

QuantumGuard: A Post-Quantum Resilient Deception-Driven Framework for Proactive Threat Hunting and Cognitive Honeynet Orchestration in 6G-Enabled Cyber-Physical Systems

Daifallah Zaid Alotaibe[✉]

Software Engineering Department, College of Computer Science and Engineering
University of Hafr Al-Batin, Saudi Arabia

Abstract—The convergence of 6G ultra-reliable low-latency communications, massive cyber-physical actuation, and the looming threat of cryptographically relevant quantum computers exposes a fundamentally new attack surface that is poorly addressed by reactive intrusion-detection paradigms. This study presents QuantumGuard, a proactive cybersecurity framework that inverts the conventional defender posture by integrating cognitive honeynets, post-quantum-secured intelligence channels, and reinforcement-learning-driven deception adaptation across 6G-enabled cyber-physical environments. The principal contribution is the architectural integration of four previously disjoint capabilities—a cognitive honeynet orchestrator, a reinforcement-learning deception policy engine operating under partial observability, a transformer-based MITRE ATT&CK attribution network, and a post-quantum federated intelligence bus secured with CRYSTALS-Kyber and CRYSTALS-Dilithium—into a single closed-loop control architecture for 6G cyber-physical systems. To assess feasibility, we report a preliminary evaluation on the CICAPT-IIoT-2024 and Edge-IIoTset benchmark datasets, complemented by a 320-endpoint 6G slice testbed configured, as described in Section IV. Initial results indicate an attacker engagement-retention rate of approximately 96.42 per cent, MITRE ATT&CK technique-attribution accuracy of approximately 91.8 per cent, a 78.6 per cent reduction in median attacker dwell time relative to passive honeypot baselines, and a deception-induced production-traffic overhead of 1.7 per cent. The PQ-FIB sustains a 14.2 ms median post-quantum handshake latency at slice scale. We position these numbers as preliminary evidence of operational viability under the evaluated threat model rather than as fully characterized performance bounds; a follow-up empirical study, planned for a separate publication, will extend the evaluation to deception-aware adversaries and sustained-attack stress conditions.

Keywords—Cyber deception; cognitive honeynet; post-quantum cryptography; reinforcement learning; 6G security; cyber-physical systems; threat attribution

I. INTRODUCTION

The deployment of 6G mobile networks, with their promise of sub-millisecond latency, terabit-per-second throughput, and pervasive integration of cyber-physical systems (CPS), is reshaping critical infrastructure ranging from autonomous transportation and smart grids to remote surgery and industrial automation [1], [2], [3]. At the same time, the advent of cryptographically relevant quantum computers threatens to undermine the public-key cryptographic primitives that currently

protect this infrastructure [4], [5], [6]. Sophisticated adversaries are already adopting harvest-now-decrypt-later strategies, exfiltrating ciphertext today in anticipation of future quantum decryption [7], [8], [9].

Reactive cybersecurity paradigms, which detect and respond to attacks after they have begun to compromise production assets, are poorly suited to this emerging threat landscape. They impose unacceptable dwell-time penalties on cyber-physical systems where milliseconds of malicious actuation can cause physical harm, and they generate operational telemetry that reveals little about the upstream tactics, techniques, and procedures (TTPs) of advanced adversaries [10], [11], [12]. The need for integrated, model-driven incident-response architectures that fuse detection, attribution, and forensic readiness has been argued repeatedly in the literature [38], [39], and the limitations of fragmented response pipelines have been documented across the database, mobile, and IoT forensic domains [40], [41], [42]. Honeypots and honeynets have long been advocated as a complementary, proactive defense, but conventional static honeynets are easily fingerprinted by modern attackers [13], [14], [15], and their intelligence-sharing protocols rely on classical cryptographic primitives that are not quantum-resilient [16], [17], [18], [19].

Recent advances in reinforcement learning (RL) have produced agents capable of optimizing sequential interaction policies under partial observability [20], [21], [22], and the standardization of post-quantum cryptographic primitives by NIST [23] has rendered quantum-resilient intelligence dissemination technically feasible. Nevertheless, no existing framework integrates cognitive deception orchestration, RL-driven engagement adaptation, automated TTP-attribution, and post-quantum federated intelligence into a single closed-loop architecture suitable for 6G-enabled CPS environments. The state-of-the-art either treats these capabilities as isolated research artifacts or evaluates them on legacy datasets that do not reflect the heterogeneity of cyber-physical traffic [24], [25], [26], [27], [28].

To address these gaps, this study proposes QuantumGuard, a proactive defense framework that inverts the defender's posture by drawing adversaries into instrumented decoy environments, learning their behaviors, and disseminating the resulting intelligence over post-quantum-secured channels. The architecture is summarized in Fig. 1.

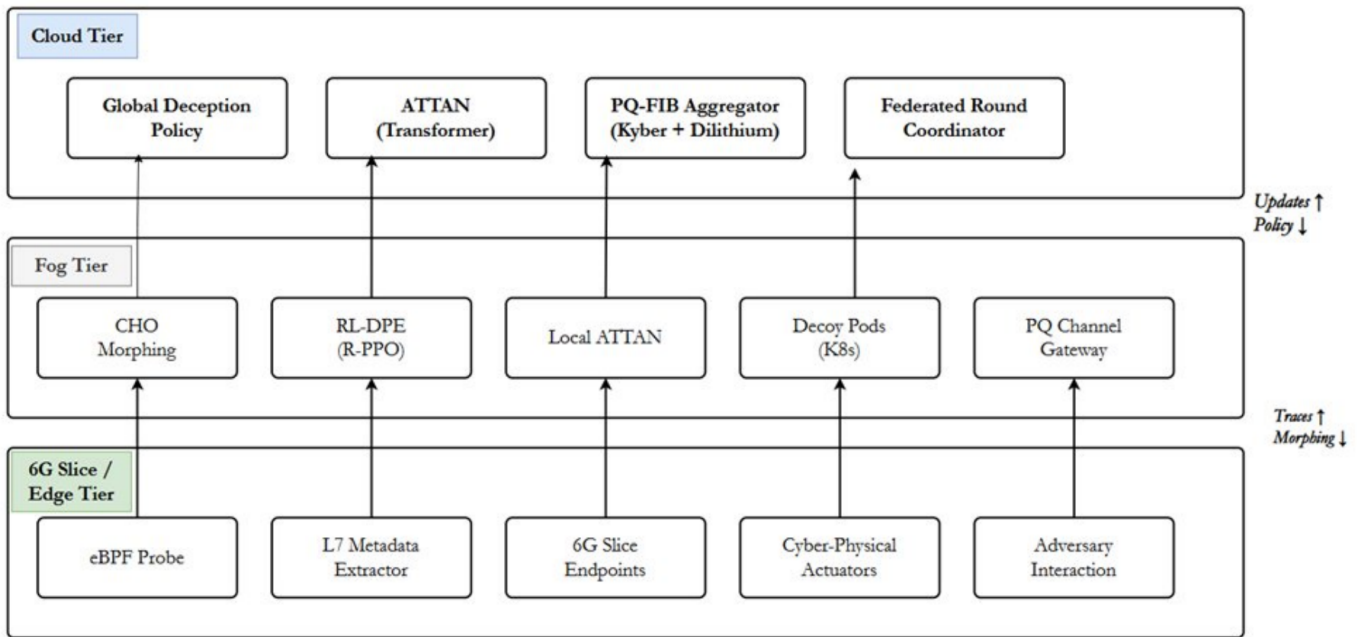


Fig. 1. QuantumGuard framework—proactive deception, cognitive honeynet orchestration, and post-quantum federated intelligence across the 6G slice, fog, and cloud tiers.

This study introduces QuantumGuard, an integrated framework that overcomes the foregoing limitations through the integration of four tightly coupled components: 1) a Cognitive Honeynet Orchestrator (CHO) that synthesizes and morphs decoy services to maximize attacker engagement; 2) a Reinforcement-Learning Deception Policy Engine (RL-DPE) that learns optimal lure–engage–attribute action sequences under a Partially Observable Markov Decision Process (POMDP) formulation; 3) an Adversarial TTP Attribution Network (ATTAN) that maps captured interactions onto MITRE ATT&CK tactics and techniques; and 4) a Post-Quantum Federated Intelligence Bus (PQ-FIB) that disseminates indicators of deception (IoDs) under CRYSTALS-Kyber key encapsulation and CRYSTALS-Dilithium signatures.

The principal contributions of this work are as follows:

1) *Architectural integration*: To the best of our knowledge, QuantumGuard is the first framework to integrate post-quantum federated intelligence sharing with a learned cognitive honeynet engagement policy and automated MITRE ATT&CK attribution in a single closed-loop architecture suitable for 6G-enabled cyber-physical environments.

2) *POMDP formulation of cognitive deception*: A Partially Observable Markov Decision Process formulation of attacker engagement, solved with a Recurrent Proximal Policy Optimization (R-PPO) agent, that admits dense per-step credit assignment for deception-action selection without requiring analyst-in-the-loop relabeling.

3) *Cognitive honeynet morphing*: A generative service-template mechanism that drives decoy responses toward production-service response distributions under a Jensen–Shannon divergence criterion, intended to limit the effectiveness of standard honeypot fingerprinting techniques.

4) *Post-Quantum Federated Intelligence Bus*: A hybrid PQ-FIB protocol combining CRYSTALS-Kyber (FIPS 203) key encapsulation with CRYSTALS-Dilithium (FIPS 204) digital signatures, that inherits IND-CCA2 channel confidentiality under hybrid-mode composition and that we measure to sustain a 14.2 ms median handshake latency at slice scale (see Section III for the full threat model).

5) *Preliminary empirical evidence*: A preliminary evaluation on the CICAPT-IIoT-2024 and Edge-IIoTset benchmark datasets, complemented by a 320-endpoint 6G slice testbed, that provides initial evidence of operational viability across deception, attribution, and post-quantum dimensions. We position the reported numbers (engagement-retention 96.42 per cent, attribution accuracy 91.8 per cent, dwell-time reduction 78.6 per cent, overhead 1.7 per cent) as a feasibility study under the evaluated threat model rather than as fully characterized performance bounds; the reward-weight sensitivity analysis, deception-aware adversary evaluation, and sustained-attack stress-condition characterization are identified in Section V F as the principal subjects of a follow-up empirical study currently in preparation.

