

AutoGF: Runtime Graph Filter Tuning for Community Node Ranking

Emmanouil Krasanakis, Symeon Papadopoulos, and Ioannis Kompatsiaris

Centre for Research and Technology Hellas, Information Technologies Institute
{maniospas,papadop,ikom}@iti.gr

Abstract. A recurring graph analysis task is to rank nodes based on their relevance to overlapping communities of shared metadata attributes (e.g. the interests of social network users). To achieve this, approaches often start with a few example community members and employ graph filters that rank nodes based on their structural proximity to the examples. Choosing between well-known filters typically involves experiments on existing graphs, but their efficacy is known to depend on the structural relations between community members. Therefore, we argue that employed filters should be determined not during algorithm design but at runtime, upon receiving specific graphs and example nodes to process. To do this, we split example nodes into training and validation sets and either perform supervised selection between well-known filters, or account for granular graph dynamics by tuning parameters of the generalized graph filter form with a novel optimization algorithm. Experiments on 27 community node ranking tasks across three real-world networks of various sizes reveal that runtime algorithm selection selects near-best AUC and NDCG among a list of 8 popular alternatives, and that parameter tuning yields similar or improved results in all cases.

Keywords: node ranking, graph signal processing, parameter tuning

1 Introduction

When graph nodes are attributed (e.g. they are social network users and attributes are their areas of interest), they can be organized into communities of shared metadata attributes [26]. By definition, these communities are not tied to specific high-level structural characteristics, such as strong connectivity between nodes. Still, it is commonly accepted that attributes could correlate to low-level dynamics leading to the creation of edges, in which case graph structure can help predict metadata. For example, nodes of social network graphs often exhibit homophilous behavior [23], a term describing their tendency to form edges with others of similar attributes. Then, tightly knit structural communities become good predictors of parts of -but not of whole- metadata communities [35].

A recurring graph analysis task, which we tackle in this work, is to rank nodes based on their relevance to communities sharing metadata attributes of interest [16, 19, 29, 32, 33]. Ranking provides greater granularity than clear-cut

predictions, for example in the scope of recommending more community members. It also respects overlaps and fuzzy boundaries between communities [21]. Furthermore, node relevance scores obtained during ranking are often the core of more sophisticated systems, such as graph neural networks for classification after initial neural estimations [13, 15] and post-processing strategies that threshold transformations of scores to predict community membership [3].

A popular use case for community node ranking, which we also follow, is to start with a few known community members serving as examples, and inferring the relatedness of all nodes to respective communities based on their structural proximity to the examples [34, 37, 19]. This task is performed independently for one or more communities. Assumptions about what constitutes proximity have coalesced under the field of graph signal processing [section 2], where they are modeled with ad-hoc graph filters and controlled by a small number of parameters [25, 8].¹ Different filters and parameters match different types of communities. For example, filter efficacy could depend on the number of community members [1, 12, 19]. As a result, deployed filters may work well in certain graphs but not necessarily in others. By extension, running filter-based tools ‘off-the-shelf’ in deployed systems risks producing node ranks of lesser quality.

In this work, we address the above issue by exploiting autotune principles [17] for runtime selection of graph filters. We explore two strategies: a) choosing the best among a list of promising filters, and b) tuning the parameters of a generalized filter form. For the second strategy, we also introduce a novel tuning algorithm that keeps examining a wide search breadth in the solution space but converges within a bounded number of filter runs. The effectiveness of our approach is corroborated on 27 community node ranking tasks across 3 real-world graphs of different domains. Results indicate that neither strategy falls significantly behind best-performing ad-hoc filters when optimizing popular node rank quality measures. Furthermore, parameter tuning frequently captures structural proximity better than ad-hoc assumptions and improves rank quality.

This paper is organized as follows. In section 2 we present graph filters as an approach for ranking nodes with respect to metadata communities, alongside a generalized literature filter form. In section 3, we describe our runtime filter selection approach and its implementation choices. We also present a novel algorithm for tuning parameters of generalized graph filters. In sections 4 and 5 we evaluate our approach in real-world data and discuss practical applicability and potential risks. Finally, in section 6 we summarize our findings and present promising research directions.

