

Fairness-Aware Graph Filter Design

O. Deniz Kose

EECS

University of California Irvine

Irvine, USA

okose@uci.edu

Yanning Shen

EECS

University of California Irvine

Irvine, USA

yannings@uci.edu

Gonzalo Mateos

ECE

University of Rochester

Rochester, USA

gmateosb@ur.rochester.edu

Abstract—Graphs are mathematical tools that can be used to represent complex real-world systems, such as financial markets and social networks. Hence, machine learning (ML) over graphs has attracted significant attention recently. However, it has been demonstrated that ML over graphs amplifies the already existing bias towards certain under-represented groups in various decision-making problems due to the information aggregation over biased graph structures. Faced with this challenge, here we take a fresh look at the problem of bias mitigation in graph-based learning by bringing to bear insights from graph signal processing. Specifically, we design a fair graph filter that can be employed in a versatile manner for graph-based learning tasks. The design of the proposed filter is based on a bias analysis and its optimality in mitigating bias compared to its fairness-agnostic counterpart is established. Experiments on real-world networks for node classification demonstrate the efficacy of the proposed filter design in mitigating bias, while attaining similar utility and better stability compared to baseline algorithms.

Index Terms—Fairness, graph filter, graph neural network, node classification, bias mitigation.

I. INTRODUCTION

We live in the era of connectivity, where the actions of humans and devices are increasingly driven by their relations to others. Concurrently, a significant amount of data that describes different interconnected systems, such as social networks, Internet of Things (IoT), the Web, and financial markets, is increasingly available. Processing and learning from such data can provide significant understanding and advancements for the corresponding networked systems [1], [2]. In this context, machine learning (ML) over graphs has attracted increasing attention [3], [4], since graphs are widely utilized to represent complex underlying relations in real-world networks [5].

These relational patterns can be captured by graph edges, while attributes of nodes (nodal features) can be interpreted as signals defined on the vertices. For example, in a social network, user ages can be modeled as a graph signal, and the friendship information can be encoded by the edges. Graph signal processing (GSP) extends the tools in classical signal processing to graph signals [6], such as frequency analysis, sampling and filtering [7]–[12]. GSP and ML over graphs are closely intertwined, where the tools in one domain can be useful in the other one [6], [13]. For instance, it has

been demonstrated that graph neural networks (GNNs) can be designed, analyzed, and improved by leveraging GSP-based insights [13]–[15], which underscores the advancements that can be made by cross-pollinating the findings in both domains. In this paper, we align with this vision and leverage GSP advances to enhance fairness in ML over graphs pipelines.

Fairness in ML. Despite the growing interest towards learning over graphs, the widespread deployment of these algorithms in real-world decision systems depends heavily on how socially responsible they are. Indeed, several studies have demonstrated that ML models propagate the historical bias within the training data and lead to discriminatory results in ensuing applications [16]. Note that algorithmic bias refers to the stereotypical correlations the ML models encode and propagate with respect to certain sensitive attributes (e.g., gender, race). Specific to graph-based learning, the utilization of graph structure in the algorithm design has been shown to amplify the already existing bias [17]. Motivated by this concern, recent works focus on fairness-aware learning over graphs and advocate different techniques, such as adversarial regularization [17], [18], fairness constraints [19], [20], and fairness-aware graph data augmentation [21]–[23].

Proposed approach and contributions. Here, we first analyze the sources of bias inherent to the graph topology and design a general purpose bias-mitigation scheme by bringing to bear GSP notions. Our previous endeavor [24] is also built upon spectral analysis of graph signals, where a fairness-aware dimensionality reduction algorithm was developed. However, in [24], the information carried in certain frequencies is completely removed, which can adversely affect the overall utility (accuracy for node classification) of the underlying ML task. Hence, in the present work, we design a more flexible bias mitigation algorithm that filters out the information coming from the sensitive attribute signal (e.g. race, gender in social networks), while also providing a better fairness-utility trade-off. To this end, we design a graph filter which selectively changes the graph Fourier coefficients [7] of input signals, while taking fairness into account. Overall, our contributions are:

- i) Based on theoretical analysis of bias in graph topologies, a graph filter is designed to mitigate bias. The novel filter is versatile and can be employed in different stages of the learning pipeline;
- ii) The proposed filter is proved to be more effective than the

Work in this paper was supported in part by the Google Research Scholar Award, and the NSF awards CCF-1750428 and CCF-1934962.

