Adaptive Anomaly Detection in Microservice Systems via Meta-Learning

Xiao Yang , Sijia Li , Ke Wu , Zhijun Wang , Yuqi Tang , Yueting Li *

Posted Date: 2 March 2026

doi: 10.20944/preprints202603.0017.v1

Keywords: microservice architecture; anomaly detection; meta-learning; intelligent operation and maintenance

 doi:10.20944/preprints202603.0017.v1

Article

Adaptive Anomaly Detection in Microservice Systems via Meta-Learning

Xiao Yang [1], Sijia Li [2], Ke Wu [3], Zhijun Wang [4], Yuqi Tang [5] and Yueting Li [6,*]

[1] Santa Clara University, Santa Clara, USA
[2] University of Michigan, Ann Arbor, USA
[3] Northeastern University, Boston, USA
[4] Rice University, Houston, USA
[5] New York University, New York, USA
[6] Purdue University, West Lafayette, USA
* Correspondence: yuetingli57@gmail.com

Abstract

This study addresses the highly dynamic runtime environment of microservice systems, the complex inter-service dependencies, and the hidden nature of anomalous behaviors. It proposes an anomaly detection method that integrates a meta learning mechanism. Based on multi-source monitoring data, the microservice execution process is modeled as a continuously evolving state sequence. A unified representation learning strategy is used to capture system evolution under normal conditions. The degree of state deviation is then adopted as the basis for anomaly discrimination. During modeling, different services or operating scenarios are treated as independent tasks. A meta learning framework is introduced to learn model initializations with strong transferability. This allows the model to adapt rapidly to new service instances and runtime environments under limited observations. It mitigates the impact of anomaly data scarcity and distribution shift. Compared with traditional methods that rely on fixed rules or single-scenario training, the proposed approach emphasizes shared runtime mechanism features. It maintains stable discrimination under noise interference and workload fluctuations. Comparative analysis under a unified data setting shows superior overall accuracy and more consistent anomaly discrimination compared with several representative methods. These results demonstrate strong robustness and generalization. The findings indicate that introducing meta learning into microservice anomaly detection improves adaptability and stability in complex cloud native environments. It provides an effective modeling strategy for anomaly identification in intelligent operations scenarios.

Keywords: microservice architecture; anomaly detection; meta-learning; intelligent operation and maintenance

CCS CONCEPTS: Computing methodologies–Machine learning–Machine learning approaches

1. Introduction

With the rapid advancement of cloud computing, containerization, and service mesh technologies, modern information systems are gradually evolving from monolithic architectures into highly decoupled and elastically scalable microservice systems. By decomposing complex business logic into multiple autonomous services, microservices enable rapid iteration, on-demand deployment, and high scalability. They have become a critical foundation for internet platforms, financial systems, and mission-critical infrastructures. At the same time, this architectural paradigm significantly increases operational complexity. A large number of service instances execute concurrently in distributed environments. Services interact frequently over networks. Their runtime

 doi:10.20944/preprints202603.0017.v1

behavior is affected by workload fluctuations, resource contention, version evolution, and scheduling strategies[1]. As a result, system behavior exhibits strong dynamics and nonlinear characteristics. Under these conditions, timely and accurate identification of abnormal behaviors has become a central challenge for ensuring system stability and business continuity[2].

In practical deployments, anomalies in microservice systems rarely manifest as explicit single-point failures. Instead, they often appear as performance degradation, amplified latency along invocation chains, imbalanced resource utilization, or breakdowns in service coordination. Such anomalies are characterized by complex triggering conditions, diverse manifestations, and continuous evolution over time. A single metric or a local observation is insufficient to capture their underlying nature. In addition, services differ substantially in functionality, workload patterns, and execution environments[3]. This heterogeneity leads to service-specific anomaly patterns. Traditional monitoring approaches based on fixed thresholds or manually defined rules struggle to cope with these characteristics. They frequently generate false alarms or miss critical issues. This increases the cognitive burden on operators and weakens early risk awareness. Therefore, developing intelligent anomaly detection methods that align with the complex runtime characteristics of microservices is of significant practical importance[4].

Recent years have witnessed growing interest in data-driven anomaly detection methods within system operations. These approaches model historical monitoring data and logs to enable automated identification of abnormal behaviors. However, they face notable limitations in real microservice environments. On the one hand, anomalies are inherently rare and highly diverse. This makes it difficult to obtain sufficient labeled samples across services and scenarios. On the other hand, microservice systems evolve continuously. Service instances are short-lived. Configurations and workloads change frequently. As a result, models are vulnerable to distribution shifts and often suffer from limited generalization. Under these conditions, anomaly detection strategies that rely solely on static training or global models struggle to maintain stable and reliable performance when confronted with new services or unseen operating conditions.Meta learning offers an alternative perspective for addressing these challenges. Its core objective is to learn how to adapt rapidly to new tasks by leveraging experience from multiple related tasks, rather than optimizing for a single fixed task. Introducing meta learning into microservice anomaly detection can alleviate the impact of data scarcity and distribution changes. It allows models to form effective decision capabilities for new services or new runtime scenarios with limited observations. By extracting shared anomaly representations and adaptation mechanisms across services and environments, such models reduce dependence on large-scale labeled data. They also maintain robustness under structural or workload variations. This emphasis on rapid adaptation aligns well with the highly dynamic and continuously evolving nature of microservice systems.

