

Graph Neural Networks for Predicting Drug–Target Interactions and Synergistic Therapeutic Combinations

Ishaan C. Prasad

Department of Computer Science, University of Central Florida, Orlando, FL, USA.

ishaan.prasad13@ucf.edu

Jean Meraleis

Department of Computer Science, Binghamton University, Binghamton, NY, USA.

moralesjean@binghamton.edu

Abstract

The identification of drug-target interactions and synergistic therapeutic combinations remains a central challenge in modern drug discovery, particularly as the complexity of biological systems and the cost of experimental screening continue to grow. Graph neural networks have emerged as a powerful framework for representing relational structure in molecular, genomic, and pharmacological data. Unlike traditional machine learning approaches that rely on independent feature vectors, graph-based methods explicitly model interactions among drugs, proteins, pathways, and diseases. This paper provides a system-level analysis of graph neural network architectures for drug-target interaction prediction and synergistic combination discovery. It examines the structural and architectural trade-offs associated with molecular graph encoding, heterogeneous biomedical networks, message passing, and multi-relational learning. The discussion further addresses data infrastructure, standardization, validation, and integration with experimental platforms. Emphasis is placed on deployment in clinical and pharmaceutical settings, including regulatory considerations, explainability, robustness, fairness, and environmental sustainability. The paper also considers policy implications for open data sharing, model governance, and equitable access. By synthesizing technical advances with socio-technical infrastructure concerns, the article offers a comprehensive perspective on the opportunities and systemic constraints that shape the translation of graph neural network models into reliable tools for precision medicine and combination therapy development.

Keywords

graph neural networks; drug-target interactions; synergistic combinations; systems pharmacology; network medicine; clinical translation; computational governance.

1. Introduction

The pharmaceutical research enterprise continues to confront a widening gap between the scale of molecular and genomic data generation and the capacity to convert these data into safe, effective therapeutic interventions. Target identification, mechanism elucidation, and the rational design of combination therapies remain constrained by the multi-scale nature of human biology, the incompleteness of interaction knowledge, and the high failure rates that characterize late-stage clinical development. Drug-target interaction prediction has traditionally relied on structural biology, ligand-based screening, or statistical associations

between molecular descriptors and experimental binding outcomes. Although such methods have produced substantial successes, they frequently fail to capture the relational dependencies that determine drug action in living systems, where a compound may bind multiple proteins, a protein may participate in numerous pathways, and therapeutic benefit often emerges from the coordinated modulation of several targets.

Graph neural networks offer a natural computational framework for addressing these relational challenges. Early graph representation learning strategies, including graph convolutional networks [1], graph attention networks [2], and inductive message passing methods [3], demonstrated that node-level and graph-level predictions can be improved by aggregating information from neighboring entities in a network. In biomedical applications, the entities may include drugs, proteins, diseases, genes, pathways, or patient-derived features, while edges may encode physical binding, functional similarity, co-expression, disease association, or pharmacological interaction. By propagating feature information across these relationships, graph-based models can learn representations that reflect both local molecular properties and broader biological context. This capacity is particularly valuable for drug-target interaction prediction, where the same compound can exhibit context-dependent selectivity, and for synergistic combination analysis, where the joint effect of two or more agents depends on network-level interference and pathway crosstalk.

Within the drug-target interaction literature, deep sequence-based architectures have been widely adopted. Models such as DeepDTA [4] process drug and target sequences through convolutional layers to predict binding affinity, while GraphDTA [5] represents drug molecules as graphs and applies graph neural operations to capture topological chemical features. The end-to-end learning framework introduced for compound-protein interaction prediction [6] further illustrates how molecular graph encoders and protein sequence encoders can be jointly optimized. More recent transformer-based systems, including MolTrans [7], have shown that capturing non-local dependencies in molecular and protein sequences can improve interaction prediction. These approaches share a common conceptual shift away from hand-crafted descriptors toward learned representations that preserve structural information.

