

When Hypergraph Meets Heterophily: New Benchmark Datasets and Baseline

Ming Li^{1,2}, Yongchun Gu¹, Yi Wang^{3,1*}, Yujie Fang^{3,1}, Lu Bai^{4*}, Xiaosheng Zhuang⁵, Pietro Lio⁶

¹Zhejiang Key Laboratory of Intelligent Education Technology and Application, Zhejiang Normal University

²Zhejiang Institute of Optoelectronics

³School of Computer Science and Technology, Zhejiang Normal University

⁴School of Artificial Intelligence, and Engineering Research Center of Intelligent Technology and Educational Application, Ministry of Education, Beijing Normal University

⁵Department of Mathematics, City University of Hong Kong

⁶Department of Computer Science and Technology, Cambridge University

{mingli, guyongchun, wangyi, yjfang}@zjnu.edu.cn, bailu@bnu.edu.cn, xzhuang7@cityu.edu.hk, pl219@cam.ac.uk

Abstract

Hypergraph neural networks (HNNs) have shown promise in handling tasks characterized by high-order correlations, achieving notable success across various applications. However, there has been limited focus on heterophilic hypergraph learning (HHL), in contrast to the increasing attention given to graph neural networks designed for graphs exhibiting heterophily. This paper aims to pave the way for HHL by addressing key gaps from multiple perspectives: measurement, dataset diversity, and baseline model development. First, we introduce metrics to quantify heterophily in hypergraphs, providing a numerical basis for assessing the homophily/heterophily ratio. Second, we develop diverse benchmark datasets across various real-world scenarios, facilitating comprehensive evaluations of existing HNNs and advancing research in HHL. Additionally, as a novel baseline model, we propose HyperUFG, a framelet-based HNN integrating both low-pass and high-pass filters. Extensive experiments conducted on synthetic and benchmark datasets highlight the challenges current HNNs face with heterophilic hypergraphs, while showcasing that HyperUFG performs competitively and often outperforms many existing models in such scenarios. Overall, our study underscores the urgent need for further exploration and development in this emerging field, with the potential to inspire and guide future research in HHL.

HHL Repository — <https://kellysylvia77.github.io/HHL>

Appendix — <https://mingli-ai.github.io/HHL.pdf>

1 Introduction

Hypergraph Neural Networks (HNNs) (Prokopchik, Benson, and Tudisco 2022; Chien et al. 2022; Duta et al. 2023; Wang et al. 2023; Li et al. 2024b) have emerged as powerful tools for capturing high-order correlations in complex data, extending the capabilities of traditional Graph Neural Networks (GNNs). By modeling relationships among multiple entities through hyperedges, HNNs are well-suited for various real-world applications, including recommendation systems, social network analysis, and bioinformatics (Antelmi

*Corresponding authors: Yi Wang, Lu Bai
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

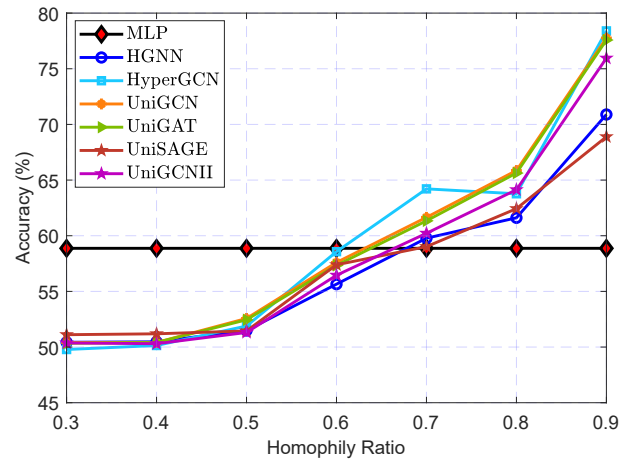


Figure 1: Performance comparison between MLP (a graph-agnostic method) and HNNs on synthetic datasets with varying homophily rates. The results show that MLP outperforms HNNs when the homophily rate is low (≤ 0.5), while HNNs excel as homophily increases.

et al. 2023; Kim et al. 2024). Despite their promise, most existing HNNs primarily focus on scenarios characterized by homophily, that is, the assumption that connected nodes share similar attributes or labels. This assumption, while effective in many cases, limits the generalizability of HNNs to settings where heterophily is predominant.

Heterophilic graph learning, where connected nodes often have dissimilar labels, has gained attention in the context of GNNs, leading to the development of models specifically designed to handle heterophilic graphs (Zhu et al. 2020; Zheng et al. 2022; Zhu et al. 2023; Gong et al. 2024; Luan et al. 2024). However, research on heterophilic hypergraph learning (HHL) remains limited, leaving a significant gap in the current literature. Our preliminary empirical investigations, as illustrated in Figure 1, reveal that existing classic HNNs, such as HGNN (Feng et al. 2019), HyperGCN (Yadati et al. 2019), and the UniGNN series (UniGCN/UniGAT/UniSAGE/UniGCNII) (Huang and Yang 2021), do not consistently outperform simpler models like Multi-Layer Perceptrons (MLPs) when applied to hypergraphs with low

homophily. This observation underscores the need to reconsider the design and application of HNNs in heterophilic settings.

This paper aims to address the challenges of heterophilic hypergraph learning (HHL) by identifying and exploring key gaps in the current body of research. First, we recognize the absence of robust metrics for quantifying homophily and heterophily in hypergraphs (**Q1: conceptual perspective**). To address this, we introduce new metrics to measure homophily/heterophily ratios, providing a more direct and systematic way to evaluate hypergraph structures. Second, we observe a scarcity of high-quality datasets that represent non-homophilous applications (**Q2: data source perspective**). In response, we develop a diverse collection of benchmark datasets that capture the complexity of real-world heterophilic hypergraphs, facilitating comprehensive evaluations of existing HNNs and driving advancements in HHL. Finally, to the best of our knowledge, few works have focused on designing HNNs specifically tailored for HHL (**Q3: methodological perspective**). To address this, we propose HyperUFG, a novel framelet-based HNN that integrates both low-pass and high-pass filters to effectively tackle the heterophily challenge. HyperUFG is designed as a strong yet simple baseline, demonstrating significant advantages over existing HNNs in heterophilic scenarios. Through extensive experiments on synthetic and benchmark datasets, we show that HyperUFG not only outperforms current models but also sets a new standard for future research in HHL.

