

GLADMamba: Unsupervised Graph-Level Anomaly Detection Powered by Selective State Space Model

Yali Fu^{1*}, Jindong Li^{2*}, Qi Wang^{1,4} (✉), and Qianli Xing³

¹ School of Artificial Intelligence, Jilin University, Changchun, China
fuy123@mails.jlu.edu.cn, qiwang@jlu.edu.cn

² Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China
jli839@connect.hkust-gz.edu.cn

³ College of Computer Science and Technology, Jilin University, Changchun, China
qianlixing@jlu.edu.cn

⁴ Engineering Research Center of Knowledge-Driven Human-Machine Intelligence,
Ministry of Education, China

Abstract. Unsupervised graph-level anomaly detection (UGLAD) is a critical and challenging task across various domains, such as social network analysis, anti-cancer drug discovery, and toxic molecule identification. However, existing methods often struggle to capture long-range dependencies efficiently and neglect the spectral information. Recently, selective state space models, particularly Mamba, have demonstrated remarkable advantages in capturing long-range dependencies with linear complexity and a selection mechanism. Motivated by their success across various domains, we propose GLADMamba, a novel framework that adapts the selective state space model into UGLAD field. We design a View-Fused Mamba (VFM) module with a Mamba-Transformer-style architecture to efficiently fuse information from different graph views with a selective state mechanism. We also design a Spectrum-Guided Mamba (SGM) module with a Mamba-Transformer-style architecture to leverage the Rayleigh quotient to guide the embedding refinement process, considering the spectral information for UGLAD. GLADMamba can dynamically focus on anomaly-related information while discarding irrelevant information for anomaly detection. To the best of our knowledge, this is the first work to introduce Mamba and explicit spectral information to UGLAD. Extensive experiments on 12 real-world datasets demonstrate that GLADMamba outperforms existing state-of-the-art methods, achieving superior performance in UGLAD. The code is available at <https://github.com/Yali-Fu/GLADMamba>.

Keywords: Unsupervised Graph-Level Anomaly Detection · Selective State Space Model · Graph Spectrum · Graph Neural Networks

* Equal contribution.




Fig. 1. The key factors related to graph anomalies.

1 Introduction

Unsupervised graph-level anomaly detection (UGLAD) is a prevalent task in numerous real-world scenarios, including social network analysis, drug discovery, and toxic molecule identification [17,24,25]. Its goal is to identify graphs that exhibit significantly different patterns from the majority, which often represent unexpected events or behaviors [22,23,24]. Unlike supervised approaches, unsupervised methods don’t require labeled data, making them more adaptable to real-world scenarios where labeled anomalies are scarce or costly to obtain. Despite the existence of many excellent methods, several challenges still exist in this field.

Most GNN-based methods are inherently limited by the over-squashing issue, which restricts their ability to effectively model long-range dependencies between nodes [1,12,31]. This limitation hampers the information propagation across distant nodes, making it challenging to capture anomaly-related patterns and ultimately weakening detection performance. Although some studies [17,40,47] have incorporated Transformer into graph anomaly detection to mitigate this issue, the quadratic computational complexity of the attention mechanism significantly restricts the scalability of these methods, particularly for large-scale graphs. Additionally, as depicted in Fig. 1(a), relying solely on a single characteristic is insufficient for comprehensively capturing anomalies [17,22]. Thus, how to efficiently integrate information from multiple aspects (e.g., attributes and topologies) poses a critical challenge.

Furthermore, we observe spectral differences between normal and abnormal graphs in Fig. 1(b). And as demonstrated in [4,35,37,43], the energy distribution in the spectral domain shifts from low-frequency to high-frequency regions as the anomaly degree increases. However, most existing methods primarily focus on anomaly information in the spatial domain, neglecting the interaction with the spectral domain and failing to account for spectral differences between normal and anomalous graphs. Although a few works have explored the use of spectral information for graph-level anomaly detection, they rely on labeled data during training [4]. In the context of UGLAD, this remains an unexplored area, underscoring the need for further research and innovation.

Recently, selective state space models, especially Mamba [7], originally designed for sequence modeling, have demonstrated remarkable advantages in capturing long-range dependencies with linear computational complexity and a selection mechanism. These advantages have been extensively validated across various domains [1,5,29,39,44], making Mamba a strong candidate for addressing the challenges in UGLAD. Building on these strengths, we take the first step in integrating Mamba into UGLAD, unlocking its potential for more effective graph modeling while significantly improving computational efficiency and anomaly detection performance.

We propose a novel unsupervised graph-level anomaly detection framework that is powered by selective state space model (Mamba), named GLADMamba. We firstly design a View-Fused Mamba (VFM) module, which is a Mamba-Transformer-style architecture and can efficiently fuse different graph views by Mamba’s selective state transition mechanism. Benefiting from the architecture and advantages of Mamba, VFM excels at integrating multi-view information and capturing long-range dependencies, showcasing powerful embedding capabilities while maintaining linear time complexity. Furthermore, we design a Spectrum-Guided Mamba (SGM) module, which is also a Mamba-Transformer-style architecture and leverages explicit spectral information to guide the embedding refinement process for UGLAD. By the intrinsic relationship between spectral energy and graph anomalies, SGM establishes interactions between the spatial and spectral domains. Concretely, it employs the Rayleigh quotient to discretize the continuous state space of SGM, making system parameters spectrum-dependent. Therefore, the Rayleigh quotient can guide the update of latent states in SGM, enabling the model to selectively focus on anomaly-related information and filter redundant information to enhance detection performance.

By the selection mechanism, GLADMamba can dynamically adjust its learning strategy based on specific characteristics of different input graphs, adaptively capturing anomaly patterns. Extensive experiments on 12 real-world datasets demonstrate the effectiveness of GLADMamba. And GLADMamba consistently outperforms state-of-the-art methods in the unsupervised graph-level anomaly detection task. Our key contributions are summarized as follows:

- We propose a novel model, named GLADMamba, which adapts the selective state space model (Mamba) for unsupervised graph-level anomaly detection. To the best of our knowledge, this is the first work to introduce Mamba to UGLAD.
- We design a View-Fused Mamba (VFM) module for efficient multi-view fusion, boosting detection accuracy. In addition, we design a Spectrum-Guided Mamba (SGM) module, the first spectrum-guided method in this field to preserve anomaly-related patterns via selective state updates. This work pioneers spectrum-guided Mamba architecture for UGLAD.
- We conduct extensive experiments on 12 real-world datasets, demonstrating that GLADMamba achieves state-of-the-art performance in the unsupervised graph-level anomaly detection task.