Synergy prediction introduces additional systemic complexity. A drug combination may produce a therapeutic effect that is greater than, equal to, or less than the sum of the individual agent effects. Deep learning models such as DeepSynergy [8] have used chemical descriptors and genomic features to predict synergy scores across cancer cell lines, while platforms such as SynergyFinder [9] provide dose-response matrix analysis and visualization for experimental combination studies. In parallel, graph-based approaches have been applied to polypharmacy side effect prediction by modeling drug-drug and drug-protein association networks [10]. Network community-based models have also been proposed for identifying synergistic combinations from herbal medicines and complex systems [11]. These methods treat synergy as a property that emerges from the organization of biological networks rather than from isolated molecular interactions.

This article examines graph neural network approaches for drug-target interaction prediction and synergistic combination identification from a systems perspective. Rather than focusing only on architectural novelty, the analysis emphasizes structural trade-offs in graph construction, multi-relational learning, data governance, deployment readiness, robustness, fairness, and sustainability. The objective is to articulate how technical design decisions interact with institutional and regulatory infrastructures, and to identify the conditions under

which graph-based models can be responsibly translated into real-world therapeutic decision support.

2. Background and System Context

The conceptual foundations of graph-based drug discovery are rooted in network medicine, which views disease as a consequence of perturbations in molecular interaction networks rather than as the result of single isolated molecular defects. Network medicine frameworks [12] propose that disease-associated proteins are not randomly distributed but tend to cluster in specific network neighborhoods, forming disease modules. This modular organization has important implications for drug-target interaction prediction because it suggests that compounds acting on proteins within or near a disease module may be more likely to produce therapeutic benefit. It also informs combination therapy design, since interventions targeting different nodes within the same module or across complementary modules can produce synergistic effects while minimizing off-target toxicity.

Biomedical knowledge bases provide the relational substrate for graph-based modeling. DrugBank [13] integrates detailed drug, target, and interaction information, including chemical structures, pharmacological data, and known drug-target pairs. These curated resources are complemented by pathway databases, protein-protein interaction maps, and genomic annotation platforms. The usefulness of these resources depends on their coverage, consistency, and update frequency. Models trained on incomplete or biased interaction graphs may inherit systematic blind spots, especially for understudied targets, rare diseases, or emerging classes of therapeutic agents. Data integration is therefore not merely a preprocessing concern but a foundational governance challenge that determines the generalizability of learned representations.

The prediction of selective drug combinations has also been pursued through target inhibition network approaches that prioritize combinations of druggable targets whose simultaneous inhibition may block cancer survival pathways [14]. Such strategies emphasize the logic of network perturbation rather than simple pairwise synergy scoring. They demonstrate that effective combination prediction requires an understanding of signaling redundancy, feedback loops, and compensatory mechanisms. Graph neural networks extend this logic by providing differentiable operators that can learn from large-scale interaction graphs, but the underlying systems reasoning remains essential for interpreting model outputs and constraining hypothesis spaces.

Machine learning applications in drug discovery have proliferated across target validation, hit identification, and preclinical optimization [15]. However, translating model predictions into robust experimental designs remains difficult. The pharmaceutical data environment is heterogeneous, noisy, and characterized by sparse positive labels. In drug-target interaction prediction, negative examples are often inferred from the absence of known interactions, which can introduce systematic uncertainty. In synergy prediction, the number of possible pairwise combinations is enormous, and experimental dose-response matrices are expensive to generate. These conditions favor model architectures that can exploit relational structure and propagate information from well-characterized regions of the network to less-characterized regions.

3. Graph-Based Representations of Drug–Target and Combination Spaces

A central design decision in graph-based drug-target interaction prediction concerns the representation of molecular entities. Drugs can be represented as molecular graphs in which

atoms are nodes and chemical bonds are edges. This representation preserves the topological structure of the molecule and allows graph neural networks to learn atomic neighborhoods, functional group patterns, and substructural motifs. Protein targets can likewise be represented as contact graphs derived from structure or as sequence graphs in which amino acid residues are connected by local and predicted long-range relationships. Alternatively, the drug-target prediction problem can be formulated on a bipartite interaction graph, where drug nodes and target nodes are connected by experimentally observed or computationally predicted interactions. This formulation permits transductive inference, in which the model predicts missing edges in the graph by aggregating signals from known interactions.

