PDNet-IDS26: A Healthcare 5.0 Multiclass Intrusion Detection Dataset with Biomedical and Network Features

Pedro H. Lui[1], Lucas P. Siqueira[1], Juliano F. Kazienko[1], Silvio E. Quincozes[3] Vagner E. Quincozes[2], Daniel Welfer[1], Douglas Rodrigues Fideles[3], and Diego Luis Kreutz[3]

[1]Universidade Federal de Santa Maria (UFSM), Santa Maria – RS, Brasil

[2]Universidade Federal Fluminense (UFF), Niterói – RJ, Brasil

[3]Universidade Federal do Pampa (UNIPAMPA), Alegrete – RS, Brasil

{pedro.lui, kazienko}@redes.ufsm.br, lucas.pittella@acad.ufsm.br

vequincozes@id.uff.br, daniel.welfer@ufsm.br

{silvioquincozes, douglasfideles.aluno, diegokreutz}@unipampa.edu.br

Abstract. *Intrusion Detection Systems (IDS) are central to securing Healthcare 5.0 environments, yet existing IDS datasets rely on network features only and exclude patient clinical data. This gap is critical in telemonitoring, where evolving physiological signals define patient-specific baselines that could enhance context-aware intrusion detection. This work introduces the Parkinson's Disease with Network Features – Intrusion Detection Systems 2026 (PDNet-IDS26), a multiclass dataset that merges network traffic with Parkinson's Disease (PD) telemonitoring records. We integrated literature-derived physiological features into network captures to enable unified intrusion detection. Results involving the merged dataset indicate that Replay and False Data Injection attacks (FDI) overcome other scenarios, achieving average F1-scores of 93% and 77%, respectively. Explainability results show that the network (Length) and biomedical (Age) features are critical for identifying FDI attacks.*

1. Introduction

The emergence of Healthcare 5.0 represents a transformative shift toward intelligent, human-centered care, evolving from the mass customization of Industry 4.0 toward mass personalization and patient wellness [Al-Hawawreh and Hossain 2025]. By integrating Artificial Intelligence (AI) and the Internet of Medical Things (IoMT), this paradigm enables continuous monitoring, personalized medicine, and predictive diagnostics through a vast network of interconnected devices [Almalki et al. 2024, Siqueira et al. 2025]. However, this rapid integration into clinical workflows significantly expands the digital attack surface of critical healthcare infrastructure. As these ecosystems become "always-on" and increasingly autonomous, Intrusion Detection Systems (IDS) have become the primary line of defense to safeguard sensitive medical data and ensure uninterrupted clinical operations [Almobaideen et al. 2025].

Recent work in IDS demonstrates that incorporating historical and relational contextual features improves detection performance over purely flow-based approaches [Chen et al. 2025]. Nevertheless, the systematic evaluation of such models in Healthcare 5.0 environments remains limited. This gap is evident in Parkinson's Disease (PD) management—affecting nearly 6 million people worldwide—where telehealth

effectively monitors motor and cognitive symptoms [Sun et al. 2026], yet benchmark datasets such as CICIoMT2024 [Dadkhah et al. 2024] focus primarily on generic network attributes and omit disease-relevant physiological telemetry (e.g., tremor or gait freezing patterns). Although integrating network flows with biomedical data has shown improved detection performance [Hady et al. 2020], existing approaches typically model these signals in isolation and do not incorporate disease-specific clinical context or longitudinal patient data (i.e., repeated time-series records collected over extended periods). As a result, Healthcare 5.0 lacks datasets that jointly represent network traffic and chronic disease evolution, limiting the ability of intelligent systems to assess whether network behavior aligns with patient-specific physiological states.

Existing benchmarks, such as WUSTL-EHMS-2020 [Hady et al. 2020] and ECU-IoHT [Ahmed et al. 2021], remain limited in scale, attack diversity, and disease-specific depth. In Healthcare 5.0 environments, connected medical devices generate longitudinal physiological data that establish patient-specific behavioral baselines. However, most datasets treat these signals as independent samples, disregarding the temporal dynamics that distinguish normal disease progression from malicious anomalies. Although recent datasets such as CICIoMT2024 [Dadkhah et al. 2024] and WUST-HDRL-2024 [Ghubaish et al. 2024] incorporate IoMT traffic, they lack detailed integration of patient-specific features. Likewise, DDoS-focused datasets [Akhi et al. 2025] emphasize network-layer characteristics without modeling time-evolving clinical states. Consequently, current IDS benchmarks do not enable verification of whether observed network activity is clinically consistent with a patient's physiological trajectory.