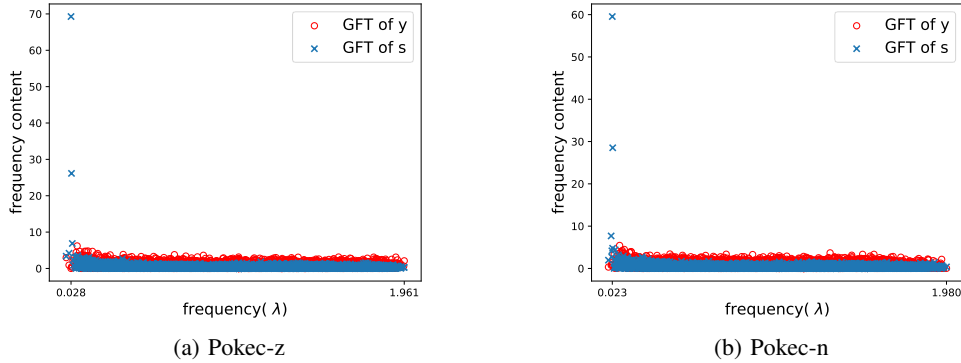


Fig. 1: Spectra of the graph signals \mathbf{s} (sensitive attributes) and \mathbf{y} (labels) over different graph frequencies. There are few low frequencies where the magnitudes of $\tilde{\mathbf{s}}$ are markedly higher than those of $\tilde{\mathbf{y}}$. Pokec-z and Pokec-n are real-world networks, whose statistics are presented in Table II.

fairness-agnostic counterpart in reducing bias; and **iii)** Experimental results for node classification on real-world networks corroborate the effectiveness of the proposed method in mitigating bias, while providing similar utility to state-of-the-art algorithms.

II. PRELIMINARIES AND PROBLEM STATEMENT

The focus of this study is to mitigate bias in graph-based learning algorithms by employing a graph filter for a given undirected graph $\mathcal{G} := (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} := \{v_1, v_2, \dots, v_N\}$ denotes the set of nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges. Connectivity of the input graph is encoded in the graph adjacency matrix $\mathbf{A} \in \{0, 1\}^{N \times N}$, where $A_{ij} = 1$ if and only if $(v_i, v_j) \in \mathcal{E}$. In addition, $\mathbf{X} \in \mathbb{R}^{N \times F}$ represents the nodal features of the input graph. The diagonal graph degree matrix is $\mathbf{D} \in \mathbb{R}^{N \times N}$, where D_{ii} denotes the the degree of v_i , and $\mathbf{L} = \mathbf{I}_N - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$ is the normalized graph Laplacian matrix. The sensitive attribute is defined to be the nodal feature on which the decisions should not be dependent on for fair decision making. Herein, the sensitive attribute is assumed to be binary and is denoted by $\mathbf{s} \in \{-1, 1\}^N$. The feature vector and the sensitive attribute of node v_i are denoted by $\mathbf{x}_i \in \mathbb{R}^F$ and $s_i \in \{-1, 1\}$, respectively.

