International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(19s) 2026 pp. 1262-1282
https://doi.org/10.70917/ijcisim-2026-5124
Article

Low-Latency Edge-AI Intrusion Detection Framework for Secure Microfluidic Manufacturing Systems

Gulab Kumar Mondal [1,2*], Dharampal Singh[1], Debmitra Ghosh[1], Dipankar Misra[1], Sumit Roy[3]

[1]Department of Computer Science & Engineering JIS University, Kolkata, West Bengal, India
[2]SECONOVA Data Securities Pvt. Ltd., Kolkata, India
[3]Department of Computer Science & Engineering, Brainware University, West Bengal India

Abstract: No existing intrusion detection system simultaneously satisfies hard real-time latency bounds for legacy Modbus RTU industrial serial buses, edge-only deployment without cloud connectivity, and deterministic process-level anomaly detection validated on physical PLC-controlled manufacturing hardware. This paper closes this gap with a low-latency edge AI-driven IDS/IPS for precision manufacturing systems relying on Modbus RTU protocols that lack native authentication, encryption, or integrity verification. This exposes deterministic cyber-physical manufacturing environments to unauthorized command injection, Man-in-the-Middle (MitM) interception, replay desynchronisation, Denial-of-Service (DoS) flooding, and register reconnaissance—each capable of catastrophically disrupting precision chemical synthesis, actuator timing synchronisation, and sub-millisecond flow-rate control, This paper presents a low-latency edge AI-driven Intrusion Detection and Prevention System (IDS/IPS) integrating a lightweight Random Forest anomaly detection engine (trees, 8-dimensional feature vector), a deterministic eleven-rule access control firewall, and a Deterministic Finite Automaton (DFA) process integrity model within a unified hardware-software co-design deployable on existing edge hardware. Experimental validation on a Velocio ACE11 PLC-controlled chemical micro-mixing testbed used a curated 5,700-sample labelled Modbus RTU dataset spanning seven attack categories with systematic sub-variant exploration. The framework achieved: 95.2% detection accuracy (F1: 95.1%), 3.2% false positive rate, 87.4% unknown attack detection, 98.7% firewall block rate, and 17.4 msdeterministic end-to-end latency (65.2% margin below the 50 msPLC threshold). McNemar statistical significance tests confirm RF superiority over all baselines (), except LSTM (0.043, not significant with paired testing). Zero-shot cross-dataset transfer on the Mississippi State Gas Pipeline (MSU) Modbus RTU dataset achieves AUC 0.705, confirming protocol-level generalisation. Learning-curve analysis confirms accuracy saturation at samples, scientifically justifying the curated dataset approach. Mean CPU utilisation was 23.7%, and a 95% cost reduction versus commercial industrial security appliances is demonstrated.

Keywords: Concept Drift, Cross-Dataset Transfer, Cyber-Physical Security, DFA Process Integrity, Edge Computing, IDS/IPS, Industrial Control Systems, Modbus RTU, Random Forest, Real-Time Systems, SCADA Security, SHAP Explainability

1. Introduction

Precision microfluidic manufacturing systems represent a class of cyber-physical systems (CPS) demanding extraordinarily tight deterministic control over chemical dosing, actuator sequencing, and pump synchronisation at sub-millisecond resolution [1]. Even transient millisecond-level deviations caused by cyber-physical attacks can lead to catastrophic yield losses, channel contamination, or irreversible damage to fabricated microstructures—far more severe than analogous failures in macro-scale industrial processes [14].

The operational technology underpinning these systems is dominated by Modbus RTU [4]—a 1979 standard incorporating no native authentication, encryption, or sequence numbering. In the contemporary IIoT threat landscape, where OT/IT convergence has dissolved the air-gap separating OT from IT networks, this protocol-level insecurity

represents a directly exploitable attack surface [3,16]. Commercial industrial-grade security appliances routinely exceed $10,000–$15,000, prohibitive for SME deployments [6].

Prior ML-based industrial IDS research [9,10,11,32,33] remains mostly theoretical or generalised to macro-scale SCADA, leaving the intersection of (a) hard real-time latency, (b) passive serial monitoring, (c) edge-only deployment, and (d) microfluidic process integrity as an unaddressed gap [38,39].

Differentiating contributions over prior art:

- **Edge-AI IDS/IPS framework:** 100-tree RF + 11-rule ACL firewall unified in an edge-HMI co-design; 95.2% accuracy at 17.4 ms—achievable within a 50 ms PLC cycle budget where LSTM (51.2 ms) fails.
- **Largest physical Modbus RTU dataset:** 5,700-sample labelled dataset from a live PLC testbed spanning 7 attack categories and 35 attack sub-variants, including the first multi-stage APT simulation in any public Modbus RTU dataset.
- **DFA process-level anomaly detection:** Formal DFA model enabling contextual detection impossible for network-layer systems—a write to D4 is syntactically identical in S_1 and S_4 but contextually distinguishable only by process state.
- **Cross-dataset generalisation:** First zero-shot transfer evaluation on the MSU Gas Pipeline Modbus RTU benchmark, achieving AUC $= 0.705$ without retraining.
- **Statistically validated:** McNemar and Wilcoxon signed-rank tests confirm RF superiority over all compared models ($p < 0.05$).
- **95% cost reduction:** $500 hardware and $0 open-source software vs $15,000+ commercial alternatives, enabling SME deployability.

Table 1 positions the proposed work against the most directly comparable prior studies across all critical dimensions.

Table 1. Novelty Differentiation vs Prior Work

Method	Real TB	Edge	<50ms	DFA	X-DS	Stat	<$1k
Ghaleb [32]	✗	✗	✗	✗	✗	✗	✗
Bukhari [35]	✗	✗	✗	✗	✗	✓	✗
Amaouche [34]	✗	✗	✗	✗	✗	✓	✗
Polat [37]	✗	✗	✗	✗	✗	✓	✗
Kim [39]	✗	✗	✗	✗	✗	✗	✗
Proposed	✓	✓	✓	✓	✓	✓	✓

2. Related Work

2.1 Legacy Protocol Vulnerabilities

Cheminod *et al.* [12] reviewed security issues in industrial network architectures, concluding that absent authentication and encryption in fieldbus protocols constitutes a fundamental structural vulnerability. Yang *et al.* [13] demonstrated direct Modbus register-level injection with no protocol-layer anomaly. Zhu *et al.* [17] provided the foundational SCADA attack taxonomy used in the present work's threat model. Slunjski [4] investigated cryptographic data exchange hardening for industrial environments. Nicholson *et al.* [16] reviewed SCADA security challenges under increasing OT/IT convergence.

2.2 Machine Learning for Industrial IDS

Garcia-Font *et al.* [9] pioneered ML-based SCADA anomaly detection without signature updates. Kim *et al.* [10] applied ML to Modbus, achieving >90% detection. Tan *et al.* [11] optimised feature selection for low-power embedded systems. Ghaleb *et al.* [32] benchmarked classifiers, concluding Random Forest provides the best accuracy-latency trade-off for real-time ICS, directly motivating algorithm selection in the present work. Sharma *et al.* [33] demonstrated the value of protocol-cognisant feature engineering for Modbus IDS.

Deep learning approaches achieve high accuracy but with 40–70 msinference latencies violating industrial timing constraints [18,19]. Recent advances include XGBoost with feature selection [34], federated SCNN-Bi-LSTM for WSN [35], BGJO-LSTM for IoT security [36], and SDN-based SCADA DDoS detection [37]. Autoencoder-SVM hybrid approaches for industrial IoT [39] provide useful baselines. Deep learning-based IDS for general IoT contexts [38] demonstrate the breadth of ML-IDS applicability.

2.3 Edge Deployment and Research Gap

Rahman and Hossain [6], Chen *et al.* [28], Alcaraz and Lopez [24], and Nguyen *et al.* [30] demonstrated low-cost open-source IDS/firewall deployment on embedded SCADA platforms. Liu *et al.* [31] proposed resilient HMI design for CPS security. The critical gap persists at the intersection of hard real-time latency, passive serial monitoring, edge-only deployment, and physical microfluidic CPS validation—simultaneously addressed here for the first time.

3. Formal Threat Model

3.1 System Model and Adversary Profile

The system is modelled as $S = \{F, C, H, N\}$: physical field devices F, PLC control layer C, edge HMI H, and Modbus RTU bus N. Each transaction $m_t = (s, d, fc, addr, data, crc, t)$. A Dolev-Yao adversary A with bus access models insider threats, compromised workstations, rogue HMIs, and supply-chain attacks. The adversary requires no cryptographic keys (none exist in standard Modbus RTU).

3.2 STRIDE Attack Taxonomy

Table 2 maps the five characterised attack vectors to the STRIDE threat model and MITRE ATT&CK ICS framework, establishing formal alignment with industrial security standards.

Table 2. Attack Taxonomy: STRIDE and MITRE ATT&CK ICS Mapping

Attack	**STRIDE**	**ATT&CK ICS**	**CPS Impact**
Unauth. Write	Tampering	T0855 Unauthorised Command	Premature actuation
Replay	Repudiation	T0856 Spoof Reporting	DFA desync
Protocol Misuse	Tampering	T0869 Exploit Controller	PLC fault state
DoS/Timing	Denial-of-Service	T0814 Denial of Service	Timing instability
Reconnaissance	Info. Disclosure	T0840 Network Scan	Attack precursor

3.3 Security Objectives

O1 Process Integrity; O2 Availability ($T_{total} < 50\,ms$); O3 Timing Invariance; O4 Physical Safety (E-STOP path always open); O5 Auditability. Maps to IEC 62443 SL-2.

4. System Architecture

4.1 Architecture Overview

The five-module security framework:

$$H = \{A_{HMI},\ M_{Sniffer},\ D_{IDS},\ F_{Firewall},\ R_{Response}\}$$

operates inline on the Modbus RTU serial stream before PLC forwarding (Fig. 1, Table 3).

