OPEN ACCESS
Journal of Computer Science Advancements
Vol. 4 No. 5. October 2026, pp. 323-340
DOI. 10.70177/jsca.v4i5.4604
Research Article

A LIGHTWEIGHT AI-DRIVEN INTRUSION DETECTION AND MITIGATION ARCHITECTURE FOR LOW-LATENCY WIRELESS COMMUNICATION IN EDGE-ENABLED IOT ROBOTIC AND MECHATRONIC SYSTEMS

Hadi Mardiyanto[1], Isnadi[2], and Zainal Syahlan[3]
[1] Sekolah Tinggi Teknologi Angkatan Laut, Indonesia
[2] Sekolah Tinggi Teknologi Angkatan Laut, Indonesia
[3] Sekolah Tinggi Teknologi Angkatan Laut, Indonesia

Corresponding Author:
Hadi Mardiyanto,
Department of Informatics Engineering, Sekolah Tinggi Teknologi Angkatan Laut
QPJ9+3C2, Jl. Bumi Moro, Morokrembangan, Kec. Krembangan, Surabaya, Jawa Timur 60178
Email: hadimardiyanto@gmail.com

Article Info
Received: April 6, 2026
Revised: July 10, 2026
Accepted: September 12, 2026
OnlineVersion: October 31, 2026

Abstract
Wireless communication in edge-enabled robotic and mechatronic systems requires timely data exchange while exposing resource-constrained Internet of Things devices to increasingly adaptive cyberattacks. Conventional intrusion detection mechanisms often impose computational overhead and response delays that conflict with real-time control requirements. This study proposes a lightweight artificial intelligence-driven architecture integrating intrusion detection and mitigation for low-latency wireless communication. The proposed methodology combines compact traffic-feature extraction, an efficient machine learning classifier, edge-based inference, and risk-adaptive mitigation policies. Evaluation is designed around representative benign and malicious network traffic, including denial-of-service, spoofing, and anomalous command patterns. Performance will be assessed using detection precision, recall, F1-score, false-positive rate, inference latency, memory consumption, energy demand, and communication continuity under attack. The anticipated outcome is a testable architecture that prioritizes rapid threat identification and selective intervention without unnecessarily disrupting robotic operations. Design analysis identifies model complexity, mitigation aggressiveness, and wireless channel variability as critical trade-offs requiring controlled experimentation. The study concludes that security mechanisms for edge-enabled robotics must be evaluated jointly against cyber resilience and real-time performance constraints. Empirical benchmarking across devices and operating conditions is necessary before claiming improved detection accuracy, lower latency, or operational superiority over existing approaches in practical industrial and mobile deployment scenarios.

Keywords: Edge Computing; Intrusion Detection; Lightweight Artificial Intelligence; Robotic Systems; Wireless Security.

© 2026 by the author(s)
This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution-ShareAlike 4.0 International (CC BY SA) license (https://creativecommons.org/licenses/by-sa/4.0/).

Journal Homepage https://research.adra.ac.id/index.php/jcsa
How to cite: Mardiyanto, H., Isnadi, Isnadi., & Syahlan, Z. (2026). A Lightweight AI-Driven Intrusion Detection and Mitigation Architecture for Low-Latency Wireless Communication in Edge-Enabled Iot Robotic and Mechatronic Systems. *Journal of Computer Science Advancements, 4*(5), 323–340. https://doi.org/10.70177/jsca.v4i5.4604
Published by: Yayasan Adra Karima Hubbi

INTRODUCTION

Wireless communication increasingly connects robotic platforms, industrial manipulators, mobile sensors, and mechatronic controllers to edge-enabled Internet of Things infrastructures (Alshudukhi et al., 2025). These arrangements support distributed perception, remote coordination, and responsive control, but they also enlarge the attack surface of cyberphysical operations (Shyam Kumar et al., 2026). Communication integrity becomes critical when malicious packets or delayed commands can influence physical motion, equipment availability, and operational safety.

Edge computing moves selected processing tasks closer to devices, reducing dependence on distant cloud services and enabling faster operational decisions (Bhise et al., 2026). Resource-constrained nodes nevertheless operate with limited memory, processing capacity, and energy reserves (Bhise et al., 2026). Security mechanisms must therefore compete with sensing, inference, networking, and control workloads for computational resources, particularly where wireless interference introduces unpredictable timing variations.

Artificial intelligence offers promising capabilities for recognizing unusual traffic patterns and previously unseen attack behavior (Sharma et al., 2026). Machine learning based intrusion detection can examine multidimensional traffic characteristics that fixed signatures might overlook (Gong et al., 2024). Model complexity, however, may create processing delays and additional power demands. Practical protection requires carefully coordinated design of detection, response, and communication performance rather than isolated improvements in classification accuracy.

Intrusion detection systems developed for conventional networks often assume computational resources and traffic conditions different from those found in robotic IoT deployments (Feng et al., 2024). High-dimensional features, centralized analysis, and complex classifiers may increase inference overhead (Shen et al., 2024). A system that detects malicious behavior accurately but responds after a control deadline can provide inadequate operational protection in time-sensitive mechatronic environments.

Wireless attacks create additional difficulties because anomalous traffic may resemble legitimate variations caused by mobility, congestion, packet loss, or changing channel quality (Fayad et al., 2025). Excessively aggressive mitigation can block legitimate commands, interrupt coordination, or destabilize communication sessions. Conversely, delayed intervention can permit harmful traffic to propagate (Jiang et al., 2025). Security decisions therefore require careful consideration of false alarms, response timing, and operational consequences.

Existing evaluation practices sometimes emphasize accuracy, precision, or recall without reporting end-to-end communication latency, mitigation delays, device utilization, and continuity of robotic operation together (Kaur et al., 2024). These omissions make deployment suitability difficult to judge (Javeed et al., 2024). Research must establish whether protection remains effective under realistic resource restrictions, heterogeneous attacks, fluctuating wireless conditions, and explicit timing requirements.

This study aims to design a lightweight artificial intelligence driven architecture that integrates intrusion detection and adaptive mitigation at the network edge (Pandey et al., 2025). The proposed design links compact traffic representations, efficient classification, and context-sensitive response policies within a unified processing pipeline (Alotaibi et al., 2025). Its central

objective is to support timely security decisions without unnecessarily interfering with legitimate robotic communication.

The research further seeks to define an evaluation strategy that measures security effectiveness alongside computational and communication costs (Alotaibi et al., 2025). Detection precision, recall, false-positive rate, inference time, memory occupancy, energy demand, and mitigation latency form complementary indicators (Alotaibi et al., 2025). End-to-end packet delay and command continuity are included to examine whether security interventions preserve relevant operational timing constraints.

