

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DYNAMIC MULTI-SAMPLE MIXUP WITH GRADIENT EX- PLORATION FOR OPEN-SET GRAPH ANOMALY DETEC- TION

006
007 **Anonymous authors**
008 Paper under double-blind review

011 ABSTRACT

013 This paper studies the problem of open-set graph anomaly detection, which aims
014 to generalize a graph neural network (GNN) trained with a small number of both
015 normal and abnormal nodes to detect unseen anomalies different from training
016 anomalies during inference. This problem is highly challenging due to both the
017 data scarcity of unseen anomalies and the label scarcity for training nodes. Towards
018 this end, we propose a novel approach named Dynamic Multi-sample Mixup with
019 Gradient Exploration (**DEMO**) for open-set graph anomaly detection. The core of
020 our proposed **DEMO** is to leverage a dynamic framework to adapt the optimization
021 procedure with high generalizability. In particular, our **DEMO** first adaptively fuses
022 multiple seen nodes to simulate the unseen anomalies, which expands the decision
023 boundary for the detection model with enhanced generalizability. Moreover, we
024 dynamically adjust sample weights based on their energy gradients to prioritize
025 uncertain and informative nodes, ensuring a robust optimization procedure. To fur-
026 ther address both label scarcity and severe class imbalance, we maintain a memory
027 bank of historical records to guide the pseudo-labeling process of unlabeled nodes.
028 Extensive experiments on various benchmark datasets validate the superiority of
029 the proposed **DEMO** in comparison to various baselines.

030 1 INTRODUCTION

031
032 Graph anomaly detection (GAD) focuses on identifying rare or malicious patterns in graph-structured
033 data that significantly deviate from expected behaviors (Akoglu et al., 2015; Ma et al., 2021; Tang
034 et al., 2023), such as fraudulent transactions in financial networks (Kim et al., 2024) or stealthy
035 attacks in Internet-of-Things (IoT) systems (Wu et al., 2021). Compared to traditional Euclidean
036 space-based anomaly detection, GAD requires the joint modeling of node attributes and topological
037 structure, where anomalies typically manifest as attribute distributional shifts coupled with structural
038 inconsistencies. This dual dependency introduces unique technical challenges: anomalies often
039 exhibit high structural diversity and strong context dependence, making them inherently more
040 difficult to define and detect (Vaska et al., 2022). As such, designing GAD approaches that are both
041 effective and generalizable remains a challenging and open research problem.

042 Current GAD methods predominantly rely on unsupervised and semi-supervised learning paradigms.
043 Unsupervised approaches exploit topological structures and attribute statistics of graphs to identify
044 anomalies in the absence of labeled data. These methods typically estimate anomaly scores based on
045 various principles (Ding et al., 2019; Liu et al., 2021; Ni et al., 2025), offering strong generalizability
046 across diverse graph domains. However, these frameworks suffer from limited precision due to the
047 absence of semantic guidance on anomaly properties. In contrast, semi-supervised frameworks (Dong
048 et al., 2025a; Chen et al., 2024; Ding et al., 2021a) leverage scarce labeled normal or anomalous
049 samples to guide learning processes, thereby achieving more discriminative representations and
050 enhanced detection performance. These methods typically incorporate consistency regularization
051 (Chen et al., 2024), generative objectives (Zhang et al., 2022a), and graph-based data augmentation
052 (Liu et al., 2022) strategies to improve model capabilities from limited supervision. Despite technical
053 diversity among existing approaches, they universally operate under a closed-set assumption, pre-
 suming training data fully represents all possible anomaly types or their underlying distributions, a
 premise that critically undermines their practicality in real-world graph analytics scenarios.

In this paper, we focus on an under-explored yet critically important problem: open-set graph anomaly detection (open-set GAD). The goal of this task is to train a GAD model using a limited number of labelled nodes, with the dual objective of detecting both seen anomalies and unseen ones that exhibit behavior patterns significantly distinct from training anomalies and absent from the training data. Achieving this goal is non-trivial, as it entails addressing several interrelated and fundamental challenges in graph-based settings. **I) How to generalize knowledge from limited seen anomalies to detect unseen anomalies?** In practice, the training set for open-set GAD often has two critical limitations: the number of anomaly classes is small, and the diversity of these classes is insufficient. To detect unseen anomalies, it is essential to fully extract and leverage all useful information from the few seen anomalies available. However, existing methods (e.g., (Wang et al., 2023b)) fail to tap into this potential, leading to a critical gap in using such supervision to support unseen anomaly detection. **II) How to perform effective graph anomaly detection under limited labels and severe class imbalance?** Open-set GAD faces two overlapping constraints: extremely scarce labeled data and severe class imbalance (normal nodes dominate the graph, while labeled/unlabeled anomalies are extremely rare). Existing semi-supervised GAD methods often struggle with this combination: they tend to overfit the dominant normal class, resulting in biased decision boundaries. These biased boundaries not only perform poorly on rare anomalies but also fail to generalize to unseen anomalies.

To address the aforementioned challenges, we propose a novel approach named Dynamic Multi-sample Mixup with Gradient Exploration (named **DEMO**) for open-set GAD. The goal of **DEMO** is to establish a dynamical adaptive training framework that improves generalization to unseen anomalies and enhances anomaly detection performance under limited supervision. Given the scarcity and homogeneity of anomaly classes in the training set, **DEMO** first adaptively fuses multiple seen anomaly samples to generate synthesized nodes with enriched representations, approximating unseen anomalies to drive the learning of broader decision boundaries. Moreover, while augmenting the training set with synthesized anomalies is beneficial, not all samples contribute positively to model generalization. Consequently, we employ an energy gradient-driven feedback mechanism to evaluate and re-weight each sample dynamically. This ensures the model prioritizes the optimization of latent uncertain samples crucial for generalization, effectively guiding the training process. Finally, to directly address the dual challenges of limited labels and severe class imbalance, **DEMO** uses a memory bank to guide pseudo-labeling with adaptive, class-specific confidence thresholds, thereby mitigating the resulting training bias. Extensive experiments on diverse benchmark datasets demonstrate **DEMO**’s consistently superior performance against state-of-the-art baselines, effectively validating its robustness and effectiveness in challenging open-set scenarios.

The contribution of this paper is summarized as follows: **① Problem Connection.** We present a new perspective that connects open-set recognition with GAD under label scarcity, emphasizing the need to generalize beyond seen anomalies and revisit GAD through the lens of open-set detection. **② Novel Framework.** We propose a novel framework named **DEMO**, which leverages a multi-sample mixup strategy and energy gradient-based feedback mechanism to fully exploit limited labeled data. Furthermore, **DEMO** incorporates a memory bank to mitigate label scarcity by guiding the pseudo-labeling process of unlabeled nodes. **③ Comprehensive Validation.** Extensive experiments on multiple real-world benchmark datasets and under various challenging evaluation settings validate the superiority of **DEMO** over a diverse range of state-of-the-art GAD baselines.

2 RELATED WORK

Graph Anomaly Detection. Existing graph anomaly detection (GAD) methods generally fall into two categories: unsupervised (Huang et al., 2023; Dong et al., 2025b; Qiao & Pang, 2023), semi-supervised (Gao et al., 2023; Huang et al., 2022; Dong et al., 2025a), and supervised (Ding et al., 2021b) methods. Unsupervised GAD does not rely on labeled anomalies, which identify outliers by modeling the graph’s inherent structural and attribute distributions (Ding et al., 2019; Fan et al., 2020). However, such methods often struggle to disentangle meaningful information from noise in the latent space. To improve discriminative ability, recent work incorporates diffusion-guided refinement of latent representations and content-preserving constraints, enhancing anomaly separability beyond what pure reconstruction can achieve (Li et al., 2024). Additionally, some unsupervised approaches integrate contrastive learning (Duan et al., 2023b;a) or generative modeling (Chen et al., 2020b) to bolster the robustness of learned representations. In contrast, semi-supervised GAD methods utilize a limited number of labeled anomalies to guide the learning process and enhance detection accuracy.

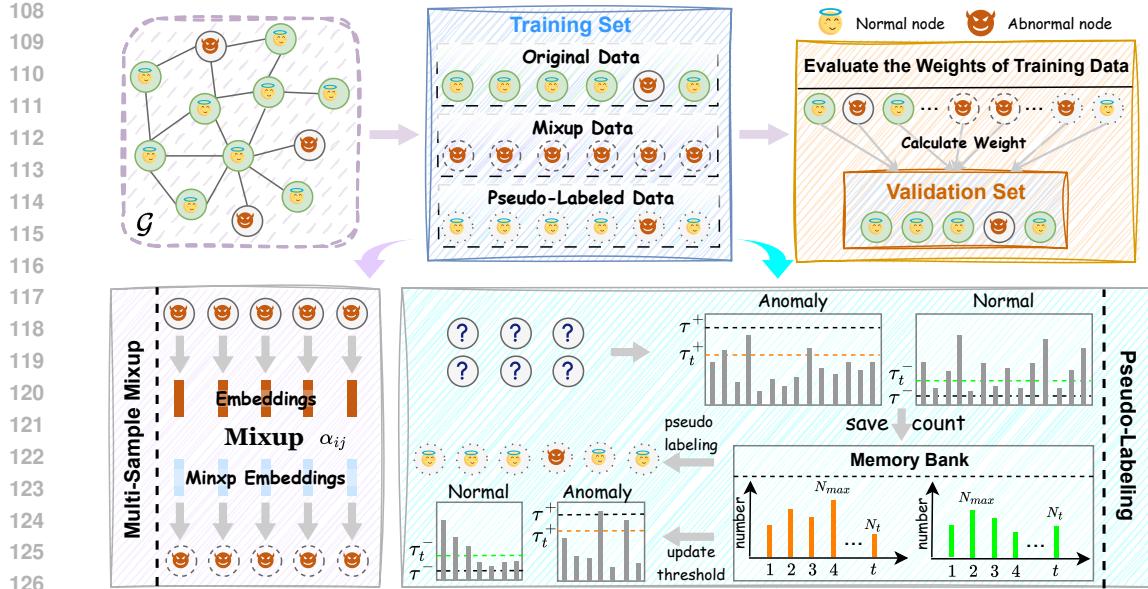


Figure 1: An overview of the proposed **DEMO**. **DEMO** first expands the training data through two parallel augmentation techniques. **Multi-Sample Mixup** generates new synthetic anomalies, while **Pseudo-Labeling** assigns labels to reliable unlabeled nodes. This augmented training data, comprising original, mixup, and pseudo-labeled data, then proceeds to a dynamic weighting stage.

Common strategies include consistency training to enforce prediction stability under perturbations (Chen et al., 2024), as well as graph-specific augmentations that improve the model’s ability to generalize from scarce supervision (Liu et al., 2022). Fully supervised GAD methods, on the other hand, assume comprehensive labels. Recent work has focused on complex generalization challenges, such as meta-learning for few-shot detection (Meta-GDN) (Ding et al., 2021b), cross-domain GAD (ACT) (Wang et al., 2023a), and generalist GAD (ARC (Liu et al., 2024), AnomalyGMF (Qiao et al., 2025)) aiming to unify performance across diverse graphs and anomaly types. However, most existing approaches, whether semi-supervised or fully supervised, assume a consistent anomaly distribution between training and testing, ignoring real-world variability and thereby limiting their ability to generalize to unseen anomaly types. To address this limitation, we propose a novel open-set GAD framework explicitly tailored for detecting diverse and previously unseen anomalies.

Open-set Classification. Open-set classification addresses a practical challenge where the model must not only accurately classify inputs from known categories but also identify instances originating from previously unseen classes (Yang et al., 2024; Wang et al., 2023b; Bisgin et al., 2024). Existing open-set classification approaches are mainly divided into discriminative (Luke et al., 2020; Chen et al., 2020a; Xu, 2024) and generative paradigms (Bao et al., 2023). Discriminative methods aim to learn well-separated decision boundaries for known classes while reserving ambiguous regions in the feature space to detect unknowns. In contrast, generative methods attempt to simulate the behavior of unknown classes by explicitly synthesizing samples, often combining generative models with discriminative classifiers in an adversarial framework. In the open-set GAD, NSReg (Wang et al., 2023b) is one of the few existing approaches that has shown promising results. It adopts a discriminative framework by introducing a regularization constraint to enforce compact semantic representations of normal nodes, thereby reducing overfitting to anomalies. However, its core idea focuses on constraining structural relations among normal nodes to enforce a strict decision boundary, while overlooking the role of anomaly nodes in shaping the model’s behavior. Therefore, our study emphasizes the significance of anomaly samples and achieves a diverse representation of anomaly samples through seen anomalies to enhance the generalization ability of the model.

