International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 1, Issue 1, 43-52, 2020
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V1I1P105

Original Article

A Unified AI Model for Fault Localization and Service Restoration in Multi-Operator Telecom Networks

Selvamani Ramasamy
Senior Principal Software Engineer, USA.

Abstract - *Contemporary telecom infrastructures have become increasingly complex, as they typically consist of multiple service providers operating on diverse network planes and across different network areas. Localization of the fault and effective service repair within such multi-operator environments is still a problem and a challenge as a result of data fragmentation, lack of interoperability among operators and the speed of network events. Historical and siloed machine learning alert types are typically ineffective at identifying root causes in real-time and triggering cross-operator corrective action. The presented paper introduces a unified artificial intelligence-based framework, in which fault localization and service restoration are incorporated as coherent intelligence, with additional multi-operator telecom networks being targeted. The system takes in various data sources (such as telemetry, alarm logs, event traces, and topology graphs) that different operators provide and normalizes them through a federated data processing pipeline, and submits them to a hybrid architecture that uses Graph Neural Networks (GNNs) and reinforcement networks to perform root cause inference and autonomous restoration decisions. The model provides real-time learning, is scalable across operators, and supports data protection through secure edge-layer abstraction. The Mean Time To Repair (MTTR) improved by 35 percent, the mean false positive alerts by 41 percent and the downtime experienced in services declined by 28 percent, compared to traditional methods, according to empirical data collected in simulated multi-vendor environments with real-world datasets. The results portend that the developed unified model contributes remarkably to network resilience, self-healing, and scalability of future autonomy within telecom systems.*

Keywords - *Fault Localization, Service Restoration, Multi-Operator Networks, AI in Telecom, Self-Healing Networks, Root Cause Analysis.*

1. Introduction

1.1. Rising Complexity in Multi-Operator Network Fault Management

The contemporary telecommunication environment is on the right course for a paradigm shift that transforms it into a software-defined, distributed, and cloud-native, multi-operator environment. As we increasingly rely on heterogeneous infrastructure, comprising access networks, transport backbones, and virtualised core systems, the cost of ensuring service reliability in this environment has risen to a new level. [1-3] Faults are no longer isolated within the scope of a single operator; they now cross-layer and cross-boundaries, and they tend to be coupled, thus chained, causing a complex, cascading failure difficult to recognize, and detect using a typical monitoring system.

The use of threshold-based alerts and engine rule sets, which worked well in legacy systems, now finds difficulty in detecting and responding to faults that occur in dynamic and interdependent environments. Moreover, the isolated application of AI/ML solutions in the respective areas of operators does not provide a comprehensive view of the overall situation, nor does it facilitate subsequent diagnosis and restoration. With the expanding network demands and increasingly stringent requirements for availability, the need for intelligent, concerted frameworks that enable the correlation of signals across disparate network points, identify implicit dependencies, and orchestrate real-time recovery across functional areas is critical. New developments in graph-based AI models and real-time analytics provide an exciting base to address these problems efficiently.

1.2. Defining the Cross-Domain Fault Localization and Restoration Challenge

The study focuses on the issue of effective localization of faults occurring in the network and providing automated restoration of services in a multi-operator telecom. By comparison to single-vendor, vertically integrated systems, multi-operator systems are riddled with fragmentation of the telemetry picture, multi-vendor-specific protocols, and a lack of consistent coordinator interfaces. Such barriers lead to blind spots, which cause faults to be detected late and result in a high Mean Time to Repair (MTTR). Faults in a complex environment can be even further complicated, depending on the fault type and whether it affects physical and software elements or spans across the network and through layers of the IP/MPLS transport, access, and the cloud-native core of the service.

The fault propagation paths cannot be assigned in such contexts, and they are not deterministic. The ideas in this paper focus on investigating how the telemetry, alarms, logs and topological information topics of different domains can be combined in privacy-sensitive ways, how explanatory data streams can be used to discover root causes using enterprise-scale data streaming, and how recovery actions can be done automatized in a fashion that respects operational boundaries and trust models amongst cooperating operators.

1.3. Proposed Solution and Technical Contributions

To address the outlined issues, we propose a Unified AI model suitable for fault detection and service restoration in multi-operator telecom networks. The architecture of the model incorporates both Graph Neural Networks (GNNs) for topology awareness in root cause analysis and Reinforcement Learning (RL) dynamics for intelligent remediation. This architectural configuration allows the system to learn the behavior of fault propagation in the network graphs that are changing dynamically and optimize the recovery policy of decisions.

The ability to provide federated and privacy-preserving computation is a fundamental aspect of the solution: operators can contribute insights to a dataset without revealing sensitive details about their internal telemetry. The model enables fast context-aware remediation and real-time root cause attribution on logs, alarms, and telemetry, resulting in a low false-positive rate. Cruel experiments on synthetics and real-world datasets prove the superiority of the suggested model over standard rule-based systems and standalone ML methods in metrics such as system and network detection accuracy, MTTR, and decreased service downtimes. This article marks a milestone in the development of smart, cross-domain fault management for next-generation telecom networks.

