Journal of Computer Technology and Software

ISSN:2998-2383

Vol. 4, No. 4, 2025

Consistency-Aware Unified Modeling for Multi-Source System Anomaly Detection

Yi Yang[1*], Markus Nagel[2], Alan Li[2]
[1]Carnegie Mellon University, Pittsburgh, USA
[2]Stevens Institute of Technology, Hoboken, USA
[3]Stevens Institute of Technology, Hoboken, USA
[*]Corresponding author: Yi Yang; yiyang4@alumni.cmu.edu

Abstract: With the large-scale deployment of cloud native and distributed systems, operating environments have become highly dynamic and structurally complex. Traditional anomaly detection methods based on a single observation source struggle to fully capture system-level abnormal behavior. This paper focuses on intelligent operations scenarios and proposes a unified anomaly modeling approach for multi-source system observations. Anomalies are characterized and identified from the perspective of overall system operating states. Logs, metrics, and call traces are taken as inputs. Heterogeneous information is fused into a shared representation space through unified mapping to form a comprehensive description of system behavior. On this basis, state evolution modeling is introduced to capture continuous temporal changes. Anomalies are then measured through state consistency deviation. This avoids information fragmentation caused by the independent modeling of different observation sources. The framework preserves the complementary nature of each source while ensuring consistent and stable anomaly decisions. Anomalies are naturally expressed as deviations from normal operating trajectories. Comparative analysis shows that the proposed method achieves stronger overall discrimination capability and higher result consistency. This confirms the effectiveness of combining unified state modeling with temporal constraints for anomaly identification in complex systems. The proposed approach provides a unified perspective for system-level anomaly modeling with multi-source observations in intelligent operations. It improves the reliability and practical applicability of anomaly detection in complex distributed systems.

Keywords: System-level anomaly modeling; multi-source observation fusion; unified state representation; intelligent operation and maintenance.

1. Introduction

With the widespread adoption of cloud computing, microservices, and distributed architectures, modern information systems have grown rapidly in scale, complexity, and dynamism. Business systems are composed of many loosely coupled components. Their operating states evolve continuously due to workload variation, configuration changes, and environmental disturbances [1]. As a result, system behavior becomes highly nonlinear and difficult to predict. Under such conditions, failures and performance anomalies rarely appear as explicit faults in a single component. Instead, they often emerge as implicit anomalies spanning multiple services and layers. This situation poses serious challenges to system stability and business continuity. Timely and accurate anomaly identification in complex runtime environments has therefore become a central problem in intelligent operations and maintenance [2].

To improve system observability, modern operations platforms usually collect multiple types of monitoring data, including logs, metrics, and distributed traces. Logs record discrete events and semantic information. Metrics reflect continuous changes in resources and performance. Traces describe interactions among components and their dependency relationships. These data sources differ significantly in format, temporal granularity, and semantic meaning. Each of them captures only a partial view of system behavior. In real operational scenarios, anomalies rarely manifest in a single observation source. They tend to emerge gradually through coordinated changes across multiple signals. Relying on one data source alone often leads to false alarms, missed detections, or an incomplete understanding of anomaly causes [3].

Conventional anomaly detection methods are mostly designed for a single data modality. They are typically based on threshold rules, statistical assumptions, or independent modeling strategies. Such approaches struggle to capture complex correlations across heterogeneous observations. They often ignore internal system dependencies and information propagation mechanisms. As a result, they cannot effectively distinguish local fluctuations from system-level anomalies. Moreover, the lack of a unified representation across data sources makes detection results difficult to interpret consistently or reason about jointly. This increases the cognitive burden on operators. As system scale continues to grow and manual operation costs rise, experience-driven integration of fragmented anomaly signals becomes increasingly impractical [4].

In intelligent operations scenarios, a unified anomaly modeling perspective is therefore of great importance. By representing and analyzing multi-source system observations within a single framework, it becomes possible to characterize

overall system states and their evolution. This improves the ability to perceive complex anomaly patterns. Unified modeling helps alleviate information fragmentation caused by data heterogeneity. It also provides a structured foundation for anomaly explanation, impact analysis, and operational decision making [5]. In systems with complex dependencies and diverse anomaly propagation paths, unified representations are especially valuable. They help reveal the internal mechanisms behind anomaly generation and diffusion, thereby improving the reliability and coherence of operational analysis [6].

