GEOMETRY-AWARE GRAPH ANOMALY DETECTION: UNIFIED HYPERBOLIC AFFINITY LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Unsupervised graph anomaly detection (UGAD) aims to identify abnormal nodes that deviate from normal patterns without labeled data and has been widely applied to fraud detection, cybersecurity, etc. However, node relationships in real-world graphs often exhibit hierarchical structure. Existing affinity-based graph anomaly detection methods mainly rely on Euclidean spaces to model node relationships, which limits their ability to capture relational differences in complex graph structures. Moreover, different relations are often modeled using different measurement spaces, making it difficult to jointly model multiple relations in complex graph structures. To address these issues, we propose a geometry-aware graph anomaly detection framework via unified hyperbolic affinity learning (UHAL-GAD). Specifically, we first map node representations into a shared Lorentz space and construct a unified hyperbolic affinity based on geodesic distances to learn more discriminative hyperbolic node representations. On this basis, the hyperbolic structural affinity learning module discriminates between normal and abnormal nodes using hyperbolic affinity, while the hyperbolic center affinity constraint module characterizes the hyperbolic center affinity discrepancy of each node with respect to the global Lorentz center. Furthermore, these modules jointly model multiple relations in a unified hyperbolic affinity measurement space. Extensive experiments on multiple graph datasets demonstrate the effectiveness and superior detection performance of UHAL-GAD.

1 INTRODUCTION

With the widespread application of graph data in domains such as social networks Wu et al. (2020), e-commerce Dou et al. (2020), and financial transactions Tang et al. (2022), anomalous behaviors, including social spammers, fake reviews, and fraudulent transactions, have attracted increasing attention. Consequently, graph anomaly detection (GAD) has become an important research direction. GAD aims to identify anomalous patterns that deviate from the majority of instances in graph data Akoglu et al. (2015). In particular, node-level anomaly detection plays an important role in practical applications, such as fraudulent user detection. In recent years, with the continuous growth in the scale of graph data, challenges such as the high cost of manual annotation have become increasingly prominent. As a result, unsupervised node-level graph anomaly detection, which does not rely on manual annotations, has received growing research attention Liu et al. (2021a); Jin et al. (2021); Pan et al. (2025).

Existing unsupervised node-level graph anomaly detection methods can be broadly categorized into three paradigms: reconstruction-based, contrastive learning-based, and graph anomaly measure-based methods Qiao et al. (2025). Reconstruction-based methods measure the degree of abnormality using reconstruction errors of node attributes and graph structures Fan et al. (2020); He et al. (2024). Contrastive learning-based methods construct positive and negative sample pairs to learn consistency between nodes and their neighborhoods or across different views, and identify anomalies based on contrastive scores Zheng et al. (2021); Zhang et al. (2022). However, the former are susceptible to noise and reconstruction overhead, while the latter rely on effective sample construction and anomaly scoring mechanisms. In comparison, graph anomaly measure-based methods directly quantify deviations from normal patterns using criteria such as one-class distance, community adherence, or local affinity. Among them, affinity-based methods typically assume that normal nodes exhibit high affinity with their neighbors, whereas anomalous nodes tend to have lower neigh-

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

borhood affinity. Accordingly, anomalies are identified by characterizing node affinities and their differences Qiao & Pang (2023); Zhang et al. (2026).

Affinity-based methods have attracted increasing attention because they directly exploit graph relationships to characterize anomalies and provide interpretable criteria for anomaly discrimination. Qiao & Pang (2023) proposed TAM, which maximizes local affinity and iteratively truncates heterophilous edges to enlarge the affinity gap between normal and anomalous nodes. Kim et al. (2023) employed Jensen–Shannon divergence to measure discrepancies between nodes and their neighbors, and combined structural and attribute information to calculate anomaly scores. Lian et al. (2026) proposed MI-GAD, which jointly exploits neighborhood-level and global anomaly signals to identify complex anomalous patterns. By characterizing affinity relationships between nodes and their associated instances, these methods provide important discriminative cues for identifying anomalous nodes.

Despite these advances, affinity modeling in complex graph structures still faces several challenges. First, existing affinity-based methods primarily rely on Euclidean space to characterize relationships between nodes, while real-world graphs often exhibit complex hierarchical structures Nickel & Kiela (2018); Chami et al. (2019). Second, recent studies have begun to jointly exploit relational information at different levels. For instance, MiLIAD integrates node, neighborhood, and global context information through multi-level information alignment and decoupling Wu & Gu (2025b). Nevertheless, such methods mainly focus on integrating information from different levels, leaving the joint characterization of multi-level node relationships under a unified affinity metric an open research problem.

To address these challenges, we propose a geometry-aware graph anomaly detection framework via unified hyperbolic affinity learning (UHAL-GAD), an unsupervised graph anomaly detection method. UHAL-GAD aims to characterize complex node relationships using hyperbolic geometry and jointly model inter-node structural relationships and global center relationships under a unified affinity metric. The overall framework is illustrated in Figure 1. First, a hyperbolic affinity learning module employs a shared encoder to obtain node representations and maps them into hyperbolic space through the Lorentz exponential map. A unified hyperbolic affinity measure is then defined based on hyperbolic geodesic distance, providing a common geometric representation for subsequent relational modeling. Next, a hyperbolic structural affinity modeling module exploits the affinities between nodes and their neighbors and non-neighbors to learn locally discriminative representations. A hyperbolic center affinity constraint module further characterizes the relationships between nodes and the Lorentz center under the same hyperbolic affinity metric, thereby capturing global anomaly information. Finally, the anomaly scores obtained from these two perspectives are fused to perform node anomaly detection. Our main contributions are summarized as follows:

- We propose UHAL-GAD, a unified hyperbolic affinity learning framework that provides a new geometric perspective for modeling node relationships in complex graph structures.
- We introduce a multi-relation collaborative modeling mechanism based on unified hyperbolic affinity, which jointly characterizes inter-node structural relationships and global center relationships to enhance the discrimination of complex anomalous patterns.
- We conduct extensive experiments on eight real-world datasets, and the results demonstrate the superior anomaly detection performance of UHAL-GAD.