Heterogeneous biomedical networks introduce additional edge types and node categories. A typical network may include drugs, proteins, diseases, pathways, and genomic variants, with edges representing binding, expression regulation, pathway membership, and disease association. The advantage of heterogeneity is that it enables the model to integrate diverse sources of evidence within a unified relational schema. The cost is architectural complexity, since different edge types may require separate message functions, attention mechanisms, or relation-specific transformations. Graph neural networks designed for multi-relational data must balance representational expressiveness against overparameterization and training instability. In the context of drug-target interaction prediction, this tension is particularly acute because the available training labels are often insufficient to constrain large numbers of relation-specific parameters.

The representation of drug combinations extends the graph perspective further. A combination can be modeled as a hypergraph in which drug nodes, target nodes, pathway nodes, and disease nodes are connected through higher-order relationships. Graph-based approaches to polypharmacy side effect prediction [10] have shown that modeling drug-drug edges alongside drug-protein edges improves the detection of adverse interactions. Similarly, synergy prediction can be formulated as a link prediction problem in a network that includes drug combination nodes connected to pharmacological outcomes. In this view, synergy is not an intrinsic property of two drugs alone but an emergent property of how their joint perturbation propagates across a network of molecular interactions and functional dependencies.

The selection of graph construction parameters has profound implications for model behavior. The radius of molecular graph neighborhoods, the inclusion or exclusion of salt and solvent atoms, the treatment of stereochemistry, and the strategy for generating protein contact graphs all influence the learned representation. In heterogeneous network modeling, the choice of edge weight thresholds, the handling of missing data, and the aggregation of multiple evidence types are equally consequential. Graph neural network models are not invariant to these choices, and performance improvements observed in benchmark studies may reflect favorable graph representations as much as architectural innovations. A systems-level evaluation should therefore report graph construction protocols and ablation analyses that isolate the contribution of the relational substrate.

Graph computational methods have also been surveyed as a general resource for drug development and discovery [16], highlighting applications in molecular property prediction, de novo design, and biomedical network analysis. These applications differ in their output targets but share common infrastructure requirements. They demand consistent molecular identifiers, standardized node features, reproducible graph partitioning, and careful handling of data leakage. In drug-target interaction prediction, data leakage is particularly problematic

when structurally similar drugs or homologous proteins appear in both training and test sets. Unless explicit splits are designed to test generalization to unseen scaffolds or novel targets, reported accuracy may overstate real-world utility.

4. Architectures and Learning Paradigms for Interaction Prediction

Graph neural network architectures for drug-target interaction prediction can be broadly categorized according to the level of graph abstraction and the learning objective. In molecular graph encoding approaches, the model first computes a graph-level drug representation from the two-dimensional molecular structure and then combines it with a protein representation to predict binding affinity or interaction probability. Graph convolutional networks [1] use learned aggregation over local neighborhoods, while graph attention networks [2] assign different importance weights to neighboring atoms or residues. GraphSAGE [3] introduced an inductive sampling and aggregation framework that is particularly useful for large molecular and biomedical graphs. These methods differ in their treatment of node degree heterogeneity, their ability to capture long-range dependencies, and their computational efficiency.

Sequence-based encoders for proteins and transformers for drug-target interaction prediction have also been integrated with graph modules. MolTrans [7] uses a molecular transformer to encode drug substructures from unlabeled chemical data and a protein transformer to encode target sequences, with an interaction module that captures residue-atom correspondence. Such hybrid architectures benefit from transfer learning but require careful alignment between the representation spaces. The interaction module must decide whether to model binding as a global similarity score, a local contact map, or a bipartite attention mechanism. These choices affect not only accuracy but also interpretability, since contact-based models can highlight the specific molecular regions that drive the predicted interaction.

The learning paradigm itself is another critical variable. Many drug-target interaction models are trained as binary classifiers, but binding affinity regression provides a more informative output for lead optimization. Multi-task learning can combine interaction prediction with auxiliary tasks such as toxicity prediction, metabolic stability, or functional class prediction. These auxiliary tasks can regularize the representation and improve generalization when the primary labels are sparse. However, multi-task training also introduces trade-offs between task-specific and shared features. Graph models must allocate parameter capacity in ways that preserve both molecular specificity and cross-task transfer.

