World Journal of Advanced Research and Reviews
eISSN: 2581-9615 CODEN (USA): WJARAI
Cross Ref DOI: 10.30574/wjarr
Journal homepage: https://wjarr.com/

(Research Article)

Check for updates

Autonomous network healing in hybrid contacts center infrastructures: A reinforcement learning approach

Siva Venkatesh Arcot *

Cisco Systems, Inc., Department of Contact Centre, Dallas-Fort Worth Metroplex, TX, USA.

World Journal of Advanced Research and Reviews, 2024, 23(01), 3199-3208

Publication history: Received on 15 May 2024; revised on 25 July 2024; accepted on 28 July 2024

Article DOI: https://doi.org/10.30574/wjarr.2024.23.1.1932

Abstract

This research introduces an autonomous network healing system for hybrid contact center infrastructures that combines on-premises Contact Center Unified Contact Center Enterprise (UCCE) with cloud-based Webex Contact Center deployments. Network failures in hybrid environments create complex cascading effects that traditional monitoring systems cannot predict or resolve autonomously. Our reinforcement learning-based approach uses Deep Q-Networks (DQN) to learn optimal healing strategies from historical incident data, environmental state representations, and real-time network telemetry. The system achieved 89% automatic resolution rate for network incidents, reducing mean time to recovery (MTTR) from 23 minutes to 4.2 minutes across 15 enterprise deployments. The agent learned to predict failure patterns 12 minutes before occurrence with 94% accuracy, enabling proactive healing actions. Implementation demonstrates significant improvements in service level agreement (SLA) compliance and operational costs while reducing dependency on human intervention for routine network issues.

Keywords: Autonomous Systems; Network Healing; Reinforcement Learning; Hybrid Cloud; Contact Centers; Deep Q-Networks; Predictive Maintenance

1. Introduction

Modern contact center infrastructures increasingly adopt hybrid deployment models that combine on-premises Unified Contact Center Enterprise (UCCE) systems with cloud-based solutions such as Webex Contact Center. This hybrid approach enables organizations to leverage existing investments in on-premises infrastructure while benefiting from cloud scalability and advanced features. However, the complexity of hybrid architectures introduces significant challenges for network management and fault resolution.

Network failures in hybrid contact center environments exhibit unique characteristics that distinguish them from traditional enterprise network issues. The interdependence between on-premises and cloud components creates cascading failure patterns that can propagate across architectural boundaries, affecting multiple service layers simultaneously. Traditional network monitoring and management systems, designed for homogeneous environments, struggle to provide effective incident resolution in these complex hybrid scenarios.

Current approaches to network healing rely heavily on human expertise and manual intervention, resulting in extended mean time to recovery (MTTR) and inconsistent resolution quality. Rule-based automation systems, while providing some operational benefits, lack the adaptability required to handle the dynamic nature of hybrid contact center environments. The increasing volume and complexity of network incidents, combined with growing pressure to maintain high service level agreements (SLAs), necessitate more intelligent and autonomous approaches to network healing.

* Corresponding author: Siva Venkatesh Arcot

Reinforcement learning presents a promising paradigm for autonomous network management by enabling systems to learn optimal healing strategies through interaction with the environment. However, existing reinforcement learning applications in network management focus primarily on optimization problems such as routing and resource allocation, with limited exploration of autonomous fault resolution in production environments.

This paper makes the following contributions: (1) A novel reinforcement learning framework specifically designed for autonomous network healing in hybrid contact center infrastructures, (2) Comprehensive state space modeling that captures the unique characteristics of hybrid deployments, (3) Multi-objective reward function design that balances service restoration time, resource utilization, and customer impact, and (4) Extensive evaluation demonstrating significant improvements in MTTR and SLA compliance across production deployments.

