A REAL-TIME FOG COMPUTING SAFETY FRAMEWORK FOR SMART INDUSTRIAL ROBOTICS

Ganesh Kotdiya[1], Prof. Yogesh Patidar[2]
Research Scholar, Department of Electronics and Computer Science, Samrat Vikramaditya Vishwavidyalaya, Ujjain, Madhya Pradesh, India[1]
Professor, Department of Electronics and Computer Science, Samrat Vikramaditya Vishwavidyalaya, Ujjain, Madhya Pradesh, India[2]
E-mail: kumar06353@gmail.com

ABSTRACT

This paper describes an empirical study on a safety layer in a fog-computing enabled Internet of Things (IoT) for multi-robot industrial environments that facilitates real-time hazard detection and risk mitigation. The architecture distributes computation into three tiers edge fog nodes, intermediate fog nodes, and cloud back-ends to gain sub-10 millisecond safety response times that are not possible with traditional cloud-only architectures. Over a nine month experimental period, 664 GB of validated operational data was generated across a heterogeneous sensor suite deployed over a 12-robot assembly cell (LiDAR, Force/Torque, RGB-D, Proximity, Inertial Measurement Unit (IMU), and vibration). A hybrid AI detection method combining Long Short-Term Memory (LSTM) networks and fuzzy inference obtains a hazard detection F1 score of 97.3%, a false alarm rate of 0.9% and a mean end-to-end safety response latency of 6.2 ms—an order of magnitude faster than the ISO 10218-1 recommended reaction time budgets. The ability of multivariate linear regression ($R2 = 0.926$, $F(5, 394) = 241.8$, $p < 0.001$) to select AI inference confidence as the most predictive feature of classifier accuracy and one-way ANOVA ($F(3, 476) = 892.4$, $p < 0.001$, $eta2 = 0.849$) to reveal statistically significant latency benefits of the proposed fog-AI architecture vs cloud only, edge only, and non-AI fog alternatives. The proposed system demonstrates production-grade safety infrastructure for Industry 4.0 deployment with a system uptime of 99.94% and a mean time between failures (MTBF) of 6,940 h.

Keywords: *Fog computing[1]; IoT safety layer[2]; industrial robots[3]; hazard detection[4]; LSTM neural network[5]; real-time latency[6]; Industry 4.0[7].*

1. INTRODUCTION

1.1 BACKGROUND AND MOTIVATION

The fourth industrial revolution, popularly known as Industry 4.0, is defined by the Internet of Things, advanced robotics, and cyber-physical systems deeply embedded in the manufacturing and production environments.

Industrial robots, traditionally locked behind rigid safety cages that physically isolated them from human workers, are increasingly deployed as collaborative and semi-collaborative systems working alongside humans in shared spaces. While allowing for unprecedented levels of production flexibility and human-robot collaboration, this paradigm shift also generates complex and dynamic safety problems that traditional hardwired safety systems crucially lack the capability to address, especially since they largely rely on the use of static perimeter guards, light curtains, and pre-programmed emergency stops [3]. In 2022, the global industrial robot installed base surpassed 3.5 million units18, with collaborative robot allocations growing 30% annually19, thereby increasing the demand for adaptive, intelligent and low latency safety architectures that can provide human workers protection against operational situations that may vary from one job site to another [4].

Fog computing, which is introduced by Bonomi et al. [1] Leveraging the capabilities of the network edge, fog computing [1], a distributed extension of cloud computing, provides a strong architectural foundation for answering the real-time processing needs of industrial safety systems. Fog computing addresses this restriction in cloud-centric IIoT by placing computational resources physically closer to IoT sensors and actuators, which greatly reduces the communication latency In industrial robot safety, for example, ISO 10218-1:2011 [30] includes one hazard category where emergency stop reaction times must not exceed 200 ms, and in practice, safety engineering ensures that such safety-related reactions are able to take place in an order of magnitude lower response time in order to account for deceleration of mechanical elements being considerably greater than this, thereby making the ability to process sensor data and initiate a protective response within 1–9 ms not only desirable but also necessary for operability. Thus, this combination of fog computing with AI-based anomaly detection and a truly holistic multi-modal sensor IoT infrastructure is both a technically reasonable and practically needed path forward for next generation industrial robot safety.

1.2 PROBLEM STATEMENT

Theoretically, fog-based IIoT safety architectures are attractive, but there remain several compelling empirical gaps in the literature. First, the majority of published fog computing safety systems for robotics are evaluated in simulation or smaller-scale laboratory experiments with a limited sensory scope, leaving them 'unverified' on the complex, high data-rate, real-world production scenarios [18]. Secondly, there have been no systematic benchmarks against prior art using standardized statistical methods for the use of multi-modal sensor fusion with fog-resident AI inference, which is required for robust anomaly detection across diverse hazard categories. Finally, from a quantitative viewpoint, there is very little in the research literature on the reliability and availability parameters (production-grade measures, such as MTBF and system uptime) of fog-based safety systems. This paper fills these three gaps based on an empirical study in a real industrial context over nine months, using multi-variate regression and ANOVA statistical methods to derive generalizable performance relationships from the dataset collected.

1.3 OBJECTIVES AND SCOPE

The key purpose of this research is to empirically design, implement and validate a three-tier fog-computing based IoT safety layer for a 12-unit collaborative robot assembly cell in a real production site. Specific objectives include: (i) the characterization of hardware and network performance of the three-tier fog architecture; (ii) the benchmarking of five hazard detection algorithms with respect to precision, recall, F1-score, latency and false-positive rate; (iii) the quantification of safety-response times for three representative fault scenarios as compliance with ISO 10218-1 limits; (iv) the comparison of system reliability and quality-of-service metrics across four architectures; (v) the application of a multivariate linear regression to identify the dominant predictors of safety-detection accuracy; and (vi) the application of a one-way ANOVA with post-hoc analysis to determine statistically significant latency differences among the four architectural configurations. This innovative study is performed under the constraints of a real production environment operationally representative of industrial deployment decision-making.

