MedTracker: A BLE-Based Indoor Localization System for Tracking Portable Medical Devices in Hospital Environments

Higor D. Oliveira[1], **Elisa M. Sarmento**[1], **Eduardo S. Saleme**[1], **Luis A. Souza Jr**[1], **Gustavo C. Vivas**[2], **Daniel R. Trindade**[3], **Rodolfo S. Villaça**[1], **André G. C. Pacheco**[1]

[1]Informatics Departament (DI/CT)
Federal University of Espirito Santo (Ufes)

{higor.d.oliveira, elisa.sarmento, eduardo.saleme}@edu.ufes.br, {la.souza, rodolfo.villaca, apacheco}@inf.ufes.br

[2]Brazilian Company of Hospital Services (Ebserh)
Cassiano Antônio de Morais University Hospital (Hucam/Ufes)

gustavo.vivas@ebserh.gov.br

[3]Instituto Federal do Espírito Santo (Ifes)

danielrt@ifes.edu.br

Abstract. *The misplacement and inefficient tracking of Portable Medical Devices (PMDs) represent a persistent operational challenge in hospital environments, leading to delays in care delivery and increased burden on medical staff. This paper presents MedTracker, an indoor localization tool based on Bluetooth Low Energy (BLE), designed for real-time asset tracking in healthcare facilities. The system employs BLE tags attached to monitored devices and strategically deployed ESP32-based gateways that forward signal data to a central server. By combining BLE signals and filtering techniques with machine learning classification and clustering algorithms, MedTracker is capable of inferring room- and floor-level device location, as well as detecting location changes over time. The proposed tool was motivated by practical demands observed at a public university hospital and validated through experimental evaluation, demonstrating its effectiveness as a scalable and deployable solution for hospital asset management.*

1. Introduction

The growing complexity of hospital environments, combined with the increasing number of medical devices in operation, has intensified the challenges associated with asset management and equipment localization. Traditional inventory systems, based on manual updates, have been proven to be inadequate for real-time control in high-demand scenarios, as evidenced during the COVID-19 pandemic, when failures in ventilator management and monitoring equipment negatively affected healthcare systems [Barach et al. 2020], [Sharma et al. 2021].

While real-time tracking systems offer a clear solution, the high costs of deployment often make it prohibitive to monitor every Portable Medical Device (PMD) in a hospital setting. Traditional solutions, such as those based on standard Wi-Fi, Ultra-Wide

Band (UWB), or RFID, typically require the installation of complex and costly fixed gateways or base stations throughout the hospital's infrastructure to receive localization data [Zafari et al. 2019, Aziz and Koo 2025]. Consequently, these traditional Real Time Localization Systems (RTLS) architectures struggle to adapt effectively to large-scale hospital scenarios without incurring massive financial and operational burdens.

To address these challenges, this work proposes a low-cost indoor localization system based on Bluetooth Low Energy (BLE). The architecture relies on strategically deploying fixed gateways, built with ESP32 microcontrollers, across hospital rooms and corridors, while standard BLE tags are attached to the PMDs requiring real-time monitoring. These gateways continuously collect signal information broadcast by the tags and forward it to a central server, which maintains a mapping between tags and their corresponding monitored devices. By leveraging this signal data with machine learning-based classification and clustering algorithms, the server is capable of inferring the location of each device within the hospital premises. Importantly, the system does not aim to determine precise geographic coordinates, as in GPS-based solutions; rather, it focuses on two clinically relevant localization objectives: (i) identifying the room or floor where a given device is currently located, and (ii) detecting whether a monitored device has undergone a change in location.

In this context, MedTracker addresses a practical demand observed at the Hospital Universitário Cassiano Antônio Moraes (HUCAM/Ufes), a public university healthcare facility that encounters persistent difficulties in tracking portable medical devices across its premises. The absence of a robust indoor localization framework leads to significant delays, increases the operational burden on medical staff, and reduces overall visibility regarding asset distribution.

MedTracker aims to bridge this gap, offering a robust framework for real-time tracking that reduces the burden on medical staff and optimizes resource distribution. The main contributions are:

- **Low-cost Infrastructure:** Utilization of standard hardware to reduce the financial barrier for public healthcare asset management;
- **Real-time Monitoring:** Continuous and reliable tracking of Portable Medical Devices (PMDs) to significantly reduce equipment search times, alleviate the operational burden on healthcare staff, and optimize clinical workflows;
- **Robust Localization Mechanism:** Implementation of RSSI filtering and processing techniques designed to mitigate severe signal interference caused by typical hospital obstacles.

