

Dual-Encoder Graph Anomaly Detection via Structural–Attribute Consistency Learning*

Hessa O. ALHARBI¹ and Hedia ZARDI²

^{1,2}Department of Computer Science, College of Computer, Qassim University, Buraydah, KSA.
, h.zardi@qu.edu.sa

Correspondence should be addressed to: Hessa O. ALHARBI, 451214900@qu.edu.sa

*Presented at the 47th IBIMA International Conference, 29-30 June 2026, Madrid, Spain

Abstract

Static graph anomaly detection is important in applications such as cybersecurity, fraud detection, and social network analysis. However, detecting anomalies in attributed graphs remains challenging because of complex graph structures and limited labeled data. Existing methods often focus on either structural information or node attributes separately, which may reduce detection performance.

This paper proposes a dual-encoder framework for anomaly detection in static graphs. The proposed method combines a GraphSAGE-based structural encoder with node-level attention and an MLP-based attribute encoder to learn complementary node representations. In addition, a hybrid anomaly scoring mechanism is introduced to capture both structural–attribute inconsistency and neighbor-aware deviations. Experiments on benchmark datasets including Cora, Citeseer, and PubMed demonstrate that the proposed method achieves competitive performance compared to existing approaches.

Keywords: Anomaly Detection, Static Graphs, Attention Mechanism, Deep Learning.

Introduction

Detecting anomalies in structured data is a key area of critical importance, given the widespread use of such data in many real-world applications, including social networks, cybersecurity, and financial analysis. Detecting anomalies in these graphs is an important task because they may indicate fraudulent activities, network intrusions, or abnormal behaviors.

Detecting anomalies in static graphs is a challenging process due to the complex relationships between nodes, the sparse graph structures, and the limited availability of labeled data. The traditional anomaly detection methods often depend on handcrafted features and statistical analysis, which limits their ability to capture complex structural patterns. Recently, Graph Neural Networks (GNNs) have shown strong performance in graph representation learning by aggregating information from neighboring nodes. However, many existing methods focus mainly on either graph structure or node attributes separately, which may reduce detection effectiveness.

To address these limitations, this paper proposes a dual-encoder framework for anomaly detection in static attributed graphs. The proposed method combines a GraphSAGE-based structural encoder with node-level attention and an MLP-based attribute encoder to learn complementary node representations.

Related Work

Graph anomaly detection aims to identify nodes or connections that significantly deviate from normal graph patterns. Traditional approaches mainly rely on statistical methods and handcrafted features to detect abnormal

behaviors. Although these methods are effective in simple scenarios, they often struggle to capture complex relationships in large graph-structured data.

Recently, deep learning techniques, especially Graph Neural Networks (GNNs), have achieved strong performance in graph representation learning (Wu et al., 2020). GNN-based methods can learn structural information by aggregating features from neighboring nodes, making them suitable for anomaly detection tasks in attributed graphs.

Several graph anomaly detection models have been proposed in recent years. DOMINANT (K. Ding et al., 2019) uses graph autoencoders to jointly reconstruct graph structure and node attributes for anomaly detection. CoLA (Y. Liu et al., 2022) introduces a contrastive self-supervised learning framework to distinguish anomalous nodes from normal ones. Other approaches incorporate attention mechanisms to improve representation learning by focusing on important nodes and connections (V. Lanko and I. Makarov, 2024).

Despite these advances, many existing methods mainly focus on either structural information or attribute information separately. In addition, some approaches do not explicitly consider neighborhood-based deviations when computing anomaly scores. Therefore, effectively combining structural consistency, attribute information, and local neighborhood patterns remains an important challenge in static graph anomaly detection.

To address these limitations, the proposed method introduces a dual-encoder framework with a hybrid anomaly scoring mechanism that jointly captures structural–attribute inconsistency and neighbor-aware deviations.

Proposed Method

The proposed method aims to detect anomalous nodes in static graphs through the combined learning of structural and attribute information. The core idea is that normal nodes typically maintain consistency between the graph structure and the node attribute, while anomalous nodes tend to disrupt this consistency.