2 Related Work

2.1 Unsupervised Graph Anomaly Detection

In recent years, unsupervised graph anomaly detection has made substantial progress in identifying anomalous nodes that deviate from the majority. Roy et al. (2024) proposed GAD-NR, which reconstructs local graph structures, node attributes, and neighbor attributes to identify anomalous nodes based on neighborhood reconstruction errors. Huang et al. (2023) introduced VGOD, combining variance-based structural anomaly detection with attribute reconstruction to capture different types of anomalous patterns. In self-supervised learning, Xu et al. (2022) developed CONAD, which employs anomaly-aware data augmentation and contrastive learning to obtain node representations and utilizes reconstruction errors for anomaly detection. Duan et al. (2023) proposed GRADATE, which

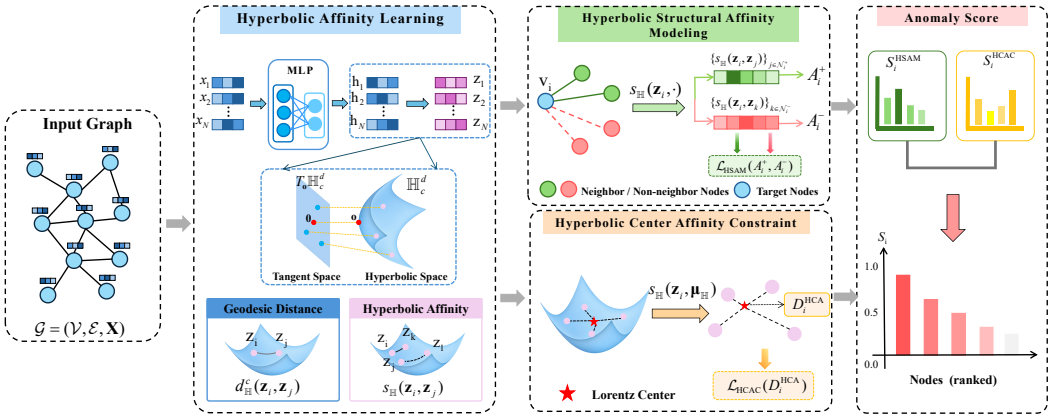

Figure 1: The overall framework of UHAL-GAD. (a) Node representations are mapped into the Lorentz space to learn unified hyperbolic affinity. (b) Hyperbolic structural affinity is modeled through neighbor and non-neighbor relationships. (c) The relationships between nodes and the Lorentz center are constrained under the unified hyperbolic affinity metric.

jointly models relationships between nodes and subgraphs through multi-view and multi-scale contrastive learning. Pan et al. (2023) introduced PREM, an efficient anomaly detection framework based on preprocessing and node-neighborhood matching. Wu & Gu (2025a) proposed TCL-GAD, integrating community information with multi-level contrastive learning to enhance the identification of complex anomalous patterns. Overall, these methods exploit anomaly signals from different perspectives, yet effectively characterizing node relationships in complex graph structures remains challenging.

2.2 Hyperbolic Graph Learning

Hyperbolic geometry has attracted increasing attention in graph representation learning due to its advantages in representing hierarchical structures. Nickel and Kiela employed the Poincaré and Lorentz models to learn hierarchical representations Nickel & Kiela (2017; 2018). Ganea et al. (2018) introduced hyperbolic neural networks, laying the foundation for feature transformations in hyperbolic space. Subsequently, Liu et al. (2019) proposed HGNN, extending graph neural networks to Riemannian manifolds. Chami et al. (2019) developed HGCN, which learns hierarchical graph representations through hyperbolic feature transformations and neighborhood aggregation. In anomaly detection, Fu et al. (2026) combined graph and hypergraph channels through hyperbolic contrastive learning to capture higher-order group information and detect graph-level anomalies. Gu & Zou (2024) revisited node-level graph anomaly detection from three perspectives: anomalous data, message passing, and hyperbolic neural networks. Collectively, these studies have established important foundations for hyperbolic graph representation learning. Nevertheless, further exploiting hyperbolic geometry to enhance affinity-based relational representations in complex graphs and jointly characterize multi-level node relationships under a unified hyperbolic affinity metric remains an open research problem.

3 The Proposed Method

3.1 Problem Definition

This paper investigates the task of unsupervised node anomaly detection on attributed graphs, aiming to identify anomalous nodes that deviate from normal patterns using only graph structure and node attributes. An attributed graph is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, where $\mathcal{V} = \{v_1, v_2, \ldots, v_N\}$ and $\mathcal{E}$ denote the sets of N nodes and M edges, respectively. The node attribute matrix is denoted as $\mathbf{X} \in \mathbb{R}^{N \times F}$, where $\mathbf{x}_i \in \mathbb{R}^F$ represents the F-dimensional attribute vector of node v_i. The graph structure is represented by a binary adjacency matrix $\mathbf{A} \in \{0,1\}^{N \times N}$, where $A_{ij} = 1$ if an edge

exists between nodes v_i and v_j, and $A_{ij} = 0$ otherwise. Accordingly, the one-hop neighborhood of node v_i is defined as $\mathcal{N}_i = \{v_j \mid A_{ij} = 1\}$.