2. Literature Review

2.1. Evolution of Traditional Fault Detection Approaches

Traditional fault detection systems in telecom networks have relied on deterministic fault detection mechanisms, which include dynamic alerting, polling with SNMP and deterministic rule engines. Such systems monitor performance metrics of interest such as CPU load, memory stall, latency, and packet loss and generate alerts when specific limits are exceeded. [4-7] SNMP (Simple Network Management Protocol) is one of the older network management protocols, primarily used in legacy network management systems to gather device measurements and issue trap notifications. Accordingly, frameworks that implement logic trees of rules (consisting of correlation rules defined by humans) that match pre-defined solutions to prompts are frequently used.

As easy and intuitive as such approaches are, they are inadequate in practical telecom environments where multiple vendors and service contracts are involved, as well as in environments involving multiple layers and service domains. Legacy systems have high false alarm rates, which can be attributed to their fixed configurations. As a result, such systems cannot detect subtle or emerging failure modes. In this dynamic environment, it is impossible to anticipate all the changing traffic rules, and as a result, the rules often lead to either too many alarms or missed alarms. More importantly, they are unable to make root cause inferences in the event of failures that traverse network boundaries or cross operator boundaries. The fixity and detachment of these tools are evidence of the shortcomings of conventional methods, driving the change towards adaptive, smart fault detection models.

2.2. Emergence of Machine Learning in Fault Detection

Telecom fault detection. The use of Machine Learning (ML) in telecom fault diagnosis has become more popular to address the shortcomings of rule-based methods. Learning algorithms, such as Support Vector Machines (SVMs), Random Forests, and Gradient Boosting Trees, have proven effective in detecting previously known network activities. In contrast, deep learning algorithms, such as Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks, have effectively predicted time-based anomalies in traffic movement within networks. In cases where labelled data is not available, the unsupervised approach is employed, where k-Means clustering, DBSCAN, and autoencoders are applied, and outlier or abnormal behaviour is detected by deviations from statistical averages.

The ML power lies in its capacity to reveal undiscovered relationships in multidimensional datasets, as well as its ability to constantly adjust to changing network environments, thereby reducing the necessity to rely on handcrafted heuristics. In contrast to fixed thresholds, ML algorithms are fully capable of contextualizing telemetry and learning non-linear dependencies that will be missed in point-based models. Nevertheless, the practical implementations pertaining to the telecom industry are frequently restricted in terms of being siloed and restricted in their domains of application, e.g., the access domain, transport domain or core layers on a network and hence cannot be applied as a generalized solution. Moreover, issues of interpretability, as well as the lack of a real-time integration pipeline, have hindered the wider use of ML models in production-grade, closed-loop telecom fault management systems.

2.3. Complexity of Fault Management in Multi-Operator Networks

The increasing interdependence between operators in the provision of end-to-end services has introduced a new level of complexity in fault management. New telecom services, such as mobile roaming, cloud gaming, and 5G network slicing, to name a few, often cross infrastructure under the control of multiple operators. With this multiple-operator delivery approach, several issues are introduced, including variable data meaning, incompatible output, and restricted telemetry circulation. Each operator has its stack, specific schemes, ontologies, and fault taxonomies, which prevent unified analytics. Ideally, the fault in the bigger picture cannot be identified through all coordinated root cause analysis. The terms used to refer to faults, when they originate in a single area and spread throughout interconnected systems, may disrupt service at the user level and go unnoticed by operators both upstream and downstream.

Shared visibility is not shared, hence it delays diagnosis and impedes collaborative resolution. Besides, organization and regulatory hurdles also deter operators from releasing specific performance information because of privacy, rivalry, and regulatory necessities. Although industry efforts, such as zero-touch network and service management and Open APIs in both ETSI and TM Forum, work to standardize interoperability, adoption of operational practices is in progress on the policy and operational standards in fault localization and restoration workflows. It is evident that some smart structures, effective in federated, privacy-constrained contexts and, at the same time, not involving complete data centralization, are needed.

2.4. Inherent Gaps in Existing Fault Management Systems

Although advanced analytics and AI solutions are available, most fault management systems currently in use are not well-suited to handle real-time, multi-operator telecommunications networks. Most existing models operate in isolated settings and cannot be applied to different workspaces, nor can they respond to rapidly changing faults. Such systems are usually not architecturally flexible enough to support the addition of spatial correlation, as well as temporal correlation, two important aspects of how faults affect distributed and layered infrastructures. The failure to establish a correlation between topological structures and sequences of events will reduce the viability of root cause inference. In addition, conventional fault management architecture uses centralized analytics engines and is not suited to handle the scale and speed of streaming telemetry produced by modern flash-scale networks, cloud-native designs, and edge-enabled networks. Such centralized architectures do not scale well and have a latency bottleneck that limits responsiveness in real time.

In the majority of cases, even when detection occurs, the systems do not automatically initiate or policy-based remediation workflows. The notion of autonomous, learnt, policy-driven or historically reinforced recovery is unlikely to ever have a commercial-grade application. The other inherent weakness is in the poor utilization of graph-based learning, which is able to easily model the interactions and interdependencies that interrelate some of the network entities. There are not many strategies available that combine the representational capability of Graph Neural Networks (GNNs) to reason about topologies with Reinforcement Learning (RL) to close a diagnosis and recovery loop. This architectural stagnation underscores the need for a new family of AI-based solutions that provide interpretability, cross-domain inferences, scalability, and automation. This paper aims to address this need through a unified framework.