2 Background

Graph edges are often represented by adjacency matrices A with elements $A[u, v] = \{1 \text{ if edge } (u, v) \text{ exists}, 0 \text{ otherwise}\}$. These are symmetrically normalized by weigh-

¹ Most non-filter node ranking algorithms, such as k-shell decomposition and variations [36], blindly rank the importance of nodes within graph structures and can not personalize ranks in terms of importance to specific communities.

ing edges to mitigate the importance of highly connected nodes per:

$$W = D^{-1/2}AD^{-1/2}$$

where D with elements $D[u, v] = \{\sum_{v'} A[u, v'] \text{ if } u = v, 0 \text{ otherwise}\}$ are diagonal matrices of node degrees. The graph's spectrum can be defined as the eigenvalues of the normalized adjacency matrix.² In detail, eigenvalue decomposition yields $W = UAU^{-1}$, where A are diagonal matrices of eigenvalues $A = \text{diag}([\lambda_1, \lambda_2, \dots, \lambda_n])$ and U are orthogonal matrices whose columns hold the corresponding eigenvectors. For connected graphs, eigenvalues of the normalized adjacency matrix are real-valued and reside in the unit range $\lambda_i \in [-1, 1]$.

Graph signal processing [27, 24, 28] manipulates signals p whose elements $p[u]$ correspond to values stored at nodes u . To do this, it defines their graph Fourier transform as $\mathcal{F}\{p\} = U^{-1}p$ and its inverse as $\mathcal{F}^{-1}\{\mathcal{F}\{p\}\} = U\mathcal{F}\{p\}$. Then, it observes that $W^n = UA^nU^{-1} \Rightarrow H(W) = UH(A)U^{-1}$ for function forms $H(\cdot)$ whose Taylor expansions exist around zero, and defines filters $H_{\mathcal{F}} = [H(\lambda_1), H(\lambda_2), \dots, H(\lambda_n)]$ in the Fourier space, whose parameters arise through transformations $H(\lambda_i)$ of eigenvalues λ_i . Graph filters can be applied on signals via an element-wise multiplication \odot on their Fourier transform $\mathcal{F}\{p\}$. The outcome of filtering in the node space becomes:

$$\mathcal{F}^{-1}(H_{\mathcal{F}} \odot \mathcal{F}\{p\}) = UH(A)U^{-1}p = H(W)p$$

During the above analysis, the function forms $H(\cdot)$ determining graph filters can be parameterized in terms of their Taylor coefficients h_0, h_1, \dots per:

$$H(W) = \sum_{k=0}^{\infty} h_k W^k$$

As $W^k p$ propagates graph signals p at k hops away through normalized adjacency matrices W , the above formula describes a weighted aggregation of multi-hop signal propagation. Filters matching different structural assumptions arise from different coefficients h_k . Two well-known filters are personalized PageRank [3, 4] and heat kernels [16]. These respectively adopt degrading hop weights $h_k = (1 - a)a^k$ and the kernel $h_k = e^{-t}t^k/k!$ for parameters $a \in [0, 1)$ and $t \in \{1, 2, 3, \dots\}$.

Given the above formulation, graph filtering can rank how nodes pertain to communities of interest [34, 37, 19]. Approaches start with signals p whose values capture whether nodes v belong to sets \mathcal{C} of known community members per:

$$p[v] = \{1 \text{ if } v \in \mathcal{C}, 0 \text{ otherwise}\}$$

Then, for graphs with normalized adjacency matrices W , graph filters $H(W)$ yield new signals $r = H(W)p$ with elements $r[u]$ corresponding to how proximate nodes u are to known member sets \mathcal{C} under some understanding of proximity. Finally, nodes are ranked by order of their proximity to known members.

² The graph's spectrum can also be defined as the eigenvalues $1 - \lambda_i$ of its normalized Laplacian $I - W$. This, too, can express filters as infinite-degree polynomials of W .

3 Tuning Graph Filters at Runtime

As previously mentioned, graph filters for community node ranking should ideally be selected at runtime, after graphs and example community members become known and therefore can be used to understand underlying structural features. We consider best-performing filters those with higher node rank quality, for instance measured with the area under curve of the receiver operating characteristics (AUC) [6] and the normalized discounted cumulative gain across all graph nodes (NDCG) [14]. Employed measures should coincide with practical objectives on unknown test data. For example, high AUC indicates higher ranks for community members than non-members, whereas high NDCG verifies the community membership of top-ranked nodes.

To optimize node rank quality at runtime, we follow an autotune paradigm that searches through the parameter space of black box algorithms to optimize validation objectives. Originally, the term was associated with specific approaches [17], but nowadays broadly describes automatic selection of machine learning model parameters. This comes at the expense of multiple algorithm runs, but there exist mature solutions for fast computation of graph filters [18].