3.2 Hyperbolic Affinity Learning

Node relations in many real-world graphs exhibit hierarchical, non-uniform, and complex community structures, which are difficult to adequately characterize across different scales in finite-dimensional Euclidean spaces. To this end, we design a Hyperbolic Affinity Learning (HAL) module that maps nodes into a shared Lorentz space and constructs a unified affinity representation based on hyperbolic geodesic relations, providing a consistent geometric basis for subsequent modeling of local neighborhood relations and global center relations.

First, to preserve individual node representations and avoid premature smoothing of anomalous features caused by neighborhood aggregation, we employ a multilayer perceptron (MLP) as a shared encoder to map the original node attributes into a low-dimensional Euclidean latent space:

$$\mathbf{h}_i = f_\theta(\mathbf{x}_i), \tag{1}$$

where $\mathbf{x}_i \in \mathbb{R}^F$ denotes the attribute vector of node v_i, $f_\theta(\cdot)$ denotes the MLP encoder parameterized by θ, $\mathbf{h}_i \in \mathbb{R}^d$ denotes the low-dimensional representation of node v_i, and d is the hidden dimension. To reduce the influence of the hidden dimension on the scale of the encoded representations, we follow the dimensional scaling strategy of Cetin et al. Cetin et al. (2022) and scale the encoded representations before the hyperbolic exponential map as $\mathbf{u}_i = \mathbf{h}_i/\sqrt{d}$. This operation reduces the dependence of the representation norm on the hidden dimension, thereby providing input representations with a consistent scale for the subsequent hyperbolic mapping.

We then map the scaled node representations into a Lorentz space with fixed negative curvature and jointly characterize local neighborhood relations and global center relations in this shared hyperbolic representation space. Given a fixed negative curvature $-c$, where $c > 0$, the d-dimensional Lorentz hyperbolic space is defined as:

$$\mathbb{H}_c^d = \left\{ \mathbf{z} \in \mathbb{R}^{d+1} : \langle \mathbf{z}, \mathbf{z} \rangle_{\mathcal{L}} = -\frac{1}{c},\ z_0 > 0 \right\}, \tag{2}$$

where $\langle \cdot, \cdot \rangle_{\mathcal{L}}$ denotes the Lorentz inner product, defined as:

$$\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}} = \mathbf{x}^\top \mathbf{J} \mathbf{y}, \qquad \mathbf{J} = \mathrm{diag}(-1, 1, \ldots, 1). \tag{3}$$

where z_0 denotes the time component of the hyperbolic representation, and the constraint $z_0 > 0$ restricts the node representations to the upper sheet of the two-sheeted hyperboloid.

Since the Euclidean encoded representations do not directly satisfy the hyperboloid constraint, we transform them into hyperbolic representations through the tangent space at the origin and the exponential map. For any $\mathbf{z} \in \mathbb{H}_c^d$, its tangent space is defined as:

$$T_{\mathbf{z}} \mathbb{H}_c^d = \left\{ \mathbf{v} \in \mathbb{R}^{d+1} : \langle \mathbf{v}, \mathbf{z} \rangle_{\mathcal{L}} = 0 \right\}.$$

Let the hyperbolic origin be $\mathbf{o} = (1/\sqrt{c}, 0, \ldots, 0)$, and extend the scaled Euclidean representation to the tangent vector $\widehat{\mathbf{u}}_i = (0, \mathbf{u}_i)$. Since $\langle \widehat{\mathbf{u}}_i, \mathbf{o} \rangle_{\mathcal{L}} = 0$, we have $\widehat{\mathbf{u}}_i \in T_{\mathbf{o}} \mathbb{H}_c^d$. The shared hyperbolic representation of node v_i is then obtained through the exponential map at the origin:

$$\mathbf{z}_i = \mathrm{Exp}_{\mathbf{o}}^c(\widehat{\mathbf{u}}_i) = \left(\frac{\cosh(\sqrt{c}\|\mathbf{u}_i\|_2)}{\sqrt{c}}, \frac{\sinh(\sqrt{c}\|\mathbf{u}_i\|_2)}{\sqrt{c}\|\mathbf{u}_i\|_2} \mathbf{u}_i \right). \tag{4}$$

After obtaining the shared hyperbolic representations, we adopt the standard geodesic distance in the Lorentz model Chami et al. (2019); Nickel & Kiela (2018) to measure the geometric difference between any two hyperbolic representations: $d_{\mathbb{H}}^c(\mathbf{z}_i, \mathbf{z}_j) = \frac{1}{\sqrt{c}} \operatorname{arcosh}\left(-c \langle \mathbf{z}_i, \mathbf{z}_j \rangle_{\mathcal{L}}\right)$, where $\operatorname{arcosh}(\cdot)$ denotes the inverse hyperbolic cosine function. This distance characterizes the intrinsic distance between two node representations on the Lorentz manifold.

Based on the above geodesic distance, we further construct a unified hyperbolic affinity that transforms the geodesic distance in hyperbolic space into a bounded and continuous affinity representation, while characterizing local neighborhood relations and global center relations under a consistent hyperbolic geometric metric.

Definition 1 *For any $\mathbf{p}, \mathbf{q} \in \mathbb{H}_c^d$, their unified hyperbolic affinity is defined as*

$$s_{\mathbb{H}}(\mathbf{p}, \mathbf{q}) = 2 \exp\left(-\gamma d_{\mathbb{H}}^c(\mathbf{p}, \mathbf{q})^2\right) - 1, \tag{5}$$

where $\gamma > 0$ is the affinity decay coefficient that controls the decay of the affinity as the hyperbolic geodesic distance increases.

