Transactions on Computational and Scientific Methods | Vol. 5, No. 8, 2025
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

A Cloud Computing Anomaly Detection Method Integrating Diffusion Graph Neural Networks with Long- and Short-Term Temporal Dependency Modeling

Bencheng Su[1]
[1]University of Cincinnati, Cincinnati, USA
*Corresponding author: Bencheng Su; benchengsu8@gmail.com

Abstract: With the continuous expansion of cloud platform scale and the widespread application of microservice architecture, container orchestration, and dynamic resource scheduling mechanisms, system operation states exhibit greater complexity, dynamism, and correlation. Traditional anomaly detection methods relying on single indicator fluctuations or local time window analysis are no longer sufficient to meet the accurate identification needs in complex cloud environments. Addressing the challenges of concealed structural propagation, intertwined temporal dependencies, and uneven influence of key nodes in cloud computing scenarios, this paper proposes a cloud computing anomaly detection method that integrates a diffusion graph neural network and a long short-term temporal dependency capture mechanism. This method first models service instances and their dependencies in the cloud environment as a dynamic graph structure and constructs a unified input representation by combining monitoring indicator sequences to enhance the joint characterization of system topology interactions and state evolution. Subsequently, it extracts high-order service association information through the diffusion graph modeling mechanism, mines potential patterns of anomaly propagation along dependency links, and utilizes the long short-term temporal dependency capture mechanism to learn the dynamic changes in system state over continuous time, thereby improving the ability to identify sudden and gradual anomalies. Based on this, an adaptive attention fusion strategy is introduced to selectively enhance key information from different service nodes and at different times, further improving the discriminative power of anomaly representation and overall detection stability. The results show that the proposed method can effectively integrate relational structure information and temporal evolution information in cloud platforms, demonstrating strong comprehensive recognition capabilities in cloud computing anomaly detection tasks, and providing an effective method for intelligent monitoring and operational status perception in complex distributed systems.

Keywords: Microservice monitoring; dynamic graph modeling; anomaly propagation detection; intelligent operation and maintenance.

1. Introduction

With the rapid popularization of cloud computing infrastructure, containerized deployment, and microservice architecture, modern information systems are continuously evolving towards high concurrency, strong coupling, dynamic scaling, and heterogeneous collaboration[1]. While cloud environments improve resource utilization and service elasticity, they also significantly increase the complexity of system operation, making anomaly patterns more concealed, propagating, and dynamically changing. In the context of multi-tenant sharing, on-demand resource scheduling, and frequent changes in service chains, performance degradation, service failures, resource contention, and call chain anomalies often intertwine, affecting not

only system stability but also potentially leading to service quality degradation, business interruptions, and increased operational costs. Therefore, building efficient, accurate, and adaptable anomaly detection methods for cloud computing scenarios has become an important research direction for ensuring the reliability, security, and intelligent operation and maintenance capabilities of cloud platforms[2,3].

Traditional anomaly detection methods typically rely on manually set thresholds, statistical rules, or shallow machine learning models. While they have some applicability to low-dimensional, static, and relatively independent monitoring indicators, they often face significant limitations in complex cloud environments[4]. On the one hand, cloud system operation data comes from a wide range of sources, including indicator sequences, log events, and call chain information, as well as service dependencies, topology evolution, and resource interaction behavior. A single-perspective modeling approach cannot fully depict the true mechanisms of anomaly occurrence. On the other hand, anomaly behavior in the cloud environment often does not occur in isolation, but rather spreads gradually along service dependency paths, resource sharing links, or temporal propagation trajectories, exhibiting significant structural correlations and temporal continuity. Focusing only on local indicator fluctuations or changes within short time windows easily overlooks the propagation patterns and cumulative effects of anomalies at the system level, thus reducing the accuracy and interpretability of detection results[5].

