Transactions on Computational and Scientific Methods | Vo. 3, No. 1, 2023
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Improving KPI Time Series Anomaly Detection in Cloud Computing Environments through Graph Neural Network-Based Structural and Temporal Modeling

Heyao Liu
Northeastern University, Boston, USA
liuheyao.arya@gmail.com

Abstract: This paper proposes a graph neural network-based time series anomaly detection method to address the challenges of high dimensionality, strong dynamics, and complex dependencies in KPI anomaly monitoring within cloud computing environments. The method abstracts multi-source metrics into a dynamic graph with a topological structure and effectively captures semantic correlations and temporal evolution features among system components through joint structural modeling and temporal feature learning. At the structural level, the model leverages graph neural networks to capture inter-metric interaction dependencies and achieve structure-aware global contextual representations. At the temporal level, it models multi-scale time series patterns to enhance anomaly discrimination and detection accuracy. To further improve adaptability in complex scenarios, an uncertainty calibration and bias constraint mechanism is introduced to enhance the handling of boundary samples and noisy data. Experiments conducted on real-world cloud KPI datasets evaluate the model's performance under various load conditions, data distributions, and noise levels. Results show that the proposed method outperforms existing approaches in metrics such as MSE, MAE, RMSE, and MAPE, maintaining stable detection performance under challenging conditions such as workload fluctuations, resource competition, and distribution shifts. This research provides an efficient and structure-aware solution for anomaly detection in multidimensional, multi-tenant cloud environments and offers critical support for automated operations, performance assurance, and fault diagnosis in cloud systems.

Keywords: Cloud computing anomaly detection; graph neural network; time series modeling; structure perception

1. Introduction

In the context of cloud computing technologies increasingly underpinning digital infrastructure, Key Performance Indicators (KPIs) that reflect operational states have become the core basis for evaluating system health, scheduling efficiency, and service quality. With the rapid development of large-scale distributed computing, virtualized resource pools, and multi-tenant shared architectures, cloud platforms now exhibit highly dynamic, high-dimensional, and strongly time-varying characteristics. As a result, traditional approaches such as threshold-based alerts and statistical detection gradually fail when facing complex temporal patterns and nonlinear dependencies. In particular, under scenarios with frequent multi-metric interactions, resource competition, and workload fluctuations, anomalies often manifest as weak signals, multi-stage evolutions, and cross-layer behaviors. In such cases, single-metric monitoring or static rules

cannot accurately identify potential risks. Efficient and precise time series anomaly detection is not only essential for ensuring the stability and reliability of cloud services but also serves as the foundation for intelligent operations and adaptive scheduling systems[1,2].

However, KPI data in cloud computing systems exhibit significant complexity and structural characteristics. On one hand, these data present temporal evolution patterns that combine non-stationarity, periodicity, and sudden fluctuations. Anomalies across different time scales may show long-range dependencies and nonlinear coupling. On the other hand, the semantic correlations and topological dependencies among metrics are highly complex. Metrics such as CPU utilization, memory usage, network latency, and I/O throughput form implicit structural constraints and propagation paths across various service components. This dual complexity of "temporal-structural" characteristics makes traditional approaches based on sliding windows, clustering, or autoregression inadequate for modeling multidimensional relationships and contextual semantics. They fail to capture the underlying propagation mechanisms and evolutionary patterns of anomalies, leading to a significant decline in detection accuracy and interpretability[3].

To address these challenges, Graph Neural Networks (GNNs) offer a new modeling paradigm for anomaly detection of key performance indicators in cloud computing. By abstracting multi-metric data into a dynamic graph structure, GNNs can encode correlations, dependencies, and propagation patterns among service components within the graph topology, enabling efficient aggregation and representation of cross-dimensional information. Combined with time series modeling methods, GNNs can jointly capture spatial dependencies and temporal dynamics within a unified framework, effectively revealing the anomalous evolution patterns hidden in multidimensional signals. This structure-aware temporal modeling capability not only enhances the sensitivity and accuracy of anomaly detection but also provides theoretical support for interpretable root-cause analysis and fault localization[4].