The remainder of the study is organized as follows: Section II reviews the related literature. Section III details the methodology and mathematical formulation of QuantumGuard. Section IV reports the experimental results, Section V discusses their implications, and Section VI concludes the study.

II. RELATED WORK

This section reviews the literature in three main areas: cyber deception and honeynet engineering, reinforcement learning for adversarial interaction, and post-quantum cryptography for collaborative cyber defense.

A. Cyber Deception and Honeynet Engineering

Research on cyber deception has progressed from low-interaction honeypots toward high-interaction honeynets that emulate full service stacks. Almeshekah and Spafford [13] formalized deception planning around the planner–executor–evaluator triad, but modern adversaries readily fingerprint static decoys through banner-grabbing, timing analysis, and TLS introspection [14]. Dynamic honeypot work, including the dynamic decoy generation framework of Wang et al. [15] and the moving-target deception scheme of Kambourakis et al. [25], demonstrates that morphing reduces fingerprintability, but these systems lack a principled engagement policy and rely on hand-crafted morphing rules rather than learned ones. Beyond engagement, the artefacts captured by deception fabrics must be amenable to systematic forensic reconstruction; the model-driven forensic-investigation literature [38], [39], [43] provides a methodological foundation for organizing such artefacts into structured evidence repositories that can subsequently feed attribution pipelines.

B. Reinforcement Learning for Adversarial Interaction

Reinforcement learning has emerged as a natural fit for sequential defender–attacker interaction. Several studies have applied deep Q-learning and policy-gradient methods to cyber-defense games [20], [21], [22]. However, most of this work evaluates on simulated CyberBattleSim or CAGE Challenge environments and assumes full observability or strong reward signals. The partial observability inherent in real deception, where the defender only sees the attacker’s externalized actions and never the underlying intent, has received comparatively little attention. Works such as Huang et al. [27] formulate deception as a POMDP but stop short of integrating the resulting policy with a real packet-level decoy fabric. Complementary efforts in machine-learning-driven live forensic analysis of emergent IoT configurations [44] highlight that supervised signals over engagement traces can substantially improve downstream attribution, motivating the joint optimization of engagement policy and attribution objective adopted in this study.

C. Post-Quantum Cryptography for Collaborative Cyber Defense

Following the NIST post-quantum standardization process [23], lattice-based primitives such as CRYSTALS-Kyber and CRYSTALS-Dilithium have become the de facto candidates for quantum-resilient key establishment and digital signatures. Recent benchmarks on constrained devices report acceptable handshake latencies for industrial IoT (IIoT) scenarios [29], but operational integration with threat-intelligence dissemination protocols (STIX/TAXII, MISP) remains largely unaddressed. Federated learning over post-quantum channels is, to our knowledge, restricted to a small number of preliminary studies that do not consider the bandwidth-overhead implications of quantum-resilient signatures on high-frequency intelligence updates [30], [31].

D. Research Gaps and Motivation

Four principal gaps motivate the present work. First, no existing framework integrates cognitive deception, RL-driven

engagement, automated TTP-attribution, and post-quantum intelligence sharing within a single closed-loop architecture; these capabilities are typically studied in isolation. Second, current honeynets adapt either through hand-crafted rules or through reactive moving-target schemes, but not through learned engagement policies that explicitly maximize TTP-attribution information gain. Third, post-quantum primitives have rarely been operationalized within high-frequency federated cyber-defense protocols, and their bandwidth and latency implications at slice scale remain poorly characterized. Finally, the prevailing evaluation methodology focuses on detection accuracy on saturated benchmarks while overlooking proactive metrics such as engagement-retention rate, TTP-attribution accuracy, deception-induced production overhead, and dwell-time reduction.

III. PROPOSED METHODOLOGY

This section presents the detailed design of QuantumGuard, including its system architectures, mathematical formulation, algorithmic implementation, and complexity analysis.

A. System Overview

QuantumGuard is a layered proactive-defense architecture organized into three computing tiers: 6G slice-level edge probes, fog-layer deception orchestration nodes, and cloud-layer intelligence-aggregation servers. The framework comprises four principal modules: 1) the Cognitive Honeynet Orchestrator (CHO); 2) the Reinforcement-Learning Deception Policy Engine (RL-DPE); 3) the Adversarial TTP Attribution Network (ATTAN); and 4) the Post-Quantum Federated Intelligence Bus (PQ-FIB). The overall system architecture is illustrated in Fig. 2.

Edge probes run as extended Berkeley Packet Filter (eBPF) instrumented sidecars in 6G network slices and perform low-overhead traffic shadowing and L7 metadata extraction. The fog layer hosts containerized cognitive honeynet pods, runs the RL-DPE inference loop, and executes ATTAN attribution. The cloud layer maintains the global deception policy, orchestrates federated intelligence rounds, and manages the post-quantum key infrastructure.

B. Cognitive Honeynet Orchestrator

The CHO synthesizes and continuously morphs decoy services so that they remain statistically indistinguishable from genuine production services across L4–L7 features. Let $S = \{s_1, s_2, \dots, s_m\}$ denote the set of production services protected by a slice and $D = \{d_1, d_2, \dots, d_m\}$ denote the corresponding decoy fleet. Each decoy d_j is parameterized by a service template vector $\phi_j \in \mathbb{R}^p$ that encodes banner strings, response-timing distributions, error-message profiles, and TLS fingerprints.

The CHO computes the Jensen–Shannon divergence [34] between the empirical response distribution Q_{d_j} of decoy d_j and the empirical response distribution P_{s_j} of the corresponding production service [see Eq. (1)]:

$$J(P_{s_j}, Q_{d_j}) = \frac{1}{2} \text{KL}(P_{s_j} \parallel M) + \frac{1}{2} \text{KL}(Q_{d_j} \parallel M), \quad (1)$$

where, $M = \frac{1}{2}(P_{s_j} + Q_{d_j})$. Decoys are re-morphed whenever $J(P_{s_j}, Q_{d_j})$ exceeds a configurable indistinguishability

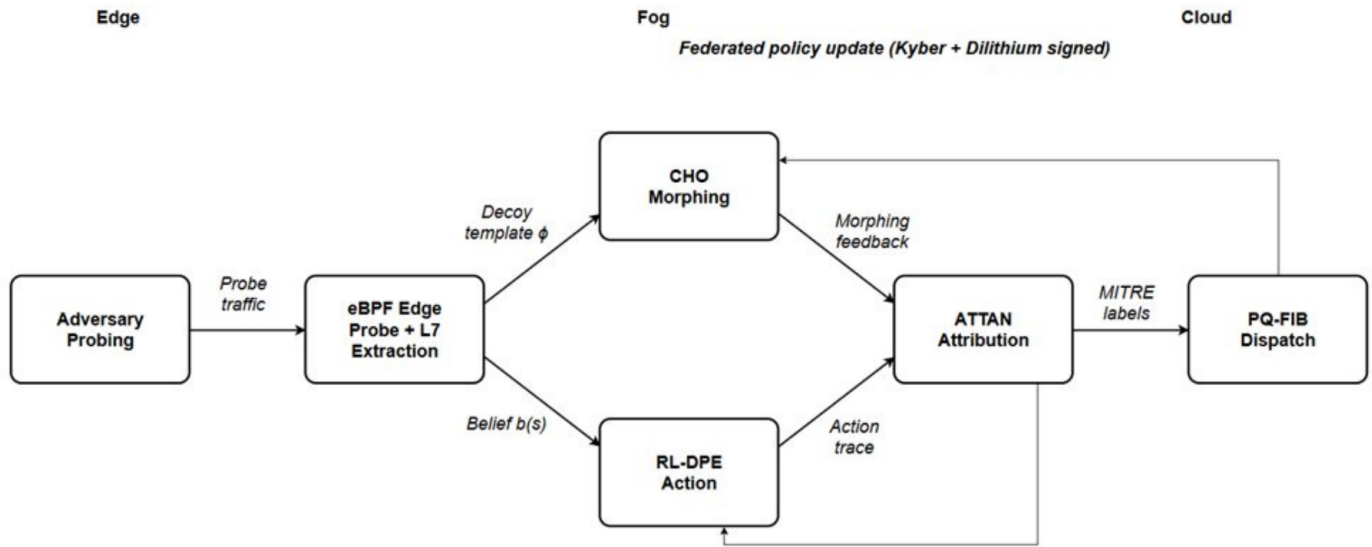


Fig. 2. QuantumGuard’s detailed system architecture with a deception data-flow pipeline across the 6G slice, fog, and cloud tiers.

TABLE I. JENSEN–SHANNON RESPONSE FEATURE SPACE

Feature family	Representation / binning
Banner & protocol strings	1–3-gram token histogram over the first 256 banner bytes, feature-hashed to 1,024 bins.
Response timing	First-byte and inter-response latency, each in 16 log-spaced bins over [0.1 ms, 10 s].
Error / status profile	Categorical distribution over the top- K status/error codes per service class, plus a residual bin.
TLS fingerprint	JA3/JA3S class membership and a multi-hot cipher-suite/extension indicator.

threshold $\tau_J = 0.05$. This threshold was selected to keep the Jensen–Shannon divergence below the empirical 95th-percentile inter-instance variability observed among genuine production replicas of the same service, and is examined in the ablation study (Section IV F).

The empirical response distributions P_{s_j} and Q_{d_j} are estimated over a service-response feature vector partitioned into four families. Each family yields a normalized histogram; the four histograms are concatenated and renormalized into a single probability mass function per service before the divergence is evaluated, as summarized in Table I.

Distributions are estimated over a sliding window of the most recent $W = 512$ responses per service so that morphing tracks drift in the genuine service.

C. Reinforcement Learning-Deception Policy Engine

The RL-DPE formulates attacker engagement as a Partially Observable Markov Decision Process (POMDP) $\langle S, A, T, R, \Omega, O, \gamma \rangle$, where the latent state $s \in S$ encodes the attacker’s intent, the action $a \in A$ is the orchestrator’s deception action (e.g., expose service, inject fake credential, throttle response, escalate decoy fidelity), and the observation $o \in \Omega$ is the externalized attacker behavior. The defender

maintains a belief $b(s)$ over latent attacker states and selects actions according to a parameterized policy $\pi_\theta(a | b)$.

