

From Subtle to Significant: Prompt-Driven Self-Improving Optimization in Test-Time Graph OOD Detection

Luzhi Wang¹, Xuanshuo Fu², He Zhang³, Chuang Liu¹, Xiaobao Wang⁴, Hongbo Liu^{1*}

¹College of Artificial Intelligence, Dalian Maritime University

²Computer Vision Center, Universitat Autònoma de Barcelona

³School of Computing Technologies, RMIT University

⁴College of Intelligence and Computing, Tianjin University

wangluzhi0@gmail.com, xuanshuo@cvc.uab.es, he.zhang@rmit.edu.au,

chuangliu@whu.edu.cn, wangxiaobao@tju.edu.cn, lhb@dlmu.edu.cn

Abstract

Graph Out-of-Distribution (OOD) detection aims to identify whether a test graph deviates from the distribution of graphs observed during training, which is critical for ensuring the reliability of Graph Neural Networks (GNNs) when deployed in open-world scenarios. Recent advances in graph OOD detection have focused on test-time training techniques that facilitate OOD detection without accessing potential supervisory information (e.g., training data). However, most of these methods employ a one-pass inference paradigm, which prevents them from progressively correcting erroneous predictions to amplify OOD signals. To this end, we propose a **Self-Improving Graph Out-of-Distribution** detector (SIGOOD), which is an unsupervised framework that integrates continuous self-learning with test-time training for effective graph OOD detection. Specifically, SIGOOD generates a prompt to construct a prompt-enhanced graph that amplifies potential OOD signals. To optimize prompts, SIGOOD introduces an Energy Preference Optimization (EPO) loss, which leverages energy variations between the original test graph and the prompt-enhanced graph. By iteratively optimizing the prompt by involving it into the detection model in a self-improving loop, the resulting optimal prompt-enhanced graph is ultimately used for OOD detection. Comprehensive evaluations on 21 real-world datasets confirm the effectiveness and out-performance of our SIGOOD method.

Introduction

Graph neural networks (GNNs) offer a powerful paradigm for graph representation learning (Wang et al. 2025b; Fu et al. 2025; Zhang et al. 2025a), which is widely applied in tasks such as sarcasm detection (Wang et al. 2025a, 2023), recommender systems (Jin et al. 2023a), and fraud detection (Pan et al. 2025). Most GNNs assume that training and test graphs are from the same distribution (in-distribution, ID) (Zhang et al. 2024a). However, when these well-trained GNNs are deployed in open-world scenarios, they inevitably encounter out-of-distribution (OOD) graphs (Shen et al. 2024), leading to misprediction risks. For instance, a GNN might misclassify a structurally similar yet distributionally

*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

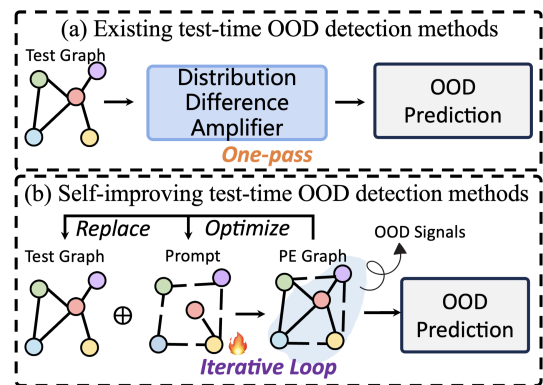


Figure 1: Comparisons between existing methods and SIGOOD. (a) Existing methods perform a one-pass pattern extraction, aiming to amplify the difference between ID and OOD graphs (Wang et al. 2024). (b) SIGOOD adopts an unsupervised self-improvement strategy, progressively refining OOD signals through iterative optimization.

distinct five-membered lactone molecule as aspirin (a typical drug molecule), resulting in prediction errors in drug identification (Wang et al. 2025c). Thus, detecting OOD graphs is essential for ensuring the reliability of GNNs.

Recent advances in graph OOD detection have attracted increasing attention (Lin et al. 2025). One line of these studies considers that pattern discrepancies exist between ID and OOD graphs (Ding et al. 2024). Consequently, these methods rely on training data to learn ID graph patterns for OOD discrimination (Hou et al. 2025), but such data may be unavailable due to privacy, storage, or deployment constraints (Wang et al. 2024). Therefore, the second line of studies has focused on test-time graph OOD detection, which aims to identify distribution shifts using only the well-trained GNN and test graphs (Jin et al. 2023b). However, these methods are suboptimal when amplifying the difference of ID and OOD core information across graphs, as their one-pass unsupervised rationales neglect to leverage the latent information in test graphs through iterative refinement.

The above limitations motivate us to develop a new test-

time graph OOD detection method, but it is not trivial due to the following challenges. Particularly, *Challenge 1*: Distribution overlapping. A well-trained GNN primarily learns the decision boundaries from ID graphs, without explicitly acquiring the ability to recognize unknown patterns. Thus, in the embedding space, the GNN tends to project OOD graph embeddings blindly toward the ID graph embeddings, leading to the distributional overlap between ID and OOD data (Wu et al. 2024). This phenomenon makes it difficult to identify OOD graphs that have subtle structural or feature differences from ID graphs. *Challenge 2*: Unknown ID/OOD labels. Under the test-time OOD detection setting, the detection model has no access to task labels or other information (e.g., training data), which raises difficulty in effectively extracting OOD patterns with unlabeled test graphs. *Challenge 3*: Effective optimization strategies. Due to the absence of task labels, it is challenging to design a self-improving optimization strategy that can continuously learn OOD signals from the test graph itself, and effectively enhance the model’s OOD detection performance.

