GRANULARITY-AWARE HIERARCHICAL INFORMATION ALIGNMENT FOR MULTI-GRANULARITY GRAPH ANOMALY DETECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-granularity graph anomaly detection is a recently emerging issue that aims to design a unified solution for the detection of multi-granularity anomalies in graphs. Existing methods usually employ subgraph sampling and unified training objectives, achieving promising performance. However, they ignore hierarchical information in latent representations, making the latent-space evolution less discriminative. Moreover, they fail to explore interdependencies across granularities in specific detection tasks, which could further regularize the representations. To address these limitations, we propose a granularity-aware hierarchical information alignment framework for multi-granularity GAD (GaHIGAD). Specifically, GaHIGAD explores latent evolutionary information about hidden representations based on Information Bottleneck (IB) theory. In addition, GaHIGAD integrates cross-granularity representations to exploit latent consistency dependencies and achieve representation regularization. Extensive experiments on comprehensive public datasets demonstrate the effectiveness and advancement of GaHIGAD compared to state-of-the-art baselines.

1 INTRODUCTION

Graph Anomaly Detection (GAD) aims to uncover graph entities that depart substantially from the prevailing patterns of normal data in terms of topology, attributes, or their joint distributions. Owing to its ability to reveal subtle yet consequential irregularities, GAD has become an important technique in a wide range of applications, including rumor identification (Zhang et al., 2026b), social media analysis (Mohawesh et al., 2026), and financial risk management (Zhang et al., 2025). Nevertheless, most existing GAD approaches are developed with a predefined granularity, such as node granularity. Such limitations result in highly specialized models with limited scalability when anomaly detection needs to operate across multiple granularities in real-world applications.

Recent studies have increasingly focused on developing unified solutions for multi-granularity GAD (Song et al., 2025; Wang et al., 2026; Han et al., 2026) to improve the scalability and adaptability of graph-based detection. Rather than specializing in a predefined objective, they unify different granularities via diverse techniques, such as energy models Song et al. (2025), Wang et al. (2026), and adversarial learning Han et al. (2026), providing promising unification solutions for multi-granularity GAD. However, these methods fail to deeply grasp the complex relationships in hierarchical representation space across granularities, neglecting the evolutionary regularization of samples, which might even destroy the inherent correlations within the data.

Therefore, although many studies have proposed promising solutions for the unification of multi-granularity GAD using both traditional and deep learning tools recently, there are still some limitations in existing methods. On the one hand, most studies achieve granularity unification through diverse sampling strategies. Although they can simply generate a consensus across different granularities, they result in limited exploration of complex hierarchical information. On the other hand, prevalent solutions can generate granularity-specific representations or integrate multi-granularity representations with unified training objectives, but they fail to explore the inter-dependencies across granularities in granularity-specific GAD tasks.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

To address these challenges, we propose a granularity-aware hierarchical information alignment framework for multi-granularity GAD (GaHIGAD). To solve the first limitation, GaHIGAD introduces Information Bottleneck (IB) theory and exploits latent evolutionary information via hidden representations. For the second limitation, GaHIGAD integrates representations from different granularities and develops consistency adjustment for representation regularization and the exploration of cross-granularity dependencies. Our contributions are as follows:

- **A new paradigm**. We make the first attempt to study multi-granularity GAD using IB theory, which fully explores the influence of consistency of hidden representations and the utility of cross-granularity representations.
- **A novel framework**. We propose GaHIGAD, developing an IB loss and consistency adjustment for optimization that considers both inter-granularity latent consistency and cross-granularity correlation on the granularity-specific task.
- **Advanced performance**. Comprehensive experiments on various public datasets demonstrate the advancement and effectiveness of GaHIGAD compared to state-of-the-art baselines.

2 Related Work

Traditional GAD. Prior research in graph anomaly detection (GAD) broadly falls into supervised and unsupervised paradigms. Supervised techniques leverage explicit label guidance to effectively isolate complex graph attributes and topological patterns. For instance, GE-GNN (Zhang & Luo, 2025) employs dual-level edge augmentation paired with a gating architecture to extract anomalous signals. Unsupervised GAD predominantly relies on contrastive frameworks (Zhang & Luo, 2026) or reconstruction objectives (Chen et al., 2026b). To enhance latent feature representation, recent advancements have integrated sophisticated mechanisms such as diffusion models (Wei et al., 2026a; Xiao et al., 2025) and adversarial paradigms (Chen et al., 2026a). *Despite this promising progress, existing methods remain largely tailored to specific tasks, particularly node-level anomaly detection, which inherently restricts their scalability toward edge- and graph-level anomaly detection.*

Multi-granularity Graph Anomaly Detection. To tackle multi-granularity GAD, various innovative frameworks have recently emerged. BOURNE (Liu et al., 2024) constructs dual hypergraph architectures to concurrently identify anomalous signals at both node and edge scales. To standardize cross-granularity analysis, UniGAD (Lin et al., 2024) projects heterogeneous inputs into a cohesive subgraph space using maximum Rayleigh quotient sampling techniques. Meanwhile, UniFORM (Song et al., 2025) deploys energy-based encoding mechanisms to capture anomaly-related uncertainty across multiple spatial scales. Beyond single-domain settings, GFM-UAD (Han et al., 2026) couples adversarial learning with fine-tuning strategies, bridging multi-granularity detection with cross-domain adaptability. More recently, GRAND (Wang et al., 2026) integrates dual diffusion mechanisms with Transformer-based denoising networks, offering a unified paradigm for identifying both node- and edge-level irregular patterns. *Although they achieve promising progress, they ignore hierarchical information in the latent representations and fail to explore the utility of cross-granularity representation for granularity-specific detection.*

3 Preliminaries

Notations and Statement. A graph is defined as $G = (X, A)$, where $X \in \mathbb{R}^{N \times d'}$ is the feature matrix, and $A \in \mathbb{R}^{N \times N}$ is the adjacency matrix. N is the number of nodes and d' is the dimension of node features. Denote the node-, edge-, and graph-granularity tasks as $\mathcal{N}$, $\mathcal{E}$, and $\mathcal{G}$. Given a training dataset $\mathcal{T}_{train}$ for a task $Y_{train} \subseteq \{\mathcal{N}, \mathcal{E}, \mathcal{G}\}$, the goal is to train a unified model using $\mathcal{T}_{train}$ to identify anomalies in a test set $\mathcal{T}_{test}$ for $Y_{test} \subseteq \{\mathcal{N}, \mathcal{E}, \mathcal{G}\}$.

Information Bottleneck. Stemming from information theory, the information bottleneck principle (IB) (Wei et al., 2026b) employs a bottleneck mechanism to regulate the flow of information. Its objective is to learn a compact feature representation that maximally compresses the original input while retaining the essential information for the target task. Therefore, the formulation of the IB theory can be defined as:

$$\mathcal{L}_{in} = I(T; X) - \beta I(T; Y), \tag{1}$$

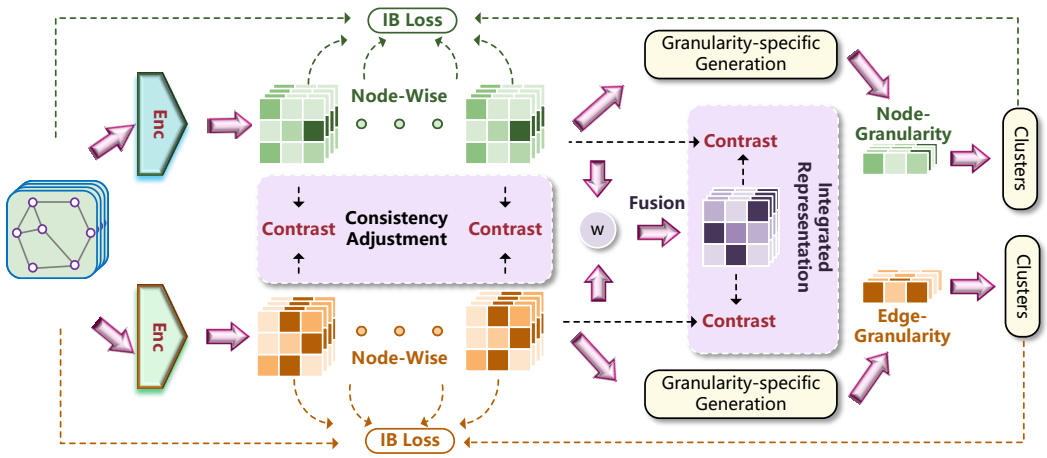


Figure 1: The framework of GaHIGAD. We take single-graph datasets as an example. With single-graph datasets, we implement node- and edge-level anomaly detection tasks. However, with multi-graph datasets, we implement node- and graph-level anomaly detection tasks, which can be easily extended from the framework with single-graph datasets.

where T is the compressed representation, X is source variable, Y is label, and β is a hyperparameter to weight compression and retention. A smaller value of $I(T;X)$ indicates a better compression ability. A larger value of $I(T;Y)$ represents a stronger retained correlation.

4 Contrastive Multi-Granularity

The framework is shown in Fig. 1. GaHIGAD utilizes individual encoders to encode the hierarchical granularity-specific representations and generates integrated representations for cross-granularity utility exploration. GaHIGAD develops an IB-based loss function to explore alignment of hierarchical information and mines the consistency information about cross-granularity representations. In addition, GaHIGAD leverages integrated representations to achieve further consistency alignment with different granularities.

4.1 Detailed Architecture

Projection. The projection module is used to individually encode the latent representations for each granularity to obtain the node-wise representations. The projection module is built with a learnable MLP. Formally, this process is formulated as:

$$H_g = f_{\theta_g}(X), \tag{2}$$

where $H_g \in \mathbb{R}^{N\times d}$ denotes the encoded representations of granularity,g and f_{θ_g} denotes the granularity-specific MLP, and d is the latent dimension.