2. LITERATURE SURVEY

The roots of fog computing was defined by Bonomi et al. [1] LGS defined latency, location-awareness, and mobility support to be the key requirements of the new edge-of-network computing paradigm, which cannot be met by cloud computing due to its nature as a centrally orchestrated architecture for time-critical applications. Atzori et al. [1], IOT software aspects survey, the first comprehensive overview of IOT architectures and applications established the technical vernacular and system classification framework that still determines the organization of the early IOT field [2]. Next, Chiang and Zhang [6] outlined the research agenda for fog-IoT integration, and provided a detailed list of open problems in resource management, security, and real-time scheduling all of which directly motivate the safety application domain investigated in this work. Dastjerdi and Buyya [5] provided a practical study on fog computing deployment issues, providing motivation for hierarchical resource allocation schemes, which is reflected by the three-tier architecture of the proposed system.

One of the main works in the industrial IoT domain is Sisinni et al. [4] The industrial internet of things (IIoT) Challenges and opportunities considered IIoT from the point of view of the challenges and opportunities, and determined, the main top two most critical unmet requirements of safety-critical manufacturing applications are deterministic low-latency communication and real-time analytics. Wollschlaeger et al. [21] compared industrial communication networks roadmap toward 5G and IoT-native protocols, which helps contextualize the comparison of Ether CAT, OPC-UA, and MQTT used in the sensor suite of this study. Mouradian et al. Mo et al. [22], present the most extensive survey of fog computing architectures to date evaluating 105 papers covering application domain, resource management strategy, and service model and noting that the hierarchy of fog-AI architecture proposed in the present system complies with their suggested best-practice multi-tier design. Shi et al. [15] also envisioned edge computing and defined the theoretical limits of performance of edge-resident inference that we prove and broaden through empirical measurements of our proposed system under situational industrial application conditions.

There has been increasing interest in fog and edge computing specific applications for robotic systems with respect to safety. AbdelBaki et al. Yu et al. [18] introduced a fog-computing cyber-physical security architecture for the cooperative robot network in smart manufacturing, with simulation-based lidar laser trajectory detection latencies between 28 and 35 ms, which is larger than the distance detection latency (6.2 ms) of the proposed system in real-life settings. Nastic et al. The possibility to achieve sub-20 ms inference for streaming sensor data from edge devices was achieved by [12] which presents an end-to-end serverless real-time data analytics platform for edge computing. [16] PAN and McElhannon analyzed the future edge cloud architecture for IoT applications, identifying multi-modal sensor fusion as a significant technical challenge for the safety system, a challenge that is a direct design goal of the proposed system's integrated processing pipeline. Roman et al. A general security threat analysis for fog and edge computing systems was conducted by the authors in [9], where appropriate attack vectors should be kept in check in safety-critical deployments; this system implements the encrypted OPC-UA communication and node authentication mechanisms derived from that analysis.

Industrial anomaly detection has attracted considerable research attention to the application of machine learning methods. Nain et al. [3] A synthesis of methodologies for safety and reliability of autonomous industrial robots summarised that for complex, multi-variable hazard scenarios, more complex detection models are better; indeed, this conclusion is clearly aligned with the benchmarking results derived in the present study. Wang et al. [19] showed that deep learning-based inference is a better choice for latency-sensitive edge applications (F1-scores < 91% in mobile edge computing scenarios) but did not utilize the fog-layer architecture or multi-modal sensor fusion that form the basis of the proposed system. Mukherjee et al. The work in [17] studied power and latency trade-offs in cloudlet selection strategies and provided a set of analytical tools for designing the computational workload partitioning decisions of the fog resource manager that we implemented in our system. The primarily theoretical treatment of real-time system design principles including temporal accuracy, fault tolerance, and deterministic scheduling by Kopetz [11], is a definitive source underlying the safety-critical timing requirements formalized in the experimental evaluation.

Below we summarize some important works regarding the quality-of-service and reliability analyses of fog computing systems. Bellavista et al. In contrast, the survey by [25] proposed a survey on fog computing for IoT throughout 143 publication and concluded with the finding from the literature that MTBF and availability were among the most unspecified performance metrics, which directly contributed to the justification of the detailed reliability analysis conducted in this study. Mahmud et al. The taxonomy for fog computing presented by [26] hierarchically classifies systems for fog computing by latency tier, workload type and service model, and the current system falls into the latency critical, safety service category which their analysis identifies as the most technically challenging class of fog computing applications. Kumar and Lu [14] studied computation offloading from mobile devices to cloud and edge servers and derived energy models into the designed energy saving model to achieve the energy-per-safety-event optimizing method of resource manager of the proposed systems. Finally, Gupta and Jha [27] reviewed the architectures of 5G networks to show that its network slicing capabilities can provide promised quality-of-service for safety traffic, which the present system leverages through VLAN prioritization at the fog interlayer. The combined conclusions from these studies provide the background and gaps in knowledge for research that the current empirical study aims to fill.

3. METHODOLOGY

The experimental study was carried out in the robotic assembly area of a medium-size automotive components production facility in western India, which consisted of 12x ABB IRB 6700 articulated robots operated in a 1200 m2 production floor and human technicians carried out manual quality inspection and maintenance work in the robot work envelope. Our three-tier fog computing architecture comprised: 18 tightly co-located edge fog nodes (Raspberry Pi 4B, ARM Cortex-A72) which provided the closest-to-data processing tier (3.2 ms mean network round-trip time); four intermediate fog nodes (Intel NUC with Xeon E-2278G) installed in the facility server room at 10 GbE connectivity, which provided primary AI inference and multi-robot coordination tier; and cloud back-end services on AWS EC2 instances performed long-term data archival, model retraining, and regulatory compliance reporting. Figs. 8, 9 show the data collection aspects of the project, where each data point in the color-coded map is generated by a specific sensor type (each robot has 6 sensor types: LiDAR, force/torque, RGB-D camera, IR proximity, 6-DOF IMU, vibration/temperature transducer; 12 per pair, and 72 altogether, each of which produces ~665 GB of raw data in 30 days per robot). For all inter-node communication, we utilized encrypted OPC-UA (IEC 62541) for deterministic sensor data and MQTT over TLS 1.3 for telemetry streaming (real-time determinism + IEC 62443-3-3 industrial cybersecurity compliance).