In the context of combination prediction, graph-based models often operate on a higher-level pharmacological network. The model may predict synergy scores by aggregating drug representations, target representations, and cell line genomic features. Relational inductive biases allow the model to share information across combinations that involve overlapping drugs or common pathway targets. This is important because the combinatorial space is too large to be explored exhaustively, and sparse experimental observations must be generalized to untested combinations. Multi-relational graph learning can explicitly model drug-drug similarity, target-pathway membership, and disease-target associations, thereby constraining the prediction space with biologically meaningful priors.

An important architectural consideration is the treatment of graph sparsity and scalability. Biomedical graphs can contain millions of nodes and edges, but the number of labeled drug-target interactions or combination experiments is comparatively small. Graph neural networks trained on such graphs must avoid over-smoothing, in which repeated message passing

renders node representations indistinguishable. Residual connections, attention mechanisms, and gating operations can mitigate over-smoothing while preserving the benefits of multi-hop information flow. At the systems level, these mechanisms introduce additional hyperparameters and memory requirements that must be managed in production environments.

The use of graph neural networks for synergy prediction also raises the question of whether synergy should be modeled as a pairwise property or as a network-level intervention effect. Pairwise models are simpler to train and validate but may miss higher-order interactions involving three or more drugs. Network-level models are more expressive but require more complex training objectives and evaluation protocols. The choice should be guided by the intended clinical use case. For example, in oncology, combinations are often tested in specific molecular subtypes, and the relevant prediction target may be a synergy score conditioned on genomic alterations. In antimicrobial combination screening, the relevant target may be fractional inhibitory concentration indices for diverse pathogen strains. These different problem settings imply different graph schemas, edge semantics, and validation criteria.

5. Synergistic Combination Prediction as a Systems Challenge

Synergy prediction cannot be reduced to a simple extension of drug-target affinity prediction. A combination may be synergistic in one cellular context but additive or antagonistic in another. The determinants of synergy include pathway redundancy, feedback activation, drug transport, metabolic crosstalk, and microenvironmental factors. Therefore, graph-based models must integrate multiple layers of biological context to produce predictions that are robust across experimental systems. The dose-response landscape itself is multidimensional, and a single synergy score may compress important information about potency shifts, efficacy ceilings, and window of selectivity.

Network community-based methods for herbal medicine combinations [11] illustrate how community structure in complex biological networks can guide the identification of synergistic formulations. These approaches rely on the intuition that agents targeting different communities within a disease network may jointly perturb complementary processes. In graph neural network terms, community detection can serve as a pre-training signal or as an auxiliary task that encourages the model to learn network neighborhoods associated with therapeutic synergy. The integration of community structure with differentiable graph learning remains an active area of investigation.

Experimental synergy datasets are often small, institutionally fragmented, and heterogeneous in their measurement protocols. Models such as DeepSynergy [8] have been evaluated on standard cell line panels, but the transferability of these models to patient-derived samples or novel drug classes remains uncertain. The statistical power of synergy prediction is further limited by the high dimensionality of the input space relative to the number of labeled combinations. Graph-based models can partially address this limitation by leveraging unlabeled network data, but they cannot fully compensate for the absence of systematic, high-quality combination screening data. Data generation infrastructure is therefore inseparable from algorithmic development.

The biological interpretation of synergy predictions is a major challenge for clinical translation. A model that accurately predicts a synergy score may not explain which molecular mechanisms drive the interaction. Graph attention mechanisms provide a partial solution by assigning importance weights to edges or nodes, but attention weights are not always causally interpretable. There is a risk that clinicians may overinterpret attention

patterns as biological mechanism, especially when visualizations appear compelling. It is important to pair graph-based predictions with pathway enrichment, network perturbation analysis, and experimental validation in relevant model systems.

The problem of combination therapy also involves regulatory and clinical dimensions. A predicted synergistic combination may involve approved drugs whose combined use is off-label, requiring rigorous evidence generation before adoption. In oncology, combination trials are often designed with biomarker stratification, and computational models may contribute to patient selection. In infectious disease, synergy prediction may inform treatment of resistant infections, but clinical breakpoints and resistance mechanisms must be considered. The graph-based model therefore operates within a larger decision infrastructure that includes clinical trial design, companion diagnostics, and post-market surveillance.

