Journal of Computing and Data Technology

https://journal.krrishpub.com/index.php/jcdt/index

RESEARCH ARTICLE

Hybrid Post-Quantum Secure Wireless Healthcare Network with Audit Logging and AI-Based Intrusion Detection

Sunnam Veera Venkata Ravi Teja [a,*], D. Latha [b]

[a] *M. Tech (2024-26), Dept. of Computer Science and Engineering, Adikavi Nannaya University, Rajamahendravaram, 533101, India.*
[b] *Associate Professor, Department of Computer Science and Engineering, Adikavi Nannaya University College of Engineering, Rajamahendravaram, 533101, India.*

Abstract

Remote patient monitoring over Wireless Body Area Networks (WBANs) continuously transports physiological measurements such as heart rate, SpO_2, blood pressure and temperature over wireless links that may be exposed to eavesdropping, replay, manipulation, injection and flooding attacks. Long-lived healthcare records also require protection against future quantum attacks on classical public-key mechanisms. This work presents a hybrid security framework that combines NIST-standardized ML-KEM for post-quantum key encapsulation, ML-DSA for digital signatures, AES-256-GCM for authenticated encryption, HMAC-SHA256 for packet-level integrity verification, a blockchain-inspired hash-chained audit ledger for tamper evidence, and a Random Forest intrusion detection subsystem for behavioral anomaly detection. The prototype integrates a Python TCP gateway, SQLite-based storage with role-based access control, a Streamlit clinician dashboard, and a React/Node.js web client. Functional evaluation covers six security test scenarios involving valid traffic, ciphertext manipulation, replay, signature forgery, malformed input, and packet flooding. The reported prototype measurements are 452 ms for ML-KEM key generation, 318 ms for ML-DSA key generation, 94 µs for AES-256-GCM encryption of a 512-byte payload, 88 ms average gateway processing time, 147 ms end-to-end latency, 342 packets/s sustained throughput, 96% packet delivery ratio, 87% five-class IDS accuracy, and 100% verified audit-chain integrity. A supplementary Random Forest experiment on the BLEWBAN dataset achieved 76.9% accuracy for twelve-class sensor-position classification. The study explicitly distinguishes functional prototype results from production security claims and identifies limitations related to the PQC backend, synthetic attack traffic, hardware platform, scalability and formal security analysis.

Keywords: Post-Quantum Cryptography, Wireless Body Area Network, Blockchain Audit Logging, Intrusion Detection, Random Forest, Healthcare Security.

Article information:
ISSN: 3107-9466 (Online)
Published by: **Krrish Scientific Publications Pvt. Ltd.**
DOI: https://doi.org/10.71426/jcdt.v2.i2.pp193-206
Received: *17 Sep. 2026* | Revised: *05 Oct. 2026* | Accepted: *10 Oct. 2026* | Published: *19 Oct. 2026*
Copyright © 2026 Author(s).
This is an open-access article distributed under the terms of the Creative Commons Attribution–NonCommercial 4.0 International License (CC BY-NC 4.0).

1. Introduction

Wearable and body-mounted sensors enable continuous collection of physiological measurements and support remote patient monitoring without requiring continuous hospital admission. WBAN architectures typically connect sensors, a personal or gateway device, backend services and clinician-facing applications [1], [2]. Because medical measurements and associated identifiers are sensitive, confidentiality, integrity, authentication, availability and accountability are central security requirements [3]- [5].

Wireless transmission creates an attack surface in which

*Corresponding author
Email addresses: svvraviteja272@gmail.com, 24886m601015@aknu.edu.in (Ravi Teja), latha.cse@aknu.edu.in (D. Latha).

an adversary may listen to traffic, inject fabricated readings, replay previously observed packets, manipulate transmitted ciphertext, or flood a gateway. Earlier wearable-health security protocols have focused on authentication and key agreement while balancing the limited computational and energy resources of sensor devices [6]. Blockchain-based WBAN research has additionally investigated integrity, privacy and decentralized audit mechanisms [14]- [16]. Surveys of quantum-era healthcare security likewise stress the need for early migration planning [19], [26]- [28].

Symmetric authenticated encryption remains appropriate for high-volume sensor payloads because it is considerably less expensive per byte than public-key operations. AES-GCM provides authenticated encryption with associated data [11], while HMAC provides a keyed integrity and authentication mechanism [12]. For the machine-learning layer, Random Forests provide an ensemble classifier capable of modeling nonlinear relationships among traffic features [13]. Recent IoMT studies further demonstrate the use of ensemble and deep-learning techniques for intrusion detection [18]- [21].

The security mechanisms above address different layers of the threat surface. Cryptography protects confidentiality, integrity and origin authentication, but a cryptographically valid packet from a compromised device may still exhibit abnormal frequency, timing or failure patterns. Conversely, an IDS can identify anomalous traffic but does not replace cryptographic protection. Audit mechanisms provide a third requirement: evidence that stored security records have not been modified after processing. The proposed framework therefore treats post-quantum cryptography, authenticated encryption, behavioral detection and audit integrity as complementary controls rather than substitutes for one another.

1.1. Problem statement

The problem addressed in this study is the absence of a single, runnable WBAN security pipeline that simultaneously provides quantum-resistant key establishment, quantum-resistant origin authentication, per-packet confidentiality and integrity, tamper-evident security logging, and machine-learning-based detection of attacks that may remain valid at the cryptographic layer. The framework must also be implementable as an end-to-end prototype so that security checks and system-level performance can be evaluated together.

1.2. Research gap

Existing healthcare and WBAN security approaches generally address individual requirements such as authentication, confidentiality, intrusion detection or auditability. Classical authentication mechanisms provide protection against conventional adversaries but are not designed for a future quantum threat. Blockchain-based approaches can improve auditability and integrity but do not, by themselves, provide complete protection for wireless communication. Machine-learning-based IDS approaches can identify anomalous traffic, but they do not inherently provide post-quantum authentication, payload confidentiality or tamper-evident storage. Recent post-quantum healthcare studies have demonstrated quantum-resistant authentication and key-agreement mechanisms [17], [18], while blockchain-oriented WBAN research has investigated decentralized authentication and integrity [14]- [16]. In parallel, IoMT IDS research has explored Random Forest, ensemble, CNN/LSTM and other machine-learning approaches [20]- [24]. However, these strands are commonly evaluated separately, with different datasets, communication models and performance measures. A unified WBAN pipeline that combines post-quantum key establishment, post-quantum signatures, authenticated encryption, packet-level integrity, hash-chained auditing and AI-based intrusion detection remains insufficiently evaluated as an integrated prototype.

1.3. Contributions

- A hybrid WBAN security architecture integrating ML-KEM, ML-DSA, AES-256-GCM, HMAC-SHA256, hash-chained audit logging and Random Forest intrusion detection.
- An end-to-end implementation connecting a patient sensor simulator, Python TCP gateway, cryptographic verification layer, IDS, audit ledger, SQLite persistence and clinician dashboards.
- Functional evaluation against six security test scenarios covering valid traffic, ciphertext manipulation, replay, signature forgery, malformed input and packet flooding.
- Measurement of cryptographic and system-level performance, including key-generation time, encryption time, gateway processing, end-to-end latency, throughput, packet delivery ratio, IDS accuracy and audit-chain integrity.
- A transparent assessment of deployment limitations, including the requirement for a genuine ML-KEM/ML-DSA backend, synthetic attack traffic, development-machine evaluation, single-patient scalability and the absence of a formal compositional security proof.

1.4. Paper organization

Section 2 presents the background and threat model. Section 3 reviews related work and compares the proposed capability set with recent healthcare security approaches. Section 4 describes the proposed methodology and its mathematical formulation. Section 5 presents the implementation, the experimental setup and a walkthrough of the coding and user interfaces. Section 6 reports functional, security and performance results and compares them with prior work. Section 7 discusses ethical and deployment considerations. Section 8 presents limitations and future

scope, and Section 9 concludes the study.