The rest of this paper is structured as follows: Section 2 presents the system architecture and implementation of the localization algorithm. Section 3 describes the experimental results from practical use cases. Section 4 provides links for the MedTracker documentation and download, while Section 5 concludes this paper with future work and extensions.

2. MedTracker Design and Implementation

The MedTracker is a low-cost BLE RSSI-based indoor localization system designed to use centralized data processing and machine learning-based classification and clustering algorithms running inside a Docker environment, using strategically deployed ESP32-based gateways distributed across the hospital as a sensing network. The system architecture is detailed in the following subsections.

2.1. System Architecture

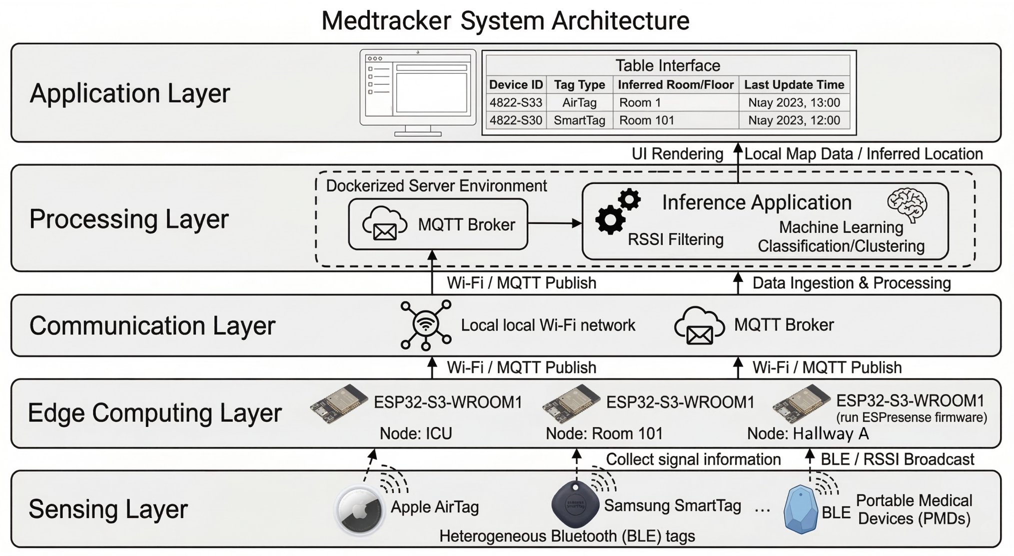

Figure 1. MedTracker system architecture illustrating the data flow from BLE tags through ESP32 gateways to the centralized inference server.

Inspired by modern IoT topologies, the system operates through a continuous pipeline of signal acquisition, edge-level filtering, and centralized processing and machine learning inference. To facilitate the deployment of the system in public healthcare facilities, the entire ecosystem relies on certified, commercially available hardware and open-source communication protocols. Figure 1 illustrates the data flow and main components of the architecture, which is divided into the hardware-centric system architecture and the containerized software infrastructure.

2.2. Hardware Infrastructure

The hardware infrastructure of the proposed system consists of a distributed set of fixed Bluetooth Low Energy (BLE) scanning nodes and mobile BLE tags attached to the assets being monitored. Figure 2 shows the fixed nodes implemented using ESP32-S3-WROOM-1 microcontroller boards that have Wi-Fi and BLE capabilities at low cost. Each node runs the open-source ESPresense firmware, configuring the device as a dedicated BLE scanner. [ESPresense Project 2024]. This setup provides a scalable and cost-effective localization infrastructure, reducing the financial barriers associated with specialized indoor positioning hardware.

