

Graph Neural Network Analysis of Gene Regulatory Networks in Cancer-Associated Transcriptional Dysregulation

Ruike Yin

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.

ruikemail@oregonstate.edu

Pavel D. Lawrence

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.

pavellawrence247@ku.edu

Abstract

The study of gene regulatory networks (GRNs) in the context of cancer-associated transcriptional dysregulation presents a complex systems-level challenge that requires advanced computational modeling. Graph neural networks (GNNs) have emerged as a powerful class of architectures for learning representations from relational data, making them particularly suited for analyzing the intricate wiring diagrams of gene interactions. This paper examines the application of GNNs to GRN analysis, focusing on the structural and functional trade-offs inherent in modeling high-dimensional, noisy, and dynamic biological networks. We discuss architectural considerations such as message-passing schemes, attention mechanisms, and graph pooling, and evaluate their implications for model interpretability, scalability, and robustness. The paper further explores the deployment of GNN-based models in translational research, including considerations of data governance, algorithmic fairness, and the sustainability of computational infrastructure in resource-constrained settings. Through a critical synthesis of recent advances, we highlight how GNNs can reveal mechanistic insights into transcriptional dysregulation while simultaneously raising important questions about reproducibility, transparency, and ethical deployment. We conclude by outlining future directions for integrating multi-omics data, incorporating temporal dynamics, and building more robust and equitable analytical pipelines for cancer genomics.

Keywords

graph neural networks, gene regulatory networks, transcriptional dysregulation, cancer genomics, systems biology, interpretability, robustness, data governance.

1. Introduction

Cancer is a disease fundamentally driven by alterations in the regulatory machinery that controls gene expression. The transcriptional program of a cell is orchestrated by a complex network of transcription factors, cofactors, chromatin remodelers, and non-coding RNAs, collectively forming a gene regulatory network (GRN). In neoplastic transformation, these networks undergo profound rewiring, leading to aberrant expression of oncogenes, silencing of tumor suppressors, and widespread changes in cellular identity [1]. Understanding how these dysregulations emerge and propagate through the network is critical for identifying

therapeutic targets and developing biomarkers. However, the sheer dimensionality, sparsity, and non-linearity of GRN data challenge traditional statistical and machine learning approaches. Graph neural networks (GNNs) have risen to prominence as a family of deep learning models that naturally operate on graph-structured data [2,3]. By aggregating information from neighboring nodes in a learned manner, GNNs can capture higher-order dependencies and contextual features that are often missed by conventional methods. This paper provides a systems-level analysis of GNN-based approaches for GRN analysis in cancer, emphasizing architectural trade-offs, deployment challenges, and broader socio-technical implications.

2. Gene Regulatory Networks and Transcriptional Dysregulation in Cancer

A gene regulatory network can be represented as a directed or undirected graph in which nodes correspond to genes or their protein products, and edges denote regulatory interactions such as transcription factor binding, chromatin looping, or co-expression relationships. In cancer, these networks are subjected to extensive perturbations. Mutations in transcription factors, copy number aberrations, epigenetic modifications, and alterations in signaling pathways collectively rewire the network, often resulting in a state of transcriptional addiction [1,4]. For instance, the MYC oncogene, a master regulator of cell growth and metabolism, is frequently overexpressed in many cancers and induces widespread changes in the transcriptional landscape [5]. Recent studies have demonstrated that MYC can undergo phase separation to selectively modulate the transcriptome, adding another layer of regulatory complexity [15]. This example illustrates that GRN dysregulation is not merely a matter of individual gene misregulation but involves emergent properties of the network as a whole. Consequently, computational models that can capture these emergent properties are essential. Traditional approaches such as correlation-based networks, Bayesian networks, and ordinary differential equation models have provided foundational insights but are often limited by assumptions of linearity, stationarity, or prior knowledge of network structure [6]. GNNs offer a more flexible framework that learns directly from data, allowing the discovery of non-linear, context-dependent regulatory relationships.