Our contributions in this work are summarized as follows:

Addressing Q1: Defining Hypergraph Homophily/Heterophily Ratios. We introduce two methods for measuring homophily and heterophily in hypergraphs, focusing on the perspectives of hypernodes and hyperedges. Using these methods, we quantify the homophily of both existing and newly proposed hypergraphs, providing a robust foundation for evaluating hypergraph structures.

Addressing Q2: Constructing Heterophilic Hypergraph Benchmarks. We develop a diverse collection of high-quality non-homophilic hypergraph datasets, characterized by their large size, broad application scope, and the intricate interplay between labels and topology. These datasets enable rigorous experimentation across a range of existing works on HNNs, advancing the evaluation and understanding of HHL.

Addressing Q3: Introducing a Novel Baseline for HHL. We propose HyperUFG, a novel HNN framework based on framelet transforms, specifically designed to tackle the heterophily challenge. Comparative experiments with eight HNN models demonstrate that HyperUFG offers a distinct advantage in handling heterophilic data, establishing it as a strong baseline for future research in HHL.

2 Challenge for Heterophilic Hypergraphs: Observations on Synthetic Data

To better understand the issues and challenges associated with heterophily in hypergraphs, this section explores the performance of classical HNNs on synthetic datasets with varying homophily levels. These datasets include both ho-

mophilic and heterophilic hypergraphs. We begin by defining appropriate measures of hypergraph homophily, followed by an introduction to the hypergraph generation process. Subsequently, we analyze the heterophily observed in the synthetic data and conclude with a discussion on the limitations of existing HNN models, emphasizing the challenges posed by hypergraph heterophily.

2.1 Hypergraph Homophily Measures

Existing studies (Wang et al. 2023; Duta et al. 2023) often apply traditional graph homophily measures, such as the CE homophily rate (Pei et al. 2020), directly to hypergraphs. However, this approach is inadequate as it fails to capture the complex structural characteristics of hypergraphs, thereby limiting our ability to accurately assess the relationship between specific information and homophily/heterophily ratios (Veldt, Benson, and Kleinberg 2023; Telyatnikov et al. 2023). To address this gap, we introduce, for the first time, metrics specifically designed to measure homophily ratios from both hyperedge and hypernode perspectives (see Definitions 1 and 2). To the best of our knowledge, these definitions are novel and have not been previously explored in the literature. Importantly, these metrics are not only tailored for hypergraphs but are also versatile enough to apply to traditional graphs, which can be considered a special case of hypergraphs where all hyperedges contain only two nodes (i.e., a 2-uniform hypergraph). When applied to traditional graphs, our methods produce results that align with established node and edge homophily rates, further demonstrating their flexibility and generalizability.

Definition 1 [Hyperedge Homophily] This ratio is the average proportion of node pairs within hyperedges that belong to the same class. It is defined as:

$$\mathcal{H}_{edge} = \frac{1}{|\mathcal{E}|} \sum_{e_j \in \mathcal{E}} \frac{|\{(u, v) \in e_j | \mathbb{1}(y_u = y_v)\}|}{C_{n_j}^2}, \quad (1)$$

where $\mathbb{1}(\cdot)$ is the indicator function (i.e., $\mathbb{1}(\cdot) = 1$ if the condition is true, and $\mathbb{1}(\cdot) = 0$ otherwise). Here, e_j ($j = 1, 2, \dots, |\mathcal{E}|$) denotes the j -th hyperedge, and n_j represents the number of nodes within e_j .

Definition 2 [Node Homophily] represents the average proportion of node pairs that belong to the same class across all hyperedges in which a given node is involved. It is defined as:

$$\mathcal{H}_{node} = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \frac{1}{|\mathcal{R}_v|} \sum_{e_j \in \mathcal{R}_v} \frac{|\{(u, v) \in e_j | \mathbb{1}(y_u = y_v)\}|}{n_j}, \quad (2)$$

where $\mathcal{R}_v = \sum_{e_j \in \mathcal{E}} \{e_j | v \in e_j\}$ is the set of all hyperedges containing node v , and $\mathbb{1}(\cdot)$ is the indicator function as defined previously. Here, e_j ($j = 1, 2, \dots, |\mathcal{E}|$) denotes the j -th hyperedge, and n_j represents the number of other nodes within e_j .

In essence, the degree of homophily in a hypergraph is determined by the proportion of samples from the same category within its hyperedges. A hypergraph is considered homophilic if its hyperedges predominantly consist of samples

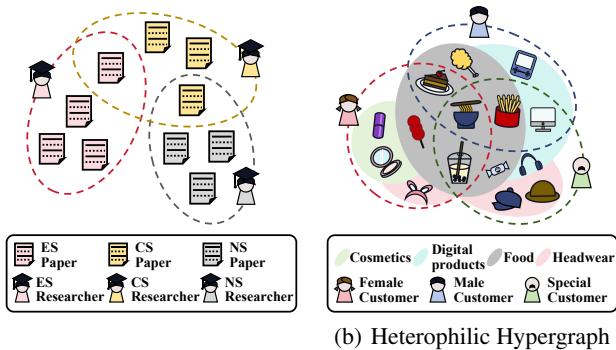


Figure 2: Schematic representation of homophilic and heterophilic hypergraphs: (a) illustrates a citation network based on co-authorship, and (b) depicts a product network based on co-purchase relationships.

from the same category, and heterophilic if this is not the case. To illustrate these concepts, we present two toy examples in Figure 2. In particular, Figure 2 (a) illustrates a homophilic hypergraph, using a citation network as an example. In such networks, hyperedges often consist of samples from the same research category, reflecting the tendency of researchers to focus on their specific domains. Although interdisciplinary collaborations have increased, leading to co-authored papers across different fields, these remain relatively rare. In contrast, Figure 2 illustrates a heterophilic hypergraph, using a product purchase network as an example. This network is characterized by heterophily, where the pursuit of a diverse lifestyle leads individuals to make a wide variety of purchases, resulting in hyperedges filled with diverse samples. Additionally, unique user profiles with specialized purchasing behaviors contribute further to the network’s diversity.

For a more comprehensive explanation of these definitions, please refer to **Appendix C**.