Encoder. Encoders aims to learn the latent representations of nodes with adequate information aggregation, which makes each node fully absorb the neighboring information. To explore the evolutionary characteristics of latent representations, the encoder has L layers ($L > 2$), and each layer is established by a GNN. Formally, this process can be formulated as follows:

$$\hat{H}_g^{l+1} = \mathrm{GNN}(\hat{H}_g^l, A), \tag{3}$$

where $\hat{H}_g^0 = H_g$ represents the input of the encoder.

Granularity-aware Fusion. In order to explore the inter-dependencies across granularities on the specific GAD problem, we design a granularity-aware fusion module that is used to achieve cross-granularity representation integration. The integrated representations and granularity-specific representations are intended to implement the concrete GAD task simultaneously, supplementing the

semantic modeling towards cross-granularity characteristics. Formally, we define the target granularity as g_t, the supplementary granularity as g_s, the corresponding node-wise representations as $H_{g_t} \in \mathbb{R}^{N\times d}$ and $H_{g_s} \in \mathbb{R}^{N\times d}$, then we generate the integrated representations $H_{int} \in \mathbb{R}^{N\times d}$:

$$\omega = \text{Sigmoid}([H_{g_t}, H_{g_s}]W), \quad H_{int} = \omega H_{g_t} + (1-\omega)H_{g_s}, \tag{4}$$

where $\text{Sigmoid}(\cdot)$ is Sigmoid activation function, $[\cdot,\cdot]$ is concatenation function, and $W \in \mathbb{R}^{2d\times d}$ is learnable parameters.

Granularity-specific Generation. This module intends to generate multi-granularity representations based on the node-wise representations. We directly utilize node-wise representations to represent node granularity, leverage the difference between two nodes as edge-granularity representations, and employ mean pooling on nodes to generate graph-granularity representations. Specifically, the edge- and graph-granularity representations can be obtained as follows:

$$z^{graph} = \text{Mean}(H^L), z_{e_{ij}} = h_i^L - h_j^L. \tag{5}$$

Through this module, we can obtain granularity-specific representations, denoted as $Z^{node} \in \mathbb{R}^{N_n\times d}$, $Z^{edge} \in \mathbb{R}^{N_e\times d}$, $Z^{graph} \in \mathbb{R}^{N_g\times d}$, where N_n, N_e, N_g represent the number of samples for different granularities and d denotes the hidden dimension. Note that $N_n = N$.

Granularity-specific Classifier. The classifier aims to predict a vector $y \in \mathbb{R}^2$ for each sample to reflect the probability that the sample belongs to each category. This process can be formulated as:

$$Y^{node} = f_c(Z^{node}), Y^{edge} = f_c(Z^{edge}), Y^{graph} = f_c(Z^{graph}) \tag{6}$$

where $Y^{node} \in \mathbb{R}^{N_n\times 2}$, $Y^{edge} \in \mathbb{R}^{N_e\times 2}$, $Y^{graph} \in \mathbb{R}^{N_g\times 2}$ denote the classification results for different granularities.

4.2 Objective Functions

Information Bottleneck. IB compresses the original input data into denser representations, obtaining more discriminative ability for the rare patterns for anomaly detection. Introducing intermediate information $H_1, H_2, \cdots, H_{L-1}$ can better regularize the evolution process of compression. Given this idea, we integrate the intermediate information with IB loss, which can be formulated as:

$$\mathcal{L}_{IB} = \sum_{i=1}^{N}\sum_{l=1}^{L} I(x_i; h_i^l) - \beta I(y_i; h_i^l), \tag{7}$$

where $I(X; H_l)$ denotes the mutual information between input data and hierarchical latent representations, $I(Y; H_l)$ represents the mutual information between hierarchical latent representations and the classification results.

Consistency Adjustment. Since the label information is not available in this study, the mutual information in the IB loss lacks imperative supervised signals to achieve convergence. Therefore, we intend to provide more unsupervised consistency regularization to model normal patterns of samples. On the other hand, for multi-granularity GAD, it is important to explore the inter-granularity consistency information to accomplish accurate classification. Although different granularities may exhibit distinctive patterns, node-wise representations should maintain consistency because they represent the same entities. On the other hand, inter-granularity adjustment can model comprehensive consistency information, while ignoring the cross-granularity correlation in the granularity-specific GAD task. To address this problem, we explore the consistency between samples in the integrated representation H_{int}, promoting the utility of cross-granularity representations in downstream detection tasks. Therefore, the consistency adjustment loss is defined as:

$$\mathcal{L}_{CA} = \gamma \sum_{l=1}^{L} I(H_{g_s}^l; H_{g_t}^l) + (1-\gamma) \sum_{g\in\{g_s,g_t\}} I(H_{int}; H_g), \tag{8}$$

where γ is a hyperparameter to control the balance of two components, $I(H_{g_s}^l; H_{g_t}^l)$ represents inter-granularity consistency, the mutual information between the representations from two different granularities, to achieve node-wise semantic regularization, and $I(H_{int}; H_g)$ evaluates the cross-granularity utility on a granularity-specific detection problem.

Total Objective Function. Consequently, the proposed method can purify the original features while achieving consistency between different granularities and the correlation between compressed representations and sample classification. Besides, unlike traditional deep IB methods, it further explores the utility of hidden information. Therefore, the total objective function can be formulated as:

$$\mathcal{L} = \alpha \mathcal{L}_{IB} + (1 - \alpha)\mathcal{L}_{CA}, \tag{9}$$

where α is a hyperparameter that is used to balance the contribution between IB and consistency information.

The hidden information is obtained by feeding the original features into encoders, exhibiting a trend of hierarchical compression. $H^1, H^2, \cdots, H^{L-1}$ are the intermediate hidden representations, indicating the evolutionary results of compression. Introducing hidden information into the final objective has the following benefits:

- **Advanced feature expressiveness**: As network depth grows, successive hidden layers progressively refine features, generating representations that are semantically richer and better organized.
- **Improved interpretability**: Hidden-layer representations expose how inputs are transformed step by step, shedding light on the learned hierarchy and thereby aiding model analysis and fine-tuning.
- **Optimization gains from the information bottleneck**: Incorporating hidden-layer information into the information bottleneck objective helps balance the compression–relevance trade-off, yielding clustering outcomes with stronger discrimination and generalization.

4.3 Optimization

In the latent representation space with high dimensionality, directly computing the mutual information becomes laborious and impractical. To achieve efficient calculation of mutual information and derive discriminative representations of samples, we deploy three different strategies to estimate the mutual information, including InfoNCE and MINE.

InfoNCE. It is used to calculate mutual information items in $\mathcal{L}_{CA}$, since the representations have the same dimensionality. The formulation is defined based on contrastive learning:

$$I(H_1; H_2) \approx \mathbb{E}_{h_1,h_2 \sim p(H_1;H_2)}[\log \frac{\exp(\cos(h_i, h_j))}{\sum_{n=1}^{N} \exp(\cos(h_i, h_j))}] + \log N_I \tag{10}$$

MINE. Since the dimension of h_i is higher than those of x_i and y_i, directly using InfoNCE to optimizing $\mathcal{L}_{IB}$ may introduce information loss and indiscrimination of contrastive pairs. In contrast, MINE provides a more robust method to evaluate the mutual information:

$$I(x, y) = -\mathbb{E}_{p(x,y)}[T(x, y)] + \log \mathbb{E}_{p(x)p(y)}[\exp(T(x, y))], \tag{11}$$

where $T(x, y)$ is the function that needs to be approximated by neural networks.