6. Data Infrastructure, Standards, and Integration

Reliable graph-based drug-target and synergy prediction depends on the quality, interoperability, and governance of biomedical data. Molecular structures, protein sequences, interaction annotations, and pharmacological outcomes must be linked across multiple databases with consistent identifiers. DrugBank [13] provides a foundational resource for drug-target interactions, but its coverage is uneven across target classes and therapeutic areas. Integrating DrugBank with genomic, transcriptomic, and proteomic datasets requires careful mapping of gene and protein identifiers. Even small identifier mismatches can create false edges or missing connections in the graph and degrade model performance.

Standardization is particularly important for combination screening data. Dose-response matrices, synergy metrics, and experimental conditions must be encoded in ways that support cross-study learning. Platforms such as SynergyFinder [9] have contributed to the standardization of dose-response analysis, but widespread adoption of common data schemas remains limited. The lack of standardized negative controls and the variability in cell line authentication further complicate model training and external validation. A systems approach should therefore include data provenance tracking, versioned graph construction, and community-level agreements on minimum information requirements.

Computational workflows for graph-based drug discovery also require integration with high-performance computing and cloud infrastructure. Training graph neural networks on large heterogeneous networks can be computationally intensive, and hyperparameter optimization amplifies this cost. Efficient data loaders, graph sampling strategies, and distributed training frameworks are critical for reproducibility and scalability. However, computational performance should not come at the expense of model interpretability or auditability. The infrastructure must support provenance records that link training data versions, model parameters, and evaluation outputs.

Data sharing and privacy constraints shape what can be learned from patient-derived networks. Federated learning and differential privacy offer partial solutions, but they introduce additional complexity in graph-based settings because graph topology itself may be sensitive. Patient similarity networks, co-expression graphs, and clinical interaction graphs may leak information about individuals even when node features are protected. System designers must consider privacy-preserving graph transformations and threat models that account for linkage attacks across datasets. The governance of biomedical graph data must balance open science with patient confidentiality and institutional data sovereignty.

7. Deployment, Governance, and Clinical Translation

Deploying graph neural network models in pharmaceutical and clinical environments requires attention to model lifecycle management, monitoring, and regulatory alignment. A model that performs well on retrospective benchmarks may degrade when applied to new chemical scaffolds, patient populations, or data distributions. Continuous monitoring is necessary to detect performance drift and to trigger recalibration. In drug-target interaction prediction, distribution shift often arises when new target classes or therapeutic modalities are introduced. In synergy prediction, differences in cell line protocols, sequencing platforms, or clinical covariates can alter model behavior.

Explainability is a central governance requirement. Regulators increasingly expect that computational tools used to inform drug development or clinical decisions provide comprehensible justifications. Graph neural networks can support interpretability through attention mechanisms, perturbation-based attribution, and concept-level explanations tied to network neighborhoods. However, these methods must be validated for biological plausibility and not merely for visual appeal. The deployment pipeline should include human review stages in which computational predictions are examined alongside experimental evidence and mechanistic knowledge.

The high-performance medicine paradigm [17] envisions the convergence of large-scale data, machine learning, and clinical decision support. In this vision, graph neural networks could contribute to individualized therapy selection by placing a patient's molecular profile within a network of disease mechanisms and therapeutic options. Yet clinical adoption remains slow because of regulatory uncertainty, workflow integration challenges, and the need for prospective validation. A model that predicts a drug combination as synergistic must be embedded in a decision process that accounts for patient preferences, comorbidities, drug-drug interactions, and economic constraints. Isolated predictive accuracy is not sufficient for clinical utility.

Governance also encompasses model versioning, auditability, and accountability. A graph-based model may influence decisions that affect patient safety, and there must be clear lines of responsibility for model errors. Model cards, data sheets, and algorithm registries can improve transparency, but they require institutional commitment. In the pharmaceutical industry, intellectual property concerns may limit the disclosure of proprietary models and datasets. Academic and regulatory stakeholders must negotiate a balance between protecting commercial innovation and ensuring independent validation and public trust.

8. Fairness, Robustness, and Sustainability

Fairness in graph-based drug-target and synergy prediction is an emerging concern. Training data are often biased toward well-studied targets, common diseases, and populations represented in genomic databases. As a result, models may perform poorly for rare diseases, under-represented populations, or novel therapeutic modalities. Graph neural networks can propagate biases through the network, reinforcing existing disparities in research attention. Fairness audits should evaluate model performance across target classes, disease areas, and demographic strata where relevant. Bias mitigation may require reweighting edges, augmenting underrepresented graph regions, or designing objectives that explicitly penalize performance disparities.