To address these gaps, this work introduces the *Parkinson Disease with Network Features - Intrusion Detection Dataset 2026* (*PDNet-IDS26*), a novel multiclass IDS dataset that enables context-aware cybersecurity by fusing network telemetry with high-fidelity longitudinal biomedical data. Unlike existing works that focus exclusively on network-layer activity, our dataset provides a disease-aware and diagnosis-specific clinical context by integrating the "Parkinson's Telemonitoring" dataset [Tsanas et al. 2010] into a dedicated IoMT testbed. By transmitting these clinical records to simultaneously capture raw network traffic and the actual biomedical patient data, *PDNet-IDS26* explicitly includes the underlying patient data within the flow records. This distinguishes our contribution from previous healthcare IDS datasets that omit the clinical features necessary for semantic verification. By aligning low-level network headers with physiological features, we enable AI-driven systems to detect subtle data manipulations that bypass conventional packet-inspection methods. Using XGBoost, KNN, and Random Forest classifiers, the merged dataset achieves average F1-scores of 93% for Replay and 77% for False Data Injection (FDI) attacks, with explainability highlighting network (*Length*) and biomedical (*Age*) as critical features for FDI detection.

The remainder of this paper is organized as follows. Section 2 reviews recent progress and examines the limitations of existing datasets. Section 3 describes the experimental testbed and the fusion process with clinical healthcare data. Section 4 presents performance results of state-of-the-art AI models on the proposed dataset. Finally, Section 5 concludes the paper and outlines directions for future work.

2. Related Works

Healthcare cybersecurity has seen many datasets aimed at addressing vulnerabilities in connected medical systems; however, most lack contextual features that reflect real clinical conditions. Work outside healthcare shows the benefit of combining cyber and physical data for security analysis. For example, cyber-physical system datasets such as the Secure Water Treatment (SWaT) testbed integrate network traffic with physical process measurements to study attacks affecting both communication and system behavior [Goh et al. 2016]. Similarly, research on smart grids studies the semantic detection of FDI attacks by analyzing inconsistencies between system measurements and communication data [Alwaisi and Soderi 2026]. These works show that correlating data from multiple domains improves the detection of complex attacks, motivating the integration of network and biomedical signals in IoMT environments.

The Table 1 compares healthcare security datasets. *Data Type* indicates whether datasets contain only network traffic or both network and biomedical data. *H5.0 Alignment* is classified as Weak (basic healthcare context), Partial (IoMT monitoring without integrated biomedical data), or Strong (human-centric monitoring with integrated biomedical data). *Disease-Specific* denotes datasets targeting a particular pathology. *Devices* represents the number of network devices in the testbed, and *Protocols* indicates whether the dataset uses a single communication protocol (Specific) or multiple protocols (Multiple).

Table 1. Comparison of Recent Healthcare Security Datasets.

Dataset	Data Type	H5.0 Alignment	Disease-Specific	Devices	Protocols
WUSTL-EHMS-2020 [Hady et al. 2020]	Net + Biomedical	Strong	No	7	Specific
ECU-IoHT [Ahmed et al. 2021]	Network Only	Partial	No	9	Multiple
BlueTack [Zubair et al. 2022]	Network Only	Partial	No	5	Multiple
CICIoMT2024 [Dadkhah et al. 2024]	Network Only	Partial	No	40	Multiple
WUST-HDRL-2024 [Ghubaish et al. 2024]	Network Only	Partial	No	9	Multiple
IoMT-TrafficData [Areia et al. 2024]	Network Only	Partial	No	11	Multiple
UL-ECE-MQTT-DDoS-2025 [Akhi et al. 2025]	Network Only	Partial	No	23	Multiple
UL-ECE-UDP-DDoS-2025 [Akhi et al. 2025]	Network Only	Partial	No	23	Multiple
MedSec-25 [Almobaideen et al. 2025]	Network Only	Partial	No	5	Multiple
Proposed Work	**Net + Biomedical**	**Strong**	**Yes**	5	Specific

As medical environments adopted diverse communication standards, datasets like CICIoMT2024 [Dadkhah et al. 2024] and IoMT-TrafficData [Areia et al. 2024] emerged to cover protocols such as MQTT and Bluetooth Low Energy (BLE). CICIoMT2024 is notable for profiling the entire lifecycle of 40 IoMT devices across various operational states, providing a robust baseline for protocol-centric threats like MQTT malformed data. Additionally, specialized datasets like BlueTack [Zubair et al. 2022] address the specific vulnerabilities of Bluetooth-enabled medical devices, providing the first dedicated resource for detecting intrusions in Classic Bluetooth and BLE healthcare networks. However, while these datasets capture the behavioral nuances of the hardware, they remain decoupled from the actual biomedical data, such as those found in Parkinson's Disease telemonitoring, leaving a gap in detecting sophisticated semantic-layer attacks.

WUSTL-EHMS-2020 [Hady et al. 2020] uses a real-time testbed capturing 35 network and 8 biomedical features, showing that modality fusion improves detection of data alteration and spoofing attacks. Similarly, ECU-IoHT [Ahmed et al. 2021] simulates attacks on a healthcare kit, enabling analysis of how network threats such as ARP spoofing

correlate with deviations in physiological readings. However, while these datasets integrate biomedical data within healthcare environments, they lack the longitudinal clinical depth needed to distinguish malicious manipulation from natural symptom progression in chronic conditions like PD. Consequently, a gap remains for diagnosis-specific systems that treat biomedical data as a continuous clinical narrative rather than a parallel feature.

