



Delft University of Technology

Online Filtering over Expanding Graphs

Das, Bishwadeep; Isufi, Elvin

DOI

[10.1109/IEEECONF56349.2022.10052045](https://doi.org/10.1109/IEEECONF56349.2022.10052045)

Publication date

2022

Document Version

Final published version

Published in

56th Asilomar Conference on Signals, Systems and Computers, ACSSC 2022

Citation (APA)

Das, B., & Isufi, E. (2022). Online Filtering over Expanding Graphs. In M. B. Matthews (Ed.), *56th Asilomar Conference on Signals, Systems and Computers, ACSSC 2022* (pp. 43-47). (Conference Record - Asilomar Conference on Signals, Systems and Computers; Vol. 2022-October). IEEE.
<https://doi.org/10.1109/IEEECONF56349.2022.10052045>

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.

Online Filtering over Expanding Graphs

Bishwadeep Das
Multimedia Computing Group
Delft Institute of Technology
Email: b.das@tudelft.nl

Elvin Isufi
Multimedia Computing Group
Delft Institute of Technology
Email: e.isufi-1@tudelft.nl

Abstract—Data processing tasks over graphs couple the data residing over the nodes with the topology through graph signal processing tools. Graph filters are one such prominent tool, having been used in applications such as denoising, interpolation, and classification. However, they are mainly used on fixed graphs although many networks grow in practice, with nodes continually attaching to the topology. Re-training the filter every time a new node attaches is computationally demanding; hence an online learning solution that adapts to the evolving graph is needed. We propose an online update of the filter, based on the principles of online machine learning. To update the filter, we perform online gradient descent, which has a provable regret bound with respect to the filter computed offline. We show the performance of our method for signal interpolation at the incoming nodes. Numerical results on synthetic and graph-based recommender systems show that the proposed approach compares well to the offline baseline filter while outperforming competitive approaches. These findings lay the foundation for efficient filtering over expanding graphs.

I. INTRODUCTION

Graph data-processing involves manipulating nodal information while obeying the underlying topology [1]. Most works towards this end consider a topology with a fixed number of nodes but several applications involve expanding graphs with a growing number of nodes [2], [3]. One such application is graph-based collaborative filtering with cold starters [4], where we first incorporate a new user into the existing user-user graph and then process the data over this new topology to predict the ratings. Another application is predicting physical phenomena such as temperature or pressure at a new sensor joining an existing sensor network.

A central approach for solving these tasks is to use graph filters [5]–[7]. These are flexible, localized, parametric operators that combine neighboring node signals with a linear complexity in the number of nodes. Designing such filters requires knowledge of the full fixed topology but in settings where a stream of new nodes attach to the existing graph, i.e., the number of nodes keep growing, these methods require a re-training for each updated topology. To avoid this, we develop a filter update rule inspired by online machine learning that operates on the fly as new nodes attach to the graph [8]–[10].

Pioneering works uniting online machine learning with graphs [11]–[13] have been focusing on online node classification on a static graph. For a dynamic setting, the two most relevant works are [14] which uses random kernels to perform online learning on graphs using only the attachment pattern of the incoming node and [15], which uses the inductive transference of graph filters. They classify a new node by

extending the kernel feature similarities between new and existing nodes to the signal similarities at these nodes via a smoothing matrix. In [14] the kernels are local, i.e., they rely on the location of the immediate neighbors and the predicted output does not rely on the graph data (e.g., signal values). Differently, graph filters rely on multi-resolution neighboring information and utilise the data over the topology to make predictions. This allows them to account for the influence of data over extended neighborhoods. In [15], the existing nodes influence the incoming node through weights learned from kernel feature similarities to perform inference. Instead, we want to adapt the learning directly to the expanding graph. Differently, the work in [16] estimates the attachment and uses it for subsequent classification. The works in [17], [18] deal with a stochastic attachment model for filter interpolation or a filter bank, but they are limited only to one incoming node and not a stream of nodes. Other works focus on mapping features to graph signals [19], node classification [20], and dynamic embedding [21]. However, none of these works discuss graph filtering on continuously expanding graphs as a principled way to process graph data.