From an application perspective, research on microservice anomaly detection integrated with meta learning plays an important role in advancing intelligent operations and autonomous system management. Improved detection timeliness and accuracy help reduce the risk of business disruptions caused by accumulated or propagated anomalies. This provides stronger operational guarantees for critical systems. At the same time, reducing reliance on manual expertise and static rules mitigates the management burden associated with large-scale systems. It supports the transition toward automated and intelligent operational workflows. More broadly, this research direction contributes to a deeper understanding of the operational mechanisms of complex distributed systems. It provides theoretical foundations and methodological support for building next generation cloud native systems with adaptive and continual learning capabilities. This has positive implications for the long-term stability of cloud infrastructures and the delivery of high-quality services.

2. Methodology Foundation

As The proposed meta-learning-based anomaly detection framework is grounded in advances in microservice anomaly modeling, structured representation learning, continual adaptation, causal robustness, and adaptive control. These methodological foundations collectively inform the design

 doi:10.20944/preprints202603.0017.v1

of unified state evolution modeling and transferable meta-initialization for rapid adaptation under dynamic cloud-native environments.

Early studies on anomaly detection in microservice systems emphasize the importance of incorporating runtime feedback loops to continuously refine detection performance [5]. Group-wise trace anomaly detection further demonstrates that modeling execution traces at aggregated structural levels improves detection accuracy compared with point-wise analysis [6]. Similarly, OpenTracing-based anomaly detection approaches highlight the necessity of leveraging distributed tracing structures to capture inter-service execution dependencies [7]. These works motivate modeling the microservice execution process as a continuously evolving state sequence rather than isolated metric points. Beyond surface-level monitoring, relational and causal modeling strategies show that embedding structured dependencies improves interpretability and robustness. AI-based causal reasoning over structured representations illustrates how latent dependency modeling supports more reliable inference under complex environments [8]. Joint cross-modal representation learning demonstrates that heterogeneous signals can be embedded into a unified latent space while preserving structural coherence [9]. These principles directly inform the unified representation learning strategy adopted in this study to model normal runtime mechanisms across services. Robustness under distribution shifts is critical in highly dynamic systems. Semantics-aware denoising through adaptive sample reweighting demonstrates that emphasizing structurally consistent samples enhances model stability under noisy or shifting distributions [10]. Semantic alignment with output constraints further shows that explicit alignment mechanisms reduce unreliable predictions in complex decision scenarios [11]. These insights align with the use of deviation-based anomaly discrimination in a structured latent space, where consistent runtime mechanisms are separated from transient perturbations. Meta-learning provides the core mechanism for rapid adaptation across heterogeneous services. Autonomous learning frameworks based on self-driven exploration highlight the importance of learning transferable structural knowledge that generalizes across tasks [12]. Continual anomaly detection under non-stationary time-series distributions demonstrates that dynamic distribution monitoring and adaptive updating significantly improve long-term stability [13]. Causal representation learning further reinforces that disentangling stable shared mechanisms from task-specific variations enhances robustness and interpretability [14]. These methods collectively support the treatment of services or runtime scenarios as independent but related tasks within a meta-learning framework.

Control-theoretic perspectives offer additional theoretical grounding. Adaptive robust control techniques show how systems can maintain stability under parameter uncertainty and environmental perturbations [15]. Analogously, meta-learned initialization provides a stable starting point that rapidly adapts to new service conditions while maintaining discrimination boundaries. Conditional generative modeling with structured control further illustrates how controlled latent transitions preserve behavioral consistency during adaptation [16]. Relational modeling approaches using graph neural networks highlight the importance of multi-hop dependency modeling in complex systems [17]. Structural generalization mechanisms further demonstrate that embedding structural priors enhances adaptability when system topology evolves [18]. Adaptive fusion strategies for heterogeneous contextual signals illustrate dynamic integration across multiple information sources [19], supporting multi-source monitoring data fusion. Controllable abstraction mechanisms emphasize adjustable representation granularity [20], which aligns with modeling state evolution at appropriate abstraction levels. Finally, uncertainty-aware modeling introduces calibrated confidence estimation into prediction pipelines [17], reinforcing stable anomaly discrimination under noise and workload fluctuation.