According to Eq. (5), $s_{\mathbb{H}}(\mathbf{p}, \mathbf{q}) \in (-1, 1]$, and its value decreases monotonically as the hyperbolic geodesic distance increases. When two hyperbolic representations coincide, their affinity reaches the maximum value of 1; as the geodesic distance between them increases, the affinity gradually approaches -1.

3.3 Hyperbolic Structural Affinity Modeling

Structural associations in graphs provide important signals for characterizing node anomalies. Existing affinity-based methods typically characterize structural relations between nodes using similarity measures in Euclidean space and identify anomalies based on differences in node affinity Qiao & Pang (2023); Zhang et al. (2026). In contrast, we leverage the unified hyperbolic affinity defined in HAL and propose Hyperbolic Structural Affinity Modeling (HSAM), which constructs positive and negative structural relation pairs in the shared Lorentz space and characterizes structural association patterns between nodes using hyperbolic affinity.

Specifically, for a node v_i, we first construct its positive structural relations according to the observed graph topology. Its positive neighborhood set is defined as $\mathcal{N}_i^+ = \{v_j \mid (v_i, v_j) \in \mathcal{E}\}$. We treat the neighbors directly connected to node v_i as positive samples and compute the hyperbolic structural affinity index of node v_i using Eq. (5):

$$A_i^+ = \frac{1}{|\mathcal{N}_i^+|} \sum_{v_j \in \mathcal{N}_i^+} s_{\mathbb{H}}(\mathbf{z}_i, \mathbf{z}_j), \tag{6}$$

where $s_{\mathbb{H}}(\mathbf{z}_i, \mathbf{z}_j)$ denotes the hyperbolic affinity between node v_i and the positive sample v_j. A_i^+ denotes the positive hyperbolic structural affinity index of node v_i, which characterizes its overall affinity with the positive samples.

Similarly, we randomly sample negative samples from the non-neighboring nodes of v_i, and denote the resulting negative neighborhood set as $\mathcal{N}_i^- = \{v_k \mid (v_i, v_k) \in \mathcal{E}^-\}$. Accordingly, the negative hyperbolic structural affinity index of node v_i is defined as

$$A_i^- = \frac{1}{|\mathcal{N}_i^-|} \sum_{v_k \in \mathcal{N}_i^-} s_{\mathbb{H}}(\mathbf{z}_i, \mathbf{z}_k), \tag{7}$$

Based on the positive and negative hyperbolic structural affinity indices, we employ binary cross-entropy to discriminate between the two types of neighborhood relations. The optimization objective of HSAM is defined as

$$\mathcal{L}_{\text{HSAM}} = -\frac{1}{2N} \sum_{i=1}^{N} \left[\log \sigma(A_i^+) + \log\left(1 - \sigma(A_i^-)\right)\right], \tag{8}$$

where $\sigma(\cdot)$ denotes the Sigmoid function. This objective enlarges the discrepancy between positive and negative structural relations in terms of their hyperbolic affinity responses, enabling the shared hyperbolic representations to preserve discriminative structural association patterns and enhancing the identification of structural anomalies in nodes.

3.4 Hyperbolic Center Affinity Constraint

Local neighborhood relations provide important cues for characterizing node anomalies. However, complex anomaly patterns in real-world graphs may lead to local consistency deception. For example, when anomalous nodes form strong local connections with each other, their local relations may exhibit high consistency, thereby weakening the anomaly signals captured solely from neighborhood information Wu & Gu (2025a). To address this issue, we propose the Hyperbolic Center Affinity Constraint (HCAC) to capture global anomaly patterns beyond local neighborhoods.

Algorithm 1 UHAL-GAD Framework

Require: Attributed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, number of training epochs T, hidden dimension d, curvature parameter c, affinity decay coefficient γ, and trade-off parameter β
Ensure: Node anomaly scores $\{S_i\}_{i=1}^N$
Initialize the shared encoder f_θ
for $t = 1$ **to** T **do**
  Obtain the shared hyperbolic representations $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^N$ via Eqs. (1) and (4)
  Compute A_i^+ and A_i^- via Eqs. (5), (6), and (7), and calculate $\mathcal{L}_{\mathrm{HSAM}}$ via Eq. (8)
  Obtain $\boldsymbol{\mu}_{\mathbb{H}}$ via Eq. (9), and compute A_i^{g} and D_i^{HCA} via Eqs. (10) and (11)
  Calculate $\mathcal{L}_{\mathrm{HCAC}}$ via Eq. (12)
  Compute the joint loss via Eq. (13), backpropagate, and update θ
end for
Compute the node anomaly scores $\{S_i\}_{i=1}^N$ via Eq. (14)
return $\{S_i\}_{i=1}^N$

Specifically, HCAC and HSAM share the hyperbolic representations $\mathbf{z}_i$ obtained by HAL. HCAC first performs Lorentz-normalized aggregation over the hyperbolic representations of all nodes to obtain a global Lorentz center. Let the sum of node representations be $\mathbf{m} = \sum_{i=1}^N \mathbf{z}_i$, then the global Lorentz center is defined as

$$\boldsymbol{\mu}_{\mathbb{H}} = \frac{\mathbf{m}}{\sqrt{-c\langle \mathbf{m}, \mathbf{m} \rangle_{\mathcal{L}}}}, \tag{9}$$

where $\boldsymbol{\mu}_{\mathbb{H}} \in \mathbb{H}_c^d$ denotes the global center in the shared Lorentz space. After normalization, $\boldsymbol{\mu}_{\mathbb{H}}$ satisfies $\langle \boldsymbol{\mu}_{\mathbb{H}}, \boldsymbol{\mu}_{\mathbb{H}} \rangle_{\mathcal{L}} = -1/c$, and thus lies on the same Lorentz manifold as the node representations.

