

Article

Contamination-Reduced Multi-View Reconstruction for Graph Anomaly Detection

Qiang Li ¹ , Peng Zhang ^{1,*} and Qingfeng Tan ^{1,2,*}¹ Cyberspace Institute of Advanced Technology, Guangzhou University, Guangzhou 510006, China; liq@e.gzhu.edu.cn² University International College, Macau University of Science and Technology, Macau

* Correspondence: p.zhang@gzhu.edu.cn (P.Z.); tqf528@gzhu.edu.cn (Q.T.)

Abstract

Graph anomaly detection (GAD) is pivotal for security-critical applications like cybersecurity and financial fraud detection. While reconstruction-based Graph Neural Networks (GNNs) are prevalent, their efficacy is often compromised by two phenomena: (1) anomaly overfitting, where expressive models capture anomalous patterns, and (2) homophily-induced attenuation, where message passing smooths localized anomaly cues. This paper proposes CLEAN-GAD, a contamination-aware framework that mitigates anomaly influence during training through multi-view robust learning. Specifically, we develop a contrastive augmentation module that utilizes local inconsistency scores to identify and suppress pseudo-anomalous nodes and edges, thereby yielding a purified augmented view. To capture diverse anomaly signals, a frequency-adaptive encoder with dual-pass channels is designed to integrate low- and high-frequency information. Furthermore, we introduce a distribution-separation regularizer and cross-view alignment to stabilize learning and resolve view shifts. Theoretical analysis confirms that reducing the contamination ratio ρ expands the reconstruction-risk gap between normal and anomalous nodes, inherently boosting detection performance. Extensive experiments on multiple benchmark datasets from various domains demonstrate the superior anomaly detection performance of CLEAN-GAD.

Keywords: graph anomaly detection; multi-view graph learning; contrastive augmentation; anomaly contamination



Academic Editors: Martin Valtierra-Rodriguez, Huijia Li, Jun Hu, Weichen Zhao and Jie Cao

Received: 25 December 2025

Revised: 26 January 2026

Accepted: 28 January 2026

Published: 1 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Graph anomaly detection refers to identifying abnormal patterns in a graph that deviate significantly from most nodes and has been widely applied in social networks [1–3], financial fraud detection [4–6], to name a few.

With recent advances in graph representation learning, reconstruction-based graph neural network (GNN) models have emerged as a prevalent paradigm for unsupervised graph anomaly detection [7,8]. These methods typically train graph autoencoders to reconstruct node attributes and graph structure, and compute anomaly scores based on reconstruction errors, under the commonly adopted assumption that normal nodes are reconstructed more accurately than anomalous ones [9]. Despite their strong empirical performance, two fundamental challenges remain and can substantially limit their reliability in practice.

The first challenge is **anomaly overfitting**. Because reconstruction objectives optimize all nodes in an undifferentiated manner, expressive GNNs may inadvertently fit anomalous patterns, yielding small reconstruction errors even for anomalous nodes and thereby weakening anomaly separability. Although a recent work proposes a denoising strategy to mitigate this issue [10], it does not provide a theoretical account of how anomaly contamination affects reconstruction-based optimization, and the resulting pseudo-supervision may be insufficiently precise to fully exploit anomaly-related signals. The second challenge is **Camouflage-Driven Anomaly Attenuation**. In many real-world graphs, anomalous nodes may connect to normal neighbors to camouflage their behavior; meanwhile, message passing tends to homogenize representations of adjacent nodes [11], which can attenuate anomaly-related cues and make anomalies difficult to distinguish. From a graph signal processing perspective [12], this homogenization effect corresponds to an excessive bias toward low-frequency components, whereas anomalous behaviors often manifest as localized high-frequency variations. This motivates a frequency-adaptive design that explicitly preserves and selectively fuses low- and high-frequency signals for anomaly scoring.

To address the above challenge, we propose **CLEAN-GAD**, a contamination-aware framework that improves reconstruction-based GAD by explicitly reducing the effective anomaly ratio seen during training and learning robust anomaly scores from complementary graph views. CLEAN-GAD consists of two components. First, we introduce a contrastive augmentation module to extract *local anomaly signals*. By contrasting a node representation against its neighborhood context, we obtain a local inconsistency score and construct an *augmented graph view* by suppressing high-confidence pseudo-anomalous nodes (and their incident edges), which better preserves dominant normal patterns. Second, we perform *multi-view reconstruction learning* on both the original graph and the augmented view. To capture anomaly cues that may reside in localized, high-frequency variations, we design a frequency-adaptive encoder with low-/high-pass channels and anomaly-aware fusion. We further align anomaly scores across views to improve consistency under view shift and employ a distribution-separation regularizer to stabilize training under noisy pseudo labels. The final anomaly score integrates the local contrastive score and the global reconstruction-based score.

Beyond the algorithmic design, we provide a focused theoretical analysis of the key mechanism behind CLEAN-GAD: *reducing the effective anomaly ratio improves reconstruction-risk separability*. Under a mixture data-generating model and mild smoothness/convexity conditions, the reconstruction-risk gap decreases approximately linearly with the anomaly ratio near zero. Therefore, the augmentation step that lowers the effective contamination ratio from ρ to ρ_{aug} preserves a larger risk gap, making reconstruction errors more discriminative and yielding improved ranking-based metrics such as AUC and AP. This analysis offers a principled explanation for the empirical gains observed with anomaly-aware augmentation. Our contributions can be summarized as follows:

- We identify and formalize the contamination issue in reconstruction-based GAD, highlighting that anomaly presence in training data reduces reconstruction-risk separability and harms anomaly ranking.
- We propose CLEAN-GAD, a two-stage framework that constructs an augmented graph view via contrastive local anomaly scoring and learns robust anomaly scores through multi-view reconstruction with frequency-adaptive encoding.
- We provide a theoretical analysis showing how reducing the effective anomaly ratio enlarges the reconstruction-risk gap, offering an explanation for the improved empirical performance of CLEAN-GAD.

- Extensive experiments and ablation studies demonstrate that CLEAN-GAD consistently outperforms representative baselines and achieves competitive or state-of-the-art performance across multiple benchmark datasets.

2. Related Works

In this section, we review the most closely related studies, focusing on graph anomaly detection and graph data augmentation.

2.1. Graph Anomaly Detection

Graph anomaly detection [3] seeks to identify nodes whose structural connectivity patterns and/or attribute distributions—at either the local neighborhood scale or the global graph level—deviate substantially from those of the majority. Early studies primarily leveraged handcrafted structural statistics or shallow embedding representations. Common indicators include node degree, clustering coefficient, PageRank, and motif frequencies; for instance, ref [13] adopts LOF-style techniques based on local density. Anomalous [14] formulates detection as a low-rank-plus-sparse decomposition of the adjacency matrix, using the residual component to quantify abnormality. Probabilistic methods [15], including mixed-membership models and stochastic block models, treat nodes with low generative likelihood as anomalies. Although these approaches offer a degree of interpretability, they often exhibit limited flexibility when confronted with high-dimensional node attributes and complex relational dependencies.

