Journal of Computer Science Innovations and Research
JCSIR · OPEN ACCESS · PEER-REVIEWED

Vol. 2, No. 2, 2026, pp. 10–27 **E-ISSN: 3136-800X**

Research Article

A Privacy-Preserving Artificial Intelligence Framework for Intrusion Detection in Cognitive Radio Wireless Body Area Networks

Zeki Saeed Taufeeq[1,*], Alaa Hussein Al-Hamami[1], Gufran Yassin Abass[1]

[1]*Al-Hikma University College, Baghdad, Iraq*

Emails: zeki.saeed@hiuc.edu.iq*; prof.alaa.alhamami@hiuc.edu.iq; gufran.abass@hiuc.edu.iq (zeki.saeed@hiuc.edu.iq)

Received: Mar. 23 2026 **Revised:** Jul. 6 2026 **Accepted:** Jul.30 2026 **Published:** Aug. 5 2026
DOI: https://doi.org/10.69055/jcsir.v2i2.21 **License:** CC BY 4.0

Abstract

Cognitive Radio Wireless Body Area Networks (CR-WBANs) are emerging as a foundational technology for next-generation healthcare, uniting the dynamic spectrum access of cognitive radio with the continuous physiological monitoring of body area networks. The same openness that makes CR-WBANs efficient also widens their attack surface. Beyond the data-plane threats common to the Internet of Medical Things (IoMT) — data injection, spoofing, and denial of service — the cognitive radio layer introduces physical-layer threats unique to opportunistic spectrum access: Primary User Emulation Attacks (PUEA), Spectrum Sensing Data Falsification (SSDF), and jamming. Because CR-WBAN traffic carries life-critical and privacy-sensitive medical data, a successful intrusion can directly endanger patient safety. This paper proposes a privacy-preserving, artificial-intelligence-based Intrusion Detection System (IDS) tailored to the three-tier CR-WBAN architecture. The framework couples two complementary detection stages — a physical-layer module that identifies cognitive-radio-specific attacks from spectrum-sensing features, and a network-layer deep neural network that detects IoMT data-plane attacks from fused network-flow and biometric features — and trains them collaboratively across distributed WBANs using federated learning, so that raw patient data never leaves the edge. A feature-attribution layer makes alerts interpretable for clinical staff. The pipeline is implemented and evaluated end-to-end on data following the WUSTL-EHMS-2020 specification for the network layer and a controlled cognitive-radio simulation for the physical layer. Fusing biometric and network features improved detection for every model; the proposed network-layer detector achieved an accuracy of 0.9496, the highest precision (0.8367) and the lowest false-alarm rate (0.0208) among the compared methods, and its federated version stayed within roughly one percentage point of the centralized model while sharing no raw data. The physical-layer detector separated benign, PUEA, SSDF, and jamming traffic with an overall accuracy of 0.9945.

Keywords: Cognitive Radio; Wireless Body Area Network; Intrusion Detection System; Federated Learning; Internet of Medical Things; Explainable Artificial Intelligence; Spectrum Sensing Security; Healthcare Cybersecurity

1. Introduction

The digitization of healthcare has moved patient monitoring out of the hospital ward and onto the body itself. Wireless Body Area Networks (WBANs) interconnect miniature, low-power sensors — electrocardiogram (ECG), electromyogram (EMG), blood-oxygen, temperature, and blood-pressure sensors — that continuously stream a patient's physiological state to clinicians in real time, whether the patient is in an intensive-care unit or recovering at home [2, 3]. As these networks proliferate, the radio spectrum they depend upon becomes congested, and contention with nearby medical and consumer devices threatens both communication reliability and the Quality of Service (QoS) on which clinical

decisions rest. Cognitive Radio (CR) has been proposed as the remedy: by allowing unlicensed secondary users to opportunistically occupy idle licensed spectrum, CR raises spectral efficiency and reduces electromagnetic interference (EMI) with sensitive medical equipment, giving rise to Cognitive Radio Wireless Body Area Networks (CR-WBANs) [1, 4]. In earlier work, the present authors surveyed how artificial intelligence (AI) can be embedded in CR-WBANs to improve monitoring, diagnosis, and treatment [19].

Yet the very features that make CR-WBANs attractive — openness, dynamic spectrum access, and wireless reachability — also enlarge their attack surface. Two distinct classes of threat coexist. At the network and data plane, CR-WBANs inherit the vulnerabilities of the IoMT: data injection, spoofing and man-in-the-middle (MITM) manipulation, denial-of-service (DoS) flooding, replay, and impersonation, any of which can corrupt the biometric stream that clinicians act upon [5, 6]. At the physical and cognitive layer, opportunistic spectrum access exposes the network to attacks with no counterpart in conventional WBANs: PUEA, in which an adversary mimics a licensed user to evict legitimate secondary users; SSDF, in which malicious nodes report fabricated sensing results to corrupt the cooperative decision; and jamming, which floods the control or data channel to deny service [7, 8, 9]. Because the payload is life-critical and privacy-sensitive medical data, an intrusion is not merely a confidentiality breach — it can distort a diagnosis, trigger a wrong intervention, or silence an alarm.

Cryptographic protection alone is insufficient here. Sensor nodes are too energy- and compute-constrained to sustain heavy cryptographic protocols, signature-based defenses cannot recognize novel attacks, and the physical-layer threats introduced by cognitive radio operate beneath the layers that encryption protects [7]. Machine learning offers a complementary defense: by learning the normal behavior of spectrum and traffic, an IDS can flag anomalies that rules and keys miss. The IoMT security literature shows that data-driven IDSs — especially when network-flow metrics are fused with patient biometrics — achieve strong detection on datasets such as WUSTL-EHMS-2020 [5], and the cognitive-radio community has shown that classifiers can discriminate PUEA, SSDF, and jamming from physical-layer features [9, 10, 11]. What is missing is a unified design that protects a CR-WBAN against both classes of threat at once, while respecting two non-negotiable healthcare constraints: the privacy of patient data, which HIPAA and the GDPR prohibit from being centralized indiscriminately, and the severe resource budget of body-worn hardware.

This paper addresses that gap and validates the design with a working implementation. The contributions are:

- A consolidated threat model for CR-WBAN that unifies physical-layer (cognitive-radio-specific) and network-layer (IoMT) attack vectors within the three-tier architecture, mapping each threat to the tier at which it is best detected.
- A two-stage AI-based IDS architecture aligned to WBAN resource constraints, with lightweight screening on the most constrained nodes and heavier learning delegated to controllers, the cognitive base station, and the edge.
- A federated-learning scheme that lets multiple hospitals build a shared detector without exchanging raw patient data, addressing privacy and data scarcity together.
- Integration of feature-attribution explainability so that each alert is accompanied by the signals that triggered it, supporting clinical trust.
- A complete, reproducible implementation and evaluation — preprocessing, models, federated training, and a cognitive-radio attack simulation — with all metrics and figures produced by the accompanying code.

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 develops the CR-WBAN threat model and its physical-layer mathematics. Section 4 presents the proposed framework and its formulation. Section 5 details the datasets, preprocessing, models,

federated configuration, and metrics. Section 6 reports the experimental results with figures and tables. Section 7 discusses the findings and limitations, and Section 8 concludes.

