
Hyperbolic Benchmarking Unveils Network Topology-Feature Relationship in GNN Performance

Roya Aliakbarisani*

Universitat de Barcelona & UBICS
roya_aliakbarisani@ub.edu

Robert Jankowski*

Universitat de Barcelona & UBICS
robert.jankowski@ub.edu

M. Ángeles Serrano

Universitat de Barcelona, UBICS & ICREA
marian.serrano@ub.edu

Marián Boguñá

Universitat de Barcelona & UBICS
marian.boguna@ub.edu

Abstract

Graph Neural Networks (GNNs) have excelled in predicting graph properties in various applications ranging from identifying trends in social networks to drug discovery and malware detection. With the abundance of new architectures and increased complexity, GNNs are becoming highly specialized when tested on a few well-known datasets. However, how the performance of GNNs depends on the topological and features properties of graphs is still an open question. In this work, we introduce a comprehensive benchmarking framework for graph machine learning, focusing on the performance of GNNs across varied network structures. Utilizing the geometric soft configuration model in hyperbolic space, we generate synthetic networks with realistic topological properties and node feature vectors. This approach enables us to assess the impact of network properties, such as topology-feature correlation, degree distributions, local density of triangles (or clustering), and homophily, on the effectiveness of different GNN architectures. Our results highlight the dependency of model performance on the interplay between network structure and node features, providing insights for model selection in various scenarios. This study contributes to the field by offering a versatile tool for evaluating GNNs, thereby assisting in developing and selecting suitable models based on specific data characteristics.

1 Introduction

Graph Neural Networks (GNNs) [29, 36, 37, 39], derived from Convolutional Neural Networks for graph-structured data, use recursive message passing between nodes and their neighbors. These models leverage graph topology and node-specific features to map nodes into a learnable embedding space. GNNs have evolved to encompass a wide variety of architectures and tasks, ranging from node and graph classifications to link prediction. Despite this growing interest in the development of GNNs, the fundamental issue of homogeneity in benchmarking datasets persists in GNN research, making it challenging to determine the most suitable GNN model for unseen datasets. In addition, since GNNs are data-driven models tailored to specific tasks, there is a potential concern of overfitting new architectures to given datasets, especially when the data have similar structural properties [27]. Thus, a fair comparison between different models in reproducible settings is required.

In this work, we propose a comprehensive benchmarking scheme for graph neural networks. Utilizing a **Hyperbolic Soft Configuration Network Model with Features (HypNF)** [3], we generate synthetic

*These two authors contributed equally

networks with realistic topological properties where node features can be correlated with the network topology. This highly flexible model allows us to evaluate GNNs in depth across various scenarios. Moreover, we suggest using the benchmark as a tool for optimal model selection by analyzing the inherent properties of the real dataset. Although the use of hyperbolic geometry might appear superfluous, it has been demonstrated that it is the simplest method for generating geometric random graphs that uniquely combine several key characteristics: They have power law degree distributions, exhibit small-world properties, and are highly clustered, meaning they have a high density of triangles [6]. These traits closely mirror those observed in real complex networks.

We aim to show the crucial factors, including the network’s structural properties and the degrees of correlation between nodes and features—controlled by the framework parameters—that influence the performance of graph machine learning models. Employing the proposed benchmark, we systematically compare the performance of well-known GNNs and evaluate models that solely utilize node features. Our study aims to evaluate machine learning models in two fundamental graph-based tasks: node classification (NC) and link prediction (LP).

Here, we highlight the main contributions of our empirical study, which provides insights into the suitability of the various models under different network conditions. It will, thus, benefit applications and the community focused on developing new GNN algorithms.

- Our framework generates benchmark networks with tunable levels of topology-feature correlation, homophily, clustering, degree distributions, and average degrees. This approach covers the most important properties of a wide range of real datasets, providing a comprehensive tool for their analysis. The code and the datasets will be publicly available at <https://github.com/networkgeometry/hyperbolic-benchmark-gnn> under the MIT License.
- GNNs exhibit varying levels of performance fluctuation under a fixed set of parameters. Notably, HGCN [9] shows less robustness compared to GCN [18] when the network’s average degree is low. However, this trend reverses in networks with homogeneous degree distributions.
- The stronger the correlation between the network topology and node features, the better GNNs and feature-based models perform in NC and LP.
- The hyperbolic-based models, specifically HGCN and HNN [12], achieve the highest AUC scores [25] in LP task. Remarkably, HNN, despite being solely a feature-based method, outperforms traditional graph-based models across various parameters.
- In the NC task, where no information about the graph data is available, emphasis should be placed on model interpretability and time complexity. This is crucial since the accuracy of graph-based models tends to be uniformly high, making these factors significant differentiators.

2 Related work

With the continuous evolution of graph machine learning, there is a growing necessity to comprehend and evaluate the performance of GNN architectures. In this respect, benchmarking can provide a fair and standardized way to compare different models. The Open Graph Benchmark (OGB) [15] stands as a versatile tool to assess the performance of the GNNs. Yet, its emphasis on a limited range of actual networks indicates that it does not encompass all network characteristics and falls short in terms of parameter manipulation. Consequently, this highlights the necessity for creating benchmarking tools based on synthetic data. Such tools would allow for the assessment of GNNs in a controlled environment and across a more extensive array of network properties [34, 24, 23]. One of them is GraphWorld [27], which is a synthetic network generator utilizing the Stochastic Block Model (SBM) to generate graphs with communities. It employs a parametrized community distribution and an edge probability matrix to randomly assign nodes to clusters and establish connections. Node features are also generated using within-cluster multivariate normal distributions. A fixed edge probability matrix in SBM prevents GraphWorld from faithfully replicating a predefined degree sequence and generating graphs with true power-law distributions. To overcome this limitation, Ref. [38] integrates Graphworld with two other generators: Lancichinetti-Fortunato-Radicchi (LFR) [20] and CABAM (Class-Assortative graphs via the Barabási-Albert Model) [32]. This integration broadens the coverage of the graph space, specifically for the NC task. In this paper, we propose an alternative network generators: a framework based on the geometric soft configuration model. This model’s underlying geometry straightforwardly couples the network topology with node features and labels. This

Table 1: Comparison of HypNF and other synthetic network generators. The \checkmark indicates control over a given property, \times indicates lack of control, and ∞ indicates indirect control. The $P(k)$ is the degree distribution, $\langle c \rangle$ is the clustering coefficient, $\langle k \rangle$ is the average degree. Whereas $P(k_n)$ and $P(k_f)$ are the degree distributions of nodes and features, respectively, $\langle k_n \rangle$ and $\langle k_f \rangle$ are the average degrees of nodes and features, respectively, and \mathcal{H} represents homophily.

	$P(k)$, Power-law	$\langle c \rangle$	$\langle k \rangle$	$P(k_n), P(k_f)$	$\langle k_n \rangle, \langle k_f \rangle$	Topology-Feature Correlation	\mathcal{H}	Underlying mechanism/ Connectivity law
HypNF	\checkmark, \checkmark	\checkmark	\checkmark	\checkmark, \checkmark	\checkmark, \checkmark	\checkmark	\checkmark	geometric
GraphWorld (SBM) [27]	\checkmark , Quasi power-law	∞	\checkmark	\times , normal	\times, \checkmark	∞	\checkmark	probability matrix
GraphWorld (CABAM) [38]	\times, \checkmark	∞	\checkmark	\times , normal	\times, \checkmark	∞	\checkmark	preferential attachment
GraphWorld (LFR) [38]	\checkmark, \checkmark	∞	\checkmark	\times , normal	\times, \checkmark	∞	\checkmark	class-aware configuration model
GenCAT [24, 23]	\checkmark, \checkmark	∞	\checkmark	\times , (normal, Bernoulli)	\times, \checkmark	∞	\checkmark	latent variables
FastSGG [34]	\checkmark, \checkmark	∞	\checkmark	\times, \times	\times, \times	\times	\checkmark	preferential attachment

capability enables independent control over the clustering coefficient in both the unipartite network of nodes and the bipartite network of nodes and features, irrespective of the degree distributions of nodes and features (see Fig. 6 in Appendix B). Table 1 presents a comparison between HypNF and several state-of-the-art benchmarking frameworks, highlighting the properties each can control.

3 HypNF Model

Graph-structured data are characterized by group of N_n nodes that create a complex network \mathcal{G}_n , along with a collection of N_f features linked to these nodes. Typically, these features are converted into binary form. Therefore, for any given node i , its feature set is depicted as a vector $\vec{f}_i \in \{0, 1\}^{N_f}$. This vector has elements that are either zero or one, depending on whether a specific feature is present or absent. An alternative approach considers that nodes and features form a bipartite graph, $\mathcal{G}_{n,f}$, with nodes defining one of the types of elements of the graph and features the other [3]. Within this approach, the full information of the data is encoded into the two networks \mathcal{G}_n and $\mathcal{G}_{n,f}$. Interestingly, in [3], the topological properties of the bipartite graph $\mathcal{G}_{n,f}$ of real graph-structured datasets were first studied, showing a rich and complex topological organization.

A remarkable finding in [3, 17] is the detection of strong correlations between the graphs \mathcal{G}_n and $\mathcal{G}_{n,f}$ within real datasets. This suggested the possibility to describe $\mathcal{G}_{n,f}$ as a geometric random graph in the same hyperbolic space used to describe the graph \mathcal{G}_n so that the shared metric space would mediate the correlation between them. Building on this concept, the study in [3] introduces a generative model that produces networks \mathcal{G}_n and $\mathcal{G}_{n,f}$, with experimental results demonstrating their ability to accurately reproduce key topological properties observed in real datasets, including degree distribution, clustering coefficient, and average nearest neighbor degree function (see Fig. 5 in Appenix A). We will now provide a detailed description of this model, which is used to create synthetic datasets for evaluating the performance of GNNs.

3.1 The $\mathbb{S}^1/\mathbb{H}^2$ model

The network \mathcal{G}_n is modeled by the $\mathbb{S}^1/\mathbb{H}^2$ model [31, 19, 30]. This geometric soft configuration model produces synthetic networks with realistic topological properties, including arbitrary degree distributions [31, 19, 14], the small-world property [1, 11, 26], self-similarity [31], and high clustering [7, 19, 14, 8, 10], to name just a few. See also [6] and references therein.