2. Background and threat model

2.1. WBAN security requirements and threat model

WBAN devices operate under constraints in computation, memory, energy and wireless bandwidth. Security mechanisms must therefore limit unnecessary communication and avoid repeated expensive public-key operations where possible [4]- [6]. The principal requirements considered in this work are confidentiality of physiological data, integrity of packets, authentication of packet origin, replay resistance, availability under traffic flooding, access control for stored records, and auditability of security events.

The threat model assumes an external or partially compromised adversary can observe wireless traffic, modify packets, inject packets, replay previously observed packets, attempt signature forgery, and generate high-rate traffic. A compromised sensor holding valid credentials is also considered capable of producing cryptographically valid but behaviorally anomalous traffic. The model does not assume compromise of the gateway's protected private keys or cryptographic primitives.

2.2. Post-quantum cryptography

NIST FIPS 203 specifies ML-KEM as a module-lattice-based key-encapsulation mechanism with ML-KEM-512, ML-KEM-768 and ML-KEM-1024 parameter sets [7]. FIPS 204 specifies ML-DSA as a module-lattice-based digital signature standard [8]. The present prototype uses the abstraction of KEM-generated shared key material followed by symmetric encryption and uses digital signatures for packet-origin authentication. The original Kyber and Dilithium publications provide the cryptographic construction background [9], [10].

The implementation must distinguish between an actual standardized PQC backend and a demonstration fallback. The prototype described in this manuscript includes a pluggable backend. If the native ML-KEM/ML-DSA libraries are unavailable, a SHA-256/HKDF demonstration construction may be used for functional development. Results produced by that fallback must not be interpreted as evidence of post-quantum security. The reported cryptographic security claims therefore apply only when the genuine ML-KEM and ML-DSA implementations are active.

2.3. Authenticated encryption and audit integrity

AES-256-GCM provides authenticated encryption with associated data [11]. A fresh 96-bit nonce is generated for each packet. The ML-KEM shared secret is converted into the AES session key using HKDF-SHA256 [12]. HMAC-SHA256 is additionally computed over the canonical packet representation as an independent integrity check. The combination is intended to provide layered protection rather than to imply that HMAC is required for the security of GCM itself.

The audit subsystem is described as blockchain-inspired rather than as a decentralized blockchain network. Each audit block contains selected packet metadata, a digest of the decrypted payload, and the hash of the previous block. No mining, consensus protocol or peer-to-peer network is used. The resulting hash chain provides tamper evidence for a single-gateway deployment and is less resource-intensive than a full distributed ledger [14]- [16].

3. Related work and state-of-the-art comparison

3.1. Classical and blockchain-based WBAN security

Earlier wearable-health protocols concentrated on mutual authentication, key agreement, anonymity and resistance to conventional network attacks. Jiang et al. developed an end-to-end authentication protocol for wearable health monitoring systems and analyzed attacks affecting earlier authentication designs [6]. Broader WBAN security surveys have emphasized that the wireless channel, constrained sensor resources and multi-tier healthcare architecture create a need for end-to-end protection [4]. These approaches remain relevant as classical baselines, but they do not address the post-quantum migration problem considered here.

Blockchain and distributed-ledger mechanisms have been investigated for healthcare and WBAN security because they can support integrity, traceability and decentralized trust [14], [15]. Singh and Kumar proposed a blockchain-aided certificateless authenticated key-agreement protocol for WBANs and evaluated computational, communication and storage overhead [16]. These studies demonstrate the value of ledger-based trust and audit mechanisms, but they do not combine blockchain-based auditing with the complete ML-KEM/ML-DSA plus AI-IDS pipeline implemented here.

3.2. Post-quantum healthcare security

Recent research has begun applying post-quantum mechanisms to healthcare and e-health systems. Bahache et al. proposed a quantum-resistant authentication and key-agreement framework for cloud-based healthcare applications using Kyber and formal verification [17]. Mansoor et al. proposed PQCAIE for IoT-based e-health systems, specifically addressing the future quantum threat to healthcare authentication [18]. SaberiKamarposhti et al. reviewed the post-quantum migration problem in healthcare and highlighted integration, compatibility and organizational challenges [19]. These works support the need for PQC migration but generally focus on authentication/key agreement rather than an integrated cryptography-plus-IDS-plus-audit pipeline. Recent surveys extend this picture.

Sabrina et al. reviewed quantum-resistant techniques for blockchain-based privacy preservation in the Internet of Medical Things and discussed open challenges [26], and Adil et al. reviewed quantum-computing threats to healthcare IoT security for the period 2017–2025 [27]. These surveys map the problem space, whereas the present work contributes an implemented and measured prototype.

3.3. AI-based intrusion detection in IoMT

Machine-learning-based IDS research has expanded rapidly in IoMT. Si-Ahmed et al. surveyed ML-based intrusion detection across IoMT layers and identified dataset, computational and deployment challenges [20]. Alsolami et al. evaluated stacking, bagging and boosting ensembles using Random Forest and SVM components on WUSTL-EHMS-2020 [21]. Shaikh et al. introduced RCLNet, combining Random Forest feature selection with CNN/LSTM and attention mechanisms [22]. Balhareth and Ilyas evaluated tree-based models including Decision Tree, Random Forest, XGBoost and CatBoost with feature selection [23]. Khan et al. investigated deep-learning-based anomaly detection for smart healthcare systems [24]. A 2024 systematic review of behavior-based IoMT IDS research also identified a shortage of IoMT-specific datasets and limited attention to direct patient-safety implications [25].

3.4. Capability comparison

Table I compares capabilities rather than ranking the studies. Direct numerical performance comparison is not appropriate because the studies use different datasets, hardware platforms, attack models and evaluation protocols. The comparison therefore focuses on whether a capability is part of the reported framework.

4. Proposed methodology

4.1. System architecture

The proposed framework is organized into three functional layers: the patient sensor node, the security gateway, and the storage/clinician layer. A sensor captures physiological measurements and constructs a canonical SecurePacket. The gateway performs structural validation, timestamp and sequence checks, HMAC verification, ML-DSA signature verification and AES-GCM decryption. The IDS operates in parallel with the security path, while the audit subsystem appends a hash-chained record after successful verification. Verified data are then persisted and exposed to authorized clinicians through role-based interfaces (Figure 1).

4.2. Registration and SecurePacket design

Patients and sensors are registered before telemetry processing. Patient identifiers (PID) and Body Sensor Identifiers (BSI) are derived using salted SHA-256 hashing and stored with role-based access-control assignments. Registered sensor identities are associated with their ML-DSA

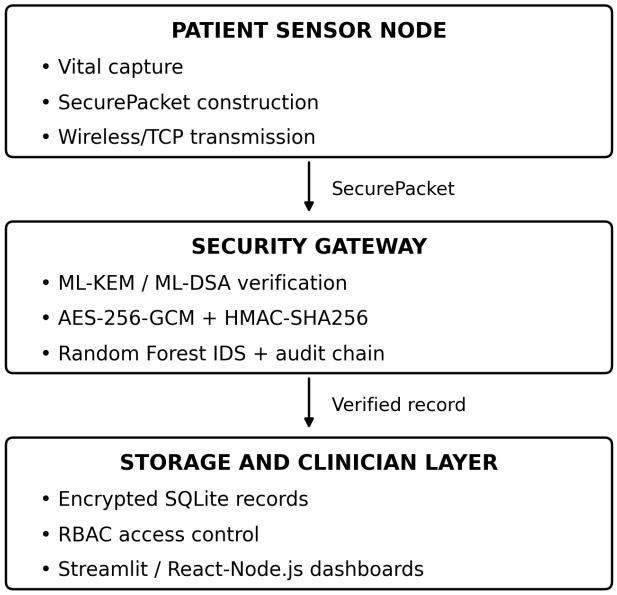

Figure 1: Architecture of the proposed hybrid post-quantum healthcare security framework.

public keys so that the gateway can perform origin verification without an external key lookup. Doctors receive patient-scoped access, analysts receive aggregate access, and administrators manage provisioning.

