

Machine Learning-Based Wi-Fi RSS Fingerprinting for Indoor Zone Classification of Biomedical Assets in Hospitals

Leonardo A. Bermeo Varón , Juan Pablo Mañunga-Mayor , and Gabriel E. Mejía-Ruiz 

Abstract—The frequent relocation of biomedical equipment in healthcare facilities delays maintenance and reduces operational efficiency. While dedicated positioning infrastructure can mitigate these delays, its deployment often incurs high financial costs. This work proposes an indoor zone classification system that leverages existing Wi-Fi received signal strength (RSS) fingerprints and machine learning. The methodology evaluates five supervised classifiers (Naive Bayes, K-Nearest Neighbors, Support Vector Machine, Random Forest, and Multilayer Perceptron) using an empirical dataset of 360 observations. These RSS measurements were collected by an ESP8266-based sensing module across six functional zones covered by four access points. Under a 5-fold stratified cross-validation protocol, the Random Forest model achieved a mean validation accuracy of 97.57%. These results demonstrate that existing Wi-Fi infrastructure combined with low-cost edge sensing provides a viable solution for biomedical asset tracking in hospitals.

Link to graphical and video abstracts, and to code:
<https://latam.ieee9.org/index.php/transactions/article/view/10843>

Index Terms—Wi-Fi RSS fingerprinting, indoor zone classification, biomedical asset tracking, healthcare facilities, machine learning.

I. INTRODUCTION

ACCURATE indoor localization is essential for biomedical asset tracking in healthcare facilities, where the location of portable equipment directly impacts operational continuity. Structural signal attenuation and multipath propagation preclude the use of Global Navigation Satellite Systems inside buildings [1], [2]. Consequently, wireless frameworks based on Wi-Fi, Bluetooth Low Energy (BLE), Radio Frequency Identification (RFID), and Ultra-Wideband (UWB) have been extensively investigated [3]. Among these, Wi-Fi Received Signal Strength (RSS) fingerprinting leverages existing access point infrastructure, reducing the capital expenditure typically associated with dedicated hardware [4], [5].

In healthcare institutions, preventive maintenance demands timely access to biomedical devices. However, assets are frequently relocated without centralized notification, causing

workflow disruptions [6], [7]. While RFID architectures have been proposed for tracking [8], [9], [10], [11], they require proximity readers and manual search procedures. A zone-level classification strategy capable of automatically identifying the functional zone of a device provides sufficient granularity for this task: preventive maintenance workflows require knowing which department holds the equipment, not its exact coordinates within a room [7]. This approach avoids the complexity of continuous coordinate estimation while directly addressing the operational search problem.

Wi-Fi RSS fingerprinting relies on constructing radio maps from measurements collected at reference coordinates, subsequently applying machine learning algorithms during the online phase [2], [12]. Conventional machine learning classifiers provide a well-documented baseline for this task. For instance, Attar et al. evaluated KNN, SVM, Random Forest (RF), Decision Trees, and Logistic Regression models across multiple datasets [3]. Similarly, Maduranga et al. validated classification accuracies near 85% using KNN and SVM over BLE RSSI signals [13]. This classical approach is supported by experimental evidence involving improved RSSI distance models [14], feature-engineered RSS fingerprinting [4], [5], [15], and multi-channel classification networks [12], [16].

Recent fingerprinting frameworks prioritize deep learning architectures to improve accuracy. Cheng et al. implemented a Stacked Autoencoder achieving 92.97% floor prediction accuracy over a multi-building benchmark [17]. Tang et al. integrated Multi-Output Gaussian Processes for database augmentation [18], while Yoo utilized Generative Adversarial Networks for feature synthesis [19]. Transformer-based models have also reported high accuracy on multi-floor datasets [20]. However, these topologies demand extensive hyperparameter tuning and high-performance processing nodes. Because models trained for a specific layout exhibit performance degradation when transferred to a different building without re-fingerprinting [21], deep architectures limit practical redeployment. Related benchmarks by Bi et al. [22] and Alhmiedat [23] have standardized this evaluation gap. To provide an empirical neural-network baseline under the same experimental conditions, this study evaluates a Multilayer Perceptron (MLP) with one hidden layer using the same RSS dataset, preprocessing pipeline, hyperparameter-search procedure, and stratified 5-fold cross-validation protocol applied to the other classifiers. The MLP is included as a shallow neural-network baseline rather than as a representative evaluation of deep-learning architectures. As reported in Section III-B, the MLP

The associate editor coordinating the review of this manuscript and approving it for publication was Mariel Alfaro-Ponce (*Corresponding author: Leonardo Antonio Bermeo Varón*).

Leonardo Antonio Bermeo Varón, and G. E. Mejía-Ruiz are with the School of Electrical and Electronics Engineering, Faculty of Engineering, Universidad del Valle, Cali, Colombia (e-mails: bermeo.leonardo@correounivalle.edu.co, and gabriel.mejia.ruiz@correounivalle.edu.co).

J. P. Mañunga-Mayor is with Bioengineering, Faculty of Engineering, Universidad Santiago de Cali, Cali, Colombia (e-mail: juan.manunga00@usc.edu.co).

achieved a mean validation accuracy of $95.13\% \pm 2.77\%$, whereas the Random Forest model achieved $97.57\% \pm 2.21\%$. Therefore, for the evaluated dataset and model configurations, RF provided the highest validation accuracy without requiring a multilayer deep-learning architecture.

Despite the accuracy of data-intensive algorithms, there is limited documentation of zone-level classification systems that operate on existing Wi-Fi infrastructure without dedicated positioning hardware. Most published results are evaluated on public multi-building benchmarks, dedicated BLE beacon deployments, or dense Wi-Fi infrastructures. For example, Narasimman *et al.* [21] evaluated a transferable framework on UJIIndoorLoc, which contains 520 available APs, whereas Maduranga *et al.* [13] used 13 BLE iBeacon transmitters in a university-library environment. These studies address different localization objectives and operate under substantially different infrastructure and dataset conditions from the hospital zone-classification setting considered here.

