

Graph Neural Networks for Detecting Relationships and Anomalies in Complex Information Networks

Clara Becker

Assistant Professor, Karlsruhe Institute of Technology, Germany

Abstract

Complex information networks represent interconnected systems in which relationships among entities contribute to the organization, exchange, and interpretation of information. Such networks appear in scientific collaboration systems, digital communication environments, knowledge repositories, and interconnected technological infrastructures. Conventional analytical approaches frequently examine individual entities or rely on predefined structural indicators. However, these procedures may encounter difficulties when network relationships are nonlinear, irregular, heterogeneous, or influenced by interactions extending beyond immediate neighbors. Graph neural networks provide computational approaches for learning representations that incorporate both node characteristics and relational structure. The present paper develops a hypothetical methodological framework for investigating graph neural networks in relationship prediction and anomaly detection within complex information networks. The proposed framework combines graph construction, neighborhood representation, neural feature learning, link prediction, and node-level anomaly assessment.

Common-neighbor analysis, random-walk-based network embeddings, graph convolutional representations, and gated graph neural networks are considered as comparative approaches. A fictional information network containing 4,000 nodes and 18,000 undirected relationships is introduced to illustrate the proposed analytical design. Two hypothetical analytical tables demonstrate how relationship prediction and anomaly detection performance could be reported using precision, recall, and F1-score. The illustrative findings suggest possible differences among conventional graph methods and neural representation approaches. However, all numerical values are manually constructed and do not represent implemented algorithms, experimental simulations, or independently verified results. The methodological discussion examines the relationship between network topology, feature representation, structural irregularities, model complexity, and predictive reliability. Particular attention is given to data leakage, graph sparsity, annotation uncertainty, and the interpretability of anomalous network behavior. The paper concludes that graph neural networks offer a relevant methodological foundation for investigating complex relational data, although their practical effectiveness must be established through reproducible experiments and independently evaluated network datasets.

Keywords: Graph Neural Networks, Complex Information Networks, Relationship Detection, Link Prediction, Anomaly Detection, Graph Representation Learning, Network Analysis, Graph Convolution, Structural Irregularities, Machine Learning.

1. Introduction

Complex information networks have become essential representations of relationships among entities in modern computational environments. Social communication systems, scientific citation databases, knowledge repositories, and digital information infrastructures contain interconnected components whose relationships contribute to their overall organization and operation. These structures can be represented mathematically through graphs consisting of nodes and edges. Nodes describe individual entities, while edges represent associations, interactions, or information exchanges between them. Analyzing such structures provides opportunities to identify meaningful relationships, understand network organization, and recognize unusual patterns that may require further investigation.

Traditional network analysis employs statistical and mathematical techniques to examine graph structure through measures such as node degree, clustering coefficients, path length, and neighborhood similarity. These methods provide useful information about network organization and the position of individual entities. However, their effectiveness may be limited when relationships involve complicated dependencies or when relevant patterns extend across multiple connected regions. A network relationship may depend on several intermediate entities rather than immediate adjacency alone. Consequently, computational methods capable of learning more complex structural representations have attracted increasing attention within network science and machine learning.

Graph neural networks provide a computational framework for learning from information organized through interconnected nodes and relationships. Unlike conventional machine learning approaches that generally process observations as independent feature vectors, graph neural models can incorporate information about neighboring entities during representation learning. Their computational procedures may combine node attributes, edge characteristics, and local network structure through repeated information propagation. This approach allows the learned representation of an entity to reflect selected characteristics of its surrounding network environment. Such capabilities are relevant to relationship prediction, node classification, and the identification of structural irregularities.

Relationship detection and anomaly identification represent two closely related problems in graph-based information analysis. Relationship prediction seeks to estimate whether an association exists or may emerge between particular entities. Anomaly detection focuses on identifying nodes, edges, or connected structures that differ meaningfully from established network patterns. These tasks present different analytical requirements because an unexpected relationship does not necessarily represent an abnormal or harmful activity. Furthermore, graph anomalies may arise through legitimate structural changes, incomplete data collection, or errors in network representation. Reliable analysis therefore requires careful differentiation between unusual statistical patterns and meaningful irregularities.

The present paper proposes a hypothetical comparative investigation of graph neural networks for relationship prediction and anomaly detection within complex information networks. The framework considers graph topology, node-level attributes, neural representation learning, and structured evaluation procedures. Conventional similarity-based approaches and graph embedding methods are included as comparative baselines. A fictional network dataset provides the basis for illustrating the intended analytical process and reporting potential classification outcomes. The study does not present completed computational experiments. Instead, it establishes a methodological foundation for future investigations involving relational learning and graph-based anomaly assessment.

