Heterophily-Aware Node Classification in Multiplex Graphs

Kamel Abdous*, Nairouz Mrabah*, and Mohamed Bouguessa(✉)

Department of Computer Science, University of Quebec at Montreal
Montreal, Quebec, Canada
abdous.kamel@courrier.uqam.ca, mrabah.nairouz@courrier.uqam.ca,
bouguessa.mohamed@uqam.ca

Abstract. In multiplex graphs, nodes are connected through multiple types of edges (referred to as dimensions), each of which may exhibit different connectivity patterns. While many graph learning models implicitly rely on homophily, real-world relations often display heterophily, particularly across different dimensions, which can lead to conflicting relational signals. This challenge is amplified in multiplex settings, where contrasting homophilic and heterophilic regimes may coexist across dimensions. We propose HAAM, an adaptive node classification approach for multiplex graphs that models per-dimension degrees of homophily and heterophily through dimension-specific compatibility matrices. Our method captures both smooth and abrupt signal variations using a composition of trainable low-pass and high-pass spectral filters, approximated with Chebyshev polynomials. These filters are integrated into a proximal-gradient optimization framework that adapts to the structural characteristics of each dimension. Experiments on real-world multiplex datasets show that HAAM achieves competitive performance compared to representative baseline methods.

Keywords: Graph Neural Networks · Multiplex Graphs · Heterophily · Node Classification.

1 Introduction

Multiplex graphs provide a flexible and expressive framework for modeling complex systems in which entities interact through multiple types of relations. Unlike unidimensional graphs, which encode a single type of interaction, multiplex graphs preserve the diversity of relational semantics by representing each interaction type as a distinct dimension [1]. This modeling paradigm has proven effective across a wide range of application domains, including social networks [2], biological systems [3], and online recommendation platforms [4], where different relational views jointly influence node behavior. By explicitly retaining multiple interaction channels, multiplex graphs enable a more faithful representation of system dynamics and support more informed learning and inference processes [5].

* The first two authors contributed equally to this work.

In this paper, we focus on multiplex graphs with homogeneous nodes, where a single set of entities is connected through multiple relational dimensions. Graphs involving multiple node types (commonly referred to as heterogeneous graphs), as well as dynamic or time-evolving graph settings, are outside the scope of this paper.

A large body of graph learning methods, including those developed for multiplex settings, is built on the homophily assumption, according to which edges tend to connect nodes sharing similar attributes or labels [6]. This premise has been instrumental in the success of neighborhood aggregation and message-passing techniques, yet it offers a limited description of systems where interactions arise from complementary roles or functional differentiation rather than similarity [7]. In such settings, heterophily plays a central role, as nodes with dissimilar characteristics are more likely to be connected [8]. Examples include professional networks that foster collaboration between individuals with complementary expertise [9] and biological networks involving interactions between entities with distinct yet interdependent roles [10].

From a structural perspective, homophily and heterophily describe graph connectivity patterns according to whether edges predominantly link nodes with similar or dissimilar attributes [11]. Dimensions dominated by homophily tend to form dense intra-class neighborhoods, whereas heterophilic dimensions are marked by a predominance of inter-class connections. Neighborhood similarity generally supports representation learning in homophilic settings [12]. By contrast, heterophilic structures complicate label propagation, as local neighborhoods may aggregate conflicting signals [7]. In multiplex graphs, this difficulty is compounded by the coexistence of multiple dimensions exhibiting distinct structural regimes. As a result, learning models must cope with multiple and potentially conflicting relational patterns within the same graph.

In multiplex graphs, the presence of heterophily introduces additional difficulties, as some dimensions may exhibit strong homophilic patterns while others are predominantly heterophilic, leading to diverse and sometimes contradictory signals for downstream learning tasks [13]. Recent approaches have made notable progress in modeling homophily within multiplex graphs [14–18], and complementary efforts have addressed heterophily in unidimensional graphs [11, 19]. However, extending these ideas to multiplex settings remains challenging, as different dimensions may induce divergent label propagation behaviors and affect overall prediction performance.

Beyond performance considerations, such structural conflicts raise practical concerns regarding the reliability of node classification in multiplex graphs, in the sense of producing consistent predictions across conflicting relational dimensions. When learning algorithms are overly influenced by specific dimensions, predictions may vary significantly depending on which interactions dominate the representation, resulting in dimension-dependent outcomes. In multi-relational environments, reliable node classification therefore requires not only the ability to capture both homophilic and heterophilic structures, but also mechanisms that promote consistent prediction behavior across diverse dimensions. Address-

ing this issue is essential for improving the reliability of graph-based models in structurally diverse settings.

To this end, we introduce HAAM (Heterophily-Aware Adaptive Multiplex model), a framework for node classification in multiplex graphs that explicitly accounts for mixed homophilic and heterophilic interactions. Unlike existing methods that predominantly emphasize homophilic structures, HAAM incorporates learnable, dimension-specific compatibility matrices [12] to model varying degrees of homophily and heterophily across dimensions. In addition, we propose a composition of trainable low-pass and high-pass spectral Chebyshev filters [20] to capture both smooth and rapidly varying signal components induced by node interactions. These filters are combined through a novel product operation, rather than the linear combinations commonly used in prior work [21–23], allowing the model to jointly account for complementary structural patterns within each dimension.

