Journal of Computer Technology and Software

ISSN: 2998-2383

Vol. 3, No. 4, 2024

Enhancing Intelligent Anomaly Detection in Cloud Backend Systems through Contrastive Learning and Sensitivity Analysis

Ziyu Cheng
University of Southern California, Los Angeles, USA
chengzy1115@gmail.com

Abstract: This study investigates anomaly detection in cloud backend systems and addresses the limitations of traditional methods under high-dimensional complex data and scarce anomaly samples. A contrastive learning-based algorithm is proposed, which constructs more discriminative latent space representations through feature mapping and representation learning and achieves effective separation of normal and abnormal patterns by jointly optimizing contrastive loss and classification loss. To validate the effectiveness of the method, comparative experiments were conducted on a public dataset, and the results show that the proposed model outperforms several mainstream approaches in terms of AUC, ACC, F1-Score, and Precision. Sensitivity experiments were also performed to analyze the effects of temperature parameter, learning rate, negative sample ratio, and environmental disturbance on model performance. The results demonstrate that proper hyperparameter selection and environmental modeling not only improve overall detection performance but also enhance robustness and stability. By combining comparative experiments with sensitivity analysis, this study comprehensively verifies the effectiveness of the contrastive learning-based anomaly detection method in cloud backend scenarios and confirms its potential application value in complex system operations.

Keywords: Cloud backend; contrastive learning; anomaly detection; sensitivity analysis

1. Introduction

In the era of rapid digitalization and intelligent development, cloud computing has become a critical foundation for the information infrastructure of enterprises and organizations. As the core that supports diverse applications, the cloud backend plays an irreplaceable role in processing massive data, handling high-concurrency requests, and ensuring service continuity. However, with the increasing complexity of business scenarios and the continuous growth of data scale, anomalies in cloud backend systems have become more prominent. These anomalies may cause performance degradation, service interruptions, and even lead to severe security risks and economic losses. Therefore, how to efficiently and accurately identify anomalies in cloud backends has become a key problem in the field of information technology[1].

Traditional anomaly detection methods mainly rely on statistical analysis or rule matching. These approaches are often inadequate when faced with complex, diverse, and high-dimensional cloud backend data. Static thresholds and fixed rules cannot adapt to dynamic environments. With the wide adoption of microservice architectures, containerization, and distributed deployment, system structures have become more flexible and complex, while anomaly patterns have become increasingly diverse[2]. This makes it difficult for traditional methods to capture hidden patterns in high-dimensional features, resulting in poor detection performance. In this context, exploring new intelligent methods to enhance the accuracy and robustness of anomaly detection in cloud backends is of great importance.

In recent years, contrastive learning has emerged as a powerful self-supervised approach in representation learning and classification tasks. Comparing features between normal and abnormal data enables models to learn discriminative representations without requiring large amounts of labeled data. Compared with traditional supervised methods, contrastive learning can better capture intrinsic structural relationships within data, improving the model's ability to recognize anomaly patterns in complex environments. In cloud backend scenarios, contrastive learning not only enhances feature discriminability but also effectively addresses data imbalance and the scarcity of abnormal samples. This provides new perspectives for anomaly detection[3,4].

At the societal level, cloud services have penetrated critical industries such as finance, healthcare, transportation, and energy. The stability and security of these systems directly affect economic development and social operations. If backend anomalies cannot be detected and resolved in time, large-scale service interruptions may occur, potentially triggering systemic risks. Research on contrastive learning-based anomaly detection in cloud backends can therefore improve the level of intelligent operations at the technical layer, while also enhancing resilience and reliability at the societal layer. Such work aligns with national strategies for secure and intelligent management of next-generation information infrastructure, carrying both strategic value and practical significance.

In summary, research on anomaly detection in cloud backends using contrastive learning not only contributes to the development of intelligent anomaly detection theory at the academic level but also addresses practical demands for ensuring cloud service quality in industry[5]. It will help build more intelligent and reliable backend operation systems, reduce risks, and improve service quality and user experience. This research is of great importance for advancing the high-quality development of cloud computing and safeguarding the stable operation of the digital economy.