Our approach starts with sets \mathcal{C} of known community members among graph nodes, which are organized into binary graph signals p per the formulation of section 2. We split known members into non-overlapping subsets $\mathcal{C}_{train}, \mathcal{C}_{valid} \subseteq \mathcal{C}$, which correspond to “training” graph signals p_{train} to be used as filter inputs, and desired output validation signals p_{valid} . We employ evaluation measures $\mathcal{M}(\cdot, \cdot)$, such as AUC or NDCG, that assess node rank quality via pairwise comparison between predictions and ground truth, and select filters with high $\mathcal{M}(r_{train}, p_{valid})$ for predicted ranking scores $r_{train} = H(W)p_{train}$. We avoid overfitting by computing measures only across non-training nodes. As long as graphs exhibit homogeneous correlations between communities and edges, filters maximizing validation evaluation are expected to also maximize $\mathcal{M}(r, p_{test})$ for $r = H(W)p$ on nodes other than known community members, where p_{test} are unknown ideal test labels. Our pipeline’s data flow is summarized in Figure 1.

We follow two strategies for graph filter selection by the autotune component of our approach. The first is to perform a simple selection among a list of popular filters, such as those we experiment with later on. The second is to start with the parameterized graph filter form presented in section 2 and tuning a vector of its parameters $h = [h_0, h_1, \dots, h_K]^T$ to optimize validation objectives. We explore only non-negative parameters to match the widespread literature practice of introducing only non-negative correlations between hops and high-quality node ranks. Then, without loss of generality, we tune all parameters in the range $[0, 1]$.

When tuning graph filter parameters on non-differentiable (potentially even non-convex) validation objectives, a first take is to adopt existing generic black box optimization algorithms [7, 10]. However, these do not guarantee convergence for all deployed system inputs. At the same time, adjusting one graph filter parameter to control the importance of propagating graph signals a fixed number of hops away could drastically affect the validity of other propagation weights. This hypothesis is also corroborated by experiments later on.

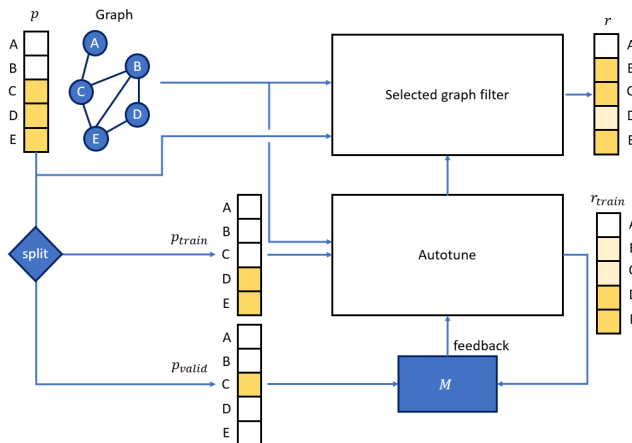


Fig. 1. Overview of graph filter autotuning under measures \mathcal{M} . Example nodes are split between training and validation graph signals, where the latter assume the role of ideal training outputs. Highlighted signal elements correspond to higher node values.

To address the above concerns, we propose an algorithm for graph filter parameter tuning that maintains a broad parameter search space while converging in finite time. This involves cycling through parameters, and progressively minimizing a loss function $\ell(h) = 1 - \mathcal{M}(H(W)p_{train}, p_{valid})$ by finding the best permutation around each parameter with coarse linear search. As tuning progresses, we shrink the search range, so that small permutations around ideal values are eventually found. Intuitively, this is equivalent to moving the center of the selected rectangle chosen for each parameter based on subsequent selections of other parameters. If shrinking is slow enough, by the time when parameter permutation breadths become small, potential combinations with drastically different permutations of other parameters have already been considered.

Conceptually, this procedure is a variation of divided rectangles (DIRECT) [9] that, instead of keeping many candidate rectangles to divide, keeps only one, though of larger width than the partition. This practice corresponds to the shrinking radius technique proposed for non-convex block coordinate optimization [22], although the two are not mathematically equivalent due to the finite sum of rectangle widths that limits the optimization within the hypercube of searched parameters instead of looking at an unconstrained range.

In detail, we start from the center of the parameter hypercube and cycle through parameters i . For each of those, we consider the range $\Delta h[i]$ in which to search for new solutions and partition it uniformly to $2P + 1$ candidate points, P of which examine higher parameter values and an equal number lower values. Values are snapped to the search bounds 0 or 1 if they subceed or exceed those respectively. Perturbations form a set H_{search} of potential parameter vectors, of which we select the one minimizing the loss. Finally, we contract the search range by division with constant $T > 1$ and move on to the next parameter.