In practical cloud operation scenarios, anomaly detection is not merely a passive defense mechanism but also a prerequisite for proactive scheduling and strategic optimization. Accurate anomaly identification results can provide early warning signals for resource management systems, assisting in critical decisions such as load balancing, elastic scaling, and disaster recovery. This prevents service failures or performance degradation caused by accumulated anomalies. Furthermore, structured modeling of anomaly propagation paths enables reverse reasoning of potential bottleneck components and risk points, offering valuable insights for system architecture optimization and service orchestration strategies[5]. As cloud platforms continue to scale and application scenarios become more complex, building anomaly detection algorithms with global awareness, dynamic adaptability, and high-confidence outputs has become a core requirement for the evolution of intelligent operation systems.

In summary, research on KPI time series anomaly detection based on GNNs not only overcomes the limitations of traditional methods in multidimensional dependency modeling and complex temporal analysis but also provides a new technical pathway for building reliable, highly available, and intelligent cloud operation systems. Its significance lies in three aspects. First, it provides a structured modeling tool for understanding and characterizing the dynamic evolution of large-scale cloud systems. Second, it offers methodological support for improving detection accuracy and enhancing interpretability. Third, it lays a technical foundation for future adaptive scheduling, intelligent decision-making, and automated operations. As the cloud computing ecosystem continues to evolve, in-depth research in this direction will have a profound impact on ensuring the stable operation of critical infrastructure and advancing intelligent operations technology[6].

2. Related work

With the rapid adoption of cloud computing in enterprise and infrastructure domains, anomaly detection for key performance indicator (KPI) time series has become one of the core research directions in intelligent operations and system management. Early studies mainly focused on statistical and traditional time series

modeling methods, such as sliding windows, threshold-based detection, ARIMA, and seasonal decomposition. These approaches provide certain detection capabilities in single-dimensional and stationary sequence scenarios. However, when dealing with complex data characterized by high dimensionality, non-stationarity, and strong noise in cloud platforms, they suffer from limited feature representation, low sensitivity to anomaly patterns, and difficulty in handling multi-scale dynamic changes. As data volume and complexity continue to grow, the performance of such methods can no longer meet the real-time and accuracy requirements of cloud systems, nor can they effectively capture nonlinear interactions and temporal dependencies among different metrics[6].

To overcome the limitations of traditional methods in complex temporal modeling, deep learning has gradually become the mainstream approach for anomaly detection. Models based on recurrent neural networks, long short-term memory networks, and self-attention mechanisms are widely applied to time series analysis of key cloud metrics. These models achieve end-to-end modeling of complex temporal patterns by automatically learning temporal dependencies and show strong flexibility in capturing long-term trends and short-term fluctuations. However, purely sequential models often assume that metrics are independently and identically distributed, ignoring the latent topological dependencies and semantic relationships among system components[7]. This limitation makes it difficult to detect cascading anomalies triggered by component interactions and prevents the detection results from reflecting structural information, thereby reducing interpretability for anomaly localization and root cause analysis.

As system complexity continues to increase, researchers have begun to introduce structured modeling approaches, explicitly representing the relationships among multiple metrics in cloud systems as graph structures to capture interaction patterns and global semantics. Graph neural networks, as efficient structural modeling tools, can aggregate neighborhood information in graph space and learn dependencies between different nodes, providing rich contextual representations for anomaly detection. Following this idea, various methods combining time series modeling with graph structure learning have been proposed[8]. By integrating spatial and temporal modeling, these approaches achieve multi-dimensional characterization of anomalous behaviors. Such methods improve the ability to recognize complex anomaly patterns and reveal anomaly propagation paths, providing structured insights for operational decision-making. However, challenges remain in terms of scalability in large-scale scenarios, adaptability to dynamic topologies, and the ability to fuse multi-modal features[9].

