Transactions on Computational and Scientific Methods | Vol. 4, No. 10, 2024
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Federated Learning and Self-Supervised Temporal Modeling for Collaborative Anomaly Detection in Distributed Tasks

Kang Yang[1]*
[1]Stevens Institute of Technology, Hoboken, USA
*Corresponding author: Kang Yang, kang.yang@alumni.stevens.edu

Abstract: To address the challenges of highly concealed anomalies, complex cross-node relationships, significant temporal dependencies, and difficulties in centralized sharing of raw data in distributed task scenarios, this paper proposes a collaborative anomaly detection method based on federated learning and self-supervised temporal modeling. This method is designed for distributed task environments involving multiple nodes, multiple stages, and multiple sources of monitoring signals. Without uploading raw data, it achieves cross-client knowledge aggregation through a federated collaboration mechanism, alleviating the limitations imposed by data silos and privacy constraints on unified modeling. At the local modeling level, the method utilizes a self-supervised mask reconstruction strategy to mine potential temporal structure information from unlabeled execution sequences, enabling the model to learn contextual dependencies and dynamic patterns during task state evolution. Combined with temporal continuity constraints, it enhances the ability to identify complex behaviors such as anomalous perturbations, state shifts, and evolutionary instability. At the global optimization level, an aggregation mechanism oriented towards client-specific differences is constructed to improve the robustness and adaptability of the shared model in heterogeneous distributed environments. In the anomaly detection stage, an anomaly scoring method is established by integrating reconstruction deviations and temporal state change information, thereby improving the ability to perceive concealed and associated anomalies. This research focuses on anomaly data in open-source microservices. Results show that the proposed method effectively improves the accuracy and stability of anomaly detection in distributed tasks and demonstrates good applicability in complex operating environments. This study provides a feasible technical framework for distributed intelligent monitoring under privacy constraints and offers new methodological support for research on the operational status analysis and anomaly identification of complex systems.

Keywords: Federated collaborative modeling; self-supervised reconstruction; temporal anomaly identification; distributed intelligent monitoring.

1. Introduction

With the rapid development of cloud computing, edge computing, and microservice architectures, distributed tasks for multi-node collaborative execution have been widely applied in complex scenarios such as intelligent manufacturing, smart transportation, financial services, industrial internet, and urban governance. These tasks typically feature long execution chains, numerous participating nodes, rapid state evolution, and highly heterogeneous operating environments[1]. This makes the system susceptible to resource contention, chain jitter, task congestion, service imbalance, and the propagation of localized faults during long-term operation, thus inducing complex anomalies. Compared to traditional centralized systems, anomalies in distributed task scenarios often exhibit cross-node propagation, cross-stage transmission, and cross-modal coupling characteristics, increasing the difficulty of anomaly identification and placing higher demands on the overall stability, security, and service continuity of the system. Therefore, researching

collaborative anomaly detection in distributed task scenarios is both a practical requirement for ensuring the reliable operation of complex intelligent systems and an important foundation for promoting intelligent operation and maintenance and autonomous decision-making capabilities[2].

Existing distributed anomaly detection methods still face several key challenges in practical applications[3,4]. On the one hand, time-series data generated by distributed tasks typically exhibit characteristics such as high dimensionality, asynchronicity, non-stationarity, and significant contextual dependencies. Traditional methods based on manual feature construction or shallow pattern matching struggle to fully characterize the deep-seated patterns of task state evolution over time, resulting in insufficient ability to identify hidden anomalies, gradual anomalies, and correlated anomalies. On the other hand, data in distributed environments is usually stored across different nodes or institutions. Limited by factors such as privacy protection, data security, and communication costs, centralized aggregation of raw data for unified modeling is often difficult to implement. This necessitates that anomaly detection not only possess the ability to model complex temporal dependencies but also must consider cross-node collaborative learning, knowledge sharing, and model optimization under privacy constraints. How to effectively fuse multi-node anomaly patterns without exposing the raw data has become a core issue in intelligent detection research for distributed tasks.