Cycling through parameters stops when loss reduction becomes smaller than a tolerance ϵ across all parameters. This process is outlined in Algorithm 1.

Algorithm 1 Parameter tuning

Inputs: parameter loss $\ell(h)$, tolerance ϵ , line search partitions P , range shrinking T
Outputs: near-optimal vector of K parameters
 $h \leftarrow [0.5] \times K$, $\Delta h \leftarrow [0.5] \times K$, $err \leftarrow [\infty] \times K$, $i \leftarrow 0$
while $\max_i err[i] > \epsilon$ **do**
 $u_i \leftarrow$ unit vector with $u_i[j] = \{1 \text{ if } i = j, 0 \text{ otherwise}\}$
 $H_{search} \leftarrow \{max(0, min(1, h + u_i \cdot \Delta h[i] \cdot (p/P - 1))) \mid p = 0, 1, \dots, 2P\}$
 $err[i] \leftarrow \ell(h) - \min_{h \in H_{search}} \ell(h)$
 $h \leftarrow \arg \min_{h \in H_{search}} \ell(h)$
 $\Delta h[i] \leftarrow \Delta h[i]/T$
 $i \leftarrow (i + 1) \bmod K$
return h

If the objective $\ell(h)$ is Lipschitz continuous with Lipschitz constant $L < \infty$ (when the loss is differentiable, this means that $\sup \|\nabla \ell(h)\| \leq L$), it is easy to see that the division of the parameter permutation radius by T every K iterations lets the algorithm run in amortized time $O(K(\text{run } \ell(h)) \log_T \frac{L}{\epsilon})$. If graph nodes are fewer than edges (as happens for connected graphs), in which case the running time of $\ell(h)$ is not dominated by node validation. Using sparse matrix multiplication to iteratively compute Krylov space elements $\{W^k p_{train} \mid k = 0, \dots, K\}$ by left-multiplying previous ones with W , graph filters run in time $O(KE)$, where E is the number of graph edges. Thus, our graph filter parameter tuning mechanism can be implemented to run in amortized time:

$$O(K^2 E (\log_T L - \log_T \epsilon))$$

Running time scales linearly with the number of edges and quadratically with the number of parameters. We recommend and employ default parameters $P = 2, T = 1.01$, which suffice to minimize the Beale and Booth functions often used in optimization benchmarks [2] to 10^{-6} parameter (instead of loss) tolerance.

4 Experiment Setup

We experiment on three publicly available real-world graphs with metadata communities. First is the Amazon co-purchasing graph [20], whose nodes and edges correspond to products and frequent co-purchases. Products are organized into metadata communities based on their type (e.g. book, movie) attribute. Second is the Citeseer citation graph [11], whose nodes and edges correspond to scientific publications and citations. Publications are organized into communities based on scientific field. Third is the Maven dependency graph [5], whose nodes and edges correspond to software projects and dependencies. Projects are organized into communities based on the organization responsible for their development.

These graphs were chosen for experimentation on merit of comprising meta-data communities with enough member nodes to conduct robust validation. To not over-represent graphs with many communities and obtain enough validation nodes later on, we experiment with the first three communities of each graph with at least 500 nodes. We treat all edges as undirected so that symmetric normalization of filters is applicable. Community details are summarized in Table 1.

Table 1. Details of communities we experiment on.

Community	Graph	Nodes	Edges	Members
amazon0	Amazon	554,789	3,577,450	280,507
amazon1	Amazon	554,789	3,577,450	64,915
amazon2	Amazon	554,789	3,577,450	17,966
citeseer0	Citeseer	3,327	9,464	596
citeseer1	Citeseer	3,327	9,464	668
citeseer2	Citeseer	3,327	9,464	701
maven0	Maven	1,965,359	19,431,302	1,687
maven1	Maven	1,965,359	19,431,302	1,043
maven2	Maven	1,965,359	19,431,302	49,883

For each of the the above-described communities, we generate three splits of known-test members by assuming that known members are uniformly sampled to comprise 10%, 20%, or 30% of total members. We remind that validation nodes can only be subsampled from known members. Sampling is seeded to ensure reproducibility and fair comparison between approaches. In total, experiments on 9 communities create $9 \cdot 3 = 27$ different known-test member splits. For each split, we consider two different node ranking objectives; optimizing AUC, and optimizing NDCG. Thus, we obtain $27 \cdot 2 = 54$ experiment setups.