3 THE PROPOSED APPROACH

Problem Definition. We consider the task of open-set GAD on attributed graphs, where only a small portion of anomalies are labeled during training and novel anomaly types may emerge at inference.

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ denote a graph with node set \mathcal{V} , edge set \mathcal{E} , and feature matrix $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$. The node set consists of a dominant group of normal nodes \mathcal{V}_n and a minority of anomalous nodes \mathcal{V}_a , i.e., $\mathcal{V} = \mathcal{V}_n \cup \mathcal{V}_a$, where $|\mathcal{V}_a| \ll |\mathcal{V}_n|$. During training, the labeled training nodes $\mathcal{V}^{\text{train}}$ only cover a partial set of anomaly classes. To formalize this, we partition \mathcal{V}_a into seen anomalies $\mathcal{V}_a^{\text{seen}}$ and unseen anomalies $\mathcal{V}_a^{\text{unseen}}$. The objective is to learn a detection model $\phi : (\mathcal{G}, \mathcal{V}) \rightarrow [0, 1]$ that assigns high anomaly scores to both $\mathcal{V}_a^{\text{seen}}$ and $\mathcal{V}_a^{\text{unseen}}$, while suppressing scores for \mathcal{V}_n . Formally, for all $v_a \in \mathcal{V}_a$ and $v_n \in \mathcal{V}_n$, the model satisfies: $\phi(\mathcal{G}, v_a) \gg \phi(\mathcal{G}, v_n)$.

3.1 FRAMEWORK OVERVIEW

Our study addresses the open-set GAD problem, aiming to detect both seen and unseen anomalies using limited labeled nodes. This involves two key challenges: (1) *Generalization to Unseen Anomalies*. (2) *Label Scarcity and Class Imbalance*. To tackle both challenges, we propose a novel framework **DEMO** that enhances the generalization of GNN-based detectors under minimal supervision. First, **DEMO** adaptively fuses multiple seen anomalies to synthesize diverse nodes that simulate unseen anomaly classes, thereby expanding the model's decision boundaries for improved generalization. Second, an energy gradient-driven feedback mechanism dynamically adjusts sample weight, enabling the model to focus on uncertain or ambiguous instances and improve generalization. Finally, to mitigate bias caused by severe class imbalance and limited anomaly supervision, we introduce a memory bank of historical records to adaptively update class-specific confidence thresholds and guide the pseudo-labeling of unlabeled nodes. The overall architecture of **DEMO** is illustrated in Figure 1, with detailed functionalities of each component elaborated in subsequent sections.

3.2 DYNAMIC MULTI-SAMPLE MIXUP FOR DECISION BOUNDARY EXPANSION

In the context of open-set GAD, existing methods (Wang et al., 2023b; Zhu et al., 2024; Isaac-Medina et al., 2024) struggle to capture the diversity of anomalies due to the limited quantity and class variety of anomalies in the training data. Consequently, it becomes crucial to introduce diversified anomaly samples to enhance the model's generalization. To this end, **DEMO** employs a dynamic multi-sample mixup strategy that generates synthesized anomalies with diverse representations, effectively prompting the model to learn broader and more robust decision boundaries. The core innovation lies in adaptively fusing multiple seen anomalies to generate more challenging and diverse anomaly representations, which not only approximate the distribution of potential unseen anomalies but also systematically encourage the model to learn more robust decision boundaries, thereby enhancing generalization to previously unseen patterns.

Definition 3.1 (Multi-sample Mixup) *Given the seen anomaly node set $\mathcal{V}_a^{\text{train}} = \{v_1^{\text{train}}, v_2^{\text{train}}, \dots, v_N^{\text{train}}\}$, which are encoded into embeddings $\mathbf{Z}_a^{\text{train}} = \{z_1^{\text{train}}, z_2^{\text{train}}, \dots, z_N^{\text{train}}\}$. For each original anomaly $z_i^{\text{train}} \in \mathbf{Z}_a^{\text{train}}$, we generate a corresponding synthetic representation \hat{z}_i . This synthetic sample is defined as a dynamic mix of all seen anomalies' embeddings: $\hat{z}_i = \sum_{j=1}^N \alpha_{ij} z_j^{\text{train}}$, where $\sum_{j=1}^N \alpha_{ij} = 1$ and $\alpha_{ij} \in [0, 1]$.*

Notably, the mixing weights α_{ij} are assigned based on the feature similarity between embeddings (Zhang et al., 2022b), defined through the following normalized form:

$$\alpha_{ij} = \frac{\exp(\mathcal{S}((z_i^{\text{train}})^\top \mathbf{w}_m, (z_j^{\text{train}})^\top \mathbf{w}_n))}{\sum_k \exp(\mathcal{S}((z_i^{\text{train}})^\top \mathbf{w}_m, (z_k^{\text{train}})^\top \mathbf{w}_n))}, \quad (1)$$

with \mathcal{S} representing a feature similarity function (e.g., inner product or cosine similarity), \mathbf{w}_m and \mathbf{w}_n are learnable weights. Since highly similar samples are more confusable to the model, assigning larger weights to them encourages the synthesized representations to stay in more ambiguous regions of the feature space. Next, we provide a theoretical justification showing that the synthesized samples can still retain high similarity with the original ones. The detailed proof is presented in the Appendix C.

Theorem 3.1 *Assume the inner-product similarity $\mathcal{S}(z_i, z_j) = z_i^\top z_j$ between original samples z_i^{train} and z_j^{train} is higher than their average similarity to all training anomalies, i.e., $\mathcal{S}(z_i^{\text{train}}, z_j^{\text{train}}) \geq \frac{1}{N} \sum_{k=1}^N \mathcal{S}(z_i^{\text{train}}, z_k^{\text{train}})$, then the synthesized sample \hat{z}_i satisfies:*

$$\mathcal{S}(\hat{z}_i, z_j) \geq \mathcal{S}(z_i, z_j) - \epsilon, \quad (2)$$

216 where $\epsilon = \sum_{k=1}^N \alpha_{ik} |\mathcal{S}(z_k, z_j) - \mathcal{S}(z_i, z_j)|$. That is, the similarity between the mixed sample \hat{z}_i
 217 and z_j is guaranteed to be no less than the original value minus a small perturbation ϵ , thereby
 218 preserving their intrinsic confusability.
 219

220 Furthermore, to avoid degenerate cases where the synthesized representation is overly biased toward
 221 the original sample z_i , we introduce a diversity regularization term to suppress its dominant influence:
 222

$$223 \quad \mathcal{L}_{\text{div}} = -\frac{1}{N} \sum_i \left\| \frac{(z_i^{\text{train}})^\top \mathbf{w}_m}{\|(z_i^{\text{train}})^\top \mathbf{w}_m\|_2} - \frac{(z_i^{\text{train}})^\top \mathbf{w}_n}{\|(z_i^{\text{train}})^\top \mathbf{w}_n\|_2} \right\|^2. \quad (3)$$

226 Upon synthesizing diverse anomalies, we further enhance the discriminative power of the model
 227 through consistency learning. The final loss function integrates both the diversity enhancement term
 228 (with a hyperparameter λ_{div}) and the consistency learning term as follows:
 229

$$230 \quad \mathcal{L}_{\text{mix}} = \mathcal{L}_{\text{cons}} + \lambda_{\text{div}} \mathcal{L}_{\text{div}}, \quad (4)$$

231 where $\mathcal{L}_{\text{cons}}$ is the consistency learning loss, calculated using the Mean Squared Error (MSE).
 232

233 3.3 GRADIENT EXPLORATION FOR GENERALIZABLE OPTIMIZATION

235 While the multi-sample mixup strategy creates a more diverse set of anomalies, not all samples
 236 in this expanded training set (both original and synthesized) are equally important for the model's
 237 generalization. Simply treating all samples equally is suboptimal, as some may be highly informative
 238 boundary cases while others might be redundant or even introduce noise. To address this, we propose
 239 an energy gradient-driven mechanism (Chen et al.; Han et al., 2022) that dynamically weights each
 240 sample based on its quantified impact on validation performance, thereby focusing optimization on
 241 the most critical nodes for more effective learning.
 242

243 We define the energy of a node as a measure of predictive uncertainty, formulated as $E_\theta(v_i) = -\log \sum \exp(z_i)$, where z_i denotes the predicted logits of node v_i . A lower energy indicates higher
 244 confidence, typically associated with normal nodes, while a higher energy suggests potentially anomalous
 245 or ambiguous behavior. To characterize how node energy affects generalization, we analyze
 246 model parameter responses to minor energy perturbations. The response direction is determined by
 247 the energy gradient to model parameters, with magnitude governed by the Hessian matrix $H_{\hat{\theta}}$:
 248

$$248 \quad \mathcal{I}_{\hat{\theta}}(v_i) = -H_{\hat{\theta}}^{-1} \nabla_\theta E_\theta(v_i). \quad (5)$$

250 The Hessian matrix evaluated at optimal parameters, quantifying sensitivity to local perturbations.
 251 However, parameter shifts alone cannot determine generalization benefits. We further evaluate their
 252 impact on validation loss by computing the inner product between validation loss gradients and
 253 energy gradient directions:
 254

$$255 \quad \mathcal{I}_{v_j^{\text{val}}}(v_i) = -\nabla_\theta \mathcal{L}(v_j^{\text{val}}, y_j^{\text{val}}; \hat{\theta})^\top H_{\hat{\theta}}^{-1} \nabla_\theta E_\theta(v_i), \quad (6)$$

256 where $\mathcal{L}(\cdot)$ denotes the binary cross-entropy loss, v_j^{val} and y_j^{val} are the nodes of the validation set
 257 and the corresponding categories. Averaging over all validation nodes yields the average influence
 258 $T_{\text{val}}(v_i)$ of training node v_i , which determines the adaptive weight coefficient:
 259

$$260 \quad \beta_{v_i} = -\frac{T_{\text{val}}(v_i)}{\max_{v_k \in \mathcal{V}^{\text{train}}} |T_{\text{val}}(v_k)|}. \quad (7)$$

262 Finally, we incorporate these weights into the energy-aware objective and define the overall loss as:
 263

$$264 \quad \mathcal{L}_{\text{energy}} = \frac{1}{n} \sum_{i=1}^n [\mathcal{L}(v_i, y_i; \theta) + \lambda_{\text{eng}} \beta_{v_i} \cdot E_\theta(v_i)], \quad (8)$$

267 where λ_{eng} is a balance coefficient. The core idea of this mechanism is that when $T_{\text{val}}(v_i) > 0$,
 268 reinforcing the energy guidance of node v_i can effectively reduce validation error; such nodes are
 269 often boundary cases that may resemble unseen anomalies. Conversely, a negative influence suggests
 that the sample may impair generalization, and its contribution should be suppressed.
 270

270 3.4 RELIABLE PSEUDO-LABELING WITH HISTORICAL GUIDANCE
271

272 After effectively leveraging labeled data to enhance the model’s generalization to unseen anomalies,
273 another crucial objective focuses on generating high-quality pseudo-labels for unlabeled samples to
274 further alleviate the training bottleneck caused by label scarcity. Traditional pseudo-labeling methods
275 employing fixed thresholds often fail to account for the dynamic evolution of model prediction
276 behaviors during training (Xie et al., 2020; Guo & Li, 2022). Meanwhile, dynamic thresholding
277 strategies (Zhang et al., 2021; Wang et al., 2023c; Chen et al., 2023) perform poorly in imbalanced
278 binary-class GAD due to their class-agnostic design, which often overlooks anomalies. To address
279 this, we propose a class-aware threshold adaptation scheme guided by historical records, enabling
280 reliable pseudo-label generation.

