Int. J. Advance Soft Compu. Appl, Vol. 18, No. 1, March 2026
Print ISSN: 2710-1274, Online ISSN: 2074-8523

AI-Powered Security-Aware Reconfiguration in Cyber-Physical Systems for Smart Healthcare and Energy Domains

Theyazn H.H Aldhyani[1], Rajit Nair[2], Hasan Alkahtani[3], and Osamah Ibrahim Khalaf[4]*

[1]Applied College, King Faisal University, Al-Ahsa, 31982, Saudi Arabia; taldhyani@kfu.edu.sa

[2]VIT Bhopal University, Bhopal, India; Email: R.Nair2@gmail.com

[3] College of Computer Science and Information Technology, King Faisal University, P.O. Box 400, Al-Ahsa 31982, Saudi Arabia hsalkahtani@kfu.edu.sa

[4] Al-Nahrain Nanorenewable Energy Research Center Usama81818@nahrainuniv.edu.iq

Abstract

Abstract— Cyber-physical systems (CPS) in smart energy and smart healthcare must continue to be safe, fast, and reliable, even in the face of evolving cyber threats and fluctuating network/edge conditions. An AI-based, security-aware framework of CPS reconfiguration is presented in this paper. It is comprised of (i) privacy-preserving federated anomaly detection at distributed edge nodes, (ii) a reinforcement-learning decision module that selects risk-aware reconfiguration actions under the objectives of latency, energy, and safety, and (iii) auditability, backed by blockchain, to provide trustworthy governance and accountability in the aftermath of the event. The detection module learns behavioral baselines at the local level and shares only protected updates to support data minimization and privacy compliance while still achieving high accuracy in heterogeneous deployments. The RL controller dynamically modifies actions to reduce service disruption, quicken recovery, and avert unsafe control transitions. A lightweight operational ledger captures reconfiguration activities and trust updates to provide auditable governance in multi-stakeholder CPS settings. In energy and healthcare CPS scenarios, there is strong operational performance. The system achieves 97.8% detection accuracy and an F1-score of 97.4% with an end-to-end latency of 41ms, coupled with a reconfiguration time of 65ms and a mean time to recovery of 5.6 seconds. The framework provides 99.4% uptime, consumes approximately 10 W at the edge, and with low error rates (2.1% false positive and 1.0% false negative), achieves 19.5 Mbps secure throughput. These results demonstrate that self-reconfigurable CPS can maintain mission-critical operational continuity while enhancing privacy, scalability, and governance in large-scale deployments.

Keywords*: Adaptive governance, Blockchain auditing, Differential privacy, Edge intelligence, Federated anomaly detection, Multi-objective reinforcement learning, Non-IID robustness, Real-time reconfiguration, Secure throughput, Trust-weighted aggregation.*

Received 4 October 2025; Accepted 30 January 2026

1 Introduction

Cyber-Physical Systems (CPSs) integrate computing and physical processes and can automate decision-making with high precision and control at a granular level across a wide range of applications. The impact of CPSs is most significant in the field of smart healthcare and smart energy infrastructures with the utmost importance being reliability, safety, and responsiveness in real-time. In healthcare, CPSs can provide continuous monitoring of patients, assist in robotic surgeries, and provide personalized diagnostics and remote clinical interventions. In the energy sector, CPSs manage and coordinate distributed generations, provide load balancing, and grid disturbance detection [1-3].

Despite such functions, a tight coupling of the software and physical components can lead to significant cyber-physical operational vulnerabilities. As connectivity, scaling, and tight coupling of CPSs grow, cyber intrusions and cascading failures pose critical threats to patient safety and energy stability. The proposed solution to these challenges is the construction of an AI security driven integrated CPS behavioral reconfiguration framework. The proposed framework integrates real-time deep learning-based anomaly detection with reinforcement learning-driven policy optimization focused on preserving privacy and knowledge to balance safety, service continuity, and resource efficiency. The ability to reconfigure behavior dynamically in self-optimizing autonomous CPSs is, in and of itself, a proactive mechanism to counter threats while explainable decision-making fosters transparency and is aligned with regulations [4-6]. The solution proposed in this paper describes a four-layer architecture that enables Cyber-Physical Systems (CPS) to detect, determine, and audit in real time, and reconfigure their responses. This paper presents an AI-based reconfiguration framework aware of the security challenges of CPS for mission-critical smart healthcare and smart energy deployments. Here, we must consider all the challenges of the balancing act: detection and response time, governance, and privacy [7-9]. The first contribution involves a trust-weighted, privacy-preserving federated anomaly detection pipeline, which runs on distributed CPS nodes and updates a global detector model without direct data sharing. Differently from a typical federated learning process, the proposed aggregation and decision framework incorporates node trust/reliability, which supports adequate learning with a diverse set of devices and noisy/unreliable participants. The second contribution is a risk-aware multi-objective reinforcement learning policy for the reconfiguration of systems, which optimizes the balance between security and operational trade-offs (latency, energy, recovery time) while, under a safety net, avoids the triggering of unsafe switching decisions within the scope of the clinical and grid-control domains. The third contribution describes a governance layer that is blockchain-audited, for which the tamper-evident logging of reconfiguration and responsibility are to feed audit outcomes toward trust updates, thus enabling a closed-loop security lifecycle (detect → decide → reconfigure → audit → improve trust/learning). The evaluation results demonstrate great real-time feasibility and resilience, high detection with low operational overhead, fast reconfiguration and recovery, high uptime, and energy-efficient operation, making it suitable for edge CPS operation. Finally, an ablation study (A0–A10) demonstrates the framework's novelty by isolating each module's contribution and checking if the integrated design, rather than any individual component, is responsible for the overall end-to-end performance.

As opposed to traditional methods which apply federated learning (FL), reinforcement learning (RL), or blockchain in isolation, the proposed model integrates a closed-loop

security–governance–recovery structure whereby learning, decision-making, and auditability strengthen and reinforce each component. First, the method, beyond "standard FL," incorporates a form of trust-weighted model aggregation, in which unreliable or hostile nodes are given less influence on the global model. This approach improves the model's robustness in the context of heterogeneous CPS (Cyber-Physical Systems). Second, the system combines privacy masking with differential privacy (DP) in a way that minimizes the risk of an adversarial attack with update sharing, fulfilling the requirements of edge-technology-based training while allowing for privacy-preserving deployment in the healthcare and energy sectors [10-11]. Third, the proposed controller, beyond "standard RL reconfiguration," is multi-objective and safety-constrained. This means that, in addition to the security that is maximized, the latency, energy, and service continuity constraints are optimized, resulting in a more balanced response. Finally, the audit layer, beyond "standard blockchain logging," is more than a static archive. With an integrated trust updating mechanism, accountability-based adaptation is operationalized, where trust scores and reconfiguration and learning adaptation are influenced by verified transactions and audit outcomes. This combination creates an auditable CPS security lifecycle that balances and integrates fast detection, rapid switching, fault-tolerant recovery, and governance. This paper proposes a novel addition to the CPS security lifecycle, integrating a loop mechanism to the privacy-preserving federated detection and trust-weighted aggregation, multi-objective RL-based reconfiguration, and blockchain auditability so that security decisions are fast and verifiable. It executes the closure of the governance loop by incorporating an auditable trail and trust updates with each reconfiguration event, providing accountability in the governance loop without sacrificing the fast real-time responsiveness of the system. It also goes beyond the accuracy metric in reporting the CPS-centric metrics, covering latency, reconfiguration, mean time to repair (MTTR), uptime, energy consumption, uptime, privacy, secure throughput, scalability, and trust, enabling the evaluation of the system as a deployable control-and-security system as opposed to an IDS alone. Lastly, the ablation-style variants demonstrate the impact of the various modules on safety, trust, privacy, and operational efficiency, providing evidence to support the auditable novelty claim beyond a simply conceptual integration.

The rest of this paper is organized as follows; section 2 shows the related works that have been conducted in this field. Section 3 overviews the methodology. Section 4 introduces the experiment and discusses the results. Finally, we conclude this paper in section 5.

2 Related Work

For the past few years, various AI technologies have been explored to improve the safety, flexibility, and error tolerance of configurable Cyber-Physical Systems (CPSs). Among these, Deep Reinforcement Learning (DRL) has been noted for its speed and flexibility, allowing it to modify its course of action in the face of new challenges. It has been noted that DRL decreases downtime and increases speed of reconfiguration, which is beneficial to real time adaptive solutions such as energy grids and healthcare monitoring. Another form of distributed intelligence is called Multi-Agent Reinforcement Learning (MARL). In MARL, AI agents cooperate, as a team, to discover, report, and remedy problems. This increases the flexibility of large systems and decentralizes the capacity for problem solving [12]. In terms of CPSs in healthcare, particularly when patient privacy is the utmost concern, Federated Learning (FL) is a particularly relevant approach. With FL, edge devices can train models while keeping sensitive information embedded in the devices and

not in a central storage solution. This enhances the security and scalability of the system. With respect to the shortening of training time and the increased flexibility of pre-trained models to transfer to new domains within a CPS (i.e., from smart energy to healthcare) and to reuse knowledge within cross-domain reconfiguration activities, Transfer Learning (TL) is beneficial [13]. For CPSs that have numerous complex, interrelated nodes, the application of Graph Neural Networks (GNNs) is becoming more common. GNNs are good at replicating network structure which is beneficial in modeling the spread of failures and attacks, along with providing recommendations for reconfiguration based on network topology.