Robustness is equally important. Graph neural networks are vulnerable to adversarial perturbations in node features and graph structure. In molecular graphs, small changes in atom types or bond connectivity can drastically alter predictions. In biomedical knowledge

graphs, missing or erroneous edges can propagate to neighboring nodes and corrupt learned representations. Robustness testing should include systematic perturbation of input graphs, noisy label simulation, and distribution shift scenarios. Calibration of uncertainty estimates is also necessary for high-stakes decisions. A model that is overconfident in a synergy prediction may lead to costly experimental failures or unsafe clinical hypotheses.

Sustainability considerations are becoming more prominent in machine learning research. The computational cost of training large graph models, tuning hyperparameters, and maintaining data infrastructure has environmental and institutional implications. Energy consumption and hardware refresh cycles contribute to the carbon footprint of biomedical machine learning, as has been documented in natural language processing [18]. While graph models for drug discovery are often smaller than large language models, their repeated retraining, ensemble construction, and deployment across institutions can accumulate substantial energy demands. Sustainable practices include model distillation, efficient graph sampling, shared benchmark infrastructure, and transparent reporting of computational costs.

The long-term sustainability of graph-based drug discovery also depends on community governance of data and models. Open benchmarks, shared evaluation protocols, and curated negative sets can reduce redundant computation and improve comparability. However, open sharing must be balanced against the need to protect patient data and proprietary chemical matter. A federated infrastructure may allow institutions to jointly train models without centralizing sensitive data, but it introduces challenges in graph alignment and communication efficiency. Sustainable progress therefore requires both technical innovation and durable institutional arrangements.

9. Policy Implications and Future Research Directions

Policy decisions shape the development and adoption of graph neural network methods in drug discovery. Public funding agencies can support the creation of high-quality, interoperable biomedical network data and open benchmarks that include explicit generalization challenges. Regulatory agencies can provide guidance on how computational models should be validated for drug-target interaction prediction and combination therapy selection. Clarifying the evidentiary status of model-generated hypotheses would help pharmaceutical sponsors decide when computational predictions can reduce the burden of experimental testing without compromising safety.

Intellectual property frameworks also influence the trajectory of the field. Graph-based models and their training data may be subject to competing proprietary claims. Patenting computational methods remains complex, and the use of public molecular databases to train commercial models raises questions about benefit sharing. Academic and industry stakeholders should develop licensing norms that encourage innovation while preserving access to foundational data resources. Without such norms, the field may become fragmented, and independent reproducibility may suffer.

Future research should address the integration of multimodal data, including single-cell genomics, spatial transcriptomics, proteomics, and electronic health records. Graph neural networks are well suited to model these modalities as interconnected relational systems, but constructing stable, interpretable multimodal graphs remains difficult. The development of causal graph learning methods may help distinguish correlative network associations from causal perturbations, which is essential for therapeutic decision-making. Research on graph foundation models for molecular and biomedical networks is also advancing rapidly, but these

models must be evaluated for their ability to generalize beyond the training distribution and to support safety-critical decisions.

The pharmaceutical productivity crisis [20] underscores the need for computational methods that can meaningfully improve attrition rates. Graph neural networks alone will not resolve this crisis, but they can contribute when embedded in a broader system of evidence generation, experimental validation, and regulatory learning. The integration of predictive models with high-throughput screening, organoid systems, and adaptive trial designs may create feedback loops that continuously improve model performance and therapeutic relevance. Achieving this integration requires not only algorithmic advances but also institutional alignment across academic, industrial, and regulatory communities.

10. Conclusion

Graph neural networks provide a flexible and powerful framework for modeling the relational complexity of drug-target interactions and synergistic therapeutic combinations. By representing drugs, targets, pathways, and diseases as interconnected graphs, these methods can learn from molecular structure and biomedical network context in ways that traditional feature-based models cannot. The architectural diversity of graph neural networks, including convolutional, attention-based, and heterogeneous relational designs, allows researchers to tailor models to different prediction tasks and data environments. However, the performance of these models is inseparable from the quality of the underlying data, the appropriateness of graph construction choices, and the rigor of validation protocols.

