

GCNH: A Simple Method For Representation Learning On Heterophilous Graphs

Original

GCNH: A Simple Method For Representation Learning On Heterophilous Graphs / Cavallo, Andrea; Grohnfeldt, Claas; Russo, Michele; Lovisotto, Giulio; Vassio, Luca. - (2023), pp. 1-8. (2023 International Joint Conference on Neural Networks, IJCNN 2023 Gold Coast (AUS) 18-23 June 2023) [10.1109/ijcnn54540.2023.10191196].

Availability:

This version is available at: 11583/3010553 since: 2026-05-05T10:03:32Z

Publisher:

Institute of Electrical and Electronics Engineers

Published

DOI:10.1109/ijcnn54540.2023.10191196

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

GCNH: A Simple Method For Representation Learning On Heterophilous Graphs

Andrea Cavallo^{*†}, Claas Grohnfeldt[†], Michele Russo[†], Giulio Lovisotto[†] and Luca Vassio^{*}
Politecnico di Torino^{*}, Huawei Munich Research Center[†]
^{*}{name}. {surname}@polito.it, [†]{name}. {surname}@huawei.com

Abstract—Graph Neural Networks (GNNs) are well-suited for learning on homophilous graphs, i.e., graphs in which edges tend to connect nodes of the same type. Yet, achievement of consistent GNN performance on heterophilous graphs remains an open research problem. Recent works have proposed extensions to standard GNN architectures to improve performance on heterophilous graphs, trading off model simplicity for prediction accuracy. However, these models fail to capture basic graph properties, such as neighborhood label distribution, which are fundamental for learning.

In this work, we propose GCN for Heterophily (GCNH), a simple yet effective GNN architecture applicable to both heterophilous and homophilous scenarios. GCNH learns and combines *separate* representations for a node and its neighbors, using one *learned importance coefficient* per layer to balance the contributions of center nodes and neighborhoods. We conduct extensive experiments on eight real-world graphs and a set of synthetic graphs with varying degrees of heterophily to demonstrate how the design choices for GCNH lead to a sizable improvement over a vanilla GCN. Moreover, GCNH outperforms state-of-the-art models of much higher complexity on four out of eight benchmarks, while producing comparable results on the remaining datasets. Finally, we discuss and analyze the lower complexity of GCNH, which results in fewer trainable parameters and faster training times than other methods, and show how GCNH mitigates the oversmoothing problem.

Index Terms—Graph Neural Networks, heterophily, disassortativity, graph representation learning

I. INTRODUCTION

GNNs are core components of current state-of-the-art methods for learning and prediction on graph-structured data across domains and applications. Their capability to encode semantic and contextual information into node embeddings is known to be particularly effective on homophilous graphs, i.e., graphs in which nodes of the same type tend to be connected [1], [2]. On the other hand, in the case of heterophilous graphs, in which neighboring nodes tend to be dissimilar in type, the achievement of competitive, or at least consistent, GNN prediction accuracy remains an open research goal [3], [4].

Among other explanations for the inconsistent GNN performance in the heterophilous case, [5] suggests that the message aggregation strategies of standard GNNs could lead to weakly representative node embeddings, due to a disproportionately high contribution of dissimilar neighbors. That is why some works attempt to improve the generalization capability of GNNs to heterophilous networks by either modifying the graph structure or tailoring aggregation strategies and network architecture for such scenario [3], [4], [6]–[18]. Nevertheless,

[19] shows that a vanilla Graph Convolutional Network (GCN) can still *outperform* more complex (heterophilous) models on some heterophilous graphs. [19] and [20] suggest that this contradiction originates from the fact that the edge homophily ratio should not be considered a representative indicator of GCN performance, and instead recommend taking into account neighborhood structure and distribution of node labels.

Based on these observations, we propose a simple yet effective GNN architecture, namely GCN for Heterophily (GCNH), to improve node representation capabilities for applications on heterophilous graphs. Differently from other GNNs, GCNH learns two different functions that encode a node and its neighbors *separately*. In addition, we allow the GCNH layer to flexibly assign different relevance to the information present in the neighborhood versus the information present in the center node. We discuss and show how this design mitigates noisy neighborhood representations from negatively influencing the learned embeddings while allowing informative neighbors to strongly influence the final node embedding. These extensions make GCNH more adaptive to heterophily than a vanilla GCN, while also improving prediction accuracy over more complex models designed for heterophily on common benchmarks.

We evaluate GCNH on the task of node classification using eight common real-world datasets and one set of synthetic graphs with varying degrees of heterophily. We show that GCNH is able to learn meaningful representations of nodes independently of the homophily level of the graph, achieving new state-of-the-art performance on heterophilous graph datasets while producing results comparable to the state-of-the-art on homophilous benchmarks.

Our main contributions are summarized below.

- We present GCNH, a simple yet effective GNN architecture that improves graph representation learning capabilities on heterophilous graphs while preserving the advantages of GCN.
- We present extensive experiments on real and synthetic datasets with varying degrees of heterophily for the node classification task. GCNH improves over the state-of-the-art on four (out of eight) real-world benchmarks while performing comparably to the state-of-the-art on all other datasets, including homophilous graphs.
- We showcase the lower complexity of GCNH compared to other state-of-the-art models, both in terms of the training time and the number of trainable parameters, and how it mitigates the oversmoothing problem.