In recent years, the development of graph structure modeling and temporal modeling methods has provided new research ideas for cloud computing anomaly detection. Graph neural networks can extract deep correlation features from service call relationships, resource dependency structures, and component interaction topologies, helping to reveal the spatial propagation paths and structural impact range of anomalies in complex systems[6]. At the same time, temporal modeling mechanisms can capture the dynamic evolution characteristics of system states from continuous observation sequences, demonstrating strong characterization capabilities for sudden anomalies, gradual degradation, and anomaly shifts under periodic fluctuations. However, relying solely on graph structure modeling can weaken the temporal dependency information during anomaly evolution, while relying solely on temporal modeling makes it difficult to fully utilize the widely existing relational topologies and dependency patterns within cloud systems. Therefore, how to collaboratively model structural dependencies and temporal evolution within a unified framework has become a key scientific issue for improving the performance and applicability of cloud computing anomaly detection[7].

Against this backdrop, research integrating the diffusion graph modeling approach with long-term and short-term temporal dependency capture mechanisms has significant theoretical and practical value. Theoretically, this direction helps to shift anomaly detection from single-point observation to relation-driven, dynamic perception, and globally collaborative modeling, further enriching the research framework of intelligent analysis methods for complex distributed systems. From an application perspective, constructing anomaly detection methods that balance structural perception and temporal characterization capabilities for cloud computing scenarios helps improve the ability to identify complex faults, potential risks, and early degradation signals, enhancing the platform's proactive response level in resource scheduling, service assurance, and fault early warning. As cloud-native technology continues to deepen, researching cloud computing anomaly detection methods that integrate diffusion graph neural networks and long-short-term time-series dependency capture mechanisms is of great significance for supporting the high-reliability operation of cloud platforms, improving intelligent operation and maintenance efficiency, and promoting the cloud service system to a higher level of adaptive and intelligent development.

Recent related studies further indicate that cloud anomaly detection benefits from integrating structural representation, risk perception, and cross-domain temporal modeling. Existing microservice anomaly detection research has expanded from distributed tracing and graph feature extraction toward service-level probabilistic inference, trace-log fusion, graph variational representation, and trace-metric-log interpretability [8], [9], [10], [11], [12], [13], [14]. Building on this line of work, self-supervised unified representation learning with structural dependency modeling strengthens early risk discovery in distributed

systems, while contrastive anomaly modeling in cross-period financial statements and microstructure-aware sequence modeling show how abnormal temporal signals can be separated from noisy, high-dimensional observations [15], [16], [17]. Frequency-guided Transformer forecasting further introduces spectral enhancement into time-series modeling, and cross-modal graph recommendation demonstrates the value of combining temporal behavior with relational fusion for heterogeneous representation learning [18], [19].

For cloud platforms, secure parameter aggregation and cross-cloud distributed anomaly detection emphasize privacy-preserving collaboration and robustness across data centers, which aligns with the distributed detection setting addressed in this study [20]. Attention-enhanced traffic prediction, topology-aware graph neural node fault identification, energy-budget-aware task scheduling, and spatiotemporal latency prediction collectively show that anomaly detection should be coupled with traffic variation, topology evolution, resource constraints, and service quality dynamics [21], [22], [23], [24]. Masked language modeling for semantic matching also provides a representation-learning perspective for improving robustness when monitoring data contains textual, log-like, or sparse contextual signals [25]. These advances support the need for a unified framework that combines graph diffusion, long-short-term temporal dependency modeling, and adaptive cloud operation awareness.

2. Methodology

Cloud anomaly detection in distributed environments requires a unified representation that can preserve service dependency, temporal fluctuation, and cross-node influence within the same modeling space. To this end, the monitoring system is first formulated as a dynamic attributed graph, where each service instance is treated as a node, each invocation or dependency relation is encoded as an edge, and each node is associated with a multivariate observation sequence describing resource usage, latency variation, and status evolution over time. This paper also presents the overall model architecture, as shown in Figure 1.