To address aforementioned challenges, we propose **SIGOOD**, an unsupervised self-improving graph OOD detection framework that iteratively enhances OOD signals at test time via energy-based feedback. According to previous studies (Zhang et al. 2024d; Wu et al. 2023), the energy based on prediction logits represents the probability of a graph being OOD, where a high energy score indicates a high likelihood of OOD. Since OOD signals are composed of key nodes that sufficiently represent OOD characteristics, we introduce prompts to amplify OOD signals by capturing the average node-level energy variations before and after prompt enhancement, thus optimizing the objective of OOD detection. Specifically, to address *Challenge 1*, SIGOOD firstly utilizes the well-trained GNN to obtain the embedding of a test graph, and then injects auxiliary prompt into it, aiming to amplify the differences between OOD patterns and ID patterns in it. To tackle *Challenge 2*, SIGOOD utilizes the energy variation between the test graph and the prompt-enhanced (PE) graph as a signal to indicate OOD or ID preference, which actively discover and iteratively refine ID and OOD patterns. To address *Challenge 3*, we propose a tailored loss function that leverages the energy variations induced by prompt injection. This loss amplifies the energy differences between OOD and ID patterns within the PE graph to facilitate more effective OOD detection. Unlike prior one-pass methods, our approach involves the prompt-enhanced graph in an iterative optimization loop, supporting a self-improving process. Fig. 1 illustrates the differences between SIGOOD and other test-time graph OOD detection methods. Our contributions are summarized as follows:

- To the best of our knowledge, we propose the first self-improving framework for test-time graph OOD detection, which progressively enhances the OOD detection performance with only relying on test data.
- We investigate the role of energy variation in capturing OOD signals and propose an energy preference optimization (EPO) loss to enhance the energy contrast between OOD and ID signals.

- Extensive experiments on 21 benchmark datasets and comparisons with 12 state-of-the-art methods demonstrate the superior performance of SIGOOD.

Related Work

Graph OOD Detection

Graph OOD detection aims to identify test graphs whose distributions deviate from the ID graphs (Lin et al. 2024; Zhang et al. 2025b). Existing methods generally use pre-trained encoders with post-hoc detectors, such as distance or energy-based scores, which are static detection methods that lack any training-time or test-time optimization process (Fuchsgruber, Wollschläger, and Günemann 2024). Several studies focus on improving ID recognition rather than detecting OOD graphs (Cao et al. 2025). Recent test-time methods amplify ID-OOD differences but lack iterative refinement mechanisms to progressively enhance OOD signals (Zhang et al. 2024c). In contrast, our SIGOOD model introduces a novel energy-based self-improving mechanism to iteratively mine OOD patterns without label or training supervision.

Test-time Training

Test-Time Training (TTT) is a general approach designed to enhance the performance of predictive models when there is a distribution shift between training and test data (Liang, He, and Tan 2025). In addition to the computer vision (Dalal et al. 2025), natural language processing (Zhang et al. 2024b), and multimodal foundation models (Bi et al. 2025), TTT technology is also widely used to improve graph OOD generalization (Zheng et al. 2024). Nevertheless, the application of TTT techniques to graph OOD detection remains largely underexplored. In our work, we leverage TTT in a self-improving manner to identify OOD signals of graphs, thereby enabling effective test-time OOD detection.

Preliminaries

Graphs. Given a graph $G = (V, E, X)$, the node set is indicated by $V = \{v_1, v_2, \dots, v_n\}$, $E = \{e_1, e_2, \dots, e_m\}$ denotes the set of edges, and $X \in \mathbb{R}^{n \times d}$ represents the d -dimensional node feature matrix. Each graph is associated with an adjacency matrix $A \in \mathbb{R}^{n \times n}$, where $A_{ij} = 1$ indicates the presence of an edge between nodes v_i and v_j , and $A_{ij} = 0$ indicates that no edge exists between v_i and v_j .

Test-time Graph OOD Detection. During the inference stage of a GNN model, the test-time graph out-of-distribution detection task aims to determine whether an input graph is drawn from a distribution different from that of the training data. At test time, neither the original training data nor auxiliary information such as task-specific labels for OOD detection is accessible. Formally, the test-time graph OOD detection can be defined as:

$$\text{Detection label} = \begin{cases} 1 \text{ (OOD)}, & \text{if } \mathcal{D}(G_t) \geq \tau \\ 0 \text{ (ID)}, & \text{otherwise} \end{cases} \quad (1)$$

where $\mathcal{D}(G_t)$ denotes a scoring function that measures the discrepancy between the test graph G_t and the in-distribution data, and τ is a predefined detection threshold.

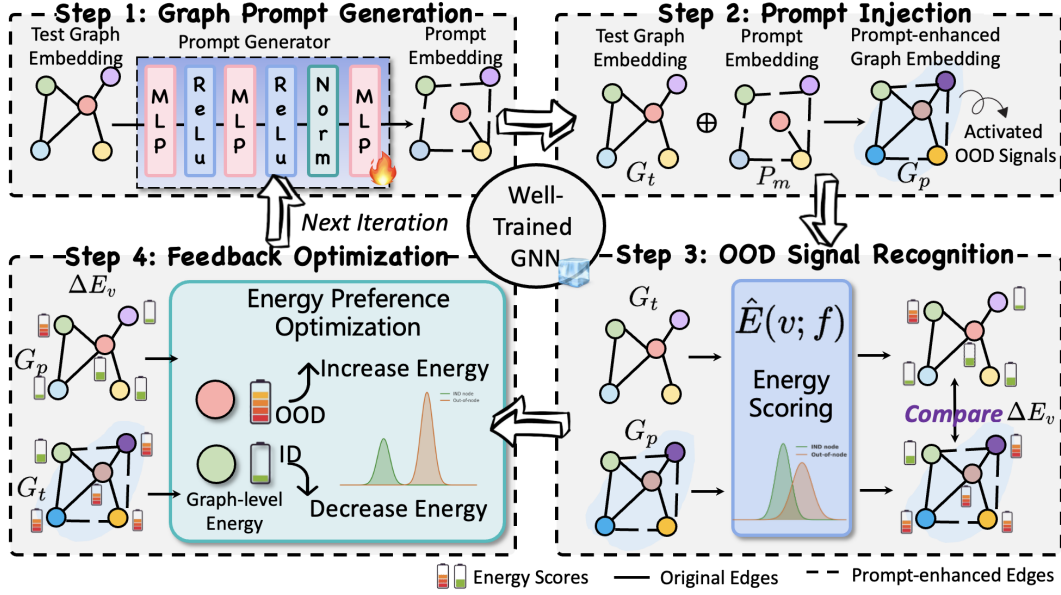


Figure 2: Overview of SIGOOD. **Step 1:** Given a test graph G_t as an input, SIGOOD first encodes it using the well-trained GNN. The obtained embedding is then passed to the prompt generator (PG), which produces a prompt P_m to enhance OOD signals of the graph. **Step 2:** The prompt P_m is applied to G_t , yielding a prompt-enhanced (PE) graph G_p with amplified OOD tendency. **Step 3:** Calculate node-wise energy variations between G_p and G_t to locate sensitive nodes of OOD signals. **Step 4:** Calculate the global energy variations between G_p and G_t to evaluate the overall OOD tendency of G_p . These energy variations are used as OOD signals to guide the optimization of the prompt generator. The updated G_p replaces G_t as the input for the next iteration. After convergence, the final G_p is used to calculate the OOD score for detection.