The resulting representations are integrated into a learning framework that balances dimension-specific predictions with a sparse consensus objective, encouraging agreement across dimensions while mitigating the influence of conflicting homophilic and heterophilic dimensions. Model training is performed using a combination of a standard cross-entropy loss and a consensus regularization term, optimized via a proximal-gradient scheme. Through this design, HAAM adapts to heterogeneous connectivity structures in multiplex graphs while maintaining competitive performance on real-world datasets. Building on these properties, our work[1] makes three key contributions: (i) a multiplex classification approach that explicitly models dimension-specific homophily and heterophily; (ii) an adaptive spectral filtering scheme based on a product of Chebyshev polynomials, enabling non-linear fusion of structural signals; and (iii) a comprehensive empirical evaluation showing that HAAM delivers competitive performance compared to representative baseline methods.

2 The HAAM Framework

We now introduce the HAAM framework, which addresses node classification in multiplex graphs. Rather than focusing solely on predictive performance under heterophily, HAAM explicitly aims to promote consistent prediction behavior across dimensions that may exhibit structurally diverse and potentially conflicting connectivity patterns. We first formalize the problem setting and necessary notations, then describe the architecture and its main components in detail.

[1] A substantially extended journal version of this work appears in [24] (DOI: `https://doi.org/10.1016/j.eswa.2026.132374`). It includes related work discussion, theoretical analyses and proofs, complexity analysis, synthetic data experiments, additional ablation studies, robustness analyses, and extended experimental results. An arXiv version is also available at `https://arxiv.org/abs/2605.12699`

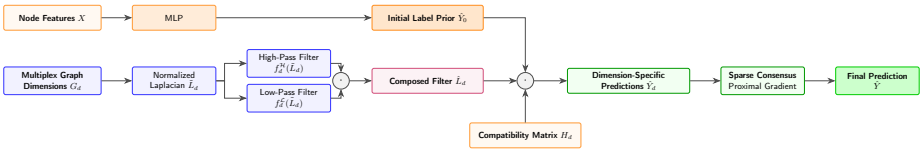

Fig. 1: Overview of the HAAM framework. Node features are first mapped to an initial label prior $\hat{Y}_0$ shared across dimensions. For each multiplex dimension, a composed spectral filter $\hat{L}_d$, combining low-pass and high-pass components, is applied to the normalized Laplacian, followed by a compatibility transformation H_d. The resulting dimension-wise predictions $\hat{Y}_d$ are then reconciled through a sparse consensus optimization to produce the final prediction $\hat{Y}$.

2.1 Preliminaries

We consider a multiplex graph G composed of D dimensions: $G = \{G_1, \ldots, G_D\}$. Each dimension $G_d = (V, A_d)$ shares the same set of N nodes $V = \{v_1, \ldots, v_N\}$ but is characterized by its own adjacency matrix $A_d \in \mathbb{R}^{N \times N}$, where $(A_d)_{ij} = 1$ if there is an edge between nodes v_i and v_j in dimension d, and 0 otherwise. The degree matrix Δ_d is diagonal with $(\Delta_d)_{ii} = \sum_{j=1}^{N} (A_d)_{ij}$, and the normalized Laplacian for dimension d is defined as:

$$L_d = I - \Delta_d^{-\frac{1}{2}} (A_d + I) \Delta_d^{-\frac{1}{2}}.$$

Node attributes are represented by a feature matrix $X \in \mathbb{R}^{N \times F}$, and node labels are encoded in a binary matrix $Y \in \{0,1\}^{N \times C}$, where C denotes the number of classes.

To characterize the class-wise connectivity structure within each dimension, we define the homophily ratio h_d. Let $\mathcal{Y}_d \in \mathbb{N}^{C \times C}$ be the matrix counting edges between classes in dimension d, where $(\mathcal{Y}_d)_{ij}$ corresponds to the number of edges connecting nodes of class i to nodes of class j. The homophily ratio is defined as:

$$h_d = \frac{\sum_{i=1}^{C} (\mathcal{Y}_d)_{ii}}{\sum_{i,j=1}^{C} (\mathcal{Y}_d)_{ij}}.$$

This quantity provides a compact description of how class interactions are distributed within a given dimension and is used throughout the paper to analyze structural diversity across multiplex dimensions.

2.2 Overall Architecture

The architecture of HAAM is designed to handle both homophilic and heterophilic patterns across multiplex dimensions while explicitly addressing the potential inconsistency of dimension-specific predictions. Fig. 1 illustrates an overview of the proposed framework. It consists of three main components: (i) an initial label prior learned independently of graph structure, (ii) dimension-specific spectral filtering combined with compatibility-aware propagation, and (iii) a sparse consensus mechanism that resolves conflicts across dimensions.

