-of-the-art methods are based on functional calculus, and define the filter coefficients using a function $g : \mathbb{R} \rightarrow \mathbb{C}$, as $g(\lambda_n)$. The parametrization of the filter coefficients by g is indifferent to the specifics of how the spectrum is indexed, and instead represents an overall response in the frequency domain, where the *value* of each frequency determines its response, and not its index. In functional calculus filters defined by (2), a small perturbation of Δ that results in a perturbation of λ_n , also results in a perturbation of the coefficients $g(\lambda_n)$. It turns out, as we prove in this paper, that the perturbation in $g(\lambda_n)$ implicitly compensates for the instability of the eigendecomposition, and functional calculus spectral filters are stable.

III. MAIN RESULTS

A. Transferability of functional calculus filters

In this paper, we define transferability as the linear robustness of the filter to re-indexing of the vertices and perturbation of the topology of the graph. Thus, to formulate transferability, we combine equivariance with stability. Since spectral filters are known to be equivariant, transferability is equivalent to stability. Thus, our goal is to prove stability.

B. Linear stability of spectral filters

Stability is proven on a dense subspace of filters in $L^p(\mathbb{R})$, which we term the Cayley smoothness space. The definition of the Cayley smoothness space is based on the Cayley transform $\mathcal{C} : \mathbb{R} \rightarrow e^{i\mathbb{R}}$, defined by $\mathcal{C}(x) = \frac{x-i}{x+i}$.

Definition 1. The Cayley smoothness space $\text{Cay}^1(\mathbb{R})$ is the subspace of functions $g \in L^2(\mathbb{R})$ of the form $g(\lambda) = q(\mathcal{C}(\lambda))$, where $q : e^{i\mathbb{R}} \rightarrow \mathbb{C}$ is in $L^2(e^{i\mathbb{R}})$, and has classical Fourier coefficients $\{c_l\}_{l=-\infty}^{\infty}$ satisfying $\|g\|_{\mathcal{C}} := \sum_{l=-\infty}^{\infty} l |c_l| < \infty$.

The mapping $g \mapsto \|g\|_{\mathcal{C}}$ is a seminorm. It is not difficult to show that $\text{Cay}^1(\mathbb{R})$ is dense in each $L^p(\mathbb{R})$ space with $1 \leq p < \infty$. Intuitively, Cayley smoothness implies decay of the filter kernel in the spatial domain, since it models smoothness in a frequency domain. This can be formulated rigorously for graph filters based on Cayley polynomials ($g(\lambda) = q(\mathcal{C}(\lambda))$) with finite expansion $\{c_l\}_{l=-L}^L$ [11, Theorem 4].

For filters in the Cayley smoothness space we have linear stability, which is our main contribution.

Theorem 2. Let $\Delta \in \mathbb{C}^{N \times N}$ be a self-adjoint matrix that we call Laplacian. Let $\Delta' = \Delta + \mathbf{E}$ be self-adjoint, such that $\|\mathbf{E}\| < 1$. Let $g \in \text{Cay}^1(\mathbb{R})$. Then

$$\begin{aligned} \|g(\Delta) - g(\Delta')\| &\leq \|g\|_{\mathcal{C}} \left((\|\Delta\| + 1) \frac{\|\mathbf{E}\|}{1 - \|\mathbf{E}\|} + \|\mathbf{E}\| \right) \\ &= O(\|\Delta - \Delta'\|). \end{aligned} \quad (4)$$

IV. EXAMPLES

ChebNets. Consider the normalized translated Laplacian $\Delta_n - \mathbf{I}$. In ChebNets [10], g is a polynomial, and since the spectrum of $\Delta_n - \mathbf{I}$ is in $[-1, 1]$, the values of g outside $[-1, 1]$ do not affect the filter $g(\Delta_n - \mathbf{I})$. Thus, we may assume that

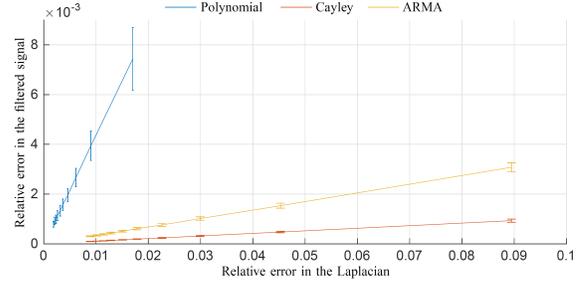


Fig. 1. Average relative error in the output of the filter as a function of the average relative error in the Laplacian, for three fixed filters: low-pass polynomial and Cayley filters of order 3, and all-pass ARMA filter of order 3.

g is a polynomial in $[-1, 1]$, and padded outside $[-1, 1]$ to obtain a smooth compactly supported function. It is easy to see that such a g is in $\text{Cay}^1(\mathbb{R})$. Thus, for two translated normalized Laplacians $\Delta_n - \mathbf{I}$ and $\Delta'_n - \mathbf{I}$ of two graphs, $\|g(\Delta_n - \mathbf{I}) - g(\Delta'_n - \mathbf{I})\| = O(\|\Delta_n - \Delta'_n\|)$.

General rational functions. The above claim is also true for general rational functions, if we assume that the spectrum of Δ, Δ' is contained in some pre-defined band $[0, M]$. Thus, the polynomial filters of [15], [12] and the ARMA rational function filters of [14], [13] are also transferable, under the assumption of uniformly bounded Laplacians.

CayleyNets. CayleyNets [11] are always transferable, since a Cayley filter is by definition in $\text{Cay}^1(\mathbb{R})$, with finite expansion.

To corroborate the proposed theoretical result, in Figure 1 we test the above three examples in the Molene weather dataset¹. The graph comprises $N = 32$ weather stations, with weights given as the Gaussian of the physical distances between stations. Each of the 744 graph signals is a temperature recording. For the polynomial filter we consider the normalized Laplacian, while for the Cayley and ARMA filters we consider the unnormalized Laplacian. The results are averaged over 100 different perturbations in the topology and the 744 graph signals. The experimental results concord with the theoretical linear stability property.

V. PROOF OF THEOREM 2

We start with a useful lemma.

Lemma 3. Suppose $\mathbf{B}, \mathbf{D}, \mathbf{E} \in \mathbb{C}^{N \times N}$ are matrices satisfying $\mathbf{B} = \mathbf{D} + \mathbf{E}$, and $\|\mathbf{B}\|, \|\mathbf{D}\| \leq C$ for some $C > 0$. Then for every $l \geq 0$

$$\|\mathbf{B}^l - \mathbf{D}^l\| \leq lC^{l-1} \|\mathbf{E}\|. \quad (5)$$

Proof. Let $l \in \mathbb{N}$.

$$\mathbf{D}^l - \mathbf{B}^l = \mathbf{D}^{l-1}(\mathbf{D} - \mathbf{B}) + (\mathbf{D}^{l-1} - \mathbf{B}^{l-1})\mathbf{B} \quad (6)$$

so

$$\|\mathbf{D}^l - \mathbf{B}^l\| \leq C^{l-1} \|\mathbf{E}\| + \|\mathbf{D}^{l-1} - \mathbf{B}^{l-1}\| C. \quad (7)$$

¹Access to the raw data is possible from <https://donneespubliques.meteofrance.fr/donneeslibres/Hackathon/RADOMEH.tar.gz>

Now, (5) follows by repeatedly using (7) with decreasing powers $l - j$, $j = 1, \dots, l - 1$. \square

Next, we cite a general property from spectral theory.

Lemma 4. *Let T be a bounded normal operator in a Hilbert space. Let σ be the spectrum of T . Define the infinity norm on the space of bounded continuous functions $f : \sigma \rightarrow \mathbb{C}$ by*

$$\|f - g\|_\infty^\sigma = \sup_{x \in \sigma} |f(x) - g(x)|.$$

Then

$$\|f(T) - g(T)\| = \|f - g\|_\infty^\sigma$$

where the norm in the left-hand-side is the operator norm.

To prove Theorem 2, we start with a version of the theorem restricted to $g = q \circ \mathcal{C} \in \text{Cay}^1(\mathbb{R})$ where q has a finite expansion with coefficients $(c_l)_{l=-L}^L$.

Proof of Theorem 2 for finite Cayley expansions. We prove for q with coefficients $(c_l)_{l=1}^L$. The part with coefficients $(c_l)_{l=-L}^0$ is treated similarly. Note that

$$\begin{aligned} \mathcal{C}(\Delta) - \mathcal{C}(\Delta') &= (\Delta - i)(\Delta + i)^{-1} - (\Delta - i)(\Delta' + i)^{-1} \\ &\quad + (\Delta - i)(\Delta' + i)^{-1} - (\Delta' - i)(\Delta' + i)^{-1} \end{aligned}$$

so

$$\|\mathcal{C}(\Delta) - \mathcal{C}(\Delta')\|$$

$$\leq \|(\Delta - i)\| \|(\Delta + i)^{-1} - (\Delta' + i)^{-1}\| + \|(\Delta' + i)^{-1}\| \|\mathbf{E}\|.$$

By the fact that the spectrum of Δ' is real, $\|(\Delta' + i)^{-1}\| \leq 1$, and we have

$$\begin{aligned} &\|\mathcal{C}(\Delta) - \mathcal{C}(\Delta')\| \\ &\leq (\|\Delta\| + 1) \|(\Delta + i)^{-1} - (\Delta' + i)^{-1}\| + \|\mathbf{E}\|. \end{aligned} \quad (8)$$

Let us bound $\|(\Delta + i)^{-1} - (\Delta' + i)^{-1}\|$ in terms of $\|\mathbf{E}\|$. Since $\|\mathbf{E}\| < 1$ we may expand

$$\begin{aligned} (\Delta + i + \mathbf{E})^{-1} &= (\Delta + i)^{-1} (\mathbf{I} + \mathbf{E}(\Delta + i)^{-1})^{-1} \\ &= (\Delta + i)^{-1} \left(\sum_{k=0}^{\infty} (-1)^k (\mathbf{E}(\Delta + i)^{-1})^k \right) \end{aligned}$$

so, by $\|(\Delta + i)^{-1}\| \leq 1$,

$$\|(\Delta + i)^{-1} - (\Delta' + i)^{-1}\| \leq \frac{\|\mathbf{E}\|}{1 - \|\mathbf{E}\|}. \quad (9)$$

Now, by (8) and (9),

$$\|\mathcal{C}(\Delta) - \mathcal{C}(\Delta')\| \leq (\|\Delta\| + 1) \frac{\|\mathbf{E}\|}{1 - \|\mathbf{E}\|} + \|\mathbf{E}\|. \quad (10)$$

Observe that $\mathcal{C}(\Delta)$ and $\mathcal{C}(\Delta')$ are unitary, so their spectrum is bounded by $C = 1$. Thus, by Lemma 3 and the triangle inequality on the polynomial expansions of $q(\mathcal{C}(\Delta)) - q(\mathcal{C}(\Delta'))$,

$$\begin{aligned} \|g(\Delta) - g(\Delta')\| &= \|q(\mathcal{C}(\Delta)) - q(\mathcal{C}(\Delta'))\| \\ &\leq \sum_{l=1}^L l c_l \|\mathcal{C}(\Delta) - \mathcal{C}(\Delta')\| \end{aligned}$$

which gives (4). \square

Proof of Theorem 2. Theorem 2 follows the above result by a simple density argument. Given $g = q \circ \mathcal{C} \in \text{Cay}^1(\mathbb{R})$, we consider the truncations $g_L = q_L \circ \mathcal{C}$, where q_L is restricted to the coefficients $(c_l)_{l=-L}^L$. We base a three-epsilon argument on the expansion, for any $L \in \mathbb{N}$,

$$\begin{aligned} \|g(\Delta) - g(\Delta')\| &\leq \|g(\Delta) - g_L(\Delta)\| \\ &\quad + \|g_L(\Delta) - g_L(\Delta')\| \\ &\quad + \|g_L(\Delta') - g(\Delta')\|. \end{aligned} \quad (11)$$

For any $\epsilon > 0$, by Lemma 4, the first and the last terms of the right-hand side of 11 can be made smaller than $\epsilon/2$ by choosing L large enough. Moreover, for any $L \in \mathbb{N}$,

$$\|g_L\|_{\mathcal{C}} = \sum_{l=-L}^L l |c_l| \leq \sum_{l=-\infty}^{\infty} l |c_l| = \|g\|_{\mathcal{C}},$$

so, by Theorem 2 for finite Cayley expansions, for every $\epsilon > 0$

$$\|g(\Delta) - g(\Delta')\| \leq \|g\|_{\mathcal{C}} \left((\|\Delta\| + 1) \frac{\|\mathbf{E}\|}{1 - \|\mathbf{E}\|} + \|\mathbf{E}\| \right) + \epsilon. \quad (12)$$

Since (12) is true for every $\epsilon > 0$, we must have (4). \square

ACKNOWLEDGMENT

E. Isufi's research is supported in part by the KAUST-MIT-TUD-Caltech consortium grant OSR-2015-Sensors-2700 Ext. 2018. G. Kutyniok acknowledges partial support by the Bundesministerium für Bildung und Forschung (BMBF) through the Berliner Zentrum for Machine Learning (BZML), Project AP4, by the Deutsche Forschungsgemeinschaft (DFG) through Grants CRC 1114 "Scaling Cascades in Complex Systems", Project B07, CRC/TR 109 "Discretization in Geometry and Dynamics", Projects C02 and C03, RTG DAEDALUS (RTG 2433), Projects P1 and P3, RTG BIOQIC (RTG 2260), Projects P4 and P9, SPP 1798 "Compressed Sensing in Information Processing", Project Coordination and Project Massive MIMO-I/II, by the Berlin Mathematics Research Center MATH+, Projects EF1-1 and EF1-4, and by the Einstein Foundation Berlin.

REFERENCES

- [1] S. Segarra, A. G. Marques, and A. Ribeiro, "Optimal graph-filter design and applications to distributed linear network operators," *IEEE Transactions on Signal Processing*, vol. 65, no. 15, pp. 4117–4131, 2017.
- [2] E. Isufi, A. Loukas, A. Simonetto, and G. Leus, "Filtering random graph processes over random time-varying graphs," *IEEE Transactions on Signal Processing*, vol. 65, no. EPFL-ARTICLE-230521, pp. 4406–4421, 2017.
- [3] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: Going beyond euclidean data," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 18–42, July 2017.
- [4] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *arXiv preprint arXiv:1901.00596*, 2019.
- [5] M. Gori, G. Monfardini, and F. Scarselli, "A new model for learning in graph domains," in *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, vol. 2, July 2005, pp. 729–734 vol. 2.
- [6] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 61–80, Jan 2009.

- [7] 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," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5425–5434, 2017.
- [8] A. Ortega, P. Frossard, J. Kovačević, J. M. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 808–828, 2018.
- [9] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, "Spectral networks and locally connected networks on graphs," *Proc. ICLR*, 2013.
- [10] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Proc. NIPS*, 2016.
- [11] R. Levie, F. Monti, X. Bresson, and M. M. Bronstein, "Cayleynets: Graph convolutional neural networks with complex rational spectral filters," *IEEE Transactions on Signal Processing*, vol. 67, no. 1, pp. 97–109, Jan 2019.
- [12] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, "Convolutional neural networks architectures for signals supported on graphs," *arXiv preprint arXiv:1805.00165*, 2018.
- [13] F. M. Bianchi, D. Grattarola, L. Livi, and C. Alippi, "Graph neural networks with convolutional arma filters," *arXiv preprint arXiv:1901.01343*, 2019.
- [14] E. Isufi, A. Loukas, A. Simonetto, and G. Leus, "Autoregressive moving average graph filtering," *IEEE Transactions on Signal Processing*, vol. 65, no. 2, pp. 274–288, Jan 2017.
- [15] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *In Proceedings of International Conference on Learning Representations (ICLR)*, 2017.