With the rapid advances in graph representation learning and graph neural networks (GNNs), research has increasingly shifted toward end-to-end deep models that jointly encode structural and attribute information and detect anomalies via reconstruction error, generative likelihood, or deviations in the learned embedding space. This trend has gradually formed a unified paradigm of deep-learning-based graph anomaly detection. Reconstruction-based methods such as DOMINANT [9] and AnomalyDAE [16] employ graph autoencoders (GAE), variational GAEs (VGAE), and related variants to reconstruct both node attributes and graph structure, treating reconstruction error as the anomaly score. In contrast, CoLA [17] and CONAD [18] introduce self-supervised contrastive learning for node-level anomaly detection, exploiting neighborhood consistency to separate normal nodes from anomalous ones. Generative approaches, exemplified by NetGAN [19], model graph connectivity through GAN-based random-walk generation, where anomalous nodes are often associated with low generation probability. Overall, these deep models better integrate structural and attribute cues, thereby improving the detection of complex anomalies.

2.2. Graph Data Augmentation

Data augmentation aims to improve data quality and/or scale by modifying the original data or generating synthetic samples [20,21]. In graph learning, augmentation is mainly realized by perturbing or replacing the graph structure and/or node attributes [22]. For example, ref [23,24] enhance graphs by adding or deleting edges, while [25] generates synthetic samples from the original data by analyzing target-oriented diversity and augmentation strength. Beyond structural perturbations, other studies optimize node representations via adversarial perturbations [26], mixup on node attributes [27], or feature generation methods [28]. In recent years, graph data augmentation has mainly focused on three aspects: first, property preservation/spectral domain controllability. Dual-Prism [29] proposes DP-Noise/DP-Mask based on spectral domain modeling, which prioritizes perturbations in the high-frequency domain while striving to preserve low-frequency and global properties. Second, task-adaptive structure shaping. TADA [30] extends features through structural embedding and performs joint topology/attribute sparsification.

SGA [31] targets signed graphs, constructing more favorable views via candidate edge addition/deletion and strategy selection. Third, generative/latent space and hybrid methods. G-CutMix [32] implements isomorphic shuffling combined with feature/label CutMix and fuses the two views after convolution.

However, merely adjusting topology and attributes to process graph data can cause information loss and degrade performance. Our augmentation strategy is denoising-based: we design a specialized anomaly-noise model to generate a new graph and align it with the original graph. This design enables our framework to capture the distributional regularities of graph data more precisely, thereby achieving significant gains on graph anomaly detection tasks.

3. Preliminaries

3.1. Problem Setup

Consider an attributed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathbf{X}\}$, where $\mathcal{V} = \{v_1, \dots, v_n\}$ is the node set, $\mathcal{E}$ is the edge set, and $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the attribute matrix. The structural information is represented by an adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$, where $A_{ij} = 1$ if an edge exists between v_i and v_j , otherwise 0. Graph anomaly detection aims to assign each node an anomaly score such that nodes deviating from the majority in structure and/or attributes receive higher scores.

3.2. Reconstruction-Based Graph Anomaly Detection

A reconstruction-based model (e.g., GAE/VGAE) learns an encoder-decoder parameterized by θ , mapping the input $(\mathbf{X}, \mathbf{A})$ to reconstructions $(\hat{\mathbf{X}}, \hat{\mathbf{A}})$: $f_\theta : (\mathbf{X}, \mathbf{A}) \rightarrow (\hat{\mathbf{X}}, \hat{\mathbf{A}})$. Training minimizes the average reconstruction loss:

$$\mathcal{L}(\theta; \mathcal{D}) = \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{x}_i, f_\theta(\mathbf{x}_i)), \quad (1)$$

where the training set $\mathcal{D} = \mathcal{D}_n \cup \mathcal{D}_a$ consists of normal and anomalous nodes. Let $N_n = |\mathcal{D}_n|$, $N_a = |\mathcal{D}_a|$, $N = N_n + N_a$, and $\rho := N_a / N$. Define the per-set average risks:

$$L_n(\theta) = \frac{1}{N_n} \sum_{i \in \mathcal{D}_n} \ell_i(\mathbf{x}_i, f_\theta(\mathbf{x}_i)), \quad L_a(\theta) = \frac{1}{N_a} \sum_{j \in \mathcal{D}_a} \ell_j(\mathbf{x}_j, f_\theta(\mathbf{x}_j)).$$

Then the empirical objective can be written as a mixture:

$$\mathcal{L}(\theta) = (1 - \rho) L_n(\theta) + \rho L_a(\theta). \quad (2)$$

Accordingly, the gradient decomposes as

$$\nabla_\theta \mathcal{L}(\theta) = (1 - \rho) \nabla_\theta L_n(\theta) + \rho \nabla_\theta L_a(\theta). \quad (3)$$

When ρ is small, updates are primarily driven by the normal part ∇L_n . Meanwhile, anomalous samples often incur larger reconstruction errors and thus may have larger per-sample gradients, but their overall influence is still weighted by ρ , potentially leading the optimizer to fit dominant normal patterns. We will later analyze how this contamination affects reconstruction-risk separability and how our model mitigates it (Section 5).

4. Methodology

Figure 1 shows the framework of CLEAN-GAD. The method contains two components. The first component is a data augmentation based on contrastive learning, which uses contrastive learning to extract local anomaly signals to remove anomalous nodes

from the graph. Then, augmented graph data can be obtained. The second component is a multi-view reconstruction learning model trained on the augmented graph and the original graph.




Figure 1. An overview of the CLEAN-GAD framework. It first uses contrastive learning (orange rectangle) to construct a clean graph and then employs multi-view alignment reconstruction learning (green rectangle) on the clean and original graphs, followed by anomaly score combination.

4.1. Graph Augmentation Learning ($\mathcal{L}_1$)

As discussed above, the anomalous nodes cause bias to the training objective. To reduce the bias introduced by anomalous nodes, we extract *local anomaly signals* via contrastive learning and construct a cleaner augmented graph by filtering high-confidence anomalies.

Local contrastive scoring. Following [33], we encode node attributes with an MLP:

$$\mathbf{H}_{local}^{(l)} = \sigma(\mathbf{H}_{local}^{(l-1)} \mathbf{W}^{(l-1)} + \mathbf{b}^{(l-1)}), \quad \mathbf{H}_{local}^{(0)} = \mathbf{X}, \quad (4)$$

where $\mathbf{h}_i$ denotes the final-layer embedding of node v_i . For each node v_i , we summarize its 1-hop neighborhood by mean pooling:

$$\mathbf{e}_p^i = \frac{1}{|\mathcal{N}(v_i)|} \sum_{t \in \mathcal{N}(v_i)} \mathbf{h}_t. \quad (5)$$

We sample M negative contexts $\mathbf{e}_{n,m}^i$ from neighborhoods of randomly selected nodes $v_j \neq v_i$. Using cosine similarity $\text{sim}(\mathbf{h}, \mathbf{e}) = \mathbf{h} \cdot \mathbf{e}^\top$ (with ℓ_2 normalization), we optimize an InfoNCE loss:

$$\mathcal{L}_1 = - \sum_{i=1}^N \log \frac{\exp(\text{sim}(\mathbf{h}_i, \mathbf{e}_p^i) / \tau)}{\exp(\text{sim}(\mathbf{h}_i, \mathbf{e}_p^i) / \tau) + \sum_{m=1}^M \exp(\text{sim}(\mathbf{h}_i, \mathbf{e}_{n,m}^i) / \tau)}. \quad (6)$$

A node is considered more likely to be anomalous if it is less consistent with its neighborhood, thus we define $s_i^{local} = 1 - \text{sim}(\mathbf{h}_i, \mathbf{e}_p^i)$. And then we rank nodes by s_i^{local} and select the top $k_{ano}\%$ as the high-confidence pseudo-anomalous set $\mathcal{V}_{ano}$. Removing $\mathcal{V}_{ano}$ and all incident edges yields a clean view $\mathcal{G}_{clean} = \{\mathcal{V}_c, \mathcal{E}_c, \mathbf{X}_c\}$, $\mathcal{V}_c = \mathcal{V} \setminus \mathcal{V}_{ano}$, $\mathcal{E}_c = \{(u, v) \in \mathcal{E} : u, v \in \mathcal{V}_c\}$.

