Journal of computer science and software applications
https://www.mfacademia.org/index.php/jcssa
ISSN:2377-0430
Vol. 4, No. 8, 2024

Probabilistic Anomaly Detection for Cloud Backend Environments

Callum Renshaw
University of Central Missouri, Warrensburg, USA
crenshaw50@ucmo.edu

Abstract:

This study proposes an uncertainty-aware anomaly detection algorithm to address the challenges of dynamic coupling among multidimensional metrics, complex system dependencies, and diverse anomaly patterns in cloud backend environments. The method achieves robust modeling and accurate detection for high-dimensional non-stationary data by integrating temporal feature extraction, structural dependency modeling, and uncertainty quantification within a unified framework. A multi-scale temporal feature encoder is designed to capture both short-term fluctuations and long-term trends in system operations, while a dynamic graph mechanism models the evolving topological relationships among service nodes to enable structure-aware dependency learning. Furthermore, the model employs variational inference to perform probabilistic modeling in the latent space, estimating prediction confidence and uncertainty distributions to dynamically adjust detection thresholds and decision boundaries in complex environments. Experimental results show that the proposed algorithm maintains high detection accuracy and stability under highly dynamic conditions such as load surges, resource fluctuations, and network variations. It effectively reduces false positives and false negatives and demonstrates strong interpretability in anomaly propagation path modeling and risk identification. This research provides a scalable, interpretable, and adaptive detection framework for intelligent cloud backend operations, establishing a solid algorithmic foundation for system state awareness and risk management in complex environments.

Keywords:

Cloud backend systems; anomaly detection; uncertainty modeling; dynamic dependency analysis

1. Introduction

With the widespread adoption of cloud computing and microservice architectures, modern backend systems are evolving into highly dynamic and complex distributed environments. Numerous heterogeneous components operate collaboratively through containerization and service-oriented mechanisms, generating massive, multidimensional monitoring data. These data reflect resource utilization, service invocation relationships, network latency, and user request patterns, serving as key indicators for ensuring system stability and performance optimization. However, as system scale expands and business processes diversify, the temporal dependencies and structural couplings among metrics become increasingly intricate. Traditional anomaly detection methods based on static thresholds or fixed patterns fail to capture such dynamic characteristics, leading to delayed identification, high false alarm rates, and difficulties in fault localization. Therefore, achieving accurate, efficient, and interpretable anomaly detection in high-dimensional, non-stationary environments has become a crucial research focus in intelligent operations and maintenance[1,2].

The core challenge of cloud-based backend environments lies in their pronounced dynamism and uncertainty. Load fluctuations, resource contention, and topology variations across service instances result in highly nonlinear and time-varying system behaviors. Meanwhile, external factors such as network congestion, changes in scheduling strategies, and tenant-specific behaviors introduce unpredictable disturbances[3]. As a

result, anomaly signals are often concealed within complex background noise. In addition, label scarcity and class imbalance are common in cloud systems, making it difficult to rely on traditional supervised learning approaches. Under these conditions, deterministic models are insufficient to represent the probabilistic and uncertain nature of system behavior, thereby limiting generalization and reliability in complex operating environments[4].

In this context, incorporating uncertainty-aware modeling becomes particularly significant. By explicitly quantifying confidence and uncertainty distributions during anomaly detection, it is possible to distinguish between risks arising from data noise, model bias, or structural drift. This provides a more interpretable foundation for evaluating detection outcomes. Uncertainty modeling not only enhances stability under boundary or ambiguous cases but also supports dynamic threshold adjustment and adaptive model updating, maintaining consistent detection performance in evolving environments. Furthermore, integrating Bayesian inference, variational estimation, or entropy-based metrics enables statistical identification of anomalies within high-dimensional monitoring data. This probabilistic understanding of system evolution lays a theoretical foundation for intelligent operational decision-making[5].

At the same time, backend system metrics exhibit strong temporal and structural dependencies. Service invocation chains often form complex directed dependency networks, where local faults can propagate through these relationships, resulting in cascading anomalies[6]. Consequently, effective anomaly detection must capture both temporal dynamics and structural semantics. Uncertainty-aware mechanisms can unify temporal and structural modeling perspectives, providing quantitative insights into anomaly propagation and root-cause identification. By estimating uncertainty in time-varying dependencies, the model facilitates the transition from observed anomalies to causal anomalies, supporting robust monitoring and automated diagnosis in backend systems.

