Insights from Network Science can advance Deep Graph Learning

Christopher Blöcker [1] **Martin Rosvall** [2] **Ingo Scholtes** [1] **Jevin D. West** [3]

Abstract

Deep graph learning and network science both analyze graphs but approach similar problems from different perspectives. Whereas network science focuses on models and measures that reveal the organizational principles of complex systems with explicit assumptions, deep graph learning focuses on flexible and generalizable models that learn patterns in graph data in an automated fashion. Despite these differences, both fields share the same goal: to better model and understand patterns in graph-structured data. Early efforts to integrate methods, models, and measures from network science and deep graph learning indicate significant untapped potential. In this position, we explore opportunities at their intersection. We discuss open challenges in deep graph learning, including data augmentation, improved evaluation practices, higher-order models, and pooling methods. Likewise, we highlight challenges in network science, including scaling to massive graphs, integrating continuous gradient-based optimization, and developing standardized benchmarks.

1. Introduction

In 1982, John Hopfield introduced a neural network model that sparked a flurry of innovations: content-addressable memory, energy dynamics, error correction, and nonlinear architecture (Hopfield, 1982). The Nobel Prize committee recently recognized the role these innovations played in the development of modern machine learning[1]. Less recognized but not any less important is the influence the paper had on the field of network science. The Hopfield network demonstrated the critical connection between network topology and the collective behaviour of complex systems—one of the enduring themes of network science and now one of the central challenges in deep graph learning.

Surprisingly, the two fields have diverged more than they have converged since Hopfield's influential paper. We see an opportunity for that to change, and argue for better integration of the two research communities. At their core, both fields model and analyze patterns in graphs. However, their needs are different. In deep graph learning, there is a need for methods that augment data to cope with limited training data, pool node representations to facilitate graph-level learning, and develop message passing schemes that incorporate higher-order interactions beyond pair-wise edges. Network science has been thinking about these issues for years, albeit often from a different perspective or with different motivation. Conversely, in network science, there is a need to scale models, measures and algorithms to massive graphs, better incorporate continuous, gradient-based methods for optimization problems, and develop standardized benchmarks to fairly compare, such as community detection algorithms. Deep learning has mastered these techniques.

We see opportunities in bridging this gap between network science and deep graph learning. In this paper, we detail these opportunities with the hope of spurring conversations across the two communities.

There is significant potential to address open challenges in deep graph learning, such as data augmentation, evaluation practices, integrating higher-order models, and modeling patterns in temporal graphs. Network science offers valuable insights by *connecting network structure with function through principled methodologies*, including probabilistic generative models that provide principled null models for complex networks, statistical inference and network reconstruction methods for noisy relation data, and community detection techniques. These approaches can *enhance the theoretical foundation and empirical insight of deep graph learning models*.

At the same time, scaling models, measures, and algorithms to massive graphs and incorporating continuous optimization remain open challenges in network science. Deep learning techniques, such as gradient-based optimization and large-scale training, offer powerful tools to advance net-

[1]Chair of Machine Learning for Complex Networks, Center for Artificial Intelligence and Data Science (CAIDAS), Julius-Maximilians-Universiät Würzburg, Germany [2]Integrated Science Lab (IceLab), Department of Physics, Umeå Universitet, Umeå, Sweden [3]Center for an Informed Public, Information School, University of Washington, Seattle, USA. Correspondence to: Christopher Blöcker <christopher.bloecker@uni-wuerzburg.de>.

[1]https://www.nobelprize.org/prizes/physics/2024/hopfield/facts/

work science.

Succeeding with these synergies requires fostering a more structured collaboration between the two fields. By bridging their complementary strengths, we aim to initiate a broader discussion on developing more interpretable, scalable, and generalizable approaches for modeling graph-structured data.

2. Principled Deep Graph Learning Modeling

Deep graph learning relies on the task-specific training of deep neural networks that allow to model patterns in graph-structured data in an end-to-end fashion. The past decade has seen rapid advances in the development of deep graph learning architectures for various tasks and applications. However, challenges to apply state-of-the-art graph neural networks to real-world problems also expose limitations that we need to address (Georgousis et al., 2021; Ju et al., 2024). These include the need for a more systematic framework to characterize which structural properties of a (temporal) graph data set contribute to the performance of a specific model, the integration of evaluation methods that are rooted in principled null models, theoretically grounded pooling methods that minimize information loss, and data augmentation techniques that model noisy or incomplete data to advance model generalizability. Insights from network science can help us to address these challenges, guiding deep graph learning toward principled architectures, improved interpretability, and more rigorous evaluation methods.

2.1. Probabilistic Generative Models

Random graph models are an important foundation of network science. Starting from the simple Erdős-Rényi model for random graphs, where edges between pairs of nodes are randomly generated with equal probability (Erdos et al., 1960), network science has developed probabilistic generative models that define a "statistical ensemble" containing all graphs that share given aggregate characteristics such as size, density, degree sequence or distribution, modular structure or motif statistics. Important examples include the Molloy-Reed model that generates random graphs with a given degree sequence or distribution (Molloy & Reed, 1995), exponential random graph models for random graphs with a given set of network statistics (Robins et al., 2007), or the stochastic block model for random graphs with given homophilic or heterophilic community patterns (Lee & Wilkinson, 2019). In network science, such statistical ensembles are the foundation to analytically study expected properties of graphs with given aggregate characteristics, for example using generating functions as a framework (Newman et al., 2001; Newman, 2009) Moreover, these models can be used to randomize the topology of empirical networks, while maintaining aggregate properties. This randomization

serves as a null model for statistical hypothesis testing, enabling us to understand which of a network's characteristics are actually due to the *topology* of the network – that is, which node is linked to which other nodes – and which characteristics can be explained based on the mere *degree distribution* or *density* of edges.

In network science, this approach has been used to understand the role of network topology in the diffusion of information or disease spreading. Unfortunately, principled null models for graph-structured data are not yet widely used in the evaluation of deep graph learning architectures, where the focus often lies on the performance of a specific model in a given task rather than on which topological features of a graph can explain the predictive power of specific architectures. Leveraging network science models could thus lead to more rigorous and meaningful evaluation practices that provide insights into the predictive capabilities of deep graph learning models.