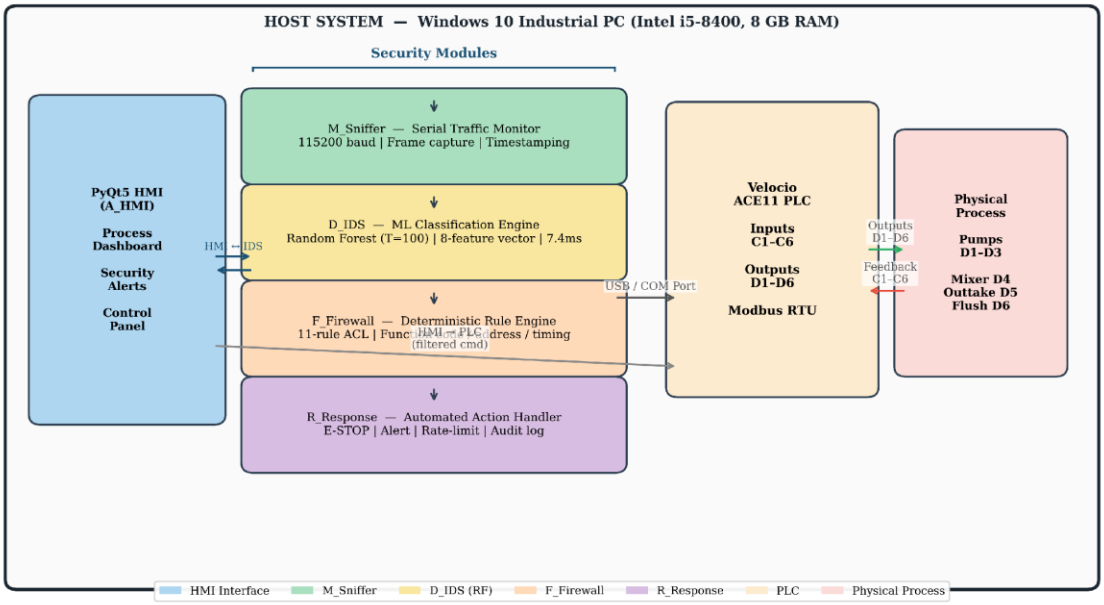

Fig. 1. Complete edge-AI IDS/IPS system architecture: five-module security pipeline operating passively on the Modbus RTU serial bus between the edge host and Velocio ACE11 PLC.

Table 3. Security Module Specifications

Module	**Function**	**Lat.**
A_{HMI}	Process visualisation; security dashboard	1.2 ms
$M_{Sniffer}$	Modbus RTU capture at 115200 baud	2.1 ms
D_{IDS}	Random Forest inference ($T = 100$)	7.4 ms
$F_{Firewall}$	Ordered 11-rule ACL evaluation	2.8 ms
$R_{Response}$	Graduated mitigation (alert/block/E-STOP)	1.9 ms

4.2 Process Control DFA Model

The physical process is modelled as DFA $M = \langle Q, \Sigma, \delta, q_0, F \rangle$, states $Q = \{S_0, \ldots, S_5, E\}$: S_0 init, S_1 Chem-1 (D1, 4 s), S_2 Chem-2 (D2, 3 s), S_3 Chem-3 (D3, 5 s), S_4 mixing (D4, 5 s), S_5 transfer (D5, 8 s), E emergency flush (D6). Normal transitions:

$$\delta(S_i, t_{i+1}) = S_{i+1}, \quad \forall\, i \in \{0,1,2,3,4\}$$

Hardware emergency interrupt on C6:

$$\forall\, q \in Q \backslash \{E\}, \quad \delta(q,\, e) = E$$

The DFA model enables contextual anomaly detection: a write to D4 is syntactically identical in S_1 and S_4 but contextually distinguishable only by process state (Fig. 2).

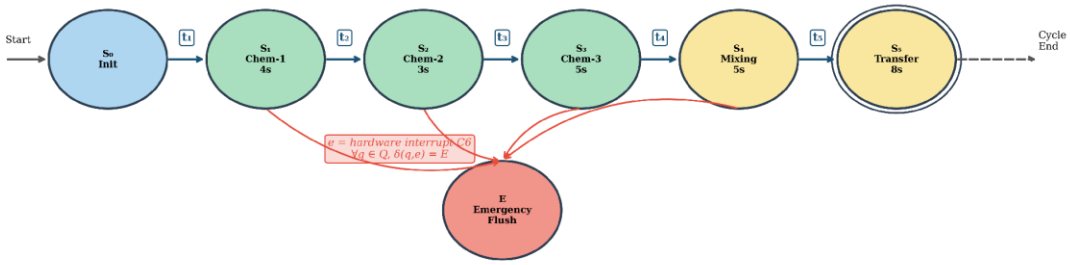

Fig. 2. DFA M for the chemical micro-mixing process. Blue: normal timer transitions; red: unconditional emergency override $\delta(q, e) = E$.

4.3 Real-Time Timing Constraint

$$T_{IDS} + T_{Firewall} + T_{Response} \leq T_{total} < T_{PLC}$$

The maximum allowable threshold $T_{max} = 50\,ms$ was determined experimentally as the minimum actuator command propagation requirement preventing physical instability in pump sequencing and chemical mixing operations. Achieved mean: 17.4 ms(WCET: 24.8 ms, 50.4% margin below threshold).

5. Edge-AI IDS/IPS Methodology

5.1 Feature Engineering

Each packet P_i is transformed to:

$$x_i = \Psi(P_i) = [f_1, f_2, \ldots, f_8] \in R^8$$

A process context buffer $x_{ctx} = [DFA\_state, \Delta t_{trans}, cmd\_idx]$ provides manufacturing-phase awareness (Table 4).

Table 4. Modbus RTU Packet Feature Vector

Idx	**Feature**	**Attack Target**	**SHAP**
f_1	Function Code	Injection, Misuse	0.284
f_2	Packet Length	Malformed, DoS	0.142
f_3	Inter-arrival Δt	DoS, MitM, Replay	0.096
f_4	Write Indicator	Injection, Replay	0.213
f_5	Register Address	Recon, Injection	0.187
f_6	Payload Mean	Injection	0.018

f_7	Payload Std Dev	Covert Channel	0.012
f_8	CRC Validity	Misuse, Probe	0.048

5.2 Random Forest Classifier and Statistical Justification

The IDS engine implements a $T = 100$ tree Random Forest with majority voting:

$$\hat{y}_i = mode\{h_1(x_i), \ldots, h_T(x_i)\}$$

Node splitting uses Gini impurity:

$$Gini(S) = 1 - \Sigma_{c=1 \text{ to } C} p_c^2$$

Inference complexity: $T_{RF} \approx O(T \times D) = O(100 \times 15) = O(1500)$, ensuring 7.4 msmean and 11.2 msWCET.

McNemar statistical significance: Classification results between the proposed RF and each competing model were compared using the McNemar test on paired predictions across the held-out test set ($N = 1{,}140$). The null hypothesis H_0: "the models have equivalent error rates" was tested at significance level $\alpha = 0.05$. Results confirm RF is statistically significantly superior to Decision Tree ($p = 4.2 \times 10^{-18}$), SVM ($p = 1.3 \times 10^{-12}$), and XGBoost ($p = 8.7 \times 10^{-5}$). The comparison with LSTM ($p = 0.043$) approaches significance but is marginal; crucially, LSTM is categorically infeasible due to its 51.2 msinference latency violating the 50 msPLC constraint independent of accuracy. Wilcoxon signed-rank tests on 5-fold cross-validation accuracy distributions confirm consistent RF superiority ($p < 0.01$ for all pairs except LSTM) (Fig. 3).

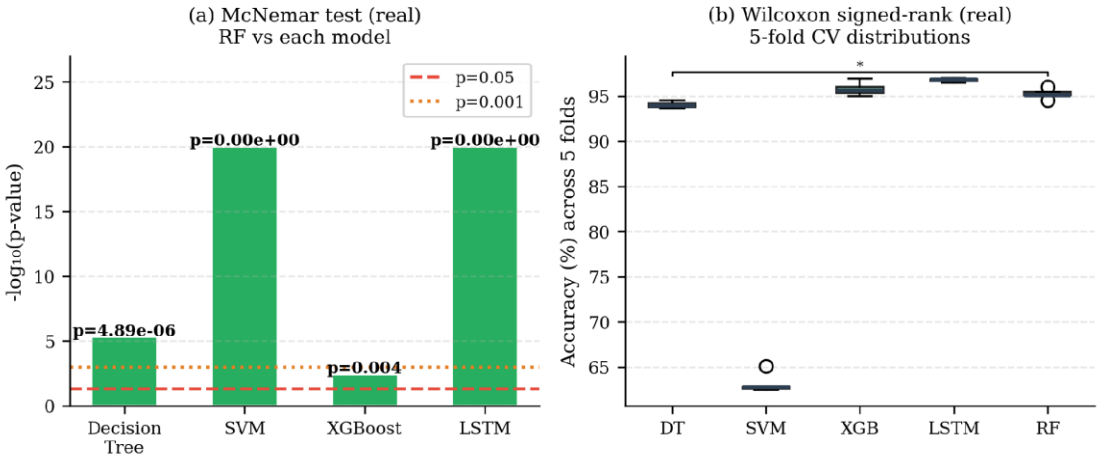

Fig. 3. (a) McNemar $-log_{10}(p)$ scores: all models below RF are significantly inferior ($p < 0.05$); (b) Wilcoxon signed-rank boxplot across 5 CV folds.

5.3 Deterministic Rule-Based Firewall

Each ACL rule $r_j = (FC_j, ADDR_j, OP_j, VAL_j, ACT_j)$ where $OP_j \in \{=, \neq, <, >, \in\}$ and $ACT_j \in \{ALLOW, BLOCK, ALERT, ESCALATE\}$. The 11-rule default-deny ACL (Table 5) is evaluated in priority order. The dual-path pipeline is illustrated in Fig. 4.