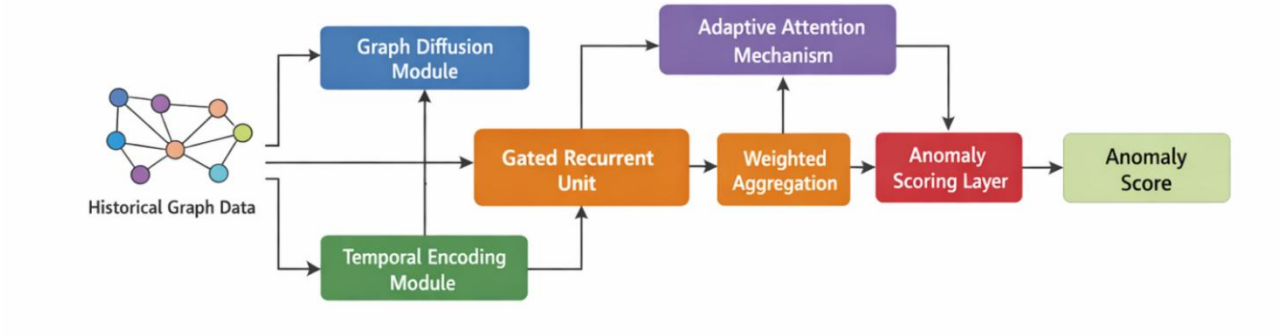


Figure 1. Overall model architecture

Such a formulation allows structural interactions and temporal dynamics to be jointly embedded instead of being processed in isolation, which is essential for characterizing cascading failures and weak anomaly propagation in cloud platforms. Let the cloud system at time step t be represented by:

$$\mathcal{G}_t = (\mathcal{V}, \mathcal{E}_t, \boldsymbol{X}_t, \boldsymbol{A}_t)$$

where $\mathcal{V}$ denotes the set of service nodes, $\mathcal{E}_t$ denotes the time-varying dependency edges, $\boldsymbol{X}_t \in \mathbb{R}^{N \times F}$ is the node feature matrix with N nodes and F monitoring attributes, and $\boldsymbol{A}_t \in \mathbb{R}^{N \times N}$ is the corresponding adjacency matrix. Rather than using the raw topology directly, diffusion normalization is introduced to suppress local noise and reveal high-order dependency transfer, thereby enabling the model to capture latent anomaly spread along service chains. The diffusion transition matrix is defined as:

$$\boldsymbol{P}_t = \boldsymbol{D}_t^{-1} \boldsymbol{A}_t$$

in which $\boldsymbol{D}_t$ is the degree matrix induced by $\boldsymbol{A}_t$. Building on this operator, the K-step diffusion representation aggregates neighborhood responses from different propagation ranges and transforms them into a structure-aware embedding,

$$\boldsymbol{S}_t = \sum_{k=0}^{K} \alpha_k \boldsymbol{P}_t^k \boldsymbol{X}_t$$

where α_k controls the contribution of the k-hop diffusion path. This design is meaningful because short-range diffusion emphasizes immediate fault correlation, while long-range diffusion captures delayed structural influence that frequently appears in large-scale cloud systems.

Beyond structural dependence, anomaly evolution in cloud platforms usually exhibits both abrupt perturbation and slow accumulation, which makes long and short temporal contexts equally important for reliable modeling. A temporal encoder is therefore introduced to map historical graph-aware features into a memory-enhanced latent space, so that transient bursts and persistent degradation can be characterized within a common recurrent process. Given the diffusion-enhanced representation $\boldsymbol{S}_t$, the input, forget, and output gates are computed as:

$$\boldsymbol{i}_t, \boldsymbol{f}_t, \boldsymbol{o}_t = \sigma\left(\boldsymbol{W}_g\left[\boldsymbol{S}_t \| \boldsymbol{h}_{t-1}\right] + \boldsymbol{b}_g\right)$$

where $\boldsymbol{h}_{t-1}$ denotes the previous hidden state, $\|$ denotes concatenation, and $\sigma(\cdot)$ denotes the sigmoid activation. Since anomaly evidence may remain weak in a single observation window, the memory cell is updated by integrating current candidate information with historical retention, which improves sensitivity to cumulative instability and avoids overreacting to isolated noise.