In the \mathbb{S}^1 model a node is endowed with a hidden degree κ , representing its popularity, influence or importance. Hidden (or expected) degrees are distributed by an arbitrary probability density $\rho(\kappa)$, with $\kappa \in (\kappa_0, \infty)$. In this way, the model has the flexibility to reproduce a variety of degree distributions. Each node is also assigned with an angular position θ in the similarity space, represented as a one-dimensional sphere ². Then, pairs of nodes are connected with a probability function taking the form of a gravity law, balancing the interplay between node angular distances and their hidden degrees:

$$p(\kappa, \kappa', \Delta\theta) = \frac{1}{1 + \chi^\beta} \text{ with } \chi \equiv \frac{R\Delta\theta}{\mu\kappa\kappa'}, \quad (1)$$

²The model with arbitrary dimensions has been defined in [16].

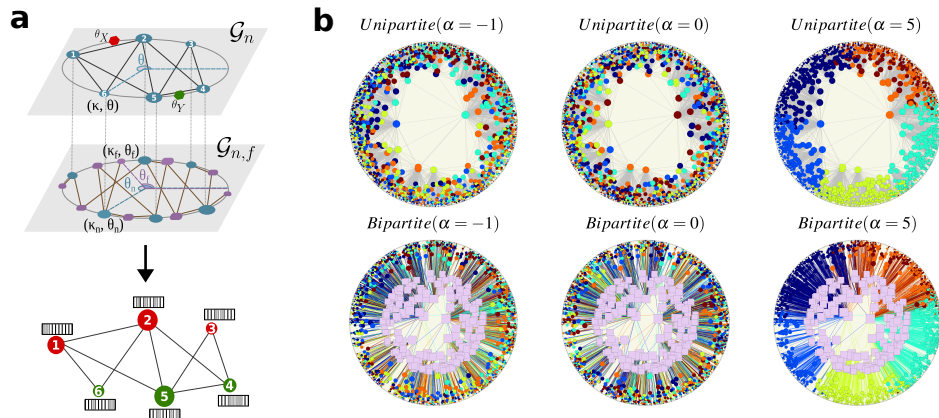


Figure 1: (a) Representation of HypNF benchmarking framework to generate rich graph-structured data. (b) Hyperbolic representation of a synthetic network with $N_n = 2000$ nodes represented as circles, where colors indicate their labels, and $N_f = 200$ features depicted as purple squares. The size of the symbols is proportional to the logarithms of the degrees of nodes and features. The parameters for the \mathbb{S}^1 model are $\beta = 3$, $\gamma = 3.5$, and $\langle k \rangle = 30$, and for the bipartite- \mathbb{S}^1 model, $\beta_b = 3$, $\gamma_n = 3.5$, $\gamma_f = 2.1$, and $\langle k_n \rangle = 3$. Only edges with an effective distance $\chi < 1$ are depicted in the figure.

where $\Delta\theta = \pi - |\pi - |\theta - \theta' ||$ is the angular separation between the nodes in the similarity space, $R = N_n/2\pi$ denotes the radius of the similarity space, $\beta > 1$ governs the level clustering, and

$$\mu = \frac{\beta}{2\pi\langle k \rangle} \sin \frac{\pi}{\beta} \quad (2)$$

controls the average degree of the network $\langle k \rangle$. Notice that the radius is chosen such that in the limit $N_n \gg 1$ the curvature of the similarity space vanishes and the process converges to a Poisson point process on \mathbb{R} of density one. With these choices, the expected degree of a node with hidden degree κ is $\bar{k}(\kappa) = \kappa$, so that by controlling the distribution of hidden degrees, we can adjust the resulting degree distribution.

Interestingly, the \mathbb{S}^1 model exhibit an isomorphism with a purely geometric model in hyperbolic space, referred to as the \mathbb{H}^2 model. The isomorphism is realized by mapping nodes' hidden degrees into radial positions within a disk of radius $R_{\mathbb{H}^2}$ in the hyperbolic plane as

$$r = R_{\mathbb{H}^2} - 2 \ln \frac{\kappa}{\kappa_0}, \quad \kappa \geq \kappa_0, \quad (3)$$

where the radius in \mathbb{H}^2 is given by $R_{\mathbb{H}^2} = 2 \ln \frac{2R}{\mu\kappa_0^2}$. Low degree nodes with hidden degree κ_0 are mapped at the boundary of the hyperbolic disk, whereas high degree nodes get located close to the center of the disk. After this mapping, the connection probability in Eq. (1) becomes

$$p(x) = \frac{1}{1 + e^{\frac{\beta}{2}(x - R_{\mathbb{H}^2})}}, \quad \text{with } x = r + r' + 2 \ln \frac{\Delta\theta}{2} \quad (4)$$

where x approximates the hyperbolic distance between two nodes at radial coordinates r and r' with angular separation of $\Delta\theta$. Thus, in this representation, the probability of connections is just a function of the hyperbolic distance.

3.2 The bipartite- $\mathbb{S}^1/\mathbb{H}^2$ model

In real-world graph-structured data, nodes' features are correlated with the topology of the graph, i.e., nodes that exhibit similarities in the network topology \mathcal{G}_n also share common features [2, 3]. The bipartite- \mathbb{S}^1 model induces this correlation by placing both \mathcal{G}_n and $\mathcal{G}_{n,f}$ in the same similarity space.

In this model, every node is given a pair of hidden variables, (κ_n, θ_n) . Here, κ_n denotes the node's expected degree in the bipartite node-feature graph, while θ_n (equal to θ) represents its angular coordinate, matching that in \mathcal{G}_n . In a similar fashion, features are assigned two hidden variables,

(κ_f, θ_f) , accounting for their expected degrees and their angular locations in a the shared similarity space. The likelihood of a link forming between a node and a feature, characterized by hidden degrees κ_n and κ_f and separated by an angular distance $\Delta\theta$, is expressed as

$$p_b(\kappa_n, \kappa_f, \Delta\theta) = \frac{1}{1 + \chi^{\beta_b}}, \quad \text{where } \chi \equiv \frac{R\Delta\theta}{\mu_b \kappa_n \kappa_f}. \quad (5)$$

In this equation,

$$\mu_b = \frac{\beta_b}{2\pi \langle k_n \rangle} \sin \frac{\pi}{\beta_b} \quad (6)$$

is a crucial parameter that defines the average degree of nodes $\langle k_n \rangle$ and features $\langle k_f \rangle = \frac{N_n}{N_f} \langle k_n \rangle$ and β_b controls bipartite-clustering, as a measure of the coupling between the resulting topology and the underlying metric space. This model, akin to the \mathbb{S}^1 model, guarantees that the expected degrees of nodes and features with hidden degrees κ_n and κ_f are $\bar{k}_n(\kappa_n) = \kappa_n$ and $\bar{k}_f(\kappa_f) = \kappa_f$, respectively. The hidden variables for both nodes and features can be generated from arbitrary distributions or tailored to mimic the topology of a specific real-world network.

As in the case of the \mathbb{S}^1 model, nodes and features in $\mathcal{G}_{n,f}$ can be mapped into the hyperbolic disc of radius $R_{\mathbb{H}^2}^b = 2 \ln \frac{2R}{\mu_b \kappa_{n,0} \kappa_{f,0}}$ through the following transformations

$$r_n = R_{\mathbb{H}^2}^b - 2 \ln \frac{\kappa_n}{\kappa_{n,0}} \quad \text{and} \quad r_f = R_{\mathbb{H}^2}^b - 2 \ln \frac{\kappa_f}{\kappa_{f,0}}, \quad (7)$$

where $\kappa_n \geq \kappa_{n,0}$ and $\kappa_f \geq \kappa_{f,0}$. This mapping leads to a connection probability between nodes and features with the same functional form as in Eq. (4), replacing β by β_b and $R_{\mathbb{H}^2}$ by $R_{\mathbb{H}^2}^b$.

3.3 Assigning labels to nodes

To complete our theoretical framework, we need to specify how labels are assigned to nodes, ensuring a correlation between node labels in \mathcal{G}_n and their features. We use the underlying similarity space to induce and tune these correlations. We introduce a node labeling strategy for any number of clusters \mathcal{N}_L , each associated with a label, and a tunable homophily level [40, 28, 21], crucial for GNNs [22, 35]. First, the cluster centroids are randomly placed in the similarity space. The probability of node i being assigned to cluster X is given by:

$$P_X(i) = \frac{\Delta\theta_{iX}^{-\alpha}}{\sum_{Y \in \mathcal{N}_L} \Delta\theta_{iY}^{-\alpha}}. \quad (8)$$

where $\Delta\theta_{iX}$ is the angular distance between node i and centroid X . The α parameter tunes homophily: a high α assigns labels based on proximity to centroids, while $\alpha = 0$ assigns labels randomly. Negative α values also induce homophily, but nodes are labeled by the farthest centroid.

Fig. 1 (a) illustrates the generation of graph-structured data with N_n nodes, represented as blue circles, and N_f features, depicted as purple rounded squares. The top part represents the unipartite network \mathcal{G}_n , connecting nodes generated with the \mathbb{S}^1 model. The positions of the cluster centroids are shown as θ_X and θ_Y . The sketch is generated with a high value of α so that nodes are assigned labels of the nearest cluster centroids. Below \mathcal{G}_n , we depict the bipartite- \mathbb{S}^1 model where nodes and features share the same similarity space. With these two models, we can generate the synthetic graphs with the class labels and the binarized feature vectors shown at the bottom.

In Fig. 1(b) we depict the hyperbolic representation of one synthetic dataset generated with our model with realistic parameters. Typically, this corresponds to highly clustered, small-world, and heterogeneous degrees in the case of \mathcal{G}_n . As for $\mathcal{G}_{n,f}$, typical topologies have quite homogeneous nodes' degree distribution and very heterogeneous features degree distribution [3]. Each column in Fig. 1(b) shows the two networks \mathcal{G}_n and $\mathcal{G}_{n,f}$ with different values of the parameter α . In the bipartite network, nodes with homogeneous degrees are closer to the edge of the disk, while features exhibiting heterogeneous degrees are closer to the center of the circle, connected to nodes within the same angular sector for higher value of $|\alpha|$.

