International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

AI-Driven Autonomous IT Operations: A Human-in-the-Loop AIOps 2.0 Framework

Samrat Mukherjee

Submitted: 17/10/2024 Revised: 20/11/2024 Accepted: 16/12/2024

Abstract: Artificial Intelligence for IT Operations (AIOps) has become a game-changer for the management of large and complex digital environments, driven by intelligent automation, predictive analytics, and decision-making. Current AIOps solutions, however, face challenges with explainability, governance, interoperability, and rely too much on fully automated remediation. This paper introduces a new framework for Human-in-the-Loop AIOps 2.0, combining anomaly detection with reinforcement learning and human validation, to ensure reliable autonomous IT operations. The proposed methodology involves utilising the Kaggle IT Incident Log Dataset, Data Preprocessing, Feature Engineering, anomaly detection using Isolation Forest (IF) algorithm and a Deep Q-Network (DQN)-based recommendation engine to create intelligent remediation actions. Human operators review AI-generated decisions for validation before execution, enhancing the trust, transparency, and governance of the operation. Through experimental assessment, the proposed framework has been proven to be the most effective with an accuracy (ACC) of 94.6%, precision (PRE) of 94.1%, recall (REC) of 94.8%, and F1 Score (F1) of 94.4% better than traditional monitoring systems and fully AI-driven monitoring systems. Additionally, the human validation process can greatly minimize false positives and mean time to resolution, and enhance downtime reduction. The findings of the study underline the need to build autonomous intelligence and integrate it with human intelligence for next-generation resilient, explainable and adaptive IT operations management.

Keywords: *AIOps, Human-in-the-Loop, Autonomous IT Operations, Intelligent Incident Management, Predictive Analytics.*

I. INTRODUCTION

As businesses navigate an era of digital transformation, AI has become more than merely a valuable tool—it's a strategic imperative for transforming IT operations [1]. Today, organisations operate large-scale heterogeneous systems of servers, applications, databases, virtual machines, microservices and network systems that are running all the time across geographically distributed environments [2]. Manual monitoring, rule-based alerting systems, and static operational rules are the foundation of traditional IT operations management strategies. These methods are unable to handle the ever-changing infrastructure and massive amounts of operational data. This has resulted in the incorporation of AI for IT Operations (AIOps) into business processes, which aids in automation, prediction, anomaly detection, and intelligent incident response [3].

Automation, big data analytics, AI, and machine learning are all part of AIOps, which aims to improve the dependability and efficiency of infrastructure systems' operations . To detect operational irregularities and minimize downtime, modern AIOps platforms employ various techniques, including anomaly detection, predictive analytics, natural language processing, reinforcement learning, and automated remediation [4]. Several technologies such as Isolation Forest, Deep Neural Networks, ARIMA forecasting, clustering algorithms and Reinforcement Learning models have been successfully incorporated in intelligent monitoring systems to enable proactive incident management [5]. These techniques can help organizations to streamline and automate the processing of large quantities of log and system data, event streams, and network telemetry data in real-time, with minimal human involvement [6].

Although, there are significant improvements, there are some critical challenges currently with AIOps systems [7]. Most solutions are black-box systems that are not easily explained, and lack enough human oversight, which diminishes organizational trust in autonomous decision-making. Furthermore other challenges for operational reliability remain, such as false positive or false negative results, model drift, interoperability with legacy systems, governance and uncontrolled automation [8][9]. In certain areas like healthcare, finance, telecommunication and cloud-native infrastructures, fully autonomous remediation strategies are sometimes enough to cause wrong actions and impact critical services [10][11]. The importance of using human expertise in autonomous operational workflows has increased over time.

This research suggests a Human-in-the-Loop AIOps 2.0 framework, which integrates intelligent anomaly detection, reinforcement learning-driven remediation, and human validation mechanisms for robust and explainable IT operations management. A deep Q-learning reinforcement learning agent is used to find operational outliers, and the

Program & Project Management
New Jersey, USA
samrat0864@gmail.com

Isolation Forest technique is used to find ways to fix them. Human operators monitor and approve AI-driven actions prior to any execution, guaranteeing transparency, governance, and safety of operations [12]. This research is important to ensure that decisions are reliably made, downtime is reduced, false positives are minimised and adaptive learning is enhanced in intelligent operational ecosystems. The significant findings of this study are listed below:

- A novel Human-in-the-Loop AIOps 2.0 framework is proposed by integrating anomaly detection, reinforcement learning, and human-supervised autonomous IT operations for reliable infrastructure management.
- An Isolation Forest-based anomaly detection model is implemented to identify abnormal operational incidents and irregular system behaviours from large-scale IT incident datasets.
- A Deep Q-Network (DQN)-based reinforcement learning engine is developed to generate adaptive remediation recommendations for intelligent incident resolution and operational optimization.
- Human validation mechanisms are incorporated into the remediation workflow to improve explainability, governance, transparency, operational trust, and decision reliability in autonomous IT environments.