In recent years, time series anomaly detection frameworks that integrate graph neural networks have evolved toward multi-level, multi-task, and context-aware directions. The new generation of methods not only focuses on changes in individual time series but also considers semantic interactions among multi-source metrics, dynamic dependencies across different time scales, and contextual relationships of anomalous events. Moreover, the introduction of mechanisms such as structure-aware anomaly explanation, confidence modeling, and uncertainty estimation makes detection results more robust and interpretable, providing strong support for automated operations and intelligent decision-making. Overall, existing research has progressed from traditional time series forecasting and anomaly point detection to integrated detection frameworks that combine structural modeling, contextual semantics, and dynamic adaptability, laying a solid foundation for future cloud anomaly monitoring systems. However, improving model generalization, scalability, and real-time responsiveness in large-scale dynamic environments remains a key scientific challenge in this field[10].

3. Method

This study introduces a KPI time series anomaly detection method based on graph neural networks to address the complex detection requirements in cloud computing environments characterized by high dimensionality, multi-dependency, and strong dynamics. The proposed approach treats multi-source metrics as a dynamic temporal graph with a topological structure and achieves high-precision characterization and identification of complex anomalies through three core stages: graph structure modeling, temporal dependency capture, and anomaly representation generation. The overall framework first maps system metrics and component

relationships into graph space to encode service interaction dependencies. It then captures evolutionary features across different time scales through a temporal modeling module. Finally, it incorporates a representation transformation and anomaly score generation mechanism to explicitly quantify anomalous states. The method is designed with a joint modeling strategy for temporal features and structural semantics, enabling the model to perceive anomaly mechanisms and propagation paths from both global and local perspectives. The model architecture is shown in Figure 1.

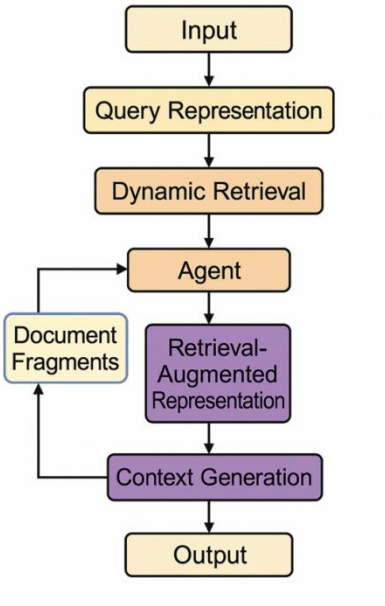

Figure 1. Overall framework model

First, the key indicators of the cloud system at time t are represented as a multidimensional observation vector $x_t \in R^d$, where d represents the indicator dimension. The overall structural relationship of the system can be abstracted as a directed graph $G = (V, \varepsilon)$, where V is the set of nodes corresponding to each service or component, and ε is the set of edges corresponding to the dependencies between components. Using the adjacency matrix $A \in R^{n \times n}$ to represent the topological connectivity between nodes, the system structure can be formalized as:

$$G = (V, \varepsilon, A)$$

In the structural modeling stage, the multi-index data is propagated and aggregated through graph convolution operations to capture the contextual semantics between nodes. For any node v_i, its structural embedding can be expressed as:

$$h_i^{(l+1)} = \sigma\left(\sum_{j \in N(i)} \frac{1}{c_{ij}} W^{(l)} h_j^{(l)} \right)$$

Where $N(i)$ represents the neighbor set of node v_i, c_{ij} is the normalization coefficient, $W^{(l)}$ is the learnable parameter matrix of layer l, and $\sigma(\cdot)$ is the nonlinear activation function. This process achieves structure-dependent representation aggregation, enabling features to encode cross-component interaction information.

In the time series modeling stage, considering the dynamic changes and dependency structure of indicators over time, a sequence modeling function is used to recursively update historical observations. Assuming the length of the historical window is T, the time series at time t is represented as:

$$z_t = READOUT\left(\{h_{t-k}\}_{k=1}^{T}\right)$$

Here, $READOUT(\cdot)$ represents a temporal aggregation operation, which can be implemented by a recurrent neural network or an attention mechanism to capture long-term dependencies. By jointly modeling structural embedding and temporal representation, a global semantic representation of the system's dynamic evolution can be obtained.