4.2. Multi-View Alignment Reconstruction Learning ($\mathcal{L}_2$)

We jointly exploit the original graph $\mathcal{G}$ and a filtered (clean) view $\mathcal{G}_{clean}$ constructed via anomaly-aware filtering. The clean view suppresses high-confidence anomalous nodes and thus better preserves dominant normal patterns, while the original view retains the full distribution with additional noise and anomalous components. This discrepancy intro-

duces cross-view heterogeneity: a model optimized on $\mathcal{G}_{clean}$ alone may yield inconsistent anomaly scores when applied to $\mathcal{G}$. To improve robustness, we propose a multi-view reconstruction framework with (i) a frequency-adaptive encoder that explicitly preserves both low- and high-frequency signals and (ii) a score-alignment objective to encourage view-invariant node-level anomaly judgments. Moreover, the high-confidence anomaly set $\mathcal{V}_{ano}$ provides a pseudo-supervised cue to calibrate frequency fusion and enhance score separability.

4.2.1. Frequency-Adaptive Encoder

Message passing in common GNNs (e.g., GCN-style propagation) is closely related to Laplacian smoothing, which tends to emphasize low-frequency components and may oversmooth node representations. However, anomalies often manifest as localized, high-frequency variations in graph signals, and excessive low-pass bias can attenuate such cues. We therefore construct a two-channel encoder to extract low-/high-frequency signals and fuse them adaptively at the node level.

Multi-channel graph filters. Let $\mathbf{A} \in \mathbb{R}^{N \times N}$ be the adjacency matrix. We add self-loops $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ and define $\hat{\mathbf{A}} = \tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2}$ with $\tilde{\mathbf{D}}_{ii} = \sum_j \tilde{\mathbf{A}}_{ij}$. The normalized Laplacian is $\mathbf{L} = \mathbf{I} - \hat{\mathbf{A}}$. We adopt a low-pass and a high-pass filter:

$$\theta_l = \mathbf{I} + \hat{\mathbf{A}}, \quad \theta_h = \mathbf{L}, \quad (7)$$

where θ_l is GCN-like and mainly captures smooth (low-frequency) signals, while θ_h preserves high-frequency variations.

Anomaly-aware adaptive fusion. Let $\mathbf{H}_l$ and $\mathbf{H}_h$ denote the embeddings produced by the low-/high-pass channels, respectively. Instead of using a fixed combination, we compute node-wise fusion weights guided by anomaly evidence. Specifically, we aggregate embeddings of nodes in the high-confidence anomaly set $\mathcal{V}_{ano}$ to form a global query vector $\mathbf{q}$, which summarizes characteristic anomaly-related patterns:

$$\mathbf{q} = \text{Readout}(\{\mathbf{h}_i : v_i \in \mathcal{V}_{ano}\}) = \frac{1}{|\mathcal{V}_{ano}|} \sum_{v_i \in \mathcal{V}_{ano}} \mathbf{h}_i,$$

where $\text{Readout}(\cdot)$ is a permutation-invariant pooling operator, and we adopt mean pooling in this work. We then compute attention weights over the two frequency channels for each node:

$$\alpha_i^{(m)} = \frac{\exp(w_i^{(m)})}{\sum_{m' \in \{0,1\}} \exp(w_i^{(m')})}, \quad w_i^{(m)} = \frac{\mathbf{q}^\top \mathbf{k}_i^{(m)}}{\sqrt{d}}, \quad m \in \{0,1\}, \quad (8)$$

where $\mathbf{k}_i^{(0)}$ and $\mathbf{k}_i^{(1)}$ are learnable key projections of the i -th node from $\mathbf{H}_l$ and $\mathbf{H}_h$ (and d is the embedding dimension).

The final representation is obtained by channel-wise weighted fusion:

$$\mathbf{z}_i = \alpha_i^{(0)} \mathbf{v}_i^{(0)} + \alpha_i^{(1)} \mathbf{v}_i^{(1)}, \quad (9)$$

where $\mathbf{v}_i^{(0)}$ and $\mathbf{v}_i^{(1)}$ are value projections from the two channels. This mechanism allows anomalous nodes to retain more high-frequency components when they better match the anomaly query, while normal nodes may prefer smoother low-frequency representations.

4.2.2. Multi-View Reconstruction Learning

We treat $\mathcal{G}$ and $\mathcal{G}_{\text{clean}}$ as two complementary views, indexed by $v \in \{\text{ori}, \text{clean}\}$. Unless otherwise specified, both views are processed by our frequency-adaptive encoder (Section 4.2.1), while reconstruction is performed using a conventional decoder following prior work (an inner-product decoder for structure and an MLP decoder for attributes), resulting in a unified anomaly scoring function:

$$\mathbf{Z}^{(v)} = \text{Enc}(\mathcal{G}^{(v)}), \quad \hat{\mathbf{X}}^{(v)} = \text{Dec}(\mathbf{Z}^{(v)}), \quad \hat{\mathbf{A}}^{(v)} = \sigma(\mathbf{Z}^{(v)}(\mathbf{Z}^{(v)})^\top). \quad (10)$$

Based on the reconstructions, we compute node-wise reconstruction errors for both attributes and structure:

$$s_i^{\text{global},(v)} = \alpha \|\hat{\mathbf{X}}_i^{(v)} - \mathbf{X}_i^{(v)}\|_2^2 + (1 - \alpha) \|\hat{\mathbf{A}}_i^{(v)} - \mathbf{A}_i^{(v)}\|_2^2, \quad (11)$$

where $\alpha \in [0, 1]$ balances attribute and structure contributions. The reconstruction objective is averaged across views:

$$\mathcal{L}_{\text{Re}} = (\mathcal{L}_{\text{Re}}^{\text{ori}} + \mathcal{L}_{\text{Re}}^{\text{clean}}) / 2. \quad (12)$$

Score Alignment. We consider the original graph $\mathcal{G}$ and the filtered, clean view $\mathcal{G}_{\text{clean}}$ as two complementary views. The clean view suppresses high-confidence anomalous nodes and thus better preserves dominant normal patterns, whereas the original view reflects the full (and potentially contaminated) data distribution with additional noise and anomalous components. This cross-view discrepancy may cause inconsistent anomaly scoring if the model is optimized on a single view. To mitigate such a distribution shift and improve scoring robustness, we explicitly align the node-level anomaly-score vectors predicted from the two views by minimizing the squared ℓ_2 distance:

$$\mathcal{L}_{\text{align}} = \left(\mathbf{s}^{(\text{ori})} - \mathbf{s}^{(\text{clean})} \right)^\top \left(\mathbf{s}^{(\text{ori})} - \mathbf{s}^{(\text{clean})} \right), \quad (13)$$

where $\mathbf{s}^{(v)} = [s_1^{\text{global},(v)}, \dots, s_N^{\text{global},(v)}]^\top$ denotes the anomaly-score vector computed on view $v \in \{\text{ori}, \text{clean}\}$. Aligning scores—rather than embeddings—encourages view-invariant anomaly judgments while allowing latent representations to remain adaptive to view-specific noise and structural variations. Moreover, the high-confidence anomaly set $\mathcal{V}_{\text{ano}}$ provides a pseudo-supervised cue that facilitates score calibration across views.