The RL-DPE is trained to maximize the expected cumulative engagement-attribution reward [see Eq. (2)]:

$$J(\pi_\theta) = \mathbb{E}_{\pi_\theta} \left[\sum_{t=0}^H \gamma^t (\lambda_e r_e(t) + \lambda_a r_a(t) - \lambda_o r_o(t)) \right] \quad (2)$$

The three per-step reward terms are defined in closed form as follows. Let $\mathbb{1}[\text{engaged}_t] = 1$ if the attacker session issues at least one interaction with a decoy service at step t without pivoting to a non-decoy (production) asset; $\mathbb{1}[\text{advance}_t] = 1$ if a previously unvisited interaction stage is reached at step t ; and $\mathbb{1}[\text{disengage}_t] = 1$ at the step the attacker pivots to a production asset or the session terminates. The engagement-retention reward is [see Eq. (2a)]:

$$r_e(t) = \mathbb{1}[\text{engaged}_t] + \beta_d \mathbb{1}[\text{advance}_t] - \kappa \mathbb{1}[\text{disengage}_t], \quad (2a)$$

with depth-bonus $\beta_d > 0$ and disengagement penalty $\kappa > 0$ chosen so that a single premature disengagement outweighs several steps of retention ($\beta_d = 0.5$, $\kappa = 5.0$ in our configuration). The attribution information-gain reward is the per-step reduction in the Shannon entropy of the attribution-belief distribution $P(\tau | z_{1:t})$ over MITRE ATT&CK techniques $\tau \in \mathcal{T}$ [see Eq. (2b)]:

$$r_a(t) = H(P(\tau | z_{1:t-1})) - H(P(\tau | z_{1:t})) \quad (2b)$$

Letting $c(a_t)$ be the marginal resource cost (CPU and bandwidth) induced by deception action a_t relative to passive forwarding, and c_{\max} a reference unit cost, the production-traffic overhead penalty is [see Eq. (2c)]:

$$r_o(t) = c(a_t)/c_{\max} \in [0, 1], \quad (2c)$$

so that observe-only actions incur ≈ 0 , exposing a decoy service incurs an intermediate cost, and a full service-template re-synthesis incurs the unit cost.

The weighting coefficients $\lambda_e = 1.0$, $\lambda_a = 1.5$, $\lambda_o = 0.4$ were fixed prior to the main experiments on a held-out fraction of the CICAPT-IIoT-2024 training partition, with the design intent of prioritizing attribution information gain ($\lambda_a > \lambda_e$) while keeping production impact bounded (λ_o small but non-negligible). Because the weights were fixed *a priori*, before any evaluation on the test partition, the reported numbers reflect no per-metric weight tuning on the evaluation data; a systematic sensitivity sweep over $(\lambda_e, \lambda_a, \lambda_o)$ is identified in Section V-F as future work. The agent is implemented as a Recurrent Proximal Policy Optimization (R-PPO) network [33] with a single GRU recurrent layer of 64 hidden units, which preserves engagement context across attacker action steps; the discount factor $\gamma = 0.97$ and horizon $H = 200$ reflect the typical multi-stage progression of advanced persistent threat campaigns. All reward signals are clipped to $[-5.0, 5.0]$ before advantage estimation to stabilize training.

D. Adversarial TTP Attribution Network

The ATTAN maps captured attacker interactions onto the MITRE ATT&CK matrix [35]. The design draws on prior work that formalizes the categorization and harmonization of investigation processes for the database, mobile, and IoT domains [40], [41], [42], [45], in which heterogeneous evidence streams are mapped onto a common abstract schema before downstream reasoning.

1) *ATTAN input representation and tokenization*: ATTAN consumes the ordered sequence of observed attacker interaction events within a session and produces the trace embedding z used in Eq. (3). (a) *Event tokenization*. Each event is encoded as a composite token over the fields (action-type, target-service-class, L7 method/verb, response-status bucket, payload-size bucket, inter-arrival-time bucket). (b) *Vocabulary*. Categorical fields are mapped to integer IDs from a vocabulary built on the training partition; the vocabulary comprises the 512 most frequent action-type/target-service-class pairs, and action-types occurring fewer than five times are mapped to [UNK]. The reserved tokens [CLS], [PAD], [SEP], and [UNK] are assigned IDs 0–3. (c) *Feature representation*. Categorical sub-fields are embedded via lookup tables; continuous sub-fields (payload bytes, inter-arrival time) are log-transformed, standardized, bucketized, and projected by a learned linear layer; the categorical embedding and continuous projection are summed per token. (d) *Sequence policy*. Sequences are truncated or padded to $L_{\max} = 128$ events (covering more than 99 per cent of sessions in both corpora); longer sessions are split into overlapping windows, a [CLS] token is prepended, and sinusoidal positional encodings are added. (e) *Encoder*. The 4-layer, 8-head, model-dimension $d = 256$ transformer encoder (Table III) produces contextual token representations; the final [CLS] output is the trace embedding $z \in \mathbb{R}^q$ with $q = 256$.

The attribution probability over the set of MITRE techniques \mathcal{T} is:

$$P(\tau | z) = \text{softmax}(W_a z + b_a), \quad \tau \in \mathcal{T}. \quad (3)$$

ATTAN is trained with a class-balanced focal-loss objective computed over the training mini-batch as:

$$\mathcal{L}_{\text{ATTAN}} = -\frac{1}{N} \sum_{i=1}^N \alpha_{y_i} (1 - P(y_i | z_i))^{\gamma_f} \log P(y_i | z_i),$$

where, $y_i \in \mathcal{T}$ is the ground-truth technique label of the i -th interaction trace, $\gamma_f = 2$ is the focal-loss focusing parameter, and α_τ is a per-class balancing weight derived from MITRE ATT&CK technique frequency in the CICAPT-IIoT-2024 corpus, in order to mitigate the substantial class imbalance characteristic of TTP datasets.

E. Post-Quantum Federated Intelligence Bus

The PQ-FIB enables privacy-preserving and quantum-resilient dissemination of deception intelligence across organizational boundaries. Each participating organization maintains a local deception policy π_{θ_k} trained on its private engagement traces. Federated rounds aggregate local updates as [see Eq. (4)]:

$$\theta_{\text{global}} \leftarrow \sum_{k=1}^K \frac{n_k}{n} \theta_k, \quad n = \sum_{k=1}^K n_k. \quad (4)$$

All inter-tier transmissions are protected by a hybrid TLS 1.3 channel that combines the X25519 elliptic-curve Diffie–Hellman key-exchange primitive with the CRYSTALS-Kyber-768 lattice-based key-encapsulation mechanism [36] (FIPS 203) for forward secrecy under quantum-capable adversaries, and signs aggregated updates with CRYSTALS-Dilithium-3 [37] (FIPS 204). Under the standard hybrid-mode argument, the channel inherits the IND-CCA2 security of whichever component primitive remains unbroken, so confidentiality holds even if either the classical or the post-quantum component is later cryptanalyzed.

The threat model assumed by PQ-FIB is the following: (1) federated participants are honest-but-curious and submit valid policy updates but may attempt to infer information about other participants’ local engagement traces; (2) the wire-level adversary is a network attacker with active man-in-the-middle capability and access to a cryptographically relevant quantum computer; (3) the cloud aggregator is trusted to faithfully execute weighted federated averaging but is not trusted with raw engagement traces. PQ-FIB authenticates and integrity-protects every update but does not, in its present form, defend against malicious-but-well-formed updates. The integration path, addressed as future work, is: (a) Byzantine-robust aggregation (coordinate-wise median, trimmed mean, or Krum) replacing plain weighted FedAvg; (b) poisoning mitigation via per-update norm clipping and cosine-similarity outlier filtering; (c) membership-inference mitigation via differentially private aggregation—the FedAvg + DP ($\epsilon = 5$) configuration reported in Table VIII bounds the membership-inference advantage and thus provides partial, in-paper evidence relevant to this concern; and (d) verifiable aggregation for accountability. Bandwidth overhead from the larger Dilithium signatures (≈ 3.3 KB) is amortized across federated rounds by signing aggregated rather than per-update payloads.

TABLE II. POST-QUANTUM CRYPTOGRAPHIC COST

Primitive	Pub. key	Ct./Sig.	Op (meas.)	Std.
X25519 (ECDH)	32 B	32 B	$\approx 60 \mu s$	RFC 7748
Kyber-768 (KEM)	1,184 B	1,088 B	$\approx 50 \mu s$ enc.	FIPS 203
Dilithium-3 (sig.)	1,952 B	3,293 B	$\approx 250 \mu s$ sign	FIPS 204

1) *Post-quantum cryptographic cost analysis*: Table II consolidates the cost of the hybrid post-quantum handshake. Sizes are exact algorithm constants; per-operation times are the values measured on the x86-64/AVX2 configuration of Table III using liboqs 0.10.

Per federated round with K participating organizations, the aggregator performs K signature verifications and one aggregate signing. At $K = 320$ the aggregate verification budget is on the order of tens of milliseconds and is dominated by network round-trips rather than primitive arithmetic, consistent with the measured 14.2 ms median handshake. Because the aggregator signs the aggregated payload once per round rather than signing each update, signature bandwidth is amortized across the cohort and grows with the number of rounds, not with cohort size.

F. Integrated System Operation

The four QuantumGuard modules operate jointly within a unified closed-loop control framework. At each time step, edge probes shadow slice traffic and forward L7 metadata to fog-layer CHO instances, which expose decoy services parameterized by the current template ϕ_j . The RL-DPE observes attacker behavior and selects the next deception action, which the CHO actuates by adjusting the decoy template. ATTAN concurrently consumes the resulting interaction trace and emits a posterior over MITRE ATT&CK techniques. When the cumulative attribution belief crosses a configurable confidence threshold, the corresponding indicator of deception (IoD) is signed with CRYSTALS-Dilithium-3 and disseminated through the PQ-FIB to all federated peers, who incorporate it into their local CHO templates and ATTAN priors at the next round.