2. Objectives of the Study

The principal objective of this study is to examine the methodological potential of graph neural networks for identifying relationships and anomalous patterns within complex information networks. The investigation emphasizes structural representation, neural classification, and evaluation reliability. Three specific objectives guide the proposed comparative framework and establish the intended scope of graph-based analysis.

2.1 Examination of Relationship Representation in Graph Networks

The first objective is to examine how graph neural networks represent relationships among interconnected entities through structural and attribute-based information. Conventional similarity measures frequently rely on shared neighbors or other predefined graph characteristics. Neural graph models offer an alternative by learning representations from connected nodes and their surrounding structural environments. The proposed study considers how these approaches may capture direct and indirect relationships. Particular attention is given to the representation of neighborhood information and its potential contribution to link prediction.

The investigation further aims to distinguish structural proximity from meaningful relational association. Two entities may occupy neighboring regions of a graph without sharing a direct relationship, while other entities may exhibit related characteristics despite limited connectivity. Consequently, the proposed framework evaluates relationships through appropriately defined positive and negative examples. This objective emphasizes the need for consistent graph construction, carefully selected link prediction criteria, and independent testing before interpreting learned network representations as evidence of meaningful associations.

2.2 Comparative Assessment of Graph-Based Anomaly Detection

The second objective focuses on establishing a comparative procedure for identifying structurally unusual nodes within complex information networks. Degree-based anomaly indicators, network embeddings, graph convolutional representations, and gated neural architectures are considered as alternative approaches. Their intended evaluation examines the ability to distinguish labeled anomalous nodes from those assigned to the normal category. The proposed comparison emphasizes precision, recall, and F1-score as complementary indicators of classification performance under a defined node-level anomaly detection problem.

A further concern involves examining how structural irregularity differs from established anomaly labels. Unusual connectivity patterns may reflect legitimate specialization, network growth, or incomplete records rather than genuinely anomalous behavior. Consequently, the proposed study treats anomaly detection as a classification problem requiring clear reference definitions and appropriate evaluation data. This objective seeks to establish how graph-based representations may contribute to anomaly recognition while recognizing that numerical irregularity alone does not establish the underlying cause or significance of an observed pattern.

2.3 Identification of Methodological and Computational Constraints

The third objective is to identify factors that may influence the reliability of relationship prediction and anomaly detection in complex graph environments. These include incomplete connections, inconsistent

node attributes, heterogeneous relationship types, sparse neighborhoods, and computational limitations. Such characteristics may affect graph representation learning and influence predictive outcomes. The proposed study considers how preprocessing, network partitioning, and appropriate evaluation procedures could address selected difficulties. Particular attention is given to methodological conditions that may produce misleading performance estimates.

The investigation also considers the importance of computational reproducibility and responsible interpretation. Graph neural models may incorporate information from connected entities, making data leakage more difficult to recognize than in conventional independent-observation datasets. Model complexity can also complicate the explanation of unusual network predictions. Consequently, this objective emphasizes transparent graph construction, carefully separated evaluation information, and appropriate documentation of modeling decisions. These considerations support a defensible assessment of graph-based predictive systems.

3. Review of Literature

3.1 Foundations of Graph Structure and Relationship Prediction

Jeh and Widom (2002) introduced SimRank as a measure of structural-context similarity, establishing an influential approach to comparing entities through network relationships. Adamic and Adar (2003) investigated social connections through web-based network information and developed insights relevant to structural similarity. Liben-Nowell and Kleinberg (2007) examined the link prediction problem and demonstrated how network proximity measures could be used to estimate potential relationships. These studies established important foundations for investigating associations among entities through graph topology and connectivity patterns.

Newman (2003) reviewed structural characteristics of complex networks and their implications for understanding connected systems. Girvan and Newman (2002) investigated community structure, while Fortunato (2010) provided a broader examination of graph communities and clustering methods. Backstrom and Leskovec (2011) subsequently investigated supervised random walks for predicting links in social networks. Together, these contributions indicate that relational prediction can benefit from structural information, although the usefulness of specific proximity measures depends on network organization and the characteristics of the underlying relationships.

