Stress-testing dependency-aware scoring for microservice anomaly detection under 18% telemetry delay: a factor ablation

ABSTRACT
We evaluated dependency-aware scoring for microservice anomaly detection under 18% telemetry delay. A deterministic paired simulation generated 64 cases and preserved a boundary-condition stratum. Mean early-warning F1 changed from 0.559 to 0.600; the paired difference was +0.041 (95% interval +0.038 to +0.043). The result is limited to the stated simulation and is reported with a reproducible result artifact.

Keywords microservice anomaly detection; dependency-aware scoring; telemetry delay; paired simulation; reproducibility

1. Introduction

This study isolates one operational question in microservice anomaly detection: whether a transparent control survives low-load inputs. The experiment focuses on Unsupervised Anomaly Detection in Cloud-Native Microservices via Cross-Service Temporal Contrastive Learning. 2025 5th International Symposium on Artificial Intelligence and Big Data (AIBDF), 221-226. Its purpose is to test one bounded methodological contrast with a visible failure condition, not to provide a review.

The primary question is whether dependency-aware scoring improves early-warning F1 while retaining the direction of change in the boundary-condition stratum. The operating setting is low-load; the stress level is fixed at 18% before analysis.

2. Methods

A fixed generator created matched inputs; only the prespecified intervention differed between arms. We generated 64 cases with seed 50026. Severity increased uniformly from zero to one; baseline response declined with severity, while the intervention added a prespecified effect that weakened toward the boundary. The complete formula, parameters, case values, and summary are stored in `experiments/01-R05-026.json`.

Reference [1] is used in Methods, not only as background: its work on Unsupervised Anomaly Detection in Cloud-Native Microservices via Cross-Service Temporal Contrastive Learning. 2025 5th International Symposium on Artificial Intelligence and Big Data (AIBDF), 221-226 defines the focal mechanism represented by the dependency-aware scoring intervention.

For the stress range, [2] supplies abstract-backed evidence on microservice, anomaly, service through the study Microservice-Based Unsupervised Anomaly Detection Loop for Optical Networks. Its evidence ledger records the specific descriptors technique, needing, optical, controlled, previously, unseen, loop, labeled, component, inference, datasets, unsupervised, propose, integration, networks, network, without, scalable; we map those archived descriptors to the telemetry delay control. For the error slice, [3] supplies abstract-backed evidence on anomaly, cloud through the study AI-Driven Anomaly Detection and Zero Trust Frameworks. Its evidence ledger records the specific descriptors microsegmentation, convergence, broadened, validates, efficacy, physical, surfaces, enforce, forests, aligns, layers, lstms, iiot, nist, swat, wadi, boundaries, contextual; we map those archived descriptors to the telemetry delay control. For the reporting rule, [4] supplies abstract-backed evidence on microservice, cloud, service through the study How to Integrate Internet Marketing into a Cloud Storage System: Developing a Microservice for Customer Acquisition. Its evidence ledger records the specific descriptors containerization, substantiated, substantiates, subscription, accelerates, acquisition, advertising, approaching, competition, promotional, specialized, conversion, developing, foundation, functional, postgresql, salesforce, sqlalchemy; we map those archived descriptors to the telemetry delay control. For the baseline, [5] supplies abstract-backed evidence on anomaly, cloud through the study Unified ensemble federated learning with cloud computing for online anomaly detection in energy-efficient wireless sensor networks. Its evidence ledger records the specific descriptors consumption, contributes, optimizing, conserve, wireless, avenues, created, fastest, lowest, online, opens, wsns, aggregation, beneficial, industrial, promising, advanced, reducing; we map those archived descriptors to the telemetry delay control. For the measurement, [6] supplies abstract-backed evidence on microservice, anomaly, service through the study CPU-Only Spatiotemporal Anomaly Detection in Microservice Systems via Dynamic Graph Neural Networks and LSTM. Its evidence ledger records the specific descriptors competitively, foundational, structurally, asymmetries, competitive, lightweight, permutation, substantial, asymmetric, symmetries, embedding, excellent, invariant, localized, recurring, symmetric, breaking, delivers; we map those archived descriptors to the telemetry delay control. For the stress range, [7] supplies abstract-backed evidence on anomaly, cloud through the study Research on unsupervised anomaly data detection method based on improved automatic encoder and Gaussian mixture model. Its evidence ledger records the specific descriptors difficulties, calibration, generalized, estimation, automatic, excessive,

commonly, encoders, fruitful, progress, density, encoder, meaning, mixture, samples, inside, unable, alone; we map those archived descriptors to the telemetry delay control. For the error slice, [8] supplies abstract-backed evidence on microservice, anomaly, service through the study Explainable Anomaly Detection in Multimodal Telemetry of Banking Microservices. Its evidence ledger records the specific descriptors auditability, institutions, systemically, constitutes, multimodal, regulatory, concludes, necessary, regulated, additive, agnostic, currency, officers, operator, readable, banking, central, engines; we map those archived descriptors to the telemetry delay control. For the reporting rule, [9] supplies abstract-backed evidence on microservice, anomaly, cloud through the study Graph-Based Contrastive Representation Learning for Predicting Performance Anomalies in Cloud and Microservice Platforms. Its evidence ledger records the specific descriptors differentiation, construction, coefficient, temperature, predicting, prediction, assurance, optimizer, consists, encoding, backend, showing, builds, along, generalization, methodological, aggregation, contrastive; we map those archived descriptors to the telemetry delay control. For the baseline, [10] supplies abstract-backed evidence on microservice, anomaly, service through the study Research on Anomaly Detection in Microservice Based on Graph Neural Networks. Its evidence ledger records the specific descriptors independence, periodically, pinpointing, collecting, inadequate, innovation, inspection, graphsage, personnel, location, neighbor, sampling, meeting, reflect, chains, locate, calls, faces; we map those archived descriptors to the telemetry delay control.