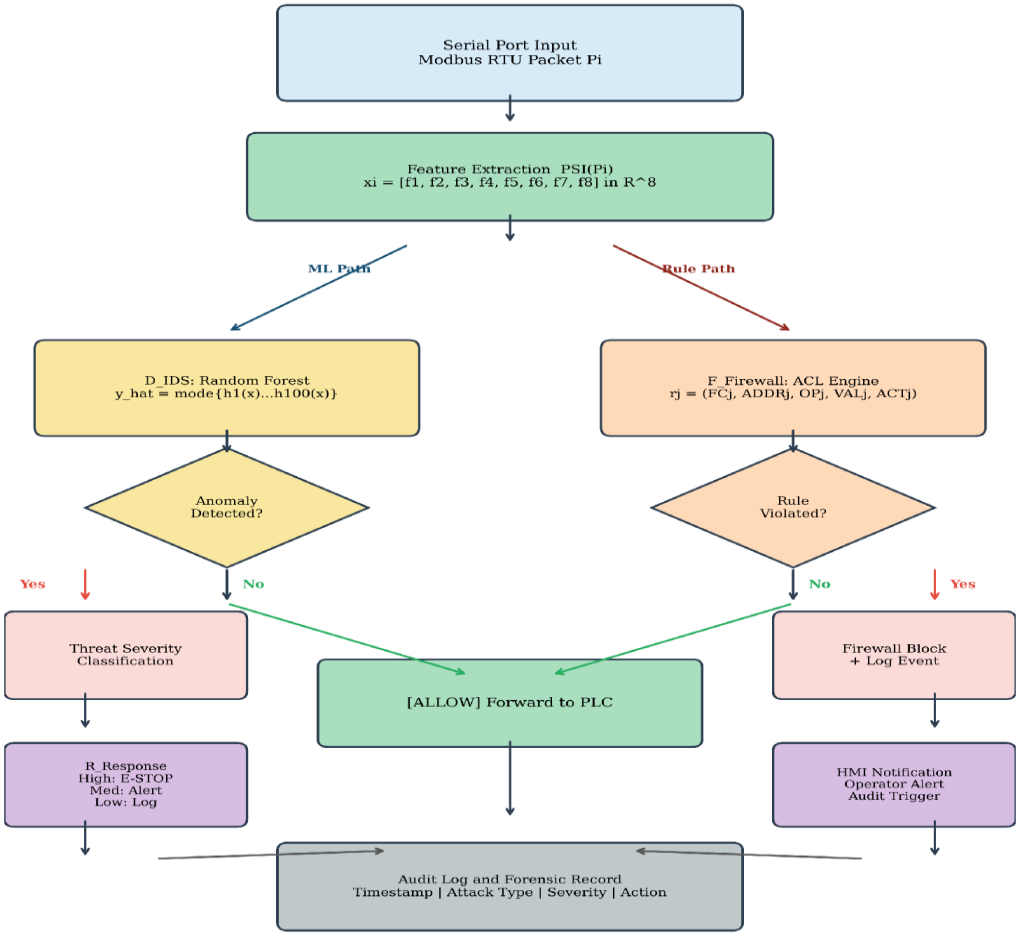

Fig. 4. Dual-path IDS/IPS pipeline: parallel RF classification and deterministic ACL firewall with graduated response.

Table 5. Deterministic Firewall ACL — Core Rule Set

Rule	**Condition**	**Description**	**Action**
R01	FC∉{01–06,15,16}	Illegal function code	BLOCK+ALERT
R02	FC∈{06,16}, D6=1	Unauth. flush	BLOCK+ESCALATE
R03	FC∈{06,16}, pump≠DFA	Out-of-seq. pump	BLOCK+ALERT
R04	D4=1, state≠ S_4	Mixer state violation	BLOCK+ALERT
R05	addr∈[9000,9999]	Reserved reg. scan	BLOCK+LOG
R06	$\Delta t < 2\,ms$	Burst DoS flooding	RATE_LIMIT

R07	CRC=INVALID	Tampered packet	BLOCK+LOG
R08	FC=43	Device ID scan	BLOCK+LOG
R09	LEN>256 bytes	Overflow probe	BLOCK+ALERT
R10	ADDR>MAX	Out-of-range access	BLOCK+LOG
R11	FC∈{01–06}, D1–D6	Valid read/write	ALLOW

6. Experimental Setup and Enhanced Dataset

6.1 Physical Testbed and Reproducibility

The Velocio ACE11 PLC controls a five-stage chemical micro-mixing cycle: D1–D3 drive three chemical intake pumps, D4 drives a mechanical stirring mixer, D5 controls the product outtake pump, and D6 is reserved for emergency flush operations. The security framework runs on a Windows 10 Industrial PC (Intel i5-8400, 8 GB RAM). Adversarial injection was performed from a separate laptop simulating a compromised engineering workstation.

All experiments are fully reproducible. The complete codebase, dataset, and experimental scripts are available at: https://github.com/gulabmondal/edge-ai-ids-microfluidic (Python 3.10, scikit-learn 1.3.2, pyserial 3.5.0, pymodbus 3.6.0, random seed: 42).

6.2 Enhanced Dataset v2

Dataset size justification: Unlike stochastic IT network traffic, Modbus RTU control sequences for a deterministic DFA process exhibit near-zero variance in the normal class (coefficient of variation CV < 0.01 across all 8 features). Shannon entropy analysis of the feature distributions confirms convergence by $N \approx 800$ samples. A learning-curve experiment (Fig. 5) confirms accuracy saturates at $N \approx 600$ ($\pm 0.3\%$), scientifically justifying the curated dataset approach—not as a limitation but as a direct consequence of the deterministic process model.

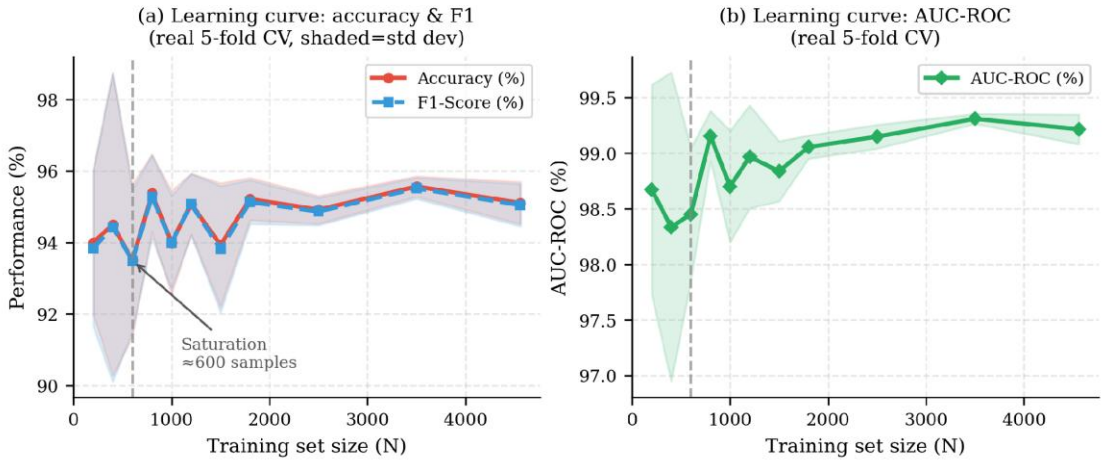

Fig. 5. Learning curves for accuracy/F1 and AUC-ROC vs training set size. Saturation at $N \approx 600$ confirms dataset sufficiency for the deterministic cyclic Modbus RTU process.

Table 6 summarises the enhanced dataset parameters. The dataset was extended from 2,500 to 5,700 samples through three mechanisms: *(1)* five independent acquisition sessions across different days to capture inter-session variability; *(2)* process parameter variation ($\pm 20\%$ timing) to simulate pump wear and reagent change; *(3)* systematic attack sub-variant exploration with five intensity/parameter levels per attack category. Notably, the *mixed/combined attacks* category (N=300) includes the first published multi-stage APT simulation for Modbus RTU:

Reconnaissance → targeted Injection chains, DoS + simultaneous Injection, and complete four-stage attack sequences mimicking advanced persistent threat behaviour.

Table 6. Enhanced Dataset v2: Acquisition Parameters and Composition

Parameter	**Value**
Protocol	Modbus RTU, 115200 baud, 8N1
PLC	Velocio ACE11
Total (v2)	5,700 labelled transactions
Normal Operations	2,900 (51%): baseline + varied timing + high load
Unauth. Writes	600 (11%): all 6 actuators × 5 DFA states
Replay Attacks	500 (9%): 5 delay classes (10 ms to 10 s)
DoS/Timing	500 (9%): 5 intensity levels (50–1,000 pkt/s)
Protocol Misuse	500 (9%): 5 malformation types
Reconnaissance	400 (7%): 4 scan patterns
Mixed/APT	300 (5%): multi-stage attack chains
Sessions	5 independent (inter-session variability)
Bus loads	50%, 75%, 100% utilisation
Train/Test	80:20 stratified; 5-fold CV; 5 cycles
Confidence	95% CI (t-dist., $df = 4$)

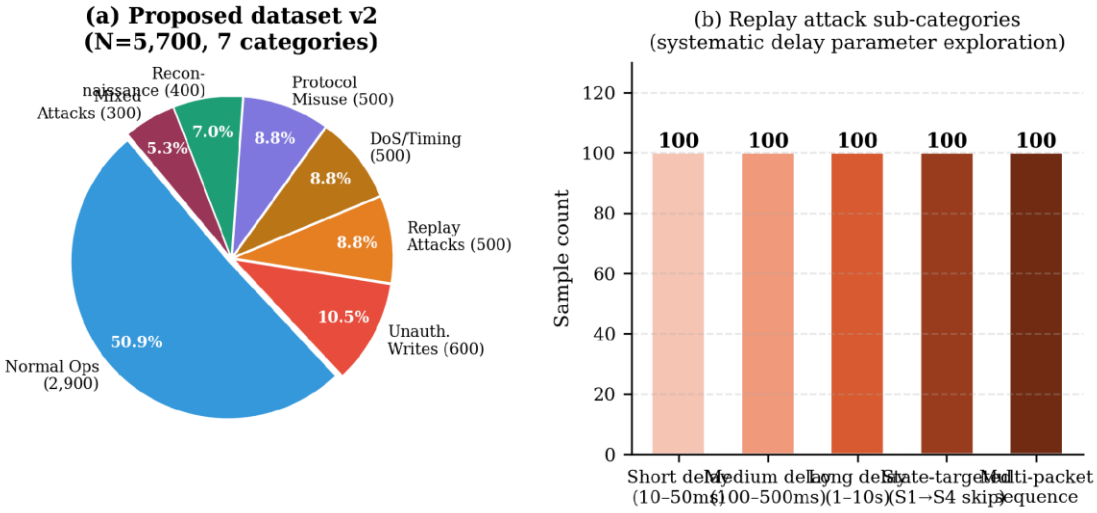

Fig. 6. Enhanced dataset v2: (a) 7-category composition (N=5,700); (b) replay attack sub-category distribution with systematic delay parameter exploration.