2 Related Work

2.1 Graph-Level Anomaly Detection

Graph-level anomaly detection aims to identify anomalous graphs within a graph set, where anomalies typically represent rare but critical patterns compared to normal graphs [6,25,45]. Conventional methods generally involve two main steps: first, a graph kernel, such as the Weisfeiler-Lehman kernel (WL) [32] or propagation kernel (PK) [28], is used to learn representations; second, an anomaly detection algorithm, such as isolation forest (iF) [18], one-class support vector machine (OCSVM) [26], or local outlier factor (LOF) [3], is applied to detect anomalous graphs based on the extracted graph representations.

In addition, graph neural networks (GNNs) [11,14,36,38] have attracted significant attention due to their remarkable performance in dealing with various graph data and tasks [16,19,20,21]. Thus, various types of GNNs are employed as the backbone to conduct graph-level anomaly detection [6,22,23,24,45]. For example, GOOD-D [22] designs a novel graph data augmentation method and employs contrastive learning at different levels for graph-level anomaly detection. CVTGAD [17] employs a lightweight Transformer with an attention mechanism to model intra-graph and inter-graph node relationships, improving detection performance.

2.2 State Space Models

State Space Models (SSMs) [13] are classical frameworks for dynamic systems, while Structured SSMs (S4) [10] enhance SSMs with efficient long-sequence modeling. Mamba [7,30] builds on S4 by introducing a selection mechanism, enabling dynamic adaptation and improved efficiency. Together, they represent an evolution from traditional SSMs to modern, high-performance sequence modeling architectures. Beyond its core advancements, Mamba has demonstrated promising applications in various domains. It has been explored in computer vision [39], multimodal learning [29], audio [5], and natural language processing [44], showcasing its versatility across different modalities. Mamba has also shown preliminary and promising applications in graph representation learning [1]. For example, DG-Mamba [42] introduces a kernelized dynamic message-passing operator and a self-supervised regularization based on the principle of relevant information to improve efficiency, expressiveness, and robustness in dynamic graph learning. MOL-Mamba [12] enhances molecular representations by integrating hierarchical structural reasoning and electronic correlation learning, designing a hybrid Mamba-Graph and Mamba-Transformer framework supported by collaborative training strategies.

3 Preliminaries

3.1 Problem Statement

A graph is represented as $G = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, where $\mathcal{V}$ denotes the set of nodes, $\mathcal{E}$ represents the set of edges, $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ denotes the adjacency matrix, and $|\mathcal{V}|$

is the number of nodes. The entry $\mathbf{A}_{i,j}$ of $\mathbf{A}$ is set to 1 if an edge exists between node v_i and node v_j ; otherwise, $\mathbf{A}_{i,j} = 0$. An attributed graph is defined as $G = (\mathcal{V}, \mathcal{E}, \mathbf{A}, \mathbf{F})$, where $\mathbf{F} \in \mathbb{R}^{|\mathcal{V}| \times d_f}$ is the feature matrix containing node attributes. Each row $\mathbf{f}$ of $\mathbf{F}$ corresponds to a feature vector of a node with d_f dimensions. The collection of graphs is denoted as $\mathcal{G} = \{G_1, G_2, \dots, G_m\}$, where m is the total number of graphs. This work addresses the GLAD problem under an unsupervised setting, where no labels are available for model training.

3.2 State Space Models & Mamba

State Space Models (SSMs) [7,10,30] model the dynamic evolution of continuous systems via latent state $h(t) \in \mathbb{R}^{N \times \tilde{L}}$, mapping input $x(t) \in \mathbb{R}^{\tilde{L}}$ to output $y(t) \in \mathbb{R}^{\tilde{L}}$ through the following state transition equation and observation equation:

$$h'(t) = \mathcal{A}h(t) + \mathcal{B}x(t), \quad y(t) = \mathcal{C}h(t), \quad (1)$$

where $\mathcal{A} \in \mathbb{R}^{N \times N}$ is the state transition matrix, $\mathcal{B} \in \mathbb{R}^{N \times 1}$ is the input matrix, $\mathcal{C} \in \mathbb{R}^{1 \times N}$ is the output matrix, N denotes the state size, and $\tilde{L}$ denotes the sequence length.

By the step size Δ , continuous parameters $(\mathcal{A}, \mathcal{B})$ are discretized into discrete forms $(\bar{\mathcal{A}}, \bar{\mathcal{B}})$ for practical application of SSMs. Discrete SSMs are described as follows:


$$h_t = \bar{\mathcal{A}}h_{t-1} + \bar{\mathcal{B}}x_t, \quad y_t = \mathcal{C}h_t. \quad (2)$$

The selective SSM (Mamba) [7] focuses on relevant information selectively by making parameters $(\bar{\mathcal{B}}, \mathcal{C}, \Delta)$ input-dependent.


3.3 Rayleigh Quotient

For a graph $G = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, let $\mathbf{D}$ be the diagonal degree matrix with $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$. Its Laplacian matrix $\mathbf{L}$ is defined as $\mathbf{L} = \mathbf{D} - \mathbf{A}$ (unnormalized) or $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ (normalized), where $\mathbf{I}$ is the identity matrix. The symmetric matrix $\mathbf{L}$ can be decomposed as $\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T$, where $\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_{|\mathcal{V}|})$ contains the corresponding eigenvalues sorted in ascending order, i.e., $0 \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{|\mathcal{V}|}$, and $\mathbf{U} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{|\mathcal{V}|})$ represents the orthonormal eigenvectors. Let $\mathbf{X} = (x_1, x_2, \dots, x_{|\mathcal{V}|})^T$ be the signal of G , and its graph Fourier transform is given by $\hat{\mathbf{X}} = (\hat{x}_1, \hat{x}_2, \dots, \hat{x}_{|\mathcal{V}|})^T = \mathbf{U}^T \mathbf{X}$. As demonstrated in [4,35,37], the Rayleigh quotient $R(\mathbf{L}, \mathbf{X})$ can represent the accumulated spectral energy, with a higher quotient value indicating more high-frequency components. We employ the following Rayleigh quotient without explicit eigenvalue decomposition for computational efficiency:

$$R(\mathbf{L}, \mathbf{X}) = \frac{\mathbf{X}^T \mathbf{L} \mathbf{X}}{\mathbf{X}^T \mathbf{X}} = \frac{\sum_{k=1}^{|\mathcal{V}|} \lambda_k \hat{x}_k^2}{\sum_{k=1}^{|\mathcal{V}|} \hat{x}_k^2} = \frac{\sum_{(i,j) \in \mathcal{E}} (x_i - x_j)^2}{2 \sum_{i \in \mathcal{V}} x_i^2}. \quad (3)$$



(a) The overall pipeline of the proposed GLADMamba.



