
Predicting Microbial Interactions Using Graph Neural Networks

Elham Gholamzadeh

Max Planck Institute for Human Cognitive and Brain Sciences
Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI)
Dresden/Leipzig, Germany
gholamzade@cbs.mpg.de

Kajal Singla

Max Planck Institute for Human Cognitive and Brain Sciences
Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI)
Dresden/Leipzig, Germany
singla@cbs.mpg.de

Nico Scherf

Max Planck Institute for Human Cognitive and Brain Sciences
Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI)
Dresden/Leipzig, Germany
nscherf@cbs.mpg.de

Abstract

Predicting interspecies interactions is a key challenge in microbial ecology, as these interactions are critical to determining the structure and activity of microbial communities. In this work, we used data on monoculture growth capabilities, interactions with other species, and phylogeny to predict a negative or positive effect of interactions. More precisely, we used one of the largest available pairwise interaction datasets to train our models, comprising over 7,500 interactions between 20 species from two taxonomic groups co-cultured under 40 distinct carbon conditions, with a primary focus on the work of Nestor et al.[17]. In this work, we propose Graph Neural Networks (GNNs) as a powerful classifier to predict the direction of the effect. We construct edge-graphs of pairwise microbial interactions in order to leverage shared information across individual co-culture experiments, and use GNNs to predict modes of interaction. Our model can not only predict binary interactions (positive/negative) but also classify more complex interaction types such as mutualism, competition, and parasitism. Our initial results were encouraging, achieving an F1-score of 80.44%. This significantly outperforms comparable methods in the literature, including conventional Extreme Gradient Boosting (XGBoost) models, which reported an F1-score of 72.76%.

1 Introduction

Understanding microbial species' functions and interactions for developing structured microbial communities and processes is experimentally challenging and remains a fundamental problem in biology and biotechnological applications [16, 25]. Gaining a deeper understanding of these interactions can significantly enhance our capacity to manage and manipulate ecological communities, offering broad applications in environmental conservation, crop health, and human health [10, 24]. For several decades, microbial communities have been used for the clean-up of contaminated water, soil, and air and the capture of renewable resources such as energy (biogas and H₂), chemicals (short-chain acids), and water [23, 4]. Studying microbial interactions exhaustively through pairwise co-culture

experiments becomes increasingly impractical with increasing number of species, since the possible combinations rise exponentially. This limits our ability to experimentally search for key positive interactions that drive community function. Designing such microbial systems for bioengineering applications is further complicated by the complexity introduced by spatial organization and temporal dynamics. This suggests that much can be gained from reducing the parameter space of the system, e.g., by supervised machine learning methods. Supervised machine learning algorithms [6, 21], including support vector machines [27], neural networks [20], random forests [3], and graph neural networks [18] to predict more complex phenotypes of microbial interactions [2, 15, 9], a task made more feasible by the rise of high-throughput techniques for generating large biological data sets. Predicting interspecific interactions is critical for both understanding and engineering microbial communities, as the collective behavior of these systems is shaped by intricate relationships between species. These interactions can be negative where one species inhibits another through nutrient competition or the production of antimicrobial compounds [5] or positive, where one species enhances the growth of another by increasing nutrient availability or creating new ecological niches [1]. Accurate prediction of these relationships becomes especially important in species-rich environments or communities composed of metabolically fastidious organisms. Traditionally, such predictions have relied on genome-scale metabolic and constraint-based modeling approaches [22, 14], which use genomic information to estimate interactions by assessing overlaps and complementarities in metabolic capabilities [7, 12].

Building on the work of Nestor et al.[17], we leverage one of the largest data sets of experimentally validated microbial interactions to evaluate how well graph neural networks can predict the sign of microbial interactions (positive/negative) compared to other standard machine learning models. Similar to [17] we use the species' phylogeny and their monoculture yield, which belong to each of the 40 carbon environments, as features to build our models. All pairwise interactions between 20 distinct soil bacteria from two taxonomic groups that were cultivated in 40 different environments [17], each containing either a single carbon source or a combination of all carbon sources, are included in this data set.