The combination of federated learning and self-supervised temporal modeling offers a promising solution to these problems. Federated learning, while preserving the local storage characteristics of data, achieves cross-node knowledge aggregation through collaborative parameter updates, providing theoretical support and a technical path for privacy-preserving modeling in distributed task scenarios. Self-supervised temporal modeling can leverage large-scale unlabeled runtime data to mine potential structural information from the perspectives of time dependence, context reconstruction, and state consistency, enhancing the model's ability to represent complex dynamic patterns[5]. Organically integrating these two approaches helps improve the ability to perceive anomalies in multi-source heterogeneous temporal signals, understand cross-node anomaly correlations, and characterize the state evolution of complex tasks, all while ensuring data security and controllable communication. Research on distributed task anomaly collaborative detection based on federated learning and self-supervised temporal modeling not only has significant theoretical value but also important application implications for building secure, reliable, robust, and highly autonomous distributed intelligent systems.

2. Background

Distributed tasks are typically completed by multiple computing nodes, service instances, or execution units, and their operation involves continuous interaction across multiple stages, including task scheduling, resource allocation, state synchronization, message passing, and result aggregation. In real-world systems, different nodes often assume different functional roles and operate in environments with varying network conditions, hardware configurations, load levels, and data sources[6]. Therefore, the system as a whole exhibits significant heterogeneity, dynamism, and coupling. As task scale and business complexity increase, the system state is no longer determined solely by a single node but is influenced by the collaborative relationships between multiple nodes, temporal dependencies, and the transmission effects of local behaviors. This complex operational mechanism means that anomalies are no longer limited to single-point failures or instantaneous fluctuations but can manifest in various forms, such as decreased task execution efficiency, accumulated state shifts, imbalances in collaborative links, amplification of local disturbances, and global performance degradation. To accurately understand the formation process of distributed task anomalies, it is necessary to construct a more systematic theoretical framework from the perspectives of system collaboration mechanisms, state evolution patterns, and the characteristics of multi-source data association.

Against the backdrop of the continuous development of intelligent detection technology, anomaly identification research has gradually shifted from the traditional paradigm relying on manual rules and supervised labels to a data-driven paradigm that places greater emphasis on representation learning and generalization capabilities. For distributed task scenarios, a large amount of runtime logs, performance metrics, behavioral trajectories, and event sequences naturally possess temporal continuity and structural correlation, providing a rich information foundation for time-series modeling. However, this type of data often suffers from problems such as scarce labeling, complex patterns, and ambiguous anomaly boundaries, making it difficult to support high-quality modeling by simply relying on explicit labels[7]. At the same time, due to practical constraints such as privacy protection, data security, and management boundaries, it is difficult for different nodes or institutions to directly share raw data, prompting collaborative learning mechanisms to become an important direction for complex system modeling. Federated learning emphasizes completing knowledge collaborative updates without data leaving the local machine, while self-supervised learning emphasizes mining inherent supervisory signals from unlabeled data. These two approaches, from the perspectives of collaborative mechanisms and representation learning, respectively, address the key needs of anomaly detection in distributed tasks and lay an important foundation for subsequently building detection methods that balance security, effectiveness, and scalability.

Recent studies further clarify the technical basis for collaborative anomaly detection in heterogeneous distributed systems. Adaptive representation learning and temporal state modeling strengthen anomaly detection by capturing evolving runtime states under distributional change[8]. Structure-aware root cause localization for microservice backends combines graph representation learning with causal inference, supporting the need to model service dependencies rather than treating monitoring signals as isolated sequences[9]. Variational autoencoder-based query plan anomaly detection extends this representation-oriented view to database systems, where cost deviation warning depends on reconstructing expected execution behavior and identifying abnormal departures from that baseline[10]. Together, these studies support this paper's emphasis on self-supervised temporal representation, cross-node structural dependency, and privacy-preserving collaboration.

