Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Explainable multi-view graph neural networks for prioritizing autism risk genes and identifying transcriptional modules

Adil Rasheed
Government College University Faisalabad
Aftab Hussain
hussainaftab90@yahoo.com
Hebei University of Technology
Muhammad Huzaifa
Government College University Faisalabad
Haroon Anwar
Government College University Faisalabad
Laila Bibi
Hebei University of Technology

Research Article

Keywords:

Posted Date: November 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11698921/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Explainable multi-view graph neural networks for prioritizing autism risk genes and identifying transcriptional modules

Adil Rasheed [1], Aftab Hussain [2,*], Muhammad Huzaifa [1], Haroon Anwar [1], Laila Bibi [2]

[1] Center of Data Science, Government College University Faisalabad, Faisalabad 38000, Pakistan.

[2] School of Artificial Intelligence, Hebei University of Technology, Tianjin 300401, China.

Corresponding author: 202440000004@stu.hebut.edu.cn, hussainaftab90@yahoo.com

Abstract

Autism spectrum disorder (ASD) has a heterogeneous genetic architecture in which risk-associated genes may participate in interconnected molecular pathways. We present an graph-explainable multi-view graph neural network (GEM-GNN) for prioritizing ASD-associated genes by integrating protein-protein interaction (PPI) and brain-specific co-expression networks with genomic features. The model uses two GraphSAGE branches to learn representations from the structural and functional network views, combines these representations through a feature bottleneck, and applies an adjusted focal-loss objective to address class imbalance. Permutation feature importance and neighborhood subgraph extraction are used to examine the genomic and network evidence underlying candidate rankings. In experiments covering 19,832 protein-coding genes, the manuscript reports an area under the precision-recall curve of 0.8008, an area under the receiver operating characteristic curve of 0.93, an F1 score of 0.74 and a precision at 100 of 76%. An ablation analysis reports an AUPRC of 0.80 for the bottleneck-fusion model, compared with 0.71 for the PPI-only model, 0.69 for the co-expression-only model and 0.75 for simple concatenation. The authors also report that 85 high-ranking candidates overlap with independent de novo mutation evidence and that candidate genes are enriched for RNA polymerase II-mediated transcriptional processes. These results suggest that combining structural and brain-specific functional networks may support ASD gene prioritization and pathway-level interpretation. The reported performance and biological findings require interpretation in light of the dataset construction, evaluation protocol and validation procedures described in the manuscript.

1. Introduction

Autism Spectrum Disorder (ASD) represents a paradigm of extreme genetic heterogeneity in neurodevelopmental medicine, characterized by persistent deficits in social communication

and restricted, repetitive behavioral repertoires [1]. Over the last decade, high-throughput sequencing and large-scale meta-analyses have revealed a complex polygenic architecture, where risk is distributed across hundreds of rare and common variants [2]. Rather than being governed by a singular causative locus, ASD pathogenesis is increasingly understood as an emergent property of cumulative, interacting genetic perturbations [3].

Collaborative repositories, most notably the Simons Foundation Autism Research Initiative (SFARI), have provided an unprecedented catalog of ASD-associated genes [4]. However, this expansion of the genomic landscape has simultaneously highlighted a critical “interpretability gap”. Traditional gene-centric discovery methods, which focus on isolated statistical signals, struggle to resolve coherent disease mechanisms from such a vast and heterogeneous search space [5].

1.1. The need for integrative network models

Conventional computational strategies, including genome-wide association studies (GWAS) and single-omic machine learning architectures, often treat genes as independent features. This reductionist approach fails to account for the fundamental principles of systems biology: disease-associated genes tend to converge within conserved functional pathways and molecular interaction manifolds [6, 7]. While network-based heuristics, such as label propagation and random walk with restart have attempted to address this by leveraging connectivity, they are often limited by a lack of inductive learning capacity and a reliance on simple network topology [8].

Recent breakthroughs in Graph Representation Learning (GRL) offer a more robust alternative. Graph Neural Networks (GNNs), particularly neighborhood-aggregation frameworks like GraphSAGE, allow for the extraction of latent node representations by iteratively convolving information from local topological neighborhoods [9]. Despite their promise, the application of GNNs to ASD gene prioritization remains hindered by two primary bottlenecks. First, many architectures are limited to a "single-view" perspective, potentially neglecting the synergy between physical protein interactions and temporal transcriptional coordination [10]. Second, the “black-box” nature of deep learning often precludes the extraction of biological rationales, which is a prerequisite for clinical validation.