The framework consists of three main components:

- A GraphSAGE-based structural encoder.
- An MLP-based attribute encoder.
- A hybrid anomaly scoring mechanism.

3.1 Structural Encoder

To capture graph topology information, a two-layer GraphSAGE encoder is used. GraphSAGE (W. L. Hamilton, 2017) learns node representations by aggregating information from neighboring nodes, allowing the model to capture local structural patterns.

To improve representation quality, a lightweight node-level attention mechanism is added after the GraphSAGE layers. The attention mechanism assigns importance weights to node embeddings, helping the model focus on informative nodes while reducing the effect of noisy representations.

The final structural embedding is obtained by scaling the node representation using the learned attention weights.

Attribute Encoder

Besides structural information, node attributes also contain important semantic information. Therefore, a two-layer Multilayer Perceptron (MLP) is used to generate attribute embeddings from node features.

The MLP applies nonlinear transformations to map node attributes into a latent representation space. These embeddings provide complementary information to the structural representations learned by GraphSAGE.

Hybrid Anomaly Scoring

The model is trained by minimizing the difference between structural embeddings and attribute embeddings for normal nodes. Mean Squared Error (MSE) loss is used to measure the consistency between the two representations. To improve training stability, embeddings are normalized before loss computation.

To detect different types of anomalies, a hybrid anomaly scoring mechanism is introduced.

Structural–Attribute Discrepancy: The first score measures the difference between the structural embedding and the attribute embedding of each node. Large differences indicate possible anomalies.

Neighbor-Aware Structural Deviation: The second score measures how much a node deviates from its neighboring nodes. The structural embedding of each node is compared with the average embedding of its neighbors. Nodes with large deviations are considered potentially anomalous.

Final Anomaly Score: The final anomaly score is computed by combining the normalized scores using weighted fusion:

$$Score(v) = w_1 \cdot Score_{orig}(v) + w_2 \cdot Score_{neigh}(v)$$

where w_1 and w_2 control the contribution of each scoring component.

This hybrid strategy enables the model to detect both attribute-based and structure-based anomalies effectively.

Experiments and Results

Datasets

The proposed method is evaluated using three widely used benchmark citation network datasets: Cora, Citeseer, and PubMed. In these datasets, nodes represent scientific publications, edges represent citation relationships, and node features are represented using bag-of-words vectors.

Table 1: The statistics of the datasets.

Dataset	# Nodes	# Edges	# Features
Cora	2,708	5,429	1,433
CiteSeer	3,327	4,732	3,703
PubMed	19,717	44,338	500

These datasets provide different graph sizes and feature distributions, making them suitable for evaluating graph anomaly detection methods.

Baseline Methods

The proposed model is compared with several representative graph anomaly detection methods, including:

AMEN – a structural anomaly detection method (B. Perozzi and L. Akoglu,2016).

CoLA – a contrastive self-supervised anomaly detection model (Y. Liu et al.,2022).

LCPool-LCUnpool – a model for ranking nodes based on their degree of abnormality (M. Mesgaran and A. Ben Hamza,2023).

These methods are selected because they represent traditional, autoencoder-based, and self-supervised graph anomaly detection approaches.

Evaluation Metrics

For **overall performance**, we used the ROC-AUC (Area Under the Receiver Operating Characteristic Curve) metrics. Based on the results compared to current baselines, the proposed method achieves the best performance on Cora and Citeseer and remains highly competitive on PubMed. These results demonstrate the effectiveness of jointly learning structural and attribute representations for anomaly detection. As shown in table 2, where the numbers in bold are the best result, while those underlined are the second-best result.

Table 2: Comparison of Anomaly Detection Performance (AUC %) Across Benchmark Datasets

Methods	Cora	Citeseer	PubMed
---------	------	----------	--------

AMEN	62.66	61.54	77.13
CoLA	87.79	89.68	95.12
LCPool-LCUnpool	<u>88.46</u>	<u>91.34</u>	92.81
Proposed Model	95.82	94.19	<u>92.91</u>

To measure other parts, we used other metrics.

