Received: 11 September 2026 | *Revised:* 5 November 2026 | *Accepted:* 15 November 2026 | *Vol. 2 | Issue 2 | December 2026*
Online ISSN: 3105-2169

International Journal of Mechatronics, Robotics, and Artificial Intelligence
Original Article

From Beamforming to Behavior: PHY-Assisted Continuous Anomaly Scoring for Zero-Trust IIoT Security in Oilfield Networks

Mustafa S. Aljumaily*[1]
[1]Research and Development (R&D) Department, Daw Alfada Company, Baghdad, Iraq

Correspondence
*Mustafa Sadiq Aljumaily
Research and Development (R&D) Department, Daw Alfada Company, Baghdad, Iraq.
Email: mustafa.s@daw-alfada.com

Abstract
Zero-trust (ZT) architectures for wireless industrial Internet of Things (IIoT) deployments usually use physical-layer (PHY) security only at authentication time. This paper asks whether channel state information (CSI) and beamspace statistics from massive multiple-input multiple-output (MIMO) and millimeter-wave (mmWave) links can instead provide a continuous behavioral signal that a ZT policy engine may consume during a session. We propose a two-tier pipeline: an edge autoencoder turns window-level CSI, beamspace and short-lag channel-autocorrelation features into a formally defined, smoothed trust score, and a gateway classifier attribute attacks on the windows the edge tier flags. We evaluate four attack classes (power-neutral spoofing, relay/man-in-the-middle insertion, physical tampering, beam-training jamming) on a clustered-multipath simulator, using sliding windows, a sequence-level train/test split, balanced metrics and 10 random seeds. The edge tier reaches an area under the ROC curve (AUC) of 0.971 and a balanced accuracy of 0.907, but is statistically indistinguishable from a Mahalanobis-distance detector. Temporal self-similarity features improve five-class accuracy by 1.7 points, whereas angular/beamspace features add nothing once they are present. The end-to-end cascade reaches 0.912 accuracy, with 0.69 relay/MITM recall. Stress tests show that a relay with its own geometry, a statistics-matched forger, uncalibrated hardware impairments and legitimate equipment relocation each degrade detection or cause false alarms, and a policy simulation over 600-step sessions shows that legitimate environmental drift produces false restrictions in most sessions. We conclude that PHY-derived scores are a useful but weak input to a ZT policy engine, not a stand-alone trigger, and we identify drift handling and adversarial robustness as open problems.

Keywords
Anomaly Detection, Beamforming, Channel State Information, Industrial IoT, Massive MIMO, mmWave, Oilfield Networks, Physical-Layer Security, Zero-Trust Security.

I. Introduction

Industrial IoT deployments in upstream oil and gas differ sharply from enterprise IT: assets are dispersed across well pads, gathering stations and pipeline corridors, backhaul is intermittent or bandwidth-constrained, and physical access control cannot be assumed at remote, unmanned sites. As private 5G and millimeter-wave (mmWave) access reach these sites, spoofed nodes, relay or man-in-the-middle (MITM) insertion, tampering and jamming become feasible against links that perimeter security was never designed to monitor continuously.

Zero trust (ZT) is the usual answer: no device, link or session is trusted by default, and trust is re-evaluated continuously rather than granted once [1]. The National Institute of Standards and Technology (NIST) SP 800-207 describes a policy decision point (PDP) and policy enforcement point (PEP) that consume signals about the monitored system's state and behavior [1], and recent surveys cover ZT for IoT and industrial IoT [2]-[4]. For wireless IIoT specifically, Lei et al. note that no concrete ZT implementation technique existed and propose a three-step framework (security-zone formation, device authentication, key negotiation) built on PHY-layer primitives [5]. That framework works at the identity layer: the PHY establishes who a device is, once. It does not produce an ongoing signal about how a link is behaving.

Separately, the PHY carries rich information about attacks. Dinh-Van et al. show that autoencoders trained on nominal mmWave beam-training behavior can detect and mitigate

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.
© 2026 The Authors.

beam-training jamming in a manufacturing setting [6]. This establishes the PHY as a learnable detection signal but within one protocol phase and with no link to access control.

This paper studies the space between these results, and states its conclusions conservatively. We ask whether CSI and beamspace statistics from a massive-MIMO/mmWave link can yield a continuous anomaly score that a ZT policy engine could consume alongside identity and device-posture signals, and how reliable such a score is under the conditions an oilfield deployment would impose (clutter, hardware impairments, benign equipment movement, environmental drift, and attackers that adapt). The contributions are:

- **A two-tier architecture and a formal trust score:** An always-on edge autoencoder produces a calibrated, smoothed trust score T_t (Eq. 3–4), and a gateway classifier attribute attacks only on windows the edge tier flags (Section III).
- **A leakage-free evaluation protocol:** Continuous streams, sliding windows, a sequence-level split, a trivial-classifier reference, balanced metrics, AUC for every method, and 10 seeds (Section IV).
- **Feature-family evidence:** Short-lag channel-autocorrelation features improve five-class accuracy by 1.7 points in 10 of 10 seeds; angular/beamspace features are redundant once temporal features are present (Section V-D).
- **Stress tests and negative results:** Relays with their own geometry, statistics-matched forgers, hardware impairments and benign relocation (Section V-F).
- **A policy simulation over long sessions:** It quantifies how drift and slow-poisoning attackers interact with re-calibration policies and shows that the score is not reliable enough to drive restrictive actions on its own (Section V-G).

The remainder of the paper is organized as follows. Section II reviews related work, Section III presents the framework, Section IV describes the simulation and protocol, Section V reports results, Section VI discusses implications and limitations, and Section VII presents the main conclusions.

II. Related Work

A. CSI-Based Anomaly and Intrusion Detection

Wi-Fi CSI amplitude and phase have long been used for device-free intrusion and presence detection [7]–[9], and CSI has been used to detect physical tampering such as transmitter displacement or an antenna-orientation change [10]. Deep-learning-based CSI anomaly detection for wireless networks has also been studied [11]. These works mostly use single- or few-antenna CSI and are not framed within an access-control loop.

B. Physical-Layer Authentication in MIMO and mmWave Systems

Physical-layer authentication (PLA) decides whether a received signal comes from the claimed transmitter. For massive MIMO, EsaNet extracts a frequency-independent fingerprint from CSI and reports robustness in time-varying environments [12]; Pham et al. use angle-of-arrival estimation for learning-based PLA [13]; and Abdrabou and Gulliver treat PLA for MIMO as a one-class problem trained on legitimate data only [14]. Online learning-based PLA has been proposed for mmWave MIMO [15], and carrier-frequency-offset features have been used for mmWave PLA [16]. Two recent preprints bear directly on the problems studied here: one combines Transformer channel prediction with correlation-based authentication to separate sustained spoofing from natural channel variation [17], and other models non-stationary MIMO channels, blockage and adaptive spoofers [18]. Our work differs in aim: it does not authenticate a claimed identity but scores link behavior over a session, and it asks how that score would feed a policy engine. Unlike [15], which combines mmWave channel and carrier-frequency-offset features and updates its authentication parameters online as new transmissions arrive, our Tier 1 is calibrated once at commissioning (adaptation is studied separately in Section V-G) and scores link behavior without verifying a claimed identity.

C. Relay and Man-in-the-Middle Detection

Relay attacks are hard to detect from channel statistics because the relayed signal itself traverses a real channel. Countermeasures include PHY/MAC-layer mechanisms with time bounds for Wi-Fi connection establishment [19], and analog MITM relays have been shown to fabricate link signatures used for source identification [20]. This is why we treat relay/MITM as the hardest class and why Section V-F reports relays that deviate from our baseline model.

D. RF Fingerprinting and Attacks on Learned PHY Classifiers

Radio-frequency (RF) fingerprinting identifies devices from hardware impairments [21], including for IIoT authentication [22]; it is an identity mechanism and answers a different question from ours. Learned PHY classifiers are themselves attack targets: white-box and universal adversarial perturbations reach high success against deep-learning RF fingerprinting [23][24], domain-shift misclassification can be exploited as a backdoor [25], and poisoned training samples degrade learning-aided PLA [26]. These results motivate the adaptive-attacker and poisoning experiments in Sections V-F and V-G.

E. Zero-Trust Architectures for IIoT

Feng and Hu propose a cyber-physical zero-trust architecture for industrial cyber-physical systems [27], and Sedjelmaci and Ansari propose a ZT-based attack-detection framework for 6G edge computing that scores network-layer behavior [28]. Recent surveys review ZT implementation challenges across domains, in IoT, and in industrial IoT [2]–[4]. Lei et al. [5] remains the closest PHY-based ZT proposal and operates at authentication and key-negotiation time.

F. mmWave Attack Detection and Drift

Jamming detection in massive MIMO has been studied through angular-domain analysis [29], Dinh-Van et al. defend beam training in manufacturing networks [6], and more recent work applies reinforcement learning to beam-stealing defense [30]. Because industrial RF environments are non-stationary,

concept drift matters: drift management has been studied for wireless network models [31][32], and drift-aware incremental anomaly detection for streams and digital-twin IoT has been proposed [33][34]. We examine the specific hazard that re-calibration can absorb a slow attacker.

TABLE I.
POSITIONING RELATIVE TO RELATED WORK

Aspect	PLA / CSI anomaly [7]–[18]	PHY-based ZT [5]	mmWave attack detection [6], [29], [30]	This work
Aim	Authenticate transmitter or detect intrusion	Identity, keys at onboarding	Detect jamming/beam attacks	Session-long behavior score for a ZT policy engine
Policy-layer simulation	No	No	No	Yes (Section V-G), with negative findings
Massive-MIMO beamspace features	Partial	No	Partial	Yes (found redundant, Section V-D)
Drift / slow-poisoning evaluation	Partial [17], [18]	No	No	Yes (Section V-G)
Evidence base	Mostly simulation and some measurements	Analysis and simulation	Simulation and some experiments	Simulation only

G. Positioning

Table I positions this work. We do not claim that no prior work touches these elements; the combination of a PHY-derived session-long behavioral score, a policy-layer simulation and an explicit account of failure modes is what we contribute.

III. THE PROPOSED FRAMEWORK

A. System Model

A gateway with an N-element uniform linear array (ULA, $N = 32$) serves IIoT nodes over a mmWave link in a topology representative of a gathering station or well pad: nodes are static or slow-moving and the link has dense multipath from metallic structures. At timestep t the channel vector is $h_t \in \mathbb{C}^N$, composed of a Rician line-of-sight (LoS) component ($K = 5$) and L scattering clusters whose angles and gains evolve slowly.

B. Feature Extraction

Each timestep yields 15 features in three families. Base features (6): total received power $P_t = \|h_t\|^2$, circular phase variance of adjacent-antenna phase differences, amplitude mean and variance, and first differences of power and mean amplitude. Angular features (5): with the DFT beam codebook $b_t = F\, h_t/\sqrt{N}$ and normalized beam powers p_t, the beamspace entropy, the sine and cosine of the dominant beam's angle, the normalized power of the strongest two beams, and of the strongest beam. Temporal self-similarity features (4): for lags $k = 1, 2, 3$,

$$r_k(t) = \frac{|h_t^H h_{t-k}|}{\|h_t\|\|h_{t-k}\|} \quad (1)$$

$$\text{corr_bump}(t) = max(0, r_2 - r_1) + max(0, r_3 - r_2) \quad (2)$$

A physically evolving channel should have $r_1 \geq r_2 \geq r_3$; $\text{corr}_{\text{bump}}$ flags local violations of that monotone decay. Windows of $W = 10$ timesteps are formed with stride 1; Tier-1 inputs are the mean and standard deviation of each feature over the window (30 values), and Tier-2 inputs are the full 10 × 15 window.

C. Tier 1 and the Trust Score

Tier 1 is a feed-forward autoencoder (tanh, 16-6-16 hidden units) trained only on benign windows after per-feature standardization. Let e_t be its mean squared reconstruction error on the window ending at t, τ the 95th percentile of benign training error, and s the standard deviation of ln e over benign training windows. A window is flagged if $e_t > \tau$. The raw trust score and its exponentially weighted moving average (EWMA) smoothed version are

$$T_t = \frac{1}{1+exp\left(\frac{ln\, e_t - ln\, \tau}{s}\right)} \quad (3)$$

$$T_t = \alpha T_t + (1-\alpha)T_{t-1}, \;\; \alpha = 0.5 \quad (4)$$

so that $\tilde{T}_t = 0.5$ exactly at the calibrated threshold and $T_t \in (0, 1)$ decreases as behavior departs from the commissioning baseline. T_t is a calibrated anomaly score, not a probability of compromise.

D. Tier 2 and the Cascade

Tier 2 is a five-class classifier (benign, spoofing, relay/MITM, tampering, jamming) operating on the full window. In the cascade, Tier 2 is evaluated only on windows flagged by Tier 1; unflagged windows are labeled benign. This is the deployment-relevant configuration, since it determines how often the gateway classifier runs and what Tier-1 misses cost.

E. Interface to the Policy Engine

The pipeline outputs T_t (and, when computed, an attack class) to a Zero-Trust PDP alongside identity, device-posture and other signals; the PDP decides and a PEP enforces (Fig. 1). For the experiments in Section V-G we instantiate a simple PDP: ALLOW; STEP-UP authentication if $T_t < 0.6$; RESTRICT if $T_t < 0.3$ for p consecutive windows; QUARANTINE (latched) if $T_t < 0.1$ for p consecutive windows; return to ALLOW after $T_t \geq 0.6$ for 10 consecutive windows unless quarantined. This PDP is a probe of the score's usefulness, not a proposed production policy.

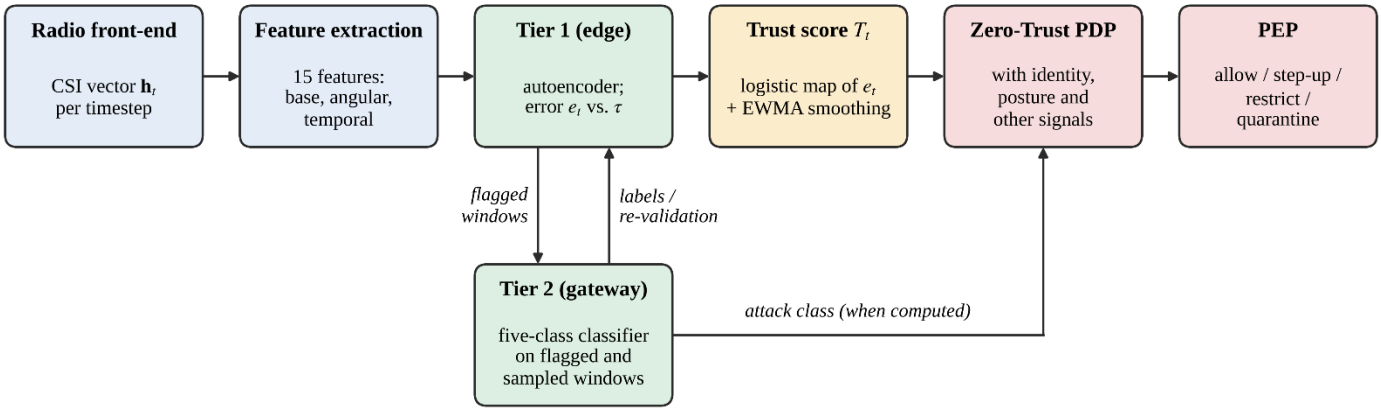

Fig. 1. Proposed pipeline and its interface to a Zero-Trust policy loop.

IV. Methods and Procedures

A. Channel Simulation

Channels are generated as clustered multipath sums of ULA steering vectors with a Rician LoS cluster ($K = 5$) and C − 1 scattering clusters. LoS angle and phase follow slow random walks ($\sigma = 0.2°$ and 0.05 rad per step); cluster gains follow AR(1) processes ($\rho = 0.97$) and cluster angles random walks ($\sigma = 0.3°$ per step), over 60-timestep streams. The cluster count C is swept from 2 (sparse) to 12 (dense), with $C = 9$ as the primary setting; we borrow only the qualitative cluster-count regime of the 3GPP TR 38.901 Indoor Factory (InF) family [35]. InF is an indoor model, so this is an approximation of dense metallic scattering and not a validated oilfield outdoor channel. Receiver noise gives a 20 dB per-element signal-to-noise ratio (SNR).

The simulator used in this revision was re-implemented from the specification given in this section. All results reported here come from that re-implementation. The code, including the simulator, feature extraction, models and evaluation scripts, is released as supplementary material. Benign streams include transient blockage (probability 0.5; LoS attenuated by a factor 0.4–0.7 for 3–8 steps) as hard negatives.

B. Attack Models

Attacks start at an onset drawn uniformly from steps 14–39 and persist to the end of the stream.

*Spoofing***:** a rogue path at a 15–50° angular offset takes a 20–60% power fraction while the legitimate path is scaled down to keep total power constant (power-neutral).

*Relay/MITM***:** a phase-rotated copy of the legitimate channel delayed by 1–4 steps is mixed in with weight $\beta \in [0.1, 0.35]$.

*Tampering***:** an abrupt LoS angle shift of 8–25° with gain scaled to 0.5–0.85.

*Jamming***:** interference at a randomized angle each step with power 0.3–1.5 times the legitimate power. A global strength factor scales all attack magnitudes.

C. Dataset and Protocol

For each of five classes, 160 streams of 60 steps are generated. Sliding windows ($W = 10$, stride 1) are extracted from every stream starting at step 3, so windows are not anchored to the attack onset. A window is labeled with the attack class if at least 5 of its 10 steps are attacked and benign if none are; windows with 1–4 attacked steps are excluded from classification metrics. Streams are split 70/30 into training and test sets by stream, so overlapping windows never straddle the split. One-class methods are trained only on benign training windows. All headline results use 10 independent seeds (new data, split and initialization each) and are reported as mean ± standard deviation; the stress, long-session and clutter experiments use 3–5 seeds as stated. With this design, about 0.534 ± 0.021 of test windows are attack windows, so we always report the accuracy and F1 of a classifier that labels every window an attack.

D. Models and Baselines

Tier 1 is the autoencoder of Section III-C. Baselines on the same 30-dimensional window summaries are Isolation Forest, One-Class SVM (RBF), PCA reconstruction error

(90% variance), Mahalanobis distance, and a four-component diagonal Gaussian mixture; a window-level autoencoder on the flattened 150-dimensional window is also included. All thresholds are the 95th percentile of benign training scores. Tier 2 in the reported experiments is a multilayer perceptron (MLP, 128-64 hidden units) on the flattened window.

E. Metrics

For binary detection we report accuracy, balanced accuracy, F1, true-positive rate (TPR), false-positive rate (FPR), ROC AUC, precision-recall AUC (PR-AUC) and TPR at 5% FPR for every method, and the same metrics for the label-everything-attack reference. For five-class results we report accuracy, macro-F1, balanced accuracy, and per-class recall, precision and FPR. For the trust score we report detection delay and sequence-level false-alarm rates. Per-sample latency was not re-measured on the re-implemented code, and no latency claim is made.

F. Hyperparameters

Table II lists the simulation and model hyperparameters used throughout.

TABLE II.
SIMULATION AND MODEL HYPERPARAMETERS

Item	Value
Array / channel	32-element ULA; $K = 5$; $C = 9$ clusters (sweep 2–12); 60-step streams; SNR 20 dB
Dynamics	LoS angle/phase walks 0.2° / 0.05 rad; cluster gain AR(1) $\rho = 0.97$; cluster angle walk 0.3°
Benign hard negatives	Blockage prob. 0.5, LoS factor 0.4–0.7, 3–8 steps
Attacks (strength 1.0)	Spoof: fraction 0.2–0.6, offset 15–50°; relay: β 0.1–0.35, lag 1–4; tamper: 8–25°, gain 0.5–0.85; jam: power 0.3–1.5; onset 14–39
Windows / labels	$W = 10$, stride 1; attack if $\geq 5/10$ steps attacked, benign if 0; 1–4 excluded
Data / split	160 streams per class; 70/30 split by stream; 10 seeds
Tier 1	MLP autoencoder 30-16-6-16-30, tanh, ≤ 300 iterations; threshold = 95th percentile benign error
Trust score	s = std of ln e (benign train); $\alpha = 0.5$; alarm thresholds as stated
Tier 2	MLP 150-128-64-5, ≤ 60 iterations, trained on every second labeled training window

V. RESULTS

A. Binary Detection and Baselines

Table III reports binary detection over about 10748 labeled test windows per seed. A classifier that labels every window an attack scores 0.534 accuracy and F1 0.696 here, and 0.5 balanced accuracy; accuracy and F1 alone are therefore not informative on this mix, and we read the balanced metrics and AUC. Tier 1 reaches a balanced accuracy of 0.907 ± 0.015 and an AUC of 0.971 ± 0.007, at a window-level FPR of 0.062 ± 0.018. It is statistically indistinguishable from Mahalanobis distance, the Gaussian mixture and the One-Class SVM, and better than Isolation Forest, PCA reconstruction and the flattened-window autoencoder. The autoencoder's advantage is therefore not detection accuracy; its appeal is its compact form, and we do not claim a detection advantage over simple Gaussian baselines.

TABLE III.
BINARY DETECTION (C = 9, 10 SEEDS, MEAN ± STD, SLIDING WINDOWS)

Method	Acc.	Bal. acc.	F1	TPR	FPR	AUC	PR-AUC
Label everything attack	0.534 ± 0.021	0.500 ± 0.000	0.696 ± 0.018	1.000 ± 0.000	1.000 ± 0.000	0.500 ± 0.000	0.534 ± 0.021
Tier 1 autoencoder	0.905 ± 0.016	0.907 ± 0.015	0.907 ± 0.016	0.876 ± 0.030	0.062 ± 0.018	**0.971 ± 0.007**	**0.976 ± 0.006**
Mahalanobis distance	0.900 ± 0.012	0.903 ± 0.011	0.902 ± 0.012	0.865 ± 0.025	0.059 ± 0.017	0.970 ± 0.006	0.975 ± 0.005
Gaussian mixture (4, diag.)	0.889 ± 0.014	0.893 ± 0.013	0.890 ± 0.014	0.840 ± 0.022	0.055 ± 0.015	0.967 ± 0.005	0.973 ± 0.005

One-Class SVM (RBF)	**0.913 ± 0.013**	**0.914 ± 0.013**	**0.918 ± 0.012**	**0.908 ± 0.019**	0.081 ± 0.023	0.965 ± 0.006	0.974 ± 0.005
Isolation Forest	0.734 ± 0.027	0.748 ± 0.021	0.687 ± 0.036	0.549 ± 0.047	**0.054 ± 0.014**	0.930 ± 0.010	0.928 ± 0.015
PCA reconstruction	0.820 ± 0.015	0.828 ± 0.014	0.810 ± 0.017	0.718 ± 0.024	0.062 ± 0.014	0.918 ± 0.012	0.938 ± 0.009
Window autoencoder (flat 150-d)	0.741 ± 0.023	0.754 ± 0.022	0.698 ± 0.036	0.564 ± 0.046	0.056 ± 0.016	0.928 ± 0.009	0.933 ± 0.012

B. Tier 2 and the End-to-End Cascade

Tier 2 applied to every window reaches 0.946 ± 0.009 accuracy and 0.939 ± 0.010 macro-F1 (Table IV). In the cascade, Tier 2 sees only flagged windows (a fraction 0.523 ± 0.021 of all windows); accuracy falls to 0.912 ± 0.018 and macro-F1 to 0.902 ± 0.019 because Tier-1 misses are never reclassified. The edge/gateway split thus costs about three accuracy points and halves the number of Tier-2 invocations; it is a compute trade-off, not an accuracy gain. Table V gives per-class results for the cascade: jamming is almost always detected (recall 0.998), tampering and spoofing are detected at 0.855 and 0.829 recall, and relay/MITM is the hardest class at 0.691 ± 0.076 recall, so roughly one relay window in three is missed. The "benign" row shows that 0.138 of non-benign windows are predicted benign.

TABLE IV.
TIER 2 ALONE VERSUS END-TO-END CASCADE (FIVE-CLASS)

Pipeline	Accuracy	Macro-F1	Balanced acc.
Tier 2 alone (all windows)	**0.946 ± 0.009**	**0.939 ± 0.010**	**0.935 ± 0.010**
Cascade (Tier 2 on Tier-1-flagged windows)	0.912 ± 0.018	0.902 ± 0.019	0.873 ± 0.025

TABLE V.
PER-CLASS RESULTS, CASCADE (ONE-VS-REST)

Class	Recall	Precision	FPR
Benign	0.990 ± 0.004	0.863 ± 0.034	0.138 ± 0.034
Spoof	0.829 ± 0.056	0.983 ± 0.012	0.002 ± 0.002
Relay	0.691 ± 0.076	0.940 ± 0.022	0.007 ± 0.003
Tamper	0.855 ± 0.063	0.949 ± 0.017	0.007 ± 0.002
Jam	0.998 ± 0.002	0.999 ± 0.001	0.000 ± 0.000

C. Trust Score Behavior

Fig. 2 shows T_t through one stream of each class, and Table VI gives the mean delay between onset and the first alarm ($T_t < 0.3$). Abrupt attacks (tampering, jamming, spoofing) alarm within about one to two steps; relay/MITM takes about 6.11 steps, and 0.213 of relay streams never alarm. Two observations matter for policy use. First, for persistent spoofing and relay, trust partly recovers after the initial drop: once the whole window lies inside the attacked regime, window summaries no longer look anomalous relative to the recent past. Second, false alarms accumulate over time: 0.219 ± 0.082 of benign 60-step streams (including blockage events) raise at least one alarm at $T_t < 0.3$, 0.528 ± 0.096 at $T_t < 0.5$, and 0.025 ± 0.035 at $T_t < 0.1$.

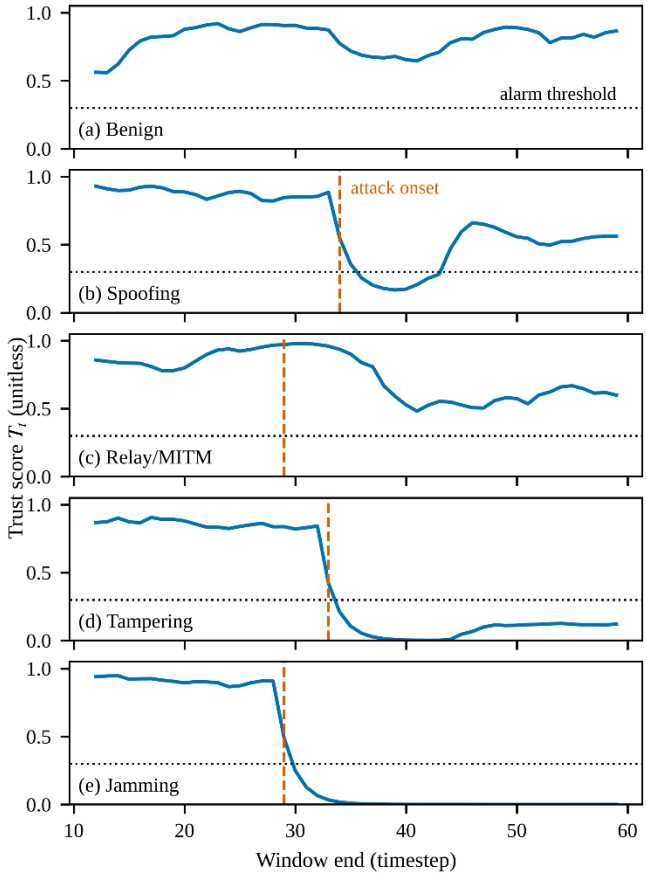

Fig. 2. EWMA trust score T_t for one stream per class (red dashed: attack onset; dotted: alarm threshold 0.3).

TABLE VI.
DETECTION DELAY AND MISSED STREAMS (ALARM AT $T_T <$ 0.3)

Attack	Mean delay after onset (steps)	Streams never alarming
Spoof	1.30 ± 0.28	0.018 ± 0.021
Relay	6.11 ± 1.34	0.213 ± 0.056
Tamper	0.94 ± 0.02	0.000 ± 0.000
Jam	0.96 ± 0.03	0.000 ± 0.000

D. Feature-Family Ablation

Table VII compares Tier-2 feature sets over the same 10 seeds. Adding temporal self-similarity features to the base set improves accuracy by +1.65 ± 0.31 points; temporal favored in 10 of 10 seeds. Adding angular/beamspace features does not help: with temporal features present, adding angular features changes accuracy by -0.42 ± 0.39 points; angular favored in 1 of 10 seeds, i.e., it does not help and slightly hurts. We therefore do not claim a contribution of angular or beamspace features, and in this simulator beamspace statistics carry little information beyond what the temporal and base features already capture. We also do not claim a relay-specific benefit of the temporal features: relay recall is 0.839 ± 0.047 for the base set and 0.846 ± 0.055 with temporal features, within noise.

TABLE VII.
FEATURE-FAMILY ABLATION, TIER 2 ALONE (10 SEEDS)

Feature set	Accuracy	Macro-F1	Relay recall	Spoof recall
Base (6)	0.933 ± 0.008	0.924 ± 0.007	0.839 ± 0.047	0.932 ± 0.033
Base + angular (11)	0.927 ± 0.007	0.916 ± 0.007	0.790 ± 0.054	0.936 ± 0.028
Base + temporal (10)	**0.950 ± 0.008**	**0.945 ± 0.008**	**0.846 ± 0.055**	0.942 ± 0.030
Base + angular + temporal (15)	0.946 ± 0.009	0.939 ± 0.010	0.821 ± 0.052	**0.948 ± 0.025**

E. Robustness to Attack Strength and Clutter

Table VIII repeats the main metrics with all attack magnitudes scaled by 0.6. Detection degrades substantially (Tier-1 AUC 0.929 ± 0.011, cascade accuracy 0.819 ± 0.025, relay recall 0.567 ± 0.105), so the main results are specific to the attack strengths chosen. Fig. 3 and Table IX sweep the cluster count. Detection improves from sparse to moderately dense clutter and plateaus at about seven clusters; relay recall rises with clutter density, which is the opposite of the sparse-channel advantage one might expect.

TABLE VIII.
ATTACK-STRENGTH ROBUSTNESS (C = 9, 10 SEEDS)

Strength	Tier-1 bal. acc.	Tier-1 AUC	Tier 2 alone acc.	Cascade acc.	Cascade relay recall	Cascade spoof recall
1.0 (primary)	0.907 ± 0.015	0.971 ± 0.007	0.946 ± 0.009	0.912 ± 0.018	0.691 ± 0.076	0.829 ± 0.056
0.6 (weaker attacks)	0.826 ± 0.021	0.929 ± 0.011	0.896 ± 0.015	0.819 ± 0.025	0.567 ± 0.105	0.505 ± 0.078

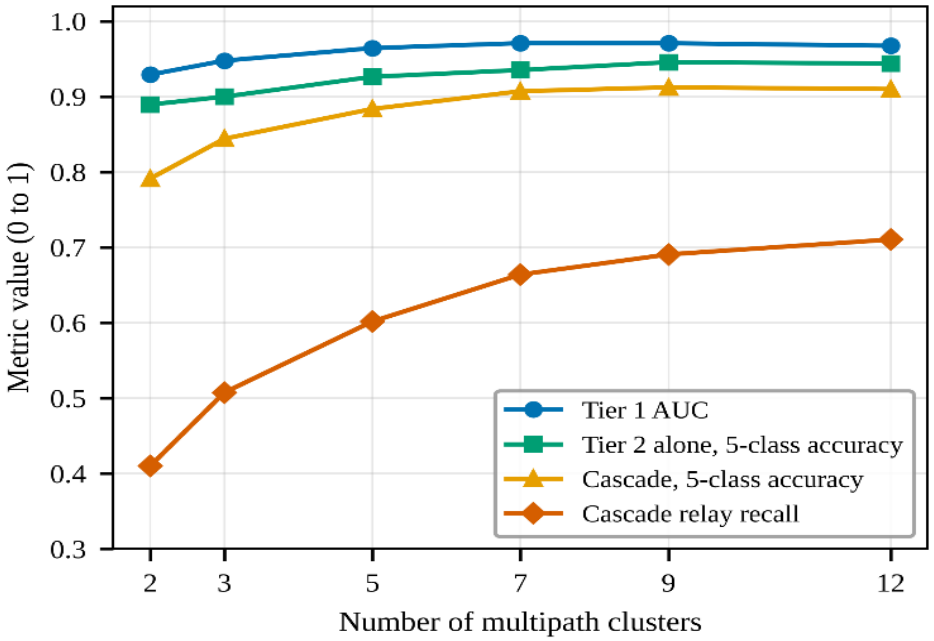

Fig. 3. Detection metrics versus multipath cluster count.

TABLE IX.
CLUTTER SWEEP (MEANS; 5 SEEDS, 10 AT C = 9)

Clusters	Tier-1 bal. acc.	Tier-1 AUC	Mahalanobis AUC	Tier 2 alone acc.	Cascade acc.	Cascade relay recall	Seeds
2	0.806	0.929	0.930	0.890	0.791	0.410	5
3	0.855	0.948	0.942	0.900	0.844	0.508	5
5	0.890	0.964	0.962	0.926	0.884	0.602	5
7	0.907	0.971	0.968	0.935	0.907	0.664	5
9	0.907	0.971	0.970	0.946	0.912	0.691	10
12	0.905	0.968	0.967	0.944	0.910	0.710	5

F. Stress Tests

Detectors trained on the standard attacks were tested on conditions that the baseline simulation omits (5 seeds; Table X).

(i) *Relay with its own geometry and delay***.** A relay with an angular offset and a 3–8 step delay is much harder: Tier-1 TPR falls from 0.74 to 0.54, cascade class recall to 0.02, and Tier-1 FPR on benign windows rises to 0.33.

(ii) *Statistics-matched forger.* An attacker that replaces the channel with an independent, benign-looking realization of the same statistics is largely undetected (Tier-1 TPR 0.36, cascade detection 0.24); only the onset transient is visible. A detector of this kind cannot be expected to catch an attacker who reproduces the expected channel statistics.

(iii) *Hardware impairments.* With per-antenna phase noise, carrier-frequency offset (CFO) drift, antenna gain/phase mismatch and 4-bit quantization, a detector trained on clean data flags almost every benign window (Tier-1 FPR 0.99); retraining on impaired data restores FPR 0.06 at TPR 0.85, compared with 0.88 without impairments.

(iv) *Benign relocation.* Legitimate relocation of the node produces a Tier-1 FPR of 0.57, and 20% of benign windows are classified as tampering; training with relocation as a benign class reduces these to 0.28 and 1.8%. Gradient-based adversarial evasion of the learned models was not evaluated.

TABLE X.
STRESS TESTS (DETECTORS TRAINED ON STANDARD ATTACKS UNLESS NOTED)

Scenario	Tier-1 TPR	Cascade detect	Cascade class recall	Tier-1 FPR	Cascade FPR	Benign → tamper
Standard relay (reference)	0.74 ± 0.04	0.69 ± 0.04	0.66 ± 0.04	0.08 ± 0.01	0.01 ± 0.01	0.00 ± 0.00
Relay with own geometry and 3–8 step delay	0.54 ± 0.02	0.39 ± 0.04	0.02 ± 0.01	0.33 ± 0.02	0.22 ± 0.03	0.01 ± 0.01
Statistics-matched forger	0.36 ± 0.01	0.24 ± 0.02	0.02 ± 0.01	0.23 ± 0.01	0.13 ± 0.01	0.09 ± 0.01
Standard spoofing (reference)	0.87 ± 0.04	0.86 ± 0.05	0.85 ± 0.05	0.08 ± 0.02	0.01 ± 0.01	0.00 ± 0.00
Hardware impairments, trained on clean data	1.00 ± 0.00	0.96 ± 0.01	0.90 ± 0.01	0.99 ± 0.00	0.12 ± 0.04	0.01 ± 0.01
Hardware impairments, trained on impaired data	0.85 ± 0.03	0.83 ± 0.03	0.81 ± 0.03	0.06 ± 0.01	0.01 ± 0.01	0.00 ± 0.00
No impairments (reference, all classes)	0.88 ± 0.01	0.86 ± 0.01	0.85 ± 0.01	0.06 ± 0.01	0.01 ± 0.00	0.00 ± 0.00
Benign relocation, trained without it	—	—	—	0.57 ± 0.01	0.43 ± 0.03	0.20 ± 0.03
Benign relocation, trained with it	—	—	—	0.28 ± 0.01	0.04 ± 0.01	0.02 ± 0.01

G. Long Sessions, Drift and the Policy Engine

We simulated 600-step sessions, with and without legitimate drift (scattering amplitude ×1.5, LoS rotated 6° and attenuated 20% between steps 150 and 350), random blockage, and attacks starting between steps 300 and 450; "slow" attacks ramp in over 150 steps to mimic a poisoning attacker. Tier 1 was trained once on standard 60-step data and run under four policies: frozen normalization; naive adaptation (the feature mean is updated on every window with $e_t \leq 2\tau$); and guarded adaptation (also requiring Tier 2 to say benign and bounding the cumulative shift to $B = 2$ or 4 standard deviations of the commissioning baseline). The PDP of Section III-E converts T_t to actions; Fig. 4 shows representative trajectories under the four policies.

At the default operating point (τ at the 95th percentile, $p = 3$), 63–70% of benign stationary sessions trigger at least one false RESTRICT, and 93–100% of drifting sessions do (Table XI). A stricter point (τ at the 99.5th percentile, $p = 5$) reduces stationary false RESTRICT to 8–17% but leaves 67–83% under drift. Adaptation lowers false restrictions modestly under drift, but it is exploitable: for slow tampering in stationary sessions at the strict point, detection is 0.58 with frozen normalization, 0.25 with naive adaptation and 0.42 with guarded adaptation (Table XII). Detection under drift is not interpretable at either operating point, because trust had usually collapsed and the session was already restricted before the attack began. In short, none of the policies separated legitimate drift from attacks.

TABLE XI.
BENIGN SESSIONS: FALSE ACTIONS (FRACTION OF SESSIONS; FRACTION OF TIME NOT IN ALLOW)

Policy	Environment	Default: false RESTRICT	Default: false QUARANTINE	Default: time not ALLOW	Strict: false RESTRICT	Strict: false QUARANTINE	Strict: time not ALLOW
frozen	stationary	0.70	0.13	0.31	0.17	**0.00**	0.04
frozen	drift	1.00	0.87	0.67	0.83	0.67	0.44
naive	stationary	**0.63**	**0.10**	**0.29**	**0.08**	**0.00**	**0.03**
naive	drift	**0.93**	**0.73**	**0.55**	**0.67**	**0.58**	**0.28**
guard2	stationary	0.67	0.13	**0.29**	0.17	**0.00**	0.04
guard2	drift	1.00	0.83	0.61	0.75	**0.58**	0.36
guard4	stationary	0.67	0.13	0.30	0.17	**0.00**	0.04
guard4	drift	1.00	0.77	0.60	0.75	**0.58**	0.36

Default: 5 seeds × 6 sessions per cell; strict: 3 seeds × 4 sessions per cell.

TABLE XII.
ATTACK DETECTION IN STATIONARY SESSIONS, STRICT OPERATING POINT (FRACTION RESTRICTED AFTER ONSET; TIME TO RESTRICT)

Attack	Frozen	Naive adaptation	Guarded (B = 4)	Time to restrict, frozen (steps)
spoof	**1.00**	**1.00**	**1.00**	7
spoof$_{slow}$	**0.92**	0.67	0.75	141
relay	**0.83**	0.58	0.75	56
tamper	**1.00**	**1.00**	**1.00**	5
tamper$_{slow}$	**0.58**	0.25	0.42	162
jam	**1.00**	**1.00**	**1.00**	5

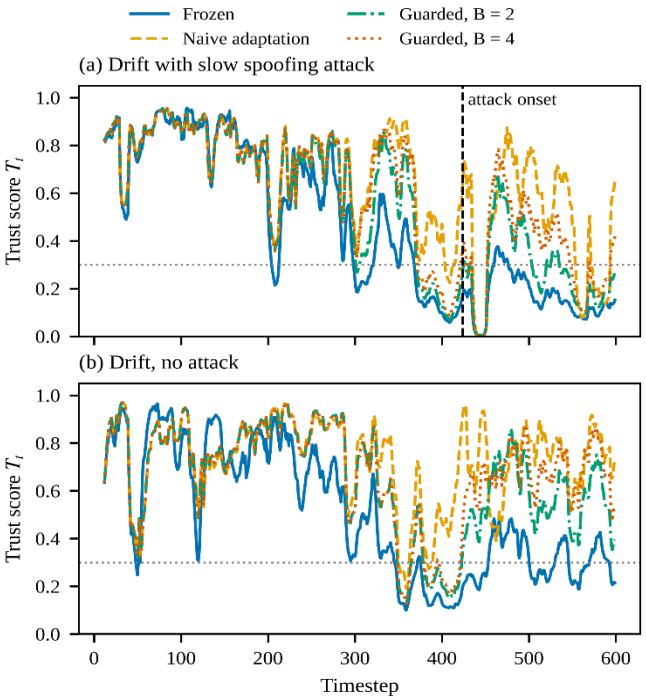

Fig. 4. Trust under the four re-calibration policies in a drifting environment, with a slow spoofing attack (top) and without an attack (bottom).

VI. DISCUSSION

A. *What the Evidence Supports*

Within the simulator, a compact autoencoder on window-level CSI statistics detects the four attack classes well when attacks are abrupt and strong, and short-lag channel autocorrelation adds a small, consistent accuracy gain. Three results limit what can be claimed. First, the autoencoder is no better than simple Gaussian detectors. Second, detection depends strongly on attack strength, relay geometry, hardware calibration and benign relocation. Third, over long sessions, legitimate drift triggers false restrictions in most sessions and no re-calibration policy cleanly separates drift from slow attacks. The accurate summary is that a PHY-derived score is a weak, noisy telemetry input, comparable to other behavioral signals NIST describes a ZT policy engine consuming [1], and that its usefulness depends on per-site calibration and on policies that combine it with identity and posture signals.

B. *Relation to Prior Work*

The score complements Lei et al. [5]: their framework establishes trust at connection time, whereas this score could inform the PDP for the remainder of a session and could in principle trigger re-authentication. It shares a methodological lineage with Dinh-Van et al. [6] but covers full-session behavior and a different application. Unlike authentication-oriented PLA [12]–[18], it scores link behavior without a claimed identity. Work on sustained spoofing and non-stationary channels [17], [18] addresses the drift-versus-attack problem we find unsolved for this pipeline, and borrowing its channel-prediction or sequential-decision ideas is the natural next step.

C. *The Relay/MITM Gap*

Relay/MITM is the weakest class: 0.69 recall at the primary setting, 0.57 at lower attack strength, and near zero for a relay with its own geometry. A relay inherits a real channel, so channel statistics alone cannot be expected to expose it; timing-based techniques (round-trip time, clock skew) target

a different property and should be combined with PHY scores. A security architecture that relied solely on this signal would under-detect relays, so the PDP must treat it as one input among several.

D. Limitations

- **Simulation only:** All data are synthetic, from a clustered-multipath model that is not 3GPP-compliant and not ray-traced. Hardware impairments, antenna coupling and real beam management are only partly represented (Section V-F), so reported accuracy should not be assumed to transfer to a deployment. Validation on measured data is the main open step.
- **Re-implemented simulator and Tier-2 model:** The revision's experiments were rerun on a re-implemented simulator with an MLP Tier 2; absolute values differ from the first submission.
- **Adversary model:** Statistics-matched forgers and slow poisoning were evaluated; gradient-based evasion of the learned models was not, although such attacks are effective against learned RF classifiers [23]–[26].
- **Policy layer:** The PDP of Section III-E is a probe, not a tuned policy; drift handling is unsolved here and the operating points were not optimized per site.
- **Latency:** was not re-measured, so no latency claim is made.

E. Deployment Considerations and Future Work

Intermittent backhaul may delay Tier-2 re-validation, edge compute is constrained, and integration with existing supervisory control and data acquisition (SCADA) and industrial control system (ICS) tooling is the realistic adoption path. Future work includes validation on measured massive-MIMO or mmWave data and ray-traced channels, per-site calibration and drift detectors tied to explicit re-attestation, adversarially trained detectors, combination with timing-based relay detection, and richer policies that fuse the score with identity and posture signals.

VII. Conclusion

We evaluated whether CSI and beamspace statistics can provide a continuous behavioral score for a Zero-Trust policy engine in oilfield IIoT networks. With a leakage-free protocol, a two-tier pipeline with a formally defined trust score reaches an AUC of 0.971 on abrupt attacks, but a Mahalanobis-distance detector matches it, angular features add nothing once temporal self-similarity features are included, relay/MITM recall is 0.69, and stress tests and long-session simulations expose sensitivity to relay geometry, forgers, hardware impairments, benign relocation and drift. The score is therefore a useful but weak input to a ZT policy engine, not a stand-alone trigger. Validation on measured data, drift-aware calibration and adversarial evaluation are the necessary next steps.

Acknowledgment

The author thanks Daw Alfada Company for supporting this research.

Conflict of Interest

The author declares that there is no conflict of interest relevant to this article.

References

[1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero Trust Architecture,” NIST Special Publication 800-207, Aug. 2020, doi: 10.6028/NIST.SP.800-207

[2] S. Mushtaq, M. Mohsin, and M. M. Mushtaq, “A systematic literature review on the implementation and challenges of zero trust architecture across domains,” Sensors, vol. 25, no. 19, art. 6118, 2025, doi: 10.3390/s25196118.

[3] C.Liu, R. Tan, Y. Wu, Y. Feng, Z. Jin, F. Zhang, and Q. Liu, “Dissecting zero trust: research landscape and its implementation in IoT,” Cybersecurity, vol. 7, art. 20, 2024, doi: 10.1186/s42400-024-00212-0.

[4] H. Wang, F. Lü, Y. Cheng, S. Lü, D. Sun, and L. Sun, “Review of zero trust security research in industrial Internet of Things,” J. Comput. Res. Dev., vol. 62, no. 11, pp. 2784–2805, 2025, doi: 10.7544/issn1000-1239.202440840.

[5] W. Lei, Z. Pang, H. Wen, W. Hou, and W. Li, “Physical layer enhanced zero-trust security for wireless industrial Internet of Things,” IEEE Trans. Ind. Informat., 2023, doi: 10.1109/TII.2023.3321106

[6] S. Dinh-Van, T. M. Hoang, B. B. Cebecioglu, D. S. Fowler, Y. K. Mo, and M. D. Higgins, “A defensive strategy against beam training attack in 5G mmWave networks for manufacturing,” IEEE Trans. Inf. Forensics Security, vol. 18, pp. 2204–2217, 2023, doi: 10.1109/TIFS.2023.3265341

[7] T. Wang, D. Yang, S. Zhang, Y. Wu, and S. Xu, “Wi-Alarm: Low-cost passive intrusion detection using WiFi,” Sensors, vol. 19, no. 10, art. 2335, 2019, doi: 10.3390/s19102335

[8] J. Xiao, K. Wu, Y. Yi, L. Wang, and L. M. Ni, “Pilot: Passive device-free indoor localization using channel state information,” in Proc. IEEE ICDCS, Philadelphia, PA, USA, 2013, pp. 236–245, doi: 10.1109/ICDCS.2013.49

[9] X. Dang, C. Li, Z. Hao, and Y. Cao, “A CSI-based indoor intrusion detection and localization method,” in Wireless Sensor Networks (CWSN 2019), Springer, 2019, doi: 10.1007/978-981-15-1785-3_24

[10] E. Dehmollaian, B. Etzlinger, N. Ballber Torres, and A. Springer, “Using channel state information for physical tamper attack detection in OFDM systems: A deep learning approach,” IEEE Wireless Commun. Lett., vol. 10, no. 7, pp. 1503–1507, 2021, doi: 10.1109/LWC.2021.3072937

[11] M. S. Aljumaily and S. Abdullah, “Deep learning for enhanced anomaly detection in wireless communication networks using channel state information (CSI),” Int. J. Mechatronics, Robot. Artif. Intell., vol. 1, no. 2, pp. 104–113, 2025, doi: 10.33971/ijmrai.1.2.13.