Each measurement is transported in a canonical JSON envelope. The packet contains a protocol version, millisecond timestamp, PID, BSI, ciphertext, 12-byte AES-GCM nonce, ML-KEM ciphertext, HMAC-SHA256 tag, ML-DSA signature and sequence number. JSON keys are sorted lexicographically and serialized without additional whitespace before signature generation. This canonicalization ensures that the sender and receiver verify identical byte sequences. Figure 2 shows the construction sequence.

4.3. Cryptographic protection pipeline

The gateway publishes an ML-KEM public key. For a new session, the sensor encapsulates a shared secret to that public key and includes the resulting KEM ciphertext in the SecurePacket. The receiver decapsulates the ciphertext using its private key. The resulting shared secret is used as input keying material for HKDF-SHA256 to derive the AES-256 session key [7], [9], [12].

The ML-KEM public and private key pair is generated according to Equation (1).

$$(pk, sk) \leftarrow \mathrm{KeyGen}_{\text{ML-KEM}}(1^{\lambda}) \tag{1}$$

The encapsulation and decapsulation operations are expressed in Equation (2).

$$(c, K) \leftarrow \mathrm{Encaps}(pk); \quad K \leftarrow \mathrm{Decaps}(sk, c) \tag{2}$$

The shared secret is processed by HKDF to obtain the

Table 1: Capability comparison of recent studies and the proposed framework.

Study	Domain	PQC	Audit / Blockchain	AI/ML IDS	Primary focus
Usman et al. [4]	WBAN	–	–	–	Classical security survey
Chandra & Kanaga Suba Raja [14]	WBAN	–	✓	–	Blockchain-oriented security
Sanjith et al. [15]	WBAN	–	✓	–	Blockchain-based security/privacy
Singh & Kumar [16]	WBAN	–	✓	–	Blockchain-aided authenticated key agreement
Bahache et al. [17]	IoMT/WBAN	✓	–	–	Quantum-resistant authentication
Mansoor et al. [18]	IoT e-health	✓	–	–	Post-quantum authentication
Alsolami et al. [21]	IoMT	–	–	✓	Ensemble ML IDS
Shaikh et al. [22]	IoMT	–	–	✓	CNN/LSTM + RF feature selection
Balhareth & Ilyas [23]	IoMT	–	–	✓	Tree-based ML IDS
Sabrina et al. [26]	IoMT	✓	✓	–	Post-quantum blockchain privacy review
Proposed framework	WBAN	✓	✓	✓	Integrated end-to-end prototype

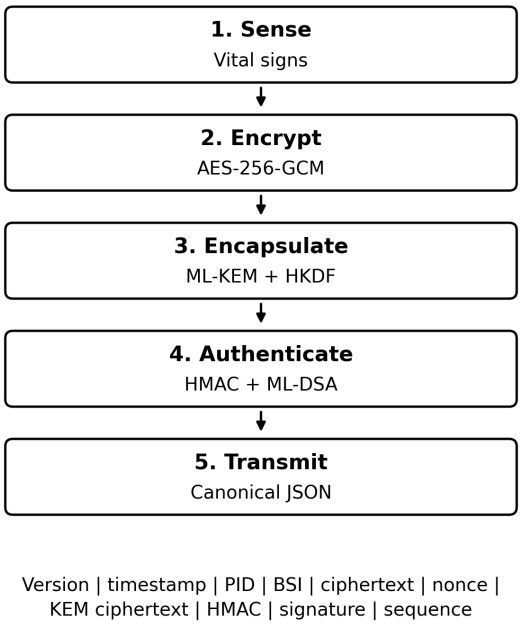

Figure 2: SecurePacket construction workflow and canonical packet fields.

AES session key, as formulated in Equation (3).

$$\begin{aligned} PRK &= \text{HKDF-Extract}(salt, K); \\ K_{\text{AES}} &= \text{HKDF-Expand}(PRK, info, 32) \end{aligned} \tag{3}$$

The sensor signs the canonical packet body with its registered ML-DSA private key. The gateway verifies the signature against the stored public key. A failed verification causes packet rejection before the payload is accepted. The ML-DSA signature generation operation is represented by Equation (4).

$$\sigma = \text{Sign}_{\text{ML-DSA}}(sk_{\text{sensor}}, m) \tag{4}$$

The corresponding signature verification condition is given in Equation (5).

$$\text{Verify}_{\text{ML-DSA}}(pk_{\text{sensor}}, m, \sigma) = \text{True} \tag{5}$$

The physiological payload is encrypted with AES-256-GCM using the derived session key and a fresh 96-bit nonce. Associated metadata can be authenticated without being encrypted. HMAC-SHA256 is computed over the canonical packet body and verified independently at the gateway. The AES-GCM encryption operation is defined in Equation (6).

$$(C, T) = \text{AES-GCM-Encrypt}(K_{\text{AES}}, N, P, A) \tag{6}$$

The authenticated decryption operation is expressed in Equation (7).

$$P = \text{AES-GCM-Decrypt}(K_{\text{AES}}, N, C, T, A) \tag{7}$$

The HMAC computation and recomputation at the receiving end are shown in Equation (8).

$$\begin{aligned} H &= \text{HMAC-SHA256}(K_H, M); \\ H' &= \text{HMAC-SHA256}(K_H, M) \end{aligned} \tag{8}$$

A packet passes the HMAC check only when H' equals H. The AES-GCM authentication tag is also required for successful decryption. These checks provide layered integrity verification, while ML-DSA provides origin authentication.

4.4. Replay protection and audit ledger

Replay resistance is implemented using a timestamp acceptance window and sequence-number monotonicity. A packet is rejected when its timestamp falls outside the configured acceptance window or when its sequence number violates the expected ordering. In the evaluated prototype, the stale replay test used a two-minute window.

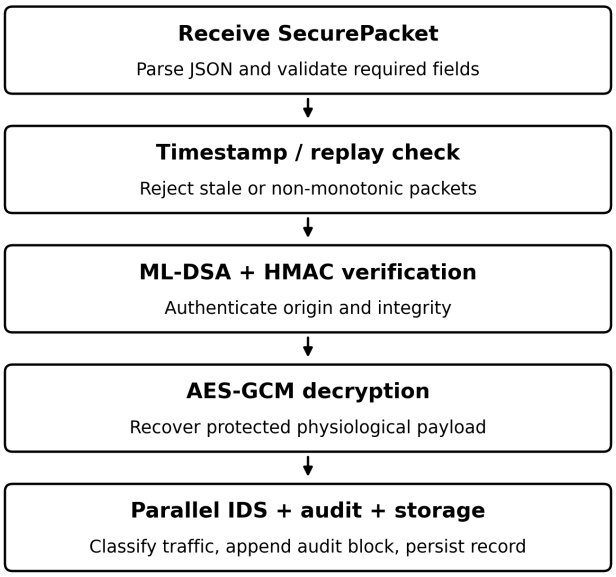

Figure 3: Gateway verification path with sequential security checks and parallel security intelligence.