(b) View-Fused Mamba (VFM).

(c) Spectrum-Guided Mamba (SGM).

Fig. 2. Overview of the proposed GLADMamba model: (a) illustrates the overall pipeline of the framework; (b) depicts the View-Fused Mamba (VFM) module, which efficiently integrates multi-view information; (c) shows the Spectrum-Guided Mamba (SGM) module, which is designed to guide embedding refinement process by Rayleigh quotient.

4 Methodology

4.1 Data Augmentation and Encoding

For data augmentation, we obtain the original view o and augmented view a by a perturbation-free graph augmentation strategy [17,22,34], which is tailored for anomaly detection. And the view o and a focus on feature and structure characteristics of the graph, respectively. After data augmentation, we employ two independent GNN encoders on two views to obtain node and graph embeddings. Taking the original view o as an example, the GNN encoder updates node embeddings in the l -th layer according to the following message passing rule:

$$\mathbf{h}_v^{(o,l)} = \text{UPDATE}^{(l-1)} \left(\mathbf{h}_v^{(o,l-1)}, \text{AGG}^{(l-1)} \left(\{\mathbf{h}_u^{(o,l-1)} : u \in \mathcal{N}(v)\} \right) \right), \quad (4)$$

where $\mathbf{h}_v^{(o,l)}$ denotes the representation of node v on the view o at the l -th layer, $\mathcal{N}(v)$ represents the set of neighboring nodes of v , AGG is the aggregation function to combine information from neighboring nodes, and UPDATE is the update function to generate a new node representation. And we have $\mathbf{h}_v^{(o,0)} = \mathbf{f}_v$. After the L -layer GNN encoder, we obtain the representation $\mathbf{h}_v^o$ of node v by concatenation operation and the representation $\mathbf{h}_G^o$ of graph G on the view o by the readout function:

$$\mathbf{h}_v^o = [\mathbf{h}_v^{(o,1)} || \dots || \mathbf{h}_v^{(o,L)}], \quad \mathbf{h}_G^o = \frac{1}{|\mathcal{V}_G|} \sum_{u \in \mathcal{V}_G} \mathbf{h}_u^o, \quad (5)$$

where $\mathcal{V}_G$ is the set of nodes in graph G . Let $\mathbf{H}^o$ denote the node representation matrix encoded from the original view o in the training/testing batch, where each row corresponds to the representation of a node. In the same manner, we obtain the representation $\mathbf{h}_v^a$ of node v , the representation $\mathbf{h}_G^a$ of graph G , and the node representation matrix $\mathbf{H}^a$ for view a .

4.2 View-Fused Mamba (VFM)

In the View-Fused Mamba (VFM) module, we propose a novel view fusion mechanism. Specifically, we deeply fuse different aspects of graph data (i.e., feature and structure information) captured by the original view o and augmented view a . Designed in a Mamba-Transformer-style architecture, the VFM efficiently captures and processes multi-view information for UGLAD. The detailed architecture of the View-Fused Mamba is depicted in Fig. 2(b).

Selective Parameterization in VF-SSM. For the parameter $\mathcal{A}$, we adopt HiPPO-LegS [8] for parameterization [7,9]. Before obtaining the input-dependent parameters $(\mathcal{B}, \mathcal{C}, \Delta)$, we sequentially process the node representations $\mathbf{H}^o$ and $\mathbf{H}^a$ from two views through the following operations: a layer normalization [2], a linear projection layer, a 1D convolutional layer, and a SiLU activation function. The formulas are as follows:

$$\mathbf{H}_{input}^o = \text{SiLU}(\text{Conv1D}(\text{Linear}(\text{LayerNorm}(\mathbf{H}^o)))), \quad (6)$$

$$\mathbf{H}_{input}^a = \text{SiLU}(\text{Conv1D}(\text{Linear}(\text{LayerNorm}(\mathbf{H}^a)))). \quad (7)$$

When parameterizing input-dependent parameters $\mathcal{B}$, $\mathcal{C}$ and Δ , the two views are fused through the following formulas:

$$\mathcal{B}^o = \mathbf{W}_{\mathcal{B}^o} \mathbf{H}_{input}^a, \quad \mathcal{C}^o = \mathbf{W}_{\mathcal{C}^o} \mathbf{H}_{input}^a, \quad \Delta^o = \text{softplus}(\mathbf{W}_{\Delta^o} \mathbf{H}_{input}^a), \quad (8)$$

$$\mathcal{B}^a = \mathbf{W}_{\mathcal{B}^a} \mathbf{H}_{input}^o, \quad \mathcal{C}^a = \mathbf{W}_{\mathcal{C}^a} \mathbf{H}_{input}^o, \quad \Delta^a = \text{softplus}(\mathbf{W}_{\Delta^a} \mathbf{H}_{input}^o), \quad (9)$$

where $\mathbf{W}$ denotes the corresponding learnable matrix.

Selective Discretization in VF-SSM. To discretize the parameters $\mathcal{A}$ and $\mathcal{B}$ into $\bar{\mathcal{A}}$ and $\bar{\mathcal{B}}$, we adopt the zero-order hold (ZOH) discretization rule, following [7,15].

$$\bar{\mathcal{A}}^o = \exp(\Delta^o \mathcal{A}^o), \quad \bar{\mathcal{B}}^o = (\Delta^o \mathcal{A}^o)^{-1} (\exp(\Delta^o \mathcal{A}^o) - I) (\Delta^o \mathcal{B}^o), \quad (10)$$

$$\bar{\mathcal{A}}^a = \exp(\Delta^a \mathcal{A}^a), \quad \bar{\mathcal{B}}^a = (\Delta^a \mathcal{A}^a)^{-1} (\exp(\Delta^a \mathcal{A}^a) - I) (\Delta^a \mathcal{B}^a). \quad (11)$$