A. Novelty of the paper

The novelty of this research lies in the development of a Human-in-the-Loop AIOps 2.0 framework that combines autonomous intelligence with human-supervised decision-making for reliable IT operations management. The framework integrates Isolation Forest-based anomaly detection with a Deep Q-Network reinforcement learning engine to generate adaptive remediation recommendations. Unlike fully automated AIOps systems, human operators validate AI-generated actions before execution, improving explainability, governance, operational trust, and reliability. The proposed framework also reduces false positives, minimizes downtime, and supports scalable intelligent remediation for modern cloud-native and distributed enterprise infrastructures.

B. Structure of the paper

The paper is organized as follows. In Section II, the literature review and comparison of current AIOps studies are provided. In Section III, the proposed Human-in-the-Loop AIOps framework is described. Experimental results and performance evaluation are discussed in Section IV. Lastly, the paper ends with the results, the limits, and future research directions in Section V.

II. LITERATURE REVIEW

The papers reviewed showed considerable progress in the capabilities of AIOps technologies such as anomaly detection, predictive analytics, and automation, and some limitations in terms of explainability, interoperability, governance, and autonomous decision-making for humans.

S. Chitta et al. (2024) offer a comprehensive overview of how AIOps can be applied in different sectors, such as financial services, healthcare, and telecom. In the financial space, AIOps tools are applied to keep an eye on transaction systems in the event of irregularities, discover potential fraud, and guarantee regulatory compliance. In healthcare, AIOps improves the reliability of electronic health records (EHR) systems, and helps manage complex IT infrastructures in healthcare facilities. By combining these AI capabilities with data from network components, AIOps helps telecommunications companies to fine-tune their network performance, allocate resources effectively, and deliver better customer experiences by ensuring high service availability and low latency. While the impact of AIOps is revolutionary, it isn't without its hurdles. Data quality is a significant issue as the accuracy and completeness of the data are crucial for the effectiveness of the AI and ML models. The integration of AIOps solutions with existing IT systems can introduce further complexities, such as interoperability and system compatibility [13].

H. Allam (2023) talks about the basic technologies that are needed for AIOps and how they can be added to current IT systems to make them easier to see and improve the ACC of decisions. It also examines why AIOps is a solution to some of the major challenges faced, such as data silos, alert fatigue, and the fact that a rule-based approach to monitoring simply isn't cutting it. A study showcasing the implementation of AIOps in a real-world hybrid infrastructure, focusing on the tangible benefits such as decreased downtime and quicker incident resolution, is presented [14].

S. Polisetty (2023) provides a number of useful methods for selecting a suitable AI model, cleaning up the data, and maintaining the algorithms. Part of the essay focuses on learning about the several kinds of algorithms that may be used for real-time data analysis, including algorithms that detect anomalies in metrics and natural language logs. Scaling models in big distributed systems and dynamic algorithm adaptation in response to operational needs are two examples of the new methods proposed for AI model management. Adapting AI models to operational needs and scaling models in dispersed, large-scale settings are two areas that might benefit from new methods for model management. Models are trained on the data continually as it enters, according to the article's innovative method. This helps AI adapt to changing IT settings and stay relevant [15].

S. M. Polisetty (2022) explains how incidents are recognised from past trends and when triggered by anomaly detection,

like change in system metrics, or in event logs. It's the use of AI-powered incident management that predicts incidents based on trend analysis and metrics that is at the heart of the article. One thing that was talked about is the ability to have runbooks that automatically resolve incidents with pre-defined policies and actions when they occur, without the need for manual intervention, and to improve the response time. Using AIOps-based feedback loops to address incidents is suggested in the paper. With each issue resolved, data is sent back into the system, which improves incident prediction and makes AIOps solutions more resilient [16].

L. Li (2022), A proposed AI search strategy utilises the K-means algorithm. Data samples are categorized in a myriad of ways based on how similar they are to one another. Before establishing a Boosting model to enhance the ACC of time series predictions, the traffic data is subjected to regression analysis using the ARIMA algorithm. In addition, shows how distinct various variables are by calculating their relative distances from one another. Fast data processing and data collection regarding user activity are two of its capabilities [17].