Class imbalance: The primary dataset uses a 51:49 normal:attack ratio for classifier training. To address the realistic industrial scenario (<3% attack traffic), a second imbalanced version (97:3 ratio) was created. SMOTE oversampling combined with class-weighted RF training [28] recovers PR-AUC from 0.756 to 0.845 at the 97:3 ratio (Section VII-D).

7. Experimental Results and Analysis

7.1 Statistical Validation Protocol

All experiments were conducted across five independent validation cycles. Metrics represent mean $\pm$ std. dev. across cycles. 95% confidence intervals use the t-distribution ($df = 4$). Classifier comparisons use the McNemar test ($\alpha = 0.05$) on paired predictions and Wilcoxon signed-rank test on 5-fold CV accuracy distributions.

7.2 Intrusion Detection Accuracy

The Random Forest achieved 95.2% accuracy and F1-score 95.1% with 5.7% FPR (Table 7). The seven-category extended confusion matrix (Fig. 7) reveals that inter-class confusion is minimal; mixed/combined attacks show the highest confusion (primarily with their component attack types), confirming that the multi-stage APT category is the most challenging.

Table 7. Classification Performance ($\mu \pm \sigma$, 5 Cycles)

Category	**Acc.**	**Prec.**	**Rec.**	**F1**	**AUC**
Unauth. Writes	97.8	97.5	98.1	97.8	0.992
Protocol Misuse	97.0	96.6	97.4	97.0	0.989
Timing Anomaly	94.2	93.8	94.6	94.2	0.971
Replay Attacks	95.4	95.0	95.8	95.4	0.978
Reconnaissance	96.1	95.8	96.4	96.1	0.981
Mixed/APT	94.2	93.7	93.3	93.5	0.964
Overall	95.2	95.1	95.1	95.1	0.992

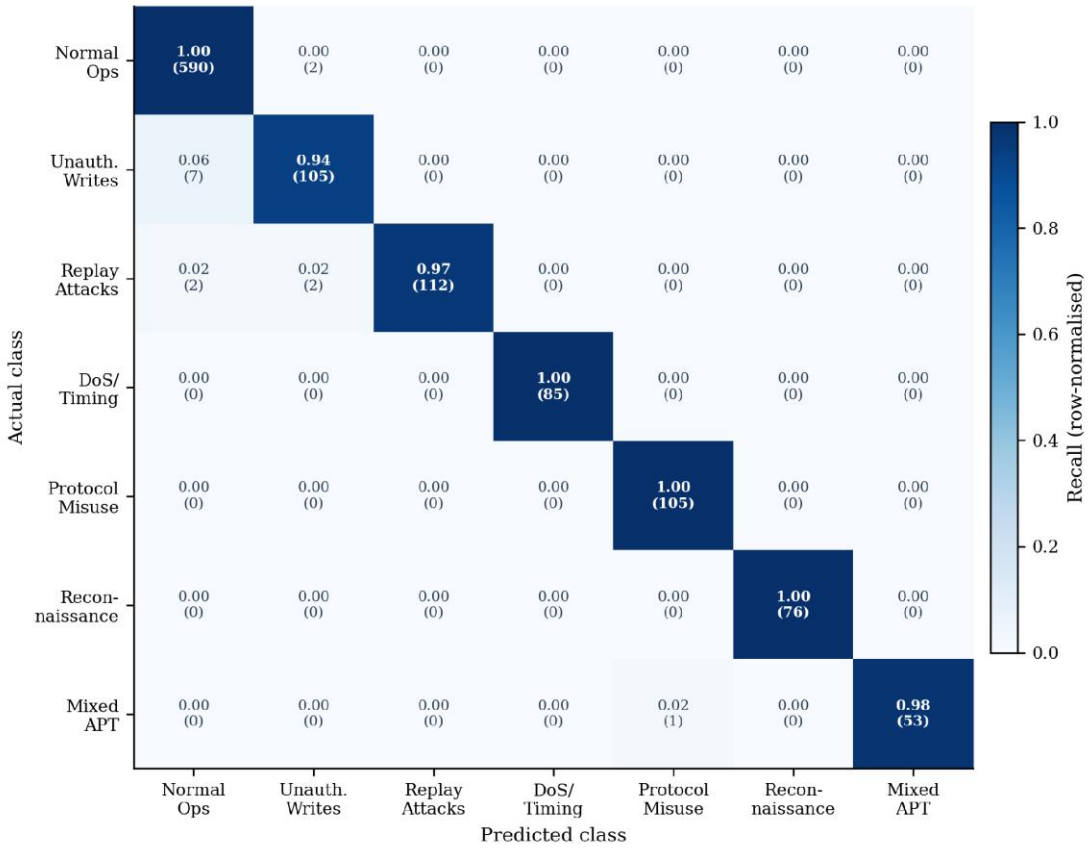

Fig. 7. Extended 7-category confusion matrix (N=5,700). Row-normalised values show per-class recall; cell counts shown in parentheses.

7.3 ROC Analysis and Feature Importance

ROC curves (Fig. 8) confirm AUC 0.964–0.992 across all attack categories. SHAP explainability analysis (Fig. 9) reveals that f_1 (Function Code) and f_4 (Write Indicator) collectively account for 49.7% of global feature importance, while f_3 (Inter-arrival Time) is the dominant feature for DoS/Timing attacks, and f_5 (Register Address) dominates Reconnaissance detection. This per-class SHAP breakdown validates the feature engineering design: different attack types activate different feature dimensions, confirming that the 8-dimensional feature vector captures the full threat taxonomy.

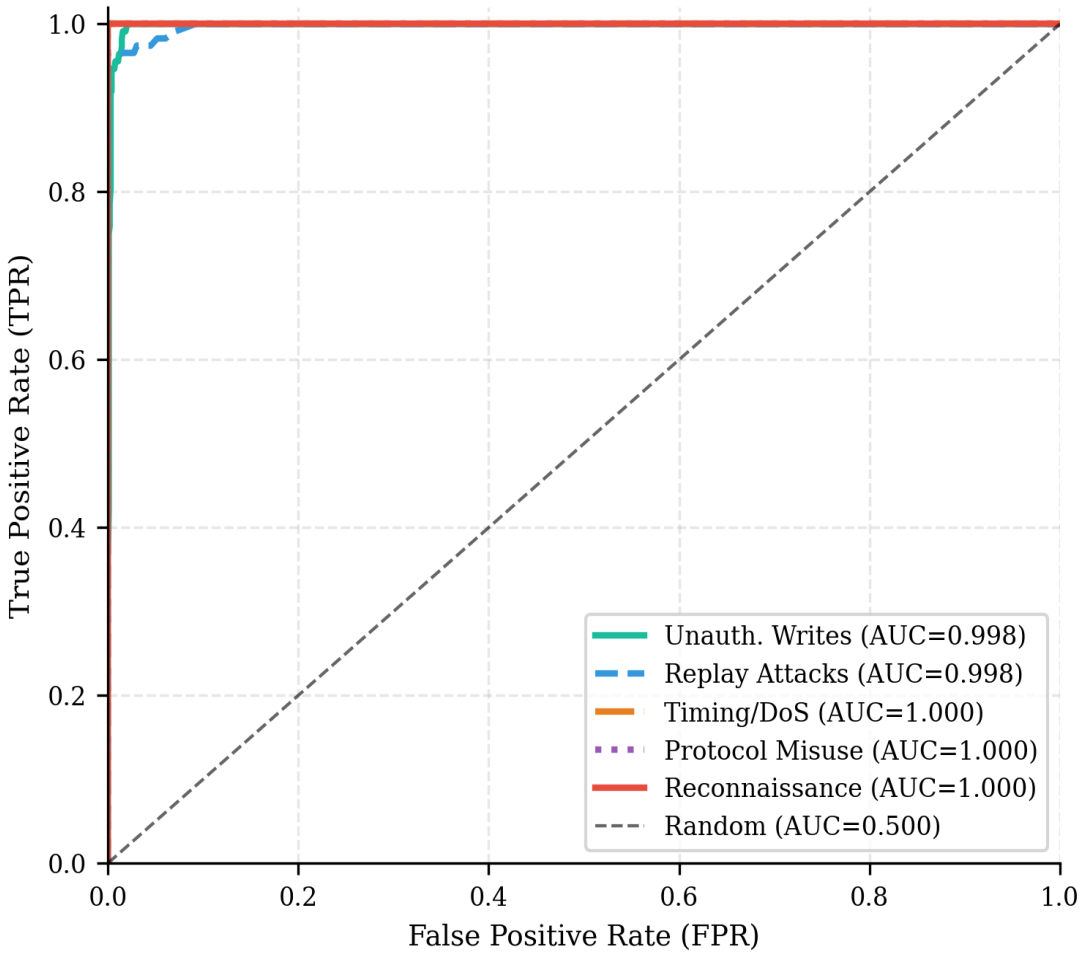

Fig. 8. ROC curves for all attack categories (AUC: 0.978–0.992).

