Tuning dependency-aware scoring for microservice anomaly detection under 37% telemetry delay: a threshold audit

ABSTRACT
We evaluated dependency-aware scoring for microservice anomaly detection under 37% telemetry delay. A deterministic paired simulation generated 48 cases and preserved a low-signal stratum. Mean early-warning F1 changed from 0.516 to 0.570; the paired difference was +0.054 (95% interval +0.051 to +0.057). 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

Reproducible stress tests can expose weaknesses in microservice anomaly detection before expensive field collection. 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 low-signal stratum. The operating setting is shifted-distribution; the stress level is fixed at 37% before analysis.

2. Methods

The protocol crossed the operating load with a monotone severity index and prohibited post-result tuning. We generated 48 cases with seed 50084. 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-R05-084.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 anomaly, cloud, service through the study Anomaly Detection on AWS cloud by using Supervised Machine Learning. Its evidence ledger records the specific descriptors affordability, functionality, industries, popularity, responses, examines, manually, variety, alerts, amazon, chosen, gained, prompt, spot, automating, cloudtrail, predefined, widespread; we map those archived descriptors to the telemetry delay control. For the error slice, [3] supplies abstract-backed evidence on microservice, anomaly, service through the study Anomaly detection in microservice environments using distributed tracing data analysis and NLP. Its evidence ledger records the specific descriptors visualization, facilitates, regressions, adequately, behaviours, occurrence, extremely, knowledge, achieved, happened, locating, besides, compass, factors, notably, popular, speeds, agile; we map those archived descriptors to the telemetry delay control. For the reporting rule, [4] supplies abstract-backed evidence on microservice, anomaly, cloud through the study Unified Graph Modeling of Multi-Source Monitoring Data with Spatiotemporal Attention for Microservice Anomaly Detection. Its evidence ledger records the specific descriptors adaptively, modalities, formulate, encoded, thereby, edges, tight, msds, discriminative, spatiotemporal, outperforming, correlations, implications, attributes, baselines, dimension, instances, intricate; we map those archived descriptors to the telemetry delay control. For the baseline, [5] supplies abstract-backed evidence on anomaly, cloud, service through the study From Detection to Action: AI-Driven Anomaly Detection and Root Cause Synthesis for Cloud Infrastructure Operations. Its evidence ledger records the specific descriptors corroborating, investigation, understanding, exclusively, progressing, proprietary, recommended, impression, navigation, prometheus, dashboard, immediate, insurance, processed, requiring, synthesis, targeting, assessed; we map those archived descriptors to the telemetry delay control. For the measurement, [6] supplies abstract-backed evidence on microservice, anomaly, cloud through the study COMPARATIVE ANALYSIS OF SEQUENCE-BASED AND GRAPH-BASED DEEP LEARNING FOR ANOMALY DETECTION IN MICROSERVICE TRACES. Its evidence ledger records the specific descriptors relationships, alternatives, paradigms, fidelity, tracevae, emerged, gnns, outperforming, implications, deeptralog, dimensions, transition, intricate, inherent, scores, tool, monolithic, awareness; we map those archived descriptors to the telemetry delay control. For the stress range, [7] supplies abstract-backed evidence on anomaly, cloud, service through the study Generative Identity Forensics and Trust System (GIFTS): Geodesic anomaly detection and manifold diffusion for cloud identity telemetry. Its evidence

ledger records the specific descriptors authentication, sessionization, vectorization, contemporary, separability, adversaries, cardinality, coordinates, invocations, permissions, constitute, blackouts, disabling, forensics, generator, geometric, inversion, nonlinear; we map those archived descriptors to the telemetry delay control. For the error slice, [8] supplies abstract-backed evidence on anomaly, cloud, service through the study Anomaly Detection in Cloud Environments. Its evidence ledger records the specific descriptors functionalities, methodologies, subsequently, virtualized, commercial, legitimate, particular, considers, exclusive, signature, attained, employed, manifest, overview, pinpoint, compose, derived, firstly; we map those archived descriptors to the telemetry delay control. For the reporting rule, [9] supplies abstract-backed evidence on microservice, cloud, service through the study Transforming Legacy .NET Architectures into Scalable Cloud-Enabled Systems via Controlled Microservice Pattern Adoption. Its evidence ledger records the specific descriptors sustainability, transformation, destabilizing, inconsistency, independently, modernization, organizations, restructuring, transactional, academically, traceability, transforming, uncontrolled, incremental, initiatives, refactoring, replacement, synthesized; we map those archived descriptors to the telemetry delay control. For the baseline, [10] supplies abstract-backed evidence on anomaly, cloud through the study AI-Driven Threat Detection: Enhancing Cloud Security with Generative Models for Real-Time Anomaly Detection and Risk Mitigation. Its evidence ledger records the specific descriptors signaturebased, strengthening, implementing, safeguarding, transformers, adversarial, unsurpassed, advantages, regulator, sensible, explore, prevent, solved, comes, gans, vaes, sophisticated, directions; 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 low-signal stratum. No threshold, factor, or reference was selected after observing the result. The threshold audit 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, shifted-distribution setting, and parameter table; it does not claim observed field measurements. This keeps the telemetry delay experiment inexpensive while preserving a rerunnable threshold audit.