4 HypNF benchmarking framework

The HypNF benchmarking framework depicted in Fig. 1 combines the $\mathbb{S}^1/\mathbb{H}^2$ and bipartite- $\mathbb{S}^1/\mathbb{H}^2$ models within a unified similarity space. Additionally, it incorporates a method for label assignment. This integration facilitates the creation of networks exhibiting a wide range of structural properties and varying degrees of correlation between nodes and their features. Specifically, our framework allows us to control the following properties:

- **Degree distributions.** The framework allows us to generate networks with arbitrary node and feature degree distributions by fixing the sets of hidden degrees in \mathcal{G}_n and $\mathcal{G}_{n,f}$. A useful choice is the power-law distribution

$$\rho(\kappa) = (\gamma - 1)\kappa_0^{\gamma-1}\kappa^{-\gamma} \quad \text{with} \quad \kappa_0 = \frac{\gamma - 2}{\gamma - 1}\langle k \rangle$$

and parameter γ controlling the heterogeneity of the degree distribution. This leaves us with three parameters $\gamma, \gamma_n, \gamma_f$ fixing the hidden degree distributions in the \mathbb{S}^1 and bipartite- \mathbb{S}^1 models, respectively.

- **Network average connectivity.** Another critical factor is the network average connectivity. We have the freedom to choose the average degree of the network $\mathcal{G}_n, \langle k \rangle$, as well as the average degree of nodes in the bipartite node-feature network $\mathcal{G}_{n,f}, \langle k_n \rangle$, by changing μ and μ_b parameters in Eqs. (2) and (6). Notice that the average degree of features is fixed by the identity $\langle k_f \rangle = \frac{N_n}{N_f}\langle k_n \rangle$.
- **Clustering and topology-features correlations.** The level of clustering is regulated by parameters β and β_b , binding the topology of both \mathcal{G}_n and $\mathcal{G}_{n,f}$ to the common similarity space through the triangle inequality. A higher value of β implies that pairs of nodes that are close in the hyperbolic space (with small χ) are more likely to get connected in \mathcal{G}_n . Likewise, a high value β_b indicates that nodes and features separated by a short hyperbolic distance have a high probability to be connected in $\mathcal{G}_{n,f}$. Consequently, when $\beta > 1$ and $\beta_b > 1$ similar nodes tend to be connected and, at the same time, share common features. Thus, β and β_b not only influence the clustering in the networks but also adjust the correlation between topology and node features. We should stress that the ability to directly control the level of clustering, and thereby topology-feature correlations, is one of the key distinctive features of HypNF as compared to other benchmarking schemes.
- **Homophily.** For the task of NC, \mathcal{N}_l identifies the number of classes and α governs the strength of the nodes' labels concentration around their centroids, allowing us to control the level of homophily in the network, defined as [28]

$$\mathcal{H} = \frac{1}{N_n} \sum_{i=1}^{N_n} \frac{n_l(i)}{k_i}, \quad (9)$$

where $n_l(i)$ is the number of neighbors of node i with the same label as i and k_i its degree. A high value of \mathcal{H} indicates a tendency for similar nodes to connect with each other. Fig. 8(a) in Appendix E depicts the relation between \mathcal{H} and the parameter α , where a monotonic dependence on $|\alpha|$ can be observed. However, this relation is not symmetric for $\mathcal{N}_l > 2$. The lower \mathcal{H} for negative α stems from the angular organization of the labeled nodes. Indeed, Fig. 8(b) in Appendix E shows that for the high positive α , the distance between two maximally distant nodes is lower compared to the same negative α value. Notice that $\alpha = 0$ corresponds to a maximally random assignment of labels, which results in the expected homophily parameter taking the value $\langle \mathcal{H} \rangle = 1/\mathcal{N}_l$.

Leveraging the HypNF model with varying parameters, our benchmarking framework generates diverse graph-structured data. This allows for the evaluation of graph machine learning models on networks with different connectivity patterns and correlations between topology and node features. For tasks like NC and LP, the framework facilitates fair model comparisons, helping to assess a novel GNN against state-of-the-art architectures and providing insights into the data's impact on performance.

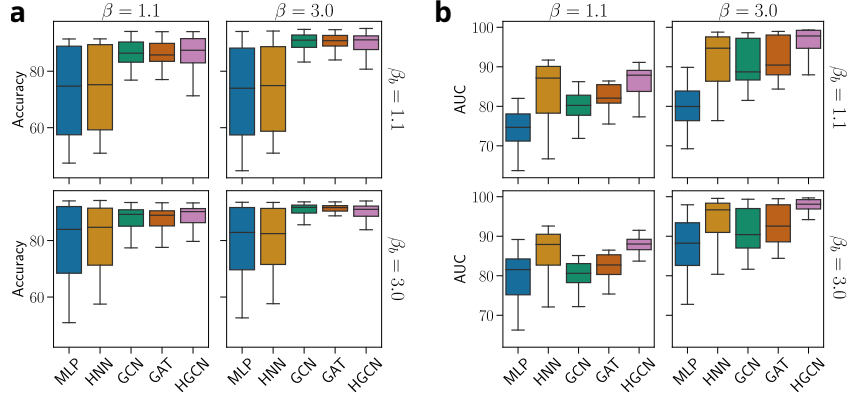


Figure 2: The impact of topology-feature correlation controlled by β and β_b on the performance of graph machine learning models in two tasks: (a) node classification and (b) link prediction. We set $\mathcal{N}_L = 6$ and $\alpha = 10$ for the NC task. The box ranges from the first quartile to the third quartile. A horizontal line goes through the box at the median. The whiskers go from each quartile to the minimum or maximum. The results are averaged over all other parameters.

5 Experiments

5.1 Parameter Space

We use the HypNF benchmarking framework introduced in Section 4 and customize its parameters to generate a collection of datasets with $N_n = 5000$ nodes and $N_f = 2000$ features, displaying a wide spectrum of properties. The degree distributions for nodes in \mathcal{G}_n and for nodes and features in $\mathcal{G}_{n,f}$ are tailored to exhibit either heterogeneity or homogeneity. This is achieved by setting $\gamma, \gamma_n, \gamma_f$ to the values $\{2.1, 3.5\}$. The average degrees vary between low and high values by adjusting $\langle k \rangle$, and $\langle k_n \rangle$ to $\{3, 30\}$. The clustering coefficient in both networks, along with the degree of correlation between them, is set to range between low and high values by fixing β and β_b to $\{1.1, 3\}$. For the task of NC, the number of labels are $\mathcal{N}_l = \{2, 3, 6, 10\}$. Homophily of node labels is varied by setting $\alpha = \{-1, 1, 5, 10\}$.

For each unique combination of parameter values, we generate 10 network realizations with node features, resulting in a comprehensive benchmark for model evaluation. Consequently, for the LP task, machine learning models are tested on a benchmark comprising $2^7 \times 10 = 1280$ instances. In the case of NC, where we introduce two additional parameters, α and \mathcal{N}_l , the size of the benchmark is $2^7 \times 4^2 \times 10 = 20480$. In addition, we evaluate the performance of GNNs using AUC for the LP task [25], and accuracy defined as the fraction of correctly classified nodes for NC task [5].

5.2 Machine learning models

In this work, we focus on two primary methodologies: feature-based methods, which entail node embedding based on their features, and GNNs, which integrate both features and network topology.

- **MLP**: A vanilla neural network transforms node feature vectors through linear layers and non-linear activations to learn embeddings in Euclidean space.
- **HNN** [12]: A variant of MLP that operates in hyperbolic space to capture complex patterns and hierarchical structures.
- **GCN** [18]: A pioneering model that averages the states of neighboring nodes at each iteration.
- **GAT** [33]: A model that uses attention mechanisms to assign different importances to different nodes in a neighborhood.
- **HGCN** [9]: A model that integrates hyperbolic geometry with graph convolutional networks to capture complex structures in graph data more effectively.

Table 2 in Appendix C lists the hyperparameters for training. In the LP task, links are split into training (85%), validation (5%), and test (10%) sets. For the NC task, nodes are distributed as 70%

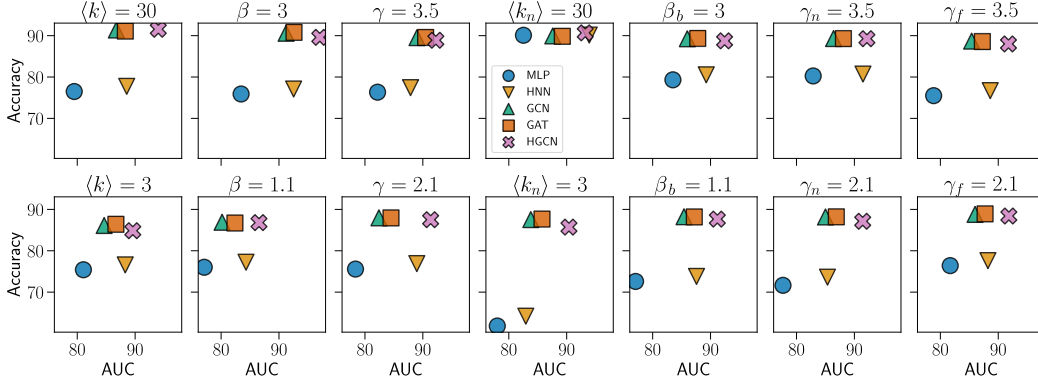


Figure 3: Impact of each individual parameter on the performance of NC and LP. In the case of NC, we set $\mathcal{N}_L = 6$ and $\alpha = 10$.

training, 15% validation, and 15% test [9]. Both tasks follow the methodology in [9], with results averaged over five test-train splits. Models were trained on an NVIDIA GeForce RTX 3080 GPU using Python 3.9, CUDA 11.7, and PyTorch 1.13.

6 Results

We begin by analyzing how topology-feature correlation affects the performance of machine learning models by varying parameters β and β_b , which control the coupling between \mathcal{G}_n , $\mathcal{G}_{n,f}$, and the shared metric similarity space. For both NC and LP tasks, Fig. 2 shows that higher topology-feature correlation ($\beta = \beta_b = 3$) improves performance for both feature-based and GNN models compared to lower correlation ($\beta = \beta_b = 1.1$). In the NC task, feature-based models do not respond to changes in β and perform better with higher β_b . Conversely, GNN models, utilizing information from both \mathcal{G}_n and $\mathcal{G}_{n,f}$, benefit from high β and β_b values, not only in terms of average performance but also in terms of the spread around this average.