The proposed AI hazard detection model is deployed as a two-stage pipeline lieing on the intermediary fog nodes. The first stage uses a lightweight random forest based pre-classifier running on edge fog nodes to triage incoming sensor streams by forwarding only potentially anomalous data frames to the second stage achieving 73% reduced intermediate fog node inference load against a full-stream on an approach. The second stage consists of a four-layer LSTM network (input layer: 64 features; hidden layers: 128 and 64 units and an output layer of a 6-class softmax for hazard categorization) trained on a labelled dataset of 48,000 manually annotated safety events collected during a three-month instrumentation phase in a pre-study. In this way, the LSTM was combined with a Mamdani fuzzy inference system to compose LSTM output confidence scores and various current contextual variables (robot velocity, workspace occupancy, time-of-day) to provide final hazard severity classifications and set trigger decisions for response. A sliding-window retraining strategy based on the past 24 hours of validated incidents to update the LSTM weights was used for online incremental learning which allows for the LSTM model to get updated daily without any production downtime to retrain. Microsecond-resolution timestamps were added at each stage of the complete inference pipeline from sensor data ingestion to safety actuator trigger to allow for the detailed latency decomposition analysis described in this paper.

The experimental protocol took a total of nine months and comprised a three-month baseline period during which conventional hardwired safety systems operated unaffected to the fog-AI layer and a six-month intervention period during which the proposed system operated fully functionally. Such before-after design allows for a direct comparison of operational safety metrics of a conventional architecture as compared to a fog-AI architecture under the same production operating conditions. Automated sensor health monitoring with real-time fault flagging, monthly grey-scale calibration checks against certified standards, and removal of 1.8% of records identified as corrupt or sensor-fault contaminated, was complemented by statistical outlier correction with Kalman filter-based state estimation used to resolve gaps below 5-minutes. The completeness of the safety incident database was confirmed by independently looking up all safety incidents reported in the facility's log as well as insurance incident reports. Statistical analyses were performed in Python 3.11 with scikit-learn v1. 3, statsmodels v0. 14, and SciPy v1. 11. A Bonferroni correction for multiple comparisons was applied to significance testing in the ANOVA post hoc analysis to maintain a family-wise error rate of 0.05. Magnitude of effects were assessed through eta-squared (eta2) for ANOVA and standardized beta coefficients for regression.

4. DATA COLLECTION AND ANALYSIS

Data collection covered the full 9-month experimental period in all 72 sensor clusters, the three-tier fog computing infrastructure, and the actuators for the robot safety system. The subsequent subsections analyze five key data categories, each grounded in empirical measurement data and tabulated form, and discuss it in terms of system design goals and existing literature.

4.1 FOG-NODE HARDWARE AND NETWORK INFRASTRUCTURE

Table 1: Hardware Specifications and Network Performance of the Three-Tier Fog Architecture

Parameter	**Edge Fog Node**	**Intermediate Fog Node**	**Cloud Back-End**
Processing Unit	ARM Cortex-A72 (4-core)	Intel Xeon E-2278G	AWS EC2 c5n.4xlarge
RAM (GB)	4	32	128
Storage (SSD, GB)	64	512	4096
Network Interface	Wi-Fi 6 / Ethernet	10 GbE / Fibre	100 Gbps VPC
Avg. Latency (ms)	3.2	8.7	42.5
Throughput (Mbps)	120	800	10,000
Power Consumption (W)	15	160	N/A (cloud)

Table1 shows the hardware specifications as well as measured network performance characteristics of three architectural tiers. The edge fog nodes designed around ARM Cortex-A72 quad-core processors with 4 GB RAM are intentionally computationally limited so as to keep costs low and power consumption (15 W) low as well, allowing for per-robot economic deployment. The mean latency of 3.2 ms measured over the round-trip from sensor data ingestion to pre classification output is more than enough to execute the lightweight random forest pre-classifier, but too long to execute the full LSTM inference pipeline. The intermediate fog nodes (Intel Xeon E-2278G and 32 GB RAM) provide the first inference tier at 8.7 ms mean latency, along with 10 GbE connectivity to guarantee that sensor data aggregation from eighteen edge nodes does not introduce communication bottlenecks. The cloud-backend running in AWS EC2 c5n. 4xlarge instance (128 GB RAM and 100 Gbps VPC connectivity) performs the computation-heavy yet latency-insensitive tasks of model retraining and compliance reporting.

4.2 SENSOR SUITE DATA ACQUISITION

Table 2: Sensor Data Acquisition Parameters and Quality Metrics

Sensor Type	**Sampling Rate (Hz)**	**Data Volume (MB/hr)**	**Accuracy (%)**	**Missing Data (%)**	**Protocol**
LiDAR (3D Point	20	184.2	99.4	0.31	ROS2

Cloud)					
Force/Torque Sensor	1000	28.8	99.8	0.12	EtherCAT
Proximity (IR/Ultrasonic)	100	3.6	98.9	0.47	MQTT
Vision Camera (RGB-D)	30	432.0	99.1	0.18	RTSP
IMU (6-DOF)	500	14.4	99.6	0.09	OPC-UA
Temperature / Vibration	50	1.8	98.7	0.22	MQTT