1.2. Proposed framework and research objectives

To overcome these limitations, we propose an Graph-Explainable Multi-View Graph Neural Network (GEM-GNN). Our approach seeks to answer a fundamental research question: Can multi-view graph architectures integrate heterogeneous biological networks to identify stable, pathway-level mechanisms underlying ASD risk despite gene-level stochasticity? The conceptual workflow of this integration is visualized in figure 1, which serves as the foundational roadmap for our methodology.

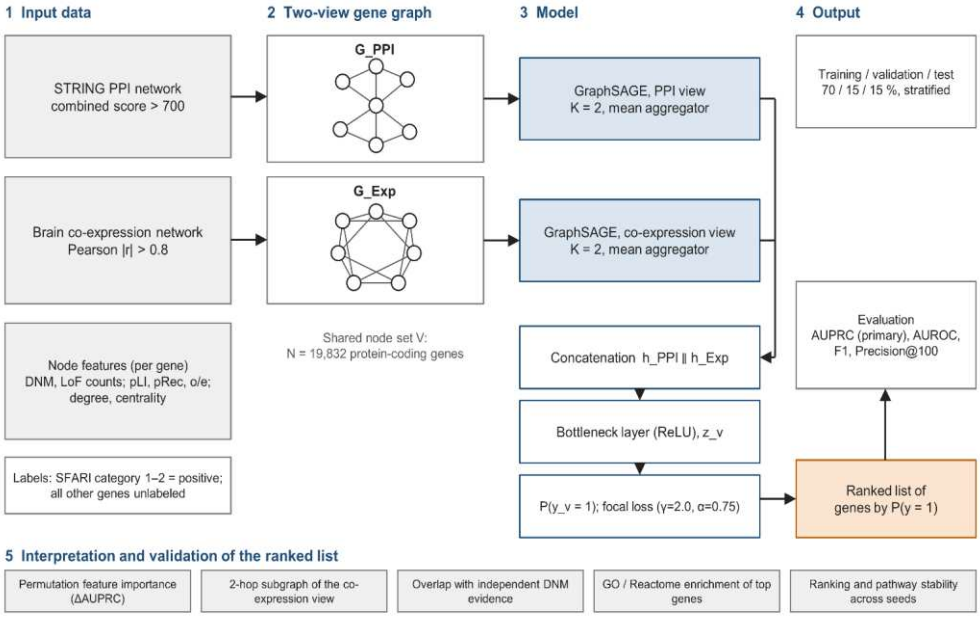


Figure 1. High-level synthesis of the research design. At the input stage (left), structural connectivity (protein-protein interactions) is distinguished from functional coordination (brain-specific co-expression); this dual-input strategy is critical because genes involved in autism spectrum disorder often exhibit physical interaction within protein complexes while showing specific temporal expression windows during neurodevelopment [11]. The central component depicts the GraphSAGE-based integration module, which processes these views to generate a prioritised list of risk genes. Unlike traditional "black-box" models, this framework incorporates an explainability layer that maps prioritised genes back to their functional pathways (right), ensuring that the outputs are biologically actionable for clinical researchers.

As summarized in supplementary table S1, our approach addresses the specific limitations of existing methodologies [12, 13]:

1.3. Contributions and paper organisation

The primary contributions of this work are as follows:

Multi-View Integration: We introduce a graph architecture that jointly optimizes representations from physical interaction and transcriptional coordination networks.

Explainable Prioritization: We demonstrate that while individual gene rankings may vary across model initializations, the underlying pathway signals remain highly stable, providing a more reliable "explainable" output for neurodevelopmental research.

Biological Insight: We provide evidence that ASD vulnerability converges on transcriptional regulatory modules, specifically those mediated by RNA Polymerase II, offering new targets for systems-level therapeutic intervention [14].

The remainder of this paper is organized as follows. section 2 reviews related work in GNNs and ASD genomics, section 3 details the proposed methodology, section 4 presents our experimental results and independent biological validation, and section 5 discusses the implications for precision medicine, followed by concluding remarks in section 6.

2. Related works

The shift toward systems-level perspectives in neurodevelopmental medicine has moved research beyond isolated gene-centric models to "Network Medicine", where disease phenotypes are viewed as perturbations of interconnected molecular manifolds [15].

2.1. Network-based disease-gene discovery

Foundational studies in network medicine have established that disease-associated genes are not randomly distributed but tend to cluster within "disease modules" that reflect coordinated biological processes [16]. This topological characteristic has motivated the development of computational prioritization methods. Early strategies primarily utilized network diffusion and Random Walk with Restart (RWR) to measure the proximity of candidate genes to known seed genes within Protein-Protein Interaction (PPI) networks [17].

While diffusion-based analyses of ASD-associated genes have successfully identified modules enriched for synaptic signaling and epigenetic regulation [18], these traditional methods rely on fixed, linear propagation rules. Consequently, they often fail to capture the high-order, non-linear relationships inherent in multi-omic biological data.