We measured the performance of machine learning models by varying the clustering level in \mathcal{G}_n (adjusted by β) and the average number of features per node in the bipartite network, $\langle k_n \rangle$. For the NC task, Fig. 9(a) in Appendix F shows significant performance differences between feature-based and GNN models. In bipartite networks with low $\langle k_n \rangle$, GNN models outperform feature-based ones. As $\langle k_n \rangle$ increases (bottom row), all models' accuracy improves, converging to almost similar levels. Here, the simplest model, MLP, performs nearly as well as the most sophisticated, HGNCN. Note that these results fix β and $\langle k_n \rangle$, averaging over other parameters. Fig. 9(b) in Appendix F compares these models in terms of AUC for the LP task. The models show a significant shift in relative performance with varying $\langle k_n \rangle$. For high $\langle k_n \rangle$, models in hyperbolic space (HNN and HGNCN) perform better. Notably, the feature-based model, HNN, matches or exceeds the performance of HGNCN, especially at low β .

In Fig. 3, we present a summary of model performance on both downstream tasks, varying parameters individually between high and low values while averaging over others. Across most cases, HGNCN outperforms or matches others in the LP task, and achieves competitive accuracy in the NC task, akin to other graph-based methods. In addition, for the NC task, the feature-based methods (MLP and HNN) are barely sensible to \mathbb{S}^1 model parameters, i.e., $\langle k \rangle$, β and γ . However, for the LP task, the topology-features correlation strongly impacts their AUC values, leading to MLP and HNN being sensitive to the parameters of the bipartite network, including β_b , γ_n , and γ_f , with a specific emphasis on $\langle k_n \rangle$. In contrast, GNNs are particularly responsive to variations in $\langle k_n \rangle$. It is worth mentioning that these observations highlight the sensitivity of graph machine learning methods to individual parameters. However, in real-world scenarios, the collective interplay of all parameters influences the overall performance of the models, see the detailed analysis in Appendix F.

Another important factor that is usually overlooked concerns the fluctuations of the performance of machine learning models, which provides insights into their robustness and reliability. We address this problem by analyzing the difference in the standard deviation of the accuracy and AUC for a given set of parameters. In Fig. 7 of Appendix D, we show how the fluctuation changes when each

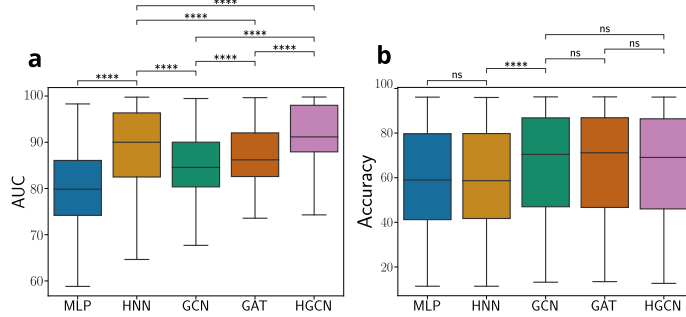


Figure 4: Aggregated results for (a) link prediction and (b) node classification tasks. The annotations denote the statistical significance levels derived from the Mann-Whitney U test, a non-parametric test suitable for comparing two independent samples, especially when the data distribution is not assumed to be normal. This test evaluates if there is a statistically significant difference in the median performance scores between two different models. The label *ns* signifies p-value ≤ 1 , **** corresponds to p-value $\leq 10^{-4}$.

parameter is set to high or low. Let us take HGCN as an example. For the high average degree $\langle k \rangle$, the accuracy and AUC display lower standard deviations as compared to low average degree, which means that HGCN is more robust to other parameters when the average degree is high. Moreover, a homogeneous degree distributions ($\gamma = 3.5$) results in a broader spread of AUC but not of accuracy. GCN and GAT display similar fluctuation behaviors, whereas for the feature-based models (MLP, HNN), the bipartite- \mathbb{S}^1 parameters dictate their sensibility to other parameters.

The global average picture of the performance of the analyzed models is shown in Fig. 4, where we averaged results over all the parameters in the HypNF benchmarking framework. This analysis sheds light on model selection when no structural information is available about the data. We observe the superiority of the hyperbolic-based models in capturing essential network properties and connectivities for the LP task (Fig. 4(a)). Yet, HGCN outperforms HNN with statistical significance. As for the NC task, results are more ambiguous, with GNNs outperforming feature-based models overall. However, the distinctions in performance within the GNNs are not substantial. Hence, when dealing with unseen data, the simplicity of the GCN model makes it a more efficient choice.

Finally, in Appendix F, we carry out a comprehensive analysis of the performance of the selected models, focusing on all parameters within the $\mathbb{S}^1/\mathbb{H}^2$ (Figs. 10(a) and 11(a)), bipartite- $\mathbb{S}^1/\mathbb{H}^2$ (Figs. 10(b) and 11(b)), and the combination of the two (Figs. 12 and 13). Moreover, Fig. 14 reveals the dependence of the number of labels and the homophily level in terms of accuracy.

7 Conclusion

In this study, we introduced HypNF benchmarking framework for graph machine learning, and in particular for graph neural networks. This framework, leveraging the $\mathbb{S}^1/\mathbb{H}^2$ and bipartite- $\mathbb{S}^1/\mathbb{H}^2$ models, enables the generation of synthetic networks with controllable properties such as degree distributions, average degrees, clustering coefficients, and homophily levels. Our findings underscore the significance of topology-feature correlations in influencing GNN performance. Specifically, models embedded in the hyperbolic space, with its intrinsic hierarchical structure, outperformed the rest of the models in the LP task. This research contributes to the broader understanding of graph machine learning, providing insights into the suitability of various models under different network conditions. Furthermore, our benchmarking framework serves as a valuable tool for the community, enabling fair and standardized comparisons of new GNN architectures against established models. A promising direction for future work is to extend the bipartite- $\mathbb{S}^1/\mathbb{H}^2$ model to incorporate weighted features. This will enable HypNF to generate networks with non-binary features, allowing for more nuanced and flexible representations [4]. Additionally, the framework can be extended to accommodate community structures by using a non-homogeneous distribution of nodes within the $\mathbb{S}^1/\mathbb{H}^2$ similarity space, as done in [41] and [13].

Acknowledgments and Disclosure of Funding

We acknowledge support from: Grant TED2021-129791B-I00 funded by MCIN/AEI/10.13039/501100011033 and the “European Union NextGenerationEU/PRTR”; Grant PID2022-137505NB-C22 funded by MCIN/AEI/10.13039/501100011033; Generalitat de Catalunya grant number 2021SGR00856. R. J. acknowledge support from the fellowship FI-SDUR funded by Generalitat de Catalunya. M. B. acknowledges the ICREA Academia award, funded by the Generalitat de Catalunya.

References

- [1] M. A. Abdullah, N. Fountoulakis, and M. Bode. 2017. Typical distances in a geometric model for complex networks. *Internet Math.* 1 (2017). <https://doi.org/10.24166/im.13.2017>
- [2] Sarwan Ali, Muhammad Haroon Shakeel, Imdadullah Khan, Safiullah Faizullah, and Muhammad Asad Khan. 2021. Predicting attributes of nodes using network structure. *ACM Transactions on Intelligent Systems and Technology (TIST)* 12, 2 (2021), 1–23.
- [3] Roya Aliakbarisani, M. Ángeles Serrano, and Marián Boguñá. 2023. Feature-enriched hyperbolic network geometry. arXiv:2307.14198 [physics.soc-ph]
- [4] Antoine Allard, M. Ángeles Serrano, Guillermo García-Pérez, and Marián Boguñá. 2017. The geometric nature of weights in real complex networks. *Nature Communications* 8, 1 (2017), 14103. <https://doi.org/10.1038/ncomms14103>
- [5] Christopher Bishop. 2006. *Pattern recognition and machine learning*. Springer, New York, NY.
- [6] M. Boguñá, I. Bonamassa, M. De Domenico, S. Havlin, D. Krioukov, and M. Á. Serrano. 2021. Network geometry. *Nat. Rev. Phys.* 3 (2021), 114–135. <https://doi.org/10.1038/s42254-020-00264-4>
- [7] M. Boguñá, D. Krioukov, P. Almagro, and M. Á. Serrano. 2020. Small worlds and clustering in spatial networks. *Phys. Rev. Res.* 2 (Apr 2020), 023040. Issue 2. <https://doi.org/10.1103/PhysRevResearch.2.023040>
- [8] E. Candellero and N. Fountoulakis. 2016. Clustering and the Hyperbolic Geometry of Complex Networks. *Internet Math.* 12, 1-2 (2016), 2–53. <https://doi.org/10.1080/15427951.2015.1067848>
- [9] I. Chami, Z. Ying, Ch. Ré, and J. Leskovec. 2019. Hyperbolic Graph Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2019/file/0415740eaa4d9decbc8da001d3fd805f-Paper.pdf
- [10] N. Fountoulakis, P. van der Hoorn, T. Müller, and M. Schepers. 2021. Clustering in a hyperbolic model of complex networks. *Electron. J. Probab.* 26 (2021), 1 – 132. <https://doi.org/10.1214/21-EJP583>
- [11] T. Friedrich and A. Krohmer. 2018. On the Diameter of Hyperbolic Random Graphs. *SIAM J. Discrete Math.* 32, 2 (2018), 1314–1334. <https://doi.org/10.1137/17M1123961>
- [12] O. Ganea, G. Bécigneul, and T. Hofmann. 2018. Hyperbolic neural networks. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc.
- [13] Guillermo García-Pérez, M. Ángeles Serrano, and Marián Boguñá. 2018. Soft Communities in Similarity Space. *Journal of Statistical Physics* 173, 3 (2018), 775–782. <https://doi.org/10.1007/s10955-018-2084-z>
- [14] L. Gugelmann, K. Panagiotou, and U. Peter. 2012. Random Hyperbolic Graphs: Degree Sequence and Clustering. In *Autom Lang Program (ICALP 2012, Part II)*, LNCS 7392. https://doi.org/10.1007/978-3-642-31585-5_51