We investigate the ability of our approach to produce high-quality community node ranks compared to ad-hoc graph filters and parameters often encountered in the literature. We compare the following alternatives, all of which we integrated in the *pygrank* Python library [18] alongside experiment setups:

- *ppr a* [25, 3, 4]. Personalized PageRank that performs stochastic random walks with restart probabilities $1 - a$ at each step [29]. We test common values $a \in \{0.5, 0.85, 0.9, 0.99\}$ and compute filters to numerical tolerance 10^{-9} .
- *hkk* [8]. Heat kernels that form bandpass windows around desired propagation hops k . We test common window centers $k \in \{2, 3, 5, 7\}$.
- *select* [this work]. Runtime selection of the best among *ppr a* and *hkk* by withholding a 10% validation subset of known community members. When graphs are unknown during algorithm selection, this becomes a baseline for tuning.
- *tune* [this work]. Tuning a generalized graph filter with 40 parameters, where the filter is obtained with non-zero Taylor coefficients $h_0 = 1$ and tuned h_1, \dots, h_{40} via Algorithm 1 towards maximizing measures of choice on the same 10% validation subset as in *select*. Optimization absolute deviation tolerance is set to $\epsilon = 10^{-6}$.

- *tuneLBFGB* [ablation study]. A variation of *tune* that substitutes our tuning algorithm with the L-BFGS-B optimizer [7] provided by the *scipy* library [31] with default parameters and 10^{-6} percentage decrease on the evaluation function as a stopping criterion to make sure that tuning does not stop early. This is a popular optimizer still used for parameter search [30] and approximates Newton’s method while limiting the number of computations to only first-order gradients. Experiments with the Nelder-Mead optimizer yielded similar or worse results that we do not report due to space constraints.

5 Experiment Results

Tables 2 and 3 present the quality of community node ranking across experiment setups in terms of AUC and NDCG respectively. Before exploring graph filter selection, we verify that individual ad-hoc filter efficacy varies across communities and training-test splits. Indeed, no explored filter outperforms the rest in all experiments. For instance, ppr0.99 is often the best in Amazon communities, but also the worst in Maven communities, where it lags behind others up to 0.035 AUC. Runtime filter selection would be useful as long as it lags less behind.

Choosing between ad-hoc filters with our validation strategy does not always retrieve the best-performing ones. We attribute this behavior to few missing examples still impacting the ideal filter propagation weights needed for high-quality node ranking. Withholding fewer nodes could degrade validation robustness and future research could investigate new mechanisms to improve generalization. For the time being, selection of best among existing alternatives at runtime chooses the best filters in 31/54 settings. But, even when this scheme fails to identify the best filter, it often retrieves near-best ones that at worst lag behind only by 0.011 in terms of AUC or NDCG, where this gap usually shrinks to 0.001.

Parameter tuning with Algorithm 1 outperforms all ad-hoc filters in 40/54 experiment settings. This induces up to 0.010 AUC and 0.033 NDCG improvements, indicating that it manages to discover nuanced notions of structural proximity. It lags behind by at worst 0.007 on account of either measure, and often by much less. Compared to selecting among filters, tuning yields better evaluation outcomes in 49/54 of experiment settings. As such, we recommend it as an out-of-the-box solution for community node ranking in new graphs, especially if structural characteristics correlating to the formation of communities are not known beforehand. Finally, comparing our optimization algorithm to L-BFGS-B, the latter induces marginal improvements in the Citeseer graph, but falls significantly behind -even compared to filter selection- in the Amazon and Maven graphs. This corroborates the need for retaining a wide parameter search space.

In relation to applying our methodology, we experimented on communities with enough example members to achieve a robust evaluation when randomly withholding 10% of them. Fewer known members may not yield robust validation strategies and we hereby caution against blindly applying our methodology when too few members are known. In principle, we expect our approach to work well

Table 2. Test set AUC of community node ranks for ad-hoc filters and those obtained through runtime tuning on the same measure. The highest value is bolded.