Initial label prior. In settings where heterophily is prominent, the relationship between graph topology and node features may be highly non-smooth and dimension-dependent. To mitigate this effect, we first encode node features independently of the graph structure using a Multilayer Perceptron (MLP). This produces an initial label distribution $\hat{Y}_0 \in \mathbb{R}^{N\times C}$ that serves as a common prior across all dimensions. This prior is subsequently refined through dimension-specific propagation mechanisms.

Dimension-specific spectral filtering. For each dimension d, we construct two learnable spectral filters: a low-pass filter $f_d^{\mathcal{L}}$ and a high-pass filter $f_d^{\mathcal{H}}$. These filters are implemented as Chebyshev polynomials applied to the rescaled Laplacian $\tilde{L}_d = 2L_d/\lambda_d^{\max} - I$, where $\lambda_d^{\max}$ denotes the largest eigenvalue of L_d. The low-pass filter emphasizes smooth variations in the graph signal, while the high-pass filter captures abrupt changes induced by heterophilic connections.

Unlike prior works that combine spectral filters through linear combinations or concatenations [21–23], we compose the two filters through a matrix product:

$$\hat{L}_d = f_d^{\mathcal{L}}(\tilde{L}_d) \cdot f_d^{\mathcal{H}}(\tilde{L}_d). \tag{1}$$

This composition enables nonlinear interactions between low-frequency and high-frequency components and introduces higher-order polynomial terms, resulting in a more expressive and controlled spectral response.

Compatibility-aware propagation. To further regularize class interactions within each dimension, HAAM incorporates a learnable compatibility matrix $H_d \in \mathbb{R}^{C\times C}$. The entry $(H_d)_{c_1c_2}$ models the likelihood that a node of class c_1 connects to a node of class c_2 in dimension d. Given the composed filter $\hat{L}_d$ and the initial label prior $\hat{Y}_0$, the dimension-specific predictions are computed as:

$$\hat{Y}_d = \hat{L}_d \cdot \hat{Y}_0 \cdot H_d. \tag{2}$$

Here, $\hat{L}_d$ propagates information through the graph topology, while H_d adjusts predictions according to observed inter-class connectivity patterns, jointly capturing homophilic and heterophilic effects.

Sparse consensus across dimensions. Predictions obtained from different dimensions may be mutually inconsistent due to structurally diverse or conflicting structural properties. To obtain a reliable final prediction, HAAM formulates consensus estimation as an optimization problem that balances agreement across dimensions with confidence in class assignments:

$$\arg\min_{\hat{Y}\in\mathbb{R}^{N\times C}} \sum_{d=1}^{D} \|\hat{Y} - \hat{Y}_d\|_2^2 + \beta\|\hat{Y}\|_1. \tag{3}$$

The first term in Eq. (3) is a sum of Frobenius norms to minimize the distance between the consensus prediction $\hat{Y}$ and the dimension-specific predictions $\hat{Y}_d$. The second term is the ℓ_1 regularization used to promote sparsity in the consensus prediction, encouraging confident decisions while mitigating the influence of conflicting dimension-specific signals. The coefficient β is a balancing hyperparameter. This convex objective is optimized using a proximal-gradient method.

2.3 Composition of Chebyshev Filters

The filters $f_d^{\mathcal{L}}$ and $f_d^{\mathcal{H}}$ are defined as Chebyshev polynomials of degree K:

$$f_d^{\mathcal{L}}(\tilde{L}_d) = \sum_{k=0}^{K} \theta_k^{\mathcal{L},d}\, T_k(\tilde{L}_d), \qquad f_d^{\mathcal{H}}(\tilde{L}_d) = \sum_{k=0}^{K} \theta_k^{\mathcal{H},d}\, T_k(\tilde{L}_d),$$

where T_k denotes the k-th Chebyshev polynomial. The coefficients $\theta_k^{\mathcal{L},d}$ and $\theta_k^{\mathcal{H},d}$ determine the contribution of different spectral components.

To ensure that these polynomials approximate genuine low-pass and high-pass responses, we adopt a double reparameterization strategy. First, each coefficient $\theta_k^{\mathcal{F},d}$, with $\mathcal{F} \in \{\mathcal{L}, \mathcal{H}\}$, is expressed as:

$$\theta_k^{\mathcal{F},d} = \frac{2}{K+1} \sum_{j=0}^{K} \gamma_j^{\mathcal{F},d}\, T_k\left(\cos\left(\frac{j+1/2}{K+1}\pi\right)\right),$$

where $\gamma^{\mathcal{F},d} \in \mathbb{R}^{K+1}$ is an intermediate vector. Second, low-pass and high-pass structures are enforced through prefix operations:

$$\gamma_i^{\mathcal{L},d} = \gamma_0^d - \sum_{j=1}^{i} \gamma_j^d, \qquad \gamma_i^{\mathcal{H},d} = \sum_{j=0}^{i} \gamma_j^d, \tag{4}$$

where γ_0^d is fixed and $\gamma_1^d, \ldots, \gamma_K^d \geq 0$ are learnable parameters. This construction guarantees that $f_d^{\mathcal{L}}$ attenuates high frequencies while $f_d^{\mathcal{H}}$ attenuates low frequencies.