By integrating structured state sequence modeling with transferable meta-initialization and adaptive robustness principles, the proposed framework enables rapid adaptation to new microservice instances and runtime scenarios while maintaining consistent anomaly discrimination in non-stationary cloud-native environments.

 doi:10.20944/preprints202603.0017.v1

3. Method

This paper addresses the challenges of highly dynamic operational states, significant service differences, and scarce anomaly samples in microservice systems by constructing an anomaly detection framework that integrates a meta-learning mechanism. First, the multi-dimensional monitoring metrics of the microservice system within a continuous time window are represented as a state sequence to characterize the system's temporal evolution. We build upon the structure-aware and semantically-enhanced graph modeling framework proposed by Lyu et al [22]. Their method fundamentally applies structure-aware graph construction together with semantic enhancement to strengthen anomaly pattern recognition in complex scheduling systems. We adopt this principle to organize multi-dimensional monitoring metrics into relational graph representations within each time window. Instead of modeling CPU usage, memory consumption, latency, and request rates independently, we incorporate dependency-aware encoding to capture their structural coupling. By leveraging semantic relationships among metrics, we extend their graph reasoning mechanism to dynamic microservice runtime modeling, enabling more stable representation under workload fluctuations and noise interference. The framework further builds upon the predictive scheduling and adaptive resource control principles introduced by Chen in FlashServe [23]. That work fundamentally applies predictive autoscaling and tiered resource management to dynamically adapt system behavior under cost performance trade-offs. We adopt its predictive adaptation philosophy to model state transitions across time windows and incorporate resource-sensitive dynamics as contextual signals in anomaly discrimination. By leveraging adaptive control awareness, we extend predictive scheduling concepts into meta-learning based runtime modeling, allowing the model to distinguish structural deviation from benign workload variation across heterogeneous services. Let the observed state of a certain service within time window t be represented as follows:

$$x_t = [x_t^1, \ldots, x_t^n] \tag{1}$$

Here, d represents the monitoring indicator dimension. By modeling the continuous state sequence $\{x_t^n\}_{t=1}^{T}$, the method can characterize the dynamic mode of system operation in a unified representation space, providing a basic representation for subsequent anomaly detection and cross-scenario adaptation. This paper also presents the overall model architecture, as shown in Figure 1.

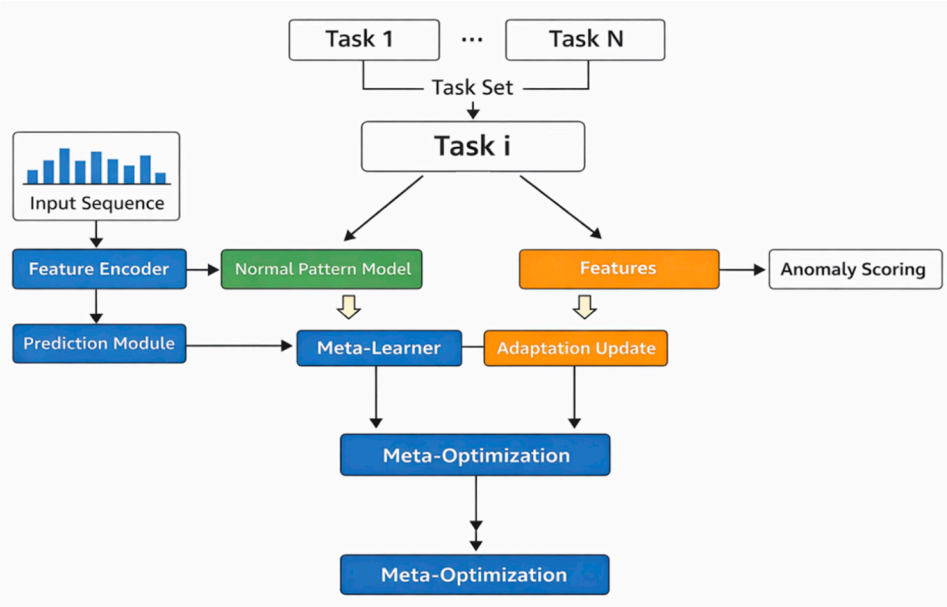

Figure 1. Overall model architecture.

Building upon this, a parameterized model is introduced to map the system state, learning the latent representation under normal operating conditions. Let the model parameters be θ, and the state mapping function be defined as:

 doi:10.20944/preprints202603.0017.v1

$$h_t = f_\theta(x_t) \tag{2}$$

where h_t is a low-dimensional latent representation used to compress redundant information in the original monitoring data and highlight key operational features. To characterize the consistency of state evolution, a prediction function is introduced to estimate the state at the next time step, with the following form:

$$\hat{h}_{t+1} = g_\theta(h_t) \tag{3}$$

Furthermore, the prediction error reflects the degree of deviation between the system behavior and the normal pattern, thus providing a quantifiable basis for anomaly detection.