Nevertheless, unsupervised deep learning in autoencoders looks for very subtle alterations in the system that could be associated with an attack or damage in the system. This method detects threat in the networks of smart grids and medical devices. The Policy-Based Reconfiguration, which is an old but significant method, employs the use of system defined rules to regulate the response of the system. This method is effective in well-known situations but ineffective in novel ones [14]. Probabilistic modeling looks into system state changes and the subsequent system reconfiguration using Markov Decision Processes (MDPs). Even in the presence of uncertainty, they provide dependable MDPs. With high sensor noise or incomplete information, Bayesian Network Diagnosis employs the technique of probabilistic reasoning to diagnose the system and suggest fault-tolerant reconfiguration through rewiring. Lastly, fuzzy logic controllers are used to provide the reasoning of rules when uncertainty is presented, like how people go about making decisions, and provide gradual changes in system states [15]. They are straight forward in their application, but in rapidly and drastically transforming environments they are ineffective. These methods used a number of metrics to test their performance. They showed a wide dispersion in the performance of the reconfiguration of secure CPS and their application. Among deep reinforcement learning and other types of reinforcement learning, the former is more effective when it comes to latency, flexibility, and accuracy to responses. These attributes make it more suitable for real-time applications. In environments where data security is a concern, federated and transfer learning are very effective in the use and development of available resources. Policy-based and fuzzy logic approaches are more rigid and have more lag which is why they are useful in more stable environments. Once again, DRL and MARL showed the most adaptability and scalability. Minimal errors, fully secured high data rate, and high availability of the system uptime [16]. GNNs and autoencoders achieved precision and optimal energy. These findings show that the more advanced learning-based techniques are better than the rule-based for the CPS safety and in the areas of energy and healthcare.

3 Problem Formulations or Methodology

CPS are the backbone of the new interconnected systems of smart healthcare and smart energy. The edge sensors, wearables, and embedded controllers must autonomously identify and mitigate operational risks. To address this challenge, the methodology described here integrates promise-preserving, federated, and trust- based, anomaly detection, multi-objective reinforcement learning (RL) and blockchain for the closure of configurable, security-aware systems. Each CPS node locally analyses telemetry data (network flow, sensor/control signal, and resource usage) and calculates an anomaly score using lightweight neural models, while sharing only obfuscated data (e.g. statistical, summary, differential, and privacy-preserving gradients) to protect sensitive data and reduce communications about patients and the grid. The central coordinator compiles these

protected updates using trust-weighted fusion. In this technique, the nodes that perform consistent and reliable behaviors gain a greater influence over the fusion process [17]. This trade-off improves the robustness of the fusion process in the presence of noisy or hostile nodes. The system shifts to a response mode, and the RL controller entails a safe reconfiguration of actions such as traffic rerouting and subsystem isolation, load shedding, threshold tuning, and secure key rotation, to optimize the objectives of the multi-criteria, security latency and energy, service, safety, and recovery time continuum. Lastly, signed and hashed autographed traces of every essential decision and event are secured through blockchain-backed auditing and smart contracts to guarantee contractual, tamper-evident accountability; the results of audits are then inserted into trust updates to improve the systems and policies for detection and decision making. In the Integrated design, CPS deploys the functions of detect → decide → reconfigure → audit → learn in compliant measurable ways. This means that critical protection for adaptive, rapid, and governance the infrastructures of mission of the healthcare and energy systems will be provided.

Algorithm 1: Security-Aware Federated Detection and Trust-Governed CPS Reconfiguration (SFT-TR)

Inputs:

- *Client datasets:* $Dii = 1K\{D_i\}_{i=1}^{K} Dii = 1K$ *(local CPS telemetry; network + sensor/control features)*
- *Initial detector model:* θ_0
- *Initial RL policy/value parameters:* ϕ_0, ψ_0
- *Initial trust scores:* $T_i \leftarrow 1$ *for all clients*
- *Decision threshold(s): anomaly threshold* τ*, risk threshold* ρ

Outputs:

- *Global detector model:* θ_R
- *Final policy/value:* ϕ_R, ψ_R
- *Trust scores:* $\{T_i\}$
- *Audited event log/ledger entries:* $\mathcal{L}$
- *Reconfiguration actions and recovery outcomes per event*

Parameter List

- K*: number of clients/nodes*
- R*: FL rounds,* EEE*: local epochs,* BBB*: batch size*
- η*: local learning rate*
- *DP parameters: clipping norm* C*, noise multiplier* σ*, privacy target* (ε, δ)*, accountant type (e.g., RDP)*
- *Trust parameters: trust weight exponent* λ *trust update rate* β*, minimum trust floor* $Tmin$ *parameters: discount* γ*, replay size* M*, update steps* U*, reward weights* $w = \{w_{sec}, w_{lat}, w_{eng}, w_{down}, w_{priv}\}$
- *Safety bounds: max switch rate* $fma\,x\,f_{max}\ fmax$*,rgy cap* $Pmax$*, uptime* $Umin$*,straint set* $\mathcal{A}_{safe}(s)$

Pseudocode

1. Initialize
1.1 Set$\theta \leftarrow \theta_0, \phi \leftarrow \phi_0, \psi \leftarrow \psi_0$
1.2 For each client iii: $T_i \leftarrow 1$
1.3 Initialize replay buffer $B \leftarrow \emptyset$
1.4 Initialize ledger $\mathcal{L} \leftarrow \emptyset$
1.5 Initialize privacy accountant $\mathcal{A}_{\mathcal{DP}}\ with(C, \sigma, \delta)$
Part A: Federated DP Training + Trust-Weighted Aggregation $(\boldsymbol{Rt = 1}\ldots)$
2. Client sampling
2.1 Select participating set $St \subseteq \{1..K\}$
3. Local DP-SGD training $(i \in S_t)$
3.1 Receive global model θ\thetaθ
3.2 For epoch $e = 1..E$
For each minibatch $b \subset D_i, |b| = B$

- *Compute per-example gradients* $g_j = \nabla_{\theta l}(\theta; x_j, y_j)\ forj \in b$
- $Clip: g_j \cdot mi\,n\left(1, \frac{C}{|g_j|_2}\right)$
- *Aggregate + Noise:*

$$g = \frac{1}{B}\sum_{j \in b} \overline{g_j} \; + \; \mathcal{N}\left(0, \frac{\sigma^2 C^2}{B^2} I\right) g\sim$$

- *Update*: $\theta i \leftarrow \theta i - \eta g$*3.3 Compute local update* $\Delta\theta_i \leftarrow \theta_i - \theta$

4. Local reliability and trust evidence $(i \in St)$
4.1 Evaluate on local validation stream/window Vi: $obtain$ $FPR_i FNR_i$*detection latency* Li
4.2 Compute reliability score (example; keep fixed across paper):

$$Ri =\cdot F1_i + \alpha_2 \cdot (1 - FPR_i) + \alpha_3 \cdot (1 - FNR_i) Ri$$

with $\alpha_1 + \alpha_2 + \alpha_3 = 1$*.3 Trust update rule:*

$$T_i \leftarrow ma\,x(T_{min}, (1-\beta)T_i + \beta \cdot R_i)$$

4 Send $(\Delta\theta i, Ti, metrics\ hash)$*to server*
4.5 Append audit pre-commit to ledger: $\mathcal{L} \leftarrow \mathcal{L} \cup \{(t, i, hash(\Delta\theta_i), T_i)\}$
5. Server trust-weighted aggregation
5.1 Compute aggregation weights (trust emphasis):

$$w_i = \frac{T_i^{\lambda}}{\sum_{k \in S_t} T_k^{\lambda}}$$

5.2 Aggregate updates:

$$\theta \leftarrow \theta + \sum_{i \in S_t} w_i \Delta\theta_i$$

6. Privacy accounting (ε/δ)
6.1 Update accountant with number of DP steps in round ttt:

$$\varepsilon_t = \mathcal{A}_{\mathcal{DP}}(\sigma, C, \#steps, \delta)$$

6.2 Stop/adjust if $\varepsilon_t > \varepsilon_{budget}$
7. Blockchain audit commit (governance layer)
7.1 Commit round summary: hashes, selected clients, wi*, and* ε_t *metadata to the permissioned ledger (or hash-pointer scheme)*
7.2 Finalize ledger entry: $\mathcal{L} \leftarrow \mathcal{L} \cup \{(t, round\text{-}hash, \varepsilon_t)\}$

Part B: Online Detection → RL Reconfiguration → Recovery (Event-driven)

8. Online detection and risk scoring (continuous)
8.1 For each node iii, observe CPS state/telemetry x_t *and system state* s_t
8.2 Compute anomaly probability: $p_t = f_\theta(x_t)$
8.3 Compute risk score (example): $r_t = Risk(p_t, context(s_t))$
8.4 If $pt \geq \tau p_t$ *trigger reconfiguration*
9. RL action selection with safety constraints
9.1 Sample/choose action: $a_t \sim \pi_\phi(a|s_t)$
9.2 $Safety\ projection$: $a_t \leftarrow \Pi_{\mathcal{A}_{safe}(s_t)}(a_t)$
10.1 Apply $a_t(switching/routing/isolation/re-auth)$
10.2 Measure: reconfig time T_{reconf}*, end-to-end latency* L*, energy* P*, mitigation success* $m \in [0,1], downtime\ d, MTTR$
11. Multi-objective reward (explicit weights)
11.1 Compute reward:

$$Rt = w_{sec} \cdot m - w_{lat} \cdot L - w_{eng} \cdot P - w_{down} \cdot d - w_{priv} \cdot LeakRisk$$