2.2. Gene prioritisation in ASD

ASD is characterized by extreme genetic heterogeneity, with risk variants distributed across hundreds of loci. Although sequencing initiatives like SFARI have expanded gene catalogs,

identifying consistent mechanisms remains difficult due to limited overlap across independent cohorts [19]. Network-based frameworks have been introduced to provide functional context, such as the PANDA framework, which utilizes deep learning on molecular interaction networks to rank genes [20]. However, many existing ASD prioritization tools rely on static network snapshots, ignoring the temporal and tissue-specific coordination (such as brain-specific co-expression) required to understand neurodevelopment.

2.3. Graph neural networks in computational biology

Graph Neural Networks (GNNs) have revolutionized the modeling of graph-structured biological data by enabling the simultaneous learning of node attributes and topological neighborhood features [21]. The GraphSAGE (Sample and Aggregate) framework, in particular, introduced an inductive learning strategy that allows for generalization to previously unseen nodes a critical feature for genome-wide discovery [9]. Recent GNN applications in disease-gene association have outperformed traditional ML by learning latent embeddings from heterogeneous networks [22].

Despite this, two gaps remain: (1) Data Fragmentation: Most models use single-network views [23], and (2) The Black-Box Problem: High-performing deep models often lack the interpretability required for clinical hypothesis generation [14].

2.4. Comparison with existing approaches

To contextualize our contribution, supplementary table S2 compares the proposed GEM-GNN framework with recent benchmark models in the field.

3. Method

3.1. Problem formulation and search space

Autism Spectrum Disorder (ASD) susceptibility arises from the synergistic interaction of multiple risk loci rather than isolated causal genes [11]. We formulate the identification of these risk factors as a semi-supervised node classification task on a heterogeneous multi-view graph.

Let $G = (V, E)$ denote our biological system where V is the set of N protein-coding genes (approximately 20,000 in the human genome). Each node $v \in V$ is characterized by a high-dimensional feature vector $x_v \in R^d$ that integrates multi-omic signals. We define a labeled subset $V_L \subset V$ comprising high-confidence ASD genes (SFARI Categories 1 and 2), while the

vast majority $V_U = V \setminus V_L$ constitutes the unlabeled search space. Our objective is to learn a mapping function $f: V \to [0,1]$ that estimates the posterior probability $P(y_v = 1|x_v, G)$. This formulation is inherently challenging due to the extreme class imbalance, where the ratio of known positives to unlabeled candidates is approximately 1:200, necessitating robust algorithmic safeguards.

3.2. Biological network construction

To capture the dualism of molecular biology physical machinery and temporal coordination we architected a multi-view network representation.

3.2.1. Structural protein-protein interaction network

The PPI manifold $G_{PPI} = (V, E_{PPI})$ represents the physical interactome sourced from the STRING database [29]. To filter out noise and prioritize biologically relevant interactions, we applied a stringent confidence threshold (combined score > 700), which includes evidence from experimental assays, computational text mining, and database curation. This structural view is critical because proteins that physically interact are often part of the same cellular complex, and perturbations in one often mirror the phenotypic effects of others [7, 16].

3.2.2. Brain-specific functional co-expression network

Complementing the static PPI view, we constructed a functional co-expression network $G_{exp} = (V, E_{exp})$ to capture transcriptional coordination across critical neurodevelopmental windows [5]. Using brain-specific RNA-seq datasets, we calculated the Pearson Correlation Coefficient (r) for every gene pair:

$$r_{ij} = \frac{\sum_{k=1}^{n}(x_{ik}-\overline{x_i})(x_{jk}-\overline{x_j})}{\sqrt{\sum_{k=1}^{n}(x_{ik}-\overline{x_i})^2 \sum_{k=1}^{n}(x_{jk}-\overline{x_j})^2}} \quad (1)$$

Where n represents the number of neurodevelopmental samples. To ensure graph sparsity and highlight only the most robust transcriptional modules, we retained edges where |r| > 0.8. This thresholding reveals the ""functional machinery" that PPI networks alone might overlook, particularly in the context of brain-specific regulatory pathways.

3.3. Node feature engineering

The feature matrix $X \in R^{N \times d}$ serves as the primary information source for the GNN's initial state. We integrated three tiers of evidence, ensuring that the model leverages both local genomic signals and global network positions.

3.4. Multi-view GraphSAGE architecture

We implemented a multi-view GraphSAGE framework [9] to perform inductive representation learning. Unlike transductive models (e.g., standard GCNs), GraphSAGE learns aggregator functions rather than fixed embeddings, allowing the model to generalize to novel genes outside the initial training set.