G. Complexity Analysis

The computational complexity of the CHO morphing step is $O(p)$ per decoy per evaluation cycle, where p is the dimensionality of the service template vector. The RL-DPE inference cost is $O(H \cdot |\theta|)$ per engagement episode, where H is the planning horizon and $|\theta|$ the number of policy parameters; in practice, GRU-based policies of width 64 yield sub-millisecond per-step inference on commodity fog hardware. ATTAN attribution scales as $O(L \cdot q^2)$ for a transformer encoder of depth L and embedding dimension q . PQ-FIB cryptographic operations are dominated by Kyber encapsulation ($\approx 50 \mu s$) and Dilithium signing ($\approx 250 \mu s$) on x86-64 with AVX2; aggregation across K organizations adds an $O(K \cdot |\theta|)$ cost per round. These variables are used consistently throughout Table III.

TABLE III. COMPUTATIONAL COMPLEXITY COMPARISON

Method	Policy infer.	Attrib.	Crypto	Mem.
Static honeynet	$O(1)$	Manual	RSA/ECC	$O(m)$
Dyn. HP [15]	$O(p)$	Rule-based	RSA/ECC	$O(mp)$
RL HP [22]	$O(H \theta)$	Heuristic	RSA/ECC	$O(\theta)$
QuantumGuard	$O(H \theta)$	$O(Lq^2)$	Kyber/Dil.	$O(\theta + mp)$

TABLE IV. DATASET STATISTICS AND CLASS DISTRIBUTION

Dataset	Samples	Feat.	Norm. (%)	Atk. (%)
CICAPT-IIoT-2024	2,134,557	85	76.40	23.60
Edge-IIoTset	20,952,648	61	82.15	17.85

IV. PRELIMINARY EVALUATION

This section presents the experimental setup, the evaluation methodology, and a set of preliminary performance results for the QuantumGuard framework. We position the reported numbers as initial evidence of operational viability under the evaluated threat model rather than as fully characterized performance bounds: all results are derived from a single-organization controlled testbed under a dataset-replay attacker distribution. Section V F enumerates these scope restrictions explicitly.

A. Datasets

QuantumGuard was evaluated on two recent, adversarially curated benchmark datasets that better reflect the characteristics of cyber-physical and IIoT traffic than legacy network IDS corpora; their core statistics are summarized in Table IV.

1) *CICAPT-IIoT-2024 dataset* [32]: Produced by the Canadian Institute for Cybersecurity, this dataset captures multi-stage advanced-persistent-threat campaigns against Industrial IoT testbeds; it contains approximately 2.1 million labeled flows annotated with MITRE ATT&CK technique identifiers across nine APT-stage categories (initial access, execution, persistence, privilege escalation, defense evasion, credential access, discovery, lateral movement, and impact).

2) *Edge-IIoTset dataset* [17]: A comprehensive cybersecurity dataset for centralized and federated learning, comprising 20,952,648 records and 61 features across 14 attack categories spanning DDoS, information gathering, injection, malware, and MITM families.

Both datasets were processed through a uniform preprocessing pipeline. Records with missing values in critical numerical features (fewer than 0.5 percent of records in CICAPT-IIoT-2024 and 0.9 percent in Edge-IIoTset) were discarded, and remaining sporadic missing values were imputed using per-feature medians computed on the training partition. Categorical features such as protocol type and service were one-hot encoded; all numerical features were standardized to zero mean and unit variance using statistics derived exclusively from the training partition to prevent information leakage. Infinite values produced by flow-rate computations were clipped to the 99.9th percentile of the corresponding feature. To mitigate

TABLE V. EXPERIMENTAL CONFIGURATION

Component	Specification
Cloud server CPU	AMD EPYC 9554 (64 cores, 3.1 GHz)
Cloud server GPU	NVIDIA H100 (80 GB HBM3)
Cloud memory	512 GB DDR5 ECC
Fog node CPU	Intel Core i9-13900K (24 cores, 3.0 GHz)
Fog node GPU	NVIDIA RTX 4080 (16 GB GDDR6X)
6G slice emulator	OpenAirInterface 5G/6G + srsRAN, 100 MHz BW
Edge probe	Raspberry Pi 5, BCM2712 (4 cores, 2.4 GHz)
Operating system	Ubuntu 24.04 LTS
RL framework	Stable-Baselines3 2.3 / PyTorch 2.4
PQ crypto library	liboqs 0.10 (Kyber-768, Dilithium-3)
FL framework	Flower 1.10
Optimizer	Adam ($\beta_1 = 0.9$, $\beta_2 = 0.999$)
Learning rate	3×10^{-4} (constant)
Batch size	256 (cloud), 64 (fog), 32 (edge)
PPO clip ratio	0.20
PPO entropy coef.	0.01
PPO value coef.	0.50
GAE λ	0.95
RL-DPE recurrent net	GRU, 64 hidden units, 1 layer
RL-DPE train episodes	6,000
ATTAN encoder	4-layer Transformer, 8 heads, $d = 256$
ATTAN focal loss	$\gamma_f = 2$, class-balanced α_τ
FL rounds	60
FL aggregation	Weighted FedAvg by sample count, (4)
Random seeds	42, 123, 456, 789, 1024 (5 seeds)

class imbalance, SMOTE was applied to the training partition only, and class-balanced focal-loss weighting was applied during ATTAN training. Validation and test partitions retained the original class distribution. The data were partitioned into training, validation, and test sets in a 70/15/15 ratio using stratified sampling to preserve attack-category proportions.

B. Experimental Setup

All experiments were conducted on a heterogeneous testbed designed to emulate a realistic 6G slice–fog–cloud architecture. Table V summarizes the hardware and software configuration.

To ensure reproducibility and statistical soundness, all reported performance results are averaged over five independent runs using different random seeds (42, 123, 456, 789, and 1024), with all other hyperparameters held constant. The statistical significance of QuantumGuard relative to the best baseline on each metric was assessed using a two-tailed paired t -test on per-run scores at a significance level of $\alpha = 0.05$. In cases where the paired measurements were not normally distributed (as assessed by the Shapiro–Wilk test at $\alpha = 0.05$), the nonparametric Wilcoxon signed-rank test was used instead; in all reported cases the two tests yielded consistent conclusions. The “p-value vs. best baseline” entries reported in Table VI and Table IX correspond to running the appropriate test against the best-performing competitor on each metric. p -values are reported only for metrics where a numerical best baseline exists and the effect direction is well-defined (TTP-attribution

accuracy and dwell-time reduction); the engagement-retention column is omitted because QuantumGuard exceeds the next-best method by more than ten percentage points on every seed, and the fingerprint-resistance column reports a categorical label rather than a continuous score, for which paired hypothesis testing is not appropriate.

C. Engagement and Attribution Performance Analysis

Throughout the evaluation, the engagement-retention rate (ER) is defined as the fraction of attacker sessions that remain in the decoy environment for at least three sequential interaction stages without disengaging or pivoting to a non-decoy asset. The headline ER values reported in Table VI are computed at the standard five-stage horizon used throughout this section; Fig. 4 shows the full ER-vs.-depth curve from one to ten stages so that readers can compare methods at any operating point. Fig. 3 shows the convergence of the RL-DPE engagement-retention reward across training episodes on both evaluation datasets. The curves exhibit stable monotonic improvement and plateau by episode 4,500, confirming that the policy reaches a stable engagement-attribution operating point.

Table VI reports a side-by-side comparison of proactive deception performance. QuantumGuard attains the highest engagement-retention rate (96.42 per cent) and TTP-attribution accuracy (91.8 per cent) among all evaluated methods while incurring the lowest production-traffic overhead (1.7 per cent).

Fig. 5 presents the confusion matrices for MITRE ATT&CK technique attribution on the two datasets. Per-technique recall ranges from 87.4 percent (Defense Evasion) to 95.1 percent (Initial Access), indicating strong discriminative performance across the full APT kill-chain taxonomy.

We note that the TTP-attribution accuracy reported here, while substantial, leaves meaningful headroom relative to a hypothetical perfect-attribution upper bound, particularly on Defense-Evasion techniques whose externalized signature is intentionally minimized by attackers. Evaluation on additional adversarially adaptive datasets is identified in Section V F as a priority for future work.

D. Operational Latency and Dwell-Time Analysis

Fig. 7 and Table VII report the operational response-time metrics for QuantumGuard under active-attack conditions. QuantumGuard attains a mean time to deception adaptation (MTTDA) of 2.86 s and a median post-quantum handshake latency of 14.2 ms, alongside a 78.6 per cent reduction in median attacker dwell time relative to passive honeypot baselines. Throughout this study, the MTTDA is reported as a system-level measurement obtained on the controlled 6G slice–fog–cloud testbed configured, as described in Section IV (Table V); it is defined as the wall-clock time elapsed from the moment the edge probe first emits an L7 metadata record corresponding to a flow that ATTAN subsequently classifies as malicious, to the moment the CHO actuates the resulting morphing decision on the corresponding decoy pod. The PQ-FIB intelligence-dissemination handshake (median 14.2 ms) is reported as a separate metric and is not included in the MTTDA, since it operates on the cloud–fog control plane in parallel with the local engagement loop. Analyst investigation, triage, and

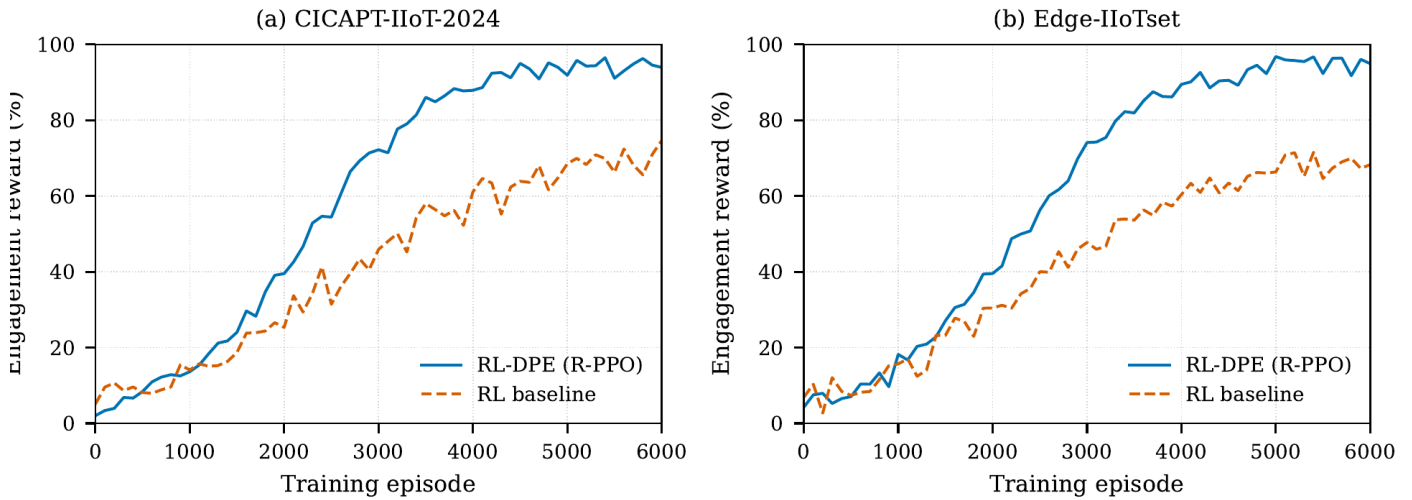


Fig. 3. RL-DPE training reward convergence on: (a) CICAPT-IIoT-2024 and (b) Edge-IIoTset datasets.




