The OSCAR PUBLISHING Services
American Journal of Applied Science and Technology

VOLUME Vol.05 Issue09 2025
PAGE NO. 136-141

# Integrating Machine Learning and Deep Reinforcement Learning for Anomaly Detection and Autonomous Resource Allocation in Distributed Microservice Ecosystems: A Comprehensive Analysis of System Reliability and Sustainability

Dr. Alistair Sterling
Department of Computer Science and Information Systems, University of Melbourne, Australia

**Received:** 15 August 2025; **Accepted:** 09 September 2025; **Published:** 30 September 2025

**Abstract:** The rapid evolution of cloud computing, transition toward 6G networking, and the proliferation of microservice architectures have necessitated a paradigm shift in how system reliability and resource efficiency are managed. Traditional monolithic monitoring systems are increasingly inadequate for the dynamic, decoupled nature of modern distributed environments. This research provides an exhaustive exploration of the integration of machine learning and deep reinforcement learning (DRL) for detecting execution anomalies and optimizing resource distribution. By synthesizing methodologies ranging from time-weighted control flow graph mining to recurrent neural network attention mechanisms, this study develops a theoretical framework for proactive system maintenance. We analyze the role of log-based anomaly detection, such as DeepLog and Logsed, in predicting failures before they impact service level agreements (SLAs). Furthermore, the paper investigates the intersection of environmental sustainability and system performance, examining how self-adaptive approaches can harness renewable energy in cloud ecosystems. The research also delves into the complexities of multi-domain service deployment within 5G and 6G frameworks, emphasizing the necessity of reliability-aware algorithms. Through an extensive review of existing literature and the proposition of a multi-level self-adaptation model, this article demonstrates that the future of resilient distributed systems lies in the convergence of automated boundary detection, intelligent workload scheduling, and adaptive flow control mechanisms. The findings suggest that while deep learning offers unprecedented accuracy in anomaly diagnosis, the integration of human-centric design-moving from Industry 4.0 toward Society 5.0-remains critical for the ethical and practical deployment of autonomous IT infrastructures.

**Keywords:** Microservices, Anomaly Detection, Deep Reinforcement Learning, Cloud Computing, Edge Computing, System Reliability, 6G Networks.

**Introduction:** The contemporary digital landscape is defined by an unprecedented scale of complexity. As organizations migrate from monolithic legacy systems to modular, microservice-based architectures, the challenges associated with monitoring, maintaining, and securing these systems have grown exponentially. In a distributed environment, a single transaction might traverse dozens of independent services, each running in its own containerized environment, often across geographically dispersed data centers. This fragmentation, while providing benefits in terms of scalability and development velocity, introduces significant "blind spots" in system observability. Traditional rule-based monitoring tools, which rely on static thresholds and manual intervention, are no longer capable of keeping pace with the ephemeral nature of modern cloud-native applications.

The fundamental problem addressed in this research is the detection of "soft failures" and performance anomalies that do not necessarily result in immediate

system crashes but degrade user experience and violate service level agreements (SLAs). As noted by Yagoub, Khan, and Jiyun (2018), IT equipment monitoring requires sophisticated forecasting and anomaly detection in log files to move from reactive to proactive maintenance. Log files serve as the "digital exhaust" of a system, capturing the granular interactions between components. However, the sheer volume of logs generated-often reaching gigabytes per hour in large-scale deployments - makes manual analysis impossible.

The literature suggests a transition toward machine learning (ML) and deep learning (DL) as the primary vehicles for making sense of this data. For instance, the use of recurrent neural networks (RNNs) with attention mechanisms has been proposed to provide interpretability in system log anomaly detection (Brown, Tuor, Hutchinson, and Nichols, 2018). This is a critical development because, in an industrial context, knowing that a system is failing is only half the battle; engineers must also understand why it is failing. Interpretable models allow for faster root cause analysis, reducing the mean time to repair (MTTR).

Furthermore, the shift toward edge computing and the impending arrival of 6G networks introduce new layers of complexity. In these environments, resource allocation is not just about CPU and memory; it involves managing latency-sensitive workloads across multiple domains and ensuring reliability in smart ecosystems (Kibalya, Serrat, Gorricho, Okello, and Zhang, 2020). The integration of deep reinforcement learning (DRL) becomes essential here, as it allows systems to learn optimal scheduling policies through continuous interaction with a dynamic environment (Zheng, Wan, Zhang, and Jiang, 2022).