The WUST-HDRL-2024 framework [Ghubaish et al. 2024] employs hybrid deep reinforcement learning within an emulated 5G testbed but focuses solely on network and host-level abstractions, omitting physiological telemetry required for context-aware detection. Similarly, the UL-ECE-2025 datasets [Akhi et al. 2025] simulate MQTT-enabled medical sensors (e.g., temperature, heart rate, oxygen saturation) in Cooja and ns-3 environments; however, the resulting datasets retain only network-layer attributes, discarding the biomedical signals that would enable longitudinal or disease-specific analysis. A comparable limitation appears in the MedSec-25 dataset [Almobaideen et al. 2025], which simulates multi-stage IoMT attack campaigns in a physical healthcare IoT lab with biomedical sensors but ultimately releases only processed network-flow features. Consequently, while these datasets effectively capture disruptions in IoMT communications, they lack the biomedical context required to determine whether transmitted clinical values remain physiologically consistent with a patient's condition.

3. Methodology

This section details the methodology of the proposed PDNet-IDS26 dataset tailored for IoMT environments. We describe the physical testbed topology, the scenario-based data acquisition pipeline, the threat modeling of specific attacks targeting the MQTT protocol, the merged feature extraction process that fuses network traffic behavior with biomedical payloads and discuss the threats to validity.

Figure 1 illustrates the end-to-end data flow and high-level conceptual architecture of the proposed IoMT system. The diagram introduces a *Malicious Intrusion* element, demonstrating that the physical testbed was subjected to four distinct *Data Collection* experiments: a benign baseline and three targeted cyberattack scenarios: FDI, Denial-of-Service (DoS) and Replay. Finally, the architecture outlines the dataset generation pipeline, illustrating how raw network traffic undergoes *Data Extraction* and *Data Cleaning* to produce the final *Merged Dataset*: the PDNet-IDS26. The dataset is publicly available in Kaggle to facilitate the results reproducibility[1]. This refined dataset consolidates both network telemetry and biomedical payloads, serving as the input for the ML evaluation in Section 4.

3.1. Biomedical Dataset Selection and Characteristics

To simulate realistic clinical traffic within the proposed testbed, the *Parkinson's Telemonitoring*[2] dataset was selected as the source of the biomedical payloads. Originally collected during a six-month clinical monitoring, as reported by Tsanas et al. [Tsanas et al. 2010]. Such a dataset comprises 5,875 instances of biomedical voice measurements extracted from 42 patients in the early stages of Parkinson's disease. The original data was gathered through a remote telemonitoring device deployed directly in patients' homes,

[1] https://www.kaggle.com/datasets/pedrohenriquelui/pdnet-ids26

[2] https://archive.ics.uci.edu/dataset/189/parkinsons+telemonitoring

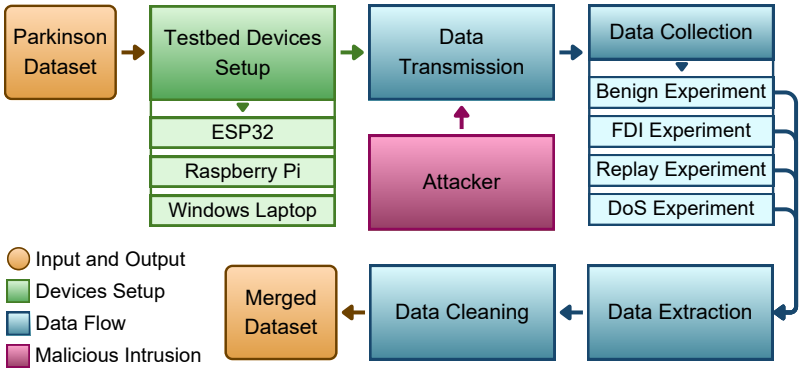

Figure 1. High-Level Conceptual Architecture of the Proposed IoMT Intrusion Detection System.

continuously transmitting biomedical signals over the internet. This aligns perfectly with the Healthcare 5.0 paradigm of digitalization and personalization of treatments.

Each transmitted sample contains 22 clinical and acoustic features, which include fundamental frequency perturbations (e.g., *Jitter*, *Shimmer*), noise-to-harmonics ratios (*NHR*, *HNR*), and nonlinear dynamical complexity measures (*RPDE*, *DFA*, *PPE*). These acoustic markers are mapped to predict the progression of the disease, quantified by the clinician-scored Unified Parkinson's Disease Rating Scale (Motor UPDRS and Total UPDRS). The UPDRS score dictates the clinical intervention and medication dosage (e.g., Levodopa) prescribed to Parkinson's patients. By using real, critical diagnostic data, the impact of the threat modeling becomes tangible. For instance, a successful FDI or Replay attack on these specific packets would seamlessly alter the UPDRS variables without breaking the network connection, directly leading to an erroneous deterioration assessment and potentially lethal medical decisions.