3. Methodology

To address the privacy sensitivity and temporal complexity of distributed task monitoring, the proposed method builds a collaborative anomaly detection framework that couples federated optimization with self-supervised sequence reconstruction. Each client node first transforms its local execution stream into ordered segments, where the observation at time step t is denoted by $\boldsymbol{x}_t \in \mathbb{R}^d$ and the corresponding window is written as $\boldsymbol{S}_i = [\boldsymbol{x}_{i,1}, \boldsymbol{x}_{i,2}, \ldots, \boldsymbol{x}_{i,T}]$. Such a formulation preserves short-term fluctuations and long-range dependency within task evolution, which is essential because abnormal behaviors in distributed environments are often reflected by delayed propagation rather than isolated deviations. This paper presents the overall model architecture, as shown in Figure 1.

A temporal encoder f_θ is then employed to map raw sequences into latent states that summarize context-dependent dynamics across heterogeneous nodes. The hidden representation of the tth step is defined as:

$$\boldsymbol{h}_{i,t} = f_\theta(\boldsymbol{x}_{i,1:t})$$

where $\boldsymbol{h}_{i,t}$ aggregates the historical information before time t and provides a compact basis for subsequent reconstruction and anomaly discrimination. To avoid over-dependence on manually labeled anomalies, a stochastic masking operator $\mathcal{M}(\cdot)$ is introduced to remove a subset of temporal positions and force the model to infer missing states from surrounding evidence, thereby encouraging the encoder to capture intrinsic temporal regularities rather than superficial patterns. The masked sequence is expressed as:

$$\widetilde{\boldsymbol{S}}_i = \mathcal{M}\left(\boldsymbol{S}_i; \rho\right)$$

with masking ratio ρ controlling task difficulty and preventing trivial identity mapping. This design allows the model to exploit abundant unlabeled operation traces and improves robustness when anomalous labels are sparse, delayed, or partially unreliable across different clients.

Beyond latent encoding, the self-supervised objective is centered on temporal reconstruction so that structurally inconsistent behaviors become measurable through prediction errors. A decoder g_ϕ receives the contextual states generated from the masked sequence and reconstructs the original observations at the removed locations, which can be written as:

$$\widehat{\boldsymbol{x}}_{i,t} = g_\phi\left(\widetilde{\boldsymbol{h}}_{i,t}\right)$$

where $\widetilde{\boldsymbol{h}}_{i,t}$ denotes the hidden state derived from $\widetilde{\mathbf{S}}_i$. Since abrupt disturbances and progressive degradation may both alter the expected temporal continuity of distributed tasks, the learning target is not limited to pointwise recovery but further regularized by a temporal smoothness constraint over adjacent latent states. The local self-supervised loss is therefore formulated as:

$$\mathcal{L}_i^{self} = \frac{1}{\left|\Omega_i\right|}\sum_{t\in\Omega_i}\left\|\boldsymbol{x}_{i,t} - \widehat{\boldsymbol{x}}_{i,t}\right\|_2^2 + \lambda\sum_{t=2}^{T}\left\|\boldsymbol{h}_{i,t} - \boldsymbol{h}_{i,t-1}\right\|_2^2$$

where Ω_i denotes the masked index set on client i and λ balances reconstruction fidelity against temporal consistency. By coupling recovery quality with latent continuity, the model becomes sensitive to both local signal corruption and abnormal state transitions, which is especially valuable for distributed tasks whose failures may emerge as subtle trajectory distortions before manifesting as explicit breakdowns.