S. Becker et al. (2021) intend to transfer processing power to the network's edge, where users reside, in order to circumvent the issues associated with centralized cloud computing. Concerns about the impact of disruptions on important use cases, such as autonomous vehicles or healthcare, and the increasing complexity of emerging infrastructures are both on the rise. With the use of machine learning techniques, AIOps aims to aid human operators in managing complex infrastructures. To facilitate its operation in distributed and heterogeneous environments, this article describes the design of an AIOps platform. Conduct experiments to find out which of three anomaly detection algorithms performs best on edge devices and figure out how much a high-frequency monitoring solution would cost. The outcomes show that tailored anomaly detection algorithms may be run directly on edge devices, data can be collected frequently, and resource consumption is appropriate [18].

Table I compares recent AIOps studies based on methodologies, contributions, advantages, limitations, and future directions, highlighting existing research gaps and the need for Human-in-the-Loop autonomous IT operations frameworks.

TABLE I. COMPARATIVE ANALYSIS AND RESEARCH GAPS IN EXISTING AIOPS STUDIES

Authors	Methods	Contributions	Advantages	Limitations	Future Work
S. Chitta et al. (2024)	Applied AI/ML-based AIOps in finance, healthcare, and telecom for anomaly detection and monitoring.	Showed improved efficiency in fraud detection, EHR reliability, and network optimization.	Reduces downtime, improves reliability, and enhances customer experience.	Depends on high-quality data and legacy system integration.	Develop HITL-AIOps with explainable AI and adaptive interoperability.
H. Allam (2023)	Integrated AIOps into IT infrastructures for observability and incident management.	Demonstrated reduced downtime and faster incident resolution.	Improves observability and automates monitoring.	Limited scalability and explainability discussion.	Research scalable and explainable HITL-AIOps for cloud-edge systems.
S. Polisetty (2023)	Used adaptive AI models, anomaly detection, and continuous learning for IT operations.	Proposed adaptive learning and distributed model management.	Supports real-time analytics and scalable monitoring.	Risk of model drift and lack of transparency.	Explore explainable AI and human-guided learning in AIOps 2.0.
S. M. Polisetty (2022)	Applied AI-driven incident prediction and automated remediation using logs and metrics.	Introduced predictive incident management and self-improving resolution.	Enables proactive maintenance and faster recovery.	Relies on historical data and predefined rules.	Add human approval and explainable remediation strategies.
L. Li (2022)	Proposed hybrid AI methods using K-means, ARIMA, and	Improved time-series prediction and behaviour analysis.	High prediction accuracy and efficient processing.	Limited focus on autonomous IT operations.	Integrate predictive analytics with

	Boosting for prediction.				HITL operational management.
S. Becker et al. (2021)	Developed distributed edge-based AIOps for monitoring and anomaly detection.	Demonstrated efficient edge monitoring with low overhead.	Reduces latency and supports critical applications.	Limited edge-cloud coordination and governance discussion.	Develop federated cloud-edge AIOps with HITL supervision.

III. Methodology

The methodology proposed combines anomaly detection, reinforcement learning and human validation, facilitating controlled autonomous IT operations, as illustrated in Fig. 1. The IT Incident Log Dataset is subject to a Data Preprocessing phase that involves managing missing values, encoding, normalization, and feature selection in order to improve the data quality. At last, for a fair evaluation of the model, the processed data set is divided into a training set and a testing set applying Data Splitting. The training data is used to implement the Isolation Forest model for anomaly detection and recognizes the abnormal incident patterns. Anomalies detected are then fed to the RL-Based Recommendation Engine that uses a Deep Q-Network to provide best remediation actions. These recommendations are also put to Human Validation and Approval for safety and reliability in operations. Lastly, the performance of the system is assessed through metrics like ACC, PRE, REC, and F1, guaranteeing comprehensive evaluation of predictive and decision-making skills in IT operations.

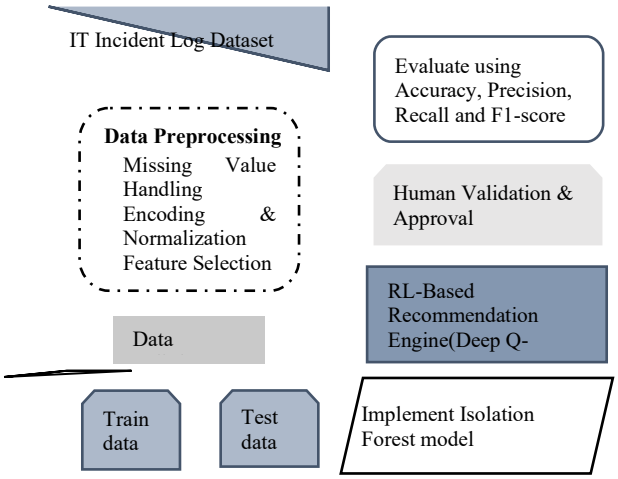

Fig. 1. Proposed Methodology Flowchart for Human-in-the-Loop AIOps Framework

A. Data Collection

The data collection phase involved using the Kaggle IT Incident Log Dataset [19] that comprises around 141,712 incident reports from enterprise IT service management operations. Some of the attributes that are important for operation include incident_id, priority and severity, category and assignment_group, opened_time, resolved_time, response_time, resolution_time, escalation_level, incident_status, and support_team. The attributes contain all the details of the lifecycle and escalation of incidents, along with details of how the operational staff carry out their work.