The main contributions of this work are:

- 1) A comparative evaluation of five supervised classifiers for indoor zone classification using only four existing Wi-Fi APs. Under the evaluated configurations, the RF model achieved the highest mean validation accuracy, $97.57\% \pm 2.21\%$, exceeding the shallow MLP baseline while requiring no dedicated positioning infrastructure.
- 2) An empirical dataset of 360 balanced RSS observations collected in a healthcare facility, providing a reference for zone-level classification under a practical hospital deployment scenario.
- 3) The operational validation and data acquisition of a battery-powered sensing node based on the ESP8266 SoC (Wemos D1 Mini), showing that zone-level tracking is feasible with commercially available hardware and standard Wi-Fi infrastructure.

The remainder of this paper is organized as follows. Section II describes the spatial constraints, sensing hardware, and classification algorithms. Section III presents the classification performance metrics and confusion matrices. Finally, Section IV discusses the findings and conclusions.

II. METHODOLOGY

A. Environment and Zone Definition

The study was conducted in a healthcare facility in Santiago de Cali, Colombia. The hospital Wi-Fi network contained additional APs beyond the four used in this study. The system was not designed as an artificial worst-case experiment. Instead, the four selected APs constituted a stable subset of the available infrastructure because they provided persistent visibility across all six functional zones during the acquisition campaign. APs with intermittent detectability, missing observations in one or more zones, or severe attenuation caused by structural shielding were excluded because they would not provide a reliable and reproducible fingerprinting feature set. The purpose was to assess whether a practical zone-classification system could operate using a small subset of pre-existing APs without deploying dedicated positioning hardware.

Because the selected APs are not uniformly distributed across the floor, the system was formulated as a functional-zone classifier rather than as a metric coordinate estimator. Zone boundaries were defined according to the architectural separation and operational function of each hospital area, including waiting rooms, diagnostic imaging, administration, outpatient care, and densitometry. Thus, each class represents a clinically and operationally meaningful destination for maintenance personnel rather than an equal-area geometric region. Unequal zone areas and AP geometry may affect class separability. In particular, adjacent areas separated by a single wall can exhibit partially overlapping RSS signatures, whereas a structurally isolated room can produce a distinctive attenuation profile.

The operational area consists of a single floor divided into six functional zones: two waiting rooms (classes 0 and 1), a diagnostic imaging sector (class 2), an administrative office (class 3), an outpatient care unit (class 4), and a densitometry room (class 5). Fig. 1 shows the spatial layout of these zones. The locations of the four reference access points are also mapped.

B. Hardware Module

The sensing module uses a Wemos D1 Mini development board based on the ESP8266 system-on-chip (SoC). This microcontroller operates at 80 MHz and includes a 2.4 GHz IEEE 802.11 b/g/n transceiver. A lithium-polymer ion battery of 1000mAh (3.7V) and a Wemos D1 battery shield provide regulated power for portable operation. Fig. 2 illustrates the hardware components.

In operation, the ESP8266 module scans the Wi-Fi spectrum to capture the RSS (in dBm) from the four reference APs. Each scan generates a four-dimensional vector $\mathbf{x} = [RSS_1, RSS_2, RSS_3, RSS_4]$. These vectors are transmitted to a MySQL database on a local server and the communication protocol used for data transmission was MQTT. The firmware was developed using the Arduino Integrated Development Environment (IDE).

To assess the feasibility of battery-powered operation, an analytical battery-life estimate was derived from the ESP8266 current specifications and the nominal capacity of the 1000-mAh Li-Po battery integrated into the prototype. The evaluated duty cycle consists of a 5-s Wi-Fi scanning and transmission interval followed by 300 s of deep-sleep operation, corresponding to one acquisition cycle every 5 min.

The average current was estimated as

$$I_{\text{avg}} = \frac{I_{\text{scan}}t_{\text{scan}} + I_{\text{sleep}}t_{\text{sleep}}}{t_{\text{scan}} + t_{\text{sleep}}}, \quad (1)$$

where $I_{\text{scan}} = 180$ mA, $t_{\text{scan}} = 5$ s, $I_{\text{sleep}} = 3.11$ mA, and $t_{\text{sleep}} = 300$ s. This yields an estimated average current of 6.01 mA. A 15% capacity derating was applied to account for practical losses associated with battery aging and power-conversion inefficiencies, resulting in an effective battery capacity of 850 mAh.

The estimated operating time is therefore

$$T_{\text{battery}} = \frac{C_{\text{effective}}}{I_{\text{avg}}} = \frac{850 \text{ mAh}}{6.01 \text{ mA}} = 141.4 \text{ h}. \quad (2)$$




Fig. 1. Distribution of access points and functional zones (measurements in m).




Fig. 2. System architecture: The sensing module Wemos D1 Mini development board based on the ESP8266 system-on-chip (SoC). A lithium-polymer ion battery of 1000mAh (3.7V) and a Wemos D1 battery shield. Access Points, Server for processing and user interface.

equivalent to approximately 5.9 days of continuous operation under the assumed duty cycle. This estimate is based on nominal current specifications rather than direct current measurements; actual autonomy may vary with Wi-Fi signal conditions, transmission retries, battery aging, and regulator efficiency.

C. Dataset Configuration

The dataset was constructed by collecting empirical RSS measurements from the four reference access points across the six functional zones identified in the study. For each zone, signal vectors were captured at a 5-second sampling interval, yielding 60 observations per class and a final dataset of 360 balanced records. Each individual observation is structured as

a feature vector $\mathbf{x} \in \mathbb{R}^4$, where the four predictors represent the signal strength (in dBm) from AP_1 to AP_4 , while the response variable is a categorical label $y \in \{0, 1, 2, 3, 4, 5\}$ identifying the functional zone.