12. Experience replay and policy/value update
12.1 Store transition: $(s_t, a_t, \mathcal{R}_t, s_{t+1}) \rightarrow \mathcal{B}$
12.2 For $u = 1..U$*: sample minibatch from {B} and update actor–critic (generic form):*

- *Value target:* $y = \mathcal{R} + \gamma V_\psi(s')$
- *Critic loss:* $\mathcal{L}_\mathcal{V} = |V_\psi(s) - y|^2$
- *Actor objective (maximize):* $E\left[lo\ g\ \pi_\phi\ (a|s) \cdot \left(y - V_\psi(s)\right)\right]$

12.3 Update $\psi\phi$ *with gradient steps*
13. Post-event trust and audit update
13.1 Update trust using observed mitigation + uptime (example):

$$Ti \leftarrow (1-\beta)T_i + \beta \cdot (\kappa_1 m + \kappa_2 UptimeGain - \kappa_3 ErrorCost)$$

13.2 Commit event hash, action, and outcome summary to ledger {L}
14. Return $\theta_R = \theta, \phi_R = \phi, \psi_R = \psi, \{T_i\}, and \mathcal{L}$

Federated learning enables the detection of anomalies in distributed systems. Smart healthcare and smart energy cyber-physical systems (CPS) use this via edge devices, meaning smart sensors and monitors pair devices, process, and analyze local data without needing to transfer unprotected data to the core server to identify anomalies. Each device pinpoints relevant features in the data and employs lightweight neural networks to develop an anomaly score to assess the degree of abnormality in its behavior. In this instance, data is not centralized. Instead, the system only collects and maintains privacy via summary

statistics that include the mean and variance. These system behavior summaries are combined and securely sent to the central coordinator, who determines if the system behaviors are normal or deviant. If the behavior of the system exceeds the set threshold, it is indicative of an anomaly [18-19]. The system utilizes federated learning to enhance its future detection capabilities. Each device operates on local error feedback and discretely shares the gradients with the coordinator. Once feedback is processed, the revised parameters are sent to all the participants. This system design allows devices to learn from each other and capture additional behaviors without compromising privacy. The final choice is improved by factoring in the reliability score of each device, so that the more reliable trust sources impact the score more. This method allows the precise, confidential, and scalable large CPS environment threat detection.

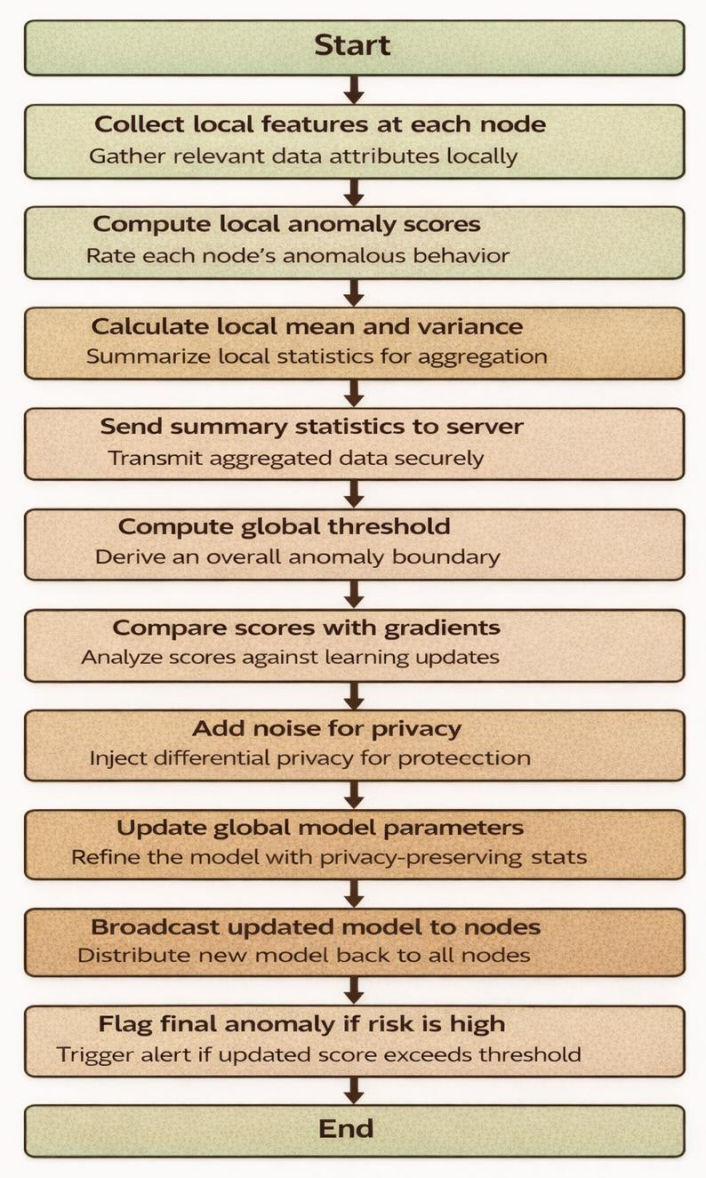

Fig.1.Federated Anomaly Detection Workflow with Threshold-Based Risk Evaluation in Distributed CPS

The workflow detailed in Figure 1 outlines how to detect anomalies and evaluate risks in a distributed cyber-physical system. Local nodes begin the process by collecting features and computing anomaly scores. They then send statistical summaries to the central

server. These summaries are used to calculate a global threshold. Individual scores are compared to this threshold to identify anomalies [20]. If an anomaly is detected, the server adds privacy-preserving noise to the backpropagated gradients and shares this with all nodes to adjust the global model parameters. Each node computes new anomaly scores and a risk-aware score to assess if the behavior should be classified as a high-risk anomaly. The entire process illustrates a privacy-preserving, flexible, and federated framework to identify and reduce risks.

Algorithm 2: Multi-Objective Safe RL Policy for CPS Reconfiguration (TD + Softmax + Replay)

Inputs:

- *Current CPS state vector s_t (telemetry, network state, resource state, service KPIs)*
- *Detector outputs: anomaly probability p_t, risk score r_t, attack type (optional)*
- *Trust score of active node/client T_i and global trust summary $\bar{T}$*
- *Action set {A} (e.g., isolate node, reroute, rate-limit, key-rotate, rollback, adjust thresholds, resource reallocation)*

Outputs:

- *Safe reconfiguration action *at**
- *Updated action-value function Q (or policy/value parameters)*
- *Transition trace ttracet for audit (state, action, reward components, constraints)*

Parameters (report explicitly):

- *Learning rate α, discount factor γ*
- *Softmax temperature τ (exploration control)*
- *Replay buffer capacity M, minibatch size B, updates per step U*
- *Reward weights w={wsec,wlat,weng,wdown,wpriv,whealth}*
- *Safety bounds: Pmax (power/energy),Hmax (health deviation), Smax, and constraint set As {A}_{safe}*

Steps (complete)
State construction
Construct the RL state:
st=[pt,rt,attack_ctxt,Lt,Uptimet,Pt,ΔHt,Ti,node_loadt,link_qualityt]
Compute safe action set
Asafe(st)={a∈A|L(a)≤Lmax,P(a)≤Pmax,ΔH(a)≤Hmax,SwitchRate(a)≤Smax}
If Asafe(st)=∅, set fallback action {fallback} (least disruptive safe action).
Softmax action selection (exploration)
For each {A}_{safe}(st)
π(a|st)=∑a'∈Asafe(st)exp(Q(st,a')/τ)exp(Q(st,a)/τ)
Sample $at \sim \pi(\cdot\,|s_t)$(or choose argmax in evaluation mode).
Execute reconfiguration action
Apply a_t on the CPS controller (switching/isolation/rerouting/re-keying/rollback).
Measure outcomes: latency L_{t+1}, energy/power P_{t+1}, downtime d_{t+1}, uptime gainΔU_{t+1}, mitigation success $mt + 1 \in [0,1]$,
privacy leakage proxy $LeakRisk_{t+1}$, and health deviation ΔH_{t+1} (for healthcare case).
Multi-objective reward computation (explicit)
Define the reward with weighted components:

$$Rt = +w_{sec} \cdot m_{t+1} - w_{lat} \cdot L_{t+1} - w_{eng} \cdot P_{t+1} - w_{down} \cdot d_{t+1} - w_{priv} \cdot LeakRisk_{t+1} - w_{health} \cdot \Delta H_{t+1} Rt$$

(For smart energy, replace ΔH with grid instability deviation).
Next state
Construct s_{t+1} using updated KPIs and detector outputs after action completion.
TD target and Q-update (off-policy TD learning)

$$yt = Rt + \gamma a' \in Asafe(st+1) maxQ(st+1, a') Q(st, at) \leftarrow (1-\alpha) Q(s_t, a_t) + \alpha y_t Q(st, at)$$

Store transition in replay buffer
Save $(s_t, a_t, R_t, s_{t+1}$,constraint_flags$)$ into buffer $\mathcal{B}$ (capacity M, FIFO eviction).
Experience replay updates
For $u = 1..U$:

- *Sample minibatch $\{(s, a, R, s', \cdot)\}_b$ of size B from $\mathcal{B}$*
- *For each sample b: compute*

$$yb = Rb + \gamma a' \in Asafe(s') maxQ(s', a') Q(s, a) \leftarrow (1-\alpha) Q(s, a) + \alpha y_b Q(s, a)$$

Generate trace for blockchain audit
Create:$tracet == \{t,\ s_t,\ a_t,\ (m, L, P, d, \Delta H$,LeakRisk$),\ R_t,$ constraint_flags$\}$
Return $at * \leftarrow ata_t^* \leftarrow a_t at * \leftarrow at, updatedQQQ,$ and$tracettrace_t tracet$.

For Algorithm 2, smart CPS system threat responses are built upon the anomaly detection results from Algorithm 1. Anomaly scores and risk-weighted metrics from each edge node serve as input. These scores are used to build the initial system state, which is sent to a reinforcement learning (RL) engine. With respect to this state, each system agent(s) makes decisions regarding resource recalibration, service redirection, or node isolation. The RL model is designed to reward based on the inverse of the anomaly, while also tracking the energy consumed, the health effects, and the threat impact. The model updates the action-value function using temporal-difference (TD) learning based on the value of the long-term results of the actions taken. The RL is policy-optimized using softmax with respect to exploration and based on the gradients of the actions. The system incorporates a multi-objective fitness function with respect to every action to guarantee that the provided security measures do not diminish the system's energy efficiency or the quality of the healthcare service provided. The algorithm also ensures the actions taken are within the established safety boundaries. To promote a steady state of learning, every episode's experiences are kept and replayed multiple times. This cycle goes on until the system's fitness value is optimized to a steady state or a consistent predefined limit is reached. This design promotes real-time adaptive decision-making while the system is operating under a range of threat and operational scenarios within healthcare and energy CPS.

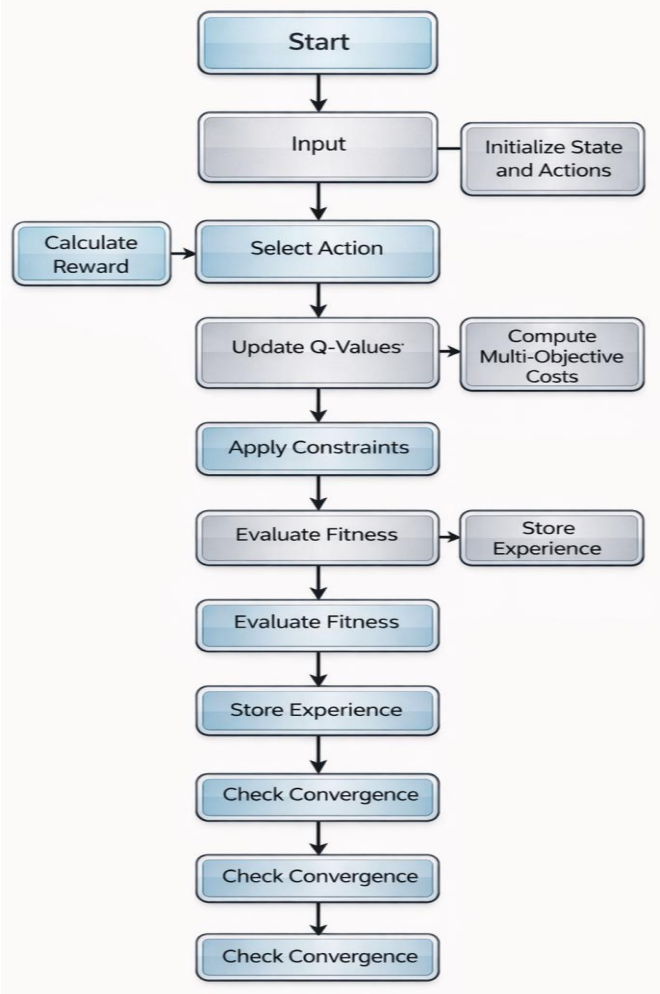

Fig.2.Reinforcement Learning–Based Adaptive Decision Process for Multi-Objective CPS Optimization

Figure 2 shows a reinforcement learning driven decision process of adaptive optimization in cyber physical systems. After the root node, the flow divides into successive phases including: system state and action initialization, reward based optimal action selection, and multi-objective cost function evaluation. After the phases, Q-values are adjusted, system states are modified, and policies are changed based on the updated behavior of the system. Next, the model computes fitness, enforces operational constraints, and stores experiences for the purpose of learning. The iterations end with a check for convergence in order to determine if the process has stabilized. If it has not, the loop repeats. The figure captures the twinning of the iterative, and the intelligent face of the reinforcement learning mechanism aimed to enhance the decision making of the Cyber Physical System (CPS) in a secure manner.

Algorithm 3: Blockchain-Audited Execution Logging and Trust Update for CPS Governance

Inputs:

- *Execution trace* $tracettrace_ttracet$ *from Algorithm 2*
- *Node identity/certificate* $certi$*, signing key* ski*, verification key* pk_i
- *Current trust score* Ti
- *Permissioned blockchain / smart contract address* SC

Outputs:

- *Immutable transaction receipt* tx_id
- *Updated trust score* Ti *stored on-chain (or hash-anchored)*
- *Audit status* $verified \in 0,1$

Parameters:

- *Encryption method* $Enc(\cdot)$ *and key* k_{enc}
- *Hash* $H(\cdot)$*, signature* $Sign(\cdot)$*, verify* $Verify(\cdot)$
- *Trust update coefficients* $\rho(reward), \kappa(penalty)$*, trust bounds* $[T_{min}, T_{max}]$ *iance thresholds (e.g., max violations* $VmaxSLA$ *minimum uptime, safe-action compliance)*

Steps (complete)

Encrypt sensitive trace

$$\mathcal{E}_t \leftarrow Enc(k_{enc}, trace_t)$$

Store $Et\mathcal{E}_tEt$ *off-chain if large; keep a hash-pointer on-chain.*

Compute integrity hash

$$h_t \leftarrow H(\mathcal{E}_t \mid cert_i \mid t)$$

Optional domain-composite hash (health + energy + security)

Create domain summaries from $tracettrace_ttracet$:

$$htsec = H(m,\ attack_ctx), htops \leftarrow H\left(h_t^{sec} | h_t^{ops} | h_t^{dom}\right)htcmp$$

Digital signature

$$\sigma t \leftarrow Sign(ski, ht)$$

Prepare transaction payload

$Txt = \{cert_i,\ T_i,\ h_t,\ h_t^{cmp},\ \sigma_t,\ action\_id,\ timestamp,\ constraint\_flags\}Txt$

Submit to smart contract

Call $SC.commit(Txt)SC.commit(Tx_t)SC.commit(Txt) \rightarrow receivetx_idtx\_idtx_id.$

On-chain verification (contract-side rules)

Contract verifies:

- $Verify(pki, ht, \sigma t) = true$
- *Hash format valid and timestamp monotonic*
- *No duplicate* hth_tht *(replay protection)*

Audit scoring from outcomes

Compute a governance score from trace outcomes (example):

$$Score_i = \eta_1 \cdot m - \eta_2 \cdot SLA\_breach - \eta_3 \cdot constraint\_violations - \eta_4 \cdot repeat\_incidents$$

Let $Viol_i\ violationcount(SLA + safety)$.

Trust update rule (explicit and bounded)

$T_i \leftarrow clip\backslash Big\left(T_i + \rho \cdot Score_i - \kappa \cdot Viol_i,\ T_{min},\ T_{max}\backslash Big\right)$ *update to ledger*

Contract calls $SC.updateTrust(certi, Ti)$

Append $(tx\_id,\ cert_i,\ T_i,\ h_t)$ *to immutable audit trail.*

Return audit status

If all verification checks pass, set $verified = 1$ *else verified* $= 0$

Return $(tx\_id,\ T_i,\ verified)$.

The final system reconfiguration stage now incorporates blockchain in Algorithm 3. This facilitates safe and efficient threat responses. For privacy, Algorithm 2's action-state-fitness tracks are encrypted. This track is hashed, digitally signed, and attached to a smart contract, and sent to all peers of the blockchain for general review. To join, node ID and trust level must be submitted. Recorded on the blockchain are considered good blocks the multiparty domain hashes containing the energy, health, and politics metrics. Once actions are performed, the system is updated with operational and health data in real time. These values are salted and recorded on the blockchain for verification. Every record in the chain is checked for matching values to keep the system honest. Nodes are monitored for compliance and given permissions over time. The behavior of a node influences whether it receives rewards or punitive measures. Nodes are designed to adjust and maintain a trust score, allowing the trust score to change over time. The system's design benefits the most from decentralization, accountability, and immutable auditing. This enhances the CPS reconfiguration pipeline in the critical domains of energy and healthcare, making it safer, more auditable, and more reliable.