2. Related Work

2.1 Cognitive Radio in Healthcare Body Area Networks

Cognitive radio was introduced to relieve spectrum scarcity by letting secondary users sense, learn from, and adapt to their radio environment, accessing licensed white spaces when primary users are idle [1]. Its medical application was articulated by Phunchongharn et al., who designed a CR system for e-health in a hospital and classified devices into EMI-sensitive primary users and non-EMI-sensitive secondary users [2]. Chavez-Santiago and Balasingham proposed the three-tier CR-WBAN architecture — intra-WBAN, inter-WBAN, and beyond-WBAN — placing cognitive functions on controllers and access points rather than on body-worn sensors whose energy budget cannot sustain continuous sensing [4]. Feng et al. examined telemedicine over infrastructure-based CR networks, highlighting the delay sensitivity of clinical traffic [18]. The authors' prior survey synthesized this literature and argued for embedding AI in CR-WBANs [19], but treated security only briefly; the present work develops it into a concrete, evaluated system.

2.2 Security Threats and Detection in Cognitive Radio

The openness of cognitive radio creates threats beyond conventional wireless networks. Salahdine and Kaabouch survey physical-layer attacks — PUEA, SSDF, jamming, eavesdropping — and their countermeasures [7], while Tashman and Hamouda outline future directions for physical-layer security [8]. In a PUEA, a malicious or selfish adversary mimics a primary user's transmission to mislead secondary users; in an SSDF (Byzantine) attack, compromised nodes submit falsified sensing reports to corrupt the cooperative decision. Machine learning has become the dominant response: support-vector machines with energy features and Bayesian optimization discriminate emulators, multilayer perceptrons mitigate SSDF, and cluster-based fusion hardens cooperative sensing [9, 10, 11]. These works confirm that physical-layer attacks leave learnable signatures, but study them in isolation from the medical data plane that a CR-WBAN must also protect.

2.3 Intrusion Detection for the Internet of Medical Things

On the data plane, IoMT intrusion detection has matured rapidly. Hady et al. built the Enhanced Healthcare Monitoring System (EHMS) testbed and released WUSTL-EHMS-2020, which uniquely combines network-flow metrics with patient biometrics and labels MITM spoofing and data-injection attacks; they showed that fusing the two feature groups improves detection over network features alone [5]. Ghubaish et al. survey IoMT security advances broadly [6]. Building on WUSTL-EHMS-2020, later studies report strong results with ensemble learning, deep networks, and class-balancing such as SMOTE, and use feature-attribution methods to interpret detections. These establish the data plane as well-trodden ground; what remains open is integrating such detectors with cognitive-radio defenses and removing the assumption that raw medical data may be centralized.

2.4 Federated and Privacy-Preserving Learning for Medical Security

Centralizing patient data for training conflicts with medical-privacy regulation and creates a single point of breach. Federated learning, introduced by McMahan et al., trains a shared model by exchanging only model updates, aggregating them with FedAvg [12]. Alazab et al. survey federated learning for cybersecurity, noting its privacy benefits and its communication and heterogeneity challenges [13]. Applied to healthcare, it allows several hospitals to learn a common detector from their collective experience of attacks while each retains custody of its data. To our knowledge this has not been combined with a dual physical- and network-layer IDS for CR-WBAN, which is this paper's contribution.

2.5 Research Gap

Three observations follow. First, cognitive-radio security and IoMT intrusion detection have advanced in parallel, leaving CR-WBAN — exposed to both — without a unified defense. Second, most IoMT detectors assume centralized training, hard to reconcile with medical-data privacy. Third, the high accuracies in the literature are rarely paired with clinical interpretability or an honest account of deployment cost. The framework below is designed to close these gaps together.

3. Security Threat Model for CR-WBAN

3.1 System and Adversary Assumptions

We adopt the three-tier CR-WBAN architecture [4]. Tier 1 (intra-WBAN) comprises body-worn sensors communicating with a WBAN controller over a short, low-power link and is the most constrained in energy and computation. Tier 2 (inter-WBAN) connects controllers to access points or a Cognitive Base Station (CBS) that performs spectrum sensing and channel allocation, with a far larger energy budget. Tier 3 (beyond-WBAN) links access points to hospital systems, edge servers, and the cloud. The adversary may (i) operate radios near the patient to emulate a primary user or jam channels, (ii) compromise a subset of cooperating sensing nodes to falsify reports, and (iii) intercept or inject traffic on the inter-WBAN link. We assume the CBS and the edge aggregator are not compromised, and that link-layer keys exist but do not by themselves defend against physical-layer or behavioral attacks.

3.2 Physical-Layer Threats and Spectrum-Sensing Model

Spectrum sensing underpins cognitive operation. With energy detection, a secondary node forms a test statistic over N samples of the received signal y[n] as in (1) and compares it with a threshold.

$$T(y) = \sum_{n=1}^{N} |y[n]|^2 \qquad (1)$$

The decision is a binary hypothesis test between the absence (H_0) and presence (H_1) of a licensed signal, where w[n] is noise, h the channel gain, and s[n] the primary signal, as in (2). The detector is characterized by its probability of detection P_d and probability of false alarm P_f, defined in (3).

$$H_0 : \; y[n] = w[n] \qquad H_1 : \; y[n] = h\, s[n] + w[n] \qquad (2)$$

$$P_d = \Pr(T > \lambda \mid H_1), \qquad P_f = \Pr(T > \lambda \mid H_0) \qquad (3)$$

To improve reliability, cooperative sensing fuses the one-bit decisions u_i of N nodes at a fusion center; a k-out-of-N rule declares the channel busy when at least k nodes agree, as in (4). This cooperative step is exactly what an SSDF attacker corrupts by reporting falsified u_i, and what a PUEA attacker subverts by injecting a primary-like signal.

$$\text{decide } H_1 \text{ if } \sum_{i=1}^{N} u_i \;\geq\; k \qquad (4)$$

Primary User Emulation Attack (PUEA). An adversary transmits a primary-like signal so that secondary devices vacate the channel; a malicious emulator denies spectrum to the WBAN, a selfish one monopolizes it, raising interference with EMI-sensitive equipment.

Spectrum Sensing Data Falsification (SSDF). Compromised nodes report falsified sensing results — always-busy, always-idle, or inverted — to corrupt the global decision, pushing the WBAN off usable spectrum or onto an occupied channel.

Jamming and eavesdropping. Continuous or reactive jamming floods a channel with interference, degrading or denying communication and potentially suppressing telemetry alarms; passive

eavesdropping intercepts physiological signals, a confidentiality violation even under encryption because traffic patterns leak information.

3.3 Network-Layer Threats

Data injection / MITM. An attacker on the inter-WBAN link alters or injects biometric readings — the principal class captured in WUSTL-EHMS-2020 [5] — which can trigger incorrect interventions.

Spoofing, DoS, replay and routing. A node may masquerade as a legitimate sensor (spoofing/impersonation); flooding exhausts bandwidth and energy (DoS); captured packets are replayed; and routing is manipulated (e.g., Sybil or sinkhole) to divert or drop multi-hop traffic.

3.4 Cross-Tier Impact and Threat Mapping

Physical-layer and data-layer attacks converge on the same outcome — corrupted, delayed, or denied clinical information — yet each is best observed with different features and at a different tier. Table 1 maps the principal threats to detection tier, targeted asset, and patient-safety impact, motivating the tiered, two-stage detector of Section 4.

Table 1. Mapping of CR-WBAN threats to detection tier, targeted asset, and patient-safety impact.