The data acquisition process was conducted under active healthcare operating conditions to ensure that the fingerprints capture realistic signal variability and multipath effects characteristic of hospital environments. To ensure this dynamism and avoid an idealized, static scenario, the 60 observations per zone were collected continuously during workdays with high operational activity. As a result, the dataset intrinsically incorporates the stochastic fluctuations and daily signal fading caused by the constant foot traffic of staff and patients, the intermittent opening of doors, and the movement of

TABLE I
ANALYTICAL BATTERY-LIFE ESTIMATE FOR THE ESP8266
SENSING NODE

Parameter	Value	Basis
Battery nominal capacity	1000 mAh	Li-Po battery specification
Capacity derating	15%	Allowance for battery aging and conversion losses
Effective battery capacity	850 mAh	1000×0.85
Wi-Fi scan and transmission current	180 mA	ESP8266 nominal specification
Scan and transmission duration	5 s	Defined acquisition duty cycle
Deep-sleep current	3.11 mA	ESP8266/Wemos standby specification
Deep-sleep duration	300 s	Defined acquisition duty cycle
Acquisition cycle duration	305 s	$5 + 300$ s
Estimated average current	6.01 mA	Duty-cycle weighted calculation
Estimated battery life	141.4 h; approx. 5.9 days	$850 \text{ mAh} / 6.01 \text{ mA}$

medical equipment. This acquisition procedure exposed the fingerprints to short-term variability associated with routine hospital activity, including pedestrian traffic, door operation, and equipment movement. It therefore improves the ecological relevance of the recorded dataset; however, it does not constitute an independent temporal or external validation of long-term robustness. This balanced dataset configuration is critical to prevent model bias during the training phase, ensuring that the classification metrics accurately reflect the system's ability to discriminate between all zones with equal statistical weighting.

D. Dimensional Sensitivity Analysis of the Data Set

To assess whether the available fingerprint density provides stable internal performance estimates, a sample-size sensitivity analysis based on learning curves was conducted. The purpose was to examine how the training and cross-validation performance of the evaluated classifiers changed as the number of RSS fingerprints per functional zone increased. This analysis evaluates internal statistical stability for the present hospital deployment and should not be interpreted as evidence of direct generalization to other buildings, Wi-Fi infrastructures, sensing devices, or long-term operating conditions.

For each sample-size level, stratified subsets containing 20%, 40%, 60%, 80%, and 100% of the 60 observations available per zone were drawn from the complete dataset. The

resulting subsets contained 12, 24, 36, 48, and 60 observations per class, respectively. Each subset was evaluated using the same stratified 5-fold cross-validation protocol adopted for the main benchmark. Within every fold, preprocessing operations were fitted exclusively on the corresponding training partition and then applied to the validation partition.

For each classifier and sample-size level, the mean training accuracy and the mean validation accuracy across the cross-validation folds were calculated. The corresponding learning curves were used to identify whether validation performance stabilized as the number of fingerprints per zone increased and to quantify the training-validation gap. A small gap and stable validation performance indicate consistency under resampling from the present dataset, but do not replace an external validation campaign collected at different times or in other hospital environments.

E. Preprocessing

To ensure numerical stability and optimal classifier performance, the raw RSS measurements were subjected to a preprocessing phase. The target variable was encoded as a categorical factor with six levels corresponding to the functional zones. Scale-sensitive algorithms, specifically KNN, SVM, and MLP, require feature standardization to ensure that distance-based metrics and gradient-based optimizations are not biased by features with larger dynamic ranges. For these models, the RSS predictors were standardized using Z-score normalization to achieve a zero mean and unit standard deviation. In contrast, Random Forest (RF) and probabilistic (NB) models utilized the raw signal magnitudes, as these architectures are naturally invariant to monotonic feature scaling.

F. Classification Models

Five supervised classifiers were evaluated: Naive Bayes (NB), K-Nearest Neighbors (KNN), Support Vector Machine (SVM), Random Forest (RF), and Multilayer Perceptron (MLP). The implementation follows the pipeline in Fig. 3. Hyperparameter test values are listed in Table II.

The MLP was included as a neural-network baseline to assess whether a nonlinear learned representation provided an advantage over the conventional machine-learning models under the limited-dimensional RSS fingerprinting setting. It was not intended to constitute a comprehensive benchmark of deep-learning architectures, which would require a broader set of multilayer models and substantially larger training datasets.

The NB implementation assumes conditional independence between predictors and uses kernel density estimation and a normal distribution for the numerical input variables. The KNN algorithm with Euclidean distance was tested for $k \in \{3, 5, 7, 9, 11, 13, 15\}$. The SVM model utilized a radial basis function (RBF) kernel with regularization $C \in \{1, 5, 8, 10\}$ and width $\gamma \in \{0.001, 0.01, 0.05\}$. The hyperparameters were selected via automated grid search within a cross-validation framework.

The RF configuration used 100 decision trees ($n_{tree} = 100$) and evaluated a split-variable range $m_{try} \in \{1, 2, 3, 4\}$. The




Fig. 3. Classification pipeline for 8th Wi-Fi RSS fingerprinting system: From Data Acquisition to Output (Prediction)

MLP model used a single-hidden-layer architecture with sigmoid activation, a maximum of 100 training iterations, hidden-layer sizes $size \in \{3, 5, 7\}$, and L2 weight decay penalties $decay \in \{0, 0.001, 0.01\}$. All algorithms were trained in Python 3.13.7.

TABLE II
HYPERPARAMETER SEARCH SPACE FOR SUPERVISED CLASSIFIERS

Model	Parameters	Values
Naive Bayes	Laplace, Adjust	0, 1 (defaults)
KNN	k	3, 5, 7, 9, 11, 13, 15
SVM (RBF)	γ	0.001, 0.01, 0.05
	C	1, 5, 8, 10
Random Forest	$mtry$	1, 2, 3, 4
MLP	size	3, 5, 7
	decay	0, 0.001, 0.01

G. Evaluation Protocol

The classification performance was evaluated using a stratified 5-fold cross-validation (CV) protocol to ensure robust performance estimation and mitigate overfitting. In each iteration, 80% of the data were utilized for model training, while the remaining 20% of the data served as the held-out validation




Fig. 4. Wi-Fi RSS fingerprints across the functional zones: spatial variability and characterization of the Access Points signals

fold. To ensure the integrity of the results and prevent data leakage, feature scaling (centering and scaling) was performed internally within each cross-validation fold. This ensures that the normalization parameters were computed solely on the training subsets and then applied to the validation folds. This iterative process ensures that each observation is used for both training and testing across the folds.