Thus, fusing this highly sensitive biomedical dataset with raw network traffic features provides the ideal environment to train the Machine Learning (ML) model to detect anomalies that threaten both network integrity and patient safety.

3.2. Testbed Architecture

The network topology relies on a centralized star architecture using the MQTT protocol, selected for its lightweight overhead and dominance in IoT ecosystems [Quincozes et al. 2019]. The MQTT protocol operates on a publish/subscribe model, where devices do not communicate directly with each other, but rather through a central broker that filters and routes messages based on topics. As illustrated in Figure 2, the testbed comprises four primary nodes:

1. **Publisher:** An ESP32 microcontroller configured to publish biomedical data, simulating a wearable Parkinson's disease monitor. It utilizes the SPIFFS/LittleFS file system to store and stream temporal data derived from the Parkinson's Telemonitoring dataset [Tsanas et al. 2010].
2. **Broker:** A Raspberry Pi hosting the Mosquitto MQTT Broker. This node acts as the central intermediary, managing the publish/subscribe messaging flow between the perception layer and the application layer. The broker's sole responsibility is to receive messages and distribute them to the actively subscribed clients.

3. **Subscriber:** A Windows-based laptop running a custom Python subscriber script designed to receive, process, and log the biomedical data routed by the broker.
4. **Attacker:** The same physical Windows laptop hosting the subscriber node, but concurrently executing various Python attack scripts to inject malicious traffic and manipulate the network environment. This architectural co-location explicitly models a severe internal threat scenario where the adversary has already compromised the clinical monitoring station or breached the local network perimeter, thereby granting direct, privileged interaction with the broker's data flow.

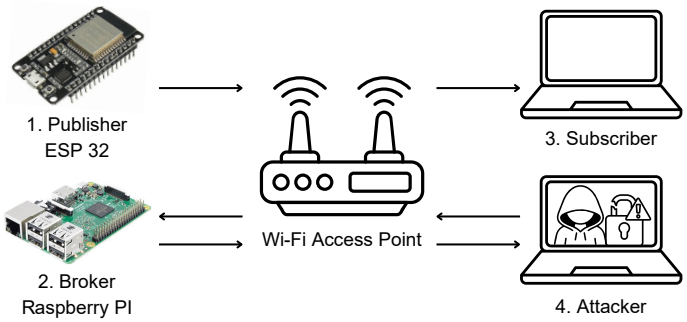

Figure 2. Testbed Architecture.

3.3. Threat and Attack Modeling

To evaluate the system's robustness, an adversarial model was implemented focusing on the core cybersecurity paradigms of availability, integrity, and confidentiality. These attacks were chosen because they represent the most critical threats in healthcare IoT, where compromised or delayed data can lead to severe real-world clinical consequences.

Figure 3 delineates the pipeline of our methodology, encompassing four main stages: (1) the physical deployment of the IoMT testbed, (2) the execution of the threat models via socket programming and the `paho-mqtt` library, (3) the packet-level data capture and lossless extraction, and (4) the data preprocessing and feature fusion that feeds the explainable XGBoost classifier. The implemented attacks are described below:

- **FDI:** This attack positions itself logically between the publisher and the broker, acting as a Man-in-the-Middle (MitM). Using a TCP Proxy approach, the attacker intercepts the direct connection from the ESP32, performs stealthy, on-the-fly modifications to the biomedical payload, and forwards the malicious payload to the legitimate broker. To emulate a highly sophisticated adversary that avoids simple threshold-based detection, the script applies a variance ranging from $\pm 15\%$ to $\pm 40\%$ to critical clinical variables (specifically targeting `motor_UPDRS`, `total_UPDRS`, and vocal metrics such as `HNR`). This randomized perturbation ensures the corrupted values fall within a plausible physiological range for a deteriorating Parkinson's patient. The FDI attack aims to deceive monitoring systems and alter medical diagnostics without disrupting the availability of the network, exploiting data integrity vulnerabilities frequently documented in Internet of Medical Things (IoMT) architectures [Elamin et al. 2024].
- **Replay:** Targets the freshness and temporal validity of the data. The adversary captures a sequence of legitimate packets representing a stable patient state and

buffers them. These packets are published again at a later time. In a clinical context, this is critical as it could mask a patient's sudden deterioration by feeding the monitoring system with outdated "healthy" metrics, a prominent threat in MQTT-based smart healthcare deployments that lack robust session freshness mechanisms [Pahlevi et al. 2021].

- **DoS:** This attack targets network availability and was modeled by flooding the broker with high-frequency MQTT messages featuring null physiological payloads. This saturates the broker's processing queue and the network bandwidth, effectively delaying or dropping the legitimate telemonitoring stream from the medical sensor, characterizing a typical MQTT publish flooding attack aimed at resource exhaustion [Vaccari et al. 2020].

3.4. Data Acquisition Pipeline

The data acquisition process mimics the behavior of a real-time telemonitoring device transmitting patient vitals. Data collection was segregated into four distinct experiments to ensure label integrity: (1) Benign, (2) FDI, (3) DoS and (4) Replay.