281 In detail, our method implements class-specific dynamic threshold adaptation via a memory bank
282 that incorporates historical information. Let \mathcal{T}_t^c denote the set of samples selected for class c
283 ($c \in \{0(\text{normal}), 1(\text{anomaly})\}$) at the t -th iteration, with its size $N_t^c = |\mathcal{T}_t^c|$ stored in the memory
284 bank to track distribution evolution. By computing the historical peak $N_{\max}^c = \max_{1 \leq k \leq t} N_k^c$,
285 the model dynamically captures the historical optimal coverage for class c . Based on the ratio
286 between the current selection count $\sigma_t(c)$ and N_{\max}^c , the model dynamically adjusts the threshold
287 $\rho_t(c) = \sigma_t(c)/N_{\max}^c$, where $\sigma_t(c)$ is determined by the following condition:

$$288 \sigma_t(c) = \sum_{i=1}^M \mathbb{I} [c = 1 \wedge \hat{p}_t(v_i) \geq \tau_t^+ \vee c = 0 \wedge \hat{p}_t(v_i) \leq \tau_t^-], \quad (9)$$

291 where M is the number of unlabeled samples, $\hat{p}_t(v_i)$ is the predicted probability of node v_i at iteration
292 t , and τ_t^+ and τ_t^- represents the dynamic thresholds for anomaly and normal classes, respectively.
293 The final asymmetric threshold update mechanism is formulated as:

$$295 \tau_t^{+/-} = \begin{cases} \rho_t(c) \cdot \tau^+, & c = \text{anomaly}, \\ \tau^- \cdot (2 - \rho_t(c)), & c = \text{normal}, \end{cases} \quad (10)$$

297 where τ^+ and τ^- denote predefined anomaly and normal class thresholds, with $\tau^+ + \tau^- = 1$.
298 This strategy ensures that anomaly thresholds τ_t^+ progressively increase with $\rho_t(\text{anomaly})$ to en-
299 hance minority-class sensitivity, while normal thresholds τ_t^- decrease through the nonlinear term
300 $(2 - \rho_t(\text{normal}))$, improving robustness against majority-class dominance.

302 3.5 SUMMARIZATION
303

304 In practice, we dynamically adjust the class-specific thresholds during training and select unlabeled
305 samples. The overall training objective is defined as:

$$307 \mathcal{L} = \mathcal{L}_{\text{energy}} + \lambda_{\text{mix}} \mathcal{L}_{\text{mix}} + \lambda_{\text{un}} \mathcal{L}_{\text{un}}. \quad (11)$$

308 where λ_{mix} and λ_{un} are the balancing coefficients for the mixup loss and the unlabeled loss, respec-
309 tively. \mathcal{L}_{un} denotes the binary cross-entropy loss applied to pseudo-labeled samples. The detailed
310 training procedure of **DEMO** is provided in the Appendix B.

312 4 EXPERIMENTS
313314 4.1 EXPERIMENTAL SETTINGS
315

316 **Datasets.** To comprehensively evaluate the performance of **DEMO** on the open-set GAD task, we
317 conduct experiments on six real-world graph datasets that vary in scale and domain characteristics.
318 Since no existing graph benchmarks are explicitly designed for open-set scenarios, we simulate by
319 partitioning normal and anomalous nodes based on their non-uniform class distributions. Following
320 (Wang et al., 2023b), we define normal nodes as majority classes and partition anomalies based on
321 class proportions: (1) For small-scale datasets with multiple classes, such as *Photo*, *Computers*, and
322 *CS* (Shchur et al., 2018), we treat categories with fewer than 5% of the total nodes as anomaly classes
323 (with at least two anomaly classes retained); (2) For large-scale datasets, including *Yelp* (Rayana
324 & Akoglu, 2015), *ogbn-arxiv* (Mikolov et al., 2013), and *ogbn-mag* (Wang et al., 2020), we apply

324
 325 Table 1: AUC-ROC and AUC-PR on three small-scale datasets. The best performance is boldfaced,
 326 with the second-best underlined.

327 328 Metrics	329 Photo		330 Computers		331 CS	
	AUC-ROC	AUC-PR	AUC-ROC	AUC-PR	AUC-ROC	AUC-PR
ANOMALOUS	0.5574 ± 0.012	0.0879 ± 0.003	0.5737 ± 0.016	0.1693 ± 0.003	0.2997 ± 0.017	0.1634 ± 0.015
DOMINANT	0.4716 ± 0.028	0.0837 ± 0.009	0.5450 ± 0.015	0.1644 ± 0.008	0.4029 ± 0.012	0.1886 ± 0.024
AnomalyDAE	0.4179 ± 0.032	0.0770 ± 0.007	0.5658 ± 0.011	0.1723 ± 0.007	0.3978 ± 0.009	0.1864 ± 0.008
GAAN	0.4346 ± 0.014	0.0710 ± 0.003	0.5595 ± 0.029	0.1796 ± 0.011	0.4646 ± 0.026	0.2111 ± 0.017
CoLA	0.5618 ± 0.008	0.0989 ± 0.006	0.4897 ± 0.010	0.1472 ± 0.006	0.4353 ± 0.014	0.2029 ± 0.015
CONAD	0.4763 ± 0.037	0.0862 ± 0.011	0.5445 ± 0.023	0.1619 ± 0.015	0.4028 ± 0.006	0.1886 ± 0.009
CONISGAD	0.8668 ± 0.021	0.5987 ± 0.002	0.6250 ± 0.017	0.3572 ± 0.009	0.7178 ± 0.034	0.5271 ± 0.019
GGAD	0.7976 ± 0.032	0.5677 ± 0.004	0.7210 ± 0.043	0.4529 ± 0.014	0.9081 ± 0.029	0.8198 ± 0.022
TAM	0.6045 ± 0.015	0.1084 ± 0.003	0.4432 ± 0.014	0.1355 ± 0.010	0.6398 ± 0.011	0.3542 ± 0.013
OCGNN	0.6279 ± 0.007	0.1323 ± 0.007	0.5049 ± 0.022	0.1505 ± 0.013	0.7819 ± 0.031	0.4926 ± 0.028
ANO-S	0.5730 ± 0.011	0.1097 ± 0.006	0.4628 ± 0.004	0.1392 ± 0.005	0.8380 ± 0.025	0.6401 ± 0.014
DOM-S	0.5785 ± 0.063	0.1107 ± 0.013	0.4488 ± 0.007	0.1330 ± 0.003	0.8445 ± 0.013	0.6382 ± 0.019
SpaceGNN	0.8030 ± 0.005	0.5271 ± 0.009	0.8296 ± 0.032	0.6439 ± 0.019	0.7784 ± 0.042	0.6587 ± 0.031
NSReg	0.8360 ± 0.012	0.4777 ± 0.018	0.7403 ± 0.021	0.5437 ± 0.013	0.9032 ± 0.035	0.8115 ± 0.016
GNN+OpenMax	0.7618 ± 0.104	0.3942 ± 0.063	0.6713 ± 0.052	0.3942 ± 0.095	0.8213 ± 0.133	0.7559 ± 0.049
DEMO	0.9023 ± 0.009	0.6330 ± 0.006	0.8439 ± 0.024	0.6458 ± 0.013	0.9448 ± 0.019	0.8857 ± 0.010

346 a similar strategy based on class proportions to distinguish normal and anomalous nodes. Further
 347 implementation details are provided in the Appendix.

348 **Baselines.** In our experiments, we compare **DEMO** with 12 representative GAD methods, which
 349 are categorized into two groups: unsupervised and semi-supervised approaches. The unsupervised
 350 methods include ANOMALOUS (Peng et al., 2018), DOMINANT (Ding et al., 2019), AnomalyDAE
 351 (Fan et al., 2020), GAAN (Chen et al., 2020b), CoLA (Liu et al., 2021), and CONAD (Xu et al.,
 352 2022). The semi-supervised group includes ConsisGAD (Chen et al., 2024), GGAD (Qiao et al.,
 353 2024), TAM (Qiao & Pang, 2023), OCGNN (Wang et al., 2021), ANOMALOUS-Semi (ANO-S),
 354 DOMINANT-Semi (DMO-S), SpaceGNN (Dong et al., 2025a), and NSReg (Wang et al., 2023b).
 355 Notably, TAM, OCGNN, ANOMALOUS, and DOMINANT are originally unsupervised models.
 356 Following the design strategy in GGAD, we adapt them into semi-supervised variants to ensure fair
 357 comparisons under the same supervision setting. Finally, we add GNN+OpenMax (Bendale & Boult,
 358 2016), which adapts the classic open-set classification framework to a GNN baseline.

359 **Implementation Details and Evaluation Metrics.** Following (Wang et al., 2023b), we adopt
 360 GraphSAGE as the backbone with 64 hidden units per layer. The model is optimized using the
 361 Adam optimizer with a learning rate of 0.001 and a weight decay of 0.0005. For training, we run
 362 200 epochs on small-scale datasets and 400 epochs on large-scale datasets. We select 50 anomalous
 363 nodes (from a single anomaly class) and 5% of normal nodes as the training set, and 30 anomalous
 364 nodes (same class as training) along with 1% of normal nodes as the validation set; all remaining
 365 nodes are used for testing. We set the hyperparameter λ_{div} to 0.5, and conduct a sensitivity analysis
 366 on the loss weights λ_{eng} , λ_{mix} , and λ_{un} , with detailed results provided in the Appendix. We evaluate
 367 performance using two widely adopted metrics: Area Under the Receiver Operating Characteristic
 368 Curve (*AUC-ROC*), which measures the model’s ability to distinguish between normal and anomalous
 369 nodes, and Area Under the Precision-Recall Curve (*AUC-PR*), which is particularly informative in
 imbalanced scenarios by emphasizing precision on the minority (anomalous) class.

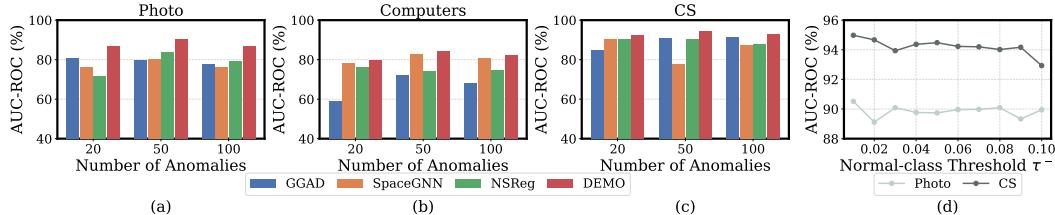
370 4.2 EMPIRICAL RESULTS

371 **Small-scale GAD Performance.** We first perform the performance comparison of **DEMO** on three
 372 small-scale graph datasets, the results are shown in Table 1. As shown, **DEMO** consistently outperforms
 373 all existing baselines across both *AUC-ROC* and *AUC-PR* metrics. **1** *Among unsupervised methods,*
 374 **DEMO achieves substantial improvements over all competitors.** For instance, on the Photo dataset,
 375 it surpasses the best-performing unsupervised baseline (CoLA) by 60.61% in *AUC-ROC*, and the
 376 gain in *AUC-PR* is even more pronounced. Across three datasets, **DEMO** achieves 80.99% higher

378

379
380
Table 2: AUC-ROC and AUC-PR on three large-scale datasets. The best performance is boldfaced,
with the second-best underlined. ‘OOM’ indicates out-of-memory.

381 Datasets	382 Metrics	383 ConsisGAD	384 GGAD	385 TAM	386 OGCNN	387 ANO-S	388 DOM-S	389 SpaceGNN	390 NSReg	391 DEMO
392 Yelp	AUC-ROC	0.6988	0.6613	0.5319	0.6410	0.6567	0.6506	0.6853	<u>0.7015</u>	0.7097
	AUC-PR	<u>0.2970</u>	0.2549	0.0977	0.1118	0.1076	0.1048	0.2916	0.3029	0.2238
394 ogbn-arxiv	AUC-ROC	<u>0.6216</u>	0.6007	OOM	OOM	0.4510	0.4505	0.6133	0.6182	0.6364
	AUC-PR	0.3148	0.2843	OOM	OOM	0.1463	0.1482	<u>0.3301</u>	0.3230	0.3329
396 ogbn-mag	AUC-ROC	<u>0.4909</u>	OOM	OOM	OOM	OOM	OOM	0.4626	0.4836	0.4967
	AUC-PR	<u>0.0043</u>	OOM	OOM	OOM	OOM	OOM	<u>0.0043</u>	0.0041	0.0054

392
393
394
395
396
Figure 2: (a), (b) and (c) Comparison of model performance under different numbers of training
397 anomalies across three datasets. (d) Sensitivity analysis of the threshold τ^- .