2.2 Synthetic Hypergraphs

This section provides details on the utilization of the Hypergraph Stochastic Block Model (HSBM) (Cole and Zhu 2020) to generate r -uniform hypergraphs, as described in Definition 3. The process of generating the necessary hypernodes and hyperedges for a hypergraph involves two specific steps:

- **Defining Nodes and Clusters:** We begin by specifying the total number of hypergraph nodes, n , and the number of clusters, z . Each cluster is composed of $s = n/z$ nodes, where n is assumed to be divisible by z ;
- **Defining Hyperedges:** A hyperedge, denoted as $e = \{v_1, v_2, \dots, v_r\}$, connects r nodes. These r nodes belong to the same cluster with a probability p .

Definition 3 [Hypergraph Stochastic Block Model (HSBM)] Let $\mathcal{C} = \{C_1, C_2, \dots, C_z\}$ represent a partition of the set $[n]$ into z clusters, each containing $s = n/z$ nodes (assuming n is divisible by z). Each cluster C_i (for $1 \leq i \leq z$) is a subset of $[n]$. For constants $0 \leq q < p < 1$, the r -uniform HSBM is defined as follows:

Datasets	Congress	Senate	House	Walmart
Hypernodes, $ \mathcal{V} $	1, 718	282	1, 290	88, 860
CE homophily	0.555	0.498	0.509	0.530
Edge hom., \mathcal{H}_{edge}	0.651	0.464	0.485	0.595
Node hom., \mathcal{H}_{node}	0.659	0.479	0.505	0.500

Table 1: Homophily ratios for existing datasets.

For any set of r distinct vertices v_1, v_2, \dots, v_r , a hyperedge $e = v_1, v_2, \dots, v_r$ is generated with probability p if the vertices v_1, v_2, \dots, v_r belong to the same cluster in \mathcal{C} . If the vertices are in different clusters, the hyperedge e is generated with probability q . This distribution of random hypergraphs is denoted as $\mathcal{H} = (n, r, \mathcal{C}, p, q)$. When $r = 2$, this model corresponds to the stochastic block models for random graphs.

By adjusting the parameters p and q , we can generate hypergraphs with varying homophily rates. For this study, we focus on hypergraphs where each hyperedge consists of 3 nodes ($r = 3$) and there are 2 distinct sample labels ($z = 2$). Under these conditions, the HSBM can generate hypergraphs with a minimum homophily rate of 0.3. Despite this constraint, the seven synthetic hypergraphs with different homophily rates that we generated are sufficient to examine and validate the challenges that HNNs face in handling heterophily.

2.3 Empirical Observations

To assess whether HNNs encounter challenges with heterophily, we conducted a statistical analysis of several classical HNNs across seven synthetic datasets with varying homophily ratios. The performance trends are depicted in Figure 1. From the observed trends, we derive two key findings:

Observation 1: HNN performance improves as the homophily ratio increases.

Observation 2: Compared to the graph-agnostic method MLP, HNNs generally underperform on hypergraphs with low homophily ratios ($\mathcal{H}_{edge} \leq 0.5$) but outperform MLP on hypergraphs with higher homophily ratios ($\mathcal{H}_{edge} > 0.5$).

These findings suggest that the prevalent HNN models indeed face challenges when dealing with heterophilic hypergraphs, an issue that has not been adequately addressed in the current research. In the following discussion, we explore the limitations of existing hypergraph representation learning approaches in the context of heterophily and highlight the key issues that need to be addressed in this field.

Motivation for Developing New Benchmarks. Recent studies (Wang et al. 2023; Tang et al. 2024; Duta et al. 2023) have utilized the CE homophily rate, i.e., a metric typically applied to graphs (Pei et al. 2020; Zhu et al. 2020; Wu et al. 2024), to assess homophily in hypergraphs. These studies tested their methods on four datasets: Congress, Senate, Walmart, and House, which are presumed to have low homophily rates. However, upon closer inspection, nearly all the CE homophily ratios for these datasets are larger than 0.5 (see Table 1 for details), which is ambiguous and therefore difficult to know *whether they are homophilic rather than truly heterophilic*. This ambiguity raises concerns about the suitability of these datasets as benchmarks for evaluating the

performance of HNNs under heterophilic conditions. Furthermore, these datasets are limited by their relatively small size and lack of diversity, as they cover similar topics. Therefore, there is a critical need to develop and introduce higher-quality datasets that can more accurately capture the heterophilic nature of hypergraphs, thereby enhancing the robustness and relevance of future research in this domain.

3 New Benchmarks of Heterophilic Hypergraphs

Motivated by the observations highlighted in the preceding section, we develop a set of novel datasets specifically tailored to evaluate the efficacy of HNNs within heterophilic contexts. In particular, our objective is to compile datasets that meet the following basic criteria:

- The datasets are expected to demonstrate heterophily, assessed through the metrics of hypergraph homophily. Detailed formal definitions can be found in Definitions 1 and 2 (see Section 2).
- The underlying hypergraph structure is expected to significantly impact task performance, highlighting the pivotal role of hypergraphs in node classification tasks. To validate this hypothesis, we perform a comparative analysis between graph-agnostic MLPs and HNNs. Our hypothesis posits that HNNs will demonstrate superior performance compared to MLPs.
- The datasets should encompass a breadth of diversity, originating from a variety of domains and exhibiting a range of distinct structural attributes. Consequently, comprehensive profiles detailing key dataset characteristics are provided for each.
- Graph sizes have been carefully chosen to strike a balance: sufficiently large to yield statistically significant results, yet manageable for evaluating heterophily-specific models outlined in the existing literature, many of which face scalability constraints. Thus, our dataset selection is confined to graphs ranging from 10K to 30K nodes.

For each dataset under consideration, we provide a detailed analysis of its fundamental characteristics, covering essential metrics such as total nodes, hyperedges, feature sets, and class labels. Additionally, we present relevant graph statistics, including hyperedge homophily (defined in Definition 1) and hypernode homophily (outlined in Definition 2). These statistical records are summarized in Table 2.

Actor (co-occurrence). The actor co-occurrence network is derived from the broader movie-actor-director-writer network¹, capturing intricate actor co-occurrences within films. This network reveals complex relationships among films, directors, actors, and writers distilled from heterogeneous information networks. Nodes represent individuals involved in film production (actors, directors, and writers), while hyperedges signify collaborations among all individuals on a single film project. Node attributes are based on keywords extracted from Wikipedia, and labels indicate the specific roles of these individuals within the network.