Despite these advancements, a significant gap remains in creating a unified framework that connects low-level log analysis with high-level resource orchestration. Most existing studies focus on either anomaly detection or resource management in isolation. This research seeks to bridge that gap by arguing that anomaly detection should serve as the primary feedback loop for self-adaptive resource managers. By modularizing legacy systems through machine learning-assisted service boundary detection (Hebbar, 2022) and employing multi-level self-adaptation (Zhang, Zhang, Ni, and Liu, 2019), organizations can build systems that are not only resilient but also sustainable and SLA-aware.

Detailed Theoretical Background and Problem Statement

To understand the necessity of the proposed integrated frameworks, one must first dissect the inherent limitations of current distributed systems. Distributed systems are characterized by partial failure modes, where a subset of components fails while the rest of the system remains operational. This often leads to "cascading failures," where a bottleneck in one microservice creates a backpressure effect that eventually topples the entire architecture.

The early work by Fu, Lou, Wang, and Li (2009) laid the groundwork for unstructured log analysis, recognizing that logs are essentially natural language streams produced by software. Unlike structured metrics (like CPU usage), logs contain rich semantic information about the execution path of a program. If a system deviates from its normal execution flow, the logs will reflect this shift. However, as systems evolved, the "unstructured" nature of logs became a hindrance. Modern deep learning approaches, such as DeepLog, have addressed this by treating log entries as a sequence of events, similar to how natural language processing (NLP) treats words in a sentence (Du, Li, Zheng, and Srikumar, 2017). By modeling the probability of the "next log event," these systems can flag anomalies when the observed event significantly deviates from the predicted one.

The problem statement also extends to the physical layer. Cloud computing is an energy-intensive industry. As the world moves toward Society 5.0, as discussed by Maier (2021), the human and environmental impact of technology must be prioritized. This leads to the challenge of managing applications while harnessing renewable energy (Xu, Toosi, and Buyya, 2020). A system that is "reliable" but environmentally destructive is no longer acceptable in the modern academic or corporate discourse. Therefore, the problem is threefold: how to detect anomalies with high precision and interpretability, how to adapt resource allocation in real-time to maintain SLAs, and how to achieve these goals while minimizing the carbon footprint of the underlying infrastructure.

## METHODOLOGY

The methodology employed in this research involves a multi-staged theoretical synthesis and an analysis of algorithmic frameworks proposed in the current literature. We categorize the methods into three primary pillars: Data Acquisition and Preprocessing, Anomaly Modeling, and Adaptive Orchestration.

Data Acquisition and Preprocessing for Log-Based Intelligence The first step in any ML-based system is the transformation of raw data into a format suitable for algorithmic ingestion. In the context of log files, this involves "log parsing." Since logs are often unstructured, researchers use techniques to extract "log templates." For example, a log line like

"Connection from 192.168.1.1 closed" is parsed into a template like "Connection from <IP> closed." This allows the system to focus on the event type rather than the specific variables. The research of Fu et al. (2009) remains foundational here, as it demonstrated the effectiveness of grouping similar log messages to identify execution patterns.

Control Flow Graph (CFG) Mining and Time-Weighted Analysis Once logs are parsed, the next methodological hurdle is understanding the sequence of events. Nandi, Mandal, Atreja, Dasgupta, and Bhattacharya (2016) introduced the concept of program control flow graph mining. By treating the execution of a program as a graph where nodes are log events and edges are transitions, researchers can identify "illegal" paths that represent bugs or security breaches. This was further refined by Jia, Yang, Chen, Li, Meng, and Xu (2017) with the introduction of "Logsed," which utilizes time-weighted control flow graphs. In this approach, the time interval between log events is treated as a critical feature. A system might follow the correct sequence of events, but if the interval between "Database Query" and "Response Received" suddenly jumps from 10ms to 5000ms, Logsed identifies this as a performance anomaly.

Deep Learning and Attention Mechanisms The core methodology for high-accuracy detection involves DeepLog (Du et al., 2017), which utilizes Long Short-Term Memory (LSTM) networks. LSTMs are uniquely suited for this task because they can remember long-term dependencies in sequential data. However, LSTMs are often criticized for being "black boxes." To address this, the methodology incorporates the work of Brown et al. (2018), which adds attention mechanisms to the RNN. The attention layer assigns weights to different parts of the log sequence, allowing the model to "point" to exactly which previous events led to the current anomaly flag. This provides the interpretability required for senior editors and system architects to trust the model's output.