Pseudo-Label Regularization. To better leverage pseudo-labeled nodes and promote score separability between normal and pseudo-anomalous nodes, we introduce a distribution-level regularizer. For each node v_i , we define the score used for regularization by averaging the local score and the reconstruction-based score:

$$s_i = \frac{s_i^{\text{local}} + s_i^{\text{global}}}{2}. \quad (14)$$

Given the score sets on the clean (normal) nodes V_c and the pseudo-anomalous nodes V_{ano} , we follow [34] and approximate their score distributions by two Gaussians, $P = \mathcal{N}(\mu_n, \sigma_n^2)$ and $Q = \mathcal{N}(\mu_a, \sigma_a^2)$, respectively. We quantify their discrepancy using the symmetric KL divergence:

$$D_{\text{KL}}^{\text{sym}}(P, Q) = \frac{1}{2} (D_{\text{KL}}(P \| Q) + D_{\text{KL}}(Q \| P)). \quad (15)$$

To encourage a sufficiently large separation, we maximize this divergence via a margin-based hinge loss:

$$\mathcal{L}_{reg} = \max(0, r - D_{KL}^{sym}(P, Q)), \quad (16)$$

where r is a margin hyperparameter controlling when the penalty becomes active. This regularizer discourages distribution collapse between V_c and V_{ano} and stabilizes training under noisy pseudo labels.

4.2.3. Joint Optimization

By integrating the reconstruction objective, the distribution-separation regularization, and the cross-view score alignment, we obtain the overall training loss:

$$\mathcal{L}_2 = \mathcal{L}_{Re} + \beta \mathcal{L}_{reg} + \gamma \mathcal{L}_{align}. \quad (17)$$

where β, γ is the weight parameter of $\mathcal{L}_{reg}$ and $\mathcal{L}_{align}$, respectively. After training based on this loss function, we obtain a new anomaly score s^{global} . Finally, we integrate the anomaly scores of the two modules to obtain the final anomaly score $s = \frac{1}{2}(s^{local} + s^{global})$.

5. Model Analysis

This section provides a focused analysis of how reducing the anomaly ratio in training data improves reconstruction-based graph anomaly detection. Our analysis isolates the effect of anomaly-aware filtering (the clean view construction) by characterizing how it reduces the effective contamination ratio and, consequently, enlarges the separability between normal and anomalous reconstruction risks.

5.1. Contaminated Risk Under a Mixture Model

Assumption 1 (Mixture data-generating distribution). *The training samples follow a two-component mixture distribution*

$$P_\rho = (1 - \rho)P_n + \rho P_a,$$

where P_n and P_a denote the normal and anomalous distributions, respectively, and $\rho \in [0, 1]$ is the anomaly ratio.

Let f_θ be a reconstruction model parameterized by the vector θ , and let $\ell(x, f_\theta(x))$ be the reconstruction loss for a vector-valued sample x . Define the population reconstruction risks

$$R_n(\theta) = \mathbb{E}_{x \sim P_n} [\ell(x, f_\theta(x))], \quad R_a(\theta) = \mathbb{E}_{x \sim P_a} [\ell(x, f_\theta(x))],$$

and the contaminated (mixture) risk

$$R(\theta; \rho) = (1 - \rho)R_n(\theta) + \rho R_a(\theta). \quad (18)$$

Denote $\theta^* = \arg \min_\theta R_n(\theta)$ as the minimizer under a clean normal training distribution ($\rho = 0$), and $\tilde{\theta}(\rho) = \arg \min_\theta R(\theta; \rho)$ as the minimizer under contamination ratio ρ .

5.2. Separability Degradation Induced by Contamination

To quantify how contamination affects anomaly detectability, we consider the reconstruction-risk separability

$$\Delta(\rho) = R_a(\tilde{\theta}(\rho)) - R_n(\tilde{\theta}(\rho)).$$

Assumption 2 (Local smoothness and strong convexity). $R(\theta; \rho)$ is twice continuously differentiable in a neighborhood of $(\theta^*, 0)$. The Hessian $\mathbf{H} = \nabla_{\theta}^2 R_n(\theta^*)$ is positive definite.

Proposition 1 (First-order shrinkage of reconstruction-risk gap). Under Assumptions 1 and 2, if $\mathbf{g}_a = \nabla_{\theta} R_a(\theta^*) \neq \mathbf{0}$, then for sufficiently small ρ ,

$$\Delta(\rho) = \Delta(0) - c\rho + O(\rho^2), \quad c = \mathbf{g}_a^{\top} \mathbf{H}^{-1} \mathbf{g}_a > 0. \quad (19)$$

Proof. A standard perturbation argument yields $\tilde{\theta}(\rho) = \theta^* - \rho \mathbf{H}^{-1} \mathbf{g}_a + O(\rho^2)$. Substituting into $R_n(\cdot)$ and $R_a(\cdot)$ via Taylor expansion around θ^* gives $R_n(\tilde{\theta}(\rho)) - R_n(\theta^*) = \frac{1}{2} \rho^2 \mathbf{g}_a^{\top} \mathbf{H}^{-1} \mathbf{g}_a + O(\rho^3)$ and $R_a(\tilde{\theta}(\rho)) - R_a(\theta^*) = -\rho \mathbf{g}_a^{\top} \mathbf{H}^{-1} \mathbf{g}_a + O(\rho^2)$. Taking the difference yields Equation (19). $\square$

Equation (19) formalizes that anomaly contamination reduces the separability gap $\Delta(\rho)$ at a first-order rate, implying that the learned reconstructor increasingly fits anomalous patterns as ρ grows.

5.3. Filtering Reduces the Effective Contamination Ratio

CLEAN-GAD constructs a clean view by removing a high-confidence pseudo-anomalous set $\mathcal{V}_{ano}$. We characterize the filtering quality by:

$$\pi_a = \Pr(v \in \mathcal{V}_{ano} \mid v \in \mathcal{V}_a) \quad (\text{TPR}), \quad \pi_n = \Pr(v \in \mathcal{V}_{ano} \mid v \in \mathcal{V}_n) \quad (\text{FPR}).$$

After removing $\mathcal{V}_{ano}$, the remaining anomaly ratio in the clean view becomes

$$\rho_{\text{clean}} = \frac{(1 - \pi_a)\rho}{(1 - \pi_a)\rho + (1 - \pi_n)(1 - \rho)}. \quad (20)$$

When the filtering is informative ($\pi_a > \pi_n$), we have $\rho_{\text{clean}} < \rho$, i.e., the training distribution is less contaminated.

Combining Equation (19) with Equation (20) yields the direct implication:

$$\Delta(\rho_{\text{clean}}) = \Delta(0) - c\rho_{\text{clean}} + O(\rho_{\text{clean}}^2) > \Delta(\rho) = \Delta(0) - c\rho + O(\rho^2), \quad (21)$$

for small ρ and informative filtering, meaning the clean view preserves a larger reconstruction-risk gap.