4. Results

The baseline mean was 0.516; the intervention mean was 0.570. The paired change was +0.054, with a 95% interval from +0.051 to +0.057. In the prespecified adverse slice, the change was +0.045.

The machine-readable artifact `experiments/01-R05-084.json` contains all 48 paired records for microservice anomaly detection. Consequently, the +0.054 result can be recomputed directly under the shifted-distribution setting rather than inferred from prose.

5. Discussion

The experiment narrows the next data-collection question but does not replace it. For microservice anomaly detection under shifted-distribution, the sign of the early-warning F1 change remained positive in both the full set and the low-signal stratum. This supports a focused follow-up on telemetry delay using measured inputs and the same threshold audit 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 low-signal stratum, and the threshold audit 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 shifted-distribution generator. A real-data study should retain telemetry delay and the low-signal 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. Jasani, T., & Padhya, M. (2025). Anomaly Detection on AWS cloud by using Supervised Machine Learning. International Journal of Next-Generation Computing. https://doi.org/10.47164/ijngc.v16i1.1837

3. Kohyarnejadfard, I., Aloise, D., Azhari, S. V., & Dagenais, M. R. (2022). Anomaly detection in microservice environments using distributed tracing data analysis and NLP. Journal of Cloud Computing, 11(1), 25. https://doi.org/10.1186/s13677-022-00296-4

4. Chen, X., Shi, G., & Cao, W. (2026). Unified Graph Modeling of Multi-Source Monitoring Data with Spatiotemporal Attention for Microservice Anomaly Detection. https://doi.org/10.2139/ssrn.6504540

5. Kasinadhuni, B. (2026). From Detection to Action: AI-Driven Anomaly Detection and Root Cause Synthesis for Cloud Infrastructure Operations. International Journal of AI, BigData, Computational and Management Studies, 7, 171-184. https://doi.org/10.63282/3050-9416.ijaibdcms-v7i2p126

6. Zhussupov, Y., Zhumadillayeva, A., Shinassylov, S., & Amirtayeva, A. (2026). COMPARATIVE ANALYSIS OF SEQUENCE-BASED AND GRAPH-BASED DEEP LEARNING FOR ANOMALY DETECTION IN MICROSERVICE TRACES. Journal of Problems in Computer Science and Information Technologies, 4(2). https://doi.org/10.26577/jpcsit4220265

7. Akugri, K. M., Agbenyo, P., & Akugri, M. (2026). Generative Identity Forensics and Trust System (GIFTS): Geodesic anomaly detection and manifold diffusion for cloud identity telemetry. World Journal of Advanced Research and Reviews, 29(1), 1536-1547. https://doi.org/10.30574/wjarr.2026.29.1.0212

8. Marnerides, A. K. (2019). Anomaly Detection in Cloud Environments. Cloud Security, 140-164. https://doi.org/10.4018/978-1-5225-8176-5.ch006

9. Boddupally, H. L. (2026). Transforming Legacy .NET Architectures into Scalable Cloud-Enabled Systems via Controlled Microservice Pattern Adoption. https://doi.org/10.2139/ssrn.6265899

10. Vadisetty, R., Polamarasetti, A., Rongali, S. K., Prajapati, S., & Butani, J. B. (2025). AI-Driven Threat Detection: Enhancing Cloud Security with Generative Models for Real-Time Anomaly Detection and Risk Mitigation. https://doi.org/10.2139/ssrn.5218294