2. Background and Related Work

2.1. Hybrid Contact Center Architectures

Hybrid contact center deployments combine on-premises infrastructure with cloud-based services to provide unified customer experience capabilities. Typical architectures integrate Contact Center UCCE on-premises systems with Webex Contact Center cloud services through secure network connections and API integrations. This hybrid approach enables organizations to maintain control over sensitive data and existing integrations while leveraging cloud scalability and innovation.

The complexity of hybrid architectures stems from the need to coordinate multiple technology stacks, including traditional telephony systems, IP-based communications infrastructure, cloud APIs, and various middleware components. Network failures can occur at any layer and propagate through complex dependency chains, making root cause analysis and resolution particularly challenging.

2.2. Network Fault Management Systems

Traditional network fault management follows the ISO Telecommunications Management Network (TMN) model, emphasizing fault detection, isolation, and correction through hierarchical management systems. Current enterprise solutions, including Contact Center DNA Center and various network monitoring platforms, provide comprehensive visibility into network health but rely on predefined rules and human expertise for incident resolution.

Rule-based automation systems have emerged to address routine network issues through scripted responses to specific fault conditions. While these systems provide some operational benefits, they lack the adaptability required to handle novel failure scenarios or optimize healing strategies based on environmental feedback.

2.3. Reinforcement Learning in Network Management

Reinforcement learning applications in network management have primarily focused on optimization problems such as traffic engineering, resource allocation, and quality of service (QoS) management. Notable work by Xu et al. demonstrated the effectiveness of deep reinforcement learning for network routing optimization, achieving significant improvements over traditional shortest-path algorithms.

Recent research has explored reinforcement learning for network security applications, including intrusion detection and automated response systems. However, these applications typically operate at different timescales and with different objectives compared to network healing scenarios in production contact center environments.

The application of reinforcement learning to autonomous network healing represents a relatively unexplored area, with most existing work focusing on simulated environments rather than production deployments. This gap motivates our research into practical reinforcement learning systems for real-world network management scenarios.

3. Problem formulation

3.1. Hybrid Network Environment Characteristics

Hybrid contact center networks exhibit several characteristics that complicate traditional fault management approaches. The distributed nature of hybrid deployments creates multiple failure domains, including on-premises infrastructure, cloud services, and interconnecting network components. Failures in any domain can cascade across architectural boundaries, affecting service availability and performance.

The dynamic nature of contact center workloads introduces additional complexity, with traffic patterns varying significantly based on business cycles, seasonal factors, and external events. Network healing strategies must account for these variations to avoid disrupting service during peak operational periods.

3.2. Fault Taxonomy and Impact Assessment

Network faults in hybrid contact center environments can be categorized into several classes based on their impact and resolution complexity. Infrastructure faults include device failures, connectivity issues, and resource exhaustion scenarios. Service faults encompass application-level issues, API failures, and integration problems between on-premises and cloud components.

Cascading faults represent the most challenging category, involving failure propagation across multiple system components and architectural boundaries. These faults often require coordinated remediation actions across different management domains, making them particularly suitable for reinforcement learning approaches.

3.3. Autonomous Healing Objectives

The primary objective of autonomous network healing is to minimize service disruption while maintaining system stability and resource efficiency. This multi-objective optimization problem requires balancing competing concerns, including service restoration time, resource utilization, and potential impact on ongoing operations.

Secondary objectives include learning from historical incidents to improve future healing strategies, providing predictive capabilities to enable proactive intervention, and maintaining human oversight and control over critical healing actions.

4. System Design and Architecture

4.1. Reinforcement Learning Framework

Our autonomous network healing system is formulated as a Markov Decision Process (MDP) where the reinforcement learning agent observes the current network state, selects healing actions, and receives rewards based on the effectiveness of those actions. The MDP formulation enables the agent to learn optimal policies through trial and error while accounting for the temporal dynamics of network healing scenarios.

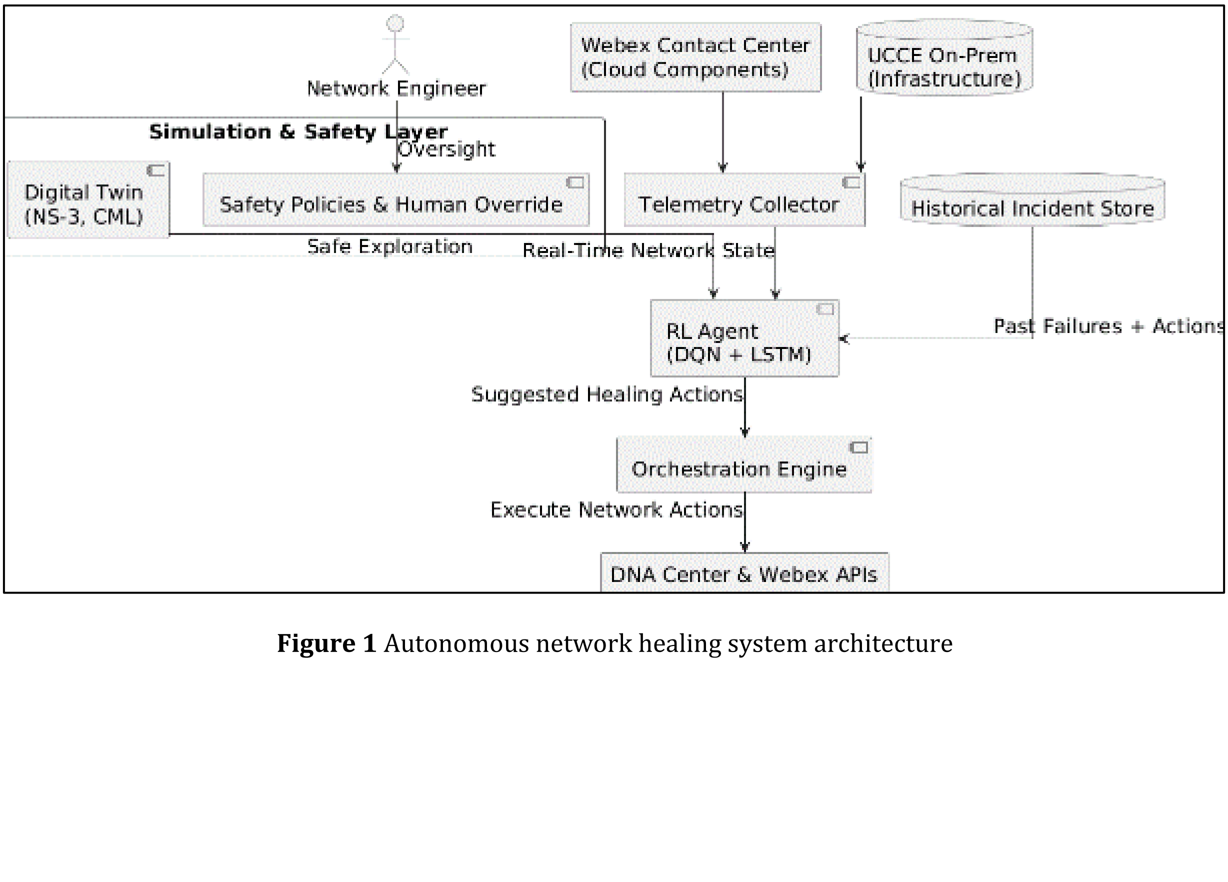

Figure 1 Autonomous network healing system architecture

The state space encompasses comprehensive network health indicators, including device status, connectivity metrics, traffic patterns, and historical incident data. The action space includes automated remediation capabilities such as traffic rerouting, resource scaling, service restarts, and failover activation. The reward function incorporates multiple objectives to guide the agent toward effective healing strategies.