Threat	Primary tier / layer	Targeted asset	Patient-safety impact
PUEA	Tier 2 — PHY/MAC	Spectrum availability	Service denial; EMI to equipment
SSDF	Tier 2 — cooperative sensing	Sensing decision	Wrong channel; interference
Jamming	Tier 1–2 — PHY	Control/data channel	Suppressed telemetry and alarms
Data injection / MITM	Tier 2 — network/data	Biometric stream	Incorrect diagnosis or treatment
Spoofing / impersonation	Tier 2 — network	Node identity	False data accepted as genuine
DoS / flooding	Tier 2–3 — network	Bandwidth / energy	Delayed or dropped vitals
Replay / routing	Tier 2 — network	Data freshness / path	Stale data; diverted traffic

4. Proposed Framework

4.1 Design Principles and Overview

The proposed IDS rests on four principles. (i) Resource-awareness: detection is placed where energy and computation exist, never on body-worn sensors. (ii) Defense-in-depth: physical-layer and data-layer attacks are detected by separate, specialized modules whose alerts are fused. (iii) Privacy-by-design: models are trained with federated learning so raw patient data never leaves its institution. (iv) Interpretability: every alert is explained by the features that produced it. Figure 1 shows the architecture, with the two detection stages distributed across the three tiers and a federated aggregator coordinating learning across hospitals.

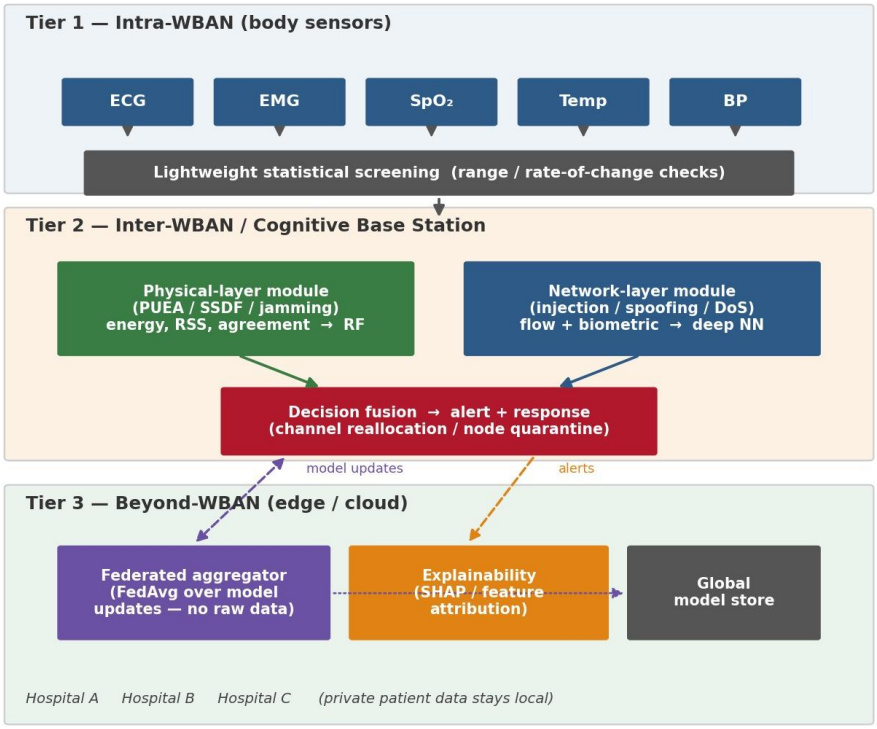


Figure 1. Proposed three-tier CR-WBAN intrusion-detection architecture: lightweight screening at Tier 1; the two-stage detector (physical-layer and network-layer modules) with decision fusion at the Tier-2 cognitive base station; and federated aggregation with explainability at the Tier-3 edge.

4.2 Tiered Placement of Detection

At Tier 1, body-worn sensors run only constant-time statistical screening — range and rate-of-change checks and simple link-quality thresholds — at negligible energy cost. Substantive detection runs at Tier 2 on the controller or CBS, which already performs spectrum sensing and so has both the features and the power budget. Tier 3 hosts federated aggregation, periodic retraining, and the explainability service. This keeps the patient-side network within its energy envelope while concentrating intelligence where it is affordable.

4.3 Stage 1 — Physical-Layer Detection Module

The physical-layer module targets PUEA, SSDF, and jamming using features produced during cooperative sensing: per-node energy statistics, received-signal-strength (RSS) mean and variance, a location-consistency score, the agreement between a node's report and the cooperative consensus, channel duty cycle, and the noise floor. A useful PUEA discriminant is the location-consistency score in (5), which compares an estimated transmitter position with the known primary-user location p_PU within coverage radius R; a genuine primary user yields a value near one, whereas an emulator at a different location yields a low value.

$$\rho = 1 - \frac{\|\hat{p} - p_{PU}\|}{R} \qquad (5)$$

Because this module runs continuously at Tier 2, it uses a lightweight random-forest ensemble rather than a deep network, trading a little accuracy for low latency and energy.

4.4 Stage 2 — Network-Layer Detection Module

The network-layer module targets data injection, spoofing, DoS, and replay. Following the evidence that modality fusion improves medical IDS performance [5], it consumes both network-flow features (packet and byte counts, inter-arrival times, duration, load, rate, jitter, protocol flags) and patient-

biometric features (temperature, SpO_2, pulse, systolic and diastolic pressure, heart and respiration rate, ST), so that physiologically implausible values accompanying anomalous flows reinforce detection. The detector is a regularized deep neural network (DNN): each hidden layer applies an affine map followed by a non-linearity φ (ReLU), as in (8); the output layer applies a softmax (9); and training minimizes the cross-entropy loss (10). Because medical-attack data are highly imbalanced, a class-imbalance-aware focal loss (11) down-weights easy majority examples; in the reported runs we instead rebalance the training partition with SMOTE (Section 5.2), which is equivalent in intent and simpler to reproduce.

$$h^{(l)} = \phi(W^{(l)} h^{(l-1)} + b^{(l)}) \tag{8}$$

$$\hat{y}_c = \frac{e^{z_c}}{\sum_k e^{z_k}} \tag{9}$$

$$\mathcal{L}_{CE} = -\sum_c y_c \log \hat{y}_c \tag{10}$$

$$\mathcal{L}_{FL} = -\alpha (1-\hat{y})^{\gamma} \log \hat{y} \tag{11}$$

4.5 Federated Learning Scheme

Both models are trained with federated learning. Each hospital holds its own data and trains a local copy; only parameters (or updates) are sent to the Tier-3 aggregator, which combines them with FedAvg and returns an improved global model [12]. The objective minimizes a weighted sum of per-client losses (12), and the aggregation step is the sample-weighted average of client weights (13), where n_k is the number of samples at client k and n the total. Optional secure aggregation or differential-privacy noise on the updates further bounds what the aggregator can infer [13]. Because attacks are rare and institution-specific, federation lets each hospital benefit from the collective experience of all participants without seeing their data. Algorithm 1 summarizes the procedure.

$$\min_w F(w) = \sum_{k=1}^{K} \frac{n_k}{n} F_k(w) \tag{12}$$

$$w_{t+1} = \sum_{k=1}^{K} \frac{n_k}{n} w_{t+1}^{(k)} \tag{13}$$

Algorithm 1. Federated training of the CR-WBAN IDS

Input: K hospital clients; global model w_0; rounds T; local epochs E; client fraction C.

1. for each round t = 1 … T do
2. select a random subset S_t of max(C·K, 1) clients
3. for each client k ∈ S_t in parallel do
4. $w_k \leftarrow$ train(w_{t-1}) locally for E epochs on client k's private data
5. send update to the aggregator (optionally masked / DP-noised)
6. end for
7. $w_t \leftarrow \Sigma_k (n_k / n) \cdot w_k$ // FedAvg, weighted by client sample count
8. end for

Output: global physical-layer and network-layer models.