II. RELATED WORK

The analysis and improvement of GNN performance on heterophilous graphs have received increasing attention in recent years [3], [4], [6]–[18]. GNNs generate node embeddings in two steps: message transformation, where node messages, i.e., features or embeddings from the previous layer, are transformed, and message aggregation, where the final node embeddings are generated by aggregating the messages of neighbors in the graph. One of the earliest and today most common GNN models is the GCN [1], which defines message transformation as a learnable linear layer and message aggregation as the average of the messages from the neighbors and from the center node.

Several works have proposed new GNN architectures or graph transformations to make the message-passing framework suitable for applications on non-homophilous graphs. We compare a number of these approaches to our method in Section V. Borrowing the categorization introduced in [5], these approaches can be subdivided as follows.

A. Non-Local Neighbor Extension

These methods selectively extend the receptive field of GNNs to include potentially important nodes located outside of the local neighborhood. These approaches are based on the assumption that neighboring nodes may be dissimilar and that information about the target node, presumably carried by nodes with the same label, is available in nodes belonging to higher-order neighborhoods. In particular, methods such as [3], [4], [7], [15] create latent representations for nodes using multiple neighborhoods at different hop distances and combine them into one overall embedding. Other approaches, such as [6], [8], [16], [18], do not consider entire neighborhoods but single nodes as potential neighbors, regardless of their location in the graph, and aggregate messages from the nodes that are estimated to be more relevant.

B. GNN Architecture Adaptation

Methods in this category modify the GNN architecture with the goal of improving representation learning capabilities on heterophilous graphs. A common approach is to estimate the relevance of a node to a given target node with respect to the prediction task at hand, and to assign individual weights to messages from neighbors based on their relevance. This approach is used in [9], [11]–[13], [17] among the others. Another approach is to learn separate embeddings for the neighborhood and the target node and merge them at a later stage. This mitigates the phenomenon of noisy embeddings, where information from potentially highly dissimilar neighbors is mixed with the features of the target node. Methods that follow this approach include [3], [18]. In addition, [3], [10], [14] adopt the strategy of treating the information captured by different GNN layers separately, instead of first aggregating the information from all layers and then using the final node embeddings for prediction.

III. PRELIMINARIES

A. Notation And Problem Statement

Let $G = (V, E)$ be an unweighted and undirected graph, where V is the set of nodes and E is the set of edges. The connectivity information in the graph is represented by the adjacency matrix $A \in \mathbb{R}^{n \times n}$, where $n = |V|$ is the number of nodes and matrix elements A_{uv} are equal to 1 if nodes u and v are adjacent and 0 otherwise. Each node is associated with a feature vector x_u of size f , and the complete set of features in the graph is denoted by $X \in \mathbb{R}^{n \times f}$. The *neighbors* of a node u are the set of nodes adjacent to it, denoted by \mathcal{N}_u . Note that \mathcal{N}_u does not include node u .

The task addressed in this work is *supervised node classification*. In this scenario, each node is associated with a label $y_u \in C$ representing the class the node belongs to, where C is the set of labels. The task corresponds to learning a mapping $\mathcal{F} : V \rightarrow C$ that uses the information of the graph G , the features X and the labels to map nodes into their ground-truth class. As is standard practice, we add a linear layer that maps the final node representations of the last network layer $H^L \in \mathbb{R}^{n \times e^L}$ (with L total number of layers and e^L size of the final node embeddings) to class probabilities with the addition of a softmax. For a node u :

$$\tilde{y}_u = \text{softmax}(h_u^L W_{cl}), \quad (1)$$

where $W_{cl} \in \mathbb{R}^{e^L \times |C|}$ is a learnable matrix. We omit the bias term from the equation for simplicity. Note that $\tilde{y}_u \in \mathbb{R}^{|C|}$ is the probability distribution over classes for node u . During inference, networks assign to the node the class with maximum probability:

$$\hat{y}_u = \text{argmax}(\tilde{y}_u). \quad (2)$$

To train the model, we minimize the negative log-likelihood loss on the training data $\mathcal{L}(y_u, \hat{y}_u)$, where $\hat{y}_u = \mathcal{F}(G, X, u)$ is the predicted label for node u . The loss function $\mathcal{L}(\cdot, \cdot)$ describes the distance between the true label y_u and the predicted label. We focus on the transductive case: we work on single-graph and separate nodes into training, validation and test set.

B. Homophily And Heterophily

Graph homophily is a social science-inspired property of graphs [21], defined as the extent to which similar nodes in a graph are connected. Although other definitions of node similarity exist [17], [19], [20], this work focuses on label similarity, i.e., we define nodes as similar if they share the same label, and we measure graph homophily using the standard *edge homophily ratio*, denoted by h :

$$h = \frac{|\{(u, v) : (u, v) \in E \wedge y_u = y_v\}|}{|E|}, \quad (3)$$

which quantifies the fraction of edges in a graph that connect nodes with the same label. Graphs with low values of h are called *heterophilous* or *disassortative*. As discussed in Section I, GNNs perform inconsistently on this category of graphs – we analyze this point in depth in Section V.