This approach is particularly suited for large-scale healthcare environments such as the HUCAM, where the distributed topology of the scanning nodes helps mitigate signal degradation in these kinds of premises. In the proposed framework, mobile assets are equipped with heterogeneous BLE tags, such as Apple AirTag, Samsung SmartTag, and standard Meatek BLE tags, which periodically broadcast connectionless advertisement

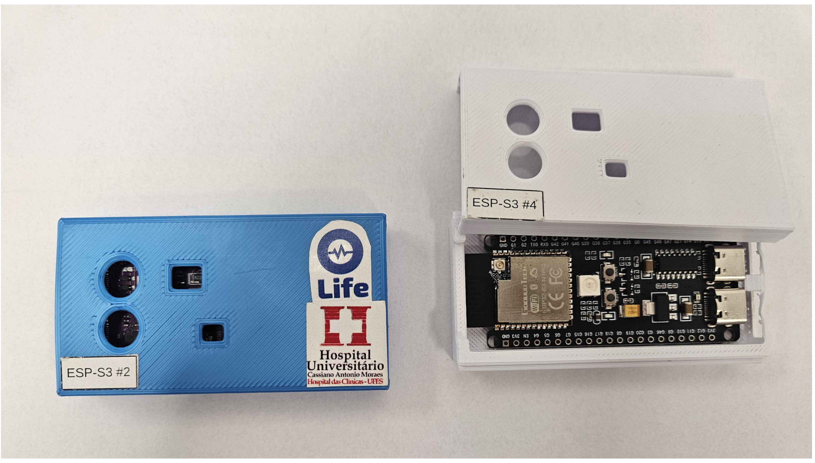

Figure 2. 3D-printed enclosure for the MedTracker node, showing the assembled case (left) and the internal ESP32-S3 module (right).

packets. These packets contain unique identifiers that allow the system to distinguish and track multiple tags simultaneously.

The ESP32 nodes passively scan the surrounding radio environment and capture these advertisement packets, extracting the Received Signal Strength Indicator (RSSI) associated with each transmission. In the context of indoor localization, the RSSI metric is commonly used as a parameter to infer the relative proximity and spatial distance between the mobile transmitter and the fixed receiving node.

However, RSSI measurements in indoor environments are often affected by multipath propagation and signal attenuation caused by walls, equipment, and other obstacles, especially in high-traffic environments such as hospitals. These effects introduce significant fluctuations in the raw RSSI values. To address these inaccuracies, the ESPresence firmware processes the incoming RSSI data using a Kalman filter combined with a median filtering approach [ESPresense Project 2024, Zhou et al. 2017]. This approach suppresses abrupt signal variations and discards outliers, resulting in the stable observations necessary for accurate localization in dense, interference-heavy environments.

Once filtered, the BLE data is transmitted via the Message Queuing Telemetry Transport (MQTT) protocol. In this architecture, each node operates as an MQTT client, publishing real-time payloads containing the tag's unique identifier, the node's ID, and the smoothed RSSI value to a central MQTT broker. These data streams serve as the primary input for the MedTracker software platform, which processes the telemetry to perform the final indoor localization of the tracked assets.

2.3. Software Infrastructure

The MedTracker software architecture, presented in Figure 3, provides an overview of the software modules, their connections, and the corresponding data flow between them.

Input: The software pipeline receives two primary Comma Separated Values (CSV) datasets as input (in violet in Figure 3) and is used for model training. The first dataset consists of raw Bluetooth Low Energy (BLE) telemetry logs acquired via MQTT, which include timestamps, receiver node identifiers, and payload metrics such as Re-

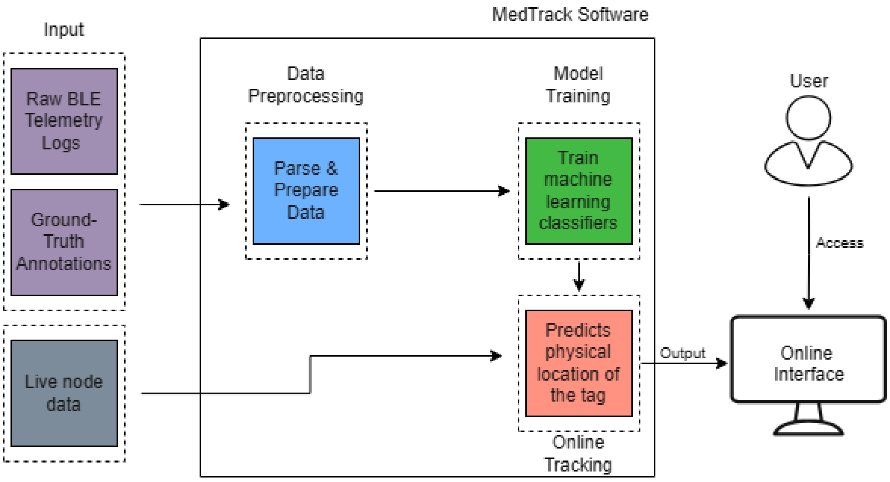

Figure 3. Software architecture of MedTracker, encompassing data preprocessing, ML model training, and online tracking modules.

ceived Signal Strength Indicator (RSSI). The second dataset provides temporal ground-truth annotations, mapping specific time intervals to the actual physical locations and floors of the tracked tags, enabling spatial validation and system calibration.