After the discretization of system parameters, the VF-SSM in VFM performs $\text{SSM}(\bar{\mathcal{A}}^o, \bar{\mathcal{B}}^o, \mathbf{H}_{input}^o)$ and $\text{SSM}(\bar{\mathcal{A}}^a, \bar{\mathcal{B}}^a, \mathbf{H}_{input}^a)$ to update state selectively following Equation (2). The output of VF-SSM is as follows:

$$\begin{aligned} \mathbf{y}_{ssm}^o &= \text{SSM}(\bar{\mathcal{A}}^o, \bar{\mathcal{B}}^o, \mathbf{H}_{input}^o), \\ \mathbf{y}_{ssm}^a &= \text{SSM}(\bar{\mathcal{A}}^a, \bar{\mathcal{B}}^a, \mathbf{H}_{input}^a). \end{aligned} \quad (12)$$

After VFM, we obtain the final node representations $\mathbf{Z}^o$ and $\mathbf{Z}^a$ for the original and augmented views, respectively, as follows:

$$\begin{aligned} \mathbf{u}^o &= \text{SiLU}(\text{Linear}(\text{LayerNorm}(\mathbf{H}^o))), \\ \mathbf{Z}^o &= \text{LayerNorm}(\text{LayerNorm}(\text{Linear}(\mathbf{y}_{ssm}^o \odot \mathbf{u}^o)) + \mathbf{H}^o), \\ \mathbf{u}^a &= \text{SiLU}(\text{Linear}(\text{LayerNorm}(\mathbf{H}^a))), \\ \mathbf{Z}^a &= \text{LayerNorm}(\text{LayerNorm}(\text{Linear}(\mathbf{y}_{ssm}^a \odot \mathbf{u}^a)) + \mathbf{H}^a), \end{aligned} \quad (13)$$

where $\odot$ denotes element-wise multiplication.

4.3 Spectrum-Guided Mamba (SGM)

To consider spectral differences between normal and anomalous graphs for the UGLAD task, we design a specialized Spectrum-Guided Mamba (SGM) module, as illustrated in Fig. 2(c). The SGM module adopts the Rayleigh quotient as a measure of spectral characteristics, which is closely related to anomalies as described in Section 3.3. Specifically, the SGM module utilizes the Rayleigh quotient to parameterize system parameters, making them spectrum-dependent and enabling spectrum-guided updates of latent states. Designed in a Mamba-Transformer-style architecture, the SGM effectively refines graph embeddings

by selectively focusing on anomaly-relevant spectral information to enhance anomaly detection performance.

Selective Parameterization in SG-SSM. Let $\mathbf{h}_G$ denote the graph representations encoded in Section 4.1. Firstly, we utilize the MLP to obtain the Rayleigh quotient representation $\mathbf{h}_{RQ}$, as follows:

$$\mathbf{h}_{RQ} = \text{MLP}(\text{diag}(R(\mathbf{L}, \mathbf{X}))), \quad (14)$$

where $\text{diag}(\cdot)$ extracts the diagonal elements of Rayleigh quotient $R(\mathbf{L}, \mathbf{X})$.

The inputs $\mathbf{h}_G$ and $\mathbf{h}_{RQ}$ are passed through a series of operations, including a layer normalization [2], a linear projection, a 1D convolutional operation, and a SiLU activation function, as follows:

$$\mathbf{h}_{input} = \text{SiLU}(\text{Conv1D}(\text{Linear}(\text{LayerNorm}(\mathbf{h}_G)))), \quad (15)$$

$$\mathbf{h}_{RQ} = \text{SiLU}(\text{Conv1D}(\text{Linear}(\text{LayerNorm}(\mathbf{h}_{RQ}))). \quad (16)$$

We then utilize the obtained $\mathbf{h}_{RQ}$ to parameterize $\mathcal{B}$, $\mathcal{C}$ and Δ , making system parameters $(\mathcal{B}, \mathcal{C}, \Delta)$ spectrum-dependent, as follows:

$$\mathcal{B} = \mathbf{W}_B \mathbf{h}_{RQ}, \quad \mathcal{C} = \mathbf{W}_C \mathbf{h}_{RQ}, \quad \Delta = \text{softplus}(\mathbf{W}_\Delta \mathbf{h}_{RQ}), \quad (17)$$

where $\mathbf{W}$ is the corresponding learnable matrix. For the parameter $\mathcal{A}$, we still adopt HiPPO-LegS [8] for parameterization [7,9].

Selective Discretization in SG-SSM. According to the ZOH rule, we discretize the parameters $\mathcal{A}$ and $\mathcal{B}$:

$$\bar{\mathcal{A}} = \exp(\Delta \mathcal{A}), \quad \bar{\mathcal{B}} = (\Delta \mathcal{A})^{-1}(\exp(\Delta \mathcal{A}) - I)(\Delta \mathcal{B}). \quad (18)$$

The SG-SSM in SGM performs $\text{SSM}(\bar{\mathcal{A}}, \bar{\mathcal{B}}, \mathbf{h}_{input})$ following Equation (2), yielding the output $\mathbf{y}_{ssm}$:

$$\mathbf{y}_{ssm} = \text{SSM}(\bar{\mathcal{A}}, \bar{\mathcal{B}}, \mathbf{h}_{input}). \quad (19)$$

After the Spectrum-Guided Mamba, we obtain the final output $\mathbf{z}_G$ as follows:

$$\begin{aligned} \mathbf{u} &= \text{SiLU}(\text{Linear}(\text{LayerNorm}(\mathbf{h}_G))), \\ \mathbf{z}_G &= \text{LayerNorm}(\text{LayerNorm}(\text{Linear}(\mathbf{y}_{ssm} \odot \mathbf{u})) + \mathbf{h}_G). \end{aligned} \quad (20)$$

According to Equation (14)-(20), we obtain graph representations $\mathbf{z}_G^o$ and $\mathbf{z}_G^a$ for the original and augmented views, respectively.