The product $\hat{L}_d = f_d^{\mathcal{L}} \cdot f_d^{\mathcal{H}}$ can be computed efficiently without explicit $N \times N$ matrix multiplications using:

$$\hat{L}_d = \frac{1}{2} \sum_{i=0}^{K} \sum_{j=0}^{K} \theta_i^{\mathcal{L},d} \theta_j^{\mathcal{H},d} \Big[T_{i+j}(\tilde{L}_d) + T_{|i-j|}(\tilde{L}_d) \Big]. \tag{5}$$

This formulation reduces computational and memory overhead while preserving the effect of nonlinear spectral composition.

2.4 Compatibility Matrices

The compatibility matrix H_d is initialized from training data to reflect observed class connectivity patterns:

$$(H_d)_{c_1 c_2} = \frac{\sum_{i \in I_{c_1}^{\text{train}},\, j \in I_{c_2}^{\text{train}}} (A_d)_{ij}}{\sum_{i,j=1}^{N} (A_d)_{ij}}. \tag{6}$$

Here, I_c^{train} denotes the set of training nodes belonging to class c. This empirical initialization provides a meaningful starting point that captures dimension-specific class interaction tendencies. During training, H_d is refined via backpropagation, allowing the model to adapt class interaction estimates to the prediction task.

2.5 Training and Optimization

HAAM is trained in a semi-supervised manner. The primary objective is an ℓ_2-regularized cross-entropy loss aggregated across all dimensions:

$$\mathcal{J}_{\text{CE}} = -\sum_{i \in I^{\text{train}}} \sum_{d=1}^{D} \sum_{c=1}^{C} Y_{ic} \log(\hat{Y}_d)_{ic} + \alpha \sum_{w \in \mathbb{W}} \|w\|_2^2, \tag{7}$$

where $\mathbb{W}$ includes the parameters of the MLP, spectral filters, and compatibility matrices.

After training, the consensus prediction $\hat{Y}$ is obtained by solving (3) using proximal-gradient iterations:

$$\hat{Y}^{(t+1)} = \text{prox}_{\beta\tau}\Big(\hat{Y}^{(t)} - 2\tau \sum_{d=1}^{D} (\hat{Y}^{(t)} - \hat{Y}_d)\Big), \tag{8}$$

with step size $\tau = 1/(4D)$ and proximal operator $\text{prox}_{\lambda}(v) = \text{sign}(v)\max(|v| - \lambda, 0)$.

3 Experiments

We conduct an empirical evaluation to show the suitability of the proposed approach for modeling heterophily and homophily for node classification in multiplex graphs. We compare HAAM against state-of-the-art multiplex graph models, namely: GATNE [25], mGCN [26], SSDCM [16], DMGI [27], HDMI [15], MGDCR [28], DMG [29], X-GOAL [14], HMGE [5], and InfoMGF [30]. In addition, we include recent unidimensional graph representation learning baselines that can handle heterophily. More precisely, our comparison includes PolyGCL [22] and TFE-GNN [23] on a unified adjacency matrix $\tilde{A}$ using the original feature matrix X. We adapt these methods to the multiplex setting following the common practice of aggregating all multiplex dimensions into one unified adjacency matrix $\tilde{A}$: $\tilde{A} = \frac{1}{D}\sum_{d=1}^{D}(A_d + I)$.

Table 1: Dataset statistics.

Dataset	arXiv	Movies	Amazon
Number of dimensions	2	3	3
Number of nodes	169,343	10,589	7,621
Number of edges	7,879,585	485,962	1,386,799
Number of attributes	128	200	2,000
Number of classes	5	4	4
Homophily ratio h_d	0.22 - 0.29	0.34 - 0.32 - 0.37	0.27 - 0.26 - 0.25

3.1 Datasets

We evaluate HAAM on three real-world multiplex graph datasets summarized in Table 1. **arXiv** [10] is derived from the OGBN-arXiv network and contains citation and co-authorship dimensions. Nodes represent scientific papers, node features are Word2Vec embeddings extracted from titles and abstracts, and labels correspond to publication years. **Movies** is constructed from the film–director–actor–writer network introduced in [31]. Nodes correspond to movies, dimensions capture shared directors, actors, and writers, node features are token-count vectors of movie descriptions, and labels indicate release years. **Amazon** [32] consists of Amazon products connected through *also viewed*, *also bought*, and *bought together* relations. Node features are bag-of-words vectors extracted from product descriptions, and labels correspond to product categories.

3.2 Evaluation Protocol & Parameter Settings

We conduct a comparative evaluation of HAAM and representative baseline methods on the node classification task, including both supervised and unsupervised approaches. For unsupervised representation learning methods, including HDMI, HMGE, GATNE, DMG, X-GOAL, InfoMGF, and PolyGCL, node embeddings are first learned without access to label information. A logistic regression classifier is then trained on top of the learned embeddings to predict node labels. In contrast, for supervised or semi-supervised methods, namely DMGI, SSDCM, MGDCR, mGCN, TFE-GNN, and HAAM, class predictions are obtained directly from the model outputs. All methods are evaluated using identical training, validation, and test splits to ensure fair comparisons.