¹<https://www.aminer.org/lab-datasets/soinf/>

Datasets	Actor	Amazon-ratings	Twitch-gamers	Pokec
Hypernodes, $ \mathcal{V} $	16, 255	22, 299	16, 812	14, 998
Hyperedges, $ \mathcal{E} $	10, 164	2, 090	2, 627	2, 406
Avg. hyperedge size	5.4 ± 2.7	3.1 ± 0.6	6.2 ± 3.4	2.3 ± 0.7
Features, d	50	111	7	65
Classes, c	3	5	2	2
Node hom., \mathcal{H}_{node}	0.4815	0.4805	0.4893	0.4952
Edge hom., \mathcal{H}_{edge}	0.4675	0.3677	0.4857	0.4529

Table 2: Statistics of the newly developed datasets.

Amazon-ratings (co-purchasing). This dataset is derived from the Amazon product co-purchasing network meta-data², sourced from SNAP Datasets (Jure 2014). It consists of nodes representing a diverse array of products such as books, music CDs, DVDs, and VHS videotapes. Hyperedges in the network connect products frequently co-purchased by individual users. The dataset poses the challenge of predicting the average rating assigned by reviewers to each product, categorized into ten distinct classes. Node features are represented using the *Bag of Words* technique (Juluru et al. 2021), which embeds textual content from product descriptions.

Twitch-gamers (co-create). The Twitch-Gamers network³ is a connected, undirected graph that models relationships among accounts on the Twitch streaming platform. Hyperedges are constructed to represent co-occurrences of users created within the same timeframe. Each node corresponds to a unique Twitch account, with hyperedges formed between accounts that mutually follow each other. Node features encompass a range of attributes, including view counts, account creation and update timestamps, language preferences, lifetime activity duration, and status indicating whether the account is inactive. The primary objective of the binary classification task is to predict the presence of explicit content on each channel.

Pokec (co-friendship). Pokec⁴, the predominant online social networking platform in Slovakia, provides a robust foundation for constructing social networks aimed at analyzing various social attributes. In this network, nodes represent individual users, while hyperedges encapsulate each user’s complete set of friends. Each node is categorized by labels indicating the reported gender of the users. To enrich the dataset, we extract node features from a comprehensive array of profile information, including age, hobbies, interests, education level, geographical region, and registration timestamp.

More detailed descriptions of these four datasets can be found in **Appendix D**.

4 New Baseline for Heterophilic HNNs

Technically, many advanced GNNs for heterophilic graphs are developed by a special design of both low-pass and high-pass filters (or more generally, homophily and heterophily filters) (Bo et al. 2021; Li, Pan, and Kang 2024; Li et al. 2024a). That is to say, both low-pass and high-pass filters play a role in heterophilic graph learning, which we assume also applies to the hypergraph scenario. Motivated by this in-

²<https://snap.stanford.edu/data/amazon-meta.html>

³http://snap.stanford.edu/data/twitch_gamers.html

⁴<https://snap.stanford.edu/data/soc-Pokec.html>

sight, we introduce a novel spectral-based hypergraph neural network for heterophilic hypergraphs.

Consider a hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consisting of a vertex set \mathcal{V} with $N = |\mathcal{V}|$ vertices, and a hyperedge set \mathcal{E} with $M = |\mathcal{E}|$ hyperedges. The vertex feature matrix is denoted by $\mathbf{X} \in \mathbb{R}^{N \times m}$. Let \mathcal{L}_h denote the hypergraph Laplacian, and $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_N]$ be the matrix of the eigenvectors of \mathcal{L}_h , with $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$ as the diagonal matrix of the eigenvalues.

Framelets over the hypergraph are constructed using a set of scaling functions $\xi = \{\delta; \eta^{(1)}, \dots, \eta^{(k)}\}$, associated with a *filter bank* $\rho = \{a; b^{(1)}, \dots, b^{(k)}\}$. These functions satisfy $\widehat{\delta}(2\vartheta) = \widehat{a}(\vartheta)\widehat{\delta}(\vartheta)$ and $\widehat{\eta}^{(r)}(2\vartheta) = \widehat{b}^{(r)}(\vartheta)\widehat{\delta}(\vartheta)$ for any $\vartheta \in \mathbb{R}$, where $\widehat{f}(\vartheta)$ denotes the Fourier transform of f , defined by $\widehat{f}(\vartheta) := \sum_{k \in \mathbb{Z}} h(k)e^{-2\pi i k \vartheta}$. Here, k denotes the number of high-pass filters involved in the construction of the framelets.

The functions $\Phi_{j,p}(\nu)$ and $\Psi_{j,p}^r(\nu)$ represent the low-pass and high-pass framelets at node ν associated with node t at scale level $j \in \{1, \dots, J\}$ respectively, defined as follows:

$$\text{Low-pass: } \Phi_{j,t}(\nu) = \sum_{p=1}^N \widehat{\delta}\left(\frac{\lambda_p}{2^j}\right) u_p(t)u_p(\nu), \quad (3)$$

$$\text{High-pass: } \Psi_{j,t}^r(\nu) = \sum_{p=1}^N \widehat{\eta}^{(r)}\left(\frac{\lambda_p}{2^j}\right) u_p(t)u_p(\nu), \quad (4)$$

$r = 1, \dots, k,$

where $u_p(t)$ represents the t -th component of \mathbf{u}_p .

These low-pass (i.e., Eq. (3)) and high-pass framelets (i.e., Eq. (4)) capture coarse-grained and fine-grained information from hypergraph signals. The framelet coefficients $V_0, W_j^r \in \mathbb{R}^{N \times m}$ are respectively the low-pass and high-pass coefficients, defined as the inner product of framelets and the hypergraph signal $\mathbf{X} \in \mathbb{R}^{N \times m}$:

$$V_0 = \langle \Phi_{0,\cdot}, \mathbf{X} \rangle = \mathbf{U}^\top \widehat{\delta}\left(\frac{\Lambda}{2}\right) \mathbf{U} \mathbf{X}, \quad (5)$$

$$W_j^r = \langle \Psi_{j,\cdot}^r, \mathbf{X} \rangle = \mathbf{U}^\top \widehat{\eta}^{(r)}\left(\frac{\Lambda}{2^{j+1}}\right) \mathbf{U} \mathbf{X}. \quad (6)$$