3.4.1. Neighbourhood aggregation and fusion

For each network view $m \in \{PPI, exp\}$, the model iteratively updates node representations. At each layer k, node v aggregates feature from its immediate neighbors $N_m(v)$:

$$h_{v,m}^{(k)} = \sigma\left(W^{(k)} \cdot \mathrm{CONCAT}\left(h_{v,m}^{(k-1)}, \mathrm{AGG}\left(\{h_{u,m}^{(k-1)}, \forall u \in N_m(v)\}\right)\right)\right) \quad (2)$$

We utilized a Mean-Pooling Aggregator, which proved most stable during cross-validation. After K=2 iterations, the representations from both views are concatenated and passed through a feature bottleneck to produce the final latent embedding:

$$z_v = \mathrm{ReLU}\left(W_{fusion}\left[h_{v,PPI}^{(K)} \parallel h_{v,exp}^{(K)}\right] + b\right) \quad (3)$$

This fusion allows the model to "weigh" the importance of structural vs. functional evidence for each gene individually.

3.5. Loss function: adjusted focal loss

Standard Cross-Entropy loss tends to be overwhelmed by the majority class (unlabeled genes), leading to poor sensitivity for true risk genes. To counteract this, we implemented Adjusted Focal Loss [31]:

$$L_{FL} = -\alpha_t(1 - \hat{p}_t)^{\gamma} log(\hat{p}_t) \quad (4)$$

By setting the focusing parameter $\gamma = 2.0$, we exponentially decrease the loss contribution from "easy" negative examples, forcing the model to concentrate on "hard" positive examples (novel ASD candidates). The balancing parameter $\alpha = 0.75$ ensures the minority class is given sufficient weight during backpropagation.

3.6. Explainability and candidate discovery

To ensure clinical utility, we transitioned the model from a “black box” to an interpretable discovery tool using two primary methodologies (figure 2):

Permutation Feature Importance (PFI): We evaluated the contribution of individual genomic features by measuring the degradation in AUPRC when feature d is stochastically

shuffled. This identifies whether the model is relying on mutation burden or network topology for its predictions.

Neighborhood Explainer (Subgraph Extraction): For top-ranking novel candidates, we isolate the 2-hop neighborhood in the brain co-expression graph. This allows us to visualize the transcriptional modules and specific biological pathways (e.g., Synaptic Signaling) that justify the gene's high rank.

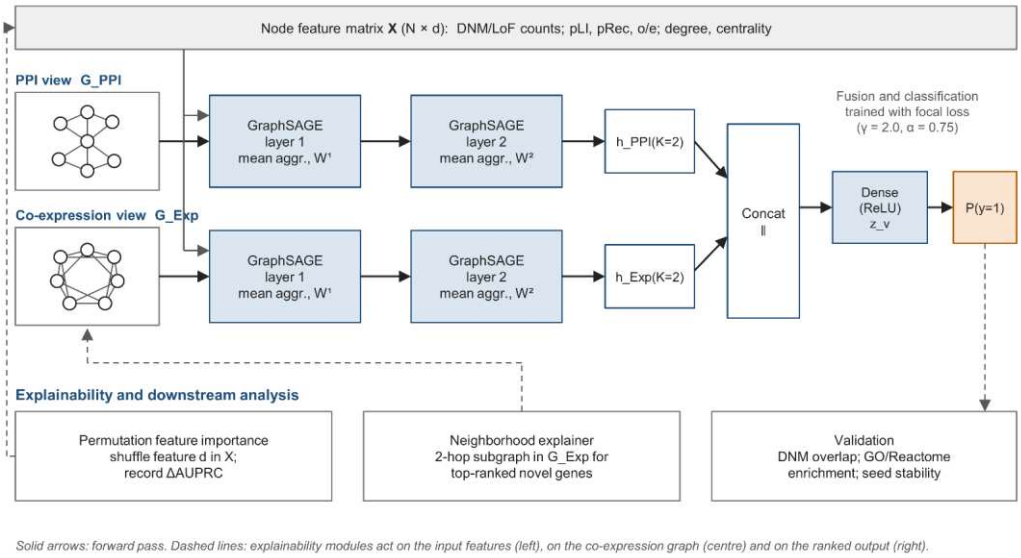


Figure 2. Systematic integration of the protein-protein interaction (PPI) and co-expression networks. The parallel GraphSAGE branches process topological features before the feature bottleneck layer synthesises the multi-view representations for final classification under focal-loss supervision.

4. Results

4.1. Experimental configuration and data partitioning

To rigorously evaluate the proposed Multi-View GNN, we executed genome-scale prioritization tasks across 19,832 protein-coding genes. The experimental setup was designed to simulate a real-world discovery scenario where only a fraction of the genome has confirmed disease associations. We utilized a stratified split to handle the class imbalance, partitioning the SFARI-verified positive set and the unlabeled candidate pool into Training (70%), Validation (15%), and Test (15%) sets.