After obtaining the global Lorentz center, we treat it as a geometric reference for the global representation and use the hyperbolic affinity defined in HAL to characterize the affinity between each node and the center. In this way, the neighborhood affinity in HSAM and the center affinity in HCAC are both measured by the same hyperbolic affinity function $s_{\mathbb{H}}(\cdot, \cdot)$. Specifically, the hyperbolic center affinity of node v_i is defined as

$$A_i^{\mathrm{g}} = s_{\mathbb{H}}(\mathbf{z}_i, \boldsymbol{\mu}_{\mathbb{H}}) = 2 \exp\left(-\gamma d_{\mathbb{H}}^c(\mathbf{z}_i, \boldsymbol{\mu}_{\mathbb{H}})^2\right) - 1, \tag{10}$$

where A_i^{g} denotes the hyperbolic center affinity of node v_i. A larger value indicates a stronger affinity between the node and the global center. We further define the Hyperbolic Center Affinity Discrepancy as

$$D_i^{\mathrm{HCA}} = 1 - A_i^{\mathrm{g}} = 2\left[1 - \exp\left(-\gamma d_{\mathbb{H}}^c(\mathbf{z}_i, \boldsymbol{\mu}_{\mathbb{H}})^2\right)\right], \tag{11}$$

Therefore, D_i^{HCA} characterizes the hyperbolic affinity discrepancy between a node and the global center in a bounded form. Based on the above hyperbolic center affinity discrepancy, HCAC imposes a global constraint on the affinity relation between each node and the global center. Its optimization objective is defined as

$$\mathcal{L}_{\mathrm{HCAC}} = \frac{1}{N} \sum_{i=1}^{N} D_i^{\mathrm{HCA}}, \tag{12}$$

This objective shares the same hyperbolic affinity function with HSAM, allowing neighborhood relations and center relations to be modeled under a unified hyperbolic affinity metric.

3.5 Joint Optimization and Anomaly Scoring

Based on HSAM and HCAC, we jointly learn neighborhood affinity relations and center affinity relations under a unified hyperbolic affinity metric to enhance the discriminability of node relations. The overall training objective is defined as

$$\mathcal{L} = \mathcal{L}_{\mathrm{HSAM}} + \beta \mathcal{L}_{\mathrm{HCAC}}, \tag{13}$$

where β is a trade-off hyperparameter that balances the two relation levels of neighborhood affinity and center affinity. During inference, the final anomaly score of node v_i is defined as

$$S_i = S_i^{\mathrm{HSAM}} + \beta S_i^{\mathrm{HCAC}} = -A_i^+ + \beta D_i^{\mathrm{HCA}}, \tag{14}$$

Table 1: Statistics of the eight benchmark datasets. R/I denotes real/injected anomalies.

Dataset	Type	R/I	Nodes	Edges	Attributes	Anomalies (Rate)
Facebook	Social Network	R	1,081	55,106	576	27 (2.5%)
Cora	Citation Network	I	2,708	5,429	1,433	150 (5.5%)
Citeseer	Citation Network	I	3,327	4,732	3,703	150 (4.5%)
Photo	Co-purchase	R	7,535	119,043	745	698 (9.3%)
Amazon	Co-review	R	10,224	175,608	25	693 (6.8%)
Reddit	Social Network	R	10,984	168,016	64	366 (3.3%)
Amazon-all	Co-review	R	11,944	4,398,392	25	821 (6.9%)
T-Finance	Transaction	R	39,357	21,222,543	10	1,803 (4.6%)

where a larger S_i indicates a higher degree of anomaly for node v_i.

Algorithm 1 summarizes the complete procedure.

4 Experiment

4.1 Experimental Settings

1) Datasets: To evaluate the effectiveness of the proposed method, we conduct experiments on eight benchmark datasets, including Facebook Xu et al. (2022), Cora and Citeseer Sen et al. (2008); Liu et al. (2021b), Photo McAuley et al. (2015); Qiao et al. (2024), Amazon Dou et al. (2020), Reddit Kumar et al. (2019), Amazon-all Dou et al. (2020); Chen et al. (2022), and T-Finance Tang et al. (2022). Detailed statistics of these datasets are presented in Table 1.

2) Comparison Methods: We select eight graph anomaly detection methods for comparison, including reconstruction-based methods DOMINANT Ding et al. (2019) and ComGA Luo et al. (2022), a contrastive learning-based method CoLA Liu et al. (2021b), affinity modeling-based methods TAM Qiao & Pang (2023) and GCTAM Zhang et al. (2026), and other graph representation learning-based detection methods GADAM Chen et al. (2024), SmoothGNN Dong et al. (2025), and MI-GAD Lian et al. (2026).

4.2 Implementation Details

All experiments are conducted on a computer equipped with an Intel Core i9-13900K CPU, an NVIDIA GeForce RTX 4090 GPU, and 64 GB of system memory. We implement the proposed method using PyTorch, employ a single-layer linear encoder, and train the model using the Adam optimizer for a maximum of 100 epochs. The search ranges for hyperparameters β and γ are $[0.1, 1]$ and $[10^{-4}, 1]$, respectively, while the candidate values for the curvature parameter c are $\{0.1, 1, 4, 16, 32\}$.