- [15] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. 2020. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems* 33 (2020), 22118–22133.
- [16] Robert Jankowski, Antoine Allard, Marián Boguñá, and M Ángeles Serrano. 2023. The D-Mercator method for the multidimensional hyperbolic embedding of real networks. *Nature Communications* 14, 1 (2023), 7585.
- [17] Robert Jankowski, Pegah Hozhabrierdi, Marián Boguñá, and M. Á Serrano. 2024. Feature-aware ultra-low dimensional reduction of real networks. *arXiv preprint arXiv:2401.09368* (2024).
- [18] T. N. Kipf and M. Welling. 2017. Semi-supervised Classification With Graph Convolutional Networks. In *International Conference on Learning Representations*.
- [19] D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguñá. 2010. Hyperbolic geometry of complex networks. *Phys. Rev. E* 82, 3 (2010), 036106. <https://doi.org/10.1103/PhysRevE.82.036106>
- [20] A. Lancichinetti, S. Fortunato, and F. Radicchi. 2008. Benchmark graphs for testing community detection algorithms. *Physical Review E* 78 (Oct 2008), 046110. Issue 4. <https://doi.org/10.1103/PhysRevE.78.046110>
- [21] Derek Lim, Xiuyu Li, Felix Hohne, and Ser-Nam Lim. 2021. New benchmarks for learning on non-homophilous graphs. *arXiv preprint arXiv:2104.01404* (2021).
- [22] Yao Ma, Xiaorui Liu, Neil Shah, and Jiliang Tang. 2021. Is homophily a necessity for graph neural networks? *arXiv preprint arXiv:2106.06134* (2021).
- [23] Seiji Maekawa, Koki Noda, Yuya Sasaki, et al. 2022. Beyond real-world benchmark datasets: An empirical study of node classification with GNNs. *Advances in Neural Information Processing Systems* 35 (2022), 5562–5574.
- [24] Seiji Maekawa, Yuya Sasaki, George Fletcher, and Makoto Onizuka. 2023. GenCAT: Generating attributed graphs with controlled relationships between classes, attributes, and topology. *Information Systems* 115 (2023), 102195. <https://doi.org/10.1016/j.is.2023.102195>
- [25] Francisco Melo. 2013. *Area under the ROC Curve*. Springer, New York, NY, 38–39. https://doi.org/10.1007/978-1-4419-9863-7_209
- [26] T. Müller and M. Staps. 2019. The diameter of KPKVB random graphs. *Adv. Appl. Probab.* 51, 2 (2019), 358–377. <https://doi.org/10.1017/apr.2019.23>
- [27] J. Palowitch, A. Tsitsulin, B. Mayer, and B. Perozzi. 2022. GraphWorld: Fake Graphs Bring Real Insights for GNNs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Washington DC, USA) (KDD '22)*. Association for Computing Machinery, New York, NY, USA, 3691–3701. <https://doi.org/10.1145/3534678.3539203>
- [28] H. Pei, B. Wei, K. C. Chang, Y. Lei, and B. Yang. 2020. Geom-GCN: Geometric Graph Convolutional Networks. In *International Conference on Learning Representations*.
- [29] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE transactions on neural networks* 20, 1 (2008), 61–80.
- [30] M. Á. Serrano and M. Boguñá. 2022. *The Shortest Path to Network Geometry: A Practical Guide to Basic Models and Applications*. Cambridge University Press. <https://doi.org/10.1017/9781108865791>
- [31] M. Á. Serrano, D. Krioukov, and M. Boguñá. 2008. Self-Similarity of Complex Networks and Hidden Metric Spaces. *Phys. Rev. Lett.* 100, 7 (2008), 078701. <https://doi.org/10.1103/PhysRevLett.100.078701>
- [32] Neil Shah. 2020. Scale-Free, Attributed and Class-Assortative Graph Generation to Facilitate Introspection of Graph Neural Networks. In *KDD Mining and Learning with Graphs*.

- [33] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- [34] Chaokun Wang, Binbin Wang, Bingyang Huang, Shaoxu Song, and Zai Li. 2021. FastSGG: Efficient Social Graph Generation Using a Degree Distribution Generation Model. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. 564–575. <https://doi.org/10.1109/ICDE51399.2021.00055>
- [35] Junfu Wang, Yuanfang Guo, Liang Yang, and Yunhong Wang. 2024. Understanding Heterophily for Graph Neural Networks. *arXiv preprint arXiv:2401.09125* (2024).
- [36] Lingfei Wu, Peng Cui, Jian Pei, Liang Zhao, and Xiaojie Guo. 2022. Graph Neural Networks: Foundation, Frontiers and Applications. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Washington DC, USA) (KDD '22)*. Association for Computing Machinery, New York, NY, USA, 4840–4841. <https://doi.org/10.1145/3534678.3542609>
- [37] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2021. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1 (2021), 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
- [38] M. Yasir, J. Palowitch, A. Tsitsulin, L. Tran-Thanh, and B. Perozzi. 2023. Examining the Effects of Degree Distribution and Homophily in Graph Learning Models. *arXiv:2307.08881*
- [39] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. *AI Open* 1 (2020), 57–81. <https://doi.org/10.1016/j.aiopen.2021.01.001>
- [40] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. 2020. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in neural information processing systems* 33 (2020), 7793–7804.
- [41] Konstantin Zuev, Marián Boguñá, Ginestra Bianconi, and Dmitri Krioukov. 2015. Emergence of Soft Communities from Geometric Preferential Attachment. *Scientific Reports* 5, 1 (2015), 9421. <https://doi.org/10.1038/srep09421>

A Empirical validation of HypNF

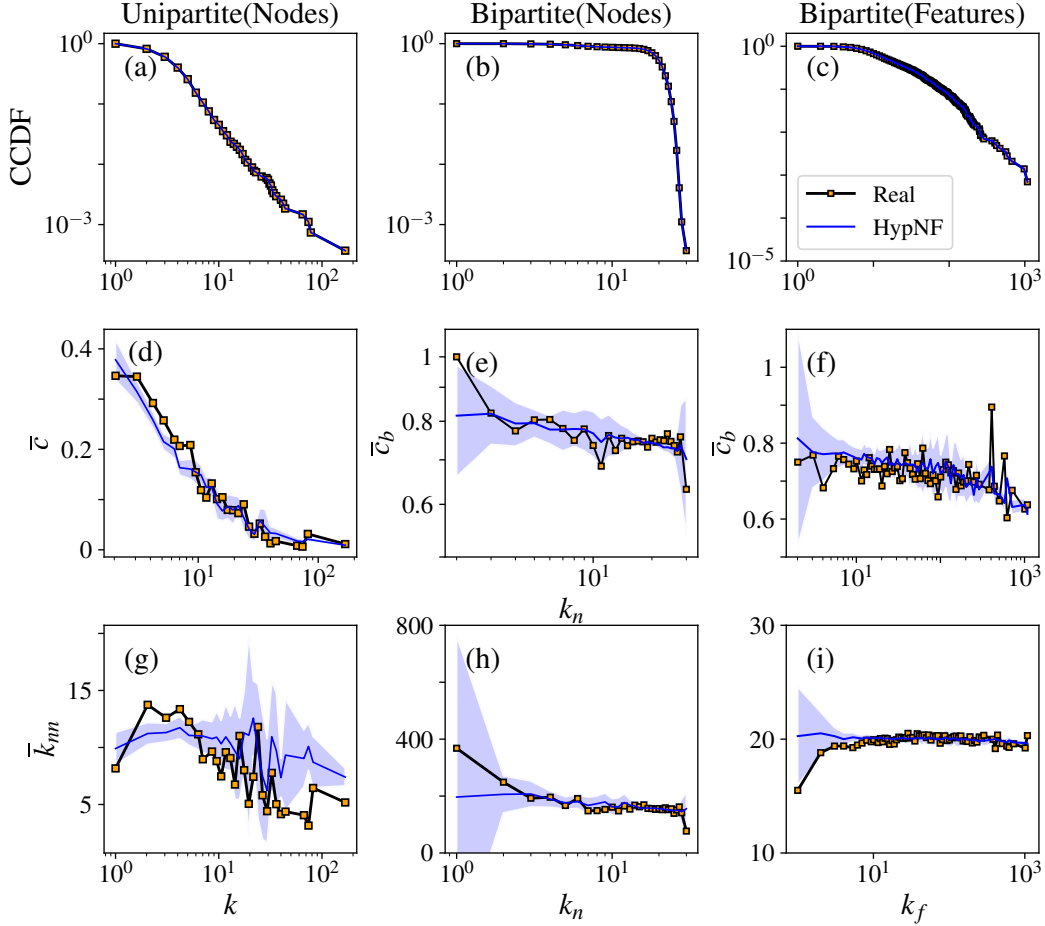


Figure 5: Topological properties of Cora dataset and its synthetic counterpart generated by HypNF. (a)-(c) Complementary cumulative degree distribution of nodes in \mathcal{G}_n and those of nodes and features in $\mathcal{G}_{n,f}$. (d)-(f) clustering spectrum of nodes as a function of node degrees in \mathcal{G}_n and bipartite clustering spectrum of nodes and features as a function of nodes and features degrees in $\mathcal{G}_{n,f}$. (g)-(i) average nearest neighbors degree function in \mathcal{G}_n and $\mathcal{G}_{n,f}$. The clustering coefficient, c , for each node in \mathcal{G}_n is calculated as the fraction of connected pairs among its neighboring nodes relative to the total possible pairs of neighbors. To compute the bipartite clustering, c_b , for a given node, each pair of its neighboring features is considered connected if they share at least one common node, apart from the node in question. The clustering coefficient is then computed similarly to the clustering coefficient in a unipartite network. This definition also applies to the bipartite clustering coefficient of features. Finally, k_{nn} in both \mathcal{G}_n and $\mathcal{G}_{n,f}$ quantifies the average degree of the neighbors for a given node or feature. Exponential binning is used in the computation \bar{c} , \bar{c}_b and \bar{k}_{nn} . The blue shaded region indicates the two- σ intervals around the mean, derived from 100 realizations of HypNF.