5.4. Connection to Empirical Metrics

Reconstruction-based detectors typically score nodes by reconstruction error (or a monotone function of it). Let random scores S_n and S_a denote the anomaly scores of a normal and an anomalous node, respectively. The AUC can be written as the pairwise ranking probability

$$\text{AUC} = \Pr(S_a > S_n).$$

A larger separation between the score distributions (e.g., a larger mean gap and/or reduced overlap) increases $\Pr(S_a > S_n)$, thereby improving AUC; similarly, reduced overlap typically improves precision-oriented metrics such as AP and Recall@K. Equation (21) predicts that learning on a view with $\rho_{\text{clean}} < \rho$ increases the normal–anomaly score separability, which explains the empirical gains observed after anomaly-aware filtering.

Overall, the above analysis indicates that anomaly-aware filtering improves reconstruction-based detection primarily by reducing the effective contamination ratio from ρ to ρ_{clean} , which preserves a larger reconstruction-risk gap between anomalous and normal nodes. Consequently, the learned reconstruction errors become more discriminative, leading to more reliable anomaly ranking and improved AUC/AP-type metrics in practice.

6. Experiments

6.1. Datasets and Baselines

We conduct experiments on seven benchmark datasets—four with synthetic (injected) anomalies and three with real anomalies. Table 1 summarizes the key statistics of these datasets. The injected-anomaly datasets are Cora, Citeseer [35], BlogCatalog [36], and Flickr [9]; the real-anomaly datasets are Amazon [37], Reddit [38], and Facebook [18]. For reproducibility, we adopt the datasets used in [39]. Following prior work [9,17], anomalies are injected in two ways: (i) attribute perturbation by swapping node features to create contextual inconsistencies, and (ii) structural perturbation by rewiring edges to introduce topological irregularities while preserving the counts and category proportions of each type.

We compare our method with three traditional graph anomaly detection methods: SCAN [40], Radar [41], and ANOMALOUS [14], as well as several GNN-based models: DOMINANT [9], GCNAE [42], DONE, AdONE [43], AnomalyDAE [16], OCGNN [44], GADNR [45], CoLA [17], CONAD [18], and ADA-GAD [10].

Table 1. Statistics of datasets.

Dataset	# Nodes	# Edges	# Anomaly	%Anomaly
Cora	2708	5429	150	5.53%
CiteSeer	3327	4732	150	4.50%
Amazon	10,244	175,608	693	6.76%
BlogCatalog	5196	171,743	300	5.77%
Flickr	7575	239,738	450	5.94%
Facebook	1081	55,104	25	2.31%
Reddit	10,984	168,016	366	3.33%

denotes the number of the corresponding items.

6.2. Evaluation Metrics and Implementation Details

We use two evaluation metrics, i.e., ROC-AUC and Average Precision (AP). We set the anomaly ratio term $k_{ano} = 1$. The trade-off for attribute vs. structural reconstruction losses is set to $\alpha = 0.2$. The weights for the KL-regularization term β and the multi-view alignment term γ are set to 0.3 and 0.7, respectively. For baselines, we rely on the PyGOD toolkit [46] and keep the default hyperparameters; when PyGOD implementations are unavailable, we use the authors' official codebases.

6.3. Comparative Study

We evaluate CLEAN-GAD against various baseline methods. The ROC-AUC results, averaged over 10 independent runs, are reported in Table 2, where we also perform Welch's *t*-test to evaluate the statistical significance of the improvements. The corresponding ROC curves are illustrated in Figure 2. Based on the results, we come to the following conclusions:

- Deep models substantially outperform non-deep baselines, and CLEAN-GAD further improves over strong deep baselines (e.g., DOMINANT/CONAD/ADONE) with consistent gains on most datasets. Notably, on datasets such as Cora, Amazon, Reddit, and Flickr, the performance gains are statistically significant ($p < 0.05$), demonstrating the clear effectiveness of our proposed framework.
- CLEAN-GAD exhibits strong cross-dataset generalization: it achieves the best or near-best performance across heterogeneous graphs, covering both highly homophilous citation networks (Cora/Citeseer) and more complex, weakly homophilous social graphs (Reddit/Facebook). On Flickr, it remains competitive with only a marginal gap to the strongest baseline.

- CLEAN-GAD is more stable across runs, with lower or comparable standard deviations, indicating improved robustness to random initialization and stochastic sampling and thus better reproducibility. Moreover, as shown in Table 3, CLEAN-GAD is the second-fastest method on both Amazon (0.0200s) and Flickr (0.0430s), offering an excellent accuracy–efficiency trade-off compared with the fastest but less accurate GADAM.

Overall, the significantly higher average ROC-AUC, low variance, and consistent performance across datasets highlight the robustness and generalization capability of CLEAN-GAD for graph anomaly detection. The efficiency results further confirm that CLEAN-GAD attains state-of-the-art AUROC with practical runtime on both medium- and large-scale graphs.




Figure 2. ROC curves of our method and several representative baselines on seven datasets.

Table 2. Performance on seven datasets (%). Bold indicates the best, and underline indicates the second-best. * indicates that the improvement of ours over the best baseline is statistically significant ($p < 0.05$) via Welch’s t -test.

	Algorithm	Cora (%)	Citeseer (%)	Amazon (%)	Reddit (%)	BlogCatalog (%)	Flickr (%)	Facebook (%)
Non-Deep	Scan	72.24±2.97	73.30±1.34	53.26±1.80	49.60±0.92	53.30±5.21	61.51±9.07	44.41±10.69
	Radar	52.13±0.84	57.65±3.48	54.92±0.62	51.88±1.43	62.98±1.94	60.87±0.50	20.91±11.69
	ANOMALOUS	49.67±1.71	42.30±0.06	56.04±3.11	50.62±0.00	63.04±10.11	62.05±1.07	29.34±20.04
Deep	DOMINANT	80.50±2.55	86.25±3.41	52.88±0.77	56.07±0.09	68.25±3.32	64.31±2.54	10.68±0.92
	GCNAE	73.25±0.00	75.78±0.00	63.98±1.60	55.38±2.82	74.53±0.00	74.76±0.01	49.81±0.00
	DONE	<u>90.09±2.73</u>	<u>92.42±1.19</u>	72.23±7.03	<u>57.64±1.44</u>	74.68±0.90	79.85±0.76	14.85±2.52
	ADONE	85.61±2.30	89.31±1.57	67.46±11.52	53.21±4.20	74.79±2.38	79.42±0.61	11.19±1.62
	AnomalyDAE	74.48±9.43	82.61±4.33	68.59±3.95	55.71±0.49	75.31±0.72	75.46±0.95	51.65±4.30
	OCCGN	50.38±0.85	50.00±0.00	49.26±5.15	46.05±3.42	50.79±1.42	<u>81.59±0.74</u>	78.98±4.34
	GANDR	72.47±1.58	74.96±0.04	<u>80.82±1.74</u>	53.45±3.94	50.57±2.42	51.60±1.33	41.83±22.73
	CoLA	51.27±3.10	63.48±2.22	50.34±11.54	50.00±0.00	51.90±1.96	81.20±1.82	77.09±2.09
	CONAD	76.13±14.10	74.01±15.17	52.81±2.86	55.67±1.54	67.20±7.99	69.77±6.66	62.63±5.88
	GADAM	84.09±1.62	91.45±1.67	75.00±9.51	53.75±0.64	72.07±1.85	69.77±6.66	<u>86.63±5.88</u>
	ADA-GAD	86.05±0.01	87.24±0.10	52.37±0.07	56.10±0.01	OOM	OOM	61.89±0.32
	Ours	92.39±0.64 *	93.01±0.78	83.21±2.82 *	60.06±2.34 *	<u>75.04±2.86</u>	82.09±1.47 *	87.31±2.88

Table 3. Efficiency and performance comparison. Bold indicates the best result, and underline indicates the second-best.