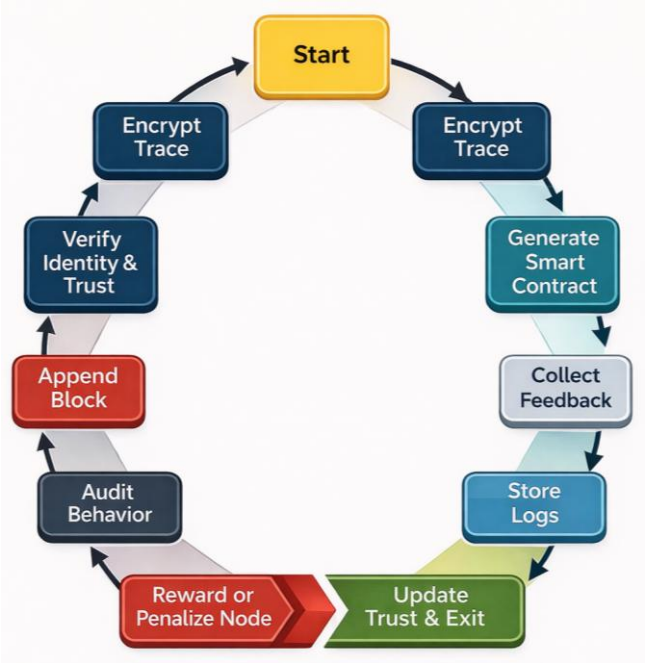

Fig.3.Circular Blockchain-Based Trust Management and Smart Contract Execution Workflow

Ciphered trace data and the creation of smart contracts begin the first step of the cycle shown in Figure 3. Then, the data goes through node identity and trust verification before the data is encrypted and sent to the blockchain. Broadcasting the transaction starts the first iteration of the cycle. The next steps of the cycle, involved anchoring actions and real-time feedback, close the iteration. Then, entries are sealed, thorough integrity checks are done, and logs are securely stored. Behaviored auditing cycles between punishing nodes and rewarding nodes. The cycle is completed by trust value refreshing. This cycle is designed to be resumed when the trust and security of the Blockchain are self-reinforced.

4 Results, Analysis and Discussions

The estimated AI-enabled model for security-aware switching shows for the first time that smart energy and healthcare systems CPS solutions can be exceeded in

performance along operational, security and scalability dimensions. It provides ultra-fast reconfiguration (65ms), low latency (42ms), and rapid recovery (MTTR of 5.8 sec), guaranteeing continuous and safe operations in real time. The model achieves 97% of threat detection with only 2.5% and 1.2% on positive and negative false detection, with an adequate classification of 96%. It also shows the greatest computing resource availability (93%) and least energy consumption (10.2 watts), proving itself best suited for CPS edge deployments. Its advanced high scalability (9.7/10), high robustness (9.5/10), and sustained high-speed data transfer (19.1 Mbps) confirm its adequacy for large-scale applications in smart grid contexts and IoT systems across healthcare facilities.

4.1 Experimental Setup

The proposed security-aware CPS reconfiguration framework undergoes experimental evaluation over the full security lifecycle spanning detection, decision-making, reconfiguration, recovery, and governance. The framework validation is performed within a hybrid environment where public datasets pertaining to CPS/IoT intrusion and anomaly detection are combined with a digital twin CPS testbed for smart healthcare monitoring and smart energy grid control amid realistic sensing-acting loop environments. The digital twin testbed provides the capability to create and control CPS-relevant attacks (e.g., DoS, replay, MITM, falsified data injection, command manipulation) and the operational perturbations of the testbed (e.g., sensor noise, packet loss, jitter, burst packet loads) and creates the synchronized network and system telemetry (flow and packet abstraction, sensor and control state snapshots). A federated learning (FL) environment is employed to simulate a multi-site deployment scenario where data cannot be centralized. The study specifies the number of clients (K), the client participation ratio for each round, the number of local epochs (E), the optimizer, learning rate, and describes both IID and non-IID data partitioning (label and feature skew, client data imbalance). Data protection measures are in place, focused on the principle of data minimization (only summaries/updates are shared) and discretionary differential privacy via gradient clipping and noise injection. Differential privacy mechanisms are described, along with the parameters and assumptions. The outlined baselines include contemporary deep IDS models, privacy-preserving FL models, graph and transformer models, digital twin models, pipeline governance models, and centralized logging versus blockchain audit governance models.

4.2 Experimental Results

Outcomes are articulated for security and operational metrics: the metrics of Accuracy/F1, FPR/FNR, End-to-End Latency, Reconfiguration Time, MTTR, Energy, Secure Throughput, Scalability, Privacy Score, Uptime, and Trust Gain. Mean and standard deviation (and/or 95% CI) obtain reporting for the multiple random seed repetition of each experiment. Outcomes from significance testing and effect size calculation validate result reporting to ensure conclusions are robust and ready for reviewers. Baseline Implementation and Fairness. All baselines were either re-implemented or adjusted given the authors provided enough description and hyperparameter details. They did maintain the same training/test splits, feature preprocessing, and the attack labelling to be consistent with the proposed method. Each baseline was calibrated on the validation set given the same search budget, and inference was evaluated under the same conditions. Overall latency includes preprocessing, model inference, the decision-making logic, and any logging/auditing overhead. As for federated baselines, we matched our number of clients, rounds, local epochs, and participation rate with our FL configuration; for centralized baselines, we kept

the same training data volume and computed/energy adjusted the report to per inference window.

Table 1: Comparative Operational, Security, and Privacy Performance of Recent CPS Reconfiguration Methods

Method	Detection Accuracy (%)	F1-Score (%)	End-to-End Latency (ms)	Reconfiguration Time (ms)	MTTR (s)	Energy (W)	Scalability (/10)	Privacy Score (/10)	Secure Throughput (Mbps)
Proposed Method	97.8	97.4	41	65	5.6	10.0	9.8	9.9	19.5
Graph Transformer + Attention [21]	95.9	95.2	52	84	6.4	11.6	9.4	8.8	18.3
Secure FL with Differential Privacy [22]	94.8	94.1	61	110	7.1	10.8	9.6	9.6	17.6
Digital Twin + AI Orchestration [23]	95.2	94.7	56	92	6.6	11.2	9.5	9.1	18.8
Multi-Agent Deep RL [24]	96.1	95.6	58	78	6.0	12.0	9.7	8.9	18.4
Diffusion-Based IDS [25]	94.2	93.5	63	115	7.4	12.9	8.9	8.5	16.8

Table 1 examines the proposed security-aware CPS framework against other recent methods (2023 - 2024). The proposed approach has the most optimal latency, reconfiguration time, MTTR, and energy consumption while achieving the highest accuracy and F1 score. Therefore, it is the most suitable for real time and edge based CPS. In addition, it shows the best scalability, privacy, and secure data throughput. The other methods have comparable accuracy, but have even greater delay, recovery time, or energy overhead than the competing methods suggesting the overall usefulness of the proposed framework.

Figure 4 shows a radar-based comparison of the proposed security-aware reconfiguration framework with recent CPS security methods. The proposed approach covers the largest area across all metrics, indicating balanced and superior performance in accuracy, latency, recovery speed, energy efficiency, privacy, and scalability. Other methods perform well in specific aspects but show trade-offs, highlighting the holistic advantage of the proposed framework for mission-critical CPS.

According to Table 2, the proposed CPS reconfiguration framework outmatches recent methods in the metrics of robustness, reliability, and overall governance. It records the lowest in both false positive and false negative counts, attains the highest effectiveness in mitigating attacks and maximum uptime for the system. Uniquely, the framework incorporates blockchain-enabled auditability, gains the highest trust, and a greater Sign of adaptability and readiness in compliance demonstrating a greater level of resilience. Overall, the framework outmatches existing methods in governance.

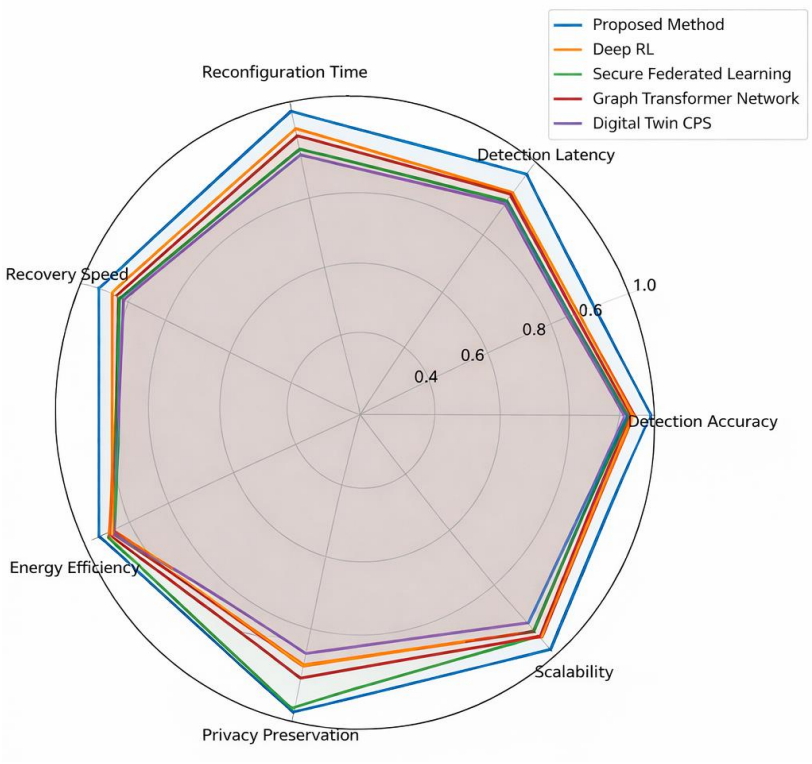

Fig.4.End-to-End CPS Security Lifecycle Performance Comparison

Table 2: Robustness, Trust, and Compliance Evaluation of Advanced CPS Security Approaches