2. Related work

In recent years, anomaly detection in cloud backends has gradually become an important research focus in both academia and industry. Existing work mainly follows three paths, including statistical modeling, machine learning, and deep learning[6]. Early studies often relied on threshold settings and statistical distribution analysis, such as detecting anomalies through fluctuations in system logs or performance indicators. These methods are simple to implement and computationally efficient. However, they are restricted to static rules. When facing high-dimensional, multi-source heterogeneous data or dynamic environments, their detection performance declines significantly, with frequent false alarms and missed detections. This makes them unsuitable for modern cloud backend systems.

With the development of machine learning, researchers introduced classification and clustering methods into anomaly detection. These approaches can learn anomaly patterns from historical data and generalize better than traditional methods. Common ideas include using support vector machines, random forests, and clustering analysis to identify anomalies in the feature space. However, their performance is limited by model capacity and heavy reliance on feature engineering. They struggle with complex nonlinear relationships. When data dimensions increase or abnormal samples are scarce, results are often unsatisfactory. Moreover, these methods usually require large amounts of labeled data, which is difficult to obtain in practical cloud backend environments[7].

The rise of deep learning has brought new perspectives for anomaly detection in cloud backends. Leveraging the strong feature extraction ability of neural networks, researchers have proposed anomaly detection methods based on time series modeling, graph neural networks, and autoencoder reconstruction. These approaches can automatically capture latent patterns in high-dimensional data and improve detection accuracy and robustness. For example, autoencoder-based methods use reconstruction errors to measure anomaly levels. Recurrent neural networks and attention mechanisms are applied to capture temporal dependencies and multidimensional feature interactions. Although these approaches address some limitations of traditional methods, challenges remain. They often rely heavily on large labeled datasets, are sensitive to changes in system structures, and lack interpretability[8].

In recent years, contrastive learning has been introduced into anomaly detection as a self-supervised paradigm, bringing breakthroughs for cloud backend scenarios. By constructing positive and negative sample pairs and emphasizing the discriminability of feature representations, it can achieve strong robustness even with limited labels. In cloud backend anomaly detection, contrastive learning alleviates the problems of insufficient labels and data imbalance. It also strengthens the model's ability to capture hidden anomaly patterns. Although initial studies have explored this direction, challenges remain in effective feature construction, adaptation to diverse anomaly categories, and real-time performance in large-scale distributed environments. Current progress provides a solid foundation for future research, while also revealing urgent issues that need to be solved.

3. Proposed Approach

The core idea of this research is to construct feature representations that can distinguish normal and abnormal patterns through contrastive learning, thereby achieving cloud backend anomaly identification. Specifically, it is necessary to first perform feature modeling on the input high-dimensional system monitoring data. This paper also gives the overall model architecture, and its experimental results are shown in Figure 1.

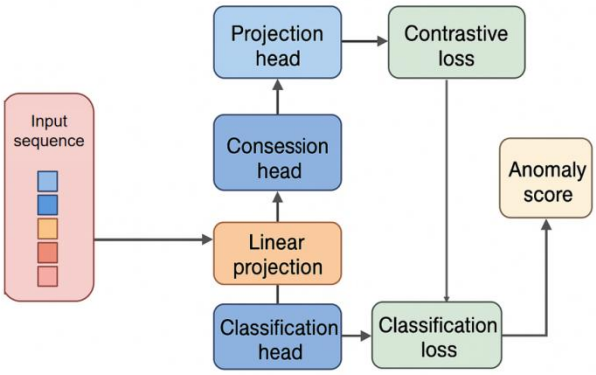

Figure 1. Overall model architecture

Let the original input sequence be $X = \{x_1, x_2, ..., x_T\}$, where each $x_t \in R^d$ represents the multidimensional features collected at time t. To obtain a stable and expressive potential representation, we use a nonlinear mapping function $f_\theta(\cdot)$ to map it to the latent space and obtain a representation vector:

$$h_t = f_\theta(x_t), \qquad h_t \in R^k$$

This representation can effectively capture the complex relationships between different features and lay the foundation for subsequent discriminative learning.