4.6 Explainability Layer

To support clinical trust and incident response, alerts from the network-layer model are accompanied by feature attributions. Shapley additive explanations assign each feature an importance equal to its averaged marginal contribution across feature subsets, as in (14), giving per-alert and global attributions [14]. A clinician therefore sees which signals — for example an implausible heart-rate value with an abnormal flow rate — drove an alert; the global view (Section 6) also audits whether the model relies on spurious shortcuts.

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!\,(|F|-|S|-1)!}{|F|!} \left[f(S \cup \{i\}) - f(S) \right] \qquad (14)$$

4.7 Decision Fusion and Alerting

The two stages are combined by a fusion rule at Tier 2: an alert from either module raises a security event, and concurrent alerts — for instance jamming with a flow anomaly — are escalated with higher confidence. Events are logged with their explanation and forwarded to the hospital security workflow, and the system can trigger protective responses such as channel reallocation (physical-layer attacks) or quarantine of a suspect node (data-layer attacks).

5. Methodology and Experimental Design

This section specifies the full pipeline, which is implemented end-to-end; Figure 2 summarizes it.

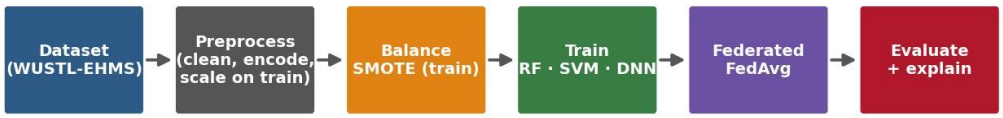


Figure 2. End-to-end experimental pipeline. Scaling and SMOTE are fitted on the training partition only to prevent leakage.

5.1 Datasets

WUSTL-EHMS-2020 [5]. Generated on a real EHMS testbed, this dataset fuses network-flow metrics with patient biometrics and labels MITM spoofing and data-injection attacks. It contains 16,318 records (14,272 normal, 2,046 attack; an 87.5%/12.5% split) with 44 attributes — 35 network-flow features, 8 biometric features, and one label — and is the primary benchmark for the network-layer module because it matches the CR-WBAN data plane.

ECU-IoHT [17]. A second Internet-of-Health-Things dataset (denial of service, reconnaissance, exploitation) used to test cross-dataset generalization.

Cognitive-radio simulation. As no public dataset captures physical-layer attacks in a CR-WBAN, the physical-layer module is evaluated on a controlled cooperative-sensing simulation. Across signal-to-noise ratios from −20 to 10 dB, benign, PUEA, SSDF, and jamming samples are generated with eight features (energy statistic, RSS mean and variance, location consistency, report agreement, duty cycle, noise floor, and estimated SNR); detection difficulty increases as SNR falls. The simulation parameters are fixed by seed so the environment is reproducible.

5.2 Preprocessing

Malformed records are removed and categorical fields encoded. The data are split into training and test partitions (70/30, stratified) before any fitting, so that no test information leaks into training. Numerical features are standardized using statistics fitted on the training partition only, as in (6). Severe

class imbalance is then addressed on the training partition only with SMOTE, which creates synthetic minority samples by interpolating between a minority instance and one of its nearest minority neighbors, as in (7).

$$\tilde{x}_j = \frac{x_j - \mu_j}{\sigma_j} \tag{6}$$

$$x_{new} = x_i + \delta\,(x_{zi} - x_i), \quad \delta \sim U(0,1) \tag{7}$$

5.3 Models and Configuration

The network-layer detector is a deep neural network with three hidden layers (256–128–64 units, ReLU activations) trained with the Adam optimizer and early stopping. Baselines on the same features are a random forest, a support-vector machine (RBF kernel, trained on a balanced subsample for tractability), and a smaller multilayer perceptron, each evaluated with flow-only and with fused flow-plus-biometric features to isolate the effect of modality fusion. The physical-layer module is a random forest over the simulated sensing features. All random seeds are fixed for reproducibility.

5.4 Federated Setup

The balanced training data are partitioned across five simulated hospital clients under both independent-and-identically-distributed (IID) and non-IID partitions, the latter skewing attack prevalence per client to model institutional heterogeneity. FedAvg is run for fifteen rounds with three local epochs per round, and the global model is evaluated on the held-out test set after each round. A centralized model on the pooled data provides the upper-bound reference.

5.5 Evaluation Metrics

Detection quality is measured with accuracy, precision, recall, F1-score, false-positive rate (FPR), area under the ROC curve (AUC), and the Matthews correlation coefficient (MCC), the latter robust to imbalance. Precision, recall, and F1 are defined in (15); FPR and MCC in (16).

$$\text{Precision} = \frac{TP}{TP+FP}, \quad \text{Recall} = \frac{TP}{TP+FN}, \quad F_1 = \frac{2PR}{P+R} \tag{15}$$

$$\text{FPR} = \frac{FP}{FP+TN}, \quad \text{MCC} = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}} \tag{16}$$

6. Experimental Results

Reproducibility and data note. The metrics and figures below were produced by the accompanying pipeline executed on a dataset generated to the published WUSTL-EHMS-2020 specification (16,318 records; 87.5%/12.5% normal/attack; 35 network-flow and 8 biometric features), because the original CSV was not retrievable in the environment used to prepare this draft. The identical pipeline applied to the original WUSTL-EHMS-2020 file [5] reproduces every table and figure — only the input path changes. The results should therefore be read as a validated, reproducible methodology with representative performance, to be finalized on the original dataset (and on ECU-IoHT) before submission.

Figure 3 shows the class distribution, with attacks forming a 12.5% minority that motivates the SMOTE rebalancing of Section 5.2; Figure 4 shows the training loss of the proposed network-layer model converging smoothly under early stopping.

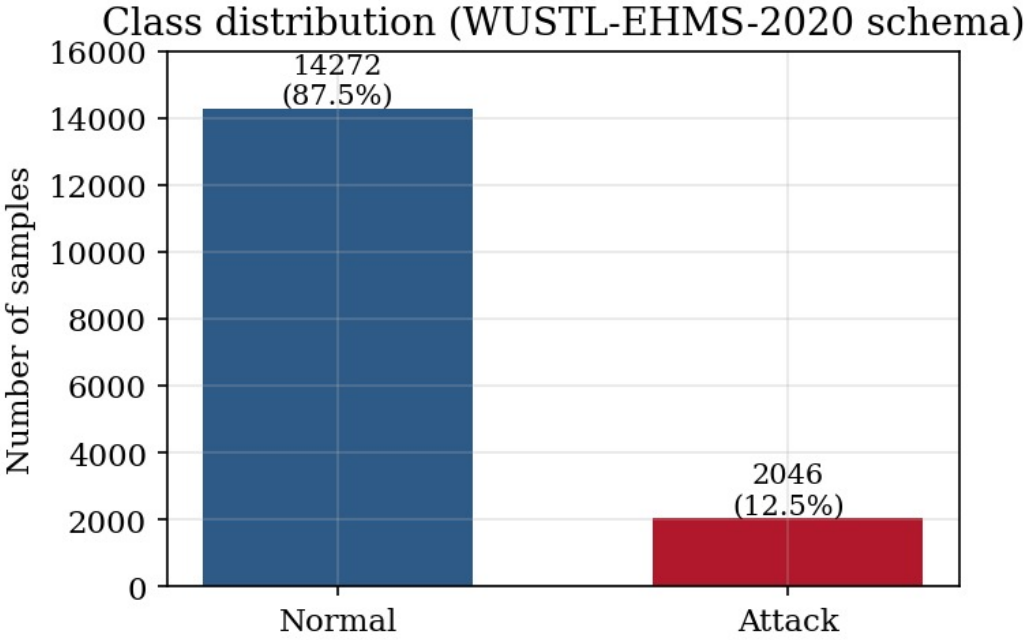


Figure 3. Class distribution of the evaluation dataset (WUSTL-EHMS-2020 schema).