Com.	Examples	Ad-hoc								Autotune		
		ppr0.5	ppr0.85	ppr0.9	ppr0.99	hk2	hk3	hk5	hk7	select	tune	tuneLFBGSB
amazon0	10%	.825	.844	.853	.901	.823	.825	.832	.844	.901	.903	.883
amazon0	20%	.820	.855	.870	.910	.817	.824	.844	.865	.910	.918	.900
amazon0	30%	.815	.868	.884	.914	.810	.824	.855	.880	.914	.924	.908
amazon1	10%	.930	.934	.935	.944	.930	.930	.931	.933	.944	.943	.940
amazon1	20%	.941	.946	.948	.954	.941	.941	.944	.946	.954	.955	.953
amazon1	30%	.946	.953	.955	.960	.946	.947	.951	.954	.960	.961	.959
amazon2	10%	.960	.964	.965	.968	.959	.960	.961	.963	.968	.968	.967
amazon2	20%	.970	.974	.974	.975	.970	.971	.972	.974	.975	.977	.976
amazon2	30%	.973	.976	.977	.976	.972	.973	.975	.976	.977	.979	.978
citeseer0	10%	.778	.789	.792	.793	.775	.777	.781	.784	.793	.795	.795
citeseer0	20%	.841	.849	.850	.845	.838	.840	.844	.847	.845	.852	.852
citeseer0	30%	.859	.867	.868	.861	.856	.859	.863	.866	.868	.869	.869
citeseer1	10%	.805	.807	.807	.798	.806	.805	.805	.806	.798	.805	.806
citeseer1	20%	.820	.823	.824	.813	.820	.820	.821	.823	.820	.819	.823
citeseer1	30%	.816	.821	.821	.811	.816	.816	.819	.820	.816	.821	.821
citeseer2	10%	.659	.669	.672	.675	.656	.657	.661	.664	.656	.670	.676
citeseer2	20%	.718	.727	.730	.731	.716	.717	.721	.724	.716	.732	.733
citeseer2	30%	.767	.773	.774	.769	.766	.767	.770	.773	.773	.775	.776
maven0	10%	.998	.997	.995	.942	.998	.998	.998	.997	.998	.994	.989
maven0	20%	.997	.995	.994	.925	.997	.997	.997	.996	.997	.998	.984
maven0	30%	.997	.995	.993	.911	.998	.997	.997	.995	.998	.998	.982
maven1	10%	.996	.993	.991	.967	.996	.996	.996	.995	.996	.995	.986
maven1	20%	.996	.992	.990	.961	.996	.996	.995	.994	.994	.995	.984
maven1	30%	.997	.992	.989	.956	.998	.997	.996	.995	.992	.992	.982
maven2	10%	.998	.998	.998	.993	.997	.998	.998	.998	.998	.998	.997
maven2	20%	.997	.997	.997	.991	.997	.997	.997	.997	.997	.997	.996
maven2	30%	.997	.997	.997	.990	.997	.997	.998	.997	.997	.997	.996
Average		.897	.903	.905	.898	.896	.897	.901	.904	.909	.912	.908

-and therefore be applicable on- community node ranking based on at least the same number of known members (at least 50) as in our experiments.

As evidence that tuning discovers non-trivial graph propagation schemes, Figure 2 shows the first 41 parameters of high-AUC filters for citeseer0 with 30% known members. There, tuning discovers a different propagation strategy than ad-hoc filters, which subsequently manages to (slightly) improve the best filter in Table 2. Moreover, Figure 3 shows that tuning is tailored not only to communities but even to specific sets of example nodes, yielding drastically different filters for the same communities. Given that tuned graph filters generally outperform others, this finding corroborates our hypothesis that filters should be selected at runtime to match the characteristics of data they are about to process. Finally, filter differences between different fractions of community examples support our practice of withholding only a small fraction of validation nodes.



**Fig. 2.** Parameters h_i of filters with high AUC on citeseer0 with 30% examples.**Fig. 3.** Parameters tuned on citeseer0 and citeseer1 with 10% and 20% examples.

Table 3. Test set NDCG of community node ranks for ad-hoc filters and those obtained through runtime tuning on the same measure. The highest value is bolded.