After feature mapping is completed, the objective function of contrastive learning needs to be constructed. The core of contrastive learning is to maximize the similarity between samples of the same class while minimizing the similarity between samples of different classes. The specific form can be expressed by the contrastive loss function with normalized temperature scaling:

$$L_{con} = -\log \frac{\exp(sim(h_i, h_j)/\tau)}{\sum_{k=1}^{N} \exp(sim(h_i, h_k)/\tau)}$$

Here, $sim(\cdot)$ represents cosine similarity, τ is the temperature parameter, h_i and h_j represent pairs of positive samples, and the denominator includes all negative samples. This loss function enables the model to automatically close the representation of similar patterns and distinguish features from different categories.

In the specific modeling of abnormality discrimination, the representation obtained by contrastive learning can be combined with the discriminant function. A linear classification head $g_\phi(\cdot)$ is defined, whose function is to map the latent representation to the binary classification space of abnormality and normality, and the output is:

$$y = g_\phi(h) = \sigma(Wh + b)$$

Here, W and b are parameters, $\sigma(\cdot)$ represents the sigmoid function, and the final output $y \in [0,1]$ represents the probability that the sample is an anomaly. This structure can further achieve explicit anomaly identification based on the representation of contrastive learning.

To enhance the generalization ability of the model in complex cloud backend scenarios, this study also introduced a joint optimization strategy that combines contrast loss and discriminant loss. The discriminant loss is usually in the form of binary cross-entropy:

$$L_{cls} = -[z \log y + (1-z)\log(1-y)]$$

Among them, $z \in [0,1]$ is the true label and y is the predicted probability. The final optimization goal is to comprehensively consider the two:

$$L = \lambda L_{con} + (1-\lambda) L_{cls}$$

Here, $\lambda \in [0,1]$ is a balancing factor that adjusts the contribution ratio between contrastive learning and discrimination tasks. This allows the model to not only learn highly discriminative representations but also directly optimize for the discrimination task, thus better adapting to the anomaly detection requirements of the cloud backend.

4. Experiment result

4.1 dataset

The dataset used in this study is the Yahoo Webscope S5 anomaly detection dataset. It consists of real server performance and business request sequences and is widely used in anomaly detection and time series modeling tasks. The dataset includes various types of indicators, such as CPU utilization, memory usage, and network throughput. Anomalous intervals are labeled, which allows for direct evaluation of model performance in cloud backend scenarios. The dataset is moderate in size. It has a certain level of complexity while maintaining experimental feasibility, making it an important benchmark in anomaly detection research.

The dataset is characterized by diverse and complex anomaly patterns. It contains both short-term burst anomalies and long-term persistent anomalies. These characteristics closely resemble the actual behavior of cloud backend systems, which often exhibit dynamic and nonlinear anomalies under different service loads and resource scheduling. Therefore, the use of this dataset not only evaluates the representation ability of models but also effectively simulates the real challenges of anomaly detection in cloud backends. It holds strong practical value.

In addition, the Yahoo Webscope S5 dataset has been widely adopted in the research community and provides a common benchmark for method comparison. Studies based on this dataset can clearly reveal the strengths and limitations of contrastive learning in anomaly detection tasks. They also lay the foundation for future extensions to larger-scale and higher-dimensional real-world cloud backend data. In summary, the choice of this dataset is both representative and scientifically sound.

4.2 Experimental Results

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

Table 1: Comparative experimental results

Method	AUC	ACC	F1-Score	Precision
MLP[9]	0.872	0.841	0.835	0.828
1DCNN[10]	0.896	0.859	0.849	0.843
LSTM[11]	0.911	0.872	0.861	0.855
Transformer[12]	0.928	0.884	0.874	0.869
Ours	0.957	0.913	0.902	0.896

From the results in Table 1, it can be observed that the traditional MLP model performs the worst across all metrics. The AUC, ACC, F1-Score, and Precision are 0.872, 0.841, 0.835, and 0.828, respectively. This indicates that relying only on shallow fully connected layers cannot effectively capture the complex nonlinear relationships and temporal dependencies in cloud backend data. As a result, its performance in anomaly detection tasks is clearly limited. The particularly low F1-Score suggests that the model struggles to balance recall and precision, making it difficult to maintain stable performance under high-dimensional and diverse anomaly patterns.