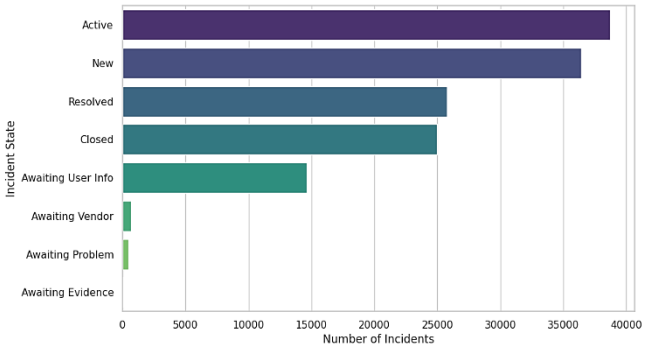

Fig. 2. Incident State Distribution

The distribution of states of the incident in the IT service management data set is shown in Fig 2. Active and New incidents make up the bulk of the incidents, with Resolved and Closed cases coming in next. Awaiting User Info, Vendor, Problem, and Evidence states show fewer cases, suggesting problems with workflows and different levels of resolution progress for incidents in each state of the operating lifecycle respectively.

B. Data Preprocessing

The purpose of Data Pre-processing is to transform the raw IT incident logs into a clean, structured and machine-readable format, ready for analysis using anomaly detection and reinforcement learning. There are several steps, as follows:

- **Missing Value Handling:** The numerical data characteristics are enhanced by using mean imputation, while the categorical data attributes are enhanced by using mode imputation, to fill in missing values and ensure consistency.
- **Duplicate Record Removal:** If multiple entries of the same record appear, they are identified and deleted from the data set to mitigate data duplication, bias and enhance the reliability of machine learning training.

- **Categorical Feature Encoding:** Machine Learning categorical attributes such as Incident Priority, Support team, escalation type and Incident category are converted to numbers using Label Encoding for machine learning compatibility.

C. Normalization

In order to have a more stable learning of the prediction model, the input data was normalized using Min-max normalization or scaling which is also called as min-max scaling [20][21]. Some input features could be overshadowed by a single input feature under a different distribution of the magnitudes of the variables. This scaling is done to make sure that bigger input features don't get in the way of smaller ones from being resized. all variables' sizes, which range from zero to one. What follows is the procedure for applying the min-max normalization. It is derived in Equation (1):

$$x' = \frac{x - x_{min}}{x_{max} - x_{min}} \quad (1)$$

The original value is represented by x, whereas the maximum and minimum values of each feature are x_{max} and x_{min}, respectively.

D. Feature Selection

The feature selection process is used to find most relevant operational attributes that affect anomaly detection, incident classification, and recommendations for remediation in IT incident management areas. Choosing significant features helps increase the ACC of the model, decreases the computational complexity and eliminates redundant or irrelevant operational information from the data set. The incident priority, incident severity, number of escalations, response time, resolution time, support group, and incident category were chosen as significant attributes because they have an impact on how operational anomalies are displayed and on remediation actions.

E. Data Splitting

Preprocessing the dataset and then splitting it into training and testing sets allowed for fair model evaluation while reducing the likelihood of overfitting. Trained the models with 80% of the incident recordings and validated and tested the anomaly detection and reinforcement learning models with the remaining 20%.

F. Implement Isolation Forest Model

An anomaly is defined as a location that is thinly distributed and remote from a dense population; such a location is more likely to be separated in an isolated forest. A dataset in a forest setting is partitioned iteratively until each sample point is completely separate from the others [22]. Abnormal incidents that showed abnormal operational characteristics were considered as anomalies and sent to the Remediation Recommendation module. The anomaly score is calculated based on the Isolation Forest scoring mechanism as given in Equation. (2)

$$s(x, n) = 2^{\frac{-E(h(x))}{c(n)}} \quad (2)$$

where the predicted path length is $E(h(x))$ and the average path length is $c(n)$. Anomaly scores are higher if there are operational irregularities and incident risks.