We propose using GNNs as a data-driven model of microbial communities and interactions (Figure.1) as GNNs are able to leverage information across separate (but related) experimental conditions. We build a graph that connects co-culture experiments through shared species or experimental conditions, and use this structure to train Graph Neural Networks (GNNs) to predict interactions between species in (experimentally) unknown combinations. In our graph structure, nodes represent a unique combination of species and condition, and edges capture the interactions between two given species under the same experimental condition. To directly encode relational information across experiments, we construct an edge graph, where each interaction (edge) in the original graph is transformed into a new node, now representing a combination of two species in a specific experimental condition. Two nodes in the edge graph are now connected if their corresponding experiments share a common species (and condition).




Figure 1: Illustration of leveraging graph neural networks to predict microbial interactions. This approach enables accurate prediction of both one-way and two-way interactions. a) Experimental measurements of microbial interaction; b) Feature creation; c) Construct graph based on microbe/condition (nodes) and interactions measured by experiment (edges); d) Convert of the graph into a dual edge graph; e) Edge graph representation; f) Prediction using a graph neural network.

2 Graph Neural Networks

2.1 Graph Neural Networks

The interactions among microbes inherently form a graph structure, enabling their representation as a network of connections. We modeled these interactions using a graph neural network (GNN) [26]. Graph Neural Networks (GNNs) extend traditional neural networks to work directly with graph-structured data, where a graph is formally represented as $G = (V, E)$, where V is the set of vertices (nodes) and E is the set of edges. The neighborhood of a node $v \in V$ is defined as $N(v) = \{u \mid \{u, v\} \in E\}$, and each node i is associated with an attribute x_i for $i \in [1, |V|]$. GNNs operate iteratively through message passing steps, where node attributes are updated based on interactions with their neighbors. At each step, the update process is as follows: For each node i , the features of its neighbors $j \in N(i)$ are retrieved. The node features x_i are updated by aggregating information from its neighbors:

$$e(i, j) = g_e(\mathbf{x}_i, \mathbf{x}_j) \quad (1)$$

$$\mathbf{x}'_i = g_v \left(\mathbf{x}_i, \text{aggr}_{j \in N(i)} (e(i, j)) \right) \quad (2)$$

where g_v and g_e are the node and edge update functions that represent an aggregation operation. Here, mean is used as the aggregation function over the features of neighboring nodes.

2.2 Edge-Graph Construction

As illustrated in Figure 1, our framework construct the relationships between different species and their available features using an edge-graph representation. This approach focuses on relationships between experiments (edges) rather than solely between individual species (nodes), allowing us to explicitly encode dependencies between experimentally measured species interactions. The edge-graph $L(G)$ provides a structural representation of edge adjacencies, making it particularly suitable for edge-centric tasks in Graph Neural Networks (GNNs).

We begin with the species interaction graph $G = (V, E)$, where:

- V denotes the set of species recorded under a shared conditions;
- E edges denote a shared species and condition between experiments.

The transformation from G to its edge-graph $L(G) = (V', E')$ is defined formally as follows:

- **Node mapping:** For each interaction $e_k = (u, v) \in E$, create a node $v'_k \in V'$ in $L(G)$. Each v'_k represents a single species interaction.
- **Adjacency mapping:** Two nodes $v'_i, v'_j \in V'$ are connected by an edge in E' if and only if their corresponding edges in G share at least one endpoint species and experimental condition.
- **Feature mapping:** Any attributes originally associated with the interaction e_k (e.g., monogrowth characteristics, phylogenetic tree-based distances, etc.) are assigned to the corresponding node v'_k in $L(G)$.

This construction turns interaction–interaction dependencies into direct connections in $L(G)$. It allows message passing to operate over interactions themselves, capturing higher-order ecological relationships.

2.3 Proposed Model

To classify microbial interactions, we implemented a two-layer GraphSAGE model using the Deep Graph Library (DGL) [8]. Each node in the graph represents a microbial species, and edges capture potential interactions. The task is to assign a binary label $y_i \in \{+, -\}$ to each node v_i .