With the increase in model complexity, both 1DCNN and LSTM demonstrate stronger detection ability than MLP. The 1DCNN leverages convolution operations to extract local patterns, leading to improvements in AUC and F1-Score. The LSTM captures temporal dependencies through memory units, which further improves ACC and Precision, reaching 0.872 and 0.855, respectively. These results highlight the importance of temporal features for anomaly detection in cloud backends. Static feature modeling alone cannot achieve optimal performance.

The Transformer model outperforms LSTM in overall performance, with consistently high results across all metrics. The AUC reaches 0.928, and the F1-Score is 0.874. Its advantage comes mainly from the multi-head self-attention

mechanism, which can model long-range dependencies and capture both global and local features. For complex anomalies that span different time ranges in cloud backends, the Transformer demonstrates strong representational power. However, although the results are superior to traditional methods, there is still room for improvement, especially when dealing with anomalies that have fuzzy boundaries.

In comparison, the proposed method achieves the best performance across all metrics. The AUC increases to 0.957, the ACC reaches 0.913, and the F1-Score and Precision rise to 0.902 and 0.896, respectively. Compared with the Transformer, the proposed method improves detection ability by 2 to 3 percentage points. This advantage is attributed to the introduction of contrastive learning, which brings normal samples closer together in the representation space while effectively distinguishing anomalies. This strengthens the discriminative power of the features. The results confirm the effectiveness of the contrastive learning-based anomaly detection approach in handling complex data and imbalanced scenarios. They also provide strong support for building more intelligent and reliable cloud backend operation systems.

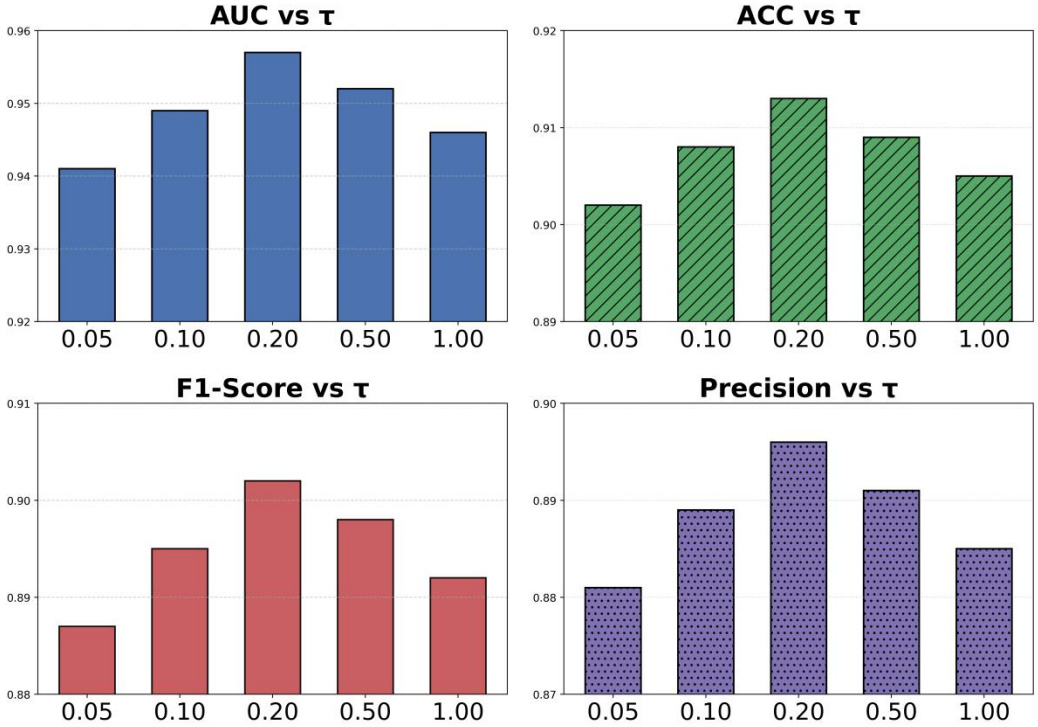

Figure 2. Comparison of the influence of the temperature parameter τ on experimental results

From the results in Figure 2, it can be seen that the temperature parameter τ has a significant impact on model performance. In terms of AUC, the model reaches its best level when τ is set to 0.2, which is clearly better than other values. This indicates that an appropriate temperature parameter can enhance the separation between positive and negative samples in contrastive learning. As a result, the model can more effectively capture hidden anomaly patterns in cloud backend data. When τ is too small or too large, the differences among samples cannot be fully utilized, which leads to weaker discrimination.