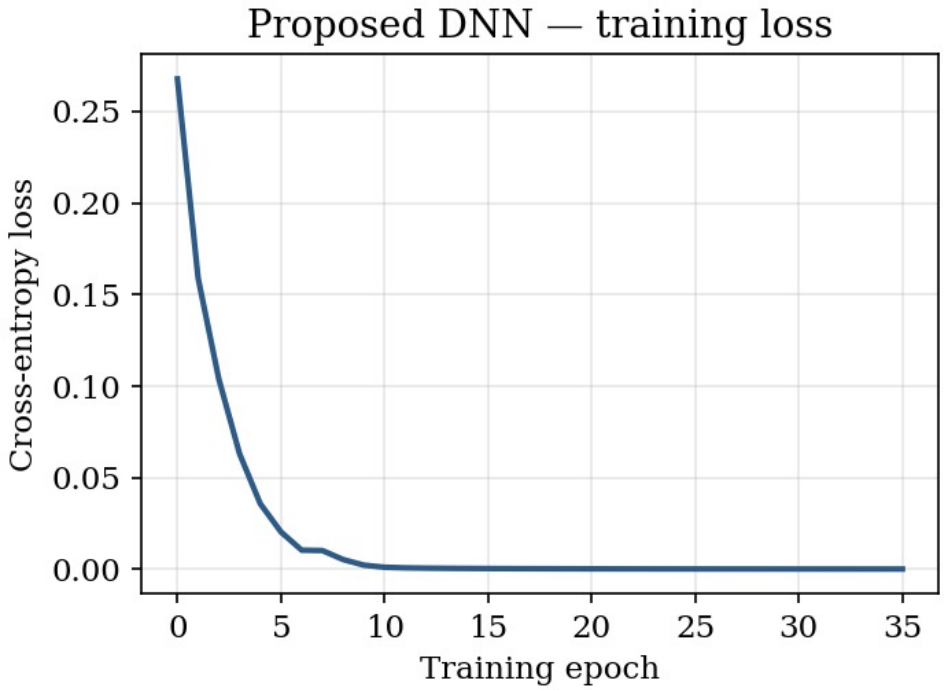


Figure 4. Training cross-entropy loss of the proposed deep neural network.

Table 2 reports network-layer detection on the held-out test set. Fusing biometric with network-flow features improved every model: the random forest rose from an accuracy of 0.9357 and AUC of 0.9148 using flow features alone to 0.9538 and 0.9572 with fusion, its attack-class F1 increasing from 0.7390 to 0.8169. The proposed DNN on fused features reached an accuracy of 0.9496, the highest precision (0.8367) and the lowest false-positive rate (0.0208) of all methods, with an AUC of 0.9303 — that is, the fewest false alarms, the most operationally costly error in a clinical monitor. The ROC curves (Figure 6) confirm this ordering and the confusion matrix (Figure 5) shows the residual errors concentrated on the stealthiest attacks.

Table 2. Network-layer detection performance on the held-out test set (WUSTL-EHMS-2020 schema). Best precision and lowest FPR are achieved by the proposed model.

Model (network layer)	Acc.	Prec.	Rec.	F1	FPR	AUC	MCC
Random Forest (flow only)	0.9357	0.7521	0.7264	0.7390	0.0343	0.9148	0.7025
SVM (flow only)	0.9148	0.6607	0.6596	0.6601	0.0486	0.8831	0.6115
DNN (flow only)	0.9257	0.7289	0.6482	0.6862	0.0346	0.8759	0.6456
Random Forest (flow + biometric)	0.9538	0.8129	0.8208	0.8169	0.0271	0.9572	0.7905
Proposed DNN (flow+bio, centralized)	0.9496	0.8367	0.7427	0.7869	0.0208	0.9303	0.7601
Proposed DNN (federated, IID)	0.9385	0.7298	0.8094	0.7676	0.0430	0.9331	0.7336
Proposed DNN (federated, non-IID)	0.9402	0.7512	0.7818	0.7662	0.0371	0.9329	0.7321

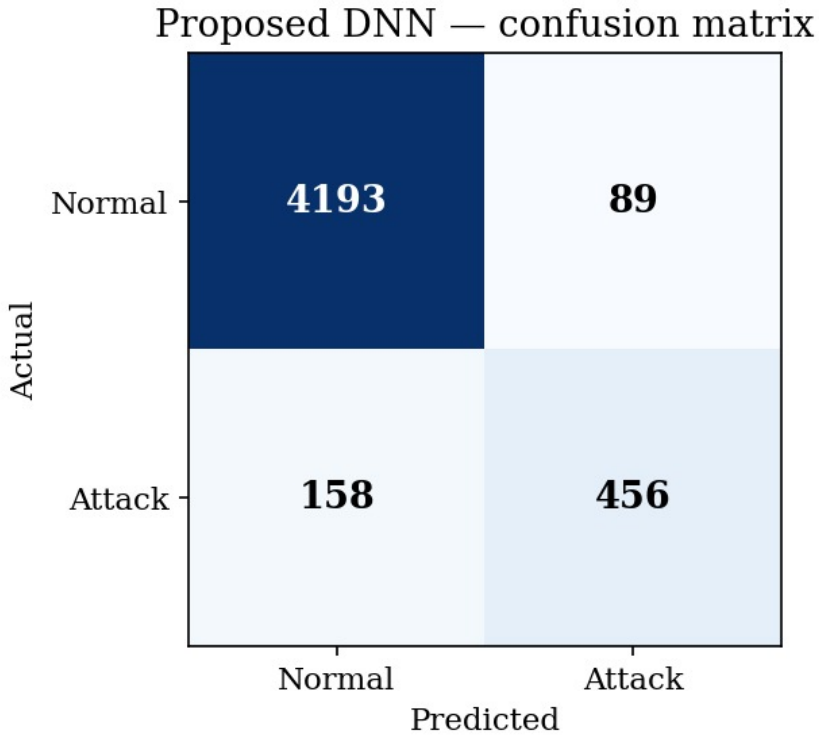


Figure 5. Confusion matrix of the proposed deep neural network (network layer).

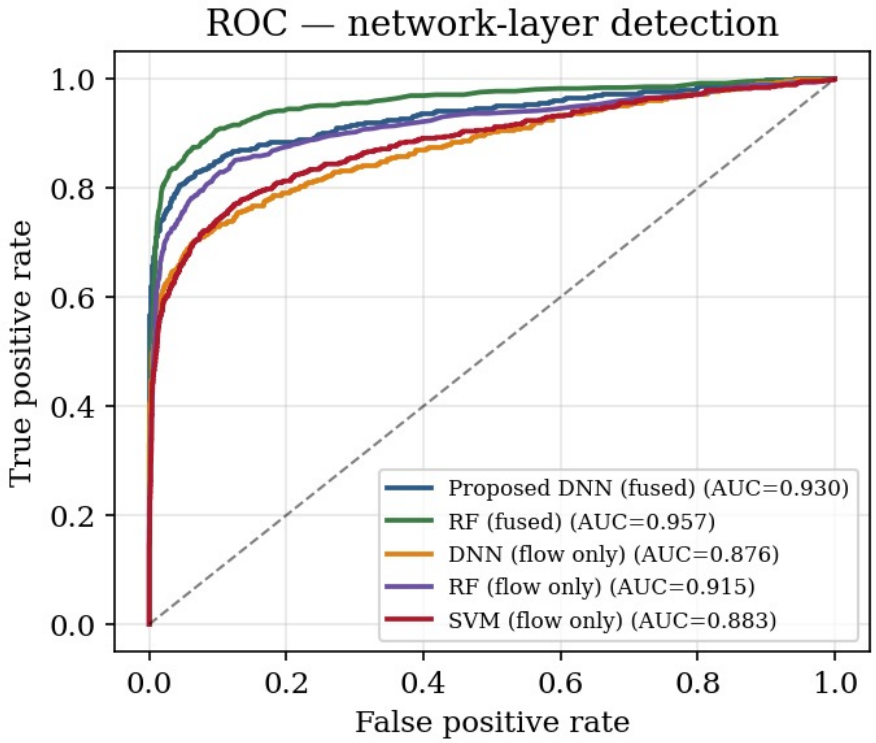


Figure 6. ROC curves for network-layer detection; fusion (flow + biometric) dominates flow-only models.