- Precision@K: Measures the proportion of correctly identified anomalous nodes among the top K nodes with the highest anomaly scores.
- Recall@K: Measures the proportion of true anomalous nodes that are included within the top K ranked nodes.
- F1-score@K: The harmonic mean of Precision@K and Recall@K, providing a balanced evaluation of detection performance.

The experimental results demonstrate that the proposed model achieves strong and consistent Precision@K performance across all benchmark datasets, indicating its effectiveness in accurately ranking anomalous nodes.




Figure 1. Precision@K (%) Performance Comparison on Benchmark Datasets.

The model also achieves competitive Recall@K results, particularly on the Cora and Citeseer datasets, showing its ability to successfully identify a large proportion of true anomalies.




Figure 2. Recall@K (%) Performance Comparison on Benchmark Datasets

In terms of F1-score@K, the proposed approach maintains a strong balance between precision@K and recall@K, outperforming several baseline methods in most evaluation settings.




Figure 3. F1-score@K (%) Performance Comparison on Benchmark Datasets

These results highlight the effectiveness of combining structural and attribute representations through the dual-encoder framework.

Overall, the experimental findings confirm that the proposed method provides robust and reliable anomaly detection performance for static attributed graphs while effectively capturing both structural irregularities and attribute inconsistencies.

Conclusion

This paper presented a dual-encoder framework for anomaly detection in static attributed graphs. The proposed method combines a GraphSAGE-based structural encoder with node-level attention and an MLP-based attribute encoder to jointly learn structural and attribute representations.

In addition, a hybrid anomaly scoring mechanism was introduced to capture both structural–attribute inconsistency and neighbor-aware deviations. Experimental results on benchmark datasets including Cora, Citeseer, and PubMed demonstrated that the proposed method achieves competitive performance compared with several existing approaches, particularly in terms of AUC and Precision@K.

Overall, the proposed framework provides an effective and robust solution for detecting anomalies in static graph data. Future work may focus on extending the method to dynamic graphs and exploring more advanced graph representation learning techniques.

References

- Ding, K., Li, J., Bhanushali, R. and Liu, H. (2019), ‘Deep anomaly detection on attributed networks (DOMINANT),’ *Proceedings of the SIAM International Conference on Data Mining (SDM)*.
- Ding, K., Zhou, Q., Tong, H. and Liu, H. (2021), ‘Deep anomaly detection in attributed networks: problem, solutions and challenges,’ *IEEE Transactions on Knowledge and Data Engineering*.
- Fotiadou, K. et al. (2021), ‘Network traffic anomaly detection via deep learning,’ *IEEE Transactions*.
- Hamilton, W. L., Ying, R. and Leskovec, J. (2017), ‘Inductive representation learning on large graphs,’ *Advances in Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA.
- Lanko, V. and Makarov, I. (2024), ‘Graph-attention diffusion for enhanced multivariate time-series anomaly detection,’ *IEEE Open Journal of the Industrial Electronics Society*, 5, 1353–1366.
- Liu, Y. et al. (2022), ‘Anomaly detection on attributed networks via contrastive self-supervised learning,’ *arXiv preprint*.
- Mesgaran, M. and Ben Hamza, A. (2023), ‘A graph encoder–decoder network for unsupervised anomaly detection,’ *arXiv preprint arXiv:2308.07774*.
- Ofori-Boateng, D. et al. (2021), ‘Nonparametric anomaly detection on time series of graphs,’ *IEEE Transactions*.
- Perozzi, B. and Akoglu, L. (2016), ‘Scalable anomaly ranking of attributed neighborhoods,’ *Proceedings of the SIAM International Conference on Data Mining (SDM)*.
- Thokar, B. et al. (2024), ‘Medical image segmentation for anomaly detection using deep learning techniques,’ *arXiv preprint*.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C. and Yu, P. S. (2020), ‘A comprehensive survey on graph neural networks,’ *IEEE Transactions on Neural Networks and Learning Systems*.