4.3 Performance Comparison

Table 2 reports the comparison results of our method against eight representative baselines on eight datasets. Based on the experimental results, we have the following three main observations. (1) Except for AUROC on the Facebook dataset, UHAL-GAD achieves the best performance on all the remaining results. In particular, on the Citeseer dataset, UHAL-GAD improves AUROC and AUPRC by 4.44 and 30.60 percentage points, respectively, over the corresponding best baselines, demonstrating its effectiveness across different graph anomaly detection scenarios. (2) Existing methods characterize anomalies from different perspectives, such as reconstruction errors, local contrast, or node affinity, and differ in how they model node relations, resulting in certain performance variations across different datasets. In contrast, UHAL-GAD jointly models structural relations and global center relations through a shared hyperbolic affinity, achieving better overall performance when compared with different types of baselines. (3) On relatively large-scale datasets such as Amazon-all and T-Finance, some baselines encounter out-of-memory (OOM) issues, whereas UHAL-GAD achieves the best results in both AUROC and AUPRC, indicating its good applicability and scalability for large-scale graph anomaly detection tasks.

Table 2: AUROC (%) and AUPRC (%) on eight datasets (mean ± std). Results are reported as mean ± std, with the best and second-best results colored red and blue, respectively. OOM denotes out-of-memory.

Metric	Method	Dataset							
		Facebook	Cora	Citeseer	Photo	Amazon	Reddit	Amazon-all	T-Finance
AUROC	DOMINANT	59.03±1.30	81.52±0.46	82.43±0.33	45.13±0.59	52.47±4.32	54.72±1.14	51.59±11.11	OOM
	CoLA	84.85±2.51	81.44±1.04	85.73±1.03	63.57±3.37	18.35±0.64	**61.83±1.02**	16.96±0.50	22.65±4.72
	ComGA	13.69±1.59	89.94±0.28	**92.62±0.93**	48.88±1.80	52.70±5.41	57.16±0.37	OOM	OOM
	TAM	89.18±0.72	77.49±0.89	86.66±0.63	58.79±0.74	73.18±0.74	57.41±0.09	80.39±0.59	OOM
	GADAM	95.31±0.67	**94.16±0.22**	90.74±0.34	**79.84±1.45**	61.17±0.67	58.51±0.48	44.71±2.86	31.61±2.08
	SmoothGNN	43.23±0.54	55.31±0.88	42.21±2.74	49.82±1.68	84.73±0.59	59.25±0.28	83.94±0.97	75.02±1.98
	GCTAM	91.14±1.25	38.62±0.85	38.22±1.42	58.12±1.85	83.49±1.42	56.40±0.18	OOM	OOM
	MI-GAD	**98.34±0.07**	91.60±0.42	89.90±0.19	70.33±1.31	**88.79±0.43**	61.48±0.23	**91.03±0.14**	**94.85±0.27**
	Ours	**97.02±0.09**	**94.72±0.14**	**97.06±0.06**	**85.62±0.67**	**90.10±0.06**	**63.39±0.08**	**91.55±0.06**	**95.47±0.34**
AUPRC	DOMINANT	3.57±0.23	38.82±1.33	37.00±1.92	8.70±0.09	6.42±0.58	4.01±0.08	6.61±1.32	OOM
	CoLA	23.53±4.62	26.67±5.12	35.63±4.08	15.57±0.84	4.14±0.05	4.59±0.22	4.30±0.09	2.89±0.18
	ComGA	1.35±0.02	48.16±2.07	44.33±3.42	9.68±0.78	6.50±0.67	3.79±0.03	OOM	OOM
	TAM	20.61±0.78	36.09±2.64	41.44±1.55	11.71±0.27	37.31±1.87	4.31±0.05	38.54±3.17	OOM
	GADAM	35.11±4.42	**70.80±0.21**	30.24±1.35	**21.13±1.08**	8.27±0.33	4.81±0.22	5.74±0.39	3.41±0.15
	SmoothGNN	2.13±0.09	6.13±1.54	3.64±0.24	9.51±1.26	36.95±2.34	4.01±0.53	36.70±2.59	13.49±0.62
	GCTAM	18.75±5.19	6.05±0.55	4.20±0.18	10.58±0.47	46.54±9.23	3.79±0.06	OOM	OOM
	MI-GAD	**53.59±1.08**	59.23±1.58	**48.76±2.50**	16.26±0.65	**71.45±3.75**	**4.87±0.04**	**78.53±0.42**	**64.81±3.79**
	Ours	**73.56±0.61**	**71.20±0.79**	**79.36±0.76**	**26.94±0.98**	**78.13±0.57**	**5.74±0.08**	**79.60±0.41**	**72.19±3.26**

Table 3: Ablation results on four graph datasets. The best results are highlighted in **red**.

Dataset	Facebook		Photo		Reddit		Amazon-all	
Metric	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
w/o HAL	95.49	50.14	75.07	18.90	55.45	4.74	55.48	19.52
w/o HSAM	45.83	2.33	50.09	10.57	54.72	3.87	60.07	8.20
w/o HCAC	93.32	64.19	84.04	25.50	58.38	4.57	64.84	11.40
Ours	**97.02**	**73.56**	**85.62**	**26.94**	**63.39**	**5.74**	**91.55**	**79.60**

4.4 Ablation Study

To evaluate the effectiveness of each component in UHAL-GAD, we design three ablation variants. As shown in Table 3, the full model achieves the best AUROC and AUPRC performance on all datasets. Specifically, replacing hyperbolic geometry with Euclidean geometry leads to performance degradation on all datasets, indicating that hyperbolic geometry provides a more effective representation space for relation modeling in anomaly detection. Removing either HSAM or HCAC also results in performance degradation. Specifically, HSAM effectively characterizes hyperbolic structural affinity differences between nodes in the shared Lorentz space, while HCAC further captures the hyperbolic affinity discrepancy between each node and the global center, thereby providing effective information for anomaly discrimination. Overall, unified hyperbolic affinity learning effectively improves the anomaly detection performance of the model.