Meanwhile, collaborative learning is realized without transferring raw task data to a central server, thereby preserving data locality and reducing cross-domain privacy exposure. After local optimization on each client, the server aggregates parameters from participating nodes to construct a shared global model that captures common temporal priors across heterogeneous environments. The federated update rule is defined as:

$$\Theta^{(r+1)} = \sum_{i=1}^{N}\alpha_i^{(r)}\,\Theta_i^{(r)}$$

where $\Theta_i^{(r)}$ denotes the local parameter set uploaded by client i at communication round r, and the aggregation weight $\alpha_i^{(r)}$ reflects the relative contribution of each client. Rather than assigning uniform importance, the framework introduces a reliability-aware weighting strategy so that nodes with more stable temporal structures and lower reconstruction uncertainty exert stronger influence on the shared model. The normalized weight is computed as:

$$\alpha_i^{(r)} = \frac{n_i \exp\left(-\beta u_i^{(r)}\right)}{\sum_{j=1}^{N} n_j \exp\left(-\beta u_j^{(r)}\right)}$$

with n_i representing the local sample size, $u_i^{(r)}$ denoting the uncertainty statistic estimated from the self-supervised loss distribution of client i, and β controlling the attenuation strength. Such an adaptive federation mechanism alleviates the negative effect of noisy or severely shifted clients and strengthens global knowledge transfer under non-independent and non-identical data distributions.

Finally, anomaly assessment is performed by measuring the mismatch between observed behaviors and self-supervised expectations in the learned temporal space. For an incoming sequence segment, the model evaluates both reconstruction discrepancy and hidden state instability, because genuine anomalies in distributed tasks often manifest as simultaneous observation distortion and dynamic inconsistency. The anomaly score at time step t is defined as:

$$a_{i,t} = \eta \|\boldsymbol{x}_{i,t} - \widehat{\boldsymbol{x}}_{i,t}\|_2^2 + (1-\eta)\|\boldsymbol{h}_{i,t} - \boldsymbol{h}_{i,t-1}\|_2^2$$

where η balances the contribution of feature-level reconstruction and temporal transition deviation. A sequence segment is then flagged when its cumulative abnormality exceeds an adaptive threshold derived from the local score distribution, which enables node-specific calibration while maintaining a shared representational basis through federation. In this way, the overall framework simultaneously supports privacy-preserving collaboration, unlabeled temporal structure mining, and fine-grained anomaly identification, making it suitable for distributed task environments in which abnormal events are sparse, evolving, and strongly coupled across execution stages.