B Degree distribution and clustering control in HypNF

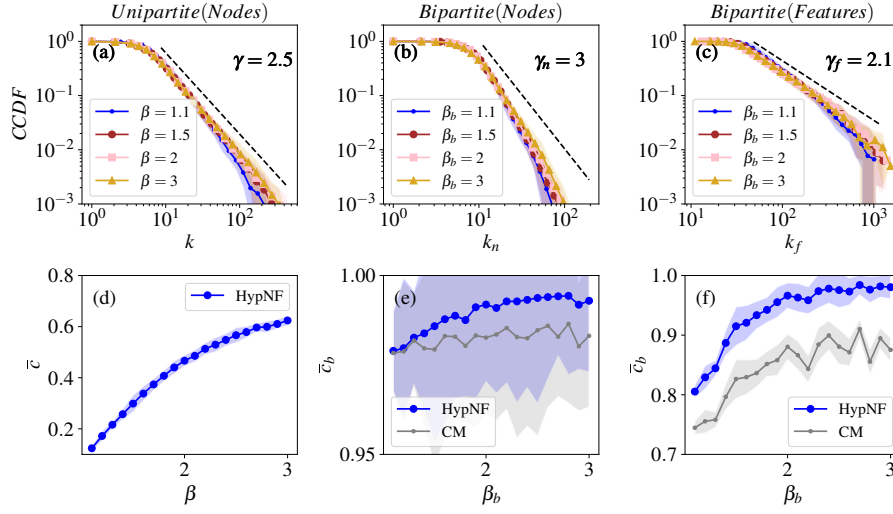


Figure 6: Independent control of degree distribution and clustering coefficient in synthetic networks generated by HypNF. (a)-(c) Complementary cumulative degree distribution of nodes in \mathcal{G}_n ($N = 2000$, $\gamma = 2.5$, and $\langle k \rangle = 10$) and those of nodes and features in $\mathcal{G}_{n,f}$ ($N_n = 2000$, $N_f = 200$, $\langle k_n \rangle = 10$, $\gamma_n = 3$, and $\gamma_f = 2.1$) for different values of β and β_b . The dashed black lines are guides for the eyes corresponding to the power-law exponents. (d)-(e) Clustering and bipartite clustering coefficients as a function of β and β_b . The gray curves in (e) and (f) illustrate the bipartite clustering in the randomized versions of the bipartite synthetic networks, generated according to the configuration model (CM).

C Hyperparameters of the machine learning models

Table 2: Hyperparameters of the machine learning models.

Model	Dropout	Weight Decay	Optimizer	Activation	Layers	Dimensions
MLP	0.2	0.001	Adam	-	2	16
HNN	0.2	0.001	Adam	-	2	16
GCN	0.2	0.0005	Adam	ReLU	2	16
GAT	0.2	0.0005	Adam	ReLU	2	16
HGCN	0.2	0.0005	Adam	ReLU	2	16

D Fluctuations in the performance of machine learning models

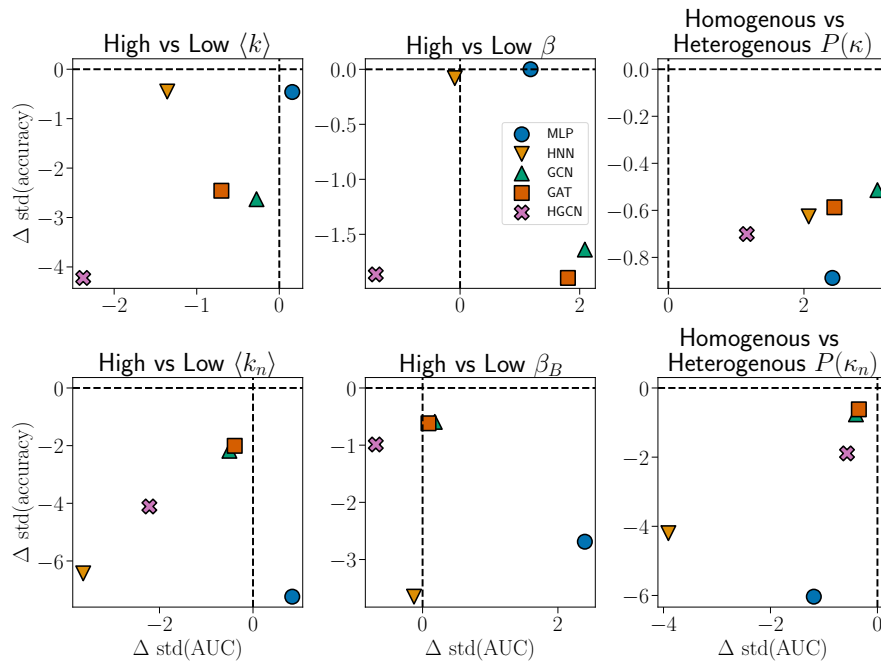


Figure 7: Fluctuations of performance of the machine learning models. The difference of the standard deviation between high and low value of each parameter. We set $\mathcal{N}_L = 6$ and $\alpha = 10$ for the NC task. The results are averaged over all other parameters.

E Homophily in the synthetic networks

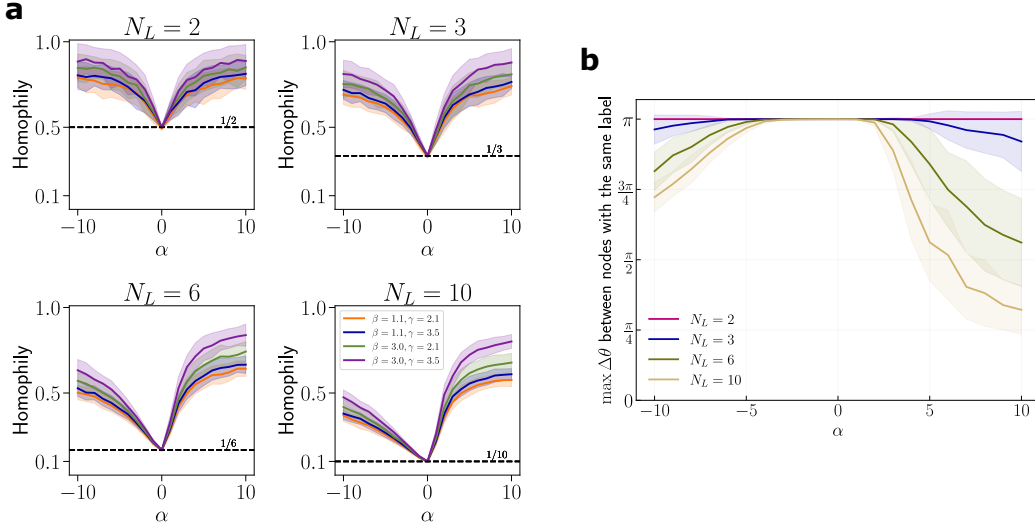


Figure 8: (a) Value of homophily in the synthetic networks. Parameters of the networks: $N = 1000$, $\langle k \rangle = 30$. The black horizontal line indicates the random case, i.e., for $\alpha = 0$ which based on Eq. 8 gives us $\mathcal{H} = 1/N_L$ where N_L is the number of labels. The results are averaged over 100 realizations. (b) Maximum angular distance between nodes in the same community in function of parameter α . The results are averaged over 100 realizations.

F Exploring the parameters' space

Here, we explore the impact of parameters within the HypNF benchmarking framework including the parameters in $\mathbb{S}^1/\mathbb{H}^2$ and bipartite- $\mathbb{S}^1/\mathbb{H}^2$ models, as well as the combinations of both on the performance of graph machine learning techniques.

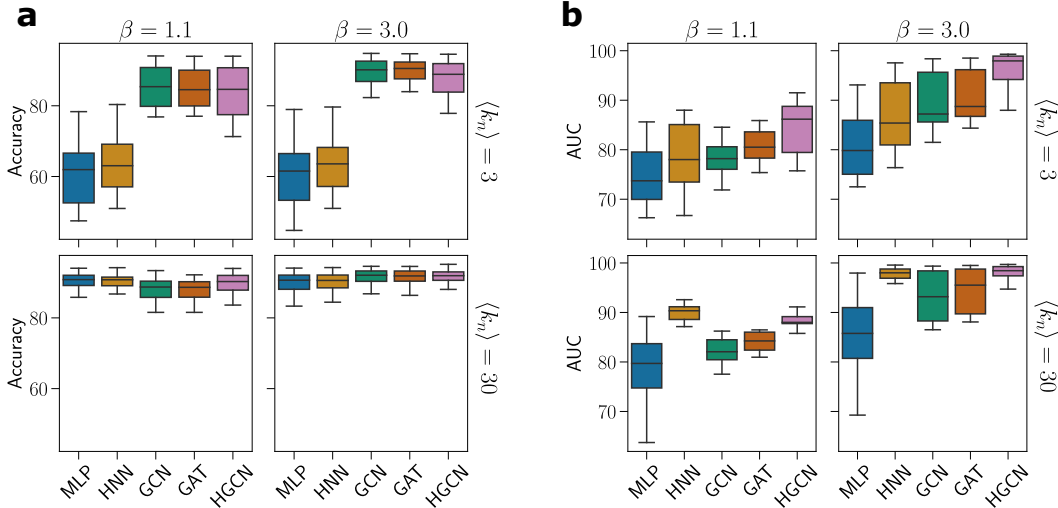


Figure 9: The impact of the level of clustering β and average degree of nodes in the bipartite network $\langle k_n \rangle$ on the performance of machine learning models in two tasks: (a) node classification and (b) link prediction. We set $N_L = 6$ and $\alpha = 10$ for the NC task. The box ranges from the first quartile to the third quartile. A horizontal line goes through the box at the median. The whiskers go from each quartile to the minimum or maximum. The results are averaged over all other parameters.