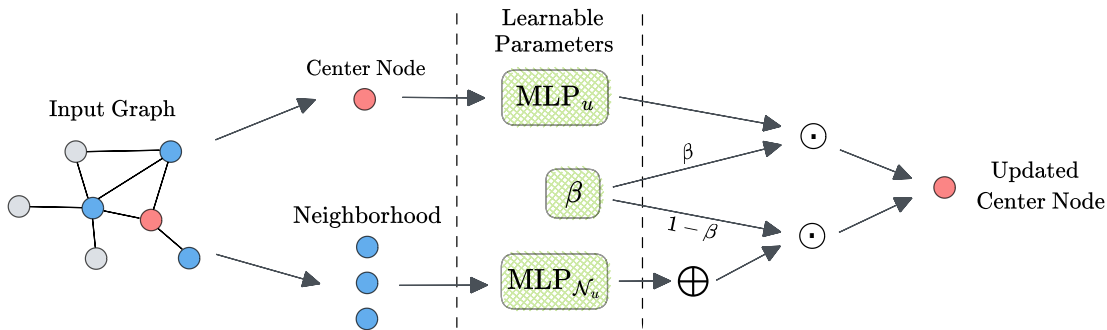


Fig. 1: Architecture of the GCNH layer. To create an updated representation for the center node (red), the GCNH layer explicitly separates the node from its 1-hop neighborhood (blue). Two different MLPs encode center node and neighborhood separately (MLP_u and $MLP_{\mathcal{N}_u}$, respectively). The encoded neighbors are then aggregated with a permutation invariant function \oplus . Finally, the aggregated neighborhood and the encoded center node are combined together with a *learnable* weighting factor β which regulates the contributions of the center node versus its neighborhood to produce the final output embedding for the node.

IV. METHOD

In this section, we describe the architecture of the model proposed in this paper, namely GCNH. Figure 1 illustrates the structure of the GCNH layer.

A. GCNH Layer Formulation

A GCNH network is composed of one or more, L , GCNH layers. The ℓ^{th} layer receives in input the adjacency matrix and the node representations computed at the previous layer $H^{\ell-1}$ and produces updated node representations H^ℓ . We set $H^0 = X$.

Within a layer, node representations are transformed through two separate 1-layer MLPs, MLP_u and $MLP_{\mathcal{N}_u}$, resulting in latent representations z_u and $z_{\mathcal{N}_u}$ of the target node and its neighborhood, respectively. For a node u , at layer ℓ , we formally describe this step as follows. First, an intermediate representation is computed for the node u :

$$z_u^\ell = MLP_u(h_u^{\ell-1}) = \sigma(h_u^{\ell-1}W_1). \quad (4)$$

Secondly, all the representations of u 's neighbors ($v \in \mathcal{N}_u$) are updated based on their current representations h_v and aggregated together to obtain a neighborhood representation:

$$z_{\mathcal{N}_u}^\ell = \bigoplus_{v \in \mathcal{N}_u} (MLP_{\mathcal{N}_u}(h_v^{\ell-1})) = \bigoplus_{v \in \mathcal{N}_u} (\sigma(h_v^{\ell-1}W_2)). \quad (5)$$

In Equations 4 and 5, $W_1, W_2 \in \mathbb{R}^{e^{\ell-1} \times e^\ell}$ are learnable matrices, $\sigma(\cdot)$ is a generic activation function, e^ℓ is the size of the embeddings created at the ℓ -th layer and \bigoplus is a permutation invariant aggregation function over the nodes $v \in \mathcal{N}_u$. We omit the bias terms of the MLPs in the equations for clarity. The final output embedding for node u is obtained as a linear combination of z_u^ℓ and $z_{\mathcal{N}_u}^\ell$, parametrized by a learnable scalar value β :

$$h_u^\ell = (1 - \beta)z_{\mathcal{N}_u}^\ell + \beta z_u^\ell \quad (6)$$

where β is normalized between 0 and 1 with a sigmoid(\cdot) function. Note that the MLPs, \bigoplus and β are different for each layer ℓ . We omit superscripts in the equations for clarity.

B. GCNH Design Choices

GCNH introduces two main design choices which differentiate it from a standard GCN: (i) the separate encoding of the target node and its neighbors (Equations 4, 5) and (ii) the explicit parameterization of the contributions of neighborhood and center node with β (Equation 6).

The positive impact of separately processing the node and its neighborhood has been previously outlined in [3], [18]. Intuitively, in homophilous settings, where neighbors are similar, aggregation from the neighborhood brings useful information. On the other hand, in heterophilous graphs, dissimilar neighbors might bring detrimental information that a GNN cannot easily ignore. Compared to [3], GCNH takes the separation principle further, while minimizing the amount of complexity that this choice adds to the network architecture. Specifically, three choices distinguish GCNH from the model proposed in [3]: (i) we only use 1-hop neighborhoods, (ii) we learn separate MLPs for center node and neighborhood and (iii) we combine separate embeddings using a learned linear combination instead of concatenation. Note that (i) leads to faster computation while retaining most of the information useful for node representation, as shown in Section V, (ii) improves flexibility and (iii) leads to models with fewer parameters. (i) and (iii) lead to lower time complexity compared to [3] (see Section IV-C for a time complexity analysis).

Combined with the embedding separation, the explicit modeling of the importance of the neighborhood informativeness with the coefficient β in GCNH allows to adaptively determine the impact of the neighborhood on the final node embeddings based on how informative the neighbors are. Note that informative neighbors might also exist in heterophilous graphs [19], [20]; therefore, the contribution of neighbors is not necessarily related to the homophily level of the network. Intuitively, modeling and learning β explicitly provides a helpful inductive bias that allows the network to directly prefer the center node versus neighborhood information. Section V shows how models that are equipped with similar, or better, flexibility (such as GAT [2]) perform poorly in comparison [22].