Performance is measured using F1-Macro and F1-Micro scores by comparing predicted labels with ground-truth annotations. Each experiment is repeated five times. For HAAM, we report the mean performance along with the corresponding standard deviation across runs, whereas for all baseline methods we report the best result obtained over five runs. For HAAM, the embedding dimension is fixed to 64. The Chebyshev polynomial filter order K is set to 4, 2, and 3 for arXiv, Movies, and Amazon, respectively. Model parameters are optimized using the Adam optimizer with a learning rate of 0.001 and an ℓ_2 weight decay coefficient $\alpha = 10^{-5}$. Training is conducted for a maximum of 1,000 epochs, with early stopping applied when the validation performance does not improve for 100 consecutive epochs. The ℓ_1 regularization parameter β is fixed to 1.0, and optimization is performed with respect to the loss function defined in Eq. (3).

3.3 Results

Table 2 reports the node classification performance on real-world multiplex datasets. The methods are grouped into *unidimensional graph approaches* (PolyGCL and TFE-GNN), which are evaluated on the aggregated adjacency matrix $\tilde{A}$, and *native multiplex graph methods*, which directly exploit all relational dimensions $\{A_d\}_{d=1}^{D}$.

Table 2: Node classification results (F1-score).

Method	arXiv		Movies		Amazon	
	Macro	Micro	Macro	Micro	Macro	Micro
Unidimensional graph methods						
PolyGCL	23.00	35.32	35.53	37.81	68.37	69.53
TFE-GNN	18.71	21.95	39.44	41.22	20.10	29.63
Multiplex graph methods						
DMGI	15.01	28.92	30.19	31.37	74.58	74.81
HDMI	–	–	<u>40.25</u>	<u>41.42</u>	84.13	84.29
SSDCM	–	–	30.14	30.53	82.20	82.40
HMGE	31.64	44.03	33.86	37.56	24.84	37.80
GATNE	23.56	34.28	29.62	31.40	19.11	30.90
MGDCR	31.91	43.42	39.49	40.10	70.10	70.19
DMG	<u>35.45</u>	44.08	39.49	40.42	55.72	57.22
X-GOAL	28.72	40.17	39.55	40.57	<u>85.70</u>	<u>85.79</u>
mGCN	33.03	<u>44.45</u>	38.63	40.61	74.86	75.13
InfoMGF	23.11	36.49	39.28	40.54	66.12	65.87
HAAM	**39.95**	**48.44**	**41.81**	**42.39**	**88.32**	**88.37**
(Standard Dev.)	± 0.26	± 0.09	± 0.37	± 0.37	± 0.25	± 0.31

Across all datasets, HAAM achieves the highest F1-Macro and F1-Micro scores, while exhibiting very low standard deviations, indicating stable and reliable performance. On arXiv, HAAM attains F1-Macro = 39.95 and F1-Micro = 48.44, outperforming the strongest multiplex baselines DMG (second-best Macro: 35.45) and mGCN (second-best Micro: 44.45) by margins of +4.50 and +3.99, respectively. Notably, several methods, such as HDMI and SSDCM, run out of memory (OOM) on this large dataset, whereas HAAM remains computationally tractable.

On the Movies dataset, HAAM achieves 41.81 (Macro) and 42.39 (Micro), improving upon the best-performing baseline HDMI (40.25/41.42) by +1.56 in F1-Macro and +0.97 in F1-Micro. On Amazon, which exhibits low homophily across all dimensions, HAAM obtains the largest absolute gains, reaching 88.32/88.37 compared to the strongest baseline X-GOAL (85.70/85.79), corresponding to improvements of +2.62 (Macro) and +2.58 (Micro). These results highlight the limitations of existing multiplex approaches that do not explicitly account for relation-dependent heterophily patterns.

Although PolyGCL and TFE-GNN are specifically designed to handle heterophily in unidimensional graphs, their adaptation to multiplex settings via the aggregated graph $\tilde{A}$ leads to substantially weaker performance, particularly on arXiv and Amazon. This suggests that collapsing multiple relations into a

single graph can obscure relation-specific structural signals that are critical for accurate classification.

Overall, Table 2 shows that HAAM effectively models dimension-specific homophilic and heterophilic interactions in multiplex graphs, leading to improved node classification performance compared to both unidimensional heterophily-aware methods and state-of-the-art multiplex baselines.

3.4 Ablation Study

Table 3: Ablation study on node classification (F1-score).