Com.	Examples	Ad-hoc								Autotune		
		ppr0.5	ppr0.85	ppr0.9	ppr0.99	hk2	hk3	hk5	hk7	select	tune	tuneLFBGSE
amazon0	10%	.970	.976	.978	.985	.970	.971	.975	.978	.985	.987	.983
amazon0	20%	.968	.976	.979	.984	.967	.970	.975	.978	.984	.987	.983
amazon0	30%	.965	.976	.978	.983	.964	.968	.974	.978	.983	.987	.982
amazon1	10%	.955	.962	.964	.972	.954	.956	.960	.963	.972	.972	.969
amazon1	20%	.954	.963	.966	.972	.953	.956	.962	.965	.972	.976	.971
amazon1	30%	.951	.963	.966	.970	.950	.955	.961	.966	.970	.976	.970
amazon2	10%	.912	.923	.927	.933	.910	.913	.920	.925	.920	.927	.934
amazon2	20%	.914	.928	.932	.933	.913	.917	.926	.931	.933	.941	.937
amazon2	30%	.912	.926	.929	.929	.910	.915	.924	.929	.929	.938	.933
citeseer0	10%	.896	.900	.902	.902	.894	.895	.897	.899	.902	.903	.904
citeseer0	20%	.922	.931	.933	.926	.920	.923	.928	.931	.926	.932	.934
citeseer0	30%	.915	.925	.926	.923	.907	.918	.923	.926	.923	.930	.929
citeseer1	10%	.905	.905	.905	.904	.905	.905	.904	.903	.904	.907	.907
citeseer1	20%	.901	.903	.904	.895	.900	.901	.903	.903	.903	.897	.905
citeseer1	30%	.892	.893	.892	.877	.891	.892	.893	.892	.892	.890	.889
citeseer2	10%	.869	.876	.878	.879	.867	.868	.872	.875	.868	.883	.881
citeseer2	20%	.887	.896	.897	.899	.886	.888	.890	.895	.886	.898	.900
citeseer2	30%	.894	.904	.905	.903	.893	.896	.902	.905	.905	.908	.908
maven0	10%	.803	.786	.778	.723	.803	.804	.796	.774	.796	.813	.745
maven0	20%	.738	.708	.693	.615	.743	.741	.724	.694	.743	.744	.637
maven0	30%	.673	.637	.625	.554	.676	.674	.651	.629	.674	.670	.581
maven1	10%	.812	.823	.821	.781	.807	.817	.828	.830	.828	.863	.814
maven1	20%	.806	.814	.812	.767	.788	.800	.817	.818	.818	.831	.801
maven1	30%	.770	.774	.773	.728	.767	.775	.780	.773	.773	.825	.762
maven2	10%	.904	.907	.906	.845	.903	.908	.911	.909	.911	.935	.889
maven2	20%	.863	.864	.861	.785	.862	.867	.869	.865	.869	.903	.840
maven2	30%	.812	.812	.808	.729	.811	.816	.818	.813	.818	.861	.786
Average		.880	.883	.883	.863	.878	.882	.885	.883	.888	.899	.877

6 Conclusions and Future Work

This work introduces a runtime graph filter selection scheme for community node ranking based on known member nodes. Selection involves either choosing between promising filters or tuning the parameters of a generalized filter form. For the latter, we introduced a novel algorithm that meshes parts of previous alternatives to satisfy both scalability and a wide parameter search breadth needed by graph filters. We verified the efficacy of our approach with experiments across real-world graph communities, where we found that, given enough example community members to satisfy robust evaluation by withholding a few of them, our methodology (especially tuning) yields filters with similar or better AUC and NDCG than alternatives. Thus, we recommend its adoption in practice.

In the future, we are interested in experimenting on more graphs, improving our tuning algorithm, and theoretically probing its optimality. More robust evaluation could also be devised to autotune from fewer known community members.

Acknowledgements

This work was partially funded by the European Commission under contract number H2020-951911 AI4Media.