2.2. Data Augmentation for Graph Neural Networks

In machine learning, we often use *data augmentation techniques* in the training phase to improve the generalizability of machine learning models and to mitigate overfitting. This typically involves enriching available training data through applying perturbations, injecting noise, or other means to augment available training data by artificial examples. Inspired by these methods, the deep graph learning community has considered various *graph augmentation* techniques that seek to manipulate edges or nodes of a graph in such a way that it improves the performance and/or generalizability of graph neural networks (Zhao et al., 2023). Recent works have considered, for example, the targeted removal of edges to increase homophilic patterns (Zhao et al., 2020) or selectively adding nodes that slow down message passing (Azabou et al., 2023). While these works on graph augmentation have made progress towards improving the generalizability of Graph Neural Networks (GNNs), researchers from the deep graph learning community have recently argued that we still lack a "theory of data augmentation" for graphs (Morris et al., 2024), highlighting that, despite theoretical results on their expressivity, GNNs have not yet matured to leverage the full potential of the data in practical settings but currently rely on preprocessing for data augmentation.

Network science provides tools that can help us develop such a theory. In particular, network scientists have long argued that it is often not desirable to directly use an observed network "as is" for network analysis or neural message passing in GNNs. Many empirical data on complex networks are unreliable insofar as they suffer from spurious or missing relationships, are incomplete in terms of observed nodes, or contain incorrect node labels or noisy attributes. For such data, network science has developed reconstruc-
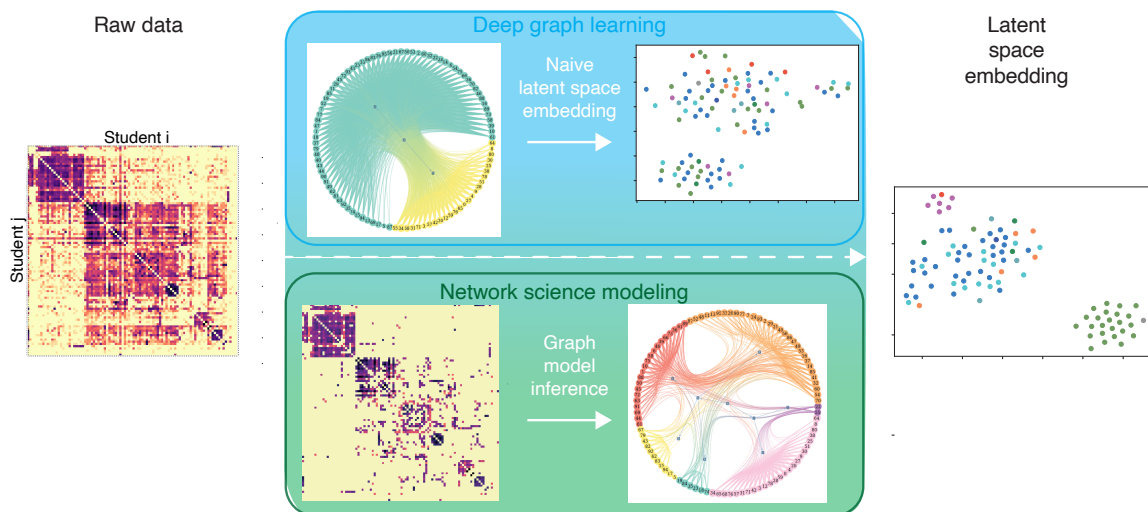
Figure 1. **Illustrative example of how network science insights can advance deep graph learning**. Deep graph learning provides tools for end-to-end representations learning for prediction tasks (top panel), while network science provides advanced statistical modeling techniques for handling noisy graph data (bottom panel). Applying deep graph learning directly to noisy co-occurrence data (left) results in a latent space representation that poorly reflects ground truth node labels (top panel). Network science modeling techniques—such as statistical ensembles of random graphs that preserve aggregate characteristics of the data—help build robust graph models that account for noise (bottom panel). Combining these techniques with deep learning methods leads to a latent space representation that better captures ground truth patterns (bottom right). Figure partly adapted from Casiraghi et al. (2017).

tion and inference techniques (Peixoto, 2019b; Newman, 2018) that use *statistical graph ensembles* to infer reasonable graph models from noisy or incomplete data. In deep learning, such methods could be used to generate a set of plausible graph models that can be used for training graph neural networks, thereby improving robustness and generalisability. An example for this is shown in fig. 1. Similarly, recent works demonstrate that insights from the modeling of dynamical processes on graphs as well as spectral graph theory provide new ways to augment graphs in a theoretically principled way, for example by tuning the spectral gap to simultaneously mitigate over-squashing and over-smoothing (Jamadandi et al., 2024; Rubio-Madrigal et al., 2025).

2.3. Community Detection and Pooling

Community detection lies at the heart of unsupervised graph learning, offering a powerful lens to uncover meaningful structures in complex systems. Identifying groups of nodes – communities – that are more densely connected internally than externally enables exploratory data analysis and unsupervised classification in diverse domains, from understanding social interactions to unraveling biological networks and optimizing information retrieval. Community detection has not only benefited from the capabilities of GNNs but has also inspired new applications, such as in graph pooling. Graph pooling techniques simplify graphs by merging nodes into clusters, creating more compact representations that enhance standard applications, including classification and

prediction. By incorporating community detection methods as a basis for merging, these techniques can replace or complement purely data-driven pooling operations, providing more interpretable reductions (Deng et al., 2024).

Neural embedding methods, such as node2vec, GraphSAGE, and autoencoder-based embeddings, have demonstrated impressive capabilities in capturing community structure (Grover & Leskovec, 2016; Hamilton et al., 2017; Kipf & Welling, 2016). These approaches often combine multiple components to construct and divide embeddings into accurate communities. For example, node2vec's random-walk approach can approximate certain spectral properties of the normalized Laplacian under controlled conditions, enabling it to approach the theoretical limits of community detection for certain graphs (Kojaku et al., 2024). However, such methods present challenges: they rely on intricate combinations of embedding generation and clustering steps, making them opaque and difficult to optimize effectively.