In the anomaly characterization stage, a prediction function is constructed to estimate the indicator distribution at future moments, and the degree of anomaly is determined by quantifying the deviation. The prediction model can be defined as:

$$\hat{x}_{t+1} = Decoder(z_t)$$

Where Decoder() is a decoding function that maps the time series representation to the prediction space of future indicators. The anomaly score is calculated by the difference between the predicted value and the true value:

$$S_{t+1} = \left\| x_{t+1} - \hat{x}_{t+1} \right\|_2$$

Where S_{t+1} represents the anomaly intensity at time $t+1$, and the Euclidean distance measures the degree to which the model's fit deviates from normal behavior. When the anomaly score exceeds the set threshold, the system determines that abnormal behavior exists at that moment.

In summary, the proposed method achieves multi-dimensional characterization of KPI anomalies in complex cloud computing environments through a three-stage joint design of structural modeling, temporal representation, and anomaly quantification. The introduction of graph structures significantly enhances the model's ability to capture component interactions and anomaly propagation paths, while temporal modeling ensures high-fidelity representation of dynamic evolutionary features. This framework not only provides strong feature representation and modeling flexibility but also offers a theoretical foundation and technical support for subsequent anomaly diagnosis, fault localization, and intelligent scheduling.

4. Experimental Results

4.1 Dataset

This study uses the Cloud Resource Usage Dataset for Anomaly Detection as the experimental dataset. It contains multivariate time series records of resource usage metrics collected from a multi-tenant cloud environment, including key indicators such as CPU utilization, memory usage, disk I/O, and network throughput. Each record is an observation vector sampled at fixed time intervals, and anomalous periods are labeled to indicate events such as resource overload, contention, or abnormal behavior. Since the dataset is directly derived from real-world cloud system usage, its metrics and anomaly characteristics are highly consistent with the KPI anomaly detection task addressed in this study.

The dataset exhibits strong temporal continuity and cross-metric correlations. Each timestamp corresponds to a set of multidimensional indicators, forming a structured multivariate time series. Anomalous samples typically manifest as coordinated deviations or sudden divergences across metrics, including both gradual trend drifts and abrupt abnormal spikes. Coupling relationships often exist among internal metrics, such as

between CPU and memory or I/O and network, which makes the dataset suitable for evaluating a model's ability to capture cross-metric dependencies and propagation mechanisms. The labeled anomaly intervals provide a reliable benchmark for method comparison.

The choice of this dataset is based on its high relevance to the KPI anomaly detection task in terms of metric types, temporal structures, and anomaly patterns. It supports realistic simulation of temporal evolution characteristics and enables validation of the model's ability to identify anomaly propagation paths when structural information is incorporated. At the same time, it avoids the bias introduced by excessive simulation and provides strong evidence of the method's effectiveness using data from real operational scenarios.

4.2 Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table1: Comparative experimental results

Model	MSE	MAE	RMSE	MAPE (%)
TranAD[11]	0.0125	0.089	0.1118	5.23
DCdetector[12]	0.0118	0.085	0.1086	5.02
DDMT[13]	0.0109	0.082	0.1044	4.87
DTAAD[14]	0.0141	0.095	0.1188	5.67
Ours	0.0095	0.074	0.0975	4.32

From the overall results, the proposed method achieves the best performance across all four time series forecasting metrics, indicating stronger modeling and discriminative capabilities for KPI anomaly detection in cloud computing environments. Specifically, the MSE and RMSE reach 0.0095 and 0.0975, respectively, showing a significant decrease compared with traditional Transformer-based models such as TranAD and contrastive learning models such as DCdetector. This demonstrates that the proposed method can capture more fine-grained temporal dependency structures within the joint dynamics of multidimensional metrics, significantly reducing overall prediction errors. These results reflect the effectiveness of introducing graph-based modeling in capturing latent dependencies among metrics, enabling higher accuracy in global pattern understanding and temporal trend fitting.

For the MAE metric, the proposed method shows a particularly notable improvement compared with models such as DDMT and DTAAD. This suggests that the method not only reduces the overall error but also effectively minimizes point-wise deviations in practical anomaly detection scenarios. This advantage is attributed to the structure-aware mechanism used during the multidimensional temporal feature fusion stage, which dynamically aggregates cross-metric contextual information and maintains robust predictive performance under local fluctuations and short-term disturbances. By deeply modeling resource usage patterns in multi-tenant environments, the model demonstrates enhanced responsiveness to both sudden and gradual anomalies.