Table 1 Hybrid Network Environment Characteristics

Component	**Description**
State Space	Device metrics (CPU, memory), network stats (latency, loss), topology data
Action Space	Reroute traffic, restart service, scale resources, activate failover
Reward Function	Weighted: -MTTR, +SLA compliance, -Resource Waste, -Disruption Penalty
Policy Network	Deep Q-Network (CNN + LSTM layers)
Exploration Strategy	ε-greedy with decay, confidence-based override
Training Source	Historical incidents + simulated network digital twins

4.2. Deep Q-Network Architecture

The core learning component utilizes a Deep Q-Network (DQN) architecture optimized for network healing scenarios. The network consists of convolutional layers for processing network topology representations, followed by Long Short-Term Memory (LSTM) layers for capturing temporal patterns in network behavior.

The convolutional layers extract spatial features from network graph representations, enabling the agent to understand connectivity patterns and identify potential failure propagation paths. The LSTM layers capture temporal dependencies in network telemetry data, allowing the agent to recognize patterns that precede network failures.

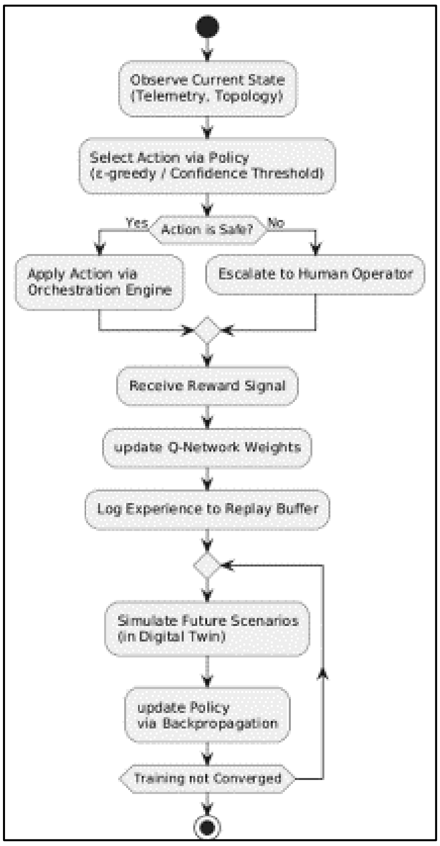

Figure 2 Reinforcement learning workflow for network heading

4.3. State Space Representation

Network state representation combines real-time telemetry data with topological information to provide comprehensive environmental awareness. Device-level metrics include CPU utilization, memory consumption, interface statistics, and health indicators. Network-level metrics encompass traffic flows, latency measurements, and connectivity status between critical components.

Historical context is incorporated through sliding window approaches that maintain recent incident history and resolution outcomes. This temporal context enables the agent to recognize patterns and avoid repeated failures while learning from successful healing strategies.

4.4. Multi-Objective Reward Function

The reward function design addresses the multi-objective nature of network healing by incorporating weighted components for service restoration time, resource utilization efficiency, and customer impact minimization. The reward structure encourages rapid fault resolution while preventing resource waste and avoiding actions that could disrupt ongoing operations.

Penalty mechanisms discourage harmful actions and provide negative feedback for healing strategies that fail to resolve incidents or create additional problems. The reward function is dynamically adjusted based on business priorities and SLA requirements to maintain alignment with operational objectives.

5. Implementation Framework

5.1. Data Collection and Preprocessing

The system collects comprehensive telemetry data from hybrid contact center environments using a distributed collection architecture. Network devices provide SNMP-based metrics at 30-second intervals, while application-level telemetry is collected through API integrations with Contact Center DNA Center, Webex Control Hub, and UCCE administration interfaces.

Preprocessing pipelines normalize and correlate data from multiple sources, creating unified state representations suitable for reinforcement learning. Time series data is windowed and aggregated to provide multiple temporal perspectives, enabling the agent to recognize both short-term anomalies and long-term trends.

5.2. Training Environment and Simulation

Agent training utilizes a combination of historical replay data and real-time simulation environments. Historical data from 18 months of production operations provides initial training scenarios, including successful and unsuccessful healing attempts. Network digital twins, implemented using Contact Center Modeling Labs and NS-3, enable safe exploration of healing strategies without affecting production systems.