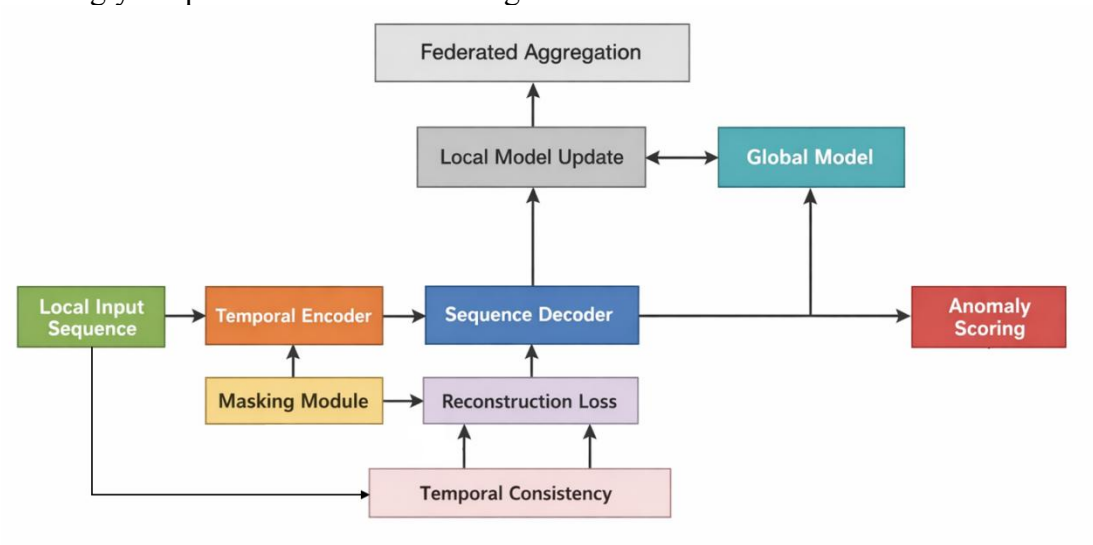


Figure 1. Overall model architecture

4. Experimental Results and Analysis

4.1 Dataset

This paper uses the LO2 distributed-task anomaly dataset. This dataset is a publicly released open-source data resource for anomaly detection in distributed systems, available as a data warehouse on Zenodo. The corresponding paper explains that its collection goal is to support anomaly detection and architectural degradation analysis in distributed systems. The data originates from the runtime of the production-grade open-source distributed system Light OAuth2, covering logs, performance metrics, and some tracing information. It can comprehensively reflect the state changes, call behaviors, and anomaly disturbance characteristics of distributed tasks during multi-node collaborative execution. Since this research focuses on collaborative anomaly detection in distributed task scenarios, the LO2 dataset is highly consistent with the

paper's theme in terms of task chain distribution, integrity of multi-source monitoring signals, and scalability of anomaly analysis.

In terms of data scale and content structure, the LO2 dataset contains approximately 657,000 log files, over 2 billion log records, and over 45 million metric files, involving 485 unique metrics, providing sufficient dynamic observation information for time-series modeling. This dataset contains runtime records for both normal and erroneous calls, labeled according to the type of error triggered. Therefore, it is suitable for anomaly detection research and for further supporting local modeling and cross-node knowledge aggregation analysis in federated scenarios. Compared to data resources containing only single logs or single metrics, LO2 better reflects the complex characteristics of distributed task anomalies in terms of temporal evolution, multimodal coupling, and local perturbation propagation, thus serving as a suitable data foundation for the construction and validation of the methods presented in this paper.

4.2 Experimental Results and Analysis

To objectively present the performance of the proposed method in distributed task anomaly collaborative detection scenarios, representative studies related to federated learning, temporal anomaly detection, microservice anomaly identification, and distributed operating environments were selected as comparison objects. The experimental results are shown in Table 1. Related work covers federated temporal anomaly detection, federated anomaly representation learning, and anomaly detection methods for microservice scenarios, providing a relatively comprehensive basis for comparison in terms of accuracy, precision, recall, and comprehensive discriminative ability.

Table 1. Experimental results compared with other models

Method	ACC	Precision	Recall	F1
Zhang et al.[11]	91.24	90.87	90.43	90.65
Zhu et al.[12]	92.01	91.55	91.12	91.33
Liu et al.[13]	92.48	92.03	91.76	91.89
Eng et al.[14]	91.73	91.26	90.95	91.10
Hao et al.[15]	93.12	92.68	92.24	92.46
Raeiszadeh et al.[16]	93.56	93.11	92.85	92.97
Xie et al.[17]	94.03	93.58	93.24	93.41
Ours	**95.62**	**95.18**	**94.86**	**95.01**

The proposed algorithm achieves superior recognition performance in collaborative anomaly detection tasks, demonstrating that the constructed method can more fully characterize the complex state changes during the operation of distributed tasks. While accurately identifying normal and abnormal behaviors, the method also maintains good discriminative stability, indicating an effective synergy between the federated learning mechanism and self-supervised temporal modeling. On the one hand, federated learning enhances the knowledge-sharing capability in multi-node environments, enabling the model to learn more universal anomaly representations without relying on aggregation from the original dataset; on the other hand, self-supervised temporal modeling further strengthens the ability to capture the evolutionary patterns of task states, giving anomaly patterns clearer, distinguishable boundaries in the latent representation space. Therefore, the proposed method not only improves overall detection performance but also provides more reliable modeling support for anomaly recognition in complex distributed scenarios.

Furthermore, the proposed algorithm exhibits strong consistency across multiple evaluation metrics, indicating that the method not only achieves improvement in a single dimension but also demonstrates good performance in terms of the completeness, stability, and comprehensive recognition capability of anomaly judgment. The main reason for this result lies in the fact that the proposed framework balances local temporal feature learning with global collaborative knowledge aggregation, enabling the model to maintain sensitivity to local anomaly details while also absorbing potential anomaly patterns shared between different nodes. Simultaneously, the reconstructive self-supervised objective helps to mine deeper dynamic structural information from unlabeled temporal data, thereby improving the ability to perceive hidden and complex anomalies. In summary, the proposed method demonstrates good effectiveness and applicability in distributed task anomaly detection, providing strong support for intelligent anomaly monitoring of real-world complex systems.