Table 1: The performance of node- and edge-level GAD on single-graph datasets

	Dataset	Reddit		Weibo		Amazon		Yelp		Tolokers		Questions		T-Finance	
	Metrics	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
Node-Level	BOURNE	$65.2_{\pm 0.7}$	$8.5_{\pm 0.3}$	$78.5_{\pm 0.9}$	$25.3_{\pm 0.7}$	$82.1_{\pm 1.0}$	$32.5_{\pm 0.9}$	$68.9_{\pm 0.8}$	$22.1_{\pm 0.7}$	$65.3_{\pm 0.8}$	$25.6_{\pm 0.8}$	$60.2_{\pm 0.7}$	$15.3_{\pm 0.6}$	$58.9_{\pm 0.9}$	$10.5_{\pm 0.5}$
	AutoGAD	$68.4_{\pm 0.9}$	$10.2_{\pm 0.6}$	$78.3_{\pm 1.2}$	$25.1_{\pm 1.0}$	$85.6_{\pm 1.3}$	$38.7_{\pm 1.2}$	$70.4_{\pm 1.1}$	$27.5_{\pm 1.0}$	$68.7_{\pm 1.0}$	$30.2_{\pm 1.0}$	$67.5_{\pm 1.0}$	$21.8_{\pm 0.9}$	$53.4_{\pm 1.1}$	$6.2_{\pm 0.2}$
	UniFORM	$70.1_{\pm 1.1}$	$11.5_{\pm 0.8}$	$80.7_{\pm 0.7}$	$28.7_{\pm 0.8}$	$84.3_{\pm 0.9}$	$36.2_{\pm 1.0}$	$70.8_{\pm 0.9}$	$26.3_{\pm 0.8}$	$67.2_{\pm 0.9}$	$28.4_{\pm 0.9}$	$62.8_{\pm 0.9}$	$18.2_{\pm 0.8}$	$60.8_{\pm 1.0}$	$12.1_{\pm 0.6}$
	GFM-UAD	$70.3_{\pm 1.4}$	$11.8_{\pm 1.0}$	$82.1_{\pm 1.4}$	$30.4_{\pm 1.3}$	$85.9_{\pm 1.4}$	$37.1_{\pm 1.4}$	$75.2_{\pm 1.3}$	$32.7_{\pm 1.3}$	$66.8_{\pm 1.3}$	$28.1_{\pm 1.4}$	$67.9_{\pm 1.3}$	$22.6_{\pm 1.2}$	$65.7_{\pm 1.4}$	$15.8_{\pm 1.2}$
	GRAND	$69.5_{\pm 0.8}$	$9.7_{\pm 0.5}$	$83.6_{\pm 1.1}$	$32.8_{\pm 1.1}$	$87.2_{\pm 1.2}$	$40.3_{\pm 1.1}$	$73.6_{\pm 1.0}$	$30.4_{\pm 1.1}$	$70.1_{\pm 1.2}$	$32.7_{\pm 1.1}$	$66.3_{\pm 1.1}$	$21.1_{\pm 1.0}$	$64.1_{\pm 1.2}$	$14.4_{\pm 1.0}$
	SGAD-DCR	$55.6_{\pm 0.6}$	$4.2_{\pm 0.1}$	$72.4_{\pm 0.6}$	$18.6_{\pm 0.5}$	$75.3_{\pm 0.7}$	$25.4_{\pm 0.5}$	$62.3_{\pm 0.6}$	$15.8_{\pm 0.5}$	$58.4_{\pm 0.6}$	$18.9_{\pm 0.5}$	$53.7_{\pm 0.5}$	$11.2_{\pm 0.5}$	$52.3_{\pm 0.7}$	$7.6_{\pm 0.1}$
	MCLASt	$53.2_{\pm 1.0}$	$3.1_{\pm 0.2}$	$74.8_{\pm 0.8}$	$20.4_{\pm 0.6}$	$77.8_{\pm 0.8}$	$28.1_{\pm 0.6}$	$64.7_{\pm 0.7}$	$18.2_{\pm 0.6}$	$60.9_{\pm 0.7}$	$21.3_{\pm 0.6}$	$56.1_{\pm 0.6}$	$13.5_{\pm 0.7}$	$54.8_{\pm 0.8}$	$9.1_{\pm 0.2}$
	MAC-GAD	$60.1_{\pm 1.2}$	$6.3_{\pm 0.4}$	$71.2_{\pm 1.3}$	$16.1_{\pm 0.9}$	$79.5_{\pm 1.1}$	$30.2_{\pm 0.7}$	$61.1_{\pm 1.2}$	$16.5_{\pm 1.2}$	$63.5_{\pm 1.4}$	$20.5_{\pm 0.7}$	$58.4_{\pm 1.2}$	$14.9_{\pm 1.1}$	$59.2_{\pm 1.3}$	$12.2_{\pm 0.4}$
	GaHIGAD	$\mathbf{73.6}_{\pm 1.3}$	$\mathbf{15.9}_{\pm 1.1}$	$\mathbf{90.7}_{\pm 1.0}$	$\mathbf{44.6}_{\pm 1.4}$	$\mathbf{94.8}_{\pm 0.6}$	$\mathbf{51.5}_{\pm 0.8}$	$\mathbf{80.8}_{\pm 0.5}$	$\mathbf{38.6}_{\pm 0.9}$	$\mathbf{76.5}_{\pm 1.1}$	$\mathbf{40.7}_{\pm 1.2}$	$\mathbf{73.1}_{\pm 0.8}$	$\mathbf{29.3}_{\pm 0.5}$	$\mathbf{71.3}_{\pm 0.6}$	$\mathbf{18.6}_{\pm 0.7}$
Edge-level	BOURNE	$62.5_{\pm 0.9}$	$6.8_{\pm 0.3}$	$80.2_{\pm 1.0}$	$28.4_{\pm 0.8}$	$78.9_{\pm 0.9}$	$30.5_{\pm 0.8}$	$58.6_{\pm 0.8}$	$4.2_{\pm 0.2}$	$62.4_{\pm 0.8}$	$20.1_{\pm 0.7}$	$66.3_{\pm 0.8}$	$25.4_{\pm 0.7}$	$65.8_{\pm 0.9}$	$20.3_{\pm 0.7}$
	AutoGAD	$61.8_{\pm 1.2}$	$5.2_{\pm 0.4}$	$83.5_{\pm 1.3}$	$32.1_{\pm 1.0}$	$82.4_{\pm 1.2}$	$35.7_{\pm 1.1}$	$62.1_{\pm 1.1}$	$5.5_{\pm 0.4}$	$61.7_{\pm 1.1}$	$18.8_{\pm 1.0}$	$69.2_{\pm 1.1}$	$28.7_{\pm 1.0}$	$69.4_{\pm 1.2}$	$24.6_{\pm 1.0}$
	UniFORM	$68.3_{\pm 1.1}$	$9.5_{\pm 0.5}$	$81.7_{\pm 0.9}$	$30.8_{\pm 0.9}$	$80.6_{\pm 1.0}$	$33.2_{\pm 0.9}$	$60.4_{\pm 1.0}$	$4.7_{\pm 0.7}$	$64.2_{\pm 0.9}$	$22.4_{\pm 0.8}$	$67.8_{\pm 1.0}$	$27.1_{\pm 0.8}$	$67.2_{\pm 1.1}$	$22.8_{\pm 0.8}$
	GFM-UAD	$70.1_{\pm 1.4}$	$10.8_{\pm 0.7}$	$85.3_{\pm 1.4}$	$35.6_{\pm 1.4}$	$84.1_{\pm 1.4}$	$38.9_{\pm 1.3}$	$64.7_{\pm 1.3}$	$8.3_{\pm 0.5}$	$67.9_{\pm 1.3}$	$26.7_{\pm 1.3}$	$71.5_{\pm 1.4}$	$30.2_{\pm 1.3}$	$71.8_{\pm 1.4}$	$26.4_{\pm 1.3}$
	GRAND	$66.9_{\pm 1.0}$	$8.9_{\pm 0.2}$	$82.8_{\pm 1.1}$	$31.5_{\pm 1.1}$	$81.5_{\pm 1.1}$	$34.1_{\pm 1.0}$	$61.2_{\pm 1.2}$	$6.6_{\pm 0.4}$	$65.1_{\pm 1.0}$	$23.2_{\pm 0.9}$	$68.4_{\pm 1.2}$	$27.9_{\pm 0.9}$	$68.3_{\pm 1.3}$	$23.5_{\pm 0.9}$
	SGAD-DCR	$53.2_{\pm 0.6}$	$4.1_{\pm 0.1}$	$76.5_{\pm 0.6}$	$22.7_{\pm 0.5}$	$74.8_{\pm 0.7}$	$25.1_{\pm 0.5}$	$55.3_{\pm 0.6}$	$2.2_{\pm 0.5}$	$63.7_{\pm 0.6}$	$21.8_{\pm 0.5}$	$62.4_{\pm 0.6}$	$21.5_{\pm 0.5}$	$60.1_{\pm 0.7}$	$17.4_{\pm 0.6}$
	MCLASt	$56.7_{\pm 0.7}$	$5.3_{\pm 0.0}$	$76.9_{\pm 0.8}$	$23.4_{\pm 0.6}$	$77.2_{\pm 0.8}$	$27.8_{\pm 0.7}$	$55.8_{\pm 0.7}$	$3.5_{\pm 0.9}$	$60.3_{\pm 1.2}$	$19.2_{\pm 0.6}$	$64.1_{\pm 0.7}$	$23.7_{\pm 0.6}$	$62.5_{\pm 0.8}$	$19.8_{\pm 1.1}$
	MAC-GAD	$59.9_{\pm 0.8}$	$6.5_{\pm 0.6}$	$80.1_{\pm 1.2}$	$27.3_{\pm 0.7}$	$79.4_{\pm 1.3}$	$30.2_{\pm 1.2}$	$51.6_{\pm 1.2}$	$1.9_{\pm 0.3}$	$62.1_{\pm 1.4}$	$21.4_{\pm 1.2}$	$65.2_{\pm 1.3}$	$25.8_{\pm 1.2}$	$63.3_{\pm 0.6}$	$21.6_{\pm 1.2}$
	GaHIGAD	$\mathbf{72.3}_{\pm 0.5}$	$\mathbf{15.1}_{\pm 1.1}$	$\mathbf{92.3}_{\pm 0.7}$	$\mathbf{47.2}_{\pm 1.2}$	$\mathbf{90.5}_{\pm 0.6}$	$\mathbf{51.8}_{\pm 0.6}$	$\mathbf{68.6}_{\pm 0.4}$	$\mathbf{13.5}_{\pm 0.8}$	$\mathbf{71.5}_{\pm 0.6}$	$\mathbf{28.9}_{\pm 1.1}$	$\mathbf{77.1}_{\pm 0.9}$	$\mathbf{36.4}_{\pm 1.1}$	$\mathbf{76.3}_{\pm 1.0}$	$\mathbf{29.6}_{\pm 0.5}$

Table 2: The performance of node- and graph-level GAD on multi-graph datasets.