Considering the differences in functionality and load patterns among various microservices, this paper treats each type of service or operational scenario as an independent task and introduces a meta-learning mechanism on the task set. For the i-th task, its task-level objective function is defined as:

$$L_i(\theta) = \frac{1}{T_i}\sum_{t=1}^{T_i} ||h_{t+1} - \hat{h}_{t+1}||_2^2 \tag{4}$$

where T_i is the time step length for this task. By jointly optimizing across multiple tasks, the model learns an initialization with good transferability in the parameter space, thereby enabling it to quickly develop effective state modeling capabilities when facing new services or new operating conditions.

At the overall optimization level, the method uses a meta-objective function to uniformly update model parameters, thereby improving cross-task adaptability and stability. The meta-optimization objective is defined as:

$$\min_\theta \sum_{i=1}^{N} L_i(\theta - \alpha\nabla_\theta L_i(\theta)) \tag{5}$$

where N represents the number of tasks, and α is the task-level update step size. In the anomaly detection phase, an anomaly scoring function is constructed based on the prediction error:

$$s_t = ||h_t - \hat{h}_t||_2 \tag{6}$$

This method is used to measure the degree of deviation between the current state and the normal evolutionary pattern. Through the above mechanism, the method achieves modeling and rapid adaptation of the operational behavior of microservice systems within a unified framework, providing an adaptive modeling approach for anomaly detection in complex environments.

4. Implementation

Dataset

This study adopts GAIA, the Generic AIOps Atlas, as the experimental data source. GAIA is an open source dataset designed for intelligent operations scenarios. It covers multi-source observations generated by microservice systems during real or near-real executions. The dataset supports tasks such as anomaly detection, fault diagnosis, and root cause analysis. It is built around microservice architectures and captures both service dependencies and dynamic runtime behaviors. These properties align well with the cross-service and cross-scenario adaptation requirements of meta learning based microservice anomaly detection.

In terms of data representation, GAIA offers multimodal inputs that align with microservice observability practices. The data typically comprises metric time series, log events, and information pertaining to service invocation traces. This structure enables the identification of anomalies not only by isolated metric changes but also by service interactions and temporal evolution. Unlike datasets that contain only single-source time series, multi-source observations better capture the subtle nature of microservice anomalies and the amplification effects along call chains. They support detection paradigms centered on state representation and consistency deviation. This makes the dataset more

 doi:10.20944/preprints202603.0017.v1

akin to the actual mechanisms of anomaly emergence and propagation in intricate cloud-native environments.

To align with the meta learning framework, data can be organized by treating different services, business scenarios, or fault types as distinct task sources. Under unified feature preprocessing and time window segmentation strategies, these tasks form a task collection. This design enables learning stable patterns of normal behavior within tasks. It also facilitates learning transferable initializations and fast adaptation across tasks. Such an organization alleviates instability caused by anomaly scarcity and distribution shifts in microservice environments. The task-oriented construction based on this dataset directly supports extracting cross-scenario commonalities from multi-task experience. It further enables more efficient adaptive anomaly identification under new services or new runtime conditions.

5. Evaluation

5.1. Evaluate Metric

This paper employs multiple evaluation metrics to measure the discriminative power and overall stability of anomaly detection methods from different perspectives, avoiding bias caused by a single metric and providing a more comprehensive reflection of the model's performance in anomaly identification scenarios within microservice systems.

Accuracy measures the model's correctness in identifying all samples, reflecting the consistency between the predicted results and the actual state. This metric comprehensively considers the predictions of both normal and anomaly samples, providing a direct characterization of the model's overall recognition ability. However, it may be affected by the majority class in scenarios with an imbalanced class distribution. Its definition is:

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \tag{7}$$

where TP, TN, FP, and FN represent the number of samples that are true anomalies, true normal samples, false anomalies, and false normal samples, respectively.

Precision characterizes the proportion of samples that the model classifies as anomalous, indicating that they are indeed anomalous, and reflects the reliability of the anomaly alarm results. In microservice systems, higher precision means fewer instances of correctly identifying normal operating conditions as anomalous, thus reducing unnecessary alarm interference. Its calculation formula is:

$$Precision = \frac{TP}{TP+FP} \tag{8}$$

Recall measures the proportion of real-world anomalies successfully identified by the model, reflecting the model's ability to cover anomalous behavior. A higher recall means the model can more fully capture potential risks and reduce missed anomalies, which is crucial for ensuring stable system operation. It is defined as:

$$Recall = \frac{TP}{TP+FN} \tag{9}$$

The F1-Score combines precision and recall using a harmonic average method, achieving a balance between the two. It is suitable for scenarios where outlier samples are relatively scarce and false positives and false negatives are equally important. This metric provides a more comprehensive reflection of the model's overall discrimination quality in anomaly detection tasks, and its expression is:

$$F1 = \frac{2\times Precision\times Recall}{Precision+Recall} \tag{10}$$

 doi:10.20944/preprints202603.0017.v1

5.2. Experimental Results

This paper first presents the experimental results compared with other models, as shown in Table 1.