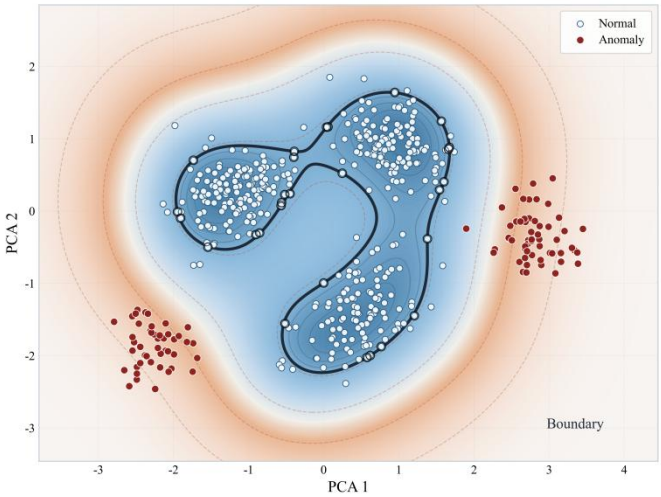


Figure 2. An experiment to visualize the decision boundary for anomaly detection.

As shown in the decision boundary visualization in Figure 2, the proposed method forms a relatively clear and continuous anomaly discrimination region in the latent representation space. Normal samples are mainly concentrated inside the boundary and have a relatively compact aggregation structure, indicating that the model can learn the inherent temporal patterns under normal operation of distributed tasks well. Anomaly samples are mostly located outside the boundary and are clearly separated from the normal sample region, indicating that the proposed federated collaborative learning and self-supervised reconstruction mechanism can effectively enhance the distinguishability of anomaly states. At the same time, the boundary shape is not a simple regular segmentation, but can adaptively adjust with the distribution of normal samples. This shows that the model can not only maintain the ability to characterize complex local structures, but also can identify nonlinear anomaly patterns, thus further demonstrating the good discriminativeness and interpretability of the proposed method in the anomaly detection scenario of distributed tasks.

The visualization results of learning rate sensitivity show that the proposed algorithm maintains good anomaly detection capability and optimization stability under reasonable parameter configuration, indicating a strong adaptability between federated collaborative updates and self-supervised temporal modeling. As shown in Figure 3, the learning rate setting not only affects the efficiency of the local model in absorbing temporal structure information, but also further affects the quality of anomaly representation after global parameter aggregation. Current results demonstrate that the proposed method can achieve a relatively reliable convergence state and discriminative ability during parameter adjustment. Overall, this figure further illustrates that the constructed model framework has good training controllability and practical application potential, providing stable and effective technical support for anomaly detection in distributed task scenarios.