Consequently, unified anomaly modeling methods for multi-source system observations play a critical role in intelligent operations and maintenance. From a theoretical perspective, this direction promotes a shift from pointwise anomaly judgment to system-level understanding. It advances research on multimodal data fusion and system behavior modeling [7]. From a practical perspective, unified anomaly modeling enhances consistency and stability in anomaly identification. It reduces operational complexity and supports automated operations and intelligent decision processes. As systems continue to scale and runtime environments become more dynamic, studying unified anomaly modeling mechanisms for multi-source observations remains essential for building reliable and adaptive intelligent operations frameworks.

Recent studies also show how unified anomaly modeling can be strengthened by memory, graph, compression, and sequence modeling mechanisms. Memory-enhanced Transformer modeling for backend microservice anomaly detection highlights the value of preserving long-range operational context when detecting abnormal service behavior [8]. Log semantic graph construction combined with convolutional event anomaly detection further supports representing logs as structured semantic dependencies rather than isolated tokens [9]. Lightweight anomaly detection for edge intelligence through compressed representation learning indicates that unified representations should remain compact enough for resource-constrained deployment [10], while Mamba-LSTM collaborative modeling for multi-table ETL anomalies demonstrates the benefit of coupling long-sequence dynamics with heterogeneous data relationships [11]. Research on cloud infrastructure task scheduling and multi-tenant distributed inference services extends this system-level view from detection to operational coordination, showing that runtime states, resource occupancy, dynamic node conditions, and service collaboration should be jointly considered in intelligent operation frameworks [12], [13]. These studies strengthen the motivation for a consistency-aware unified anomaly model that fuses multi-source observations, tracks temporal state evolution, and supports stable downstream operational decisions.

2. Method

This paper's approach starts with modeling the overall system operational state, utilizing multi-source system observation data such as logs, metrics, and call chains to construct a unified anomaly representation and inference framework. First, different observation sources are considered as multiple perspectives depicting the same system state. Through time alignment and semantic mapping, the original observation sequence is transformed into a state description on a shared time axis. Let z_t represent the potential operational state of the system at time t. This state is not directly observable but is indirectly reflected through multi-source observation data. The differences between multi-source observations are absorbed by independent mapping functions, thus forming a comprehensive characterization of the system state in a unified space. This modeling approach avoids over-reliance on a single observation source, enabling the system operational state to be continuously and stably described at the overall level. Its overall model architecture is shown in Figure 1.

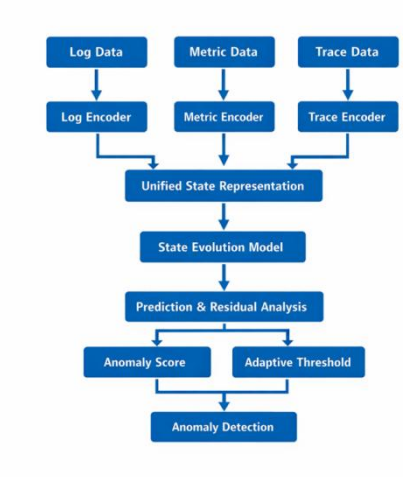

Figure 1. Overall model architecture

In the process of constructing the unified representation, for each observation source m, a mapping function $f_m(\cdot)$ from the original observation $x_t^{(m)}$ to the shared state space is defined, and the system state representation is obtained through weighted fusion:

$$z_t = \sum_{m=1}^{m} w_m f_m(x_t^{(m)})$$

Here, w_m represents the relative importance of different observation sources in state modeling, used to balance the differences in expressive power among logs, metrics, and link information. In this way, multi-source heterogeneous data are uniformly projected onto the same representation space, forming a holistic description of the system's operational state. This representation not only preserves the key information of each observation source but also provides a unified input for subsequent time-series modeling and anomaly detection.