The complete verification sequence is shown in Figure 3. The combined packet acceptance condition is formalized in Equation (9).

$$\begin{aligned} \text{Accept}_{\text{packet}} = \text{TimestampValid} \wedge \text{SequenceValid} \\ \wedge \text{SignatureValid} \wedge \text{HMACValid} \\ \wedge \text{GCMValid} \end{aligned} \tag{9}$$

Each successfully verified packet creates an audit block containing sequence number, timestamp, PID, BSI, a SHA-256 digest of the decrypted payload and the previous block hash. The current block hash is calculated over the block fields and the previous hash. Because the current hash depends on the preceding block, modification of an earlier record causes a chain inconsistency during recomputation. The hash-chaining operation used for audit-block integrity is presented in Equation (10).

$$H_i = \text{SHA256}(B_i \,\|\, H_{i-1}) \tag{10}$$

The ledger provides tamper evidence and chronological integrity within the gateway's audit store. It does not use consensus, mining or a peer-to-peer network (Section 2.3).

4.5. AI-based intrusion detection

A Random Forest classifier operates asynchronously beside the cryptographic pipeline. Four network features are computed per packet: packet size, arrival frequency within a 10-second sliding window, gateway processing delay, and a rolling count of HMAC/signature failures over the previous 50 packets. The classifier produces one of five labels: Normal, Replay, MITM, Injection or DoS. The four-dimensional input feature vector for the classifier is defined in Equation (11).

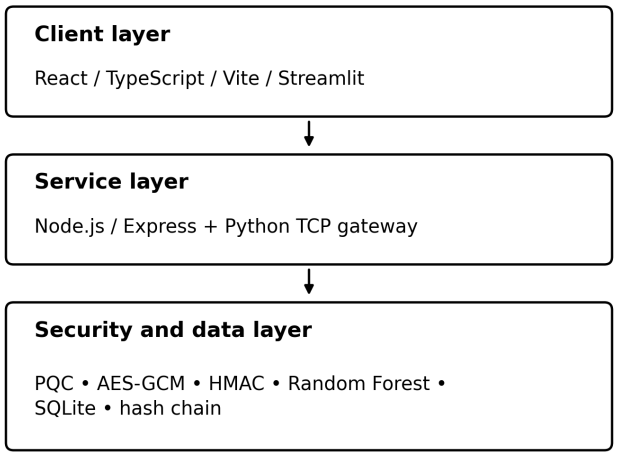

Figure 4: Implementation and deployment stack used by the prototype.

$$\begin{aligned} X = [\text{packet_size}, \text{arrival_frequency}, \\ \text{delay}, \text{failure_count}] \end{aligned} \tag{11}$$

The final Random Forest prediction is determined by majority voting, as expressed in Equation (12).

$$\hat{y} = \text{mode}\{T_1(X), T_2(X), \ldots, T_n(X)\} \tag{12}$$

For the reported prototype, n is the number of decision trees configured in the Random Forest model. The implementation also uses class-weighting so that attack classes are not ignored when traffic proportions differ.

4.6. Storage, RBAC and clinician interfaces

Verified records are persisted in an AES-protected SQLite database with separate logical entities for patients, sensors, RBAC users, vital records, audit blocks and alerts. Query-level role checks restrict clinicians to authorized patient records. The prototype exposes the same gateway security path to a Streamlit dashboard and a React/Node.js web client. Figure 4 summarizes the implementation stack.

5. Implementation and experimental setup

5.1. Implementation environment

The prototype is implemented primarily in Python 3.10, with scikit-learn used for the Random Forest models and NumPy used for feature processing. The gateway, IDS inference loop, audit ledger and SQLite persistence communicate through a TCP socket on port 5055. A Node.js/Express backend accepts HTTP requests from the React client and forwards SecurePacket submissions to the Python gateway. A launch script starts the gateway, patient simulator and dashboard components for repeatable local experiments. The pluggable cryptographic interface and

Table 2: Threat-to-control mapping of the proposed framework.

Threat	**Attack surface**	**Primary control**	**Expected effect**
Eavesdropping	Wireless interception	AES-256-GCM	Payload confidentiality
Man-in-the-middle	Packet substitution or origin spoofing	ML-DSA + HMAC	Origin and integrity verification
Replay	Old valid packet resent	Timestamp + sequence	Stale packets rejected
Packet injection	Fabricated or modified packet	ML-DSA + HMAC + GCM	Packet rejected
DoS / flooding	High-rate packet bursts	Random Forest IDS	Behavioral anomaly alert
Audit tampering	Stored audit record modification	Hash chain	Chain inconsistency detected
Unauthorized access	Improper record retrieval	RBAC	Role-scoped queries
Quantum attack	Future compromise of classical public-key layer	ML-KEM + ML-DSA	Post-quantum primitives when genuine backends are active

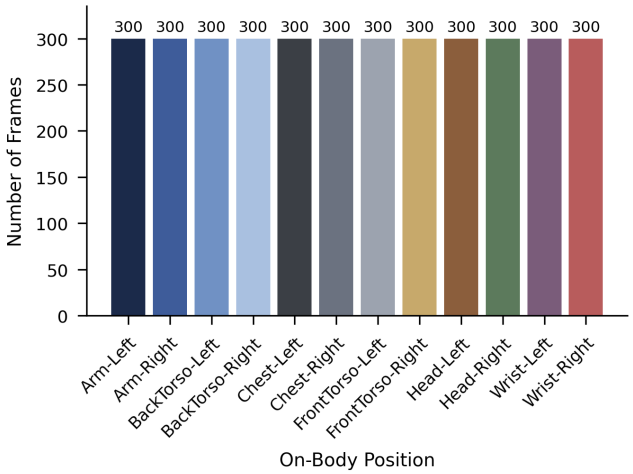

Figure 5: BLEWBAN class distribution: 300 frames for each of 12 on-body sensor positions.

its fallback condition are described in Section 2.2; native liboqs or pqcrypto implementations are required for any post-quantum security claim.

5.2. Data sources

Two data sources are used. Fitbit biometric time series drive the patient simulator, which produces realistic waveform and vital-sign inputs for functional testing. The supplementary sensor-position classification experiment uses the BLEWBAN dataset, which contains 3,600 synthesized Bluetooth Low Energy frames distributed across twelve on-body positions, with 300 frames per position (Figure 5). The carrier-frequency offsets of these frames are summarized in Figure 6. The network IDS uses labeled synthetic WBAN traffic generated by the patient simulator across five classes: Normal, Replay, MITM, Injection and DoS.

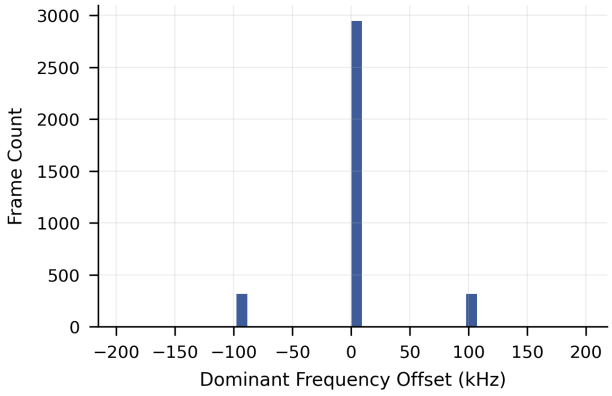

Figure 6: Distribution of the dominant carrier-frequency offset across the 3,600 BLEWBAN frames.

5.3. Feature extraction and training

For BLEWBAN frames, raw I/Q samples are processed through an FFT pipeline to obtain thirteen signal features, including dominant carrier-frequency offset, spectral centroid, phase-gradient statistics, I/Q amplitude statistics, three sub-band energy measures, SNR and spectral entropy. For the network IDS, four packet-level features are used as described in Section 4.5. Random Forest training uses class weights inversely proportional to class frequency.

The BLEWBAN classifier is evaluated on a held-out test split. The network IDS is evaluated on a held-out five-class test set. Because the network attacks are synthetically generated, the resulting 87% IDS accuracy is treated as prototype evidence rather than a deployment-ready performance guarantee.