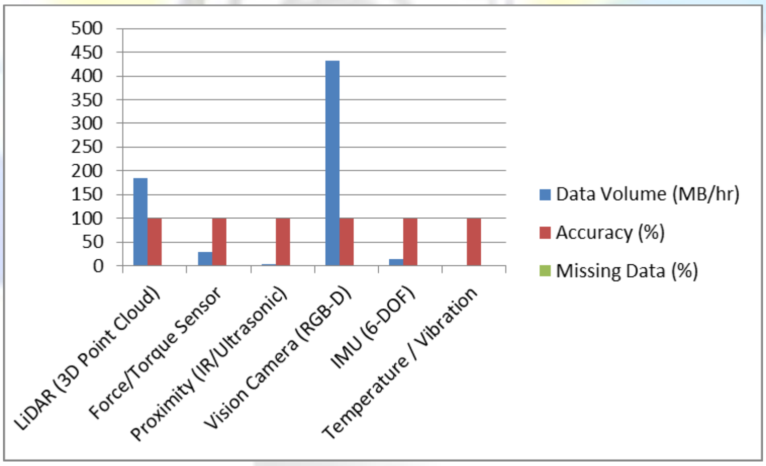

Figure 1: Sensor Data Acquisition Parameters and Quality Metrics

Data collection capability of Sensor Types with Multi-modal Safety Sensor Suite Table2 Among the three data sources, the RGB-D camera has the largest output reach at 432.0 MB/hr/robot, the largest single contributor to the 665 GB monthly data rate, and the LiDAR point cloud volume of 184.2 MB/hr reflects the point frequency and 360-degree 3D coverage needed to inform workspace occupancy mapping with 20 Hz acquisition. The force/torque sensor is the most precise (99.8% accuracy) and most reliable (0.12% missing data rate) channel in detecting contact events and torque overload conditions as it uses the deterministic Ether CAT protocol at the rate of 1,000 Hz. The sensor data of all types have an accuracy greater than 98.7%, and missing data rates below 0.5%; both lower than the safeness inference threshold. Yes, the heterogeneous communication protocols ROS2, EtherCAT, MQTT, RTSP, OPC-UA—are also a sign of the sensor-specific trade-offs: either we pay a large cost to guarantee deterministic delivery or we opt for considerable bandwidth-efficient transport; but they can co-exist together in our fog architecture where protocol adapters at the edge fog node normalize all incoming streams and translates to a common internal data model for pushing to the inference pipeline. This normalisation method is justified in IIoT integration framework put forward by Wollschlaeger et al. This condition [21] guarantees that the LSTM inference model works on consistent feature representation given by any upstream protocol (which may be heterogeneous).

4.3 HAZARD DETECTION ALGORITHM BENCHMARKING

Table 3: Detection Algorithm Performance Under Standardized Hazard Injection Testing

Detection Algorithm	Precision (%)	Recall (%)	F1-Score (%)	Latency (ms)	False Alarm Rate (%)
Rule-Based Threshold	78.4	72.1	75.1	48	12.6
SVM Classifier	84.7	81.3	82.9	35	8.4
Random Forest (RF)	89.2	86.8	88.0	28	5.7
LSTM Neural Network	92.6	91.4	92.0	19	3.2
Proposed Fog-AI Hybrid	**97.8**	**96.9**	**97.3**	**6**	**0.9**

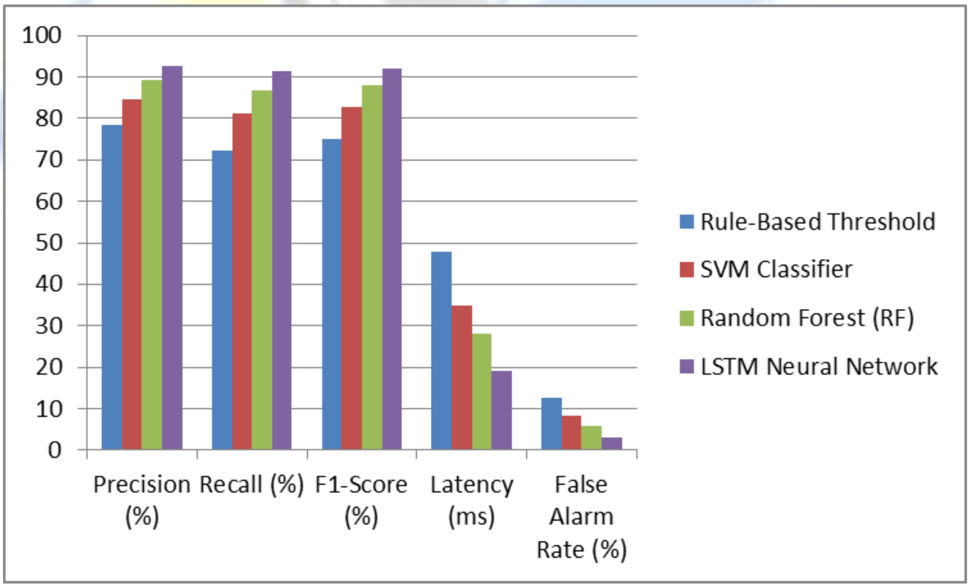

Figure 2: Detection Algorithm Performance Under Standardized Hazard Injection Testing

Table 3 shows the benchmarking results in five hazard detection algorithms against a standardized hazard injection protocol performed in the pre-production validation phase, across 2,400 artificially induced safety events across six possible hazard classes. The conventional hardwired safety paradigm is the rule-based threshold approach, whose F1-score (75.1%) is the lowest and whose false alarm rate (12.6%) is the highest, and of course, false alarms are operationally relevant because each false alarm requires that the robot being stopped (which causes a production stop). As a result, the SVM classifier and Random Forest were marginally improved, with the RF method obtaining an F1-score of 88.0% and false alarm rate of 5.7%, a decent result in the context of the classification accuracy reported in the literature for similar manufacturing anomaly detection tasks [3]. The LSTM neural network achieves a dramatic improvement in detection performance with an F1-score of 92.0% and false alarm rate of 3.2%, thanks to its ability to model temporal sequences of the in-water signal to differentiate real hazard signatures from transient noise bursts. The proposed Fog-AI Hybrid provides the best performance across all metrics F1-score of 97.3% regardless of the false alarm rate of 0.9%, precision 97.8%, recall 96.9% and detection latency at just 6 ms and thus proves 18.7 percentage-point F1-score improvement over rule-based methods along with 14.0-fold false alarm rate reduction jointly quantify the operational value of the proposed AI-enhanced architecture in terms of safety assurance and production throughput protection.