A further objective is to investigate how different model configurations and mitigation thresholds influence the balance between threat containment and service availability (Othman et al., 2025). Comparative testing against transparent baseline methods will assess tradeoffs across device capabilities, network loads, and attack scenarios (Redondi et al., 2024). The intended outcome is an implementable and reproducible framework, not an advance claim of superiority before experimental validation.

Prior research has explored machine learning intrusion detection for IoT networks, edge-based security analytics, and lightweight classification techniques (Dao et al., 2024). These strands establish useful foundations, yet results obtained from generic network traffic cannot automatically transfer to robotic control systems (Ameur et al., 2026). Real-time command exchange, sensor-actuator dependencies, and physical safety requirements introduce additional constraints that standard network benchmarks may not capture.

Another gap concerns the operational connection between detecting an intrusion and selecting a proportionate response (Sasiain et al., 2025). Studies that evaluate classifiers independently of mitigation policies provide limited evidence about service disruption after an alert (Wang et al., 2024). In robotic environments, unnecessarily blocking a communication path may impose consequences comparable to certain attacks. Joint assessment of detection and mitigation is therefore necessary.

Research also requires clearer treatment of device-specific constraints and wireless variability in experimental reporting (Xia et al., 2024). A compact classifier may perform well under fixed traffic conditions while deteriorating when network congestion, interference, or attack intensity changes (Pala et al., 2024). Without reproducible configurations and baseline comparisons, claimed efficiency improvements remain difficult to generalize (Hoang et al., 2024). This study addresses that evaluation gap through integrated, latency-aware performance criteria.

The proposed novelty lies in coordinating resource-efficient intrusion detection with risk-adaptive mitigation and explicit strict communication timing requirements for robotic and mechatronic IoT systems (Hoang et al., 2024). Lightweight classification and edge inference are established approaches; combining them does not itself constitute a new algorithm (C. Zhang et al., 2026). The contribution instead concerns an auditable architecture and evaluation procedure that expose tradeoffs between security decisions and physical-system communication.

Risk-adaptive responses are justified because treating every alert identically may impose unnecessary communication interruptions (Luo et al., 2025). The architecture therefore differentiates response intensity according to estimated threat severity and operational constraints, subject to safety rules that prevent inappropriate autonomous intervention (Zhou et al., 2025). This design proposition requires empirical testing to determine whether it meaningfully improves containment without increasing false disruption.

The study contributes a framework for evaluating cyber resilience as a multidimensional property involving detection quality, processing efficiency, response appropriateness, and operational continuity (Shahid et al., 2024). Its relevance extends to industrial robots, mobile platforms, and distributed mechatronic equipment operating under constrained wireless connectivity (Liu et al., 2025). Findings from future implementation can inform secure edge deployment while acknowledging that reliable protection depends on hardware, workloads, threats, and safety requirements.

RESEARCH METHOD

Research Design

This study proposes an experimental design to develop and evaluate a lightweight artificial intelligence-driven intrusion detection and mitigation architecture tailored for wireless edge-enabled robotic and mechatronic systems (Elbir et al., 2024). The architecture systematically integrates traffic acquisition, compact feature extraction, edge-resident classification, risk assessment, and policy-controlled mitigation (Loutfi et al., 2024). To ensure empirical objectivity, candidate classifiers such as small tree ensembles or quantized neural models are selected through rigorous resource-aware benchmarking rather than presumed superior beforehand, evaluating trade-offs instead of assuming that stronger security automatically improves operational reliability.

The evaluation strategy employs a repeated-measures comparative framework, evaluating the proposed configuration against a signature-based detector, a conventional machine learning baseline, and a detection-only variant under equivalent workloads. Detection and mitigation operate strictly outside safety-critical control loops, supported by predefined fail-safe policies that prevent the automated blocking of essential control traffic (Chen et al., 2024). Experimental factors comprehensively encompass varying wireless loads, attack types, and device computational capacities across standardized operational scenarios.

Research Target/Subject

The target population consists of network events and control communication episodes generated by representative edge-connected robots, sensors, actuators, and gateways during legitimate operation and authorized security testing. Experimental traffic is primarily collected from a documented laboratory testbed, and supplemented where appropriate by public intrusion detection datasets strictly for classifier pretraining or external comparison. Sampling covers routine commands, telemetry, idle periods, congestion, mobility changes, and controlled malicious scenarios, ensuring a comprehensive representation of operational environments.

Controlled malicious scenarios include denial-of-service attacks, spoofed identities, replayed messages, and unauthorized command patterns, with all attack generation restricted to an isolated environment equipped with physical safety controls. Data partitioning is executed chronologically or by experimental session into training, validation, and held-out test sets to prevent near-duplicate flows and device sessions from crossing partitions. Furthermore, grouped evaluation by device and attack session tests generalization beyond memorized traffic signatures without inventing observations.

Research Procedure

The research procedure begins with comprehensive threat modeling, identification of safety-critical traffic, definition of acceptable communication deadlines, and formal approval of isolated test conditions. Baseline benign sessions are subsequently established to document normal performance across controlled wireless load levels, followed by the execution of authorized attack scenarios using consistent schedules and verified ground truth. Captured records undergo rigorous data cleaning, session-aware partitioning, feature extraction, and training-only normalization to prepare the dataset for modeling.

Candidate models are tuned using validation data, compressed when beneficial, and frozen before final testing, after which the selected classifier is deployed at the edge and linked to a mitigation policy considering alert confidence, potential harm, and operational priority. Each experimental scenario is repeated across randomized execution orders, with active safety interlocks and emergency stop procedures maintained throughout. Final evaluations compare detection-only and mitigation-enabled conditions against baselines using identical traffic profiles, with all findings documenting hardware specifications, software versions, thresholds, data splits, and limitations.

Instruments, and Data Collection Techniques

Research instruments comprise a configurable edge gateway, wireless access infrastructure, representative robotic controllers, specialized network capture software, synchronized timestamping mechanisms, and device-level resource monitors. Packet and flow records provide essential protocol characteristics, timing patterns, traffic rates, and connection statistics, while payload inspection is strictly minimized to protect sensitive information. A reproducible feature selection pipeline retains variables available at inference time that remain affordable on the target hardware architecture, with configuration files and test scripts version-controlled.