Table 1. Comparative experimental results.

Method	Acc	Precision	Recall	F1-Score
TraceGra [24]	0.7321	0.7214	0.7098	0.7155
ServiceAnomaly [25]	0.7546	0.7423	0.7351	0.7387
ReconRCA [26]	0.7684	0.7565	0.7482	0.7523
Iadcps [27]	0.7819	0.7701	0.7627	0.7663
EasyAD [28]	0.7957	0.7844	0.7736	0.7789
Ours	0.8342	0.8217	0.8153	0.8185

The overall comparison shows that the proposed method consistently achieves leading performance across all four core metrics, improving anomaly discrimination while balancing coverage and reliability in dynamic microservice environments. By modeling runtime evolution rather than relying solely on static indicators or local patterns, it more effectively captures structural differences between normal and abnormal states, resulting in clearer anomaly boundaries, reduced false alarms without increased missed detections, and stronger F1-Score performance that aligns with practical alert quality requirements. Compared with graph-only, reconstruction-based, or rule-based strategies that are vulnerable to scaling changes and workload drift, the integration of task-oriented modeling and fast adaptation enhances robustness and cross-scenario generalization, supporting the value of meta-learning in microservice anomaly detection. In addition, sensitivity analysis on the sliding step size shows that while excessively small steps amplify short-term noise and overly large steps weaken fine-grained anomaly capture, Figure 2 demonstrates stable performance within a reasonable range, confirming the method's robustness to this key preprocessing parameter.

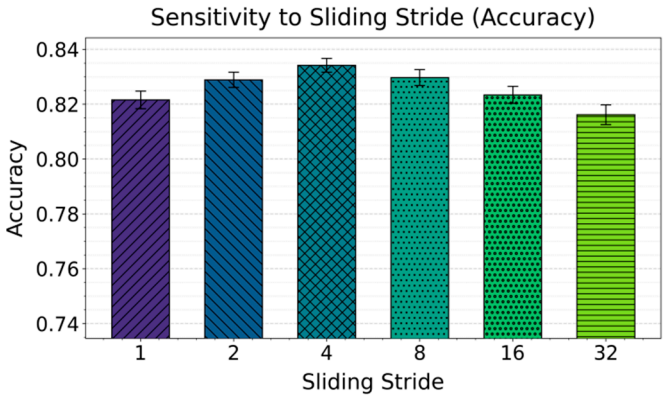

Figure 2. Sensitivity analysis of sliding step size to accuracy.

The trend shows that detection accuracy first increases and then declines as the sliding step grows. This indicates that this preprocessing parameter directly affects how the model represents the evolution of microservice runtime states. When the step increases gradually, redundant correlations between adjacent segments are moderately reduced. The model can focus more on stable patterns and abnormal deviations across time intervals. This improves the identification of hidden defect behaviors. When the step becomes too large, time segmentation becomes overly sparse. Critical anomaly cues may be skipped by window sampling. Information representation becomes incomplete. Accuracy, therefore, decreases. This behavior is consistent with the fact that anomaly signals in microservice systems often exhibit stage-wise accumulation and chain amplification. An appropriate step better aligns with this temporal structure.

 doi:10.20944/preprints202603.0017.v1

When the step is too small, high overlap between adjacent samples introduces strong sample correlation. The training data become statistically close to repeated sampling. This amplifies the influence of short-term noise and transient jitter on the decision boundary. In microservice environments, metric fluctuations are strongly affected by scheduling and resource contention. Local spikes do not necessarily indicate real anomalies. With overly dense steps, the model is more likely to treat transient fluctuations as part of normal patterns. The representation of anomaly concepts becomes blurred. In meta learning based frameworks, such highly correlated inputs also weaken the representation of inter-task differences. Fast adaptation then tends to memorize short-term details rather than learn transferable state evolution patterns. This reduces overall robustness.

When the step is too large, insufficient window coverage reduces the ability to capture fine-grained anomaly signals. Early cues that gradually amplify along call chains are especially likely to be missed by sparse sampling. Microservice anomalies usually do not erupt instantly. They evolve from minor deviations into global risks. Excessive steps cause observed sequences to become fragmented. It becomes difficult to establish consistent state transition relationships. For detection strategies centered on state prediction and consistency deviation, this directly degrades the modeling quality of continuous evolution patterns. Sensitivity to critical stages is reduced. Accuracy consequently declines.

The feature embedding dimension determines the model's ability to compress and represent the microservice's operational state, directly affecting the balance between preserving key information and suppressing noise. Too small a dimension may lead to insufficient representation of complex interactions and temporal dependencies, while too large a dimension may introduce redundancy and amplify sensitivity to environmental fluctuations. To verify the stability of the method under different representation capacities, different embedding dimensions were set, and the response trend of accuracy with changing dimensions was observed. The experimental results are shown in Figure 3.