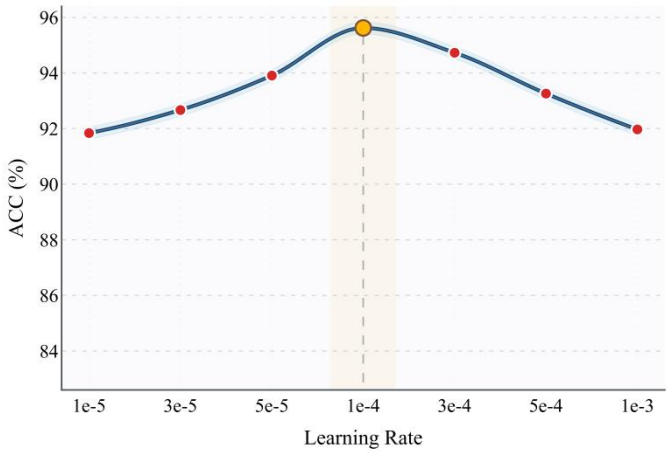


Figure 3. The impact of the learning rate on ACC

5. Conclusion

This paper addresses the challenges of anomaly detection in distributed task scenarios, including strong privacy constraints, complex temporal dependencies, covert anomaly propagation, and difficulties in cross-node collaboration. It proposes a collaborative anomaly detection method based on federated learning and self-supervised temporal modeling. This method leverages the widely available multi-node runtime data in distributed environments, achieving cross-client knowledge aggregation without directly sharing original data. Furthermore, it utilizes a self-supervised reconstruction mechanism to mine dynamic structural information in unlabeled temporal data, thereby enhancing the ability to identify complex anomaly patterns. To achieve this goal, the paper systematically studies problem modeling, framework design, collaborative optimization, and anomaly judgment mechanisms, ensuring the method adapts to the heterogeneity of distributed task data while maintaining the robustness, generalization, and scalability required for anomaly detection. Overall, this work provides a feasible technical path for anomaly detection in distributed tasks that balances data security and intelligent modeling, and also offers new research ideas for analyzing the operational state of complex systems.

From a methodological perspective, the significance of this research lies not only in introducing federated learning into distributed task anomaly detection but also in its organic integration with self-supervised temporal modeling, enabling collaborative learning capabilities and temporal representation capabilities to complement each other within a unified framework. Federated learning ensures data locality and privacy in multi-node environments, while self-supervised temporal modeling reduces reliance on high-quality anomaly labels and improves the model's applicability to real-world complex scenarios. Based on this design, the model can extract more discriminative latent state information from multi-source execution sequences and form a more stable anomaly identification boundary during global knowledge aggregation. Since distributed tasks are widely present in cloud computing platforms, microservice systems, industrial internet, intelligent transportation, edge computing networks, and fintech infrastructure, the proposed method has strong application extension value. Its research results not only help improve the intelligence level of anomaly monitoring but also provide important support for task scheduling optimization, service stability assurance, fault early warning, and autonomous system operation and maintenance.

From an application impact perspective, this work has significant practical implications for promoting the secure operation of complex digital infrastructure. In modern distributed systems, anomaly events often no longer manifest as isolated failures but can spread across multiple nodes through task chains, resource scheduling relationships, and service call dependencies, making traditional methods relying on single-point modeling or centralized analysis insufficient for practical needs. This paper's method emphasizes cross-node

collaborative learning and dynamic temporal understanding, which more closely approximates the operational mechanisms of real-world distributed tasks. Therefore, it has strong potential for application in high-reliability computing platforms, intelligent manufacturing process monitoring, online business service assurance, and the operation and maintenance of digital systems in critical industries. Especially in the context of strict data privacy requirements, limited cross-domain data flow, and continuously increasing demand for anomaly monitoring, this research provides a valuable framework for intelligent anomaly analysis under privacy protection conditions and lays the methodological foundation for building a secure, reliable, autonomous, and collaborative intelligent monitoring system.

Looking ahead, there is still significant room for research expansion in distributed task anomaly detection. As system scale continues to expand, task relationships evolve, and operating environments become more open, anomaly patterns will exhibit greater dynamism, complexity, and scenario dependence. Therefore, future research can further consider online federated update mechanisms for continuously evolving environments to enhance the model's adaptability to long-term concept drift and new anomalies. Simultaneously, combining richer multimodal operational signals and deeply integrating the correlation information between logs, metrics, call chains, and event sequences can improve the understanding and early warning capabilities of complex anomaly causes. Furthermore, at the practical deployment level, how to further reduce communication overhead, improve training efficiency, enhance model interpretability, and achieve deep integration with automated operation and maintenance platforms will also become important directions worthy of attention. Overall, this research not only provides theoretically valuable and practically significant methodological support for distributed task anomaly collaborative detection but also actively promotes the development of related fields towards greater security, intelligence, and autonomy.