Model performance is reported as the mean and standard deviation across all CV folds for three multiclass metrics: Accuracy (the ratio of correct predictions to total observations), precision (the ability to ensure that positive identifications per zone are truly correct), and the F1-score (the harmonic mean of precision and sensitivity). The standard deviation serves as a measure of model stability across different data partitions.

III. RESULTS AND DISCUSSION

A. Hardware and RSS Capture

The system architecture of the sensing module is shown in Fig. 2. The collected Wi-Fi RSS fingerprints for the six functional zones are plotted in Fig. 4. While there is observable separation between some classes, the raw signal distributions exhibit non-linear overlap, supporting the use of supervised classification.

This overlap is particularly relevant between “Waiting Room 1” (Class 1) and “Diagnostic Imaging” (Class 2), which are adjacent areas separated by a common partition wall. The RSS vectors observed from the four reference APs exhibit partially overlapping profiles in these zones, consistent with the boundary-related misclassifications observed in the validation results. In contrast, the “Densitometry Room” (Class 5) exhibits a distinctive attenuation pattern. Despite its distance from the APs, the radiation-shielded walls of this unit produce a consistently lower RSS profile that remains distinguishable from the remaining zones. This distinctive fingerprint is consistent with the perfect classification observed for this class in the reported confusion matrices.

B. Classification Performance

The training phase results are summarized in Table III. The Random Forest (RF) model ($n_{tree} = 100$, $mtry = 2$) achieved the highest mean training accuracy ($98.80\% \pm 0.32\%$) with the lowest standard deviation. Other classifiers




Fig. 5. Distribution of training accuracy across the cross-validation folds for the evaluated classifiers.

showed the following accuracies: KNN ($k = 5$, *Minkowski metric*) ($98.38\% \pm 0.34\%$), SVM ($\gamma = 0.001$) ($98.07\% \pm 0.35\%$), NB ($86.65\% \pm 1.17\%$), and MLP (decay = 0.01, hidden layer size = 3) ($95.02\% \pm 0.32\%$). The KNN, SVM and RF models achieved a precision, recall, and F1-score of more than 98%, indicating a high level of accuracy and a strong balance between the detection rate and prediction reliability across all zones.

TABLE III
TRAINING PHASE PERFORMANCE METRICS

Model	Accuracy (%)	Precision (%)	F1-Score (%)	Recall (%)
NB	86.65 ± 1.17	90.92 ± 0.46	85.42 ± 1.58	86.65 ± 1.17
KNN	98.38 ± 0.34	98.45 ± 0.33	98.38 ± 0.34	98.38 ± 0.34
SVM	98.07 ± 0.35	98.12 ± 0.34	98.07 ± 0.35	98.07 ± 0.35
RF	98.80 ± 0.32	98.83 ± 0.31	98.80 ± 0.32	98.80 ± 0.32
MLP	95.02 ± 0.32	95.16 ± 0.30	95.00 ± 0.32	95.02 ± 0.32

Fig. 5 shows the quartile distribution of training accuracy for the five evaluated classifiers. This boxplot depicts the central tendency, dispersion, and robustness of each model under the 5-fold stratified cross-validation scheme. Among the evaluated methods, Random Forest achieves the highest median accuracy and one of the narrowest interquartile ranges, supporting its selection as the most accurate and stable classifier for the proposed zone-level classification task.

Validation results are reported in Table IV. The RF model sustained the highest validation accuracy at $97.57\% \pm 2.21\%$, with a precision of 97.94%, a Recall of 97.57% and an F1-score of 97.55%. KNN followed with an accuracy of 97.23%, while SVM and MLP reached 96.52% and 95.13%, respectively. The NB model recorded the lowest validation accuracy at 86.39%. These results indicate that RSS measurements from four APs provide sufficient discriminatory information for zone-level classification with all models except the NB.

Fig. 6 shows the quartile distribution of validation accuracy for the five evaluated classifiers. We observed that these results present the same behavior as training results, i.e., the central tendency, dispersion, and robustness of each model under the 5-fold stratified cross-validation scheme present a high median accuracy and a narrow interquartile range, supporting the selection of the models as the most accurate and stable classifier for the proposed zone-level classification task, except for the NB model

TABLE IV
VALIDATION PHASE PERFORMANCE METRICS

Model	Accuracy (%)	Precision (%)	F1-Score (%)	Recall (%)
NB	86.39 ± 7.02	86.25 ± 10.50	84.10 ± 9.05	86.39 ± 7.02
KNN	97.23 ± 3.38	97.72 ± 2.71	97.18 ± 3.46	97.23 ± 3.38
SVM	96.52 ± 3.08	97.02 ± 2.69	96.50 ± 3.10	96.53 ± 3.08
RF	97.57 ± 2.21	97.94 ± 1.96	97.55 ± 2.23	97.57 ± 2.21
MLP	95.13 ± 2.77	95.92 ± 2.27	95.03 ± 2.90	95.13 ± 2.77




Fig. 6. Distribution of validation accuracy across the cross-validation folds for the evaluated classifiers.

C. Error Analysis

Figs. 7–11 present the confusion matrices for a random single test run. The RF model produced two misclassifications out of 72 test observations (97.22% discrete accuracy): one instance of class 1 predicted as class 2, and one instance of class 2 predicted as class 5. Zones 0, 3, 4, and 5 were classified without error.

The primary misclassification pattern across all models involves classes 1 (waiting room) and 2 (diagnostic imaging). As shown in Fig. 1, these classes correspond to physically adjacent zones sharing a boundary. This pattern is therefore consistent with spatial boundary confusion rather than with random classification failure. While the SVM model recorded the highest failure rate (7 errors out of 72), the misclassifications consistently favored spatially adjacent areas. For the intended application, such errors would direct the search to a neighboring zone, which is operationally less critical than a misclassification between distant sectors. In the RF confusion matrix (Fig. 9), no errors were observed between clearly non-adjacent zones.