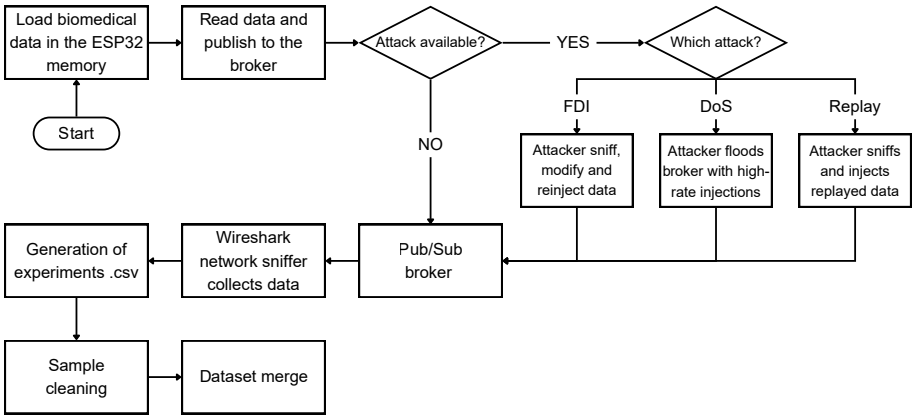

Figure 3. Dataset Generation, Feature Extraction, and ML/XAI Evaluation Flowchart.

In each experiment, the ESP32 initiates the transmission of clinical samples at a frequency of 5Hz (200ms intervals). Concurrently, the subscriber node connects to the broker, while Wireshark, running on the same laptop, sniffs all network packets using the `tcp port 1883` capture filter. To ensure the ML model evaluates the attacks from the perspective of the targeted clinical application, the traffic was strictly filtered to isolate communications arriving from the Broker to the Subscriber.

To extract the data reliably and bypass Wireshark's graphical interface limitations—which severely truncates payloads when multiple MQTT messages are aggregated into a single TCP packet by the network—the `tshark` tool was utilized, ensuring a lossless conversion of the raw network captures into `.csv` format, decoding hexadecimals and properly separating grouped payloads.

It is crucial to clarify the structure of the final proposed dataset: each row represents a single instance combining the physical network frame with its corresponding biomedical payload. The application payload is transmitted as comma-separated plain

text within the MQTT messages. To ensure accurate synchronization, a custom Python parser was developed to hex-decode the `mqtt.msg` field from the `tshark` output. For each packet, the parser reversed this encoding to recover the original string, splitting it into its 20 distinct biomedical variables. These variables were then horizontally concatenated with the corresponding network features of that exact same row.

The extraction process yielded a dataset fusing 12 network-level features with 22 application-level biomedical features. The selection of a reduced network feature set is intentional; it serves to shift the analytical focus directly onto the biomedical payload, aiming to demonstrate that semantic clinical context is the decisive factor for detecting integrity anomalies. Finally, to address the extreme class imbalance caused by the DoS flood, the resulting merged dataset was subjected to a random undersampling technique using the Pandas[3] `.sample()` function with a fixed `random_state=42`. This ensured reproducibility and established a balanced baseline for the subsequent evaluations. The final dataset consists in a distribution of 5,380 samples per class (Benign, FDI, DoS, Replay), totaling 21,520 instances. Although this balanced setting does not fully reflect real-world class distributions, we deliberately adopted it to enable a less biased and more controlled evaluation of the model's discriminative ability.

3.5. Threats to Validity

Internal validity is affected by the controlled experimental design. First, FDI attacks rely on predefined perturbations that maintain local physiological plausibility, but may fail to emulate adversaries capable of forging complex longitudinal correlations over extended periods. Second, the explicit use of undersampling to artificially balance the dataset impacts validity; while it guarantees an unbiased algorithmic comparison across classes, it inherently overestimates detection performance compared to real operational environments where cyberattacks are extremely rare. Finally, external validity is constrained by the testbed's specific focus on Parkinson's disease and the MQTT protocol, limiting immediate generalization to other medical conditions or communication standards, though PDNet-IDS26 establishes a robust baseline for future IoMT security expansion.

4. Results and Discussion

In this section, we present the experimental evaluation of the proposed PDNet-IDS26 Dataset tailored for IoMT environments. The assessment aims to validate the hypothesis that integrating biomedical context with network telemetry improves the detection of sophisticated attacks, particularly those targeting data integrity.

The results are based on the merged dataset constructed in the methodology, employing the XGBoost classifier to distinguish between benign traffic and three specific attack vectors: FDI, DoS, and Replay attacks. The discussion is organized into two primary subsections. First, in Subsection 4.1 we present a Scenario Evaluation and Performance Study, conducting an analysis to quantify the individual contributions of network flow metrics and biomedical payload features to the overall detection accuracy. Subsequently, in Subsection 4.2 we provide an Analysis of Biomedical Features Importance, utilizing SHAP (SHapley Additive exPlanations) values to interpret the model's decision-making process and empirically demonstrate the necessity of biomedical features in identifying attacks [Lundberg and Lee 2017].