4.5 Anomaly Score Distribution Visualization

To intuitively illustrate the anomaly detection capability of UHAL-GAD, Fig. 2 presents the anomaly score distributions of normal and abnormal nodes on four datasets. We observe that the anomaly scores of normal nodes are mainly concentrated in the low-score region, whereas abnormal nodes generally receive higher scores. Overall, UHAL-GAD can effectively distinguish normal nodes from abnormal nodes, demonstrating its anomaly detection capability.

4.6 Parameter Analysis

We investigate the effects of the trade-off parameter β in the overall objective function, affinity decay coefficient γ in hyperbolic affinity and the hyperbolic curvature parameter c on model performance. As shown in Fig. 3 and Fig. 4, we make the following observations: (1) Different datasets

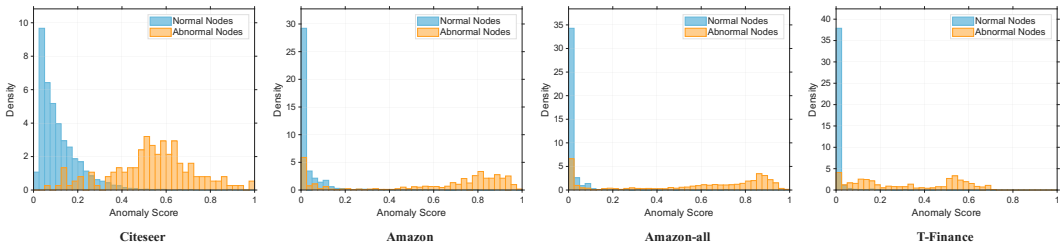

Figure 2: Visualization of anomaly score distributions.

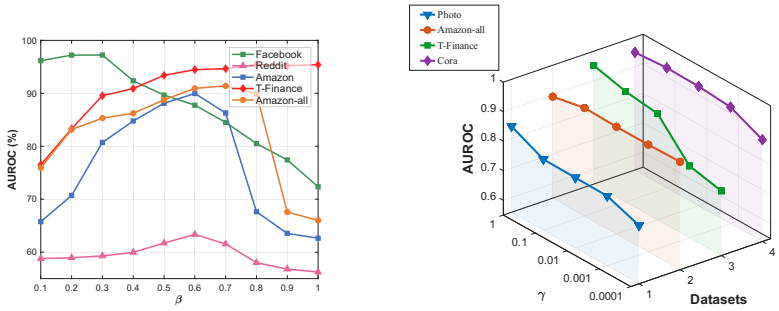

Figure 3: Parameter sensitivity analysis of β and γ.

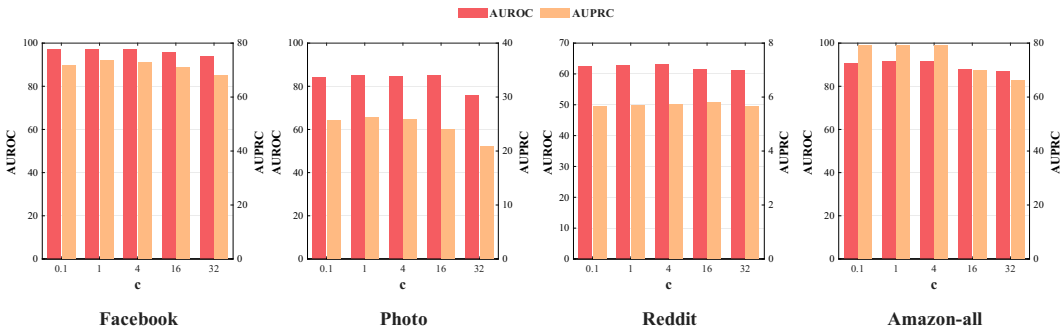

Figure 4: Curvature sensitivity analysis.

exhibit varying sensitivities to β, indicating that properly balancing the local and global losses is crucial for anomaly detection. (2) A small γ weakens the ability to distinguish node relationships. As γ increases, AUROC gradually stabilizes. (3) when c is small, hyperbolic space gradually approaches Euclidean space, which may limit the representational advantages of hyperbolic geometry. In contrast, an excessively large c can lead to numerical instability during training.

5 Conclusion

In this work, we proposed UHAL-GAD, a unified hyperbolic affinity learning framework for unsupervised node-level graph anomaly detection. By learning shared Lorentz representations and a unified hyperbolic affinity, the framework jointly characterizes structural relationships between nodes and their relationships with the global center. The two modules are jointly optimized, and their anomaly scores are fused to identify complex anomalous patterns. Extensive experiments on eight real-world datasets demonstrate the effectiveness of UHAL-GAD.

LLM USAGE STATEMENT

During the preparation of this manuscript, we used Large Language Models (LLMs) to assist with language organization and text refinement, aiming to improve the accuracy, readability, and overall coherence of the writing. Their primary uses included optimizing sentence structures, checking grammar, and refining academic wording.

The authors take full responsibility for the research content and the final manuscript. All text generated or revised with LLM assistance was manually reviewed and edited by the authors to ensure accuracy and compliance with academic ethics and integrity standards.

REFERENCES

Leman Akoglu, Hanghang Tong, and Danai Koutra. Graph based anomaly detection and description: a survey. *Data mining and knowledge discovery*, 29(3):626–688, 2015.

Edoardo Cetin, Benjamin Chamberlain, Michael Bronstein, and Jonathan J Hunt. Hyperbolic deep reinforcement learning. *arXiv preprint arXiv:2210.01542*, 2022.

Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. *Advances in neural information processing systems*, 32, 2019.