398
399 average AUC-ROC (vs. 0.4956 of CoLA) and 368.11% improvement in average AUC-PR (vs.
400 GAAN’s 0.1539). These results indicate that unsupervised approaches relying solely on structural
401 reconstruction or embedding learning are limited in capturing the diversity of anomaly types under
402 open-set conditions, resulting in poor generalization. **2** *In the semi-supervised setting, DEMO*
403 *maintains dominant advantages, notably outperforming strong baselines such as NSReg*. On average,
404 across the three datasets, **DEMO** achieves a 7.86% improvement in AUC-ROC over the second-
405 best method (NSReg), and a 17.60% gain in AUC-PR compared to GGAD. This highlights that
406 beyond effectively leveraging limited labeled data, **DEMO** benefits from its diverse anomaly modeling
407 and adaptive mechanisms, yielding stronger recognition of unseen anomalies. **3** *In summary, the*
408 *comprehensive performance of DEMO on small-scale datasets validates its effectiveness in the open-*
409 *set GAD task, particularly under conditions of sparse anomaly distribution and limited supervision,*
410 *where unified and robust modeling of both seen and unseen anomalies is essential.*

411 **Large-scale GAD Performance.** Furthermore, we present the performance evaluation results of
412 **DEMO** on three large-scale graph datasets in Table 2. **1** *As observed, DEMO demonstrates strong*
413 *robustness and broad adaptability in large-scale open-set scenarios.* On the Yelp dataset, **DEMO**
414 achieves an AUC-ROC of 0.7097, outperforming the second-best method NSReg by 1.16%. Although
415 its AUC-PR is slightly lower than NSReg, it still exhibits reliable and stable performance. On the two
416 extremely large datasets, **DEMO** delivers consistently strong results, improving AUC-ROC by 2.33%
417 and 1.17%, and boosting AUC-PR by 0.85% and 25.58%, respectively, compared to the strongest
418 baselines. **2** *These results confirm that DEMO maintains reliable performance even on large-scale*
419 *graphs, and is particularly effective in handling complex topologies and extreme class imbalance by*
420 *robustly expanding decision boundaries and accurately identifying previously unseen anomalies.*

421 **Data Efficiency.** To further evaluate the impact of anomaly sample quantity on model performance,
422 we conducted experiments by varying the number of anomalous nodes in the training set across
423 three levels: 20, 50, and 100. These experiments were performed on three datasets. Additional
424 results and comparisons are provided in the Appendix. Figure 2 (a-c) illustrates how each method’s
425 AUC-ROC score changes with different amounts of anomaly supervision. As shown in the figure,
426 **DEMO** consistently maintains the best or near-best performance across all datasets as the number of
427 anomalies increases, and notably demonstrates clear advantages in low-resource settings, such as
428 when 20 anomalous nodes are available. On the CS dataset, even with just 20 anomaly nodes, **DEMO**
429 achieves an AUC-ROC of approximately 0.90, showing a substantial gap. On the Computers dataset,
430 **DEMO** reaches its peak performance with 50 anomalies, significantly outperforming all baselines.
431 In summary, **DEMO** exhibits strong anomaly detection capabilities under severe data scarcity and
432 maintains excellent scalability and training efficiency as the number of anomaly samples grows,
433 confirming its advantage in efficient use of limited supervision.

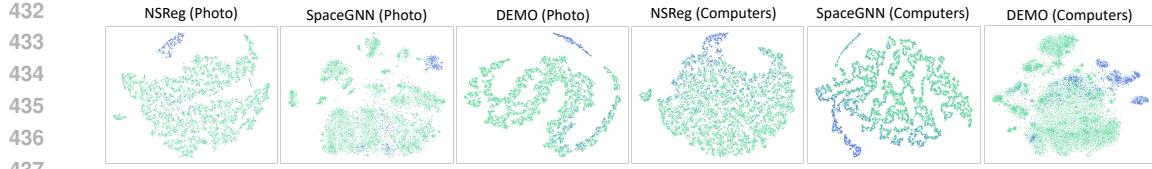


Figure 3: Visualization results of the features of the three algorithms on the two datasets.

Sensitivity Analysis. In this part, we conduct a sensitivity analysis on the hyperparameters τ^+ and τ^- . Figure 2 (d) illustrates the impact of the predefined normal-class threshold τ^- on the detection performance under our proposed class-aware threshold adjustment strategy. Since the anomaly-class threshold is defined as $1 - \tau^-$, this experiment effectively evaluates the joint influence of both class thresholds. We vary τ^- within the range of $[0.01, 0.1]$ and examine its effect on AUC-ROC across two datasets. The results show that both datasets achieve optimal performance when $\tau^- = 0.01$, indicating a well-balanced separation between normal and anomalous classes. As τ^- decreases (*i.e.*, the anomaly threshold increases), the model performance begins to fluctuate or decline, suggesting that an overly relaxed anomaly selection criterion may introduce low-confidence samples into training and compromise learning. This demonstrates that setting a relatively lower normal-class threshold (or, equivalently, a higher anomaly-class threshold) enables the model to more reliably select high-confidence pseudo-labeled samples, effectively mitigating label bias and improving overall detection performance. Therefore, we set τ^+ and τ^- to 0.99 and 0.01, respectively.

Ablation Study. We introduce several variants of **DEMO** to analyze the contribution of each component: (1) **DEMO** w/o All removes all three components; (2) **DEMO** w/o Mix removes the multi-sample fusion component; (3) **DEMO** w/o EG removes the energy gradient-guided optimization strategy; and (4) **DEMO** w/o PL disables the pseudo-labeling component. As shown in Table 3, the ablation results reveal that each component plays a vital role in enhancing model performance. First, removing the PL component results in the most substantial performance drop across both datasets, demonstrating the crucial impact of class-aware pseudo-labeling in mitigating label scarcity and improving anomaly detection. Second, the exclusion of the EG module consistently reduces performance, validating the effectiveness of gradient-based reweighting in prioritizing informative and uncertain samples during training. Third, the absence of the Mix module also causes noticeable degradation, highlighting the benefit of synthesizing diverse anomaly representations to better expand the decision boundary. Additionally, the variant **DEMO** w/o All, which disables all three components, yields the lowest overall performance, indicating that the full combination of modules is essential for achieving the best anomaly detection results. Overall, these results confirm that all three modules contribute synergistically to the robustness and generalization of **DEMO**.

Visualization Results. To further assess the representation quality of node embeddings learned by different models, we visualize the 2D t-SNE projections of node features on the Photo and Computers datasets, as shown in Figure 3. Blue and green points correspond to anomaly and normal nodes, respectively. On both datasets, **DEMO** produces clearly separated clusters between anomalous and normal nodes. The anomalies (in blue) are more tightly grouped and distinctly detached from the normal class, indicating that **DEMO** effectively enhances anomaly-specific feature encoding and decision boundary clarity. In contrast, the embeddings generated by NSReg and SpaceGNN exhibit substantial overlap between the two classes, particularly on the Computers dataset, where anomalies appear largely entangled with normal nodes.

5 CONCLUSION

This paper presents **DEMO**, a novel framework for open-set graph anomaly detection. By combining multi-sample fusion with gradient-guided weight adjustment, **DEMO** enhances model generalization to unseen anomalies. To mitigate the effects of label scarcity, it further incorporates a class-aware threshold adaptation scheme guided by historical records for reliable pseudo-label generation. Extensive experiments validate its effectiveness across diverse datasets. In addition to strong empirical perfor-

Table 3: Ablation Study on two benchmarks.

Datasets Metrics	Photo		Computers	
	AR	AP	AR	AP
DEMO w/o All	0.8300	0.5692	0.7576	0.5325
DEMO w/o Mix	0.8750	0.6023	0.8197	0.6292
DEMO w/o EG	0.8849	0.6171	0.8100	0.5998
DEMO w/o PL	0.8616	0.6150	0.8094	0.5949
DEMO	0.9023	0.6330	0.8439	0.6458

486 mance, **DEMO** demonstrates robustness across varying data scales and supervision levels, making it
 487 broadly applicable to real-world graph anomaly detection tasks.
 488

489 **ETHICS STATEMENT**
 490

491 In conducting this research, we have adhered to the ICLR Code of Ethics. Our work utilizes six
 492 publicly available and widely used benchmark datasets: Photo, Computers, CS, Yelp, ogbn-arxiv, and
 493 ogbn-mag. These datasets were collected by third parties for academic research, and we have used
 494 them in accordance with their intended purpose. Our study did not involve the collection of new data
 495 from human subjects. Furthermore, this work does not involve any other ethical issues.
 496

497 **REPRODUCIBILITY STATEMENT**
 498

499 To ensure the reproducibility of our work, we provide detailed descriptions of our methodology and
 500 experimental setup. The complete training algorithm for our proposed **DEMO** framework is outlined
 501 in Algorithm 1. All datasets used are publicly available, and comprehensive processing details are
 502 provided in Appendix D. Our implementation details, including hyperparameter settings and the
 503 experimental protocol, are described in Sec. 4.1 and Appendix F. The theoretical justification for
 504 our multi-sample mixup strategy is supported by a full proof of Theorem 3.1 in Appendix C. We
 505 will make our source code publicly available at: <https://anonymous.4open.science/r/DEMO-7615/>.
 506

507 **REFERENCES**
 508

510 Leman Akoglu, Hanghang Tong, and Danai Koutra. Graph based anomaly detection and description:
 511 a survey. *Data mining and knowledge discovery*, 29:626–688, 2015.

512 Wentao Bao, Qi Yu, and Yu Kong. Latent space energy-based model for fine-grained open set
 513 recognition. *arXiv preprint arXiv:2309.10711*, 2023.

515 Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *Proceedings of the IEEE
 516 conference on computer vision and pattern recognition*, pp. 1563–1572, 2016.

517 Halil Bisgin, Andres Palechor, Mike Suter, and Manuel Günther. Large-scale evaluation of open-set
 518 image classification techniques. *arXiv preprint arXiv:2406.09112*, 2024.

520 Guangyao Chen, Limeng Qiao, Yemin Shi, Peixi Peng, Jia Li, Tiejun Huang, Shiliang Pu, and
 521 Yonghong Tian. Learning open set network with discriminative reciprocal points. In *Computer
 522 Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings,
 523 Part III 16*, pp. 507–522. Springer, 2020a.

524 Hao Chen, Ran Tao, Yue Fan, Yidong Wang, Jindong Wang, Bernt Schiele, Xing Xie, Bhiksha Raj,
 525 and Marios Savvides. Softmatch: Addressing the quantity-quality trade-off in semi-supervised
 526 learning. *International Conference on Learning Representations (ICLR)*, 2023.

528 Nan Chen, Zemin Liu, Bryan Hooi, Bingsheng He, Rizal Fathony, Jun Hu, and Jia Chen. Consistency
 529 training with learnable data augmentation for graph anomaly detection with limited supervision.
 530 In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024.

531 Yiting Chen, Qitian Wu, and Junchi Yan. Regularizing energy among training samples for out-of-
 532 distribution generalization. In *The Thirteenth International Conference on Learning Representations
 533 (ICLR)*.

534 Zhenxing Chen, Bo Liu, Meiqing Wang, Peng Dai, Jun Lv, and Liefeng Bo. Generative adversarial
 535 attributed network anomaly detection. In *Proceedings of the 29th ACM International Conference
 536 on Information & Knowledge Management*, pp. 1989–1992, 2020b.

538 Kaize Ding, Jundong Li, Rohit Bhanushali, and Huan Liu. Deep anomaly detection on attributed
 539 networks. In *Proceedings of the 2019 SIAM international conference on data mining*, pp. 594–602.
 SIAM, 2019.

540 Kaize Ding, Jundong Li, Nitin Agarwal, and Huan Liu. Inductive anomaly detection on attributed
 541 networks. In *Proceedings of the twenty-ninth international conference on international joint*
 542 *conferences on artificial intelligence*, pp. 1288–1294, 2021a.

543

544 Kaize Ding, Qinghai Zhou, Hanghang Tong, and Huan Liu. Few-shot network anomaly detection via
 545 cross-network meta-learning. In *Proceedings of the web conference 2021*, pp. 2448–2456, 2021b.

546

547 Xiangyu Dong, Xingyi Zhang, Lei Chen, Mingxuan Yuan, and Sibo Wang. Spacegnn: Multi-space
 548 graph neural network for node anomaly detection with extremely limited labels. *The Thirteenth*
 549 *International Conference on Learning Representations (ICLR)*, 2025a.