The two layers of the model apply GraphSAGE convolutions with mean aggregation, allowing each node to iteratively incorporate feature information from its local neighborhood. In GraphSAGE, the edge function $e(i, j)$ outputs the attributes of a neighboring node $j \in \mathcal{N}(i)$:

$$g_e(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_j, \quad (3)$$

and the node update function g_v is defined as

$$\mathbf{x}'_i = W_1 \mathbf{x}_i + W_2 \cdot \text{mean}_{j \in \mathcal{N}(i)} \mathbf{x}_j, \quad (4)$$

where W_1 and W_2 are learnable weight matrices.

The ReLU activation function is applied after first layer to introduce non linearity. No explicit edge features are used in the model learning is driven entirely by node attributes and graph connectivity.

The input feature size is 13 (see Table A.1), and ReLU is used as the non-linear activation function:

$$\text{ReLU}(x) = \max(0, x) \quad (5)$$

We optimize the model using cross-entropy loss:

$$\mathcal{L}(\hat{y}, y) = - \sum_{i=1}^N y_i \log(\hat{y}_i), \quad (6)$$

where $\hat{y}$ is the predicted label and y is the ground-truth.

This configuration allows for inductive learning and generalization to unseen nodes or graphs, making it well-suited for microbial community analysis.

3 MATERIALS

3.1 Dataset

The dataset comprises over 7,500 pairwise interactions among 20 species from two taxonomic groups across 40 carbon environments. It is primarily based on the works of Nestor *et al.* [17] and Kehe *et al.* [11], which employ the kChip platform a high-throughput, nanodroplet-based system for combinatorial screening to evaluate species growth in monoculture and coculture settings.

3.1.1 Features creation

By primarily focusing on the [17], to enhance the predictive power of our model, monoculture growth yields for each species and supplementary features, such as the species' phylogenetic relationships and their metabolic profiles across the carbon environments, are incorporated. The phylogenetic data was summarized using principal component analysis (PCA), where the first two principal components captured 95% of the variance. Similarly, the metabolic profiles were represented using the first four principal components, which accounted for 90% of the variation. These reduced-feature phylogenetic data and metabolic profiles were used to construct a feature set that effectively captures both the evolutionary and metabolic similarities between species. Additionally, the Euclidean distance between the monoculture growth profiles of interacting species was computed to further describe their metabolic relationship. This feature set, which is presented in (Table.A.1 in A.1 Appendix), allowed us to model and predict the sign effect of interactions, enabling better understanding of species relationships in diverse carbon environments.

3.2 Model implementation

We evaluated our method using standard classification metrics, including accuracy (Acc), sensitivity (Sen or Recall), precision (Prec), and F1 score (equations are provided in Appendix A.2) We implemented the proposed solution in Python (version 3) using the GraphSAGE model from the PyTorch framework [19]. During training, we employed the Adam optimizer [13] to perform first-order gradient-based optimization on the stochastic objective function, with a learning rate set to 1×10^{-2} . The dataset was split into 80 % for training and 20 % for testing. The performance of the GraphSAGE model was evaluated over 300 training epochs. All experiments were conducted on the RAVEN HPC system, equipped with Intel Xeon IceLake-SP processors and NVIDIA A100 GPU nodes interconnected via NVLink.

4 Results

We compare our method with the standard machine learning algorithms used in [17] for predicting the sign of one- and two-way growth yield effects. Specifically, we evaluated the predictive performance of our approach against two baselines: k-Nearest Neighbors (kNN) and Extreme Gradient Boosting (XGBoost). The Graph Neural Network achieved significantly higher accuracy in predicting the signs of both one- and two-way effects. The results for one-way prediction of different models are reported in Table.1. GraphSAGE learns these node embeddings in an end-to-end manner, optimizing them specifically for the classification objective and explicitly models the connection across experiments. Figure.2 shows the confusion matrix of the GraphSAGE model and XGBoost for one-sign prediction. Our results show a class imbalance as positive interactions appear to be rare in both models for the one-way interaction sign prediction. Thus, the F1-score (harmonic mean of precision and recall) is a more informative metrics for an imbalanced class distribution. Our edge-graph-based models demonstrated strong performance, achieving an F1-score of 80.44%, which outperformed XGBoost (72.76%) and k-NN (66.13%).