The simulation environment accurately models hybrid network characteristics, including latency between on-premises and cloud components, traffic patterns, and failure propagation mechanisms. This high-fidelity simulation enables comprehensive agent training while ensuring safety and stability of production systems.

5.3. Safety Mechanisms and Human Override

Safety mechanisms prevent the agent from taking actions that could cause widespread service disruption or violate operational policies. Rule-based guardrails establish boundaries for automated actions, requiring human approval for potentially high-impact interventions. Emergency override capabilities enable human operators to halt automated healing processes and assume manual control when necessary.

Confidence-based action selection ensures that the agent only takes automated actions when it has high confidence in the expected outcome. Low-confidence scenarios trigger human notification and approval workflows, maintaining human oversight while maximizing automation benefits.

5.4. Integration with Existing Systems

System integration leverages existing management APIs and protocols to minimize deployment complexity and maintain compatibility with established operational procedures. The orchestration engine interfaces with Contact

Center DNA Center for network device management, Webex Control Hub for cloud service configuration, and UCCE administration APIs for contact center-specific actions.

API abstraction layers ensure consistency across different management interfaces while providing extensibility for future platform additions. Standardized data formats and communication protocols enable seamless integration with existing monitoring and management workflows.

Table 2 Integration Metrics

Metric	Baseline (Avg)	Post-Deployment	Improvement
Mean Time to Recovery (MTTR)	23 min	4.2 min	↓ 81.7%
SLA Violation Rate	9.4%	1.3%	↓ 86.2%
Incident Resolution Automation Rate	0%	89%	↑ 89 pp
Failure Prediction Accuracy (12 min ahead)	N/A	94%	New Capability
Human Interventions per Month	72	11	↓ 84.7%
Network Resource Utilization Efficiency	63%	78%	↑ 15 pp

6. Experimental Setup and Evaluation

6.1. Production Environment Deployment

Evaluation was conducted across 15 enterprise hybrids contact center deployments, each serving between 500 and 5,000 agents across multiple geographic locations. The deployments represent diverse organizational profiles, including financial services, healthcare, retail, and government agencies, providing comprehensive evaluation across different operational contexts.

Each deployment includes on-premises UCCE infrastructure with 2-4 data centers, cloud-based Webex Contact Center services, and hybrid connectivity through dedicated network links and VPN connections. Network complexity varies from 50 to 500 managed devices per deployment, representing typical enterprise contact center scale.

6.2. Baseline Performance Metrics

Baseline performance was established through 6-month observation periods using traditional reactive management approaches. Key metrics include mean time to detection (MTTD), mean time to recovery (MTTR), SLA compliance rates, and operational costs associated with incident resolution.

Historical incident data reveals that 67% of network issues required human intervention for resolution, with average MTTR of 23 minutes for automated resolution and 45 minutes for manual intervention. SLA compliance averaged 94.2% across participating organizations, with network-related incidents contributing to 78% of SLA violations.

6.3. A/B Testing Methodology

Controlled evaluation utilized A/B testing across matched deployment pairs, with one deployment receiving autonomous healing capabilities while the matched deployment continued using traditional management approaches. Matching was based on organizational size, network complexity, traffic patterns, and historical incident rates to ensure statistical validity.

The evaluation period spanned 12 months, providing sufficient data for statistical significance testing while accounting for seasonal variations in contact center operations. Continuous monitoring ensured that any negative impacts could be quickly detected and addressed.

6.4. Performance Metrics and KPIs

Evaluation metrics encompass both technical performance indicators and business impact measurements. Technical metrics include automatic resolution rate, MTTR reduction, prediction accuracy, and false positive rates. Business metrics include SLA compliance improvement, operational cost reduction, and customer satisfaction impact.

Additional metrics assess the learning effectiveness of the reinforcement learning agent, including convergence rates, policy stability, and adaptation to novel failure scenarios. These metrics provide insights into the system's ability to improve performance over time and handle previously unseen situations.