In contrast, network science provides a well-established and interpretable framework for community detection. Over the years, methods such as the stochastic block model (SBM) (Peixoto, 2019a) and the map equation (Rosvall & Bergstrom, 2008) have been developed and refined, each grounded in clear mathematical objectives. While the SBM identifies latent groups by modeling the plain topology of a network, the map equation highlights modular regularities in network flows, providing a complementary perspective by focusing on how information, resources, or behaviors prop-

agate through the system. Like node2vec, this flow-based approach is especially valuable in applications where understanding dynamic processes is critical, such as transportation, communication, and biological networks. In contrast to opaque embedding pipelines, both Bayesian inference of the SBM and the map equation directly quantify how well the underlying process generating a network and flows on the network are compressed in different partitions, making the resulting community assignments highly interpretable.

The rigour of network science extends beyond methodology to evaluation. Generative models create benchmark networks with known ground truth and allow theoretical frameworks to systematically compare community detection techniques, offering a solid foundation against which machine learning approaches can be measured (Lancichinetti et al., 2008; Blöcker & Scholtes, 2024). Incorporating these insights into machine learning evaluation workflows would strengthen the theoretical and practical reliability of community detection and pooling in GNNs.

The convergence of network science and machine learning offers significant opportunities to enhance community detection. By leveraging the scalability and pattern-recognition capabilities of deep learning alongside the theoretical depth and interpretability of methods such as the SBM and the map equation, we can create tools that address both the static and dynamic complexities of real-world networks. Such synergies will advance graph representation learning and deepen our understanding of complex systems.

2.4. Modeling Patterns in Temporal Graphs

Due to the growing availability of time series data, the application of graph neural networks to *temporal graph data*, where the edges and/or nodes are changing over time has recently seen a surge of interest. Many recently proposed architectures are snapshot- or event-based adaptations of existing deep learning methods that model (sequential) patterns in time-evolving batches of edges, which capture the time-varying topology of a temporal graph (Longa et al., 2023). Examples of this approach include the event-based Temporal Graph Network model (Rossi et al., 2020) or the snapshot-based EvolveGCN architecture (Pareja et al., 2020). While the performance of these models has been evaluated for several tasks in empirical temporal graphs provided by, for example the Temporal Graph Learning Benchmark (TGB) (Huang et al., 2023), it is often unclear which patterns in temporal graphs they are actually able to learn. In network science, the analysis and modeling of patterns in temporal networks have been a topic of major interest for almost two decades (Holme, 2015). Network scientists have developed measures and models that capture different temporal, topological, and temporal-topological patterns found in real-world temporal graphs, as well as

generative models that selectively reproduce specific patterns. Examples include measures and models that capture bursty activation patterns of nodes or edges (Moinet et al., 2015; Takaguchi et al., 2013), models that capture the temporal evolution of community patterns (Peixoto & Rosvall, 2017), or approaches that capture sequential patterns in the causal ordering of time-stamped edges that influence *time-respecting paths* (Scholtes et al., 2014; Rosvall et al., 2014). Selectively applying such models to empirical temporal graph data allows us to "disentangle" the different temporal, topological, and temporal-topological patterns present in the data. Much work remains to be done in applying such models to recently proposed architectures for temporal graph neural networks in order to better understand which of those patterns are actually captured by a specific deep temporal graph neural network architecture.

2.5. Higher-Order Models

An important insight obtained by the network science community over the past few years is that it is not enough to model dyadic interactions in complex systems (Lambiotte et al., 2019). Dyadic links can capture essential aspects of many real systems but they cannot directly represent higher-order interactions such as multi-body interactions or stateful aspects of the data. For example, sequence data, such as click streams, require modeling dependencies of paths consisting of several links. Social and bio-molecular systems often involve groups of three or more entities in interactions. Public transport systems provide more than one mode of transportation; social media users are active on several social media platforms and communicate with the same friends via different platforms. Consequently, network science has developed different higher-order modeling frameworks that help us gain insights into the structure and dynamics of such complex systems (Battiston et al., 2020; Torres et al., 2021). Crucially, the choice of higher-order model determines what aspects of a complex system can be faithfully represented and discovered in downstream analyses.

Non-Markovian dynamics, such as click streams or the spread of information in social networks, can be modelled with so-called sparse memory networks (Rosvall et al., 2014) or De Bruijn graphs (De Bruijn, 1946). Essentially, they perform a state-space expansion, introducing higher-order nodes where a node of n-th order represents first-order paths of length $n-1$, thus modeling interactions with memory. Interactions involving an arbitrary number of entities, that is, zero or more entities, can be modelled with hypergraphs where edges encompass a set of nodes (Battiston et al., 2020). Simplicial complexes, similar to hypergraphs, capture scenarios where edges contain sets of nodes but in addition, an edge with n nodes implies the existence of every edge involving any subset of those n nodes (Torres et al., 2021). Complex systems involving different modes

of interaction can be represented with multilayer networks where each layer represents a specific mode and the same node can exist across different layers (De Domenico et al., 2013; Kivelä et al., 2014). However, while they may be more convenient to work with, multilayer networks can be converted to memory networks and, thus, have the same modeling capabilities as memory networks.

Recent integrations of higher-order network models with graph neural networks highlight the potential and benefits of encoding domain-specific knowledge via these modeling techniques. For example, several works have integrated hypergraphs with neural networks for applications including recommender systems, bioinformatics, time-series analysis, and computer vision (Kim et al., 2024; Antelmi et al., 2023). An adaptation of De Bruijn graphs for De Bruijn graph neural networks (DBGNN) captures causal relationships in temporal graphs (Qarkaxhija et al., 2022; Heeg & Scholtes, 2024). Similarly, there is a growing interest in graph neural network architectures that leverage higher-order modeling frameworks like simplicial complexes that have been intensely studied in network science (Frantzen & Schaub, 2024). We expect to see an increasing number of future works leveraging the modeling capabilities developed by network science for learning better representations.