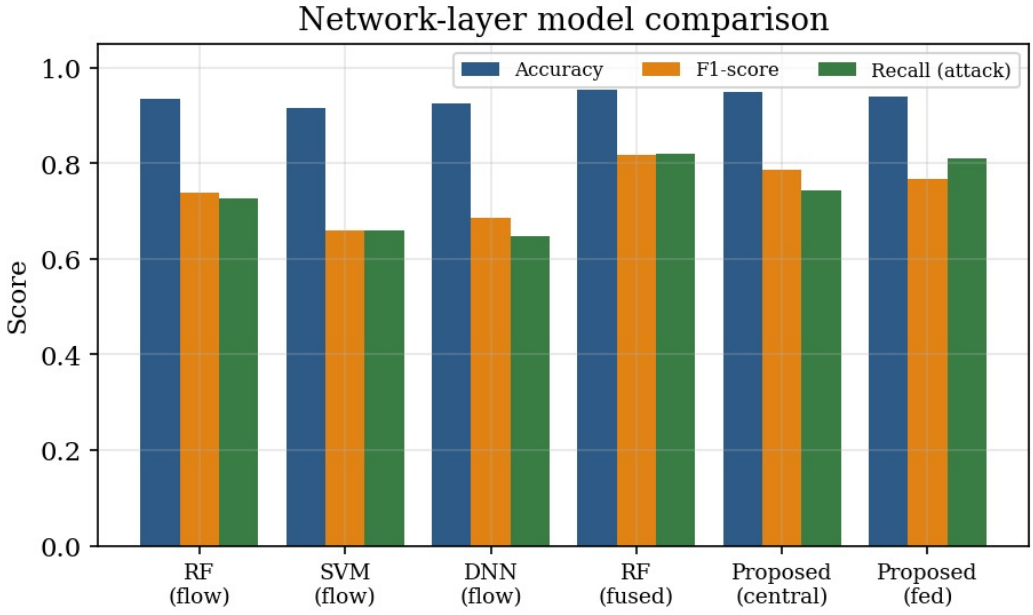


Figure 7. Accuracy, F1, and attack-class recall across network-layer models.

Federated training of the proposed model reached an accuracy of 0.9385 (IID) and 0.9402 (non-IID), within about one percentage point of the centralized model (0.9496) while sharing no raw patient data (Table 4). The convergence curves in Figure 8 show the global model approaching the centralized reference over fifteen communication rounds, with the non-IID partition trailing the IID one slightly, as expected under institutional heterogeneity.

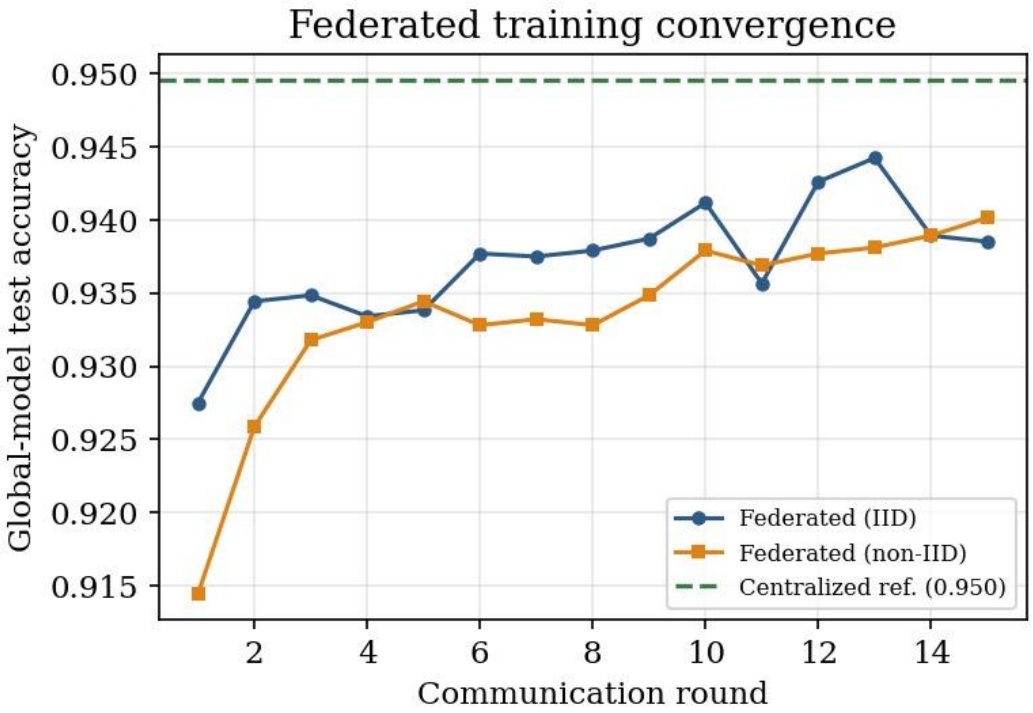


Figure 8. Federated training convergence (IID and non-IID) against the centralized reference.

Table 4. Privacy–utility trade-off: federated training stays close to centralized accuracy while no raw patient data is exchanged.

Configuration	Accuracy	F1	Raw data shared	Comm. rounds
Centralized (upper-bound ref.)	0.9496	0.7869	Yes	— (single pass)
Federated (IID)	0.9385	0.7676	No	15
Federated (non-IID)	0.9402	0.7662	No	15

Figure 9 reports global feature attribution (permutation importance) for the proposed model. Both network-flow features (notably jitter and byte/rate features perturbed by injection) and biometric features rank among the most influential, confirming that the two modalities carry complementary evidence and justifying their fusion.