References

1. Abu-El-Haija, S., Kapoor, A., Perozzi, B., Lee, J.: N-gcn: Multi-scale graph convolution for semi-supervised node classification. In: *uncertainty in artificial intelligence*. pp. 841–851. PMLR (2020)
2. Al-Roomi, A.R.: Unconstrained Single-Objective Benchmark Functions Repository (2015), <https://www.al-roomi.org/benchmarks/unconstrained>
3. Andersen, R., Chung, F., Lang, K.: Local partitioning for directed graphs using pagerank. *Internet Mathematics* 5(1-2), 3–22 (2008)
4. Bahmani, B., Chowdhury, A., Goel, A.: Fast incremental and personalized pagerank. *Proceedings of the VLDB Endowment* 4(3) (2010)
5. Benelallam, A., Harrand, N., Valero, C.S., Baudry, B., Barais, O.: Maven central dependency graph (2018)
6. Bradley, A.P.: The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition* 30(7), 1145–1159 (1997)
7. Byrd, R.H., Lu, P., Nocedal, J., Zhu, C.: A limited memory algorithm for bound constrained optimization. *SIAM Journal on scientific computing* 16(5), 1190–1208 (1995)
8. Chung, F.: The heat kernel as the pagerank of a graph. *Proceedings of the National Academy of Sciences* 104(50), 19735–19740 (2007)
9. Finkel, D.E., Kelley, C.: Additive scaling and the direct algorithm. *Journal of Global Optimization* 36(4), 597–608 (2006)
10. Galántai, A.: Convergence of the nelder-mead method. *Numerical Algorithms* pp. 1–30 (2021)
11. Getoor, L.: Link-based classification. In: *Advanced methods for knowledge discovery from complex data*, pp. 189–207. Springer (2005)
12. Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 855–864 (2016)
13. Huang, Q., He, H., Singh, A., Lim, S.N., Benson, A.R.: Combining label propagation and simple models out-performs graph neural networks. *arXiv preprint arXiv:2010.13993* (2020)
14. Järvelin, K., Kekäläinen, J.: Ir evaluation methods for retrieving highly relevant documents. In: *ACM SIGIR Forum*. vol. 51, pp. 243–250. ACM New York, NY, USA (2017)
15. Klicpera, J., Bojchevski, A., Günnemann, S.: Predict then propagate: Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997* (2018)
16. Kloster, K., Gleich, D.F.: Heat kernel based community detection. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 1386–1395 (2014)
17. Koch, P., Golovidov, O., Gardner, S., Wujek, B., Griffin, J., Xu, Y.: Autotune: A derivative-free optimization framework for hyperparameter tuning. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. pp. 443–452 (2018)
18. Krasanakis, E., Papadopoulos, S., Kompatsiaris, I., Symeonidis, A.: pygrank: A python package for graph node ranking. *arXiv preprint arXiv:2110.09274* (2021)
19. Krasanakis, E., Schinas, E., Papadopoulos, S., Kompatsiaris, Y., Symeonidis, A.: Boosted seed oversampling for local community ranking. *Information Processing & Management* 57(2), 102053 (2020)

20. Leskovec, J., Adamic, L.A., Huberman, B.A.: The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* 1(1), 5–es (2007)
21. Leskovec, J., Lang, K.J., Dasgupta, A., Mahoney, M.W.: Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics* 6(1), 29–123 (2009)
22. Lyu, H.: Convergence of block coordinate descent with diminishing radius for non-convex optimization. arXiv preprint arXiv:2012.03503 (2020)
23. McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of a feather: Homophily in social networks. *Annual review of sociology* 27(1), 415–444 (2001)
24. Ortega, A., Frossard, P., Kovačević, J., Moura, J.M., Vandergheynst, P.: Graph signal processing: Overview, challenges, and applications. *Proceedings of the IEEE* 106(5), 808–828 (2018)
25. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: Bringing order to the web. Tech. rep., Stanford InfoLab (1999)
26. Papadopoulos, S., Kompatsiaris, Y., Vakali, A., Spyridonos, P.: Community detection in social media. *Data Mining and Knowledge Discovery* 24(3), 515–554 (2012)
27. Shuman, D.I., Narang, S.K., Frossard, P., Ortega, A., Vandergheynst, P.: The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE signal processing magazine* 30(3), 83–98 (2013)
28. Stanković, L., Daković, M., Sejdić, E.: Introduction to graph signal processing. In: *Vertex-Frequency Analysis of Graph Signals*, pp. 3–108. Springer (2019)
29. Tong, H., Faloutsos, C., Pan, J.Y.: Fast random walk with restart and its applications. In: *Sixth international conference on data mining (ICDM'06)*. pp. 613–622. IEEE (2006)
30. Tooley, R.: Auto-tuning spark with bayesian optimisation (2021)
31. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17, 261–272 (2020)
32. Whang, J.J., Gleich, D.F., Dhillon, I.S.: Overlapping community detection using seed set expansion. In: *Proceedings of the 22nd ACM international conference on information & knowledge management*. pp. 2099–2108 (2013)
33. Whang, J.J., Gleich, D.F., Dhillon, I.S.: Overlapping community detection using neighborhood-inflated seed expansion. *IEEE Transactions on Knowledge and Data Engineering* 28(5), 1272–1284 (2016)
34. Wu, F., Huberman, B.A.: Finding communities in linear time: a physics approach. *The European Physical Journal B* 38(2), 331–338 (2004)
35. Yang, J., Leskovec, J.: Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems* 42(1), 181–213 (2015)
36. Zareie, A., Sheikahmadi, A.: A hierarchical approach for influential node ranking in complex social networks. *Expert Systems with Applications* 93, 200–211 (2018)
37. Zhang, T., Wu, B.: A method for local community detection by finding core nodes. In: *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. pp. 1171–1176. IEEE (2012)