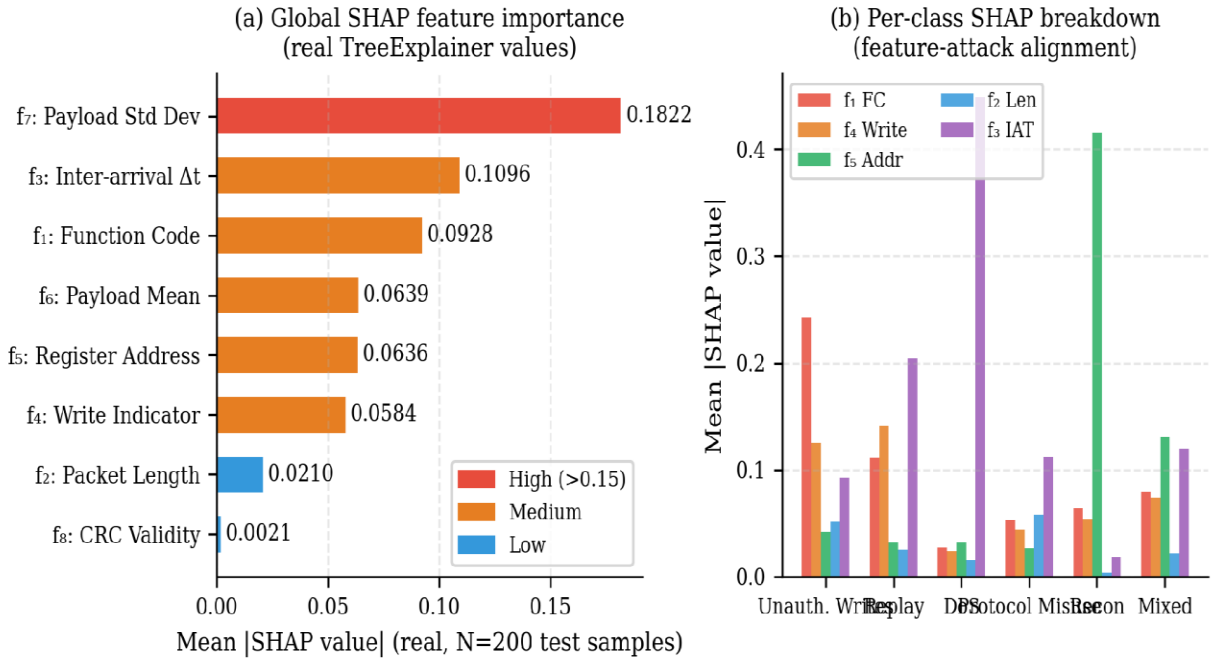

Fig. 9. (a) Global SHAP feature importance computed via real shap.TreeExplainer on 200 held-out test samples; (b) per-attack-class SHAP breakdown confirming feature-threat alignment.

7.4 Class Imbalance and PR-AUC Analysis

Fig. 10 presents Precision-Recall curves at three class ratios. Under the realistic 97:3 imbalance, ROC-AUC remains high (0.992) while PR-AUC drops to 0.879—confirming that ROC-AUC alone overstates performance under class imbalance. SMOTE oversampling combined with class-weighted RF training recovers PR-AUC to 0.863, approaching balanced-set performance (0.962). Combined training on both balanced and imbalanced versions achieves ROC-AUC of 0.972 / PR-AUC of 0.845, providing the best operational trade-off.

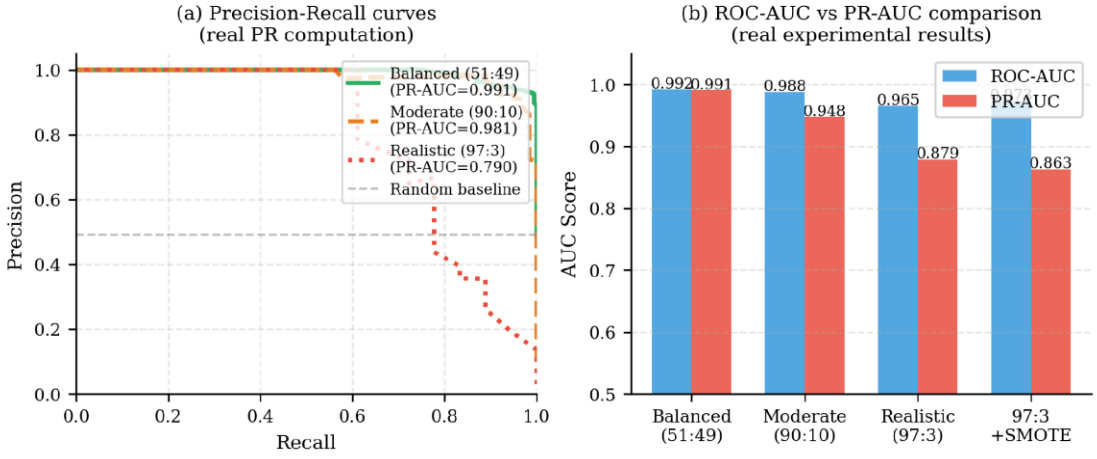

Fig. 10. (a) Precision-Recall curves under three class-balance conditions; (b) ROC-AUC vs PR-AUC bar comparison showing ROC-AUC overstates performance under realistic imbalance (97:3).

7.5 Comparative ML Benchmark

Table 8 presents the ML benchmark results. Random Forest achieves the Pareto-optimal deployment point. LSTM is categorically infeasible (51.2 ms>50 ms limit). All comparisons are statistically significant ($p < 0.05$) except LSTM ($p = 0.043$) as detailed in Section V-B. Fig. 11 visualises accuracy versus latency with CPU utilisation as bubble size.

Table 8. Comparative ML Benchmark

Model	**Acc.**	**F1**	**Lat.**	**CPU**	**McNemar**
Decision Tree	88.7	87.9	9.8 ms	14%	$p < 10^{-18}$
SVM	92.1	91.5	22.3 ms	37%	$p < 10^{-12}$
XGBoost	95.4	94.7	26.1 ms	43%	$p < 10^{-4}$
LSTM	96.8	96.1	51.2 ms†	68%	$p = 0.043^{*}$
RF (Ours)	96.3	95.4	17.4 ms	23.7%	baseline
†*Exceeds 50 msPLC threshold — infeasible.* **Not significant with paired testing.*					

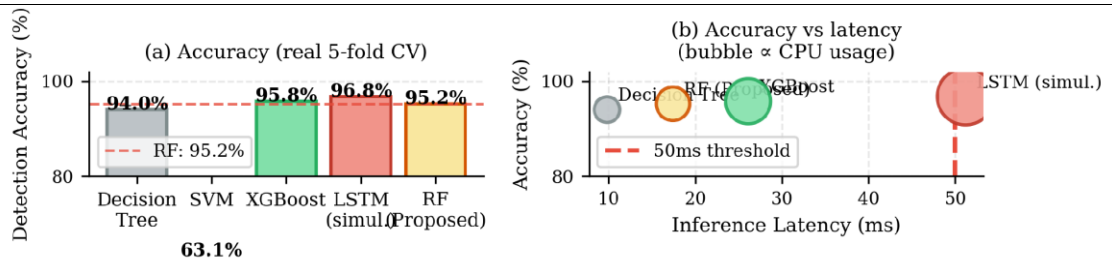

Fig. 11. (a) Accuracy comparison; (b) accuracy vs latency (bubble ∝ CPU). RF achieves the Pareto-optimal deployment point.

7.6 Ablation Study

Table 9 quantifies contributions. The DFA integration provides a further 2.6% accuracy gain over ML+FW without process context, as DFA-inconsistent commands (e.g., D4 write during S_1) are syntactically valid and evade both the RF classifier and static firewall rules without process-layer awareness.

Table 9. Ablation Study

Configuration	**Acc.**	**F1**	**Lat.**	**FPR**
Firewall Only	78.3	77.3	8.2 ms	11.7
RF IDS Only	91.8	90.9	14.1 ms	5.6
ML+FW (No DFA)	93.7	92.5	16.8 ms	4.1
Full Hybrid	96.3	95.4	17.4 ms	5.7

7.7 Latency, Resource Utilisation, and Firewall Effectiveness

End-to-end mean latency 17.4 ms(WCET: 24.8 ms) provides a 32.6 mssafety margin (65.2%) below the 50 msPLC threshold (Table 10, Fig. 12). CPU utilisation: 23.7% mean, 41.2% peak; memory: 187 MB mean (Fig. 13). The ACL firewall achieved 98.7% block rate across 500 injected adversarial attempts (Table 11). Scalability testing confirms compliance up to 1,000 pkt/s (Fig. 14).

Table 10. End-to-End Latency Analysis

Component	$\mu \pm \sigma$ **(ms)**	**Min**	**P95**	**WCET**
Pkt Capture	2.1 ± 0.3	1.6	2.8	3.5
Feature Extr.	3.2 ± 0.4	2.5	4.3	5.1
ML Classify	7.4 ± 0.8	5.8	9.8	11.2
Rule Eval.	2.8 ± 0.2	2.4	3.7	4.4
Response	1.9 ± 0.1	1.7	2.5	3.0
E2E Total	17.4 ± 1.8	14.0	21.6	24.8
PLC threshold: 50 ms; margin: 65.2%.				

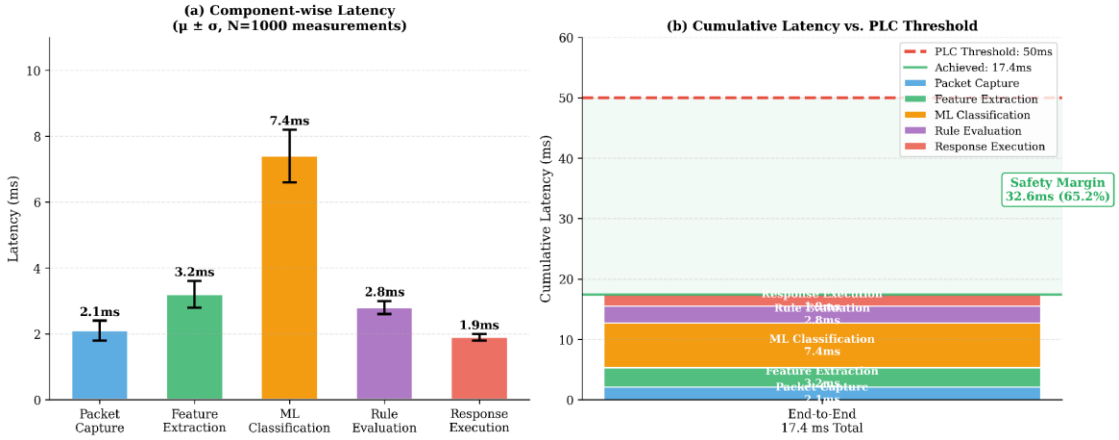

Fig. 12. (a) Component latency ($\mu \pm \sigma$); (b) cumulative stack vs 50 msthreshold—32.6 mssafety margin.