Graph signal processing fundamentals. The graph Fourier transform (GFT) is an orthonormal transform that provides the representation of a graph signal $\mathbf{z} \in \mathbb{R}^N$ in the graph spectral domain [7], [25], [26]. Specifically, taking the GFT of a graph signal amounts to projecting the signal onto a space spanned by the orthogonal eigenvectors of the positive semi-definite (PSD) normalized graph Laplacian matrix \mathbf{L} [7]. Let the eigendecomposition of the normalized Laplacian be $\mathbf{L} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$, where $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_N)$ collects the non-negative eigenvalues and \mathbf{V} is the matrix of Laplacian eigenvectors. Then, the GFT of the graph signal $\mathbf{z} \in \mathbb{R}^N$ is given by $\tilde{\mathbf{z}} = \mathbf{V}^T \mathbf{z}$. Graph frequencies correspond to the eigenvalues of the Laplacian (a measure of smoothness of the eigenvectors with respect to the graph), meaning that

the GFT decomposes signals into frequency modes (i.e., the eigenvectors of \mathbf{L}) of different variability over \mathcal{G} .

In classical signal processing, filters are utilized to manipulate signals such that their, e.g., unwanted components are attenuated or removed. Similarly, graph filters can be used to modify graph signals for different purposes including graph signal classification [27], [28], smoothing and denoising [29], [30]. Filtering an input graph signal $\mathbf{z}_{\text{in}} \in \mathbb{R}^N$ via a filter with frequency response $\tilde{\mathbf{h}} := [\tilde{h}_1, \dots, \tilde{h}_N]^T$ can be mathematically expressed as (e.g., [6], [7], [14])

$$\mathbf{z}_{\text{out}} = \mathbf{V} \underbrace{\text{diag}(\tilde{h}_1, \dots, \tilde{h}_N)}_{\text{Frequency domain filtering}} \tilde{\mathbf{z}}_{\text{in}}. \quad (1)$$

Therefore, filtering in the frequency domain corresponds to point-wise multiplication of the input signal's GFT $\tilde{\mathbf{z}}_{\text{in}}$ with the frequency response of graph filter $\tilde{\mathbf{h}}$. In this paper, given \mathcal{G} and \mathbf{s} , we address the problem of designing a graph filter with frequency response $\tilde{\mathbf{h}} \in \mathbb{R}^N$, so that the bias caused by the graph topology can be attenuated when the filter is applied to input/output graph signals in the learning algorithm.

III. BIAS MITIGATING GRAPH FILTER DESIGN

A. Spectrum Analysis

The homophily principle suggests that nodes with similar attributes are more likely to connect in networks, which hints at denser connectivity between the nodes with the same sensitive attributes and also with the same label [31]. Accordingly, both the sensitive attribute \mathbf{s} and node labels \mathbf{y} are expected to be smooth signals over \mathcal{G} . This implies higher energy concentration for $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{y}}$ over lower frequencies. However, the extent of the overlap between the spectra of $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{y}}$ plays an important role in fairness-aware filter design, since we want to preserve the necessary information for a downstream task (node classification in this paper) after “filtering out” traces of the sensitive attribute. Indeed, the spectra of $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{y}}$ should not match completely for the feasibility of our main idea.

To examine this, the GFT coefficients in $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{y}}$ over different frequencies are depicted for two real-world networks

in Figure 1. Notice how the spectra of $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{y}}$ exhibit similar characteristics. However, there are certain frequencies where the magnitudes of $\tilde{\mathbf{s}}$ take significantly higher values than those of $\tilde{\mathbf{y}}$. Such observation inspires us to design a graph filter which attenuates the sensitive information while preserving the data structure necessary for downstream ML tasks.

B. Bias Analysis

Features which are correlated with the sensitive attribute lead to intrinsic bias, even when the sensitive attribute is not utilized in learning [32]. The correlation between the input to a learning algorithm and sensitive attributes is thus a measure of the resulting bias. Motivated by this, the linear correlation between the sensitive attribute signal \mathbf{s} and graph topology $\hat{\mathbf{A}} = \mathbf{D}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}^{-\frac{1}{2}}$ is considered for the ensuing bias analysis.