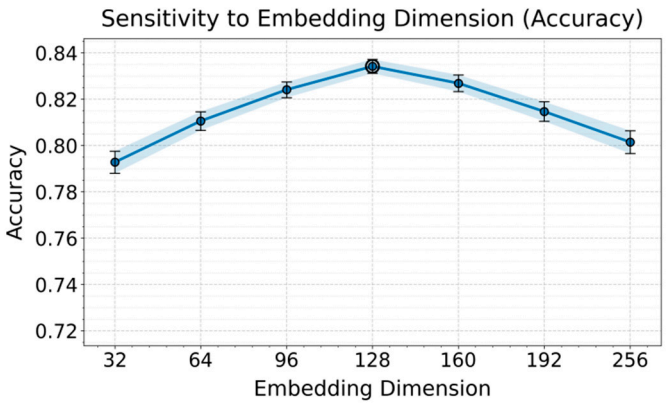

Figure 3. Sensitivity analysis of feature embedding dimension to accuracy.

The sensitivity curve shows an overall pattern in which performance improves as the embedding dimension increases and then declines at larger dimensions. This indicates that embedding dimensionality has a direct impact on representation quality in microservice anomaly detection. Within a certain range, higher-dimensional embeddings allow the model to accommodate richer state information. Key correlations and temporal dependencies from multi-source monitoring data can be encoded into a unified representation space. This enables a clearer distinction between normal evolution and abnormal deviation. This observation is consistent with the focus on understanding dynamic mechanisms and modeling state sequences in microservice systems. An appropriate representation capacity helps form clearer anomaly decision boundaries.

When the embedding dimension is too small, the capacity of the representation space is limited. Complex service interactions and workload variations are compressed into overly coarse representations. Fine-grained signals related to anomalies may be smoothed out. Anomalies in

 doi:10.20944/preprints202603.0017.v1

microservice systems often involve cross-service propagation and chain amplification. They require joint preservation of local fluctuations and global trends. With insufficient dimensionality, the model struggles to express such hierarchical structures. This reduces stability in identifying hidden defect behaviors. In meta learning based frameworks, small dimensions also weaken the capacity to carry transferable knowledge across tasks. Fast adaptation then relies more on superficial patterns rather than shared mechanism-level representations.

As the embedding dimension continues to increase to higher ranges, accuracy declines. This suggests that overly large representation spaces introduce redundancy and amplify sensitivity to noise and scenario differences. In microservice runtime data, short-term jitter, sampling noise, and scheduling disturbances are common. High-dimensional embeddings tend to encode these unstable factors together with meaningful signals. This reduces focus on core anomaly patterns. For meta learning, this increases uncertainty during task-specific updates. The adaptation process may drift toward incidental features. Cross-scenario generalization consistency is then affected.

The overall narrow and smoothly varying error bands indicate that the method has a certain degree of robustness to embedding dimensionality. Performance differences are mainly driven by the trade-off between representation capacity and noise suppression rather than random fluctuations. This behavior shows that the proposed method maintains stable detection behavior under different capacity settings. It also suggests the existence of a moderate dimensional range that better matches the modeling requirements of microservice temporal mechanisms. Overall, the sensitivity results support the main argument of the paper. After obtaining transferable state representations through task-oriented learning, a reasonable embedding capacity more effectively highlights factors relevant to anomaly discrimination. It reduces interference from redundancy and noise in complex dynamic environments.

Finally, this paper presents the impact of the weight decay coefficient on the experimental results, which are shown in Figure 4.

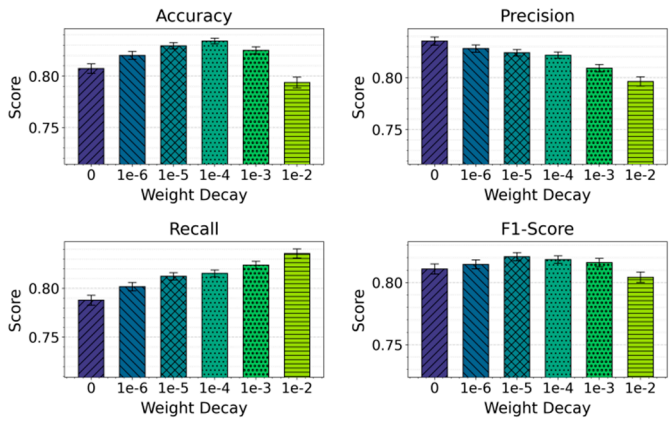

Figure 4. The impact of the weight decay coefficient on experimental results.

The overall trend indicates that the weight decay coefficient exerts a systematic influence on the model's decision behavior. Different evaluation metrics respond differently as the regularization strength changes. This shows that regularization does not merely perform simple parameter shrinkage. It directly participates in how the model selects and emphasizes microservice runtime patterns and anomaly characteristics. In anomaly detection for complex microservice systems, appropriate weight constraints help the model form more stable representations under high-dimensional monitoring features and noise perturbations. This improves overall decision consistency.