D. Dimensional Sensitivity Analysis

The learning curves in Figs. 12–16 show that the smallest subsets, corresponding to 12 and 24 fingerprints per zone, produce unstable validation estimates. In particular, isolated validation accuracies close to 100% are observed for some classifiers at the 20% and 40% sampling levels. These values should be interpreted cautiously because the corresponding validation folds contain a limited number of observations and therefore have high sampling variability.

As the number of fingerprints per zone increased to 48 and 60, the validation curves became more stable and approached the corresponding training curves. Training accuracy was calculated using the samples available to the model during



Fig. 7. Confusion Matrix of Naive Bayes for a random single test run (Laplace = 0 and Adjust = 1): Accuracy = 88.89%




Fig. 10. Confusion Matrix of RF for a random single test run ($ntree = 100$, $mtry = 2$): Accuracy = 97.22%




Fig. 8. Confusion Matrix of KNN for a random single test run ($k = 5$, Minkowski metric): Accuracy = 94.44%




Fig. 11. Confusion Matrix of MLP for a random single test run (decay = 0.01, hidden layer size = 3): Accuracy = 91.66%



Fig. 9. Confusion Matrix of SVM for a random single test run ($\gamma = 0.001$): Accuracy = 90.27%

fitting, whereas validation accuracy was calculated using the held-out fold in each cross-validation iteration. Therefore, a nonzero training-validation gap is expected. This gap is more variable for the smallest subsets because each fold contains fewer observations per class and the fitted decision boundaries are more sensitive to the specific fingerprints included in the training partition.

At the full dataset size, the training-validation gap remained below 1.5 percentage points for most evaluated classifiers.

The reduced gap and the stabilization of the validation curves support internal consistency under resampling of the available dataset. However, this behavior should not be interpreted as evidence of external, temporal, cross-device, or cross-building generalization.

Nevertheless, this result does not establish universal sample sufficiency or long-term generalizability. The reported evidence is limited to resampling-based validation within one hospital deployment, one sensing device, and the acquisition conditions represented in the collected dataset. Performance under different hospitals, changes in access-point configuration, receiver heterogeneity, or temporal RSS drift requires independent validation and periodic radio-map reassessment.

E. Feasibility Analysis in the Healthcare Sector: Accuracy, Cost, and Maintenance

The practical value of an indoor localization system in healthcare depends not only on predictive performance, but also on the additional infrastructure required and the operational effort needed to sustain reliable service. In hospital environments, these factors are closely coupled: systems designed for fine spatial resolution generally require denser infrastructure, dedicated hardware, calibration, and more demanding maintenance, whereas lower-complexity alternatives may provide zone-level information that is nevertheless sufficient for




Fig. 12. Comparison of training and validation classification accuracy for Naive Bayes across varying proportions of dataset size (20% to 100%)




Fig. 15. Comparison of training and validation classification accuracy for Random Forest across varying proportions of dataset size (20% to 100%)




Fig. 13. Comparison of training and validation classification accuracy for KNN across varying proportions of dataset size (20% to 100%)




Fig. 16. Comparison of training and validation classification accuracy for Multilayer Perceptron across varying proportions of dataset size (20% to 100%)




Fig. 14. Comparison of training and validation classification accuracy for Support Vector Machine across varying proportions of dataset size (20% to 100%)

specific operational workflows. Table V summarizes these deployment trade-offs for representative indoor localization technologies.

The proposed Wi-Fi RSS approach prioritizes functional-

zone identification rather than continuous coordinate estimation. Under the evaluated four-AP configuration, the RF classifier achieved a mean validation accuracy of $97.57\% \pm 2.21\%$ across six hospital zones. This level of granularity is appropriate for preventive-maintenance workflows in which personnel need to identify the department or functional area containing an asset, rather than its exact position within a room.

The confusion-matrix analysis indicates that most observed classification errors occurred between physically adjacent areas. Thus, under the evaluated conditions, the system can support zone-level management of biomedical assets when approximate departmental location is operationally sufficient. This interpretation should not be extended to applications requiring continuous tracking, sub-meter accuracy, or safety-critical real-time positioning. Compared with UWB, BLE, or RFID-based systems, the proposed approach deliberately trades spatial precision for lower incremental infrastructure requirements and reduced deployment complexity.

The principal deployment advantage is the reuse of four pre-existing Wi-Fi APs together with a low-cost ESP8266

TABLE V
DEPLOYMENT TRADE-OFFS AMONG INDOOR LOCALIZATION TECHNOLOGIES FOR HEALTHCARE FACILITIES

Technology	Localization granularity	Additional infrastructure requirement	Incremental deployment cost	Maintenance complexity	Healthcare deployment implication
UWB	Continuous coordinates; potentially sub-meter resolution	Dedicated anchors, tags, synchronization, and site configuration	High	High	Appropriate when precise real-time asset positioning is required; however, it requires dedicated hardware deployment, calibration, and maintenance.
BLE beaconing	Room-, zone-, or proximity-level localization	BLE beacon deployment and periodic battery replacement, or powered beacon infrastructure	Moderate	Moderate to high	Suitable for room-level tracking, but requires management of beacon density, battery life, replacement cycles, and periodic recalibration.
RFID	Presence detection or checkpoint-level tracking	Tags and reader infrastructure at selected doors, corridors, or storage locations	Moderate to high	Moderate	Useful for inventory control and checkpoint-based monitoring, but does not inherently provide continuous indoor position estimates.
Wi-Fi RSS fingerprinting (this work)	Functional-zone classification	No dedicated positioning infrastructure when stable Wi-Fi coverage already exists; ESP8266 sensing nodes and a local server are required	Low	Moderate	Appropriate for department- or zone-level biomedical asset identification when AP coverage is stable and periodic radio-map verification is feasible.

sensing node. Nevertheless, the approach is not cost-free. Practical operation still requires sensing-node acquisition, battery replacement or charging, local-server availability, network connectivity, database maintenance, and periodic RSS auditing and radio-map verification. These requirements are generally lower than those associated with dedicated dense-anchor systems, but they remain necessary for sustained deployment.