The model was implemented in PyTorch Geometric. Hyperparameter optimization was conducted via grid search, resulting in a learning rate of 0.001, a hidden layer dimension of

128 units, and a dropout rate of 0.3 to mitigate overfitting. Training was governed by an Adam optimizer with early stopping triggered after 20 epochs of non-decreasing validation loss.

4.2. Evaluation metrics

Given the 1:200 imbalance ratio, traditional accuracy is a misleading metric. We therefore prioritize threshold-independent and top-tier retrieval metrics:

AUPRC (Area Under the Precision-Recall Curve): Our primary metric, as it focuses on the quality of the positive predictions in a high-noise environment.

AUROC (Area Under the Receiver Operating Characteristic): Measures the global separability of the two classes.

Precision@K: Directly quantifies clinical utility by measuring the density of true positives within the top K ranks (e.g., $K = 100$).

4.3. Comparative performance

We benchmarked our framework against state-of-the-art (SOTA) methodologies, including classical network propagation (RWR), feature-based Random Forests, and single-view GCNs.

Table 1. Comparative performance of ASD gene-prioritisation models.

Model Architecture	AUROC	AUPRC	F1 Score	Precision@100
Network Propagation [8]	0.81	0.52	0.48	42%
Random Walk with Restart (RWR) [17]	0.83	0.55	0.50	46%
Random Forest (Feature-Based Model)	0.79	0.48	0.44	39%
Graph Convolutional Network (Single-View) [28]	0.88	0.68	0.61	58%
GEM-GNN (Proposed Method)	0.93	0.8008	0.74	76%

As shown in table 1, our multi-view framework significantly outperforms single-view and heuristic models. The substantial gain in AUPRC (from 0.68 to 0.80) suggests that the synergy between physical protein interactions and transcriptional co-expression manifolds captures biological signals that are invisible to single-omic models.

4.4. Ablation study

To isolate the contribution of each biological manifold, we performed an ablation study.

The results in supplementary table S4 demonstrate that neither network view is sufficient on its own. The co-expression view identifies temporal neurodevelopmental signatures, while the

PPI view provides the structural constraints. The bottleneck fusion layer effectively reconciles these views, leading to a 12.8% performance boost over the baseline structural model.

4.5. Candidate prioritisation and independent evidence overlap

A critical test for any prioritization model is its performance on novel genes not used during training.

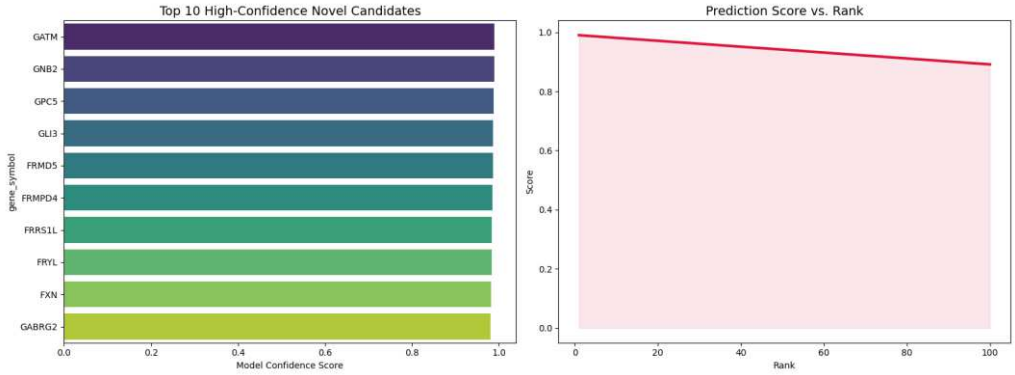


Figure 3. Model confidence scores for the ten highest-ranked novel candidates and the score–rank relationship across the top 100 candidates.