	Dataset	BM-MN		BM-MS		BM-MT		MUTAG		MNIST0		MNIST1		T-Group	
	Metrics	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
Node-level	BOURNE	$78.2_{\pm 0.8}$	$32.5_{\pm 0.6}$	$76.4_{\pm 1.1}$	$38.7_{\pm 0.9}$	$72.3_{\pm 0.7}$	$28.5_{\pm 0.5}$	$82.1_{\pm 0.9}$	$40.2_{\pm 0.7}$	$88.5_{\pm 1.2}$	$72.1_{\pm 0.9}$	$85.6_{\pm 0.6}$	$62.4_{\pm 1.1}$	$80.2_{\pm 1.0}$	$42.3_{\pm 0.8}$
	AutoGAD	$75.6_{\pm 1.0}$	$28.9_{\pm 0.8}$	$72.8_{\pm 0.9}$	$35.2_{\pm 1.2}$	$69.4_{\pm 1.0}$	$25.1_{\pm 0.7}$	$79.8_{\pm 1.1}$	$37.6_{\pm 0.9}$	$86.2_{\pm 0.8}$	$68.9_{\pm 1.1}$	$83.1_{\pm 1.0}$	$58.7_{\pm 0.7}$	$77.5_{\pm 0.9}$	$39.1_{\pm 1.1}$
	UniFORM	$72.3_{\pm 1.2}$	$25.4_{\pm 0.5}$	$70.1_{\pm 1.3}$	$31.8_{\pm 0.7}$	$66.7_{\pm 0.8}$	$22.6_{\pm 0.9}$	$77.5_{\pm 0.7}$	$34.1_{\pm 1.2}$	$83.7_{\pm 1.1}$	$65.3_{\pm 0.8}$	$80.4_{\pm 0.7}$	$54.2_{\pm 1.0}$	$74.8_{\pm 1.2}$	$36.5_{\pm 0.9}$
	GFM-UAD	$70.1_{\pm 0.6}$	$22.1_{\pm 1.1}$	$67.5_{\pm 0.8}$	$29.4_{\pm 1.0}$	$63.9_{\pm 1.1}$	$20.8_{\pm 0.6}$	$74.9_{\pm 1.0}$	$31.5_{\pm 0.8}$	$81.2_{\pm 0.9}$	$62.7_{\pm 1.2}$	$78.0_{\pm 1.1}$	$50.8_{\pm 0.6}$	$72.1_{\pm 0.7}$	$33.9_{\pm 1.2}$
	GRAND	$80.4_{\pm 0.9}$	$35.2_{\pm 1.0}$	$74.6_{\pm 0.7}$	$36.9_{\pm 0.8}$	$70.8_{\pm 1.2}$	$26.4_{\pm 1.1}$	$81.3_{\pm 0.8}$	$38.9_{\pm 1.0}$	$87.1_{\pm 0.7}$	$70.5_{\pm 0.6}$	$84.2_{\pm 0.9}$	$60.1_{\pm 1.2}$	$79.0_{\pm 1.1}$	$40.8_{\pm 0.7}$
	SGAD-DCR	$62.5_{\pm 1.1}$	$15.6_{\pm 0.4}$	$58.3_{\pm 0.9}$	$20.7_{\pm 0.5}$	$54.2_{\pm 0.7}$	$12.8_{\pm 0.3}$	$65.8_{\pm 1.2}$	$22.4_{\pm 0.6}$	$70.3_{\pm 0.8}$	$48.2_{\pm 0.5}$	$68.9_{\pm 1.0}$	$38.5_{\pm 0.4}$	$63.7_{\pm 0.9}$	$25.1_{\pm 0.5}$
	MCLASt	$65.8_{\pm 0.7}$	$18.2_{\pm 0.6}$	$61.2_{\pm 1.0}$	$23.5_{\pm 0.8}$	$57.6_{\pm 0.9}$	$15.3_{\pm 0.5}$	$68.4_{\pm 0.8}$	$25.7_{\pm 1.1}$	$73.1_{\pm 1.2}$	$52.8_{\pm 0.7}$	$71.5_{\pm 0.6}$	$42.3_{\pm 0.9}$	$66.4_{\pm 1.1}$	$28.4_{\pm 0.7}$
	MAC-GAD	$68.9_{\pm 1.3}$	$20.3_{\pm 0.7}$	$64.1_{\pm 0.6}$	$26.9_{\pm 1.1}$	$60.5_{\pm 1.0}$	$18.2_{\pm 0.8}$	$71.2_{\pm 1.1}$	$28.6_{\pm 0.7}$	$76.4_{\pm 0.6}$	$56.7_{\pm 1.0}$	$74.8_{\pm 1.2}$	$46.1_{\pm 0.8}$	$69.3_{\pm 0.8}$	$31.6_{\pm 1.0}$
	GaHIGAD	$\mathbf{83.5}_{\pm 0.6}$	$\mathbf{37.8}_{\pm 0.8}$	$\mathbf{89.9}_{\pm 1.2}$	$\mathbf{56.4}_{\pm 0.5}$	$\mathbf{85.5}_{\pm 0.6}$	$\mathbf{48.1}_{\pm 1.0}$	$\mathbf{86.7}_{\pm 1.1}$	$\mathbf{43.6}_{\pm 0.4}$	$\mathbf{96.2}_{\pm 0.5}$	$\mathbf{86.7}_{\pm 0.8}$	$\mathbf{95.4}_{\pm 1.3}$	$\mathbf{78.9}_{\pm 0.5}$	$\mathbf{88.4}_{\pm 0.9}$	$\mathbf{51.6}_{\pm 1.1}$
Graph-level	BOURNE	$76.4_{\pm 0.9}$	$38.5_{\pm 0.7}$	$82.1_{\pm 1.0}$	$40.2_{\pm 0.8}$	$78.6_{\pm 0.7}$	$38.9_{\pm 0.6}$	$89.2_{\pm 1.1}$	$52.4_{\pm 0.9}$	$88.7_{\pm 0.8}$	$44.1_{\pm 0.7}$	$86.3_{\pm 1.0}$	$39.5_{\pm 0.8}$	$76.9_{\pm 0.9}$	$33.2_{\pm 0.7}$
	AutoGAD	$74.1_{\pm 1.1}$	$35.8_{\pm 0.9}$	$79.5_{\pm 0.8}$	$37.6_{\pm 1.1}$	$75.8_{\pm 1.0}$	$35.4_{\pm 0.8}$	$86.7_{\pm 0.7}$	$49.8_{\pm 0.6}$	$86.2_{\pm 1.1}$	$41.5_{\pm 0.9}$	$83.9_{\pm 0.8}$	$36.7_{\pm 1.0}$	$74.2_{\pm 1.1}$	$30.5_{\pm 0.6}$
	UniFORM	$71.8_{\pm 0.7}$	$32.4_{\pm 1.0}$	$76.9_{\pm 1.2}$	$34.1_{\pm 0.7}$	$73.2_{\pm 0.8}$	$32.7_{\pm 1.0}$	$83.9_{\pm 1.0}$	$46.5_{\pm 0.8}$	$83.6_{\pm 0.7}$	$38.2_{\pm 1.2}$	$81.1_{\pm 1.1}$	$33.9_{\pm 0.7}$	$71.5_{\pm 0.8}$	$27.8_{\pm 1.0}$
	GFM-UAD	$69.2_{\pm 1.2}$	$29.7_{\pm 0.6}$	$74.3_{\pm 0.9}$	$31.5_{\pm 1.0}$	$70.6_{\pm 1.1}$	$29.8_{\pm 0.7}$	$81.2_{\pm 0.8}$	$43.7_{\pm 1.1}$	$80.9_{\pm 1.0}$	$35.6_{\pm 0.8}$	$78.4_{\pm 0.7}$	$30.2_{\pm 1.1}$	$68.9_{\pm 1.0}$	$24.6_{\pm 0.8}$
	GRAND	$78.9_{\pm 0.8}$	$41.2_{\pm 1.1}$	$84.6_{\pm 0.7}$	$43.8_{\pm 0.9}$	$80.3_{\pm 0.9}$	$41.5_{\pm 1.1}$	$91.5_{\pm 0.6}$	$55.1_{\pm 0.7}$	$90.8_{\pm 1.2}$	$46.7_{\pm 0.6}$	$88.6_{\pm 0.9}$	$41.8_{\pm 0.7}$	$78.7_{\pm 1.2}$	$35.4_{\pm 0.9}$
	GUL	$66.5_{\pm 0.7}$	$25.8_{\pm 0.5}$	$70.1_{\pm 0.8}$	$27.4_{\pm 0.6}$	$67.3_{\pm 0.9}$	$25.6_{\pm 0.5}$	$78.4_{\pm 1.1}$	$39.2_{\pm 0.9}$	$77.8_{\pm 0.6}$	$31.4_{\pm 0.7}$	$75.2_{\pm 1.0}$	$26.8_{\pm 0.6}$	$65.3_{\pm 0.8}$	$20.7_{\pm 0.5}$
	M-GLAD	$63.9_{\pm 1.0}$	$22.4_{\pm 0.8}$	$67.8_{\pm 1.1}$	$24.1_{\pm 0.9}$	$64.5_{\pm 0.6}$	$22.3_{\pm 0.8}$	$75.1_{\pm 0.7}$	$35.6_{\pm 1.0}$	$74.6_{\pm 1.1}$	$27.9_{\pm 0.9}$	$72.4_{\pm 0.8}$	$23.5_{\pm 1.0}$	$62.1_{\pm 1.2}$	$17.4_{\pm 0.7}$
	MoEGAD	$72.8_{\pm 0.6}$	$34.2_{\pm 0.9}$	$77.5_{\pm 1.0}$	$36.8_{\pm 0.7}$	$74.9_{\pm 0.8}$	$34.5_{\pm 1.2}$	$85.3_{\pm 0.9}$	$48.1_{\pm 1.1}$	$84.7_{\pm 0.7}$	$40.3_{\pm 0.8}$	$82.5_{\pm 1.2}$	$35.6_{\pm 0.9}$	$73.6_{\pm 0.7}$	$29.1_{\pm 1.1}$
	GaHIGAD	$\mathbf{80.6}_{\pm 0.5}$	$\mathbf{45.7}_{\pm 1.9}$	$\mathbf{86.3}_{\pm 0.7}$	$\mathbf{45.7}_{\pm 0.4}$	$\mathbf{81.9}_{\pm 1.0}$	$\mathbf{44.1}_{\pm 1.1}$	$\mathbf{93.8}_{\pm 0.9}$	$\mathbf{57.2}_{\pm 0.6}$	$\mathbf{92.6}_{\pm 0.5}$	$\mathbf{48.1}_{\pm 1.0}$	$\mathbf{90.7}_{\pm 0.7}$	$\mathbf{43.5}_{\pm 0.8}$	$\mathbf{80.3}_{\pm 1.1}$	$\mathbf{37.6}_{\pm 0.5}$

5 EXPERIMENTS

5.1 EXPERIMENTAL SETTINGS

Datasets. We leverage 14 graph datasets (Lin et al., 2024), including 7 single-graph datasets for node- and edge-level anomaly detection and 7 multi-graph datasets for node- and graph-level anomaly detection. The single-graph datasets contain Reddit, Weibo, Amazon, Yelp, Tolokers, Questions, and T-finance. The multi-graph datasets contain BM-MN, BM-MS, BM-MT, MUTAG, MNIST0, MNIST1, and T-Group. More details about the datasets are in Section A.