Reinforcement Learning for Resource Allocation For the resource allocation component, the methodology shifts toward Markov Decision Processes (MDPs). As described by Yang, Nguyen, Jin, and Nahrstedt (2019) in the MIRAS framework, model-based reinforcement learning can be used to allocate resources to microservices within scientific workflows. The "Agent" (the resource manager) observes the "State" (current CPU, memory, and latency), takes an "Action" (scaling up or down), and receives a "Reward" (based on SLA compliance and cost). In multi-domain environments, such as those discussed by Kibalya et al. (2019) for 5G network slicing, this methodology is expanded to handle constraints across different administrative boundaries, ensuring that reliability is maintained even as data passes between different providers.

Self-Adaptive Heartbeat Detection Finally, for real-time fault diagnosis, we analyze the Multi-Factor Self-Adaptive Heartbeat Detection algorithm (Zang, Chen, Zou, Zhou, Lisong, and Ruigang, 2018). Unlike traditional heartbeats that have a fixed frequency, this adaptive method changes its checking interval based on network conditions and historical reliability data, preventing "false positives" in unstable network environments.

## RESULTS

The descriptive analysis of the integrated frameworks reveals a significant improvement in system resilience when ML-based anomaly detection is paired with DRL-based resource management. We categorize these results into four major findings.

Finding 1: Superiority of Sequence-Based Models over Threshold-Based Systems The analysis shows that sequence-based models like DeepLog and Logsed consistently outperform traditional threshold-based monitoring. In complex microservice environments, a "normal" CPU usage level for one service might be an "abnormal" level for another. Furthermore, many failures occur within "normal" resource parameters. By focusing on the execution path rather than just metrics, sequence-based models can detect logic errors and deadlocks that do not immediately spike resource usage. The use of time-weighted CFGs (Jia et al., 2017) is particularly effective in cloud environments where network jitter is common, as it can distinguish between a transient network hiccup and a systemic failure.

Finding 2: The Efficacy of Attention Mechanisms in Root Cause Analysis The inclusion of attention mechanisms (Brown et al., 2018) provides a measurable decrease in MTTR. In qualitative terms, when an anomaly is detected, the system provides a "heat map" of the log sequence. Our analysis suggests that this allows junior developers to perform at the level of senior researchers when diagnosing distributed system failures. The interpretability of the model bridges the gap between AI-driven automation and human oversight, a key tenet of the Society 5.0 vision (Maier, 2021).

Finding 3: Impact of Smart Deployment Frameworks on SLA Compliance The implementation of SLA-aware frameworks like SmartVM (Zheng, Zheng, Zhang, Deng, Dong, Zhang, and Liu, 2019) demonstrates that microservice deployment can be optimized to meet specific performance targets. By treating the deployment as an optimization problem, these frameworks ensure that latency-critical services are placed on high-performance nodes, while background tasks are relegated to lower-cost or "green" energy

nodes. The results indicate that DRL-based workload scheduling (Zheng et al., 2022) can reduce SLA violations by up to 30% compared to round-robin or least-connection scheduling algorithms.

Finding 4: Reliability in Multi-Domain and Edge Ecosystems Research into 5G and 6G ecosystems (Giordani, Polese, Mezzavilla, Rangan, and Zorzi, 2020) highlights the success of multi-domain service deployment. Using deep reinforcement learning, systems can navigate the complexities of network slicing, where different slices (e.g., one for autonomous vehicles, one for general internet) require different reliability guarantees. The work by Kibalya et al. (2020) proves that reliability-aware algorithms can maintain service continuity even when moving between different administrative domains, which is crucial for the global scalability of smart city technologies.

## DISCUSSION

The results presented above suggest a promising future for autonomous system management, but they also raise several theoretical and practical questions that merit deep interpretation.

The Convergence of 6G and Human-Centric Systems As we look toward 6G, the discussion must move beyond raw speed and latency. Maier (2021) argues for a shift "from Industry 4.0 toward Society 5.0." This means that the anomaly detection and resource allocation systems we build must consider the human element. For example, if an AI-driven system decides to shut down a low-priority service to save energy or maintain the reliability of a high-priority service, what are the ethical implications if that "low-priority" service is a communication tool for a marginalized community? The "intelligence" of the system must be tempered by policy and human values.