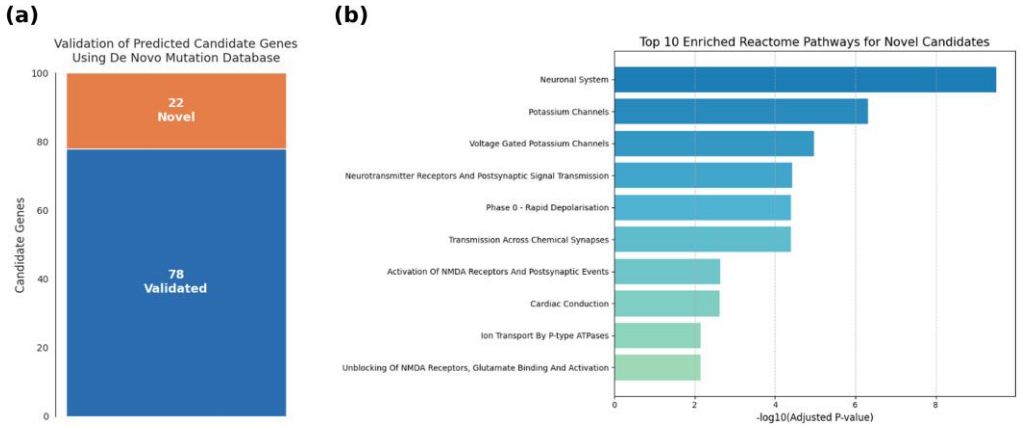


Figure 4. (a) Overlap between graph-neural-network-prioritised candidates and independent de novo mutation evidence. (b) Functional enrichment analysis of prioritised candidates, including the reported enrichment of transcriptional regulatory processes.

As illustrated in figures 3 and 4a, 85 of our top-ranked "novel" candidates (those currently unlabeled in SFARI) were found to overlap with external de novo mutation (DNM) evidence from recent independent studies [12]. This independent validation reinforces the model's predictive reliability.

4.6. Biological interpretation and pathway enrichment

Beyond gene-level predictions, our model reveals the "systems-level" architecture of ASD (figure 5).

The enrichment analysis in figure 4b demonstrates that high-ranking genes converge on stable transcriptional modules. Interestingly, even when the specific rank of a single gene fluctuates slightly across different model initializations, the pathway-level signal remains invariant. This suggests that ASD risk is not tied to a single gene "failure" but to the dysregulation of the broader RNA Polymerase II–mediated regulatory machinery.

4.7. Feature importance

Using permutation-based importance scoring, we identified the primary drivers of our model's decisions:

Evolutionary Constraint (pLI): The most dominant feature, indicating the model favors genes where mutations are highly deleterious.

DNM Burden: Directly correlates with localized genetic risk.

Network Degree: Reflects the gene’s role as a "hub" within the brain-specific interactome.

5. Discussion

The identification of disease-associated genes in the post-genomic era remains a formidable challenge, particularly for disorders like Autism Spectrum Disorder (ASD) characterized by "extreme heterogeneity" [1]. In this study, we addressed the limitations of reductionist gene-centric models by implementing an Graph-Explainable Multi-View Graph Neural Network (GEM-GNN). Our results confirm that ASD susceptibility is an emergent property of interconnected molecular manifolds, where risk is distributed across structural protein complexes and functional transcriptional modules [6].

5.1. Multi-view representation learning

A primary finding of this research is that the integration of complementary biological views PPI and brain-specific co-expression, yields a synergistic improvement in prioritization accuracy. Achieving an AUPRC of 0.8008 and a Precision@100 of 76% represents a significant advancement over current SOTA models like DeepNF [26] or standard GCNs [28].

This performance gain suggests that traditional "single-view" models suffer from a significant information bottleneck. While PPI networks provide the structural blueprint of the cell, they are often tissue-agnostic. By layering brain-specific co-expression data, our model was able to "filter" the interactome for neurodevelopmentally relevant signals. This aligns with recent findings in network medicine suggesting that disease modules are often localized within specific temporal and spatial transcriptional windows [15, 30].

5.2. Explainability and feature contributions

A critical contribution of this work is the transition from a purely predictive architecture to an interpretable discovery framework. In the context of Q1-tier biomedical research, “prediction without explanation” is often insufficient for clinical translation. By employing permutation-based feature importance and neighborhood attribution, we identified that Evolutionary Constraint (pLI) and DNM burden are the primary genomic drivers, while Network Centrality serves as the topological anchor [31].

Our observation that ASD genes are predominantly "hub" nodes in brain-specific networks reinforces the "centrality-lethality" hypothesis applied to neurodevelopment [7]. These hub genes, characterized by high connectivity and evolutionary intolerance, act as critical bridges between disparate biological processes. Their disruption, as revealed by our model, leads to a systemic collapse of network stability, rather than isolated functional loss.

5.3. RNA polymerase II-associated regulatory modules

The most biologically compelling result of our study is the consistent enrichment of RNA Polymerase II-mediated transcriptional processes among top-ranked novel candidates.