Theoretically, the tightness of the framelet system is ensured by the condition:

$$\left| \widehat{\delta}\left(\frac{\lambda_p}{2^j}\right) \right|^2 + \sum_{r=1}^k \left| \widehat{\eta}^{(r)}\left(\frac{\lambda_p}{2^j}\right) \right|^2 = 1. \quad (7)$$

Let $\mathcal{F}_{r,j}$ denote the decomposition operators such that $V_0 = \mathcal{F}_{0,J} \mathbf{X}$ and $W_j^r = \mathcal{F}_{r,j} \mathbf{X}$. Referring to Eq. (5) and Eq. (6), the framelet transform operators for decomposition are defined as follows:

$$\begin{aligned} \mathcal{F}_{0,J} &= \mathbf{U}^\top \widehat{a}(2^{-C+J-1}\Lambda) \dots \widehat{a}(2^{-C}\Lambda) \mathbf{U} := \mathbf{U}^\top \Lambda_{0,J} \mathbf{U}, \\ \mathcal{F}_{r,1} &= \mathbf{U}^\top \widehat{b}^{(r)}(2^{-C}\Lambda) \mathbf{U} := \mathbf{U}^\top \Lambda_{r,1} \mathbf{U}, \\ \mathcal{F}_{r,j} &= \mathbf{U}^\top \widehat{b}^{(r)}(2^{-C+j-1}\Lambda) \widehat{a}(2^{-C+j-2}\Lambda) \dots \widehat{a}(2^{-C}\Lambda) \mathbf{U} \\ &:= \mathbf{U}^\top \Lambda_{r,j} \mathbf{U}. \end{aligned}$$

Here, C is chosen to be sufficiently large such that the largest eigenvalue λ_{max} satisfies $\lambda_{max} \leq 2^C \pi$.

This tightness (Eq. (7)) guarantees the invertibility of framelet decomposition and reconstruction, that is,

$$\mathcal{F}_{0,J}^\top \mathcal{F}_{0,J} \mathbf{X} + \sum_{r,j} \mathcal{F}_{r,j}^\top \mathcal{F}_{r,j} \mathbf{X} = \mathbf{X}.$$

The incorporation of Haar-type filters ensures an efficient multi-scale system computationally (Dong 2017; Zheng et al. 2021). For example, for a framelet system with two scale levels ($j = 1, 2$) and one high-pass filter ($r = 1$), the filters are constructed as described: $\widehat{\delta}(\frac{\Lambda}{2}) = \cos(\frac{\Lambda}{8})\cos(\frac{\Lambda}{16})$, $\widehat{\eta}(\frac{\Lambda}{2}) = \sin(\frac{\Lambda}{8})\cos(\frac{\Lambda}{16})$ and $\widehat{\eta}(\frac{\Lambda}{4}) = \sin(\frac{\Lambda}{16})$.

To tackle the computational complexities associated with the eigendecomposition of the hypergraph Laplacian, an approximation strategy employing Chebyshev polynomials is adopted, inspired by the methods outlined in (Dong 2017). Specifically, fixed-degree Chebyshev polynomials $\mathcal{T}_0, \dots, \mathcal{T}_s$ are utilized, where filters approximate where filters $a \approx \mathcal{T}_0$ and $b^{(r)} \approx \mathcal{T}_r$. Consequently, the operators defined in Eq. (4) are approximated as follows:

$$\begin{aligned} \mathcal{F}_{0,J} &\approx \mathbf{U}^\top \mathcal{T}_0(2^{-C+J-1}\Lambda) \dots \mathcal{T}_0(2^{-C}\Lambda) \mathbf{U} \\ &= \mathcal{T}_0(2^{-C+J-1}\mathcal{L}) \dots \mathcal{T}_0(2^{-C}\mathcal{L}), \end{aligned} \quad (8)$$

$$\mathcal{F}_{r,1} \approx \mathbf{U}^\top \mathcal{T}_r(2^{-C}\Lambda) \mathbf{U} = \mathcal{T}_r(2^{-C}\mathcal{L}), \quad (9)$$

$$\begin{aligned} \mathcal{F}_{r,j} &\approx \mathbf{U}^\top \mathcal{T}_r(2^{-C+j-1}\Lambda) \mathcal{T}_0(2^{-C+j-2}\Lambda) \dots \mathcal{T}_0(2^{-C}\Lambda) \mathbf{U} \\ &= \mathcal{T}_r(2^{-C+j-1}\mathcal{L}) \mathcal{T}_0(2^{-C+j-2}\mathcal{L}) \dots \mathcal{T}_0(2^{-C}\mathcal{L}). \end{aligned} \quad (10)$$

Based on the hypergraph framelets transform system derived in Eqs. (8, 9, 10), we introduce a novel framelet-based spectral hypergraph convolution termed HyperUFG. This formulation incorporates initial residual and identity mapping techniques as introduced in (Chen et al. 2020). The convolution operation is defined as follows:

$$\begin{aligned} \mathbf{X}^{(\ell+1)} &= \sigma \left(\left((1 - \alpha_\ell) \sum_{(r,j) \in \Gamma} \mathcal{F}_{r,j}^\top \text{diag}(\theta_{r,j}) \mathcal{F}_{r,j} \mathbf{X}^{(\ell)} + \alpha_\ell \mathbf{X}^{(0)} \right) \right. \\ &\quad \left. \cdot \left((1 - \beta_\ell) \mathbf{I} + \beta_\ell \Theta^{(\ell)} \right) \right), \end{aligned}$$

where $\theta_{r,j} \in \mathbb{R}^N$ are the learnable filter, and $\Gamma = \{(r, j) : r = 1, \dots, R, j = 0, 1, \dots, J\} \cup \{(0, J)\}$ represents the index set for all framelet decomposition matrices.

Specifically, HyperUFG employs both low-pass and high-pass filters within the hypergraph convolution layer, effectively addressing the potential limitations of spatial-based message-passing that can hinder effective neighbor aggregation in HHL scenarios. The low-pass components enhance ‘similarity’ among homophilic neighbors, while the high-pass components emphasize ‘distinguishable’ information among heterophilic neighbors.