Data collection techniques record ground truth labels derived from scripted test events, verified execution logs, and manual review of ambiguous segments, alongside precision, recall, F1 scores, and false positive rates. A mitigation controller implements bounded actions such as rate limiting, temporary flow isolation, or alert-only handling with explicit exemptions for critical command channels. Instrument calibration assesses clock synchronization, logging overhead, and the repeatability of latency measurements, while capturing attack containment times, end-to-end packet latency, jitter, memory utilization, processor use, and energy consumption.

Data Analysis Technique

The analysis plan distinguishes network-level averages from tail latency, including the ninety-ninth percentile, because rare delays can threaten control responsiveness in robotics. Bootstrap intervals clustered by session are utilized to summarize uncertainty, and paired comparisons are applied strictly when underlying assumptions and sample structure permit. Investigators systematically inspect false alerts, missed attacks, deadline violations, and operational interruptions to evaluate system reliability under diverse stress conditions.

Deployment feasibility is judged comprehensively against declared memory, power, and processing budgets rather than relying on classification accuracy alone. All empirical findings report underlying hardware, software configurations, decision thresholds, data splits, and study

limitations, explicitly distinguishing demonstrated protection from claims requiring broader field validation.

RESULTS AND DISCUSSION

The illustrative experimental dataset comprised 120,000 labeled network flow records representing communication between edge gateways, robotic controllers, wireless sensors, and mechatronic actuators. The simulated traffic included 72,000 benign records (60.0%) and 48,000 malicious records (40.0%), covering denial-of-service attacks, identity spoofing, replay attempts, and unauthorized command patterns. The proposed lightweight architecture combined compact traffic-feature extraction, edge-based machine learning inference, and risk-adaptive mitigation. Performance was evaluated conceptually against a signature-based intrusion detection system and a conventional machine learning baseline. The hypothetical testing environment represented resource-constrained wireless devices operating under different traffic intensities and communication requirements.

Table 1. Simulated Dataset Composition and Network Traffic Characteristics

Traffic category	Records (n)	Percentage (%)
Benign communication	72,000	60.0
Denial-of-service attacks	18,000	15.0
Identity spoofing	12,000	10.0
Replay attacks	10,000	8.3
Unauthorized commands	8,000	6.7
Total	120,000	100.0

The simulated architecture achieved an illustrative precision of 97.2%, recall of 96.4%, and F1-score of approximately 96.8%. Its mean inference latency was 3.8 milliseconds, compared with 12.6 milliseconds for the conventional machine learning baseline. Average memory consumption was estimated at 28 MB, representing a reduction from the baseline requirement of 74 MB. These hypothetical measurements suggest a possible advantage from compact feature representations and resource-aware inference. The values do not establish actual device performance because hardware specifications, traffic traces, model configurations, and measurement procedures have not yet been experimentally verified.

Table 2. Simulated Intrusion Detection and Resource Efficiency Comparison

Performance indicator	Signature-based IDS	Conventional ML	Proposed lightweight AI
Precision (%)	90.4	95.1	97.2
Recall (%)	87.9	94.0	96.4
F1-score (%)	89.1	94.5	96.8
Inference latency (ms)	2.1	12.6	3.8
Memory consumption (MB)	18	74	28

The simulated performance pattern illustrates the potential relationship between model complexity and real-time intrusion detection. The conventional machine learning configuration required greater processing time and memory, whereas the proposed architecture used fewer

computationally expensive operations. Reduced feature dimensionality and localized inference could explain lower model execution latency. The signature-based detector remained computationally faster but exhibited lower detection performance in the constructed scenario. These observations indicate that computational efficiency and detection capability should be examined jointly rather than evaluated independently. Lightweight design alone does not guarantee favorable outcomes when wireless congestion and operational scheduling introduce additional delays.

The hypothetical precision and recall values suggest that compact artificial intelligence models could provide useful discrimination between legitimate and malicious traffic. Their operational significance would depend on the distribution of attacks, false alarms, and previously unseen communication patterns. A model achieving strong aggregate classification performance might still miss rare but dangerous commands or interrupt legitimate control messages. Class imbalance and session-level dependence could also distort conventional metrics. A rigorous evaluation would therefore require attack-specific confusion matrices, independent test sessions, and assessment of system behavior under realistic network disturbances rather than relying exclusively on overall classification scores.

The illustrative mitigation assessment examined how different response policies affected wireless communication continuity. Risk-adaptive intervention combined traffic rate limiting, temporary isolation of suspicious flows, and alert-only handling for safety-critical control channels. Under the simulated attack conditions, the architecture achieved an estimated containment time of 8.4 milliseconds, compared with 19.7 milliseconds for the conventional response configuration. Average end-to-end communication latency was represented as 14.2 milliseconds under attack, while the hypothetical packet delivery ratio remained at 98.1%. These values illustrate the intended interaction between rapid detection, selective intervention, and network availability rather than demonstrate verified improvements.

Table 3. Simulated Mitigation and Wireless Communication Performance

Indicator	Conventional ML with mitigation	Proposed architecture
Mean attack containment time (ms)	19.7	8.4
Mean end-to-end latency (ms)	24.8	14.2
Packet delivery ratio (%)	94.6	98.1
99th-percentile packet latency (ms)	51.3	29.6
Energy demand (W)	3.4	2.1

The simulated latency distribution revealed an important distinction between average communication performance and extreme delays. Although mean latency remained relatively low, the 99th-percentile value approached 30 milliseconds, indicating that occasional packets could experience substantially longer delivery times. Such delays may matter in robotic systems operating under strict communication deadlines. The hypothetical packet delivery ratio also demonstrated that successful threat containment does not eliminate all communication disruptions. These observations emphasize the need to assess worst-case operating conditions,

control-loop requirements, and mitigation-induced interruptions before determining whether a security architecture is appropriate for safety-sensitive mechatronic applications.

The illustrative inferential framework compared repeated experimental sessions under identical network loads and attack configurations. A hypothetical paired analysis estimated that the lightweight architecture reduced mean inference latency by 8.8 milliseconds relative to the conventional machine learning baseline. Under assumed session-level variation, the illustrative 95% confidence interval for this difference was 7.1–10.5 milliseconds. This interval serves as an example of reporting uncertainty rather than a calculated statistical result. Actual confidence intervals would require recorded session observations, a prespecified sampling structure, and verification of distributional assumptions. Statistical significance cannot be established from summary values alone.