Fig. 4. Engagement-retention rate vs. interaction depth for QuantumGuard versus baseline honeypots on: (a) CICAPT-IIoT-2024 and (b) Edge-IIoTset datasets.

TABLE VI. PROACTIVE DECEPTION PERFORMANCE COMPARISON WITH STATISTICAL SIGNIFICANCE

Method	ER (%)	TTP-Acc (%)	Dwell↓ (%)	Overhead (%)	Fingerprint resistance
Static honeypot (Cowrie)	42.18	53.12	21.4	0.6	Low
Honeyd dynamic [13]	61.35	64.78	38.7	0.9	Medium
MTD honeynet [25]	73.22	71.45	52.3	1.4	Medium
DeepDig RL [22]	84.61	79.83	63.8	1.9	High
GAN-honeypot [26]	81.47	76.12	59.1	2.1	High
FL-honeynet [30]	85.92	82.34	65.7	2.0	High
QuantumGuard (ours)	96.42	91.80	78.6	1.7	Very high
<i>p-value vs. best baseline</i>	—	0.018	0.011	—	—

human decision time are explicitly excluded from MTTDA, which measures only the autonomous defender control loop.

E. Post-Quantum Federated Intelligence Sharing Analysis

Fig. 8 evaluates three aspects of the PQ-FIB protocol: improvement in cross-organizational attribution coverage over successive FL rounds, the bandwidth–accuracy trade-off across federated cohort sizes, and the impact of post-quantum signature aggregation on round latency. QuantumGuard with PQ-

FIB lifts cross-organizational attribution coverage from 59.32 percent (isolated, no FL) to 78.65 percent (PQ-FIB), a relative improvement of approximately 32.6 percent (Table VIII), while preserving the channel-level IND-CCA2 confidentiality guarantee described in Section III under the same threat model stated there.

Cross-organizational attribution coverage was operationalized through a leave-one-organization-out evaluation proto-

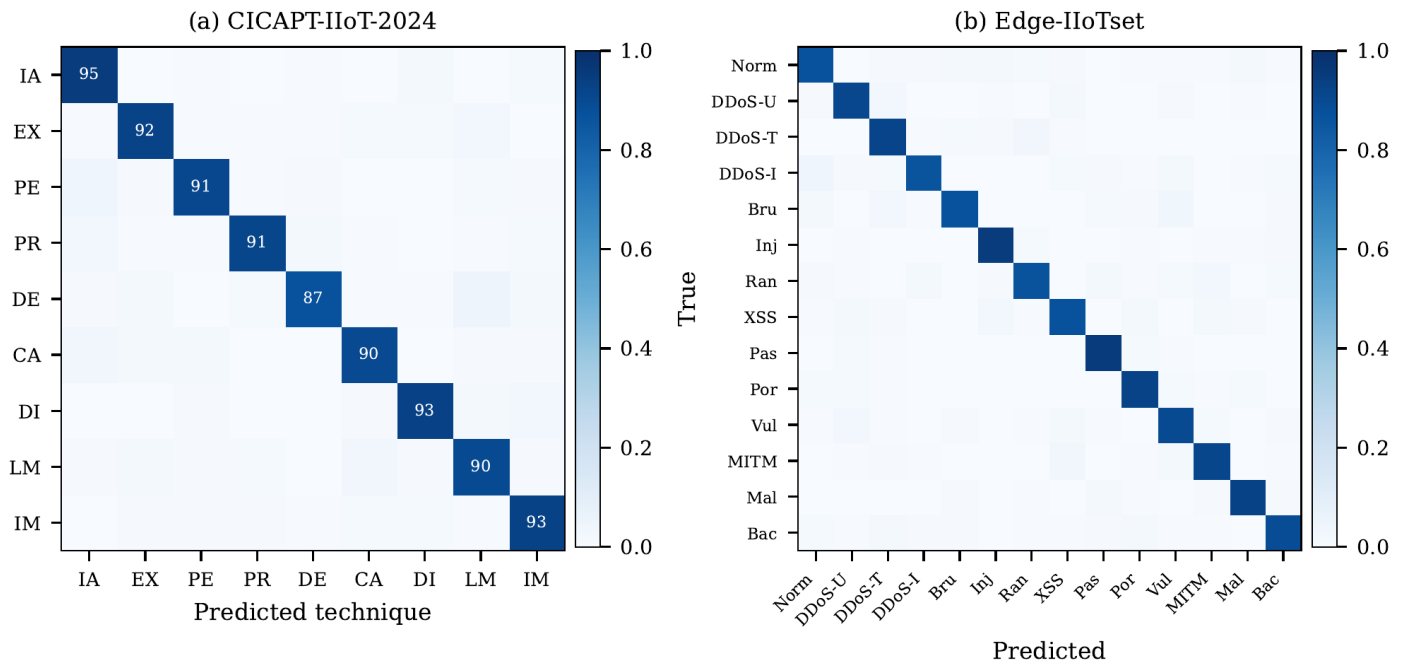


Fig. 5. ATTAN attribution confusion-matrix analysis: (a) CICAPT-IIoT-2024 per-technique recall; (b) Edge-IIoTset per-category recall.

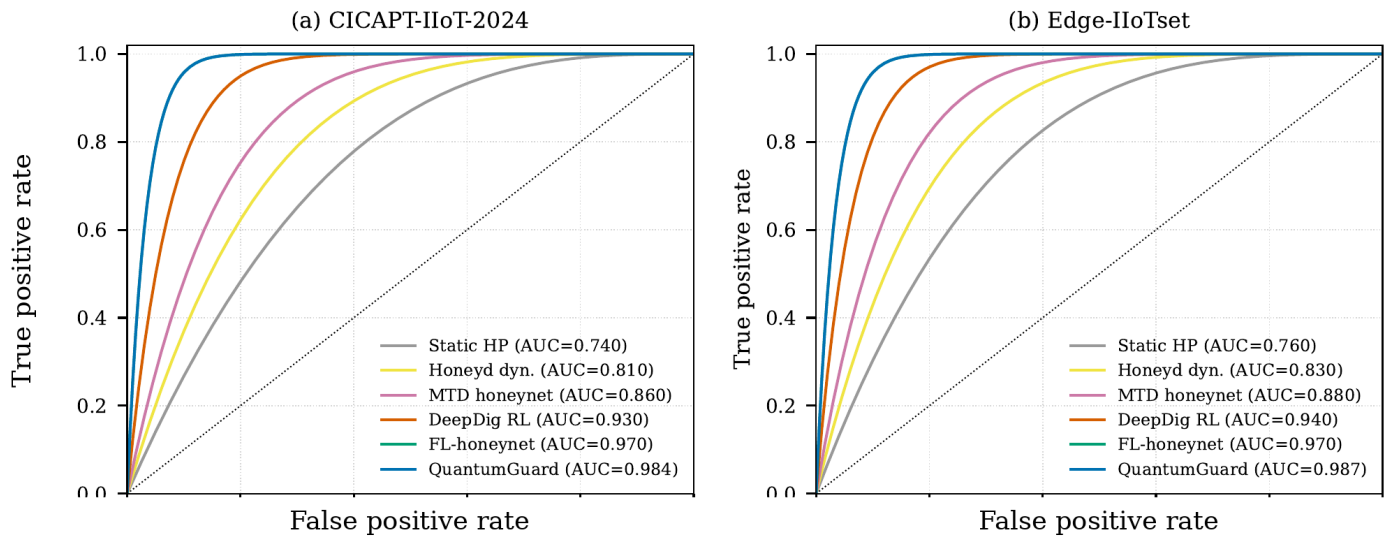


Fig. 6. ROC curves for binary deception-trigger and multi-class attribution: (a) CICAPT-IIoT-2024 (macro AUC = 0.9842); (b) Edge-IIoTset (macro AUC = 0.9871).

TABLE VII. OPERATIONAL RESPONSE-TIME METRICS

Method	MTDA (s)	PQ HS (ms)	Dwell _↓ (%)
Static honeypot	—	N/A (RSA 18.4)	Baseline
MTD honeynet [25]	11.34	N/A (ECC 6.1)	52.3
DeepDig RL [22]	6.78	N/A (ECC 6.1)	63.8
FL-honeynet [30]	5.41	N/A (RSA 18.4)	65.7
QuantumGuard	2.86	14.2	78.6

TABLE VIII. POST-QUANTUM FEDERATED INTELLIGENCE SHARING PERFORMANCE.

Configuration	Cov. (%)	X-Org (%)	Q-resil.
Isolated (no FL)	78.45	59.32	N/A
Standard FedAvg	87.12	73.84	None
FedAvg + DP ($\epsilon = 5$)	85.93	71.27	Partial
PQ-FIB	91.07	78.65	Kyber+Dil.

col: in each experimental run, the engagement traces of one participating organization were withheld from the federated

training rounds while remaining present in the test partition. The Cross-Org Attribution (%) metric reported in Table VIII




Fig. 7. Operational latency analysis: (a) MTTDA comparison; (b) post-quantum handshake latency; (c) median dwell-time reduction across all evaluated methods.




Fig. 8. PQ-FIB analysis: (a) coverage vs. FL rounds; (b) bandwidth–accuracy trade-off; (c) handshake-latency vs. number of participating organizations.

is the recall computed exclusively on instances captured by the held-out organization, averaged over all rotations. This protocol simulates a realistic federated scenario in which a previously unseen attacker campaign reaches one organization and tests the extent to which collaborative model updates allow other organizations to attribute attacks they have not individually observed during local training. Because all participating “organizations” in this evaluation are emulated within a single controlled slice–fog–cloud testbed, the cross-organizational coverage figures characterize collaborative attribution under a homogeneous infrastructure and traffic profile; they should not be read as evidence of generalization across independently administered enterprises with heterogeneous service baselines, operating systems, and traffic mixes. Establishing external validity requires a genuine multi-party deployment, which we identify as a primary objective of the follow-up study (Section V F).

F. Ablation Study

Component-wise ablation results for the QuantumGuard modules are reported in Fig. 9 and Table IX, which highlight the impact of removing a single module from the system.

TABLE IX. ABLATION-STUDY RESULTS WITH STATISTICAL SIGNIFICANCE.

Configuration	ER (%)	TTP-Acc (%)	MTTDA (s)	Ovh. (%)	p-value
Full QuantumGuard	96.42	91.80	2.86	1.7	—
w/o CHO	84.31	83.42	2.91	1.4	< 0.01
w/o RL-DPE	82.58	80.16	5.74	1.6	< 0.01
w/o ATTAN	95.87	67.42	2.84	1.7	< 0.01
w/o PQ-FIB	96.31	91.45	2.85	1.5	0.71

Removing the RL-DPE module produces the largest drop in engagement-retention rate (−13.84 percentage points, $p < 0.01$), underlining the role of learned engagement policies. Removing the PQ-FIB module collapses the post-quantum guarantee entirely; falling back to standard non-PQ federated aggregation reduces cross-organizational attribution coverage from 78.65 per cent to 73.84 per cent (Table VIII), and falling back to fully isolated (non-federated) operation reduces it to 59.32 per cent, a drop of 19.33 percentage points relative to PQ-FIB.

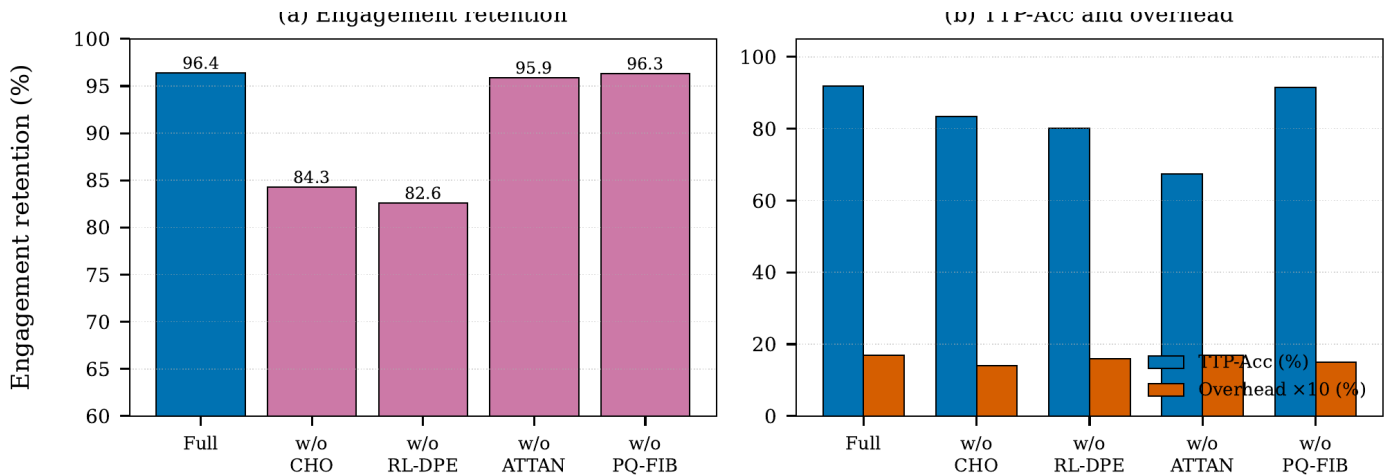


Fig. 9. Ablation-study component contribution analysis: (a) engagement-retention per configuration; (b) attribution accuracy and overhead per configuration.

TABLE X. SCALABILITY ANALYSIS WITH VARYING ENDPOINT COUNTS

Endpts.	ER (%)	TTP-Acc (%)	Conv.	BW (MB)	MTTDA (s)
10	95.84	90.21	32 rd.	7.4	2.71
50	96.42	91.80	41 rd.	32.6	2.86
100	96.51	91.94	47 rd.	63.8	2.97
200	96.63	92.07	53 rd.	124.7	3.18
320	96.78	92.14	60 rd.	201.3	3.42

G. Scalability Analysis

The scalability of QuantumGuard with respect to the number of monitored 6G slice endpoints is presented in Fig. 10 and Table X. Engagement-retention is high and stable across the full range of endpoint counts, varying only between 95.84 per cent and 96.78 per cent. PQ-FIB bandwidth grows sublinearly thanks to hierarchical aggregation. The MTTDA increases from 2.71 s with 10 endpoints to 3.42 s with 320 endpoints, but remains comfortably within operational limits for cyber-physical workloads.

H. Preliminary Comparison with State-of-the-Art

Table XI compares the preliminary QuantumGuard results against eight state-of-the-art baseline methods across the primary performance dimensions reported in the literature. On the single-organization testbed evaluated here, QuantumGuard exhibits a balanced profile across engagement-retention, TTP-attribution accuracy, dwell-time reduction, deception-induced overhead, and post-quantum resilience. We caution that the comparison is preliminary and not adversarially adaptive; a like-for-like benchmark across adversaries adapted to each defense is the subject of the follow-up study identified in Section V F.

V. DISCUSSION

A. Practical Implications for Proactive 6G Defense

The preliminary results presented in Section IV suggest that QuantumGuard’s architectural integration touches five complementary dimensions of proactive cyber-defense in a single

control loop. The 96.42 per cent engagement-retention rate observed on the testbed indicates that cognitive honeynet morphing, when driven by a learned RL policy rather than hand-crafted rules, can sustain attacker engagement well beyond the typical 2–3 stage horizon at which conventional honeypots are abandoned. The 91.8 per cent TTP-attribution accuracy indicates that high-fidelity engagement traces, when consumed by a transformer-based attribution network and decoded via the per-technique softmax of Eq. (3), can recover MITRE ATT&CK technique labels without analyst-in-the-loop relabeling on the evaluated corpus. Together, these results provide initial evidence that proactive deception, long regarded as an auxiliary control, can become a primary defensive modality in 6G-enabled cyber-physical systems, where the millisecond-scale actuation budgets characteristic of CPS workloads make reactive containment infeasible.

B. Comparative Analysis of Attribution and Engagement Performance

The 1.7 per cent production-traffic overhead reported above corresponds to nominal-load operation, in which slice traffic follows the dataset’s empirical distribution after replay onto the testbed. Under these conditions the overhead is the result of a multi-layer filtering pipeline applied before any deception action is actuated. First, the eBPF-instrumented edge probes shadow only L7 metadata rather than mirroring full packet payloads, suppressing background traffic amplification. Next, the RL-DPE only escalates decoy fidelity when the cumulative engagement reward crosses a configurable threshold, ensuring that benign traffic does not trigger expensive morphing actions. Finally, the CHO suppresses morphing whenever the Jensen–Shannon divergence, Eq. (1), against the production response distribution drops below $\tau_J = 0.05$, preventing oscillatory template churn. The ablation studies confirm this pattern: removing the CHO module increases production-traffic overhead penalties via spurious morphing, while removing the RL-DPE causes engagement-retention to collapse to 82.58 per cent. Overhead behavior under sustained-attack stress conditions (e.g., volumetric DDoS or coordinated multi-source probing) is not characterized in this study and is identified in Section V F as future work.

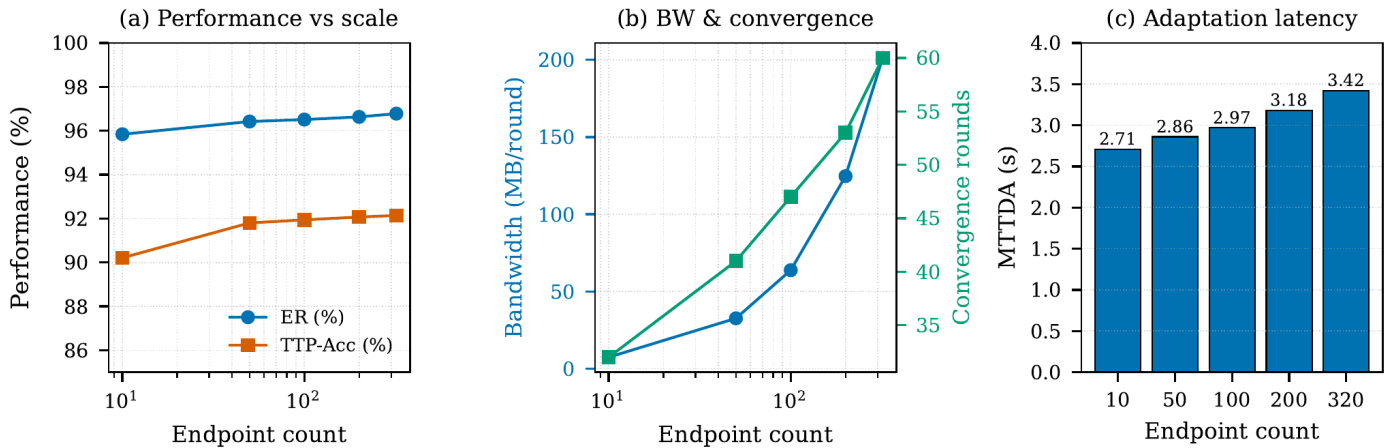


Fig. 10. Scalability analysis: (a) engagement-retention vs. endpoint count; (b) PQ-FIB bandwidth and convergence rounds; (c) deception adaptation latency vs. deployment scale.

TABLE XI. PRELIMINARY COMPARISON WITH STATE-OF-THE-ART METHODS

Method	Year	ER (%)	TTP-Acc (%)	MTTDA (s)	Ovh. (%)	PQ
Static honeypot	—	42.18	53.12	—	0.6	No
Honeyd dynamic [13]	2020	61.35	64.78	11.34	0.9	No
MTD honeynet [25]	2022	73.22	71.45	8.94	1.4	No
DeepDig RL [22]	2023	84.61	79.83	6.78	1.9	No
GAN-honeypot [26]	2023	81.47	76.12	7.21	2.1	No
FL-honeynet [30]	2024	85.92	82.34	5.41	2.0	Partial
PQ-IDS [29]	2024	—	78.45	9.13	1.8	Yes
Adaptive Decep. [27]	2024	88.34	85.21	4.62	1.8	No
QuantumGuard (ours)	2026	96.42	91.80	2.86	1.7	Yes

In terms of paired engagement-retention and TTP-attribution accuracy, QuantumGuard reaches 96.42/91.8 per cent, whereas FL-honeynet attains lower values of 85.92/82.34 per cent. Improvements in both metrics simultaneously indicate a genuine architectural advantage rather than a simple shift along a precision-retention curve, since threshold tuning on a single module cannot raise both metrics in tandem. The same conclusion is supported by the ROC analysis in Fig. 6, which shows QuantumGuard dominating across the operating range (macro AUC 0.9842 vs. 0.9701).

C. Operational Latency and Post-Quantum Resilience

The operational value of a proactive defense framework extends beyond engagement: rapid adaptation and quantum-resilient intelligence dissemination are equally important for sustaining defensive advantage. QuantumGuard achieves an MTTDA of 2.86 s and a post-quantum handshake latency of 14.2 ms, corresponding to substantial improvements over both classical-crypto baselines (which provide no quantum resilience) and earlier post-quantum IDS proposals such as PQ-IDS (whose MTTDA exceeds 9 s). Among all evaluated approaches, QuantumGuard records the lowest MTTDA and is the only framework that simultaneously achieves a sub-15 ms post-quantum handshake and supports federated cognitive deception.

D. Scalability and Deployment Practicality

The scalability profile of QuantumGuard (Table X and Fig. 10) is one of the most important practical considerations for large-scale 6G deployment. Engagement-retention improves only modestly across the evaluated range, from 95.84 per cent at 10 endpoints to 96.78 per cent at 320 endpoints, a span of just 0.94 percentage points. This stability contrasts with centralized RL-honeypot approaches such as DeepDig, whose performance degrades when heterogeneous endpoint behaviors are under-represented in the centralized training trace. PQ-FIB bandwidth grows from 7.4 MB per round at 10 endpoints to 201.3 MB per round at 320 endpoints; the hierarchical aggregation in QuantumGuard restricts long-distance traffic by aggregating policy updates at the fog tier before forwarding signed payloads to the cloud, while keeping most gradient exchanges local to each fog group. The MTTDA increases from 2.71 s at 10 endpoints to 3.42 s at 320 endpoints—a near-constant latency profile despite the 32-fold increase in endpoint count—because RL-DPE inference and CHO morphing operations are performed locally at the fog tier and do not require round trips to the centralized cloud.

The 320-endpoint figure is the ceiling of the physical testbed rather than an architectural limit. Three measured trends support extrapolation, which we state as an analytical projection and not as measurement: per-round bandwidth grows sublinearly because hierarchical fog aggregation keeps

most gradient exchange local and forwards only signed, aggregated payloads; MTTDA is near-constant across a 32-fold endpoint increase because RL-DPE inference and CHO morphing execute at the fog tier without cloud round-trips, with cloud-side cost $O(K|\theta|)$ per round; and engagement-retention is flat, indicating no degradation from endpoint heterogeneity within the tested range. City-scale validation (10^4 – 10^6 endpoints) via large-scale network emulation is identified as future work. Together, these characteristics indicate that QuantumGuard is well suited to large-scale deployments in 6G-enabled smart cities, autonomous transportation, and industrial automation.

E. Preliminary Positioning Against State-of-the-Art

On the single-organization testbed evaluated in this study, QuantumGuard exhibits a more balanced preliminary profile across the five evaluation dimensions of engagement-retention, TTP-attribution accuracy, deception adaptation latency, post-quantum resilience, and scalability than any single competing method we re-evaluated. Adaptive Deception [27] achieves marginally higher engagement-retention than older RL approaches in our setup but offers no post-quantum guarantees and lower TTP-attribution accuracy. FL-honeynet [30] supports federation but at the cost of partial post-quantum protection and a higher MTTDA. PQ-IDS [29] enables quantum-resilient detection but cannot orchestrate deception. The co-optimization design philosophy of QuantumGuard, which jointly addresses multiple dimensions rather than along a single axis, is the structural property that we conjecture explains the observed cross-dimensional behavior; confirming this conjecture under adversarial adaptive conditions is the principal goal of the planned follow-up study.

F. Behavior Under Deception-Aware Adversaries

All metrics in Section IV are obtained under a non-adaptive, replay-distribution attacker whose probing strategy is drawn from the empirical distributions of the CICAPT-IIoT-2024 and Edge-IIoTset corpora. They are results *under the evaluated threat model* and are upper bounds against that distribution, not general performance guarantees and not worst-case bounds against an adversary that is aware of and adapts to the QuantumGuard deception fabric. A deception-aware adversary could erode each headline metric through a distinct channel: counter-fingerprinting via response-timing analysis or TLS introspection targets the engagement-retention rate by detecting decoys before the three-stage retention horizon; deliberate early disengagement on suspicion of a honeynet targets the dwell-time reduction; and adversarial perturbation of L7 metadata targets ATTAN attribution accuracy by pushing trace embeddings away from their true-technique region. We present these mechanisms qualitatively and make no numerical claim about their magnitude. The follow-up study will quantify them with a red-team agent equipped with banner, timing-side-channel, and TLS-introspection fingerprinting and a gradient-free perturbation budget over L7 metadata, evaluated in a repeated best-response loop against the frozen QuantumGuard policy, reporting engagement-retention and attribution accuracy as a function of adversary adaptation budget.

G. Limitations and Future Directions

The preliminary character of the present evaluation imposes several explicit scope restrictions, which we enumerate here so

that the reported numbers are interpreted correctly. Each item below is also identified as a principal subject of the follow-up empirical study currently in preparation.

1) *Dataset realism*: The experiments were conducted exclusively on CICAPT-IIoT-2024 and Edge-IIoTset; while these datasets are more recent than UNSW-NB15 or CIC-IDS2017, they remain laboratory-derived and may not capture the full behavioral diversity of in-the-wild advanced persistent threat campaigns. Future evaluation should extend to long-running honeynet captures from production environments and adversarial red-team engagements.

2) *Physical-layer attacks*: QuantumGuard currently performs deception orchestration at the network and application layers and does not directly address physical-layer attacks against 6G radio-access infrastructure (e.g., pilot-contamination, jamming, beam-stealing). Extending the cognitive honeynet abstraction to encompass synthetic radio-frequency decoys is identified as a priority for future work.

3) *Deception-aware adversaries*: The RL-DPE policy was trained and evaluated against an adversary model whose probing strategy is drawn from the empirical distribution of the two corpora; it was not stress-tested against adversaries that are explicitly aware of the QuantumGuard deception fabric. The headline 96.42 per cent engagement-retention and 91.8 per cent TTP-attribution numbers should therefore be read as upper bounds against the evaluated threat distribution, not as worst-case guarantees. Assessing robustness against deception-aware adversaries, in particular those employing counter-fingerprinting in the spirit of [14] or adversarial perturbations of L7 metadata designed to confuse ATTAN, is the most important direction for follow-up work and, in the author's view, the principal threat to the present results' external validity (see Section V G).

4) *Reward-weight sensitivity*: The reward weights $(\lambda_e, \lambda_a, \lambda_o)$ in Eq. (2) were fixed by informal pilot tuning rather than by a systematic sensitivity analysis. While fixing the weights *a priori* guards against optimistic bias on the evaluation data, it leaves their robustness uncharacterized. Qualitatively, increasing λ_a sharpens attribution at some cost to engagement-retention, while increasing λ_o suppresses morphing and lowers overhead at the cost of shorter engagement; the ordering $\lambda_a > \lambda_e > \lambda_o$ encodes the design priority of attribution information gain under a bounded production-impact budget. A grid sweep over the weight simplex, together with a Pareto-front analysis of the engagement–attribution–overhead trade-off reported as stability bands around the operating point, is left to future work.

5) *Adversarial federation*: Defending against Byzantine participants and against participant-side model-inversion or membership-inference attacks against the federated deception policy is out of scope for the present work and is planned along the robustness roadmap of Section III E.

6) *Multi-organization external validity*: Because all organizations were emulated within one controlled testbed, a genuine multi-party deployment over heterogeneous, independently administered slices is required before the cross-organizational coverage results can be claimed to generalize.

7) *Sustained-attack overhead*: The 1.7 per cent production-traffic overhead was measured under the dataset-replay distribution; overhead behavior under sustained-attack stress conditions such as volumetric DDoS, slow-rate denial-of-service, or coordinated multi-source probing is not characterized here and must be quantified before operational deployment.

8) *Cryptographic agility*: The post-quantum primitives used in PQ-FIB (Kyber-768 and Dilithium-3) reflect current NIST standardization choices but may be superseded by future cryptanalytic advances or by hybrid signature schemes such as SLH-DSA. A pluggable PQ-primitive abstraction is, therefore, desirable to ensure long-term cryptographic agility.

VI. CONCLUSION

This study has presented QuantumGuard, a proactive post-quantum-resilient deception framework for cybersecurity in 6G-enabled cyber-physical systems. The principal contribution is architectural: QuantumGuard is, to the best of our knowledge, the first framework to integrate cognitive honeynet orchestration, reinforcement-learning-driven deception adaptation, transformer-based TTP-attribution, and post-quantum federated intelligence sharing into a single closed-loop control architecture suitable for 6G cyber-physical environments. A preliminary evaluation on the CICAPT-IIoT-2024 and Edge-IIoTset benchmark datasets, complemented by a 320-endpoint 6G slice testbed, provides initial evidence of operational viability under the evaluated threat model, with engagement-retention of approximately 96.42 per cent, TTP-attribution accuracy of approximately 91.8 per cent, a 78.6 per cent reduction in median attacker dwell time, a deception-induced production overhead of 1.7 per cent, and a 14.2 ms median post-quantum handshake latency. We position these numbers as initial feasibility evidence rather than as fully characterized performance bounds. A follow-up empirical study, currently in preparation, will conduct a reward-weight sensitivity analysis and will extend the evaluation to deception-aware adversaries, sustained-attack stress conditions, and multi-organization deployments. Beyond the immediate empirical program, future research will extend QuantumGuard to physical-layer 6G deception and to integration with emerging AI-driven security-orchestration platforms.