Baselines. To keep fairness, we only introduce unsupervised learning baselines. We introduce three groups of state-of-the-art baselines. (1) Multi-granularity GAD methods: BOURNE (Liu et al., 2024), AutoGAD (Yu et al., 2025), UniFORM (Song et al., 2025), GFM-UAD (Han et al., 2026), and GRAND (Wang et al., 2026). (2) Traditional baselines: SGAD-DCR (Chen et al., 2026b), MCLASt (Zhang & Luo, 2026), and MAC-GAD (Zhang et al., 2027). (3) GUL (Yang et al., 2026), M-GLAD (Zhang et al., 2026a), and MoEGAD (Cai et al., 2026). For node- and edge-level detection on single-graph datasets, we leverage the first and second groups as baselines. For node- and graph-level detection on multi-graph datasets, we second and third groups as baselines. More details are in Section B.

Table 3: The ablation study results of IB loss on node- and edge-level GAD on single-graph datasets with respect to AUROC.

Method	Reddit		Weibo		Amazon		Yelp		Tolokers		Questions		T-finance	
	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge
Rec	$62.5_{\pm 1.2}$	$60.8_{\pm 0.9}$	$78.4_{\pm 1.4}$	$80.1_{\pm 0.7}$	$82.3_{\pm 1.1}$	$78.2_{\pm 1.3}$	$68.5_{\pm 0.8}$	$58.6_{\pm 1.0}$	$64.8_{\pm 1.5}$	$60.3_{\pm 0.6}$	$61.5_{\pm 1.2}$	$65.4_{\pm 0.9}$	$59.7_{\pm 1.1}$	$64.2_{\pm 0.8}$
Con	$58.2_{\pm 0.8}$	$56.4_{\pm 1.1}$	$72.6_{\pm 0.6}$	$75.9_{\pm 1.3}$	$76.5_{\pm 0.9}$	$72.4_{\pm 1.2}$	$62.7_{\pm 1.4}$	$56.2_{\pm 0.7}$	$60.1_{\pm 1.0}$	$56.8_{\pm 1.5}$	$57.2_{\pm 0.8}$	$60.8_{\pm 1.1}$	$55.4_{\pm 1.3}$	$59.6_{\pm 0.9}$
IB (Proposed)	$73.6_{\pm 1.3}$	$72.3_{\pm 0.5}$	$90.7_{\pm 1.0}$	$92.3_{\pm 0.7}$	$94.8_{\pm 0.6}$	$90.5_{\pm 0.6}$	$80.8_{\pm 0.5}$	$68.6_{\pm 0.4}$	$76.5_{\pm 1.1}$	$71.5_{\pm 0.6}$	$73.1_{\pm 0.8}$	$77.1_{\pm 0.9}$	$71.3_{\pm 0.6}$	$76.3_{\pm 1.0}$

Implementations. The hidden dimension d' is set to 64. For unified experiments, α, β, and γ are respectively set to 0.5, 0.001, and 0.5. The training epoch is set to 200. The learning rate is 1e-4. All methods are executed in a Python 3.9.12 environment, utilizing a single NVIDIA A40 GPU, 40GB of RAM, and a 2.60GHz Xeon (R) Gold 6240 CPU. We utilize two metrics to evaluate the performance: AUROC and AUPRC. Each experiment is run 5 times with 5 different random seeds, and the average performance is recorded with standard deviation. More details are in Section C.

5.2 EXPERIMENTAL ANALYSIS

Performance. To examine how well GaHIGAD performs, we evaluate it on anomaly detection at multiple granularities under two scenarios compared with state-of-the-art baselines. In the first scenario, detection is carried out at both the node and edge levels on single-graph benchmarks. The results reported in Table 1 show that GaHIGAD attains the highest AUROC and AUPRC among all competitors, which cover both traditional GAD methods and existing multi-granularity GAD frameworks. This suggests that GaHIGAD extracts more granularity-specific semantics and distinguishes anomalies more reliably across different levels. Specifically, for node-level detection, GaHIGAD achieves relative AUROC improvements of no less than 4.7%, 8.5%, 8.7%, 7.4%, 9.1%, 7.7%, and

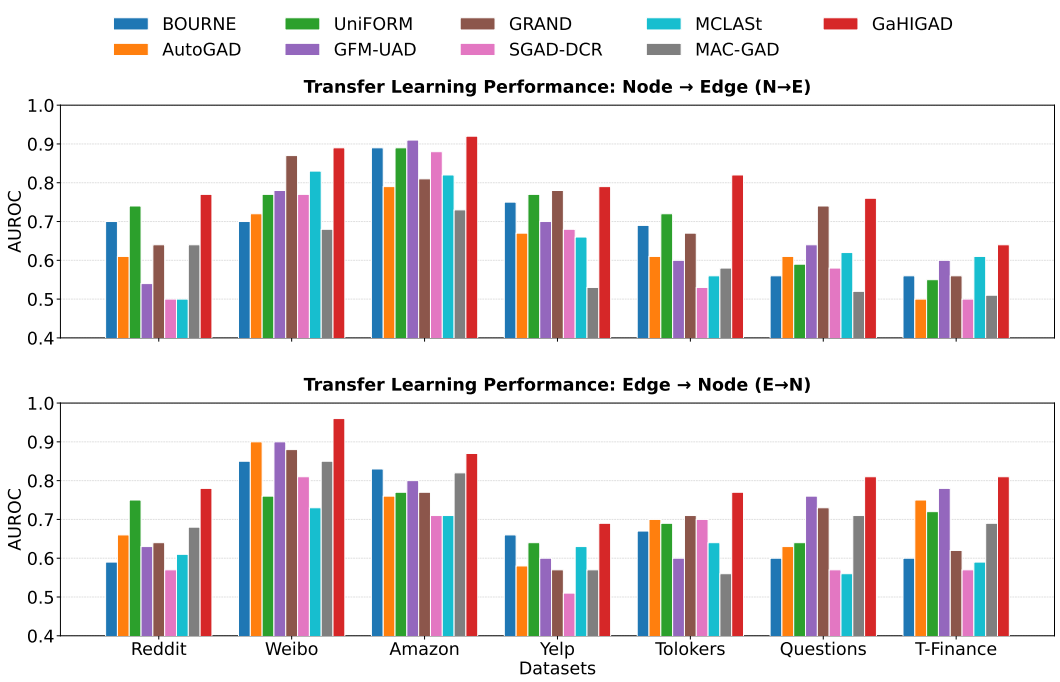


Figure 2: The zero-shot transfer detection experiments on single-graph datasets.

8.5% over the baselines on Reddit, Weibo, Amazon, Yelp, Tolokers, Questions, and T-Finance, respectively. In the second scenario, we perform node- and graph-level detection on multi-graph benchmarks. As presented in Table 2, GaHIGAD maintains its superiority over all baselines across the datasets. On the graph-level task, it achieves an average relative AUPRC improvement of 5.52% over the best-performing competitor, which shows that GaHIGAD is also effective for graph-level anomaly detection.

Table 4: The ablation study results of IB loss on node- and edge-level GAD on single-graph datasets with respect to AUROC.

Method	BM-MN		BM-MS		BM-MT		MUTAG		MNIST0		MNIST1		T-Group	
	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge
Rec	$68.9_{\pm 1.0}$	$65.3_{\pm 0.6}$	$74.2_{\pm 1.3}$	$76.1_{\pm 0.7}$	$74.8_{\pm 0.9}$	$71.3_{\pm 1.2}$	$76.4_{\pm 0.6}$	$83.2_{\pm 1.0}$	$86.1_{\pm 0.5}$	$82.4_{\pm 1.3}$	$85.3_{\pm 1.5}$	$80.2_{\pm 0.9}$	$78.1_{\pm 1.4}$	$69.7_{\pm 0.8}$
Con	$73.2_{\pm 1.1}$	$70.4_{\pm 0.8}$	$79.5_{\pm 1.4}$	$71.5_{\pm 0.9}$	$70.1_{\pm 1.5}$	$66.8_{\pm 0.7}$	$71.4_{\pm 1.1}$	$78.3_{\pm 0.8}$	$81.2_{\pm 0.5}$	$77.6_{\pm 1.2}$	$80.5_{\pm 0.9}$	$75.4_{\pm 1.4}$	$73.1_{\pm 0.7}$	$64.9_{\pm 1.0}$
IB (Proposed)	$83.5_{\pm 0.6}$	$80.6_{\pm 0.5}$	$89.9_{\pm 1.2}$	$86.3_{\pm 0.7}$	$85.5_{\pm 0.6}$	$81.9_{\pm 1.0}$	$86.7_{\pm 1.1}$	$93.8_{\pm 0.9}$	$96.2_{\pm 0.5}$	$92.6_{\pm 0.5}$	$95.4_{\pm 1.3}$	$90.7_{\pm 0.7}$	$88.4_{\pm 0.9}$	$80.3_{\pm 1.1}$

Zero-shot Transfer Detection. Fig. 2 shows the zero-shot detection performance. We can see that although the performance on the $N \rightarrow E$ task exhibits fluctuations compared with the performance on the $E \rightarrow N$ task, which is because the granularity-specific characteristics inevitably have discrepancies, the zero-shot transfer detection exhibits performance fluctuations compared with granularity-specific detection in Table 1 and Table 2. GaHIGAD consistently outperforms all baselines on single-graph datasets with respect to zero-shot transfer learning tasks. This indicates that GaHIGAD can capture some cross-granularity semantic information, which enhances zero-shot detection performance. For baselines, multi-granularity GAD baselines comprehensively achieve better performance than traditional GAD baselines because they have powerful scalability beyond node granularity. More zero-shot experiments on multi-graph datasets are shown in Section D.

Utility of IB loss. We explore the utility of IB loss. We introduce reconstruction (Rec) and traditional contrastive learning (Con) as variants. For reconstructions, we leverage the latent representations to achieve reconstruction of original features. For contrastive learning, we regard different granularities as contrastive views. Table 3 shows the ablation performance on single-graph datasets with respect to node- and edge-level detection. Table 4 shows the ablation performance on multi-graph datasets with respect to node- and graph-level detection. From these observations, we can

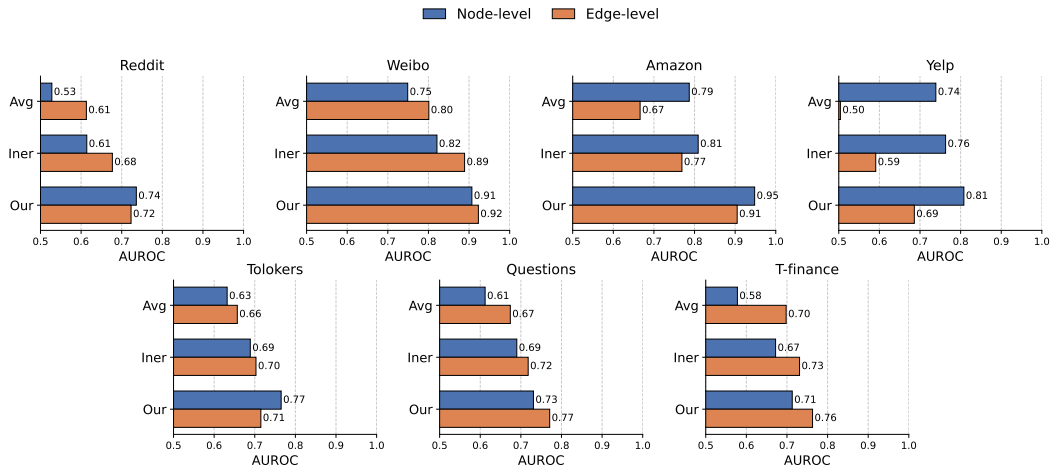


Figure 3: The ablation of fusion strategies on single-graph datasets.

see that using reconstruction and contrastive training paradigms to replace IB loss brings significant performance degradation, which is probably because they fail to filter noise information and capture discriminative features of granularity-specific representations.

Table 5: The ablation study results of consistency learning on node- and edge-level GAD on single-graph datasets with respect to AUROC.

inter	cross	Reddit		Weibo		Amazon		Yelp		Tolokers		Questions		T-finance	
		Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge
✓		$63.2_{\pm 1.1}$	$61.4_{\pm 0.7}$	$78.5_{\pm 1.4}$	$80.6_{\pm 0.9}$	$82.3_{\pm 0.8}$	$78.4_{\pm 1.2}$	$70.1_{\pm 1.3}$	$58.2_{\pm 0.6}$	$65.8_{\pm 0.9}$	$61.2_{\pm 1.0}$	$62.7_{\pm 1.2}$	$66.4_{\pm 0.5}$	$60.9_{\pm 0.8}$	$65.7_{\pm 1.1}$
	✓	$58.7_{\pm 0.9}$	$57.8_{\pm 1.2}$	$72.4_{\pm 1.1}$	$74.1_{\pm 0.6}$	$76.8_{\pm 1.3}$	$72.6_{\pm 0.7}$	$64.5_{\pm 0.8}$	$56.4_{\pm 1.0}$	$60.3_{\pm 1.2}$	$57.9_{\pm 0.9}$	$57.5_{\pm 0.6}$	$61.8_{\pm 1.1}$	$56.2_{\pm 1.0}$	$60.4_{\pm 0.8}$
✓	✓	$73.6_{\pm 1.3}$	$72.3_{\pm 0.5}$	$90.7_{\pm 1.0}$	$92.3_{\pm 0.7}$	$94.8_{\pm 0.6}$	$90.5_{\pm 0.6}$	$80.8_{\pm 0.5}$	$68.6_{\pm 0.4}$	$76.5_{\pm 1.1}$	$71.5_{\pm 0.6}$	$73.1_{\pm 0.8}$	$77.1_{\pm 0.9}$	$71.3_{\pm 0.6}$	$76.3_{\pm 1.0}$

Table 6: The ablation study results of consistency learning on node- and graph-level GAD on single-graph datasets with respect to AUROC.

inter	cross	BM-MN		BM-MS		BM-MT		MUTAG		MNIST0		MNIST1		T-Group	
		Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge	Node	Edge
✓		$68.7_{\pm 0.8}$	$62.4_{\pm 0.7}$	$69.8_{\pm 1.1}$	$72.1_{\pm 0.9}$	$73.2_{\pm 0.6}$	$59.7_{\pm 1.2}$	$70.5_{\pm 0.8}$	$75.6_{\pm 1.0}$	$81.3_{\pm 0.7}$	$72.8_{\pm 0.9}$	$82.1_{\pm 1.1}$	$73.4_{\pm 0.6}$	$69.2_{\pm 0.8}$	$69.5_{\pm 1.0}$
	✓	$60.4_{\pm 0.9}$	$68.3_{\pm 0.5}$	$74.6_{\pm 1.0}$	$65.2_{\pm 0.8}$	$60.8_{\pm 1.2}$	$67.7_{\pm 0.7}$	$64.5_{\pm 0.6}$	$80.9_{\pm 1.1}$	$76.1_{\pm 0.8}$	$76.7_{\pm 0.5}$	$71.3_{\pm 1.0}$	$71.6_{\pm 0.9}$	$71.5_{\pm 0.6}$	$55.8_{\pm 1.1}$
✓	✓	$83.5_{\pm 0.6}$	$80.6_{\pm 0.5}$	$89.9_{\pm 1.2}$	$86.3_{\pm 0.7}$	$85.5_{\pm 0.6}$	$81.9_{\pm 1.0}$	$86.7_{\pm 1.1}$	$93.8_{\pm 0.9}$	$96.2_{\pm 0.5}$	$92.6_{\pm 0.5}$	$95.4_{\pm 1.3}$	$90.7_{\pm 0.7}$	$88.4_{\pm 0.9}$	$80.3_{\pm 1.1}$

Utility of Consistency. We explore the utility of consistency loss. Consistency loss $\mathcal{L}_{CA}$ contain two components, $\sum_{l=1}^{L} I(H_{g_s}^l; H_{g_t}^l)$ denotes the inter-granularity consistency, while $\sum_{g \in \{g_s, g_t\}} I(H_{int}; H_g)$ denotes the cross-granularity utility. Specifically, inter-granularity consistency mainly regularizes the mutual information of latent representations from different granularities, while cross-granularity consistency focuses on achieving alignment between granularity-specific representations and the integrated representations. Table 5 shows the ablation performance on single-graph datasets with respect to node- and edge-level detection. Table 6 shows the ablation performance on multi-graph datasets with respect to node- and graph-level detection. From the observations, both modules have a significant contribution to the performance, which demonstrates the effectiveness and rationality of the consistency adjustment.

Granularity-aware Fusion. We explore different strategies for generating integrated representations in Eq. (4). We introduce two variants, including Avg and Iner. Avg denotes $H_{int} = 0.5 * H_{gt} + 0.5 * H_{gs}$, while Iner denotes $H_{int} = H_{gt} \odot H_{gs}$, where $\odot$ is element-wise product. The observations are shown in Fig. 3. The variants have relatively poor performance compared with the proposed fusion strategy, demonstrating the effectiveness of using learnable parameters to dynamically integrate granularity-specific representations. More results are shown in Section E.

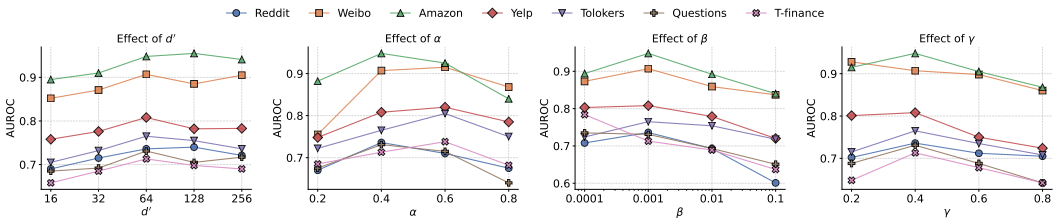


Figure 4: The sensitivity experiments on single-graph datasets with respect to the node-level task.

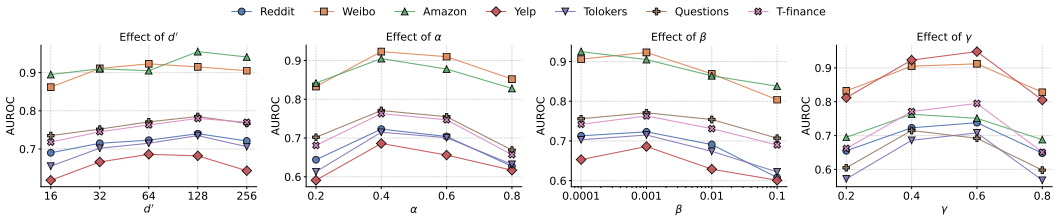


Figure 5: The sensitivity experiments on single-graph datasets with respect to the edge-level task.

Sensitivity. We conduct sensitivity experiments for hyperparameters, including hidden dimension d', α, β, and γ. Fig. 4 and Fig. 5 exhibit observations of GaHIGAD on single-graph datasets, respectively, for node- and edge-level tasks. Specifically, the size of the hidden dimension has a slight influence on the expression of latent representations. α controls the balance between $\mathcal{L}_{IB}$ and $\mathcal{L}_{CA}$. β adjusts the information compression ability in IB. γ controls the balance in Eq. (8). The sensitivity of α, β, and γ is more significant than that of the hidden dimension d', because they directly impact the characteristic extraction; an extreme value means that the contribution of a certain loss function is greatly compromised. More results are shown in Section F.

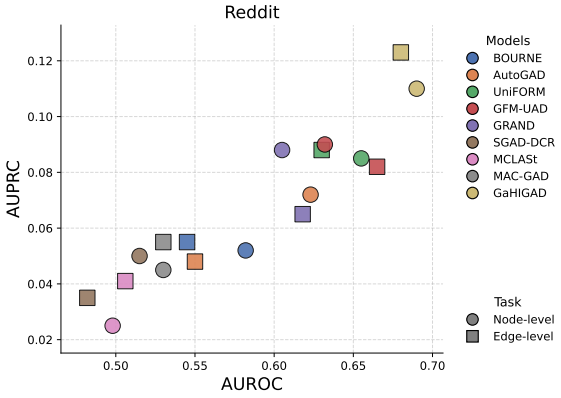


Figure 6: The robustness experiment on Reddit.

Robustness. We explore the robustness of GaHIGAD against interference by adding Gaussian noise to the Gaussian noise to 10% of the nodes. As shown in Fig. 6, we can see that when injecting noise, GaHIGAD still maintains superiority on both node-level and edge-level tasks over state-of-the-art baselines, demonstrating its strong and advanced robustness. MCLASt has the poorest performance among all models with respect to the node-level task, while SGAD-DCR exhibits the worst performance with respect to the node-level task. More experiments are shown in Section G.

6 Conclusion

This paper studies the multi-granularity graph anomaly detection (GAD) problem and proposes a novel framework, GaHIGAD. GaHIGAD tackles twofold challenges. On the one hand, GaHIGAD explores complex hierarchical information by introducing information bottleneck theory. On the other hand, GaHIGAD designs consistency learning to enhance representation regularization and discriminate cross-granularity dependencies. Extensive experiments demonstrate the effectiveness and advancement of GaHIGAD compared with state-of-the-art baselines. However, this framework ignores the utility of spectral signals. In the future, we will explore more effective spectral graph algorithms for unifying multi-granularity GAD.

AI USE STATEMENT

During manuscript preparation, the authors used ChatGPT solely for grammatical correction and English-language polishing. The tool was not used for research design, methodological development, code generation, experimentation, data analysis, or the generation of scientific results or conclusions. The authors take full responsibility for the content, accuracy, and integrity of the manuscript.

ETHICS STATEMENT

This study does not involve human subjects, practices related to dataset releases, potentially harmful insights, methodologies and applications, potential conflicts of interest and sponsorship, discrimination/bias/fairness concerns, privacy and security issues, legal compliance, and research integrity issues (e.g., IRB, documentation, research ethics).

REPRODUCIBILITY STATEMENT

Section 3 defines the problem, the proposed framework, and the corresponding objective. The supplement specifies more details about experimental configuration and results. All evaluated datasets are public and cited; no private data are used.

REFERENCES

Jinyu Cai, Yunhe Zhang, Pengyang Wang, and See-Kiong Ng. Moegad: A mixture-of-experts framework with pseudo-anomaly generation for graph-level anomaly detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 48(4):4473–4489, 2026. doi: 10.1109/TPAMI.2025.3646069.

Qian Chen, Huiying Xu, Ruidong Wang, Yue Liu, and Xinzhong Zhu. Structure adversarial augmented graph anomaly detection via multi-view contrastive learning. *Knowl. Based Syst.*, 338: 115455, 2026a. doi: 10.1016/J.KNOSYS.2026.115455.

Yujia Chen, Jian Li, Linfei Sun, Tongcun Liu, and Guanjun Liu. Semi-supervised graph anomaly detection via dual-channel reconstruction. *Neurocomputing*, 668:132401, 2026b. doi: 10.1016/J.NEUCOM.2025.132401.

Renda Han, Xiaobao Wang, Luzhi Wang, Wenxin Zhang, Guangzhen Yao, and Hongxiang Liang. A graph foundation model for unified anomaly detection. In *Proceedings of the ACM Web Conference 2026*, pp. 523–534, 2026. doi: 10.1145/3774904.3792108.

Yiqing Lin, Jianheng Tang, Chenyi Zi, H Vicky Zhao, Yuan Yao, and Jia Li. Unigad: Unifying multi-level graph anomaly detection. *Advances in neural information processing systems*, 37: 136120–136148, 2024.

Jie Liu, Mengting He, Xuequn Shang, Jieming Shi, Bin Cui, and Hongzhi Yin. Bourne: Bootstrapped self-supervised learning framework for unified graph anomaly detection. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pp. 2820–2833, 2024.

Rami Mohawesh, Sumbal Maqsood, Haythem Bany Salameh, and Ghaleb A. El Refae. Leveraging graphsage and large language models with cross-attention for transaction fraud detection: Evidence from credit card and paysim datasets. *Inf. Process. Manag.*, 63(7):104860, 2026. doi: 10.1016/J.IPM.2026.104860.

Chuancheng Song, Xixun Lin, Hanyang Shen, Yanmin Shang, and Yanan Cao. Uniform: Towards unified framework for anomaly detection on graphs. In *the Association for the Advancement of Artificial Intelligence*, pp. 12559–12567, 2025.

Maolin Wang, Beining Bao, Hongyu Chen, Zichun Liu, Lang Fu, Jun Chu, Langzhang Liang, and Zenglin Xu. GRAND: A robust diffusion framework for multi-granularity graph anomaly detection in web platforms. In *Proceedings of the ACM Web Conference 2026*, pp. 9572–9580, 2026. doi: 10.1145/3774904.3793012.

Chunyu Wei, Siyuan He, Yu Wang, Yueguo Chen, Yunhai Wang, Bing Bai, Yidong Zhang, Yong Xie, Shunming Zhang, and Fei Wang. 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*, pp. 1542–1553, 2026a. doi: 10.1145/3770854.3780240.

Hao Wei, Wen Wang, and Wanli Ni. Multimodal information bottleneck guided task-agnostic semantic communications. In *IEEE Wireless Communications and Networking Conference*, pp. 1–6, 2026b. doi: 10.1109/WCNC65185.2026.11555719.

Chunjing Xiao, Yao Tong, Jiahui Lu, Yuxia Xue, Wenxin Tai, Zhangtao Cheng, and Fan Zhou. Graph-enhanced multi-scale contrastive learning for graph anomaly detection with adaptive diffusion models. *IEEE Trans. Big Data*, 11(6):2946–2957, 2025.

Zhenyu Yang, Ge Zhang, Shan Xue, Xiaoxiao Ma, Jian Yang, Hao Peng, Amin Beheshti, and Jia Wu. Revisiting graph-level anomaly detection: From partially to fully unsupervised learning. In *Proceedings of the ACM Web Conference 2026*, pp. 901–912, 2026. doi: 10.1145/3774904.3792307.

Jiaqi Yu, Yang Gao, Hong Yang, Zhihong Tian, Peng Zhang, and Xingquan Zhu. Automated graph anomaly detection with large language models. *Knowledge-based Systems*, 324:113809, 2025.

Ge Zhang, Huimei Li, Guohao Sun, Xiu Fang, Xixun Lin, Xiaobao Wang, Pengfei Jiao, and Liang Yang. Unsupervised graph-level anomaly detection via multi-granular graph structure learning. In *Proceedings of the Thirty-Fifth International Joint Conference on Artificial Intelligence and the 29th European Conference on Artificial Intelligence*, pp. 3196–3204, 2026a. doi: 10.24963/IJCAI.2026/355.

Muhan Zhang. Graph neural networks: link prediction. In *Graph Neural Networks: Foundations, Frontiers, and Applications*, pp. 195–223. Springer, 2022.

Wenxin Zhang and Cuicui Luo. GE-GNN: gated edge-augmented graph neural network for fraud detection. *IEEE Trans. Big Data*, 11(4):1664–1676, 2025.

Wenxin Zhang and Cuicui Luo. Mclast: Multi-hierarchy contrastive learning graph anomaly detection with structure-awareness. *Neurocomputing*, 669:132480, 2026.

Wenxin Zhang, Yuhao Zhu, Cuicui Luo, Desheng Wu, Xi Xuan, Ljupco Kocarev, Renda Han, and Xiangxiang Lang. Feature and structural uncertainty modeling for risk identification in supply chain networks. *Chaos, Solitons & Fractals*, 201:117238, 2025.

Wenxin Zhang, Xi Xuan, Renda Han, Zonghao Ying, Cuicui Luo, Desheng Wu, and Ljupco Kocarev. Robust rumor detection against noise. *Neurocomputing*, 671:132741, 2026b. doi: 10.1016/J.NEUCOM.2026.132741.

Wenxin Zhang, Xi Xuan, Guangzhen Yao, Renda Han, Chuang Zhao, Xiangxiang Lang, Feng Zhou, and Cuicui Luo. Multi-scale asymmetric graph contrastive anomaly detection. *Inf. Process. Manag.*, 64(1):105026, 2027.

A Datasets

The statistics of the datasets are shown in Table 7.

B Baselines

In this section, we introduce the details of baselines:

- BOURNE (Liu et al., 2024) extracts a subgraph centered on each target node as node context and transforms it into a dual hypergraph as edge context, and enables the mutual detection of node and edge anomalies by swapping the context embeddings between nodes and edges and measuring the agreement in the embedding space.

Table 7: Detailed statistics of the datasets used in our experiments.

Dataset	Train%	# Graphs	# Edges	# Nodes	# Dims	Nodes$_{ab}$	Edges$_{ab}$	Graphs$_{ab}$
Reddit	40%	1	168,016	10,984	64	3.33%	2.72%	/
Weibo	40%	1	416,368	8,405	400	10.33%	5.71%	/
Amazon	70%	1	8,847,096	11,944	25	6.87%	2.49%	/
Yelp	70%	1	7,739,912	45,954	32	14.53%	13.89%	/
Tolokers	50%	1	530,758	11,758	10	21.82%	33.44%	/
Questions	50%	1	202,461	48,921	301	2.98%	7.50%	/
T-Finance	40%	1	21,222,543	39,357	10	4.58%	2.77%	/
BM-MN	40%	700	40,032	12,911	1	48.91%	/	14.29%
BM-MS	40%	700	30,238	9,829	1	31.99%	/	14.29%
BM-MT	40%	700	32,042	10,147	1	34.49%	/	14.29%
MUTAG	40%	2,951	179,732	88,926	14	4.81%	/	34.40%
MNIST0	10%	70,000	41,334,380	4,939,668	5	35.46%	/	9.86%
MNIST1	10%	70,000	41,334,380	4,939,668	5	35.46%	/	11.25%
T-Group	40%	37,402	93,367,082	11,015,616	10	0.64%	/	4.26%