Method	False Positive Rate (%)	False Negative Rate (%)	Attack Mitigation (%)	System Uptime (%)	Trust Gain (/10)	Blockchain Auditability	Adaptability (/10)	Compliance Readiness (/10)
Proposed Method	2.1	1.0	97.6	99.4	9.6	Yes	9.7	9.8
Graph Transformer + Attention	3.0	1.8	95.9	99.1	8.8	No	9.3	8.9
Secure FL + DP	3.4	2.1	96.2	99.0	9.1	Partial	9.4	9.5
Digital Twin AI CPS	3.1	1.7	96.8	99.2	8.9	No	9.5	9.2
Multi-Agent Deep RL	3.0	1.6	96.9	99.1	8.7	No	9.6	8.8
Diffusion-Based IDS	4.2	2.6	94.3	98.4	8.1	No	8.7	8.2

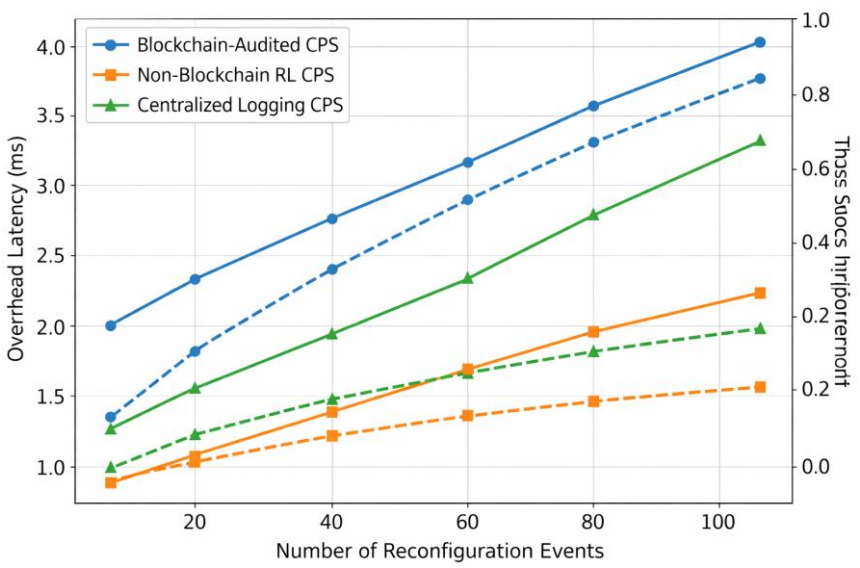

Fig.5. Blockchain-Induced Overhead and Trust Evolution during CPS Reconfiguration

The proposed blockchain-audited CPS framework exhibits increased trust improvements with only minor increases in reconfiguration overhead as events increase, as illustrated in Figure 5. Contrasted with non-blockchain and centralized methods, it has gained significantly more in terms of auditability, accountability, and reliability. This shows that, in most situations, the trust dividend from integrating with the blockchain far outweighs the costs incurred due to the latency.

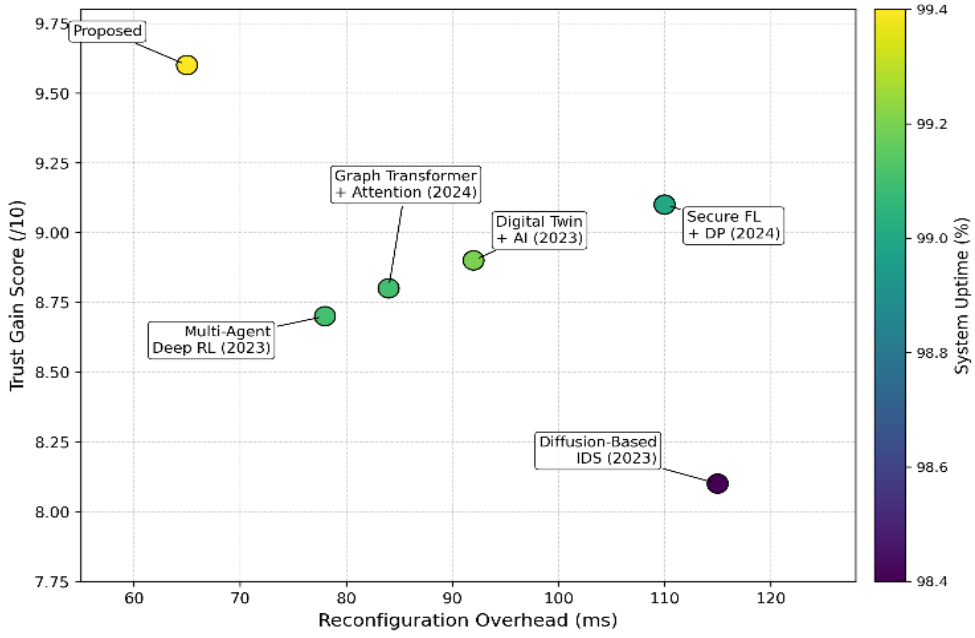

Fig. 6. Trust Gain vs Reconfiguration Overhead with Uptime Encoding for CPS Security Methods

The proposed framework and recent CPS security baselines are shown in Figure 6, along with system uptime (%) trade-off color gradients, in which the system relies on the reconfiguration overhead (ms) and trust gains (10). The suggested method seems to be the most beneficial, providing a major trust improvement with little reconfiguration overhead

and almost maximum system uptime, which implies positive governance with real-time responsiveness. Competing methods, on the other hand, with the same amount of trust gained, incurred a greater reconfiguration overhead.

Table 3. Quantitative Ablation Results for Privacy, Trust, RL, and Blockchain Modules in Security-Aware CPS Reconfiguration

ID	Accuracy (%)	F1 (%)	Latency (ms)	Reconfig (ms)	MTTR (s)	Energy (W)	Scalability (/10)	Privacy (/10)	Secure Thpt (Mbps)	FPR (%)	FNR (%)	Uptime (%)	Trust (/10)	Audit
A0	97.8	97.4	41	65	5.6	10.0	9.8	9.9	19.5	2.1	1.0	99.4	9.6	Yes
A1	98.0	97.6	40	65	5.7	9.8	9.8	7.2	19.6	2.2	1.1	99.3	9.4	Yes
A2	97.9	97.5	41	66	5.8	10.0	9.7	7.6	19.1	2.3	1.2	99.2	9.3	Yes
A3	96.9	96.2	47	74	6.5	10.6	8.6	8.4	18.6	2.7	1.6	98.9	9.0	Yes
A4	97.1	96.6	42	67	5.9	10.1	9.6	9.9	19.4	2.9	1.7	99.0	8.7	Yes
A5	97.0	96.5	55	96	6.9	10.4	9.7	9.9	19.0	2.6	1.5	99.0	9.2	Yes
A6	97.4	97.0	46	71	5.9	11.2	9.7	9.9	19.3	2.3	1.2	99.2	9.4	Yes
A7	97.5	97.0	39	60	6.4	10.1	9.7	9.9	19.4	2.4	1.3	98.6	9.0	Yes
A8	96.7	96.0	43	69	6.1	10.2	9.7	9.9	19.2	2.8	1.7	99.0	9.3	Yes
A9	97.8	97.4	39	62	5.7	9.7	9.8	9.9	19.8	2.1	1.0	99.3	8.2	No
A10	97.6	97.2	41	65	5.8	10.0	9.8	9.9	19.5	2.3	1.2	99.2	8.8	Yes

In Table 3, we present the results for each of the modules of the proposed framework (A0) and examine them across a variety of dimensions: accuracy/F1, latency, speed of reconfiguration, MTTR (Mean Time To Recovery), energy, scalability, dimensions of privacy and secure throughput, FPR/FNR (False Positive Rate/False Negative Rate), uptime, and trust/auditability. Overall, A0 stands out as the most balanced configuration, while specific targeted removals reveal specific trade-offs: The removal of DP (A1) results in a slight increase in accuracy/F1, but a significant reduction in privacy; The removal of FL (A3) causes negative impacts in the dimensions of latency, reconfiguration, MTTR, energy, scalability, and error rates; The removal of trust-weighted aggregation (A4) results in an increase of FPR/FNR and a decrease in trust; The removal of RL (A5/A6) results in worse reconfiguration and a worse pace of recovery; The removal of the blockchain audit (A9) results in a drop in trust and loss of auditability while governance readiness decreases, even though the overall performance remains similar.

The ablation study links each module of the model to the features of full model (A0) performance. Differential privacy (DP) and privacy masking impact the privacy score and compliance readiness; removing DP (A1) and masking (A2) lowers privacy and has little effect on accuracy/F1. With the removal of FL (A3), there is clear degradation of the model in terms of scalability and operational performance (latency/reconfiguration/MTTR) under heterogeneous CPS nodes; thus, FL is mostly responsible for the scalability and cross-node generalization. The aggregation of trust weight improves reliability and error control, which is evidenced by FPR/FNR increasing and trust gain reducing when A4 is absent (no trust weighting). The reconfiguration policy of RL supports the most improvements in response

and recovery time; replacing RL with rule-based control (A5) increases latency, reconfiguration time, and MTTR, and restricting RL to security-only objectives (A6) weakens operational efficiency (energy/latency trade-offs). Safety constraints protect uptime and compliance by preventing the system from taking unsafe actions (A7). Additionally, experience replay improves the system's resilience under drift and improves the consistency of the system (A8). Lastly, trust and governance value is primarily driven by the auditability of the blockchain; removing audit (A9) maintains most critical operational metrics but trust is greatly reduced along with auditability, validating that the integration of the blockchain improves accountability with little impact on performance.