In regimes with weak regularization, the model has strong fitting capacity on the training data. It can capture fine-grained state variations. At the same time, it becomes more sensitive to short-term fluctuations and local noise. This effect is particularly evident in microservice environments. Resource scheduling, transient workload changes, and invocation jitter frequently introduce non-anomalous disturbances. Without sufficient constraints, the model may incorporate these unstable

 doi:10.20944/preprints202603.0017.v1

factors into anomaly decisions. This weakens its ability to abstract stable patterns across time and services.

As weight decay increases, model parameters are guided toward smoother solution spaces. Sensitivity to incidental fluctuations is reduced. Persistent and structural anomaly deviations become more prominent. This behavior aligns with the mechanism-level modeling objective emphasized in this work. The goal is to distinguish real defects from environmental noise by learning stable runtime evolution patterns. However, when regularization becomes too strong, excessive parameter constraints compress the model's expressive capacity. The model then struggles to capture complex interactions and temporal dependencies in microservice systems. Discrimination of certain anomaly patterns is therefore suppressed.

The fact that different metrics do not follow identical trends under changing weight decay further reflects that regularization affects the balance between false alarms and missed detections in different ways. Weight decay influences not only performance levels but also decision preferences. The model may shift between conservative and aggressive detection behaviors.

6. Conclusions

This work addresses the highly dynamic runtime environment of microservice systems, the hidden nature of anomalies, and the pronounced differences across services. It proposes an anomaly detection method that integrates meta learning to improve stability and generalization from the perspective of runtime mechanism modeling. By representing microservice execution as continuously evolving state sequences and learning transferable initialization and fast adaptation within a task-oriented framework, the method accounts for heterogeneity across services and operating scenarios in a unified manner. This approach overcomes the limitations of traditional methods that rely heavily on static patterns or single scenarios. It provides a more robust modeling paradigm for anomaly detection in complex cloud native environments. Comparative analysis under a unified data setting shows consistent advantages across multiple core evaluation metrics. The results indicate strong overall performance in detection accuracy, stability, and risk discrimination. These gains do not arise from emphasizing local features or specific rules. They stem from a holistic modeling of system runtime evolution. The model remains sensitive to abnormal deviations under noise interference and workload fluctuations. For hidden defects and progressive anomalies common in microservice systems, the method better distinguishes real risks from incidental variations. This provides more reliable decision support for subsequent system handling and risk control.

From an application perspective, the study offers direct implications for intelligent operations and autonomous system management. Reducing reliance on manual experience and static threshold configuration helps mitigate the rapid growth of operational complexity as system scale increases. It promotes a shift toward automated and intelligent anomaly detection. In addition, the fast adaptation capability makes the model well-suited for deployment in continuously evolving cloud platforms and microservice architectures. Stable performance can be maintained as the business expands and systems change. This property is critical for ensuring the continuous operation of key services and reducing the risk of fault propagation.

Looking ahead, anomaly detection frameworks that integrate meta learning can be further combined with microservice governance and decision mechanisms. Beyond detection, they can support higher-level analysis and scheduling capabilities. For example, closed-loop coordination can be formed with automated recovery, resource scheduling optimization, and risk assessment modules. This would enable stronger adaptability and self-regulation. As cloud computing and distributed systems continue to scale, the application potential of such methods will expand in large cloud platforms, critical infrastructures, and complex business systems. They will provide important support for building next generation cloud native systems with long-term stability and intelligent decision-making capabilities.