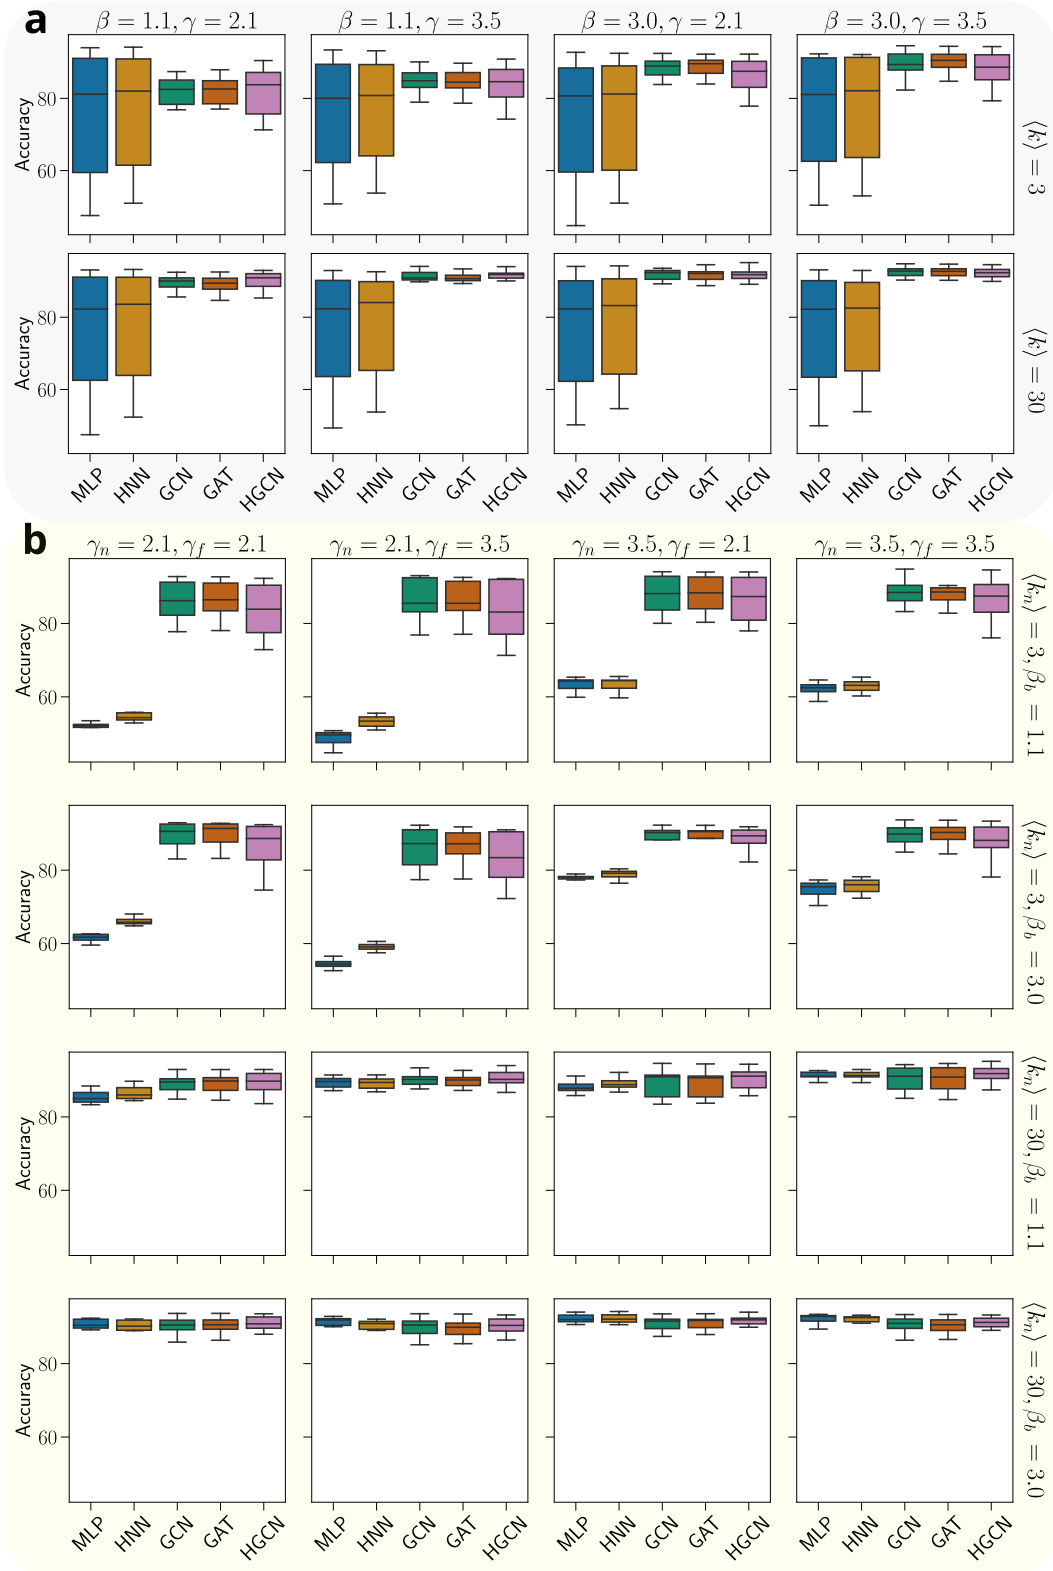


Figure 10: Performance comparison of machine learning models on NC task with respect to (a) $\mathbb{S}^1/\mathbb{H}^2$ model parameters and (b) bipartite- $\mathbb{S}^1/\mathbb{H}^2$ model parameters. We set $\mathcal{N}_L = 6$ and $\alpha = 10$.

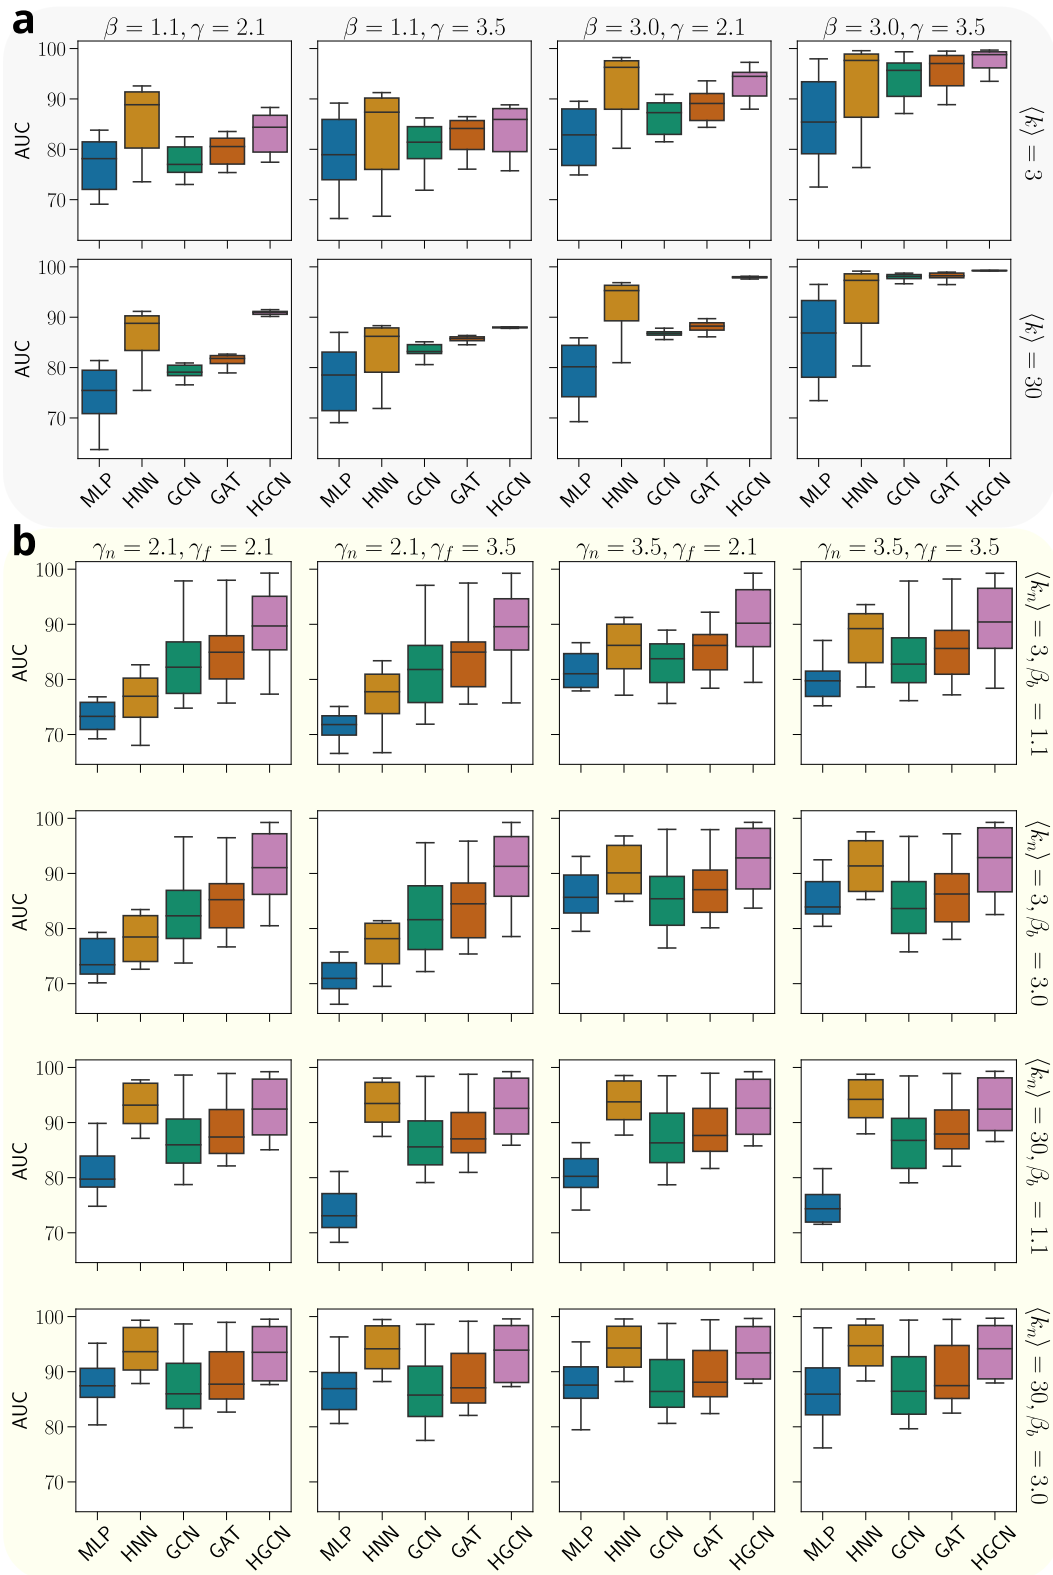


Figure 11: Performance comparison of machine learning models on LP task with respect to (a) $\mathbb{S}^1/\mathbb{H}^2$ model parameters and (b) bipartite- $\mathbb{S}^1/\mathbb{H}^2$ model parameters.