Benchmark h	Cornell 0.30	Texas 0.11	Wisconsin 0.21	Film 0.22	Chameleon 0.23	Squirrel 0.22	Cora 0.81	Citeseer 0.74
MLP	81.89±6.40	80.81±4.75	85.29±3.31	36.53±0.70	46.21±2.99	28.77±1.56	75.69±2.00	74.02±1.90
GCN	60.54±5.30	55.14±5.16	51.76±3.06	27.32±1.10	64.82±2.24	53.43±2.01	86.98±1.27	76.50±1.36
GAT	61.89±5.05	52.16±6.63	49.41±4.09	27.44±0.89	60.26±2.50	40.72±1.55	87.30±1.10	76.55±1.23
Geom-GCN	60.54±3.67	66.76±2.72	64.51±3.66	31.59±1.15	60.00±2.81	43.80±1.48	85.35±1.57	78.02±1.15
H ₂ GCN	82.70±5.28	84.86±7.23	87.65±4.98	35.70±1.00	60.11±2.15	36.48±1.86	87.87±1.20	77.11±1.57
GPRGNN	80.27±8.11	78.38±4.36	82.94±4.21	34.63±1.22	46.58±1.71	31.61±1.24	87.95±1.18	77.13±1.67
GGCN	85.68±6.63	84.86±4.55	86.86±3.29	37.54±1.56	71.14±1.84	55.17±1.58	87.95±1.05	77.14±1.45
O(d)-SD	84.86±4.71	85.95±5.51	89.41±4.74	37.81±1.15	68.04±1.58	56.34±1.32	86.90±1.13	76.70±1.57
GCNH	86.49±6.98	87.84±3.87	87.65±3.59	36.89±1.50	71.56±1.86	61.85±1.54	86.88±1.04	75.81±1.14

TABLE I: Mean classification accuracy and standard deviation for GCNH on real-world datasets, on the 10 splits taken from [6]. Best results are in **red**, second best results in **blue** and third best in **violet**. The results for the baselines are taken from [12] and [13]. We use sum as the aggregation function in GCNH; other parameters of GCNH are selected from the best-performing configuration (see Table VI in the Appendix for details on the hyperparameters).

C. Time Complexity

We compute the time complexity of a generic GCNH layer ℓ by analyzing the individual processing steps of the model.

First, the node representations $H^{\ell-1}$ are transformed, separately for the center node and the neighbors, as described in Equations (4) and (5). This step has a time complexity of $\mathcal{O}(ne^{\ell-1}e^\ell)$. Subsequently, the neighbors’ representations are aggregated and merged with the center-node embedding according to Equation (6). All the neighbor aggregation functions tested (see Section V-F) have a time complexity of $\mathcal{O}(|E|e^\ell)$, whereas the weighted sum with the self-node representation has a time complexity of $\mathcal{O}(ne^\ell)$. This last term is dominated by the complexity of the transformation. In total, the overall time complexity of the GCNH layer is $\mathcal{O}(ne^{\ell-1}e^\ell + |E|e^\ell)$. The linear dependency of the complexity of the GCNH layer on the number of nodes n makes it easier to scale on large graphs compared to attention-based models (e.g., GGCN [12]), whose complexity depends quadratically on n .

V. EXPERIMENTS

In this section, we analyze the learning capability of GCNH on the task of supervised node classification, commonly used for models dealing with heterophilous graphs. We perform the evaluation both on synthetic datasets with different levels of homophily ratio and on the real-world graphs widely used in related works. Further information about the datasets is reported in Appendix A.

A. Baselines

The baselines used for comparison belong to two main categories. The first category includes well-known methods:

- **MLP**: Multi-Layer Perceptron.
- **GCN** [1]: Graph Convolutional Network.
- **GAT** [2]: Graph Attention Network.

Methods in the second category, instead, are specifically designed for heterophilous graphs. In particular:

- **Geom-GCN** [6] maps nodes to a latent space and defines a new graph based on embedding similarity.

- **H₂GCN** [3] introduces three specific designs to boost performance on heterophilous graphs: ego and neighbor-embedding separation, higher-order neighborhoods and combination of intermediate representations.
- **GPRGNN** [14] uses PageRank to determine relations between nodes.
- **GGCN** [12] applies two designs to extend GNNs: degree correction and signed messages.
- **O(d)-SD** [13] is based on sheaf diffusion.

B. Experimental Setting

We evaluate model performance in terms of classification accuracy, i.e., the percentage of nodes in the test set that are assigned to the correct class. We test different values for hyperparameters and we select the best ones. We report further information about the grids for the hyperparameters in Appendix B. The aggregation function \oplus in Equation (5) is an element-wise sum unless differently specified (see Section V-F for a comparison of aggregation functions). We compute classification accuracies on 10 different train/validation/test splits provided by [6] for each dataset, and we report mean and standard deviation. The sizes of the splits are 48%/32%/20%. For each dataset, we train the models on the training sets and the model performing best on average across the validation sets is used on the test sets, on which we compute the accuracy values.

We run experiments on an NVIDIA Tesla V100 PCIE with 16 GB, except for the experiments in Table IV that we run on an NVIDIA Tesla T4. The code is implemented in Python and PyTorch is used for deep learning models. Our code is available at <https://github.com/SmartData-Polito/GCNH>.