Method

Overview

According to existing studies (Liu et al. 2020), energy scores (e.g., the negative log-likelihood of sample predictions) reflect model confidence during the inference phase. For example, a higher energy value (i.e., a lower likelihood) indicates lower confidence in model predictions, suggesting a greater probability that the test sample is OOD data (Chen et al. 2023). In this paper, our intuition is that the subtle difference between ID and OOD data can be amplified by introducing additional prompts, thereby increasing the energy-based distinction for OOD detection. Due to the inductive bias of a well-trained GNN on ID graphs (Wu et al. 2020), introducing prompts into a test graph generally leads to a greater energy change in its OOD components (e.g., nodes) than in its ID components. This difference in energy change is then used to iteratively identify potential OOD signals in the test graph and guide the prompt generator toward optimal detection performance.

Fig. 2 illustrates the overview of our SIGOOD method. Given a test graph, it is first embedded by a well-trained GNN, and then an optimizable prompt is utilized to generate a prompt-enhanced graph that amplifies the OOD signal. Under the guidance of energy-based loss function (EPO loss), the prompt generator is optimized to enhance OOD signals and the corresponding PE graph is regarded as the updated test graph in the self-improving loop. The specific steps are described in detail below.

Step 1: Graph Prompt Generator

Given a test graph G_t , SIGOOD leverages a well-trained GNN to obtain its node embedding h . However, due to the distribution overlapping between ID and OOD data, it is not trivial to conduct test-time OOD detection by directly using the raw embeddings of test graph. Therefore, we use a graph prompt generator to dynamically guides and improves the distribution distinguishability of ID and OOD data. Specifically, the prompt generator $PG(\cdot)$ is implemented as a lightweight three-layer MLP module (Rumelhart, Hinton, and Williams 1986), which rapidly generates OOD prompts during the testing phase of GNNs. It is formally defined as :

$$v^* = \text{ReLU}(\mathbf{W}_2 \cdot \text{ReLU}(\mathbf{W}_1 v + b_1) + b_2),$$

$$PG(v) = \mathbf{W}_3 \left(\gamma \cdot \frac{v^* - \mu}{\sqrt{\sigma^2 + \epsilon}} + \lambda \right) + b_3, \quad (2)$$

where v is the node embedding, $\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3$ are trainable weight matrices, and $\gamma, \lambda, \mu, \sigma$ represent learnable parameters for normalization. The prompt P_m is composed of all nodes re-embedded by the prompt generator $PG(v)$. These parameters allow the generator to effectively guide the distinguishing on graphs.

Step 2: Prompt Injection

Previous studies have shown that graph components, such as certain nodes or subgraphs, typically contain specific signals that represent their ID or OOD distribution patterns (Yu, Liang, and He 2023). Motivated by this, SIGOOD uses the

generated prompt to facilitate the recognition of OOD signals included in a test graph G_t . Specifically, SIGOOD integrates the generated prompt P_m with G_t to construct a prompt-enhanced graph G_p :

$$G_p = G_t \oplus P_m, \quad (3)$$

where \oplus denotes element-wise addition between the corresponding node embeddings of G_t and its prompt P_m . By optimizing the prompt generator, G_p is obtained by improving the OOD signals while suppressing the ID signals in G_t , thus enhancing the separability between ID and OOD data.

Step 3: OOD Signal Recognition

A well-trained GNN exhibits strong inductive bias toward ID graphs (Wu et al. 2020). Therefore, if a test graph is ID, the representation generated by the well-trained GNN has a high probability to align with the training distribution. Otherwise, its embeddings are less aligned with the training distribution. Given that energy is a transformed form of negative log-probability, existing OOD detection methods often leverage energy as a scoring function to distinguish ID and OOD graphs (Jiang et al. 2025). However, directly observing the energy provides only a static information and fails to capture the model’s dynamic response after prompt injection. To further exploit energy scores, SIGOOD analyzes node-wise energy variations between the original test graph G_t and the prompt-enhanced graph G_p .

The calculation of energy variations begins by computing the energy scores of each node in both G_t and G_p . SIGOOD adds a lightweight scoring head on the well-trained GNN to produce a 2-dimensional logit for representing the ID and OOD preference. For a node embedding v , let $f(v) \in \mathbb{R}^2$ denote logit function, the energy of node v is defined as:

$$\hat{E}(v) = -\log \sum_{i=1}^2 \exp(f_i(v)). \quad (4)$$

After computing the energy of each node in both the original graph G_t and the PE graph G_p , we quantify the energy variation of each corresponding node to emphasize relative changes. Specifically, the energy variation is defined as:

$$\begin{aligned} \Delta E_v &= \log \hat{E}(v; G_p) - \log \hat{E}(u; G_t) \\ &= \log \frac{\hat{E}(v; G_p)}{\hat{E}(u; G_t)}, \end{aligned} \quad (5)$$

where $u \in G_t$ is the corresponding node of $v \in G_p$. The energy variation ΔE_v provides a directional indicator for identifying and amplifying potential OOD signals. Following the intuition in the overview, SIGOOD identifies nodes with $\Delta E > 0$ as OOD nodes v_{ood} , whereas nodes with $\Delta E < 0$ are regarded as ID nodes v_{id} .

Step 4: Feedback Optimization

To optimize SIGOOD, we amplify the energy variation between the prompt-enhanced graph and the original test graph

to strengthen the OOD signal. We proposed an energy preference optimization (EPO) loss \mathcal{L}_{EPO} , which is defined as:

$$\begin{aligned} \mathcal{L}_{EPO} &= -\log \sigma \left(\beta \log \mathbb{E}_{v_{ood} \sim G_p} \left[\frac{\hat{E}(v_{ood}; G_p)}{\hat{E}(u_{ood}; G_t)} \right] \right. \\ &\quad \left. - \beta \log \mathbb{E}_{v_{id} \sim G_p} \left[\frac{\hat{E}(v_{id}; G_p)}{\hat{E}(u_{id}; G_t)} \right] \right). \end{aligned} \quad (6)$$

The EPO loss consists with an energy-based Bradley-Terry model (Bradley and Terry 1952) and Kullback–Leibler (KL) divergence (Kullback and Leibler 1951). The implementation of EPO is as follows.