In the time dimension, the system's operating state is considered a continuously evolving process, with inherent dependencies between states at adjacent time points. To characterize this dynamic characteristic, a state evolution function is introduced to recursively model the system state:

$$z_t = g(z_{t-1}) + \epsilon_t$$

In this modeling approach, $g(\cdot)$ describes the system's state evolution trend under normal operating conditions, while

ϵ_t represents the random disturbance term, used to characterize environmental fluctuations and instantaneous uncertainties. Under this modeling assumption, the system's normal behavior corresponds to the state evolving along a stationary trajectory, while anomalies manifest as significant deviations from this trajectory. This time-consistent modeling approach helps distinguish between short-term noise and persistent anomalous changes.

Anomaly detection is based on a unified metric of the consistency offset of the system state over time. Specifically, the deviation between the current state and its predicted state is defined as the anomaly strength index:

$$s_t = //z_t - \widehat{z}_t//_2$$

Where $\widehat{z}_t = g(z_{t-1})$ represents the prediction result obtained based on historical states. When the deviation exceeds the normal fluctuation range of the system, the system can be considered to have an anomaly at that moment. To ensure the stability of the discrimination process, an adaptive threshold based on historical statistical characteristics is introduced:

$$\theta = \mu_s - \lambda\sigma_s$$

Where μ_s and σ_s represent the mean and standard deviation of the anomaly intensity sequence, respectively, and λ is the adjustment coefficient. Through the above unified modeling and discrimination mechanism, the anomalous information in multi-source system observation data is naturally aggregated into the same state space, realizing the overall characterization and consistent judgment of system-level anomalies.

3. Experimental Evaluation

3.1 Dataset

This study adopts OpenRCA as a unified open source dataset to support the task setting of integrated anomaly detection and root cause analysis in distributed systems. OpenRCA targets failure scenarios in software systems and organizes multi-source observability data around each fault case. This design enables anomaly detection and root cause analysis to be conducted within the same data context and validated in a closed-loop manner. The dataset uses fault events as the core indexing unit. This structure makes it convenient to align multiple source observations during anomaly periods on a unified time axis. It also incorporates system dependency information into the analysis. As a result, it naturally fits the log metric trace fusion and reasoning paradigm.

In terms of data modalities, OpenRCA provides three key types of observations. The first type is logs, which are usually semi-structured or textual event streams. They are used to characterize component behavior and error semantics. The second type is KPIs or metrics, represented as time series that reflect continuous changes in performance and resource states. The third type is traces or dependency trace graphs. They describe cross-service invocation chains and inter-component dependency structures. These data sources are inherently complementary. Logs provide semantic cues. Metrics provide dynamic intensity signals. Trace graphs provide structural paths and propagation channels. Together, they enable graph-based integrated methods to achieve unified modeling of anomaly representation, structural propagation, and cause attribution within a single dataset.

With respect to annotation and task definition, OpenRCA provides root cause-related labels for each fault case, with root cause elements or locations as the target. It also offers corresponding evaluation protocols and data organization schemes. This design supports treating anomaly detection and root cause localization as continuous reasoning tasks within a single pipeline. Since the dataset contains KPI time series, dependency trace graphs, and log texts, and is organized around real fault cases, it is well-suited for constructing a log metric trace multimodal graph reasoning framework. It also enables reproducible research without relying on private production data.

3.2 Performance Results

This article first presents the results of the comparative experiments, as shown in Table 1.

Table 1. Comparative experimental results

Method	Acc	Precision	Recall	AUC
SIAS[14]	0.835	0.828	0.812	0.871
TraceDAE[15]	0.857	0.849	0.834	0.889
POSTER[16]	0.872	0.866	0.851	0.903
Automap[17]	0.888	0.881	0.869	0.916
Ours	0.914	0.907	0.898	0.943

The comparative results show a clear progressive improvement from traditional methods to unified modeling approaches. This trend indicates that, in intelligent operations scenarios, advancing anomaly identification from independent judgment on single signals to joint characterization of multi-source observations enables more stable capture of system-level anomaly patterns. Compared with approaches that rely only on local evidence, unified representations more easily form consistent anomaly criteria under complex dependency structures. As a result, they achieve stronger overall discrimination performance and demonstrate the benefits brought by multi-source information synergy.