5.4. Experimental procedure

The functional evaluation exercises valid packets and adversarial cases. Ciphertext manipulation is created by

flipping a bit in the protected payload; replay is tested using a stale timestamp; signature forgery uses an invalid ML-DSA signature; malformed JSON is used to test parser robustness; and a 500 packets/s burst is used to exercise DoS detection. Performance measurements include one-time key generation, per-packet encryption, gateway verification/decryption, end-to-end latency, throughput and packet delivery ratio.

5.5. Coding and interface walkthrough

To document the executable prototype, Figures 7–10 show the interfaces and code paths used during the experiments. Figure 7 shows the Python gateway after start-up. The gateway reports the SQLite database path in use, loads the serialized scikit-learn estimators (Decision Tree Classifier and Random Forest Classifier) used for intrusion detection, and then listens on 127.0.0.1:5055, the TCP endpoint to which the sensor simulator and the Node.js backend submit Secur ePackets. The Inconsistent Version Warning messages indicate that the stored estimators were serialized with scikit-learn 1.9.0 but loaded under version 1.6.0. The models loaded and the gateway started, but the library versions should be pinned in the released environment to guarantee reproducible IDS behavior.

Figure 8 shows the React client used for patient registration and login. In panel (a), the user enters a PID and password and selects Generate Keys; the interface displays a public key and a MAC value and confirms generation with a notification. The short key and MAC strings shown are demonstration values rather than full-length ML-KEM or ML-DSA keys. In panel (b), successful registration returns the assigned PID (PID8888 in this run) and asks the user to store it, because it is required for later login. These identifiers follow the PID and BSI scheme described in Section 4.2.

Figure 9 shows the sender-side routine of the smoke-test script. Before the excerpt, the script registers a patient and a sensor, generates a KEM key pair, encapsulates a shared secret and takes a 32-byte AES session key from it. The excerpt then serializes simulated vital signs (heart rate, temperature and blood pressure) as JSON and encrypts them with AES-GCM. It builds the SecurePacket, removes the signature field, serializes the remainder as canonical JSON with sorted keys and compact separators, computes the HMAC-SHA256 tag over that byte string and signs it. This follows the canonicalization rule in Section 4.2. The script sets `PQC_MODE` to demo, so this run exercises the functional fallback described in Section 2.2 and must not be read as evidence of post-quantum security.

Figure 10(a) shows the transmission and verification part of the script. The packet is sent to the gateway over a TCP socket and the gateway response is printed. The script then queries SQLite and counts the AuditBlock rows for the patient–sensor pair and the CloudRecord rows for the patient, which confirms that an accepted packet produces both an audit block and a stored record. Figure 10(b) shows a second probe that sends an empty JSON object to the gateway to examine how it reacts to a packet with missing required fields, which is related to the malformed-input scenario in Section 5.4. Together, these views show how the prototype is started, exercised and checked: registration and key generation in the web client, packet construction and submission by the sender, verification and persistence at the gateway, and confirmation through audit and storage records.

Table 3: Functional security test scenarios and observed outcomes.

Scenario	Test condition	Observed outcome
Valid packet	HMAC and ML-DSA correct	Accepted; audit block written; record persisted
Ciphertext tampering	Single bit modified	AES-GCM tag mismatch; rejected before decryption
Replay	Packet resent outside two-minute window	Timestamp validation failed; packet dropped
Forged ML-DSA signature	Invalid signature	Signature verification failed; packet rejected
Malformed JSON	Structural injection	Parse exception handled; gateway remained operational
Packet flood	500 packets/s burst	IDS flagged DoS on affected sliding windows

6. Results and discussion

6.1. Functional security testing

All six functional/security test cases produced the expected behavior in the reported prototype. Cryptographic failures were rejected before the payload was accepted, while the flood test was handled by the asynchronous IDS path. These tests demonstrate functional integration; they do not constitute a formal proof of security against all possible attack variants.

6.2. Cryptographic and system performance

The reported 147 ms end-to-end latency is below the 200 ms target used for the prototype's normal-load configuration. The 94 µs AES-GCM encryption time is small relative to the 88 ms gateway processing time, indicating that verification, signature processing, packet handling and related gateway work dominate the measured path. The one-time ML-KEM and ML-DSA key-generation costs are substantially larger than symmetric encryption, supporting the architectural choice to use PQC primarily for session establishment and origin authentication rather than encrypting every payload directly with public-key cryptography.

The 342 packets/s throughput and 96% packet delivery ratio demonstrate operation under the tested local configu-

```
[Gateway] Runtime DB_PATH=e:\mtech pt\mtech-healthcare-pqc\server\..\data\healthcare.db
C:\Users\DELL\Python\Lib\site-packages\sklearn\base.py:380: InconsistentVersionWarning: Trying to unpickle estimator
 DecisionTreeClassifier from version 1.9.0 when using version 1.6.0. This might lead to breaking code or invalid res
ults. Use at your own risk. For more info please refer to:
https://scikit-learn.org/stable/model_persistence.html#security-maintainability-limitations
  warnings.warn(
C:\Users\DELL\Python\Lib\site-packages\sklearn\base.py:380: InconsistentVersionWarning: Trying to unpickle estimator
 RandomForestClassifier from version 1.9.0 when using version 1.6.0. This might lead to breaking code or invalid res
ults. Use at your own risk. For more info please refer to:
https://scikit-learn.org/stable/model_persistence.html#security-maintainability-limitations
  warnings.warn(
[Gateway] Listening on 127.0.0.1:5055
```

Figure 7: Python gateway start-up showing the database path, the loading of the serialized IDS estimators and the TCP listener on 127.0.0.1:5055.

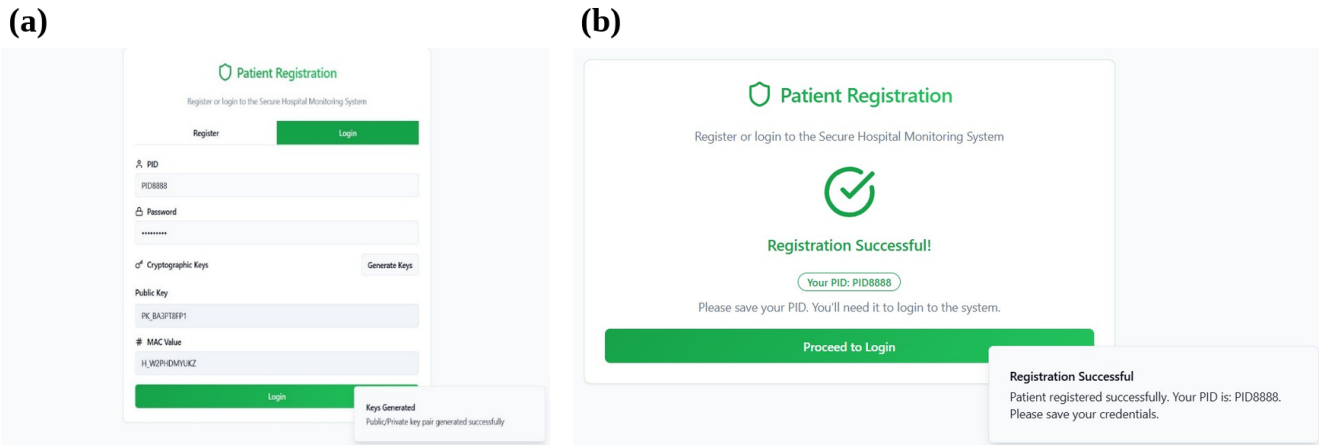

Figure 8: React client interface: (a) patient login with key generation; (b) successful patient registration returning the assigned PID.

Table 4: Cryptographic and system performance measurements.