Table 1: Model performance metrics of one-way prediction on test dataset.

Model	Accuracy	Sensitivity	Precision	F1-Score
GraphSAGE	0.9528	0.7392	0.8822	0.8044
K-nearest neighbors [17]	0.8593	0.6613	0.6613	0.6613
XGBoost [17]	0.8835	0.8017	0.7276	0.7276




Figure 2: Comparison of the confusion matrices for XGBoost and GraphSAGE models of one-way prediction.

We further implemented the GraphSAGE framework to predict two-way interactions between species: mutualism (+,+), parasitism (+,-), and competition (-,-). This approach leverages the structural properties of ecological networks to model complex interspecies relationships. By integrating both direct and indirect effects of species on one another, the GNN facilitates accurate predictions of interaction types.

In our study, we observed that mutualistic interactions (+,+) were particularly challenging to predict accurately, often being misclassified as parasitic interactions (+,-). This challenge in predicting positive one-way effects is consistent with ecological research, which suggests that mutualistic interactions are complex and difficult to model because of their intricate dynamics. Table.2 summarizes the results for different two-way interactions. Figure.3 shows the comparison of confusion matrix of node classification of two-way interaction classes. The proposed edge-graph-based models demonstrated good performance in two-way sign prediction, achieving an F1-score of of 43.1% and 68.4% slightly outperforming XGBoost in predicting mutualism and parasitism, while XGBoost is slightly better at predicting competition. Among the two-way sign predictions, competition (-,-) is the most frequently identified interaction type (57.39%). In contrast, mutualistic interactions (+,+) are particularly challenging to predict. This difficulty likely arises from the inherent challenge of detecting positive one-way effects, which may cause the model to misclassify them as parasitism

(+,-) (20.21%) or, less frequently, as mutualism itself (3.48%). These misclassifications may also reflect the imbalance in class distributions, as mutualistic cases are relatively underrepresented in the dataset.

Table 2: Classification performance for microbial two-way interactions using GraphSAGE and XGBoost.

Class	GraphSAGE			XGBoost[17]		
	Precision	Recall	F1-score	Precision	Recall	F1-score
Mutualism	0.611	0.333	0.431	0.430	0.440	0.430
Competition	0.855	0.933	0.893	0.880	0.920	0.900
Parasitism	0.707	0.663	0.684	0.640	0.570	0.600
Accuracy	0.811			0.780		




Figure 3: Comparison of confusion matrices for XGBoost [17] and GraphSAGE in predicting two-way interaction signs.

5 Limitations

Microbial interactions offer important insights into community dynamics but are challenging to measure. While our results show that microbial interactions can be predicted under controlled laboratory conditions, it is unclear how well this generalizes to natural environments. The model focuses on pairwise interactions and does not capture higher-order effects, which may be influenced by additional species in complex communities. Performance also depends on the availability and quality of input features, such as monoculture growth yield. Moreover, graph neural networks are computationally intensive and require access to HPC resources, potentially limiting reproducibility. We might also need to consider whether there are better ways to structure the graphs to encode the experimental information more effectively. These limitations suggest directions for future work, including testing on broader datasets, incorporating higher-order interactions through approaches such as Topological Deep Learning (TDL), improving computational efficiency, and developing lightweight methods for resource-limited settings.

6 Conclusion

In this work, we illustrate that an edge-focused graph learning approach can reveal both the structure and direction of interactions in complex biological systems. By converting microbial networks into edge-graphs and applying inductive GNNs like GraphSAGE, we capture the properties of interactions directly, rather than only the traits of species. We tested this approach with data of bacterial communities, where it outperforms XGBoost as a strong baseline. The method is generalizable and could be applied to other ecological systems, epidemiological networks, or any network with diverse,

directional links. In the long run, by embedding biologically meaningful features at each edge, geometric deep learning models can achieve high predictive accuracy while reflecting real-world interaction patterns, bridging the gap between data-driven prediction and biological understanding.