References

[1] M. Raeiszadeh, A. Ebrahimzadeh, A. Saleem, et al., "Real-time anomaly detection using distributed tracing in microservice cloud applications," in Proc. IEEE 12th Int. Conf. Cloud Networking (CloudNet), 2023, pp. 36–44.

[2] R. Xu, H. Miao, S. Wang, et al., "PeFAD: A parameter-efficient federated framework for time series anomaly detection," in Proc. 30th ACM SIGKDD Conf. Knowledge Discovery and Data Mining, 2024, pp. 3621–3632.

[3] I. Iwan, T. A. Bukit, B. N. Yahya, et al., "Federated learning framework for collaborative time series anomaly detection on distributed machines," in Proc. IEEE 48th Annu. Computers, Software, and Applications Conf. (COMPSAC), 2024, pp. 1665–1670.

[4] J. Huang, D. Xu, and T. Yang, "Fed-SMAE: Federated-Learning Based Time Series Anomaly Detection with Shared Memory Augmented Autoencoder," in Proc. IEEE 7th Int. Conf. Industrial Cyber-Physical Systems (ICPS), 2024.

[5] F. Chen, Y. Zhang, Z. Qin, et al., "Learning multi-pattern normalities in the frequency domain for efficient time series anomaly detection," in Proc. IEEE 40th Int. Conf. Data Engineering (ICDE), 2024, pp. 747–760.

[6] Y. Li, M. A. Al Mamun, C. Wu, et al., "Self-Supervised Learning for Time Series Analysis: Taxonomy, Progress, and Prospects," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 46, no. 10, pp. 6775–6794, 2024.

[7] M. AL-Naday, V. Dobre, M. Reed, et al., "Federated deep Q-learning networks for service-based anomaly detection and classification in edge-to-cloud ecosystems," Annals of Telecommunications, vol. 79, pp. 165–178, 2024.

[8] Y. Zhan, "Adaptive representation learning and temporal state modeling for robust anomaly detection in distributed systems," 2024.

[9] L. Wang and C. Y. Chang, "Structure-aware root cause localization for microservice backends via joint graph representation learning and causal inference," 2024.
[10] L. Ren, S. Li, and Z. Liu, "Variational autoencoder-based query plan anomaly detection and cost deviation warning for database systems," 2024.
[11] K. Zhang, Y. Jiang, L. Seversky, et al., "Federated variational learning for anomaly detection in multivariate time series," in Proc. IEEE Int. Performance, Computing, and Communications Conf. (IPCCC), 2021, pp. 1–9.
[12] W. Zhu, D. Song, Y. Chen, et al., "Deep federated anomaly detection for multivariate time series data," in Proc. IEEE Int. Conf. Big Data (Big Data), 2022, pp. 1–10.
[13] F. Liu, C. Zeng, L. Zhang, et al., "FedTADBench: Federated time-series anomaly detection benchmark," in Proc. IEEE 24th Int. Conf. High Performance Computing & Communications; 8th Int. Conf. Data Science & Systems; 20th Int. Conf. Smart City; 8th Int. Conf. Dependability in Sensor, Cloud & Big Data Systems & Application (HPCC/DSS/SmartCity/DependSys), 2022, pp. 303–310.
[14] T. Q. Eng, H. K. Pao, and C. C. Liao, "Self-supervised federated learning for anomaly detection," in Proc. IEEE Int. Conf. Big Data (BigData), 2023, pp. 5770–5779.
[15] J. Hao, P. Chen, J. Chen, et al., "Multi-task federated learning-based system anomaly detection and multi-classification for microservices architecture," Future Generation Computer Systems, vol. 159, pp. 77–90, 2024.
[16] S. Zhang, T. Xu, J. Zhu, et al., "Privacy-preserving MTS anomaly detection for network devices through federated learning," Information Sciences, p. 121590, 2024.
[17] Z. Xie, H. Xu, W. Chen, et al., "Unsupervised anomaly detection on microservice traces through graph VAE," in Proc. ACM Web Conf. 2023, 2023, pp. 2874–2884.