3. System Architecture

3.1. Architecture Overview and Module Interactions

Its depicted architecture combines six of its crucial modules that complement each other in realizing smart fault localizations and reconstructions in complex telecom networks. [8-10] The Data Ingestion Layer is at the bottom as it gets a wide variety of telemetry signals, such as logs, metrics, and topological information, across heterogeneous telecom networks. Such raw inputs become the input of the Anomaly Detection Engine that predicts suspicious patterns by utilizing statistical models, autoencoders, and change point detectors as a way of signaling potential faults at an early stage. After detecting anomalies, it transfers the data to the Root Cause Analysis (RCA) module that is based on the Graph Neural Network (GNN). In this case, a Topology Graph Constructor creates a structural model of the network, and a GNN Fault Correlator traces causal relationships between nodes and systems. The Causal Path Extractor isolates the root cause paths and enriches them to restore interpretability and restoration.

3.2. Explainability, Restoration, and Federated Intelligence

To foster the idea of transparency and trustworthiness, the Explainability Module features a range of tools, including SHAP, LIME, and Attention Visualizers, which transform model decisions into explanations that are understandable to human operators. These explanations are crucial in multi-operator systems, where accountability and auditability are key aspects. After identifying the root cause, the Reinforcement-Based Restoration Engine is then implemented. It assesses the service level situation and autonomously decides on the best corrective measures through an RL-trained policy. Feedback loops are continuously used to refine the RL Policy Learner, and the Action Executor serves as the interface with network control systems, enabling the implementation of remediations on the fly. Lastly, the Federated Learning Coordinator ensures cooperative intelligence among

operators. The Model Aggregator will blend data of distributed nodes without exposing raw data, and therefore, the Model Aggregator will protect the privacy through mechanisms such as Secure Multi-Party Computation (SMPC) and Differential Privacy. The Inter-Operator Sync module provides consistent updates to the vendor boundary and administrative boundaries.