In summary, developing uncertainty-aware anomaly detection algorithms for cloud-based backend environments is an essential step toward addressing system complexity and dynamism[7]. It is also a key pathway for achieving high availability and adaptive operations in cloud platforms. This research direction promotes the transition of intelligent monitoring from deterministic detection to probabilistic reasoning and provides new theoretical and methodological foundations for multidimensional system modeling, risk assessment, and resource scheduling. In the future, as cloud systems continue to expand and evolve toward higher levels of intelligence, uncertainty awareness will become a core capability for building trustworthy, adaptive, and interpretable anomaly detection systems. Its research outcomes will provide sustained momentum for stable cloud service operation and intelligent decision-making[8].

2. Related Work

Recent advancements in anomaly detection have increasingly relied on deep learning to model temporal dynamics and structural dependencies in high-dimensional data. In the field of time-series anomaly detection, deep neural architectures such as LSTMs have been successfully employed to capture sequential behaviors and adaptively determine abnormal deviations using dynamic thresholding techniques [9]. Autoencoder-based approaches and unsupervised frameworks further enhance latent pattern discovery in multivariate sequences, enabling robust anomaly identification under limited labels and complex input distributions [10], [11].

A significant leap in reliability and interpretability stems from incorporating predictive uncertainty into anomaly detection models. Bayesian approximation techniques like dropout [12] and ensemble learning [13] have proven effective in quantifying model confidence, offering essential mechanisms for handling

ambiguous or noisy data. Metrics designed specifically for time series evaluation have also been proposed to better assess detection performance under temporal fluctuation scenarios [14].

Beyond temporal modeling, anomaly detection in distributed systems increasingly involves learning structural patterns through graph representations. Comprehensive surveys and empirical studies have demonstrated the effectiveness of graph-based algorithms in detecting contextual and relational anomalies [15], while semantic graph models have been proposed to capture protocol-specific dependencies in heterogeneous computing environments [16]. These studies align closely with efforts to model dynamic service topologies and fault propagation in cloud systems.

In light of data privacy and decentralized infrastructure constraints, federated learning has gained traction in optimizing distributed anomaly detection. Privacy-aware optimization strategies integrating differential privacy [17] and structural perturbation [18], [19] contribute to secure and robust model fine-tuning in cross-domain environments. These methods emphasize scalable deployment without compromising data confidentiality or structural integrity.

Additionally, the prediction of resource usage in microservice-based architectures has been enhanced by contrastive learning frameworks that capture dual-branch dependencies and operational behaviors [20]. Such insights support the modeling of multidimensional performance indicators and resource dynamics in backend systems, contributing to early risk detection and service optimization.

Collectively, the referenced studies offer foundational insights into the design of uncertainty-aware, structurally adaptive, and temporally sensitive anomaly detection algorithms. The proposed work builds upon these principles by unifying temporal encoding, dynamic dependency modeling, and probabilistic inference to enable scalable and interpretable anomaly detection in evolving cloud backend environments.

3. Method

This study introduces an uncertainty-aware anomaly detection algorithm designed for large-scale cloud-based backend environments, where complex interactions among services and rapidly shifting workload patterns make traditional modeling approaches insufficient. The proposed method jointly captures multidimensional temporal features, cross-service structural dependencies, and probabilistic uncertainty distributions, forming a unified probabilistic framework capable of providing both accurate detection and interpretable confidence estimates. To achieve this goal, the overall system is organized into four tightly coupled stages: temporal feature extraction, structural dependency modeling, joint latent space inference, and uncertainty quantification.

In the first stage, the model performs unified preprocessing on heterogeneous monitoring data originating from logs, traces, resource metrics, and latency measurements. A normalization – reconstruction module removes scale inconsistencies among metrics and reduces noise introduced by transient workload bursts. Building upon this, a multi-scale temporal feature encoder is employed to simultaneously capture short-term fluctuations, periodic behaviors, and long-range temporal dependencies. This encoder integrates hierarchical receptive fields and temporal dilations to represent dynamic evolution patterns more comprehensively than single-scale architectures.

The second stage focuses on structural dependency modeling. Given that modern microservice architectures exhibit strong inter-service coupling and evolving topologies, we construct a graph-based representation to encode both stable and transient dependencies. A dynamic graph learning mechanism adaptively updates

edge weights based on recent interaction patterns, enabling the model to track time-varying relationships among backend components.