5 Experiments

5.1 Datasets and Experimental Settings

This section explores the potential challenges posed by HHL on the four newly developed heterophilic hypergraph benchmarks and evaluates the effectiveness of HyperUFG. We

Methods Edge hom. ratio, \mathcal{H}_{edge}	Cora 0.7462	Citeseer 0.6814	Pubmed 0.7765	Cora-CA 0.7797	DBLP-CA 0.8656	Rank
MLP	75.16 ± 1.41	71.71 ± 1.01	87.20 ± 0.34	75.17 ± 1.41	84.37 ± 0.33	10
HGNN	79.39 ± 1.36	72.45 ± 1.16	86.44 ± 0.44	82.64 ± 1.65	91.03 ± 0.20	6
HyperGCN	78.45 ± 1.26	71.28 ± 0.82	82.84 ± 8.67	79.48 ± 2.08	89.38 ± 0.25	9
UniGCNII	78.81 ± 1.05	73.05 ± 2.21	88.25 ± 0.40	83.60 ± 1.14	91.69 ± 0.19	4
HyperND	79.20 ± 1.14	72.62 ± 1.49	86.68 ± 0.43	80.62 ± 1.32	90.35 ± 0.26	8
AllDeepSets	76.88 ± 1.80	70.83 ± 1.63	88.75 ± 0.33	81.97 ± 1.50	91.27 ± 0.27	7
AllSetTransformer	78.58 ± 1.47	73.08 ± 1.20	88.72 ± 0.37	83.63 ± 1.47	91.53 ± 0.23	5
ED-HNN	80.31 ± 1.35	73.70 ± 1.38	89.03 ± 0.53	83.97 ± 1.55	91.90 ± 0.19	2nd
SheafHyperGNN	81.30 ± 1.70	74.71 ± 1.23	87.68 ± 0.60	85.52 ± 1.28	91.59 ± 0.24	3rd
HyperUFG	81.51 ± 0.99	74.72 ± 2.10	88.73 ± 0.42	85.18 ± 0.69	91.67 ± 0.31	1st

Table 3: Comparative performance of graph-agnostic MLP and various HNNs on homophilic hypergraphs. The best-performing model is highlighted in lilac, the second-best in blue, and the third-best in grey.

Datasets Edge hom. ratio, \mathcal{H}_{edge}	Actor 0.4675	Amazon-ratings 0.3677	Twitch-gamers 0.4857	Pokec 0.4529	Senate 0.4642	House 0.4851	Rank
MLP	85.45 ± 1.21	26.70 ± 2.82	52.77 ± 1.81	56.92 ± 2.46	52.25 ± 5.17	51.86 ± 2.34	4
HGNN	74.47 ± 0.32	23.79 ± 0.24	51.88 ± 0.26	49.82 ± 0.27	48.59 ± 4.52	61.39 ± 2.96	8
HyperGCN	68.67 ± 4.38	22.53 ± 3.94	51.32 ± 1.02	52.43 ± 3.68	42.45 ± 3.67	48.32 ± 2.93	10
UniGCNII	80.48 ± 1.13	26.63 ± 1.32	50.84 ± 0.76	54.25 ± 2.70	49.30 ± 4.25	67.25 ± 2.57	7
HyperND	92.52 ± 0.81	26.08 ± 0.33	51.44 ± 0.67	55.94 ± 0.45	52.82 ± 3.20	51.70 ± 3.37	5
AllDeepSets	82.00 ± 2.33	18.60 ± 0.17	50.72 ± 0.96	51.11 ± 1.04	48.17 ± 5.67	67.82 ± 2.40	9
AllSetTransformer	83.39 ± 1.73	18.60 ± 0.17	50.45 ± 0.76	58.40 ± 0.42	51.83 ± 5.22	69.33 ± 2.20	6
ED-HNN	91.86 ± 0.43	26.21 ± 0.36	50.86 ± 0.88	59.11 ± 0.57	64.79 ± 5.14	72.45 ± 2.28	2nd
SheafHyperGNN	80.09 ± 2.45	26.93 ± 3.04	51.03 ± 0.76	55.34 ± 4.39	68.73 ± 4.68	73.84 ± 2.30	3rd
HyperUFG	89.32 ± 0.75	40.53 ± 2.25	52.35 ± 0.04	62.30 ± 0.12	67.61 ± 7.00	72.82 ± 2.22	1st

Table 4: Comparative performance of graph-agnostic MLP and various HNNs on heterophilic hypergraphs. The best-performing model is highlighted in lilac, the second-best in blue, and the third-best in grey.

perform extensive performance comparisons across various baselines for HHL, utilizing both the newly introduced datasets and seven existing benchmark datasets. To facilitate this, we categorize the datasets into two groups, i.e., homophily and heterophily, based on the homophily ratios \mathcal{H}_{node} and \mathcal{H}_{edge} (see Definitions 1 and 2 in Section 2). Specifically, datasets with homophily ratios exceeding 0.5 are classified as homophilic, while those with lower ratios are deemed heterophilic. For all new benchmarks, we utilize feature vectors, class labels, and ten random splits (40%/20%/40% of nodes per class for training/validation/testing, respectively). The dataset details are summarized in Tables 2 (for our developed datasets) and Table S-1 (in Appendix A) for the existing benchmark datasets. We compare HyperUFG with MLP, HGNN (Feng et al. 2019), HyperGCN (Yadati et al. 2019), UniGCNII (Huang and Yang 2021), HyperND (Prokopchik, Benson, and Tudisco 2022), AllDeepSets and AllSetTransformer (Chien et al. 2022), ED-HNN (Wang et al. 2023), and SheafHyperGNN (Duta et al. 2023). We train the model for a total of 1,000 epochs, employing early stopping with a patience threshold of 200 epochs. The experiments are conducted on a NVIDIA RTX A6000 GPU with 48GB of memory. The baseline results are reproduced using their publicly available code, with hyperparameters set according to the original papers. We utilize grid search to fine-tune the key hyperparameters

via the lightweight but powerful toolkit NNI (<https://nni.readthedocs.io/en/stable/>). Additional details can be found in Appendix A and Appendix B.

5.2 Results and Discussion

We conduct a thorough evaluation of HyperUFG’s performance across hypergraphs characterized by varying degrees of homophily, encompassing both homophilic and heterophilic hypergraphs. This evaluation includes a detailed comparative analysis of the performance of established HNNs specifically on heterophilic hypergraphs. The results of this comparison are presented in Table 4, which highlights the effectiveness and limitations of these models when applied to heterophilic structures. In addition to this, we extend our study by incorporating findings from peer-reviewed research on HNNs in the context of homophilic hypergraphs. These findings, drawn from (Wang et al. 2023), provide a benchmark for understanding the performance of HNNs under conditions of high homophily and are detailed in Table 3. By comparing the performance data from both homophilic and heterophilic settings, we gain valuable insights into how the structural properties of hypergraphs influence the efficacy of these neural networks.