$$\boldsymbol{c}_t = \boldsymbol{f}_t \odot \boldsymbol{c}_{t-1} + \boldsymbol{i}_t \odot tanh\left(\boldsymbol{W}_c\left[\boldsymbol{S}_t \| \boldsymbol{h}_{t-1}\right] + \boldsymbol{b}_c\right)$$

where $\odot$ denotes element-wise multiplication. On this basis, the hidden representation is generated as:

$$\boldsymbol{h}_t = \boldsymbol{o}_t \odot tanh\left(\boldsymbol{c}_t\right)$$

thereby preserving temporal continuity while highlighting time steps that carry stronger anomaly-related responses. The combination of diffusion encoding and gated recurrence is adopted not merely for feature fusion, but for aligning structural propagation with temporal memory, which is crucial for identifying failures that unfold across multiple services and multiple time horizons.

Another key consideration is that different cloud services do not contribute equally to anomaly formation at every moment, especially under workload bursts, scheduling changes, or dynamic resource contention. For this reason, an adaptive fusion mechanism is constructed to reweight node-level representations according to their current structural relevance and temporal saliency, allowing the model to emphasize influential services and suppress redundant responses. Given the temporal state $\boldsymbol{h}_t^{(i)}$ of node i, the anomaly-aware attention coefficient is formulated as:

$$\beta_t^{(i)} = \frac{exp\left(\boldsymbol{q}^\top tanh\left(\boldsymbol{W}_a \boldsymbol{h}_t^{(i)} + \boldsymbol{b}_a\right)\right)}{\sum_{j=1}^{N} exp\left(\boldsymbol{q}^\top tanh\left(\boldsymbol{W}_a \boldsymbol{h}_t^{(j)} + \boldsymbol{b}_a\right)\right)}$$

where $\boldsymbol{W}_a$ and $\boldsymbol{b}_a$ are trainable parameters and $\boldsymbol{q}$ is the attention vector. Such weighting enables the representation to reflect the uneven impact of heterogeneous service components, which is particularly important when anomalies originate from a small subset of critical nodes but gradually affect the broader system. The global anomaly context can then be obtained through weighted aggregation,

$$\boldsymbol{z}_t = \sum_{i=1}^{N} \beta_t^{(i)} \boldsymbol{h}_t^{(i)}$$

so that both local abnormal responses and system-level interactions are preserved in a compact descriptor. In practical terms, this fusion strategy enhances interpretability because higher attention values naturally indicate services that exert a stronger influence on the current system state.

Finally, anomaly scoring is performed in the latent space to quantify the deviation between the learned system representation and its normality-oriented projection, thereby transforming complex spatiotemporal behavior into a discriminative decision signal. Instead of relying on manually designed thresholds over raw metrics, the learned representation $\mathbf{z}_t$ is projected by a lightweight detection head to estimate the abnormal confidence, which improves robustness under heterogeneous monitoring scales and changing operational conditions. The anomaly score is defined as:

$$\widehat{y}_t = \sigma\left(\boldsymbol{W}_o \boldsymbol{z}_t + b_o\right)$$

where $\widehat{y}_t \in [0,1]$ measures the anomaly probability at time step t. During optimization, the objective function jointly constrains prediction consistency and parameter regularization, so that the model can preserve expressive capability without excessive fitting to short-term fluctuations.

$$\mathcal{L} = -\frac{1}{T}\sum_{t=1}^{T}\left[y_t log\widehat{y}_t + (1-y_t)\, log\left(1-\widehat{y}_t\right)\right] + \lambda \|\Theta\|_2^2$$

where y_t denotes the ground-truth anomaly label, Θ denotes the trainable parameter set, and λ controls the regularization strength. Through this formulation, diffusion-based relational perception, long and short temporal dependency capture, adaptive service reweighting, and end-to-end anomaly discrimination are integrated into a coherent framework, making the method suitable for modeling the intricate propagation, accumulation, and manifestation mechanisms of anomalies in cloud computing environments.