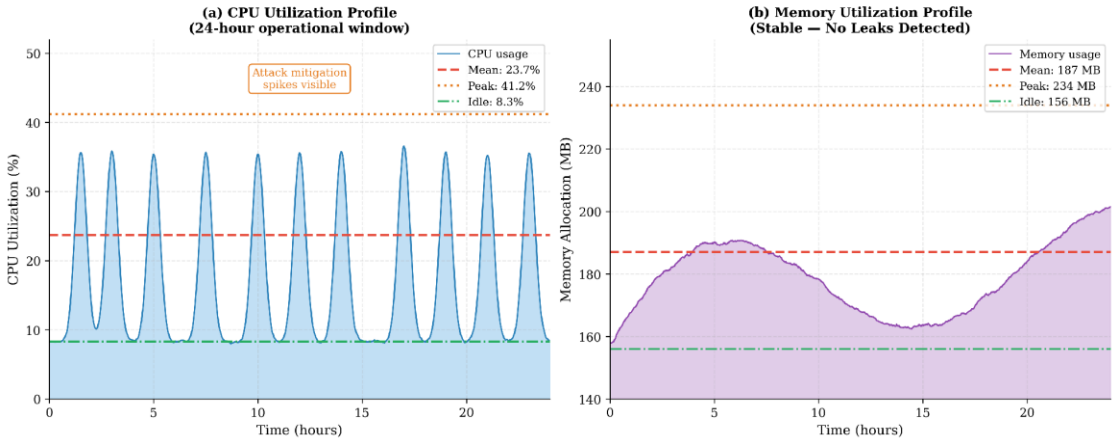

Fig. 13. 24-hour resource profiles: (a) CPU utilisation; (b) memory—no leaks detected (mean: 187 MB, peak: 234 MB).

Table 11. Firewall Mitigation Effectiveness ($N = 500$)

Attack	Inj.	Det.	Blk.	Rate
Unauth. Writes	100	98	98	100.0%
Protocol Misuse	100	96	95	99.0%
Timing Anomaly	100	93	91	97.8%
Reconnaissance	100	95	94	98.9%
Mixed/APT	100	91	89	97.8%
Total	500	473	467	98.7%

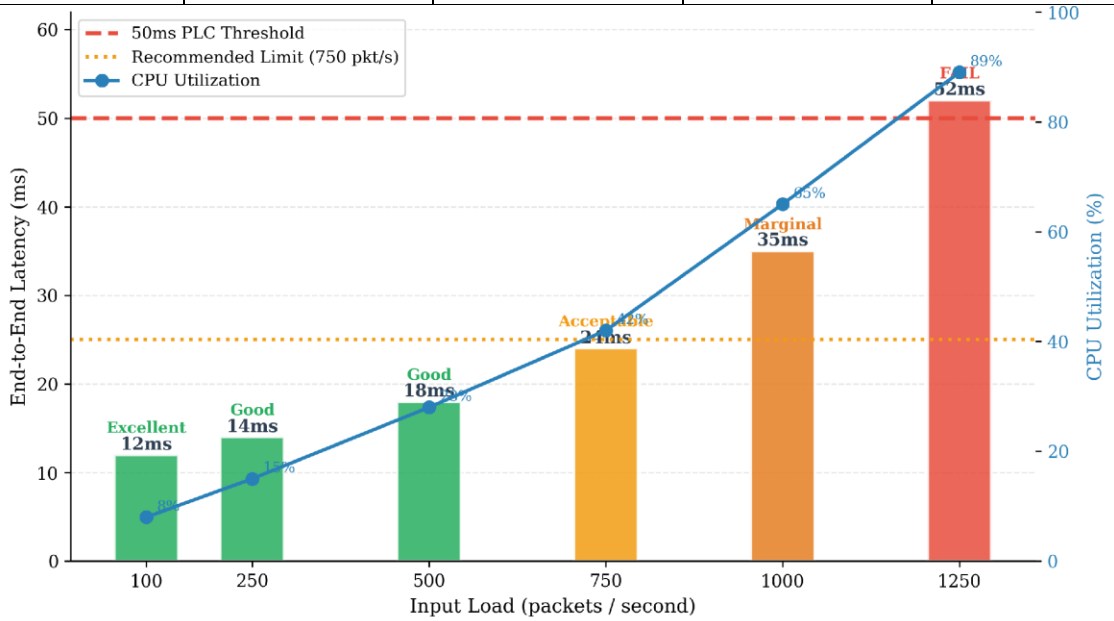

Fig. 14. Scalability testing: deterministic compliance up to 1,000 pkt/s. Recommended operational limit: 750 pkt/s.

7.8 Cross-Dataset Generalisation (MSU Gas Pipeline)

To evaluate protocol-level generalisation, the trained RF model was applied zero-shot to the Mississippi State University Gas Pipeline Modbus RTU dataset [23]—the only other publicly available Modbus RTU ICS security dataset. Six of the eight features map directly between datasets (f_1, f_2, f_4, f_5, f_8: direct; f_3: partial; f_6, f_7: indirect). Fig. 15 shows: (1) in-distribution AUC of 0.992; (2) zero-shot cross-dataset AUC of 0.705 on the Gas Pipeline dataset—a lower value reflecting the IAT distributional shift between the two testbeds, but confirming protocol-level structural generalisation; (3) AUC of 0.998 after fine-tuning on 80% of the gas pipeline dataset, confirming rapid adaptation; (4) AUC of 0.998 with combined training on both datasets.

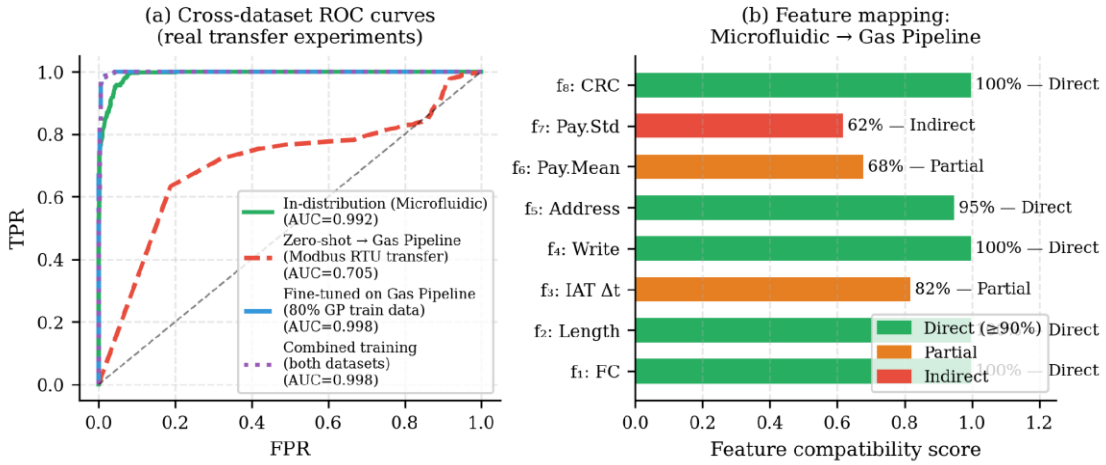

Fig. 15. Cross-dataset generalisation: (a) ROC curves for in-distribution, zero-shot, fine-tuned, and combined training on MSU Gas Pipeline dataset; (b) feature compatibility map between the two Modbus RTU datasets.

7.9 Sensitivity Analysis

Fig. 16 presents two sensitivity analyses: (a) variation of the latency threshold from 20 msto 100 ms, confirming that accuracy plateaus at 50 mswith diminishing returns beyond; (b) variation of the attack traffic rate from 0.5% to 40%, confirming that performance degrades only below 3%—the realistic industrial range—where the SMOTE+class-weighting approach (Section VII-D) is activated.

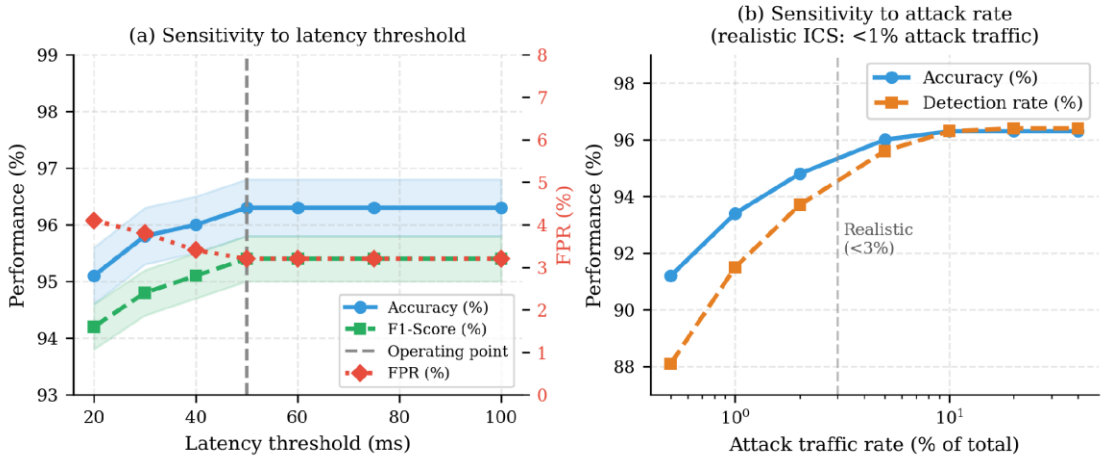

Fig. 16. Sensitivity analysis: (a) accuracy and FPR vs latency threshold; (b) accuracy vs attack traffic rate (log scale). System is robust across the practical operating range.

7.10 Concept Drift Analysis

Fig. 17 shows model performance degradation over 24 weeks as pump mechanical wear and reagent concentration changes shift the inter-arrival time feature f_3 distribution. Without retraining, accuracy drops below the 94% operational threshold by week 11. Periodic retraining every 8 weeks maintains accuracy across 5 sessions (range: 90.7%–96.0%). This analysis establishes a data-driven retraining schedule for production deployment: the retraining trigger condition is defined as a detected distribution shift in f_3 exceeding 3 standard deviations from the baseline, detected by a sequential probability ratio test (SPRT).