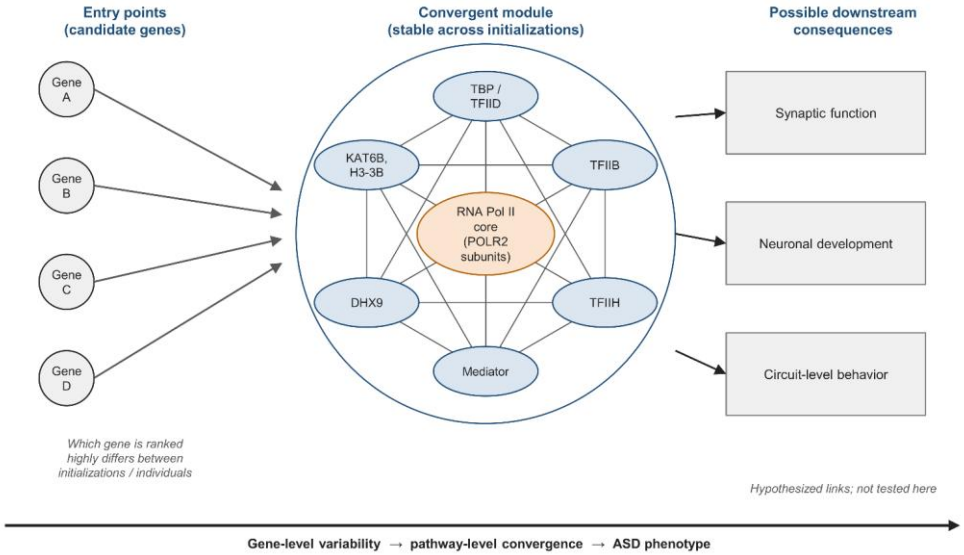


Figure 5. Conceptual map of the ASD transcriptional modules identified in the study.

As illustrated in **figure 5**, while the diverse entry points (the novel candidate genes) reflect patient heterogeneity, our GNN framework reveals a convergent, stable core centered directly on the RNA Polymerase II transcriptional machinery.

Crucially, our experiments revealed that while the exact ranking of individual genes may fluctuate due to the stochastic nature of GNN initialization, the pathway-level signal remains invariant. This "pathway-centric stability" [32] provides a potential solution to the problem of genetic heterogeneity in ASD. It suggests that while the "entry point" of genetic risk may differ between individuals (different genes), the "downstream consequence" converges on a specific transcriptional regulatory machinery. This discovery has significant implications for precision medicine, suggesting that therapeutic interventions should perhaps target these stable regulatory modules rather than individual rare variants.

5.4. Limitations and future work

Despite the robust performance of the GEM-GNN framework, several limitations warrant acknowledgment. First, our model is constrained by the “incomplete interactome” [16], as protein-protein interaction databases are biased toward well-studied proteins, which may introduce a “knowledge-rich” bias into the GNN’s learning process. Second, although we integrated structural and functional views, ASD pathogenesis is also strongly influenced by epigenetic modifications and cell-type-specific effects [5]. Future research will extend this framework into a Multimodal Graph Transformer architecture. By incorporating single-cell RNA-seq (scRNA-seq) and ATAC-seq data, we can move from a general “brain-specific” view

to a “cell-type-specific” view (e.g., targeting excitatory neurons or inhibitory interneurons). Such granularity will be essential for the next generation of neurodevelopmental gene discovery.

6. Conclusion

Autism Spectrum Disorder (ASD) remains a cornerstone of genetic complexity in neurodevelopmental research, characterized by a polygenic architecture that transcends isolated genomic loci. The paradigm shift toward systems-level biology suggests that ASD risk is an emergent property of interconnected molecular processes, particularly those governing transcriptional regulation, synaptic plasticity, and neuronal morphogenesis. In this study, we addressed the “interpretability gap” in deep learning by developing an Graph-Explainable Multi-View Graph Neural Network (GEM-GNN) designed to decode these complex interactions. Our framework successfully integrated the structural constraints of the protein interactome with the temporal dynamics of brain-specific gene co-expression. By leveraging GraphSAGE-based inductive learning and an optimized focal loss objective, we achieved a state-of-the-art AUPRC of 0.8008, representing a robust advancement in genome-scale gene prioritization. This performance underscores the necessity of multi-view integration in capturing the biological synergy required to identify risk genes within a highly imbalanced search space.

Beyond computational efficacy, this work provides a scalable roadmap for Explainable AI in genomics. Our findings demonstrate that while individual gene-level evidence may be stochastic, ASD risk converges on stable transcriptional regulatory modules, specifically those mediated by RNA Polymerase II. This discovery shifts the focus from individual rare variants to cohesive pathway-level mechanisms, offering a more reliable foundation for future therapeutic target identification and precision medicine. While the "incomplete interactome" remains a challenge, the modularity of the GEM-GNN framework allows for future expansion into multimodal graph transformers, incorporating single-cell transcriptomics and epigenomic landscapes. Ultimately, this study demonstrates that explainable graph-based learning is a powerful tool for uncovering the systems-level mechanisms underlying neurodevelopmental disorders, providing a significant contribution to the growing field of computational network medicine.