Several graph-based learning approaches rely on node representations obtained via local aggregation of information, a process that can be summarized as

$$\mathbf{R} = \hat{\mathbf{A}}\mathbf{X}\mathbf{W}, \quad (2)$$

where \mathbf{R} denotes the obtained node representations, \mathbf{X} is the input graph signal, and \mathbf{W} represents the learnable weight matrix; see e.g., [3], [4]. Hence, if a filtered graph signal $\mathbf{X}^f = \mathbf{V}\text{diag}(\tilde{\mathbf{h}})\mathbf{V}^\top\mathbf{X}$ is input, the obtained representation becomes

$$\begin{aligned} \mathbf{R}^f &= \hat{\mathbf{A}}\mathbf{X}^f\mathbf{W} \\ &= \mathbf{V}(\mathbf{I}_N - \mathbf{\Lambda})\mathbf{V}^\top\mathbf{X}^f\mathbf{W} \\ &= \mathbf{V}(\mathbf{I}_N - \mathbf{\Lambda})\mathbf{V}^\top\mathbf{V}\text{diag}(\tilde{\mathbf{h}})\mathbf{V}^\top\mathbf{X}\mathbf{W} \\ &= \mathbf{V}(\mathbf{I}_N - \mathbf{\Lambda})\text{diag}(\tilde{\mathbf{h}})\mathbf{V}^\top\mathbf{X}\mathbf{W} \\ &= \mathbf{A}^f\mathbf{X}\mathbf{W}, \end{aligned} \quad (3)$$

where $\mathbf{A}^f := \mathbf{V}(\mathbf{I}_N - \mathbf{\Lambda})\text{diag}(\tilde{\mathbf{h}})\mathbf{V}^\top$. Therefore, if we feed the aggregation process with a filtered signal \mathbf{X}^f , the effective graph topology that is utilized in the information aggregation becomes \mathbf{A}^f . Building on this key observation, the linear correlation between the sensitive attributes \mathbf{s} and \mathbf{A}^f is employed as a bias measure, which is proportional to $\mathbf{s}^\top\mathbf{A}^f_{:,i}$ for the i th column of \mathbf{A}^f . The following Proposition reveals the sources of bias and provides an upper bound on the total correlation between \mathbf{s} and \mathbf{A}^f , where total correlation is given by $\rho := \|\mathbf{s}^\top\mathbf{A}^f\|_1$. The proof is omitted due to lack of space.

Proposition 1: For a filtered input graph signal \mathbf{X}^f by a graph filter with frequency response $\tilde{\mathbf{h}}$, $\rho := \|\mathbf{s}^\top\mathbf{A}^f\|_1$ can be bounded by

$$\rho \leq \sqrt{N} \sum_{i=1}^N |\tilde{s}_i|(1 - \lambda_i)|\tilde{h}_i|. \quad (4)$$

Proposition 1 shows that the linear correlation between the effective graph topology and sensitive attributes is a function of $\sum_{i=1}^N |\tilde{s}_i|(1 - \lambda_i)|\tilde{h}_i|$. In the sequel we design a ‘‘matched’’ graph filter to reduce this term and hence the bias.

C. Fair Filter Design

Here we design a fair filter $\tilde{\mathbf{h}}^{\text{fair}}$ to ‘‘filter-out’’ the sensitive information from the bias-amplifying graph connectivity. The novel filter can be applied to input/output graph signals used in general purpose learning algorithms; see also Remark 1.

From Proposition 1, it follows that the intrinsic bias is bounded above by $\sum_{i=1}^N |\tilde{s}_i|(1 - \lambda_i)|\tilde{h}_i|$. Hence, a filter can be designed that suppresses frequencies where $m_i := |\tilde{s}_i|(1 - \lambda_i)$ are largest. To this end, let $\mathcal{C} := \{i | m_i > \tau m_{\max}\}$, where τ is a hyperparameter and $m_{\max} = \max_i\{m_i\}$. The frequency response corresponding to larger m_i are designed to reduce the value of $|\tilde{s}_i|(1 - \lambda_i)|\tilde{h}_i| = m_i|\tilde{h}_i|$ in the upper bound of ρ in (4). Accordingly, the frequency response of the fairness-aware graph filter $\tilde{\mathbf{h}}^{\text{fair}}$ is given by