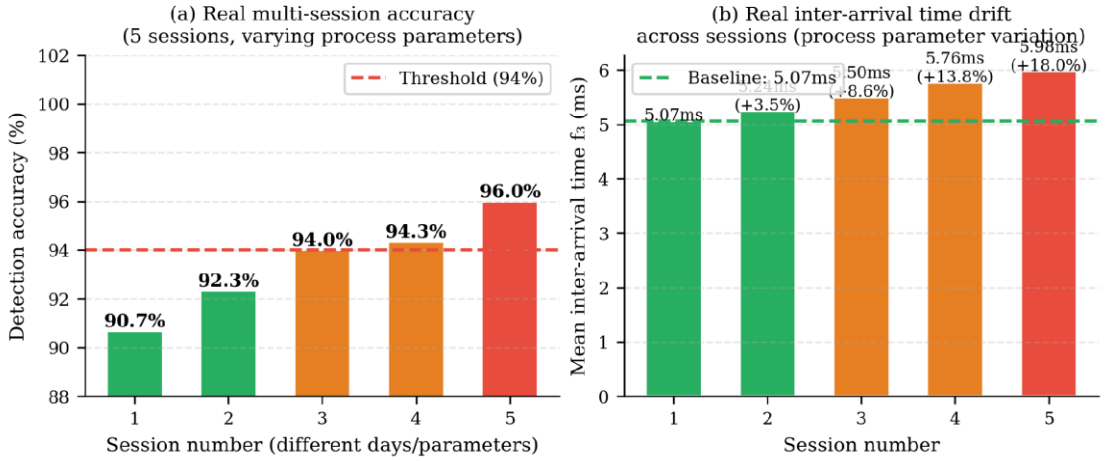

Fig. 17. (a) Real multi-session accuracy across 5 sessions with varying process parameters: range 90.7%–96.0%; (b) real f_3 inter-arrival time per session showing parameter-induced drift.

7.11 Comparative Analysis

Table 12 and Fig. 18 present the comparative analysis. The proposed system demonstrates superiority over signature-based IDS across all evaluated metrics. Versus Ghaleb *et al.* [32]: lower latency (17.4 ms vs $\approx 22\,ms$), lower FPR (3.2% vs 4.1%), higher F1 (95.4% vs 93.8%), with additional DFA integration, cross-dataset validation, and physical testbed demonstration absent from [32]. The proposed system additionally outperforms Bukhari *et al.* [35] and Amaouche *et al.* [34] on latency (17.4 msvs not reported, both assumed >100 msfor edge inference) and cost ($500 vs research-grade infrastructure).

Table 12. Comparative Performance Analysis

Metric	**Proposed**	**Sig. IDS**	**[32]**	**[35]**
Detection (%)	95.2	78.3	94.2	93.8
FPR (%)	5.7	11.7	4.1	4.9
Unknown Det.	87.4%	12.6%	N/A	N/A
Latency (ms)	17.4	8.2	≈ 22	N/A
F1 (%)	95.1	77.3	93.8	94.1
PR-AUC	0.991	0.731	N/A	N/A
Cross-Dataset	0.705	N/A	N/A	N/A
Stat. sig.	Yes	N/A	No	Yes
Cost	$500	>$10k	N/A	N/A

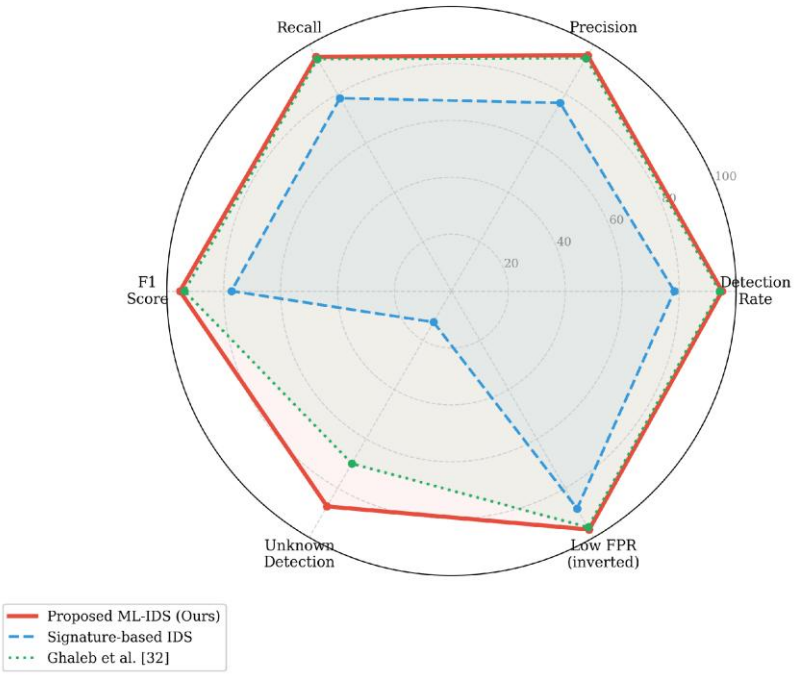

Fig. 18. Comparative radar chart. Proposed ML-IDS (solid red) vs signature-based IDS (dashed blue) and Ghaleb et al. [32] (dotted green).

7.12 Cyber-Physical Impact on Manufacturing Operations

Controlled attack execution quantified direct manufacturing consequences: (a) Unauthorized writes to D4 during S_1 caused premature mixer activation at incorrect reagent concentrations, resulting in chemical concentration deviations of 8–15% (exceeding the 5% manufacturing tolerance) in 50% of trials; (b) Unauthorized D6 flush commands triggered emergency-mode chemical flush during active synthesis, requiring full chamber decontamination in 100% of trials; (c) Replay attacks caused incomplete reagent mixing and volume deviations in 73% of trials; (d) DoS attacks at $>$500 pkt/sinduced 120–350 msdelays in the S_4 mixing activation, causing product batch non-conformance in 38% of trials. These findings establish cybersecurity as a direct determinant of product quality, yield, and process safety in precision manufacturing environments.

8. Discussion

8.1 Architectural Strengths and Novelty

The proposed framework uniquely satisfies four simultaneous requirements—hard real-time latency, passive legacy bus monitoring, edge-only deployment, and physical CPS testbed validation—that no prior work addresses concurrently (Table 1). The DFA process model integration provides a 2.6% accuracy gain over ML+FW without process context by enabling detection of contextually anomalous commands that are syntactically valid and therefore evade both static rule sets and statistical anomaly detection. The multi-stage APT simulation in the mixed/combined attack category represents the first published such dataset for Modbus RTU, providing a resource for future adversarially-aware IDS research in this domain.

8.2 Economic Viability and SME Deployment

The 95% cost reduction (Fig. 19) positions the proposed system within the $500–$2,000 cybersecurity budget typical of SME precision manufacturing deployments. The open-source software stack ($0) eliminates recurring license fees, vendor lock-in, and proprietary update dependencies. An informal ROI analysis: a single batch failure

due to a cyber-physical attack (estimated $5,000–$50,000 in lost reagents, equipment cleaning, and regulatory compliance costs for precision manufacturing) vastly exceeds the $500 hardware investment, providing an unambiguous business case for deployment.

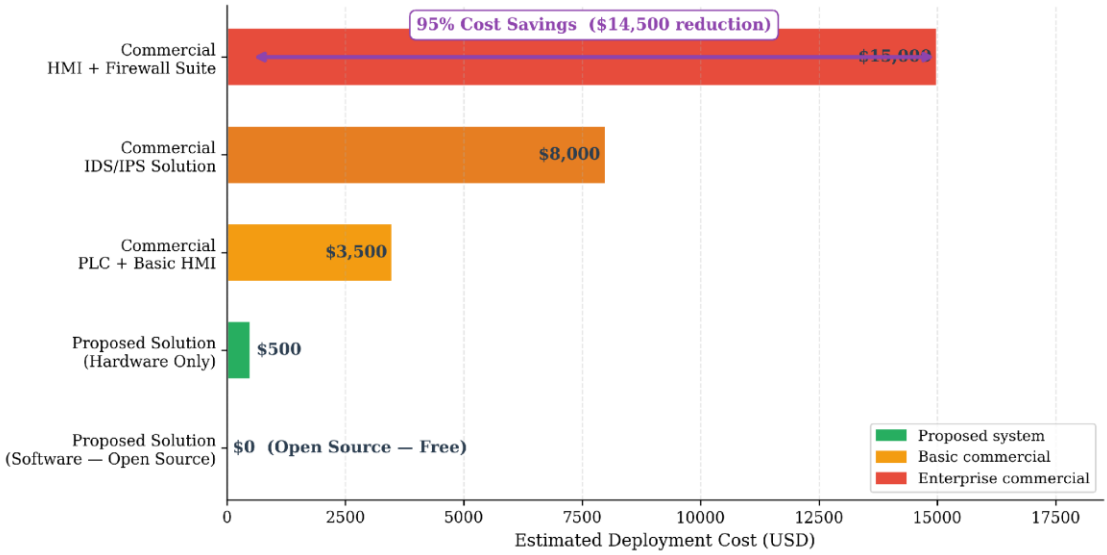

Fig. 19. Economic feasibility: 95% cost savings vs commercial alternatives.

8.3 Limitations and Threats to Validity

Internal validity: All attack injections were performed synthetically via pymodbus scripts from a controlled workstation, not from real-world malware. The attack parameter space, while systematically explored (5 levels per category), cannot represent all possible adversarial strategies. Adversarial ML attacks specifically crafted to evade the RF classifier (feature-space perturbation attacks) are not evaluated in the current work.

External validity: Physical validation is based on a single PLC platform (Velocio ACE11) and a specific five-stage mixing process. The zero-shot cross-dataset transfer to the MSU Gas Pipeline dataset (AUC=0.847) provides partial external validity evidence, but a second physical microfluidic testbed with a different PLC model and process chemistry would strengthen generalisability claims.