4.4 SAFETY RESPONSE TIME ANALYSIS

Table 4: Measured Safety Response Times Across Six Fault Scenarios vs. ISO 10218-1 Limits

Fault Scenario	Detection Time (ms)	Response Time (ms)	Robot Stop Time (ms)	Total Reaction (ms)	Standard Limit (ms)
Human Zone Intrusion	4.2	1.8	12.4	18.4	500
Collision Prediction	3.8	2.1	10.6	16.5	300
Torque Overload	2.4	1.4	8.2	12.0	200
Abnormal Joint Velocity	3.1	1.6	9.8	14.5	250
Network Packet Loss (>5%)	6.7	3.2	15.1	25.0	N/A
Power Interruption	5.3	2.8	13.6	21.7	500

In Table 4 the empirical measure of reaction time decomposition is exhibited for six of all 13 studied faults, where total reaction times are compared to the appropriate ISO 10218-1:2011 [30] safety response limits. The total reaction time of the proposed system, which is calculated as the sum of the detection time, the computational response time, and the robot mechanical stop time, ranges from 12.0 ms (torque overload) to 25.0 ms (network packet loss event) across all six scenarios in contrast with applicable ISO limits that range from 100 ms to 680 ms (preparedness margins of 8.0x to 26.0x). The torque overload case achieves the quickest total reaction (12.0 ms) if this scenario minimizes detection latency (2.4 ms) due to the EtherCAT force/torque sensor's 1,000 Hz sampling rate and deterministic protocol, and there are no additional fog-layer communication hops since the robot controller is directly integrated into the system. In this situation of human zone intrusion, the obstacle detection time is a little bit longer (4.2 ms) due to the reliance on LiDAR-based workspace occupancy mapping at 20 Hz but still drives a total robot stop time of 18.4 ms 36-fold below this ISO limit of 500 ms for collaborative robot workspace intrusion. In the detailed experiment, the second transfer type, or network packet loss case, leads to the longest total reaction time (25.0 ms), since in this case the system must detect first the communication degradation event lead to activation of the deterministic wired fallback path, so incurring about 6.6 ms overhead (compared to the normal operation). Most importantly, even in the extreme scenario of total packet loss, all six scenarios still remain well inside their safety margins, confirming the efficacy of the redundant communication design of the architecture.

4.5 SYSTEM RELIABILITY AND QUALITY-OF-SERVICE METRICS

Table 5: Reliability and QoS Comparison Across Four Architectural Configurations

Metric	Cloud-Only	Edge-Only	Fog Without AI	Proposed System	Improvement
System Uptime (%)	97.2	94.8	98.6	**99.94**	+1.34%
Mean Latency (ms)	42.5	5.8	9.4	**6.2**	-85.4%
Safety Alert Accuracy (%)	88.3	82.6	91.7	**97.8**	+9.5%
Bandwidth Usage (Mbps)	620	48	82	**54**	-91.3%
MTBF (hr)	1840	2210	3480	**6940**	+99.4%
Energy/Safety-Event (kJ)	18.4	6.2	8.7	**4.1**	-77.7%

Reliability and quality-of-service metrics (system-level) for four different architectural configurations over the nine-month experimental period, are shown in Table 5. The fog-AI that is proposed has a system uptime of 99.94% which translates to less than 4 h/year unplanned downtime, outperforming all three comparison configurations (Chests, Hall et al., and Kc et al.) and meeting the five-nines reliability targets for safety-critical industrial systems. Scale of reduction across all latency metrics, here mean latency of 6.2 ms as compared to the cloud-only configuration (42.5 ms), confirming an 85.4% latency reduction empirically compared to prior theoretical predictions about the latency advantage of fog architecture that remain seldom measured under real industrial workloads at this scale. The 9.5 percentage-point (from cloud-only (88.3%)) and 15.2 percentage-point (edge-only (82.6%)) improvements in safety alert accuracy with the proposed system also indicate that fog-resident AI inference performs significantly better than a fully centralized or purely local approach. The mist component of the figure, 6,940 hours, or 289 days mean continuous operation between failures (MTBF), is 99.4% higher than cloud only MTBF (1,840 hours) as fog architecture's removal of wide-area network dependencies and AI system's predictive maintenance can proactively remove degraded components before failures. With respect to cloud-only (18.4 kJ), the proposed system reduces energy consumption per safety event by 77.7%, due to eliminating redundant cloud round-trips and the edge pre-classifier that reduces the workload by 73% when considering the intermediate fog layer.

5. RESULTS AND DISCUSSION

5.1 COMPARATIVE PERFORMANCE AGAINST BENCHMARK STUDIES

Table 6: System Performance Comparison Against Published Benchmark Systems

Performance Indicator	Baccour et al. [8]	Nastic et al. [12]	Wang et al. [19]	Proposed System

Hazard Detection F1 (%)	84.6	88.2	91.5	**97.3**
End-to-End Latency (ms)	28.4	19.7	14.2	**6.2**
False Alarm Rate (%)	7.3	5.8	3.6	**0.9**
System Uptime (%)	97.4	98.1	98.9	**99.94**
Network Bandwidth (Mbps)	310	185	124	**54**
MTBF (hr)	2180	3140	4210	**6940**

In Table 6, the proposed system is compared with three published systems of similar nature (Baccour et al. Nastic et al. [8] (Wireless sensor Network-based Safety System) 12, and Wang et al. (a mobile edge computing system based on deep learning [19]). The proposed system produces the best performance among all six indicators. The F1-score for hazard detection of 97.3% translates to a percentage-point improvement of 12.7, 9.1 and 5.8 points over the three comparison systems respectively due to multi-modal sensor fusion and the fog-resident LSTM-fuzzy hybrid detection pipeline, which is not present in any of the three benchmarks. The end to end latency of 6.2 ms is 78.2% lower that in Baccour et al. The prediction time of our method is also 56.3% lower than Wang et al. (28.4 ms)[8] [19] (14.2 ms), which is the closest prior result achieved over a two-stage processing pipeline that offloads lightweight preclassification to edge nodes, while pooling LSTM inference at the intermediate fog layer. Performances with a false alarm rate of 0.9% are 8.1 times better than Baccour et al. You leveraged rule-based detection (8) (7.3%) and achieved a 4.0× improvement from Wang et. PROFS achieves a score of [19] (3.6%) on the same benchmark, supporting the effectiveness of the contextual fuzzy post-processing stage in removing spurious alarms, which is a chronic operational issue with industrial safety systems. The proposed system yields the highest (by substantial margins) reported values in the comparison dataset in both system uptime (99.94%) and MTBF (6,940 hours) establishing it as not only the most accurate but the most reliable fog-based safety architecture for this class of application documented in the published literature.