3. Rethinking Deep Graph Learning Architectures

Graph neural networks are at the heart of deep graph learning, relying heavily on message passing and training via backpropagation with gradient descent. Despite their success, recent advances have been incremental rather than groundbreaking, and the flexibility of GNNs makes many aspects of their inner workings hard to interpret. Integrating insights from network science offers a path to rethink deep graph learning architectures, for example by integrating new objective functions, taking a dynamical systems view on message passing, explaining when and why message passing works for certain tasks, or helping to make graph neural networks more interpretable and generalizable.

3.1. From Discrete to Continuous

With roots in graph theory, network science models and algorithms typically build on discrete mathematics and involve discrete objective functions, discrete data models, and combinatorial optimization algorithms. In contrast, due to the use of backpropagation, deep learning is largely based on continuous and differentiable loss functions with data models building on matrix representations, which allows effectively utilizing gradient-based optimisation techniques on GPUs. Consequently, integrating existing discrete approaches from network science with the gradient-based optimization framework requires making their objective func-

tions continuous and differentiable.

For example, community detection traditionally considers hard partitions where nodes belong to exactly one community. Detecting communities often utilizes discrete and stochastic search algorithms that move one node per iteration to optimize their objective function (Blondel et al., 2008; Edler et al., 2017; Traag et al., 2019). Several recent deep graph clustering works have adapted community detection approaches and made them differentiable by considering nodes' community memberships as continuous rather than discrete: Tsitsulin et al. (2024) adapted the modularity criterion, Shchur & Günnemann (2019) used the Poisson random process model, Bianchi et al. (2020) built on minimum cuts, and Blöcker et al. (2024) adapted the map equation. With soft community assignments, these approaches naturally produce overlapping communities which often better capture the characteristics of real systems where nodes can belong to several groups, such as in social networks.

Continuous adaptations of network science objectives present an immense opportunity for deep graph learning to integrate established methods into their architectures. This way, principled unsupervised loss formulations can become part of a composite loss to guide learning in supervised tasks as done, for example, in recent graph pooling works (von Pichowski et al., 2024; Castellana & Bianchi, 2025). However, this comes with challenges as the continuous versions of such losses often require explicit regularization to avoid trivial solutions (Bianchi et al., 2020; Tsitsulin et al., 2024; Shchur & Günnemann, 2019).

3.2. Message Passing vs. Network Analysis

Message passing is a core mechanism in GNNs. It allows nodes to gather information from their neighbors and update their embedding, supporting tasks such as link prediction and node classification. While the specifics regarding message, update, and aggregation functions differ between different GNN architectures, they follow the same principle, often choosing neural networks, such as MLPs, for message and update, and aggregations such as mean, sum, or max. Recently, several works have proposed using other activations to improve robustness in situations with outliers or to learn aggregations that better align with specific tasks (Hamilton et al., 2017; Corso et al., 2020; Geisler et al., 2020; Jin et al., 2022). A deeper understanding of how different aggregation and activation choices interact with graph structure remains an open challenge.

Insights from network science could guide the development of more adaptive, interpretable, and task-specific message-passing mechanisms, helping the deep graph learning community design models that better capture structural patterns and improve generalization across diverse graph-based tasks. Most message-passing architectures rely on training to opti-

mize their message and update functions for specific objectives. But even without training, message passing captures structural patterns in graphs, including common neighbors, triadic closure, and PageRank. These measures have long helped network scientists understand graph topology and can inform GNN methods because they reflect key structural features of graphs. For example, common neighbors quantify the number of shared neighbors between two nodes. Without training, a single iteration of message passing aggregates information from neighboring nodes, enabling a GNN to approximate common neighbors directly. Similarly, multiple layers of message passing can capture more complex patterns, such as PageRank, by propagating information along paths of increasing length.

This overlap highlights a promising connection between GNNs and network science, where network science provides interpretable and well-established tools for efficiently analyzing graph structure. By integrating these insights, the deep graph learning community can develop architectures that not only improve predictive performance but also enhance interpretability and robustness across diverse graph-based applications.

3.3. Explaining Deep Learning Models

Breaking open the "black box" is one of the grand challenges in deep learning. Network analysis techniques offer an avenue for doing this. Using a new loss function and simple network measures such as modularity, clustering coefficients, and centrality, Bonifazi et al. (2024) and colleagues gain insight into the network structure of various GNNs. They then use these insights to improve the node classification performance of various GNNs.

As the authors note themselves (Bonifazi et al., 2024), this is only the tip of the iceberg. Many other network analysis techniques could be used to gain further insight into the structural elements, and hopefully the functioning, of these complex collections of nodes and edges. For example, one could use multi-layer networks as a way of capturing additional attributes of the embeddings. Others have used comparisons of eigenvectors across various layers of Convolutional Neural Networks to investigate universal encodings (Guth & Ménard, 2024). Could similar techniques be applied to GNNs as a way of tracking performance and representation changes as the models increase in scale?

3.4. From Example-based to Unsupervised Learning

Deep graph learning methods focus on learning data patterns in a supervised, semi-supervised, or self-supervised fashion. Supervised and semi-supervised learning rely on labeled data to quantify performance and train models' weights via backpropagation to minimize training loss. Consequently, they require that their training data contain examples cov-

ering all labels that will be considered during prediction. If this is not the case, the trained model cannot identify instances belonging to the unseen label. A way around this is to apply self-supervised techniques such as graph auto-encoders that encode and then decode the data, compressing it into a lower-dimensional representation, followed by decompressing it to reconstruct the original data (Kipf & Welling, 2016; Liu et al., 2023). The training criterion in this case is to minimize the discrepancy between the original and reconstructed data, which can be measured via the L_2 norm or other measures that quantify the "distance" between matrices. However, this approach allows the user very limited control over the exact patterns that are learnt during the embedding process.

In contrast, network science methods typically do not assume labeled training data and employ models designed to capture specific patterns. This makes them more widely applicable in a world where most data is unlabeled and research questions require focusing on a specific aspect of the data in a controllable way. Integrating such model-based approaches from network science with deep graph learning methods would make it possible to align GNNs with transparent and interpretable models and train them in an unsupervised fashion.