Data Preprocessing: The module that parses, analyzes, and prepares the input data for model training. This phase of the process uses a time window within which the raw RSSI data are grouped, and a mean filter is applied to mitigate signal noise and fluctuations. The main output of this stage is a matrix containing the location of the tag, the time interval, and the mean RSSI values read from each node that exists in the system.

Model Training: Using the preprocessed data, machine learning models are trained to predict the physical location of the tracking tags. Specifically, three distinct classifiers are implemented using the scikit-learn library [Pedregosa et al. 2011]: Random Forest (RF), Logistic Regression (LR), and K-Nearest Neighbors (KNN). All models are evaluated using cross-validation, and their respective accuracies are displayed in the terminal. Ultimately, a single model selected by the user is saved and deployed for the online tracking phase.

Online Tracking from Live Node Data: The system collects live node data from the receiver nodes to predict the physical location of each tag in real time. Upon receiving the data payload, the system extracts the tag identifier, the receiving node's ID, and the RSSI value. The RSSI readings for each tag are then grouped within a defined time window and smoothed using a moving average filter. Finally, this processed real-time data serves as the direct input for the previously trained machine learning model, which outputs the estimated location of the tag. The current location, the last known location, and the active nodes, along with other system information, are displayed on an online interface, as shown in Figure 4.

3. MedTracker Use Cases

The performance of MedTracker was assessed under two operational scenarios: floor-level and room-level localization. The experiments were conducted in a three-floor build-

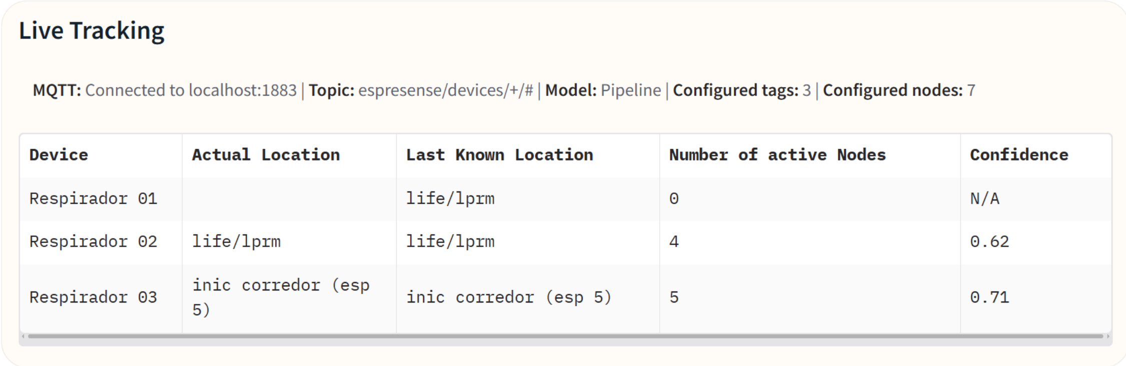
Live Tracking

MQTT: Connected to localhost:1883 | **Topic:** espresense/devices/+/# | **Model:** Pipeline | **Configured tags:** 3 | **Configured nodes:** 7

Device	Actual Location	Last Known Location	Number of active Nodes	Confidence
Respirador 01		life/lprm	0	N/A
Respirador 02	life/lprm	life/lprm	4	0.62
Respirador 03	inic corredor (esp 5)	inic corredor (esp 5)	5	0.71

Figure 4. Screenshot of the MedTracker web interface showing live tracking status for monitored BLE tags.

ing at the Federal University of Espírito Santo. While the floor-level tests encompassed all three floors of the facility, the room-level evaluations were conducted in selected rooms on the third floor.

3.1. Floor-Level Localization Test

To evaluate the system's ability to identify the correct floor, we deployed six scanning nodes across the building's three levels. Figure 5 illustrates the architectural layout and the spatial distribution of these nodes, with two units placed on opposite sides of each floor to ensure coverage.