[3] `https://pandas.pydata.org/docs/index.html`

4.1. Scenario Evaluation and Performance Study

To isolate the contribution of each data modality to the intrusion detection capabilities, we conducted an study comparing three distinct feature sets: `Bio-Only`, `Net-Only`, and the proposed `Merged` approach. Furthermore, to prevent "Shortcut Learning", session identifiers and patient-specific identifiers were dropped prior to training. This ensures the classifier learns genuine threat behaviors rather than memorizing isolated network sessions or individual patient baselines. The evaluation protocol was strictly based on a 5-Fold Stratified Cross-Validation. This stratified approach guarantees that each 20% fold maintains the exact class distribution (25% for each of the four classes) of the original dataset, mitigating the risk of overfitting to a specific data split and ensuring the robustness of the reported metrics. Finally, to demonstrate that the performance gains provided by the merged dataset are not dependent on a specific algorithm, the evaluation was expanded to benchmark three distinct ML classifiers: XGBoost, Random Forest (RF), and K-Nearest Neighbors (KNN). Figure 4 illustrates the mean F1-Score isolated by class (Benign, DoS, FDI, and Replay), alongside their respective 95% confidence intervals (error bars), across the 5 folds for each evaluated scenario.

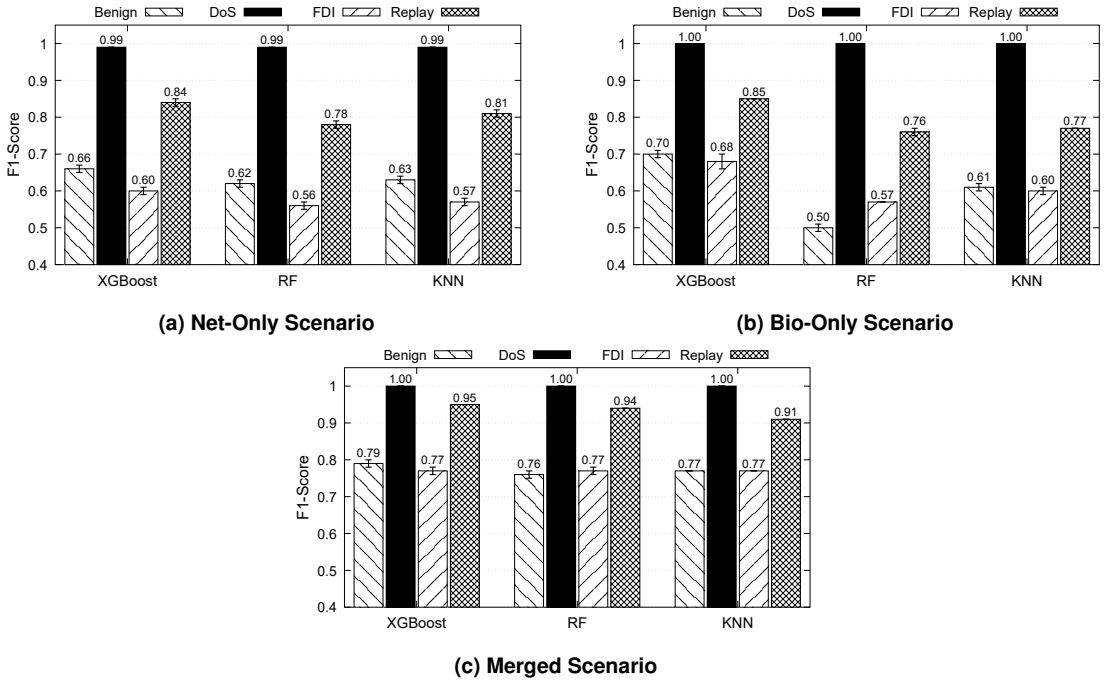

Figure 4. PDNet-IDS26 Scenario Based Performance Per Class Comparison.

As observed in Figure 4b, the `Bio-Only` model, trained exclusively on clinical payload features, exhibited the better performance for FDI attacks. The model could differentiate the manipulated FDI values from genuine physiological fluctuations even without network context, achieving an FDI F1-Score of 0.68 with XGBoost. Also, it achieved perfect detection for DoS attacks, as the flooding payload contained formatting anomalies distinct from valid medical data. The model also showed good performance in detecting Replay attacks (F1-Score 0.85), likely memorizing the exact replicated physiological values, which act as a signature in the absence of network features.

The `Net-Only` model (Figure 4a) showed good performance for identifying at-

tacks with distinct traffic patterns (DoS and Replay, with F1-Scores of 0.99 and 0.84, respectively). Nevertheless, the model's critical limitation was observed in the FDI class, where the F1-Score dropped to 0.60. Since the implemented FDI attack is a stealthy MitM proxy, the network-based classifier lacked the insight required to detect the underlying integrity violation, relying solely on subtle delays.