Algorithm	Amazon (Medium Scale)			Flickr (Large Scale)		
	Time (s)	RAM (MB)	AUROC	Time (s)	RAM (MB)	AUROC
GADAM	0.0014	72.2	75.00	0.0024	<u>798.0</u>	69.77
DONE	0.0210	2587.4	72.23	3.0636	631.4	<u>79.85</u>
ADA-GAD	19.0876	4019.5	52.37	OOM	OOM	OOM
CLEAN-GAD (Ours)	<u>0.0200</u>	<u>953.0</u>	83.21	<u>0.0430</u>	3347.5	82.09

6.4. Parameter Analysis

6.4.1. Embedding Dimension

Figure 3 investigates the impact of embedding dimensionality in the reconstruction and contrastive branches. Overall, both branches exhibit a clear performance optimum. The reconstruction branch performs best with a moderate dimensionality (around 64), whereas the contrastive branch is comparatively more stable in the 64–128 range. Increasing the dimensionality beyond these regimes yields diminishing returns and can even degrade AUROC, indicating that larger representational capacity does not necessarily improve anomaly separability.

The sensitivity to dimensionality is also dataset-dependent. Flickr is notably prone to over-parameterization, where performance decreases once the embedding size exceeds 128. In contrast, Amazon and BlogCatalog show a smoother “improve–plateau–decline” trend as the dimension increases. This suggests that the optimal dimension is governed by both structural characteristics (e.g., homophily/heterophily) and noise levels. In practice, compact representations tend to generalize more reliably by mitigating noise amplification and reducing the risk of overfitting.

From a modeling perspective, an oversized reconstruction branch may capture incidental patterns and noise, thereby shrinking the reconstruction-error gap between normal and anomalous nodes. Similarly, in the contrastive branch, high-dimensional spaces can make effective negatives sparse; under fixed temperature and negative-sampling settings, this may weaken discriminative margins due to an imbalance between alignment and uniformity. Taken together, we recommend fixing the reconstruction dimension to 64 and selecting the contrastive dimension within 64–128 to balance accuracy, stability, and computational efficiency.




Figure 3. Effect of embedding dimensionality on AUROC.

6.4.2. Weight Parameters

We further analyze the sensitivity to the loss weights β and γ (Figure 4), which control the pseudo-label regularization and the cross-view score alignment, respectively. Across datasets, β primarily determines the overall performance regime, while γ mainly provides local refinement within a suitable β range. Specifically, when β is small to moderately small, the model typically operates in a high-performance region, and tuning γ yields a clear local optimum with relatively stable gains. In contrast, overly large β drives the model into a lower-performance regime, within which the influence of γ becomes marginal and may even lead to minor degradations.

This “region-of-optimality” pattern is consistently observed on both Amazon and BlogCatalog: a moderate β is crucial for placing the model in a favorable baseline regime, after which γ can be used to fine-tune performance. Accordingly, a practical tuning strategy is to first identify a moderate β that yields strong performance, and then conduct a local search over γ within that regime. This joint tuning procedure is more effective than increasing either parameter in isolation and tends to produce more robust performance across datasets.




Figure 4. AUROC under different combinations of β and γ . Lighter colors represent higher detection accuracy (higher AUROC values).

6.5. Ablation Study

6.5.1. Effect of Components on Detection Performance

To assess the contribution of each module, we construct three ablated variants under the same backbone and training protocol and compare them with the full model: (i) **w/o denoise**: remove the stage-1 anomaly-aware filtering/denoising augmentation and train directly on the original graph; (ii) **w/o align**: remove the cross-view score-alignment loss and optimize only the reconstruction objective; (iii) **w/o reg**: remove the pseudo-label distribution regularization (implemented via symmetric KL divergence); (iv) **CLEAN-GAD**: the complete CLEAN-GAD.

Figure 5 yields three main observations.

- **Anomaly-aware denoising is the most influential component.** Removing denoising leads to the largest AUROC drop on multiple datasets (e.g., Flickr and Amazon), indicating that constructing a cleaner training view effectively reduces anomaly contamination and helps preserve a larger reconstruction gap between normal and anomalous nodes. This trend is consistent with our theoretical analysis: reducing contaminated samples mitigates the parameter shift induced by anomalous gradients during optimization.
- **Score alignment consistently improves cross-view robustness.** Disabling the alignment term degrades AUROC across datasets (with the largest drop observed on Reddit), suggesting that aligning node-level anomaly scores encourages view-invariant scoring under cross-view distribution shift and alleviates anomaly-signal attenuation caused by neighborhood homogenization.
- **Distribution regularization enhances separability, especially in AUPRC.** Removing the distribution regularizer results in a clear degradation in AUPRC (notably on Amazon and Flickr), implying that explicitly enforcing score separation between clean nodes and pseudo-anomalies improves the detection of rare anomalies and stabilizes training under noisy pseudo labels.

Overall, the full model achieves the best performance by jointly leveraging denoising augmentation, cross-view score alignment, and pseudo-label distribution regularization, validating the effectiveness of our two-stage design in mitigating anomaly overfitting and homophily-induced anomaly attenuation.




Figure 5. Ablation results of CLEAN-GAD and its variants.

6.5.2. Effect of k_{ano} on Anomaly Ratio and Detection Performance

We further examine how the anomaly-filtering budget k_{ano} affects the effective anomaly contamination in the clean view and its impact on detection performance. Let AR denote the anomaly ratio in the remaining node set after filtering: $AR = \frac{|\mathcal{V}'_a|}{|\mathcal{V}'|}$, where $\mathcal{V}'$ is the node set of the filtered graph (i.e., $\mathcal{G}_{clean}$) and $\mathcal{V}'_a \subseteq \mathcal{V}'$ denotes the anomalous nodes within $\mathcal{V}'$.

As shown in Figure 6, increasing k_{ano} generally decreases AR on all datasets, confirming that the proposed filtering mechanism progressively reduces anomaly contamination in the training view. Regarding detection performance, AUROC exhibits dataset-dependent sensitivity to k_{ano} . On Amazon, reducing AR yields a clear performance gain, supporting our motivation that lowering contamination mitigates anomaly overfitting and enlarges the reconstruction gap. On Flickr and Reddit, AUROC changes within a relatively narrow range as AR decreases, suggesting that once the training view becomes sufficiently clean, further filtering brings limited marginal benefit and the remaining errors may be dominated by dataset-specific noise or ambiguity.




Figure 6. Sensitivity analysis of k_{ano} : effective anomaly ratio (AR) and AUROC.

Overall, a larger k_{ano} reliably produces a cleaner effective training distribution, while the resulting AUROC improvement depends on how strongly each dataset is affected by anomaly contamination and cross-view distribution shift.

6.5.3. Effect of Frequency-Adaptive Encoder on Anomaly Overfitting

To verify the effectiveness of our model in alleviating the *anomaly overfitting* issue, we conduct a reconstruction-error analysis experiment on **Flickr** by tracking the reconstruction errors of *normal* and *anomalous* nodes over training epochs. We consider two variants under an identical pipeline: **CLEAN-GAD** (our full method) and a controlled baseline in which we only replace the encoder with a standard **GCN** while keeping all other components and the reconstruction objective unchanged. For both variants, we visualize the reconstruction-error distributions using boxplots at representative epochs $\{1, 41, 81, 120, 160, 200\}$, as shown in Figure 7.

As shown in Figure 7b, the GCN-encoder variant exhibits a rapid decrease in reconstruction error for both normal and anomalous nodes. After the early training stage (around epoch 41), the reconstruction errors of the two groups collapse to near-zero values with highly similar medians and interquartile ranges. This behavior indicates that the model becomes increasingly capable of reconstructing anomalous nodes almost as well as normal nodes, thereby shrinking the reconstruction-error contrast that reconstruction-based detectors rely on.

In contrast, Figure 7a demonstrates that CLEAN-GAD maintains a clear and persistent separation between the reconstruction errors of normal and anomalous nodes throughout training. While the reconstruction error of normal nodes remains within a relatively narrow range, anomalous nodes consistently exhibit higher errors and substantially larger dispersion, with noticeably heavier upper tails. The sustained gap suggests that CLEAN-GAD fits normal patterns without indiscriminately reducing the reconstruction error of anomalies to the same level, thus preserving the discriminative signal for anomaly detection during optimization.

Overall, this learning-curve analysis reveals distinct training dynamics between the two variants: the GCN-encoder baseline tends to equalize reconstruction errors for normal and anomalous nodes, whereas CLEAN-GAD consistently preserves reconstruction-error separability, which is essential for robust reconstruction-based graph anomaly detection.




Figure 7. Learning-curve analysis of reconstruction error on Flickr. Boxplots show reconstruction errors for normal vs. anomalous nodes at epochs {1, 41, 81, 120, 160, 200} for (a) CLEAN-GAD and (b) the GCN-encoder variant (all other components unchanged).

7. Conclusions

This paper identifies and addresses two fundamental challenges in reconstruction-based unsupervised graph anomaly detection (GAD): anomaly overfitting and homophily-induced attenuation. We proposed CLEAN-GAD, a contamination-aware framework that mitigates these issues through a dual-pronged approach. By integrating a contrastive augmentation module with a multi-view reconstruction scheme, CLEAN-GAD effectively reduces the influence of anomalies during training while preserving localized high-frequency signals via a frequency-adaptive encoder. Our theoretical analysis provides a principled foundation for this approach, demonstrating that reducing the effective anomaly ratio (ρ) linearly expands the reconstruction-risk gap, thereby enhancing separability. Empirical evaluations across multiple benchmarks confirm that CLEAN-GAD achieves superior anomaly detection performance. Future research will focus on extending this contamination-aware logic to complex structural and subgraph anomalies, as well as adapting the framework for large-scale dynamic graph environments.

Author Contributions: Q.L.: writing—original draft, software, visualization, methodology, formal analysis, investigation. P.Z.: writing—supervision, funding acquisition. Q.T.: writing—review and editing, validation, resources, data curation. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported in part by the National Natural Science Foundation of China (U2336202).

Data Availability Statement: The data and code presented in this study are openly available on GitHub at <https://github.com/user172894254/CLEAN-GAD>, accessed on 27 January 2026.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Li, R.; Liu, Z.; Ma, Y.; Yang, D.; Sun, S. Internet financial fraud detection based on graph learning. *IEEE Trans. Comput. Soc. Syst.* **2022**, *10*, 1394–1401. [[CrossRef](#)]
- Fan, H.; Zhang, F.; Li, Z. Anomalydae: Dual autoencoder for anomaly detection on attributed networks. In *Proceedings of the ICASSP 2020—2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*; IEEE: New York, NY, USA, 2020; pp. 5685–5689.
- Duan, J.; Wang, S.; Zhang, P.; Zhu, E.; Hu, J.; Jin, H.; Liu, Y.; Dong, Z. Graph anomaly detection via multi-scale contrastive learning networks with augmented view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Washington, DC, USA, 7–14 February 2023; Volume 37, pp. 7459–7467.
- Bian, T.; Xiao, X.; Xu, T.; Zhao, P.; Huang, W.; Rong, Y.; Huang, J. Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, New York, NY, USA, 7–12 February 2020; Volume 34, pp. 549–556.
- Cheng, D.; Wang, X.; Zhang, Y.; Zhang, L. Graph neural network for fraud detection via spatial-temporal attention. *IEEE Trans. Knowl. Data Eng.* **2020**, *34*, 3800–3813. [[CrossRef](#)]
- Dou, Y.; Liu, Z.; Sun, L.; Deng, Y.; Peng, H.; Yu, P.S. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, Virtual Event, Ireland, 19–23 October 2020; pp. 315–324.
- Ma, X.; Liu, F.; Wu, J.; Yang, J.; Xue, S.; Sheng, Q.Z. Rethinking unsupervised graph anomaly detection with deep learning: Residuals and objectives. *IEEE Trans. Knowl. Data Eng.* **2024**, *37*, 881–895. [[CrossRef](#)]
- Liang, S.; Xie, Q.; Feng, J.; Yue, Y.; Li, H.; Ye, J.; Jie, W. Feature Reconstruction for Anomaly Detection on Directed Multigraphs: A Preprocessing Framework for GNNs. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, Toronto, ON, Canada, 3–7 August 2025; pp. 1707–1717.
- Ding, K.; Li, J.; Bhanushali, R.; Liu, H. Deep anomaly detection on attributed networks. In *Proceedings of the 2019 SIAM International Conference on Data Mining*; SIAM: Bangkok, Thailand, 2019; pp. 594–602.
- He, J.; Xu, Q.; Jiang, Y.; Wang, Z.; Huang, Q. Ada-gad: Anomaly-denoised autoencoders for graph anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*; AAAI Press: Washington, DC, USA, 2024; Volume 38, pp. 8481–8489.
- Chai, Z.; You, S.; Yang, Y.; Pu, S.; Xu, J.; Cai, H.; Jiang, W. Can abnormality be detected by graph neural networks? In *Proceedings of the IJCAI*, Vienna, Italy, 23–29 July 2022; pp. 1945–1951.
- Tang, J.; Li, J.; Gao, Z.; Li, J. Rethinking graph neural networks for anomaly detection. In *Proceedings of the International Conference on Machine Learning*; PMLR: New York, NY, USA, 2022; pp. 21076–21089.
- Akoglu, L.; McGlohon, M.; Faloutsos, C. Oddball: Spotting anomalies in weighted graphs. In *Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining*; Springer: Berlin/Heidelberg, Germany, 2010; pp. 410–421.
- Peng, Z.; Luo, M.; Li, J.; Liu, H.; Zheng, Q. ANOMALOUS: A joint modeling approach for anomaly detection on attributed networks. In *Proceedings of the IJCAI*, Stockholm, Sweden, 13–19 July 2018; Volume 18, pp. 3513–3519.
- Henderson, K.; Gallagher, B.; Eliassi-Rad, T.; Tong, H.; Basu, S.; Akoglu, L.; Koutra, D.; Faloutsos, C.; Li, L. Rolx: Structural role extraction & mining in large graphs. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Beijing, China, 12–16 August 2012; pp. 1231–1239.
- Sakurada, M.; Yairi, T. Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the MLSDA 2014 2nd Workshop on Machine Learning for Sensory Data Analysis*, Gold Coast, Australia, 2 December 2014; pp. 4–11.
- Liu, Y.; Li, Z.; Pan, S.; Gong, C.; Zhou, C.; Karypis, G. Anomaly detection on attributed networks via contrastive self-supervised learning. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *33*, 2378–2392. [[CrossRef](#)] [[PubMed](#)]
- Xu, Z.; Huang, X.; Zhao, Y.; Dong, Y.; Li, J. Contrastive attributed network anomaly detection with data augmentation. In *Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining*; Springer: Berlin/Heidelberg, Germany, 2022; pp. 444–457.
- Bojchevski, A.; Shchur, O.; Zügner, D.; Günnemann, S. Netgan: Generating graphs via random walks. In *Proceedings of the International Conference on Machine Learning*; PMLR: New York, NY, USA, 2018; pp. 610–619.
- Ding, K.; Xu, Z.; Tong, H.; Liu, H. Data augmentation for deep graph learning: A survey. *ACM SIGKDD Explor. Newsl.* **2022**, *24*, 61–77. [[CrossRef](#)]
- Liu, F.; Ma, X.; Wu, J.; Yang, J.; Xue, S.; Beheshti, A.; Zhou, C.; Peng, H.; Sheng, Q.Z.; Aggarwal, C.C. Dagad: Data augmentation for graph anomaly detection. In *Proceedings of the 2022 IEEE International Conference on Data Mining (ICDM)*; IEEE: New York, NY, USA, 2022; pp. 259–268.