From the perspectives of precision and completeness, the results indicate that the proposed approach reduces false alarms while also lowering the risk of missed detections. This behavior reflects the ability of the fusion strategy to better suppress noise and short-term fluctuations. The complementarity of logs, metrics, and traces is particularly important in this task. Logs provide event semantics and anomaly context. Metrics convey continuous state variations. Traces describe cross-service invocation relationships and propagation paths. After being jointly mapped into a shared state space, these sources help distinguish local jitters from genuine anomaly propagation. This improves the credibility and consistency of detection outcomes.

From an operational semantics perspective, the simultaneous improvement of related metrics implies that the

model not only detects anomalies more easily but also tends to produce more actionable anomaly decisions. This makes it suitable for use as an upstream filter in alarm reduction and automated response pipelines. In microservice and distributed systems, anomaly signals often span multiple components and propagate along call chains. Without a structured dependency view, models may misidentify symptom nodes as root causes or produce overly scattered candidates. Unified anomaly modeling aggregates multiple source observations at the system state level. This aggregation improves decision stability under propagation consistency constraints.

The performance gap with baseline methods also highlights a key fact. Single modality approaches or weak fusion strategies are prone to information fragmentation when facing heterogeneous sources and strong coupling dependencies. This leads to unclear decision boundaries. The results support the central claim of this work. Anomaly detection for intelligent operations should move beyond data-level fusion toward state-level unified modeling. It should also incorporate temporal evolution and consistency shift to form a unified anomaly measure. The overall superior performance indicates that the proposed framework aligns more closely with the generation mechanisms and manifestation patterns of real operational data. It therefore provides more effective support for robust anomaly identification in complex systems.

This sensitivity experiment was used to examine the impact of changes in the temporal aggregation scale on the stability of anomaly detection within a unified modeling framework based on multi-source observations. By adjusting the observation window length, the model's ability to balance short-term fluctuations with long-term trends can be evaluated, thus providing a more robust basis for parameter selection for different operational load scenarios. The experimental results are shown in Figure 2.

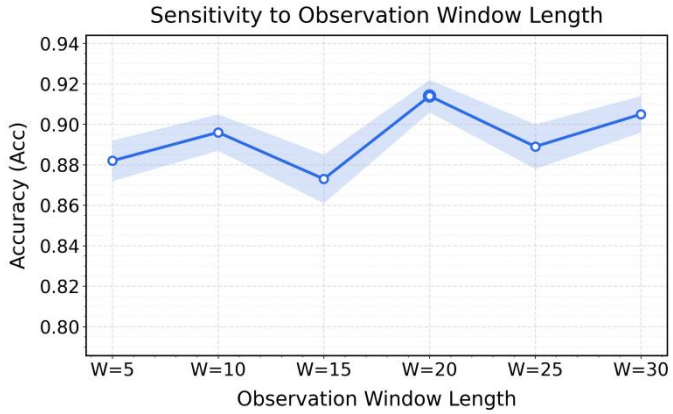

Figure 2. Sensitivity experiment of observation window length to Acc

The sensitivity results highlight the critical role of observation window length in the unified anomaly modeling framework. As the temporal window scale varies, the model shows clear differences in its ability to characterize system operating states. This indicates that anomalies are not triggered solely by instantaneous signals. They are closely related to coordinated changes in multiple source observations over a certain time range. An appropriate window length preserves anomaly evolution patterns while suppressing short-term noise. This allows the model to better reflect system-level abnormal states. Such behavior is well aligned with the overall objectives of intelligent operations-oriented modeling.

In terms of overall trends, overly short observation windows are easily affected by transient fluctuations and random disturbances. This makes the unified state representation unstable and weakens the consistency of anomaly decisions. In contrast, excessively long windows may introduce too much historical information. This can smooth or dilute anomaly characteristics. These findings suggest that, in multi-source observation fusion scenarios, the temporal aggregation scale directly influences the sensitivity of state evolution modeling to anomaly signals. The window length effectively balances local variations and global trends.

When combined with the unified state representation and temporal evolution assumptions in this work, the results indicate that the model does not rely on anomaly intensity at a single time point. Instead, it determines abnormality through temporal continuity. When the window length matches the system operating rhythm, the relationships among multiple source observations in the shared representation space become clearer. Anomaly deviations then appear in a more stable form. This improves the reliability of overall discrimination. The findings also indirectly confirm the importance of introducing state evolution mechanisms for anomaly modeling in complex systems.