550

551 Xiangyu Dong, Xingyi Zhang, Yanni Sun, Lei Chen, Mingxuan Yuan, and Sibo Wang. Smoothgnn:
 552 Smoothing-aware gnn for unsupervised node anomaly detection. In *Proceedings of the ACM on*
 553 *Web Conference 2025*, pp. 1225–1236, 2025b.

554

555 Jingcan Duan, Siwei Wang, Pei Zhang, En Zhu, Jingtiao Hu, Hu Jin, Yue Liu, and Zhibin Dong.
 556 Graph anomaly detection via multi-scale contrastive learning networks with augmented view. In
 557 *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 7459–7467, 2023a.

558

559 Jingcan Duan, Pei Zhang, Siwei Wang, Jingtiao Hu, Hu Jin, Jiaxin Zhang, Haifang Zhou, and Xinwang
 560 Liu. Normality learning-based graph anomaly detection via multi-scale contrastive learning. In
 561 *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 7502–7511, 2023b.

562

563 Haoyi Fan, Fengbin Zhang, and Zuoyong Li. Anomalydae: Dual autoencoder for anomaly detection
 564 on attributed networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech*
 565 *and Signal Processing (ICASSP)*, pp. 5685–5689. IEEE, 2020.

566

567 Yuan Gao, Xiang Wang, Xiangnan He, Zhenguang Liu, Huamin Feng, and Yongdong Zhang. Ad-
 568 dressing heterophily in graph anomaly detection: A perspective of graph spectrum. In *Proceedings*
 569 *of the ACM web conference 2023*, pp. 1528–1538, 2023.

570

571 Lan-Zhe Guo and Yu-Feng Li. Class-imbalanced semi-supervised learning with adaptive thresholding.
 572 In *International Conference on Machine Learning*, pp. 8082–8094. PMLR, 2022.

573

574 Andi Han, Dai Shi, Zhiqi Shao, and Junbin Gao. Generalized energy and gradient flow via graph
 575 framelets. *arXiv preprint arXiv:2210.04124*, 2022.

576

577 Mengda Huang, Yang Liu, Xiang Ao, Kuan Li, Jianfeng Chi, Jinghua Feng, Hao Yang, and Qing He.
 578 Auc-oriented graph neural network for fraud detection. In *Proceedings of the ACM web conference*
 579 2022, pp. 1311–1321, 2022.

580

581 Yihong Huang, Liping Wang, Fan Zhang, and Xuemin Lin. Unsupervised graph outlier detection:
 582 Problem revisit, new insight, and superior method. In *2023 IEEE 39th International Conference*
 583 *on Data Engineering (ICDE)*, pp. 2565–2578. IEEE, 2023.

584

585 Brian KS Isaac-Medina, Yona Falinie A Gaus, Neelanjan Bhowmik, and Toby P Breckon. Towards
 586 open-world object-based anomaly detection via self-supervised outlier synthesis. In *European*
 587 *Conference on Computer Vision*, pp. 196–214. Springer, 2024.

588

589 Yejin Kim, Youngbin Lee, Minyoung Choe, Sungju Oh, and Yongjae Lee. Temporal graph networks
 590 for graph anomaly detection in financial networks. *arXiv preprint arXiv:2404.00060*, 2024.

591

592 Jinghan Li, Yuan Gao, Jinda Lu, Junfeng Fang, Congcong Wen, Hui Lin, and Xiang Wang. Diffgad:
 593 A diffusion-based unsupervised graph anomaly detector. *The Thirteenth International Conference*
 594 *on Learning Representations (ICLR)*, 2024.

595

596 Fanzhen Liu, Xiaoxiao Ma, Jia Wu, Jian Yang, Shan Xue, Amin Beheshti, Chuan Zhou, Hao Peng,
 597 Quan Z Sheng, and Charu C Aggarwal. Dagad: Data augmentation for graph anomaly detection.
 598 In *2022 IEEE international conference on data mining (ICDM)*, pp. 259–268. IEEE, 2022.

599

600 Yixin Liu, Zhao Li, Shirui Pan, Chen Gong, Chuan Zhou, and George Karypis. Anomaly detection on
 601 attributed networks via contrastive self-supervised learning. *IEEE transactions on neural networks*
 602 *and learning systems*, 33(6):2378–2392, 2021.

594 Yixin Liu, Shiyuan Li, Yu Zheng, Qingfeng Chen, Chengqi Zhang, and Shirui Pan. Arc: A generalist
 595 graph anomaly detector with in-context learning. *Advances in Neural Information Processing*
 596 *Systems*, 37:50772–50804, 2024.

597

598 Ditría Luke, Benjamin J Meyer, and Tom Drummond. Opengan: Open set generative adversarial
 599 networks. In *Proceedings of the Asian Conference on Computer Vision*, 2020.

600

601 Xiaoxiao Ma, Jia Wu, Shan Xue, Jian Yang, Chuan Zhou, Quan Z Sheng, Hui Xiong, and Le-
 602 man Akoglu. A comprehensive survey on graph anomaly detection with deep learning. *IEEE*
 603 *transactions on knowledge and data engineering*, 35(12):12012–12038, 2021.

604

605 Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. Image-based
 606 recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR*
 607 *conference on research and development in information retrieval*, pp. 43–52, 2015.

608

609 Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations
 610 of words and phrases and their compositionality. *Advances in neural information processing*
 611 *systems*, 26, 2013.

612

613 Hang Ni, Jindong Han, Nengjun Zhu, and Hao Liu. Unsupervised graph anomaly detection via
 614 multi-hypersphere heterophilic graph learning. *arXiv preprint arXiv:2503.12037*, 2025.

615

616 Zhen Peng, Minnan Luo, Jundong Li, Huan Liu, Qinghua Zheng, et al. Anomalous: A joint modeling
 617 approach for anomaly detection on attributed networks. In *IJCAI*, volume 18, pp. 3513–3519,
 618 2018.

619

620 Hezhe Qiao and Guansong Pang. Truncated affinity maximization: One-class homophily modeling for
 621 graph anomaly detection. *Advances in Neural Information Processing Systems*, 36:49490–49512,
 622 2023.

623

624 Hezhe Qiao, Qingsong Wen, Xiaoli Li, Ee-Peng Lim, and Guansong Pang. Generative semi-
 625 supervised graph anomaly detection. *Advances in neural information processing systems*, 2024.

626

627 Hezhe Qiao, Chaoxi Niu, Ling Chen, and Guansong Pang. Anomalygfm: Graph foundation model
 628 for zero/few-shot anomaly detection. In *Proceedings of the 31st ACM SIGKDD Conference on*
 629 *Knowledge Discovery and Data Mining V. 2*, pp. 2326–2337, 2025.

630

631 Shebuti Rayana and Leman Akoglu. Collective opinion spam detection: Bridging review networks
 632 and metadata. In *Proceedings of the 21th acm sigkdd international conference on knowledge*
 633 *discovery and data mining*, pp. 985–994, 2015.

634

635 Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls
 636 of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.

637

638 Jianheng Tang, Fengrui Hua, Ziqi Gao, Peilin Zhao, and Jia Li. Gadbench: Revisiting and bench-
 639 marking supervised graph anomaly detection. *Advances in Neural Information Processing Systems*,
 640 36:29628–29653, 2023.

641

642 Nathan Vaska, Kevin Leahy, and Victoria Helus. Context-dependent anomaly detection with knowl-
 643 edge graph embedding models. In *2022 IEEE 18th International Conference on Automation*
 644 *Science and Engineering (CASE)*, pp. 2020–2027. IEEE, 2022.

645

646 Kuansan Wang, Zhihong Shen, Chiyuan Huang, Chieh-Han Wu, Yuxiao Dong, and Anshul Kanakia.
 647 Microsoft academic graph: When experts are not enough. *Quantitative Science Studies*, 1(1):
 648 396–413, 2020.

649

650 Qizhou Wang, Guansong Pang, Mahsa Salehi, Wray Buntine, and Christopher Leckie. Cross-domain
 651 graph anomaly detection via anomaly-aware contrastive alignment. In *Proceedings of the AAAI*
 652 *Conference on Artificial Intelligence*, volume 37, pp. 4676–4684, 2023a.

653

654 Qizhou Wang, Guansong Pang, Mahsa Salehi, Xiaokun Xia, and Christopher Leckie. Open-set graph
 655 anomaly detection via normal structure regularisation. *The Thirteenth International Conference on*
 656 *Learning Representations*, 2023b.

648 Xuhong Wang, Baihong Jin, Ying Du, Ping Cui, Yingshui Tan, and Yupu Yang. One-class graph
 649 neural networks for anomaly detection in attributed networks. *Neural computing and applications*,
 650 33:12073–12085, 2021.

651 Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, , Zhen Wu, Jindong Wang, Marios
 652 Savvides, Takahiro Shinozaki, Bhiksha Raj, Bernt Schiele, and Xing Xie. Freematch: Self-adaptive
 653 thresholding for semi-supervised learning. *International Conference on Learning Representations*
 654 (*ICLR*), 2023c.

655 Yulei Wu, Hong-Ning Dai, and Haina Tang. Graph neural networks for anomaly detection in industrial
 656 internet of things. *IEEE Internet of Things Journal*, 9(12):9214–9231, 2021.

657 Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation
 658 for consistency training. *Advances in neural information processing systems*, 33:6256–6268, 2020.

659 Jiawen Xu. Know yourself better: Diverse discriminative feature learning improves open set
 660 recognition. *arXiv preprint arXiv:2404.10370*, 2024.

661 Zhiming Xu, Xiao Huang, Yue Zhao, Yushun Dong, and Jundong Li. Contrastive attributed network
 662 anomaly detection with data augmentation. In *Pacific-Asia conference on knowledge discovery*
 663 and data mining, pp. 444–457. Springer, 2022.

664 Chen Yang, Zhixi Feng, Shuyuan Yang, and Qiukai Pan. Open-icl: Open-set modulation classification
 665 via incremental contrastive learning. *IEEE Transactions on Neural Networks and Learning Systems*,
 666 2024.

667 Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. Graph-
 668 saint: Graph sampling based inductive learning method. *International Conference on Learning*
 669 *Representations (ICLR)*, 2019.

670 Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro
 671 Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling.
 672 *Advances in neural information processing systems*, 34:18408–18419, 2021.

673 Jiaqiang Zhang, Senzhang Wang, and Songcan Chen. Reconstruction enhanced multi-view contrastive
 674 learning for anomaly detection on attributed networks. *Proceedings of the Thirty-First International*
 675 *Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2376–2382, 2022a.

676 Shaofeng Zhang, Meng Liu, Junchi Yan, Hengrui Zhang, Lingxiao Huang, Xiaokang Yang, and
 677 Pinyan Lu. M-mix: Generating hard negatives via multi-sample mixing for contrastive learning. In
 678 *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pp.
 679 2461–2470, 2022b.

680 Jiawen Zhu, Choubo Ding, Yu Tian, and Guansong Pang. Anomaly heterogeneity learning for
 681 open-set supervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer*
 682 *vision and pattern recognition*, pp. 17616–17626, 2024.