Interpretable AI vs. Performance Trade-offs A recurring theme in the research is the trade-off between model complexity and interpretability. While deep neural networks provide the highest accuracy, they are often the hardest to explain. The attention mechanisms proposed by Brown et al. (2018) are a step in the right direction, but they add computational overhead. In an edge computing context, where resources are limited (Cao, Zhang, Li, Feng, and Cao, 2019), the energy cost of running a complex deep learning model for anomaly detection might outweigh the benefits. This suggests a need for "lightweight" ML models that can run on the edge while still providing enough diagnostic information to be useful.

The Challenge of Legacy Modularization A significant hurdle for many organizations is the presence of monolithic legacy systems. Hebbar (2022) proposes ML-assisted service boundary detection as a solution. This is a critical discussion point because even the best anomaly detection system will struggle with a "spaghetti-code" monolith where everything is tightly coupled. The ability to automatically identify where one service should end and another should begin allows for the gradual modernization of infrastructure, making the subsequent application of DeepLog or SmartVM much more effective.

Sustainability and Renewable Energy The work of Xu et al. (2020) on harnessing renewable energy represents a vital shift in the cloud computing discourse. Future research must explore how to make anomaly detection "energy-aware." If the system detects a potential failure but correcting it requires a massive burst of energy during a period of low renewable availability, should the system delay the fix? This creates a multi-objective optimization problem where reliability, cost, and sustainability must be balanced.

Limitations and Future Scope Despite the robustness of current ML models, "data drift" remains a significant limitation. Software is constantly updated; a log pattern that was "normal" yesterday might be "abnormal" after a new version is deployed. Current models require frequent retraining, which is resource-intensive. Future scope includes the development of "online learning" models that can adapt to new log patterns in real-time without needing a complete overhaul. Additionally, the security of the anomaly detection systems themselves is a burgeoning field. If an attacker knows how the ML model works, they could potentially "poison" the logs to hide their activities.

## CONCLUSION

This research has synthesized a wide array of perspectives on the monitoring and management of modern distributed systems. From the granular analysis of log files using time-weighted control flow graphs and deep learning to the high-level orchestration of resources using reinforcement learning in 5G and 6G environments, a clear picture emerges: the manual management of IT infrastructure is at its end.

The integration of machine learning into the system lifecycle-from initial modularization to daily operations and long-term scaling-is not just an optimization but a necessity for survival in the digital age. We have shown that deep learning models like DeepLog, when enhanced with attention mechanisms, provide the dual benefits of high accuracy and human-readable interpretability. Furthermore, we have highlighted that system reliability must be viewed through the lenses of both SLA compliance and environmental sustainability.

As we move toward Society 5.0, the goal is to create "self-healing" systems that can detect anomalies,

diagnose root causes, and reallocate resources autonomously, all while minimizing their carbon footprint. The path forward involves refining these models to be more energy-efficient, resilient to data drift, and deeply integrated into the cross-domain realities of future telecommunications networks. By adopting the multi-level self-adaptation approaches discussed in this article, the academic and industrial communities can ensure that the next generation of cloud and edge computing is as reliable as it is revolutionary.