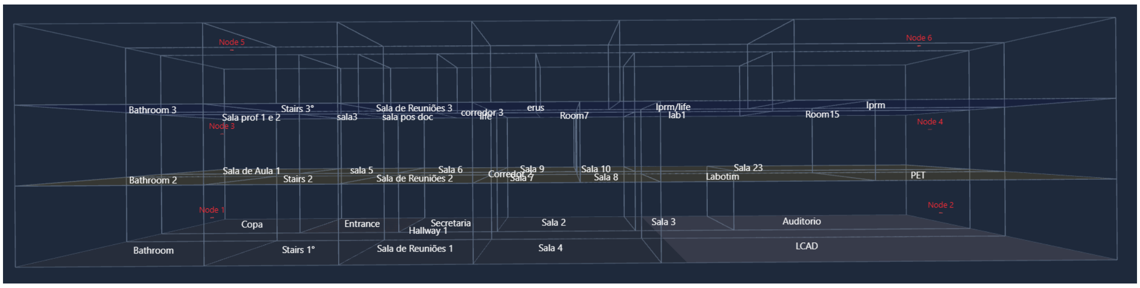

Figure 5. Experimental setup for floor-level localization. Red markers indicate the two nodes placed on opposite sides of each floor.

3.2. Room-Level Localization

The second phase of testing focused on distinguishing between neighbors' rooms on the third floor. For this scenario, five scanning nodes were strategically distributed to optimize spatial coverage. Figure 6 presents the floor plan and the specific placement of the nodes used to evaluate the system's room-level precision.

3.3. Experimentation Results

The evaluation results, obtained through cross-validation, demonstrate an overall accuracy of 84.5% for the floor-level localization test and 81.0% for the room-level localization test. The confusion matrices detailing the performance of each test are presented in Figures 7a and 7b.

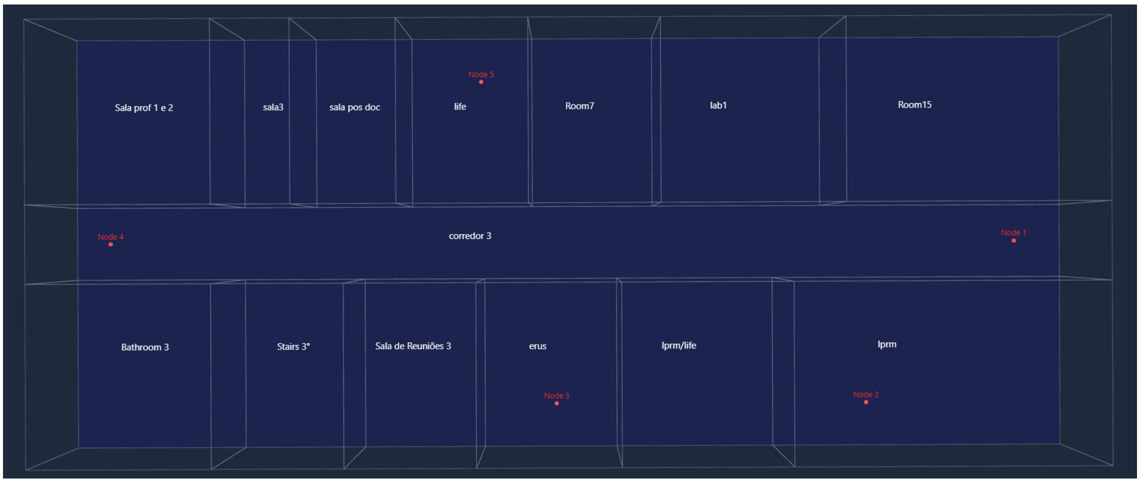

Figure 6. Experimental setup for room-level localization on the third floor. Nodes (in red) were positioned to maximize spatial distribution.