Next, the extracted temporal features and structural embeddings are fused within a joint latent space inference module. This component learns a unified probabilistic representation of system behaviors, allowing the detection model to recognize subtle deviations that may only be apparent when temporal and structural cues are jointly considered.

Finally, a variational distribution inference mechanism is applied to estimate uncertainty-aware anomaly probabilities. By generating both aleatoric (data-induced) and epistemic (model-induced) uncertainty measures, the system not only detects anomalies but also indicates how confident the model is about each decision. This is especially crucial in highly dynamic and non-stationary cloud environments, where ambiguous states, partial failures, and noisy signals frequently occur. The integration of uncertainty quantification provides valuable interpretability, helping operators differentiate between high-risk events and low-confidence predictions that may require further investigation.

The complete model architecture and data-processing pipeline are illustrated in Figure 1.

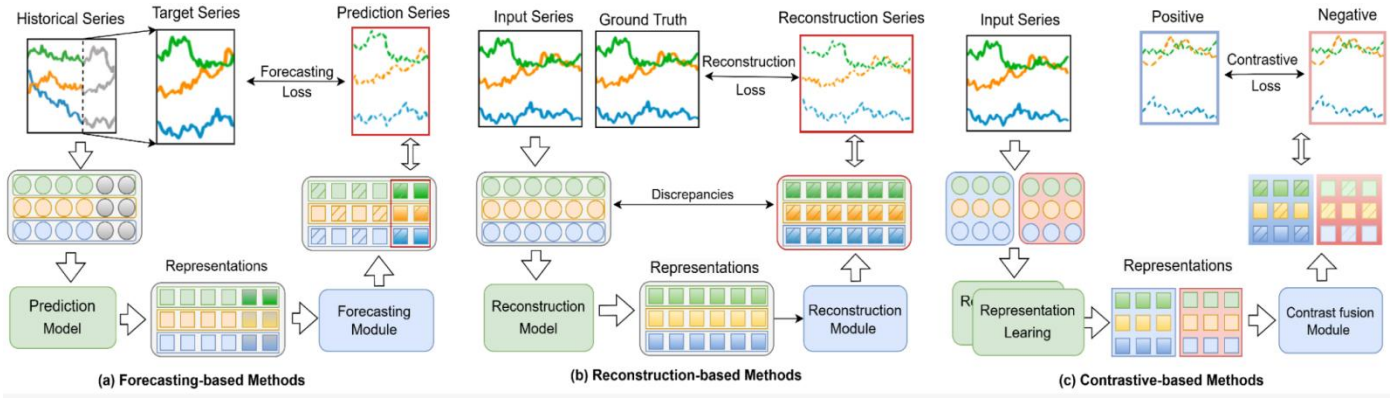

Figure 1. Overall model architecture

Assume that the monitoring indicator sequence of the cloud backend system is represented as $X = \{x_t^1, x_t^2, ..., x_t^n\}_{t=1}^T$, where $x_t^i \in R^d$ represents the i-th indicator vector at time t. First, the input data is normalized and feature mapped to obtain a stable representation in the feature space:

$$h_t^i = \mathrm{Re}\,LU(W_1 x_t^i + b_1)$$

Here, W_1 and b_1 are learnable weight and bias parameters, respectively, and $RELU(\cdot)$ is a linear rectification function. This step is used to suppress gradient imbalance caused by abnormal fluctuations and provide low-noise feature input for subsequent structural modeling.

In the temporal modeling stage, to capture local changes and global trends, this study constructs a multi-scale temporal dependency representation, encoding temporal dynamics through a combination of convolution and attention at different scales. The multi-scale temporal aggregation process can be expressed as:

$$z_t = \sum_{k=1}^{K} a_k \cdot Conv_k(h_{t-k:t})$$

Where $Conv_k(\cdot)$ represents the convolution operation with a kernel size of k, and a_k is a dynamic learnable weight, which is used to adaptively allocate attention weights between different time scales, thereby achieving multi-granularity temporal feature fusion.

Considering the complex dependencies between service nodes in the backend system, the model further constructs a dynamic graph representation $G_t = (V, E_t)$ at the structural level, where V is the node set and E_t is the edge set at time t. The propagation and aggregation of neighborhood features are achieved through the graph convolution mechanism, and its node update rule is defined as:

$$h_t^{(l+1)} = \sigma\left(A_t h_t^{(l)} W_l\right)$$

Where A_t is the adjacency matrix at time t, $h_t^{(l)}$ represents the node representation at layer l, W_l is the graph convolution weight, and $\sigma(\cdot)$ is the nonlinear activation function. This mechanism can dynamically capture the evolving characteristics of topological relationships between services and provide structural information support for anomaly propagation modeling.