3.5. Balancing Interpretability and Flexibility

Deep graph learning and network science take different approaches to modeling graph data. Network science provides interpretable frameworks based on explicit assumptions about structure and relationships, but these assumptions can limit flexibility in complex graphs. In contrast, deep learning methods prioritize flexibility, accommodating diverse inputs and intricate dependencies, often at the cost of interpretability.

This trade-off between interpretability and flexibility is an opportunity for deep graph learning to incorporate insights from network science while maintaining flexibility. For example, GNNs struggle in heterophilic settings, where connected nodes differ in features or labels (Zheng et al., 2024). In contrast, statistical models like the SBM can capture both homophilic and heterophilic structures. Integrating such principled models into GNNs could improve generalization and robustness across varying graph topologies.

Meta-learning emerges as a promising approach to balancing interpretability and flexibility. It enables models to dynamically adjust their assumptions or learning strategies based on the characteristics of the input data. Hybrid frameworks can combine the interpretable structure of probabilistic network models with the scalability and flexibility of deep graph learning techniques. For example, integrating automatic differentiation and Bayesian inference enables scalable, interpretable models capable of handling diverse

network configurations (Contisciani et al., 2025). Similarly, GNNs inspired by structured models like the SBM could better capture heterophilic patterns, improving their applicability across different types of graphs (Wang et al., 2024).

3.6. Towards Foundation Models for Graphs

Finally, a major open challenge in deep graph learning is that current tasks are still largely domain-centered and models need to be trained for a specific task in a given graph or, for graph-level learning tasks, a specific set of graphs from a given domain. Different from recent advances in computer vision and natural language processing, deep graph learning still largely lacks *foundation models* that could generalize to new, previously unseen graphs (Morris et al., 2024). While deep graph learning has now taken the first steps in this direction (Zhao et al., 2024; Mao et al., 2024), network science has a long tradition of identifying and modeling *universal organizing principles* that govern the structure and evolution of networks across domains like social networks, information systems, biology, and large-scale infrastructures. To this end, many insights have been generated regarding how simple (local) mechanisms shape collective patterns and characteristics of large-scale networks and/or dynamical processes across domains. Concrete examples include phase transitions in the connectivity of large graphs that can be explained based on the ratio between moments of their degree distributions (Newman et al., 2001), simple growth rules like preferential attachment that lead to similar scale-invariant patterns in very different networks (Barabási & Albert, 1999), or common characteristics in financial and social networks that result in similar propagation and consensus dynamics across different systems (Lorenz et al., 2009). Despite these strong results in network science, we lack a theoretical understanding of whether and how these common organization principles in networks from across different domains affect the generalization capabilities of commonly used deep graph learning architectures, and how we could build architectures able to capture such principles.

4. Alternative Views

While the sections above outline how insights from network science can advance deep graph learning, one should also ask whether these arguments apply in the opposite direction. Given that deep graph learning already achieves high performance on its own, is integrating insights from network science unlikely or unnecessary? Will deep graph learning instead advance network science?

4.1. Scalability Challenges in Network Science

Deep graph learning emphasizes scalability and promotes efficient methods that scale to large datasets. This is particularly important in applications considering large and dynamic networks that evolve over time. While scalability is also an important aspect in network science, its focus lies more often than not on understanding the fundamental principles and properties of complex systems.

Many breakthroughs in computer vision, natural language processing, and graph learning have only been possible due to the computational power provided by GPUs. With thousands or even tens of thousands of special-purpose cores, modern GPUs have become essential in deep learning and enable processing large-scale data in a fraction of the time required on general-purpose CPUs. As large-scale network datasets with millions to billions of nodes and links become increasingly available, network science approaches struggle but have an opportunity to scale with the data by leveraging GPUs. This requires shifting from implementing algorithms with "traditional" control structures such as loops and conditions to expressing them in terms of (sparse) tensor operations and for computations on GPUs. Moreover, developers need to ensure that the data fits into GPU memory or employ batch or stream processing techniques.

Recently, accelerating network analyses with GPU-based implementations has received attention and brought forward tools that scale to large networks with substantial speedups: `nx-cugraph`[2] provides a GPU-accelerated backend for the popular Python network analysis library `networkx`, scaling computations such as node centrality via PageRank or betweenness centrality, and Louvain community detection to networks with millions of nodes and edges. `pathpyG`[3], a tool for analysing and modeling temporal graphs builds on `pytorch` (Paszke et al., 2019) and `pyG` (Fey & Lenssen, 2019), facilitating, for example, the efficient GPU-based calculation of time-respecting paths, temporal node centralities and higher-order graph models.

4.2. Evaluation Practices and Research Infrastructures

Deep learning and network science can be seen as two ends of the spectrum when it comes to the role of benchmarks and standardized evaluation practices. At the one end, method development in deep learning is often driven by a relatively small set of tasks at the node, link, and graph level: node classification, link prediction, and graph classification (Chami et al., 2022). To evaluate the performance of new methods, the deep learning community places a strong focus on utilizing standardized benchmarks with common datasets, frameworks, and leaderboards, which has greatly improved our ability to compare methods (Hu et al., 2020; Huang et al., 2023). At the other end, network science works are typically driven by a specific application or the limitations of existing methods to capture specific aspects of a complex system. While standardized methods for gen-

---

[2] https://github.com/rapidsai/nx-cugraph
[3] https://github.com/pathpy/pathpyG

erating benchmark graphs exist (Karrer & Newman, 2011; Lancichinetti et al., 2008; Lee & Wilkinson, 2019; Kamiński et al., 2021), they are used with application-specific parameters, and, thus, different methods are seldom evaluated under the exact same setting. Several online repositories host collections of networks often used for benchmarking methods (Kunegis, 2013; Leskovec & Krevl, 2014; Clauset et al., 2016; Peixoto, 2020), however, altogether there are hundreds of networks available but without any subset of datasets that would be considered a "standard benchmark".