683
 684
 685
 686
 687
 688
 689
 690
 691
 692
 693
 694
 695
 696
 697
 698
 699
 700
 701

702 A STATEMENT ON THE USE OF LARGE LANGUAGE MODELS
703704 During the preparation of this manuscript, we used the Large Language Model (LLM) to polish the
705 language and correct grammatical errors to improve readability. The LLM was not involved in any
706 core research aspects of the paper, such as research ideation, experimental design, or analysis of
707 results.709 B ALGORITHM
710711 The complete training algorithm of our proposed **DEMO** framework is outlined in Algorithm 1. We
712 divide the training process into three successive parts. Each part corresponds to a core component of
713 **DEMO**, including mixup anomalies generation, adaptive weighting, and pseudo-label generation.
714715 **Algorithm 1** Training Procedure of the Proposed **DEMO** Framework

716 **Require:** Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, trianing set $\mathcal{V}^{\text{train}}$, training anomalies $\mathcal{V}_a^{\text{train}}$, unlabeled set \mathcal{V}_u ,
717 training epochs T

718 **Ensure:** Trained model parameters θ

719 1: Initialize model parameters θ

720 2: Initialize memory bank $\mathcal{M} \leftarrow \emptyset$

721 3: **for** t in $1, 2, \dots, T$ **do**

722 4: Part I: Multi-sample Mixup for Unseen Anomaly Simulation

723 5: **for all** anomaly node $v_i^{\text{train}} \in \mathcal{V}_a^{\text{train}}$ **do**

724 6: Generate mixup embeddings $\hat{z}_i = \sum_j \alpha_{ij} z_j$ using similarity-based weights α_{ij} (Eq. 1)

725 7: Obtain augmented views \hat{z}_i^a, \hat{z}_i^b via node augmentation

726 8: Calculate consistency loss $\mathcal{L}_{\text{cons}}$ between views

727 9: **end for**

728 10: Part II: Energy Gradient-Based Weight Adaptation

729 11: **for all** training node $v_i^{\text{train}} \in \mathcal{V}^{\text{train}}$ **do**

730 12: Compute energy $E_\theta(v_i)$ and influence score $T_{\text{val}}(v_i)$

731 13: Derive adaptive weight β_{v_i} based on energy gradient (Eq. 7)

732 14: **end for**

733 15: Part III: Pseudo-Labeling Generation

734 16: Update memory bank \mathcal{M} with selection history from epoch t

735 17: **for all** unlabeled node $v_j \in \mathcal{V}_u$ **do**

736 18: Predict probability $\hat{p}_t(v_j)$ for each class

737 19: Compute class-specific adaptive thresholds τ_t^+ and τ_t^- (Eq. 10)

738 20: **if** $\hat{p}_t(v_j) \geq \tau_t^+$ or $\hat{p}_t(v_j) \leq \tau_t^-$ **then**

739 21: Assign pseudo-label \hat{y}_j to v_j based on threshold rule

740 22: Add (v_j, \hat{y}_j) to \mathcal{V}_l

741 23: **end if**

742 24: **end for**

743 25: Compute total loss in Eq. 11

744 26: Update θ using optimizer

745 27: **end for**

746 28: **return** Final model parameters θ

748 C PROOF OF THEOREM
749750 Since more similar samples are inherently more challenging for the model to distinguish, leveraging
751 such samples can enhance the robustness of the model. Therefore, we desire that the synthesized
752 samples generated through the mixup operation retain similarity to the original seen anomaly samples,
753 yet differ sufficiently to approximate the distribution of unseen anomalies. Specifically, by using
754 similarity as weights in the mixup process, we aim to achieve:
755

$$\mathcal{S}(\hat{z}_i, z_j) \approx \mathcal{S}(z_i, z_j), \quad (12)$$

756 where \mathcal{S} denotes the inner-product similarity function. Thus, the synthesized samples maintain similar
 757 levels of similarity to original samples after mixup. Note that both z_i and z_j are anomaly training
 758 samples (z_i^{train} and z_j^{train}), abbreviated here for clarity.
 759

760 **Proof of Theorem 3.1:** The similarity between the synthesized sample \hat{z}_i and an original sample z_j
 761 can be expressed as:

$$762 \quad 763 \quad 764 \quad \mathcal{S}(\hat{z}_i, z_j) = \left(\sum_{k=1}^N \alpha_{ik} z_k \right)^\top z_j = \sum_{k=1}^N \alpha_{ik} \mathcal{S}(z_k, z_j). \quad (13)$$

765 We separate the contribution from z_j and other samples:

$$766 \quad 767 \quad 768 \quad \mathcal{S}(\hat{z}_i, z_j) = \alpha_{ij} \mathcal{S}(z_j, z_j) + \sum_{k \neq j} \alpha_{ik} \mathcal{S}(z_k, z_j) \quad (14)$$

770 Since the weights α_{ik} are obtained via a softmax function, it follows that if $\mathcal{S}(z_i, z_j) \geq \mathcal{S}(z_i, z_k)$,
 771 then $\alpha_{ij} \geq \alpha_{ik}$. Thus, the weights concentrate on highly similar samples. Assuming the similarity
 772 between z_i and z_j is the highest among all other samples, it follows that α_{ij} significantly exceeds the
 773 remaining weights.

774 Further, observing that $\mathcal{S}(z_j, z_j) = \|z_j\|^2$, and letting $\mathcal{S}(z_k, z_j) = \mathcal{S}(z_i, z_j) + \delta_{kj}$, where $\delta_{kj} =$
 775 $\mathcal{S}(z_k, z_j) - \mathcal{S}(z_i, z_j)$, we substitute into the above equation:

$$776 \quad 777 \quad 778 \quad \mathcal{S}(\hat{z}_i, z_j) = \alpha_{ij} \|z_j\|^2 + \sum_{k \neq j} \alpha_{ik} [\mathcal{S}(z_i, z_j) + \delta_{kj}]. \quad (15)$$

779 Rearranging terms, we have:

$$780 \quad 781 \quad 782 \quad \mathcal{S}(\hat{z}_i, z_j) = \mathcal{S}(z_i, z_j) \cdot \sum_{k \neq j} \alpha_{ik} + \alpha_{ij} \|z_j\|^2 + \sum_{k \neq j} \alpha_{ik} \delta_{kj} \quad (16)$$

783 Given that $\sum_{k=1}^N \alpha_{ik} = 1$, this simplifies to:

$$784 \quad 785 \quad 786 \quad \mathcal{S}(\hat{z}_i, z_j) = \mathcal{S}(z_i, z_j) \cdot (1 - \alpha_{ij}) + \alpha_{ij} \|z_j\|^2 + \sum_{k \neq j} \alpha_{ik} \delta_{kj}. \quad (17)$$

787 Since δ_{kj} can be positive or negative, we cannot directly determine the sign of the summation
 788 $\sum_{k \neq j} \alpha_{ik} \delta_{kj}$. However, we can bound its magnitude by applying the triangle inequality:

$$789 \quad 790 \quad 791 \quad \sum_{k \neq j} \alpha_{ik} \delta_{kj} \geq - \sum_{k \neq j} \alpha_{ik} |\delta_{kj}|. \quad (18)$$

792 Thus, we have:

$$793 \quad 794 \quad 795 \quad \mathcal{S}(\hat{z}_i, z_j) \geq \mathcal{S}(z_i, z_j) \cdot (1 - \alpha_{ij}) + \alpha_{ij} \|z_j\|^2 - \sum_{k \neq j} \alpha_{ik} |\delta_{kj}|. \quad (19)$$

796 Letting $\epsilon = \sum_{k \neq j} \alpha_{ik} |\delta_{kj}|$, we rewrite this as:

$$797 \quad 798 \quad 799 \quad \mathcal{S}(\hat{z}_i, z_j) \geq \mathcal{S}(z_i, z_j) \cdot (1 - \alpha_{ij}) + \alpha_{ij} \|z_j\|^2 - \epsilon. \quad (20)$$

800 Despite z_j being highly similar to z_i , it is evident that the similarity between identical samples is
 801 greater than between different samples, thus, $\alpha_{ij} \|z_j\|^2 \geq \alpha_{ij} \mathcal{S}(z_i, z_j)$. Consequently, we have:

$$802 \quad 803 \quad 804 \quad \mathcal{S}(\hat{z}_i, z_j) \geq \mathcal{S}(z_i, z_j) \cdot (1 - \alpha_{ij}) + \alpha_{ij} \mathcal{S}(z_i, z_j) - \epsilon = \mathcal{S}(z_i, z_j) - \epsilon. \quad (21)$$

805 **Interpretation of ϵ :** The value ϵ quantifies the cumulative disturbance caused by non-target samples
 806 ($k \neq j$) in the synthesized sample's similarity. It is influenced by two factors: 1) Concentration of the
 807 weights α_{ik} (greater concentration implies smaller ϵ). 2) Similarity differences $|\delta_{kj}|$ among samples
 808 (smaller differences imply smaller ϵ).

810
 811 Table 4: Statistics of 6 datasets including the number of nodes, edges, and dimensions, the number of
 812 classes, the number of anomalous classes, the number of anomalous nodes, and the selection interval
 813 for anomalous nodes.

Dataset	#Nodes	#Edge	#Dimension	#Classes	#Ano. Classes	# Ano.	#Ano. Prop
Photo	7,650	238,162	745	8	2	700	0%-5%
Computers	13,752	491,722	767	10	5	2,064	0%-5%
CS	18,333	163,788	6,805	15	8	4,159	0%-5%
Yelp	45,954	3,846,979	32	3	2	6,677	0%-5%
ogbn-arxiv	169,343	1,166,243	128	40	4	27,830	3%-5%
ogbn-mag	736,389	5,416,271	128	349	15	3,135	0%-0.03%

D DATASETS DETAILS

825 We evaluate the performance of **DEMO** on open-set graph anomaly detection using six real-world graph
 826 datasets. Notably, there are currently no dedicated benchmark datasets tailored for open-set GAD
 827 tasks. Therefore, we adapt existing datasets through customized processing to simulate the open-set
 828 scenario. Specifically, we conduct experiments on three small-scale datasets (Photo, Computers, and
 829 CS) and three large-scale datasets (Yelp, ogbn-arxiv, and ogbn-mag). Table 4 summarizes the dataset
 830 processing details, where “Ano. Classes” denotes the number of anomalous classes, “Ano” indicates
 831 the total number of anomalous nodes, and “Ano. Prop” reflects the proportion range of node count
 832 in each class used to determine whether it is selected as an anomaly class. To emulate the open-set
 833 setting, we designate classes with low sample counts as anomaly classes, creating a mismatch between
 834 the class distribution of the training and testing phases. This better reflects real-world scenarios
 835 where unknown or unseen anomaly types may emerge at test time, posing a significant challenge for
 836 generalization. Additional descriptions of each dataset are provided below.

- 837 • **Photo and Computers** (McAuley et al., 2015). Amazon Computers and Amazon Photo are subsets
 838 of the Amazon co-purchase graph, where nodes represent products, and edges indicate co-purchase
 839 relationships. The node features are encoded as bag-of-words from product reviews, and class
 840 labels correspond to the product category.
- 841 • **CS**. The Computer Science (CS) section of the Coauthor dataset is used for node classification
 842 tasks. In this dataset, nodes represent authors connected by an edge if they have co-authored a paper.
 843 The node features capture the keywords of the authors’ papers, while the class labels represent the
 844 most active research fields of each author.
- 845 • **Yelp** (Zeng et al., 2019). The task of Yelp dataset is categorizing types of businesses based on
 846 customer reviewers and friendship. The Yelp dataset is a heterogeneous graph with three distinct
 847 views. For our experiments, we focus on the “Review-User-Review (RUR)” edge subset.
- 848 • **ogbn-arxiv** (Mikolov et al., 2013). The ogbn-arxiv dataset is a citation network of Computer
 849 Science (CS) arXiv papers, where each node represents a paper and each directed edge indicates a
 850 citation. The task is to predict the 40 subject areas of these papers, using a 40-class classification
 851 approach, with a realistic data split based on publication dates for training, validation, and testing.
- 852 • **ogbn-mag** (Wang et al., 2020). The ogbn-mag dataset is a heterogeneous network derived from the
 853 Microsoft Academic Graph (MAG), consisting of four types of entities: papers, authors, institutions,
 854 and fields of study. The task is to predict the venue (conference or journal) of each paper, given
 855 its content, references, authors, and affiliations, which is formulated as a 349-class classification
 856 problem. For our experiments, we focus on the “paper” nodes, using the (“paper”, “citeps”, “paper”)
 857 edges to represent citations between papers.