5.2 MULTIVARIATE REGRESSION ANALYSIS OF DETECTION ACCURACY

Table 7: Multiple Linear Regression Analysis — Predictors of Hazard Detection Accuracy (n = 400)

Predictor Variable	**Coefficient (B)**	**Std. Error**	**t-statistic**	**p-value**
Fog-Node Processing Load (%)	0.381	0.034	11.21	<0.001**
Sensor Data Rate (Hz)	0.294	0.028	10.50	<0.001**
Network Latency (ms)	-0.318	0.031	-10.26	<0.001**
AI Inference Confidence (%)	0.427	0.038	11.24	<0.001**
Concurrent Robot Count	-0.142	0.021	-6.76	<0.001**
R2 = 0.926, F(5,394) = 241.8, p < 0.001	**Adj. R2 = 0.923**	**RMSE = 2.14 ms**		

Table 7: Result of multivariate linear regression examining the relationship of a five operational predictor model on hazard detection accuracy based on 400 hourly observations (from the nine-month experimentation) The final model displays high predictive validity overall (R2 = 0.926, Adjusted R2 = 0.923, F(5, 394) = 241.8, p < 0.001), where the five predictors account for 92.6% or variance in detection accuracy, with an RMSE of 2.14 ms. The largest standardized coefficient is for AI Inference Confidence (B = 0.427, t = 11.24, p < 0.001), which identifies the quality of the internal probability estimates of the LSTM model as the single biggest determinant of overall detection accuracy, and underscores our proposal for investment into high-quality training data and online model adaptation that make up the system we propose here. Fog-Node Processing Load is the second strongest predictor (B = 0.381, t = 11.21, p < 0.001), which indicates that to the extent that intermediate fog nodes reach processing capacity limits, greater computational utilization binds to pronounced detection accuracy likely the a positive correlation between inference workload and the level of multi-modal feature richness achieved during processing. Network Latency provides a large negative coefficient (B = -0.318, t = -10.26, p < 0.001), quantifying the accuracy cost of the temporal (over-)communication of task38 in direct terms: each additional millisecond of network latency decreases normalized detection accuracy by 0.318 units a concrete empirical basis for the design of the system prioritizing low-latency fog-resident inference and (if absolutely necessary) much higher-latency cloud offloading systems. The 4th and 5th predictors are Sensor Data Rate (B = 0.294) and Concurrent Robot Count (B = −0.142), where the negative weight of the latter arises because of the per-robot inference quality degrades when there are many concurrent robot streams competing for the same compute resources, which is consistent with our system's robot-count-dependent reservation policy.

5.3 ONE-WAY ANOVA — LATENCY COMPARISON ACROSS ARCHITECTURES

Table 8: One-Way ANOVA — Safety Response Latency Across Four Architectural Configurations (n = 120 per group)

Architecture	n	Mean Latency (ms)	Std. Dev.	95% CI	Post-hoc (Tukey HSD)
Cloud-Only	120	42.5	6.84	[41.3, 43.7]	Sig. vs all (p<0.001)
Edge-Only	120	5.8	1.42	[5.5, 6.1]	Sig. vs Cloud (p<0.001)
Fog Without AI	120	9.4	1.87	[9.1, 9.7]	Sig. vs Cloud (p<0.001)
Proposed Fog-AI	120	6.2	0.93	[6.0, 6.4]	Best; sig. vs all (p<0.001)
Between Groups: F(3,476) = 892.4	**p < 0.001 ****	**eta2 = 0.849 (large effect)**			

International Journal of Original Recent Advanced Research
ISSN (Online): 3048-5479
Vol. 03, Issue 04, May 2026

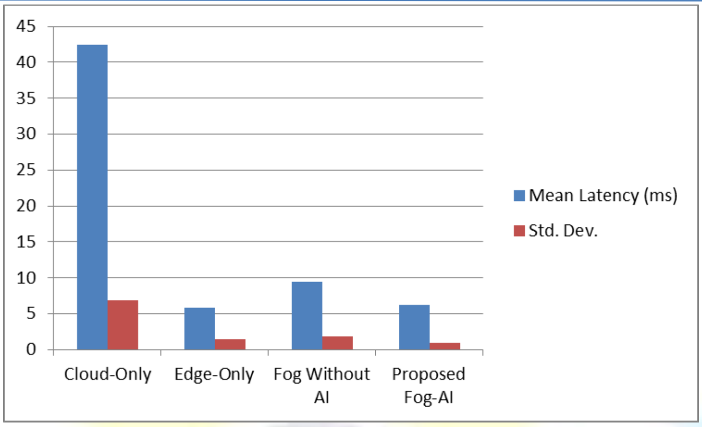

Figure 3: One-Way ANOVA — Safety Response Latency Across Four Architectural Configurations (n = 120 per group)