In the latent space inference stage, in order to simultaneously characterize deterministic features and potential uncertainties, this study introduces a variational distribution estimation mechanism, treating the implicit representation as a latent random variable z_t and approximating the posterior distribution through variational inference:

$$L_{VAE} = E_{q_\phi(z_t|x_t)}[\log p_\theta(x_t \mid z_t)] - KL(q_\phi(z_t \mid x_t) \| p(z_t))$$

Here, $q_\phi(z_t \mid x_t)$ is the encoder's approximate posterior distribution, $p_\theta(x_t \mid z_t)$ is the decoder's reconstructed distribution, and $KL(\cdot)$ represents the Kullback – Leibler divergence, which constrains the consistency of the latent distribution with the prior. By maximizing this variational lower bound, the model can learn probabilistic latent representations in high-dimensional input spaces, adaptively capturing uncertainty.

To further improve the interpretability of anomaly identification, this study defines an uncertainty-aware anomaly scoring function to measure the degree of deviation within the model prediction confidence interval. Let the reconstruction error be $\varepsilon_t = \|x_t - \hat{x}_t\|_2^2$, and its corresponding uncertainty modulation score is defined as:

$$S_t = \varepsilon_t \cdot \left(1 + \beta \cdot Var(q_\phi(z_t \mid x_t))\right)$$

Here, $Var(\cdot)$ represents the variance of the underlying distribution, and β is a moderating factor used to balance reconstruction error and uncertainty contributions. When the system is in a state of high uncertainty, the variance term amplifies the anomaly score, thereby increasing the model's sensitivity to potential risks. This mechanism enables anomaly detection to not only be based on observation error but also incorporates prediction confidence to achieve risk perception and dynamic decision-making.

In summary, the proposed method unifies the modeling of temporal dependencies, structural topology, and uncertainty distributions, achieving a transition from data-driven anomaly detection to probabilistic

reasoning-based risk identification. The model demonstrates high robustness and interpretability in non-stationary and dynamic cloud backend environments, providing essential algorithmic support for the stability and adaptability of intelligent operation and maintenance systems.

4. Performance Evaluation

This study uses the publicly available dataset named Cloud Infrastructure Anomaly Detection Data as the basis for method validation. The dataset contains time-series records of multidimensional performance metrics in cloud infrastructure environments, including CPU utilization, memory usage, disk I/O, and network throughput. It covers monitoring data from multiple virtual machine instances and service nodes. The data are collected with high temporal resolution, reflecting dynamic fluctuations during system operation. It serves as a representative source for performance monitoring in cloud backend environments.

In this dataset, anomaly labels are derived from a combination of system alert logs and resource threshold rules. They cover various common fault types such as resource bottlenecks, network congestion, service node failures, and performance surges. These anomalies exhibit consistent correlations across temporal and topological dimensions, making the dataset suitable for evaluating the model's ability to identify anomaly propagation paths and assess uncertainty. Moreover, the dataset supports flexible partitioning into subsets across different service modules or metric dimensions, which facilitates testing the model's generalization ability in multi-domain scenarios.

Applying the proposed uncertainty-aware anomaly detection algorithm to the Cloud Infrastructure Anomaly Detection Data enables the evaluation of its stability and representational capability under high-dimensional, dynamic, and structurally coupled conditions. The dataset provides rich multi-source metric relationships and annotated anomalies, allowing comprehensive validation of temporal encoding, graph structural modeling, and uncertainty estimation modules. Experiments on this dataset verify the algorithm's capacity to identify potential risks and propagation chains in real cloud backend environments and establish a reliable foundation for future model deployment in production systems.