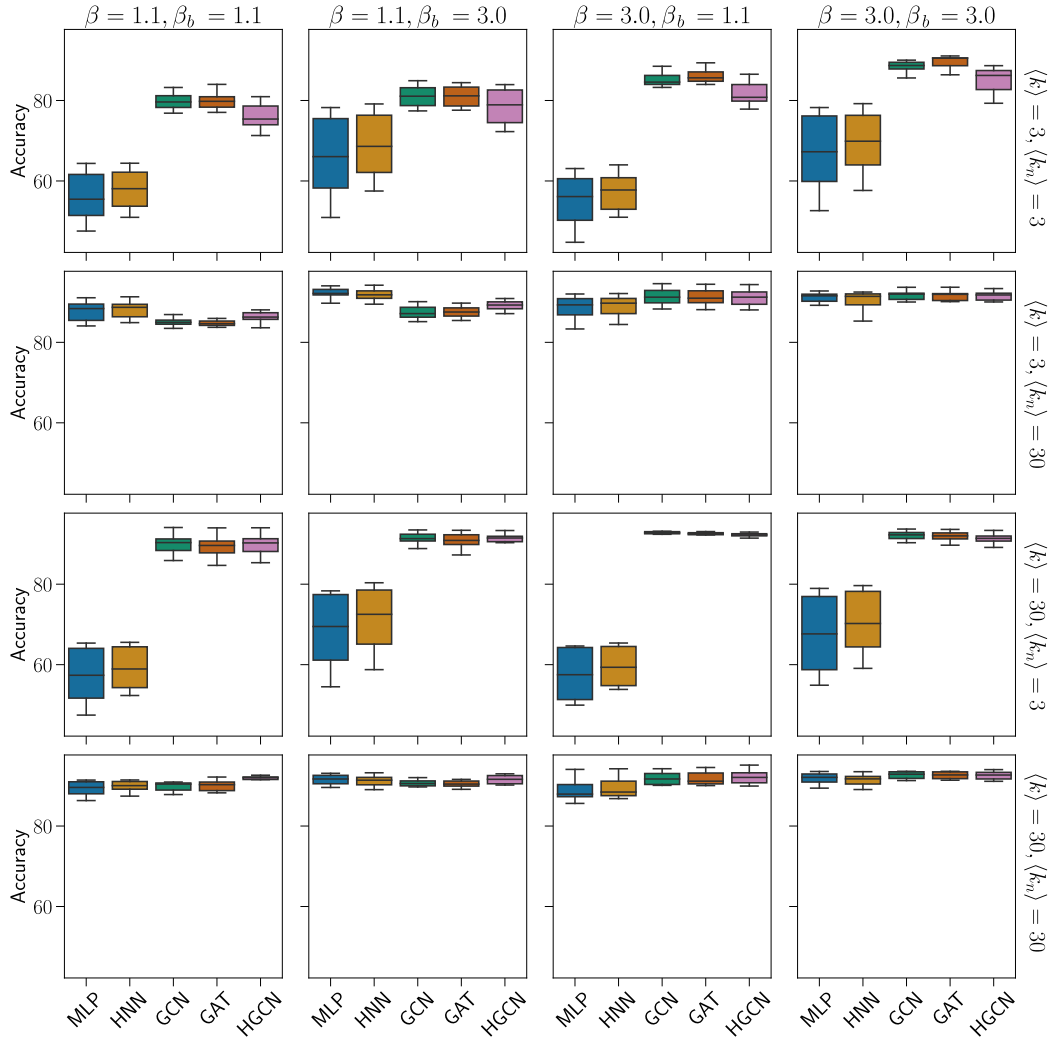


Figure 12: Performance comparison of machine learning models on NC task with respect to a combination of parameters in the $\mathbb{S}^1/\mathbb{H}^2$ and bipartite- $\mathbb{S}^1/\mathbb{H}^2$ models. We set $\mathcal{N}_L = 6$ and $\alpha = 10$.

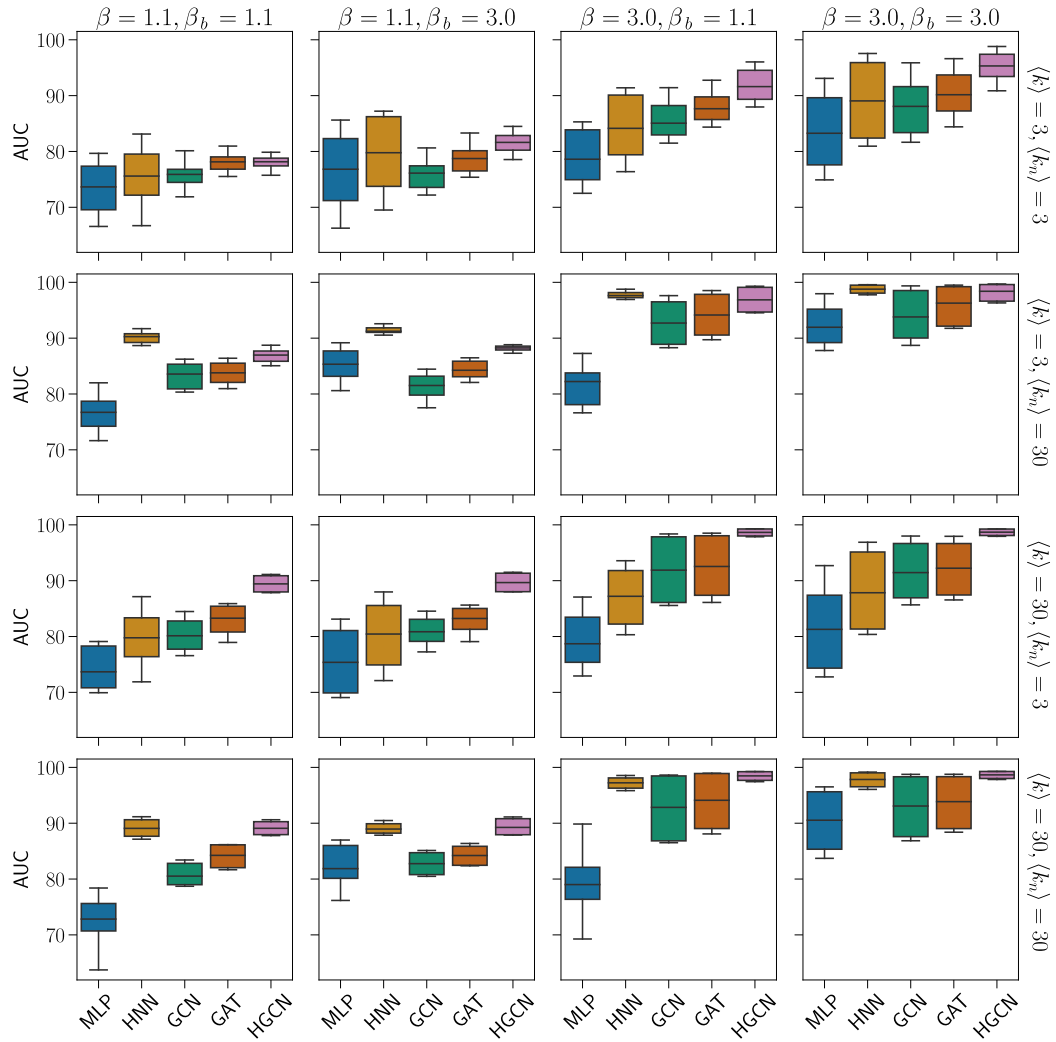


Figure 13: Performance comparison of machine learning models on LP task with respect to a combination of parameters in the $\mathbb{S}^1/\mathbb{H}^2$ and bipartite- $\mathbb{S}^1/\mathbb{H}^2$ models.

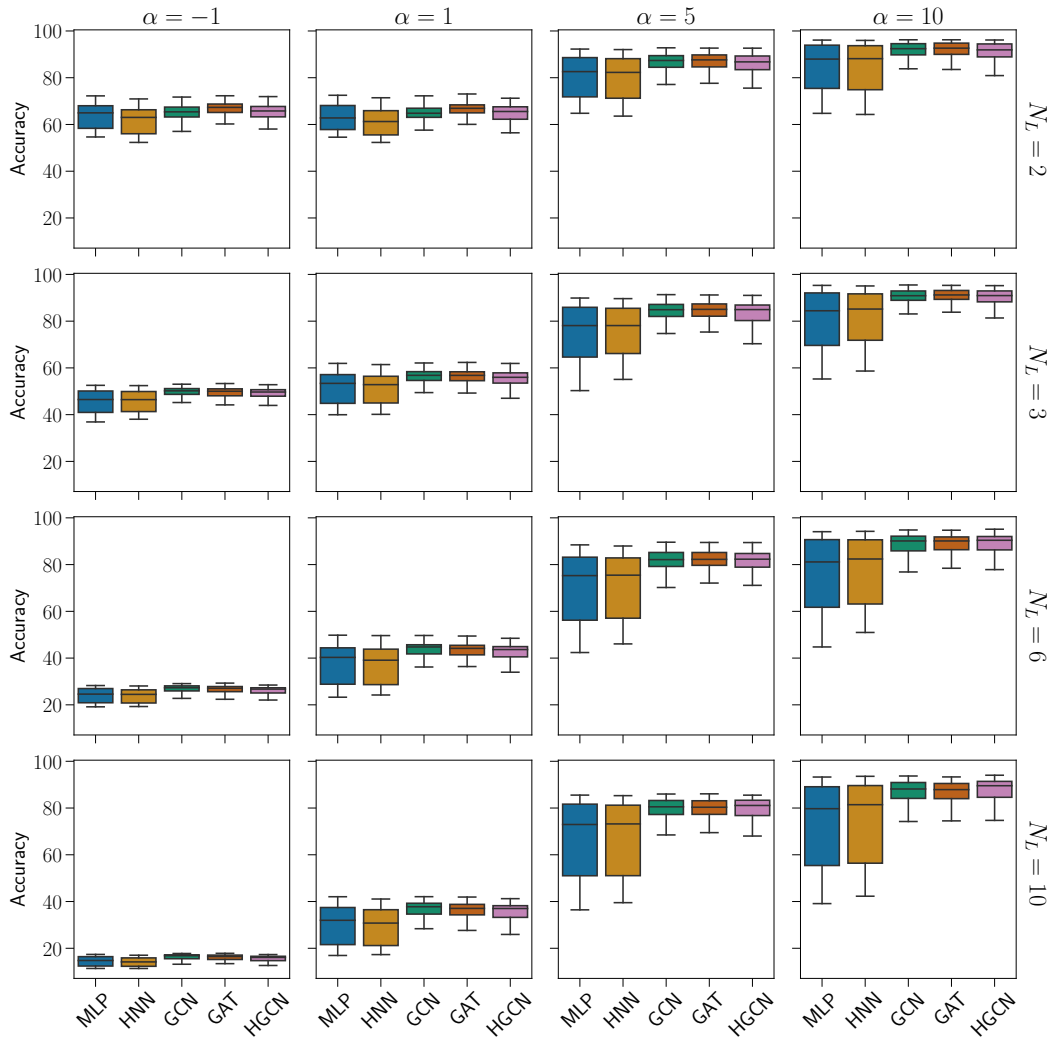


Figure 14: Performance comparison of machine learning models on the NC task with respect to a combination of N_l and α parameters averaged over all other parameters.