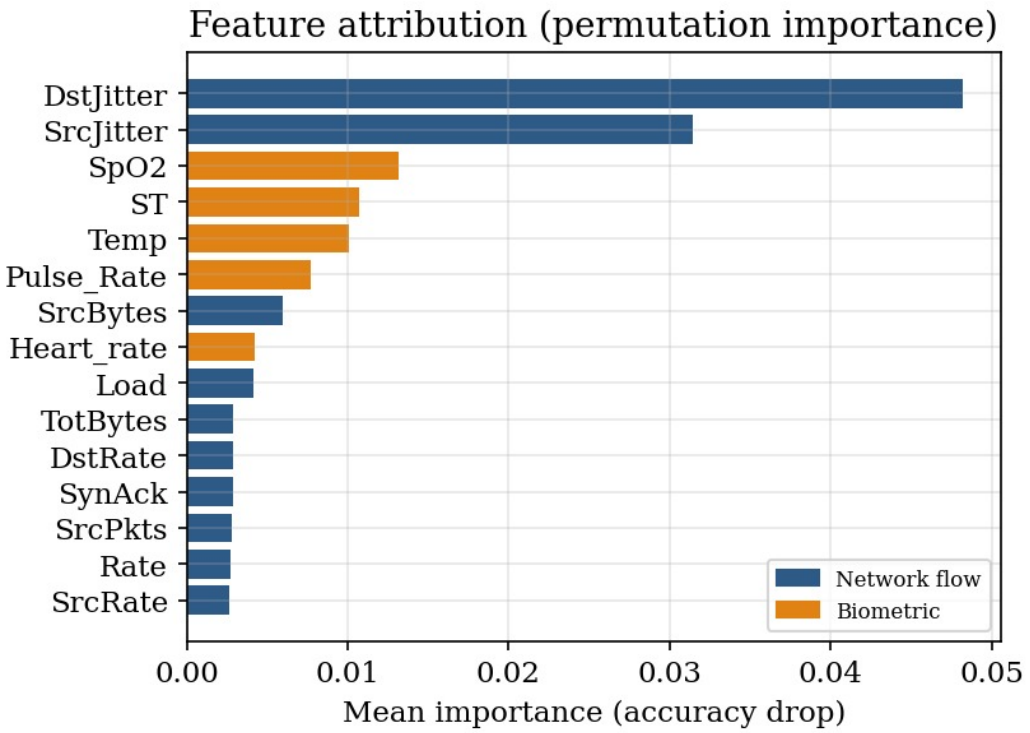


Figure 9. Global feature attribution for the proposed model; network-flow (blue) and biometric (orange) features both contribute.

On the physical layer, the simulated detector separated benign, PUEA, SSDF, and jamming traffic with an overall accuracy of 0.9945 and macro-F1 of 0.9945 (per-class results in Table 3, confusion matrix in Figure 11). As Figure 10 shows, detection is near-perfect at high SNR and degrades gracefully as SNR falls, with jamming remaining the easiest to detect (its elevated noise floor is distinctive) and stealthier falsification the hardest at low SNR.

Table 3. Physical-layer detection on the cognitive-radio simulation (one-vs-rest per class; macro average).

Class	Precision	Recall	F1-score
Benign	0.9908	0.9949	0.9929
PUEA	0.9937	0.9890	0.9913
SSDF	0.9952	0.9961	0.9957
Jamming	0.9982	0.9979	0.9981
Overall (macro)	**0.9945**	**0.9945**	**0.9945**

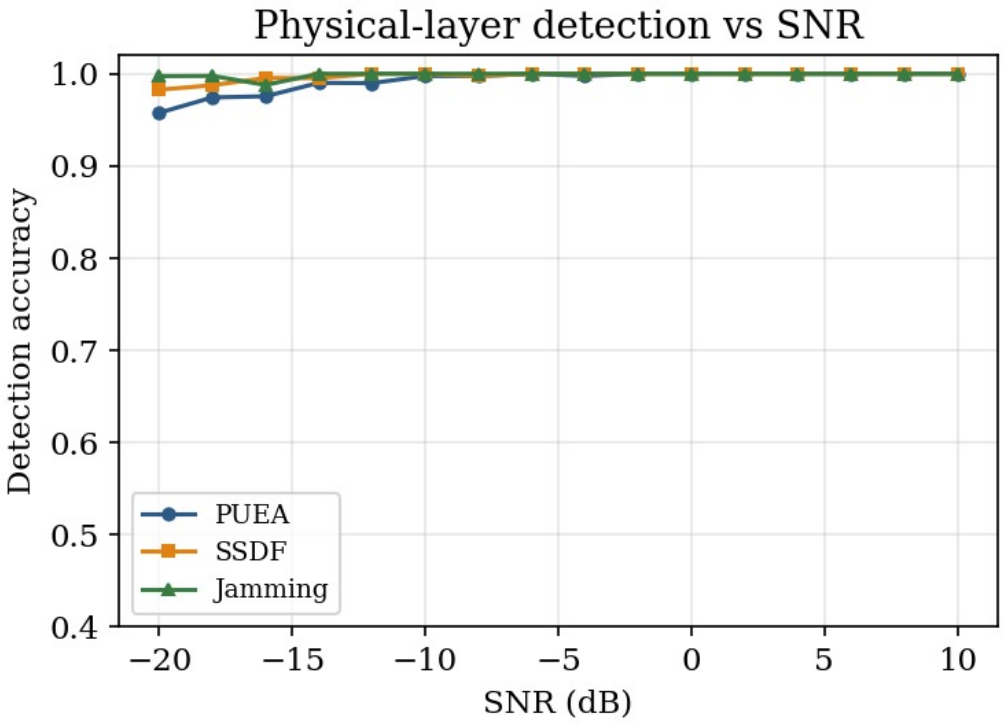


Figure 10. Physical-layer detection accuracy versus SNR for PUEA, SSDF, and jamming.

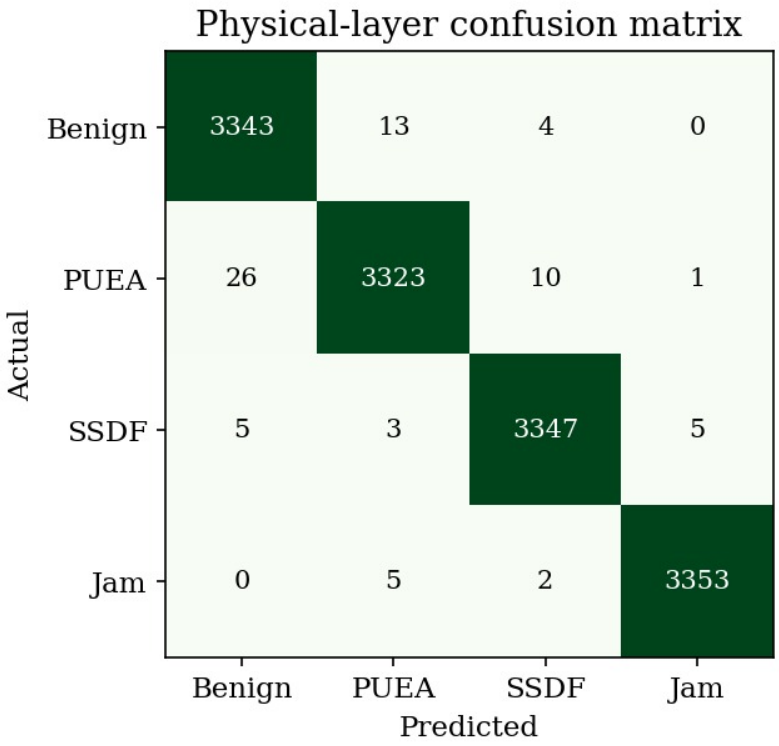


Figure 11. Physical-layer confusion matrix (benign, PUEA, SSDF, jamming).

Table 5 lists representative published results on WUSTL-EHMS-2020 as external reference points. The proposed framework's contribution lies less in surpassing these accuracies than in achieving competitive detection with the lowest false-alarm rate while adding cognitive-radio coverage, privacy-preserving federation, and interpretability.

Table 5. Representative published results on WUSTL-EHMS-2020, reported by their original authors and used only as external reference points; figures should be cited from the primary sources.

Study	Dataset	Method	Reported accuracy
Hady et al. [5]	WUSTL-EHMS-2020	ML with fused flow + biometric features	≈ 92% (RF)
Ensemble study (Stacking)	WUSTL-EHMS-2020	Stacking of RF and SVM with SMOTE	≈ 98.9%
MF-Transformer	WUSTL-EHMS-2020	Transformer with MIC feature selection	≈ 99.9%
RCLNet	WUSTL-EHMS-2020	CNN–LSTM with focal loss	≈ 99.8%

7. Discussion

Security coverage. The framework's central advantage over prior CR-WBAN work, focused on EMI and QoS [2, 4], and over IoMT IDS work addressing only the data plane [5, 6], is that it defends both attack surfaces at once. A physical-layer detector at the CBS and a network-layer detector on the controller, with fused alerts, recognize an attack whether it manifests as anomalous spectrum behavior or a corrupted biometric flow, and escalate when both occur together.