3. Experimental Results and Analysis

3.1 Dataset Introduction

This paper uses an open-source anomaly monitoring dataset built on the Train-Ticket microservice benchmark system as the source of experimental data. This dataset is generated from a realistic microservice ticketing system containing 41 microservices, providing inter-service dependencies, call chain information, log data, and key performance indicators collected using Prometheus. Therefore, it possesses the information foundation required for graph structure association modeling and time-series anomaly analysis. The publicly released version of the data includes monitoring records under multiple anomaly scenarios and describes the anomalous microservices and their related operational states, making it suitable for anomaly detection, anomaly propagation analysis, and multimodal operation and maintenance modeling research in cloud computing environments. Because this dataset retains service topology interaction characteristics and contains continuous time-series monitoring signals, it exhibits high consistency with cloud computing anomaly detection tasks that integrate diffused graph neural networks and long-short-term time-series dependency capture mechanisms.

3.2 Experimental Results and Analysis

To further pinpoint the technical characteristics of the proposed method within the broader research context, it is necessary to examine it alongside representative works in cloud computing anomaly detection, microservice call chain anomaly identification, and multi-source operational data modeling. Table 1 selects works highly relevant to the research topic of this paper. These methods explore distributed tracing, graph structure modeling, log and metric fusion, real-time detection, and interpretable anomaly analysis, comprehensively reflecting the main technical routes and performance levels in this field.

Table 1. Experimental results compared with other models

Method	ACC	PRE	REC	F1
Liu et al.[8]	90.12	88.74	89.03	88.88
Zhang et al.[9]	91.45	90.17	90.62	90.39
Lee et al.[10]	92.08	91.12	90.84	90.98
Xie et al.[11]	92.36	91.75	91.43	91.59
Panahandeh et al.[12]	91.84	90.96	91.28	91.12
Raeiszadeh et al.[13]	90.73	89.85	90.14	89.99
Ren et al.[14]	92.21	91.43	91.67	91.55
Ours	**94.68**	**93.97**	**94.12**	**94.04**

Experimental results show that the proposed algorithm demonstrates outstanding overall classification performance and can reliably complete the task of identifying cloud computing anomalies. This method exhibits strong advantages in accuracy, precision, recall, and comprehensive discriminative ability, indicating that the constructed model can not only effectively mine structural correlation information during anomaly propagation but also effectively characterize the continuous evolution of system states over time. Because the diffusion graph modeling mechanism enhances the representation of service dependencies and potential anomaly propagation paths, and the long-short-term time-series dependency capture mechanism further improves the perception of complex dynamic changes, this method demonstrates higher recognition quality and stronger overall robustness in anomaly detection tasks, reflecting significant application value.

To further analyze the impact of each key component on the overall model performance, Table 2 presents the performance evaluation results under the ablation settings. The settings strictly adhere to the diffusion graph modeling, long-short-term time-series dependency capture mechanism, and adaptive attention fusion mechanism in our proposed method, thereby examining the specific contributions of different modules to structural correlation modeling, dynamic evolution perception, and global discrimination capabilities in cloud computing anomaly detection tasks.

Table 2. Ablation study results of different model settings

Ablation Setting	ACC	PRE	REC	F1
w/o Graph Diffusion Modeling	91.74	90.86	91.08	90.97
w/o Long Short-Term Dependency Capture	92.16	91.33	91.57	91.45
w/o Adaptive Attention Fusion	92.48	91.62	91.85	91.73
Ours	**94.68**	**93.97**	**94.12**	**94.04**

The results in Table 2 show that each core component of the proposed method plays a crucial role in improving overall performance. Removing any key module leads to a decrease in the model's anomaly detection capability to varying degrees. This indicates that diffusion graph modeling plays an irreplaceable role in characterizing service dependencies and anomaly propagation in the cloud environment. The long-short-term dependency capture mechanism is also critical for representing the dynamic accumulation and

evolution of anomalous states, while adaptive attention fusion further enhances the model's ability to focus on important service nodes and effective temporal information. The complete model can simultaneously consider structural relationship modeling, temporal dependency perception, and discriminative information enhancement, thus exhibiting stronger comprehensive recognition capabilities and more stable feature representation effects in cloud computing anomaly detection tasks.