The `Merged` model (Figure 4c) yielded the best overall performance, validating the proposed hybrid architecture. Notably, the detection of FDI improved substantially reaching an F1-Score of 0.77, while the F1-Score for the Benign class jumped to 0.79, indicating a significant reduction in false positives. This structural advantage of the `Merged` dataset proved to be algorithm-agnostic. Regardless of whether the system employed XGBoost, Random Forest, or KNN, fusing the clinical data with network telemetry universally produced the highest detection rates.

Furthermore, XGBoost consistently outperformed the other classifiers. Its gradient boosting architecture proved highly effective in capturing the complex, non-linear correlations between network flow anomalies (e.g., IAT deviations) and subtle biomedical perturbations (e.g., stochastic variations in UPDRS), solidifying it as the optimal classification engine for the proposed IoMT IDS architecture.

4.2. Analysis of Biomedical Features Importance

The results presented in Figure 4, combined with the interpretability analysis provided by SHAP [Lundberg and Lee 2017], offer empirical evidence regarding the critical role of biomedical features in detecting sophisticated integrity attacks like FDI.

While network-based features proved sufficient for identifying DoS and Replay attacks, they failed to reliably detect FDI. The ML classifiers, in the `Net-Only` scenario, achieved an average F1-Score of only 0.58 for the FDI class. This limitation occurs because the implemented FDI attack was designed to be protocol-compliant and stealthy, maintaining legitimate packet sizes and routing behaviors, thereby rendering it nearly invisible to traditional traffic analysis. In contrast, the inclusion of biomedical features in the Merged model raised the FDI detection performance significantly (F1-Score 0.78). This quantitative improvement is corroborated by the SHAP feature importance analysis. As illustrated in the SHAP summary plot (Figure 5), biomedical features such as `age` and `HNR` emerged as critical contributors for the FDI class classification, alongside network scars left by the proxy (e.g., `Length` and `TCP_WinSize`).

This findings confirms that the model learned to distinguish attacks not just by "how" the packet arrived (Network), but by "what" the packet contained (Biomedical). Specifically, the model identified that the manipulated values injected by the adversary deviated from the physiological consistency expected in genuine Parkinson's disease progression patterns. The classifier models leveraged the underlying biological manifold of Parkinson's disease to detect the FDI attacks. When the adversarial script applied variations to the payload to mimic a sudden spike in disease severity, it perturbed variables mathematically, but without respecting the complex physiological correlations shared among the vocal markers. The XGBoost classifier, having learned the authentic biological footprint of PD patients during training, recognized these decoupled alterations as mathematically possible but biologically impossible anomalies. Therefore, we conclude that semantic features act as a necessary validation layer for ensuring data integrity in

IoMT environments, preventing attackers from synthesizing medical deterioration simply by randomizing payload bytes.

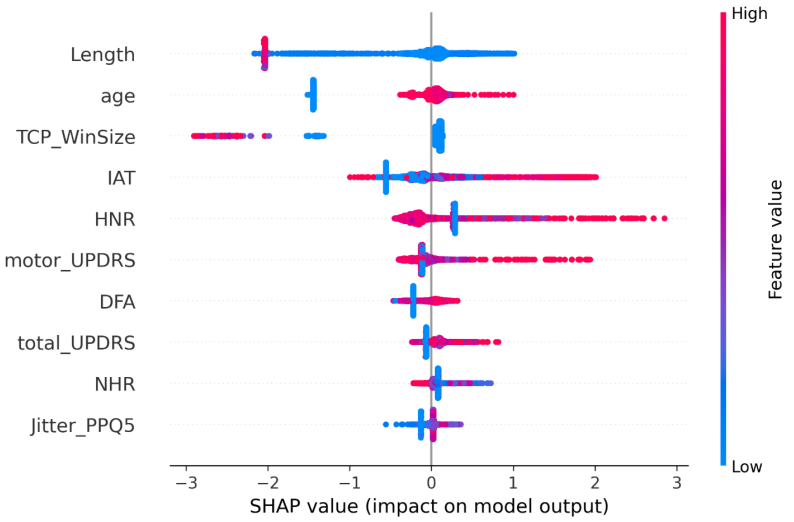

Figure 5. XGBoost SHAP Summary Plot for FDI Detection in the Merged Scenario.

5. Conclusion and Future Works

This work introduced *PDNet-IDS26*, a dataset for Healthcare 5.0 intrusion detection that integrates Parkinson's telemetry with raw network traffic. By linking clinical signals with network events, the dataset enables IDS evaluation beyond packet-level anomalies toward clinically consistent attack detection. Results show that adding biomedical features improves performance across classifiers. Replay and FDI attacks achieved average F1-scores of 93% and 77%. Explainability analysis shows that network (Length) and biomedical (Age) features are key for detecting FDI attacks. These results highlight the value of contextual data for securing telemonitoring systems. Future work will expand the dataset and explore synthetic data and sampling strategies (oversampling and undersampling).