4.4 Training and Inference

Training. We adopt the InfoNCE loss [46] as contrastive objective to maximize the agreement between the representations from two views at node and graph scales:

$$\begin{aligned}\mathcal{L}'_{node} &= \frac{1}{|B|} \sum_{G_j \in B} \frac{1}{2|V_{G_j}|} \sum_{v_i \in V_{G_j}} [\ell(\mathbf{z}_i^o, \mathbf{z}_i^a) + \ell(\mathbf{z}_i^a, \mathbf{z}_i^o)], \\ \ell(\mathbf{z}_i^o, \mathbf{z}_i^a) &= -\log \frac{e^{(\cos(\mathbf{z}_i^o, \mathbf{z}_i^a)/\tau)}}{\sum_{v_k \in V_{G_j} \setminus v_i} e^{(\cos(\mathbf{z}_i^o, \mathbf{z}_k^a)/\tau)}},\end{aligned}\tag{21}$$

where B is the training batch, $\mathbf{z}_i^o$ and $\mathbf{z}_i^a$ are the embeddings of node v_i on two views, $\cos(\cdot)$ is the cosine similarity function, and τ is the temperature parameter.

$$\begin{aligned}\mathcal{L}'_{graph} &= \frac{1}{2|B|} \sum_{G_i \in B} [\ell(\mathbf{z}_{G_i}^o, \mathbf{z}_{G_i}^a) + \ell(\mathbf{z}_{G_i}^a, \mathbf{z}_{G_i}^o)], \\ \ell(\mathbf{z}_{G_i}^o, \mathbf{z}_{G_i}^a) &= -\log \frac{e^{\cos(\mathbf{z}_{G_i}^o, \mathbf{z}_{G_i}^a)/\tau}}{\sum_{G_j \in B \setminus G_i} e^{\cos(\mathbf{z}_{G_i}^o, \mathbf{z}_{G_j}^a)/\tau}},\end{aligned}\tag{22}$$

where $\mathbf{z}_{G_i}^o$ and $\mathbf{z}_{G_i}^a$ are the embeddings of graph G_i on two views, and other notations are analogous to those in Equation (21).

During the training phase, we adopt an adaptive loss to consider different sensitivities of node and graph scales for different datasets [17,22], as follows:

$$\begin{aligned}\mathcal{L}_{node} &= (\sigma_{node})^\alpha \mathcal{L}'_{node}, \quad \mathcal{L}_{graph} = (\sigma_{graph})^\alpha \mathcal{L}'_{graph}, \\ \mathcal{L} &= \mathcal{L}_{node} + \mathcal{L}_{graph},\end{aligned}\tag{23}$$

where α is the hyper-parameter and σ is the standard deviation of predicted errors on the corresponding scale.

Inference. By minimizing $\mathcal{L}$ in Equation (23) during the training, the model learns common patterns of normal graphs. When testing an anomalous graph, the loss $\mathcal{L}$ tends to be significantly higher; therefore, we utilize $\mathcal{L}$ as the anomaly score. Additionally, the z-score standardization is adopted to balance anomaly scores from different scales. The final anomaly score is formulated as:

$$S = \text{Std}(\mathcal{L}_{node}) + \text{Std}(\mathcal{L}_{graph}),\tag{24}$$

where $\text{Std}(\mathcal{L}) = (\mathcal{L} - \mu)/\sigma$, and μ is the mean value of predicted errors of training samples at the corresponding scale.

5 Experiments

5.1 Experiment Settings

Datasets. We conduct experiments on 12 public real-world datasets from Tu-Dataset benchmark [27], which involve small molecules, bioinformatics, and social networks. The statistics of the datasets are presented in Table 1. Following the setting in [17,22,24], the samples in the minority class or real anomalous class are viewed as anomalies, while the rest are viewed as normal data. And only normal samples are used during training under the unsupervised setting.

Table 1. Statistics of the datasets [27] used in our experiments.

Category	Dataset	Graphs	Avg. Nodes	Avg. Edges	Node Attr.
Bioinformatics	ENZYMES	600	32.63	62.14	18
Small molecules	AIDS	2000	15.69	16.20	4
	DHFR	467	42.43	44.54	3
	BZR	405	35.75	38.36	3
	COX2	467	41.22	43.45	3
	NCI1	4110	29.87	32.30	-
	HSE	8417	16.89	17.23	-
	MMP	7558	17.62	17.98	-
	p53	8903	17.92	18.34	-
	PPAR-gamma	8451	17.38	17.72	-
Social networks	IMDB-B	1000	19.77	96.53	-
	REDDIT-B	2000	429.63	497.75	-

Baselines. To evaluate the effectiveness of GLADMamba, we compare it with 9 competitive baselines, spanning both earlier and recent approaches. These include the two-stage methods PK-iF [18,28], WL-OCSVM [26,32], WL-iF [18,32], InfoGraph-iF [18,33] and GraphCL-iF [18,41], as well as end-to-end methods OCGIN [45], GLocalKD [24], GOOD-D [22] and CVTGAD [17].

- **Two-stage methods:** These methods first generate graph embeddings by graph kernels (e.g., propagation kernel (PK) [28] or Weisfeiler-Lehman kernel (WL) [32]) or graph representation learning methods (e.g., InfoGraph [33] or GraphCL [41]), and then apply traditional algorithms (e.g., isolation forest (iF) [18] or one-class SVM (OCSVM) [26]) to identify anomalies.
- **End-to-end methods:** These approaches integrate graph representation learning and anomaly detection into a unified framework, enabling joint optimization. OCGIN [45] adapts GIN for end-to-end graph anomaly detection by one-class classification objective. GLocalKD [24] employs random distillation to learn normal patterns by training one GNN to predict another randomly-initialized GNN. GOOD-D [22] uses hierarchical contrastive learning to detect anomalies through semantic inconsistency. CVTGAD [17] employs a simplified Transformer for UGLAD.

Evaluation Metrics. Following [17,22,24], we use the area under the receiver operating characteristic curve (AUC) as the evaluation metric for UGLAD, where a higher AUC reflects better performance.

Implementation Details. For the baselines, we report their public results from [17,22]. We implement GLADMamba by PyTorch¹ on NVIDIA L40 and A40 GPUs, and ensure reproducibility by explicitly setting random seeds following [17,22]. We employ GCN as the default GNN encoder. For AIDS, DHFR, HSE, and MMP datasets, the encoder employs GIN.

¹ <https://pytorch.org/>

Table 2. Overall performance comparison in terms of AUC (% , mean $\pm$ std). The best and second-best results are highlighted in **bold** and underlined, respectively.