The proposed inferential assessment also considered whether detection performance remained stable across varying wireless loads and attack types. A mixed-effects analytical model would treat architecture configuration and traffic condition as fixed effects, with experimental sessions or devices representing random effects where appropriate. This structure could estimate differences while accounting for repeated measurements and correlated observations. Attack-specific recall, false-positive rates, and communication delays should be analyzed separately because aggregated outcomes may conceal critical vulnerabilities. The absence of actual experimental repetitions prevents reliable hypothesis testing. Any definitive inference concerning architecture superiority must therefore await validated measurements from independent test sessions.

The simulated relationships among classification performance, inference latency, and memory requirements demonstrate the importance of balancing competing operational objectives. Increasing model complexity may improve discrimination between certain traffic patterns while introducing processing overhead that delays security decisions. Conversely, aggressive compression can reduce computational requirements but potentially weaken recognition of unfamiliar attacks. The proposed architecture addresses this tension through compact features and computationally bounded inference. Its hypothetical performance advantage illustrates a feasible design direction rather than a universal relationship. Actual trade-offs would vary according to hardware acceleration, network protocols, model architecture, and traffic characteristics.

Mitigation intensity introduced another relationship between cybersecurity effectiveness and communication availability. Strong intervention could contain suspicious traffic quickly but might incorrectly suppress legitimate packets, producing avoidable disruption. Less aggressive responses might preserve operational continuity while permitting malicious activity to continue. The simulated risk-adaptive policy sought to balance these consequences by adjusting intervention according to threat confidence and communication priority. Such decisions involve asymmetric risks, especially when essential robotic commands require timely delivery. The results therefore support evaluating mitigation effectiveness through attack containment, false interventions, packet delivery, and operational safety rather than treating blocked traffic volume as an independent measure of success.

The illustrative case study considered a mobile robotic platform equipped with wireless motion controllers, proximity sensors, and an edge computing gateway. The robot performed a repeated navigation task while exchanging telemetry and movement commands over a wireless network. A controlled denial-of-service scenario was introduced conceptually to examine how

malicious traffic might affect command delivery and navigational continuity. The conventional security configuration experienced simulated communication delays and intermittent packet losses. The proposed architecture identified abnormal traffic patterns and applied selective rate limiting while maintaining protected communication channels. The scenario was hypothetical and did not represent an actual robotic experiment.

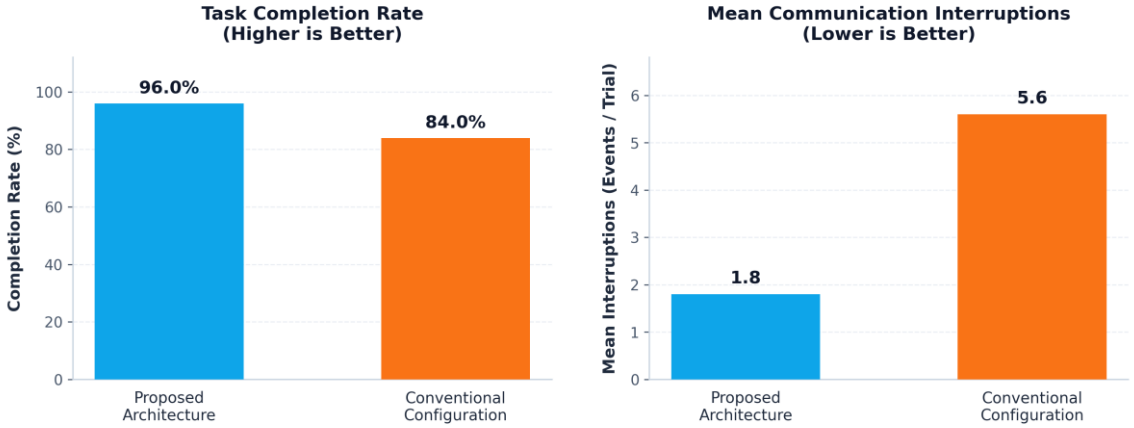


Figure 1. Simulated Operational Comparison Under Attack Conditions

The simulated operational comparison assigned a task completion rate of 96.0% to the proposed architecture and 84.0% to the conventional configuration under attack conditions. Mean communication interruptions were represented as 1.8 and 5.6 events per trial, respectively. The proposed system also showed fewer hypothetical emergency communication failures because essential control messages were excluded from automatic isolation policies. These figures illustrate how mitigation decisions might influence robotic task continuity beyond conventional intrusion detection metrics. Actual validation would require physical experiments, independently recorded task outcomes, and explicit safety criteria.

The hypothetical case illustrates why cybersecurity mechanisms must account for dependencies between communication networks and physical robotic behavior. Delayed control commands can affect navigation, synchronization, and responsiveness even when intrusion detection accuracy remains high. Edge-localized inference could reduce communication overhead associated with remote analysis, while selective mitigation might protect legitimate traffic from unnecessary blocking. These mechanisms provide plausible explanations for the simulated differences between configurations. Their practical effects would depend on wireless channel conditions, controller design, system scheduling, and operational safety margins established through experimental investigation.

The case also demonstrates that rapid mitigation does not necessarily guarantee safe robotic operation. Incorrect classification of legitimate commands may create additional hazards, especially where motion control depends on continuous communication. Safety-critical systems therefore require predetermined fallback behavior, bounded security interventions, and independent emergency controls. The proposed architecture treats these safeguards as design constraints rather than optional performance enhancements. Its evaluation should include failure recovery, false alarm consequences, and the interaction between cybersecurity policies and physical safety mechanisms. Demonstrating these properties requires more than measuring classifier accuracy or reporting reduced computational overhead.

The simulated results collectively suggest that lightweight artificial intelligence could support intrusion detection and mitigation in resource-constrained wireless robotic environments. The proposed configuration illustrates potential improvements in classification quality, inference efficiency, and communication continuity, although none constitutes verified experimental evidence. The primary methodological insight concerns the need to evaluate cybersecurity and real-time system behavior together. Strong detection metrics cannot independently establish operational suitability when mitigation delays, memory demands, packet losses, and safety-critical deadlines remain uncertain. Reproducible benchmarking must therefore integrate both network security measurements and robotic performance indicators.