[12] N. Gao, Q. Huang, C. Li, S. Jin, and M. Matthaiou, “EsaNet: Environment semantics enabled physical layer authentication,” IEEE Wireless Commun. Lett., 2023, doi: 10.1109/LWC.2023.3324981.

[13] T. M. Pham, L. Signialesi, M. Baldi, G. P. Fettweis, and A. Chorti, “Machine learning-based robust physical layer authentication using angle of arrival estimation,” in *Proc.* IEEE GLOBECOM, 2023, doi: 10.1109/GLOBECOM54140.2023.10437915.

[14] M. Abdrabou and T. A. Gulliver, “Secure authentication in MIMO systems: exploring physical limits,” Front. Commun. Netw., 2024, doi: 10.3389/frcmn.2024.1370496.

[15] Y. Liu, P. Zhang, Y. Shen, L. Peng, and X Jiang, “Online machine learning-based physical layer authentication for mmWave MIMO systems,” Ad Hoc Networks, vol. 131, art. 102864, 2022, doi: 10.1016/j.adhoc.2022.102864

[16] Y. Teng, H. Shi, P. Zhang, J. Dong, and F. Xiao, “PHY-layer authentication exploiting CFO for smart healthcare systems with mmWave communication technology,” Ad Hoc Networks, vol. 140, art. 103075, 2023, doi: 10.1016/j.adhoc.2022.103075

[17] Y. Guo, J. Zhang, and Y.-W. P. Hong, “Channel prediction-based physical layer authentication under consecutive spoofing attacks,” in Proc. IEEE ICC, Glasgow, U.K., 2026, pp. 1–6, doi: 10.1109/ICC59461.2026.11588016

[18] A. Khandan Boroujeni, G. Bagheri, K. R. R. Ranasinghe, G. T. Freitas de Abreu, S. Köpsell, and R. F. Schaefer, “Temporal disambiguation physical layer authentication for non-stationary MIMO channels,” in Proc. IEEE PIMRC, 2026, pp. 1–6, doi: 10.1109/PIMRC69400.2026.11711840

[19] N. Hoque, and H. Rahbari, “Countering relay and spoofing attacks in the connection establishment phase of Wi-Fi systems,” in Proc. ACM WiSec, 2023, doi: 10.1145/3558482.3590185

[20] Y. Tung, K. G. Shin, and K. Kim, “Analog man-in-the-middle attack against link-based packet source identification,” in Proc. ACM MobiHoc, 2016, doi: 10.1145/2942358.2942361

[21] A. Jagannath, J. Jagannath, and P. S. Pattanshetty Vasanth Kumar, “A comprehensive survey on radio frequency (RF) fingerprinting: Traditional approaches, deep learning, and open challenges,” Computer Networks, vol. 219, art. 109455, 2022, doi: 10.1016/j.comnet.2022.109455

[22] E. Lomba, R. Severino, and A. Fernández Vilas, “Work in progress: Towards adaptive RF fingerprint-based authentication of IIoT devices,” in Proc. IEEE ETFA, Stuttgart, Germany, 2022, doi: 10.1109/ETFA52439.2022.9921575

[23] J. Ma, J. Zhang, G. Shen, A. Marshall, and C.-H. Chang, “White-box adversarial attacks on deep learning-based radio frequency fingerprint identification,” in *Proc. IEEE ICC*, 2023, doi: 10.1109/ICC45041.2023.10278927.

[24] J. Ma, J. Zhang, G. Shen, A. Marshall, and C.-H. Chang, “Adversarial attacks against deep learning-based radio frequency fingerprint identification,” arXiv:2512.12002, 2025 (preprint), doi: 10.48550/arXiv.2512.12002

[25] X. Cao, B. Adhikari, S. Zhao, J. Wu and Y. Pan, "An Adversarial-Driven Experimental Study on Deep Learning for RF Fingerprinting," MILCOM 2025 - 2025 IEEE Military Communications Conference (MILCOM), Los Angeles, CA, USA, 2025, pp. 739-744, doi: 10.1109/MILCOM64451.2025.11309917.

[26] Y. Huang, W. Liu, and H. Wang., “Hidden backdoor attack: A new threat to learning-aided physical layer authentication,” in Proc. Int. Conf. Ubiquitous Commun. (Ucom), 2023, doi: 10.1109/Ucom59132.2023.10257584.

[27] X. Feng and S. Hu, “Cyber-physical zero trust architecture for industrial cyber-physical systems,” IEEE Trans. Ind. Cyber-Phys. Syst., vol. 1, pp. 394–405, 2023, doi: 10.1109/TICPS.2023.3333850

[28] H. Sedjelmaci and N. Ansari, “Zero trust architecture empowered attack detection framework to secure 6G edge computing,” IEEE Network, vol. 38, no. 1, pp. 196–202, 2024, doi: 10.1109/MNET.131.2200513

[29] S. Bagherinejad and S. M. Razavizadeh, “Direction-based jamming detection and suppression in mmWave massive MIMO networks,” IET Commun., vol. 15, no. 14, pp. 1780–1790, 2021.doi: 10.1049/cmu2.12190

[30] S. B. H. Natanzi, H. Mohammadi, B. Tang and V. Marojevic, "Secure mmWave Beamforming with Proactive-ISAC Defense Against Beam-Stealing Attacks," 2026 IEEE 23rd Consumer Communications & Networking Conference (CCNC), Las Vegas, NV, USA, 2026, pp. 1-4, doi: 10.1109/CCNC65079.2026.11366415.

[31] A. Palathumveettil Jagadeesan, I. Yilmaz, A. Palaios, G. T. Nguyen and F. H. P. Fitzek, "First Steps in Concept Drift Management for Resilient Wireless Networks," European Wireless 2024; 29th European Wireless Conference, Brno, Czech Republic, 2024, pp. 85-90.

[32] A. Tziouvaras, C. Fortuna, G. Floros, K. Kolomvatsos, P. Sarigiannidis, Marko Grobelnik, and Blaž Bertalanič., “Towards reliable AI in 6G: Detecting concept drift in wireless network,” arXiv:2508.00042, 2025 (preprint).

[33] D. Li, N.-K. C. Nair, and K. I.-K. Wang, “ADDAEIL: Anomaly detection with drift-aware ensemble-based incremental learning,” Algorithms, vol. 18, no. 6, art. 359, 2025, doi: 10.3390/a18060359.

[34] R. Gupta, B. Tian, Y. Wang, and K. Nahrstedt., “TWIN-ADAPT: Continuous learning for digital twin-enabled online anomaly classification in IoT-driven smart labs,” Future Internet, vol. 16, no. 7, art. 239, 2024, doi: 10.3390/fi16070239.

[35] T. Jiang, J. Zhang, P. Tang, L. Tian, Y. Zheng, J. Dou, H. Asplund, L. Raschkowski, R. D’Errico, and T. Jämsä, “3GPP standardized 5G channel model for IIoT scenarios: A survey,” IEEE Internet Things J., vol. 8, no. 11, pp. 8799–8815, 2021, doi: 10.1109/JIOT.2020.3048992.