- AutoGAD (Yu et al., 2025) first designs a novel search space tailored for GAD tasks based on the characteristics of the graph autoencoder framework and then leverages a large language model as the controller of GNAS, guiding the LLM to rapidly generate architectures suitable for GAD within the search space through well-designed prompts.
- UniFORM (Song et al., 2025) unifies node-, edge-, and graph-level tasks from a subgraph perspective, leveraging an energy-based GNN for iterative multi-granular anomaly detection, and enhances meta-learning through contrastive learning and employs Langevin dynamics to generate phantom samples as substitutes for anomalous data, reducing reliance on labeled data.
- GFM-UAD (Han et al., 2026) utilizes adversarial learning to generate negative samples for pretraining and introduces KL divergence to achieve alignment.
- GRAND (Wang et al., 2026) integrates dual diffusion mechanisms with Transformer-based denoising networks, offering a unified paradigm for identifying both node- and edge-level irregular patterns.
- SGAD-DCR (Chen et al., 2026b) employs a dual-filter attention encoder to construct dual aggregation channels for low-frequency and high-frequency signals and mitigates the homophily trap via semantic-level alignment.
- MCLASt (Zhang & Luo, 2026) introduces node-node, node-subgraph, and subgraph-subgraph contrastive learning modules to capture multi-hierarchy consistency information, and integrates both attribute and topological anomaly signals for anomaly detection.
- MAC-GAD (Zhang et al., 2027) develops a multi-level dominant view generation mechanism to enhance the contrastive learning performance. Based on the dominant view, an asymmetric contrastive learning paradigm is introduced for local and global consistency to capture comprehensive consistency of anomaly nodes.
- GUL (Yang et al., 2026) is an end-to-end framework that iteratively optimizes anomaly detection and uncertainty modeling via an Expectation-Maximization algorithm, with a dedicated loss that utilizes potential anomalies to enhance robustness.
- M-GLAD (Zhang et al., 2026a) performs unsupervised GLAD via multi-granular graph structure learning, where multiple levels of structure granularity of graphs and prototypes are abstracted by a graph structure abstraction module. Granularity-aware anomaly scoring is enabled to consider anomalous characteristics across different levels of structure granularity.
- MoEGAD (Cai et al., 2026) introduces an iterative anomalous graph generation module to produce pseudo-anomalous graphs, and proposes a latent MoE to promote cross-task adaptability for diverse GLAD problems.

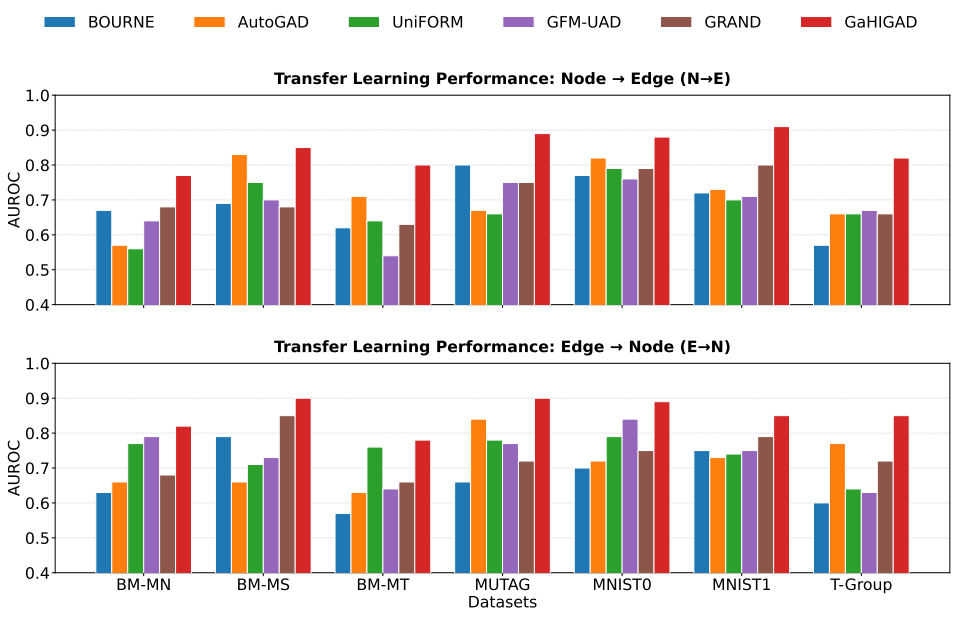


Figure 7: The zero-shot transfer detection experiments on single-graph datasets.

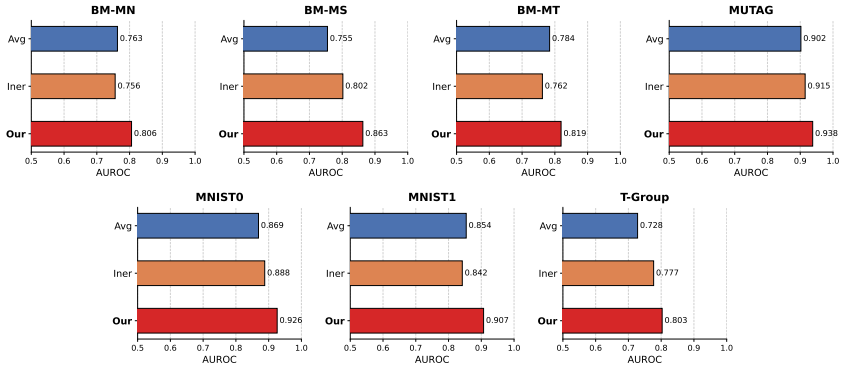


Figure 8: The ablation of fusion strategies on multi-graph datasets with respect to graph-level detection.

C Implementation

Considering the limited studies for edge-level detection, we integrate a concatenated strategy (Zhang, 2022) based on link prediction into the above baselines, and we generate abnormal edges based on the probability of nodes $p(e_{ij}) = \mathrm{Avg}(p_i, p_j)$ (Lin et al., 2024). For edge-level anomaly detection, we ignore edges between the training nodes and the testing nodes to prevent data leakage; for graph-level anomaly detection,

D Zero-Shot Transfer Learning

Since different tasks employ distinct baseline models on multi-graph datasets for zero-shot detection, we only adopt the first group of baselines, namely the multi-granularity GAD baselines, for validation. Fig. 7 exhibits zero-shot transfer learning on multi-graph datasets with respect to node- and graph-level detection tasks. From the results, we can see that GaHIGAD consistently and significantly outperforms other baselines.

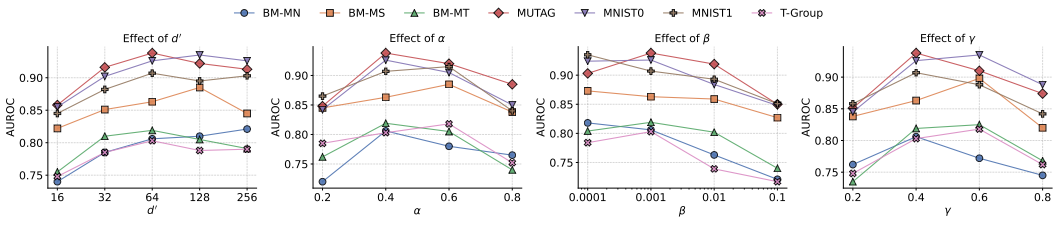


Figure 9: The sensitivity experiments on multi-graph datasets with respect to the graph-level task.

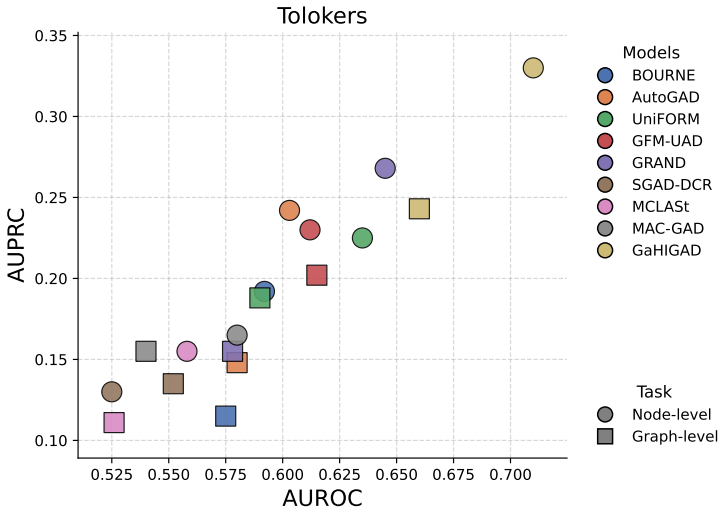


Figure 10: The robustness experiments on the Tolokers dataset.

E GRANULARITY-AWARE FUSION

Fig. 8 exhibits the ablation of fusion strategies on multi-graph datasets with respect to node- and graph-level detection. From the observations, we can see that the proposed fusion strategy displays superior performance consistently compared with variants, demonstrating the effectiveness of using learnable parameters to integrate the granularity-specific representations.

F SENSITIVITY

We conduct additional sensitivity experiments on multi-graph datasets. According to Fig. 9, for the hidden dimension, the performance exhibits slight fluctuations because the hidden size has an impact on the model expression. For the hyperparameters of the loss function, the performance on graph-level detection also demonstrates the utility of hyperparameters in the loss function.

G ROBUSTNESS

We also exhibit the robustness performance on the Tolokers dataset. From the observations in Fig. 10, we can see that GaHIGAD also achieves the best performance compared with state-of-the-art baselines, although it shows a slight performance degradation with the injection of noise.

H PSEUDO-CODE

The algorithm is shown in Algorithm 1.

Algorithm 1 Algorithm of GaHIGAD.

Input: Graph $\mathcal{G} = (\mathcal{V}, \mathbf{X}, \mathbf{A})$; Epochs E.

Initialize the parameters of GaHIGAD.
for $e \in (1, 2, \cdots, E)$ **do**
    Obtain node-wise hierarchical latent representations via Eq. (2)-Eq. (3);
    Achieve granularity-aware fusion using node-wise representations via Eq. (4);
    Obtain granularity-specific representations via Eq. (5);
    Obtain the detection results via Eq. (6);
    Calculate hierarchical IB loss via Eq. (8);
    Calculate consistency loss via Eq. (9);
    Calculate the total loss functions via Eq. (10);
    Update the model's parameters.
end for
return Scores