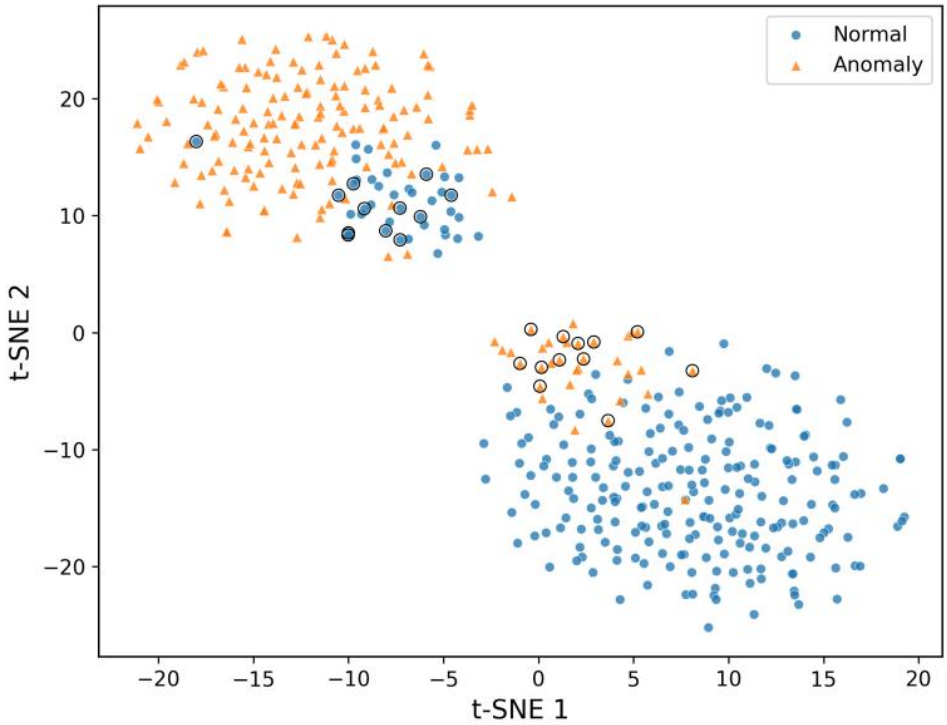


Figure 2. Experimental results of t-SNE

As shown in Figure 2, the proposed algorithm can effectively distinguish between normal and abnormal samples in the feature space, exhibiting strong class separability and clear structure in its overall distribution. Normal and abnormal states form relatively compact clusters with defined boundaries, indicating that the constructed model can effectively extract key discriminative information corresponding to abnormal behavior in the cloud computing environment, while preserving important features of the system's operational state in terms of structural correlation and temporal changes. Although a small number of overlapping samples still exist in local areas, this reflects the complexity of abnormal patterns in the real cloud environment. The proposed method can still learn highly discriminative representations in this complex context, demonstrating its good feature modeling capabilities and practical application value in cloud computing anomaly detection tasks.

4. Conclusion

This paper addresses the challenges of complex anomaly behavior, hidden propagation paths, significant temporal evolution, and close interrelationships among multiple services in cloud computing environments. It proposes an anomaly detection method that integrates a diffusion graph neural network and a long-short-term temporal dependency capture mechanism. Starting from the widespread service dependencies and continuous monitoring sequences in cloud platforms, this method incorporates structural correlation modeling and temporal dynamic awareness into a unified framework. This allows the model to retain system topology interaction information while further characterizing the accumulation, diffusion, and change of anomaly states over time. Compared to modeling methods relying solely on a single structural or temporal perspective, this proposed method more closely approximates the intrinsic mechanisms of anomaly

generation and propagation in real-world cloud environments, contributing to improved accuracy, stability, and adaptability of anomaly identification in complex distributed systems. From an overall research perspective, this paper not only focuses on the anomaly detection results themselves but also emphasizes the joint representation of multi-level operational behavior in cloud computing systems, providing relatively complete methodological support for intelligent operation and maintenance analysis of complex service architectures.