Given the ID/OOD region signal in Step 3, SIGOOD optimizes the prompt generator to enhance the distinguishability of OOD data. An expected loss is that which can increase the difference between ID and OOD signals to facilitate the detection of OOD data. An intuitive optimization strategy is the probability \mathcal{P} of OOD signals is more preferred by the model than ID signals. Inspired by the Bradley-Terry model, the probability \mathcal{P} is formulated as $\mathcal{P}(v_{ood} \succ v_{id}) = \frac{\exp(r(P_m, v_{ood}))}{\exp(r(P_m, v_{ood})) + \exp(r(P_m, v_{id}))}$, where $r(P_m, v)$ represents a reward function modeling the joint effect of the prompt P_m and node v , designed to enhance OOD signals in graphs. Then we maximize the log-likelihood of such pairwise comparisons, the optimization goal is defined as:

$$\text{Max}(\log \sigma(r(P_m, v_{ood}) - r(P_m, v_{id}))). \quad (7)$$

SIGOOD aims to optimize the OOD signals in the original graph G_p under the guidance of a reward function $r(\cdot)$, while ensuring that the PE graph G_p does not deviate excessively from G_t . To achieve this, SIGOOD introduces the Kullback–Leibler (KL) divergence constraint to regularize the difference between G_p and G_t . Specifically, SIGOOD defines an energy-based KL-divergence as follows: $D_{KL} = (\hat{E}(v; G_p) || \hat{E}(u; G_t))$. Inspired by DPO (Rafailov et al. 2023) reward mechanisms, the reward function of SIGOOD can be defined as:

$$r(P_m, v) = \beta \log \frac{\hat{E}(v; G_p)}{\hat{E}(u; G_t)}, \quad (8)$$

where u is the correspond node of v in graph G_t . The detailed derivation is provided in the *Appendix*. Combined with Eq.(5), the reward function $r(P_m, v)$ can be interpreted as the energy variation induced by parameter amplification. This observation suggests that energy variations can serve as an effective optimization objective for enhancing OOD signals. Furthermore, by incorporating the reward function into Eq.(7), the optimization goal can be formulated as:

$$\text{Max} \left(\log \sigma \left(\beta \log \frac{\hat{E}(v_{ood}; G_p)}{\hat{E}(u_{ood}; G_t)} - \beta \log \frac{\hat{E}(v_{id}; G_p)}{\hat{E}(u_{id}; G_t)} \right) \right). \quad (9)$$

To optimize the prompt generator, we contrast the energy variation of OOD and ID nodes between the PE graph G_p and the input graph G_t . The first term encourages energy of OOD nodes to increase after prompt injection, while the second term penalizes excessive energy variation in ID nodes.

By maximizing their difference, SIGOOD promotes OOD-signal responses while maintaining stability for ID nodes.

While the aforementioned optimization objectives primarily focus on identifying and amplifying local OOD signals, the core task of SIGOOD is graph-level OOD detection. Therefore, relying solely on local signals is insufficient. To comprehensively assess whether a test graph deviates from the training distribution, we further incorporate graph-level energy variation as a global optimization signal, aligning with the model’s preference for global energy shifts. Specifically, we compute the average energy increase and decrease across all nodes in graphs, capturing the overall energy variation to reflect global OOD tendencies. The final EPO loss, is shown in Eq.(6). Guided by energy-based feedback, the prompt generator is iteratively updated to amplify the OOD signal in G_p . The refined G_p replaces G_t as the new input for subsequent iterations, forming a self-improving loop that continuously strengthens SIGOOD’s ability to detect OOD patterns. After training, the input graph G_t is combined with the optimized prompt to produce a prompt-enhanced graph. This PE graph is then evaluated using the EPO loss, which serves as the OOD detection scores. A predefined threshold τ is applied to determine whether the graph is ID or OOD.

Why SIGOOD is effective? Well-trained GNNs exhibit strong expressive power in modeling ID patterns, resulting in high confidence when classifying ID graphs. However, they often show low confidence and unstable predictions for OOD graphs. Energy, which serves as a transformed representation of the GNNs output probability, is widely used as an indicator for OOD detection, where higher energy values suggest a greater likelihood of being OOD. Leveraging this property, SIGOOD initially identifies potential OOD signals through energy variations. Guided by energy-based feedback, the prompt generator is iteratively optimized to produce prompt-enhanced graphs that better expose OOD patterns. These enhanced graphs are then used as new test inputs G_t in the next iteration. Through this iterative refinement, SIGOOD progressively corrects prior mispredictions, forming a self-improving optimization loop that enhances OOD detection performance over time.

Experiments

We evaluate the effectiveness of SIGOOD on both graph OOD/anomaly detection tasks. Due to space limitations, the complete results are provided in the *arXiv* version.

Experimental Setups

Datasets. We selected multiple datasets from diverse domains included in the widely-used UB-GOLD benchmark (Wang et al. 2025c). The distribution of OOD detection datasets comes from drug chemical formulas, protein structures, and etc. During testing, ID and OOD samples are mixed in a 1 : 1 ratio. The anomaly detection datasets span biological, chemical, and social domains. Instances belonging to the minority or ground-truth anomaly class are designated as anomalies, whereas others are treated as normal.

Baselines. We compare SIGOOD with 12 competitive baseline methods, detailed as follows:

- **Traditional OOD Detectors.** These methods employ pre-trained encoders to extract representations, followed by classical OOD or anomaly detectors. Common encoders include the Weisfeiler-Lehman (WL) kernel (Shervashidze et al. 2011) and the propagation kernel (PK) (Neumann et al. 2016). The downstream OOD detectors include Local Outlier Factor (LOF) (Breunig et al. 2000), One-Class SVM (OCSVM) (Manevitz and Yousef 2001), Isolation Forest (iF) (Liu, Ting, and Zhou 2008), and Mahalanobis Distance-based detector (MD) (Sehwag, Chiang, and Mittal 2021).
- **GNNs with Post-hoc Detectors.** These methods generate graph representations using GNNs and perform detection in a post-hoc detectors. The GNN encoders include InfoGraph (Sun et al. 2020) and GraphCL (You et al. 2020), while the post-hoc detector typically involves Isolation Forest.
- **Test-time Training Methods.** These methods perform OOD detection during test-time without training datasets. Notable methods include GTrans (Jin et al. 2023b) and GOODAT (Wang et al. 2024).