 doi:10.20944/preprints202603.0017.v1

References

1. V. Ramamoorthi, "Machine learning models for anomaly detection in microservices," Quarterly Journal of Emerging Technologies and Innovations, vol. 5, no. 1, pp. 41-56, 2020.
2. J. Nobre, E. J. S. Pires and A. Reis, "Anomaly detection in microservice-based systems," Applied Sciences, vol. 13, no. 13, p. 7891, 2023.
3. P. P. Naikade, "Automated anomaly detection and localization system for a microservices based cloud system," Ph.D. dissertation, The University of Western Ontario, Canada, 2020.
4. Q. Du, T. Xie and Y. He, "Anomaly detection and diagnosis for container-based microservices with performance monitoring," Proceedings of the International Conference on Algorithms and Architectures for Parallel Processing, pp. 560-572, 2018.
5. A. Hrusto, E. Engström and P. Runeson, "Optimization of anomaly detection in a microservice system through continuous feedback from development," in Proceedings of the 10th IEEE/ACM International Workshop on Software Engineering for Systems-of-Systems and Software Ecosystems, pp. 13-20, 2022.
6. Z. Xie, C. Pei, W. Li et al., "From point-wise to group-wise: A fast and accurate microservice trace anomaly detection approach," in Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, pp. 1739-1749, 2023.
7. M. Khanahmadi, A. Shameli-Sendi, M. Jabbarifar et al., "Detection of microservice-based software anomalies based on OpenTracing in cloud," Software: Practice and Experience, vol. 53, no. 8, pp. 1681-1699, 2023.
8. R. Ying, Q. Liu, Y. Wang and Y. Xiao, "AI-based causal reasoning over knowledge graphs for data-driven and intervention-oriented enterprise performance analysis," 2025.
9. X. Zhang, Q. Wang and X. Wang, "Joint cross-modal representation learning of ECG waveforms and clinical reports for diagnostic classification," Transactions on Computational and Scientific Methods, vol. 6, no. 2, 2026.
10. X. Yang, Y. Wang, Y. Li and S. Sun, "Semantics-aware denoising: A PLM-guided sample reweighting strategy for robust recommendation," arXiv preprint arXiv:2602.15359, 2026.
11. J. Yang, S. Sun, Y. Wang, Y. Wang, X. Yang and C. Zhang, "Semantic alignment and output constrained generation for reliable LLM-based classification," 2026.
12. F. Wang, Y. Ma, T. Guan, Y. Wang and J. Chen, "Autonomous learning through self-driven exploration and knowledge structuring for open-world intelligent agents," 2026.
13. Y. Ou, S. Huang, F. Wang, K. Zhou and Y. Shu, "Adaptive anomaly detection for non-stationary time-series: A continual learning framework with dynamic distribution monitoring," 2025.
14. J. Li, Q. Gan, R. Wu, C. Chen, R. Fang and J. Lai, "Causal representation learning for robust and interpretable audit risk identification in financial systems," 2025.
15. X. T. Li, X. P. Zhang, D. P. Mao and J. H. Sun, "Adaptive robust control over high-performance VCM-FSM," Optics and Precision Engineering, vol. 25, pp. 2428-2436, 2017.
16. R. Liu, L. Yang, R. Zhang and S. Wang, "Generative modeling of human-computer interfaces with diffusion processes and conditional control," arXiv preprint arXiv:2601.06823, 2026.
17. K. Cao, Y. Zhao, H. Chen, X. Liang, Y. Zheng and S. Huang, "Multi-hop relational modeling for credit fraud detection via graph neural networks," 2025.
18. C. Hu, Z. Cheng, D. Wu, Y. Wang, F. Liu and Z. Qiu, "Structural generalization for microservice routing using graph neural networks," in Proceedings of the 2025 3rd International Conference on Artificial Intelligence and Automation Control (AIAC), pp. 278-282, 2025.
19. X. Hu, Y. Kang, G. Yao, T. Kang, M. Wang and H. Liu, "Dynamic prompt fusion for multi-task and crossdomain adaptation in LLMs," in Proceedings of the 2025 10th International Conference on Computer and Information Processing Technology (ISCIPT), pp. 483-487, 2025.
20. X. Song, Y. Liu, Y. Luan, J. Guo and X. Guo, "Controllable abstraction in summary generation for large language models via prompt engineering," arXiv preprint arXiv:2510.15436, 2025.
21. S. Pan and D. Wu, "Trustworthy summarization via uncertainty quantification and risk awareness in large language models," in Proceedings of the 2025 6th International Conference on Computer Vision and Data Mining (ICCVDM), pp. 523-527, 2025.

 doi:10.20944/preprints202603.0017.v1

22. N. Lyu, J. Jiang, L. Chang, C. Shao, F. Chen and C. Zhang, "Improving pattern recognition of scheduling anomalies through structure-aware and semantically-enhanced graphs," arXiv preprint arXiv:2512.18673, 2025.
23. B. Chen, "FlashServe: Cost-efficient serverless inference scheduling for large language models via tiered memory management and predictive autoscaling," 2025.
24. J. Chen, F. Liu, J. Jiang et al., "TraceGra: A trace-based anomaly detection for microservice using graph deep learning," Computer Communications, vol. 204, pp. 109-117, 2023.
25. M. Panahandeh, A. Hamou-Lhadj, M. Hamdaqa et al., "ServiceAnomaly: An anomaly detection approach in microservices using distributed traces and profiling metrics," Journal of Systems and Software, vol. 209, p. 111917, 2024.
26. Z. Zhang, J. Wang, B. Li et al., "ReconRCA: Root cause analysis in microservices with incomplete metrics," Proceedings of the 2025 IEEE International Conference on Web Services (ICWS), pp. 116-126, 2025.
27. J. Tian, M. Li, L. Chen et al., "iADCPS: Time series anomaly detection for evolving cyber-physical systems via incremental meta-learning," arXiv preprint arXiv:2504.04374, 2025.
28. Q. Liu, S. Lee and J. Paparrizos, "EasyAD: A demonstration of automated solutions for time-series anomaly detection," Proceedings of the VLDB Endowment, vol. 18, no. 12, pp. 5431-5434, 2025.