G. Reinforcement Learning-Based Recommendation Engine

An intelligent reinforcement learning (RL) agent using a deep-learning approach, called Deep Q-Network (DQN), has been designed and created to suggest intelligent remediation actions based on the discovered operational anomalies. A DQN-based RL agent was developed to learn from previous incidents and feedback to recommend remediation actions for operational anomalies. The state contains priority, response delay, severity, anomaly score and escalation count. Actions may involve restarting a service, reallocating resources, reassigning tickets, managing escalation and creating an automated recovery. The reward function is shown in Equation. (3)

$$R = \alpha(DR) - \beta(FP) + \gamma(HA) \quad (3)$$

where DR is downtime reduction, FP is false positives and HA is human approval. The agent receives positive reward for a good recovery and a penalty if they fail to recover or take unhelpful actions.

H. Human Validation

The framework features a HITL layer where operators approve AI remediation actions before implementation. The operators have the ability to approve, reject or tweak recommendations, thereby enhancing transparency, trust and governance, and minimizing the chance of wrongful automation. The reinforcement learning policy is also refined using feedback from humans. When approved remediation actions are then performed in a controlled IT environment, such as restarting service, reallocating resources, escalating downfalls, or automated recovery. Outcomes are evaluated and monitored for effectiveness and to reduce downtime, and returned to learning agent for ongoing policy improvement and adaptive optimization.

I. Evaluation Metrics

This section outlines a set of metrics to assess anomaly detection, reinforcement learning-based remediation, and HITL validation, with a particular emphasis on ACC, remediation effectiveness, system reliability, reduction in downtime and overall system performance improvement.

Accuracy: The categorization process relies on ACC as its primary metric. It measures how many samples were accurately identified relative to how many samples were totalled [23]. As demonstrated in Equation (4) the ACC rate quantifies the frequency with which the classifier produces accurate predictions.

$$Accuracy = \frac{TP+TN}{TP+FN+FP+TN} \quad (4)$$

Precision: The ACC of a classifier is defined by its PRE. The degree to which predictions across all classes are accurate is called PRE. As shown in Equation (5), it is the proportion of samples that have been positively labelled to those that have been positively classified.

$$Precision = \frac{TP}{TP+FP} \quad (5)$$

Recall: A sensitivity (or REC) test measures how many samples were actually positive out of all the samples that were positively tagged. The sensitivity of a positive forecast is a measure of how accurate it is. The percentage of accurate predictions is displayed. The sensitivity of a good classifier, as shown in Equation (6), should be 1.

$$Recall(Rc) = \frac{TP}{TP+FN} \quad (6)$$

F1 score: The F-measure is determined by utilizing measures for sensitivity and PRE. PRE or sensitivity can be achieved by optimizing the system with its help. A reliable indicator of the classifier's efficacy is the F score. Like in Equation (7), the F-score is commonly used in the literature to compare classifier success.

$$F1\ score(F1) = \frac{2\times Precision\times Recall}{Precision+Recall} \quad (7)$$

All of each model's capabilities are thoroughly assessed using these measures.

IV. Results and Discussion

The comparative performance analysis of Traditional Monitoring, Autonomous AIOps, and the proposed HITL-AIOps framework is shown in Table II. The proposed model demonstrated the best ACC (94.6%), PRE (94.1%), REC (94.8%), and F1 (94.4%), demonstrating its excellent performance in terms of anomaly detection and decision-making. Moreover, the success rate of incident resolution increased dramatically to 94.1%, showcasing the ability of combining human-in-the-loop validation with autonomous AIOps operations for reliable infrastructure management.

TABLE II. Omparative Performance Evaluation of Monitoring Approaches

Metric	Traditional Monitoring	Autonomous AIOps	Proposed HITL-AIOps
Accuracy (%)	82.7	90.8	94.6
Precision (%)	81.9	91.4	94.1
Recall (%)	79.8	92.2	94.8
F1-Score (%)	80.8	91.8	94.4
Incident Resolution Success (%)	76.4	87.5	94.1

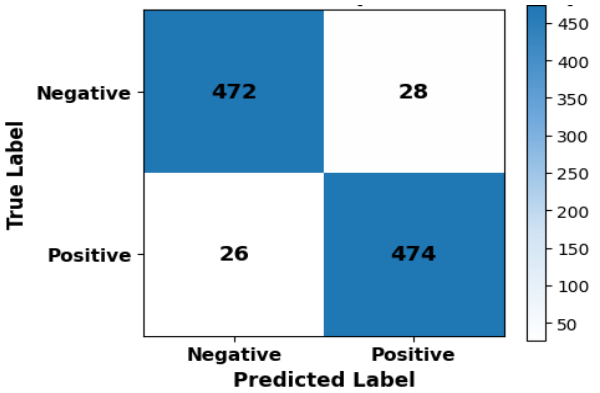

Fig. 3. Confusion Matrix of the Proposed HITL-AIOps Model

Fig 3 shows the suggested HITL-AIOps model's confusion matrix, which shows the classification result utilizing true positive, true negative, false positive, and false negative. There are very few misclassifications for the model, with very low misclassifications in most of the instances. The amount of correct predictions is relatively high, indicating high prediction ability and anomaly detection in operational environments.