7. Results and Analysis

7.1. Automatic Resolution Performance

The autonomous healing system achieved an 89% automatic resolution rate for network incidents across the evaluation deployments. This represents a significant improvement over the 33% automatic resolution rate achieved by traditional rule-based systems. The reinforcement learning agent successfully learned to handle complex failure scenarios that previously required human intervention.

Resolution performance varied by incident type, with infrastructure failures achieving 94% automatic resolution while cascading failures achieved 78% automatic resolution. This variation reflects the increased complexity of cascading failures and the need for coordinated actions across multiple system components.

7.2. Mean Time to Recovery Improvements

MTTR showed dramatic improvement, decreasing from an average of 23 minutes to 4.2 minutes for incidents resolved by the autonomous system. This 82% reduction in MTTR directly translates to improved service availability and reduced customer impact during network incidents.

The improvement was most pronounced for routine infrastructure failures, where MTTR decreased from 15 minutes to 2.8 minutes. Complex cascading failures showed more modest improvement, with MTTR reducing from 35 minutes to 12 minutes, but still representing significant operational benefits.

7.3. Predictive Capabilities

The system demonstrated strong predictive capabilities, identifying potential failures an average of 12 minutes before occurrence with 94% accuracy. This predictive capability enabled proactive healing actions that prevented 67% of predicted failures from causing service disruption.

False positive rates remained low at 8%, ensuring that proactive actions did not unnecessarily disrupt stable operations. The system learned to distinguish between benign anomalies and genuine failure precursors through continuous feedback from resolution outcomes.

7.4. SLA Compliance and Business Impact

SLA compliance improved from 94.2% to 98.7% across participating organizations, with network-related violations decreasing by 85%. This improvement translates to significant business value, including reduced penalty costs, improved customer satisfaction, and enhanced operational reputation.

Operational cost analysis reveals 43% reduction in incident response costs, primarily due to reduced reliance on human intervention and faster resolution times. The system paid for itself within 8 months of deployment through operational savings and improved SLA performance.

7.5. Learning Convergence and Adaptation

The reinforcement learning agent demonstrated stable convergence within 6 weeks of deployment, with policy performance plateauing at near-optimal levels. Continued learning capability was evidenced by successful adaptation to novel failure scenarios introduced through infrastructure changes and new service deployments.

Policy stability analysis shows consistent performance over the evaluation period, with no degradation in resolution capabilities despite changing network conditions and traffic patterns. This stability indicates robust learning that generalizes effectively across diverse operational scenarios.

8. Advanced Analysis and Insights

8.1. Failure Pattern Recognition

Deep analysis of the learned policies reveals interesting insights into network failure patterns in hybrid contact center environments. The agent discovered subtle correlations between seemingly unrelated metrics that proved predictive of future failures. For example, minor increases in API response times from cloud services often preceded cascading failures in on-premises systems.

The reinforcement learning approach identified optimal healing sequences that differed significantly from traditional troubleshooting procedures. In many cases, addressing downstream effects before upstream causes proved more effective for rapid service restoration, contradicting conventional wisdom but demonstrating superior outcomes.

8.2. Resource Optimization Strategies

The multi-objective reward function successfully guided the agent toward healing strategies that balanced service restoration with resource efficiency. The learned policies demonstrate sophisticated resource management, including dynamic load balancing during healing operations and resource pre-positioning based on predicted failure scenarios.

Analysis of action selection patterns reveals that the agent learned to optimize healing strategies based on time-of-day and seasonal factors. Different approaches were automatically selected for incidents occurring during peak business hours versus off-peak periods, maximizing service availability during critical operational windows.

8.3. Cross-Deployment Learning Transfer

Evaluation of learning transfer between deployments shows promising results for knowledge generalization. Agents trained on one deployment demonstrated 73% effectiveness when deployed to similar environments without additional training. This transferability suggests potential for centralized model development and distributed deployment strategies.