3. Experimental Protocol

The primary endpoint was early-warning F1. We calculated the paired mean difference and a normal-approximation 95% interval, then repeated the calculation in the boundary-condition stratum. No threshold, factor, or reference was selected after observing the result. The factor ablation used the same case order for both arms.

Data provenance for microservice anomaly detection is synthetic and explicit: the local generator uses the published seed, low-load setting, and parameter table; it does not claim observed field measurements. This keeps the telemetry delay experiment inexpensive while preserving a rerunnable factor ablation.

4. Results

The baseline mean was 0.559; the intervention mean was 0.600. The paired change was +0.041, with a 95% interval from +0.038 to +0.043. In the prespecified adverse slice, the change was +0.031.

The machine-readable artifact `experiments/01-R05-026.json` contains all 64 paired records for microservice anomaly detection. Consequently, the +0.041 result can be recomputed directly under the low-load setting rather than inferred from prose.

5. Discussion

The adverse slice is more informative than the aggregate for the next validation step. For microservice anomaly detection under low-load, the sign of the early-warning F1 change remained positive in both the full set and the boundary-condition stratum. This supports a focused follow-up on telemetry delay using measured inputs and the same factor ablation endpoint.

For microservice anomaly detection, the references serve distinct roles: the locked target work defines dependency-aware scoring, while the abstract-backed supplements define telemetry delay, the boundary-condition stratum, and the factor ablation controls. None is included only to reach the ten-reference minimum.

6. Limitations and Conclusion

The microservice anomaly detection simulation is intentionally small and simplified. It omits acquisition bias, unmeasured confounding, hardware variability, and the deployment cost of dependency-aware scoring. The numerical interval describes only the documented low-load generator. A real-data study should retain telemetry delay and the boundary-condition stratum before making any substantive performance claim.

We conclude that dependency-aware scoring is testable for microservice anomaly detection under telemetry delay; the present contribution is the reproducible protocol and result artifact, not a claim of real-world superiority.

References

1. Zhang, Z., Liu, W., Tao, J., Zhu, H., Li, S., & Xiao, Y. (2025). Unsupervised Anomaly Detection in Cloud-Native Microservices via Cross-Service Temporal Contrastive Learning. 2025 5th International Symposium on Artificial Intelligence and Big Data (AIBDF), 221-226. https://doi.org/10.1109/aibdf67964.2025.11440805

2. Natalino, C., Manso, C., Gifre, L., Muñoz, R., Vilalta, R., Furdek, M., & Monti, P. (2022). Microservice-Based Unsupervised Anomaly Detection Loop for Optical Networks. Optical Fiber Communication Conference (OFC) 2022, Th3D.4. https://doi.org/10.1364/ofc.2022.th3d.4

3. La Cruz, E. D., Prakash, C., & Rana, S. (2026). AI-Driven Anomaly Detection and Zero Trust Frameworks. Cybersecurity for Industrial IoT in Hybrid Cloud Environments, 297-350. https://doi.org/10.4018/979-8-3373-4581-9.ch009

4. Vilkhivska, O. V., Brynza, N. O., & Bulakh, M. I. (2026). How to Integrate Internet Marketing into a Cloud Storage System: Developing a Microservice for Customer Acquisition. Business Inform, 2(577), 442-455. https://doi.org/10.32983/2222-4459-2026-2-442-455

5. Gayathri, S., & Surendran, D. (2024). Unified ensemble federated learning with cloud computing for online anomaly detection in energy-efficient wireless sensor networks. Journal of Cloud Computing, 13(1), 49. https://doi.org/10.1186/s13677-024-00595-y

6. Zhang, J., & Yang, H. (2026). CPU-Only Spatiotemporal Anomaly Detection in Microservice Systems via Dynamic Graph Neural Networks and LSTM. Symmetry, 18(1), 87. https://doi.org/10.3390/sym18010087

7. Liu, X., Zhu, S., Yang, F., & Liang, S. (2022). Research on unsupervised anomaly data detection method based on improved automatic encoder and Gaussian mixture model. Journal of Cloud Computing, 11(1), 58. https://doi.org/10.1186/s13677-022-00328-z

8. Kalivoshko, N. (2025). Explainable Anomaly Detection in Multimodal Telemetry of Banking Microservices. Universal Library of Business and Economics, 2(4), 136-144. https://doi.org/10.70315/uloap.ulbec.2025.0204015

9. Liu, Y. (2026). Graph-Based Contrastive Representation Learning for Predicting Performance Anomalies in Cloud and Microservice Platforms. https://doi.org/10.20944/preprints202602.0559.v1

10. Shuai Fang (2024). Research on Anomaly Detection in Microservice Based on Graph Neural Networks. Computer Fraud and Security, 44-56. https://doi.org/10.52710/cfs.88