The combined data from Tables 3 and 4 allow us to make several significant observations. These include variations in model performance depending on the degree of homophily

present in the hypergraph, the relative strengths and weaknesses of different HNN architectures in handling diverse hypergraph structures, and the implications of these findings for future research and application of HNNs in complex graph domains. These observations are discussed in detail in the following, providing a comprehensive understanding of the behavior of HyperUFG and related models across different types of hypergraphs.

Observation 3: Hypergraph neural networks perform effectively in handling tasks on homophilic hypergraph data. The performance of all comparative models on homophilic hypergraphs is markedly superior to that of the graph-agnostic model MLP, as evidenced in Table 3. This superiority can be largely attributed to the benefits of hypergraph filtering, which effectively aggregates information from homophilic relationships, thereby enhancing the model’s ability to capture relevant structural features. This finding is consistent with empirical observations reported in related works on HHL, where similar performance gains have been noted in homophilic settings.

Observation 4: Most existing hypergraph neural networks face challenges when dealing with heterophilic hypergraph data. Table 4 reveals that the majority of HNNs underperform relative to MLP, particularly on the Twitchgamers dataset, where MLP outperforms all other baselines. This underperformance is likely due to the detrimental effect of heterophilic relationships on the efficacy of hypergraph filtering, which may result in the incorporation of information that is misaligned with the inherent features of the data, leading to less effective inductive representations. This new finding has not received sufficient attention in related HHL studies, as these works predominantly restrict their analysis to homophilic settings, overlooking the challenges posed by heterophilic hypergraph data.

Observation 5: HyperUFG shows relatively strong performance across both homophilic and heterophilic hypergraph data. As shown in Tables 3 and 4, HyperUFG consistently outperforms other baseline models on both homophilic and heterophilic hypergraphs. Notably, HyperUFG shows a significant performance advantage on heterophilic hypergraphs, indicating its ability to effectively manage the complexities associated with diverse hypergraph structures.

5.3 Comparison of Homophily/Heterophily Metrics

In this study, we introduce two sets of synthetic datasets, each constructed using different homophily measurement approaches, namely, hyperedge homophily and node homophily, to evaluate the performance of MLP, various HNNs, and HyperUFG. The statistics of these synthetic datasets, categorized by hyperedge homophily and node homophily, are summarized in Table 5. The comparative performance results are illustrated in Figure 3.

As shown, HyperUFG generally outperforms other baseline models across a range of homophily settings. A notable observation is that most baseline models perform worse than MLP when homophily levels are low (i.e., when $\mathcal{H}_{edge}/\mathcal{H}_{node} < 0.5$). However, methods such as ED-HNN and SheafHyperGNN, which incorporate strategies to ad-

Setting in HSBM: $n = 10000, r = 3, z = 2$							
Hyperedge hom. ratio, \mathcal{H}_{edge}	0.3	0.4	0.5	0.6	0.7	0.8	0.9
p	0.01	0.1	0.3	0.5	0.6	0.5	0.9
q	0.99	0.6	0.8	0.6	0.5	0.2	0.2
Node hom. ratio, \mathcal{H}_{node}	0.3	0.4	0.5	0.6	0.7	0.8	0.9
p	0	0.1	0.3	0.5	0.6	0.5	0.9
q	1	0.6	0.8	0.6	0.5	0.2	0.1

Table 5: Statistics of synthetic hypergraphs.

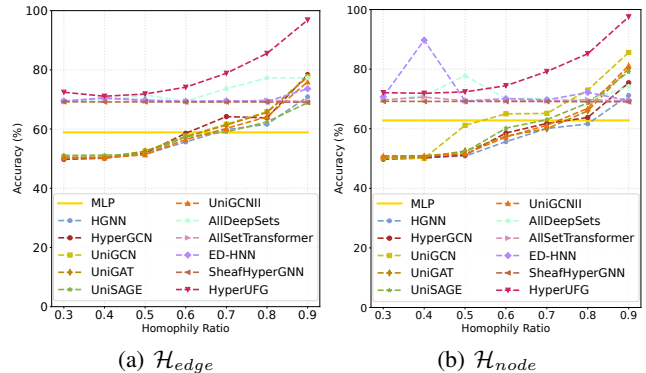


Figure 3: Baseline performance comparison under different homophily ratios.

dress heterophily, are able to manage the challenges posed by low homophily, though they experience some performance degradation as homophily increases, indicating inherent limitations. Overall, HyperUFG demonstrates good performance across various homophily conditions, highlighting its adaptability and effectiveness in diverse settings.

6 Conclusion

In this paper, we have explored the underrepresented area of heterophilic hypergraph learning (HHL) within the broader context of hypergraph neural networks (HNNs). While HNNs have shown significant promise in tasks involving high-order correlations, our study highlights the critical gaps that exist when these models are applied to heterophilic hypergraphs. To address these gaps, we first introduce new metrics to quantify homophily/heterophily ratio for hypergraphs. Then, a diverse set of benchmark datasets across various real-world scenarios, which serve as essential tools for evaluating existing HNNs and advancing research in HHL, are developed. Additionally, we propose HyperUFG, a novel framelet-based HNN that leverages both low-pass and high-pass filters, offering a new baseline for future studies on HHL. Our extensive experiments on synthetic and benchmark datasets reveal that many current HNNs struggle with heterophilic hypergraphs, often underperforming compared to simpler models like MLP. However and surprisingly, HyperUFG demonstrates strong and competitive performance across various settings, frequently outperforming existing models, especially in heterophilic scenarios. Overall, our work sets the stage for further exploration and development in this emerging field, and we expect our findings to inspire and guide future research efforts towards advancing HHL.

Acknowledgements

This work was supported in part by the “Pioneer” and “Leading Goose” R&D Program of Zhejiang (No. 2024C03262), the National Natural Science Foundation of China (No. U21A20473, No. 62172370) and the Jinhua Science and Technology Plan (No. 2023-3-003a). L. Bai was supported by the National Natural Science Foundation of China (No. T2122020). X. Zhuang was supported in part by the Research Grants Council of the Hong Kong Special Administrative Region, China, under Project CityU 11309122, CityU 11302023, and CityU 11301224.