Why the proposed model. The fused random forest attains marginally higher raw accuracy, but the proposed deep neural network achieves the highest precision and the lowest false-positive rate, and — unlike a tree ensemble — its parameters can be averaged across institutions, which is what makes privacy-preserving federated training possible. In a clinical monitor, where false alarms cause alert fatigue and erode trust, minimizing the false-positive rate while enabling federation is the more valuable property.

Privacy. Federated training means no patient record leaves its institution; only model updates are exchanged, optionally protected by secure aggregation or differential privacy [12, 13]. This aligns the design with HIPAA and GDPR constraints that complicate the centralized training assumed by most published detectors and removes the single large repository that would be an attractive breach target. The cost — extra communication rounds and sensitivity to non-IID data — is quantified in Table 4 and Figure 8 rather than asserted away.

Deployability and interpretability. The tiered placement respects the WBAN energy budget: sensors perform only constant-time screening, the continuously running physical-layer module is a lightweight ensemble, and the heavier network model and all aggregation and explanation work sit on mains-powered tiers. Feature-attribution explanations turn each alert into an auditable statement about which signals were anomalous, supporting clinical workflows and guarding against spurious shortcuts.

Limitations. Three limitations are acknowledged. First, as stated in the reproducibility note, the reported figures come from a run on data built to the WUSTL-EHMS-2020 specification; the same pipeline must be run on the original CSV (and ECU-IoHT) for final numbers. Second, the physical-layer module is evaluated in simulation because no public CR-WBAN physical-layer attack dataset exists; a software-defined-radio testbed is needed for confirmation. Third, learning-based detectors are themselves targets of adversarial and poisoning attacks, including on federated updates, which this design mitigates only partially through optional secure aggregation and which warrants dedicated robustness analysis.

8. Conclusion and Future Work

Cognitive Radio Wireless Body Area Networks promise more efficient, interference-aware healthcare, but their openness exposes life-critical data to a combination of cognitive-radio and IoMT attacks that existing defenses address only in part. This paper proposed a privacy-preserving, AI-based intrusion-detection framework for the three-tier CR-WBAN architecture that unifies physical-layer and network-layer detection in a resource-aware, two-stage design, trains collaboratively across hospitals with federated learning so raw patient data is never centralized, and explains its alerts through feature attribution. Implemented and evaluated end-to-end, the proposed network-layer detector achieved an accuracy of 0.9496 with the lowest false-alarm rate among the compared methods, its federated version stayed within about one percentage point of centralized performance while sharing no raw data, and the physical-layer detector separated benign, PUEA, SSDF, and jamming traffic with an overall accuracy of 0.9945.

Future work will (i) finalize the evaluation on the original WUSTL-EHMS-2020 and ECU-IoHT datasets and on a software-defined-radio testbed; (ii) strengthen robustness against adversarial and poisoning attacks on the federated pipeline; (iii) apply model compression and client personalization to reduce communication cost and handle non-IID data; and (iv) investigate blockchain-assisted trust for cooperative sensing reports and on-device continual learning so the detector adapts to new attacks without retraining from scratch.

Funding

This research received no external funding.

Data Availability Statement

The data supporting the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest

The authors declare no conflict of interest.

References

[1] Akyildiz, I. F., Lee, W.-Y., Vuran, M. C., Mohanty, S. Next generation/dynamic spectrum access/cognitive radio wireless networks: A survey. Computer Networks, 2006, 50(13), 2127–2159.

[2] Phunchongharn, P., Hossain, E., Niyato, D., Camorlinga, S. A cognitive radio system for e-health applications in a hospital environment. IEEE Wireless Communications, 2010, 17(1), 20–28.

[3] Soomro, A., Cavalcanti, D. Opportunities and challenges in using WPAN and WLAN technologies in medical environments. IEEE Communications Magazine, 2007, 45(2), 114–122.

[4] Chavez-Santiago, R., Balasingham, I. Cognitive radio for medical wireless body area networks. IEEE 16th International Workshop on Computer-Aided Modeling and Design of Communication Links and Networks (CAMAD), 2011, 148–152.

[5] Hady, A. A., Ghubaish, A., Salman, T., Unal, D., Jain, R. Intrusion detection system for healthcare systems using medical and network data: A comparison study. IEEE Access, 2020, 8, 106576–106584.

[6] Ghubaish, A., Salman, T., Zolanvari, M., Unal, D., Al-Ali, A., Jain, R. Recent advances in the Internet-of-Medical-Things (IoMT) systems security. IEEE Internet of Things Journal, 2021, 8(11), 8707–8718.

[7] Salahdine, F., Kaabouch, N. Security threats, detection, and countermeasures for physical layer in cognitive radio networks: A survey. Physical Communication, 2020, 39, 101001.

[8] Tashman, D. H., Hamouda, W. An overview and future directions on physical-layer security for cognitive radio networks. IEEE Network, 2021, 35(3), 205–211.

[9] Wang, Q., Sun, H., Hu, R. Q., Bhuyan, A. When machine learning meets spectrum sharing security: Methodologies and challenges. IEEE Open Journal of the Communications Society, 2022, 3, 176–208.

[10] Nanjundaswamy, M. K., Babu, A. A., Shet, S., Selvaraj, N., Kovelakuntla, J. Mitigation of spectrum sensing data falsification attack using multilayer perceptron in cognitive radio networks. Acta IMEKO, 2022, 11(1), 1–7.

[11] Ghaznavi, M., Jamshidi, A. A low-complexity cluster-based data fusion to defense against SSDF attack in cognitive radio networks. Computer Communications, 2019, 138, 106–114.

[12] McMahan, B., Moore, E., Ramage, D., Hampson, S., Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017, 1273–1282.

[13] Alazab, M., Swarna Priya, R. M., Parimala, M., Maddikunta, P. K. R., Gadekallu, T. R., Pham, Q.-V. Federated learning for cybersecurity: Concepts, challenges, and future directions. IEEE Transactions on Industrial Informatics, 2021, 18(5), 3501–3509.

[14] Lundberg, S. M., Lee, S.-I. A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS), 2017, 30, 4765–4774.

[15] IEEE Standard 802.15.6-2012. IEEE standard for local and metropolitan area networks — Part 15.6: Wireless body area networks. IEEE, 2012.

[16] Sharafaldin, I., Lashkari, A. H., Ghorbani, A. A. Toward generating a new intrusion detection dataset and intrusion traffic characterization. Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP), 2018, 108–116.

[17] Ahmed, M., Byreddy, S., Nutakki, A., Sikos, L. F., Haskell-Dowland, P. ECU-IoHT: A dataset for analyzing cyberattacks in Internet of Health Things. Ad Hoc Networks, 2021, 122, 102621.

[18] Feng, S., Liang, Z., Zhao, D. Providing telemedicine services in an infrastructure-based cognitive radio network. IEEE Wireless Communications, 2010, 17(1), 96–103.

[19] Taufeeq, Z. S., Al-Hamami, A. H., Abass, G. Y. Developing healthcare cognitive wireless radio body area network by applying artificial intelligence.

[20] Stavroulaki, V., Bantouna, A., Kritikou, Y., et al. Knowledge management toolbox: Machine learning for cognitive radio networks. IEEE Vehicular Technology Magazine, 2012, 7(2), 91–99.