The maintenance burden is also linked to the stability of the radio environment. Changes in AP placement, transmit power, channel configuration, shielding conditions, furniture distribution, or hospital occupancy patterns may alter RSS fingerprints and reduce classification reliability. Therefore, the proposed approach is best suited to hospital floors where a small number of APs provide stable and spatially diverse coverage. In such settings, periodic radio-map reassessment represents a lower operational burden than maintaining dedicated positioning infrastructure, while still being essential to preserve reliable zone-level classification.

Accordingly, the proposed system should be interpreted as a cost-aware functional-zone classification solution for biomedical asset management. It is not intended to replace high-precision technologies where continuous or sub-meter positioning is required; rather, it provides a practical compromise among classification performance, incremental infrastructure requirements, and maintenance complexity for the evaluated hospital deployment.

F. Contextualized Comparison with Prior Studies

Table VI contextualizes the present study against representative indoor localization works cited in this manuscript. The comparison includes the sensing technology, localization target, dataset scale, environmental characteristics, infrastructure density, evaluation protocol, and reported metric because reported performance is strongly dependent on these experimental conditions. Closely matched studies that combine hospital functional-zone classification, classical machine-learning

models, and no more than five pre-existing Wi-Fi access points are limited. Therefore, the comparisons in Table VI are intended to contextualize infrastructure requirements and evaluation conditions rather than to establish a direct ranking of reported performance.

Maduranga et al. [13] formulated indoor zone identification as a four-class BLE RSSI classification problem using 13 iBeacon transmitters deployed on one floor of a university library. Their dataset contained 5191 RSS records and was divided into 70% training and 30% testing subsets; the best reported kNN accuracy was approximately 85%. In contrast, the present work evaluates six hospital functional zones using four existing Wi-Fi APs and 360 balanced RSS fingerprints under stratified 5-fold cross-validation.

Narasimman et al. [21] addressed floor prediction and continuous coordinate estimation on UJIIndoorLoc and ten additional public datasets. Their main UJIIndoorLoc experiment involved a multi-building and multi-floor benchmark with 19,938 training samples, 1112 test samples, and 520 available APs, although their framework retained fewer than 5% of the AP signals for one reported configuration. Their reported 94.15% floor-classification accuracy and 8.45 m mean positioning error therefore correspond to a substantially larger and more heterogeneous coordinate-level localization problem than the present zone-classification task.

Bi et al. [22] introduced the SODIndoorLoc benchmark for continuous Wi-Fi positioning across three buildings and multiple floors. The dataset covers approximately 8000 m², includes 105 pre-installed APs, 23,925 RSS samples, and dense sampling with adjacent reference points separated by approximately 0.5–1.2 m. Their best reported average positioning accuracy was approximately 2.3 m under dedicated training and testing sheets. This metric is not directly comparable with zone-classification accuracy because it evaluates coordinate error under a dense infrastructure and high-density radio map.

Alhmiedat [23] evaluated a fingerprint-based WSN localization system in two real indoor environments with obstacles and

TABLE VI
CONTEXTUAL COMPARISON WITH REPRESENTATIVE INDOOR LOCALIZATION STUDIES

Study	Signal and method	Localization target	Environment	Infrastructure	Dataset and protocol	Reported result and contextual relevance
This work	Wi-Fi RSS; NB, KNN, SVM, RF, and MLP	Six functional-zone classes	Single hospital floor: waiting rooms, diagnostic imaging, administration, outpatient care, and densitometry	Four pre-existing Wi-Fi APs selected for stable visibility across all zones	360 balanced fingerprints, 60 samples per zone; stratified 5-fold CV	97.57% $\pm$ 2.21% RF validation accuracy. Direct reference case for hospital functional-zone classification with four APs.
Maduranga et al. [13]	BLE RSSI; KNN and SVM	Four-zone classification	University library under regular operating conditions	13 BLE iBeacon transmitters	5191 RSS records; 70/30 train-test split	85.11% kNN classification accuracy. Comparable at the zone-classification level, but not in sensing technology or infrastructure density.
Wang et al. [24]	Wi-Fi CSI; Random Forest, KNN, WKNN, and decision-tree methods	Reference-point classification and coordinate localization	Microwave anechoic chamber and office environment under LOS/NLOS conditions	Four Wi-Fi APs evaluated	Multiple experiments with varying AP count; RF evaluated with four APs	93.12% classification accuracy with four APs in the office LOS experiment; 84.18% in office NLOS. Relevant as a four-router RF deployment, but CSI-based and not hospital zone classification.
Amirisoori et al. [25]	Wi-Fi RSS; KNN	Reference-point localization	Real university environment	Four and six existing Wi-Fi APs evaluated	70/30 train-test split and 10-fold CV	For four APs, the best result was 74% accuracy with $k = 7$ and 0.83 m MSE; 69% under 10-fold CV. Relevant low-infrastructure RSS baseline, but it addresses point localization rather than hospital zone classification.
Narasimman et al. [21]	Wi-Fi RSS; transfer-learning framework	Floor classification and continuous coordinate estimation	Public multi-building and multi-floor benchmarks	UJIIndoorLoc with 520 available APs; fewer than 5% retained in one configuration	19,938 training and 1112 testing samples in UJIIndoorLoc; train-test and cross-dataset transfer tests	94.15% floor-classification accuracy and 8.45 m mean positioning error. Relevant for transferability, but not a sparse four-AP comparison.
Bi et al. [22]	Wi-Fi RSS fingerprinting; multiple ML regressors and classifiers	Continuous coordinate estimation	Three buildings with corridors, offices, and meeting rooms	105 pre-installed Wi-Fi APs	23,925 samples; dedicated training and testing sets	Best reported mean positioning error of approximately 2.3 m. Dense-infrastructure coordinate-localization benchmark; not directly comparable with zone-classification accuracy.
Alhmiedat [23]	WSN RSS fingerprinting; KNN and other ML models	Continuous coordinate estimation	Two real indoor environments with walls and obstacles	Four reference nodes	Three datasets with different reference-point densities; experimental validation and 20-point study-room test	Best reported mean localization error of approximately 1.4 m. Strong infrastructure-density analogue, but based on WSN nodes and coordinate estimation.

walls. The system used four reference nodes and reported an average continuous-positioning error near 1.4 m for the best KNN configuration, followed by a separate validation using 20 testing points in a study room. This work is closer to the present study in infrastructure density, but it still addresses continuous coordinate estimation rather than six-class hospital zone classification.