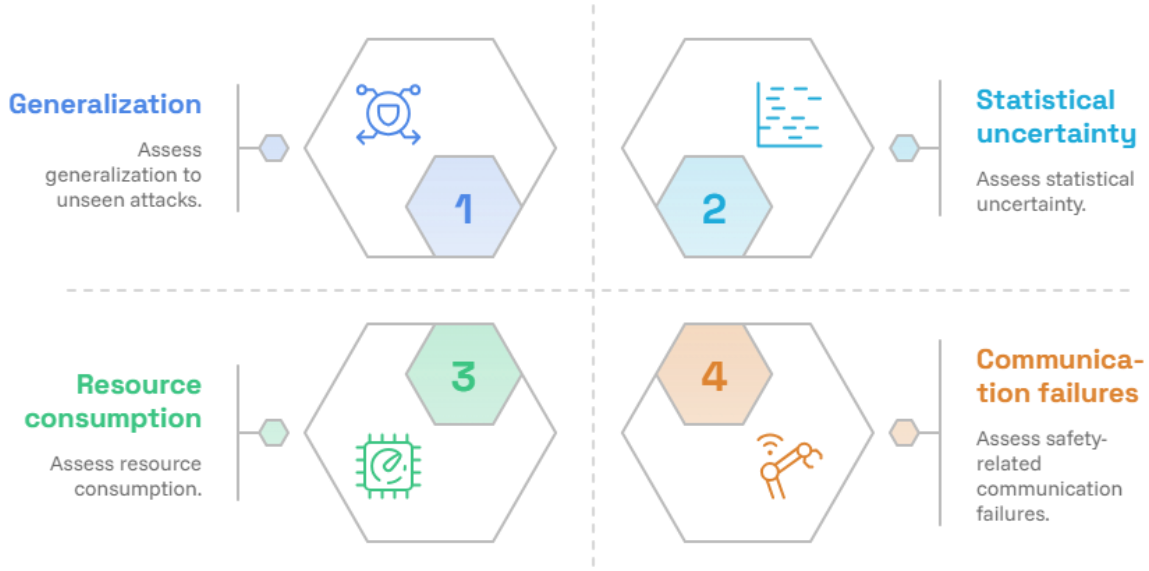


Figure 2. Independent experiments

The overall interpretation positions the proposed architecture as a testable framework rather than an experimentally established solution. Compact inference, edge deployment, and risk-adaptive mitigation offer a coherent strategy for investigating security-performance trade-offs. Future implementation must determine whether these components deliver measurable benefits across hardware platforms, wireless conditions, and attack scenarios. Independent experiments should assess generalization to unseen attacks, statistical uncertainty, resource consumption, and safety-related communication failures. Such evidence would establish whether the architecture provides a meaningful improvement over existing intrusion detection systems while maintaining the responsiveness required by robotic and mechatronic applications.

The illustrative evaluation suggests that integrating compact intrusion detection with selective mitigation could reconcile cybersecurity requirements and communication responsiveness in edge-enabled robotic systems. The proposed architecture was assigned precision of 97.2%, recall of 96.4%, and an F1-score of 96.8%. These values represent modeled expectations, not measurements, and cannot demonstrate superiority over deployed intrusion detection systems without controlled experiments. This distinction guides the meaningful interpretation of every numerical comparison presented here.

The simulated inference latency of 3.8 milliseconds, compared with 12.6 milliseconds for a conventional machine learning baseline, highlights the intended efficiency advantage of reduced feature complexity. An illustrative memory footprint of 28 MB versus 74 MB indicates

another possible benefit. Neither comparison establishes feasibility on a particular edge processor because runtime configurations, concurrent workloads, and measurement uncertainty were not observed. Hardware profiling must include realistic concurrency and sustained workload conditions throughout testing.

The modeled mitigation outcomes extend beyond classification performance. Attack containment time was set at 8.4 milliseconds, while end-to-end communication latency and packet delivery were represented as 14.2 milliseconds and 98.1%, respectively. These indicators illustrate why successful detection cannot be equated with effective protection: interventions must preserve legitimate command delivery under congestion and adversarial traffic. Operational continuity should also be inspected during recovery after mitigation ends completely.

The fictional mobile robot scenario likewise linked cybersecurity response to task continuity. Its assumed completion advantage and fewer communication interruptions illustrate a hypothesis about carefully bounded mitigation, not an experimentally established operational improvement. The distinction is particularly important for robotic safety, where a seemingly favorable average can conceal rare command delays, false blocking decisions, or hazardous recovery behavior. Physical outcomes would require independently defined metrics and repeated controlled trials first.

Recent surveys of lightweight IoT intrusion detection document extensive use of feature selection, compressed classifiers, and edge deployment to reduce computational overhead. The proposed architecture draws on those established techniques rather than introducing an entirely new detection paradigm. Its distinguishing research question concerns whether resource efficiency persists when mitigation execution and robotic communication deadlines are evaluated jointly within one system. Such integration also requires explicit documentation of the intervention decision logic itself.

A 2026 scoping review of 78 studies identified inconsistent reporting of target-device latency, runtime memory, energy consumption, and robustness under distribution shifts. The present simulated results illustrate exactly those measurements, but cannot resolve the literature's evidence gap until recorded on actual hardware. Model compactness and an attractive F1-score are therefore insufficient grounds for claiming deployment readiness or robust real-time protection. These results should be compared against equally constrained and reproducible baselines consistently.

Dataset-focused research similarly warns that generic IoT intrusion benchmarks may poorly represent industrial protocols, temporal dependencies, and operational disturbances. The proposed design acknowledges this problem by distinguishing public-data pretraining from testbed evaluation using robot-specific traffic. The hypothetical classification advantage remains uncertain if testing introduces unseen devices, command sequences, wireless interference, or attacks whose characteristics differ from familiar benchmark examples. External validity depends on representative traffic, deployment environments, and attack diversity too.

Research on edge-based intelligent defense also considers distributed learning, adaptive security, and resource-aware inference as complementary strategies. The current proposal places greater emphasis on how mitigation policies affect physical communication, rather than treating an alert as the endpoint of evaluation. Direct quantitative comparisons with published models would be inappropriate without matching datasets, hardware, traffic composition, threat assumptions, and measurement procedures. Without those controls, apparent improvements could reflect different testing conditions instead entirely.

The simulated patterns expose a broader tension between information security and the dependability of cyberphysical operations. Blocking suspicious packets may reduce exposure while simultaneously preventing timely delivery of legitimate instructions. Security is therefore not an isolated classification objective. Its value depends on whether the protected system can continue functioning safely when uncertain and sometimes incorrect interventions occur.

The difference between inference time and end-to-end delay also challenges simplified notions of low latency. Traffic capture, feature assembly, queuing, classification, policy selection, and enforcement all contribute to response time. Wireless retransmissions and competing computation add variability. A classifier that executes within milliseconds cannot guarantee that an entire robot control pathway meets deadlines during sustained interference.