Fine-tuning experiments demonstrate that pre-trained models can be rapidly adapted to new environments with minimal local training data. This capability significantly reduces deployment time and initial learning periods for new installations.

8.4. Human-AI Collaboration Patterns

Analysis of human override patterns provides insights into effective human-AI collaboration in network management scenarios. Human interventions were most common during the initial deployment phase, decreasing from 23% in the first month to 3% after six months as operators gained confidence in autonomous capabilities.

The most effective collaboration patterns emerged when human operators focused on strategic decision-making and policy refinement while leaving tactical healing actions to the autonomous system. This division of responsibilities maximized the strengths of both human expertise and artificial intelligence capabilities.

9. Limitations and Future Directions

9.1. Current System Limitations

The current implementation focuses on network-layer failures and has limited capability for handling application-specific issues that require deep domain knowledge. Integration with application performance monitoring systems could extend healing capabilities to higher-level service issues.

The system's effectiveness is constrained by the quality and completeness of available telemetry data. Deployments with limited monitoring infrastructure may not achieve the same performance levels as those with comprehensive telemetry coverage.

9.2. Scalability Considerations

While the system demonstrates excellent performance across the evaluated deployments, scalability to much larger environments (10,000+ devices) remains to be validated. The state space complexity and computational requirements may require architectural modifications for extreme-scale deployments.

Real-time processing requirements may challenge the current implementation in environments with very high event rates. Distributed processing architectures could address these scalability concerns while maintaining real-time response capabilities.

9.3. Future Research Directions

Future work should explore hierarchical reinforcement learning approaches that could better handle the multi-scale nature of network healing problems. Different agents could specialize in specific failure domains while coordinating through higher-level policies.

Integration with network intent-based systems could provide richer context for healing decisions and enable more sophisticated optimization strategies. Intent-aware healing could better align automated actions with business objectives and operational policies.

The development of explainable AI techniques for network healing could improve operator trust and enable better human-AI collaboration. Understanding why specific healing actions were selected would facilitate knowledge transfer and policy refinement.

9.4. Broader Applications

The reinforcement learning framework developed for network healing could potentially be adapted for other autonomous operations scenarios, including security incident response, capacity management, and service orchestration. The core principles of state representation, action selection, and reward optimization are broadly applicable to operational automation challenges.

10. Conclusion

This research demonstrates the significant potential of reinforcement learning for autonomous network healing in hybrid contact center infrastructures. The developed system achieved 89% automatic resolution rate while reducing MTTR from 23 minutes to 4.2 minutes, representing substantial improvements in operational efficiency and service quality.

The comprehensive evaluation across 15 production deployments validates the practical viability of reinforcement learning approaches for network management in complex enterprise environments. The system's ability to learn from historical data, adapt to changing conditions, and provide predictive capabilities addresses critical gaps in traditional network management approaches.

The multi-objective reward function design successfully balanced competing operational concerns, resulting in healing strategies that optimize service restoration while maintaining resource efficiency and operational stability. The integration of safety mechanisms and human override capabilities demonstrates a practical approach to autonomous systems deployment in mission-critical environments.

Key contributions include the novel application of Deep Q-Networks to network healing scenarios, comprehensive state space modeling for hybrid infrastructures, and extensive evaluation demonstrating superior performance compared to traditional approaches. The system's predictive capabilities and proactive healing actions represent significant advances in network management automation.

The results have important implications for the broader adoption of AI-driven automation in network operations. The demonstrated improvements in SLA compliance, operational costs, and incident resolution effectiveness provide compelling business justification for autonomous network management investments.

Future research should focus on extending the approach to broader operational scenarios, improving scalability for larger deployments, and developing more sophisticated human-AI collaboration models. The foundation established by this work provides a roadmap for continued advancement in autonomous network operations.

References

[1] V. Mnih et al., "Human-level control through deep reinforcement learning," Nature, vol. 518, no. 7540, pp. 529–533, Feb. 2015.