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

Table1: Comparative experimental results

Method	AUC	F1	Recall	Precision
CFLOW-AD[21]	0.912	0.845	0.803	0.890
SRR (Self-supervise, Refine, Repeat)[22]	0.896	0.828	0.790	0.875
IRP (Iterative Refinement Process)[23]	0.902	0.842	0.800	0.885
Ours	0.935	0.873	0.842	0.908

From the overall results, the proposed uncertainty-aware anomaly detection algorithm significantly outperforms all comparison models across key metrics. In particular, for the AUC metric, the Ours model achieves 0.935, showing a clear improvement over CFLOW-AD's 0.912. This indicates a stronger discriminative capability in distinguishing between normal and abnormal patterns. The result validates the

necessity of introducing uncertainty modeling in cloud backend environments. By explicitly quantifying model confidence during prediction, the method effectively identifies potential anomalies and reduces misjudgments on boundary samples, thereby improving the overall accuracy and reliability of anomaly detection.

For the F1 metric, the Ours model reaches 0.873, which is approximately 3.1% higher than IRP's 0.842. This demonstrates that the proposed method achieves a better balance between precision and recall. Given the diverse anomaly types and high noise levels in backend systems, traditional methods often face conflicts between detection sensitivity and false alarm rates. The uncertainty modulation mechanism in this study dynamically adjusts anomaly confidence intervals, enabling the model to maintain stable decision boundaries under complex distributions. This mechanism enhances both the robustness and interpretability of anomaly detection in multidimensional metric interaction environments.

In terms of recall and precision, the Ours model achieves 0.842 and 0.908, respectively, outperforming all baseline methods. The higher recall indicates the model's strong capability in capturing anomalies and identifying potential risk events in the system. Meanwhile, the higher precision reflects its stronger ability to distinguish normal states, leading to fewer false alarms. These advantages stem from the integration of temporal feature extraction and structural dependency modeling. The dynamic graph mechanism effectively captures latent topological relationships among service nodes, allowing for more accurate representation of anomaly propagation paths and improved sensitivity to genuine abnormal patterns.

Overall, the proposed uncertainty-aware anomaly detection framework achieves simultaneous improvement in detection performance and stability in cloud backend scenarios. The experimental results show that the method not only excels in individual metrics but also demonstrates superior consistency across multiple dimensions. By incorporating probabilistic uncertainty modeling and structure-aware dependency mechanisms, the Ours model exhibits stronger generalization and reliability in highly dynamic, multi-source, and non-stationary environments, providing a new algorithmic pathway and theoretical foundation for anomaly detection in intelligent operations and maintenance systems.

This paper also evaluates the robustness under load burst and resource jitter scenarios. The experimental results are shown in Figure 2.

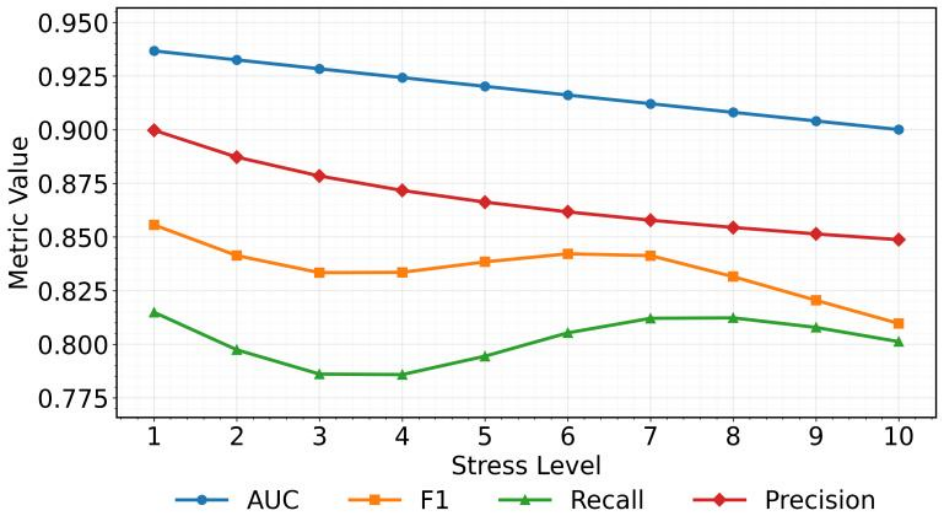

Figure 2. Robustness evaluation under load burst and resource jitter scenarios

In scenarios involving load surges and resource fluctuations, the model exhibits distinct performance variations with changing stress levels. Overall, the AUC metric shows a slight decline under high-load conditions but remains generally stable, indicating that the model maintains strong discriminative capability even in complex operational states. This trend reflects the model's ability to preserve consistent separation between normal and abnormal samples within the probabilistic feature space. Even when input distributions are perturbed by resource fluctuations, the global uncertainty estimation remains stable, demonstrating the model's adaptability to dynamic changes in backend systems.