References

- [1] John F Bruno, John J Stachowicz, and Mark D Bertness. Inclusion of facilitation into ecological theory. *Trends in ecology & evolution*, 18(3):119–125, 2003.
- [2] Alfredo Cesario, Marika D’Oria, Francesco Bove, Giuseppe Privitera, Ivo Boškoski, Daniela Pedicino, Luca Boldrini, Carmen Erra, Claudia Loreti, Giovanna Liuzzo, et al. Personalized clinical phenotyping through systems medicine and artificial intelligence. *Journal of personalized medicine*, 11(4):265, 2021.
- [3] Alison J Dicker, Mike Lonergan, Holly R Keir, Alexandria H Smith, Jennifer Pollock, Simon Finch, Andrew J Cassidy, Jeffrey TJ Huang, and James D Chalmers. The sputum microbiome and clinical outcomes in patients with bronchiectasis: a prospective observational study. *The Lancet Respiratory Medicine*, 9(8):885–896, 2021.
- [4] Alexander Eng and Elhanan Borenstein. Microbial community design: methods, applications, and opportunities. *Current opinion in biotechnology*, 58:117–128, 2019.
- [5] Melanie Ghoul and Sara Mitri. The ecology and evolution of microbial competition. *Trends in microbiology*, 24(10):833–845, 2016.
- [6] Stephen J Goodswen, Joel LN Barratt, Paul J Kennedy, Alexa Kaufer, Larissa Calarco, and John T Ellis. Machine learning and applications in microbiology. *FEMS microbiology reviews*, 45(5):fuab015, 2021.
- [7] Changdai Gu, Gi Bae Kim, Won Jun Kim, Hyun Uk Kim, and Sang Yup Lee. Current status and applications of genome-scale metabolic models. *Genome biology*, 20:1–18, 2019.
- [8] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- [9] Rebecca E Hughes, Richard JR Elliott, John C Dawson, and Neil O Carragher. High-content phenotypic and pathway profiling to advance drug discovery in diseases of unmet need. *Cell chemical biology*, 28(3):338–355, 2021.
- [10] M Kavino and SK Manoranjitham. In vitro bacterization of banana (musa spp.) with native endophytic and rhizospheric bacterial isolates: novel ways to combat fusarium wilt. *European journal of plant pathology*, 151:371–387, 2018.
- [11] Jared Kehe, Anthony Kulesa, Anthony Ortiz, Cheri M Ackerman, Sri Gowtham Thakku, Daniel Sellers, Seppe Kuehn, Jeff Gore, Jonathan Friedman, and Paul C Blainey. Massively parallel screening of synthetic microbial communities. *Proceedings of the National Academy of Sciences*, 116(26):12804–12809, 2019.
- [12] Tae Yong Kim, Seung Bum Sohn, Yu Bin Kim, Won Jun Kim, and Sang Yup Lee. Recent advances in reconstruction and applications of genome-scale metabolic models. *Current opinion in biotechnology*, 23(4):617–623, 2012.
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [14] Elena G Lara, Isabelle van der Windt, Douwe Molenaar, Marjon GJ de Vos, and Chrats Melkonian. Using functional annotations to study pairwise interactions in urinary tract infection communities. *Genes*, 12(8):1221, 2021.
- [15] Laura Judith Marcos-Zambrano, Kanita Karaduzovic-Hadziabdic, Tatjana Loncar Turukalo, Piotr Przymus, Vladimir Trajkovic, Oliver Aasmets, Magali Berland, Aleksandra Gruca, Jasminka Hasic, Karel Hron, et al. Applications of machine learning in human microbiome studies: a review on feature selection, biomarker identification, disease prediction and treatment. *Frontiers in microbiology*, 12:634511, 2021.
- [16] Carey D Nadell, Knut Drescher, and Kevin R Foster. Spatial structure, cooperation and competition in biofilms. *Nature Reviews Microbiology*, 14(9):589–600, 2016.
- [17] Einat Nestor, Gal Toledano, and Jonathan Friedman. Interactions between culturable bacteria are predicted by individual species’ growth. *Msystems*, 8(2):e00836–22, 2023.