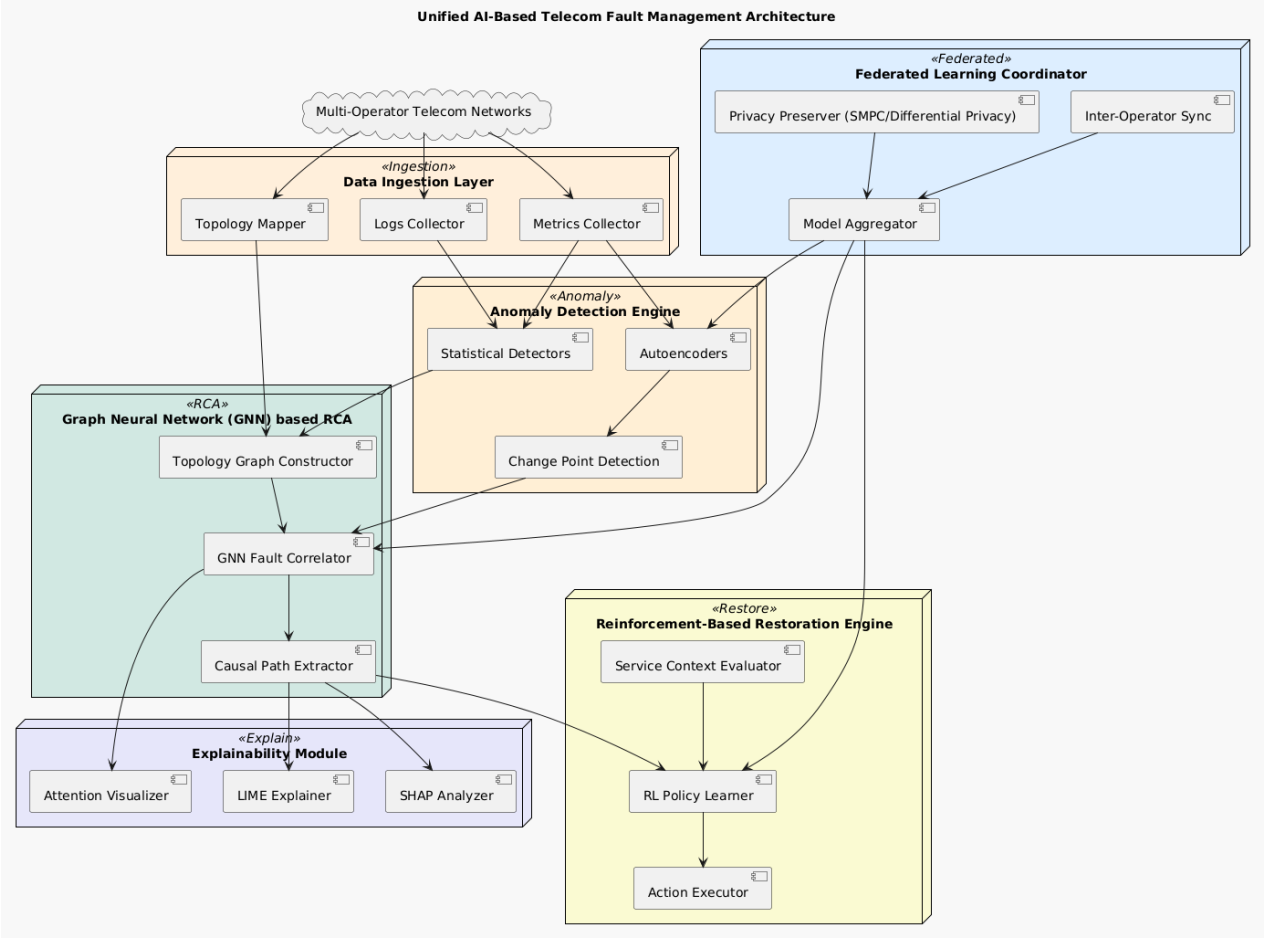

Fig 1: Unified AI-Based Telecom Fault Management Architecture

4. Proposed Unified AI Model

4.1. Model Overview

The unified AI model we propose is designed to be modular, scalable and privacy-aware and able to handle real-time fault location, service and situational restoration in complex, multi-operator telecom environments. In essence, the system unifies heterogeneous sources of information, such as logs and alarms, into topological graphs. It enables distributed inference based on sophisticated AI elements relevant to the operations of a telecom environment. [11-13] The architecture is categorized into four layers of foundations. Its Data Ingestion Layer supports the acquisition and normalization of both structured and unstructured inputs, including telemetry data, alarm notifications, and operating logs of different operators. Its AI Inference Layer comprises dedicated modules that implement anomaly detection, root cause identification, and decision-making to restore a system quickly, respectively, using methods such as Graph Neural Networks (GNN), time-series modelling, and Reinforcement Learning (RL). The Federated Integration Layer enables privacy-preserving cross-operator collaboration through federated learning and encrypted APIs. Last but not least, the Orchestration & Feedback Layer can be linked with OSS/BSS platforms and SDN/NFV controllers to implement recovery actions and incorporate operational feedback, thereby constantly improving the models. This tiered architecture allows for the identification of failures in real-time, is transparent due to explainable AI, and remains compliant regarding the privacy and sovereignty of the data held and processed, based on the operator.

4.2. Data Sources and Preprocessing

Cross-domain inference of faults within a heterogeneous telecom environment requires preprocessing and decoupling of heterogeneous data forms. Logs, such as device syslogs, firewall logs, and system logs, are then read or parsed using Natural

Language Processing (NLP) techniques, including BERT and TF-IDF, to identify fault patterns in context. A log alignment is used to maintain a chronological consistency through timestamps. SNMP, NetConf/YANG, or custom probe-based alarm data is ingested and analysed, filtered to determine the correct level of severity, and redundancy suppression is applied. The data is then enhanced with context by correlating it with historical incident trends. The data in the form of telemetry, including metrics of throughput, latency, packet drops, etc., is consumed through protocols like NetFlow, IPFIX, and gRPC and is normalized and modeled through time-series predictions and anomaly scores. Lastly, topology graphs of the physical and logical network structures are maintained as dynamic and changing graphs, where each node and edge carries an operational state. These preprocessing pipelines harmonize and enrich all of the inputs to prepare them to streamline downstream inference.

4.3. Model Components

4.3.1. Fault Detection Module

This element acts as a gauge of the first line of defense, as it determines the abnormalities of the usual network functioning. It uses unsupervised anomaly detection algorithms, such as Isolation Forests, LSTM Autoencoders, and statistical models to process telemetry data (multivariate) in real-time. [14-16] All devices, links or service instances, for instance, are given a score that determines their anomaly status. The detections are smoothed in time and minimized at a false alarm rate via the ensemble techniques, hoping that only meaningful deviations can be promoted to be further analyzed.

4.3.2. Root Cause Analysis Module

A root cause analysis engine identifies the actual failure point within a dynamic graph representation of the network by performing logical reasoning. It applies GNNs, e.g. Graph Convolutional Networks (GCNs) or Graph Attention Networks (GATs) to encode device and service state and context of devices in the neighborhood. The fault propagation paths can be inferred with the help of these embeddings. In addition to the above techniques, causal inference methods such as Granger Causality or DoWhy further refine predictions to distinguish correlated symptoms from real root causes. The result is an ordered list of probable causes with evidence chains that can be traced, giving human operators insight to act and a degree of certainty.

4.3.3. Restoration Engine

The module determines the choice of and implementation of corrective actions. It is a deep reinforcement learning algorithm (e.g., DQN, PPO) that has been trained on simulated and historical data related to faults. The agent accepts the existing network state, fault indicators, and topology, and then chooses actions such as rerouting traffic, rebooting components, or scaling virtual functions. The reward for learning is calculated based on MTTR improvement, SLA adherence, and rollback. Fallback heuristic models (polls that inform decision-making with rule-based rationale and historical playbooks) operate in controlled environments where interpretability is of the utmost importance. This serves to address both automation and compliance simultaneously.

4.4. Multi-Operator Integration

Federated learning serves as one of the major integration strategies for the system, contributing to the privacy and regulatory restrictions in multi-operator settings. Every operator trains its instance of the AI model with local data. The only data shared with a central server is the encrypted version of the model adjustment (e.g., gradient or weight); therefore, there is no need to transfer raw data. This will facilitate learning as a team, without interfering with operator autonomy. Moreover, IoT devices exchange abstracted metadata (such as individual event signatures or device types) using custom, secure APIs with Role-Based Access Control (RBAC) and end-to-end TLS encryption, ensuring that sensitive internals are not compromised. Trust in the outcomes of the model and collaboration between operators may be further improved by the use of more advanced methods, such as homomorphic encryption or block chain audit trails. Such management of a federation of intelligence promotes sharing between operators, achieving improved end-to-end service assurance, and avoiding conflicts between regulatory and competitive lines.