C. Results On Real-World Datasets

Table I reports classification accuracies for GCNH and baselines on real-world datasets. GCNH achieves state-of-the-art performance on four out of the eight datasets used; it ranks second and third on Wisconsin and Film and performs slightly worse on the two homophilous graphs Cora and Citeseer. The results on all the heterophilous graphs prove the effectiveness of the design choices of GCNH. On Cornell, Texas, Wisconsin




Fig. 2: Comparison of the classification accuracy achieved by a set of models on the syn-cora synthetic datasets. Each point shows the average accuracy on three datasets generated with a specific homophily ratio.

and Film, GCNH avoids the detrimental influence of neighbors on the final embeddings observed, for example, in GCN and GAT, with respect to which improvements are large (up to 35%). On Chameleon and Squirrel, GCNH manages to encode the useful neighborhood information present in the graph thanks to its separate processing of node and 1-hop neighbors. Indeed, on these two graphs, GCN performs quite well (it ranks fourth on both), meaning that information contained in the 1-hop neighborhood is useful for node classification, whereas more complex models (e.g. Geom-GCN, H₂GCN and GPRGNN) fail to capture this property and perform worse (from 5 to 20% accuracy drop with respect to GCN). On the homophilous graphs Cora and Citeseer, GCNH performs comparably to the other models, although slightly worse. These results show that GCNH flexibly adapts to various homophily settings, achieving consistent performance.

D. Results On Synthetic Datasets

To better understand how GCNH deals with different levels of homophily, we evaluate it on synthetic graphs covering a wide range of values of h , while keeping node features and other graph properties unchanged. We provide additional details about the dataset used in Appendix A-B.

Figure 2 shows the accuracy of GCNH and three other simple baselines: MLP, whose performance is not affected by the value of h as it is a graph-unaware method, and GCN and GAT, which achieve significantly worse results on heterophilous settings. The hyperparameter configurations for these baselines are reported in Appendix B-A. GCNH’s performance is significantly less dependent on the homophily level of the graph than the performances of GAT and GCN. On heterophilous graphs, GCNH improves by almost 50% over GCN and GAT and by $\sim 2/8\%$ over MLP, meaning that the separate encoding of the neighbors has a positive impact. On homophilous graphs, GCNH performs comparably to GCN and GAT, achieving perfect accuracy on perfectly homophilous graphs ($h = 1$).




Fig. 3: Values of learned parameter β for GCNH on real and synthetic datasets. Results are obtained using the best-performing hyperparameter configuration for GCNH models with one layer.

E. Analysis of the GCNH Design

To measure the impact of the design choices we made for GCNH, we perform an ablation study by removing the two main components of GCNH, namely, the separate MLPs and the learned importance coefficient β , to see how they affect the performance. We report the results in Table II, where we use a standard GCN as a baseline. The table shows how both components bring an improvement in classification accuracy on heterophilous graphs, with a small tradeoff on the performance on homophilous graphs. The MLPs separation leads to the largest gain, and its combination with the β coefficient consistently brings further improvements.

We analyze further the behavior of β . Since β balances the contribution of messages of self-node and its neighbors, we expect heterophilous graphs to lead to larger β s given that neighborhoods are generally less informative in those cases. Figure 3 shows the values of the parameter β on different graphs. On synthetic graphs, β is evidently correlated with the edge homophily ratio measured by h . On most real graphs, the values follow a trend similar to that observed in synthetic graphs. Chameleon and Squirrel are exceptions: we find that, while these datasets are heterophilous according to the h metric, GCNH tends to learn low values of β , corresponding to highly informative neighborhoods. This corroborates the findings outlined in [19], which pointed out a similar contradiction: how a standard GCN performs unexpectedly well on the heterophilous Chameleon and Squirrel, even outperforming heterophily-specific methods. In fact, this result further suggests that edge homophily ratio is not suited to describe neighborhood informativeness in general.

F. Aggregation Functions

We test three different aggregation functions for $\bigoplus_{v \in \mathcal{N}_u}$ in Equation (5): element-wise sum, element-wise mean across the neighbors and element-wise max. Table III reports a comparison of the node classification results for these aggregation functions.

	Separate MLPs	Learned β	Cornell	Texas	Wisconsin	Film	Chameleon	Squirrel	Cora	Citeseer
GCN	\times	\times	60.54	55.14	51.76	27.32	64.82	53.43	86.98	76.50
GCNH	\checkmark	\times	83.78(+23.2)	86.49(+31.3)	85.49(+33.7)	36.01(+8.7)	70.22(+5.4)	59.74(+6.3)	86.90(-0.1)	75.65(-0.8)
GCNH	\checkmark	\checkmark	86.49(+25.9)	87.84(+32.7)	87.65(+35.9)	36.89(+9.6)	71.56(+6.7)	61.85(+8.4)	86.88(-0.1)	75.81(-0.7)

TABLE II: Mean classification accuracy of GCN and ablated GCNH. In parentheses, we report the performance **improvement** or **degradation** from the GCN baseline. ‘‘Separate MLPs’’ refers to whether we learn separate linear layers for node and neighborhoods (W_1, W_2 in Equation 4 and 5, respectively). ‘‘Learned β ’’ refers to whether we learn β as a parameter or not; if not, it is fixed at $\beta=0.5$. Best results are in **bold**.