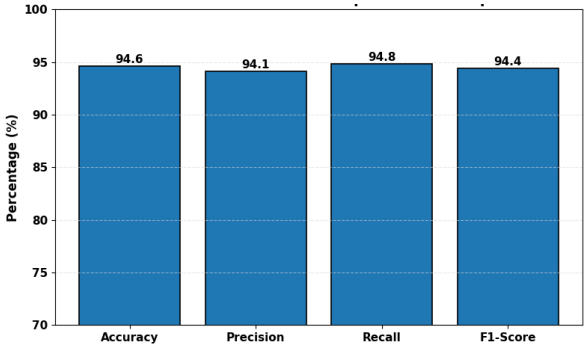

Fig. 4. Performance Metrics of the Proposed HITL-AIOps Framework

Figure 4 shows the results of evaluating the proposed HITL-AIOps framework using a number of measures, including REC, ACC, PRE, and F1. The performance in detecting anomalies and managing incidents was found to be very reliable with all the metrics being above 94%. The results show that the suggested framework is strong and works well in intelligent IT operations settings; they are also balanced.

The results of the different human validation strategies for the operational efficiency metrics are shown in Table III. The lowest MTTR of 41 minutes and the lowest false positive of 5.7% are obtained when using the Full Human-in-the-Loop validation. Furthermore, the reduction in downtime rose to 58.3%, showing that improved human control led to more

accurate decisions, less system disruption, and increased reliability in intelligent IT operations environments.

TABLE III. IMPACT OF HUMAN VALIDATION STRATEGIES ON OPERATIONAL PERFORMANCE

Validation Strategy	MTTR (minutes)	False Positive Rate (%)	Downtime Reduction (%)
No Human Validation	67	9.6	43.8
Partial Human Validation	52	7.4	50.6
Full Human-in-the-Loop	41	5.7	58.3

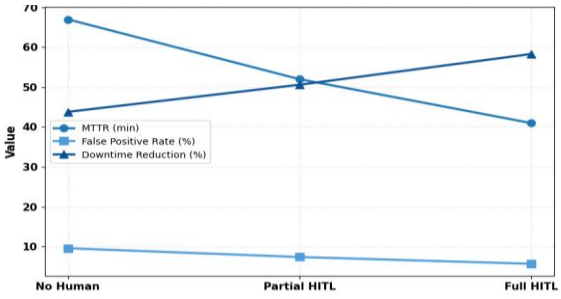

Fig. 5. Impact of Human Validation Strategies on Operational Metrics

Fig 5 shows the impact of various human validation approaches on the operational metrics: MTTR, false positive rate and downtime reduction. MTTR and False Positives dramatically decrease as the level of human involvement rises and Downtime Reduction steadily increases. The findings illustrate the benefits of Human-in-the-Loop validation for improving the reliability of operations and decision-making.

The reinforcement learning training performance across various ranges of episodes is summarized in Table IV. The cumulative and average rewards always increased as training went on and the length of the episode shortened from 24 to 12. The success rate also jumped significantly from 61.4% to 91.4%, showing good convergence of learning and optimization of policy. The results further validate the ability of the proposed HITL-AIOps framework to enhance adaptive decision-making over time.

TABLE IV. REINFORCEMENT LEARNING TRAINING PERFORMANCE ACROSS EPISODES

Episode Range	Cumulative Reward	Average Reward	Episode Length	Success Rate (%)
1–50	1840	36.8	24	61.4
51–100	2485	49.7	21	70.2
101–150	3160	63.2	18	79.1
151–200	3825	76.5	15	86.3
201–250	4355	87.1	12	91.4