For ACC, a similar trend can be observed, with a peak at $\tau = 0.2$. This suggests that when the temperature parameter is set to a moderate level, the model achieves the best overall classification accuracy. It can better balance the discrimination ability across different classes. Too low a τ value forces features to be overly concentrated, making it difficult to cover diverse anomaly patterns. Too high a τ value dilutes the differences among features, resulting in blurred decision boundaries. Therefore, a moderate temperature setting can steadily improve detection performance in the complex environment of cloud backends.

The trend of the F1-Score further confirms this observation, with the highest value obtained at $\tau = 0.2$. This means that at this parameter value, the model achieves the best balance between recall and precision. In cloud backend scenarios, anomalies are often scarce and diverse. Maintaining high recall while preserving precision is therefore critical. The experimental results show that a reasonable temperature parameter helps the model remain robust when facing sparse anomaly samples. It also prevents performance degradation caused by overfitting or underfitting.

The variation in Precision also shows a peak at $\tau = 0.2$. This demonstrates that at this temperature, the model not only identifies anomalies effectively but also reduces false alarms. This is especially important for practical cloud backend operations. Both overly high and overly low τ values lead to a drop in Precision, indicating that the model is prone to misclassification in anomaly detection. Overall, the experimental results verify the sensitivity and dependence of the proposed method on the temperature parameter within the

contrastive learning framework. They also highlight the importance of parameter selection for improving anomaly detection performance in cloud backends.

This paper also provides a detailed analysis of the impact of the learning rate on the experimental outcomes, with a particular focus on how different values of this parameter influence the stability and effectiveness of the training process. The investigation highlights the role of the learning rate as a crucial hyperparameter that directly affects the convergence speed of the model and its ability to capture complex patterns within the data. To clearly illustrate this aspect, the corresponding results are systematically presented in Figure 3, offering an intuitive representation of the relationship between learning rate settings and the overall model behavior.