Benchmark	Cornell	Texas	Wisconsin	Film	Chameleon	Squirrel	Cora	Citeseer
GCNH (sum)	86.49 \pm 6.98	87.84 \pm 3.87	87.65 \pm 3.59	36.89\pm1.50	71.56 \pm 1.86	61.85 \pm 1.54	86.88\pm1.04	75.81 \pm 1.14
GCNH (mean)	86.49\pm4.41	88.92\pm3.24	89.80\pm5.29	36.83 \pm 1.44	55.64 \pm 1.95	39.24 \pm 1.02	85.75 \pm 1.57	76.00\pm0.99
GCNH (max)	85.41 \pm 8.18	88.65 \pm 4.56	88.04 \pm 3.26	36.07 \pm 1.26	71.64\pm1.85	63.35\pm1.99	86.68 \pm 1.10	75.96 \pm 1.33

TABLE III: Mean classification accuracy and standard deviation for GCNH with different aggregation functions in Equation 5. Note that results for ‘‘max’’ are obtained full-batch training only (see Table VI and discussion in Section V-H), as it allows for more efficient implementation. Best results are in **bold**.

	N. Params	Cora	Citeseer	Chameleon	Squirrel	Film
GGCN	118 k	96.05	99.94	74.83	331.64	628.36
O(d)-SD	46 k	19.64	20.15	48.87	275.90	44.22
GCNH (sum)	30 k	8.79	10.28	8.56	12.61	15.59
GCNH (max)	30 k	11.26	13.19	10.71	17.20	12.82

TABLE IV: Number of trainable parameters and training times (sec) for two state-of-the-art methods and GCNH on several graphs. Training is performed for 200 epochs on 10 splits for each dataset. Models have one layer and hidden size 16. Shortest times for each dataset are in **bold**.

G. GCNH And Oversmoothing

The performance of GNNs is known to gradually decrease when increasing the number of layers. This decay is partly attributed to oversmoothing, i.e., repeated graph convolutions eventually making node embeddings indistinguishable from each other [23], [24]. We show experimentally that the design choices of GCNH alleviate the oversmoothing problem. As shown in Figure 4, GCNH’s accuracy decreases just slightly when increasing the number of layers, whereas increasing the layers of GCN leads to a larger drop in performance; see Section B-C in Appendix for the experimental details.

H. Training Times And Trainable Parameters

We report in Table IV the training times required by GCNH and two of its main competitors GGCN [12] and O(d)-SD [13], as well as their number of trainable parameters. We take the implementations of these methods from the repositories of the authors [25], [26]. We report the result for GCNH using sum and max aggregation. We omit results for mean since sum and mean require the same amount of computation as they can be implemented efficiently with a matrix multiplication, i.e., their training times are the same. Max, instead, requires a scattered max pooling operation for which we use the implementation in [27]. GCNH is noticeably faster and has fewer trainable parameters than both GGCN and O(d)-SD.




Fig. 4: Accuracy for GCNH and GCN with different numbers of layers. With a large number of layers, GCNH prevents oversmoothing.

VI. CONCLUSIONS

We introduced GCNH, a simple yet effective GNN architecture that improves representation capabilities on heterophilous graphs. GCNH leverages two design choices: (i) two distinct learnable mapping functions that separately encode the center node and its neighbors into intermediate representations and (ii) the importance coefficients β , one per layer, to balance the contributions of these two representations on the final node embeddings. We analyze and demonstrate how these two components enhance the representation capabilities of GCNH compared to GCN, as, together, they reduce the detrimental contributions of noisy neighborhoods to the updated node representations, while making effective use of basic structural information beneficial for classification. Through extensive experiments on real and synthetic graphs, we show that GCNH performs competitively on the node classification task, outperforming state-of-the-art methods on four out of eight common real-world datasets. In addition to its effectiveness, GCNH’s simple design results in significantly faster training times and fewer trainable parameters compared to competing methods.