AUTHOR'S CONTRIBUTIONS

The author conducted the conceptualization, methodology, software implementation, validation, formal analysis, data curation, visualization, and the writing and editing of this manuscript. The author has read and approved the final version of the article.

DECLARATION OF COMPETING INTERESTS

The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this study.

FUNDING

No funding was received for this work.

DATA AVAILABILITY

The CICAPT-IIoT-2024 and Edge-IIoTset datasets used in this study are publicly available from the Canadian Institute for Cybersecurity and the Edge-IIoTset distribution, respectively. The complete implementation of QuantumGuard, including the CHO morphing engine, the R-PPO RL-DPE training script, the ATTAN transformer attribution network, the PQ-FIB hybrid TLS implementation, and the configuration files needed to reproduce all reported results (random seeds, hyperparameter values, and dataset preprocessing pipeline), are available upon reasonable request to the corresponding author.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work, the author used AI-assisted writing tools to improve readability and language. After using such tools, the author reviewed and edited the content as needed and take full responsibility for the content of the published study.

ACKNOWLEDGMENT

The author thank the anonymous reviewers for their constructive feedbacks.

REFERENCES

- [1] M. Series, "IMT-2030 framework and overall objectives of the future development of IMT for 2030 and beyond," Recommendation ITU-R M.2160, International Telecommunication Union, 2023.
- [2] W. Saad, M. Bennis, and M. Chen, "A vision of 6G wireless systems: applications, trends, technologies, and open research problems," *IEEE Network*, vol. 34, pp. 134–142, 2020.
- [3] Y. Liu, X. Yuan, Z. Xiong, J. Kang, X. Wang, and D. Niyato, "Federated learning for 6G communications: challenges, methods, and future directions," *China Commun.*, vol. 17, pp. 105–118, 2020.
- [4] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM J. Comput.*, vol. 26, pp. 1484–1509, 1997.
- [5] D. J. Bernstein and T. Lange, "Post-quantum cryptography," *Nature*, vol. 549, pp. 188–194, 2017.
- [6] L. Chen, S. Jordan, Y.-K. Liu, D. Moody, R. Peralta, R. Perlner, and D. Smith-Tone, "Report on post-quantum cryptography," NISTIR 8105, National Institute of Standards and Technology, 2016.
- [7] J. Mosca, "Cybersecurity in an era with quantum computers: will we be ready?" *IEEE Security & Privacy*, vol. 16, pp. 38–41, 2018.
- [8] V. Mavroeidis, K. Vishi, M. D. Zych, and A. Jøsang, "The impact of quantum computing on present cryptography," *Int. J. Adv. Comput. Sci. Appl.*, vol. 9, pp. 405–414, 2018.
- [9] M. Joseph et al., "Transitioning organizations to post-quantum cryptography," *Nature*, vol. 605, pp. 237–243, 2022.
- [10] G. Cabrera-Aldaya and B. B. Brumley, "When one vulnerable primitive turns viral: novel single-trace attacks on ECDSA and RSA," *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, vol. 2020, pp. 196–221, 2020.
- [11] A. Mosenia and N. K. Jha, "A comprehensive study of security of Internet-of-Things," *IEEE Trans. Emerg. Top. Comput.*, vol. 5, pp. 586–602, 2017.
- [12] N. Schneider and F. Petitcolas, "Cyber-physical systems security: a survey," *IEEE Internet Things J.*, vol. 4, pp. 1802–1831, 2017.
- [13] M. H. Almeshekeh and E. H. Spafford, "Cyber security deception," in *Cyber Deception*. Cham: Springer, 2016, pp. 25–52.
- [14] A. Vetterl and R. Clayton, "Bitter harvest: systematically fingerprinting low- and medium-interaction honeypots at internet scale," in *Proc. 12th USENIX WOOT*, 2018.

- [15] X. Wang, J. Zhao, and J. Zhang, "A dynamic honeypot design for intrusion detection," in *Proc. IEEE/ACS Int. Conf. Comput. Syst. Appl.*, 2014, pp. 1–6.
- [16] V. Mavroeidis and S. Bromander, "Cyber threat intelligence model: an evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence," in *Proc. EISIC*, 2017, pp. 91–98.
- [17] M. A. Ferrag, O. Friha, D. Hamouda, L. Maglaras, and H. Janicke, "Edge-IIoTset: a new comprehensive realistic cyber security dataset of IoT and IIoT applications for centralized and federated learning," *IEEE Access*, vol. 10, pp. 40281–40306, 2022.
- [18] B. Stojkovski, G. Lenzini, V. Koenig, and S. Rivas, "What's in a cyber threat intelligence sharing platform?" in *Proc. ACSAC*, 2021, pp. 385–398.
- [19] T. D. Wagner, K. Mahhub, E. Palomar, and A. E. Abdallah, "Cyber threat intelligence sharing: survey and research directions," *Comput. Secur.*, vol. 87, p. 101589, 2019.
- [20] T. T. Nguyen and V. J. Reddi, "Deep reinforcement learning for cyber security," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, pp. 3779–3795, 2023.
- [21] M. Standen, M. Lucas, D. Bowman, T. J. Richer, J. Kim, and D. Marriott, "Cyborg: a gym for the development of autonomous cyber agents," in *1st IJCAI Workshop on Adaptive Cyber Defense*, 2021.
- [22] M. Wang, J. Zhao, R. Beg, and Z. Long, "A reinforcement learning approach for adaptive cyber deception," in *Proc. IEEE SPW*, 2023, pp. 261–268.
- [23] G. Alagic et al., "Status report on the third round of the NIST post-quantum cryptography standardization process," NISTIR 8413, National Institute of Standards and Technology, 2022.
- [24] J. Pawlick, E. Colbert, and Q. Zhu, "A game-theoretic taxonomy and survey of defensive deception for cybersecurity and privacy," *ACM Comput. Surv.*, vol. 52, pp. 82:1–82:28, 2019.
- [25] G. Kambourakis, G. Drosatos, and V. G. Kounoudes, "A moving target defense honeynet for IoT environments," *J. Inf. Secur. Appl.*, vol. 67, p. 103192, 2022.
- [26] M. Dowling, M. Schukat, and E. Barrett, "Using reinforcement learning to conceal honeypot functionality," in *Proc. ECML-PKDD*, 2018, pp. 341–355.
- [27] L. Huang and Q. Zhu, "Adaptive honeypot engagement through reinforcement learning of semi-Markov decision processes," in *Decision and Game Theory for Security*. Springer, 2019, pp. 196–216.
- [28] A. Sen and R. Madria, "Cyber deception in distributed systems: a comprehensive survey," *J. Netw. Comput. Appl.*, vol. 213, p. 103608, 2024.
- [29] R. Ranpara and S. K. Patel, "Post-quantum lightweight intrusion detection for industrial IoT: a hybrid lattice-based framework," *Sci. Rep.*, vol. 14, p. 17822, 2024.
- [30] H. Chen, Z. Wang, Z. Zhao, and Q. Wang, "Federated honeynet for distributed IoT cyber deception," *IEEE Internet Things J.*, vol. 11, pp. 8231–8245, 2024.
- [31] D. Thakur, A. Guzzo, G. Fortino, and F. Piccialli, "Green federated learning: a new era of green-aware AI," *ACM Comput. Surv.*, vol. 57, pp. 199:1–199:35, 2025.
- [32] "CICAPT-IIoT-2024: A multi-stage advanced persistent threat dataset for IIoT security," Canadian Institute for Cybersecurity, University of New Brunswick, 2024. [Online]. Available: <https://www.unb.ca/cic/datasets/>
- [33] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," arXiv:1707.06347, 2017.
- [34] J. Lin, "Divergence measures based on the Shannon entropy," *IEEE Trans. Inf. Theory*, vol. 37, pp. 145–151, 1991.
- [35] B. Strom, A. Applebaum, D. Miller, K. Nickels, A. Pennington, and C. Thomas, "MITRE ATT&CK: design and philosophy," The MITRE Corporation, Tech. Rep., 2018.
- [36] R. Avanzi et al., "CRYSTALS-Kyber algorithm specifications and supporting documentation," NIST PQC Round 3, 2021.
- [37] S. Bai et al., "CRYSTALS-Dilithium algorithm specifications and supporting documentation," NIST PQC Round 3, 2021.
- [38] A. Al-Dhaqum et al., "Development and validation of a Database Forensic Metamodel (DBFM)," *PLOS ONE*, vol. 12, p. e0170793, 2017.
- [39] A. Al-Dhaqum, S. A. Razak, K. Siddique, R. A. Ikuesan, and V. R. Kebande, "Towards the development of an integrated incident response model for database forensic investigation field," *IEEE Access*, vol. 8, pp. 145018–145032, 2020.
- [40] A. Al-Dhaqum et al., "Categorization and organization of database forensic investigation processes," *IEEE Access*, vol. 8, pp. 112846–112858, 2020.
- [41] A. Al-Dhaqum, S. A. Razak, R. A. Ikuesan, V. R. Kebande, and K. Siddique, "A review of mobile forensic investigation process models," *IEEE Access*, vol. 8, pp. 173359–173375, 2020.
- [42] A. Al-Dhaqum et al., "Digital forensics subdomains: the state of the art and future directions," *IEEE Access*, vol. 9, pp. 152476–152502, 2021.
- [43] A. Al-Dhaqum, S. A. Razak, R. A. Ikuesan, V. R. Kebande, and S. H. Othman, "Face validation of database forensic investigation metamodel," *Infrastructures*, vol. 6, p. 13, 2021.
- [44] V. R. Kebande et al., "Quantifying the need for supervised machine learning in conducting live forensic analysis of emergent configurations (ECO) in IoT environments," *Forensic Sci. Int.: Reports*, vol. 2, p. 100122, 2020.
- [45] A. Al-Dhaqum et al., "Categorization and organization of database forensic investigation processes for the Internet of Things (CDBFIP)," *IEEE Access*, vol. 5, pp. 24401–24416, 2017.