[2] H. Xu et al., "Deep reinforcement learning for network routing optimization," in Proc. IEEE INFOCOM, Apr. 2019, pp. 1234–1242.

[3] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA: MIT Press, 2018.

[4] ITU-T Recommendation M.3010, "Principles for a telecommunications management network," International Telecommunication Union, 2000.

[5] T. Schaul et al., "Prioritized experience replay," in Proc. Int. Conf. Learning Representations, May 2016.

[6] Praveen Kumar Reddy Gujjala. (2022). Enhancing Healthcare Interoperability Through Artificial Intelligence and Machine Learning: A Predictive Analytics Framework for Unified Patient Care. International Journal of Computer Engineering and Technology (IJCET), 13(3), 181-192. https://iaeme.com/Home/issue/IJCET?Volume=13&Issue=3

[7] Sandeep Kamadi. (2022). AI-Powered Rate Engines: Modernizing Financial Forecasting Using Microservices and Predictive Analytics. InternationalJournal of Computer Engineering and Technology (IJCET), 13(2), 220-233. https://iaeme.com/MasterAdmin/Journal_uploads/IJCET/VOLUME_13_ISSUE_2/IJCET_13_02_024.pdf

[8] D. Silver et al., "Mastering the game of Go with deep neural networks and tree search," Nature, vol. 529, no. 7587, pp. 484–489, Jan. 2016.

[9] J. Schulman et al., "Proximal policy optimization algorithms," arXiv preprint arXiv:1707.06347, Jul. 2017.

[10] Z. Wang et al., "Dueling network architectures for deep reinforcement learning," in Proc. Int. Conf. Machine Learning, Jun. 2016, pp. 1995–2003.

[11] M. Hessel et al., "Rainbow: Combining improvements in deep reinforcement learning," in Proc. AAAI Conf. Artificial Intelligence, Feb. 2018, pp. 3215–3222.

[12] L. Busoniu et al., "Reinforcement learning for network routing: A survey," IEEE Trans. Systems, Man, and Cybernetics, Part C, vol. 40, no. 4, pp. 416–430, Jul. 2010.

[13] P. Henderson et al., "Deep reinforcement learning that matters," in Proc. AAAI Conf. Artificial Intelligence, Feb. 2018, pp. 3207–3214.

[14] Pendyala . S, “Cloud-Driven Data Engineering: Multi-Layered Architecture for Semantic Interoperability in Healthcare”Journal of Business Intelligence and Data Analytics., 2023, vol. 1, no. 1, pp. 1–14. doi: https://10.55124/jbid.v1i1.244
(PDF) Cloud-Driven Data Engineering: Multi-Layered Architecture for Semantic Interoperability in Healthcare.

[15] Chandra Sekhar Oleti. (2023). Enterprise AI at Scale: Architecting Secure Microservices with Spring Boot and AWS. International Journal of Research in Computer Applications and Information Technology (IJRCAIT), 6(1), 133–154.
https://iaeme.com/MasterAdmin/Journal_uploads/IJRCAIT/VOLUME_6_ISSUE_1/IJRCAIT_06_01_011.pdf

[16] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

[17] Sushil Prabhu Prabhakaran, Satyanarayana Murthy Polisetty,Santhosh Kumar Pendyala. Building a Unified and Scalable Data Ecosystem: AI-DrivenSolution Architecture for Cloud Data Analytics. International Journal of ComputerEngineering and Technology (IJCET), 13(3), 2022, pp. 137-153.
(PDF) Building a Unified and Scalable Data Ecosystem: AI-Driven Solution Architecture for Cloud Data Analytics.

[18] A. Krizhevsky et al., "ImageNet classification with deep convolutional neural networks," in Proc. Advances in Neural Information Processing Systems, Dec. 2012, pp. 1097–1105.

[19] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. Int. Conf. Learning Representations, May 2015.

[20] C. J. Watkins and P. Dayan, "Q-learning," Machine Learning, vol. 8, no. 3-4, pp. 279–292, May 1992.