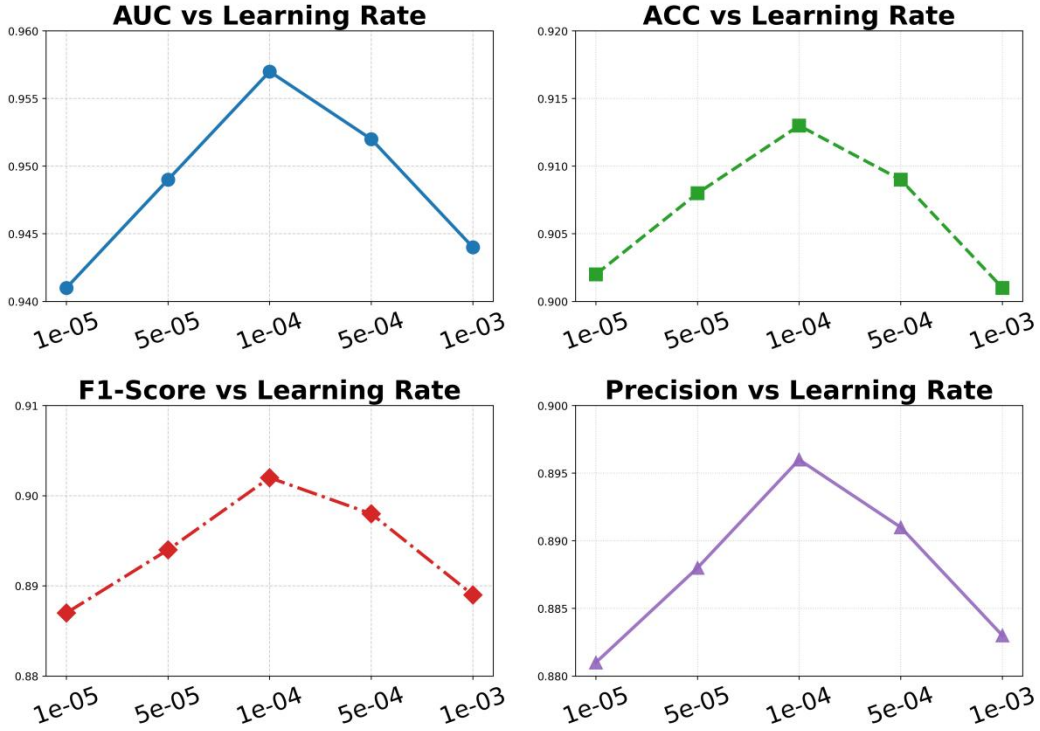

Figure 3. The impact of the learning rate on experimental results

From the results in Figure 3, it can be observed that the learning rate has a clear impact on model performance. In terms of AUC, the highest value is achieved when the learning rate is 1×10^{-4} . This indicates that under this setting, the model can better separate normal and abnormal patterns. When the learning rate is too small, such as 1×10^{-5} , the model converges slowly, leading to poor performance. When the learning rate is too large, such as 1×10^{-3}, the training process becomes unstable and performance decreases. These findings show that a proper learning rate is critical for contrastive learning to achieve strong representation ability in cloud backend anomaly detection.

For ACC, a similar trend is observed, with the best results around 1×10^{-4} . This further confirms the importance of a moderate learning rate. Too low a learning rate results in very small updates, making it difficult for the model to capture complex anomaly patterns. Too high a learning rate causes oscillations and reduces classification accuracy. In cloud backend scenarios, higher ACC means that anomaly detection systems can maintain strong discriminative ability more consistently, with fewer false alarms and missed detections.

The trend of the F1-Score also shows that learning rate selection affects the balance between precision and recall. At 1×10^{-4} , the model reaches its peak F1-Score. This suggests that this learning rate maintains high recall while avoiding a drop in precision. This is critical in anomaly detection, since anomalies in cloud backends are often rare. Low recall would cause serious missed detections, while low precision would result in many false alarms. A reasonable learning rate ensures the best balance between the two.

The results for Precision further support this point. The model achieves its highest Precision at 1×10^{-4} , showing that this learning rate reduces misclassification and improves the reliability of detection. When the learning rate is too high or too low, the Precision decreases significantly. This indicates that parameter tuning is essential for the deployment of anomaly detection models in cloud backends. A reasonable learning rate not only improves performance but also enhances robustness and usability in real complex environments.

This paper also presents a sensitivity experiment on the negative sample ratio to F1-Score, and the experimental results are shown in Figure 4.

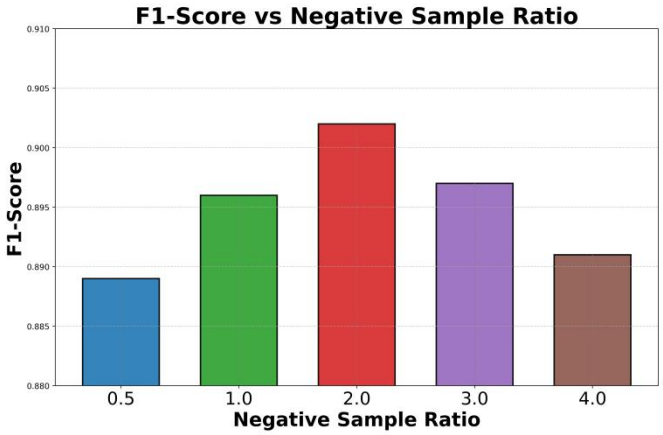

Figure 4. Sensitivity experiment of the negative sample ratio to F1-Score

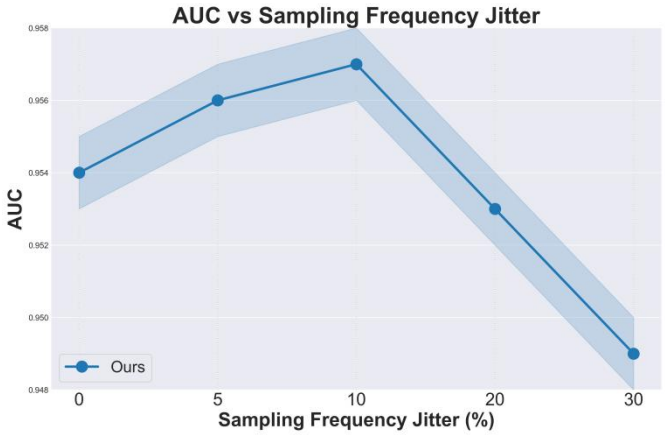

Figure 5. Experiment on the environmental sensitivity of sampling frequency jitter to AUC

From the results in Figure 4, it can be observed that the ratio of negative samples has a clear effect on the F1-Score. When the ratio of negative samples is 0.5, the F1-Score is relatively low, around 0.889. This indicates that an insufficient number of negative samples makes it difficult for the model to learn the boundary between normal and abnormal data. In the contrastive learning framework, too few negative samples weaken the contrastive signal and reduce the overall anomaly detection ability.