Implementation Details. Following prior studies, we adopt the Area Under the ROC Curve (AUC) as the primary evaluation metric (Wang et al. 2024). All experiments are conducted on an NVIDIA RTX 4090 GPU with 24GB of memory. Each experiment is repeated five times to ensure stability. For baseline methods, we use the results reported in their original papers, such as GOODAT.

Performance on OOD Detection

Table 1 reports the AUC scores (%) for graph OOD detection across eight ID/OOD dataset pairs. Overall, SIGOOD consistently achieves the best performance, outperforming 12 state-of-the-art (SOTA) baselines across all benchmarks. In particular, (1) Compared to traditional OOD detectors, SIGOOD demonstrates a clear advantage over traditional graph OOD methods such as PK-iF and WL-OCSVM. This illustrates the limitations of traditional shallow OOD detectors in capturing complex distribution shifts in graphs. (2) Compared to GNN-based baselines, SIGOOD exhibits steady improvements, highlighting the effectiveness of test-time enhancement. This indicates that refining the OOD signal at inference time enhances the model’s ability to capture the patterns of OOD graphs. (3) Compared with other test-time OOD detection methods, SIGOOD maintains robust performance on biochemical datasets. On Tox21/ToxCast and ClinTox/LIPO, it achieves AUCs of 69.97% and 71.33%, respectively, surpassing the previous best method GOODAT by 1.52% and 14.20%. These results confirm the effectiveness of SIGOOD, particularly in different domains graph OOD detection scenarios.

Performance on Anomaly Detection

Table 2 reports the AUC scores for anomaly detection on 13 benchmark datasets with 8 baselines. Overall, SIGOOD

ID dataset	BZR	PTC-MR	AIDS	ENZYMES	Tox21	ClinTox	Esol	Avg.
OOD dataset	COX2	MUTAG	DHFR	PROTEIN	SIDER	LIPO	MUV	Rank
PK-LOF	42.22 ± 8.39	51.04 ± 6.04	50.15 ± 3.29	50.47 ± 2.87	51.33 ± 1.81	50.00 ± 2.17	50.82 ± 1.48	11.1
PK-OCSVM	42.55 ± 8.26	49.71 ± 6.58	50.17 ± 3.30	50.46 ± 2.78	51.33 ± 1.81	50.06 ± 2.19	51.00 ± 1.33	11.0
PK-iF	51.46 ± 1.62	54.29 ± 4.33	51.10 ± 1.43	51.67 ± 2.69	49.87 ± 0.82	50.81 ± 1.10	50.85 ± 3.51	8.5
WL-LOF	48.99 ± 6.20	53.31 ± 8.98	50.77 ± 2.87	52.66 ± 2.47	51.92 ± 1.58	51.29 ± 3.40	51.26 ± 1.31	8.0
WL-OCSVM	49.16 ± 4.51	53.31 ± 7.57	50.98 ± 2.71	51.77 ± 2.21	51.08 ± 1.46	50.77 ± 3.69	50.97 ± 1.65	9.0
WL-iF	50.24 ± 2.49	51.43 ± 2.02	50.10 ± 0.44	51.17 ± 2.01	50.25 ± 0.96	50.41 ± 2.17	50.61 ± 1.96	10.0
InfoGraph-iF	63.17 ± 9.74	51.43 ± 5.19	93.10 ± 1.35	60.00 ± 1.83	56.28 ± 0.81	48.51 ± 1.87	54.16 ± 5.14	6.4
InfoGraph-MD	86.14 ± 6.77	50.79 ± 8.49	69.02 ± 11.67	55.25 ± 3.51	59.97 ± 2.06	48.12 ± 5.72	77.57 ± 1.69	6.1
GraphCL-iF	60.00 ± 3.81	50.86 ± 4.30	92.90 ± 1.21	61.33 ± 2.27	56.81 ± 0.97	47.84 ± 0.92	62.12 ± 4.01	7.0
GraphCL-MD	83.64 ± 6.00	73.03 ± 2.38	93.75 ± 2.13	52.87 ± 6.11	58.30 ± 1.52	51.58 ± 3.64	78.73 ± 1.40	3.8
GTrans	55.17 ± 5.04	62.38 ± 2.36	60.12 ± 1.98	49.94 ± 5.67	61.67 ± 0.73	58.54 ± 2.38	76.31 ± 3.85	6.5
GOODAT	82.16 ± 0.15	81.84 ± 0.57	96.43 ± 0.25	66.29 ± 1.54	68.92 ± 0.01	62.46 ± 0.54	85.91 ± 0.27	2.3
Ours	87.36 ± 0.17	85.70 ± 0.03	97.38 ± 0.01	67.88 ± 0.04	69.97 ± 0.68	71.33 ± 0.14	87.72 ± 0.05	1
Improve	1.41%	4.72%	0.99%	2.40%	1.52%	14.20%	2.11%	—

Table 1: OOD detection results in terms of AUC score (%). The best results are highlighted with **bold**.