Metric	Measured result
ML-KEM key generation (per session)	452 ms
ML-DSA key generation (per session)	318 ms
AES-256-GCM encryption (512-byte payload)	94 µs
Gateway processing including signature and HMAC verification	88 ms
End-to-end latency (sensor → gateway confirmation)	147 ms
Sustained throughput	342 packets/s
Packet delivery ratio	96%
IDS accuracy (5-class held-out test set)	87%
Audit-chain integrity	100% verified

ration. They should not be interpreted as clinical-network capacity figures because the evaluation was performed on a development machine and a single-patient configuration.

The latency breakdown supports this interpretation. Gateway processing accounts for about 60% of the 147 ms end-to-end latency (88 ms), whereas AES-GCM encryption contributes less than 0.1% (94 µs); the remaining 59 ms or so covers the other stages of the sensor-to-confirmation path. Optimization effort should therefore target signature verification and gateway handling rather than symmetric encryption.

6.3. BLEWBAN position classification

The supplementary Random Forest model trained on thirteen FFT-derived BLEWBAN features achieved 76.9% held-out accuracy for twelve-class sensor-position identification. Figure 11 shows that phase-gradient standard deviation and mean amplitude provide the largest Gini importance, while dominant carrier-frequency offset contributes the least relative importance once richer signal statistics are available.

The confusion matrix in Figure 12 is computed on a held-out test split of 900 frames (75 per class), of which 692 are classified correctly (76.89%), about nine times the chance level of 8.3% for twelve classes. The matrix shows strong diagonal behavior for several chest and front-torso positions. The most noticeable confusion occurs among bilateral arm and wrist positions, which can have similar propagation and multipath characteristics from the gateway viewpoint. In fact, 207 of the 208 misclassified frames (99.5%) are left–right confusions at the same body site, so the classifier identifies the body region reliably but struggles to separate the two sides. This auxiliary classification task does not

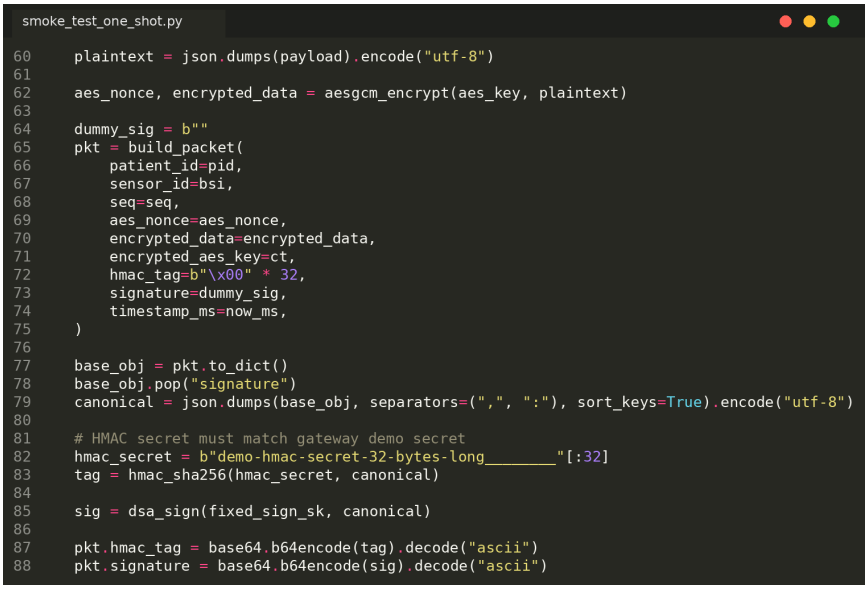

Figure 9: Sender-side routine of the smoke-test script (`smoke_test_one_shot.py`, lines 60–88), rendered from the project repository: AES-GCM encryption, SecurePacket construction, canonical JSON, HMAC-SHA256 and signature generation.

determine whether a vital-sign record is accepted; records remain bound to their registered BSI.

6.4. Comparison with state-of-the-art frameworks

The capability comparison in Table 1 shows that recent research tends to specialize in one or two security dimensions. Blockchain-oriented WBAN studies emphasize auditability, integrity or authenticated key agreement [14]- [16]. Post-quantum healthcare studies emphasize quantum-resistant authentication and key establishment [17], [18]. IoMT IDS studies focus on classification accuracy and feature engineering using public healthcare datasets [20]- [24]. The proposed prototype combines these dimensions in one runnable pipeline.

Relative to the post-quantum healthcare studies, the proposed pipeline extends quantum-resistant authentication and key agreement [17], [18] with per-packet authenticated encryption, replay protection, a tamper-evident audit chain and behavioral detection, and it reports measured gateway-level latency and throughput. Its limitations are the absence of the formal protocol verification used by Bahache et al. [17] and of a certified native PQC backend. Relative to blockchain-aided WBAN protocols [14]- [16], the hash-chained ledger is lighter and avoids consensus overhead, at the cost of decentralized trust: it protects a single-gateway deployment and does not remove the gateway as a trust anchor. Singh and Kumar [16] reported computational, communication and storage overhead for their protocol, and a comparable cost analysis of the audit chain at larger scale remains future work.

The numerical values should not be compared as direct performance rankings because the underlying datasets, hardware, attack distributions and experimental procedures differ. In particular, published IoMT IDS studies [21]- [23] may report substantially higher accuracy on public datasets than the 87% obtained here on synthetic five-class WBAN traffic; this difference is not evidence that one system is universally better. Instead, it highlights the trade-off between a narrowly optimized classifier and an end-to-end security pipeline whose evaluation includes cryptographic processing, audit logging and communication behavior. The systematic review in [25] reports a scarcity of IoMT-specific datasets and limited attention to patient safety; the present study meets the same constraint in the form of synthetic attack traffic.

Survey literature on post-quantum healthcare security [19], [26], [27] emphasizes integration, compatibility and open-challenge themes. The present measurements contribute an implemented data point for these discussions, namely the cost of PQC primitives relative to symmetric encryption and the share of gateway processing in end-to-end latency, but they come from a single development machine and should be read as indicative.

6.5. Threats to validity

Internal validity is limited by the use of synthetic attack traffic for the network IDS. The 87% accuracy should therefore be re-evaluated on captured adversarial traffic before deployment. External validity is constrained by development-machine measurements; performance on ARM-class gateway hardware may differ. The audit-chain integrity result is stronger within the tested storage configuration because recomputation produced 100% verification across the evaluated chains. The auxiliary biomarker outputs in the dashboard have not been clinically validated

smoke_test_one_shot.py

```
99          "encrypted_data",
100         "aes_nonce",
101         "encrypted_aes_key",
102         "hmac_tag",
103         "signature",
104         "seq",
105     }
106     missing = sorted(list(required_keys - set(pkt_obj.keys())))
107     if missing:
108         raise RuntimeError(f"Smoke test built invalid packet JSON, missing keys: {missing}")
109
110     with socket.create_connection((host, port), timeout=10) as s:
111         s.sendall(packet_json.encode("utf-8"))
112         resp = s.recv(50_000)
113
114     resp_str = resp.decode("utf-8")
115     print("[SMOKE_TEST] Gateway response:", resp_str)
116
117     # Verify DB side effects
118     conn = connect(db_path)
119     audit_rows = conn.execute(
120         "SELECT COUNT(*) AS c FROM AuditBlock WHERE patient_pid=? AND sensor_bsi=?",
121         (pid, bsi),
```

smoke_test_dbg_recv.py

```
6  def main():
7      host = '127.0.0.1'
8      port = 5055
9
10     # Send a packet-shaped JSON with missing required fields on purpose
11     # to validate what the gateway expects.
12     test_packet: dict[str, Any] = {}
13
14     payload = json.dumps(test_packet, separators=(',', ':'))
15     with socket.create_connection((host, port), timeout=5) as s:
16         s.sendall(payload.encode('utf-8'))
17         resp = s.recv(50_000)
18
19     print(resp.decode('utf-8'))
```

Figure 10: (a) Packet submission to the gateway and database verification (`smoke_test_one_shot.py`, lines 99–121); (b) malformed-input probe (`smoke_test_dbg_recv.py`, lines 6–22). Both excerpts are rendered from the project repository.

and should not be interpreted as medical measurements.