From a methodological design perspective, the diffusion graph modeling mechanism enhances the model's ability to perceive service dependencies, neighborhood propagation impacts, and potential anomaly propagation paths. This prevents anomalies from being treated as isolated local events, but rather analyzed within a more realistic system interaction network. Meanwhile, the long-short time-series dependency capture mechanism further enhances the dynamic representation capabilities of sudden disturbances, continuous degradation, and the accumulation of staged anomalies, enabling the model to more fully understand the temporal continuity and memory characteristics in changes in the cloud environment state. Building upon this, the introduction of an adaptive attention fusion mechanism further strengthens the model's ability to focus on key service nodes and important moment information, resulting in a more discriminative anomaly representation. These designs demonstrate the strong comprehensive modeling advantages of our method in cloud computing anomaly detection tasks, and also indicate that anomaly analysis for complex cloud platforms requires a collaborative approach across multiple levels, including relational structure, time dependency, and information filtering. This research has significant reference value for cloud-native applications, microservice operations and maintenance, intelligent resource scheduling, fault early warning, and service reliability assurance, helping to promote the further development of cloud platform operation and maintenance systems from passive response to proactive perception, and from single-point monitoring to global collaborative analysis.

Future research can continue to focus on more complex and open cloud computing scenarios. On the one hand, as cloud-native architectures continue to evolve, dynamic scaling of service instances, cross-cluster collaboration, heterogeneous resource hybrid scheduling, and multi-tenant shared environments will further increase the complexity of anomaly patterns. Maintaining the robustness, real-time performance, and generalization ability of anomaly detection models under more dynamic and larger-scale conditions remains a problem worthy of in-depth research. On the other hand, cloud platform operational data typically exhibits multi-source heterogeneous characteristics. Future work can further explore deep collaborative modeling mechanisms among call chains, indicator sequences, log semantics, and resource topology information to enhance the model's ability to identify complex anomalies and hidden risks. Furthermore, combining interpretable analysis, online learning, and continuous update mechanisms is expected to further improve the deployability and long-term adaptability of the method in real-world industrial scenarios. Overall, this research provides an effective approach to cloud computing anomaly detection that balances structural propagation relationships and temporal evolution characteristics, and lays a valuable theoretical and methodological foundation for the continued development of intelligent operation and maintenance, cloud service governance, and highly reliable distributed system management.

References

[1] I. Kohyarnejadfard, D. Aloise, S. V. Azhari, et al., "Anomaly detection in microservice environments using distributed tracing data analysis and NLP," Journal of Cloud Computing, vol. 11, no. 1, p. 25, 2022.

[2] M. Cinque, R. Della Corte, and A. Pecchia, "Micro2vec: Anomaly detection in microservices systems by mining numeric representations of computer logs," Journal of Network and Computer Applications, vol. 208, p. 103515, 2022.

[3] J. Chen, F. Liu, J. Jiang, et al., "TraceGra: A trace-based anomaly detection for microservice using graph deep learning," Computer Communications, vol. 204, pp. 109-117, 2023.