Acknowledgments

The authors would like to thank the *Fundação de Amparo à Pesquisa do Estado do Rio Grande do Sul* (FAPERGS) and the UFSM/PRPGP *Edital N.50/2024 - Programa de Fortalecimento e Redução de Assimetrias da Pós-Graduação da UFSM* for their financial support to this research effort.

References

Ahmed, M., Byreddy, S., Nutakki, A., Sikos, L. F., and Haskell-Dowland, P. (2021). Ecu-ioht: A dataset for analyzing cyberattacks in internet of health things. *Ad Hoc Networks*, 122:102621.

Akhi, M., Eising, C., and Dhirani, L. L. (2025). Datasets for distributed denial-of-service detection in healthcare internet of things environments. *Data in Brief*, 63:112222.

Al-Hawawreh, M. and Hossain, M. S. (2025). A human-centered quantum machine learning framework for attack detection in iot-based healthcare industry 5.0. *IEEE Internet of Things Journal*, 12(22):46065–46074.

Almalki, J., Alshahrani, S. M., and Khan, N. A. (2024). A comprehensive secure system enabling healthcare 5.0 using federated learning, intrusion detection and blockchain. *PeerJ Computer Science*, 10:e1778.

Almobaideen, W., Abdullah, M., Alam, U., Hussain, S. B., and Bouharrat, A. (2025). Medsec-25: Creating an iomt dataset for a healthcare iot environment. In *2025 7th International Conference on Blockchain Computing and Applications (BCCA)*, pages 628–634. IEEE.

Alwaisi, Z. and Soderi, S. (2026). Semantic communication-based detection of false data injection attacks in 6g-enabled smart grids. *International Journal of Electrical Power & Energy Systems*, 175:111649.

Areia, J., Bispo, I. A., Santos, L., and Costa, R. L. d. C. (2024). Iomt-trafficdata: Dataset and tools for benchmarking intrusion detection in internet of medical things. *IEEE Access*, 12:115370–115385.

Chen, Z., Zou, H., Hu, T., Yuan, X., Fang, X., Pan, Y., and Li, J. (2025). Hc-nids: Historical contextual information based network intrusion detection system in internet of things. *Computers & Security*, 152:104367.

Dadkhah, S., Neto, E. C. P., Ferreira, R., Molokwu, R. C., Sadeghi, S., and Ghorbani, A. A. (2024). Ciciomt2024: A benchmark dataset for multi-protocol security assessment in iomt. *Internet of Things*, 28:101351.

Elamin, U. M. B. E. et al. (2024). Security analysis for smart healthcare systems. *Sensors*, 24(11):3375.

Ghubaish, A., Yang, Z., and Jain, R. (2024). Hdrl-ids: a hybrid deep reinforcement learning intrusion detection system for enhancing the security of medical applications in 5g networks. In *2024 International Conference on Smart Applications, Communications and Networking (SmartNets)*, pages 1–6. IEEE.

Goh, J., Adepu, S., Junejo, K. N., and Mathur, A. (2016). A dataset to support research in the design of secure water treatment systems. In *International conference on critical information infrastructures security*, pages 88–99. Springer.

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

Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *31st International Conference on Neural Information Processing Systems*, NIPS'17, page 4768–4777.

Pahlevi, R. R., Sukarno, P., and Erfianto, B. (2021). Secure mqtt puf-based key exchange protocol for smart healthcare. *Jurnal Rekayasa Elektrika*, 17(2).

Quincozes, S., Emilio, T., and Kazienko, J. F. (2019). MQTT protocol: fundamentals, tools and future directions. *IEEE Latin America Transactions*, 17(9):1439–1448.

Siqueira, L. P., Batista, C. L., Lui, P. H., Kazienko, J. F., Quincozes, S. E., Quincozes, V. E., Welfer, D., and Nomura, S. (2025). A comprehensive survey on intrusion detection systems for healthcare 5.0: Concepts, challenges, and practical applications. *Sensors*, 25(20):6261.

Sun, M., Tang, F., Wen, S., Wang, S., Jiang, H., et al. (2026). Effects of telehealth interventions for people with parkinson disease: Systematic review and meta-analysis of randomized controlled trials. *JMIR mHealth and uHealth*, 14(1):e70994.

Tsanas, A., Little, M. A., McSharry, P. E., and Ramig, L. O. (2010). Accurate telemonitoring of parkinson's disease progression by noninvasive speech tests. *IEEE Transactions on Biomedical Engineering*, 57(4):884–893.

Vaccari, I., Aiello, M., and Cambiaso, E. (2020). MQTTset, a new dataset for machine learning techniques on MQTT. *Sensors*, 20(22):6578.

Zubair, M., Ghubaish, A., Unal, D., Al-Ali, A., Reimann, T., Alinier, G., Hammoudeh, M., and Qadir, J. (2022). Secure bluetooth communication in smart healthcare systems: A novel community dataset and intrusion detection system. *Sensors*, 22(21):8280.