3.2 Evolution of Graph Neural Representation Learning

Scarselli et al. (2009) introduced a graph neural network model capable of processing information represented through interconnected graph structures. Their approach established a foundation for learning representations that reflect relationships among nodes. Perozzi et al. (2014) developed DeepWalk, which applies random-walk-based representation learning to network data. Tang et al. (2015) introduced LINE for large-scale information network embedding. These studies expanded computational approaches for transforming graph connectivity information into representations suitable for predictive analysis.

Duvenaud et al. (2015) investigated convolutional networks operating directly on graph structures. Defferrard et al. (2016) developed localized spectral filtering for graph convolutional networks, while Li et al. (2016) introduced gated graph sequence neural networks. Grover and Leskovec (2016) proposed node2vec for learning flexible structural representations, and Wang et al. (2016)

examined structural deep network embedding. Kipf and Welling (2016) introduced a variational graph autoencoder framework. Collectively, these contributions established alternative mechanisms for learning graph representations and predicting relationships.

3.3 Anomaly Identification in Connected Information Structures

Noble and Cook (2003) examined graph-based anomaly detection and demonstrated the relevance of structural patterns to identifying unusual network observations. Chandola et al. (2009) reviewed anomaly detection methods and established a broader methodological classification of approaches used across different data environments. Akoglu et al. (2010) introduced OddBall for detecting structural irregularities in weighted graphs. These studies established important foundations for identifying deviations from expected network patterns through graph structure, neighborhood characteristics, and statistical relationships.

Akoglu et al. (2015) provided a comprehensive survey of graph-based anomaly detection and discussed approaches involving static, dynamic, attributed, and unattributed networks. Their review emphasized the importance of identifying unusual graph structures and examining why particular observations receive anomalous classifications. Fawcett (2006) and Davis and Goadrich (2006) contributed important principles for evaluating binary classification performance. These studies indicate that reliable network anomaly analysis requires suitable reference standards, appropriate metrics, and careful differentiation between structural irregularity and meaningful anomalous behavior.

4. Research Methodology

4.1 Research Design and Graph Representation Framework

The present paper adopts a hypothetical comparative methodology involving relationship prediction and node-level anomaly detection. The proposed framework combines structural similarity methods, random-walk embeddings, graph convolutional representations, and gated graph neural architectures. Each approach would operate on a common network representation, although task-specific prediction procedures would differ. Relationship prediction would evaluate candidate node pairs, whereas anomaly detection would assess individual nodes. No computational experiments have been conducted, and the methodology is presented exclusively as an illustrative research design.

The proposed analytical process consists of graph preparation, attribute normalization, representation learning, task-specific prediction, and independent evaluation. For link prediction, model embeddings would support the estimation of relationships between candidate node pairs. For anomaly detection, graph-based scores would reflect deviations in neighborhood structure or reconstruction behavior. Model parameters would be selected using validation information. The independent test sets would remain excluded from model optimization. This structure is intended to provide a reproducible basis for comparing graph-based predictive approaches.

4.2 Hypothetical Network Dataset and Preprocessing

For illustrative purposes, the proposed study assumes a fictional undirected information network containing 4,000 nodes and 18,000 observed edges. Each node is assigned a 32-dimensional attribute vector representing hypothetical informational characteristics. The graph is assumed to contain several connected communities with different densities and interaction patterns. A subset of 200 nodes is

designated as anomalous for the intended evaluation. These specifications are manually constructed and do not describe actual collected or computationally generated network observations.

The proposed preprocessing procedure would include removal of duplicate edges, verification of node identifiers, attribute normalization, and examination of disconnected components. Relationship prediction would require withholding selected observed edges and constructing appropriate comparison pairs. Test edges would be excluded from the message-passing graph during model development. Nonedge samples would be treated as assumed negatives rather than universally verified absent relationships. For anomaly detection, reference labels would remain separated from unsupervised representation learning unless explicitly required by a supervised baseline.

4.3 Graph Model Development and Performance Assessment

The relationship prediction component would compare common-neighbor similarity, DeepWalk-based representations, a graph convolutional encoder, and a gated graph neural encoder. Neural representations would be transformed into pairwise relationship scores through a consistent decoding procedure. The anomaly detection component would compare a degree-based structural baseline, node2vec embeddings with an outlier detector, and reconstruction-based neural graph models. These approaches would provide contrasting methods for representing neighborhood structure and identifying unexpected network patterns.