From a systems perspective, the successful deployment of graph-based drug-target and synergy prediction requires attention to data infrastructure, standardization, governance, fairness, robustness, and sustainability. Models must be interpreted in light of biological mechanism, clinical context, and regulatory expectations. The enthusiasm for computational prediction should be tempered by an awareness of the inherent uncertainty in biomedical networks and the biases that may be encoded in training data. At the same time, the potential for graph neural networks to contribute to more rational combination therapy design is substantial, particularly when integrated with experimental and clinical workflows.

Future progress will depend on the development of open, interoperable data resources, reproducible evaluation frameworks, and multidisciplinary collaborations that bridge machine learning, systems biology, clinical pharmacology, and regulatory science. Graph neural networks are not a substitute for biological knowledge or clinical judgment, but they can serve as valuable tools for generating hypotheses, prioritizing experiments, and informing the design of therapeutic strategies. A responsible systems-level approach will treat these models as components within a larger socio-technical infrastructure for improving human health.

References

1. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*.
2. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. *International Conference on Learning Representations*.
3. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 1024–1034.

4. Öztürk, H., Özgür, A., & Ozkirimli, E. (2018). DeepDTA: Deep drug-target binding affinity prediction. *Bioinformatics*, 34(17), i821–i829.
5. Nguyen, T., Le, H., Quinn, T. P., Nguyen, T., Le, T. D., & Venkatesh, S. (2021). GraphDTA: Predicting drug-target binding affinity with graph neural networks. *Bioinformatics*, 37(8), 1140–1147.
6. Tsubaki, M., Tomii, K., & Sese, J. (2019). Compound-protein interaction prediction with end-to-end learning of neural networks for graphs and sequences. *Bioinformatics*, 35(2), 309–318.
7. Huang, K., Xiao, C., Glass, L. M., & Sun, J. (2021). MolTrans: Molecular interaction transformer for drug-target interaction prediction. *Bioinformatics*, 37(6), 830–836.
8. Preuer, K., Lewis, R. P. I., Hochreiter, S., Bender, A., Bulusu, K. C., & Klambauer, G. (2018). DeepSynergy: Predicting anti-cancer drug synergy with deep learning. *Bioinformatics*, 34(9), 1538–1546.
9. Ianevski, A., He, L., Aittokallio, T., & Tang, J. (2017). SynergyFinder: A web application for analyzing drug combination dose-response matrix data. *Bioinformatics*, 33(15), 2413–2415.
10. Zitnik, M., Agrawal, M., & Leskovec, J. (2018). Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*, 34(13), i457–i466.
11. Wang, Y., Yao, J., Sui, Y., Jiang, H., Ma, B., Lai, S., ... & Tan, N. (2026). HerbSyner_Finder: a network community-based model for identifying synergistic combinations from herbal medicines and complex systems. *Targetome*, 2(2).
12. Barabási, A. L., Gulbahce, N., & Loscalzo, J. (2011). Network medicine: A network-based approach to human disease. *Nature Reviews Genetics*, 12(1), 56–68.
13. Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., ... & Wilson, M. (2018). DrugBank 5.0: A major update to the DrugBank database for 2018. *Nucleic Acids Research*, 46(D1), D1074–D1082.
14. Tang, J., Karhinen, L., Xu, T., Szwajda, A., Yadav, B., Wennerberg, K., & Aittokallio, T. (2014). Target inhibition networks: Predicting selective combinations of druggable targets to block cancer survival pathways. *PLoS Computational Biology*, 10(9), e1003826.
15. Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., ... & Zhao, S. (2019). Applications of machine learning in drug discovery and development. *Nature Reviews Drug Discovery*, 18(6), 463–477.
16. Sun, M., Zhao, S., Gilvary, C., Elemento, O., Zhou, J., & Wang, F. (2020). Graph convolutional networks for computational drug development and discovery. *Briefings in Bioinformatics*, 21(3), 919–935.
17. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.
18. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.

19. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650.
20. Paul, S. M., Mytelka, D. S., Dunwiddie, C. T., Persinger, C. C., Munos, B. H., Lindborg, S. R., & Schacht, A. L. (2010). How to improve R&D productivity: The pharmaceutical industry's grand challenge. *Nature Reviews Drug Discovery*, 9(3), 203–214.