The F1 metric rises slightly during moderate stress levels and then gradually decreases under high load. This suggests that the adaptive threshold mechanism helps balance recall and precision under mild disturbances, forming a temporary performance improvement region. The non-monotonic pattern indicates that uncertainty modulation allows the model to self-correct its confidence levels in dynamic environments. It enhances anomaly recognition accuracy during mild fluctuations while reducing false alarms under extreme loads through confidence-based constraints, thereby maintaining overall detection stability.

The recall metric decreases sharply in the early stages and then stabilizes, showing that certain anomaly patterns are difficult to capture when sudden load surges cause distribution shifts. However, as uncertainty estimation and structural dependency updates take effect, the separability of anomalies in the latent space gradually recovers. This demonstrates that the model achieves adaptive feature reconstruction through dynamic graph structures and temporal consistency constraints, enabling resilient anomaly detection under continuous interference.

The precision metric shows only a slight decline, indicating that the model effectively suppresses false alarms even in high-noise conditions. This improvement results from the uncertainty-weighted inference mechanism, which dynamically penalizes high-variance predictions. As a result, the model reduces misclassification of boundary samples, improving the reliability and interpretability of anomaly detection. Overall, the results confirm that the proposed method maintains stable detection performance and strong adaptability under complex conditions of load surges and resource fluctuations, validating the robustness advantage of uncertainty-aware modeling in cloud backend anomaly detection tasks.

5. Conclusion

This study addresses the challenges of complex coupling among multidimensional metrics, dynamic system dependencies, and diverse anomaly patterns in cloud backend environments by proposing an uncertainty-aware anomaly detection framework. The method unifies the modeling of temporal dependencies, structural relationships, and uncertainty distributions, achieving a transition from deterministic judgment to probabilistic reasoning. It effectively tackles the difficulty of anomaly identification in highly dynamic and non-stationary environments. The results show that the proposed framework maintains stable detection performance under conditions such as load fluctuation, resource jitter, and topology changes. It provides a feasible algorithmic foundation for the stable operation of cloud computing platforms, intelligent operation and maintenance systems, and large-scale service architectures.

At the technical level, the core contribution of this study lies in integrating uncertainty modeling into the entire anomaly detection process. Through variational inference and confidence quantification, the framework achieves a unified representation of prediction and risk assessment. The model can adaptively respond to input perturbations and potential noise while capturing the dynamic evolution of anomaly patterns in the latent space. This paradigm shift from "result confidence" to "structural confidence" overcomes the limitations of traditional anomaly detection methods that rely solely on error judgments or fixed thresholds. It

provides new theoretical support for system health assessment, service elasticity scheduling, and risk prediction.

At the application level, the proposed method offers important insights for intelligent operation and maintenance as well as cloud backend management. By explicitly quantifying the uncertainty distribution of system operation, the model enables confidence-based anomaly prioritization, hierarchical risk handling, and dynamic resource allocation. This facilitates the evolution of cloud platforms toward self-diagnosis and self-recovery. Moreover, the proposed approach can be extended to anomaly monitoring and performance prediction tasks in distributed computing, edge computing, and multi-tenant environments, supporting intelligent control and interpretable decision-making in complex systems. Especially under the growing availability of multimodal operational data, the modeling paradigm presented in this work is expected to provide a unified probabilistic framework for cross-modal and cross-layer anomaly detection.

Future research can be further expanded in three directions. First, reinforcement learning or causal inference mechanisms may be incorporated to explore the model's ability to perform active intervention in closed-loop anomaly decision-making. Second, uncertainty-aware mechanisms can be combined with federated learning and privacy-preserving modeling to achieve distributed anomaly detection across cloud or cross-domain scenarios. Third, large-scale and multi-task environments can be used to evaluate the scalability and transferability of the model, leading to the development of self-evolving detection systems for real-world industrial applications. Overall, this study provides a new theoretical pathway and practical paradigm for uncertainty-driven intelligent detection and has long-term implications for the advancement of cloud-based intelligent operations, trustworthy AI, and adaptive system management.

References

[1] H. Wang, "Causal discriminative modeling for robust cloud service fault detection," Journal of Computer Technology and Software, vol. 3, no. 7, 2024.