The proposed evaluation would use independently reserved observations for both analytical tasks. Link prediction would be assessed using precision, recall, and F1-score on candidate node pairs. Anomaly detection would employ the same measures for a labeled node-level evaluation set. Additional ranking-based metrics and uncertainty estimates would be appropriate in a genuine experiment. Hyperparameters and detection thresholds would be selected using development information. Repea

5. Analysis

5.1 Illustrative Evaluation of Relationship Prediction

Relationship prediction involves identifying whether associations between selected node pairs can be inferred from available network information. Traditional proximity measures estimate associations using shared neighbors, whereas learned representations can incorporate broader structural patterns. Table 1 presents hypothetical classification outcomes for four graph-based prediction approaches. The constructed evaluation set contains 600 positive relationships and 600 sampled nonedges. The illustrative values demonstrate how precision, recall, and F1-score could be calculated from assumed true-positive and false-positive classifications.

The hypothetical comparison assigns progressively favorable outcomes to the representation-learning approaches, with the gated graph neural encoder receiving the highest illustrative F1-score. However, these numerical differences do not demonstrate actual predictive superiority because none of the models were trained or evaluated. A genuine investigation would require appropriate relationship labels, independent edge withholding, consistent negative sampling, and repeated performance assessment. The interpretation of li

Table 1: Hypothetical Relationship Prediction Performance

Relationship Prediction Model	True Positives	False Positives	Precision	Recall	F1-Score
Common Neighbors	390	135	0.743	0.650	0.693
DeepWalk-Based Model	450	115	0.796	0.750	0.773
Graph Convolutional Encoder	480	92	0.839	0.800	0.819
Gated Graph Neural Encoder	498	84	0.856	0.830	0.843

The illustrative table demonstrates how learned network representations might be evaluated against conventional proximity measures. The assumed outcomes are not empirical findings, and the apparent differences have no established statistical significance.

5.2 Illustrative Assessment of Node-Level Anomaly Detection

Node-level anomaly detection involves identifying entities whose structural or attribute-related characteristics differ from those represented by an established reference distribution. Such deviations may include unusual connectivity, inconsistent neighborhood characteristics, or unexpected reconstruction behavior. Table 2 presents hypothetical classification outcomes for four anomaly detection approaches. The constructed evaluation set contains 200 anomalous nodes and 1,800 normal nodes. These labels and model predictions are fictional and have not been derived from an actual information network.

Table 2: Hypothetical Node-Level Anomaly Detection Performance

Anomaly Detection Model	True Positives	False Positives	Precision	Recall	F1-Score
Degree-Based Anomaly Rule	90	115	0.439	0.450	0.444
node2vec with Outlier Detection	116	92	0.558	0.580	0.569
Graph Convolutional Reconstruction	142	76	0.651	0.710	0.679
Gated Graph Reconstruction	150	66	0.694	0.750	0.721

The example highlights the importance of examining precision and recall together. An approach identifying more anomalous nodes may also generate additional false alerts, requiring interpretation within the intended information network application.

6. Discussion

6.1 Structural Representation and Relationship Discovery

Graph neural networks offer a methodological approach for examining relational information through the interaction of node attributes and network connectivity. Their computational procedures can incorporate information from neighboring entities, allowing learned representations to reflect selected structural characteristics. This approach may provide advantages when relationships are influenced by more than immediate graph proximity. However, useful representation learning depends on appropriate neighborhood selection, attribute quality, and a network structure that contains relevant information for the prediction task.

The proposed framework also demonstrates why conventional structural methods remain important comparative baselines. Common-neighbor measures provide relatively transparent estimates of relational proximity and can be computationally inexpensive. Neural approaches introduce additional flexibility but require model optimization and more extensive evaluation. Consequently, improved structural representation should not be assumed merely because a model contains neural components. A meaningful comparison must determine whether learned representations provide reproducible benefits beyond simpler network statistics under equivalent data conditions.

6.2 Network Anomalies and Contextual Interpretation

Anomalous nodes may differ from surrounding entities in terms of connectivity, attribute values, or the consistency of their relationships. Graph representation methods can support the identification of such differences by incorporating neighborhood information into node-level predictions. However, unusual network behavior is not necessarily undesirable or erroneous. Highly connected entities may occupy legitimate coordinating positions, while nodes with uncommon attributes may represent specialized roles. Consequently, anomaly identification requires context-sensitive interpretation rather than automatic assumptions about the significance of structural irregularities.