REFERENCES

- [1] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. In: *International Conference on Learning Representations (ICLR)*, 2017.
- [2] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. In: *International Conference on Learning Representations (ICLR)*, 2018.
- [3] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs. In: *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [4] Di Jin, Zhizhi Yu, Cuiying Huo, Rui Wang, Xiao Wang, Dongxiao He, and Jiawei Han. Universal Graph Convolutional Networks. In: *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [5] Xin Zheng, Yixin Liu, Shirui Pan, Miao Zhang, Di Jin, and Philip S. Yu. Graph Neural Networks for Graphs with Heterophily: A Survey. In: *arXiv:2202.07082*, 2022.
- [6] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-GCN: Geometric Graph Convolutional Networks. In: *International Conference on Learning Representations (ICLR)*, 2020.
- [7] Yu Wang and Tyler Derr. Tree Decomposed Graph Neural Network. In: *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*, 2021.
- [8] Wei Jin, Tyler Derr, Yiqi Wang, Yao Ma, Zitao Liu, and Jiliang Tang. Node Similarity Preserving Graph Convolutional Networks. In: *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 2021.
- [9] Deyu Bo, Xiao Wang, Chuan Shi, and Huawei Shen. Beyond Low-frequency Information in Graph Convolutional Networks. In: *Conference on Artificial Intelligence (AAAI)*, 2021.
- [10] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, Yaliang Li. Simple and Deep Graph Convolutional Networks. In: *International Conference on Machine Learning (ICML)*, 2020.
- [11] Lun Du, Xiaozhou Shi, Qiang Fu, Xiaojun Ma, Hengyu Liu, Shi Han, and Dongmei Zhang. 2022. GBK-GNN: Gated Bi-Kernel Graph Neural Networks for Modeling Both Homophily and Heterophily. In: *Proceedings of the ACM Web Conference 2022 (WWW '22)*, 2022.
- [12] Yujun Yan, Milad Hashemi, Kevin Swersky, Yaoqing Yang, and Danai Koutra. Two Sides of the Same Coin: Heterophily and Oversmoothing in Graph Convolutional Neural Networks. In: *International Conference on Data Mining (ICDM)*, 2022.
- [13] Cristian Bodnar, Francesco Di Giovanni, Benjamin Paul Chamberlain, Pietro Lio, and Michael M. Bronstein. Neural Sheaf Diffusion: A Topological Perspective on Heterophily and Oversmoothing in GNNs. In: *ICLR 2022 Workshop on Geometrical and Topological Representation Learning*, 2022.
- [14] Eli Chien, Jianhao Peng, Pan Li, and Olgica Milenkovic. Adaptive Universal Generalized PageRank Graph Neural Network. In: *International Conference on Learning Representations (ICLR)*, 2021.
- [15] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. MixHop: Higher-Order Graph Convolutional Architectures via Sparsified Neighborhood Mixing. In: *International Conference on Machine Learning (ICML)*, 2019.
- [16] Dongxiao He, Chungong Liang, Huixin Liu, Mingxiang Wen, Pengfei Jiao, and Zhiyong Feng. Block modeling-guided graph convolutional neural networks. In: *Conference on Artificial Intelligence (AAAI)*, 2022.
- [17] Liang Yang, Mengzhe Li, Liyang Liu, Bingxin Niu, Chuan Wang, Xiaochun Cao, and Yuanfang Guo. Diverse Message Passing for Attribute with Heterophily. In: *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [18] Susheel Suresh, Vinith Budde, Jennifer Neville, Pan Li, and Jianzhu Ma. Breaking the Limit of Graph Neural Networks by Improving the Assortativity of Graphs with Local Mixing Patterns. In: *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2021.
- [19] Yao Ma, Xiaorui Liu, Neil Shah, and Jiliang Tang. Is Homophily a Necessity for Graph Neural Networks? In: *International Conference on Learning Representations (ICLR)*, 2022.
- [20] Andrea Cavallo, Claas Grohnfeldt, Michele Russo, Giulio Lovisotto, and Luca Vassio. 2-Hop Neighbor Class Similarity (2NCS): A Graph Structural Metric Indicative of Graph Neural Network Performance. In: *arXiv:2212.13202*, 2022.
- [21] Miller McPherson, Lynn Smith-Lovin, and James Cook. Birds of a Feather: Homophily in Social Networks. In: *Annual Review of Sociology* 27, 2001.
- [22] Derek Lim, Felix Matthew Hohne, Xiuyu Li, Sijia Linda Huang, Vaishnavi Gupta, Omkar Prasad Bhalerao, and Ser-Nam Lim. Large Scale Learning on Non-Homophilous Graphs: New Benchmarks and Strong Simple Methods. In: *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [23] Kenta Oono and Taiji Suzuki. Graph Neural Networks Exponentially Lose Expressive Power for Node Classification. In: *International Conference on Learning Representations (ICLR)*, 2020.
- [24] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper Insights Into Graph Convolutional Networks for Semi-Supervised Learning. In: *Conference on Artificial Intelligence (AAAI)*, 2018.
- [25] Yujun Yan. Public code for “Two Sides of the Same Coin: Heterophily and Oversmoothing in Graph Convolutional Neural Networks”, [Online]: https://github.com/yujun-yan/heterophily_and_oversmoothing, Accessed on 03-02-2023.
- [26] Twitter Research. Public code for “Neural Sheaf Diffusion”, [Online]: <https://github.com/twitter-research/neural-sheaf-diffusion>, Accessed on 03-02-2023.
- [27] Matthias Fey et al. PyTorch Scatter, [Online]: https://github.com/rusty1s/pytorch_scatter, Accessed on 03-02-2023.
- [28] Sunil Kumar Maurya. Public code for “Improving Graph Neural Networks with Simple Architecture Design”, [Online]: <https://github.com/sunilkmaurya/FSGNN>, Accessed on 03-02-2023.
- [29] Graph-structured Data Mining and Machine Learning at University of Illinois at Urbana-Champaign (UIUC) and Jilin University (JLU). Public code for “Geom-GCN: Geometric Graph Convolutional Networks”, [Online]: <https://github.com/graphdml-uiuc-jlu/geom-gcn>, Accessed on 03-02-2023.
- [30] Carnegie Mellon University, CMU World Wide Knowledge Base, [Online]: <http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-11/www/wwkb/>, Accessed on 03-02-2023.
- [31] Jie Tang, Jimeng Sun, Chi Wang, and Zi Yang. Social Influence Analysis in Large-Scale Networks. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD*, 2009.
- [32] Oleg Platonov, Denis Kuznetsov, Michael Diskin, Artem Babenko and Liudmila Prokhoronkova. A critical look at evaluation of GNNs under heterophily: Are we really making progress? In: *International Conference on Learning Representations (ICLR)* (2023).
- [33] Galileo Mark Namata, Ben London, Lise Getoor, and Bert Huang. Query-driven Active Surveying for Collective Classification. In: *Workshop on Mining and Learning with Graphs*, 2012.
- [34] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective Classification in Network Data. In: *AI Magazine* 29.3, 2008.