Method	WL-OCSVM	WL-iF	InfoGraph-iF	GraphCL-iF	GTrans	GOODAT	Ours	Improve
PROTEINS-full	51.35 ± 4.35	61.36 ± 2.54	57.47 ± 3.03	60.18 ± 2.53	60.16 ± 5.06	77.92 ± 2.37	79.54 ± 0.60	2.07%
ENZYMES	55.24 ± 2.66	51.60 ± 3.81	53.80 ± 4.50	53.60 ± 4.88	38.02 ± 6.24	52.33 ± 4.74	76.80 ± 0.89	39.02%
DHFR	50.24 ± 3.13	50.29 ± 2.77	52.68 ± 3.21	51.10 ± 2.35	61.15 ± 2.87	61.52 ± 2.86	65.17 ± 0.45	5.93%
BZR	50.56 ± 5.87	52.46 ± 3.30	63.31 ± 8.52	60.24 ± 5.37	51.97 ± 8.15	64.77 ± 3.87	75.42 ± 3.87	16.44%
COX2	49.86 ± 7.43	50.27 ± 0.34	53.36 ± 8.86	52.01 ± 3.17	53.56 ± 3.47	59.99 ± 9.76	77.78 ± 1.85	29.65%
DD	47.99 ± 4.09	70.31 ± 1.09	55.80 ± 1.77	59.32 ± 3.92	76.73 ± 2.83	77.62 ± 2.88	72.59 ± 1.84	—
NCI1	50.63 ± 1.22	50.74 ± 1.70	50.10 ± 0.87	49.88 ± 0.53	41.42 ± 2.16	45.96 ± 2.42	59.07 ± 0.47	16.41%
IMDB-B	54.08 ± 5.19	50.20 ± 0.40	56.50 ± 3.58	56.50 ± 4.90	45.34 ± 3.75	65.46 ± 4.34	68.96 ± 0.05	5.07%
REDDIT-B	49.31 ± 2.33	48.26 ± 0.32	68.50 ± 5.56	71.80 ± 4.38	69.71 ± 2.21	80.31 ± 0.85	86.64 ± 1.97	7.88%
HSE	62.72 ± 10.13	53.02 ± 5.12	53.56 ± 3.98	51.18 ± 2.71	58.49 ± 2.68	63.05 ± 0.90	64.68 ± 0.72	2.58%
MMP	55.24 ± 3.26	52.68 ± 3.34	54.59 ± 2.01	54.54 ± 1.86	48.19 ± 3.74	69.41 ± 0.04	70.17 ± 0.11	1.09%
p53	54.59 ± 4.46	50.85 ± 2.16	52.66 ± 1.95	53.29 ± 2.32	53.74 ± 2.98	63.27 ± 0.04	60.51 ± 1.95	—
PPAR-gamma	57.91 ± 6.13	49.60 ± 0.22	51.40 ± 2.53	50.30 ± 1.56	56.20 ± 1.57	68.23 ± 1.54	72.59 ± 0.04	6.39%
Avg. Rank	5.3	6.2	4.9	5.4	5.5	2.6	1.2	—

Table 2: Anomaly detection results in terms of AUC score (%). The best results are highlighted with **bold**.

consistently delivers the top performance, ranking first on 11 out of 13 datasets and achieving the best average rank (1.2) among 8 competing methods. Specifically, (1) SIGOOD significantly outperforms traditional baselines. These results indicate that SIGOOD is also well-suited for graph anomaly detection. (2) SIGOOD also surpasses GNN-based detectors. On BZR dataset, SIGOOD improves over the strongest neural baseline (e.g., GraphCL-iF) by 19.12%. It suggests that, beyond pretraining or contrastive learning, the energy-based self-improvement mechanism enhances the SIGOOD to distinguish subtle OOD signals that not captured by existing detectors. (3) Compared to SOTA test-time methods, SIGOOD shows notable gains, with improvements of 39.02% on ENZYMES and 29.65% on COX2, demonstrating the effectiveness of self-improvement mechanism across biochemical and social graphs. (4) While SIGOOD does not achieve the top score on DD and p53, these deviations likely stem from dataset-specific variance rather than fundamental limitations. Overall, although SIGOOD fails to achieve the highest score on a few datasets, it achieves highly competi-

tive average rankings across a variety of benchmarks.

Ablation Study

The Fig. 3 (a) visualizes the performance (AUC) of three model variants across different dataset pairs. As shown, SIGOOD consistently achieves the highest AUC on all three tasks (BZR/COX2, PTC-MR/MUTAG, and Esol/MUV), demonstrating the effectiveness of the full model. In contrast, removing the energy preference optimization loss (W/O L_{EPO}) leads to the most severe performance drop, particularly on PTC-MR/MUTAG and Esol/MUV, indicating that L_{EPO} is crucial for enhancing OOD sensitivity. Replacing the prompt generator with an optimizable parameter matrix (W/O PG) also results in a significant performance degradation. This suggests that the prompt generator plays a crucial role in amplifying ood signals. These results collectively confirm that both the energy-guided optimization and prompt-based modulation are essential to SIGOOD’s superior generalization capability.

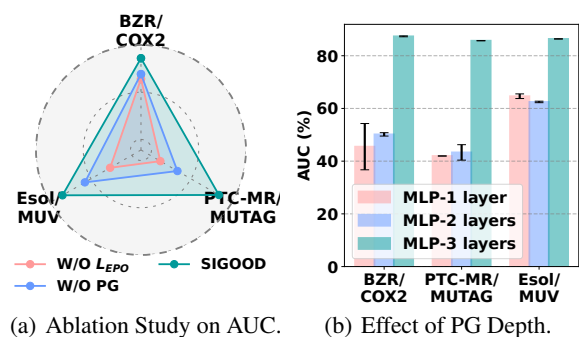


Figure 3: Ablation Study and PG Depth Analysis.




Figure 4: Parameter Sensitivity Analysis.

Parameter Sensitivity Analysis

Effect of Prompt Generator (PG) Depth. To evaluate the impact of the prompt generator depth on OOD detection performance, we experiment with 1-layer, 2-layer, and 3-layer MLPs on three datasets: BZR/COX2, PTC-MR/MUTAG, and Esol/MUV. As shown in Fig. 3 (b), using a 3-layer MLP significantly improves detection accuracy, aligning with common empirical practices in deep learning. In contrast, shallow configurations yield noticeably lower AUCs, especially on BZR/COX2 and PTC-MR/MUTAG. These results suggest that a deeper prompt generator (with-in 3-layers) is more effective in capturing complex semantic patterns and refining energy-based representations for reliable OOD detection.

The effects of Self-improving Iterations. To intuitively demonstrate the impact of self-improvement iterations on the OOD detection performance, we conducted experiments on the Esol/MUV dataset with large-scale and high-density features. As shown in Fig. 4 (a), increasing the number of iterations from 1000 to 10000 consistently improves the performance of SIGOOD, with the AUC reaching a peak of 87.72%. This trend demonstrates the effectiveness of iterative on self-learning. However, when the iterations exceed 10000, the performance gain diminishes slightly, indicating that SIGOOD has approached convergence. Therefore, a moderate iterations offers a favorable trade-off between performance and computational efficiency.

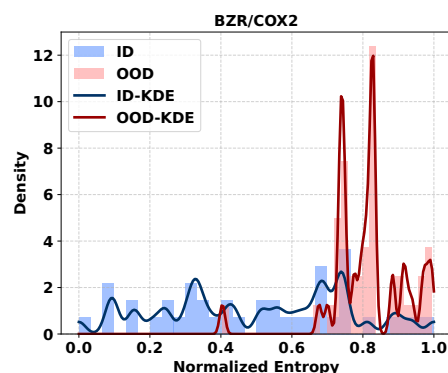


Figure 5: Visualization of Graph Distribution.