E BASELINES DETAILS

859 In our experiments, we compare **DEMO** with 12 representative GAD methods, which are categorized
 860 into two groups: unsupervised and semi-supervised approaches. Notably, TAM, OCGNN, ANOMA-
 861 LOUS, and DOMINANT are originally unsupervised models; to incorporate supervision for fair
 862 comparison, we adopt the semi-supervised design strategy proposed in GGAD. Specifically, for

864 TAM, we refine its affinity maximization objective to focus exclusively on labeled normal nodes. In
 865 OCGNN, the one-class center optimization is constrained to labeled normal instances. For DOMI-
 866 NANT and AnomalyDAE, we restrict their auto-encoding loss computation to labeled normal nodes
 867 during training. Below, we provide brief descriptions of these 12 methods.
 868

- 869 • **ANOMALOUS** (Peng et al., 2018). ANOMALOUS proposes a joint framework for anomaly
 870 detection on attributed networks that integrates attribute selection and anomaly detection using CUR
 871 decomposition and residual analysis. It filters out noisy and irrelevant node attributes, improving
 872 detection performance by focusing on the most representative attributes.
- 873 • **DOMINANT** (Ding et al., 2019). DOMINANT proposes a deep learning model that combines
 874 Graph Convolutional Networks (GCN) and autoencoders for anomaly detection in attributed
 875 networks, addressing challenges like sparsity and nonlinearity by measuring reconstruction errors
 876 from both structural and attribute perspectives.
- 877 • **AnomalyDAE** (Fan et al., 2020). AnomalyDAE incorporates an attention mechanism to capture
 878 structure patterns and uses both node and attribute embeddings to model cross-modality interactions
 879 during reconstruction, enabling effective anomaly detection.
- 880 • **GAAN** (Chen et al., 2020b). GAAN generates fake graph nodes using Gaussian noise and employs
 881 an encoder to map both real and fake nodes into a latent space, using a discriminator to distinguish
 882 between real and fake nodes. Anomaly detection is then performed by evaluating reconstruction
 883 errors and identification confidence.
- 884 • **CoLA** (Liu et al., 2021). CoLA uses a contrastive self-supervised learning framework for anomaly
 885 detection in attributed networks, which captures local information by sampling contrastive instance
 886 pairs and utilizing a GNN-based contrastive learning model. The framework targets anomaly
 887 detection through a specific learning objective and adapts to large networks by training on batches
 888 of instance pairs rather than the full graph.
- 889 • **CONAD** (Xu et al., 2022). CONAD introduces a framework that integrates human knowledge
 890 of different anomaly types into attributed network anomaly detection. It employs a novel data
 891 augmentation strategy to model prior human knowledge, which is then incorporated into a Siamese
 892 graph neural network encoder with contrastive loss. Anomalies are detected by ranking nodes
 893 based on their reconstruction error.
- 894 • **ConsisGAD** (Chen et al., 2024). ConsisGAD is a model designed for graph anomaly detection
 895 in settings with limited supervision and class imbalance. It leverages unlabeled data through
 896 consistency training and a learnable data augmentation mechanism, while utilizing homophily
 897 distribution variance to improve class distinction using a simplified GNN backbone.
- 898 • **GGAD** (Qiao et al., 2024). GGAD introduces a semi-supervised generative approach for graph
 899 anomaly detection, aiming to better utilize known normal nodes by generating pseudo anomaly
 900 nodes (“outlier nodes”) for training a one-class classifier. It leverages priors on asymmetric local
 901 affinity and egocentric closeness to generate reliable outlier nodes, thus enhancing detection
 902 performance in the absence of ground truth for real anomalies.
- 903 • **TAM** (Qiao & Pang, 2023). TAM introduces a novel unsupervised anomaly scoring measure
 904 based on local node affinity, where normal nodes exhibit stronger connections with each other than
 905 abnormal nodes. It employs Truncated Affinity Maximization (TAM) to learn node representations
 906 that maximize local affinity, optimizing on truncated graphs to remove non-homophily edges and
 907 improve anomaly detection performance.
- 908 • **OCGNN** (Wang et al., 2021). OCGNN introduces a one-class classification framework for graph
 909 anomaly detection, combining the powerful representation ability of graph neural networks with
 910 the classical one-class support vector machine objective. This approach addresses the limitations
 911 of traditional anomaly detection methods in graph data, achieving significant improvements over
 912 existing baselines.
- 913 • **SpaceGNN** (Dong et al., 2025a) SpaceGNN introduces a model for node anomaly detection with
 914 limited labels, using a Learnable Space Projection function and a Distance Aware Propagation
 915 module to enhance node representations and information propagation. It outperforms data augmentation
 916 techniques and achieves better results than existing methods.
- 917 • **NSReg** (Wang et al., 2023b). NSReg introduces a novel approach for open-set graph anomaly
 918 detection by adding a regularization term to enforce compact, semantically-rich representations of

918
 919 Table 5: AUC-ROC and AUC-PR on the unseen anomaly classes on three small-scale datasets. The
 920 best performance is boldfaced, with the second-best underlined.

Datasets Metrics	Photo		Computers		CS	
	AUC-ROC	AUC-PR	AUC-ROC	AUC-PR	AUC-ROC	AUC-PR
Anomalous	0.4841 ± 0.015	0.0435 ± 0.010	0.5640 ± 0.075	0.1435 ± 0.008	0.2364 ± 0.018	0.1323 ± 0.004
DOMINANT	0.4721 ± 0.021	0.0548 ± 0.004	0.5597 ± 0.078	0.1474 ± 0.003	0.3836 ± 0.026	0.1601 ± 0.003
AnomalyDAE	0.4798 ± 0.019	0.0495 ± 0.007	0.5793 ± 0.045	0.1530 ± 0.006	0.3724 ± 0.019	0.1566 ± 0.006
GAAN	0.4708 ± 0.031	0.0432 ± 0.002	0.5739 ± 0.020	0.1620 ± 0.012	0.4636 ± 0.015	0.1845 ± 0.009
CoLA	0.7064 ± 0.024	0.0808 ± 0.006	0.4565 ± 0.018	0.1149 ± 0.010	0.4411 ± 0.029	0.1792 ± 0.005
CONAD	0.4746 ± 0.024	0.0564 ± 0.011	0.5560 ± 0.027	0.1436 ± 0.011	0.3838 ± 0.035	0.1602 ± 0.010
CosisGAD	0.5571 ± 0.059	0.0563 ± 0.012	0.5057 ± 0.026	0.1325 ± 0.019	0.7364 ± 0.026	0.4594 ± 0.018
GGAD	0.6274 ± 0.043	0.0679 ± 0.004	0.6636 ± 0.097	0.2471 ± 0.054	0.8931 ± 0.059	0.7596 ± 0.039
TAM	0.5979 ± 0.027	0.0604 ± 0.003	0.4458 ± 0.049	0.1159 ± 0.007	0.6692 ± 0.018	0.3396 ± 0.023
OCGNN	0.5208 ± 0.015	0.0563 ± 0.004	0.5112 ± 0.013	0.1308 ± 0.005	0.7662 ± 0.056	0.4406 ± 0.027
ANO-S	0.5315 ± 0.032	0.0573 ± 0.007	0.4508 ± 0.061	0.1168 ± 0.002	0.8181 ± 0.082	0.5727 ± 0.052
DOM-S	0.5705 ± 0.003	0.0589 ± 0.001	0.4424 ± 0.046	0.1118 ± 0.008	0.8239 ± 0.058	0.5645 ± 0.033
SpaceGNN	0.6808 ± 0.031	0.0733 ± 0.007	0.8108 ± 0.042	0.5010 ± 0.025	0.8145 ± 0.065	0.6337 ± 0.041
NSReg	0.7369 ± 0.014	0.0956 ± 0.006	0.6846 ± 0.013	0.3547 ± 0.036	0.8863 ± 0.079	0.7460 ± 0.027
DEMO	0.8202 ± 0.019	0.1303 ± 0.003	0.8119 ± 0.036	0.5032 ± 0.015	0.9354 ± 0.039	0.8417 ± 0.017

937
 938 normal nodes. This helps the model generalize to unseen anomalies while reducing false negatives
 939 and improving detection performance.

F IMPLEMENTATION DETAILS

F.1 EXPERIMENTAL PROTOCOL

946 In our open-set graph anomaly detection (GAD) experiments, we follow a process that alternates
 947 between treating different anomaly classes as “seen” anomalies and considering the remaining
 948 anomaly classes as “unseen” anomalies. For each dataset, we first select a class to be treated as
 949 seen and train the model using this class along with normal nodes. The training set consists of
 950 the normal nodes and the selected anomaly class. Next, we evaluate the model’s performance by
 951 measuring AUC-ROC and AUC-PR scores on all anomaly classes (results in Section 4.2) and only on
 952 the unseen anomaly classes (results can be found in Section G). Additionally, for smaller datasets
 953 (Photo, Computers, CS), we repeat this process five times and record the average and variance. For
 954 larger datasets (Yelp, ogbn-arxiv, ogbn-mag), we conduct a single experiment due to the large volume
 955 of data. We use the same data split for training the baselines to ensure fairness.

F.2 HYPERPARAMETER SETTING

956 In addition to the implementation details described in Section 4.1, we made the following hyperpa-
 957 rameter settings. According to (Wang et al., 2023b), we set the number of neighbors for GraphSAGE
 958 aggregation to 25 for the first layer and 10 for the second layer, except for the ogbn-mag dataset,
 959 to improve computational efficiency. We configured the GraphSAGE model with 2 layers, set the
 960 dropout rate to 0.5, and the batch size to 512. For the reproduction of unsupervised methods, we
 961 referred to the [code](#) and would like to thank the authors for making their implementation publicly
 962 available.

G MORE EXPERIMENTS

G.1 GAD PERFORMANCE EXTENSION EXPERIMENT

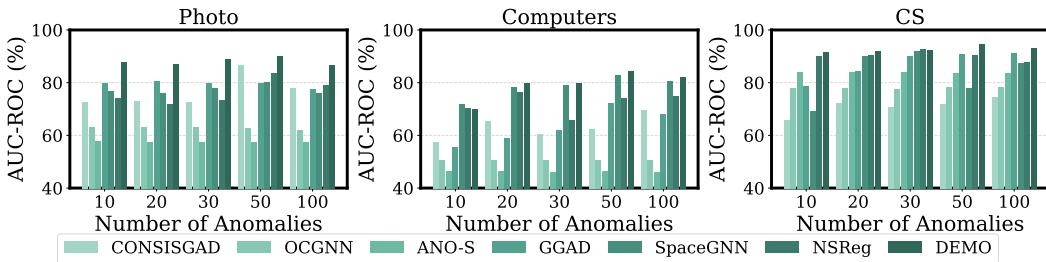
970 We evaluate the detection capabilities of all methods on unseen anomaly classes. Specifically, Tables 5
 971 and 6 present the performance of each algorithm on three small-scale datasets (Photo, Computers, and
 972 CS) and three large-scale datasets (Yelp, ogbn-arxiv, and ogbn-mag), respectively, where seen anomaly

972

973
974
Table 6: AUC-ROC and AUC-PR on the unseen anomaly classes on three large-scale datasets. The
best performance is boldfaced, with the second-best underlined.

Datasets	Metrics	ConsisGAD	GGAD	TAM	OCGNN	ANO-S	DOM-S	SpaceGNN	NSReg	DEMO
Yelp	AUC-ROC	0.5238	0.6819	0.6016	<u>0.7185</u>	0.7168	0.7056	0.5212	0.4941	0.7235
	AUC-PR	0.0205	0.0776	0.0351	0.0625	0.0538	0.0381	0.0220	0.0178	<u>0.0635</u>
ogbn-arxiv	AUC-ROC	0.5254	0.5101	OOM	OOM	0.4462	0.4474	0.5180	<u>0.5356</u>	0.5643
	AUC-PR	<u>0.1655</u>	0.1473	OOM	OOM	0.1187	0.1204	0.1611	0.1566	0.1808
ogbn-mag	AUC-ROC	<u>0.4879</u>	OOM	OOM	OOM	OOM	OOM	0.4425	0.4628	0.4995
	AUC-PR	0.0039	OOM	OOM	OOM	OOM	OOM	0.0036	0.0036	0.0053

981



982

983

984

985

986

987

988

989

990

991

992

Figure 4: The AUC-ROC under different numbers of training anomalies across three datasets.

993

994

995

nodes are excluded from the test set. Only normal and unseen anomaly nodes are considered in the computation of AUC-ROC and AUC-PR. Overall, **DEMO** consistently achieves the best performance across both categories of datasets, further confirming its strong generalization ability to unknown anomalies in realistic open-set GAD scenarios. Across all six datasets, **DEMO** not only outperforms all unsupervised baselines but also significantly exceeds other competitive semi-supervised methods, especially in AUC-PR. Specifically, **DEMO** achieves the highest AUC-ROC and AUC-PR scores on all three small-scale datasets. For example, on the Photo dataset, **DEMO** reaches an AUC-ROC of 0.8202, notably surpassing the second-best method NSReg (0.7369), while also achieving an AUC-PR of 0.1303, substantially outperforming all baselines. Similar trends are observed on Computers and CS, indicating that **DEMO** effectively detects diverse and distribution-shifted unseen anomalies. On large-scale datasets, **DEMO** again obtains the best and near-best results on Yelp and ogbn-arxiv, with AUC-ROC scores of 0.7235 and 0.5643, showing significant improvements over all other methods. Even on the most challenging ogbn-mag dataset, characterized by extreme class imbalance and a large graph size, **DEMO** maintains its advantage, achieving an AUC-ROC of 0.4995 and an AUC-PR of 0.0053, demonstrating high sensitivity to low-frequency unseen anomalies. In summary, **DEMO** exhibits strong robustness and generalizability even when seen anomalies are excluded from evaluation, validating its effectiveness in identifying unknown patterns in open-set graph anomaly detection.