3. Graph Neural Networks: Architectural Considerations for GRN Analysis

The core operation of a GNN is message passing, where each node aggregates feature information from its neighbors and updates its own representation through a learnable function [2]. Several architectural variants have been applied to GRN analysis, each with distinct trade-offs. Graph convolutional networks (GCNs) approximate spectral graph convolutions using a first-order neighborhood aggregation scheme, providing computational efficiency but limited ability to distinguish between different edge types [3]. Graph attention networks (GATs) introduce attention mechanisms that learn the relative importance of neighboring nodes, enabling the model to focus on biologically relevant interactions [7]. Message-passing neural networks (MPNNs) generalize the message-passing framework, allowing for flexible edge features and advanced aggregation functions [8]. For GRN analysis, the choice of architecture depends on the specific goal: node-level tasks such as predicting gene expression or classifying gene function; edge-level tasks such as inferring regulatory interactions; or graph-level tasks such as predicting cancer subtype or drug response. In each case, the model must contend with the high sparsity and noise of GRN data, where many edges are either false positives from indirect correlations or false negatives due to technical limitations. Moreover, the scale of the human genome—approximately 20,000 protein-coding genes—results in graphs with hundreds of millions of potential edges, necessitating efficient

sampling and mini-batching strategies [9]. Graph pooling techniques, such as hierarchical clustering or differential pooling, can reduce dimensionality and highlight community structure within the network, but they introduce additional hyperparameters that affect interpretability [10].

A key trade-off exists between model expressiveness and interpretability. Highly expressive GNNs with multiple layers and complex aggregation functions can capture intricate non-linear dependencies but become black boxes, making it difficult to attribute predictions to specific regulatory interactions. For biomedical applications, interpretability is crucial for generating testable hypotheses and gaining regulatory approval. Recent work has explored explainability methods for GNNs, including gradient-based attribution, attention weight analysis, and graph-subnetwork extraction [11]. However, these methods often provide only local explanations and may be sensitive to input perturbations. An alternative approach involves designing inherently interpretable GNNs, such as those that enforce sparsity or incorporate prior biological knowledge through graph constraints. For example, incorporating known transcription factor binding sites as edge features can guide the model toward biologically plausible interactions while reducing the search space [12]. The architectural decisions also affect model robustness. GNNs are known to be vulnerable to adversarial attacks that add or remove edges, which could be exploited if GRN data are contaminated by batch effects or measurement errors [13]. Robustness can be improved through regularization, adversarial training, or using more stable aggregation schemes such as mean pooling instead of sum pooling. However, these measures may reduce capacity to detect genuine but subtle regulatory signals.

4. System-Level Trade-offs: Scalability, Interpretability, and Robustness

Deploying GNNs for large-scale GRN analysis requires careful consideration of computational infrastructure. Training a GNN on a genome-wide graph with millions of nodes and edges demands substantial memory and processing power, often necessitating distributed computing frameworks or graph partitioning strategies [14]. Cloud-based solutions offer elasticity but raise concerns about data privacy and sovereignty, particularly when dealing with patient-derived genomic data. Federated learning has been proposed as a mechanism to train models across multiple institutions without centralizing data, but it introduces communication overhead and potential heterogeneity in data distributions [9]. From a governance perspective, the reliance on large-scale compute resources can exacerbate inequities between well-funded research centers and resource-limited institutions. Moreover, the carbon footprint of training deep learning models is a growing concern; sustainable AI practices, such as efficient model architectures, pruning, and use of renewable energy, should be integrated into the research pipeline.

Interpretability is not only a technical requirement but also a policy imperative. If GNN models are to inform clinical decision-making, healthcare providers and patients must trust the model's outputs. Regulatory frameworks such as the European Union's General Data Protection Regulation (GDPR) include provisions for explainability of automated decisions, which directly applies to AI-assisted diagnostics [16]. In the context of GRN analysis, this means that a model predicting that a particular gene is a driver of transcriptional dysregulation should be able to provide the regulatory evidence supporting that claim. Current explainability methods, while promising, do not yet meet the standards required for clinical deployment. There is a need for rigorous validation of explanation fidelity, stability, and relevance. Additionally, fairness considerations arise when GNNs are trained on genomic data

that predominantly come from populations of European ancestry. Models may learn regulatory patterns that are not generalizable to other ethnic groups, leading to biases in biomarker discovery and treatment recommendations [17]. Addressing this requires intentional inclusion of diverse data sources and the development of domain adaptation techniques that can transfer knowledge across populations.