Method	PK-iF	WL-OCSVM	WL-iF	InfoGraph-iF	GraphCL-iF
ENZYMES	51.30 $\pm$ 2.01	55.24 $\pm$ 2.66	51.60 $\pm$ 3.81	53.80 $\pm$ 4.50	53.60 $\pm$ 4.88
AIDS	51.84 $\pm$ 2.87	50.12 $\pm$ 3.43	61.13 $\pm$ 0.71	70.19 $\pm$ 5.03	79.72 $\pm$ 3.98
DHFR	52.11 $\pm$ 3.96	50.24 $\pm$ 3.13	50.29 $\pm$ 2.77	52.68 $\pm$ 3.21	51.10 $\pm$ 2.35
BZR	55.32 $\pm$ 6.18	50.56 $\pm$ 5.87	52.46 $\pm$ 3.30	63.31 $\pm$ 8.52	60.24 $\pm$ 5.37
COX2	50.05 $\pm$ 2.06	49.86 $\pm$ 7.43	50.27 $\pm$ 0.34	53.36 $\pm$ 8.86	52.01 $\pm$ 3.17
NCI1	50.58 $\pm$ 1.38	50.63 $\pm$ 1.22	50.74 $\pm$ 1.70	50.10 $\pm$ 0.87	49.88 $\pm$ 0.53
IMDB-B	50.80 $\pm$ 3.17	54.08 $\pm$ 5.19	50.20 $\pm$ 0.40	56.50 $\pm$ 3.58	56.50 $\pm$ 4.90
REDDIT-B	46.72 $\pm$ 3.42	49.31 $\pm$ 2.33	48.26 $\pm$ 0.32	68.50 $\pm$ 5.56	71.80 $\pm$ 4.38
HSE	56.87 $\pm$ 10.51	62.72 $\pm$ 10.13	53.02 $\pm$ 5.12	53.56 $\pm$ 3.98	51.18 $\pm$ 2.71
MMP	50.06 $\pm$ 3.73	55.24 $\pm$ 3.26	52.68 $\pm$ 3.34	54.59 $\pm$ 2.01	54.54 $\pm$ 1.86
p53	50.69 $\pm$ 2.02	54.59 $\pm$ 4.46	50.85 $\pm$ 2.16	52.66 $\pm$ 1.95	53.29 $\pm$ 2.32
PPAR-gamma	45.51 $\pm$ 2.58	57.91 $\pm$ 6.13	49.60 $\pm$ 0.22	51.40 $\pm$ 2.53	50.30 $\pm$ 1.56
Avg.Rank	8.83	7.50	8.58	6.83	7.42

Method	OCGIN	GLocalKD	GOOD-D†	CVTGAD	GLADMamba
ENZYMES	58.75 $\pm$ 5.98	61.39 $\pm$ 8.81	63.90 $\pm$ 3.69	<u>67.79$\pm$5.43</u>	68.39$\pm$4.55
AIDS	78.16 $\pm$ 3.05	93.27 $\pm$ 4.19	97.28 $\pm$ 0.69	99.39$\pm$0.55	<u>99.29$\pm$0.47</u>
DHFR	49.23 $\pm$ 3.05	56.71 $\pm$ 3.57	62.67 $\pm$ 3.11	<u>62.95$\pm$3.03</u>	63.79$\pm$4.16
BZR	65.91 $\pm$ 1.47	69.42 $\pm$ 7.78	75.16 $\pm$ 5.15	<u>75.92$\pm$7.09</u>	77.25$\pm$4.62
COX2	53.58 $\pm$ 5.05	59.37 $\pm$ 12.67	62.65 $\pm$ 8.14	<u>64.11$\pm$3.22</u>	66.38$\pm$1.40
NCI1	<u>71.98$\pm$1.21</u>	68.48 $\pm$ 2.39	61.12 $\pm$ 2.21	69.07 $\pm$ 1.15	73.06$\pm$1.87
IMDB-B	60.19 $\pm$ 8.90	52.09 $\pm$ 3.41	65.88 $\pm$ 0.75	70.97$\pm$1.35	<u>69.63$\pm$2.70</u>
REDDIT-B	75.93 $\pm$ 8.65	77.85 $\pm$ 2.62	88.67$\pm$1.24	84.97 $\pm$ 2.41	<u>86.12$\pm$0.41</u>
HSE	64.84 $\pm$ 4.70	59.48 $\pm$ 1.44	69.65 $\pm$ 2.14	<u>70.30$\pm$2.90</u>	71.19$\pm$2.68
MMP	<u>71.23$\pm$0.16</u>	67.84 $\pm$ 0.59	70.57 $\pm$ 1.56	70.96 $\pm$ 1.01	73.19$\pm$3.22
p53	58.50 $\pm$ 0.37	64.20 $\pm$ 0.81	62.99 $\pm$ 1.55	<u>67.58$\pm$3.31</u>	68.38$\pm$1.62
PPAR-gamma	71.19$\pm$4.28	64.59 $\pm$ 0.67	67.34 $\pm$ 1.71	68.25 $\pm$ 4.66	<u>69.21$\pm$3.46</u>
Avg.Rank	4.50	4.58	3.25	<u>2.17</u>	1.33

5.2 Overall Performance Comparison

We evaluate the performance of GLADMamba against several baselines in terms of AUC across 12 datasets. As summarized in Table 2, GLADMamba achieves the highest average rank, outperforming all baselines on 8 datasets and ranking second on the remaining datasets. The results also show that unified models surpass two-stage methods, highlighting the advantages of end-to-end optimization for representation learning and anomaly detection. These results underscore the effectiveness of GLADMamba in the UGLAD task across diverse domains.

5.3 Ablation Study

To validate the effectiveness of key components of GLADMamba, we conduct extensive ablation experiments on 8 representative datasets in Fig. 3. The variant




Fig. 3. Ablation study on key components across representative datasets.

w/o VF-SSM modifies the selective parameterization process in the VF-SSM using corresponding single-view inputs. The variant *w/o* SG-SSM replaces the Rayleigh quotient with corresponding graph-level representations. The variants *w/o* VFM, *w/o* SGM, and *w/o* Mamba denote the removal of the View-Fused Mamba, the Spectrum-Guided Mamba, and both, respectively.

The results show that GLADMamba consistently outperforms all variants. The variants *w/o* VFM, *w/o* SGM, and *w/o* Mamba significantly degrade performance, validating the effectiveness of Mamba in the UGLAD task and its superiority in capturing long-range dependencies. And the *w/o* VF-SSM variant, which relies solely on single-view information, suffers a performance decline due to its limited ability to identify complex anomalies. Furthermore, the significant performance drop in the *w/o* SG-SSM variant demonstrates the importance of spectral differences between normal and anomalous graphs for guiding GLADMamba’s latent state updates. This also underscores that the SG-SSM component enhances GLADMamba’s sensitivity to anomaly-related information by introducing Rayleigh quotient.