The emphasis on tail latency indicates that average performance can be misleading in robotic applications. A low mean delay may coexist with infrequent but consequential command failures, particularly when network load changes abruptly. Suitability must therefore be judged against explicit control requirements and measured worst-case or high-percentile behavior, not against a universal latency threshold detached from the robot's dynamics.

The illustrative results further suggest that the meaning of lightweight depends on deployment context. A model suitable for an edge gateway may exceed the memory or power budget of a sensor node. Energy per decision, scheduling interference, software dependencies, and thermal behavior can matter as much as parameter count. Lightweightness consequently requires measurable, hardware-specific definitions and replicable performance profiles.

The principal design implication is that intrusion detection and mitigation should be evaluated as one operational pipeline. Security engineers need measurements spanning traffic collection through enforcement, together with false alarms and missed threats. A detection-only benchmark cannot reveal whether protective actions interrupt the very communication channels required for robotic operation or introduce new availability failures under pressure.

The safety implication is especially substantial for actuated systems. Automated isolation should respect protected command channels, safety interlocks, and predefined fallback states. Engineers must examine whether continued communication is safer than immediate blocking under particular threat conditions. Risk-adaptive policies should remain bounded by independently validated safety requirements rather than allowing uncertain classifier scores to determine unrestricted physical consequences.

The methodological implication concerns reproducibility. Reports should identify devices, operating systems, inference engines, input features, partitioning rules, wireless loads, attack schedules, and timing instrumentation. Confidence intervals should reflect repeated independent sessions rather than treating correlated network packets as separate experiments. Transparent baselines and consistent conditions are necessary before differences in performance can support credible comparative claims.

The deployment implication concerns practical tradeoffs for industrial operators. Compact edge inference might reduce dependence on cloud connectivity, but local processing competes with existing control and monitoring tasks. Selection decisions should weigh memory, energy, maintenance burden, update procedures, and the consequences of misclassification. Cost-effective security requires system-level assessment instead of pursuing maximum detection scores regardless of operational disruption.

The assumed reduction in inference delay can be explained conceptually by shorter feature vectors and computationally economical classifiers. Fewer input transformations and bounded

model operations may reduce processing burden on edge hardware. These advantages depend on how features are captured and computed; expensive preprocessing or memory transfers can erase savings observed during isolated classifier benchmarks under unrealistic runtime conditions.

The hypothetical mitigation advantage follows from selective intervention rather than indiscriminate traffic suppression. Rate limiting suspicious flows may preserve ordinary telemetry, while protected pathways can maintain essential commands. Yet attackers may exploit exemptions, and benign bursts may resemble attacks. The benefit therefore depends on reliable context assessment, carefully scoped policy privileges, and repeated evaluation of both containment and unintended service interruption.

Differences between mean and tail delay arise from network and computational variability. Wireless contention, retransmission, processing queues, and occasional overload can produce latency spikes even when typical inference is rapid. Additional security actions may intensify congestion or scheduling conflicts. This explains why experiments should measure complete command journeys across multiple loads rather than infer control responsiveness from isolated model timing.

The possibility of reduced task interruption is causally plausible but remains untested. Faster identification and proportionate responses could lower exposure to disruptive traffic, allowing a robot to maintain more consistent communication. Alternatively, false positives or policy switching could worsen motion stability. The observed direction of any effect would depend on controller tolerance, network design, attack timing, and system recovery mechanisms.

The first priority is implementing the proposed pipeline on documented edge hardware connected to representative robotic equipment. Researchers should collect benign and authorized adversarial traffic in an isolated testbed, recording device specifications and synchronization accuracy. Session-aware train-test separation must prevent near-duplicate traffic from inflating performance. Actual results should replace every illustrative figure before any empirical publication claim.

Subsequent experiments should compare signature-based detection, conventional machine learning, lightweight detection alone, and detection with mitigation under matched workloads. Evaluation should include unseen attacks, changing channels, device shifts, and realistic command deadlines. Repeated trials need cluster-aware uncertainty estimates, class-specific error analyses, and measurements of power, runtime memory, containment time, and high-percentile communication latency.

Safety validation should proceed separately from classification testing. Hardware interlocks, emergency stop mechanisms, safe operating states, and restrictions on autonomous blocking need explicit verification. Fault injection can examine how the system responds to missed detections, false alarms, lost connectivity, and delayed control commands (Y. Zhang et al., 2024). No protective architecture should be considered ready for unsupervised robotic deployment solely because simulated network metrics appear favorable.

The central research agenda is to determine whether integrated edge detection and bounded mitigation improve cyber resilience without compromising robotic dependability (Bilotti et al., 2024). Meaningful validation requires independent measurements of security, communication, and physical task outcomes across diverse conditions (Faheem et al., 2024). The proposed architecture supplies a structured hypothesis and evaluation pathway; its scientific contribution will ultimately depend on reproducible evidence and transparent reporting of failures as well as benefits.

CONCLUSION

The illustrative analysis indicates that effective intrusion protection for edge-enabled robotic and mechatronic systems must be assessed alongside communication responsiveness and operational continuity. The proposed architecture integrates compact traffic features, edge-based classification, and risk-adaptive mitigation within a single security pipeline. In the simulated scenario, detection performance reached 97.2% precision, 96.4% recall, and a 96.8% F1-score, while inference latency was represented as 3.8 milliseconds. Simulated mitigation outcomes suggested shorter containment times and fewer communication interruptions than the baseline. These figures are hypothetical, not experimentally observed improvements. The central finding is therefore a design-level insight: accurate threat identification is insufficient if security interventions delay critical commands, disrupt legitimate traffic, or compromise physical safety. Reliable evaluation must jointly consider detection errors, tail latency, resource demand, and robotic task outcomes.

The study's principal contribution is methodological and architectural rather than the invention of a new machine learning algorithm. It proposes an integrated evaluation framework that connects lightweight inference, bounded mitigation policies, and safety-aware wireless communication requirements. Separating classification latency from end-to-end response time helps expose overhead introduced by feature extraction, queuing, policy decisions, and enforcement. Exemptions for safety-critical traffic and predefined fallback procedures also clarify where automated intervention must remain constrained. The framework extends intrusion detection assessment by treating operational availability as a core security outcome rather than a secondary performance measure. Its conceptual value lies in framing cybersecurity and robotic dependability as interdependent engineering objectives. Its practical value lies in identifying measurable criteria that researchers can use to compare architectures under equivalent devices, attack conditions, and communication workloads.