When the ratio of negative samples increases to 1.0, the F1-Score rises significantly to 0.896, and the model's discriminative ability improves. This change shows that an appropriate increase in negative samples helps the model capture the differences between normal and abnormal data more effectively. It also improves the balance between recall and precision. In complex cloud backend scenarios, diversity in negative samples supports the detection of different types of anomalies.

When the ratio of negative samples further increases to 2.0, the F1-Score reaches its highest value of 0.902. At this point, the model receives the strongest contrastive signal. The model achieves the best trade-off between precision and recall, which enhances overall detection performance. This suggests that increasing the number of negative samples is beneficial up to a certain point, but improvements do not continue beyond the optimal ratio.

When the ratio continues to increase to 3.0 and 4.0, the F1-Score begins to decline, dropping to 0.897 and 0.891, respectively. This indicates that too many negative samples introduce noise and weaken the effectiveness of contrastive learning. As a result, the model's decision boundary becomes blurred. For anomaly detection in cloud backends, this finding highlights the importance of choosing an appropriate negative sample ratio. A reasonable ratio can significantly improve the effectiveness of contrastive learning models, while ratios that are too high or too low will harm performance.

This paper also presents an experiment on the environmental sensitivity of sampling frequency jitter to AUC, and the experimental results are shown in Figure 5.

From the results in Figure 5, it can be seen that sampling frequency jitter has a clear effect on the AUC performance of the model. When the jitter is 0 percent, the AUC is 0.954, which shows that the model can maintain high anomaly detection ability in an interference-free environment. When the jitter increases to 5 percent, the AUC rises to 0.956, showing a slight improvement. This indicates that moderate sampling disturbance can enhance the model's ability to adapt to uncertainty and improve detection performance to some extent.

When the jitter increases further to 10 percent, the AUC reaches a peak of 0.957, which is the best result in this experiment. This shows that moderate environmental disturbance does not weaken stability. Instead, it encourages the model to learn more robust feature representations. This is important for cloud backend scenarios, as system operation often involves signal fluctuations and noise. A model that performs better under moderate disturbance demonstrates stronger adaptability in real applications.

However, when the jitter increases to 20 percent, the AUC drops significantly to 0.953. This indicates that excessive sampling jitter begins to distort the original data distribution, making it harder for the model to distinguish between normal and abnormal patterns. In complex cloud backend environments, excessive fluctuations weaken the clarity of the decision boundary, which leads to performance degradation.

When the jitter reaches 30 percent, the AUC decreases further to 0.949. This shows that under high-intensity disturbances, the model's detection ability is strongly affected. The result verifies the vulnerability of the model under extreme environmental disturbance and emphasizes the importance of reasonable environmental modeling. Overall, the findings show that moderate jitter can improve the robustness of contrastive learning models, but excessive disturbance significantly reduces performance. This reveals the key issue of balancing robustness and sensitivity to noise in cloud backend anomaly detection.

5. Conclusion

This study focuses on anomaly detection in cloud backends and proposes a contrastive learning-based framework. The

effectiveness and robustness of the framework are verified from different perspectives. The study first analyzes the complexity and diversity of cloud backend systems and highlights the limitations of traditional methods under high-dimensional data, dynamic environments, and scarcity of anomaly samples. By introducing contrastive learning, the model can automatically learn discriminative feature representations without relying on large labeled datasets, which significantly improves anomaly detection performance. The results show that the proposed method not only outperforms traditional and mainstream approaches on overall metrics but also demonstrates strong adaptability and stability under different parameters and environmental disturbances.