Model variant	arXiv		Movies		Amazon	
	Macro	Micro	Macro	Micro	Macro	Micro
Naive model	33.03	44.45	38.63	40.61	74.86	75.13
f_d^h	38.97	47.05	40.12	41.32	86.62	86.79
f_d^h + prox	39.51	47.66	40.53	41.55	87.28	87.51
$H_d + f_d^h$	39.37	47.57	39.40	40.32	87.43	87.51
H_d + prox	36.83	46.92	38.46	40.56	86.08	86.43
$H_d + f_d^h$ + prox	40.63	48.38	41.19	42.21	88.26	88.28
Full model (RAMC)	**39.95**	**48.44**	**41.81**	**42.39**	**88.32**	**88.37**

The ablation study in Table 3 explores the impact of the components of HAAM on node classification performance. Starting with the naive model, which only runs a standard GCN on each dimension, the addition of various components brings notable improvements. The heterophily-aware spectral component f_d^h is the strongest single addition to the naive per-dimension GCN. The gain is especially large on Amazon, whose dimensions have low homophily, showing the benefit of explicitly capturing high-frequency heterophilic signals. The proximal consensus step further improves f_d^h on all datasets. By contrast, variants that use H_d without f_d^h are weaker, indicating that compatibility modeling and consensus cannot substitute for heterophily-aware spectral filtering. Combining H_d, f_d^h, and prox gives the strongest ablated variant. The full HAAM model, which adds the composed low-pass/high-pass spectral filtering, achieves the best result in five of the six metrics. Its only exception is arXiv Macro, where it remains second-best. Overall, the results show that f_d^h provides the largest individual gain, the proximal consensus stabilizes predictions across dimensions, and the full spectral composition yields the most balanced performance across datasets.

3.5 Sensitivity Analysis

Fig. 2 analyzes the sensitivity of HAAM with respect to the embedding dimension M and the Chebyshev filter order K on the node classification task. The

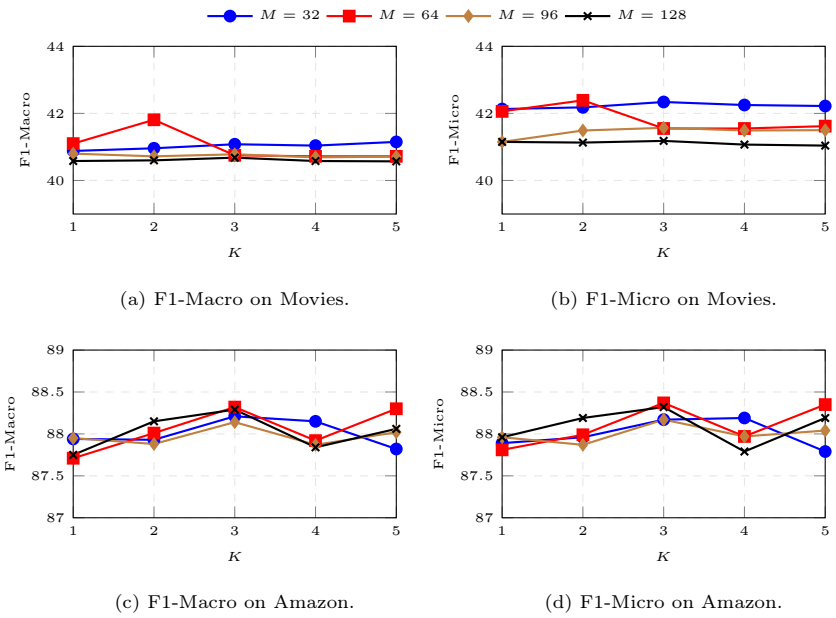

(a) F1-Macro on Movies.

(b) F1-Micro on Movies.

(c) F1-Macro on Amazon.

(d) F1-Micro on Amazon.

Fig. 2: Sensitivity analysis of RAMC.

results are reported for both F1-Macro and F1-Micro metrics on the Movies and Amazon datasets. Across all settings, the performance curves exhibit limited variation, indicating that HAAM is not overly sensitive to the choice of K or M. For Movies, both metrics show mild fluctuations. The best performance is typically achieved with lower filter orders ($K \in \{2, 3\}$), particularly for $M = 64$, where a clear peak is observed at $K = 2$ followed by a slight decline. In contrast, Amazon displays a remarkably consistent behavior across all configurations. The performance remains stable for moderate filter orders, with a slight optimal consensus around $K = 3$ for most embedding dimensions. Overall, these results demonstrate that HAAM maintains stable performance over a wide range of hyperparameter configurations. This suggests that the proposed model does not rely on tedious fine-tuning.

3.6 Spectral Filter Responses

To provide an interpretable view of the spectral behavior learned by HAAM, we visualize the frequency responses of the learned low-pass $f_d^{\mathcal{L}}(\lambda)$, high-pass $f_d^{\mathcal{H}}(\lambda)$, and composed filters over the normalized Laplacian spectrum $\lambda \in [0, 2]$. These responses illustrate how different portions of the spectrum are emphasized or attenuated by the model across datasets.