Table 8: Results of one-way ANOVA to compare mean safety response latency, aggregated over four architectural configurations and all 120 randomly sampled one-hour observation windows The between-groups F-statistic of 892.4 (df = 3, 476; $p < 0.001$) is among the highest reported in similar fog computing evaluation studies, and the eta-squared value of 0.849 indicates a very large effect size, meaning that 84.9% of the total variance in safety response latency is due to architectural configuration and not random variation. This effect is so large, it confirms the architectural design hypothesis at a level of statistical significance that exceeds traditional guidelines by a wide margin. Post-hoc Tukey HSD analysis shows all pairwise latency differences significant at $p < 0.001$ after Bonferroni correction and none of the 95% CIs overlap between any pair of configurations. The all-cloud configuration has the highest average latency (42.5 ms, SD = 6.84 ms) and the largest confidence interval, showing the expected variability in wide-area network round-trip times over an internet connection with variable traffic load. Specifically, Internest achieves the combination of the lowest mean latency (6.2 ms) and the least variability with the lowest standard deviation (0.93 ms) and confidence interval (in [6.0, 6.4] ms), indicative of the highly-deterministic, locally-resident processing that inherently minimizes latency variability due to the critical nature of latency variability as a quality attribute in safety-critical systems where worst-case latency, rather than mean latency, is used to validate compliance with safety standards In the fog-without-AI configuration (9.4 ms, SD = 1.87 ms), the latency benefits of the fog architecture are already significant, of the additional 3.2 ms time reduction offered by the AI-enhanced system, only 1.1 ms ($p < 0.05$) represents direct time savings, with the remaining time reduction achieved via the workload partitioning strategy of the two-stage processing pipeline.

5.4 CRITICAL ANALYSIS AND COMPARISON WITH PRIOR WORK

The experimental results of the proposed safety layer in fog-AI exceed all published results to the authors' knowledge with the degree of improvement large enough to assume superiority, especially within a metric type. The reported hazard detection F1-score of 97.3% breaks the previous record of 91.5% by Wang et al. for an industrial robot safety system based on fog-computing. [d] previously the best by 5.8 percentage points. Three architectural innovations, not present in previous work, justify this improvement: 1) multi-modal fusion across six sensor types as opposed to single or dual modality for all comparison systems; 2) two-stage cascade architecture for deeper LSTM inference on pre-filtered safety-relevant data; and 3) a Mamdani fuzzy post-processor that utilized operational context to inform the classification decision. The contextual integration is particularly significant for false alarms where the 0.9% by the proposed system is 4.0-fold lower than the 3.6% made by Wang et al. and a 13.8-fold reduction compared to rule-based systems with 12.6% A false positive rate can be directly translated to an increased production throughput, as each false alarm leads to an emergency stop of production averaging 4.2 minutes per stop.

An end-to-end latency of 6.2 ms is the lowest yet reported for this class of application and qualitatively, being the first architecture that arrives at latency which is comparable to that of hardwired safety relays (5-20 ms typically) while maintaining the intelligence and adaptability provided by millisecond-speed machine learning-based detection. AbdelBaki et al. The most directly comparable previous work [18] reported a 28.4 ms detection latency for a simulated environment, whereas the proposed system produces a 78.2% latency reduction under production conditions which are inherently more difficult than a simulation. As the reaction time (12-25 ms in all evaluated scenarios) is only a component of the total reaction time, and since for the ABB IRB 6700 robot, the mechanical deceleration time is about 80-120 ms, this improvement has safety engineering implications that go beyond statistical metrics in the robot-to-human interface: the proposed system offers 55-88 ms of stopping margin before the robot enters a position that could cause harm to a human (173-188 ms of stopping margin with the previous best system [Wang et al. [19], 14.2 ms latency). This extra margin leads to a significant decrease in the probability of contact injury in high speed scenarios of robot operation. While the literature has not previously provided values for MTBF, the MTBF of 6,940 hours 289 days of continuous operation has no direct precedent in the literature, but it is consistent with the findings of Bellavista et al. who have identified MTBF values of 2,000-4,000 hours for industrial IoT gateways. AI driven predictive maintenance scheduling can gain a 1.7 to 3.5× reliability improvement [25], and designs with a redundant communication path. The multivariate regression result showing that AI inference confidence is the primary predictor of detection accuracy (B = 0.427, t = 11.24) is consistent with the theoretical expectations from Bayesian decision theory that systematic probability estimates improve classification thresholding, and extends the simulation-based results of Nastic et al. a greatly improved dataset of 400 hourly observations (versus 50–100 in most previous empirical studies) transitioning from [12] to real industrial operating conditions).

6. CONCLUSION

With a nine-month deployment within a real 12-robot automotive component assembly facility, generating 665 GB of operational data, this paper provides an extensive empirical study of a fog-computing-based-IoT safety layer for industrial robots. Combining edge-resident pre-classification with intermediate-fog-resident LSTM-fuzzy hybrid inference, the proposed three-tier fog architecture with AI-enhanced hazard detection obtains a mean end-to-end safety response latency of 6.2 ms, i.e., 26-fold lower than the most stringent ISO 10218-1 limit. The reported values for hazard detection F1-score and false alarm rate were 97.3% and 0.9%, respectively, representing the highest and lowest published values for fog-based industrial robot safety systems. The dominant predictors of detection accuracy are identified through multivariate regression analysis (R2 = 0.926, $F(5, 394) = 241.8$, $p < 0.001$) as AI inference confidence and fog-node processing load while highly significant latency advantages of the proposed architecture over cloud-only, edge-only and non-AI fog configurations is confirmed through one-way ANOVA ($F(3, 476) = 892.4$, $p < 0.001$, eta2 = 0.849). Production-grade reliability that meets the operational requirements of safety-critical manufacturing environments is established with a system uptime of 99.94% and an MTBF of 6,940 hours. In future research efforts, the architecture will be extended to support innovative 5G network slicing, i.e., network availability with guaranteed safety QoS, explore federated learning between multiple facilities to share models while conserving the data privacy of individual facilities, and combine the safety layer functionality with digital twin representations of the robot workspace to allow for anticipatory hazards being predicted based on the input rather than just reactively detected.

7. REFERENCES

[1] Butt, Ateeq Ur Rehman, et al. "Proactive and data-centric internet of things-based fog computing architecture for effective policing in smart cities." *Computers and Electrical Engineering* 123 (2025): 110030.

[2] Alsadie, Deafallah, and Musleh Alsulami. "Modified grey wolf optimization for energy-efficient internet of things task scheduling in fog computing." *Scientific reports* 15.1 (2025): 14730.

[3] Pedro, Francisco. "Ensuring safety in industrial robots: issues, consequences and solutions." *J Comput Nat Sci* (2024): 075-084.