The effects of parameter β . In SIGOOD, β serves as a weighting factor in the energy preference optimization (EPO) objective, controlling the trade-off between preserving original graph semantics and encouraging energy-based differentiation between in-distribution and OOD samples. We investigate the sensitivity of our model to the hyperparameter β on the BZR/COX2 dataset. As shown in Fig. 4 (b), the highest value of 87.36% is observed at $\beta = 80$. This indicates that moderate values of β yield slightly better performance. A small β may under-emphasize the energy gap, resulting in weaker OOD separation, while a large β may distort the embedding space.

Visualization

To further illustrate the effectiveness of our method in distinguishing OOD samples, we visualize the normalized entropy distribution of the BZR/COX2 dataset in Fig. 5. The figure clearly shows a separation between ID and OOD samples, with OOD instances exhibiting higher entropy values. The kernel density estimation (KDE) curves further reveal that the OOD distribution is skewed towards the high-entropy region, while the ID distribution remains concentrated in the low-entropy area. The distinct distributional shift highlights the reliability of entropy as a scoring metric for OOD detection in our framework.

Conclusion

In this paper, we propose SIGOOD, a self-improving framework for test-time graph OOD detection, which leverages energy-based feedback to iteratively refine OOD detection. By integrating a lightweight prompt generator with a well-trained GNN, SIGOOD constructs PE graphs that amplify potential OOD signals. Through energy variations between both the PE graph and the original test graph, the model identifies and enhances OOD-relevant signals. SIGOOD also proposes a novel energy preference optimization loss to guide prompt updates, enabling the framework to form a closed-loop self-improving process without requiring additional labels at test time. Extensive experiments demonstrate that SIGOOD achieves superior OOD detection performance across various graph datasets, highlighting its effectiveness in real-world scenarios.

Acknowledgments

This work is supported by National Natural Science Foundation of China (62502065, 62176036, 62302333); the Beatriu de Pinós del Departament de Recerca i Universitats de la Generalitat de Catalunya (2022 BP-00256); the predoctoral program AGAUR-FI ajuts (2025 FI-200470) Joan Oró, which is backed by the Secretariat of Universities and Research of the Department of Research and Universities of the Generalitat of Catalonia, as well as the European Social Plus Fund.

References

- Bi, X.; Lu, J.; Liu, B.; Cun, X.; Zhang, Y.; Li, W.; and Xiao, B. 2025. Customttt: Motion and Appearance Customized Video Generation via Test-time Training. In *AAAI 2025*, 1871–1879.
- Bradley, R. A.; and Terry, M. E. 1952. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika*, 39(3/4): 324–345.
- Breunig, M. M.; Kriegel, H.; Ng, R. T.; and Sander, J. 2000. LOF: Identifying Density-Based Local Outliers. In *SIGMOD 2000*, 93–104.
- Cao, Y.; Shi, F.; Yu, Q.; Lin, X.; Zhou, C.; Zou, L.; Zhang, P.; Li, Z.; and Yin, D. 2025. IBPL: Information Bottleneck-based Prompt Learning for Graph Out-of-distribution Detection. *Neural Networks*, 188: 107381.
- Chen, S.; Huang, L.-K.; Schwarz, J. R.; Du, Y.; and Wei, Y. 2023. Secure Out-of-distribution Task Generalization with Energy-based Models. In *NeurIPS 2023*, 67007–67020.
- Dalal, K.; Koceja, D.; Xu, J.; Zhao, Y.; Han, S.; Cheung, K. C.; Kautz, J.; Choi, Y.; Sun, Y.; and Wang, X. 2025. One-minute Video Generation with Test-time Training. In *CVPR 2025*, 17702–17711.
- Ding, Z.; Shi, J.; Shen, S.; Shang, X.; Cao, J.; Wang, Z.; and Gong, Z. 2024. Sgood: Substructure-enhanced Graph-level Out-of-distribution Detection. In *CIKM 2024*, 467–476.
- Fu, L.; Deng, B.; Huang, S.; Liao, T.; Zhang, C.; and Chen, C. 2025. Learn from Global Rather Than Local: Consistent Context-Aware Representation Learning for Multi-View Graph Clustering. In *IJCAI 2025*, 5145–5153.
- Fuchsgruber, D.; Wollschläger, T.; and Günnemann, S. 2024. Energy-based Epistemic Uncertainty for Graph Neural Networks. In *NeurIPS 2024*, 34378–34428.
- Hou, Y.; Zhu, H.; Liu, R.; Su, Y.; Xia, J.; Wu, J.; and Xu, K. 2025. Structural Entropy Guided Unsupervised Graph Out-Of-Distribution Detection. In *AAAI 2025*, 17258–17266.
- Jiang, Z.; Lu, J.; Fan, H.; Wang, T.; and Yan, J. 2025. Learning Structured Universe Graph with Outlier OOD Detection for Partial Matching. In *ICLR 2025*, 1–16.
- Jin, D.; Wang, L.; Zheng, Y.; Song, G.; Jiang, F.; Li, X.; Lin, W.; and Pan, S. 2023a. Dual Intent Enhanced Graph Neural Network for Session-based New Item Recommendation. In *WWW 2023*, 684–693.
- Jin, W.; Zhao, T.; Ding, J.; Liu, Y.; Tang, J.; and Shah, N. 2023b. Empowering Graph Representation Learning with Test-Time Graph Transformation. In *ICLR 2023*, 1–27.
- Kullback, S.; and Leibler, R. A. 1951. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1): 79–86.
- Liang, J.; He, R.; and Tan, T. 2025. A Comprehensive Survey on Test-time Adaptation under Distribution Shifts. *International Journal of Computer Vision*, 133(1): 31–64.
- Lin, X.; Cao, Y.; Sun, N.; Zou, L.; Zhou, C.; Zhang, P.; Zhang, S.; Zhang, G.; and Wu, J. 2025. Conformal Graph-level Out-of-distribution Detection with Adaptive Data Augmentation. In *WWW 2025*, 4755–4765.
- Lin, X.; Zhang, W.; Shi, F.; Zhou, C.; Zou, L.; Zhao, X.; Yin, D.; Pan, S.; and Cao, Y. 2024. Graph Neural Stochastic Diffusion for Estimating Uncertainty in Node Classification. In *ICML 2024*, 30457–30478.
- Liu, F. T.; Ting, K. M.; and Zhou, Z. 2008. Isolation Forest. In *ICDM 2008*, 413–422.
- Liu, W.; Wang, X.; Owens, J.; and Li, Y. 2020. Energy-based Out-of-distribution Detection. In *NeurIPS 2020*, 21464–21475.
- Manevitz, L. M.; and Yousef, M. 2001. One-Class SVMs for Document Classification. *Journal of Machine Learning Research*, 2: 139–154.
- Neumann, M.; Garnett, R.; Bauckhage, C.; and Kersting, K. 2016. Propagation Kernels: Efficient Graph Kernels from Propagated Information. *Machine Learning*, 102(2): 209–245.
- Pan, J.; Liu, Y.; Zheng, X.; Zheng, Y.; Liew, A. W.-C.; Li, F.; and Pan, S. 2025. A Label-free Heterophily-guided Approach for Unsupervised Graph Fraud Detection. In *AAAI 2025*, 12, 12443–12451.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct Preference Optimization: Your Language Model is Secretly A Reward Model. In *NeurIPS 2023*, volume 36, 53728–53741.
- Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1986. Learning Representations by Back-Propagating Errors. *Nature*, 323(6088): 533–536.
- Sehwag, V.; Chiang, M.; and Mittal, P. 2021. SSD: A Unified Framework for Self-Supervised Outlier Detection. In *ICLR 2021*, 1–17.
- Shen, X.; Wang, Y.; Zhou, K.; Pan, S.; and Wang, X. 2024. Optimizing OOD Detection in Molecular Graphs: A Novel Approach with Diffusion Models. In *SIGKDD 2024*, 2640–2650.
- Shervashidze, N.; Schweitzer, P.; van Leeuwen, E. J.; Mehlhorn, K.; and Borgwardt, K. M. 2011. Weisfeiler-Lehman Graph Kernels. *Journal of Machine Learning Research*, 12: 2539–2561.
- Sun, F.; Hoffmann, J.; Verma, V.; and Tang, J. 2020. InfoGraph: Unsupervised and Semi-supervised Graph-level Representation Learning via Mutual Information Maximization. In *ICLR 2020*, 1–16.
- Wang, L.; He, D.; Zhang, H.; Liu, Y.; Wang, W.; Pan, S.; Jin, D.; and Chua, T.-S. 2024. Goodat: Towards Test-time Graph Out-of-distribution Detection. In *AAAI 2024*, 15537–15545.