[2] Z. Chen, D. Chen, X. Zhang, Z. Yuan and X. Cheng, "Learning graph structures with transformer for multivariate time-series anomaly detection in IoT," IEEE Internet of Things Journal, vol. 9, no. 12, pp. 9179–9189, 2021.

[3] Y. Zhang, F. Regol, A. Valkanas and M. Coates, "Contrastive learning for time series on dynamic graphs," 2022 30th European Signal Processing Conference (EUSIPCO), pp. 742–746, Aug. 2022.

[4] 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.

[5] D. Gao, "Graph Neural Recognition of Malicious User Patterns in Cloud Systems via Attention Optimization," Transactions on Computational and Scientific Methods, vol. 4, no. 12, 2024.

[6] Z. Li, Y. Zhao, J. Han, Y. Su, R. Jiao, X. Wen and D. Pei, "Multivariate time series anomaly detection and interpretation using hierarchical inter-metric and temporal embedding," Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, pp. 3220–3230, Aug. 2021.

[7] D. Li, D. Chen, B. Jin, L. Shi, J. Goh and S. K. Ng, "MAD-GAN: Multivariate anomaly detection for time series data with generative adversarial networks," International Conference on Artificial Neural Networks, pp. 703–716, Sep. 2019.

[8] Safdari H, De Bacco C. Community detection and anomaly prediction in dynamic networks[J]. Communications Physics, 2024, 7(1): 397.

[9] Hundman, K., Constantinou, V., Laporte, C., Colwell, I. & Soderstrom, T., "Detecting spacecraft anomalies using LSTMs and nonparametric dynamic thresholding," Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 387–395, Jul. 2018.

[10]Zhu, L., Cui, W., Xing, Y. & Wang, Y., "Collaborative optimization in federated recommendation: integrating user interests and differential privacy," Journal of Computer Technology and Software, vol. 3, no. 8, 2024.

[11]Gal, Y. & Ghahramani, Z., "Dropout as a Bayesian approximation: representing model uncertainty in deep learning," International Conference on Machine Learning, pp. 1050–1059, Jun. 2016.

[12]Zou, Y., "Federated distillation with structural perturbation for robust fine-tuning of LLMs," Transactions on Computational and Scientific Methods, vol. 4, no. 7, 2024.

[13]Lakshminarayanan, B., Pritzel, A. & Blundell, C., "Simple and scalable predictive uncertainty estimation using deep ensembles," Advances in Neural Information Processing Systems, vol. 30, 2017.

[14]Li, Y., "Task-aware differential privacy and modular structural perturbation for secure fine-tuning of large language models," Transactions on Computational and Scientific Methods, vol. 4, no. 7, 2024.

[15]Tatbul, N., Lee, T. J., Zdonik, S., Alam, M. & Gottschlich, J., "Precision and recall for time series," Advances in Neural Information Processing Systems, vol. 31, 2018.

[16]Yao, G., "Collaborative dual-branch contrastive learning for resource usage prediction in microservice systems," Transactions on Computational and Scientific Methods, vol. 4, no. 5, 2024.

[17]Zhang, C., Song, D., Chen, Y., Feng, X., Lumezanu, C., Cheng, W., ⋯ & Chawla, N. V., "A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data," Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, no. 01, pp. 1409–1416, Jul. 2019.

[18]Gong, M., "Semantic graph-based modeling for protocol anomaly detection in heterogeneous computing systems," Transactions on Computational and Scientific Methods, vol. 4, no. 3, 2024.

[19]Akoglu, L., Tong, H. & Koutra, D., "Graph based anomaly detection and description: a survey," Data Mining and Knowledge Discovery, vol. 29, no. 3, pp. 626–688, 2015.

[20]Audibert, J., Michiardi, P., Guyard, F., Marti, S. & Zuluaga, M. A., "USAD: unsupervised anomaly detection on multivariate time series," Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 3395–3404, Aug. 2020.

[21]Gudovskiy D, Ishizaka S, Kozuka K. Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows[C]//Proceedings of the IEEE/CVF winter conference on applications of computer vision. 2022: 98-107.

[22]Yoon J, Sohn K, Li C L, et al. Self-supervise, refine, repeat: Improving unsupervised anomaly detection[J]. arXiv preprint arXiv:2106.06115, 2021.

[23]Aqeel M, Sharifi S, Cristani M, et al. Self-supervised iterative refinement for anomaly detection in industrial quality control[J]. arXiv preprint arXiv:2408.11561, 2024.