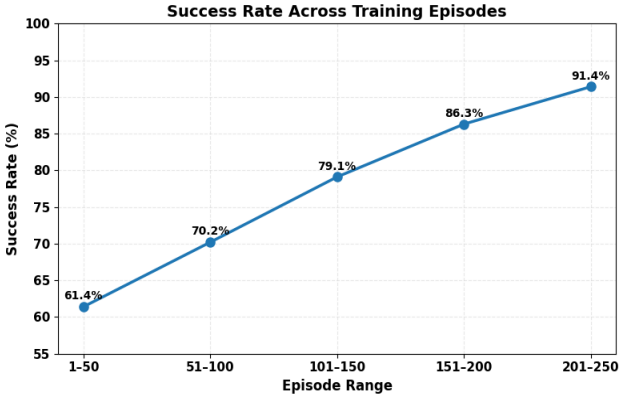

Fig. 6. Success Rate Evolution Across Training Episodes

The improvement of the model's success rate over 250 training episodes, averaged over five ranges of episodes, is shown in Fig. 6. The success rate rose steadily in the first 50 episodes and in the last 50 episodes (episodes 201–250) the success rate was 91.4%. The steady rise shows successful learning convergence and policy optimization, as well as a greater ability to make decisions during the reinforcement learning process.

A. Discussion

The results of the experiments prove that AIOps technologies enable intelligent automation, detection of anomalies, predictive analyses, and proactive management of incidents, which significantly enhance the modern IT process. Traditional techniques like Isolation Forest, clustering, and reinforcement learning boost service reliability and minimize downtime in cloud and enterprise environments. But numerous autonomous AIOps systems are still plagued by a myriad of issues such as false positive rates, lack of explainability, governance problems, scalability limitations, model drift, and the overreliance on historical operational data and automated decision-making. AIOps in the proposed framework comes with a combination of intelligent automation and expert human validation to overcome these shortcomings. The proposed approach of anomaly detection, using Isolation Forest algorithm and adaptive remediation using Deep Q-Network reinforcement learning, increases the ACC of anomaly identification and adaptive remediation. However, the framework might need high computational resources and frequent human inputs for its optimal performance in large-scale dynamic environments.

V. CONCLUSION AND FUTURE STUDY

The framework of a Human-in-the-Loop AIOps 2.0 was proposed to make intelligent and reliable IT operations

management possible by combining anomaly detection, reinforcement learning, and Human-supervised remediation. A Deep Q-Network reinforcement learning engine was used to create adaptive remediation recommendations, whereas for operational anomaly identification, the algorithm used was Isolation Forest. Autonomous decision-making processes enhanced explainability, governance, transparency and operational trust using human validation mechanisms. The results of the experiments showed the ACC, PRE, REC, F1, incident resolution success rate, and downtime reduction to be significantly better than the traditional monitoring and fully autonomous AIOps systems. The results highlight the benefits of seamlessly integrating AI with human expertise to enhance operational efficiency, reduce false alarms, and mitigate system risks in large-scale enterprise environments. However, there are some drawbacks, such as relying on historical operational data, the complexity of the computations during reinforcement learning training, and scalability issues in very distributed infrastructures. Future research involves the development of more scalable, transparent and resilient next-generation autonomous IT operations systems for explainable AI, federated cloud-edge learning, generative AI-assisted remediation, adaptive cybersecurity intelligence and self-healing infrastructure mechanisms.

REFERENCES

[1] S. Tatineni, "AIOps in Cloud-native DevOps: IT Operations Management with Artificial Intelligence," *J. Artif. Intell. Cloud Comput.*, vol. 2, pp. 1–7, 2023, doi: 10.47363/JAICC/2023(2)154.

[2] Tyagi, "Intelligent DevOps: Harnessing artificial intelligence to revolutionize CI/CD pipelines and optimize software delivery lifecycles," J. Emerg. Technol. Innov. Res., vol. 8, pp. 367–385, 2021.

[3] G. C. Kakaraparthi, "Integrating serverless architectures and Kubernetes for scalable and high-availability AI workflows," *Int. J. Intell. Syst. Appl. Eng.*, vol. 12, no. 4, pp. 5896–5905, 2024, doi: 10.6084/m9.figshare.30445046.

[4] S. K. Malaraju and R. Bondalapati, "Least Outstanding Requests (LOR) Algorithm in Application Load Balancer," *Int. J. Sci. Technol.*, vol. 14, no. 3, p. 7, 2023, doi: 10.71097/ijsat.v14.i3.3171.

[5] D. P. Guda and C. Appani, "Graph Neural Networks for Dynamic Malware Behaviour Analysis and Classification in Advanced Persistent Threats (APT)," *Int. J. Commun. Networks Inf. Secur.*, vol. 14, no. 3, 2022.

[6] S. Chatterjee, "Disaster Recovery Plan in Utility Industry for Virtual Asset Management - A Comprehensive Overview to Avoid Cyber Attacks," *Int. J. Sci. Res.*, vol. 13, no. 12, pp. 1163–1171, 2024, doi: 10.21275/sr241217215432.