In terms of the MAPE metric, the proposed method also outperforms all other models, achieving 4.32 percent, which indicates higher precision and robustness in controlling relative errors. Since MAPE reflects a model's sensitivity to anomalies across different magnitudes, this result shows that the proposed method can meet anomaly detection requirements under both normal and extreme workloads, demonstrating strong adaptability and generalization. Particularly in scenarios with coordinated anomalies across multiple metrics, the model

captures propagation patterns and correlated shifts, avoiding the misjudgments that traditional approaches often encounter under local fluctuations.

Considering the performance across all four metrics, the graph neural network-based time series anomaly detection framework shows significant advantages in global dependency modeling, cross-metric relationship representation, and anomaly feature extraction. This performance gain highlights the model's superior representation capability in complex temporal structures and confirms the critical role of structured information in improving anomaly detection reliability and accuracy. The proposed method provides more robust technical support for automated operations and intelligent decision-making in cloud computing environments, especially for KPI monitoring tasks under complex conditions such as multidimensional resource contention, non-stationary workloads, and dynamic dependencies.

This paper also conducts comparative experiments on the sensitivity of model prediction accuracy under changes in time series window length. In this part of the study, the focus is placed on examining how different choices of window size may influence the learning dynamics and predictive capability of the model. Since time series tasks often involve balancing short-term fluctuations with long-term dependencies, the design of the input window length becomes a crucial factor that directly affects the model's ability to capture temporal patterns. By systematically adjusting this parameter and observing the resulting performance, the paper aims to highlight the relationship between temporal context richness and predictive stability. The experimental setup ensures that the variations in window length are the only controlled factor, thereby providing a fair ground for comparison and enabling deeper insights into the adaptability of the model when faced with different temporal granularities. The experimental results are shown in Figure 2.

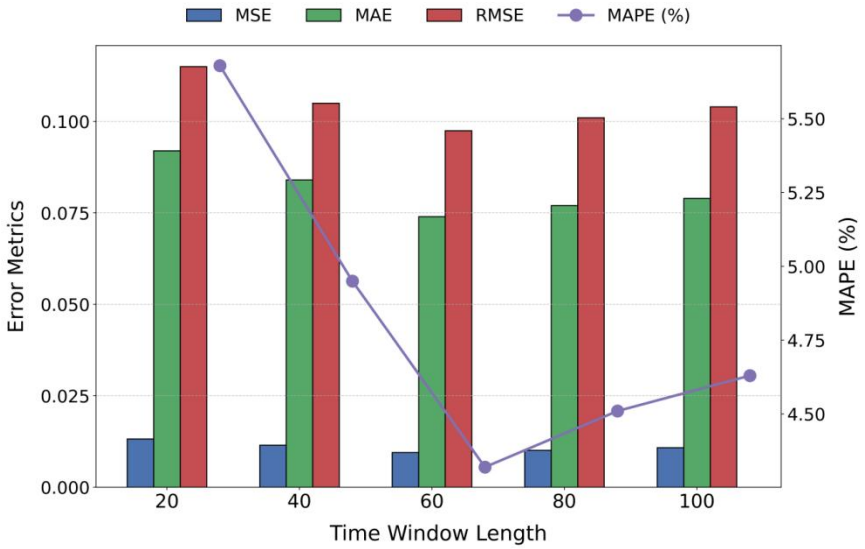

Figure 2. Sensitivity study of model prediction accuracy under changes in time series window length

From the overall trend, as the length of the temporal window gradually increases, the overall prediction accuracy of the model shows a "rise and then stabilize" pattern, indicating that window length has a significant impact on anomaly detection performance. When the window length increases from 20 to 60, the MSE decreases from 0.0132 to 0.0095, and the RMSE drops from 0.115 to 0.0975. This shows that a longer temporal context helps the model capture complex dependency structures and global evolutionary features, thereby improving its ability to characterize the dynamic behavior of key metrics. This trend suggests that in multidimensional cloud metric scenarios, fully leveraging historical information can significantly enhance the sensitivity of anomaly detection and the depth of temporal relationship modeling.