[4] J. Huang, Y. Yang, H. Yu, et al., "Twin graph-based anomaly detection via attentive multi-modal learning for microservice system," in 2023 38th IEEE/ACM International Conference on Automated Software Engineering, 2023, pp. 66-78.
[5] H. He, X. Li, P. Chen, et al., "Efficiently localizing system anomalies for cloud infrastructures: a novel dynamic graph transformer based parallel framework," Journal of Cloud Computing, vol. 13, no. 1, p. 115, 2024.
[6] P. Wang, X. Zhang, Y. Chen, et al., "Unsupervised microservice system anomaly detection via contrastive multi-modal representation clustering," Information Processing & Management, vol. 62, no. 3, p. 104013, 2025.
[7] H. Gao, R. Xin, P. Chen, et al., "Memory-augment graph transformer based unsupervised detection model for identifying performance anomalies in highly-dynamic cloud environments," Journal of Cloud Computing, vol. 14, no. 1, p. 40, 2025.
[8] P. Liu, H. Xu, Q. Ouyang, et al., "Unsupervised detection of microservice trace anomalies through service-level deep bayesian networks," in 2020 IEEE 31st International Symposium on Software Reliability Engineering, 2020, pp. 48-58.
[9] C. Zhang, X. Peng, C. Sha, et al., "Deeptralog: Trace-log combined microservice anomaly detection through graph-based deep learning," in Proceedings of the 44th International Conference on Software Engineering, 2022, pp. 623-634.
[10] C. Lee, T. Yang, Z. Chen, et al., "Eadro: An end-to-end troubleshooting framework for microservices on multi-source data," in 2023 IEEE/ACM 45th International Conference on Software Engineering, 2023, pp. 1750-1762.
[11] Z. Xie, H. Xu, W. Chen, et al., "Unsupervised anomaly detection on microservice traces through graph vae," in Proceedings of the ACM Web Conference 2023, 2023, pp. 2874-2884.
[12] M. Panahandeh, A. Hamou-Lhadj, M. Hamdaqa, et al., "ServiceAnomaly: An anomaly detection approach in microservices using distributed traces and profiling metrics," Journal of Systems and Software, vol. 209, p. 111917, 2024.
[13] M. Raeiszadeh, A. Ebrahimzadeh, A. Saleem, et al., "Real-time anomaly detection using distributed tracing in microservice cloud applications," in 2023 IEEE 12th International Conference on Cloud Networking, 2023, pp. 36-44.
[14] R. Ren, Y. Wang, F. Liu, et al., "Triple: the interpretable deep learning anomaly detection framework based on trace-metric-log of microservice," in 2023 IEEE/ACM 31st International Symposium on Quality of Service, 2023, pp. 1-10.
[15] T. Ying and C. Ding, "Self-Supervised Unified Representation Learning with Structural Dependency Modeling for Early Risk Detection in Distributed Systems," 2024.
[16] Z. Peng, "Cross-Period Financial Statement Anomaly Detection via Contrastive Representation Learning," 2024.
[17] Z. Su, "Microstructure-Aware Risk Detection in High-Frequency Trading Using Embedding and Sequence Modeling," 2024.
[18] Y. Lyu, J. Jiang, and J. Hu, "Frequency-Guided Transformer Modeling with Spectral Enhancement for Accurate Time Series Forecasting," 2024.
[19] Z. Ke, "MCA-TAN: Cross-Modal Graph Neural Recommendation with Temporal Behavior and Social Network Fusion," 2025.
[20] S. Zhang, "Secure Parameter Aggregation and Distributed Anomaly Detection in Cross-Cloud Data Center Networks," Journal of Computer Technology and Software, vol. 3, no. 4, 2024.
[21] Z. Qing and F. Huang, "Attention-Enhanced Network Traffic Prediction for Intelligent Load Balancing Optimization," 2025.
[22] F. Xue, K. Wu, and A. Abdulla, "Topology-Aware Graph Neural Networks for Distributed Node Fault Identification," 2025.

[23] J. Sun, "Reinforcement Learning-Based Task Scheduling Under Energy Budget Constraints in Edge-Cloud Systems," 2024.
[24] Z. Xiao, K. Ying, and B. Venkatesan, "Spatiotemporal Transformer-Based Latency Prediction for Large-Scale Cloud Computing Clusters," 2025.
[25] Z. Bian and X. Su, "Enhancing Semantic Matching and Text Robustness through Masked Language Modeling," 2024.