Accordingly, the 97.57% validation accuracy reported here should not be interpreted as superior to the accuracy or positioning errors reported in the cited studies. Its contribution is instead the demonstration that, for the evaluated hospital floor, a four-AP Wi-Fi RSS configuration can support reliable

functional-zone identification without deploying dedicated positioning infrastructure. Cross-study comparison is therefore qualitative and context-dependent rather than a ranking of algorithmic superiority.

The low-infrastructure studies included in Table VI confirm that indoor localization can be implemented with four access points, routers, or reference nodes. However, these studies predominantly address CSI-based fingerprinting, reference-point or coordinate estimation, and non-healthcare environments; therefore, they provide infrastructure-level context rather than direct performance baselines for the present hospital functional-zone classification task.

G. Practical Implications and Limitations

The research demonstrates that a low-cost ESP8266-based sensing node can reliably identify device locations at the functional zone level. From an operational perspective, the observed misclassifications primarily occurred between physically adjacent zones sharing a boundary wall (e.g., classes 1 and 2 in Fig. 1). For biomedical asset tracking, such boundary errors carry a substantially lower operational cost than failures between distant sectors, as they still direct the search to the correct vicinity for maintenance personnel.

Several limitations apply to the current scope. This study models a single healthcare facility with six zones; generalization to different building materials or layouts requires further validation. The system provides discrete zone identification rather than continuous coordinates. Additionally, the 360-sample dataset is modest compared to synthetic benchmarks. Environmental dynamics, such as furniture repositioning or varying human occupancy, may affect RSS stability and require periodic radio map updates. The battery-life result is an analytical estimate derived from nominal current specifications and the assumed scanning duty cycle. Direct measurements of current consumption, instantaneous power, and energy use under varying Wi-Fi conditions, including weak signals and transmission retries, were not conducted. Therefore, the reported autonomy should be interpreted as an engineering estimate rather than as experimentally measured battery performance.

The present study also did not include direct measurements of inference latency, computational complexity, processor utilization, memory footprint, or end-to-end response time. Although the evaluated Random Forest configuration used 100 trees and only four RSS inputs, its real-time feasibility on resource-constrained hardware requires direct processor-level profiling under deployment conditions.

IV. CONCLUSION

This work validated a Wi-Fi RSS fingerprinting system for biomedical asset tracking in healthcare facilities. By utilizing RSS vectors from four reference access points, five supervised classifiers were evaluated for zone-level classification across six functional zones.

Performance analysis using a 5-fold stratified cross-validation protocol on a 360-sample dataset showed that the Random Forest model achieved the highest classification accuracy ($97.57\% \pm 2.21\%$). KNN, SVM, and MLP exceeded 90% validation accuracy, whereas NB achieved $86.39\% \pm 7.02\%$. The error analysis identified spatial adjacency as the primary factor in misclassification, particularly between waiting and diagnostic imaging sectors.

These results confirm that existing Wi-Fi infrastructure and low-cost sensing hardware provide a viable solution for zone-level biomedical asset tracking in healthcare facilities. Future work will address three directions: energy-efficient firmware to extend the battery life of the sensing modules during continuous operation, methods to compensate for temporal RSS drift caused by environmental changes such as furniture repositioning and variable occupancy, and validation across

multiple floors or facilities to assess the generalization of the fingerprinting approach. Integration with computerized maintenance management systems (CMMS) will also be explored to close the loop between zone identification and maintenance scheduling.

ACKNOWLEDGMENTS

This work was supported by Universidad del Valle under Internal Research Call No. 171-2025, Projects CI 21300 and CI 21308.


REFERENCES

- [1] M. A. C. Alvarado and J. C. S. Delgado, "Sistemas y tecnologías que facilitan el posicionamiento indoor," *Pensamiento Actual*, vol. 17, no. 29, pp. 132–144, 2017. DOI: 10.15517/pa.v17i29.31585
- [2] N. Singh, S. Choe, and R. Punmiya, "Machine learning based indoor localization using Wi-Fi RSSI fingerprints: An overview," *IEEE Access*, vol. 9, pp. 127 150–127 174, 2021. DOI: 10.1109/ACCESS.2021.3111083
- [3] H. Attar et al., "Machine learning in indoor localization prediction using received signal strength indicator and wi-fi network," *International Journal of Intelligent Networks*, vol. 6, pp. 233–243, 2025. DOI: 10.1016/j.ijin.2025.10.001
- [4] B. Ezhumalai, M. Song, and K. Park, "An efficient indoor positioning method based on Wi-Fi RSS fingerprint and classification algorithm," *Sensors*, vol. 21, no. 10, p. 3418, 2021. DOI: 10.3390/s21103418
- [5] L. Polak, S. Rozum, M. Slanina, T. Bravenec, T. Fryza, and A. Pikrakis, "Received signal strength fingerprinting-based indoor location estimation employing machine learning," *Sensors*, vol. 21, no. 13, 2021. DOI: 10.3390/s21134605
- [6] M. F. S. van der Ham, S. Zlatanova, E. Verbree, and R. Voûte, "Real time localization of assets in hospitals using Quappa indoor positioning technology," *Remote Sensing and Spatial Information Sciences*, vol. IV-4/W1, pp. 105–106, 2016. DOI: 10.5194/isprs-annals-IV-4-W1-105-2016
- [7] J. Wichmann, "Indoor positioning systems in hospitals: A scoping review," *Digital Health*, vol. 8, 2022. DOI: 10.1177/20552076221081696
- [8] D. Zhang, L. T. Yang, M. Chen, S. Zhao, M. Guo, and Y. Zhang, "Real-time locating systems using active RFID for internet of things," *IEEE Systems Journal*, vol. 10, no. 3, pp. 1226–1235, 2016. DOI: 10.1109/JSYST.2014.2346625
- [9] A. Kapoor et al., "Estimating physical work-load on ed clinicians and staff using real-time location systems," *IEEE Journal of Radio Frequency Identification*, vol. 5, no. 3, pp. 331–334, 2021. DOI: 10.1109/JRFID.2021.3091903
- [10] J. Maneesilp, C. Wang, H. Wu, and N.-F. Tzeng, "RFID support for accurate 3D localization," *IEEE Transactions on Computers*, vol. 62, no. 7, pp. 1447–1459, 2013. DOI: 10.1109/TC.2012.83
- [11] Y. Álvarez López, J. Franssen, G. Álvarez Narcandi, J. Pagnozzi, I. G.-P. Arrillaga, and F. L.-H. Andrés, "RFID technology for management and tracking: E-health applications," *Sensors*, vol. 18, no. 8, p. 2663, 2018. DOI: 10.3390/s18082663
- [12] A.-M.-S. Chong, B.-C. Yeo, and W.-S. Lim, "Integration of uwb rss to Wi-Fi RSS fingerprinting-based indoor positioning system," *Cogent Engineering*, vol. 9, no. 1, 2022. DOI: 10.1080/23311916.2022.2087364
- [13] M. W. P. Maduranga, V. Tilwari, and R. Abeysekera, "Improved rss indoor localization in iot systems with machine learning algorithms," *Signals*, vol. 4, no. 4, pp. 651–668, 2023. DOI: 10.3390/signals4040036
- [14] G. Li, E. Geng, Z. Ye, Y. Xu, J. Lin, and Y. Pang, "Indoor positioning algorithm based on the improved RSSI distance model," *Sensors*, vol. 18, no. 9, 2018. DOI: 10.3390/s18092820
- [15] C. Du, B. Peng, Z. Zhang, W. Xue, and M. Guan, "KF-KNN: Low-cost and high-accurate fm-based indoor localization model via fingerprint technology," *IEEE Access*, vol. 8, pp. 197 523–197 531, 2020. DOI: 10.1109/ACCESS.2020.3031089
- [16] S.-H. Lee, C.-H. Cheng, T.-H. Huang, and Y.-F. Huang, "Machine learning-based indoor positioning systems using multi-channel information," *Journal of Engineering and Technological Sciences*, vol. 55, no. 4, pp. 373–383, 2023. DOI: 10.5614/j.eng.technol.sci.2023.55.4.2