The significance of this work lies not only at the algorithmic level but also in its practical value. In critical industries such as finance, healthcare, transportation, and energy, the accuracy of backend anomaly detection directly affects service quality and system security. The proposed contrastive learning framework reduces false alarms and missed detections and enhances system robustness in complex environments. This means that in future intelligent operation scenarios, the method has the potential to provide enterprises with more reliable technical support, reduce risks, improve user experience, and promote the healthy development of the cloud computing industry.

At the same time, this study also reveals both the potential and limitations of contrastive learning in anomaly detection tasks. On one hand, the results show that proper parameter settings and moderate environmental disturbances can further improve robustness and discriminative ability. On the other hand, the findings also indicate that excessive disturbances or severe imbalance in sample ratios may still degrade model performance. Therefore, how to adaptively adjust parameters according to business requirements and how to construct more diverse and high-quality datasets remain important challenges for future work. These findings provide new directions for further research and also lay the foundation for cross-domain applications.

Looking ahead, the contributions of this study can be extended to broader intelligent system operation scenarios. For example, combining contrastive learning with federated learning frameworks in large-scale distributed systems can enable anomaly detection across data centers. Incorporating graph neural networks and time series modeling techniques can further enhance the model's ability to understand multi-source heterogeneous data. In addition, integrating anomaly detection with automated decision-making systems would allow not only recognition of anomalies but also the triggering of intelligent responses and optimization strategies. In summary, this study advances the theoretical development of cloud backend anomaly detection and provides practical support for real-world applications, offering significant implications for future intelligent operations and cloud computing security.

References

[1] M. Jin, Y. Liu, Y. Zheng, L. Chi, Y. F. Li, and S. Pan, "Anemone: Graph anomaly detection with multi-scale contrastive learning," Proceedings of the 2021 30th ACM International Conference on Information & Knowledge Management, pp. 3122-3126, 2021.

[2] Yu H, Yang W, Cui B, et al. Enhanced anomaly traffic detection framework using BiGAN and contrastive learning[J]. Cybersecurity, 2024, 7(1): 71.

[3] Kong S, Ai J, Lu M. CL-MMAD: A contrastive learning based multimodal software runtime anomaly detection method[J]. Applied Sciences, 2023, 13(6): 3596.

[4] H. Sun, Y. Huang, L. Han, C. Fu, and C. Zhou, "HCL-MTSAD: Hierarchical contrastive consistency learning for accurate detection of industrial multivariate time series anomalies," arXiv preprint arXiv:2404.08224, 2024.

[5] H. C. V. Ngu and K. M. Lee, "CL-TAD: A contrastive-learning-based method for time series anomaly detection," Applied Sciences, vol. 13, no. 21, p. 11938, 2023.

[6] Y. Tian, G. Pang, Y. Chen, R. Singh, J. W. Verjans, and G. Carneiro, "Weakly‐supervised video anomaly detection with contrastive learning of long and short‐range temporal features," arXiv preprint arXiv:2101.10030, 2021.

[7] T. Wu, Q. Chen, D. Zhao, J. Wang, and L. Jiang, "Domain adaptation of time series via contrastive learning with task-specific consistency," Applied Intelligence, vol. 54, no. 23, pp. 12576-12588, 2024.

[8] Kong X, Zhang W, Wang H, et al. Federated graph anomaly detection via contrastive self-supervised learning[J]. IEEE Transactions on Neural Networks and Learning Systems, 2024, 36(5): 7931-7944.

[9] Farzad A, Gulliver T A. Log message anomaly detection with fuzzy C-means and MLP[J]. Applied Intelligence, 2022, 52(15): 17708-17717.

[10] Athar A, Mozumder M A I, Ali S, et al. Deep learning-based anomaly detection using one-dimensional convolutional neural networks (1D CNN) in machine centers (MCT) and computer numerical control (CNC) machines[J]. PeerJ Computer Science, 2024, 10: e2389.

[11] A. Iqbal, R. Amin, F. S. Alsubaei, and A. Alzahrani, "Anomaly detection in multivariate time series data using deep ensemble models," PLOS ONE, vol. 19, no. 6, p. e0303890, 2024.

[12] Y. Zhang, "A multi-scale temporal feature extraction approach for network traffic anomaly detection," International Journal of Information Security and Privacy (IJISP), vol. 18, no. 1, pp. 1-20, 2024.