1012

1013

G.2 DATA EFFICIENCY EXTENSION EXPERIMENT

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

We investigate the impact of the number of seen anomalous nodes in the training set on model performance. Specifically, we conduct experiments on three small-scale datasets (Photo, Computers, and CS) by varying the number of training anomalies across five levels: 10, 20, 30, 50, and 100. We evaluate the performance of all methods using AUC-ROC (Figure 4) and AUC-PR (Figure 5) to analyze their adaptability and robustness under both low-resource and high-resource settings. Overall, **DEMO** consistently outperforms all baselines across different settings, particularly exhibiting a clear advantage under low-resource scenarios (*e.g.*, with only 10 or 20 anomalies), indicating its strong training efficiency and generalization ability in anomaly-scarce situations. More specifically, on the Photo dataset, **DEMO** achieves the best results on both AUC-ROC and AUC-PR in most configurations, with stable performance especially at 20 and 30 anomalies. On the Computers dataset, performance improves steadily as more anomalies are added. On the CS dataset, **DEMO** demonstrates high overall stability, maintaining the best performance even with as few as 10 anomalies, and consistently achieving top scores across different anomaly scales. In conclusion, **DEMO** demonstrates stable and

1026

1027

Table 7: Loss weights of six datasets

Datasets	λ_{un}	λ_{mix}	λ_{eng}
Photo	0.5	0.1	0.1
Computers	0.1	0.1	0.5
CS	0.3	0.1	0.3
Yelp	0.3	0.1	0.1
ogbn-arxiv	0.5	0.1	0.1
ogbn-mag	0.5	0.1	0.1

1035

1036

1037

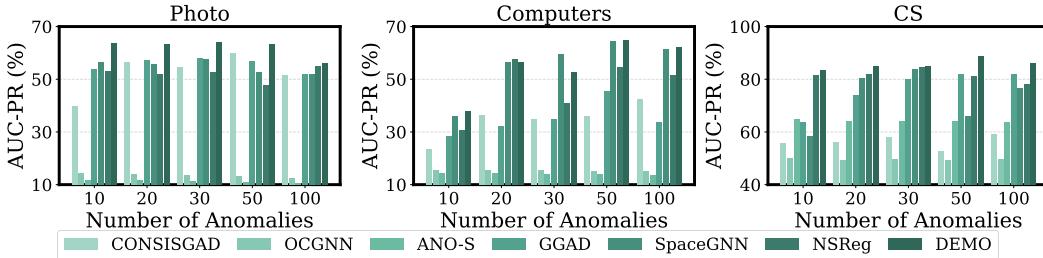
Table 8: Ablation Study on the other three benchmarks.

Metrics	CS		Yelp		ogbn-arxiv	
	AR	AP	AR	AP	AR	AP
w/o All	0.7815	0.6394	0.5942	0.1063	0.4219	0.1944
w/o Mix	0.8841	0.8151	0.6519	0.1927	0.5061	0.2455
w/o EG	0.9027	0.7964	0.6637	0.1833	0.5337	0.2849
w/o PL	0.8364	0.8005	0.6362	0.1624	0.4868	0.2517
DEMO	0.9448	0.8857	0.7097	0.2238	0.6364	0.3329

1046

1047

1048 superior detection capability across different levels of anomaly supervision, with its exceptional
1049 performance under limited anomaly samples further validating its data efficiency and robustness in
1050 open-set GAD.



1061

1062

Figure 5: The AUC-PR under different numbers of training anomalies across three datasets.

1063

1064

1065 G.3 MORE ABLATION STUDY

1066

1067

1068 To further validate our method, we present the ablation study results on three additional benchmarks in
1069 Table 8. The results consistently corroborate the conclusions from our main analysis. First, removing
1070 the PL component results in a substantial performance drop across all three datasets, particularly on
1071 CS and ogbn-arxiv, demonstrating the critical impact of our class-aware pseudo-labeling in complex
1072 scenarios. Second, the exclusion of the EG and Mix modules also leads to a clear degradation
1073 in performance, validating their respective effectiveness in prioritizing informative samples and
1074 synthesizing diverse anomaly representations. Additionally, the variant DEMO w/o All consistently
1075 yields the lowest performance, confirming that the full combination of modules is essential for
1076 achieving optimal results. Overall, these extensive results provide further evidence that all three
1077 components contribute synergistically to the robustness and superior performance of DEMO.

1078

1079

1077 G.4 PARAMETERS SENSITIVITY

1079 We conduct a sensitivity analysis of the loss weights on two datasets (Photo and Computers) to evaluate
1080 the contribution and robustness of each component to both overall performance and unseen anomaly

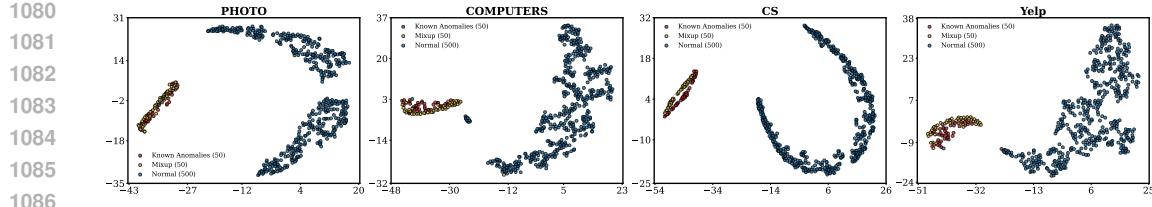


Figure 6: Visualization of sample embeddings after Mixup

detection. Specifically, we vary the loss weight $\lambda_{un} \in \{0.1, 0.3, 0.5, 0.7\}$ and $\lambda_{eng} \in \{0.1, 0.3, 0.5\}$. The results are shown in Figure 7 and 8. Notably, although multi-sample mixup enhances anomaly representation diversity, an overly large mixup ratio may introduce distributional shifts or generate overly blurred features, thereby disrupting the learning of original decision boundaries. This issue becomes more pronounced when labeled anomalies are scarce, often leading to unstable training. To mitigate this, we treat mixup as a lightweight augmentation strategy to provide auxiliary distributional support while minimizing structural disruption, and fix its weight at 0.1. As shown in the Figure, on the Photo dataset, the model achieves optimal performance when the $\lambda_{un} = 0.5$ and the $\lambda_{eng} = 0.1$. For the Computers dataset, the best results are observed when $\lambda_{un} = 0.1$ and $\lambda_{eng} = 0.5$. In summary, the optimal loss weight configuration varies across datasets, which can be attributed to differences in graph structure complexity, anomaly separability, and feature distribution. For example, datasets with more entangled representations or higher intra-class variability (*e.g.*, Computers) may benefit from stronger sample reweighting to focus on informative nodes. Finally, Table 7 shows the loss weight values across the different datasets.

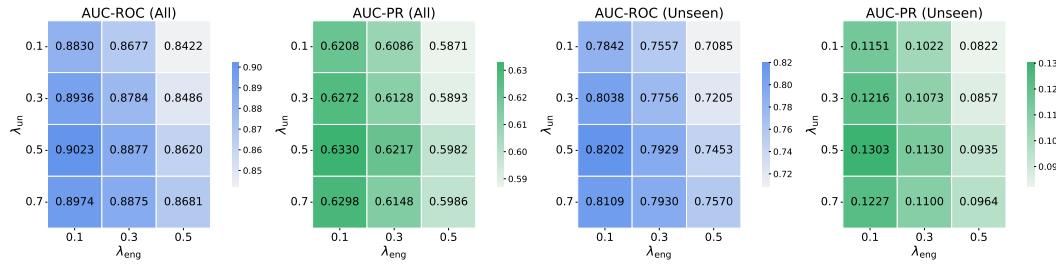


Figure 7: Performance impact of loss weights on the Photo dataset.

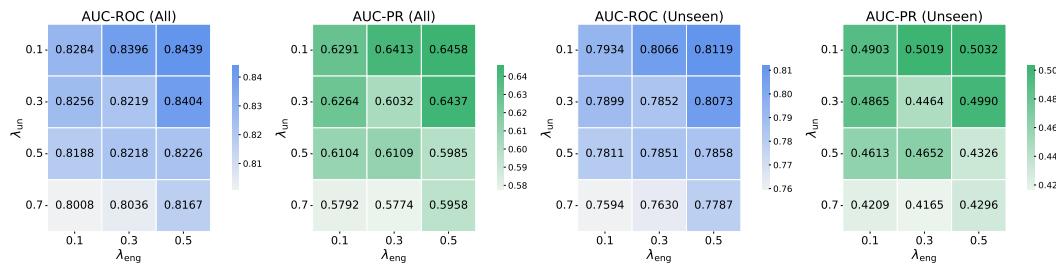


Figure 8: Performance impact of loss weights on the Computers dataset.

G.5 VISUALIZATION OF MIXUP EMBEDDINGS

To investigate the distribution of synthesized anomalies relative to both seen anomalies and normal data, we conducted an extended t-SNE visualization experiment across datasets with varying characteristics (Photo, Computers, CS, and Yelp). We projected the embeddings of normal nodes, original ‘Seen Anomaly’ samples, and the synthetic ‘Mixup Anomaly’ samples generated by our

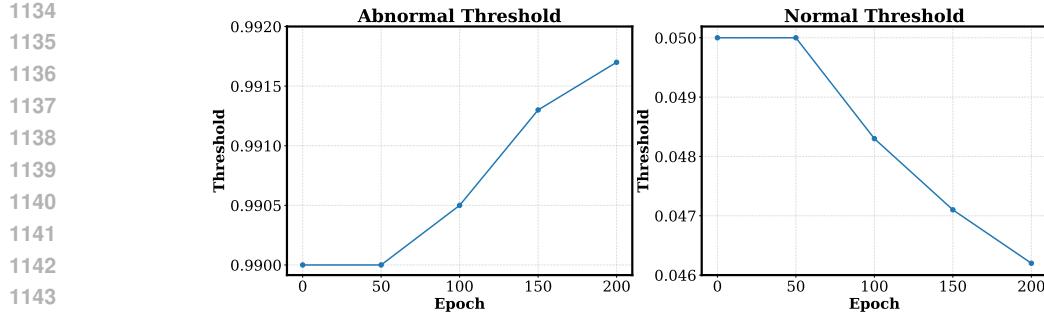


Figure 9: Threshold variation over epochs.

method. As illustrated in Figure 6, the visualization reveals two key observations. First, the ‘Mixup Anomaly’ embeddings (yellow) do not merely overlap with the ‘Seen Anomaly’ cluster (red) but form a broader, diffused distribution that expands outward. This confirms that our strategy effectively explores the potential feature space beyond the limited training samples. Second, and crucially, across all datasets—including those with large scale (e.g., Yelp), there remains a clear and distinct separation between the synthesized Mixup anomalies and the normal node clusters. This empirical evidence demonstrates that **DEMO** successfully expands the anomaly decision boundary to enhance generalization without encroaching upon the normal data distribution, thereby ensuring both the diversity of anomaly representations and the stability of model training.

G.6 VISUALIZATION OF DYNAMIC PSEUDO-LABELING THRESHOLDS

To depict the dynamic behavior of our reliable pseudo-labeling mechanism, we experimented to visualize the evolution of the class-specific thresholds. We tracked the anomaly-class threshold (τ_t^+) and the normal-class threshold (τ_t^-) throughout the training process on the Photo datasets. The results in Figure 9 show a clear “diverging” trend: the anomaly threshold τ_t^+ progressively increases towards 1.0, while the normal threshold τ_t^- progressively decreases towards 0.0. This visualization empirically confirms our claim in Section 3.4. It demonstrates that our memory bank-guided strategy effectively adapts as the model’s predictions become more confident, ensuring the pseudo-labeling process becomes increasingly stringent.