Following the example set by the deep learning community regarding evaluation practices carries opportunities for the network science community. The most direct benefit would be to make methods more easily comparable by standardizing tasks, datasets, and evaluation frameworks. This would provide clearer guidance for practitioners when choosing a method to answer a research question. It would also make it easier for researchers entering the field to choose a starting point for evaluating their methods. However, as with all things, there are several open challenges that need to be considered: Typically, there is no one-size-fits-all benchmark setup that suits all scenarios, especially considering the wide range of applications in network science. Choosing a relatively small number of benchmark datasets may lead to reduced exploration and diversity, limiting the community to a prescribed set of designated "standard networks". Works that consider off-mainstream research questions may be forced to shoehorn their evaluation into existing frameworks and argue for their non-standard evaluation in order to avoid rejection.

4.3. Bridging Scientific Cultures

A final challenge that could hinder a better integration is that network science and deep learning largely differ in terms of scientific culture. The attentive reader may have noticed that the majority of network science works cited in this paper were either published in (interdisciplinary) journals or in statistical physics outlets, that are largely outside the scope of machine learning researchers. In contrast, the deep learning community largely focuses on (a small set of) major machine learning venues. This difference in publication practices not only makes it difficult to communicate results, it also affects the strategies and research challenges adopted by junior researchers who rely on the prestige of publication venues to secure a tenured position.

Second, the two communities use largely different terminology, sometimes for the same concepts. This is due to the fact that many concepts in network science, such as statistical ensembles, phase transitions, etc., originally come from the statistical physics study of complex systems.

These two aspects create challenges not only for network science researchers to publish in deep learning and deep

learning researchers publishing in network science, it also creates challenges for the visibility of the large body of results in network science, which – since it has been published in physics journals - remains largely unknown to the deep learning community.

Facilitated by the previous point that limits the amount of deep learning works that utilize network science models, in the network science community, the recent rapid growth of works on deep graph learning is partly seen as a competitive rather than a collaborative effort.

5. Conclusion

We offer a pathway for integrating methods and insights from network science into deep graph learning and vice versa. By bridging the theoretical frameworks of network science with the scalability and adaptability of deep learning, researchers in the deep graph learning community can address challenges such as robust and principled data augmentation, effective pooling techniques, and higher-order modeling of complex interactions. We also see potential advancements for network science by incorporating standardized benchmarks, evaluation practices, and GPU-based implementations of algorithms that scale to massive graphs.

To realize this potential, we call for integrating network science methods—such as random graph theory, generating function analysis of random graph ensembles, and models of dynamical processes like diffusion, synchronization, or consensus dynamics—into computer science and deep learning academic curricula. Similarly, network and data science curricula would benefit from incorporating deep learning methods, including GPU-based tensor operations, fast gradient-based optimization, end-to-end learning approaches, and common evaluation procedures such as hyperparameter tuning and ablation studies.

At the scholarly level, fostering collaboration between network science and deep graph learning is essential. Strengthening this connection requires bringing researchers from different backgrounds and publication cultures together, whether through dedicated journal-first tracks in deep learning conferences or by organizing workshops and special journal issues that bridge the two fields.

Ultimately, we see this convergence as more than just a set of technical advances. By embracing the interpretability of network science and the flexibility of deep learning, researchers from both fields will benefit—just as Hopfield did by integrating ideas across statistical physics, neuroscience, and complex systems.

Acknowledgements

References

Antelmi, A., Cordasco, G., Polato, M., Scarano, V., Spagnuolo, C., and Yang, D. A survey on hypergraph representation learning. *ACM Comput. Surv.*, 56(1), August 2023. ISSN 0360-0300. doi: 10.1145/3605776.

Azabou, M., Ganesh, V., Thakoor, S., Lin, C.-H., Sathidevi, L., Liu, R., Valko, M., Veličković, P., and Dyer, E. L. Half-hop: A graph upsampling approach for slowing down message passing. In *International Conference on Machine Learning*, pp. 1341–1360. PMLR, 2023.

Barabási, A.-L. and Albert, R. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

Battiston, F., Cencetti, G., Iacopini, I., Latora, V., Lucas, M., Patania, A., Young, J.-G., and Petri, G. Networks beyond pairwise interactions: Structure and dynamics. *Physics Reports*, 874:1–92, 2020. ISSN 0370-1573. doi: https://doi.org/10.1016/j.physrep.2020.05.004. Networks beyond pairwise interactions: Structure and dynamics.

Bianchi, F. M., Grattarola, D., and Alippi, C. Spectral clustering with graph neural networks for graph pooling. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 874–883. PMLR, 13–18 Jul 2020.

Blöcker, C., Tan, C., and Scholtes, I. The map equation goes neural: Mapping network flows with graph neural networks. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, oct 2008. doi: 10.1088/1742-5468/2008/10/P10008.

Blöcker, C. and Scholtes, I. Flow divergence: Comparing maps of flows with relative entropy, 2024. URL https://arxiv.org/abs/2401.09052.

Bonifazi, G., Cauteruccio, F., Corradini, E., Marchetti, M., Ursino, D., and Virgili, L. A network analysis-based framework to understand the representation dynamics of graph neural networks. *Neural Computing and Applications*, 36(4):1875–1897, 2024.

Casiraghi, G., Nanumyan, V., Scholtes, I., and Schweitzer, F. From relational data to graphs: Inferring significant links using generalized hypergeometric ensembles. In Ciampaglia, G. L., Mashhadi, A., and Yasseri, T. (eds.), *Social Informatics*, pp. 111–120, Cham, 2017. Springer International Publishing. ISBN 978-3-319-67256-4.

Castellana, D. and Bianchi, F. M. Bn-pool: a bayesian nonparametric approach to graph pooling, 2025.

Chami, I., Abu-El-Haija, S., Perozzi, B., Ré, C., and Murphy, K. Machine learning on graphs: A model and comprehensive taxonomy. *Journal of Machine Learning Research*, 23(89):1–64, 2022.

Clauset, A., Tucker, E., and Sainz, M. The colorado index of complex networks, 2016. URL https://icon.colorado.edu.

Contisciani, M., Hobbhahn, M., Power, E. A., Hennig, P., and De Bacco, C. Flexible inference in heterogeneous and attributed multilayer networks. *PNAS Nexus*, pp. pgaf005, 2025.