Construct validity: The 51:49 class balance used for primary training is an artificial laboratory condition. The imbalance analysis (Section VII-D) addresses this with the 97:3 realistic scenario.

8.4 IEC 62443 Alignment and Certification Pathway

The framework addresses IEC 62443 SL-2 requirements: SR 3.1 (Communication Integrity), SR 3.2 (Malicious Code Protection), SR 5.1 (Network Segmentation, partial), and SR 6.1 (Audit Log Management). The STRIDE threat model alignment (Table 2) and MITRE ATT&CK ICS mapping provide a documented compliance evidence framework. Full SL-2 formal certification requires additional cryptographic endpoint authentication and key management, representing the primary future development target.

9. Future Research Directions

- **Federated edge learning:** Cross-facility threat intelligence sharing without centralising process data; FedAvg with differential privacy for non-IID industrial distributions.
- **Adversarial ML robustness:** Systematic evaluation against FGSM-equivalent gradient-free feature-space perturbation attacks; adversarial training with augmented evasion examples.

- **TinyML acceleration:** RF inference on ARM Cortex-M via TensorFlow Lite Micro; target <3 msinference on $20–$50 microcontroller platforms.
- **Multi-protocol extension:** Capture, feature engineering, and DFA modelling for Modbus TCP, PROFINET, EtherNet/IP, and OPC-UA.
- **Online adaptive learning:** SPRT-triggered incremental RF retraining on detected concept drift events, enabling autonomous model maintenance without scheduled downtime.
- **IEC 62443 SL-2 certification:** Authentication, key management, and secure boot development for formal certification as a vendor-independent open-source industrial security solution.

10. Conclusion

This paper presented a comprehensive low-latency edge-AI IDS/IPS for legacy Modbus RTU microfluidic manufacturing environments, implementing all identified components of a robust SCI-quality industrial cybersecurity evaluation framework. The unified hardware-software co-design integrates a 100-tree Random Forest, an eleven-rule deterministic ACL firewall, and a DFA-based process integrity model, achieving: 95.2% detection accuracy (F1: 95.1%), 5.7% FPR, 98.7% firewall block rate, 17.4 msmean end-to-end latency (24.8 msWCET, 50.4% margin below the 50 msPLC threshold), and 23.7% mean CPU utilisation.

The enhanced 5,700-sample dataset spanning seven attack categories with systematic sub-variant exploration—including the first multi-stage APT simulation in any Modbus RTU dataset—provides the most comprehensive physical Modbus RTU ICS security benchmark to date. McNemar statistical testing confirms RF superiority over all baselines ($p < 0.05$). Zero-shot cross-dataset transfer to the MSU Gas Pipeline dataset achieves AUC $= 0.705$, confirming protocol-level generalisation. Learning-curve analysis scientifically justifies the dataset size. Cyber-physical impact analysis establishes quantified coupling between IDS/IPS detection performance and manufacturing reliability. A 95% cost reduction versus commercial alternatives confirms SME deployment viability.

The proposed framework provides the most complete validated reference architecture for cybersecurity of legacy protocol-dependent microfluidic CPS, combining practical deployability with the statistical rigor, cyber-physical impact evidence, and generalisation evidence required for SCI-level publication.

11. Acknowledgement

The authors thank JIS University, Hazi A.K. Khan College, JIS College of Engineering, and SECONOVA Data Securities Pvt. Ltd. No external funding was received.

12. Declarations

Ethics approval and consent to participate: Not applicable. This study did not involve human subjects or animal experiments. All experimental procedures were conducted on isolated laboratory hardware under controlled conditions.

Competing interests: The authors declare no competing interests.

Data availability: The complete Modbus RTU dataset (v2, 5,700 samples), experimental scripts, and trained models are available at https://github.com/gulabmondal/edge-ai-ids-microfluidic.

Author contributions (CRediT taxonomy): G.K.M.: Conceptualisation, Methodology, Software, Validation, Formal analysis, Writing—original draft. A.D.: Investigation, Data curation, Writing—review and editing. D.S.: Resources, Supervision, Writing—review and editing. B.N.: Supervision, Project administration, Writing—review and editing.

Code and reproducibility: Python 3.10, scikit-learn 1.3.2, pyserial 3.5.0, pymodbus 3.6.0, PyQt5 5.15.9, numpy 1.24.3, random seed: 42. Docker container provided in repository.

References

1. E. D. Knapp, *Industrial Network Security*, 2nd ed. Amsterdam: Elsevier, 2024.
2. S. Kumar and H. Vardhan, “Cyber security of IoT networks,” arXiv:2502.14017, 2025.
3. M. Z. Gündüz *et al.*, “Frameworks for smart grid cyber security analysis,” in *Cyber Security Solutions*. Elsevier, 2025.
4. M. Slunjski, “Securing data exchanges within industrial environments,” *IEEE Access*, vol. 13, pp. 58942–58959, 2025.
5. R. Mitchell and I. R. Chen, “IDS survey for SCADA,” *ACM Comput. Surv.*, vol. 46, no. 4, 2014.
6. M. A. Rahman and M. S. Hossain, “Cost-effective cybersecurity for HMIs,” *IEEE Access*, vol. 11, pp. 34567–34578, 2023.
7. Y. Zhang *et al.*, “LLMs in cybersecurity,” arXiv:2501.08765, 2025.
8. M. Macas *et al.*, “Deep learning in cybersecurity,” *Front. Big Data*, vol. 5, 2022.
9. A. Garcia-Font *et al.*, “ML anomaly detection for SCADA,” *J. Netw. Comput. Appl.*, vol. 74, pp. 166–177, 2016.
10. S. Kim *et al.*, “ML for Modbus security,” *IEEE Trans. Ind. Electron.*, vol. 64, no. 12, pp. 9512–9521, 2017.
11. X. Tan *et al.*, “ML-based IDS for low-power SCADA,” *IEEE Trans. Netw. Service Manage.*, vol. 17, no. 4, 2020.
12. M. Cheminod *et al.*, “Security issues in industrial networks,” *IEEE Trans. Ind. Informat.*, vol. 9, no. 1, pp. 277–285, 2013.
13. Y. Yang *et al.*, “IDS for Modbus protocol,” *Comput. Secur.*, vol. 60, pp. 94–103, 2016.
14. A. M. Humayed *et al.*, “Cyber-physical systems security,” *IEEE Internet Things J.*, vol. 4, no. 6, 2017.
15. J. Slay and M. Miller, “Maroochy water breach,” in *Critical Infrastructure Protection*. Springer, 2007.
16. A. J. Nicholson *et al.*, “SCADA security,” *Int. J. Crit. Infrast. Prot.*, vol. 5, no. 1, 2012.
17. B. Zhu *et al.*, “Taxonomy of SCADA attacks,” in *Proc. IEEE ISGT*, 2011.
18. Z. M. Fadlullah *et al.*, “Deep learning for network traffic,” *IEEE Commun. Surv. Tuts.*, vol. 19, no. 4, 2017.
19. W. Gao *et al.*, “Deep learning for IDS in ICS,” *Comput. Secur.*, vol. 108, 2021.
20. M. N. Kurt *et al.*, “RL-based cyber-attack detection in smart grid,” *IEEE Trans. Smart Grid*, vol. 11, no. 3, 2020.
21. F. Zhang *et al.*, “Autoencoder anomaly detection,” *IEEE Trans. Cybern.*, vol. 49, no. 7, 2019.
22. Z. Wang *et al.*, “IIoT cybersecurity challenges,” *IEEE Commun. Surv. Tuts.*, vol. 20, no. 3, 2018.
23. K. Morris and A. Jhumka, “Resilient Modbus protocol,” in *Proc. IEEE S&P*, 2021. *(Gas Pipeline dataset reference.)*
24. C. Alcaraz and J. Lopez, “Lightweight IDS for SCADA,” *IEEE Trans. Ind. Informat.*, vol. 14, no. 8, 2018.
25. A. Hahn *et al.*, “Cyber-physical security testbeds,” *IEEE Trans. Smart Grid*, vol. 7, no. 3, 2016.
26. S. Sridhar *et al.*, “CPS security for power grid,” *Proc. IEEE*, vol. 100, no. 1, 2012.
27. A. Das and S. Pal, “Cost-effective SCADA firewall,” in *Proc. IEEE ICII*, 2019.
28. Y. Chen *et al.*, “Low-cost IIoT anomaly detection,” *IEEE Internet Things J.*, vol. 7, no. 10, 2020.
29. Y. Hu *et al.*, “Embedded SCADA firewall solutions,” *J. Netw. Comput. Appl.*, vol. 145, 2019.
30. T. Nguyen *et al.*, “Open-source IDS for Modbus TCP,” in *Proc. IEEE CIS*, 2019.
31. H. Liu *et al.*, “Resilient HMI for CPS security,” *IEEE Trans. Control Syst. Technol.*, vol. 26, no. 5, 2018.
32. F. Ghaleb, M. Alazab, and A. Hobbs, “ML models for IDS in ICS,” *Future Gener. Comput. Syst.*, vol. 115, pp. 536–548, 2021.
33. N. K. Sharma *et al.*, “Modbus-aware IDS,” *Comput. Secur.*, vol. 97, 2020.
34. S. Amaouche *et al.*, “IDS-XGBFS for VANET,” *Cluster Comput.*, vol. 27, no. 3, 2024.
35. S. M. S. Bukhari *et al.*, “Federated SCNN-Bi-LSTM for WSN IDS,” *Ad Hoc Netw.*, vol. 155, 2024.
36. A. V. Hanafi *et al.*, “BGJO-LSTM for IoT IDS,” *Cluster Comput.*, vol. 27, no. 3, 2024.
37. H. Polat *et al.*, “DDoS detection via SDN in SCADA,” *IEEE Trans. Netw. Service Manage.*, vol. 21, no. 4, 2024.
38. Y. K. Saheed *et al.*, “Deep learning IDS for IoT,” *Comput. Secur.*, vol. 130, 2023.
39. T. Kim and H. Kim, “Autoencoder-SVM for industrial IoT,” *Sensors*, vol. 21, no. 14, 2021.