The hypothetical anomaly detection comparison illustrates how different modeling approaches could produce distinct balances between correct identifications and false alerts. These differences may reflect the information considered by each model and the selected decision threshold. However, the numerical patterns do not establish that one algorithm identifies meaningful anomalies more effectively. A completed investigation would require clearly defined anomaly categories and appropriate validation information. Analysts would also need to examine why particular nodes receive unusual scores.

6.3 Comparative Reliability and Model Generalization

Graph-based predictive evaluation presents methodological challenges because observations within a network are not generally independent. Nodes may share neighbors, attributes, or community memberships that influence their learned representations. Consequently, conventional random splitting procedures may create information leakage when the intended deployment requires predictions about unseen relationships or future network states. The proposed framework emphasizes explicit separation of

withheld edges and careful control of graph information available during model development. Such safeguards are essential for interpreting relationship prediction performance.

Model generalization also depends on whether the learned representations remain useful across networks with different structural characteristics. A graph neural model developed using a relatively dense information network may behave differently when applied to sparse or heterogeneous graphs. Similarly, anomaly thresholds selected for one network may not transfer reliably to another. Therefore, future research should investigate performance across distinct topologies, attribute distributions, and network sizes. The illustrative results presented here cannot establish such generalization properties.

7. Challenges

7.1 Graph Sparsity and Incomplete Relationship Information

Complex information networks frequently contain incomplete or unevenly distributed relationship information. Some entities may have numerous recorded connections, while others possess only a small number of observed links. Sparse neighborhoods can limit the information available to graph neural models and reduce the reliability of learned representations. Missing edges may also complicate link prediction because an unobserved relationship is not necessarily absent. Consequently, graph construction and negative-sample selection require careful methodological justification before model performance can be interpreted.

Incomplete relationship records also create difficulties for anomaly detection. A node with unusually few observed connections may appear anomalous because of limited data collection rather than an underlying structural irregularity. Conversely, missing information can conceal meaningful deviations from expected patterns. The proposed methodology therefore emphasizes data completeness assessment and transparent documentation of graph construction. Future experiments should evaluate whether apparent anomalies remain consistent when additional relationship information becomes available or when alternative network representations are considered.

7.2 Computational Scalability and Representation Complexity

Graph neural networks can require substantial computational resources when processing large information networks with numerous nodes, edges, and attributes. Repeated neighborhood aggregation may increase memory requirements and training costs, particularly when multiple connected layers are used. Dense regions of a graph may also produce extensive computational dependencies. These conditions can restrict the practical application of complex neural architectures. Consequently, model development should consider computational efficiency alongside the quality of relationship prediction and anomaly detection outcomes.

Representation complexity introduces additional methodological concerns because deeper message-passing procedures may not always preserve useful distinctions among neighboring entities. Repeated aggregation can produce increasingly similar node representations, potentially weakening discrimination in selected architectures. Computational strategies must therefore balance neighborhood coverage, model depth, and the preservation of informative structural characteristics. The proposed framework considers these trade-offs conceptually. Actual evaluation would require measurements of training duration, memory consumption, inference efficiency, and predictive stability across different network sizes.

7.3 Interpretability, Evaluation Bias, and Anomaly Validation

Interpretability remains a significant challenge because neural graph models may produce relationship scores and anomaly classifications through complex interactions among node attributes and network topology. Identifying the specific connections responsible for a prediction can be difficult, particularly when information propagates across several neighboring regions. In sensitive analytical environments, unexplained classifications may limit confidence in automated decisions. Consequently, future research should consider methods for examining influential nodes, relevant attributes, and structural evidence associated with individual predictions.

Evaluation bias also requires attention because graph relationships can introduce dependencies between training and testing observations. Inappropriate use of withheld edges, inconsistent anomaly labels, or repeated threshold selection may generate misleading performance estimates. Reliable experimental procedures should establish clear prediction targets, restrict information leakage, and preserve independent evaluation data. Furthermore, anomaly classifications require suitable reference standards because structural unusualness is not equivalent to verified abnormal behavior. These safeguards are necessary for establishing defensible conclusions about graph-based predictive effectiveness.

8. Conclusion

Graph neural networks provide a methodological approach for examining complex information systems through the combined analysis of node attributes and relational structure. The present paper developed a hypothetical framework for relationship prediction and node-level anomaly detection using conventional proximity measures, network embeddings, graph convolutional representations, and gated neural architectures. Its principal contribution is the organization of a comparative evaluation procedure that recognizes the distinct analytical requirements of identifying relationships and detecting structural irregularities.

The illustrative tables demonstrate possible methods for reporting true-positive classifications, false-positive classifications, precision, recall, and F1-score. However, all numerical values were manually constructed and do not represent completed computational experiments or verified network observations. Meaningful performance comparisons would require authentic or reproducibly generated datasets, implemented learning algorithms, clearly defined anomaly labels, and independent testing. Consequently, the hypothetical findings cannot establish the superiority of graph neural networks over conventional graph analytical methods.

Overall, the study identifies graph sparsity, structural dependency, computational complexity, interpretability, and evaluation reliability as important considerations in graph-based predictive analysis. Future empirical research should examine these factors across networks with different topological characteristics and information requirements. Graph neural networks may offer useful capabilities for investigating interconnected data, but their practical value must be established through transparent computational procedures and independently verified evidence. The proposed framework therefore serves as a foundation for subsequent experimental investigation.

References

1. Akoglu, L., Tong, H., & Koutra, D. (2015). Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29, 626–688. <https://doi.org/10.1007/s10618-014-0365-y>
2. Bandyopadhyay, S., N, L., Vivek, S. V., & Murty, M. N. (2020). Outlier resistant unsupervised deep architectures for attributed network embedding. In *Proceedings of the 13th ACM International Conference on Web Search and Data Mining* (pp. 25–33). ACM. <https://doi.org/10.1145/3336191.3371808>
3. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15. <https://doi.org/10.1145/1541880.1541882>
4. Ding, K., Li, J., Bhanushali, R., & Liu, H. (2019). Deep anomaly detection on attributed networks. In *Proceedings of the 2019 SIAM International Conference on Data Mining* (pp. 594–602). Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9781611975673.67>
5. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, 30.
6. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1609.02907>
7. Kim, H., Lee, B. S., Shin, W.-Y., & Lim, S. (2022). Graph anomaly detection with graph neural networks: Current status and challenges. *IEEE Access*, 10, 111820–111829. <https://doi.org/10.1109/ACCESS.2022.3211306>
8. Koutra, D., Vogelstein, J. T., & Faloutsos, C. (2013). Deltacon: A principled massive-graph similarity function. In *Proceedings of the 2013 SIAM International Conference on Data Mining* (pp. 162–170). SIAM. <https://doi.org/10.1137/1.9781611972832.18>
9. Ma, X., Wu, J., Xue, S., Yang, J., Zhou, C., Sheng, Q. Z., Xiong, H., & Akoglu, L. (2021). A comprehensive survey on graph anomaly detection with deep learning. *arXiv*. <https://arxiv.org/abs/2106.07178>
10. Newman, M. (2010). *Networks: An introduction*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199206650.001.0001>
11. Perozzi, B., Al-Rfou, R., & Skiena, S. (2014). DeepWalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 701–710). ACM. <https://doi.org/10.1145/2623330.2623732>
12. Ren, J., Xia, F., Lee, I., Hoshyar, A. N., & Aggarwal, C. (2023). Graph learning for anomaly analytics: Algorithms, applications, and challenges. *ACM Transactions on Intelligent Systems and Technology*, 14(2), Article 28. <https://doi.org/10.1145/3570906>
13. Kipf, T. N., & Welling, M. (2016). Variational graph auto-encoders. *arXiv*. <https://arxiv.org/abs/1611.07308>
14. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1710.10903>

15. Grover, A., & Leskovec, J. (2016). node2vec: Scalable feature learning for networks. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 855–864). ACM. <https://doi.org/10.1145/2939672.2939754>
16. Zhang, M., & Chen, Y. (2018). Link prediction based on graph neural networks. In Advances in Neural Information Processing Systems, 31.
17. Bei, Y., Zhou, S., Tan, Q., Xu, H., Chen, Z., Li, J., & Bu, J. (2023). Reinforcement neighborhood selection for unsupervised graph anomaly detection. In 2023 IEEE International Conference on Data Mining (ICDM) (pp. 11–20). IEEE.
18. Bhatia, S., Wadhwa, M., Kawaguchi, K., Shah, N., Yu, P. S., & Hooi, B. (2023). Sketch-based anomaly detection in streaming graphs. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 93–104). ACM.
19. Ekle, O. A., & Eberle, W. (2024). Anomaly detection in dynamic graphs: A comprehensive survey. arXiv. <https://arxiv.org/abs/2406.00134>
20. A brief survey on graph anomaly detection. (2024). Procedia Computer Science, 242, 1263–1270. <https://doi.org/10.1016/j.procs.2024.08.145>