5.4 Efficiency Analysis

We assess the efficiency of GLADMamba on representative datasets in terms of FLOPs, parameter size, and GPU usage in Table 3. On the small-scale molecular dataset AIDS, the complexity differences between CVTGAD and GLADMamba are tolerable. However, on larger-scale datasets REDDIT-B and p53, the complexity differences become significant, particularly in terms of FLOPs and parameter size. This demonstrates the superiority of GLADMamba over Transformer-based approaches in scaling to large-scale graph anomaly detection.




Fig. 4. The hyper-parameter analysis on representative datasets.

5.5 Hyper-Parameter Analysis

The state size of VFM and SGM. To investigate the impact of state size in Mamba, we conduct hyper-parameter analysis experiments in Fig. 4(a). The results indicate that performance variations with state size are not entirely consistent across datasets, likely due to the dependency of state size on dataset scale. Overall, the model can achieve satisfactory performance at smaller values (e.g., 4, 8).

The size of Δ projection of VFM and SGM. The parameter Δ is constructed by linear projection of the input, controlling the strength of state updates. To investigate the impact of Δ projection size, we conduct experiments in Fig. 4(b). The results show that the model generally performs worst when the size is 2, except for the COX2 dataset. In general, the model can achieve respectable performance when the size is 4. As the projection size increases further, the performance remains relatively stable overall. These findings indicate that a moderate projection size (e.g., 4) is sufficient for optimal performance, avoiding suboptimal performance or unnecessary resource overhead.

The local convolution width of VFM and SGM. The local convolution width refers to the kernel size of the Conv1D layer, controlling the model’s receptive field. We conduct experiments to investigate its impact in Fig. 4(c). We observe that when the width is 2, the performance degrades due to limited expressiveness of the Conv1D. In most cases, the model achieves optimal perfor-

Table 3. Efficiency comparison on FLOPs, parameter size, and GPU memory usage.

Dataset	Model	FLOPs (M)↓	Params (MB)↓	GPU (GB)↓
AIDS	CVTGAD	33.27	1.88	8.58
	GLADMamba	48.97	0.07	5.12
REDDIT-B	CVTGAD	602.46	13.19	30.70
	GLADMamba	431.59	0.11	23.79
p53	CVTGAD	701.31	20.61	4.98
	GLADMamba	226.19	0.08	3.96




Fig. 5. Visualization analysis on AIDS dataset. (● normal graph, ● anomalous graph.)

mance when the width is 4. Overall, GLADMamba remains relatively insensitive to this parameter, demonstrating robust performance.

5.6 Visualization Analysis

We utilize t-SNE to visualize the graph embeddings from two views and anomaly scores learned by GLADMamba in Fig. 5. It can be observed that normal and anomalous graphs are not well-separated under a single view, as their distributions exhibit considerable overlap. However, when multiple views are integrated, normal and anomalous graphs exhibit a clear boundary in terms of anomaly scores, demonstrating the effectiveness of GLADMamba in UGLAD.

6 Conclusion

In this work, we introduce GLADMamba, a novel framework for UGLAD that effectively integrates the selective state space model and explicit spectral information. By leveraging the proposed View-Fused Mamba (VFM) and Spectrum-Guided Mamba (SGM) modules, GLADMamba dynamically selects and refines anomaly-related information, significantly improving detection performance. As far as we know, this is the first work to introduce Mamba and explicit spectral information to UGLAD. Experimental results on 12 real-world datasets demonstrate that GLADMamba outperforms existing state-of-the-art methods, highlighting its potential for advancing the UGLAD field. This fundamental architecture advancement establishes new possibilities for both graph learning and anomaly detection research.

Acknowledgments. This work is supported by the Youth Fund of the National Natural Science Foundation of China (No. 62206107).

Disclosure of Interests. We declare that our research adheres to ethical guidelines, ensuring no harm or misuse of data and respecting privacy and confidentiality. We have conducted the study with integrity and disclosed any potential conflicts of interest.

References

1. Atitallah, S.B., Rabah, C.B., Driss, M., Boulila, W., Koubaa, A.: Exploring graph mamba: A comprehensive survey on state-space models for graph learning. arXiv preprint arXiv:2412.18322 (2024)
2. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
3. Breunig, M.M., Kriegel, H.P., Ng, R.T., Sander, J.: Lof: identifying density-based local outliers. In: Proceedings of the 2000 ACM SIGMOD international conference on Management of data. pp. 93–104 (2000)
4. Dong, X., Zhang, X., Wang, S.: Rayleigh quotient graph neural networks for graph-level anomaly detection. arXiv preprint arXiv:2310.02861 (2023)
5. Erol, M.H., Senocak, A., Feng, J., Chung, J.S.: Audio mamba: Bidirectional state space model for audio representation learning. IEEE Signal Processing Letters (2024)
6. Fu, Y., Li, J., Liu, J., Xing, Q., Wang, Q., King, I.: Hc-glad: Dual hyperbolic contrastive learning for unsupervised graph-level anomaly detection. arXiv preprint arXiv:2407.02057 (2024)
7. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752 (2023)
8. Gu, A., Dao, T., Ermon, S., Rudra, A., Ré, C.: Hippo: Recurrent memory with optimal polynomial projections. Advances in neural information processing systems **33**, 1474–1487 (2020)
9. Gu, A., Goel, K., Gupta, A., Ré, C.: On the parameterization and initialization of diagonal state space models. Advances in Neural Information Processing Systems **35**, 35971–35983 (2022)
10. Gu, A., Goel, K., Re, C.: Efficiently modeling long sequences with structured state spaces. In: International Conference on Learning Representations (2022)
11. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. Advances in neural information processing systems **30** (2017)
12. Hu, J., Guo, D., Si, Z., Liu, D., Diao, Y., Zhang, J., Zhou, J., Wang, M.: Mol-mamba: Enhancing molecular representation with structural & electronic insights. arXiv preprint arXiv:2412.16483 (2024)
13. Kalman, R.E.: A new approach to linear filtering and prediction problems (1960)
14. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
15. Li, D., Tan, S., Zhang, Y., Jin, M., Pan, S., Okumura, M., Jiang, R.: Dymamba: Continuous state space modeling on dynamic graphs. arXiv preprint arXiv:2408.06966 (2024)
16. Li, J., Liu, Y., Xing, Q., Wang, Q., Pan, S.: No fear of representation bias graph contrastive learning with calibration and fusion. Available at SSRN 4774833 (2024)
17. Li, J., Xing, Q., Wang, Q., Chang, Y.: Cvtgad: simplified transformer with cross-view attention for unsupervised graph-level anomaly detection. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases. pp. 185–200. Springer (2023)
18. Liu, F.T., Ting, K.M., Zhou, Z.H.: Isolation forest. In: 2008 eighth IEEE international conference on data mining. pp. 413–422. IEEE (2008)
19. Liu, J., Fournier-Viger, P., Zhou, M., He, G., Nouioua, M.: Cspm: Discovering compressing stars in attributed graphs. Information Sciences **611**, 126–158 (2022)