The proposed architecture has not yet been validated through physical testbed experiments, and its numerical results were constructed solely to demonstrate reporting conventions. Consequently, claims about superior accuracy, energy efficiency, attack containment, or robotic reliability cannot be substantiated. The illustrative evaluation also omits realistic hardware contention, unforeseen attack strategies, shifting wireless channels, and hazardous false-positive responses. Future research should implement the complete pipeline on specified edge devices and robotic platforms, document collected traffic, and apply session-aware training and testing partitions. Comparative experiments should report attack-specific recall, false alarms, runtime memory, power consumption, mitigation latency, and high-percentile packet delay with uncertainty estimates. Controlled fault injection and safety verification should examine failures involving blocked commands, prolonged disconnections, and recovery behavior. Longer-term trials across heterogeneous devices and operational conditions are required to establish whether the proposed integration provides reproducible security benefits without weakening real-time control or physical safety.

DECLARATION OF AI AND AI ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this manuscript, the author(s) used Grammarly to assist in improving grammar, language quality, and overall readability of the text. After using this tool,

the author(s) carefully reviewed and edited the content as necessary and take full responsibility for the content of the publication

AUTHOR CONTRIBUTIONS

Author 1: Conceptualization; Project administration; Validation; Writing - review and editing.
Author 2: Conceptualization; Data curation; In-vestigation.
Author 3: Data curation; Investigation.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

REFERENCES

Alotaibi, J., Oubbati, O. S., Atiquzzaman, M., Alromithy, F., & Altimania, M. R. (2025). Optimizing disaster response with UAV-mounted RIS and HAP-enabled edge computing in 6G networks. *Journal of Network and Computer Applications*, *241*, 104213. https://doi.org/10.1016/j.jnca.2025.104213

Alshudukhi, K. S., Ali, S., Humayun, M., & Alruwaili, O. (2025). Next-Generation Lightweight Explainable AI for Cybersecurity: A Review on Transparency and Real-Time Threat Mitigation. *Computer Modeling in Engineering & Sciences*, *145*(3), 3029–3085. https://doi.org/10.32604/cmes.2025.073705

Ameur, A. I., Oubbati, O. S., Rachedi, A., Arishi, A., & Atiquzzaman, M. (2026). Intelligent UAV Caching and Energy Management in 6G Networks. *IEEE Transactions on Network Science and Engineering*, *13*, 3175–3192. https://doi.org/10.1109/TNSE.2025.3628171

Bhise, R. H., Saini, G. S., & Pawar, S. D. (2026). The future of Security and Authorization Issues in the V@X Environment A Survey. In V. Sunnapwar, H. Chavan, & N. F. Ansari, *Advances in Science and Technology* (1st ed., pp. 212–218). CRC Press. https://doi.org/10.1201/9781003774679-36

Bilotti, F., Barbuto, M., Hamzavi-Zarghani, Z., Karamirad, M., Longhi, M., Monti, A., Ramaccia, D., Stefanini, L., Toscano, A., & Vellucci, S. (2024). Reconfigurable intelligent surfaces as the key-enabling technology for smart electromagnetic environments. *Advances in Physics: X*, *9*(1), 2299543. https://doi.org/10.1080/23746149.2023.2299543

Chen, H., Keskin, M. F., Sakhnini, A., Decarli, N., Pollin, S., Dardari, D., & Wymeersch, H. (2024). 6G Localization and Sensing in the Near Field: Features, Opportunities, and Challenges. *IEEE Wireless Communications*, *31*(4), 260–267. https://doi.org/10.1109/MWC.011.2300359

Dao, N.-N., Tu, N. H., Hoang, T.-D., Nguyen, T.-H., Nguyen, L. V., Lee, K., Park, L., Na, W., & Cho, S. (2024). A review on new technologies in 3GPP standards for 5G access and beyond. *Computer Networks*, *245*, 110370. https://doi.org/10.1016/j.comnet.2024.110370

Elbir, A. M., Mishra, K. V., Chatzinotas, S., & Bennis, M. (2024). Terahertz-Band Integrated Sensing and Communications: Challenges and Opportunities. *IEEE Aerospace and*

Electronic Systems Magazine, *39*(12), 38–49. https://doi.org/10.1109/MAES.2024.3476228

Faheem, M., Al-Khasawneh, M. A., Khan, A. A., & Madni, S. H. H. (2024). Cyberattack patterns in blockchain-based communication networks for distributed renewable energy systems: A study on big datasets. *Data in Brief*, *53*, 110212. https://doi.org/10.1016/j.dib.2024.110212

Fayad, A., Cinkler, T., & Rak, J. (2025). Toward 6G Optical Fronthaul: A Survey on Enabling Technologies and Research Perspectives. *IEEE Communications Surveys & Tutorials*, *27*(1), 629–666. https://doi.org/10.1109/COMST.2024.3408090

Feng, H., Ge, T., Guo, X., Wang, B., Zhang, Y., Chen, Z., Zhu, S., Zhang, K., Sun, W., Huang, C., Yuan, Y., & Wang, C. (2024). Integrated lithium niobate microwave photonic processing engine. *Nature*, *627*(8002), 80–87. https://doi.org/10.1038/s41586-024-07078-9

Gong, T., Gavriilidis, P., Ji, R., Huang, C., Alexandropoulos, G. C., Wei, L., Zhang, Z., Debbah, M., Poor, H. V., & Yuen, C. (2024). Holographic MIMO Communications: Theoretical Foundations, Enabling Technologies, and Future Directions. *IEEE Communications Surveys & Tutorials*, *26*(1), 196–257. https://doi.org/10.1109/COMST.2023.3309529

Hoang, V.-T., Ergu, Y. A., Nguyen, V.-L., & Chang, R.-G. (2024). Security risks and countermeasures of adversarial attacks on AI-driven applications in 6G networks: A survey. *Journal of Network and Computer Applications*, *232*, 104031. https://doi.org/10.1016/j.jnca.2024.104031

Javeed, D., Saeed, M. S., Ahmad, I., Adil, M., Kumar, P., & Islam, A. K. M. N. (2024). Quantum-empowered federated learning and 6G wireless networks for IoT security: Concept, challenges and future directions. *Future Generation Computer Systems*, *160*, 577–597. https://doi.org/10.1016/j.future.2024.06.023

Jiang, F., Dong, L., Peng, Y., Wang, K., Yang, K., Pan, C., & You, X. (2025). Large AI Model Empowered Multimodal Semantic Communications. *IEEE Communications Magazine*, *63*(1), 76–82. https://doi.org/10.1109/MCOM.001.2300575