$$\tilde{h}_i^{\text{fair}} = \begin{cases} \frac{\frac{1}{N-k} \sum_{1 \leq j \leq N, j \notin \mathcal{C}} m_j}{m_i}, & \text{if } i \in \mathcal{C} \\ 1, & \text{otherwise} \end{cases}, \quad (5)$$

where k is the cardinality of set \mathcal{C} . It can be observed that with the designed filter the resulting term $m_i|\tilde{h}_i^{\text{fair}}| = \frac{1}{N-k} \sum_{1 \leq j \leq N, j \notin \mathcal{C}} m_j \leq m_i, \forall i \in \mathcal{C}$. While for $i \notin \mathcal{C}$, $m_i|\tilde{h}_i^{\text{fair}}| = m_i$ is unchanged for frequencies that are less relevant to the bias in order to preserve the information in the original graph structure.

It is important to emphasize that ρ can be minimized by setting $\tilde{\mathbf{h}}^{\text{fair}} = \mathbf{0}$, which is equivalent to filtering out all information. However, such trivial design is uninteresting, since it maximally sacrifices the utility in the ensuing task. Therefore, there needs to be a trade-off between the utility and fairness. This trade-off can be empirically adjusted via the design parameter τ . Furthermore, we also theoretically demonstrate in Proposition 2 that the proposed fair filter $\tilde{\mathbf{h}}^{\text{fair}}$ decreases the upper bound on ρ more effectively than a uniform fairness-agnostic filter $\tilde{\mathbf{h}}^u$ incurring the same amount of information loss, i.e., $\|\tilde{\mathbf{h}}^{\text{fair}}\|_1 = \|\tilde{\mathbf{h}}^u\|_1$.

Proposition 2: The proposed fair filter $\tilde{\mathbf{h}}^{\text{fair}}$ results in a lower upper bound for the correlation measure ρ between \mathbf{s} and \mathbf{A}^f , when compared to the fairness-agnostic counterpart $\tilde{\mathbf{h}}^u = \frac{1}{N} \sum_{i=1}^N \tilde{h}_i^{\text{fair}}, \forall j = 1, \dots, N$, meaning

$$\sum_{i=1}^N |\tilde{s}_i|(1 - \lambda_i)|\tilde{h}_i^{\text{fair}}| \leq \sum_{i=1}^N |\tilde{s}_i|(1 - \lambda_i)|\tilde{h}_i^u|. \quad (6)$$

Remark 1: The designed fair filter $\tilde{\mathbf{h}}^{\text{fair}}$ can be employed in a flexible way to mitigate bias for different graph-based learning algorithms. It can be applied to the graph signals that are input to, or, output from the learning algorithms. Models that are designed for attributed graphs generally utilize the information coming from both the nodal features and graph topology. Thus, the proposed filter $\tilde{\mathbf{h}}^{\text{fair}}$ can be applied to the nodal features before they are fed to the learning pipeline, in order to prevent the amplification of bias due to the graph connectivity. Alternatively, for any algorithm that outputs a graph signal (e.g., node labels in node classification), $\tilde{\mathbf{h}}^{\text{fair}}$ can be employed on the output graph signal as a fairness-aware post-processing operation. Overall, impact of the proposed fair

TABLE I: Proposed filter as fairness-aware pre-processing operator for a GNN model.