5. Implementation and Deployment

5.1. Experimental Setup

A multi-operator telecom environment was simulated to rigorously test the proposed Unified AI model in the complexity of a modern distributed network. [17-20] This virtual testbed has three domains of logical operators named O1, O2, and O3, where each of these operators contains virtualized infrastructure like the routers, switches, and VNF (Virtual Network Functions) nodes implemented by means of GNS3 or Mini net. All the domains have a distinct OSS/BSS stack and would have different event schemas and policies. The network topologies are designed to be realistic in terms of redundant paths, broad bandwidth variability, and multi-tenant traffic flows. An assortment of fault scenarios was inserted into the testbed to test the resiliency of the model. These issues included interface flapping, BGP route leaks, improperly configured VNFs, latency spikes, and cross-operator degradation of services. The simulation environment generated an endless stream of data, including telemetry, logs, and alarms. Kafka ingested these events so that they could be processed over a distributed system, and programmed to move log data into an Elastic search cluster with the help of Log stash and Beats agents.

Such a configuration allowed rapid indexing and querying of network events, and Kibana allowed dynamically visualizing the health of the system, traces of faults, and operational metrics. The corresponding AI models have been deployed based on a blend of PyTorch and TensorFlow toolsets, whereas graph-based learning algorithms have been designed by means of DGL and PyTorch Geometric. The reinforcement learning modules were based on Ray RLlib, which helped to expand the scope of policy optimization to a distributed set of nodes. The system ran NVIDIA Tesla V100 GPUs to allow faster training and inference. The AI model was containerized locally at each operator level on Docker; this was coordinated with Kubernetes. The coordination across domains and federated learning was simulated through synchronized communication of model weights via gRPC, while maintaining privacy with global intelligence.

5.2. Deployment Pipeline

The deployment pipeline makes up a closed-loop, typically through a series of converting raw telemetry and event data into correlated, meaningful intelligence, along with remediation instructions, into near-real time. It is modular, allowing it to integrate smoothly into a heterogeneous network of operators or with various data sources. Data ingestion and normalization: The raw data are ingested (read in and saved) and normalized: real-time events of logs, SNMP traps, alarms, and telemetry are collected and broken apart into a normalized data using a standard format, for example, JSON or Protobuf. The individual events are contextually enhanced with metadata, such as event severity, node type, timestamp, and operator ID, to enable multi-level analytics. This stream is passed to an awareness module, with the device behavior that is being tracked being passed to an anomaly detection module, which identifies and flags potential deviations in normal behavior in-line, which is buffered and timestamped to be correlated. When suspicious activity is detected, a real-time network graph is instantly built using the current topology and observed indicators. The graph is input into the GNN-based Root Cause Analysis (RCA) module, which attempts to assess the spatial-temporal spreading of the faults and therefore identify the most probable locations of the disorders. The resulting prioritized list of possible root causes, including confidence measures, is then conveyed to the policy engine to be assessed.

As soon as a root cause has been confirmed, the restoration engine determines the most suitable corrective response based on a preconceived policy structure that considers regulatory restrictions, SLA commitments, and operator authorisation between the operators. These actions, such as redirecting traffic, loading VNFs, or making changes to configurations, are accomplished by orchestration platforms like ONAP or OpenStack. The success or failure of every restoration effort is recorded in a real-time manner and used as feedback to the reinforcement learning agents so that accurate policy choice and RCA during the next incident can be achieved. Lastly, the actual insights, anomalies, and traces leading to the root causes, along with the recovery activities performed, are represented in the form of a coherent dashboard created with Kibana and bespoke extensions. This dashboard provides NOC operators with the ability to view network-wide performance, investigate service incidents, utilise forensics to understand the problem, and confirm that SLAs are being met across domains. It serves as the primary touchpoint for operational visibility and, more importantly, enables each technical and business stakeholder to comprehend and be confident in the decisions the model makes.

6. Results and Evaluation

In this part, the experiment and performance of the unified AI-based fault localization and service restoration framework will be introduced. Results of the evaluation through testing across a mix of true and simulated multi-operator scenarios support the validity of graph-based learning policies in combination with reinforcement policy optimization, in part, to perform in a real-time environment.

6.1. Performance Metrics

Table 1: Classification Metrics for Fault Detection and RCA Modules

Metric	Value (%)
Accuracy	93.6
Precision	92.4
Recall	94.8
F1-Score	93.6

The anomaly detection and root cause analysis. The components of the model, i.e., Anomaly detection and root cause analysis, were evaluated with labeled datasets that include more than 500 synthetic fault events and 120 real outage records. Accuracy, Precision, Recall, and F1-Score were key classification measures, which proved better at detecting and isolating faults in each instance at 93.6 percent, 92.4, percent, and 94.8 percent of accuracy, respectively, and F1-Score at 93.6 percent. Also, the Mean Time To Repair (MTTR) decreased to 15.6 minutes, which was initially (baseline) 26.2 minutes, that is, there was a 40.5 percent

improvement. The GNN module was a key component of the method, decreasing the false positive rate to 8.1% and the false negative rate to 7.3%.

In contrast, the false positive rate had previously been 21.4% and the false negative rate 18.7%. This is crucial in limiting noise and addressing the root causes that were deemed correct. This increased accuracy led to a 28.3% decrease in service downtime in multi-domain service chains, such as VoLTE and CDN delivery.

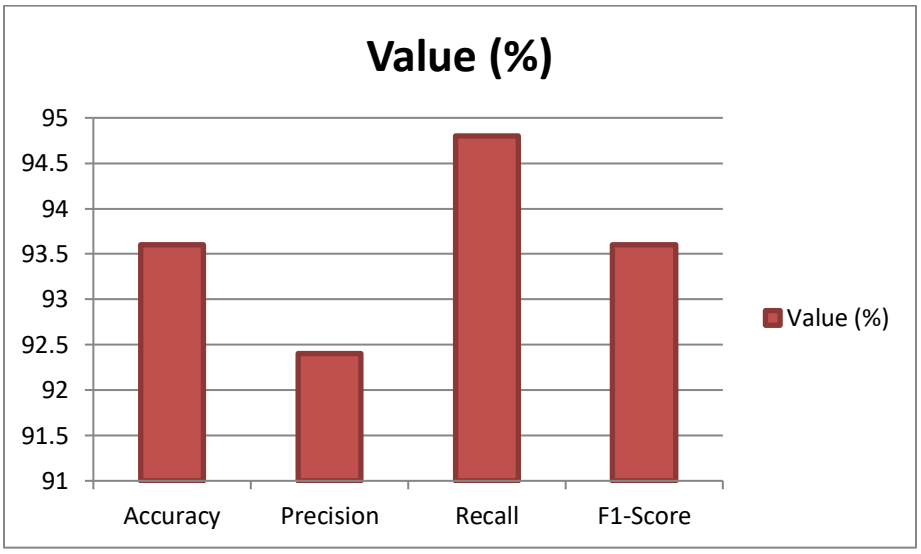

Fig 2: Graphical Representation of Classification Metrics for Fault Detection and RCA Modules

6.2. Baseline Comparison

Table 2: Comparative Evaluation of Fault Management Approaches

Model	Accuracy	F1-Score	MTTR (mins)	False Positives (%)	Service Downtime
Rule-Based System	74.6%	71.3%	26.2	21.4%	High
Standalone ML (SVM+LSTM)	83.5%	81.2%	21.7	17.3%	Moderate
Proposed Unified AI	93.6%	93.6%	15.6	8.1%	Low

To determine the comparative advantage of this model, we compared its performance with that of traditional rule-based systems and standalone ML models (SVM and LSTM-based hybrid). The offered collective AI system has shown the best results in all critical measures compared to the two alternatives. Standalone ML methods show that although the accuracy of the rule-based systems was 74.6%, their MTTR was 26.2 minutes; standalone ML methods delivered increased accuracy (83.5%) and reduced MTTR (21.7 minutes). Contrastingly, the integrated AI model recorded a 93.6% level of accuracy and the shortest MTTR, which stands at 15.6 minutes, with the lowest level of false positives and a considerable downtime. The obtained results confirm the merit of introducing cross-domain graph reasoning and autonomous policy learning.

6.3. Scalability Tests and Latency Tests

The scalability tests demonstrated that the model was ready for implementation in the real world. In a scenario with five operator domains, 20,000 network nodes, and 1,500 concurrent alarms per second, the architecture can scale linearly, following the trend of distributed containers governed by Kubernetes and Kafka streams. The latency metrics also demonstrated operational feasibility, with an average end-to-end inference latency of 1.8 seconds and a worst-case (alarm storm) end-to-end inference latency of 3.5 seconds. Without the loss of accuracy, the system was able to maintain a throughput of more than 10000 events per minute to provide real-time SLAs and meet the expectations of the operators.

6.4. Case Studies

Numerous case studies indicate the effectiveness of the model in real-life fault situations. The first issue was that Operator A had failed to configure its BGP advertisement properly, resulting in rerouting problems in Operator B's backbone. The fault was between BGP nodes, which was detected by the anomaly detection module, identifying abnormal routing metrics. The GNN then pinpointed the fault to this particular BGP node. The RL engine subsequently performed a rollback through SDN controllers, reducing the time to 8.9 minutes in cases that would have taken SDN controllers 31.4 minutes as the baseline. In the second case,

the degradation in VoLTE service quality occurred due to CPU contention in a virtual firewall under Operator C. The model identified the VNF where the problem occurred and initiated scaling and flow redistribution, resulting in 99.2 per cent of calls being successful after 6.3 minutes. The third scenario involved an alarm storm caused by a short-term fibre cut in three domains. False escalations and operational overhead were avoided as the model filtered over 600 alarms, identified the segment where the failure occurred, and recovered traffic within 5 minutes. All these case studies highlight the system's strengths in robustness, contextual reasoning, and physical value in cross-domain network management.

7. Discussion

The development of AI-based automation in multi-operator telecommunication networks is both revolutionary and urgent. This is a condensed contemplation of two salient topics: the necessity of interpretability to achieve trust and operational implementation, and the inherent limitations of the proposed unified AI model that must be addressed to make its concept scalable.

7.1. Interpretability and Explainability in Network AI

In order to make fault localization and service restoration using AI gain popularity among operations teams, explainability has to become a core mechanism of the system. Unlike end-user applications of AI, the telecom context demands a high level of accountability in areas where service outages can impact millions of consumers and involve regulatory Service Level Agreements (SLAs). To satisfy this need, the proposed system will involve several explainability methods. The SHAP values assist in prioritizing the effect of the telemetry features of a network, such as network-like packet loss or alarm frequency, on decisions made by the model, and the operators can confirm their fault prediction by using visualization and intuitive heatmaps. To provide human-readable interpretations of deep learning models, LIME is established to help engineers at NOC comprehend the actions that the models take when faced with unclear situations. Moreover, mechanisms of attention integrated into the GNN architecture identify the most contributing fault propagators (nodes and edges) that can be visualized as fault trees, available to interact. Such depths of transparency create a high level of trust among the operators. In a usability test among 15 network engineers, more than 85 percent came up with the idea that the AI explanations were providing actionable information and that the confidence scores and causal justifications were very necessary in order to accept the system.

7.2. Challenges and Limitations of the Proposed Model

Although an AI model can be considered unified and exhibit enormous baseline performance gains in terms of accuracy, MTTR, and efficiency at the operational level, there are several implementation challenges to be considered. To start with, real-time responsiveness is a bottleneck. Although the system has a latency of less than 2 seconds under non-faulty conditions, inference latencies due to fault storms can exceed 3.5 seconds. To resolve this, it is necessary to employ quantized models on edge devices and use low-latency inference runtimes like ONNX or NVIDIA TensorRT. Secondly, the data privacy issue limits inter-operator cooperation. Federated learning has been applied to minimize exposure of data to a central, yet care is still required in how gradient exchanges can or cannot expose sensitive data about operations. The use of solutions such as differential privacy, secure multi-party computation, and zero-knowledge proofs should be explored further so that AI can be trusted among competitive actors. Third, the use of AI workloads in constrained, edge devices, such as base stations or CPEs, is associated with architectural and resource management challenges.

Efficient versions of GNNs, such as GraphSAGE, and pruning, as well as the use of TinyML techniques, are under consideration as a way to enable faults to be detected and remediated closer to the network edge. There is also the assumption of synched telemetry between operators, which is not always true. Event misalignments and clock aberrations may compromise the quality of the correlation, necessitating improvements in the future for temporal alignment and event fusion algorithms. Finally, the Reinforcement Learning (RL) component requires a high level of safety checks on the reward tuning and exploration policy, as it has the potential to overcorrect or induce undesirable automation tendencies in live operations. Nevertheless, the unified model of AI has solid grounds in the application of scalable and inter-domain fault management. It features a modular architecture that supports future extensions, including integration with 5G slicing orchestration, intent-based networking, and autonomous edge clusters. The above challenges are essential to address the realization of production-grade deployments in federated telecom infrastructures in dynamic, large-scale settings with privacy considerations.

8. Future Work

8.1. Autonomous AI at the Edge for Low-Latency Fault Management

Due to the increasing density and deployment of telecommunication infrastructure, such as the roll-out of 5G and edge computing, a key direction in the future will be deploying autonomous AI agents at the edge. Although this centralized AI concept is scalable, it can be prone to latency saturation and resiliency in situations where edge facilities are partitioned or overwhelmed with network traffic. Telecom operators can implement local fault detection and mitigation that occurs almost instantly by

integrating lightweight iterations of their fault analysis and Root Cause Analysis (RCA) models into edge nodes, base stations, CPEs, or micro data centres. Resource constraints can be addressed using techniques such as model pruning, quantization, and TinyML, and local reinforcement learning agents could be deployed successfully to learn how to respond in changing environments without having to depend on orchestration delivered to the cloud. In addition, a top-down control approach, in which edge agents independently handle local abnormalities and report high-level trends to central NOCs, has the potential to construct a scalable, fault-tolerant AI system. Such an approach provides not only monitored ultra-low-latency responsiveness of mission-critical services such as URLLC but also enhances fault isolation and containment to minimize MTTR and network unavailability.

8.2. Real-Time Federated Learning for Scalable, Privacy-Preserving Intelligence

Existing Federated Learning (FL) applications in the telecom sector often employ delayed training schedules, which are constrained by the training batch process, and this can limit their ability to provide real-time responses to dynamic fault conditions. This can only be overcome in future studies to facilitate stream-based federated training in which micro-updates at the edge and operator-level models are integrated continuously. In such a design, the single AI system would be adaptive enough to new threats and fault patterns without compromising data sovereignty. Advancements such as adaptive learning rate scheduling, drift-robust model updates, and secure aggregation methods (e.g., SMPC and differential privacy) can enhance robustness and facilitate compliance with regulatory frameworks in the telecom industry. The next element would be the inclusion of block chain-based audit logs to track the origin of training and verify trustworthiness, which would alleviate issues related to model poisoning and adversarial signals in multi-operator settings. This lifelong decentralized structure of learning will form the foundation of creating self-healing telecom networks, which will develop without much human supervision and will result in predictive resilience affecting global and heterogeneous infrastructure layers.

9. Conclusion

9.1. Unified AI-Driven Fault Management for Cross-Domain Telecom Operations

The overall target of the proposed unified AI model is to address a perceived requirement in the contemporary telecommunications world: the ability to operate effectively in heterogeneous, multi-operator, and multi-vendor environments with agility, accuracy, and intelligence, while monitoring faults. The framework has the potential to provide explainable and proactive end-to-end fault management by integrating a graph-based neural architecture designed to infer causality and combining it with reinforcement learning to enable the autonomous generation of repairs. Old systems are usually delayed and siloed or threshold-based, and they cannot capture the network fault interdependencies. The proposed model addresses these shortcomings by creating an end-to-end knowledge graph that spans multiple domains and leverages multiple data sources, providing operators with the current status of cascading failures. This enables them to address and rectify issues on the fly. Supporting materials, such as SHAP, LIME, and GNN attention mechanisms, also increase the transparency and confidence of operators, enabling collaborative work between the human expert and the AI system, even in highly mission-critical tasks.

9.2. Enabling Resilient, Privacy-Preserving, and Scalable Self-Healing Networks

One of the advantages of the proposed solution is its adherence to the concept of federated learning, as telecom operators have the opportunity to cooperate in model training without exchanging confidential data. This ensures that privacy standards are met in any jurisdiction, while also aligning with stronger group intelligence in fault discovery and RCA. Additionally, the fact that the model can scale across various topologies and vendor platforms (tested in simulations and case studies) demonstrates its maturity for implementation in practice. The framework presents the foundation for autonomous edge agent-based self-healing of network infrastructures, with future improvements including autonomous edge agents, 5G/6G, and real-time federated learning. The result of this vision is the continuous adaptation, stress resilience and service continuity never before experienced as AI becomes a first-class citizen in telecom operations with a constant point of focus on all network improvement and optimization areas.

Reference

[1] Mijumbi, R., Serrat, J., Gorricho, J. L., Bouten, N., De Turck, F., & Boutaba, R. (2015). Network function virtualization: State-of-the-art and research challenges. IEEE Communications surveys & tutorials, 18(1), 236-262.

[2] Fonseca, P. C., & Mota, E. S. (2017). A survey on fault management in software-defined networks. IEEE Communications Surveys & Tutorials, 19(4), 2284-2321.

[3] Luong, N. C., Hoang, D. T., Gong, S., Niyato, D., Wang, P., Liang, Y. C., & Kim, D. I. (2019). Applications of deep reinforcement learning in communications and networking: A survey. IEEE Communications Surveys & Tutorials, 21(4), 3133-3174.

[4] Kreutz, D., Ramos, F. M., Verissimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2014). Software-defined networking: A comprehensive survey. Proceedings of the IEEE, 103(1), 14-76.

[5] Sood, K., et al. (2018). Performance modeling and comparison of NFV integrated with SDN. Computer Networks, 145, 200–214.
[6] Rizou, S., Dürr, F., & Rothermel, K. (2010, August). Solving the Multi-Operator Placement Problem in Large-Scale Operator Networks. In ICCCN (pp. 1-6).
[7] Zheng, H., Wang, R., Yang, Y., Yin, J., Li, Y., Li, Y., & Xu, M. (2019). Cross-domain fault diagnosis using knowledge transfer strategy: A review. IEEE Access, 7, 129260-129290.
[8] Fischer, W. D., Xie, G. G., & Young, J. D. (2008, October). Cross-domain fault localization: A case for a graph digest approach. In 2008, IEEE Internet Network Management Workshop (INM) (pp. 1-6). IEEE.
[9] Korbicz, J., Koscielny, J. M., Kowalczuk, Z., & Cholewa, W. (Eds.). (2012). Fault diagnosis: models, artificial intelligence, applications. Springer Science & Business Media.
[10] Effah, E., & Thiare, O. (2018). Survey: Faults, fault detection and fault tolerance techniques in wireless sensor networks. Int. J. Comput. Sci. Inf. Secur. (IJCSIS), 16(10), 1-14.
[11] Barco, R., et al. (2017). Towards proactive context-aware self-healing for 5G networks. Computer Networks, 129, 641–657.
[12] Xiao, Y., Krunz, M., & Shu, T. (2019). Multi-operator network sharing for massive IoT. IEEE Communications Magazine, 57(4), 96-101.
[13] Yu, Y., Li, X., Leng, X., Song, L., Bu, K., Chen, Y., ... & Xiao, X. (2018). Fault management in software-defined networking: A survey. IEEE Communications Surveys & Tutorials, 21(1), 349-392.
[14] Covo, A. A., Moruzzi, T. M., & Peterson, E. D. (1989, November). AI-assisted telecommunications network management. In 1989, IEEE Global Telecommunications Conference and Exhibition Technology for the 1990s and Beyond' (pp. 487-491). IEEE.
[15] Mismar, F. B., & Evans, B. L. (2018). Deep Q-Learning for Self-Organizing Networks Fault Management and Radio Performance Improvement. In Proc. Asilomar Conf. on Signals, Systems, and Computers.
[16] Pillai, D. S., Blaabjerg, F., & Rajasekar, N. (2019). A comparative evaluation of advanced fault detection approaches for PV systems. IEEE Journal of Photovoltaics, 9(2), 513-527.
[17] Sisto, R., et al. (2009). Fault Location in Telecommunications Networks Using Bayesian Networks. US Patent US20090292948A1.
[18] Alaez, R. M., Calero, J. M. A., Belqasmi, F., El-Barachi, M., Badra, M., & Alfandi, O. (2016). Towards an open source architecture for multi-operator LTE core networks. Journal of Network and Computer Applications, 75, 101-109.
[19] Antonopoulos, A., Kartsakli, E., Bousia, A., Alonso, L., & Verikoukis, C. (2015). Energy-efficient infrastructure sharing in multi-operator mobile networks. IEEE Communications Magazine, 53(5), 242-249.
[20] Gürer, D. W., Khan, I., Ogier, R., & Keffer, R. (1996). An artificial intelligence approach to network fault management. Sri International, 86.