Robustness extends beyond adversarial attacks to encompass the inherent variability in GRN data. Different experimental platforms (e.g., RNA-seq, ChIP-seq, ATAC-seq) produce data with distinct biases and noise levels. A GNN trained on one data modality may fail when applied to another. Multi-modal integration is an active area of research, where graph-based models can fuse information from various omics layers to produce more stable representations [18]. However, the lack of ground truth for regulatory interactions complicates model evaluation. Researchers often rely on known databases (e.g., Encode, Roadmap Epigenomics) for validation, but these databases are incomplete and biased toward well-studied genes. Therefore, model performance metrics must be interpreted cautiously, and efforts to create standardized benchmarks for GRN inference are essential.

5. Deployment, Governance, and Ethical Implications

Translating GNN-based GRN models from academic research to clinical or industrial settings involves multiple layers of governance. Data access agreements, institutional review boards, and patient consent processes must ensure that genomic data are used ethically and that results are communicated responsibly. The predictive capabilities of these models could lead to commercial applications such as drug target identification or patient stratification, which carry financial incentives that may conflict with open science principles [19]. There is a risk of creating proprietary algorithms that are not subject to peer review, potentially perpetuating hidden biases.

Moreover, the high cost of computational resources may concentrate GNN expertise in a small number of institutions, creating a digital divide. To democratize access, open-source implementations, pre-trained models, and cloud credits should be provided to under-resourced groups. The development of lightweight GNN architectures that can run on standard hardware also contributes to equity. From a sustainability perspective, the energy consumption of large model training can be reduced through techniques such as quantization, knowledge distillation, and early stopping. Research institutions and funding agencies should consider environmental impact statements for large-scale computational projects.

The ethical implications of using GNNs to predict gene regulation extend to the interpretation of results. A model that identifies a novel regulatory interaction may lead to preclinical studies that require animal testing or human trials. The decision to pursue such studies should be based not only on model confidence but also on biological plausibility and clinical urgency. Over-reliance on black-box predictions without mechanistic validation could result in wasted resources or harmful interventions. Therefore, GNN analyses should be embedded in a broader workflow that includes experimental validation, systematic review of existing literature, and collaboration with domain experts [20].

6. Future Directions and Sustainability

The future of GNN-based GRN analysis lies in integrating temporal dynamics and multi-omics data. Cancer is a dynamic disease, and GRNs evolve over time due to treatment pressures, immune interactions, and clonal evolution. Temporal GNNs that incorporate time-series gene expression data could capture the trajectory of dysregulation, offering insights into

therapeutic resistance [21]. Additionally, spatial transcriptomics and single-cell sequencing provide high-resolution snapshots of gene expression within tissue contexts. GNNs can be extended to handle such heterogeneous data by using heterogeneous graph representations that differentiate between cell types, spatial locations, and time points [22]. These advances will require new architectures that can scale to billions of nodes and edges while maintaining interpretability.

Sustainability also involves the long-term maintenance of computational tools and databases. Many GNN models are developed as standalone projects without ongoing support, leading to reproducibility crises. The establishment of community standards, such as the use of common data formats, benchmark datasets, and model registries, can mitigate this issue. Funding agencies should require that software and models be deposited in public repositories with clear documentation and versioning. Furthermore, the field must address the reproducibility of GNN results, which are sensitive to random seeds, hyperparameters, and platform differences. Systematic benchmarking and robust statistical testing are needed to ensure that reported findings are not artifacts [23].