- [17] Y. Cheng et al., “Indoor wifi fingerprint localization based on dual population pso of stacked autoencoder and multi label classification,” *Scientific Reports*, vol. 15, no. 32941, 2025. DOI: 10.1038/s41598-025-14990-1
- [18] Z. Tang, S. Li, K. S. Kim, and J. S. Smith, “Multi-dimensional wi-fi received signal strength indicator data augmentation based on multi-output gaussian process for large-scale indoor localization,” *Sensors*, vol. 24, no. 1026, 2024. DOI: 10.3390/s24031026
- [19] J. Yoo, “Wi-fi fingerprint indoor localization by semi-supervised generative adversarial network,” *Sensors*, vol. 24, no. 5698, 2024. DOI: 10.3390/s24175698
- [20] M. Nasir et al., “Hytra: Hyperclass transformer for wifi fingerprinting-based indoor localization,” *Transactions on Energy Systems and Engineering Applications*, vol. 5, no. 1, 1–24, 2024. DOI: 10.32397/tesea.vol5.n1.542
- [21] S. C. Narasimman and A. Alphones, “Dumbloc: Dumb indoor localization framework using wi-fi fingerprinting,” *IEEE Sensors Journal*, vol. 24, no. 9, pp. 14 623–14 630, 2024. DOI: 10.1109/JSEN.2024.3374415
- [22] J. Bi, Y. Wang, B. Yu, H. Cao, C. Liu, and Z. Yuan, “Supplementary open dataset for WiFi indoor localization based on received signal strength,” *Satellite Navigation*, vol. 3, no. 1, p. 25, 2022.
- [23] T. Alhmiedat, “Fingerprint-based localization approach for wsn using machine learning models,” *Applied Sciences*, vol. 13, no. 5, 2023. DOI: 10.3390/app13053037
- [24] Y. Wang, C. Xiu, X. Zhang, and D. Yang, “WiFi Indoor Localization with CSI Fingerprinting-Based Random Forest,” *Sensors*, vol. 18, no. 9, p. 2869, 2018. DOI: 10.3390/s18092869
- [25] S. Amirsoori, S. Mohd Daud, N. A. Ahmad, N. S. N. Abd Aziz, N. I. Mohd Sa’at, and N. Q. Mohd Noor, “WI-FI Based Indoor Positioning Using Fingerprinting Methods (KNN Algorithm) in Real Environment,” *International Journal of Future Generation Communication and Networking*, vol. 10, no. 9, pp. 23–36, 2017. DOI: 10.14257/ijfgcn.2017.10.9.03



Leonardo A. Bermeo Varón received his B.Tech and M.Tech degrees in Electronics Engineering from Universidad del Valle, Cali, Colombia, in 2004 and 2011, respectively, and the Ph.D degree in Mechanical Engineering from Federal University of Rio de Janeiro, Brazil, in 2015. Currently, he is an Associate Professor of the Schools of Electrical and Electronics Engineering, Universidad del Valle, Cali, Colombia. His research interests include industrial informatics, artificial intelligence, Bayesian framework and IoT systems.



Juan Pablo Mañunga-Mayor received his B.Tech degree in Bioengineering from Universidad Santiago de Cali, Colombia, in 2023. Currently, he is a maintenance engineer at Chistus Sinergia, Cali, Colombia. His main interests include management and maintenance of biomedical equipment.



Gabriel Eduardo Mejía-Ruíz is a Full-Time Associate Professor at Schools of Electrical and Electronics Engineering, Universidad del Valle, Colombia. He received his B.Sc. in Control Engineering (Universidad Nacional de Colombia, 2007), M.Sc. in Electrical Engineering (Universidad de Antioquia, 2015), and Ph.D. in Electrical Engineering (UNAM, 2023). His research focuses on the design, real-time simulation, and lab-scale prototyping of grid-connected power electronic converters for renewable energy systems.