Corso, G., Cavalleri, L., Beaini, D., Liò, P., and Veličković, P. Principal neighbourhood aggregation for graph nets. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 13260–13271. Curran Associates, Inc., 2020.

De Bruijn, N. G. A combinatorial problem. *Proceedings of the Section of Sciences of the Koninklijke Nederlandse Akademie van Wetenschappen te Amsterdam*, 49(7):758–764, 1946.

De Domenico, M., Solé-Ribalta, A., Cozzo, E., Kivelä, M., Moreno, Y., Porter, M. A., Gómez, S., and Arenas, A. Mathematical formulation of multilayer networks. *Phys. Rev. X*, 3:041022, Dec 2013. doi: 10.1103/PhysRevX.3.041022.

Deng, S., Yang, G., Yang, Y., Gong, Z., Chen, C., Chen, X., and Hao, Z. Module-based graph pooling for graph classification. *Pattern Recognition*, 154:110606, 2024.

Edler, D., Bohlin, L., and Rosvall, M. Mapping higher-order network flows in memory and multilayer networks with infomap. *Algorithms*, 10(4), 2017. ISSN 1999-4893. doi: 10.3390/a10040112.

Erdos, P., Rényi, A., et al. On the evolution of random graphs. *Publ. math. inst. hung. acad. sci*, 5(1):17–60, 1960.

Fey, M. and Lenssen, J. E. Fast graph representation learning with pytorch geometric, 2019.

Frantzen, F. and Schaub, M. T. Learning from simplicial data based on random walks and 1d convolutions, 2024. URL https://arxiv.org/abs/2404.03434.

Geisler, S., Zügner, D., and Günnemann, S. Reliable graph neural networks via robust aggregation. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 13272–13284. Curran Associates, Inc., 2020.

Georgousis, S., Kenning, M. P., and Xie, X. Graph deep learning: State of the art and challenges. *IEEE Access*, 9:22106–22140, 2021. doi: 10.1109/ACCESS.2021. 3055280.

Grover, A. and Leskovec, J. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 855–864, 2016.

Guth, F. and Ménard, B. On the universality of neural encodings in cnns. In *ICLR 2024 Workshop on Representational Alignment*, 2024.

Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

Heeg, F. and Scholtes, I. Using time-aware graph neural networks to predict temporal centralities in dynamic graphs. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

Holme, P. Modern temporal network theory: a colloquium. *The European Physical Journal B*, 88:1–30, 2015.

Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 22118–22133. Curran Associates, Inc., 2020.

Huang, S., Poursafaei, F., Danovitch, J., Fey, M., Hu, W., Rossi, E., Leskovec, J., Bronstein, M., Rabusseau, G., and Rabbany, R. Temporal graph benchmark for machine learning on temporal graphs. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA, 2023. Curran Associates Inc.

Jamadandi, A., Rubio-Madrigal, C., and Burkholz, R. Spectral graph pruning against over-squashing and over-smoothing. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

Jin, D., Wang, R., Ge, M., He, D., Li, X., Lin, W., and Zhang, W. Raw-gnn: Random walk aggregation based graph neural network. In Raedt, L. D. (ed.), *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pp. 2108–2114. International Joint Conferences on Artificial Intelligence Organization, 7 2022. doi: 10.24963/ijcai.2022/293. Main Track.

Ju, W., Yi, S., Wang, Y., Xiao, Z., Mao, Z., Li, H., Gu, Y., Qin, Y., Yin, N., Wang, S., Liu, X., Luo, X., Yu, P. S., and Zhang, M. A survey of graph neural networks in real world: Imbalance, noise, privacy and ood challenges, 2024. URL https://arxiv.org/abs/2403.04468.

Kamiński, B., Prałat, P., and Théberge, F. Artificial benchmark for community detection (abcd)—fast random graph model with community structure. *Network Science*, 9(2):153–178, 2021. doi: 10.1017/nws.2020.45.

Karrer, B. and Newman, M. E. J. Stochastic blockmodels and community structure in networks. *Phys. Rev. E*, 83:016107, Jan 2011. doi: 10.1103/PhysRevE.83.016107.

Kim, S., Lee, S. Y., Gao, Y., Antelmi, A., Polato, M., and Shin, K. A survey on hypergraph neural networks: An in-depth and step-by-step guide. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, pp. 6534–6544, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704901. doi: 10.1145/3637528.3671457.

Kipf, T. N. and Welling, M. Variational graph auto-encoders. In *Proceedings of the Bayesian Deep Learning Workshop, 30th Conference on Neural Information Processing Systems (NeurIPS)*, Centre Convencions Internacional Barcelona, Barcelona, Spain, 2016.

Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., and Porter, M. A. Multilayer networks. *Journal of Complex Networks*, 2(3):203–271, 07 2014. ISSN 2051-1310. doi: 10.1093/comnet/cnu016.

Kojaku, S., Radicchi, F., Ahn, Y.-Y., and Fortunato, S. Network community detection via neural embeddings. *Nature Communications*, 15(1):9446, 2024.

Kunegis, J. Konect: the koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web*, WWW '13 Companion, pp. 1343–1350, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450320382. doi: 10.1145/2487788. 2488173.

Lambiotte, R., Rosvall, M., and Scholtes, I. From networks to optimal higher-order models of complex systems. *Nature physics*, 15(4):313–320, 2019.

Lancichinetti, A., Fortunato, S., and Radicchi, F. Benchmark graphs for testing community detection algorithms. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 78(4):046110, 2008.

Lee, C. and Wilkinson, D. J. A review of stochastic block models and extensions for graph clustering. *Applied Network Science*, 4(1):1–50, 2019.

Leskovec, J. and Krevl, A. SNAP Datasets: Stanford large network dataset collection, June 2014. URL http://snap.stanford.edu/data.

Liu, Y., Jin, M., Pan, S., Zhou, C., Zheng, Y., Xia, F., and Yu, P. S. Graph self-supervised learning: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 35(6):5879–5900, 2023. doi: 10.1109/TKDE.2022.3172903.