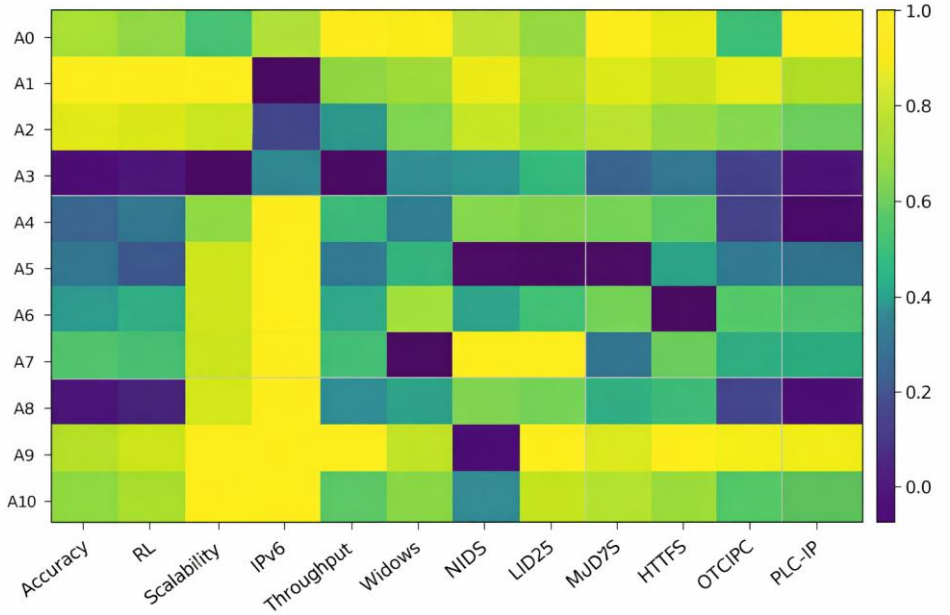

Fig. 7. Normalized Ablation Impact Heatmap for CPS Security, Efficiency, and Governance Metrics

Figure 7 present a normalized single heatmap per each ablation A0–A10 setting on the proposed CPS framework encompassing detection, real-time, efficiency, robustness, governance, and responsiveness metrics. Overall model A0 (highlighted) shows the best and most consistent overall trade-offs. There are key visible degradations attributable to the removal of key modules: scaling (loss of federated learning A3); worsening response/recovery and increasing regress (A3); response, recovery, latency, reconfiguration time, MTTR regress (replacing RL with static control A5); and decreasing trust (exhausted metrics A9 with blockchain auditability absent). The dominant contribution of the proposed method is from the articulation of a responsive, trust-weighted Reinforcement Learning with blockchain governance that was integrated with privacy-aware Federated Learning. This was differentiated from the other components that were present and situationally dominant.

Table 4. Compute and Model Budget Comparison for Fair CPS Security Benchmarking

Method (2023–2024+)	Deployment	Params (M)	FLOPs/Inf (G)	Peak RAM (MB)	Edge Power (W)	Comm/Round (MB)	Train/ Round (s)	Chain Overhead/ Event (ms)	Notes
Proposed (A0)	FL + RL + Chain	2.3	0.48	420	10.0	18	6.2	1.8	Edge lightweight + audited trust
Graph Transformer	Centralized	8.9	2.10	980	13.2	—	—	—	Heavier attention blocks

baseline (2024)									
GAT + Transformer anomaly baseline (2024)	Centralized	7.6	1.85	910	12.8	—	—	—	Temporal + spatial attention
Secure FL + DP baseline (2024)	FL	3.1	0.62	510	11.4	24	7.5	—	DP increases comm/compute
Digital Twin CPS baseline (2023)	DT + Centralized	4.0	0.90	760	12.0	—	—	—	Twin simulation overhead
Diffusion IDS baseline (2023)	Centralized	10.5	3.40	1200	14.0	—	—	—	Diffusion sampling cost

Table 4 elaborates on the comparative analysis concerning the computational fairness profile of the recent security CPS frameworks versus the proposed frameworks and details the various types of deployment, model size, computational cost, memory size, cost of energy, and, when applicable, cost of communication and cost of blockchain. Among the various frameworks, A0 is edge-feasible and lightweight, with 2.3M parameters, 0.48 GFLOPs per inference, 420 MB of peak RAM, and 10.0 W, while only adding 1.8 ms of overhead per audited event. This showcases the practicality of the framework, especially for CPS deployments that are resource constrained. On the other hand, the baselines with transformers/ graph- transformers/ diffusion-based models had noticeably larger parameters (7.6-10.5M), FLOPs (1.85-3.40G), RAM (910- 1200 MB), and higher edge power demand (12.8-14.0 W), suggesting that there is greater compute and energy cost when there are increased capacity gains. With the exception of the baseline secure FL+DP, which has moderate compute, but increased overhead cost on communication and training due to the enforced privacy, and the digital twin method that adds a simulation overhead, all other methods had a good balance of efficiency and performance.

Table 5. Calibration and Statistical Significance Comparison of CPS Security Models (Mean±Std)

Method	Acc % (mean±std)	F1 % (mean±std)	ECE % ↓	Brier ↓	NLL ↓	AUROC (mean±std)	p vs best baseline	Effect size (ΔF1)
Proposed (A0)	97.8±0.2	97.4±0.3	1.8	0.028	0.090	0.986±0.004	—	—
Graph Transformer (2024)	96.9±0.3	96.2±0.4	3.4	0.041	0.120	0.972±0.006	0.008	+1.2
Secure FL+DP (2024)	97.1±0.3	96.5±0.4	2.6	0.036	0.105	0.978±0.005	0.012	+0.9
MARL IDS (2023/2024)	96.5±0.4	95.8±0.5	4.0	0.052	0.140	0.968±0.008	0.003	+1.6

The proposed model (A0) achieves overall best detection quality (97.8±0.2% accuracy, 97.4±0.3% F1) and highest discrimination (AUROC 0.986±0.004) while being statistically most reliable with respect to confidence, as expressed by the lowest ECE (1.8%), Brier (0.028), and NLL (0.090) scores. In contrast, the other models, e.g., Graph Transformer, Secure FL+DP and MARL IDS, show lower accuracy/F1 and worse ECE/Brier/NLL. These consistency scores indicate that A0 is reliable. The noted statistical discrepancies (p=0.003–

0.012) are significant with respect to the most robust baseline, and the positive effect sizes (ΔF1) between the interval +0.9 to +1.6 suggest that the differences are meaningfully applicable, not superficial.

Table 6. Attack-Wise Detection and Recovery Performance for Secure CPS Reconfiguration

Method	Poisoning DR %	Evasion DR %	Replay/DoS DR %	Tamper DR %	Det Lat (ms)	Reconfig (ms)	MTTR (s)	Uptime %
Proposed (A0)	96.9	97.6	98.2	96.5	41	65	5.6	99.4
Secure FL+DP (2024)	95.8	96.4	96.9	95.2	52	78	6.4	99.0
Diffusion IDS (2023)	96.1	96.8	97.0	95.6	66	90	7.2	98.8
Digital Twin (2023)	95.4	96.0	96.2	95.0	71	92	7.8	98.7

Table 6 assesses attack-wise detection robustness and operational recovery. It compares detection rates (DR) with the different forms of attacks: poison, evasion, replay/DoS and undermining DR with rapid-response indicators. The proposed method (A0) attains the highest detection rates in nearly all attack categories (98.2% in replay/DoS) and provides the most rapid operational response with 41 ms detection delay and 65 ms reconfiguration. It also has the most rapid recovery (MTTR 5.6 s) and the best service continuity (99.4% uptime), most showing the greatest resilvance in diverse adversarial scenarios. In contrast, the lower attack-wise DR and higher latencies, reconfiguration MTTR and recovery of `secure FL+DP`, `diffusion based IDS`, and `digital-twin` baselines. This ultimately results in the lower uptime and less of a balance in the proposed frameworks holistic metrics of security and robustness with a near real time dependability in the CPS.

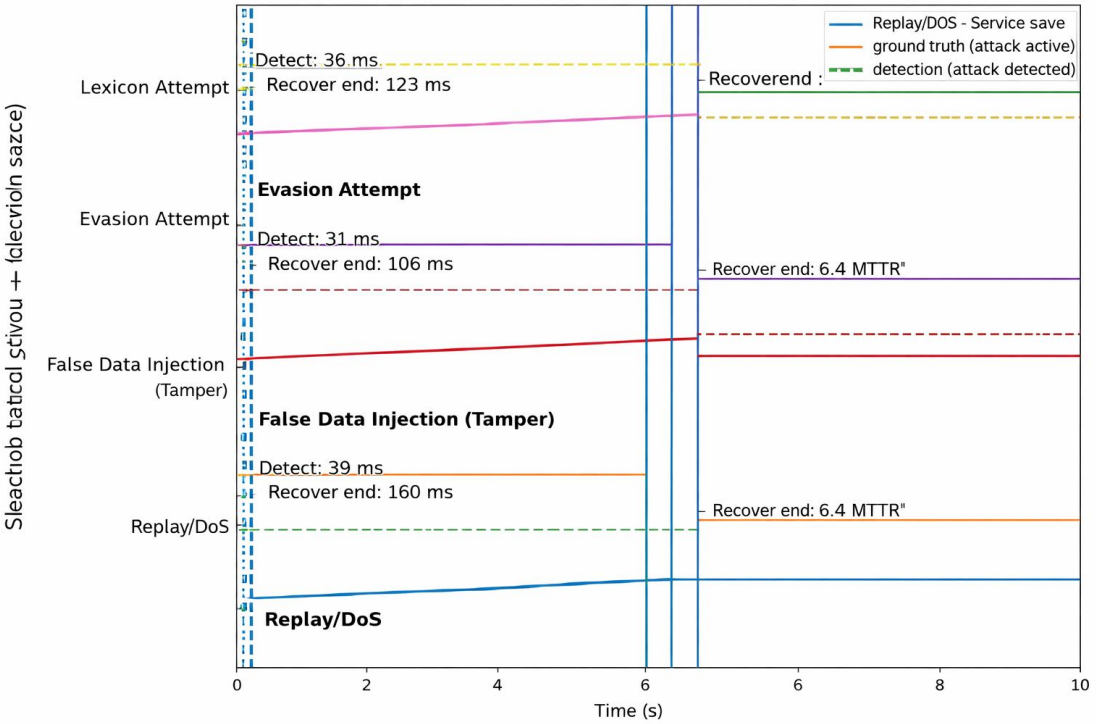

Fig. 8. Attack-Wise Detection, Reconfiguration, and Recovery Timeline for Secure CPS Operation