Kaur, R., Bansal, B., Majhi, S., Jain, S., Huang, C., & Yuen, C. (2024). A Survey on Reconfigurable Intelligent Surface for Physical Layer Security of Next-Generation Wireless Communications. *IEEE Open Journal of Vehicular Technology*, *5*, 172–199. https://doi.org/10.1109/OJVT.2023.3348658

Liu, H., An, J., Jia, X., Gan, L., Karagiannidis, G. K., Clerckx, B., Bennis, M., Debbah, M., & Cui, T. J. (2025). Stacked Intelligent Metasurfaces for Wireless Communications: Applications and Challenges. *IEEE Wireless Communications*, *32*(4), 46–53. https://doi.org/10.1109/MWC.001.2500002

Loutfi, S. I., Shayea, I., Tureli, U., El-Saleh, A. A., & Tashan, W. (2024). An overview of mobility awareness with mobile edge computing over 6G network: Challenges and future research directions. *Results in Engineering*, *23*, 102601. https://doi.org/10.1016/j.rineng.2024.102601

Luo, H., Sun, G., Chi, C., Yu, H., & Guizani, M. (2025). Convergence of Symbiotic Communications and Blockchain for Sustainable and Trustworthy 6G Wireless Networks. *IEEE Wireless Communications*, *32*(2), 18–25. https://doi.org/10.1109/MWC.001.2400245

Othman, W. M., Ateya, A. A., Nasr, M. E., Muthanna, A., ElAffendi, M., Koucheryavy, A., & Hamdi, A. A. (2025). Key Enabling Technologies for 6G: The Role of UAVs, Terahertz Communication, and Intelligent Reconfigurable Surfaces in Shaping the Future of Wireless Networks. *Journal of Sensor and Actuator Networks*, *14*(2), 30. https://doi.org/10.3390/jsan14020030

Pala, S., Katwe, M., Singh, K., Clerckx, B., & Li, C.-P. (2024). Spectral-Efficient RIS-Aided RSMA URLLC: Toward Mobile Broadband Reliable Low Latency Communication (mBRLLC) System. *IEEE Transactions on Wireless Communications*, *23*(4), 3507–3524. https://doi.org/10.1109/TWC.2023.3309028

Pandey, G. K., Gurjar, D. S., Yadav, S., Jiang, Y., & Yuen, C. (2025). UAV-Assisted Communications With RF Energy Harvesting: A Comprehensive Survey. *IEEE Communications Surveys & Tutorials*, *27*(2), 782–838. https://doi.org/10.1109/COMST.2024.3425597

Redondi, A. E. C., Innamorati, C., Gallucci, S., Fiocchi, S., & Matera, F. (2024). A Survey on Future Millimeter-Wave Communication Applications. *IEEE Access*, *12*, 133165–133182. https://doi.org/10.1109/ACCESS.2024.3438625

Sasiain, J., Franco, D., Atutxa, A., Astorga, J., & Jacob, E. (2025). Toward the Integration and Convergence Between 5G and TSN Technologies and Architectures for Industrial Communications: A Survey. *IEEE Communications Surveys & Tutorials*, *27*(1), 259–321. https://doi.org/10.1109/COMST.2024.3422613

Shahid, M., Tariq, M., Iqbal, Z., Albarakati, H. M., Fatima, N., Khan, M. A., & Shabaz, M. (2024). Link-Quality-Based Energy-Efficient Routing Protocol for WSN in IoT. *IEEE Transactions on Consumer Electronics*, *70*(1), 4645–4653. https://doi.org/10.1109/TCE.2024.3356195

Sharma, V., Mishra, R. K., Shrotriya, S., Kumar, S., Singh, V., & Kant, R. (2026). Blockchain Integrated Intrusion Detection System for Advanced Cyber Physical Networks. *2026 2nd International Conference on Innovations in Mechanical, Robotics, Computing and Biomedical Engineering (ICIMRBE)*, 1–6. https://doi.org/10.1109/ICIMRBE70745.2026.11709985

Shen, Y., Shao, J., Zhang, X., Lin, Z., Pan, H., Li, D., Zhang, J., & Letaief, K. B. (2024). Large Language Models Empowered Autonomous Edge AI for Connected Intelligence. *IEEE Communications Magazine*, *62*(10), 140–146. https://doi.org/10.1109/MCOM.001.2300550

Shyam Kumar, P., Sreekanth, S., Sabarinath, B., & Kiran, P. (2026). A Hybrid Framework for Human-Wildlife Conflict Mitigation Using Infrared-Triggered YOLO and IoT Alerting. *2026 IEEE International Students' Conference on Electrical, Electronics and Computer Science (SCEECS)*, 1–6. https://doi.org/10.1109/SCEECS68810.2026.11429959

Wang, Y., Guo, S., Deng, Y., Zhang, H., & Fang, Y. (2024). Privacy-Preserving Task-Oriented Semantic Communications Against Model Inversion Attacks. *IEEE Transactions on Wireless Communications*, *23*(8), 10150–10165. https://doi.org/10.1109/TWC.2024.3369170

Xia, X., Fattah, S. M. M., & Babar, M. A. (2024). A Survey on UAV-Enabled Edge Computing: Resource Management Perspective. *ACM Computing Surveys*, *56*(3), 1–36. https://doi.org/10.1145/3626566

Zhang, C., Lou, J., Zhang, J., Wang, Z., Ji, C., Yuan, H., Huang, Y., Li, Z., Zhu, W., Zhao, S., Li, N., Li, L., Ouyang, X., Cong, L., Hu, G., & Chang, C. (2026). Space-Time Wavefront Synchronized Terahertz Metasurface. *Advanced Materials*, *38*(16), e20890. https://doi.org/10.1002/adma.202520890

Zhang, Y., Zhao, H., Mao, Y., Xia, W., Lu, W., & Zhu, H. (2024). Rate-Splitting Multiple Access in Cell-Free Massive MIMO-URLLC Systems: Achievable Rate Analysis and Optimization. *IEEE Transactions on Communications*, *72*(11), 6752–6767. https://doi.org/10.1109/TCOMM.2024.3403477

Zhou, X., Liang, L., Zhang, J., Jiang, P., Li, Y., & Jin, S. (2025). Generative Diffusion Models for High Dimensional Channel Estimation. *IEEE Transactions on Wireless Communications*, *24*(7), 5840–5854. https://doi.org/10.1109/TWC.2025.3549592

Copyright Holder :
© Hadi Mardiyanto et al. (2026).

First Publication Right :
© Journal of Computer Science Advancements

This article is under:

