



Balanced Anomaly-guided Ego-graph Diffusion Model for Inductive Graph Anomaly Detection

Chunyu Wei*
Renmin University of China
Beijing, China
weichunyu@ruc.edu.cn

Siyuan He*
Renmin University of China
Beijing, China
river0307@ruc.edu.cn

Yu Wang
Independent Researcher
Beijing, China
feather1014@gmail.com

Yueguo Chen†
Renmin University of China
Beijing, China
chenyueguo@ruc.edu.cn

Yunhai Wang
Renmin University of China
Beijing, China
cloudseawang@gmail.com

Bing Bai
Microsoft
Beijing, China
baibing12321@163.com

Yidong Zhang
Renmin University of China
Beijing, China
ydzhang@ruc.edu.cn

Yong Xie
Nanjing University of Posts and
Telecommunications
Nanjing, China
yongxie@njupt.edu.cn

Shunming Zhang
Renmin University of China
Beijing, China
szhang@ruc.edu.cn

Fei Wang
Cornell University
Storrs, United States
few2001@med.cornell.edu

Abstract

Graph anomaly detection (GAD) is crucial in applications like fraud detection and cybersecurity. Despite recent advancements using graph neural networks (GNNs), two major challenges persist. At the model level, most methods adopt a transductive learning paradigm, which assumes static graph structures, making them unsuitable for dynamic, evolving networks. At the data level, the extreme class imbalance, where anomalous nodes are rare, leads to biased models that fail to generalize to unseen anomalies. These challenges are interdependent: static transductive frameworks limit effective data augmentation, while imbalance exacerbates model distortion in inductive learning settings. To address these challenges, we propose a novel data-centric framework that integrates dynamic graph modeling with balanced anomaly synthesis. Our framework features: (1) a discrete ego-graph diffusion model, which captures the local topology of anomalies to generate ego-graphs aligned with anomalous structural distribution, and (2) a curriculum anomaly augmentation mechanism, which dynamically adjusts synthetic data generation during training, focusing on underrepresented anomaly patterns to improve detection and generalization. Experiments on five datasets demonstrate that the effectiveness of our framework.

*Both authors contributed equally to this research.

†Corresponding author. He works at BRAIN of RUC.



This work is licensed under a Creative Commons Attribution 4.0 International License. *KDD '26, Jeju Island, Republic of Korea*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2258-5/2026/08
<https://doi.org/10.1145/3770854.3780240>

CCS Concepts

• Information systems → Data mining; • Computing methodologies → Anomaly detection; *Neural networks*.

Keywords

Graph Neural Networks, Anomaly Detection, Diffusion Model

ACM Reference Format:

Chunyu Wei, Siyuan He, Yu Wang, Yueguo Chen, Yunhai Wang, Bing Bai, Yidong Zhang, Yong Xie, Shunming Zhang, and Fei Wang. 2026. Balanced Anomaly-guided Ego-graph Diffusion Model for Inductive Graph Anomaly Detection. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1 (KDD '26)*, August 09–13, 2026, Jeju Island, Republic of Korea. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3770854.3780240>

Resource Availability:

The source code of this paper has been made publicly available at <https://github.com/OaxKnud/BAED>.

1 Introduction

Graph anomaly detection (GAD) is essential in various applications, as networks are often modeled as graphs with nodes representing entities and edges representing relationships [53, 69]. It has gained significant attention in fields like fraud detection [4, 34, 49], social network analysis [54, 66, 70], and cybersecurity [47]. Unlike traditional anomaly detection tasks, where anomalies are typically identified by observing the features of individual nodes, GAD requires consideration not only of node attributes but also the graph structure [57, 59]. This is because the relationships between nodes (i.e., the topology) often hold critical information that can reveal

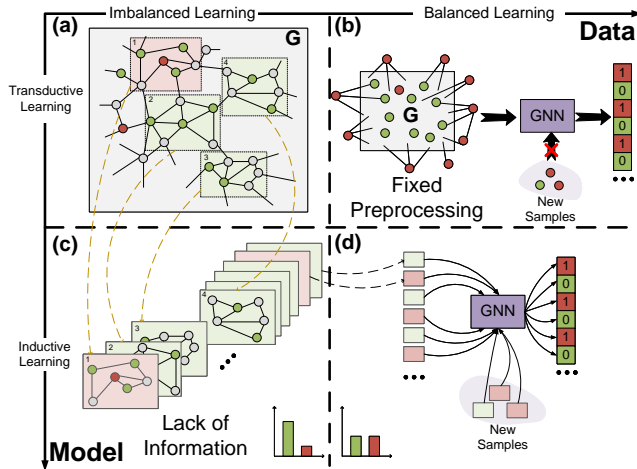


Figure 1: (a) Vanilla GAD paradigm: transductive learning is performed on the whole graph with imbalanced labels. (b) Graph-level data augmentation: fixed preprocessing that require retraining the model whenever new samples are introduced, limiting adaptability during the training process. (c) Inductive GAD on ego-graphs: inductive learning on ego-graphs inherits the label imbalance issue from the original graph. (d) Dynamic and balanced ego-graph augmentation: our proposed inductive framework dynamically adjusts the type and ratio of generated samples during training to address imbalance and enhance model adaptability in real-time.

anomalous behaviors. For example, a node may disguise itself by modifying its features, but the interaction patterns represented by its topology can still highlight its anomalies.

Various Graph Neural Network (GNN)-based methods have been proposed to address graph anomaly detection by learning node representations that capture both feature and topological information. However, these methods face two main limitations: one at the **model level** and another at the **data level**, as illustrated in Figure 1(a).

At the **model level**, most approaches rely on transductive learning, where models are trained on a fixed set of labeled graph data, and the learned representations are subsequently used to detect anomalies. However, this approach assumes static graph structures, which is often unrealistic as real-world networks are dynamic [58]. When the graph structure changes, the model must be retrained from scratch, a process that is both computationally expensive and impractical for large-scale or real-time applications [19, 20, 55]. This limitation highlights a critical mismatch between the static assumptions of transductive learning and the inherently dynamic nature of real-world networks [52, 56].

At the **data level**, class imbalance is a significant challenge, as anomalous nodes are rare and their distribution is skewed. For instance, in the DGraph [18] dataset, only 1.3% of nodes are anomalous. This imbalance leads to biased models that fail to capture the distinguishing characteristics of anomalous nodes, ultimately resulting in poor generalization during downstream prediction.

These two challenges, transductive learning and imbalanced learning, are intrinsically linked. On one hand, the shift from transductive to inductive learning exacerbates the imbalance issue as illustrated in Figure 1(c). This is because inductive learning requires generalization to unseen graph data, and the scarcity of anomalous nodes greatly hampers this ability. On the other hand, existing methods to address the imbalance through data synthesis often rely on pre-processing techniques, which do not dynamically adapt during model training, as illustrated in Figure 1(b). In transductive settings, the model’s inability to adapt to graph structure changes exacerbates the complexity of the problem.

To simultaneously tackle these two challenges, we propose a novel data-centric framework, **Balanced Anomaly-guided Ego-graph Diffusion (BAED)** model, designed to address both the dynamic nature of graphs and the imbalance issue in GAD tasks, as illustrated in Figure 1(d). Specifically, we introduce an anomalous ego-graph synthesizer aimed at enhancing inductive graph anomaly detection by generating realistic, anomaly-specific ego-graphs that can be used to augment the training process. Our framework operates on the principle that effective anomaly detection requires not only a robust model but also the generation of informative samples that guide the model towards better generalization.

To achieve this, we focus on two key goals:

Better Fitting: We propose a discrete ego-graph diffusion model that generates ego-graphs more aligned with the distribution of anomalous nodes. Unlike traditional diffusion models that map the entire graph to a latent space and overlook critical local neighborhood information, our model operates directly on the graph’s topology, preserving the local structure of anomalous nodes.

Better Adaptation: We introduce a curriculum anomaly augmentation mechanism that dynamically adjusts during training. Unlike traditional static pre-processing methods, our approach adapts to the training process itself. At each iteration, based on the model’s previous loss, the synthesis process adjusts the distribution of generated anomaly samples to focus on underrepresented types of anomalies. This enhances the model’s ability to detect specific anomalies, improving both accuracy and generalization.

Our main contributions can be summarized as follows:

- We propose a data-centric framework that simultaneously addresses the challenges of transductive learning and imbalanced learning in graph anomaly detection.
- We introduce a novel discrete ego-graph diffusion model that effectively generates realistic ego-graphs for curriculum anomaly augmentation, which enabling dynamic adaptation of the synthetic data generation during training.
- Through experiments on benchmarks for graph anomaly detection, BAED achieves state-of-the-art results.

2 Related Works

2.1 Graph Anomaly Detection

With the advent of graph neural networks (GNNs) [40, 48, 60, 61, 64], significant progress has been made in research on anomaly detection using node features and graph topology [6, 29, 38]. Traditional model-centric approaches predominantly rely on discriminative techniques, where representations are learned over the graph, followed by the construction of decision boundaries based on node

embeddings to classify anomalies [57]. These methods often make use of partial label information and are typically framed in a semi-supervised or supervised learning setting. Notable examples include FdGars [50], which employs multi-layer GNNs to classify users based on their content and behavior, and CARE-GNN [7], which adjusts neighbor aggregation thresholds via reinforcement learning to address inconsistencies. Similarly, approaches such as FRAUDRE [67] and PC-GNN [28] mitigate class imbalance by adopting specialized loss functions or label-balanced sampling strategies.

Spectral-based methods, such as BWGNN [43] and AMNet [2], explore the graph’s frequency spectrum to capture anomalies, using techniques like Beta-graph wavelets and Bernstein polynomials to separate normal and abnormal frequency bands. In parallel, GHRN [10] eliminates harmful heterogeneous edges to enhance model robustness, and SEC-GFD [63] tackles heterophily and label imbalance through spectral filtering.

Unsupervised learning techniques, essential due to the scarcity of labeled data, have gained prominence in graph anomaly detection (GAD). These include traditional outlier detection methods like LOF [1] and Isolation Forest (IF) [26], as well as autoencoder-based models like MAPLE [41] and variational graph autoencoders (VGAEs) [22, 68], which improve anomaly detection by learning robust node representations. Structural similarity-based methods like SCAN [65] and Radar [25] incorporate both attribute and structural information, while BOND [27] benchmarks various unsupervised anomaly detection methods.

To address the challenges of data scarcity and extreme imbalance in anomaly detection scenarios, various methods have adopted data augmentation strategies. For instance, some approaches enhance node features, generate contrastive views, or modify egonet structures [37, 71], often tailored for specific algorithms [71, 72]. Recently, novel data-centric augmentation methods have emerged, directly tackling the issue of label scarcity by generating additional training data [31]. Our work falls into this category. However, while these methods generate additional node features, the graph’s fixed topological structure, which is constrained by transductive approaches, remains unchanged throughout the augmentation process. Our method is designed to address this issue by considering both dynamic graph structures and newly added nodes, ensuring better adaptability and generalization to GAD tasks.

2.2 Inductive Learning

Inductive learning on graphs generalizes to unseen nodes or new graphs, essential for dynamic graph data.

Network Embedding-Based Methods. Embedding methods extend static techniques like Skip-gram to dynamic settings, capturing temporal dynamics (e.g., Dynamic Skip-gram [8], temporal embeddings [35, 73]). Autoencoder-based methods like DynGEM [12] and dyngraph2vec [11] excel in link prediction and anomaly detection but struggle to generalize to new graphs and lack scalability in highly dynamic environments.

Graph Neural Network-Based Methods. GNNs support inductive learning by utilizing node features and structural information. Methods like GraphSAGE [14] introduce sampling-aggregation for new nodes and graphs, while TGAT [62] captures dynamic patterns with temporal encoding. Recent pre-training strategies (e.g.,

L2P-GNN [30], graph-level pre-training [17]) improve inductive learning by learning transferable representations, though they still face challenges in aligning pre-training with fine-tuning objectives.

3 Preliminaries

Attributed Graph. An *attributed graph* is defined as a structure consisting of nodes, their associated features, and the relations (edges) between them. Formally, we represent an attributed graph as $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathbf{X}\}$, where:

- $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$ denotes the set of N nodes.
- $\mathcal{E} = \{e_{ij} \mid v_i \text{ and } v_j \text{ are connected, } v_i, v_j \in \mathcal{V}\}$ represents the set of edges between nodes.
- $\mathbf{X} \in \mathbb{R}^{N \times d}$ represents the feature matrix of the graph, where each row $\mathbf{x}_i \in \mathbb{R}^d$ corresponds to the d -dimensional feature vector of node v_i .

To facilitate matrix-based computations, we define the graph’s *adjacency matrix* $\mathbf{A} \in \mathbb{N}^{N \times N}$, where each entry A_{ij} is binary:

$$A_{ij} = \begin{cases} 1, & \text{if there exists an edge } e_{ij} \in \mathcal{E}, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Ego-graph. An *Ego-graph* is a specialized subgraph structure centered around a specific node v_i , referred to as the *ego*, and its neighborhood $\mathcal{N}^K(i)$ within a predefined number of hops K . Formally, the node set of the ego-graph can be expressed as: $\mathcal{V}_K^i = \{v_i\} \cup \mathcal{N}^K(i)$, where $\mathcal{N}^K(i)$ denotes the set of nodes reachable from v_i within K hops. The ego-graph \mathcal{G}_K^i is then defined as:

$$\mathcal{G}_K^i = \{\mathcal{V}_K^i, \{e_{jk} \mid e_{jk} \in \mathcal{E}, v_j, v_k \in \mathcal{V}_K^i\}, [x_k, \dots]\} \text{ where } v_k \in \mathcal{V}_K^i,$$

where $[x_k, \dots]$ represents the set of feature vectors corresponding to each node $v_k \in \mathcal{V}_K^i$.

Anomaly Detection on Graph. Given an *attributed graph* \mathcal{G} , the node set \mathcal{V} can be partitioned into two disjoint subsets, namely the anomalous node set \mathcal{V}_a and the normal node set \mathcal{V}_n , such that: $\mathcal{V}_a \cap \mathcal{V}_n = \emptyset$ and $\mathcal{V}_a \cup \mathcal{V}_n = \mathcal{V}$.

However, only a subset of nodes in the graph are labeled. Let $\mathcal{V}_{train} \subset \mathcal{V}$ denote the set of labeled nodes, and their corresponding labels are given as: $Y_{train} = \{y_i \mid y_i \in \{0, 1\}, v_i \in \mathcal{V}_{train}\}$, where $y_i = 1$ indicates that the node v_i is anomalous ($v_i \in \mathcal{V}_a$), and $y_i = 0$ indicates that the node is normal ($v_i \in \mathcal{V}_n$).

The objective of *graph anomaly detection* (GAD) is to infer the anomaly labels Y_{test} of the unlabeled nodes $\mathcal{V}_{test} = \mathcal{V} \setminus \mathcal{V}_{train}$ by estimating the probability:

$$p(v_i \in \mathcal{V}_a \mid \mathcal{G}, v_i \in \mathcal{V}_{test}), \quad (2)$$

which represents the likelihood that node v_i is anomalous given the graph \mathcal{G} and the partial node labels in \mathcal{V}_{train} .

Inductive Anomaly Detection on Graph. Unlike traditional transductive approaches, which utilize the entire graph structure and node features, *inductive anomaly detection* focuses on local information within the node’s ego-graph to compute the anomaly probability. The goal of *inductive graph anomaly detection* is to estimate the probability:

$$p(v_i \in \mathcal{V}_a \mid \mathcal{G}_K^i, v_i \in \mathcal{V}_{test}), \quad (3)$$

which represents the likelihood that the node v_i is anomalous, conditioned only on its local ego-graph \mathcal{G}_K^i .

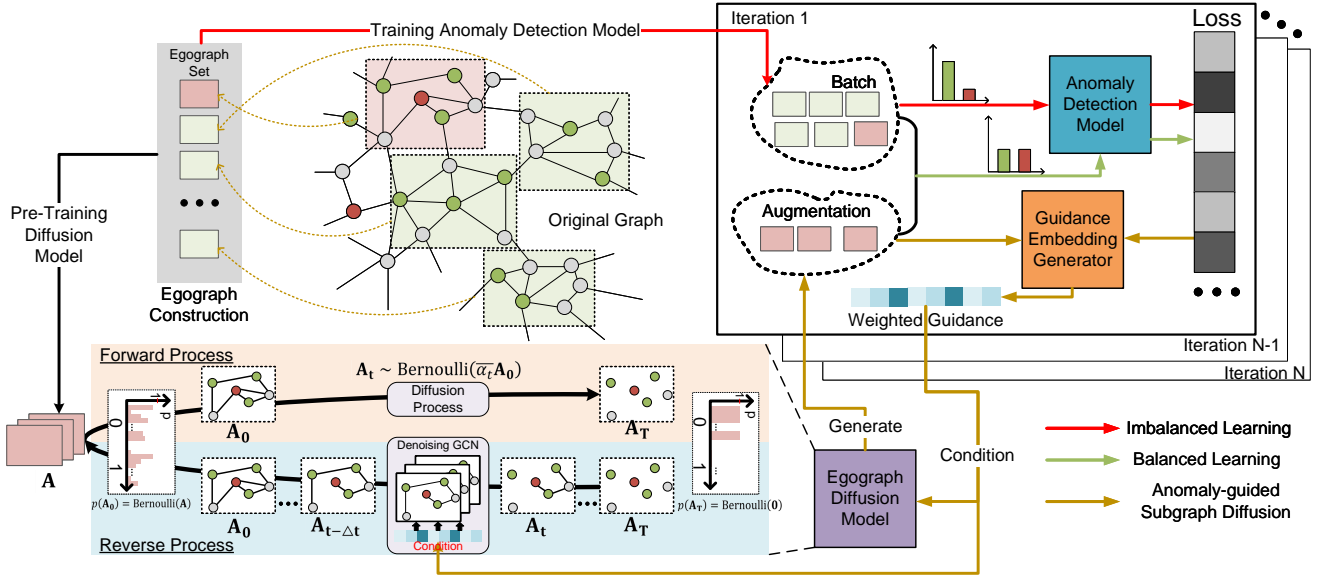


Figure 2: The BAED framework. Left: Pre-training the discrete ego-graph diffusion model with forward noise addition and reverse denoising processes. Right: Training iterations where the anomaly detection model processes both original imbalanced batch and augmented samples. The Guidance Embedding Generator (GIN) encodes anomalous ego-graphs into guidance embeddings, which are dynamically weighted based on previous losses to focus on underrepresented anomaly types.

4 Methodology

We propose the Balanced Anomaly-guided Ego-graph Diffusion (BAED) framework, as shown in Figure 2. The core idea is to use a discrete **Ego-graph Diffusion Model** to generate anomalous samples during the inductive GAD training, ensuring a balanced ratio of anomalous to normal samples. To improve adaptability, we introduce an **Anomaly-Guided Graph Generation** mechanism, controlled by a guidance embedding to align generated samples with specific anomaly patterns. Additionally, our **Curriculum Anomaly Augmentation** dynamically adjusts the anomaly patterns and proportions based on prior training results. Finally, we provide theoretical analyses demonstrating BAED’s effectiveness.

4.1 Inductive anomaly detection model

In Equation 3, the goal is to estimate the anomaly probability $p(v_i \in \mathcal{V}_a \mid \mathcal{G}_K^i, v_i \in \mathcal{V}_{test})$. Our approach offers plug-and-play flexibility, allowing the adaptation of any graph operator-based transductive GNN. For training, we use the GAE (Graph Autoencoder) framework to train our selected GNN on the entire graph [9, 61]. During inference, the model is applied only to the local ego-graph for node v_i : $p(v_i \in \mathcal{V}_a \mid \mathcal{G}_K^i) = GNN(\mathcal{G}_K^i)$.

In the l -th layer of the GNN, node representations are updated based on node features and graph topology:

$$\mathbf{E}^{(l)} = GCN(\mathbf{E}^{(l-1)}, \mathcal{G}_K^i), \quad (4)$$

where $\mathbf{E}^{(l)}$ is the node embedding at layer l , and $GCN(\cdot)$ is the graph convolution operation. Initially, $\mathbf{E}^{(0)} = [x_k, \dots]$, where $v_k \in \mathcal{V}_K^i$. After L layers of convolution, the final node representations for the ego-graph \mathcal{G}_K^i are aggregated:

$$\mathbf{H}_K^i = AGG(\mathbf{E}^{(1)}, \mathbf{E}^{(2)}, \dots, \mathbf{E}^{(L)}), \quad (5)$$

where $AGG(\cdot)$ can be implemented using functions like mean or concatenation.

To improve anomaly detection, we propose a deviation-based representation for the final input:

$$h_{\mathcal{G}_K^i} = \mathbf{H}_K^i[v_i] - \text{mean_pooling}(\mathbf{H}_K^i), \quad (6)$$

where $\mathbf{H}_K^i[v_i]$ is the node embedding of v_i , and $\text{mean_pooling}(\cdot)$ is the mean of all node embeddings. This deviation highlights differences between the ego node v_i and its neighbors, reflecting its anomaly degree.

PROPOSITION 4.1. *The magnitude of the deviation $\|h_{\mathcal{G}_K^i}\|_2$ increases as the central node v_i exhibits abnormal behavior compared to its local neighborhood. Formally, given the ego-graph \mathcal{G}_K^i , the deviation satisfies:*

$$\|h_{\mathcal{G}_K^i}\|_2 \propto \delta(v_i, N^K(i)), \quad (7)$$

where $\delta(v_i, N^K(i))$ denotes the degree of deviation of v_i from the distribution of its neighbors $N^K(i)$.

Thus, the deviation operation provides a monotonically increasing measure of node v_i ’s anomaly degree within its local ego-graph. We then use a Multi-Layer Perceptron (MLP)-based discriminator to map the representation $h_{\mathcal{G}_K^i}$ to a probability in the range $[0, 1]$, representing the likelihood that node v_i is anomalous:

$$p(v_i \in \mathcal{V}_a \mid \mathcal{G}_K^i) = MLP(h_{\mathcal{G}_K^i}). \quad (8)$$

4.2 Ego-graph Diffusion Model

We focus on generating ego-graphs by manipulating the adjacency matrix through a discrete diffusion process. The process begins with the initial adjacency matrix A^0 , and noise is incrementally

added at each step, represented as $q(\mathbf{A}^t | \mathbf{A}^{t-1})$, forming a trajectory $(\mathbf{A}^0, \mathbf{A}^1, \dots, \mathbf{A}^T)$. The diffusion process models the joint distribution $q(\mathbf{A}^1, \dots, \mathbf{A}^T | \mathbf{A}^0) = \prod_{t=1}^T q(\mathbf{A}^t | \mathbf{A}^{t-1})$. To model the reverse process, we introduce a denoising process $p_\theta(\mathbf{A}^{t-1} | \mathbf{A}^t)$, parameterized by θ , which is modeled as:

$$p_\theta(\mathbf{A}^{0:t}) = p(\mathbf{A}^T) \prod_{t=1}^T p_\theta(\mathbf{A}^{t-1} | \mathbf{A}^t), \quad (9)$$

where $p(\mathbf{A}^T)$ is the distribution the adjacency matrix converges to after the forward diffusion. An efficient diffusion model must satisfy three key properties:

4.2.1 Closed-form Conditional Distribution. The conditional distribution $q(\mathbf{A}^t | \mathbf{A}^0)$ should have a closed-form solution to avoid recursive noise application during training. Following [3, 46], we use $\mathcal{B}(x; \mu)$ to denote a Bernoulli distribution with parameter μ . Each entry of the adjacency matrix, representing an edge, is treated as an independent random variable, with $\mathcal{B}(\mathbf{A}; \mu)$ representing the entire adjacency matrix.

We define the forward process as resampling with probability β_t , following a Bernoulli distribution with parameter p . Thus, $q(\mathbf{A}^t | \mathbf{A}^{t-1})$ is an independent Bernoulli distribution for each edge:

$$\begin{aligned} q(\mathbf{A}^t | \mathbf{A}^{t-1}) &= \prod_{i,j:i < j} \mathcal{B}\left(\mathbf{A}_{i,j}^t; (1 - \beta_t) \mathbf{A}_{i,j}^{t-1} + \beta_t p\right) \\ &:= \mathcal{B}\left(\mathbf{A}^t; (1 - \beta_t) \mathbf{A}^{t-1} + \beta_t p\right). \end{aligned} \quad (10)$$

This leads to:

$$q(\mathbf{A}^t | \mathbf{A}^0) = \mathcal{B}\left(\mathbf{A}^t; \bar{\alpha}_t \mathbf{A}^0 + (1 - \bar{\alpha}_t) p\right), \quad (11)$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{\tau=1}^t \alpha_\tau$.

4.2.2 Limiting Distribution. The second key property is that the limiting distribution $q_\infty = \lim_{T \rightarrow \infty} q(\mathbf{A}^T | \mathbf{A}^0)$ should be independent of \mathbf{A}^0 , enabling its use as a prior. To achieve this, we use noise scheduling [16] to adjust β_1, \dots, β_t such that $\bar{\alpha}_t$ approaches 0 as t increases. This ensures $\lim_{T \rightarrow \infty} q(\mathbf{A}^T | \mathbf{A}^0) = \mathcal{B}(\mathbf{A}^T; p) = q(\mathbf{A}^T)$, where $q(\mathbf{A}^T)$ is independent of \mathbf{A}^0 . For simplicity, we set $p = 0$.

4.2.3 Closed-form Posterior Distribution. The third key property is that the posterior distribution $q(\mathbf{A}^{t-1} | \mathbf{A}^t, \mathbf{A}^0)$ should have a closed-form solution to serve as the target for the neural network in the reverse denoising process. The posterior of the forward transition conditioned on \mathbf{A}^0 can be computed as:

$$q(\mathbf{A}^{t-1} | \mathbf{A}^t, \mathbf{A}^0) = \frac{q(\mathbf{A}^t | \mathbf{A}^{t-1}) q(\mathbf{A}^{t-1} | \mathbf{A}^0)}{q(\mathbf{A}^t | \mathbf{A}^0)}. \quad (12)$$

Since the posterior are trackable, it can be used to train the denoising model $p_\theta(\mathbf{A}^{t-1} | \mathbf{A}^t)$ by maximizing the evidence lower bound (ELBO) on the log-likelihood $\log p_\theta(\mathbf{A}^0)$.

4.3 Anomaly-guided graph generation

Directly applying the denoising reverse process from a random noise distribution $\mathcal{B}(\mathbf{A}^T; p)$ generates data that aligns with the original anomalous distribution. However, real-world anomalies often follow a multimodal distribution due to diverse anomaly patterns [45]. When aiming to generate samples around a specific mode (e.g., a particular anomaly pattern) to enhance model training, a random reverse process is insufficient. To address this, we propose

an **Anomaly-Guided Graph Generation** framework to generate augmented samples for the desired anomaly pattern.

For a given anomalous sample represented by an ego-graph \mathcal{G}_K^i , we hypothesize that a suitable graph embedding can capture its anomalous characteristics.

PROPOSITION 4.2. Distributional Separability of Local Ego-graph Embedding Space. For node embeddings \mathbf{z}_v in the embedding space, if the Ego-graph distributions $p(S_v)$ of different anomaly pattern satisfy separability conditions, then the embedding distribution $p(\mathbf{z}_v)$ also satisfies:

$$D_{KL}(p(\mathbf{z} | C_i) \| p(\mathbf{z} | C_j)) > \delta, \quad \forall i \neq j, \quad (13)$$

where $\delta > 0$ is the threshold for distributional separability.

Graph Isomorphism theory [64] demonstrates that GNNs can approximate the ability of graph isomorphism testing. They aggregate information from neighboring nodes to capture the structural properties of graphs. This implies that embedding a node's local neighborhood can effectively capture subtle structural differences. Consequently, we employ the Graph Isomorphism Network (GIN) as an encoder to model the anomalous sample into a guidance embedding. Formally, we define: $\mathbf{h}^g = \text{GIN}(\mathcal{G}_K^i)$, where $\mathbf{h}^g \in \mathbb{R}^d$.

Conditional Reverse Process. Given the noise addition process defined in Equation 10, we define the discrete conditional denoising reverse process as:

$$p_\theta(\mathbf{A}^{t-1} | \mathbf{A}^t, \mathbf{h}^g) = \prod_{i,j:i < j} p_\theta(\mathbf{A}_{i,j}^{t-1} | \mathbf{A}_{i,j}^t, \mathbf{h}^g). \quad (14)$$

Training the Denoising Network. Following [16, 42], we optimize the denoising model $p_\theta(\mathbf{A}^{t-1} | \mathbf{A}^t, \mathbf{h}^g)$ by maximizing the evidence lower bound (ELBO) of the log-likelihood $\log p_\theta(\mathbf{A}^0 | \mathbf{h}^g)$. The ELBO is formulated as:

$$\log p_\theta(\mathbf{A}^0 | \mathbf{h}^g) \geq \mathbb{E}_{q(\mathbf{A}^{1:T} | \mathbf{A}^0, \mathbf{h}^g)} \left[\log \frac{p_\theta(\mathbf{A}^{0:T} | \mathbf{h}^g)}{q(\mathbf{A}^{1:T} | \mathbf{A}^0, \mathbf{h}^g)} \right]. \quad (15)$$

This leads to the following training loss:

$$\begin{aligned} \mathcal{L}(\mathbf{A}^0; \theta) &= \mathbb{E}_q \left[-\log p_\theta(\mathbf{A}^0 | \mathbf{A}^1, \mathbf{h}^g) \right] + \mathbb{D}_{KL} \left(q(\mathbf{A}^T | \mathbf{A}^0, \mathbf{h}^g) \| p(\mathbf{A}^T | \mathbf{h}^g) \right) \\ &+ \sum_{t=2}^T \mathbb{E}_q \left[\mathbb{D}_{KL} \left(q(\mathbf{A}^{t-1} | \mathbf{A}^t, \mathbf{A}^0, \mathbf{h}^g) \| p_\theta(\mathbf{A}^{t-1} | \mathbf{A}^t, \mathbf{h}^g) \right) \right]. \end{aligned} \quad (16)$$

Here, the first term is the reconstruction loss, denoted as \mathcal{L}_{recon} ; the second term is the prior matching loss, \mathcal{L}_{prior} ; and the third term is the consistency loss, \mathcal{L}_{con} .

Anomaly-Guided Inference. After training the denoising network $p_\theta(\mathbf{A}^{t-1} | \mathbf{A}^t, \mathbf{h}^g)$, the inference process guided by a specific guidance embedding \mathbf{h}^g is outlined in Algorithm 1.

This anomaly-guided framework enables the generation of augmented samples that focus on specific anomalous patterns, improving the robustness and diversity of anomaly detection models.

4.4 Curriculum Anomaly Augmentation

In Section 4.3, we introduced how the Ego-graph Diffusion Model generates Ego-graphs representing specific anomaly patterns guided by a guidance embedding. However, during training, the learning

Algorithm 1 Inference for Anomaly Detection

-
- 1: Initialize $\mathbf{A}_T \sim \mathcal{B}(\mathbf{A}^T; p)$, the Anomaly-Guidance Embedding \mathbf{h}^g , and diffusion steps T .
 - 2: **for** $t = T$ to 1 **do**
 - 3: Draw $\mathbf{A}_{t-1} \sim p_\theta(\mathbf{A}^{t-1} | \mathbf{A}^t, \mathbf{h}^g)$.
 - 4: **end for**
 - 5: **return** \mathbf{A}_0 .
-

ability for different anomaly patterns varies. Thus, we propose **Curriculum Anomaly Augmentation** strategy, which dynamically adjusts the generation of guidance embeddings to focus on specific anomaly patterns. This enables the anomaly-guided graph generation to adaptively generate augmented samples that match the model's curriculum learning needs at each training iteration.

Based on the principle above, we use the loss values to compute importance weights for the guidance embeddings. As discussed in Section 4.3, we use the GIN encoder to compute the guidance embedding for each anomalous ego-graph \mathcal{G}_K^i during training iteration t . The importance weight for each embedding is computed using a heuristic sigmoid function:

$$w_i = \frac{1}{1 + e^{-\alpha(l_i - \beta)}} \cdot h(t), \quad (17)$$

where α and β are hyperparameters that control the sensitivity and shift of the weight in response to the loss values l_i (the anomaly detection loss of the i -th sample in the batch). The term $h(t)$ is a temporal factor, designed as a monotonically increasing function of the training iteration: $h(t) = \frac{t}{T}$, where T is the total number of training iterations. This design ensures that the generation model focuses less on hard-to-learn samples in the early training stages, gradually shifting its attention to these underfit samples in later stages.

Using the computed weights, we aggregate the guidance embeddings for all samples in the batch. The final guidance embedding $\tilde{\mathbf{h}}^g$ for the anomaly-guided generation is computed as:

$$\tilde{\mathbf{h}}^g = \sum_{i=1}^N \frac{w_i}{\sum_{j=1}^N w_j} \cdot \text{GIN}(\mathcal{G}_K^i), \quad (18)$$

where w_i is the normalized weight for the i -th sample, N is the number of anomaly in the batch, and $\text{GIN}(\mathcal{G}_K^i)$ represents the guidance embedding computed by the GIN encoder for the corresponding ego-graph \mathcal{G}_K^i .

This weighting and aggregation process enables the dynamic generation of augmented samples that are tailored to the model's learning progress at each training stage. By adaptively focusing on underfit anomaly patterns in later training stages, this strategy integrates a curriculum learning paradigm into the anomaly-guided graph generation framework. Furthermore, it ensures stable training by gradually increasing the difficulty of the generated samples as the model progresses through training.

4.5 Theoretical Analysis

We provide theoretical justification for how BAED addresses the critical challenge of class imbalance in graph anomaly detection. Our analysis demonstrates that balanced anomaly augmentation

through ego-graph diffusion enables optimal decision boundaries for anomaly detection.

Consider the binary classification where $X \in \mathcal{X}$ represents the input features and $Y \in \{0, 1\}$ denotes the anomaly label. Let $\Theta = \{\theta : \mathcal{X} \rightarrow [0, 1]\}$ be the set of all classifiers, where $\theta(x)$ represents the predicted probability of anomaly. Following the framework in [23], we assume the data is generated i.i.d. from distribution \mathbb{P} , with $\pi = \mathbb{P}(Y = 1)$ denoting the prior probability of anomalies.

For any classifier θ , we define the true positive rate as $\text{TP}(\theta) = \mathbb{P}(Y = 1, \theta = 1)$ and the classification rate as $\gamma(\theta) = \mathbb{P}(\theta = 1)$. The performance metric takes the general form:

$$\mathcal{L}(\theta) = \frac{c_0 + c_1 \cdot \text{TP}(\theta) + c_2 \cdot \gamma(\theta)}{d_0 + d_1 \cdot \text{TP}(\theta) + d_2 \cdot \gamma(\theta)}, \quad (19)$$

where c_i, d_i for $i \in \{0, 1, 2\}$ are constants determined by π and the specific performance metric (e.g., F1-score, precision-recall).

THEOREM 4.1 (OPTIMAL DECISION BOUNDARY UNDER IMBALANCE). *Under mild regularity conditions on \mathbb{P} , the optimal decision boundary δ^* for maximizing $\mathcal{L}(\theta)$ is given by:*

$$\delta^* = \frac{d_2 \mathcal{L}^* - c_2}{c_1 - d_1 \mathcal{L}^*}, \quad (20)$$

where $\mathcal{L}^* = \sup_{\theta \in \Theta} \mathcal{L}(\theta)$. The Bayes optimal classifier takes:

$$\theta^*(x) = \begin{cases} \mathbb{1}[\eta(x) > \delta^*] & \text{if } c_1 > d_1 \mathcal{L}^* \\ \mathbb{1}[\eta(x) < \delta^*] & \text{if } c_1 < d_1 \mathcal{L}^* \end{cases} \quad (21)$$

where $\eta(x) = \mathbb{P}(Y = 1 | X = x)$ is the posterior probability of anomaly.

This theorem from [23] above reveals a critical insight: when π is small (severe class imbalance), the optimal threshold δ^* deviates significantly from the conventional 0.5, leading to suboptimal performance if not properly addressed.

PROPOSITION 4.3 (BAED'S BALANCING EFFECT). *Let $\tilde{\mathbb{P}}$ denote the augmented distribution after applying BAED's ego-graph generation, with $\tilde{\pi} = \tilde{\mathbb{P}}(Y = 1)$ being the augmented anomaly ratio. BAED's curriculum augmentation strategy ensures:*

$$|\tilde{\pi} - 0.5| < |\pi - 0.5|, \quad (22)$$

leading to an optimal threshold $\tilde{\delta}^*$ closer to 0.5, thereby improving the stability and generalization of the anomaly detector.

PROOF. Let N_a and N_n denote the number of anomalous and normal samples in the original dataset, respectively, with $\pi = N_a / (N_a + N_n)$. At training iteration t , BAED generates M_t synthetic anomalous ego-graphs based on the importance weights w_i defined in Eq. (18). The effective anomaly ratio becomes:

$$\tilde{\pi}_t = \frac{N_a + M_t}{N_a + N_n + M_t}. \quad (23)$$

Since $w_i = \frac{1}{1 + e^{-\alpha(l_i - \beta)}} \cdot \frac{t}{T}$ increases with the loss l_i and training progress t/T , underrepresented anomaly types receive higher generation priority. The curriculum mechanism ensures M_t increases monotonically, driving $\tilde{\pi}_t$ towards 0.5. As $\pi < 0.5$ in imbalanced settings, we have $\tilde{\pi}_t > \pi$, thus $|\tilde{\pi}_t - 0.5| < |\pi - 0.5|$. \square

By dynamically generating anomalous ego-graphs that target underrepresented patterns, BAED effectively transforms the learning problem from one with extreme imbalance ($\pi \ll 0.5$) to a more balanced scenario ($\tilde{\pi} \approx 0.5$), thereby achieving superior detection performance as validated in our experiments.

Table 1: Overall performance comparison on five datasets within inductive setting, with the bold numbers representing the best results and the underlined numbers indicating the most competitive results (RQ1)

Method	Elliptic			Reddit			Photo			T-Finance			Dgraph			
	AUROC	AUPRC	F1	AUROC	AUPRC	F1	AUROC	AUPRC	F1	AUROC	AUPRC	F1	AUROC	AUPRC	F1	
GCN	0.517±0.004	0.065±0.003	0.491±0.009	0.508±0.007	0.040±0.005	0.492±0.005	0.499±0.001	0.087±0.004	0.512±0.017	0.357±0.011	0.156±0.016	0.488±0.007	0.529±0.014	0.013±0.005	0.497±0.009	
+ AEGIS	0.541±0.014	0.098±0.004	0.516±0.007	0.531±0.007	0.039±0.003	0.498±0.015	0.585±0.008	0.124±0.005	0.523±0.012	0.579±0.005	0.048±0.003	0.537±0.018	0.550±0.009	0.005±0.001	0.495±0.012	
+ GGAD	0.603±0.005	0.112±0.004	0.532±0.017	0.525±0.015	0.041±0.002	0.505±0.007	0.560±0.015	0.122±0.005	0.497±0.019	0.823±0.008	0.183±0.007	0.587±0.016	0.558±0.007	0.006±0.001	0.531±0.014	
+ CGenGA	0.595±0.011	0.075±0.004	0.523±0.006	0.681±0.016	0.047±0.003	0.532±0.017	0.503±0.014	0.094±0.004	0.496±0.005	0.847±0.019	0.776±0.019	0.618±0.019	0.549±0.015	0.013±0.002	0.513±0.015	
+ BAED	0.671±0.007	0.154±0.006	0.566±0.007	0.732±0.004	0.077±0.005	0.543±0.011	0.703±0.009	0.306±0.007	0.581±0.015	0.908±0.008	0.856±0.003	0.713±0.009	0.566±0.006	0.014±0.002	0.521±0.007	
BWGNN	0.637±0.006	0.064±0.003	0.505±0.007	0.520±0.017	0.036±0.007	0.517±0.007	0.613±0.004	0.165±0.008	0.558±0.019	0.208±0.008	0.133±0.006	0.301±0.015	0.580±0.006	0.015±0.002	0.519±0.009	
+ AEGIS	0.564±0.014	0.100±0.005	0.512±0.017	0.545±0.007	0.037±0.005	0.563±0.012	0.649±0.006	0.149±0.006	0.561±0.003	0.528±0.014	0.042±0.003	0.551±0.003	0.445±0.014	0.006±0.001	0.488±0.014	
+ GGAD	0.554±0.012	0.099±0.004	0.507±0.005	0.553±0.014	0.040±0.003	0.558±0.011	0.646±0.007	0.208±0.007	0.567±0.011	0.770±0.018	0.098±0.005	0.572±0.018	0.594±0.016	0.008±0.001	0.511±0.015	
+ CGenGA	0.585±0.019	0.073±0.004	0.502±0.002	0.570±0.014	0.045±0.001	0.587±0.024	0.751±0.008	0.497±0.013	0.613±0.004	0.627±0.016	0.212±0.008	0.581±0.019	0.635±0.017	0.012±0.002	0.601±0.018	
+ BAED	0.669±0.016	0.102±0.005	0.563±0.009	0.654±0.016	0.082±0.004	0.562±0.002	0.971±0.020	0.866±0.003	0.911±0.011	0.853±0.012	0.909±0.019	0.848±0.009	0.702±0.021	0.599±0.016	0.016±0.002	0.522±0.018
GraphSAGE	0.505±0.012	0.054±0.003	0.507±0.015	0.546±0.014	0.037±0.003	0.508±0.002	0.501±0.013	0.085±0.004	0.503±0.026	0.554±0.007	0.238±0.008	0.532±0.017	0.464±0.014	0.012±0.001	0.497±0.008	
+ AEGIS	0.576±0.024	0.103±0.005	0.521±0.017	0.553±0.014	0.048±0.003	0.525±0.017	0.628±0.016	0.128±0.006	0.562±0.030	0.533±0.003	0.037±0.003	0.521±0.011	0.540±0.012	0.005±0.001	0.509±0.017	
+ GGAD	0.579±0.015	0.105±0.005	0.534±0.011	0.533±0.013	0.043±0.003	0.509±0.013	0.669±0.012	0.175±0.007	0.572±0.011	0.689±0.017	0.074±0.004	0.562±0.005	0.537±0.015	0.006±0.001	0.502±0.016	
+ CGenGA	0.531±0.033	0.061±0.003	0.529±0.015	0.624±0.015	0.053±0.004	0.522±0.019	0.544±0.014	0.088±0.004	0.513±0.011	0.596±0.015	0.322±0.009	0.563±0.009	0.529±0.015	0.009±0.001	0.498±0.013	
+ BAED	0.593±0.015	0.134±0.006	0.583±0.013	0.667±0.016	0.082±0.005	0.562±0.002	0.971±0.020	0.866±0.003	0.763±0.015	0.750±0.018	0.425±0.012	0.612±0.017	0.668±0.019	0.020±0.002	0.522±0.019	
BernNet	0.495±0.012	0.053±0.007	0.402±0.014	0.469±0.013	0.068±0.004	0.482±0.018	0.587±0.015	0.119±0.005	0.509±0.028	0.224±0.009	0.135±0.006	0.292±0.014	0.437±0.014	0.003±0.001	0.431±0.015	
+ AEGIS	0.666±0.016	0.139±0.005	0.522±0.011	0.525±0.013	0.037±0.007	0.511±0.022	0.618±0.017	0.123±0.005	0.552±0.013	0.543±0.015	0.044±0.003	0.492±0.017	0.469±0.014	0.005±0.001	0.472±0.016	
+ GGAD	0.653±0.029	0.134±0.006	0.513±0.014	0.515±0.013	0.036±0.003	0.529±0.011	0.617±0.017	0.137±0.006	0.501±0.015	0.653±0.016	0.067±0.004	0.542±0.019	0.487±0.014	0.006±0.001	0.492±0.016	
+ CGenGA	0.575±0.019	0.071±0.004	0.462±0.016	0.620±0.015	0.053±0.004	0.581±0.016	0.916±0.019	0.712±0.018	0.682±0.015	0.825±0.019	0.421±0.011	0.632±0.018	0.501±0.015	0.012±0.002	0.501±0.006	
+ BAED	0.702±0.017	0.141±0.006	0.370±0.006	0.715±0.017	0.075±0.005	0.652±0.015	0.982±0.020	0.886±0.021	0.713±0.017	0.901±0.010	0.814±0.017	0.722±0.009	0.666±0.017	0.016±0.002	0.539±0.019	

Table 2: Overall comparison across two datasets in the inductive setting under dynamic graphs, with bold numbers indicating the best performance and underlined numbers highlighting the most competitive results(RQ2).

Dataset	Method	T=1		T=2		T=3		T=4	
		AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
Dgraph	GCN	0.4327±0.0105	0.0092±0.0113	0.4882±0.0127	0.0104±0.0142	0.5091±0.0139	0.0115±0.0128	0.5199±0.0116	0.0118±0.0097
	+AEGIS	<u>0.4684±0.0132</u>	0.0028±0.0087	0.5046±0.0115	0.0045±0.0091	<u>0.5352±0.0128</u>	0.0042±0.0083	<u>0.5492±0.0143</u>	0.0055±0.0096
	+GGAD	0.4471±0.0109	0.0019±0.0084	0.4986±0.0127	0.0037±0.0089	0.5169±0.0136	0.0043±0.0094	0.5441±0.0118	0.0057±0.0086
	+CGenGA	0.4524±0.0145	0.0095±0.0102	0.5001±0.0139	0.0109±0.0116	0.5136±0.0122	0.0118±0.0108	0.5377±0.0147	0.0124±0.0121
	+BAED	0.4817±0.0082	0.0102±0.0051	0.5249±0.0069	0.0115±0.0047	0.5417±0.0073	0.0128±0.0039	0.5623±0.0091	0.0136±0.0054
Elliptic	GCN	0.4927±0.0091	0.0317±0.0135	0.5123±0.0123	0.0525±0.0148	0.5087±0.0118	0.0478±0.0132	0.5139±0.0106	0.0496±0.0129
	+AEGIS	0.4977±0.0093	0.0542±0.0128	0.5275±0.0104	0.0701±0.0135	0.5128±0.0117	0.0694±0.0142	0.5318±0.0098	0.0867±0.0126
	+GGAD	0.5083±0.0125	0.0643±0.0149	0.0542±0.0136	<u>0.0721±0.0153</u>	<u>0.5727±0.0141</u>	<u>0.0804±0.0138</u>	<u>0.5974±0.0150</u>	<u>0.1084±0.0147</u>
	+CGenGA	0.5017±0.0108	0.0438±0.0119	<u>0.5429±0.0143</u>	0.0508±0.0132	0.5875±0.0135	0.0621±0.0124	0.5911±0.0116	0.0742±0.0139
	+BAED	0.5818±0.0043	0.1271±0.0026	0.6036±0.0038	0.1442±0.0031	0.5924±0.0029	0.1303±0.0025	0.6083±0.0035	0.1469±0.0028

5 Experiments

In this section, we answer the following five questions.

- **RQ1:** What is BAED’s overall performance compared to state-of-the-art methods?
- **RQ2:** How does BAED perform on dynamic graph datasets compared to state-of-the-art methods?
- **RQ3:** How does Anomaly-Guidance Embedding improve the performance of graph anomaly detection?
- **RQ4:** How does the weighting factor in Curriculum Anomaly Augmentation affect the performance of denoising models?
- **RQ5:** What advantages does BAED exhibit in generative data augmentation?
- **RQ6:** Is the distribution of the generated samples similar to the distribution of the original samples?

5.1 Experimental Setups

5.1.1 Datasets. In the experiments, we evaluate the performance of our approach on five large-scale, real-world graph datasets from diverse domains. Specifically, the datasets used include the transaction network from T-Finance [44], the social interaction network from Reddit [24], the Bitcoin transaction network from Elliptic [51],

the co-purchase network from Photo [33], and the financial network from DGraph [18]. Detailed statistics are provided in Appendix B.

5.1.2 Competing Methods. To validate the effectiveness of the BAED model, we compared it with three state-of-the-art generative anomaly detection models: AEGIS [5], GGAD [39], and CGenGA [32]. These models represent the latest advancements in GAD from the perspective of anomalous sample generation. Detailed information about these models can be found in Appendix C.

5.1.3 Evaluation Metric. Based on [36, 39], three commonly used complementary evaluation metrics are employed in this study: the Area Under the ROC Curve (AUROC), the Area Under the Precision-Recall Curve (AUPRC), and the F1-score. Higher AUROC, AUPRC, and F1-score values indicate better model performance.

5.1.4 Implementation Details. Experiments are conducted on four baseline GNN models (GCN [21], BWGNN [44], GraphSAGE [13], and BernNet [15]) and four negative sample generators (GGAD [39], AEGIS [5], CGenGA [32], and BAED). The datasets are split into training (70%), validation (15%), and test (15%) sets. Unlike traditional transductive graph learning, our approach adopts an inductive learning framework, where the model does not have access

to the test set during training. The graph is transformed into ego-graph for each node, with the ego-graph of test nodes excluded from the training set. This design promotes better generalization to unseen data, which is a key advantage of inductive learning. The Adam optimizer is used, and the learning rate is tuned within the range of $[10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}]$.

5.2 Improvement on Anomaly Detection (RQ1)

Table 1 demonstrates BAED’s superior performance over existing data augmentation methods. BAED-generated synthetic data enhances traditional GNN models by up to twice their original performance, showcasing strong generalization capabilities.

BAED exhibits exceptional robustness on challenging datasets. On the sparse Elliptic dataset, while AEGIS and GGAD degraded BWGNN’s AUROC by 12.99% and 15.01% respectively, BAED achieved a 4.97% improvement. More dramatically, on the large T-Finance dataset, BAED boosted BWGNN’s AUROC and AUPRC by 90.93% and 84.84%, substantially outperforming competing methods. These results indicate that BAED generates data better aligned with true anomaly distributions, enhancing detector robustness in noisy and sparse environments.

5.3 Performance on Dynamic Graph Datasets (RQ2)

Table 2 presents the experimental results of BAED and baseline methods under an incremental learning paradigm on dynamic graph datasets, which simulates real-world streaming data scenarios by sequentially training and testing models on consecutive temporal segments to mimic non-stationary graph environments. BAED demonstrates superior and stable performance across all time steps on both datasets, exhibiting significantly lower variance in its results compared to other methods, which highlights its robust adaptability to evolving graph structures.

5.4 The Impact of Anomaly-Guidance (RQ3)

To evaluate the Anomaly-Guidance Embedding in BAED, we compared it with two variants: BUED, a balanced unconditional ego-graph diffusion model, and the **original** model, which lacks data augmentation. Experiments on the Reddit and Photo datasets (Figure 3) demonstrate that both BAED and BUED, with curriculum anomaly augmentation, outperform the original model trained on imbalanced data. This highlights the crucial role of specific anomaly self-graphs in enhancing detection performance. Furthermore, BAED outperforms BUED, validating the effectiveness of our dynamic anomaly-guided graph generation technique, which customizes augmented samples based on specific anomaly patterns and improves model performance.

5.5 Ablation Study on Curriculum Anomaly Augmentation (RQ4)

To evaluate the Curriculum Anomaly Augmentation strategy, we conducted an ablation study on four baseline models, focusing on the impact of different anomaly-guided embedding methods. We tested three weight assignment strategies: random, average, and curriculum-based, which dynamically adjusts embeddings for

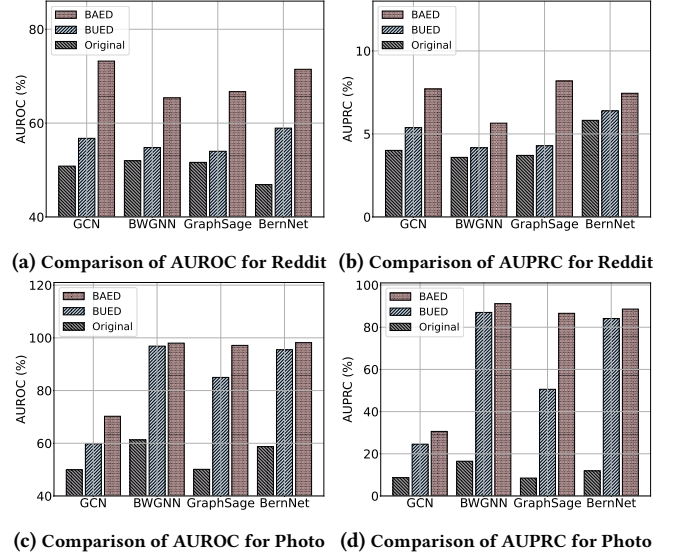


Figure 3: Impact of Anomaly-Guidance Embedding.

Table 3: Performance on Different Guidance Embedding

Model	Weighting Strategy	Reddit		Elliptic	
		AUROC	AUPRC	AUROC	AUPRC
GCN	random	0.6075	0.0405	0.5745	0.1329
	average	0.6225	0.0696	0.6016	0.1437
	curriculum-based	0.7322	0.0771	0.6706	0.1537
BWGNN	random	0.5316	0.0425	0.5713	0.0852
	average	0.5642	0.0401	0.5878	0.0761
	curriculum-based	0.6538	0.0565	0.6689	0.1019
GraphSage	random	0.5954	0.0406	0.5729	0.0957
	average	0.5979	0.0405	0.5898	0.1312
	curriculum-based	0.6669	0.0819	0.5927	0.1340
BernNet	random	0.5565	0.0490	0.5150	0.0828
	average	0.5930	0.0384	0.5875	0.1150
	curriculum-based	0.7147	0.0745	0.7016	0.1412

specific anomaly patterns. The results in Table 3 show that the curriculum-based strategy outperforms both random and average methods across all metrics. The latter failed to prioritize anomaly patterns at each training stage, resulting in less effective guidance. In contrast, the dynamic weight adjustment significantly improved performance by emphasizing core anomalies.

Additionally, training loss curves for the Reddit dataset (Figure 4) show that the curriculum-based method achieves smoother and lower loss convergence, demonstrating more effective guidance during training.

5.6 The Impact of BAED on data diversity (RQ5)

To evaluate the diversity of anomalous samples generated by BAED, we compared BAED-generated ego-graphs with randomly sampled ones for balancing positive and negative samples. Results in Figure 5 show that BAED-generated ego-graphs significantly outperform

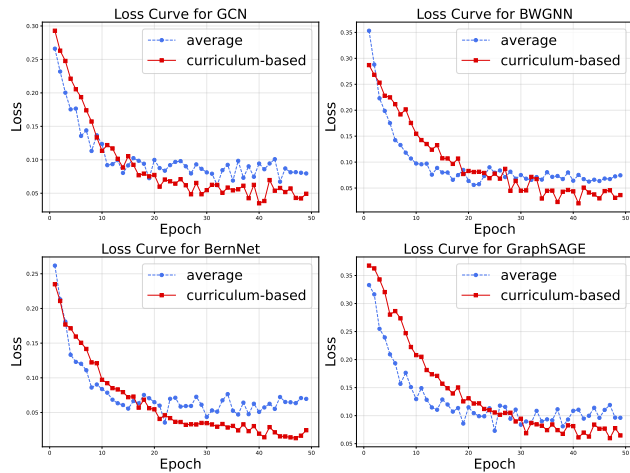


Figure 4: Comparison of model loss curves for different weighting strategy on reddit dataset

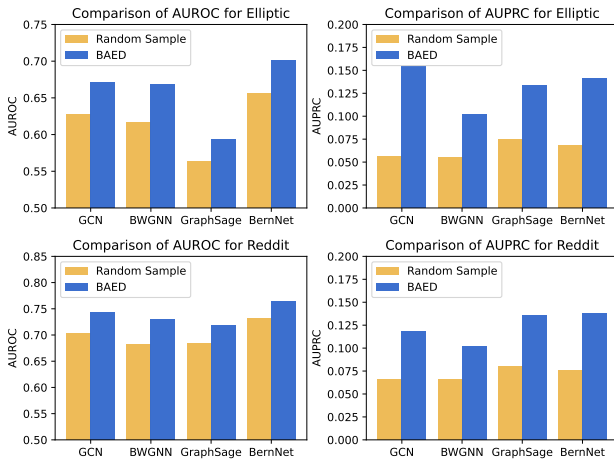


Figure 5: Comparison of different augmentation methods

random samples in anomaly detection, offering greater sample diversity. This highlights the effectiveness of BAED’s diffusion-based mechanism in generating anomaly-specific ego-graphs that capture the local structure of anomalous nodes, enhancing sample diversity and improving model generalization.

5.7 Visualization of Data Distribution (RQ6)

To compare the distribution of embeddings between generated and original anomalous ego-graphs, we used t-SNE for dimensionality reduction (Figure 6). The results show that the embeddings of BAED-generated anomalous samples closely match those of the original samples, with nearly identical distributions. This indicates that BAED effectively preserves the distributional characteristics of the original data through its discrete diffusion process, ensuring the validity of the generated samples.

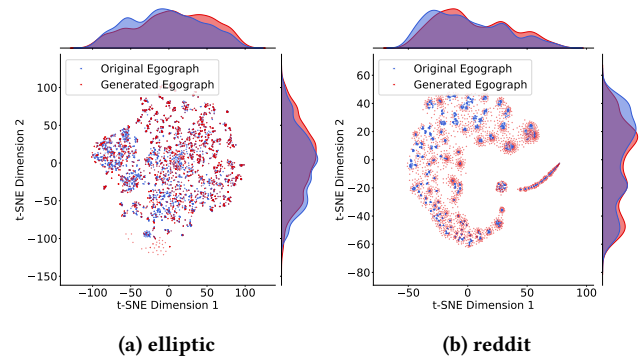


Figure 6: Comparison of data distribution.

6 Conclusions and Limitations

This paper presented **Balanced Anomaly-Guided Ego-graph Diffusion** (BAED), a novel framework that addresses the dual challenges of dynamic graph structures and class imbalance in inductive graph anomaly detection. By combining a discrete ego-graph diffusion model with curriculum anomaly augmentation, BAED achieves dynamic adaptability and enhanced generalization to unseen graph data. Extensive experiments demonstrate its state-of-the-art performance, underscoring the potential of integrating topology-aware augmentation with adaptive learning strategies for anomaly detection tasks.

While BAED introduces significant advancements, it has a few limitations that warrant future exploration. First, the generated ego-graphs primarily focus on structural anomalies. Expanding the framework to address more nuanced feature-based anomalies could broaden its applicability. Second, BAED has been validated in the context of extreme class imbalance for anomaly detection tasks. Exploring its potential to handle other imbalance scenarios, such as imbalanced ego-graph classification or rare event prediction, could further demonstrate its versatility.

Acknowledgments

This research was supported by the National Key R&D Program of China (No. 2023YFC3304701), NSFC (No.6250072448, No.62272466, U24A20233). It was also supported by the Big Data and Responsible Artificial Intelligence for National Governance, Renmin University of China, and the Joint Academy on Future Humanity, Renmin University of China and Westlake University.

References

- [1] Sambaran Bandyopadhyay, N. Lokesh, and M. Narasimha Murty. 2019. Outlier Aware Network Embedding for Attributed Networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*. AAAI Press, 12–19.
- [2] Ziwei Chai, Siqi You, Yang Yang, Shiliang Pu, Jiarong Xu, Haoyang Cai, and Weihao Jiang. 2022. Can Abnormality be Detected by Graph Neural Networks?. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, Luc De Raedt (Ed.). ijcai.org, 1945–1951.
- [3] Xiaohui Chen, Jiaying He, Xu Han, and Liping Liu. 2023. Efficient and Degree-Guided Graph Generation via Discrete Diffusion Modeling. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*

- (*Proceedings of Machine Learning Research, Vol. 202*), Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 4585–4610.
- [4] Dawei Cheng, Xiaoyang Wang, Ying Zhang, and Liqing Zhang. 2022. Graph Neural Network for Fraud Detection via Spatial-Temporal Attention. *IEEE Trans. Knowl. Data Eng.* 34, 8 (2022), 3800–3813.
 - [5] Kaize Ding, Jundong Li, Nitin Agarwal, and Huan Liu. 2021. Inductive anomaly detection on attributed networks. In *Proceedings of the twenty-ninth international conference on international joint conferences on artificial intelligence*. 1288–1294.
 - [6] Kaize Ding, Jundong Li, Rohit Bhanushali, and Huan Liu. 2019. Deep Anomaly Detection on Attributed Networks. In *Proceedings of the 2019 SIAM International Conference on Data Mining, SDM 2019, Calgary, Alberta, Canada, May 2-4, 2019*, Tanya Y. Berger-Wolf and Nitesh V. Chawla (Eds.). SIAM, 594–602. doi:10.1137/1.9781611975673.67
 - [7] Yingtong Dou, Zhiwei Liu, Li Sun, Yutong Deng, Hao Peng, and Philip S. Yu. 2020. Enhancing Graph Neural Network-based Fraud Detectors against Camouflaged Fraudsters. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, Mathieu d'Aquin, Stefan Dietze, Claudia Hauff, Edward Curry, and Philippe Cudré-Mauroux (Eds.). ACM, 315–324.
 - [8] Lun Du, Yun Wang, Guojie Song, Zhicong Lu, and Junshan Wang. 2018. Dynamic Network Embedding : An Extended Approach for Skip-gram based Network Embedding. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, Jérôme Lang (Ed.). ijcai.org, 2086–2092.
 - [9] Haoyi Fan, Fengbin Zhang, and Zuyong Li. 2020. Anomalydae: Dual Autoencoder for Anomaly Detection on Attributed Networks. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*. IEEE, 5685–5689. doi:10.1109/ICASSP40776.2020.9053387
 - [10] Yuan Gao, Xiang Wang, Xiangnan He, Zhenguang Liu, Huamin Feng, and Yongdong Zhang. 2023. Addressing Heterophily in Graph Anomaly Detection: A Perspective of Graph Spectrum. In *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*, Ying Ding, Jie Tang, Juan F. Sequeda, Lora Aroyo, Carlos Castillo, and Geert-Jan Houben (Eds.). ACM, 1528–1538.
 - [11] Palash Goyal, Sujit Rokka Chhetri, and Arquimedes Canedo. 2020. dyngraph2vec: Capturing network dynamics using dynamic graph representation learning. *Knowl. Based Syst.* 187 (2020).
 - [12] Palash Goyal, Nitin Kamra, Xinran He, and Yan Liu. 2018. DynGEM: Deep Embedding Method for Dynamic Graphs. *CoRR abs/1805.11273* (2018).
 - [13] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).
 - [14] William L. Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 1024–1034.
 - [15] Mingguo He, Zhewei Wei, Hongteng Xu, et al. 2021. Bernnet: Learning arbitrary graph spectral filters via bernstein approximation. *Advances in Neural Information Processing Systems* 34 (2021), 14239–14251.
 - [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.).
 - [17] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay S. Pande, and Jure Leskovec. 2020. Strategies for Pre-training Graph Neural Networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
 - [18] Xuanwen Huang, Yang Yang, Yang Wang, Chunping Wang, Zhisheng Zhang, Jiarong Xu, Lei Chen, and Michalis Vazirgiannis. 2022. Dgraph: A large-scale financial dataset for graph anomaly detection. *Advances in Neural Information Processing Systems* 35 (2022), 22765–22777.
 - [19] Zhixuan Jia, Yushun Fan, Chunyu Wei, and Ruyi Yan. 2023. A spatial-temporal hypergraph based method for service recommendation in the Mobile Internet of Things-enabled service platform. *Adv. Eng. Informatics* 57 (2023), 102038.
 - [20] Zhixuan Jia, Yushun Fan, Jia Zhang, Chunyu Wei, Ruyi Yan, and Xing Wu. 2023. Improving Next Location Recommendation Services With Spatial-Temporal Multi-Group Contrastive Learning. *IEEE Trans. Serv. Comput.* 16, 5 (2023), 3467–3478.
 - [21] Thomas N Kipf and Max Welling. 2016. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv preprint arXiv:1609.02907* (2016).
 - [22] Thomas N. Kipf and Max Welling. 2016. Variational Graph Auto-Encoders. *CoRR abs/1611.07308* (2016).
 - [23] Oluwasanmi O Koyejo, Nagarajan Natarajan, Pradeep K Ravikumar, and Inderjit S Dhillon. 2014. Consistent Binary Classification with Generalized Performance Metrics. *Advances in Neural Information Processing Systems* 27 (2014).
 - [24] Srijan Kumar, Xikun Zhang, and Jure Leskovec. 2019. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 1269–1278.
 - [25] Jundong Li, Harsh Dani, Xia Hu, and Huan Liu. 2017. Radar: Residual Analysis for Anomaly Detection in Attributed Networks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, Carles Sierra (Ed.). ijcai.org, 2152–2158.
 - [26] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. 2012. Isolation-Based Anomaly Detection. *ACM Trans. Knowl. Discov. Data* 6, 1 (2012), 3:1–3:39.
 - [27] Kay Liu, Yingtong Dou, Yue Zhao, Xueying Ding, Xiyang Hu, Ruitong Zhang, Kaize Ding, Canyu Chen, Hao Peng, Kai Shu, Lichao Sun, Jundong Li, George H. Chen, Zhihao Jia, and Philip S. Yu. 2022. BOND: Benchmarking Unsupervised Outlier Node Detection on Static Attributed Graphs. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (Eds.).
 - [28] Yang Liu, Xiang Ao, Zidi Qin, Jianfeng Chi, Jinghua Feng, Hao Yang, and Qing He. 2021. Pick and Choose: A GNN-based Imbalanced Learning Approach for Fraud Detection. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, Jure Leskovec, Marko Grobelnik, Marc Najork, Jie Tang, and Leila Zia (Eds.). ACM / IW3C2, 3168–3177.
 - [29] Yixin Liu, Zhao Li, Shirui Pan, Chen Gong, Chuan Zhou, and George Karypis. 2022. Anomaly Detection on Attributed Networks via Contrastive Self-Supervised Learning. *IEEE Trans. Neural Networks Learn. Syst.* 33, 6 (2022), 2378–2392.
 - [30] Yuanfu Lu, Xunqiang Jiang, Yuan Fang, and Chuan Shi. 2021. Learning to Pre-train Graph Neural Networks. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*. AAAI Press, 4276–4284.
 - [31] Xiaoxiao Ma, Ruikun Li, Fanzhen Liu, Kaize Ding, Jian Yang, and Jia Wu. 2024. Graph Anomaly Detection with Few Labels: A Data-Centric Approach. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2024, Barcelona, Spain, August 25-29, 2024*, Ricardo Baeza-Yates and Francesco Bonchi (Eds.). ACM, 2153–2164.
 - [32] Xiaoxiao Ma, Ruikun Li, Fanzhen Liu, Kaize Ding, Jian Yang, and Jia Wu. 2024. Graph anomaly detection with few labels: A data-centric approach. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2153–2164.
 - [33] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*. 43–52.
 - [34] Soroor Motie and Bijan Raahemi. 2024. Financial fraud detection using graph neural networks: A systematic review. *Expert Syst. Appl.* 240 (2024), 122156.
 - [35] Giang Hoang Nguyen, John Boaz Lee, Ryan A. Rossi, Nesreen K. Ahmed, Eunye Koh, and Sungchul Kim. 2018. Continuous-Time Dynamic Network Embeddings. In *Companion of the The Web Conference 2018 on The Web Conference 2018, WWW 2018, Lyon, France, April 23-27, 2018*, Pierre-Antoine Champin, Fabien Gandon, Mounia Lalmas, and Panagiotis G. Ipeirotis (Eds.). ACM, 969–976.
 - [36] Guansong Pang, Anton van den Hengel, Chunhua Shen, and Longbing Cao. 2021. Toward deep supervised anomaly detection: Reinforcement learning from partially labeled anomaly data. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*. 1298–1308.
 - [37] Joonhyung Park, Jaeyun Song, and Eunho Yang. 2022. GraphENS: Neighbor-Aware Ego Network Synthesis for Class-Imbalanced Node Classification. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
 - [38] Zhen Peng, Minnan Luo, Jundong Li, Huan Liu, and Qinghua Zheng. 2018. ANOMALOUS: A Joint Modeling Approach for Anomaly Detection on Attributed Networks. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, Jérôme Lang (Ed.). ijcai.org, 3513–3519.
 - [39] Hezhe Qiao, Qingsong Wen, Xiaoli Li, Ee-Peng Lim, and Guansong Pang. 2024. Generative Semi-supervised Graph Anomaly Detection. *arXiv preprint arXiv:2402.11887* (2024).
 - [40] Jiezhong Qiu, Jian Tang, Hao Ma, Yuxiao Dong, Kuansan Wang, and Jie Tang. 2018. DeepInf: Social Influence Prediction with Deep Learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, Yike Guo and Faisal Farooq (Eds.). ACM, 2110–2119.
 - [41] Mayu Sakurada and Takehisa Yairi. 2014. 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, QLD, Australia, December 2, 2014*, Ashfaqur Rahman, Jeremiah D. Deng, and Juyong Li (Eds.). ACM, 4.

- [42] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6–11 July 2015 (JMLR Workshop and Conference Proceedings, Vol. 37)*, Francis R. Bach and David M. Blei (Eds.). JMLR.org, 2256–2265.
- [43] Jianheng Tang, Jiajin Li, Ziqi Gao, and Jia Li. 2022. Rethinking Graph Neural Networks for Anomaly Detection. In *International Conference on Machine Learning, ICML 2022, 17–23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.). PMLR, 21076–21089.
- [44] Jianheng Tang, Jiajin Li, Ziqi Gao, and Jia Li. 2022. Rethinking graph neural networks for anomaly detection. In *International Conference on Machine Learning, PMLR, 21076–21089*.
- [45] Atiq ur Rehman and Samir Brahim Belhaouari. 2021. Unsupervised outlier detection in multidimensional data. *J. Big Data* 8, 1 (2021), 80.
- [46] Clément Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. 2023. DiGress: Discrete Denoising diffusion for graph generation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1–5, 2023*. OpenReview.net.
- [47] Cheng Wang and Hangyu Zhu. 2022. Wrongdoing Monitor: A Graph-Based Behavioral Anomaly Detection in Cyber Security. *IEEE Trans. Inf. Forensics Secur.* 17 (2022), 2703–2718.
- [48] Daixin Wang, Peng Cui, and Wenwu Zhu. 2016. Structural Deep Network Embedding. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13–17, 2016*, Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi (Eds.). ACM, 1225–1234.
- [49] Daixin Wang, Yuan Qi, Jianbin Lin, Peng Cui, Quanhui Jia, Zhen Wang, Yanming Fang, Quan Yu, Jun Zhou, and Shuang Yang. 2019. A Semi-Supervised Graph Attentive Network for Financial Fraud Detection. In *2019 IEEE International Conference on Data Mining, ICDM 2019, Beijing, China, November 8–11, 2019*, Jianyong Wang, Kyuseok Shim, and Xindong Wu (Eds.). IEEE, 598–607.
- [50] Jianyu Wang, Rui Wen, Chunming Wu, Yu Huang, and Jian Xiong. 2019. FdGars: Fraudster Detection via Graph Convolutional Networks in Online App Review System. In *Companion of The 2019 World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13–17, 2019*, Sihem Amer-Yahia, Mohammad Mahdian, Ashish Goel, Geert-Jan Houben, Kristina Lerman, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (Eds.). ACM, 310–316.
- [51] Mark Weber, Giacomo Domeniconi, Jie Chen, Daniel Karl I Weidele, Claudio Bellei, Tom Robinson, and Charles E Leiserson. 2019. Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics. *arXiv preprint arXiv:1908.02591* (2019).
- [52] Chunyu Wei, Bing Bai, Kun Bai, and Fei Wang. 2022. GSL4Rec: Session-based Recommendations with Collective Graph Structure Learning and Next Interaction Prediction. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 – 29, 2022*, Frédérique Laforest, Raphaël Troncy, Elena Simperl, Deepak Agarwal, Aristides Gionis, Ivan Herman, and Lionel Médini (Eds.). ACM, 2120–2130.
- [53] Chunyu Wei, Yushun Fan, Zhixuan Jia, and Jia Zhang. 2024. Cross-View Graph Alignment for Mashup Recommendation. *IEEE Trans. Serv. Comput.* 17, 5 (2024), 2151–2164.
- [54] Chunyu Wei, Yushun Fan, and Jia Zhang. 2022. High-Order Social Graph Neural Network for Service Recommendation. *IEEE Trans. Netw. Serv. Manag.* 19, 4 (2022), 4615–4628.
- [55] Chunyu Wei, Yushun Fan, and Jia Zhang. 2023. Time-Aware Service Recommendation With Social-Powered Graph Hierarchical Attention Network. *IEEE Trans. Serv. Comput.* 16, 3 (2023), 2229–2240.
- [56] Chunyu Wei, Yushun Fan, Jia Zhang, Zhixuan Jia, and Ruyi Yan. 2024. Dynamic Relation Graph Learning for Time-Aware Service Recommendation. *IEEE Trans. Netw. Serv. Manag.* 21, 2 (2024), 1503–1517.
- [57] Chunyu Wei, Wenji Hu, Xingjia Hao, Yunhai Wang, Yueguo Chen, Bing Bai, and Fei Wang. 2025. Graph Evidential Learning for Anomaly Detection. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining, V.2, KDD 2025, Toronto ON, Canada, August 3–7, 2025*, Luiza Antonie, Jian Pei, Xiaohui Yu, Flavio Chierichetti, Hady W. Lauw, Yizhou Sun, and Srinivasan Parthasarathy (Eds.). ACM, 3122–3133.
- [58] Chunyu Wei, Jian Liang, Bing Bai, and Di Liu. 2022. Dynamic Hypergraph Learning for Collaborative Filtering. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, Atlanta, GA, USA, October 17–21, 2022*, Mohammad Al Hasan and Li Xiong (Eds.). ACM, 2108–2117.
- [59] Chunyu Wei, Jian Liang, Di Liu, Zehui Dai, Mang Li, and Fei Wang. 2023. Meta Graph Learning for Long-tail Recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA, USA, August 6–10, 2023*, Ambuj K. Singh, Yizhou Sun, Leman Akoglu, Dimitrios Gunopoulos, Xifeng Yan, Ravi Kumar, Fatma Ozcan, and Jieping Ye (Eds.). ACM, 2512–2522.
- [60] Chunyu Wei, Jian Liang, Di Liu, and Fei Wang. 2022. Contrastive Graph Structure Learning via Information Bottleneck for Recommendation. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 – December 9, 2022*, Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (Eds.).
- [61] Chunyu Wei, Yu Wang, Bing Bai, Kai Ni, David Brady, and Lu Fang. 2023. Boosting Graph Contrastive Learning via Graph Contrastive Saliency. In *International Conference on Machine Learning, ICML 2023, 23–29 July 2023, Honolulu, Hawaii, USA (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 36839–36855.
- [62] Da Xu, Chuanwei Ruan, Evren Körpeoglu, Sushant Kumar, and Kannan Achan. 2020. Inductive representation learning on temporal graphs. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*. OpenReview.net.
- [63] Fan Xu, Nan Wang, Hao Wu, Xuezhi Wen, Xibin Zhao, and Hai Wan. 2024. Revisiting Graph-Based Fraud Detection in Sight of Heterophily and Spectrum. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20–27, 2024, Vancouver, Canada*, Michael J. Wooldridge, Jennifer G. Dy, and Sriaram Natarajan (Eds.). AAAI Press, 9214–9222.
- [64] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*. OpenReview.net.
- [65] Xiaowei Xu, Nurcan Yuruk, Zhidan Feng, and Thomas A. J. Schweiger. 2007. SCAN: a structural clustering algorithm for networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Jose, California, USA, August 12–15, 2007*, Pavel Berkhin, Rich Caruana, and Xindong Wu (Eds.). ACM, 824–833.
- [66] Rose Yu, Huidia Qiu, Zhen Wen, Ching-Yung Lin, and Yan Liu. 2016. A Survey on Social Media Anomaly Detection. *SIGKDD Explor.* 18, 1 (2016), 1–14.
- [67] Ge Zhang, Jia Wu, Jian Yang, Amin Beheshti, Shan Xue, Chuan Zhou, and Quan Z. Sheng. 2021. FRAUDRE: Fraud Detection Dual-Resistant to Graph Inconsistency and Imbalance. In *IEEE International Conference on Data Mining, ICDM 2021, Auckland, New Zealand, December 7–10, 2021*, James Bailey, Pauli Miettinen, Yun Sing Koh, Dacheng Tao, and Xindong Wu (Eds.). IEEE, 867–876.
- [68] Le Zhang, Wei Cheng, Ji Xing, Xuefeng Chen, Zelin Nie, Shuo Zhang, Junying Hong, and Zhao Xu. 2023. Self-Supervised Variational Graph Autoencoder for System-Level Anomaly Detection. *IEEE Trans. Instrum. Meas.* 72 (2023), 1–11.
- [69] Xuan Zhang, Chunyu Wei, Ruyi Yan, Yushun Fan, and Zhixuan Jia. 2024. Large Language Model Ranker with Graph Reasoning for Zero-Shot Recommendation. In *Artificial Neural Networks and Machine Learning - ICANN 2024 - 33rd International Conference on Artificial Neural Networks, Lugano, Switzerland, September 17–20, 2024, Proceedings, Part V (Lecture Notes in Computer Science, Vol. 15020)*, Michael Wand, Kristina Malinová, Jürgen Schmidhuber, and Igor V. Tetko (Eds.). Springer, 356–370.
- [70] Xuan Zhang, Ruyi Yan, Yushun Fan, Jia Zhang, Hongxing Yuan, and Chunyu Wei. 2025. Cross-domain attention transfer network for recommendation. *Adv. Eng. Informatics* 68 (2025), 103667.
- [71] Tianxiang Zhao, Xiang Zhang, and Suhang Wang. 2021. GraphSMOTE: Imbalanced Node Classification on Graphs with Graph Neural Networks. In *WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, March 8–12, 2021*, Liane Lewin-Eytan, David Carmel, Elad Yom-Tov, Eugene Agichtein, and Evgeniy Gabrilovich (Eds.). ACM, 833–841.
- [72] Yu Zheng, Ming Jin, Yixin Liu, Lianhua Chi, Khoa Tran Phan, and Yi-Ping Phoebe Chen. 2023. Generative and Contrastive Self-Supervised Learning for Graph Anomaly Detection. *IEEE Trans. Knowl. Data Eng.* 35, 12 (2023), 12220–12233.
- [73] Yuan Zuo, Guannan Liu, Hao Lin, Jia Guo, Xiaoqian Hu, and Junjie Wu. 2018. Embedding Temporal Network via Neighborhood Formation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19–23, 2018*, Yike Guo and Faisal Farooq (Eds.). ACM, 2857–2866.

A Algorithm

The implementation of BAED consists of three core phases: **Pre-trained GAE** for representation learning; **Diffusion Model Training and Sample Generation** as detailed in Algorithms 1 and 2; and **Anomaly Detection** according to the procedure in Algorithm 3.

B Dataset Overview

We utilize five public datasets from diverse domains: Dgraph, T-Finance, Photo, Reddit, and Elliptic. Each exhibits a significant class

Algorithm 2 BAED: Training of Diffusion

```

1: Input: node features  $X^0$ , topological structure  $A^0$ , diffusion steps  $T$ , forward transition distributions  $q(A^t|A^{t-1})$  and  $q(A^t|A^0)$ , ego-graph embedding function  $e^t = f_\theta(A^t)$ .
2: while training do
3:   Compute ego-graph embedding  $e^0 = f_\theta(X^0, A^0)$ 
4:   Sample  $t \sim \mathcal{U}[1, T]$ , and set  $t - 1$  immediately
5:   Draw  $A^{t-1} \sim q(A^{t-1}|A^0)$ , then draw  $A^t \sim q(A^t|A^{t-1})$ 
6:   Perform gradient descent over  $\theta$  with gradient  $\nabla_\theta \log p_\theta(A^{t-1}|A^t, e^0)$ 
7: end while

```

Algorithm 3 Training of Graph Anomaly Detection

```

1: Initialize training parameters: learning rate lr, number of epochs  $E$ , optimizer Adam, BAED denoising model.
2: Split dataset into training set, validation set, and test set.
3: for  $e = 1$  to  $E$  do
4:   for each batch  $\mathcal{B}$  in training set do
5:     Compute the difference in the number of positive and negative samples  $N$ 
6:     Compute the anomaly-guidance embedding for negative samples in the batch.
7:     if  $e > 0$  then
8:       Compute weight factor  $w_i$  and time factor  $h(t)$  for each sample  $i$  in the batch:  $w_i = \frac{1}{1+e^{-\alpha(l_i-\beta)}} \cdot h(t)$ .
9:     else
10:      Set weight factor  $w_i$  to the average value for all samples.
11:    end if
12:    Compute weighted embedding  $\tilde{\mathbf{h}}^g = \sum_{i=1}^N \frac{w_i}{\sum_{j=1}^N w_j} \cdot \text{GIN}(\mathcal{G}_K^i)$ .
13:    Generate  $N$  negative samples using BAED denoising model with  $\tilde{\mathbf{h}}^g$ .
14:    Combine the original negative samples with the generated samples in the batch.
15:    Compute cross-entropy loss  $\nabla_\theta \mathcal{L}$  to update model.
16:  end for
17: end for
18:
19: return trained model.

```

Table 4: The statistics of the graph anomaly detection datasets

Dataset	#Nodes	#Edges	#Features	#Degree	#Anomalies	Anomaly Ratio
Dgraph	3,700,550	4,300,999	17	1.16	15,509	1.3%
T-Finance	39,357	21,222,543	10	539.23	1,803	4.6%
Photo	7,535	119,043	745	15.8	698	9.2%
Reddit	10,984	168,016	64	15.3	366	3.3%
Elliptic	203,769	234,355	166	1.15	4,545	9.8%

imbalance between positive and negative samples, aligning with our research focus. Details are summarized in Table 4.

C Baseline Methods

In this section, we present and compare the properties of the baseline methods used in this paper.

- **AEGIS** [5]: AEGIS introduces a graph discretization layer to learn anomaly-aware node representations through an attentional mechanism. It employs generative adversarial learning to enable inductive anomaly detection without retraining.
- **GGAD** [39]: GGAD is a semi-supervised method that enhances detection by generating pseudo-outlier nodes as negative samples. It optimizes these nodes using asymmetric local affinity and egocentric proximity priors to align with real anomalies.
- **CGenGA** [32]: CGenGA utilizes a denoising diffusion model to generate synthetic nodes matching the original graph distribution. It enhances existing detectors by optimizing the generation process through classifier guidance.

D Experimental Details

All experiments run on an NVIDIA Tesla A800 GPU using PyTorch and Adam optimizer. Results are averaged over five runs with random seeds 15, 42, 63, 87, 94.

Diffusion Model. Our conditional-guided generator uses node degree and an 8-head self-attention mechanism. It trains for 300 epochs with a diffusion dimension of 64, 128 steps, a learning rate of $1e-4$, linear noise scheduling, and a dropout rate of 0.1.

Training. Batches are balanced using generated negative samples, while the test set retains the original imbalance. Synthetic node features combine original values with Gaussian noise $\mathcal{N}(0, 1)$. Downstream models train for 50 epochs, with learning rates tuned in 10^{-3} – 10^{-6} .

E Computational Efficiency for Real-time Graph Anomaly Detection

We analyze the time complexity of BAED’s training phase. The per-epoch complexity is:

$$O(N \cdot L \cdot \bar{e} \cdot d + B \cdot T \cdot \bar{e} \cdot d) \quad (24)$$

where N is the number of ego-graphs, L the GNN layers, \bar{e} the average edges per ego-graph, d the feature dimension, B the batch size, and T the diffusion steps.

Therefore, BAED based on diffusion model demands high computational efficiency, especially for large-scale dynamic networks. We evaluate the runtime performance of all methods across five datasets under GPU acceleration, with results presented in Table 5.

Table 5: Runtime comparison (seconds) across datasets on GPU

Model	Elliptic	Dgraph	Reddit	Photo	T-Finance
AEGIS	5,761	2,247	261	498	15,258
GGAD	5,271	2,713	463	487	9,345
CGenGA	11,233	14,474	1,598	2,986	30,650
BAED	8,641	10,802	1,142	2,489	27,864

While BAED requires more time than lighter baselines (AEGIS, GGAD), it achieves a 27.98% average runtime reduction compared to CGenGA. This gain is due to our anomaly-guided generation, which focuses the diffusion process. The improved detection performance justifies the computational cost for accuracy-critical applications like fraud detection.