7. Ethical and deployment considerations

The framework is a cybersecurity prototype and not a diagnostic system. IDS outputs are attack-type labels rather than clinical decisions. Because an 87% IDS accuracy implies nonzero false positives and false negatives, security alerts should remain subject to human review in any clinical deployment.

The distinction between a genuine PQC implementation and the demonstration fallback is particularly important: a deployment must verify that the actual ML-KEM and ML-DSA libraries are active before claiming quantum-resistant protection (Section 2.2). Similarly, because the attack dataset is synthetic, the IDS should initially be deployed in monitoring-only mode while its alerts are compared against real network traces.

The audit mechanism is intentionally lightweight and centralized. It provides tamper evidence but should not be represented as equivalent to a decentralized blockchain with independent consensus. A production deployment may require stronger key management, hardware-backed protection, centralized security monitoring, incident response integration and a formal privacy assessment.

8. Limitations and future scope

The current constraints fall into six groups, which also define the evidence still required before production use.

- Synthetic attack traffic: the 87% IDS result is based on synthetic labeled traffic and should be validated using captured adversarial sessions and broader attack scenarios.
- Prototype deployment: the 147 ms latency result must be re-measured on representative ARM-class gateway hardware and constrained sensor platforms.
- Scalability: current measurements represent a single-patient configuration; concurrent multi-sensor and multi-patient testing is required.
- Real-world and clinical validation: dashboard auxiliary outputs require expert annotation and clinical study before they can be used beyond an investigational inter-

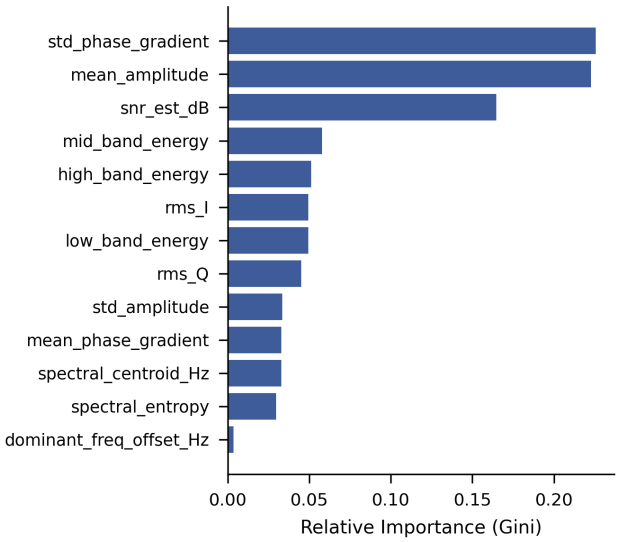

Figure 11: Random Forest Gini feature importance for BLEWBAN sensor-position classification.

face.

- Native PQC backend validation: quantum-resistance claims require genuine ML-KEM and ML-DSA implementations rather than the functional fallback.
- Formal security analysis: the composition of ML-KEM, ML-DSA, AES-GCM, HMAC and the audit-chain mechanism has not been formally proven as a combined protocol.

Future work will therefore focus on native PQC backend certification, embedded/ARM benchmarking, real-world WBAN traffic collection, formal protocol analysis, multi-patient scalability testing, and independent clinical validation of any non-security dashboard outputs.

9. Conclusion

This paper presented a hybrid post-quantum security architecture for wireless healthcare networks and evaluated it through an end-to-end prototype. ML-KEM provides the post-quantum key-establishment layer, ML-DSA provides packet-origin authentication, AES-256-GCM protects payload confidentiality and authenticated encryption, HMAC-SHA256 provides an additional packet-integrity check, a blockchain-inspired hash chain protects audit-record integrity, and Random Forest provides behavioral intrusion detection.

In this work, the measured values included 452 ms ML-KEM key generation, 318 ms ML-DSA key generation, 94 µs AES-GCM encryption for a 512-byte payload, 88 ms gateway processing, 147 ms end-to-end latency, 342 packets/s sustained throughput, 96% packet delivery ratio, 87% five-class IDS accuracy and 100% audit-chain verification. The BLEWBAN supplementary classifier achieved 76.9% accuracy on twelve sensor-position classes.

The principal contribution is the integration of com-

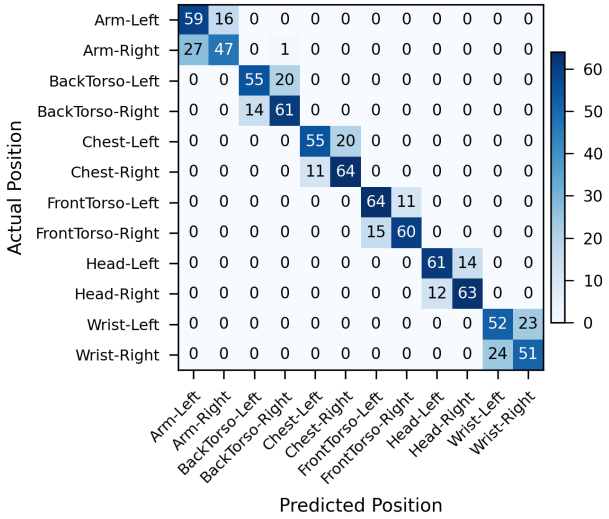

Figure 12: Twelve-class confusion matrix for BLEWBAN sensor-position classification (held-out accuracy: 76.89%).

plementary controls in a single runnable WBAN pipeline rather than a claim that any one component is sufficient by itself. The results establish functional feasibility under the tested conditions, while the stated limitations define the evidence still required for production deployment: genuine PQC backend validation, representative hardware testing, real adversarial traffic, formal security analysis and multi-patient scalability.

Declarations and ethical statements

Conflict of interest: The authors declare that they have no conflicts of interest.

Funding statement: The authors declare that no specific funding was received for this research.

Artificial Intelligence usage statement: During the preparation of this manuscript, the authors used Claude solely for language refinement and grammatical correction. The authors reviewed and revised all generated content and take full responsibility for the accuracy, integrity, and originality of the final manuscript.

Preprint disclosure statement: An earlier version of this article was made publicly available as a preprint on Research Square (https://www.researchsquare.com/article/rs-10558164/v1). The manuscript subsequently underwent peer review and editorial evaluation by the Journal of Computing and Data Technology (JCDT). The present article represents the peer-reviewed and published Version of Record.

Data availability statement: The source code supporting the prototype, cryptographic modules, gateway, intrusion-detection components, audit-chain implementa-

tion, database schema, patient simulator and experimental scripts is publicly available through a version-controlled repository at https://github.com/Venkata99990/HealthCare-PQC.git. The supplementary BLEWBAN experiment uses the dataset described in the original study setup, while network attack traffic is synthetically generated by the patient simulator. No new clinical intervention or diagnostic study was performed as part of this prototype.

CRediT authorship contribution statement

Sunnam Veera Venkata Ravi Teja: Conceptualization, Methodology, Visualization, Writing – original draft. **D. Latha:** Supervision, Validation & Writing- editing.