Bo Chen, Jing Zhang, Xiaokang Zhang, Yuxiao Dong, Jian Song, Peng Zhang, Kaibo Xu, Evgeny Kharlamov, and Jie Tang. Gccad: Graph contrastive coding for anomaly detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(8):8037–8051, 2022.

Jingyan Chen, Guanghui Zhu, Chunfeng Yuan, and Yihua Huang. Boosting graph anomaly detection with adaptive message passing. In *International Conference on Learning Representations*, 2024.

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

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

Yingtong Dou, Zhiwei Liu, Li Sun, Yutong Deng, Hao Peng, and Philip S Yu. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pp. 315–324, 2020.

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

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

Yali Fu, Jindong Li, Jiahong Liu, Qianli Xing, Qi Wang, and Irwin King. Hc-glad: dual hyperbolic contrastive learning for unsupervised graph-level anomaly detection. *Neural Networks*, pp. 109009, 2026.

Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. *Advances in neural information processing systems*, 31, 2018.

Jing Gu and Dongmian Zou. Three revisits to node-level graph anomaly detection: Outliers, message passing and hyperbolic neural networks. In *Learning on Graphs Conference*, pp. 14–1. PMLR, 2024.

Junwei He, Qianqian Xu, Yangbangyan Jiang, Zitai Wang, and Qingming Huang. Ada-gad: Anomaly-denoised autoencoders for graph anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 8481–8489, 2024.

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

Ming Jin, Yixin Liu, Yu Zheng, Lianhua Chi, Yuan-Fang Li, and Shirui Pan. Anemone: Graph anomaly detection with multi-scale contrastive learning. In *Proceedings of the 30th ACM international conference on information & knowledge management*, pp. 3122–3126, 2021.

Junghoon Kim, Yeonjun In, Kanghoon Yoon, Junmo Lee, and Chanyoung Park. Class label-aware graph anomaly detection. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pp. 4008–4012, 2023.

Srijan Kumar, Xikun Zhang, and Jure Leskovec. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1269–1278, 2019.

Jie Lian, Zhihao Wu, Jielong Lu, Jiajun Yu, Qianqian Shen, and Haishuai Wang. Beyond local patterns: Multiscale inconsistency learning for graph anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 15216–15224, 2026.

Qi Liu, Maximilian Nickel, and Douwe Kiela. Hyperbolic graph neural networks. *Advances in neural information processing systems*, 32, 2019.

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

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

Xuexiong Luo, Jia Wu, Amin Beheshti, Jian Yang, Xiankun Zhang, Yuan Wang, and Shan Xue. Comga: Community-aware attributed graph anomaly detection. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, pp. 657–665, 2022.

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

Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. *Advances in neural information processing systems*, 30, 2017.

Maximillian Nickel and Douwe Kiela. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *International conference on machine learning*, pp. 3779–3788. PMLR, 2018.

Junjun Pan, Yixin Liu, Yizhen Zheng, and Shirui Pan. Prem: A simple yet effective approach for node-level graph anomaly detection. In *2023 IEEE International Conference on Data Mining (ICDM)*, pp. 1253–1258. IEEE, 2023.

Junjun Pan, Yixin Liu, Xin Zheng, Yizhen Zheng, Alan Wee-Chung Liew, Fuyi Li, and Shirui Pan. A label-free heterophily-guided approach for unsupervised graph fraud detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 12443–12451, 2025.

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

Hezhe Qiao, Qingsong Wen, Xiaoli Li, Ee-Peng Lim, and Guansong Pang. Generative semi-supervised graph anomaly detection. *Advances in neural information processing systems*, 37: 4660–4688, 2024.

Hezhe Qiao, Hanghang Tong, Bo An, Irwin King, Charu Aggarwal, and Guansong Pang. Deep graph anomaly detection: A survey and new perspectives. *IEEE Transactions on Knowledge and Data Engineering*, 37(9):5106–5126, 2025.

Amit Roy, Juan Shu, Jia Li, Carl Yang, Olivier Elshocht, Jeroen Smeets, and Pan Li. Gad-nr: Graph anomaly detection via neighborhood reconstruction. In *Proceedings of the 17th ACM international conference on web search and data mining*, pp. 576–585, 2024.

Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–106, 2008.

Jianheng Tang, Jiajin Li, Ziqi Gao, and Jia Li. Rethinking graph neural networks for anomaly detection. In *International conference on machine learning*, pp. 21076–21089. PMLR, 2022.

Wensen Wu and Yijun Gu. Advancing unsupervised graph anomaly detection: A multi-level contrastive learning framework to mitigate local consistency deception. *Neurocomputing*, 646: 130507, 2025a.

Wensen Wu and Yijun Gu. Graph anomaly detection via multi-level information alignment and decoupling. *Knowledge-Based Systems*, 326:114045, 2025b.

Yongji Wu, Defu Lian, Yiheng Xu, Le Wu, and Enhong Chen. Graph convolutional networks with markov random field reasoning for social spammer detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 1054–1061, 2020.

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

Jiaqiang Zhang, Senzhang Wang, and Songcan Chen. Reconstruction enhanced multi-view contrastive learning for anomaly detection on attributed networks. *arXiv preprint arXiv:2205.04816*, 2022.

Xiong Zhang, Hong Peng, Zhenli He, Cheng Xie, Xin Jin, and Hua Jiang. Gctam: Global and contextual truncated affinity combined maximization model for unsupervised graph anomaly detection. *arXiv preprint arXiv:2603.01806*, 2026.

Yu Zheng, Ming Jin, Yixin Liu, Lianhua Chi, Khoa T Phan, and Yi-Ping Phoebe Chen. Generative and contrastive self-supervised learning for graph anomaly detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12220–12233, 2021.