As shown in Fig. 3, the low-pass and high-pass components exhibit complementary trends. Their composition, $f_d^{\mathcal{L}}(\lambda) \cdot f_d^{\mathcal{H}}(\lambda)$, yields a dataset-dependent spectral profile. On Amazon (Fig. 3b), the composed response remains relatively stable across the spectrum, indicating a balanced integration of low- and

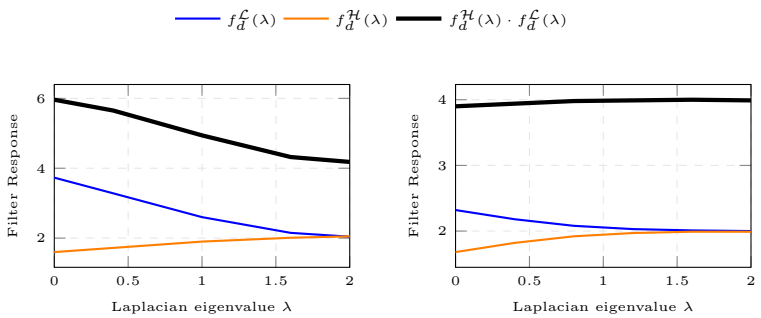

(a) Movies Filter Response. (b) Amazon Filter Response.

Fig. 3: Spectral frequency responses of the learned low-pass, high-pass, and composed filters on Movies and Amazon. The responses are evaluated over the spectrum of the normalized Laplacian. The plots illustrate how the model distributes spectral emphasis across frequencies and how the composed filter yields a stable behavior while combining low- and high-frequency components.

high-frequency signals. In contrast, on Movies (Fig. 3a), the composed response places more emphasis on lower frequencies, reflecting a stronger contribution of smoothing effects, consistent with the higher homophily of this dataset, while still preserving non-negligible high-frequency information. Overall, these spectral patterns demonstrate the ability of HAAM to flexibly adapt its propagation behavior to different homophily regimes. Accordingly, the visualized filter responses provide insight into how different frequency components are emphasized, reflecting the role of homophilic and heterophilic interactions across relations.

4 Conclusion

This paper introduced HAAM, an adaptive framework for node classification in settings exhibiting both homophilic and heterophilic relations. Unlike many existing approaches that implicitly rely on homophily, HAAM explicitly models diverse relational patterns through the combination of low-pass and high-pass Chebyshev filters, dimension-specific compatibility matrices, and a sparse consensus formulation. This design enables the model to adapt to heterogeneous connectivity structures while reducing over-reliance on any single relational view.

Several directions for future work remain promising. One direction is to extend the framework to multiplex graphs with a larger number of dimensions, which may require more expressive aggregation or geometric representations to preserve structural information at scale. Another direction is to investigate unsupervised or self-supervised learning settings for transferable node representations beyond semi-supervised classification. Extending HAAM to dynamic multiplex graphs or graphs involving multiple node types also constitutes an important direction for future work, since the current framework focuses on static multiplex graphs with homogeneous nodes. We believe these extensions could further strengthen the applicability of reliable multiplex graph learning methods.

References

1. M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, "Multilayer networks," *Journal of Complex Networks*, vol. 2, no. 3, pp. 203–271, 2014.
2. M. Berlingerio, M. Coscia, F. Giannotti, A. Monreale, and D. Pedreschi, "Multidimensional networks: foundations of structural analysis," *World Wide Web*, vol. 16, no. 5, pp. 567–593, 2013.
3. R. Oughtred, J. Rust, C. Chang, B.-J. Breitkreutz, C. Stark, A. Willems, L. Boucher, G. Leung, N. Kolas, F. Zhang *et al.*, "The biogrid database: A comprehensive biomedical resource of curated protein, genetic, and chemical interactions," *Protein Science*, vol. 30, no. 1, pp. 187–200, 2021.
4. J. Sun, Y. Zhang, C. Ma, M. Coates, H. Guo, R. Tang, and X. He, "Multi-graph convolution collaborative filtering," in *International Conference on Data Mining*. IEEE, 2019, pp. 1306–1311.
5. K. Abdous, N. Mrabah, and M. Bouguessa, "Hierarchical aggregations for high-dimensional multiplex graph embedding," *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 4, pp. 1624–1637, 2024.
6. H. Sun, X. Li, Z. Wu, D. Su, R.-H. Li, and G. Wang, "Breaking the entanglement of homophily and heterophily in semi-supervised node classification," in *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 2024, pp. 2379–2392.
7. J. Zhu, Y. Yan, L. Zhao, M. Heimann, L. Akoglu, and D. Koutra, "Beyond homophily in graph neural networks: Current limitations and effective designs," *Advances in Neural Information Processing Systems*, vol. 33, pp. 7793–7804, 2020.
8. S. Luan, C. Hua, Q. Lu, J. Zhu, M. Zhao, S. Zhang, X.-W. Chang, and D. Precup, "Revisiting heterophily for graph neural networks," *Advances in Neural Information Processing Systems*, vol. 35, pp. 1362–1375, 2022.
9. O. Barranco, C. Lozares, and D. Muntanyola-Saura, "Heterophily in social groups formation: a social network analysis," *Quality & Quantity*, vol. 53, no. 2, pp. 599–619, 2019.
10. D. Lim, F. Hohne, X. Li, S. L. Huang, V. Gupta, O. Bhalerao, and S. N. Lim, "Large scale learning on non-homophilous graphs: New benchmarks and strong simple methods," *Advances in Neural Information Processing Systems*, vol. 34, pp. 20 887–20 902, 2021.
11. X. Zheng, Y. Wang, Y. Liu, M. Li, M. Zhang, D. Jin, P. S. Yu, and S. Pan, "Graph neural networks for graphs with heterophily: A survey," *arXiv preprint arXiv:2202.07082*, 2022.
12. J. Zhu, R. A. Rossi, A. Rao, T. Mai, N. Lipka, N. K. Ahmed, and D. Koutra, "Graph neural networks with heterophily," in *AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, 2021, pp. 11 168–11 176.
13. O. Boutemine and M. Bouguessa, "Mining community structures in multidimensional networks," *ACM Transactions on Knowledge Discovery from Data*, vol. 11, no. 4, pp. 1–36, 2017.
14. B. Jing, S. Feng, Y. Xiang, X. Chen, Y. Chen, and H. Tong, "X-goal: multiplex heterogeneous graph prototypical contrastive learning," in *International Conference on Information & Knowledge Management*, 2022, pp. 894–904.
15. B. Jing, C. Park, and H. Tong, "Hdmi: High-order deep multiplex infomax," in *The Web Conference*, 2021, pp. 2414–2424.
16. A. Mitra, P. Vijayan, R. Sanasam, D. Goswami, S. Parthasarathy, and B. Ravindran, "Semi-supervised deep learning for multiplex networks," in *International Conference on Knowledge Discovery and Data Mining*, 2021, pp. 1234–1244.