Publisher's note

Krrish Scientific Publications Pvt. Ltd. and the ***Journal of Computing and Data Technology*** remain neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

[1] Reetha. SmartBio: An AI-Enabled Smart Medical Device for Early Cancer Detection using Variational Autoencoders and Multimodal Sensor Integration. *Journal of Modern Technology*. 2025 Jun. 28 ;2(1):292-301. Available from: https://doi.org/10.71426/jmt.v2.i1.pp292-301

[2] Mogili VB. Designing Trustworthy AI in Healthcare: Experiences with Copilot Agents, Agentic Models, and RAG Integration. *IPHO-Journal of Advance Research in Science And Engineering*. 2026 Mar.3;4(2):19-25. Available from: https://doi.org/10.5281/zenodo.18851338

[3] Patel R, Pothireddy SR, Sharma N. Scalable and secure cloud architectures for IoT-enabled healthcare systems. In *2026 2nd International Conference on Big Data & Machine Learning (ICBDML)* 2026 Feb 27 (pp. 1-6). IEEE. Available from: https://doi.org/10.1109/ICBDML68582.2026.11544363

[4] Usman M, Asghar MR, Ansari IS, Qaraqe M.Security in Wireless Body Area Networks: From In-Body to Off-Body Communications. *IEEE Access*. 2018;6:58064–58074. Available from: https://doi.org/10.1109/ACCESS.2018.2873825

[5] Wang J, Tabassum AT, Li JP. Systematic survey on data security in wireless body area networks in IoT healthcare system. *Frontiers in Medicine*. 2024;11:1422911. Available from: https://doi.org/10.3389/fmed.2024.1422911

[6] Jiang Q, Ma J, Yang C, Ma X, Shen J, Chaudhry SA. Efficient end-to-end authentication protocol for wearable health monitoring systems. Computers & Electrical Engineering. 2017;63:182–195. Available from: https://doi.org/10.1016/j.compeleceng.2017.03.016

[7] Avanzi R, Bos J, Ducas L, Kiltz E, Lepoint T, Lyubashevsky V, et al. Module-lattice-based key-encapsulation mechanism standard. 2024 Aug. Available from: https://doi.org/10.6028/nist.fips.203

[8] Avanzi R, Bos J, Ducas L, Kiltz E, Lepoint T, Lyubashevsky V, et al. Module-lattice-based digital signature standard. 2024 Aug. Available from: https://doi.org/10.6028/nist.fips.204

[9] Bos J, Ducas L, Kiltz E, Lepoint T, Lyubashevsky V, Schanck JM, et al. CRYSTALS - Kyber: a CCA-Secure Module-Lattice-Based KEM. *IEEE European Symposium on Security and Privacy*. 2018 Apr 1;353–67. Available from: https://doi.org/10.1109/eurosp.2018.00032

[10] Ducas L, Kiltz E, Lepoint T, Lyubashevsky V, Schwabe P, Seiler G, et al. CRYSTALS-Dilithium: a Lattice-Based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*. 2018 Feb 14;238–68. Available from: https://doi.org/10.46586/tches.v2018.i1.238-268

[11] Dworkin MJ. Recommendation for block cipher modes of operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication 800-38D. Gaithersburg (MD): NIST; 2007. Available from: https://doi.org/10.6028/NIST.SP.800-38D

[12] Krawczyk H, Eronen P. HMAC-based Extract-and-Expand Key Derivation Function (HKDF). RFC 5869. *Internet Engineering Task Force*; 2010. Available from: https://doi.org/10.17487/RFC5869

[13] Breiman L. Random forests. *Machine Learning*. 2001;45(1):5–32. Available from: https://doi.org/10.1023/A:1010933404324

[14] Chandra B, S KSR. Security in Wireless Body Area Network (WBAN) using blockchain. *IETE Journal of Research*. 2023 Jul 20;70(1):487–98. Available from: https://doi.org/10.1080/03772063.2023.2233472

[15] Sanjith S, Thangaiah PRJ, Navamani JMA, Venkataramana A. Integration of blockchain technology for security and privacy enhancement in wireless body area network systems. *Indian Journal of Science and Technology*. 2023 Nov 12;16(41):3583–90. Available from: https://doi.org/10.17485/ijst/v16i41.1610

[16] Singh AK, Kumar S. An efficient and secure CLAKA protocol for blockchain-aided wireless body area networks. *Expert Systems With Applications*. 2024;242:122740. Available from: https://doi.org/10.1016/j.eswa.2023.122740

[17] Bahache AN, Chikouche N, Akleylek S. Securing cloud-based healthcare applications with a quantum-resistant authentication and key agreement framework. *Internet Things*. 2024;26:101200. Available from: https://doi.org/10.1016/j.iot.2024.101200

[18] Mansoor K, Afzal M, Iqbal W, Abbas Y, Mussiraliyeva S, Chehri A. PQCAIE: Post quantum cryptographic authentication scheme for IoT-based e-health systems. *Internet of Things*. 2024 May 23;27:101228. Available from: https://doi.org/10.1016/j.iot.2024.101228

[19] SaberiKamarposhti M, Ng KW, Chua FF, Abdullah J, Yadollahi M, Moradi M, et al. Post-quantum healthcare: a roadmap for cybersecurity resilience in medical data. *Heliyon*. 2024;10(10):e31406. Available from: https://doi.org/10.1016/j.heliyon.2024.e31406

[20] Si-Ahmed A, Al-Garadi MA, Boustia N. Survey of Machine Learning based intrusion detection methods for Internet of Medical Things. *Applied Soft Computing*. 2023 Mar 22;140:110227. Available from: https://doi.org/10.1016/j.asoc.2023.110227

[21] Alsolami T, Alsharif B, Ilyas M. Enhancing cybersecurity in healthcare: evaluating ensemble learning models for intrusion detection in the Internet of Medical Things. *Sensors*. 2024;24(18):5937. Available from: https://doi.org/10.3390/s24185937

[22] Shaikh JA, Wang C, Muhammad WUS, Arshad MA, Owais M, Alnashwan RO, et al. RCLNet: an effective anomaly-based intrusion detection for securing the IoMT system. Front Digit Health. 2024;6:1467241. Available from: https://doi.org/10.3

389/fdgth.2024.1467241

[23] Balhareth G, Ilyas M. Optimized intrusion detection for IoMT networks with tree-based machine learning and filter-based feature selection. *Sensors*. 2024;24(17):5712. Available from: https://doi.org/10.3390/s24175712

[24] Khan A, Rizwan M, Bagdasar O, Alabdulatif A, Alamro S, Alnajim A. Deep learning-driven anomaly detection for IoMT-based smart healthcare systems. *Computer Modeling in Engineering & Sciences*. 2024;141(3):2121–2141. Available from: https://doi.org/10.32604/cmes.2024.054380

[25] Doménech Fons J, Martín Faus IV, Mhiri S, Pegueroles Vallés JR. Ensuring patient safety in IoMT: a systematic literature review of behavior-based intrusion detection systems. Internet Things. 2024;28:101420. Available from: https://doi.org/10.1016/j.iot.2024.101420

[26] Sabrina F, Sohail S, Tariq UU. A review of post-quantum privacy preservation for IoMT using blockchain. *Electronics*. 2024;13(15):2962. Available from: https://doi.org/10.3390/electronics13152962

[27] Adil M, Ali A, Tin TT, Abulkasim H, Farouk A, Al-Kuwari S, et al. Quantum Computing and the future of healthcare Internet of things security: Challenges and opportunities. *IEEE Internet of Things Journal*. 2025 Sep 1;12(22):46316–46. Available from: https://doi.org/10.1109/jiot.2025.3605040

[28] Chaudhary AK, Patel RB, Jatav DS, Patel A, Mogili VB. IoT Based Deep Learning Framework for Continuous Healthcare Monitoring of Vital Signs. In *2025 International Conference on Intelligent and Secure Engineering Solutions (CISES)* 2025 Aug 11 (pp. 1089-1094). IEEE. Available from: https://doi.org/10.1109/CISES66934.2025.11265584