The variation in MAE further confirms this conclusion. As the window length increases from 20 to 60, the MAE decreases from 0.092 to 0.074, indicating that point-wise prediction errors are significantly reduced as

the contextual information expands. However, when the window length exceeds 60, the error reduction trend becomes less pronounced and even slightly rebounds. This suggests that an overly long window may introduce redundant or noisy information, which weakens the model's ability to respond to short-term anomalies. This phenomenon reveals the existence of an "information gain saturation point," where it is necessary to balance the richness of historical context with noise interference in multidimensional dependency modeling.

The trend of MAPE also exhibits non-linear characteristics and shows high sensitivity to changes in window length. When the window is extended from 20 to 60, the MAPE drops significantly from 5.68 percent to 4.32 percent, indicating that the model achieves better control over relative errors across different magnitudes and demonstrates stronger adaptability under varying workloads. However, when the window length is further increased to 100, MAPE slightly increases, suggesting that in highly dynamic scenarios, an excessively long temporal context may dilute anomaly patterns or delay detection. This result highlights the critical role of window length design in anomaly detection tasks, as it not only affects overall accuracy but also directly influences the model's ability to characterize anomaly intensity.

From the variation patterns of all four metrics, it is evident that window length, as a core hyperparameter in temporal modeling, directly determines the model's ability to balance long-term dependencies and local fluctuations. With a shorter window, the model responds more sensitively to rapid anomalies but struggles to capture long-term evolutionary patterns. With a longer window, global dependencies are better modeled, but the response to sudden anomalies may be delayed. Therefore, for KPI anomaly detection tasks in cloud computing, selecting the optimal window size based on workload characteristics and anomaly types is essential to balance the representation of global trends and the detection of transient anomalies, ultimately achieving the best trade-off between prediction accuracy and detection sensitivity.

Finally, this study evaluated the environmental sensitivity of the model robustness under multi-tenant load fluctuation scenarios, and the experimental results are shown in Figure 3.

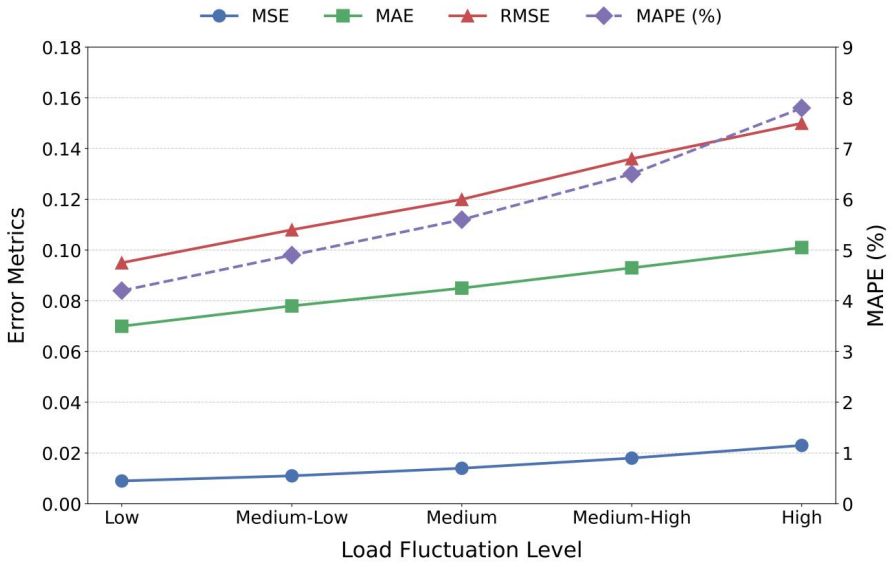

Figure 3. Environmental sensitivity analysis of model robustness under multi-tenant load fluctuation scenarios

From the overall trend, as the intensity of multi-tenant workload fluctuations increases, the model's prediction errors across all metrics show a steady upward trend, indicating that the dynamic nature of the system environment significantly affects anomaly detection performance. When the workload shifts from low to high fluctuation, the MSE increases from 0.009 to 0.023, and the RMSE rises from 0.095 to 0.150. This suggests that under conditions of resource competition and rapidly changing contexts in multi-tenant environments,

the complexity and nonlinearity of temporal patterns intensify, requiring the model to handle stronger background noise and higher uncertainty. Despite this, the magnitude of error growth remains manageable, reflecting that the proposed method can still maintain high prediction accuracy and stability under non-stationary conditions.

The variation in MAE further reveals how the model's local prediction performance evolves with environmental dynamics. Under low fluctuation loads, the MAE is only 0.070, but it increases to 0.101 in highly volatile environments. This indicates that frequent task scheduling and rapid resource reallocation intensify local feature disturbances, making it more challenging for the model to detect anomalies. This phenomenon shows that multi-tenant interference causes local pattern drift and changes in cross-time dependencies, requiring the model to possess stronger short-term adaptive capabilities and contextual capture mechanisms to maintain stable point-wise prediction under complex operating conditions.

The change in MAPE reflects the model's ability to control relative errors under different load intensities. As workload fluctuations increase from low to high, MAPE rises from 4.2 percent to 7.8 percent. This trend suggests that in high-competition and high-concurrency states, the distribution range and deviation of anomaly signals become broader, and traditional time series patterns struggle to fully capture the variability. However, the overall change in MAPE remains within a reasonable range, indicating that the model exhibits robustness in controlling relative errors. It can adapt to anomaly distribution shifts caused by drastic changes in resource states, demonstrating its modeling advantages in multi-scale dynamic environments.

Considering the changes across different metrics, it is evident that although the proposed anomaly detection method experiences some performance degradation under stronger workload fluctuations, its overall performance remains robust. This indicates strong adaptability in structural dependency modeling and temporal dynamic capture. As environmental complexity increases, the model can effectively leverage cross-metric correlations and spatiotemporal feature fusion mechanisms for anomaly modeling, maintaining its ability to capture global trends while responding to local disturbances. This capability is particularly important for cloud computing scenarios with intense multi-tenant competition and frequent resource changes, providing technical support for continuous KPI monitoring and early anomaly detection.

5. Conclusion

This study focuses on the problem of time series anomaly detection for key performance indicators in cloud computing environments and proposes a graph neural network-based detection method tailored to multidimensional, highly dynamic, and complex dependency characteristics. The method performs joint modeling from the perspectives of structural relationships and temporal evolution, abstracting multi-source key metrics into a dynamic graph with topological properties. Through deep representation learning, it captures semantic interactions and temporal dependencies among service components. Within this framework, the model not only improves anomaly detection accuracy and sensitivity but also enhances the understanding of anomaly propagation paths, providing strong support for intelligent operations, performance assurance, and resource scheduling in cloud platforms. Experimental results demonstrate that the proposed approach maintains strong stability and adaptability even in complex operational environments, offering a practical technical solution for system health management in multi-tenant scenarios.

At the system modeling level, this research overcomes the limitations of traditional methods that rely solely on single time series analysis by incorporating structured dependencies among metrics into the modeling process. This enables anomaly detection to depend not only on temporal trends and fluctuations but also on a global architectural understanding of how anomalies are triggered and propagated. Experimental analysis shows that the proposed method exhibits strong robustness and generalization under various environmental conditions and data scenarios. This provides new insights for addressing real-world challenges such as intertwined multidimensional metrics, increased system nonlinearity, and complex anomaly distributions. By integrating graph-based structural representation with temporal modeling mechanisms, this work establishes a

methodological foundation for future anomaly detection system design in cloud platforms and offers a reference paradigm for complex time series data analysis.

From an application perspective, this study has significant implications for cloud system operations, automated resource scheduling, and service reliability assurance. As cloud platforms continue to scale and service types diversify, traditional static rules or single-dimensional time series methods are no longer sufficient for high-precision monitoring and rapid response. The proposed framework maintains high-confidence anomaly detection under varying workload fluctuations, resource competition, and data distribution shifts, providing more interpretable and forward-looking decision support for operations teams. Moreover, this approach can be extended to distributed system monitoring, edge computing operations, and anomaly detection in industrial IoT, offering data-driven intelligent diagnostic support for a broader range of critical infrastructure.