From the perspective of intelligent operations applications, the sensitivity analysis emphasizes the impact of parameter selection on practical deployment. The observation window length is not merely a technical hyperparameter. It links data collection frequency, system dynamic characteristics, and anomaly response timeliness. The results show that, within a unified anomaly modeling framework, properly setting the temporal window enhances model adaptability to real system anomalies. It also provides more robust anomaly identification support for operational scenarios with different load levels and operating rhythms.

This sensitivity experiment was used to evaluate the impact of changes in the proportion of anomalous samples in the data on the model's discriminative stability within a unified anomaly modeling framework for multi-source system observations. By controlling the anomaly ratio, the adaptability and robustness of the unified state representation and temporal consistency constraints under different anomaly density scenarios can be tested. The experimental results are shown in Figure 3.

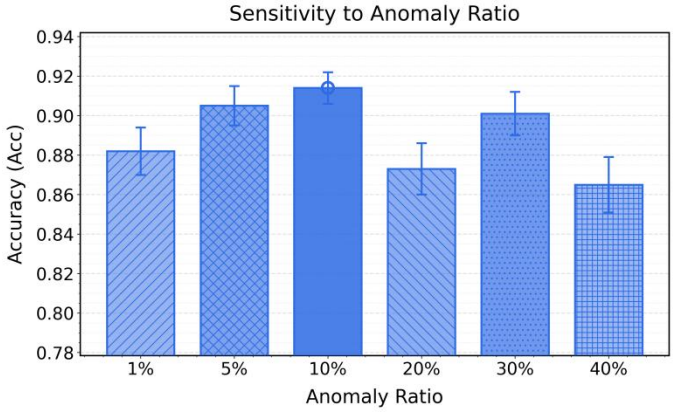

Figure 3. Sensitivity experiment of abnormal proportions to Acc

The results reveal how changes in the proportion of anomalous samples affect the stability of unified anomaly modeling. As the anomaly ratio gradually increases, the model's ability to characterize system states does not change linearly. Instead, it is jointly influenced by the degree of coordination among multiple source observations and by temporal consistency constraints. When the anomaly proportion remains within a moderate range, abnormal correlations across different signals are more likely to form clear structures in the unified state space. This helps the model stably distinguish normal evolution from anomalous deviations. Such behavior is consistent with system-level modeling objectives in intelligent operations.

When the anomaly proportion increases further, the overall system operating state becomes more strongly perturbed. The coverage of abnormal signals in multi-source observations expands. Normal behavior patterns are partially obscured. Under these conditions, unified modeling can still maintain responsiveness to anomalies, but the difficulty of discrimination increases. This reflects that model sensitivity to changes in anomaly density mainly stems from the stability of the state evolution assumption. The findings indicate that, in complex operational environments, the anomaly proportion affects not only data distributions but also the separability of system state modeling. This further highlights the necessity of improving anomaly identification robustness through multi-source fusion and temporal constraints.

4. Conclusion

This work addresses practical challenges in intelligent operations, including growing system scale, highly dynamic environments, and increasingly complex observations. It proposes a unified anomaly modeling approach for multi-source system observations. Starting from the overall system operating state, heterogeneous data such as logs, metrics, and call traces are mapped into a shared representation space. State evolution and consistency deviation are then used to characterize anomalies. This approach enables system-level understanding of how anomalies emerge and change, rather than relying on isolated signal-based judgments. It helps reduce bias caused by fragmented observations and provides more stable and consistent support for anomaly identification in complex distributed systems.

From a methodological perspective, this study emphasizes that anomaly detection should serve system-level cognition rather than focus only on pointwise recognition performance. By introducing unified representations and temporal constraints, anomalies are naturally modeled as deviations from normal operating trajectories. Relationships among multiple source observations are explicitly reflected in the state space. This improves the consistency of anomaly decisions and provides a structured view for understanding anomaly propagation and impact within the system. In operational environments with complex dependencies, the method reduces misjudgments caused by local fluctuations and offers more reliable evidence for operational decision-making.