	Pokec-z			Pokec-n		
	Accuracy (%)	Δ_{SP} (%)	Δ_{EO} (%)	Accuracy (%)	Δ_{SP} (%)	Δ_{EO} (%)
GNN	66.52 \pm 0.27	6.79 \pm 2.45	7.26 \pm 3.29	64.96 \pm 0.19	6.79 \pm 2.45	7.26 \pm 3.29
Adversarial	64.26 \pm 1.79	4.85 \pm 2.16	5.99 \pm 2.71	64.22 \pm 0.71	4.34 \pm 3.87	3.84 \pm 2.71
EDITS	62.67 \pm 2.64	3.17 \pm 2.49	4.54 \pm 2.99	62.67 \pm 0.51	4.40 \pm 2.41	5.38 \pm 1.92
FairDrop	66.79 \pm 0.65	9.11 \pm 1.89	8.35 \pm 3.81	64.33 \pm 0.44	4.46 \pm 1.67	5.02 \pm 1.84
$\bar{\mathbf{h}}^{\text{fair}}$ + GNN	66.34 \pm 0.27	1.23 \pm 1.43	2.15 \pm 1.96	65.05 \pm 0.21	2.13 \pm 0.93	2.39 \pm 1.78

filter can permeate several GNN-based learning frameworks in a versatile manner.

IV. EXPERIMENTAL RESULTS

A. Dataset and Experimental Setup

TABLE II: Dataset statistics.

Dataset	$ \mathcal{S}_{-1} $	$ \mathcal{S}_1 $	Inter-edges	Intra-edges	F
Pokec-z	4851	2808	1730	39370	59
Pokec-n	4040	2145	1422	29220	59

Datasets. The performance of the proposed fair filter design is evaluated on the node classification task over real-world social networks Pokec-z and Pokec-n [17]. Pokec-z and Pokec-n are the sampled versions of the 2012 Pokec network [33], which is a Facebook-like social network in Slovakia [17]. The region of the users is utilized as the sensitive attributes, where the users are from two major regions. Labels for the node classification task are the binarized working field of the users. Statistical information for the utilized datasets are presented in Table II, where \mathcal{S}_i represents the set of nodes with sensitive attribute i and inter-edges connect different sensitive attributes, while intra-edges are the edges between the same sensitive attribute. Note that $N = |\mathcal{S}_{-1}| + |\mathcal{S}_1|$.

Evaluation metrics. For the utility metric of node classification, accuracy is utilized. For fairness assessment, two quantitative measures of group fairness metrics are reported, namely *statistical parity*: $\Delta_{SP} = |P(\hat{y} = 1 | s = -1) - P(\hat{y} = 1 | s = 1)|$ and *equal opportunity*: $\Delta_{EO} = |P(\hat{y} = 1 | y = 1, s = -1) - P(\hat{y} = 1 | y = 1, s = 1)|$, where y represents the ground truth label, and \hat{y} is the predicted label. Lower values for Δ_{SP} and Δ_{EO} indicate better fairness performance [17] and are more desirable.

Implementation details. In the experiments, the designed filter is employed as a pre-processing operator on the features that are input to GNN layers in a two-layer graph convolutional network [4]. The model is trained over 40% of the nodes, while the remaining nodes are equally divided to validation and test sets. The hyperparameter τ is selected as 0.05 for Pokec-z and 0.04 for Pokec-n via grid search among the values $\{0.04, 0.05, 0.06\}$. Furthermore, adversarial regularization [17], EDITS [23], and FairDrop [22] are employed as fairness-aware baselines in the experiments. For adversarial regularization, the multiplier of the regularizer is tuned via a grid

search among the values $\{0.1, 1, 10, 100, 1000\}$ (the multiplier of classification loss is assigned to be 1). Furthermore, for EDITS, the threshold proportion is tuned among the values $\{0.015, 0.02, 0.29\}$, where these values are the optimized thresholds for other datasets used in the corresponding work. Finally, δ in FairDrop algorithm is tuned among the values $\{0.7, 0.8, 0.9\}$. For all experiments, results are obtained for five random data splits, and their average along with the standard deviations are reported.

B. Results

Results for GNN-based node classification are presented in Table I. For the proposed scheme, the natural baseline is to employ the exactly same GNN model without the fair filter as a pre-processing operator (denoted as GNN in Table I). Moreover, Adversarial, EDITS, and FairDrop in Table I stand for the employment of adversarial regularization in training [17], and state-of-the-art fairness-aware baselines EDITS [23], and FairDrop [22], respectively. Results in Table I demonstrate that the proposed fair filter consistently achieves better fairness measures compared to the fairness-aware baselines Adversarial and EDITS [23], along with better utility values. While our approach outperforms FairDrop [22] in terms of both fairness and utility on Pokec-z, FairDrop [22] achieves lower Δ_{EO} on Pokec-n. However, this better Δ_{EO} result of FairDrop on Pokec-n is accompanied by a drop in the utility. Furthermore, it can be observed that the employment of the proposed filter leads to the lowest standard deviation values, and therefore enhances the stability of the results. Overall, the results corroborate the efficacy of the proposed filter design in mitigating bias while also providing similar utility measures compared to the state-of-the-art fairness-aware baselines.

V. CONCLUSION

We developed a fairness-aware graph filter that can be flexibly employed in various graph-based ML and signal processing algorithms. Theoretical analysis on the source of bias is provided, which guides the design of the fairness-aware filter. The novel filter is provably more effective in terms of mitigating bias with the same amount of filtered information compared to the fairness-agnostic counterpart. Node classification experiments on real-world networks demonstrate that the proposed fair filter consistently provides better fairness and robustness together with similar utility compared to baselines. This work opens up exciting future directions. For instance,

the proposed design requires a Laplacian eigendecomposition, which incurs a complexity of $\mathcal{O}(N^3)$ and may not be ideal for large-scale graphs. Computationally-efficient designs are part of our ongoing research agenda.

REFERENCES

- [1] E. D. Kolaczyk and G. Csárdi, *Statistical analysis of network data with R*. Springer, 2014, vol. 65.
- [2] W. L. Hamilton, “Graph representation learning,” *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 14, no. 3, pp. 1–159, 2020.
- [3] I. Chami, S. Abu-El-Haija, B. Perozzi, C. Ré, and K. Murphy, “Machine learning on graphs: A model and comprehensive taxonomy,” *Journal of Machine Learning Research*, vol. 23, no. 89, pp. 1–64, 2022.
- [4] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2016.
- [5] G. Mateos, S. Segarra, A. G. Marques, and A. Ribeiro, “Connecting the dots: Identifying network structure via graph signal processing,” *IEEE Signal Processing Magazine*, vol. 36, no. 3, pp. 16–43, 2019.
- [6] 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.
- [7] 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 Processing Magazine*, vol. 30, no. 3, pp. 83–98, 2013.
- [8] A. Sandryhaila and J. M. Moura, “Discrete signal processing on graphs,” *IEEE Transactions on Signal Processing*, vol. 61, no. 7, pp. 1644–1656, 2013.
- [9] A. G. Marques, S. Segarra, G. Leus, and A. Ribeiro, “Stationary graph processes and spectral estimation,” *IEEE Transactions on Signal Processing*, vol. 65, no. 22, pp. 5911–5926, 2017.
- [10] 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, 2016.
- [11] E. Isufi, F. Gama, D. I. Shuman, and S. Segarra, “Graph filters for signal processing and machine learning on graphs,” *arXiv preprint arXiv:2211.08854*, 2022.
- [12] D. Romero, M. Ma, and G. B. Giannakis, “Kernel-based reconstruction of graph signals,” *IEEE Transactions on Signal Processing*, vol. 65, no. 3, pp. 764–778, 2016.
- [13] X. Dong, D. Thanou, L. Toni, M. Bronstein, and P. Frossard, “Graph signal processing for machine learning: A review and new perspectives,” *IEEE Signal Processing Magazine*, vol. 37, no. 6, pp. 117–127, 2020.
- [14] F. Gama, E. Isufi, G. Leus, and A. Ribeiro, “Graphs, convolutions, and neural networks: From graph filters to graph neural networks,” *IEEE Signal Processing Magazine*, vol. 37, no. 6, pp. 128–138, 2020.
- [15] F. Gama, J. Bruna, and A. Ribeiro, “Stability properties of graph neural networks,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 5680–5695, 2020.
- [16] A. Beutel, J. Chen, Z. Zhao, and E. H. Chi, “Data decisions and theoretical implications when adversarially learning fair representations,” in *Fairness, Accountability, and Transparency in Machine Learning Workshop (FAT/ML)*, 2017, pp. 1–5.
- [17] E. Dai and S. Wang, “Say no to the discrimination: Learning fair graph neural networks with limited sensitive attribute information,” in *Proc. ACM Int. Conf. on Web Search and Data Mining (WSDM)*, March 2021, pp. 680–688.
- [18] A. Bose and W. Hamilton, “Compositional fairness constraints for graph embeddings,” in *Int. Conf. on Machine Learning*. PMLR, 2019, pp. 715–724.
- [19] M. Buyl and T. D. Bie, “The KL-divergence between a graph model and its fair I-projection as a fairness regularizer,” in *Joint European Conf. on Machine Learning and Knowledge Discovery in Databases*. Springer, 2021, pp. 351–366.
- [20] O. D. Kose and Y. Shen, “Fairnorm: Fair and fast graph neural network training,” *arXiv preprint arXiv:2205.09977*, 2022.
- [21] —, “Fair contrastive learning on graphs,” *IEEE Transactions on Signal and Processing over Networks*, vol. 8, pp. 475–488, 2022.
- [22] I. Spinelli, S. Scardapane, A. Hussain, and A. Uncini, “Fairdrop: Biased edge dropout for enhancing fairness in graph representation learning,” *IEEE Transactions on Artificial Intelligence*, vol. 3, no. 3, pp. 344–354, 2021.
- [23] Y. Dong, N. Liu, B. Jalaian, and J. Li, “EDITS: Modeling and mitigating data bias for graph neural networks,” in *Proc. ACM Web Conference*, 2022, pp. 1259–1269.
- [24] O. D. Kose and Y. Shen, “Fairness-aware selective sampling on attributed graphs,” in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, 2022, pp. 5682–5686.
- [25] D. K. Hammond, P. Vandergheynst, and R. Gribonval, “Wavelets on graphs via spectral graph theory,” *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129–150, 2011.
- [26] F. R. Chung and F. C. Graham, *Spectral graph theory*. American Mathematical Soc., 1997, vol. 92.
- [27] X. Zhu, Z. Ghahramani, and J. D. Lafferty, “Semi-supervised learning using Gaussian fields and harmonic functions,” in *Int. Conf. on Machine Learning*. PMLR, 2003, pp. 912–919.
- [28] M. Belkin and P. Niyogi, “Semi-supervised learning on Riemannian manifolds,” *Machine Learning*, vol. 56, no. 1, pp. 209–239, 2004.
- [29] F. Zhang and E. R. Hancock, “Graph spectral image smoothing using the heat kernel,” *Pattern Recognition*, vol. 41, no. 11, pp. 3328–3342, 2008.
- [30] D. I. Shuman, P. Vandergheynst, and P. Frossard, “Chebyshev polynomial approximation for distributed signal processing,” in *Int. Conf. on Distributed Computing in Sensor Systems and Workshops (DCOSS)*. IEEE, 2011, pp. 1–8.
- [31] B. Hofstra, R. Corten, F. Van Tubergen, and N. B. Ellison, “Sources of segregation in social networks: A novel approach using facebook,” *American Sociological Review*, vol. 82, no. 3, pp. 625–656, 2017.
- [32] S. Hajian and J. Domingo-Ferrer, “A methodology for direct and indirect discrimination prevention in data mining,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 7, pp. 1445–1459, July 2013.
- [33] L. Takac and M. Zabovsky, “Data analysis in public social networks,” in *Int. Scientific Conf. and Int. Workshop. 'Present Day Trends of Innovations'*, vol. 1, no. 6, May 2012.