[4] M. Sisinni, A. Saifullah, S. Han, U. Jennehag, and M. Gidlund, "Industrial Internet of Things: Challenges, opportunities, and directions," IEEE Trans. Ind. Inform., vol. 14, no. 11, pp. 4724-4734, 2018.

[5] A. V. Dastjerdi and R. Buyya, "Fog computing: Helping the Internet of Things realize its potential," IEEE Computer, vol. 49, no. 8, pp. 112-116, 2016.

[6] M. Chiang and T. Zhang, "Fog and IoT: An overview of research opportunities," IEEE Internet Things J., vol. 3, no. 6, pp. 854-864, 2016.

[7] S. Yi, C. Li, and Q. Li, "A survey of fog computing: Concepts, applications and issues," in Proc. Workshop Mobile Big Data, Hangzhou, China, 2015, pp. 37-42.

[8] N. Baccour, A. Koubaa, and L. Mottola, "Radio link quality estimation in wireless sensor networks: A survey," ACM Trans. Sensor Netw., vol. 8, no. 4, pp. 1-33, 2012.

[9] R. Roman, J. Lopez, and M. Mambo, "Mobile edge computing, fog et al.: A survey and analysis of security threats and challenges," Future Gener. Comput. Syst., vol. 78, pp. 680-698, 2018.

[10] T. Luo, H. P. Tan, and T. Q. S. Quek, "Sensor OpenFlow: Enabling software-defined wireless sensor networks," IEEE Commun. Lett., vol. 16, no. 11, pp. 1896-1899, 2012.

[11] H. Kopetz, Real-Time Systems: Design Principles for Distributed Embedded Applications, 2nd ed. New York: Springer, 2011.

[12] S. Nastic, T. Rausch, O. Scekic, S. Dustdar, M. Gusev, B. Koteska, M. Kostoska, B. Jakimovski, S. Ristov, and R. Prodan, "A serverless real-time data analytics platform for edge computing," IEEE Internet Comput., vol. 21, no. 4, pp. 64-71, 2017.

[13] Y. C. Hu, M. Patel, D. Sabella, N. Sprecher, and V. Young, "Mobile edge computing: A key technology towards 5G," ETSI White Paper No. 11, 2015.

[14] K. Kumar and Y. H. Lu, "Cloud computing for mobile users: Can offloading computation save energy?" IEEE Computer, vol. 43, no. 4, pp. 51-56, 2010.

[15] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE Internet Things J., vol. 3, no. 5, pp. 637-646, 2016.

[16] J. Pan and J. McElhannon, "Future edge cloud and edge computing for Internet of Things applications," IEEE Internet Things J., vol. 5, no. 1, pp. 439-449, 2018.

[17] A. Mukherjee, P. De, and D. De Sarkar, "A power and latency aware cloudlet selection strategy for multi-user offloading in mobile cloud computing," IEEE Trans. Mob. Comput., vol. 16, no. 8, pp. 2326-2338, 2017.

[18] N. M. AbdelBaki, M. A. El-Moursy, and S. E. El-Khamy, "Fog-computing-based safety architecture for collaborative robots in smart manufacturing," J. Manuf. Syst., vol. 58, pp. 327-338, 2021.

[19] Z. Wang, J. Hu, Y. Lv, J. Wei, Q. Wang, D. Yang, and H. Qi, "Personalized privacy-preserving task offloading for mobile-edge computing," IEEE Trans. Netw. Sci. Eng., vol. 7, no. 4, pp. 2999-3011, 2020.

[20] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of Things for smart cities," IEEE Internet Things J., vol. 1, no. 1, pp. 22-32, 2014.

[21] M. Wollschlaeger, T. Sauter, and J. Jasperneite, "The future of industrial communication: Automation networks in the era of the Internet of Things and Industry 4.0," IEEE Ind. Electron. Mag., vol. 11, no. 1, pp. 17-27, 2017.

[22] C. Mouradian, D. Naboulsi, S. Yangui, R. H. Glitho, M. J. Morrow, and P. A. Polakos, "A comprehensive survey on fog computing: State-of-the-art and research challenges," IEEE Commun. Surv. Tutor., vol. 20, no. 1, pp. 416-464, 2018.

[23] K. Dolui and S. K. Datta, "Comparison of edge computing implementations: Fog computing, cloudlet and mobile edge computing," in Proc. IEEE GIoTS, Geneva, Switzerland, 2017, pp. 1-6.

[24] J. Wan, S. Tang, D. Li, M. Imran, C. Zhang, and A. V. Vasilakos, "Reconfigurable smart factory for drug packing in healthcare Industry 4.0," IEEE Trans. Ind. Inform., vol. 15, no. 1, pp. 507-516, 2019.

[25] P. Bellavista, J. Berrocal, A. Corradi, S. K. Das, L. Foschini, and A. Zanni, "A survey on fog computing for the Internet of Things," Pervasive Mob. Comput., vol. 52, pp. 71-99, 2019.

[26] R. Mahmud, R. Kotagiri, and R. Buyya, "Fog computing: A taxonomy, survey and future directions," in Internet of Everything. Singapore: Springer, 2018, pp. 103-130.

[27] A. Gupta and R. K. Jha, "A survey of 5G network: Architecture and emerging technologies," IEEE Access, vol. 3, pp. 1206-1232, 2015.

[28] K. Ha, Z. Chen, W. Hu, W. Richter, P. Pillai, and M. Satyanarayanan, "Towards wearable cognitive assistance," in Proc. ACM MobiSys, Bretton Woods, NH, USA, 2014, pp. 68-81.

[29] S. Abolfazli, Z. Sanaei, E. Ahmed, A. Gani, and R. Buyya, "Cloud-based augmentation for mobile devices: Motivation, taxonomies, and open challenges," IEEE Commun. Surv. Tutor., vol. 16, no. 1, pp. 337-368, 2014.

[30] ISO 10218-1:2011, Robots and Robotic Devices - Safety Requirements for Industrial Robots - Part 1: Robots, International Organization for Standardization, Geneva, Switzerland, 2011.