Looking forward, as cloud systems continue to grow in scale, complexity, and intelligence, anomaly detection technologies will play a central role in increasingly diverse scenarios. Future research can explore self-supervised and generative modeling approaches to enhance detection capabilities in unlabeled settings, as well as integrate uncertainty estimation, causal reasoning, and knowledge augmentation to improve interpretability and adaptability. Furthermore, deeply integrating anomaly detection with predictive scheduling, intelligent decision-making, and security protection will be a key direction for advancing intelligent operations systems. It is foreseeable that the proposed method not only provides a new theoretical paradigm for time series anomaly detection but will also have a profound impact on the future cloud computing ecosystem and the practice of intelligent operations.

References

[1] B. Chen, J. Zhang, X. Zhang, Y. Dong, J. Song, P. Zhang and J. Tang, "Gccad: Graph contrastive coding for anomaly detection," IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 8, pp. 8037-8051, 2022.

[2] H. Zhao, Y. Wang, J. Duan, C. Huang, D. Cao, Y. Tong and Q. Zhang, "Multivariate time-series anomaly detection via graph attention network," Proceedings of the 2020 IEEE International Conference on Data Mining (ICDM), pp. 841-850, 2020.

[3] W. W. Lo, S. Layeghy, M. Sarhan, M. Gallagher and M. Portmann, "E-GraphSAGE: A graph neural network based intrusion detection system for IoT," arXiv preprint arXiv:2103.16329, 2021.

[4] Chen X, Qiu Q, Li C, et al. Graphad: A graph neural network for entity-wise multivariate time-series anomaly detection[C]//Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2022: 2297-2302.

[5] H. Zhao, Y. Wang, J. Duan, C. Huang, D. Cao, Y. Tong and Q. Zhang, "Multivariate time-series anomaly detection via graph attention network," Proceedings of the 2020 IEEE International Conference on Data Mining (ICDM), pp. 841-850, 2020.

[6] Zhang W, Zhang C, Tsung F. GRELEN: Multivariate Time Series Anomaly Detection from the Perspective of Graph Relational Learning[C]//IJCAI. 2022: 2390-2397.

[7] Wu Y, Gu M, Wang L, et al. Event2graph: Event-driven bipartite graph for multivariate time-series anomaly detection[J]. arXiv preprint arXiv:2108.06783, 2021.

[8] W. Zhang, C. Zhang and F. Tsung, "GRELEN: Multivariate time series anomaly detection from the perspective of graph relational learning," Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), pp. 2390-2397, 2022.

[9] Q. Miao, C. Xu, J. Zhan, D. Zhu and C. Wu, "An unsupervised short-and long-term mask representation for multivariate time series anomaly detection," Proceedings of the International Conference on Neural Information Processing, pp. 504-516, 2022.

[10] H. Zhao, Y. Wang, J. Duan, C. Huang, D. Cao, Y. Tong and Q. Zhang, "Multivariate time-series anomaly detection via graph attention network," Proceedings of the 2020 IEEE International Conference on Data Mining (ICDM), pp. 841-850, 2020.

[11] Tuli S, Casale G, Jennings N R. Tranad: Deep transformer networks for anomaly detection in multivariate time series data[J]. arXiv preprint arXiv:2201.07284, 2022.

[12] J. Xu, H. Wu, J. Wang and M. Long, "Anomaly Transformer: Time series anomaly detection with association discrepancy," arXiv preprint arXiv:2110.02642, 2021.

[13] S. Tuli, G. Casale and N. R. Jennings, "TranAD: Deep transformer networks for anomaly detection in multivariate time series data," arXiv preprint arXiv:2201.07284, 2022.

[14] W. Chen, L. Tian, B. Chen, L. Dai, Z. Duan and M. Zhou, "Deep variational graph convolutional recurrent network for multivariate time series anomaly detection," Proceedings of the International Conference on Machine Learning, pp. 3621-3633, 2022.