17. L. Pio-Lopez, A. Valdeolivas, L. Tichit, É. Remy, and A. Baudot, "Multiverse: a multiplex and multiplex-heterogeneous network embedding approach," *Scientific Reports*, vol. 11, no. 1, pp. 1–20, 2021.
18. Y. Sadikaj, J. Rass, Y. Velaj, and C. Plant, "Semi-supervised embedding of attributed multiplex networks," in *The Web Conference*, 2023, pp. 578–587.
19. E. Pan and Z. Kang, "Beyond homophily: Reconstructing structure for graph-agnostic clustering," in *International Conference on Machine Learning*. PMLR, 2023, pp. 26 868–26 877.
20. M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," *Advances on Neural Information Processing Systems*, vol. 29, p. 3844–3852, 2016.
21. Y. Liu, Y. Zheng, D. Zhang, V. C. Lee, and S. Pan, "Beyond smoothing: Unsupervised graph representation learning with edge heterophily discriminating," in *AAAI Conference on Artificial Intelligence*, vol. 37, no. 4, 2023, pp. 4516–4524.
22. J. Chen, R. Lei, and Z. Wei, "Polygcl: Graph contrastive learning via learnable spectral polynomial filters," in *The Twelfth International Conference on Learning Representations*, 2024.
23. R. Duan, M. Guang, J. Wang, C. Yan, H. Qi, W. Su, C. Tian, and H. Yang, "Unifying homophily and heterophily for spectral graph neural networks via triple filter ensembles," *Advances in Neural Information Processing Systems*, vol. 37, pp. 93 540–93 567, 2024.
24. K. Abdous, N. Mrabah, and M. Bouguessa, "Modeling heterophily in multiplex graphs: An adaptive approach for node classification," *Expert Systems with Applications*, vol. 323, p. 132374, 2026. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S095741742601287X
25. Y. Cen, X. Zou, J. Zhang, H. Yang, J. Zhou, and J. Tang, "Representation learning for attributed multiplex heterogeneous network," in *International Conference on Knowledge Discovery and Data Mining*, 2019, pp. 1358–1368.
26. Y. Ma, S. Wang, C. C. Aggarwal, D. Yin, and J. Tang, "Multi-dimensional graph convolutional networks," in *International Conference on Data Mining*. SIAM, 2019, pp. 657–665.
27. C. Park, D. Kim, J. Han, and H. Yu, "Unsupervised attributed multiplex network embedding," in *AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 5371–5378.
28. Y. Mo, Y. Chen, Y. Lei, L. Peng, X. Shi, C. Yuan, and X. Zhu, "Multiplex graph representation learning via dual correlation reduction," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 12, pp. 12 814–12 827, 2023.
29. Y. Mo, Y. Lei, J. Shen, X. Shi, H. T. Shen, and X. Zhu, "Disentangled multiplex graph representation learning," in *International Conference on Machine Learning*, 2023, pp. 24 983–25 005.
30. Z. Shen, S. Wang, and Z. Kang, "Beyond redundancy: Information-aware unsupervised multiplex graph structure learning," *Advances in Neural Information Processing Systems*, vol. 37, pp. 31 629–31 658, 2024.
31. J. Tang, J. Sun, C. Wang, and Z. Yang, "Social influence analysis in large-scale networks," in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2009, pp. 807–816.
32. R. He and J. McAuley, "Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering," in *The Web Conference*, 2016, pp. 507–517.