[7] S. Chatterjee, "A Data Governance Framework for Big Data Pipelines: Integrating Privacy, Security, and Quality in Multitenant Cloud Environments," *Tech. Int. J. Eng. Res.*, vol. 10, no. 5, 2023, doi: 10.56975/tijer.v10i5.158181.

[8] L. S. Surisetty, "Modernizing Legacy Systems with AI Orchestration: From Monoliths to Autonomous Micro services," *Int. J. Adv. Res. Comput. Sci. \& Technol.*, vol. 5, no. 6, pp. 7299–7306, 2022.

[9] M. R. Mohammed, "Enhancing the Reliability of Cloud-Based Software Systems Using AI-Driven Fault Prediction and Auto-Remediation Techniques," *Am. Int. J. Comput. Sci. Technol.*, vol. 3, no. 5, pp. 1–13, 2021.

[10] Patel, "A Review of Multi-Channel CRM Strategies Using Big Data and Cloud Integration," Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol., vol. 8, no. 1, pp. 577–588, 2022.

[11] Parupalli and H. Kali, "An In-Depth Review of Cost Optimization Tactics in Multi-Cloud Frameworks," Int. J. Adv. Res. Sci. Commun. Technol., vol. 3, no. 5, pp. 1043–1052, Jun. 2023, doi: 10.48175/IJARSCT-11937Q.

[12] Methuku, S. Kamatala, P. Naayini, and P. R. Vontela, "From Ethical Principles to Technical Safeguards: A Unified Framework for Safe and Human-Centred Artificial Intelligence," *Am. Int. J. Comput. Sci. Technol.*, vol. 4, no. 5, pp. 26–34, Sep. 2022, doi: 10.63282/3117-5481/AIJCST-V4I5P103.

[13] S. Chitta, C. Ravi, V. K. R. Vangoor, and S. M. Yellepeddi, "AIOps: Integrating AI and Machine Learning into IT Operations," *Aust. J. Mach. Learn. Res. Appl.*, vol. 4, no. 1, pp. 288--305, Jan. 2024.

[14] H. Allam, "From Monitoring to Understanding: AIOps for Dynamic Infrastructure," *Int. J. AI, BigData, Comput. Manag. Stud.*, vol. 4, no. 2, pp. 77–86, jun. 2023, doi: 10.63282/3050-9416.IJAIBDCMS-V4I2P109.

[15] S. Polisetty, "Training Ai Models: Preparing and Managing Ai Algorithms for Aiops," *Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol.*, vol. 9, no. 5, 2023.

[16] S. M. Polisetty, "Resolving Incidents and Alerts in AIOps with Predictive Analytics," *Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol.*, vol. 8, no. 5, pp. 375–387, 2022, doi: 10.32628/CSEIT23902182.

[17] L. Li, "The Application of AI-Based Technology in Computer Network Operation and Maintenance," *Mob. Inf. Syst.*, vol. 2022, no. 1, p. 2971393, 2022, doi: 10.1155/2022/2971393.

[18] S. Becker, F. Schmidt, A. Gulenko, A. Acker, and O. Kao, "Towards AIOps in Edge Computing Environments," *Publ. 2020 IEEE Int. Conf. Big Data (Big Data)*, 2021, doi: 10.1109/BigData50022.2020.9378038.

[19] S. I. Shifat, "IT Incident Log Dataset," *Kaggle.*, 2020, [Online]. Available: https://www.kaggle.com/datasets/shamiulislamshifat/it-incident-log-dataset

[20] Ampountolas, "Comparative Analysis of Machine

Learning, Hybrid, and Deep Learning Forecasting Models: Evidence from European Financial Markets and Bitcoins," Forecasting, vol. 5, no. 2, pp. 472–486, 2023, doi: 10.3390/forecast5020026.

[21] D. Song, A. M. Chung Baek, and N. Kim, "Forecasting Stock Market Indices Using Padding-Based Fourier Transform Denoising and Time Series Deep Learning Models," *IEEE Access*, vol. 9, pp. 83786–83796, 2021, doi: 10.1109/ACCESS.2021.3086537.

[22] M. Jin *et al.*, "An Anomaly Detection Algorithm for Microservice Architecture Based on Robust Principal Component Analysis," in *IEEE Access*, IEEE, 2020, pp. 226397–226408. doi: 10.1109/ACCESS.2020.3044610.

[23] M. Kayakuş, F. Yiğit Açikgöz, M. N. Dinca, and O. Kabas, "Sustainable Brand Reputation: Evaluation of iPhone Customer Reviews with Machine Learning and Sentiment Analysis," *Sustainability*, vol. 16, no. 14, 2024, doi: 10.3390/su16146121.