7. Conclusion

Graph neural networks offer a principled framework for modeling the complex architecture of gene regulatory networks and understanding their dysregulation in cancer. The trade-offs between expressiveness, interpretability, robustness, and scalability demand careful architectural choices and a systems-level perspective. As these models move toward translational applications, considerations of data governance, algorithmic fairness, and computational sustainability become paramount. The integration of temporal and multi-modal data, along with community-driven reproducibility standards, will shape the next generation of GNN-based cancer genomics. Ultimately, the success of these approaches depends on a collaborative ecosystem that bridges computer science, biology, medicine, ethics, and policy.

References

1. Hanahan, D., & Weinberg, R. A. (2011). Hallmarks of cancer: The next generation. *Cell*, 144(5), 646-674.
2. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. (2017). Neural message passing for quantum chemistry. *Proceedings of the 34th International Conference on Machine Learning*, 70, 1263-1272.
3. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*.
4. Bradner, J. E., Hnisz, D., & Young, R. A. (2017). Transcriptional addiction in cancer. *Cell*, 168(4), 629-643.
5. Dang, C. V. (2012). MYC on the path to cancer. *Cell*, 149(1), 22-35.
6. Marbach, D., Costello, J. C., Küffner, R., Vega, N. M., Prill, R. J., Camacho, D. M., ... & Stolovitzky, G. (2012). Wisdom of crowds for robust gene network inference. *Nature Methods*, 9(8), 796-804.
7. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. *International Conference on Learning Representations*.

8. Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., ... & Pascanu, R. (2018). Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
9. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Stoyanov, D. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3(1), 119.
10. Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., & Leskovec, J. (2018). Hierarchical graph representation learning with differentiable pooling. *Advances in Neural Information Processing Systems*, 31.
11. Ying, Z., Bourgeois, D., You, J., Zitnik, M., & Leskovec, J. (2019). GNNExplainer: Generating explanations for graph neural networks. *Advances in Neural Information Processing Systems*, 32.
12. Zitnik, M., & Leskovec, J. (2017). Predicting multicellular function through multi-layer tissue networks. *Bioinformatics*, 33(14), i190-i198.
13. Dai, H., Li, H., Tian, T., Huang, X., Wang, L., Zhu, J., & Song, L. (2018). Adversarial attack on graph structured data. *Proceedings of the 35th International Conference on Machine Learning*, 80, 1115-1124.
14. Wang, M., Zheng, D., Ye, Z., Gan, Q., Li, M., Song, X., ... & Zhang, C. (2019). Deep graph library: A graph-centric, highly-performant package for graph neural networks. *arXiv preprint arXiv:1909.01315*.
15. Yang, J., Chung, C. I., Koach, J., Liu, H., Navalkar, A., He, H., ... & Shu, X. (2024). MYC phase separation selectively modulates the transcriptome. *Nature Structural & Molecular Biology*, 31(10), 1567-1579.
16. Goodman, B., & Flaxman, S. (2017). European Union regulations on algorithmic decision-making and a "right to explanation". *AI Magazine*, 38(3), 50-57.
17. Martin, A. R., Kanai, M., Kamatani, Y., Okada, Y., Neale, B. M., & Daly, M. J. (2019). Clinical use of current polygenic risk scores may exacerbate health disparities. *Nature Genetics*, 51(4), 584-591.
18. Ma, T., & Zhang, A. (2021). Integrate multi-omics data with graph neural networks: A review. *Briefings in Bioinformatics*, 22(6), bbab151.
19. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44-56.
20. Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. *JAMA*, 319(13), 1317-1318.
21. Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., & Bronstein, M. (2020). Temporal graph networks for deep learning on dynamic graphs. *arXiv preprint arXiv:2006.10637*.
22. Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., ... & Leskovec, J. (2020). Open graph benchmark: Datasets for machine learning on graphs. *Advances in Neural Information Processing Systems*, 33, 22118-22133.

23. Pineau, J., Vincent-Lamarre, P., Sinha, K., Larivière, V., Beygelzimer, A., d'Alché-Buc, F., ... & Larochelle, H. (2021). Improving reproducibility in machine learning research (a report from the NeurIPS 2019 reproducibility program). *Journal of Machine Learning Research*, 22(164), 1-20.