20. Liu, J., Yang, M., Zhou, M., Feng, S., Fournier-Viger, P.: Enhancing hyperbolic graph embeddings via contrastive learning. arXiv preprint arXiv:2201.08554 (2022)
21. Liu, J., Zhou, M., Fournier-Viger, P., Yang, M., Pan, L., Nouioua, M.: Discovering representative attribute-stars via minimum description length. In: 2022 IEEE 38th International Conference on Data Engineering (ICDE). pp. 68–80. IEEE (2022)
22. Liu, Y., Ding, K., Liu, H., Pan, S.: Good-d: On unsupervised graph out-of-distribution detection. In: Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining. pp. 339–347 (2023)
23. Luo, X., Wu, J., Yang, J., Xue, S., Peng, H., Zhou, C., Chen, H., Li, Z., Sheng, Q.Z.: Deep graph level anomaly detection with contrastive learning. Scientific Reports **12**(1), 19867 (2022)
24. Ma, R., Pang, G., Chen, L., van den Hengel, A.: Deep graph-level anomaly detection by glocal knowledge distillation. In: Proceedings of the fifteenth ACM international conference on web search and data mining. pp. 704–714 (2022)
25. Ma, X., Wu, J., Xue, S., Yang, J., Zhou, C., Sheng, Q.Z., Xiong, H., Akoglu, L.: A comprehensive survey on graph anomaly detection with deep learning. IEEE transactions on knowledge and data engineering **35**(12), 12012–12038 (2021)
26. Manevitz, L.M., Yousef, M.: One-class svms for document classification. Journal of machine Learning research **2**(Dec), 139–154 (2001)
27. Morris, C., Kriege, N.M., Bause, F., Kersting, K., Mutzel, P., Neumann, M.: Tudataset: A collection of benchmark datasets for learning with graphs. arXiv preprint arXiv:2007.08663 (2020)
28. Neumann, M., Garnett, R., Bauckhage, C., Kersting, K.: Propagation kernels: efficient graph kernels from propagated information. Machine learning **102**, 209–245 (2016)
29. Qiao, Y., Yu, Z., Guo, L., Chen, S., Zhao, Z., Sun, M., Wu, Q., Liu, J.: Vlmamba: Exploring state space models for multimodal learning. arXiv preprint arXiv:2403.13600 (2024)
30. Qu, H., Ning, L., An, R., Fan, W., Derr, T., Liu, H., Xu, X., Li, Q.: A survey of mamba. arXiv preprint arXiv:2408.01129 (2024)
31. Qureshi, S., et al.: Limits of depth: Over-smoothing and over-squashing in gnns. Big Data Mining and Analytics **7**(1), 205–216 (2023)
32. Shervashidze, N., Schweitzer, P., Van Leeuwen, E.J., Mehlhorn, K., Borgwardt, K.M.: Weisfeiler-lehman graph kernels. Journal of Machine Learning Research **12**(9) (2011)
33. Sun, F.Y., Hoffman, J., Verma, V., Tang, J.: Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. In: International Conference on Learning Representations (2020)
34. Tan, Y., Liu, Y., Long, G., Jiang, J., Lu, Q., Zhang, C.: Federated learning on non-iid graphs via structural knowledge sharing. In: Proceedings of the AAAI conference on artificial intelligence. vol. 37, pp. 9953–9961 (2023)
35. Tang, J., Li, J., Gao, Z., Li, J.: Rethinking graph neural networks for anomaly detection. In: International conference on machine learning. pp. 21076–21089. PMLR (2022)
36. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017)
37. Xu, F., Wang, N., Wu, H., Wen, X., Zhao, X., Wan, H.: Revisiting graph-based fraud detection in sight of heterophily and spectrum. In: Proceedings of the AAAI conference on artificial intelligence. vol. 38, pp. 9214–9222 (2024)
38. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? arXiv preprint arXiv:1810.00826 (2018)

39. Xu, R., Yang, S., Wang, Y., Du, B., Chen, H.: A survey on vision mamba: Models, applications and challenges. *arXiv e-prints* pp. arXiv-2404 (2024)
40. Yang, X., Zhao, X., Shen, Z.: A generalizable anomaly detection method in dynamic graphs. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 39, pp. 22001–22009 (2025)
41. You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., Shen, Y.: Graph contrastive learning with augmentations. *Advances in neural information processing systems* **33**, 5812–5823 (2020)
42. Yuan, H., Sun, Q., Wang, Z., Fu, X., Ji, C., Wang, Y., Jin, B., Li, J.: Dg-mamba: Robust and efficient dynamic graph structure learning with selective state space models. *arXiv preprint arXiv:2412.08160* (2024)
43. Zang, Y., Ren, L., Li, Y., Wang, Z., Selby, D.A., Wang, Z., Vollmer, S.J., Yin, H., Song, J., Wu, J.: Rethinking cancer gene identification through graph anomaly analysis. *arXiv preprint arXiv:2412.17240* (2024)
44. Zhao, H., Zhang, M., Zhao, W., Ding, P., Huang, S., Wang, D.: Cobra: Extending mamba to multi-modal large language model for efficient inference. *arXiv preprint arXiv:2403.14520* (2024)
45. Zhao, L., Akoglu, L.: On using classification datasets to evaluate graph outlier detection: Peculiar observations and new insights. *Big Data* **11**(3), 151–180 (2023)
46. Zhu, Y., Xu, Y., Yu, F., Liu, Q., Wu, S., Wang, L.: Graph contrastive learning with adaptive augmentation. In: *Proceedings of the web conference 2021*. pp. 2069–2080 (2021)
47. Zou, D., Peng, H., Liu, C.: A structural information guided hierarchical reconstruction for graph anomaly detection. In: *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. pp. 4318–4323 (2024)