- Wang, X.; Dong, Y.; Jin, D.; Li, Y.; Wang, L.; and Dang, J. 2023. Augmenting Affective Dependency Graph via Iterative Incongruity Graph Learning for Sarcasm Detection. In *AAAI 2023*, 4, 4702–4710.
- Wang, X.; Wang, Y.; He, D.; Yu, Z.; Li, Y.; Wang, L.; Dang, J.; and Jin, D. 2025a. Elevating Knowledge-enhanced Entity and Relationship Understanding for Sarcasm Detection. *IEEE Transactions on Knowledge and Data Engineering*, 37(6): 3356–3371.
- Wang, Y.; Liu, Y.; Liu, N.; Miao, R.; Wang, Y.; and Wang, X. 2025b. AdaGCL+: An Adaptive Subgraph Contrastive Learning Towards Tackling Topological Bias. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8073–8087.
- Wang, Y.; Liu, Y.; Shen, X.; Li, C.; Miao, R.; Ding, K.; Wang, Y.; Pan, S.; and Wang, X. 2025c. Unifying Un-supervised Graph-Level Anomaly Detection and Out-of-Distribution Detection: A Benchmark. In *ICLR 2025*, 1–27.
- Wu, Q.; Chen, Y.; Yang, C.; and Yan, J. 2023. Energy-based Out-of-Distribution Detection for Graph Neural Networks. In *ICLR 2023*.
- Wu, Q.; Nie, F.; Yang, C.; Bao, T.; and Yan, J. 2024. Graph Out-of-distribution Generalization via Causal Intervention. In *WWW 2024*, 850–860.
- Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Yu, P. S. 2020. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1): 4–24.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph Contrastive Learning with Augmentations. In *NeurIPS 2020*, 5812–5823.
- Yu, J.; Liang, J.; and He, R. 2023. Mind the Label Shift of Augmentation-based Graph OOD Generalization. In *CVPR 2023*, 11620–11630.
- Zhang, H.; Wu, B.; Yang, X.; Yuan, X.; Liu, X.; and Yi, X. 2025a. Dynamic Graph Unlearning: A General and Efficient Post-processing Method via Gradient Transformation. In *WWW 2025*, 931–944.
- Zhang, H.; Wu, B.; Yuan, X.; Pan, S.; Tong, H.; and Pei, J. 2024a. Trustworthy Graph Neural Networks: Aspects, Methods, and Trends. *Proceedings of the IEEE*, 112(2): 97–139.
- Zhang, J.; Wang, Y.; Yang, X.; Wang, S.; Feng, Y.; Shi, Y.; Ren, R.; Zhu, E.; and Liu, X. 2024b. Test-time Training on Graphs with Large Language Models (LLMs). In *ACM MM 2024*, 2089–2098.
- Zhang, J.; Wang, Y.; Yang, X.; and Zhu, E. 2024c. A Fully Test-time Training Framework for Semi-supervised Node Classification on Out-of-distribution Graphs. *ACM Transactions on Knowledge Discovery from Data*, 18(7): 1–19.
- Zhang, Q.; Shi, Z.; Pan, S.; Chen, J.; Wu, H.; and Chen, X. 2024d. EGonc: Energy-based Open-set Node Classification with Substitute Unknowns. In *NeurIPS 2024*, volume 37, 66147–66177.
- Zhang, S.; Zhou, C.; Liu, Y.; Zhang, P.; Lin, X.; and Pan, S. 2025b. Conformal Anomaly Detection in Event Sequences. In *ICML 2025*, 1–17.
- Zheng, X.; Song, D.; Wen, Q.; Du, B.; and Pan, S. 2024. Online GNN Evaluation Under Test-time Graph Distribution Shifts. In *ICLR 2024*, 1–22.