From an application perspective, the proposed unified anomaly modeling framework demonstrates strong generality and extensibility. Its core principles do not depend on a specific system implementation or fixed observation format. It can be readily adapted to cloud native systems and distributed platforms of different scales and architectures. In practical operations, the framework can serve as a foundational component for anomaly analysis and alarm management. It can be integrated with automated scheduling, resource management, and failure recovery processes. This integration enhances system stability and self-healing capability. As reliability and continuity requirements continue to rise, system-level anomaly modeling is expected to play a more critical role in intelligent operations platforms.

Looking ahead, unified anomaly modeling offers substantial room for further development as observation techniques expand and runtime environments continue to evolve. More refined state representations can be explored to better capture long-term evolution patterns and latent risks in complex systems. The framework can also be more deeply integrated with higher-level operational decision mechanisms. This integration can support a shift from reactive handling to proactive warning and intelligent control. By continuously extending its applicability in real operations scenarios, unified anomaly modeling can provide a stronger foundation for building reliable and adaptive intelligent operations systems.

References

[1] H. Chen, P. Chen, and G. Yu, "A framework of virtual war room and matrix sketch-based streaming anomaly detection for microservice systems," IEEE Access, vol. 8, pp. 43413-43426, 2020.

[2] W. Zheng, H. Lu, Y. Zhou, et al., "iFeedback: Exploiting user feedback for real-time issue detection in large-scale online service systems," in Proc. 2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE), 2019, pp. 352-363.

[3] X. Zhang, F. Meng, P. Chen, et al., "TaskInsight: A fine-grained performance anomaly detection and problem locating system," in Proc. 2016 IEEE 9th International Conference on Cloud Computing (CLOUD), 2016, pp. 917-920.

[4] M. Weiss, J. Stumpfle, F. Dettinger, et al., "Simulating cloud environments of connected vehicles for anomaly detection," arXiv preprint arXiv:2404.11740, 2024.

[5] G. Baye, F. Hussain, A. Oracevic, et al., "API security in large enterprises: Leveraging machine learning for anomaly detection," in Proc. 2021 International Symposium on Networks, Computers and Communications (ISNCC), 2021, pp. 1-6.

[6] G. Yu, Z. Huang, and P. Chen, "TraceRank: Abnormal service localization with disaggregated end-to-end tracing data in cloud native systems," Journal of Software: Evolution and Process, vol. 35, no. 10, Art. no. e2413, 2023.

[7] L. L. Corbisier, "Adding Kalman Filter into ESI: An anomaly detection approach for middleware metadata," Ph.D. dissertation, Universidade NOVA de Lisboa, Portugal, 2020.

[8] Y. Yang, "Memory-enhanced Transformer for backend microservice anomaly detection," 2024.

[9] Y. Yang and A. Xu, "Log semantic graph construction and system event anomaly detection via convolutional neural networks," 2024.

[10] X. Zhang and A. He, "Lightweight anomaly detection for edge intelligence through compressed representation learning," 2025.

[11] Y. Zhang and E. Cheng, "Anomaly detection in e-commerce multi-table ETL processes through Mamba-LSTM collaborative modeling," 2025.

[12] R. Wei and D. Qiao, "Task scheduling research for cloud computing infrastructures: A reinforcement learning method integrating resource

occupancy prediction, dynamic node state encoding, and placement profit maximization," 2025.

[13] Z. Huang, J. Luo, and C. Chen, "Learning-driven collaborative scheduling for multi-tenant distributed inference services," 2025.

[14] L. D. Aprillia, H. Fakhrurroja, N. Ahmad, et al., "SIAS: Backend development for predictive satellite anomaly analysis," in Proc. International Conference on Data Analytics & Management. Cham, Switzerland: Springer Nature Switzerland, 2025, pp. 132-151.

[15] J. Li, S. Ying, T. Li, et al., "TraceDAE: Trace-based anomaly detection in micro-service systems via dual autoencoder," IEEE Transactions on Network and Service Management, 2025.

[16] O. Berlin, A. Held, M. Matousek, et al., "POSTER: Anomaly-based misbehaviour detection in connected car backends," in Proc. 2016 IEEE Vehicular Networking Conference (VNC), 2016, pp. 1-2.

[17] M. Ma, J. Xu, Y. Wang, et al., "Automap: Diagnose your microservice-based web applications automatically," in Proc. The Web Conference 2020, 2020, pp. 246-258.