Longa, A., Lachi, V., Santin, G., Bianchini, M., Lepri, B., Lio, P., Scarselli, F., and Passerini, A. Graph neural networks for temporal graphs: State of the art, open challenges, and opportunities, 2023. URL https://arxiv.org/abs/2302.01018.

Lorenz, J., Battiston, S., and Schweitzer, F. Systemic risk in a unifying framework for cascading processes on networks. *The European Physical Journal B*, 71:441–460, 2009.

Mao, H., Chen, Z., Tang, W., Zhao, J., Ma, Y., Zhao, T., Shah, N., Galkin, M., and Tang, J. Position: graph foundation models are already here. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Moinet, A., Starnini, M., and Pastor-Satorras, R. Burstiness and aging in social temporal networks. *Physical review letters*, 114(10):108701, 2015.

Molloy, M. and Reed, B. A critical point for random graphs with a given degree sequence. *Random structures & algorithms*, 6(2-3):161–180, 1995.

Morris, C., Frasca, F., Dym, N., Maron, H., Ceylan, I. I., Levie, R., Lim, D., Bronstein, M. M., Grohe, M., and Jegelka, S. Position: Future directions in the theory of graph machine learning. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 36294–36307. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/morris24a.html.

Newman, M. E. Random graphs with clustering. *Physical review letters*, 103(5):058701, 2009.

Newman, M. E. Network structure from rich but noisy data. *Nature Physics*, 14(6):542–545, 2018.

Newman, M. E., Strogatz, S. H., and Watts, D. J. Random graphs with arbitrary degree distributions and their applications. *Physical review E*, 64(2):026118, 2001.

Pareja, A., Domeniconi, G., Chen, J., Ma, T., Suzumura, T., Kanezashi, H., Kaler, T., Schardl, T., and Leiserson, C. Evolvegcn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 5363–5370, 2020.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Peixoto, T. P. Bayesian stochastic blockmodeling. *Advances in network clustering and blockmodeling*, pp. 289–332, 2019a.

Peixoto, T. P. Network reconstruction and community detection from dynamics. *Physical review letters*, 123(12):128301, 2019b.

Peixoto, T. P. The netzschleuder network catalogue and repository, August 2020.

Peixoto, T. P. and Rosvall, M. Modelling sequences and temporal networks with dynamic community structures. *Nature communications*, 8(1):582, 2017.

Qarkaxhija, L., Perri, V., and Scholtes, I. De bruijn goes neural: Causality-aware graph neural networks for time series data on dynamic graphs. In Rieck, B. and Pascanu, R. (eds.), *Proceedings of the First Learning on Graphs Conference*, volume 198 of *Proceedings of Machine Learning Research*, pp. 51:1–51:21. PMLR, 09–12 Dec 2022. URL https://proceedings.mlr.press/v198/qarkaxhija22a.html.

Robins, G., Pattison, P., Kalish, Y., and Lusher, D. An introduction to exponential random graph (p*) models for social networks. *Social networks*, 29(2):173–191, 2007.

Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., and Bronstein, M. M. Temporal graph networks for deep learning on dynamic graphs. *CoRR*, abs/2006.10637, 2020. URL https://arxiv.org/abs/2006.10637.

Rosvall, M. and Bergstrom, C. T. Maps of random walks on complex networks reveal community structure. *Proceedings of the national academy of sciences*, 105(4): 1118–1123, 2008.

Rosvall, M., Esquivel, A. V., Lancichinetti, A., West, J. D., and Lambiotte, R. Memory in network flows and its effects on spreading dynamics and community detection. *Nature communications*, 5(1):4630, 2014.

Rubio-Madrigal, C., Jamadandi, A., and Burkholz, R. Gnns getting comfy: Community and feature similarity guided rewiring. In *The Thirteenth International Conference on Learning Representations*, 2025.

Scholtes, I., Wider, N., Pfitzner, R., Garas, A., Tessone, C. J., and Schweitzer, F. Causality-driven slow-down and speed-up of diffusion in non-markovian temporal networks. *Nature communications*, 5(1):5024, 2014.

Shchur, O. and Günnemann, S. Overlapping community detection with graph neural networks. *Deep Learning on Graphs Workshop, KDD*, 2019.

Takaguchi, T., Masuda, N., and Holme, P. Bursty communication patterns facilitate spreading in a threshold-based epidemic dynamics. *PloS one*, 8(7):e68629, 2013.

Torres, L., Blevins, A. S., Bassett, D., and Eliassi-Rad, T. The why, how, and when of representations for complex systems. *SIAM Review*, 63(3):435–485, 2021. doi: 10.1137/20M1355896.

Traag, V. A., Waltman, L., and van Eck, N. J. From louvain to leiden: guaranteeing well-connected communities. *Scientific Reports*, 9(1):5233, Mar 2019. ISSN 2045-2322. doi: 10.1038/s41598-019-41695-z.

Tsitsulin, A., Palowitch, J., Perozzi, B., and Müller, E. Graph clustering with graph neural networks. *J. Mach. Learn. Res.*, 24(1), March 2024. ISSN 1532-4435.

von Pichowski, J., Blöcker, C., and Scholtes, I. Hierarchical graph pooling based on minimum description length, 2024.

Wang, J., Guo, Y., Yang, L., and Wang, Y. Understanding heterophily for graph neural networks. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 50489–50529. PMLR, 21–27 Jul 2024.

Zhao, J., Mostafa, H., Galkin, M., Bronstein, M., Zhu, Z., and Tang, J. Graphany: A foundation model for node classification on any graph, 2024. URL https://arxiv.org/abs/2405.20445.

Zhao, T., Liu, Y., Neves, L., Woodford, O. J., Jiang, M., and Shah, N. Data augmentation for graph neural networks. *CoRR*, abs/2006.06830, 2020. URL https://arxiv.org/abs/2006.06830.

Zhao, T., Jin, W., Liu, Y., Wang, Y., Liu, G., Günnemann, S., Shah, N., and Jiang, M. Graph data augmentation for graph machine learning: A survey, 2023. URL https://arxiv.org/abs/2202.08871.

Zheng, Y., Luan, S., and Chen, L. What is missing for graph homophily? disentangling graph homophily for graph neural networks. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.