22. Wang, Y.; Wang, W.; Liang, Y.; Cai, Y.; Liu, J.; Hooi, B. NodeAug: Semi-Supervised Node Classification with Data Augmentation. In Proceedings of the KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual, 23–27 August 2020.
23. Zhao, T.; Liu, Y.; Neves, L.; Woodford, O.; Jiang, M.; Shah, N. Data augmentation for graph neural networks. In Proceedings of the AAAI Conference on Artificial Intelligence, Online, 2–9 February 2021; Volume 35, pp. 11015–11023.
24. Zhu, Y.; Xu, Y.; Yu, F.; Liu, Q.; Wu, S.; Wang, L. Graph contrastive learning with adaptive augmentation. In Proceedings of the Web Conference 2021, Ljubljana, Slovenia, 19–23 April 2021; pp. 2069–2080.
25. Park, H.; Lee, S.; Kim, S.; Park, J.; Jeong, J.; Kim, K.M.; Ha, J.W.; Kim, H.J. Metropolis-hastings data augmentation for graph neural networks. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 19010–19020.
26. Kong, K.; Li, G.; Ding, M.; Wu, Z.; Zhu, C.; Ghanem, B.; Taylor, G.; Goldstein, T. Robust Optimization as Data Augmentation for Large-Scale Graphs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 19–24 June 2022; IEEE: New York, NY, USA, 2022; pp. 60–69.
27. Wang, Y.; Wang, W.; Liang, Y.; Cai, Y.; Hooi, B. Mixup for node and graph classification. In Proceedings of the Web Conference 2021, Ljubljana Slovenia, 19–23 April 2021; pp. 3663–3674.
28. Liu, S.; Ying, R.; Dong, H.; Li, L.; Xu, T.; Rong, Y.; Zhao, P.; Huang, J.; Wu, D. Local augmentation for graph neural networks. In *Proceedings of the International Conference on Machine Learning*; PMLR: New York, NY, USA, 2022; pp. 14054–14072.
29. Xia, Y.; Yu, R.; Liang, Y.; Bresson, X.; Wang, X.; Zimmermann, R. Through the Dual-Prism: A Spectral Perspective on Graph Data Augmentation for Graph Classifications. In Proceedings of the AAAI Conference on Artificial Intelligence, Philadelphia, PA, USA, 25 February–4 March 2025; Volume 39, pp. 21635–21643.
30. Lai, Y.; Lin, X.; Yang, R.; Wang, H. Efficient topology-aware data augmentation for high-degree graph neural networks. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Barcelona, Spain, 25–29 August 2024; pp. 1463–1473.
31. Zhang, Z.; Li, L.; Wan, S.; Wang, W.; Wang, Z.; Lu, Z.; Hao, D.; Li, W. DropEdge not foolproof: Effective augmentation method for signed graph neural networks. *Adv. Neural Inf. Process. Syst.* **2024**, *37*, 117041–117069.
32. Li, Y.; Shi, S.; Guo, X.; Zhou, C.; Hu, Q. G-CutMix: A CutMix-based graph data augmentation method for bot detection in social networks. *PLoS ONE* **2025**, *20*, e0331978. [[CrossRef](#)] [[PubMed](#)]
33. Chen, J.; Zhu, G.; Yuan, C.; Huang, Y. Boosting graph anomaly detection with adaptive message passing. In Proceedings of the The Twelfth International Conference on Learning Representations, Vienna, Austria, 7–11 May 2024.
34. Pang, G.; Shen, C.; Van Den Hengel, A. Deep anomaly detection with deviation networks. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, AK, USA, 4–8 August 2019; pp. 353–362.
35. Sen, P.; Namata, G.; Bilgic, M.; Getoor, L.; Galligher, B.; Eliassi-Rad, T. Collective classification in network data. *AI Mag.* **2008**, *29*, 93. [[CrossRef](#)]
36. Tang, L.; Liu, H. Relational learning via latent social dimensions. In Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, 28 June–1 July 2009; pp. 817–826.
37. Qiao, H.; Pang, G. Truncated affinity maximization: One-class homophily modeling for graph anomaly detection. *Adv. Neural Inf. Process. Syst.* **2023**, *36*, 49490–49512.
38. Kumar, S.; Zhang, X.; Leskovec, J. Predicting dynamic embedding trajectory in temporal interaction networks. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, AK, USA, 4–8 August 2019; pp. 1269–1278.
39. Liu, Y.; Li, S.; Zheng, Y.; Chen, Q.; Zhang, C.; Pan, S. Arc: A generalist graph anomaly detector with in-context learning. *Adv. Neural Inf. Process. Syst.* **2024**, *37*, 50772–50804.
40. Xu, X.; Yuruk, N.; Feng, Z.; Schweiger, T.A. Scan: A structural clustering algorithm for networks. In Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Jose, CA, USA, 12–15 August 2007; pp. 824–833.
41. Li, J.; Dani, H.; Hu, X.; Liu, H. Radar: Residual analysis for anomaly detection in attributed networks. In Proceedings of the IJCAI, Melbourne, Australia, 19–25 August 2017; Volume 17, pp. 2152–2158.
42. Kipf, T.N.; Welling, M. Variational graph auto-encoders. *arXiv* **2016**, arXiv:1611.07308. [[CrossRef](#)]
43. Bandyopadhyay, S.; N, L.; Vivek, S.V.; Murty, M.N. Outlier resistant unsupervised deep architectures for attributed network embedding. In Proceedings of the 13th International Conference on Web Search and Data Mining, Houston, TX, USA, 3–7 February 2020; pp. 25–33.
44. Wang, X.; Jin, B.; Du, Y.; Cui, P.; Tan, Y.; Yang, Y. One-class graph neural networks for anomaly detection in attributed networks. *Neural Comput. Appl.* **2021**, *33*, 12073–12085. [[CrossRef](#)]

45. Roy, A.; Shu, J.; Li, J.; Yang, C.; Elshocht, O.; Smeets, J.; Li, P. Gad-nr: Graph anomaly detection via neighborhood reconstruction. In Proceedings of the 17th ACM International Conference on Web Search and Data Mining, Merida, Mexico, 4–8 March 2024; pp. 576–585.
46. Liu, K.; Dou, Y.; Zhao, Y.; Ding, X.; Hu, X.; Zhang, R.; Ding, K.; Chen, C.; Peng, H.; Shu, K.; et al. Bond: Benchmarking unsupervised outlier node detection on static attributed graphs. *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 27021–27035.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.