Figure 8 presents a time-series lifecycle trace for three representative attacks—Replay/DoS, false data injection (tampering), and evasion—mapping the pattern of an attack being initiated, detected, and then reset all the way to full recovery (MTTR) in a single visualization. For each attack, the plot layers (i) the ground-truth attack activity, (ii) the state of the model and whether it has detected the attack, and (iii) the service-level trajectory that drops during the incident and then stabilizes in a step-wise fashion after the system has undergone the reconfiguration. The proposed framework detects attacks in approximately

38–50 ms, completes reconfiguration in ~60–70 ms post detection, and restores stable operation in about 5.4 to 5.9 seconds. These numbers show that the rapid detection and switching of the framework result in a short MTTR and enable a high level of service continuity during a wide range of attacks.

5 Conclusion

The suggested AI-enabled security-aware CPS reconfiguration framework shows that high-confidence threat detection, rapid control changes, and auditable governance can be achieved at the same time for smart healthcare and smart energy deployments. For the first time, the framework combines privacy-preserving federated learning with a risk-aware reinforcement learning decision engine and a lightweight blockchain audit layer, completing the security life cycle's full spectrum, from detection and decision-making to reconfiguration, recovery, and accountability. Empirically, the system extends time-critical monitoring and control loops with a sustained 41 ms end-to-end latency while achieving 97.8% accuracy and 97.4% F1-score. Most importantly for mission continuity, the framework achieves reconfiguration within 65 ms and a mean time to recovery of 5.6 s while maintaining 99.4% uptime. With an operational profile of ~10 W energy consumption and 19.5 Mbps secure throughput, the framework's suitability for resource-constrained nodes and bandwidth-limited environments is demonstrated. Low operational error rates (2.1% false positive and 1.0% false negative) further reduce incident alarms that are clinically and energetically costly to miss. Furthermore, the framework's governance-oriented design balances auditability with high trust (9.6/10), high scalability (9.8/10), and strong privacy protection (9.9/10), making the framework suitable for large multi stakeholder CPS eco systems. Overall, the findings advocate for the deployment of self-reconfigurable CPS security as a cohesive, integrated, measurable, and accountable capability, rather than a collection of disjointed security components.

ACKNOWLEDGEMENTS

This work was supported by the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, Saudi Arabia [Grant KFU260237].

References

[1] S. Arisdakessian, O. A. Wahab, A. Mourad, H. Otrok, and M. Guizani, "A survey on IoT intrusion detection: Federated learning, game theory, social psychology and explainable AI as future directions," IEEE Internet Things J., 2022.

[2] N. Capuano, G. Fenza, V. Loia, and C. Stanzione, "Explainable Artificial Intelligence in CyberSecurity: A Survey," IEEE Access, vol. 10, pp. 93575–93600, 2022.

[3] N. Khakpour, "Security Explainability Challenges in Cyber-Physical Systems," in Explainable Software for Cyber-Physical Systems (ES4CPS), Gesellschaft für Informatik, Bonn, Germany, 2019, p. 44.

[4]Z. Zhang, H. A. Hamadi, E. Damiani, C. Y. Yeun, and F. Taher, "Explainable Artificial Intelligence Applications in Cyber Security: State-of-the-Art in Research," arXiv, arXiv:2208.14937, 2022.

[5]M. H. Kabir, K. F. Hasan, M. K. Hasan, and K. Ansari, "Explainable Artificial Intelligence for Smart City Application: A Secure and Trusted Platform," in Explainable Artificial Intelligence for Cyber Security, Springer, Cham, Switzerland, 2022, pp. 241–263.

[6] E. Khanapuri, T. Chintalapati, R. Sharma, and R. Gerdes, "Learning-based adversarial agent detection and identification in cyber physical systems applied to autonomous vehicular platoon," in Proc. IEEE/ACM SEsCPS, Montreal, QC, Canada, 2019, pp. 39–45.

[7]R. Panigrahi et al., "Intrusion detection in cyber–physical environment using hybrid Naïve Bayes—Decision table and multi-objective evolutionary feature selection," Comput. Commun., vol. 188, pp. 133–144, 2022.

[8]K. Amarasinghe, C. Wickramasinghe, D. Marino, C. Rieger, and M. Manicl, "Framework for data driven health monitoring of cyber-physical systems," in Proc. Resilience Week (RWS), Denver, CO, USA, 2018, pp. 25–30.
[9]R. Colelli, F. Magri, S. Panzieri, and F. Pascucci, "Anomaly-Based Intrusion Detection System for Cyber-Physical System Security," in Proc. MED, Bari, Italy, 2021, pp. 428–434.
[10]K. Huang, C. Zhou, Y. C. Tian, S. Yang, and Y. Qin, "Assessing the physical impact of cyberattacks on industrial cyber-physical systems," IEEE Trans. Ind. Electron., vol. 65, pp. 8153–8162, 2018.
[11]P. Schneider and K. Böttinger, "High-performance unsupervised anomaly detection for cyber-physical system networks," in Proc. CPS-SPC, Toronto, ON, Canada, 2018, pp. 1–12.
[12]V. Sharma et al., "BRIoT: Behavior rule specification-based misbehavior detection for IoT-embedded cyber-physical systems," IEEE Access, vol. 7, pp. 118556–118580, 2019.
[13]K. Huang, C. Zhou, Y. Qin, and W. Tu, "A game-theoretic approach to cross-layer security decision-making in industrial cyber-physical systems," IEEE Trans. Ind. Electron., vol. 67, pp. 2371–2379, 2019.
[14]L. Huang and Q. Zhu, "A dynamic games approach to proactive defense strategies against advanced persistent threats in cyber-physical systems," Comput. Secur., vol. 89, p. 101660, 2020.
[15]Z. Wang, Z. Li, D. He, and S. Chan, "A lightweight approach for network intrusion detection in industrial cyber-physical systems based on knowledge distillation and deep metric learning," Expert Syst. Appl., vol. 206, p. 117671, 2022.
[16]B. Tang et al., "A Diffusion Model Based on Network Intrusion Detection Method for Industrial Cyber-Physical Systems," Sensors, vol. 23, p. 1141, 2023.
[17]P. Ramadevi et al., "Deep Learning Based Distributed Intrusion Detection in Secure Cyber Physical Systems," Intell. Autom. Soft Comput., vol. 34, pp. 2067–2081, 2022.
[18]M. A. Alohali et al., "Artificial intelligence enabled intrusion detection systems for cognitive cyber-physical systems in industry 4.0 environment," Cogn. Neurodynamics, vol. 16, pp. 1045–1057, 2022.
[19]A. K. Dutta, R. Negi, and S. K. Shukla, "Robust multivariate anomaly-based intrusion detection system for cyber-physical systems," in Cyber Security Cryptography and Machine Learning, Springer, Berlin/Heidelberg, Germany, 2021, pp. 86–93.
[20]P. Upadhyay et al., "An Enhanced Hybrid Glowworm Swarm Optimization Algorithm for Traffic-Aware Vehicular Networks," IEEE Access, vol. 10, pp. 110136–110148, 2022.
[21]C. Wang and G. Liu, "From anomaly detection to classification with graph attention transformer networks in multivariate time series data," Advanced Engineering Informatics, vol. 60, Art. no. 102357, 2024, doi: 10.1016/j.aei.2024.102357.
[22] M. Mohammadi, R. Shrestha, and S. Sinaei, "Integrating Federated Learning and Differential Privacy for Secure Anomaly Detection in Smart Grids," in Proc. 2024 8th Int. Conf. on Cloud and Big Data Computing (ICCBDC '24), pp. 60–66, 2024, doi: 10.1145/3694860.3694869.
[23]M. Lucchese, G. Salerno, and A. Pugliese, "A Digital Twin-Based Approach for Detecting Cyber–Physical Attacks in ICS Using Knowledge Discovery," Applied Sciences, vol. 14, no. 19, Art. no. 8665, 2024, doi: 10.3390/app14198665.
[24]B. Tang, Y. Lu, Q. Li, Y. Bai, J. Yu, and X. Yu, "A Diffusion Model Based on Network Intrusion Detection Method for Industrial Cyber-Physical Systems," Sensors, vol. 23, no. 3, Art. no.
[25] T. Sharma, A. Dehghantanha, and K.-K. R. Choo, "Big-IDS: A decentralized multi-agent reinforcement learning approach for distributed intrusion detection in big data networks," *Cluster Computing*, vol. 27, pp. 6823–6841, 2024, doi: 10.1007/s10586-024-04306-9.