We focus on online graph filter learning by accounting for the expanding graph topology. This yields filters that are localized over the growing graphs, inheriting the transference properties, and enjoy interpretability in terms of their frequency response. Our contributions are: 1) We propose a framework for the online learning of graph filters over continually expanding graphs; 2) We evaluate the learning for signal interpolation at each incoming node; 3) We corroborate the proposed model and provide interpretability via spectral analysis with the offline solution, while also comparing with the related literature for synthetic and real data.

The rest of this paper is as follows. Section 2 contains the problem formulation. Section 3 contains the online learning method. Section 4 contrasts our approach with others on synthetic and real data and Section 5 concludes the paper.

II. FILTERING OVER EXPANDING GRAPHS

Consider a starting graph $\mathcal{G}_0 = (\mathcal{V}_0, \mathcal{E}_0)$ of N_0 nodes, comprising a node set $\mathcal{V}_0 = \{v_1, \dots, v_{N_0}\}$ and an edge set \mathcal{E}_0 . Let $\mathbf{A}_0 \in \mathbb{R}^{N_0 \times N_0}$ be its adjacency matrix and $\mathbf{x}_0 \in \mathbb{R}^{N_0}$ the respective graph signal mapping each vertex in \mathcal{V}_0 to \mathbb{R} . Let $\{v_1, \dots, v_T\}$ be a sequence of T incoming nodes where v_i arrives at time $t=i$, connects to the existing graph and stays attached for $t \geq i$. Under this rule, graph $\mathcal{G}_t = \{\mathcal{V}_t, \mathcal{E}_t\}$ at time

$t < T$ has $N_t = N_0 + t$ nodes, adjacency matrix $\mathbf{A}_t \in \mathbb{R}^{N_t \times N_t}$, and graph signal $\mathbf{x}_t \in \mathbb{R}^{N_t}$. Node v_{t+1} with signal x_{t+1} attaches to \mathcal{G}_t with attachment vector $\mathbf{a}_{t+1} = [a_1, \dots, a_{N_t}]^\top$, where each non-zero element of \mathbf{a}_{t+1} implies an edge between v_{t+1} and $v \in \mathcal{V}_t$. We consider each edge is directed from $v \in \mathcal{V}_t$ to v_{t+1} . Such a setting is suited to inference tasks where existing nodes exert their influence over incoming ones, and not the other way around. The adjacency matrix $\mathbf{A}_{t+1} \in \mathbb{R}^{N_{t+1} \times N_{t+1}}$ and graph signal $\mathbf{x}_{t+1} \in \mathbb{R}^{N_{t+1}}$ have the respective¹ forms

$$\mathbf{A}_{t+1} = \begin{bmatrix} \mathbf{A}_t & \mathbf{0} \\ \mathbf{a}_{t+1}^\top & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{x}_{t+1} = \begin{bmatrix} \mathbf{x}_t \\ x_{t+1} \end{bmatrix}. \quad (1)$$

To process signal \mathbf{x}_{t+1} , we employ graph filters, which process graph data locally by combining successive shifts of an input signal over the topology. The output \mathbf{y}_{t+1} of a filter of order K on \mathbf{x}_{t+1} with shift operator \mathbf{A}_{t+1} is

$$\mathbf{y}_{t+1} = \sum_{k=0}^K h_{t+1}^k \mathbf{A}_{t+1}^k \mathbf{x}_{t+1} \quad (2)$$

where h_{t+1}^k is the weight given to the k th shift $\mathbf{A}_{t+1}^k \mathbf{x}_{t+1}$. The vector $\mathbf{h}_{t+1} = [h_{t+1}^0, \dots, h_{t+1}^{K+1}]^\top$ contains the filter parameters at time $t+1$. By substituting the k th adjacency matrix power

$$\mathbf{A}_{t+1}^k = \begin{bmatrix} \mathbf{A}_t^k & \mathbf{0} \\ \mathbf{a}_{t+1}^\top \mathbf{A}_t^{k-1} & 0 \end{bmatrix} \quad (3)$$

in (2), we can write the filter output

$$\mathbf{y}_{t+1} = \begin{bmatrix} \sum_{k=0}^K h_{t+1}^k \mathbf{A}_t^k \mathbf{x}_t \\ \mathbf{a}_{t+1}^\top \sum_{k=1}^K h_{t+1}^k \mathbf{A}_t^{k-1} \mathbf{x}_t \end{bmatrix}. \quad (4)$$

The output at v_{t+1} is $[\mathbf{y}_{t+1}]_{N_{t+1}} = \mathbf{a}_{t+1}^\top \sum_{k=1}^K h_{t+1}^k \mathbf{A}_t^{k-1} \mathbf{x}_t$.

We learn the filter coefficients from a training set $\mathcal{T} = \{v_t, x_t, \mathbf{a}_t\}_{t=1:T}$ where each element of \mathcal{T} comprises an incoming node v_t , its signal x_t , and its attachment vector \mathbf{a}_t . We measure our performance at time $t+1$ relative to the signal x_{t+1} , revealed online, through a loss function $l_{t+1}([\mathbf{y}_{t+1}]_{N_{t+1}}, x_{t+1})$. Since the output is a function of the filter parameter, we consider the loss as $l_{t+1}(\mathbf{h}_{t+1}, x_{t+1})$. The cumulative loss over the incoming node sequence is $L_T = \sum_{t=1}^T l_t(\mathbf{h}_t, x_t)$, where \mathbf{h}_t is the filter generated by the online algorithm at time t .

Problem statement. Given $(\mathcal{G}_0, \mathbf{A}_0)$, signal \mathbf{x}_0 , and the training set \mathcal{T} , our goal is to predict a sequence of fixed-order FIR graph filters $\{\mathbf{h}_t\}_{t=1:T}$ generated online for minimizing a sequence of loss functions l_t at the incoming nodes.

We solve this with the tools of online machine learning, specifically online convex optimization to update the filter coefficients via online gradient descent (OGD) and characterize the regret of the approach relative to the optimal batch filter.

¹The undirected attachment case extends trivially to this update. As we focus only on inference at the incoming node, we assume the edges should not capture its influence over the existing vertices.

III. ONLINE FILTER LEARNING

Online filter learning implies updating parameters \mathbf{h}_t based on \mathbf{h}_{t-1} and the new node information (i.e., connections \mathbf{a}_t and signal \mathbf{x}_t). Focusing on signal reconstruction at the incoming node v_t , the loss function takes the form

$$l_t(\mathbf{h}_t, x_t) = \frac{1}{2}(\mathbf{a}_t^\top \mathbf{A}_{x,t-1} \mathbf{h}_t - x_t)^2 + \mu \|\mathbf{h}_t\|_2^2 \quad (5)$$

where $\mathbf{A}_{x,t-1} = [\mathbf{0}, \mathbf{x}_{t-1}, \mathbf{A}_t \mathbf{x}_{t-1}, \dots, \mathbf{A}_t^K \mathbf{x}_{t-1}]$ contains the shifted signals. The term $\mu \|\mathbf{h}_t\|_2^2$ acts as a regulariser to ensure there is no overfitting and $\mu > 0$ is the respective weight. We perform the online update via gradient descent, i.e.,

$$\mathbf{h}_{t+1} = \mathbf{h}_t - \eta_t \nabla_t \quad (6)$$

where the gradient ∇_t has the expression

$$\nabla_t = (\mathbf{a}_t^\top \mathbf{A}_{x,t-1} \mathbf{h}_t - x_t) \mathbf{A}_{x,t-1}^\top \mathbf{a}_t + 2\mu \mathbf{h}_t \quad (7)$$

and $\eta_t > 0$ is the learning rate at time t . The loss is differentiable, convex in \mathbf{h}_t , and Lipschitz. The latter can be seen from the gradient expression, whose first term $\frac{1}{2}(\mathbf{a}_t^\top \mathbf{A}_{x,t-1} \mathbf{h}_t - x_t) \mathbf{A}_{x,t-1}^\top \mathbf{a}_t$ is bounded as long as the graph signals have finite energy, the adjacency matrix is bounded (e.g., normalized adjacency matrix [22]), and the filter \mathbf{h} belongs to a closed set \mathcal{H} such that $\|\mathbf{h}\|_2 \leq L_h$. Let the first term be bounded by a constant C . This makes the function Lipschitz with a factor of $(C + 2\mu L_h)$, i.e., $\|\nabla_t\|_2 \leq (C + 2\mu L_h)$ for all t . We denote the constant $(C + 2\mu L_h)$ as L .

To measure the performance of the online algorithm w.r.t. the offline batch filter \mathbf{u} computed upon observing the entire sequence, we pursue a standard regret analysis. We define the static regret of our algorithm w.r.t. the batch filter \mathbf{u} over a T -length sequence as

$$R_T(\mathbf{u}) = \sum_{t=1}^T l_t(\mathbf{h}_t, x_t) - l_t(\mathbf{u}, x_t). \quad (8)$$

The regret measures how much worse off we are relative to an offline filter \mathbf{u} . We usually want the regret to grow at a sub-linear rate with time [8]. This implies $\frac{1}{T} \lim_{T \rightarrow \infty} R_T(\mathbf{u}) = 0$, i.e., the average performance of the algorithm asymptotically approaches the average performance of the offline filter \mathbf{u} . We can bound the regret of our algorithm $R_T(\mathbf{u})$ as follows.

Proposition 1. Given the sequence of L Lipschitz losses $l_t(\mathbf{h}_t, x_t)$ [cf. (5)] with $t = 1, \dots, T$ and a fixed gradient step size η [cf. (6)], the static regret $R_T(\mathbf{u})$ [cf. (8)] for the online algorithm generating filters $\{\mathbf{h}_t\} \in \mathcal{H}$ relative to any other filter $\mathbf{u} \in \mathcal{H}$ is upper bounded as

$$R_T(\mathbf{u}) \leq \frac{\|\mathbf{u}\|_2^2}{2\eta} + \frac{\eta}{2} L^2 T \quad (9)$$

Proof. The proof is similar to the regret for OGD [Thm. 2.13, [10]]. \square

The bound holds with strict equality when \mathbf{u} is the optimal batch filter. It is also a function of the step size η and attains a minimum value of $\|\mathbf{u}\|_2 L \sqrt{T}$ for $\eta = \frac{\|\mathbf{u}\|_2}{L \sqrt{T}}$. However, such a step size is unattainable in the online setting as it

depends on the batch solution \mathbf{u} . As we focus more on filters and expanding graphs, we do not optimize the regret bound. However, there are approaches allied with OGD which do the same. Interested readers can refer to [10].

Computational complexity. Let us consider graph \mathcal{G}_0 having M_0 edges and that each incoming node forms P edges w.l.o.g. Graph \mathcal{G}_t has $M_0 + Pt$ edges. The complexity of obtaining the filter output at time t is $\mathcal{O}((M_0 + Pt)K)$, where K is the filter order [7]. The gradient and filter update step both share the same complexity. Thus at time t the online algorithm has a complexity of order $\mathcal{O}((M_0 + Pt)K)$. The total complexity over a sequence of length T summed over $t = 1, \dots, T$ is of the order $\mathcal{O}((M_0T + PT^2)K)$.

Remark 1. We discuss the static regret concerning a fixed offline filter \mathbf{u} , available once all the nodes in the sequence are observed. For our attachment method, this is the same as training the filter at all incoming nodes in the final graph \mathcal{G}_T . One can alternatively consider the dynamic regret, where at each time instant we would have an optimal \mathbf{u}_t . Here, the online algorithm can be evaluated on how efficiently it tracks the optimal solution at each time step [23]. \square

IV. NUMERICAL RESULTS

We corroborate the proposed approach and compare it with both heuristic baselines and state-of-the-art alternatives across synthetic and real data. These comparisons are focused to answer the following research questions:

RQ1. *How far off from the offline/ batch solution is the proposed method?*

We observe how much we lose in performance over the filter solved with complete knowledge of \mathcal{T} . We compare with:

1. **Batch Solution (BS):** This is the solution given by

$$\mathbf{u} = \underset{\mathbf{h} \in \mathbb{R}^{K+1}}{\operatorname{argmin}} \sum_{t=1}^T (\mathbf{a}_t^\top \mathbf{A}_{x,t-1} \mathbf{h} - x_t)^2 + \mu_b \|\mathbf{h}\|_2^2 \quad (10)$$

which has a closed form least squares expression for $\mu_b > 0$.

RQ2. *How does online filtering for interpolation compare to online random kernel-based prediction?*

We compare with the state-of-the-art alternative:

2. **Online Kernel Learning (OKL)** [14]: We consider a Gaussian kernel with variance $\sigma^2 = \{1, 10\}$.

RQ3 *How does online filtering for interpolation compare to pre-trained / inductive filtering?*

Following the inductive bias capabilities of graph filters [24], we want to contrast online learning to a filter pre-trained over \mathcal{G}_0 and deployed on each \mathcal{G}_t . Depending on the availability of signals at node v_t , we considered the alternatives

3. **Inductive Transfer without features (IT1):** We learn a filter on \mathcal{G}_0 using $\mathbf{A}_0, \mathbf{x}_0$ and deploy it on $\mathcal{G}_1, \dots, \mathcal{G}_T$ to predict the signals at v_1, \dots, v_T , respectively without retraining. We train the filter through a regularized least square problem [5].
4. **Inductive Transfer with features (IT2)** [15]: First, we learn the filter in **IT1**. The output at v_t is predicted

as $y_t = \mathbf{m}^\top \mathbf{A}_{x,t-1} \mathbf{h}$, where \mathbf{h} is the pre-trained filter, $\mathbf{m} \in \mathbb{R}^{N_t}$ is a smoothing vector which contains the relational dependencies between the signals at the existing and incoming nodes. This vector is obtained as $\mathbf{k}'(\mathbf{K} + \lambda \mathbf{I})^{-1}$, where \mathbf{k} contains the kernel similarities between the features of v_t and \mathcal{V}_{t-1} , and \mathbf{K} the same between the nodes in \mathcal{V}_{t-1} [15]. We report the better of the performance between a linear and a Gaussian kernel with $\sigma^2 = 1, 5, 10$.

We also compare with the standard baseline:

5. **Weighted K-Nearest Neighbor Prediction (WKNN):** The predicted signal at v_t is the weighted mean of the signal at the nodes to which v_t attaches, with the corresponding edge weights normalized to have unit sum.

We evaluate all approaches via the root Normalized Mean Square Error (rNMSE) over the sequence

$$\text{rNMSE} = \sqrt{\frac{\sum_{t=1}^T (\hat{x}_t - x_t)^2}{\sum_{t=1}^T x_t^2}} \quad (11)$$

where \hat{x}_t and x_t are the predicted and true signal at v_t , respectively. All filters have order $K = 5$. Both the online methods (proposed and **OKL**) have parameters of equal dimension and have the same hyper-parameters η and μ selected via grid search over $[10^{-6}, 10]$. For **BS**, **IT1**, and **IT2**, the regularisation hyper-parameter is chosen over $[10^{-4}, 10]$.

A. Synthetic

Experimental setup. We consider an Erdos-Renyi graph $\mathcal{G}_0(N, p)$ with $N = 100$ and $p = 0.2$ [2]. We normalize the adjacency matrix by the largest absolute eigenvalue to stabilise the filter output as in [5]. We construct the training set \mathcal{T} with each v_t forming $P = 5$ edges *uniformly-at-random* relative to the existing graph \mathcal{G}_t . Each edge has a weight equal to the median of the edge weights in \mathcal{E}_0 . To generate x_t at v_t , we consider three options:

- 1) **Kernel**, where all x_t s are generated following [14], which favors kernel solutions.
- 2) **Filter**, where x_t is generated using a pre-trained filter with $K = 5$ on \mathcal{G}_0 using data at 80 percent of its nodes, which favors filtering solutions.
- 3) **WMean**, where x_t is the **WKNN** prediction, which is a neutral choice.

Our motivation here is to see how each of the methods compares in settings that may or may not be suitable for it. We iterate all methods over 20 initial graphs $\mathcal{G}_0(N, p)$, with $T = 1000$ for each realization. Since v_t has no features, we consider **IT1** for **RQ3**, whose performance is averaged over 50 sets of sampled signals for pre-training.

Results. Table I showcases the rNMSE along with the standard deviation for all approaches averaged across all graphs in all three scenarios. Concerning **RQ1**, the online solution underperforms **BS** as expected with the gap being the most for **Filter** data. This is because the setting is biased towards filtering. For **RQ2**, the proposed outperforms **OKL** except in

TABLE I
rNMSE AND STANDARD DEVIATION OF ALL APPROACHES AVERAGED OVER TEN SYNTHETIC DATA-SETS FOR EACH SIGNAL GENERATION SCENARIO.

Method	Filter					WMean					Kernel				
	Prop.	BS	OKL	IT1	WKNN	Prop.	BS	OKL	IT1	WKNN	Prop.	BS	OKL	IT1	WKNN
rNMSE	0.09	3×10^{-4}	0.82	0.26	0.27	0.09	0.034	0.71	0.27	0	0.50	0.39	0.39	4.79	4.03
sdev	0.03	1×10^{-4}	0.19	0.06	0.13	0.02	0.007	0.2	0.07	0	0.05	0.04	0.04	0.38	0.47

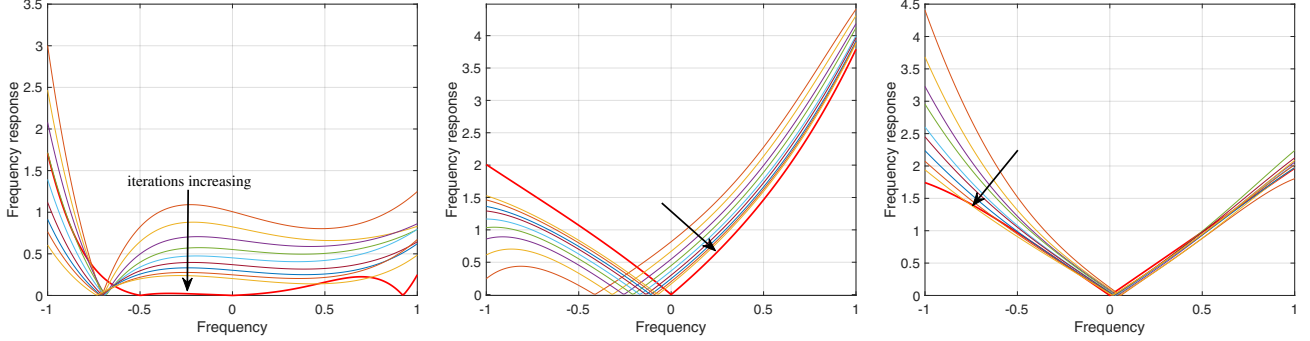


Fig. 1. Frequency response of the online filter (colored lines) every 100 iterations compared with that of batch solution **BS** (in bold red) for (left) *Kernel*, (centre) *Filter*, and (right) *WMean* generated data.

the case for *Kernel* data, which favours **OKL** but still the filtering solution is close to it. This indicates the advantage of graph filters that incorporate higher-order node interactions through successive shifts, as opposed to the kernel model which relies only on the incoming attachment. For **RQ3**, the proposed method outperforms **IT1** as online approaches adapt to the incoming data and topology change better than pre-trained filters which are biased to one signal and topology. Note that the zero error obtained by **WKNN** for *WMean* data is due to them being the same method, and is of no importance.

TABLE II
NORMALIZED REGRET FOR ALL THREE DATA GENERATIONS

Method	Kernel	Filter	wMean	MovieLens
$R_T(\mathbf{u})/T$	0.29	0.0021	0.0043	0.12

Fig. 1 showcases how the online filter processes data over the existing graph in the frequency domain. We plot the frequency response of the online filter every 100 iterations and that of the batch over the frequencies $[-1, 1]$ following the frequency response interpretation w.r.t. the adjacency matrix [22]. Frequencies approaching -1 and 1 are the high and low frequencies, respectively. The batch frequency response is in bold red. For *Filter* and *wMean* data, the online frequency response tends towards the batch response. Table II contains the normalized regret and we see that for *Filter* and *wMean*, the values are quite low for $T = 1000$, as suggested by the respective frequency plots. For *Kernel* data, the online filter approaches the batch response for the middle and low frequencies while attenuating the high frequencies more. This could be because of an insufficient number of incoming nodes. This gap is also reflected in the regret value in Table II.

B. Collaborative filtering

In this section, we predict ratings for a sequence of incoming users using collaborative filtering. We considered the MovieLens100K data-set, comprising 100,000 ratings provided by 943 users over 1152 items [25] with at least ten ratings. We use graph collaborative filters for this task [26].

Experimental setup. We consider $N_0 = 500$ starting users at random and treat the remaining 443 as online. We build the 15-NN directed user-user graph \mathcal{G}_0 with normalized adjacency matrix \mathbf{A}_0 for the starting users using the cosine similarity of their rating vectors². For each new user v_t , we use 50 percent of the ratings to build \mathbf{a}_t , which has 15 edges connected to nearest neighbors via cosine similarity. We predict the remaining ratings. We evaluated all methods over 20 realizations of the basic setup. Since each incoming user has available features, i.e., a subset of ratings, we compare also with **IT2**. We train this filter over the same number of existing ratings as the number of online ratings.

TABLE III
rNMSE OF ALL APPROACHES AVERAGED OVER 20 REALIZATIONS FOR USER RATING PREDICTION IN MOVIELENS100K

Method	Prop.	BS	OKL	IT2	WKNN
rNMSE	0.56	0.55	0.69	0.6	0.61
sdev	0.02	0.006	0.006	0.03	0.008

Results. Table III showcases the rNMSE of all approaches for this experiment. Concerning **RQ1**, the proposed is close to **BS**, while regarding **RQ2** and **RQ3**, the proposed does better than **OKL** and **IT2**, by improving the rNMSE by up to 18.8% and 8.2%, respectively.

²All online methods performed better at this range of nearest neighbors.

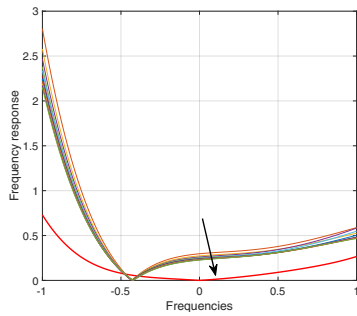


Fig. 2. Frequency response of the online filter every 2000 updates compared with that of **BS** (bold red) for Movielens-100K

Fig. 2 plots the frequency response of the online filter every 2000 steps over roughly 23,000 samples in one run of an experiment. It also contains the response of **BS**. Despite achieving an rNMSE close to **BS**, the frequency responses are not close, with a normalized regret of 0.12 [cf. Table II]. The online filter has a higher amplitude for higher frequencies than the batch filter. For recommender systems, this implies a greater diversity in the recommendation (i.e., prioritizing good ratings to diverse items and not only to a few popular ones) [27]. To test this, we evaluate the average ten item aggregated diversity AD@10 for both approaches. This shows what fraction of the total items appear in a list containing the top ten recommendation over all users, with a higher value showing a better diversity. We evaluate this online for the proposed and with the batch filter for all incoming users. The proposed and **BS** have average AD@10 as 0.17 ± 0.01 and 0.12 ± 0.01 , respectively. The proposed achieves a higher value, thus showing that the online filter promotes while retaining a similar rNMSE.

V. CONCLUSION

We proposed online graph filtering for expanding graphs over a sequence of incoming nodes. With each incoming node, the filter processes the existing signal to make a prediction. We update the filter parameters based on an online gradient descent update over the loss. Numerical results over synthetic and real data for signal interpolation at the incoming node show that the proposed approach is competitive w.r.t. the offline filter in the rNMSE while also exhibiting similar behaviour in the frequency domain, also providing more diverse recommendations. Results also show the improvement in rNMSE over online random kernels and inductive filtering. For future work, we will consider the scenario where the existing graph data obeys a time varying model and adapt the online algorithm to be graph-time aware.

REFERENCES

- [1] A. Ortega, P. Frossard, J. Kovačević, J. M. F. Moura, and P. Vandergheynst. Graph Signal Processing: Overview, Challenges, and Applications. *Proceedings of the IEEE*, 106(5):808–828, May 2018.
- [2] P. Erdos. On the evolution of random graphs. *Bulletin of the Institute of International Statistics*, 38:343–347, 1961.
- [3] A. L. Barabási and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439), October 1999. Publisher: American Association for the Advancement of Science.
- [4] S. Liu, I. Ounis, C. Macdonald, and Z. Meng. A heterogeneous graph neural model for cold-start recommendation. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 2029–2032, 2020.
- [5] A. Sandryhaila and J. M. F. Moura. Discrete signal processing on graphs. *IEEE transactions on signal processing*, 61(7):1644–1656, 2013.
- [6] S. K. Narang, A. Gadde, E. Sanou, and A. Ortega. Localized iterative methods for interpolation in graph structured data. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 491–494. IEEE, 2013.
- [7] M. Coutino, E. Isufi, and G. Leus. Advances in distributed graph filtering. *IEEE Transactions on Signal Processing*, 67(9):2320–2333, 2019.
- [8] S. Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- [9] Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- [10] Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- [11] M. Herbster, M. Pontil, and L. Wainer. Online learning over graphs. In *Proceedings of the 22nd international conference on Machine learning*, pages 305–312, 2005.
- [12] M. Herbster and M. Pontil. Prediction on a graph with a perceptron. In *NIPS*, volume 21, pages 577–584, 2006.
- [13] Q. Gu, C. Aggarwal, J. Liu, and J. Han. Selective sampling on graphs for classification. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 131–139, 2013.
- [14] Y. Shen, G. Leus, and G. B. Giannakis. Online Graph-Adaptive Learning With Scalability and Privacy. *IEEE Transactions on Signal Processing*, 67(9):2471–2483, May 2019.
- [15] S. Chen, F. Cerda, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević. Semi-supervised multiresolution classification using adaptive graph filtering with application to indirect bridge structural health monitoring. *IEEE Transactions on Signal Processing*, 62(11):2879–2893, 2014.
- [16] F. Dornaika, R. Dahbi, A. Bosaghzadeh, and Y. Ruichek. Efficient dynamic graph construction for inductive semi-supervised learning. *Neural Networks*, 94:192–203, 2017.
- [17] B. Das and E. Isufi. Learning expanding graphs for signal interpolation. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2022.
- [18] B. Das and E. Isufi. Graph filtering over expanding graphs. In *IEEE Data Science Learning Workshop Processing (IEEE-DSLW)*, May 2022.
- [19] A. Venkitaraman, S. Chatterjee, and B. Wahlberg. Recursive Prediction of Graph Signals With Incoming Nodes. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5565–5569, May 2020.
- [20] L. Jian, J. Li, and H. Liu. Toward online node classification on streaming networks. *Data Mining and Knowledge Discovery*, 32(1):231–257, 2018.
- [21] J. Li, H. Dani, X. Hu, J. Tang, Y. Chang, and H. Liu. Attributed network embedding for learning in a dynamic environment. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 387–396, 2017.
- [22] A. Sandryhaila and J. M. F. Moura. Discrete signal processing on graphs: Frequency analysis. *IEEE Transactions on Signal Processing*, 62(12):3042–3054, 2014.
- [23] A. Simonetto, A. Mokhtari, A. Koppel, G. Leus, and A. Ribeiro. A class of prediction-correction methods for time-varying convex optimization. *IEEE Transactions on Signal Processing*, 64(17):4576–4591, 2016.
- [24] 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*, 37(6):128–138, 2020.
- [25] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *Acem transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
- [26] W. Huang, A. G. Marques, and A. R. Ribeiro. Rating Prediction via Graph Signal Processing. *IEEE Transactions on Signal Processing*, 66(19):5066–5081, October 2018.
- [27] E. Isufi, M. Pocchiari, and A. Hanjalic. Accuracy-diversity trade-off in recommender systems via graph convolutions. *Information Processing & Management*, 58(2):102459, 2021.