- [18] Jie Pan, Zhen Zhang, Ying Li, Jiaoyang Yu, Zhuhong You, Chenyu Li, Shixu Wang, Minghui Zhu, Fengzhi Ren, Xuexia Zhang, et al. A microbial knowledge graph-based deep learning model for predicting candidate microbes for target hosts. *Briefings in Bioinformatics*, 25(3):bbae119, 2024.
- [19] A Paszke. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.
- [20] Simone Rampelli, Marco Fabbri, Marco Candela, Elena Biagi, Patrizia Brigidi, and Silvia Turrioni. G2s: a new deep learning tool for predicting stool microbiome structure from oral microbiome data. *Frontiers in Genetics*, 12:644516, 2021.
- [21] Parminder S Reel, Smarti Reel, Ewan Pearson, Emanuele Trucco, and Emily Jefferson. Using machine learning approaches for multi-omics data analysis: A review. *Biotechnology advances*, 49:107739, 2021.
- [22] Saeed Shoaie, Fredrik Karlsson, Adil Mardinoglu, Intawat Nookaew, Sergio Bordel, and Jens Nielsen. Understanding the interactions between bacteria in the human gut through metabolic modeling. *Scientific reports*, 3(1):2532, 2013.
- [23] Ryan Tsoi, Zhuojun Dai, and Lingchong You. Emerging strategies for engineering microbial communities. *Biotechnology advances*, 37(6):107372, 2019.
- [24] Bret D Wallace and Matthew R Redinbo. The human microbiome is a source of therapeutic drug targets. *Current opinion in chemical biology*, 17(3):379–384, 2013.
- [25] Lothar Wondraczek, Georg Pohnert, Felix H Schacher, Angela Köhler, Michael Gottschaldt, Ulrich S Schubert, Kirsten Küsel, and Axel A Brakhage. Artificial microbial arenas: materials for observing and manipulating microbial consortia. *Advanced Materials*, 31(24):1900284, 2019.
- [26] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.
- [27] Longxiu Yang, Yuan Qin, and Chongdong Jian. Screening for core genes related to pathogenesis of alzheimer’s disease. *Frontiers in Cell and Developmental Biology*, 9:668738, 2021.

A Technical Appendices and Supplementary Material

A.1 Features names information

Table A.1 summarizes the features used for microbial interaction prediction, including monoculture growth measurements, metabolic dissimilarity, and principal components derived from carbon source utilization and phylogenetic embeddings.

Table A.1: Features used for microbial interaction prediction.

Feature	Description
Monoculture growth features	
monoGrow_x	Monoculture growth of species x
monoGrow_y	Monoculture growth of species y
monoGrow24_x	24-hour monoculture growth of species x
monoGrow24_y	24-hour monoculture growth of species y
metDis	Metabolic dissimilarity between species x and y
Carbon source PCA components (90% variance)	
carbon_component_0	Principal component 0 of metabolic profiles
carbon_component_1	Principal component 1 of metabolic profiles
carbon_component_2	Principal component 2 of metabolic profiles
carbon_component_3	Principal component 3 of metabolic profiles
Phylogenetic PCA components (95% variance)	
phy_strain_component_0_x	Principal component 0 of phylogenetic embedding (species x)
phy_strain_component_1_x	Principal component 1 of phylogenetic embedding (species x)
phy_strain_component_0_y	Principal component 0 of phylogenetic embedding (species y)
phy_strain_component_1_y	Principal component 1 of phylogenetic embedding (species y)

A.2 Assessment metrics

The evaluation metrics used in this study are defined as follows:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (7)$$

$$\text{Sensitivity} = \frac{TP}{TP + FN} \quad (8)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (9)$$

$$\text{F1} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (10)$$

Where True Positive (TP) denotes the positive instances correctly identified as positive, True Negative (TN) indicates the negative instances correctly identified as negative, False Positive (FP) represents the negative instances incorrectly identified as positive, and False Negative (FN) refers to the positive instances incorrectly identified as negative.