References

- Antelmi, A.; Cordasco, G.; Polato, M.; Scarano, V.; Spagnuolo, C.; and Yang, D. 2023. A survey on hypergraph representation learning. *ACM Computing Surveys*, 56(1): 1–38.
- Bo, D.; Wang, X.; Shi, C.; and Shen, H. 2021. Beyond low-frequency information in graph convolutional networks. In *AAAI*, 3950–3957.
- Chen, M.; Wei, Z.; Huang, Z.; Ding, B.; and Li, Y. 2020. Simple and deep graph convolutional networks. In *ICML*, 1725–1735.
- Chien, E.; Pan, C.; Peng, J.; and Milenkovic, O. 2022. You are AllSet: A multiset function framework for hypergraph neural networks. In *ICLR*.
- Cole, S.; and Zhu, Y. 2020. Exact recovery in the hypergraph stochastic block model: A spectral algorithm. *Linear Algebra and its Applications*, 593: 45–73.
- Dong, B. 2017. Sparse representation on graphs by tight wavelet frames and applications. *Applied and Computational Harmonic Analysis*, 42(3): 452–479.
- Duta, I.; Cassarà, G.; Silvestri, F.; and Liò, P. 2023. Sheaf hypergraph networks. In *NeurIPS*, 12087–12099.
- Feng, Y.; You, H.; Zhang, Z.; Ji, R.; and Gao, Y. 2019. Hypergraph neural networks. In *AAAI*, 3558–3565.
- Gong, C.; Cheng, Y.; Yu, J.; Xu, C.; Shan, C.; Luo, S.; and Li, X. 2024. A Survey on Learning from Graphs with Heterophily: Recent Advances and Future Directions. *arXiv preprint arXiv:2401.09769*.
- Huang, J.; and Yang, J. 2021. UniGNN: A unified framework for graph and hypergraph neural networks. In *IJCAI*, 2563–2569.
- Juluru, K.; Shih, H.-H.; Keshava Murthy, K. N.; and El-najjar, P. 2021. Bag-of-words technique in natural language processing: A primer for radiologists. *RadioGraphics*, 41(5): 1420–1426.
- Jure, L. 2014. SNAP Datasets: Stanford large network dataset collection. [Online] Available: <http://snap.stanford.edu/data>.
- Kim, S.; Lee, S. Y.; Gao, Y.; Antelmi, A.; Polato, M.; and Shin, K. 2024. A survey on hypergraph neural networks: An in-depth and step-by-step guide. 6534–6544.
- Li, B.; Pan, E.; and Kang, Z. 2024. PC-Conv: Unifying homophily and heterophily with two-fold filtering. In *AAAI*, 13437–13445.
- Li, J.; Zheng, R.; Feng, H.; Li, M.; and Zhuang, X. 2024a. Permutation equivariant graph framelets for heterophilous graph learning. *IEEE Transactions on Neural Networks and Learning Systems*, 35(9): 11634–11648.
- Li, M.; Li, Z.; Huang, C.; Jiang, Y.; and Wu, X. 2024b. EduGraph: Learning path-based hypergraph neural networks for MOOC course recommendation. *IEEE Transactions on Big Data*, 10(6): 706–719.
- Luan, S.; Hua, C.; Lu, Q.; Ma, L.; Wu, L.; Wang, X.; Xu, M.; Chang, X.-W.; Precup, D.; Ying, R.; et al. 2024. The heterophilic graph learning handbook: Benchmarks, models, theoretical analysis, applications and challenges. *arXiv preprint arXiv:2407.09618*.
- Pei, H.; Wei, B.; Chang, K. C. C.; Lei, Y.; and Yang, B. 2020. Geom-GCN: Geometric graph convolutional networks. In *ICLR*.
- Prokopchik, K.; Benson, A. R.; and Tudisco, F. 2022. Non-linear feature diffusion on hypergraphs. In *ICML*, 17945–17958.
- Tang, B.; Liu, Z.; Jiang, K.; Chen, S.; and Dong, X. 2024. Simplifying hypergraph neural networks. *arXiv preprint arXiv: 2402.05569*.
- Telyatnikov, L.; Bucarelli, M. S.; Bernardez, G.; Zaghen, O.; Scardapane, S.; and Lio, P. 2023. Hypergraph neural networks through the lens of message passing: a common perspective to homophily and architecture design. *arXiv preprint arXiv:2310.07684*.
- Veldt, N.; Benson, A. R.; and Kleinberg, J. 2023. Combinatorial characterizations and impossibilities for higher-order homophily. *Science Advances*, 9(1): eabq3200.
- Wang, P.; Yang, S.; Liu, Y.; Wang, Z.; and Li, P. 2023. Equivariant hypergraph diffusion neural operators. In *ICLR*.
- Wu, L.; Lin, H.; Hu, B.; Tan, C.; Gao, Z.; Liu, Z.; and Li, S. Z. 2024. Beyond homophily and homogeneity assumption: Relation-based frequency adaptive graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 35(6): 8497–8509.
- Yadati, N.; Nimishakavi, M.; Yadav, P.; Nitin, V.; Louis, A.; and Talukdar, P. 2019. HyperGCN: A new method for training graph convolutional networks on hypergraphs. In *NeurIPS*, 1511–1522.
- Zheng, X.; Wang, Y.; Liu, Y.; Li, M.; Zhang, M.; Jin, D.; Yu, P. S.; and Pan, S. 2022. Graph neural networks for graphs with heterophily: A survey. *arXiv preprint arXiv:2202.07082*.
- Zheng, X.; Zhou, B.; Gao, J.; Wang, Y. G.; Lio, P.; Li, M.; and Montufar, G. 2021. How framelets enhance graph neural networks. In *ICML*, 12761–12771.
- Zhu, J.; Yan, Y.; Heimann, M.; Zhao, L.; Akoglu, L.; and Koutra, D. 2023. Heterophily and graph neural networks: Past, present and future. *IEEE Data Engineering Bulletin*, 47(2): 10–32.
- Zhu, J.; Yan, Y.; Zhao, L.; Heimann, M.; Akoglu, L.; and Koutra, D. 2020. Beyond homophily in graph neural networks: Current limitations and effective designs. In *NeurIPS*, 7793–7804.