**REFERENCES**


1. Brown, A.; Tuor, A.; Hutchinson, B.; Nichols, N. Recurrent neural network attention mechanisms for interpretable system log anomaly detection. In Proceedings of the First Workshop on Machine Learning for Computing Systems, Tempe, AZ, USA, 12 June 2018; pp. 1–8.
2. Cao B, Zhang L, Li Y, Feng D, Cao W (2019) Intelligent offloading in multi-access edge computing: a state-of-the-art review and framework. IEEE Commun Magaz 57(3):56–62.
3. Du, M.; Li, F.; Zheng, G.; Srikumar, V. Deeplog: Anomaly detection and diagnosis from system logs through deep learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, Dallas, TX, USA, 30 October–3 November 2017; pp. 1285–1298.
4. Fu, Q.; Lou, J.G.; Wang, Y.; Li, J. Execution anomaly detection in distributed systems through unstructured log analysis. In Proceedings of the 2009 Ninth IEEE International Conference on Data Mining, Miami Beach, FL, USA, 6–9 December 2009; pp. 149–158.
5. Giordani M, Polese M, Mezzavilla M, Rangan S, Zorzi M (2020) Toward 6g networks: use cases and technologies. IEEE Commun Magaz 58(3):55–61.
6. K. S. Hebbar, “MACHINE LEARNING-ASSISTED SERVICE BOUNDARY DETECTION FOR MODULARIZING LEGACY SYSTEMS,” International Journal of Applied Engineering & Technology, vol. 04,no.02, pp. 401-414, Sep. 2022, https://romanpub.com/resources/ijaet-v4-2-2022-48.pdf
7. Jia, T.; Yang, L.; Chen, P.; Li, Y.; Meng, F.; Xu, J. Logsed: Anomaly diagnosis through mining time-weighted control flow graph in logs. In Proceedings of the 2017 IEEE 10th International Conference on Cloud Computing (CLOUD), Honololu, HI, USA, 25–30 June 2017; pp. 447–455.
8. Kibalya G, Serrat J, Gorricho J-L, Okello D, Zhang P (2020) A deep reinforcement learning-based algorithm for reliability-aware multi-domain service deployment in smart ecosystems, Neural Computing and Applications 1–23.
9. Kibalya G, Serrat J, Gorricho J-L, Pasquini R, Yao H, Zhang P (2019) A reinforcement learning based approach for 5g network slicing across multiple domains, In: 2019 15th International Conference on Network and Service Management (CNSM), IEEE, pp. 1–5.
10. Maier M (2021) 6g as if people mattered: From industry 4.0 toward society 5.0, In: 2021 International Conference on Computer Communications and Networks (ICCCN), IEEE, pp. 1–10.
11. Mao Y, You C, Zhang J, Huang K, Letaief KB (2017) A survey on mobile edge computing: the communication perspective. IEEE Commun Surveys Tutor 19(4):2322–2358.
12. Nandi, A.; Mandal, A.; Atreja, S.; Dasgupta, G.B.; Bhattacharya, S. Anomaly detection using program control flow graph mining from execution logs. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; pp. 215–224.
13. Sharma, B.; Jayachandran, P.; Verma, A.; Das, C.R. CloudPD: Problem determination and diagnosis in shared dynamic clouds. In Proceedings of the 2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Budapest, Hungary, 24–27 June 2013; pp. 1–12.
14. Xu, M., Toosi, A. N., and Buyya, R. (2020). A Self-adaptive Approach for Managing Applications and Harnessing Renewable Energy for Sustainable Cloud Computing. IEEE Transactions on Sustainable Computing.
15. Yagoub, I.; Khan, M.A.; Jiyun, L. IT equipment monitoring and analyzing system for forecasting and detecting anomalies in log files utilizing machine learning techniques. In Proceedings of the 2018 International Conference on Advances in Big Data, Computing and Data Communication Systems (icABCD), Durban, South Africa, 6–7 August 2018; pp. 1–6.
16. Yang, Z., Nguyen, P., Jin, H., and Nahrstedt, K. (2019). MIRAS: Model-based Reinforcement Learning for Microservice Resource Allocation over Scientific Workflows. In 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), pp. 122–132.
17. Yudong, L., Yuqing, Z., and Zhangbin, Z. (2020). Service Availability Guarantee with Adaptive

Automatic Flow Control. In 2020 IEEE World Congress on Services (SERVICES), pp. 101–105.

**18.** Zang, X., Chen, W., Zou, J., Zhou, S., Lisong, H., and Ruigang, L. (2018). A Fault Diagnosis Method for Microservices Based on Multi-Factor Self-Adaptive Heartbeat Detection Algorithm. In 2018 2nd IEEE Conference on Energy Internet and Energy System Integration (EI2), pp. 1–6.

**19.** Zhang, S., Zhang, M., Ni, L., and Liu, P. (2019). A Multi-Level Self-Adaptation Approach For Microservice Systems. In 2019 IEEE 4th International Conference on Cloud Computing and Big Data Analysis (ICCCBDA), pp. 498–502.

**20.** Zheng, T., Wan, J., Zhang, J., Jiang C (2022) Deep reinforcement learning-based workload scheduling for edge computing. J Cloud Comput 11(1):3.

**21.** Zheng, T., Zheng, X., Zhang, Y., Deng, Y., Dong, E., Zhang, R., and Liu, X. (2019). SmartVM: a SLA-aware microservice deployment framework. World Wide Web, 22(1):275–293.