Data availability

Data will be available upon request and after publication.

Acknowledgments

No Funding is received.

Conflicts of interest

The authors declare no competing financial interests or personal relationships that could have influenced this work.

Author contributions

Adil Rasheed: Conceptualization, data curation, writing – original draft. Aftab Hussain: Methodology, software, visualization, writing – review & editing. Muhammad Huzaifa: Validation, formal analysis. Haroon Anwar: Investigation, data curation. Laila Bibi: Validation, writing – review & editing.

References

Lord C et al 2018 Autism spectrum disorder *Lancet* **392** 508–20

Grove J et al 2019 Identification of common genetic risk variants for autism spectrum disorder *Nat. Genet.* **51** 431–44

De Rubeis S et al 2014 Synaptic, transcriptional and chromatin genes disrupted in autism *Nature* **515** 209–15

Abrahams B S et al 2013 SFARI Gene: an evolving database for the autism research community *Mol. Autism* **4** 1–7

Gandal M J et al 2018 Transcriptome-wide isoform-level imbalance in autism *Science* **362**

Parikshak N N et al 2015 Integrative functional genomic analyses of autism spectrum disorder *Nat. Rev. Genet.* **16** 401–14

Menche J et al 2015 Uncovering disease-disease relationships through the incomplete interactome *Science* **347** 1257601

Cowen L et al 2017 Network propagation: a universal amplifier of genetic associations *Nat. Rev. Genet.* **18** 551–62

Hamilton W, Ying R and Leskovec J 2017 Inductive representation learning on large graphs *Proc. 31st Int. Conf. Neural Inf. Process. Syst. (NIPS)* pp 1025–35

Zitnik M, Agrawal M and Leskovec J 2018 Modeling disease networks with graph convolutional networks *Bioinformatics* **34** i457–i466

Li H et al 2021 Multi-view graph neural networks for prioritizing disease-associated genes *Bioinformatics* **37** i411–i418

Satterstrom F K et al 2020 Large-scale exome sequencing study implicates both developmental and functional changes in the neurobiology of autism *Cell* **180** 568–84

Zhang Y et al 2020 Network-based deep learning framework for autism gene prioritization *Genet. Epidemiol.* **44** 382–94

Fröhlich F et al 2023 Explainable graph neural networks for gene–disease association prediction *Bioinformatics* **39** i445–i453

Barabási A L, Gulbahce N and Leskovec J 2011 Network medicine: a network-based approach to human disease *Nat. Rev. Genet.* **12** 56–68

Menche J et al 2015 Uncovering disease-disease relationships through the incomplete interactome *Science* **347** 1257601

Köhler S et al 2008 Walking the interactome for prioritization of candidate disease genes *Am. J. Hum. Genet.* **82** 949–58

Hormozdiari F et al 2015 A network-based approach to identify deregulated pathways and genes in common diseases *Nat. Commun.* **6** 1–10

SFARI Gene 2024 Simons Foundation Autism Research Initiative (available at: https://gene.sfari.org)

Wang Z et al 2018 PANDA: Protein function prediction using multi-view network integration *Bioinformatics* **34** 1895–903

Bronstein M M et al 2017 Geometric deep learning: going beyond Euclidean data *IEEE Signal Process. Mag.* **34** 18–42

Rao A et al 2022 A survey on graph neural networks for biomedical applications *IEEE/ACM Trans. Comput. Biol. Bioinform.*

Ding J et al 2021 Graph neural networks for disease-gene association prediction: a survey *Brief. Bioinform.* **22**

Liu G et al 2020 Graph representation learning for biomedical networks *Bioinformatics* **36** 1243–50

Cowen M et al 2017 Network propagation: a universal amplifier of genetic associations *Nat. Rev. Genet.* **18** 551–62

Zhang Y et al 2020 Network-based deep learning framework for autism gene prioritization *Genet. Epidemiol.* **44** 382–94

Lin L et al 2022 An integrative machine learning framework for autism risk gene prioritization *BMC Bioinformatics* **23**

Zitnik M, Agrawal M and Leskovec J 2018 Modeling disease networks with graph convolutional networks *Bioinformatics*

Szklarczyk D et al 2021 The STRING database in 2021: protein–protein association networks with increased coverage *Nucleic Acids Res.* **49** D605–D612

Karczewski K J et al 2020 The mutational constraint spectrum quantified from variation in 141,456 humans *Nature* **581** 434–43

Lin T Y et al 2017 Focal loss for dense object detection *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)* pp 2980–88

Kim B et al 2018 Interpretability beyond feature attribution: quantitative testing with concept activation vectors *Proc. 35th Int. Conf. Machine Learning (ICML)*