Benchmark	Cornell	Texas	Wisconsin	Film	Chameleon	Squirrel	Cora	Citeseer	syn-cora
#Nodes	183	183	251	7,600	2,277	5,201	2,708	3,327	1490
#Edges	280	295	466	26,752	31,421	198,493	1,433	3,703	2965 to 2968
#Classes	5	5	5	5	5	5	7	6	5
#Features	1,703	1,703	1,703	931	2,325	2,089	1,433	3,703	1,433
Homophily h	0.30	0.11	0.21	0.22	0.23	0.22	0.81	0.74	[0, 0.1, ..., 1]

TABLE V: Statistics of real-world and synthetic datasets.

Dataset	#Layers	Batch size	#Epochs	Hidden size	Dropout rate	β
Cornell	{1, 2, 3}	{50, V }	{100, 200, 300 }	{16, 32, 64}	{0.0, 0.25 , 0.5}	0.75
Texas	{1, 2, 3}	{50, V }	{100, 200, 300 }	{16, 32 , 64}	{0.0, 0.25 , 0.5}	0.62
Wisconsin	{1, 2, 3}	{50, V }	{100, 200, 300 }	{16, 32 , 64}	{0.0, 0.3 , 0.6}	0.70, 0.69
Film	{1, 2, 3}	{300, 500 , V }	{ 150 }	{16, 32 , 64}	{0.0, 0.3, 0.6 }	0.72, 0.75
Chameleon	{1, 2, 3}	{ 300 , 600, V }	{500, 1000 , 1500}	{16, 32 , 64}	{ 0.0 , 0.25, 0.5}	0.12
Squirrel	{1, 2, 3}	{300, 1400 , V }	{500, 1000, 1500 }	{16, 32 , 64}	{ 0.0 , 0.25, 0.5}	0.35
Cora	{1, 2, 3}	{ 150 , 300, V }	{ 300 }	{16, 32, 64 }	{0.0, 0.25, 0.5, 0.75 }	0.69, 0.61
Citeseer	{1, 2, 3}	{150, 300 , V }	{ 300 }	{ 16 , 32, 64}	{0.0, 0.25 , 0.5, 0.75}	0.59
syn-cora	{1, 2, 3}	{300, V }	{300}	{16, 32}	{0.0, 0.25, 0.5}	

TABLE VI: Hyperparameters and learned values of β for GCNH. For each dataset, the best hyperparameters are in **bold**. For 2-layer models, the values of β are in the order β^0, β^1 . For syn-cora, the best hyperparameters and the values of β are not reported as they depend on the homophily value of the graph.

APPENDIX A DATASETS

This section and Table V report details about the datasets.

A. Real-World Datasets

These datasets are commonly used in most of the works dealing with heterophilous graphs. Chameleon and Squirrel are taken from [28], the others from [29]. We present further details in the following:

- **Texas, Wisconsin and Cornell** are webpage datasets collected from different universities [30]. Nodes represent web pages and edges are hyperlinks between them.
- **Film (Actor)**, is the actor-only induced subgraph of the film-director-actor-writer network [31]. Nodes correspond to Wikipedia pages of actors and edges denote the co-occurrence of two actors on the same page.
- **Chameleon and Squirrel** are Wikipedia pages on the specific topics of chameleons and squirrels. Nodes are Wikipedia pages and edges are mutual links between them. *Note*: contemporary to this work, [32] highlights drawbacks of these two datasets.
- **Cora and Citeseer** are citation networks where nodes represent papers and edges represent citations of one paper by another [33], [34]. These datasets are treated as undirected.

B. Synthetic Datasets

The synthetic datasets used in this work are taken from [3], which defines a graph generation strategy similar to [15]. The homophily ratio of the graph is defined *a priori*. Then, the graph is generated such that the degree distribution follows a power law. Classes are assigned randomly and features are sampled from nodes of the corresponding class in the real graph Cora. For each of the specified homophily levels, three different graphs are generated.

APPENDIX B

HYPERPARAMETERS AND EXPERIMENTAL DETAILS

We report here additional details about the experiments in this work. All models are trained with Adam optimizer, learning rate $5 \cdot 10^{-3}$ and weight decay $5 \cdot 10^{-3}$.

A. Baselines On Synthetic Datasets

We obtain the results for MLP, GCN and GAT on the synthetic graphs in Figure 2 through hyperparameter optimization among the following values: 100 epochs, {1,2,3} layers, hidden size {16,32}. For all experiments on the synthetic datasets, we randomly generate train/evaluation/test splits with sizes 50%/20%/30%. We compute results on a single split for each dataset and we average the results of three different graphs for each level of homophily ratio reported in the figure.

B. GCNH

The hyperparameters for GCNH on real and synthetic graphs are optimized among the values reported in Table VI. The activation function $\sigma(\cdot)$ in Equations 4 and 5 is LeakyReLU. We perform batching by forwarding, for each batch, the complete graph through the model and computing the loss only on the nodes in the current batch.

C. Experiments On Oversmoothing

The hyperparameters of the models in Section V-G are optimized among the following values: batch size 300, epochs {300,500}, hidden size {16,32,64}, dropout rate {0.0,0.5}. Note that the results for GCN in Table I are taken from [12], while those in Figure 4 are obtained from our own experiments, and they slightly differ because the hyperparameter grid search is different.