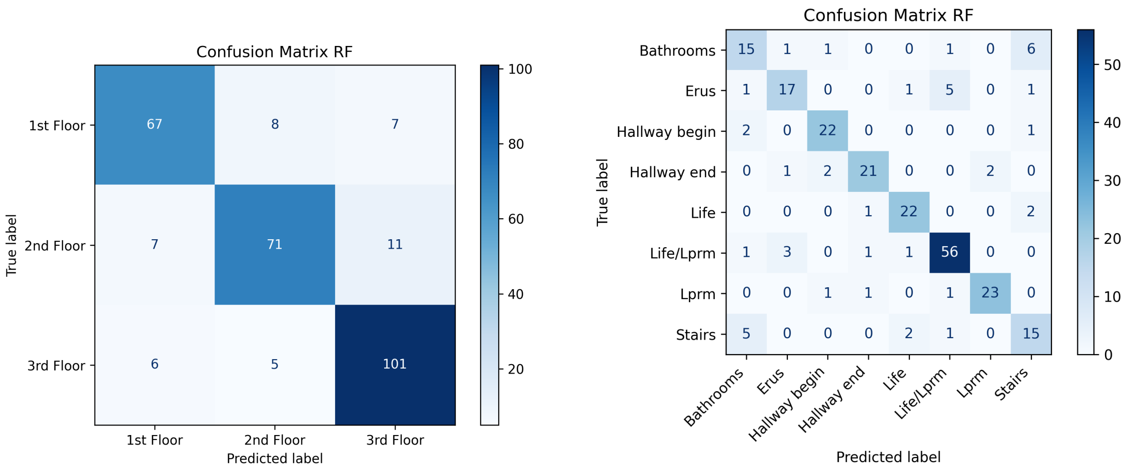

(a) Floor-level localization. Overall accuracy: 84.5%.

(b) Room-level localization. Overall accuracy: 81.0%.

Figure 7. Confusion matrices for the MedTracker localization tests using Random Forest classifier.

A closer analysis of both confusion matrices reveals that the decrease in accuracy is primarily driven by misclassifications between adjacent spaces, such as confusing the ”Erus” room with the ”Life/lprm” room or the second floor with the third floor. Despite this boundary overlap, the primary utility of the tool remains highly effective. As previously discussed, the core motivation of MedTracker is to establish a highly scalable, low-cost RTLS to facilitate real-time equipment tracking. Even though the ESP32 hardware is known to exhibit high RSSI variance as a function of distance [Al-Maktary et al. 2025], the achieved results are highly promising when considering the trade-off between deployment cost and localization accuracy. Furthermore, these results underscore the capability of the underlying machine learning model to provide a robust framework that significantly reduces the operational burden on medical staff.

4. MedTracker Documentation and Download

MedTracker project development is ongoing, with bug fixes, new feature releases, and expansions. The code is open-source and available on GitHub[1] page, alongside documentation. The hardware's firmware is available on the official ESPresense's GitHub[2]. A demonstration video focusing on reproducing the use case described in 3 can be found on YouTube[3] and on the project's GitHub page.

5. Conclusion

This paper presented MedTracker, a low-cost indoor localization tool designed to address the critical challenge of tracking Portable Medical Devices (PMDs) in hospital environments. Using a decoupled architecture composed of heterogeneous commercial BLE tags, ESP32-based edge gateways running ESPresense, and a containerized central server, MedTracker provides a highly scalable alternative to traditional, cost-prohibitive RTLS solutions.

The experimental validation demonstrated the practical viability of the system, achieving an overall accuracy of 84.5% for floor-level detection and 81.0% for room-level inference. While the inherent RSSI fluctuations of BLE technology resulted in some misclassifications between adjacent spaces, the tool successfully fulfills its primary objective: significantly reducing the equipment search space and the operational burden on healthcare professionals. The favorable trade-off between deployment cost, hardware accessibility, and localization accuracy makes MedTracker a highly suitable framework for public healthcare facilities, such as HUCAM/Ufes.

Future work will focus on mitigating the boundary overlap issues observed in the confusion matrices. We plan to implement temporal smoothing techniques and more advanced filtering algorithms to reduce the impact of transient signal reflections and improve the stability of the machine learning inferences. Additionally, we intend to deploy and evaluate the system in a fully operational clinical environment at HUCAM to assess the impact of dynamic obstacles, such as moving beds and high human density. Finally, the tool's user interface will be expanded to include dynamic 2D floor plan visualizations, further enhancing real-time asset visibility for the medical staff.

Acknowledgment

This study was financed in part by the Fundação de Amparo à Pesquisa e Inovação do Espírito Santo (FAPES) – Grant number 201/2025, under the Call for Proposals No. 13/2023 – PROCAP 2024.

Fapes (2026-B74MN, 2023-RWXSZ); CNPq e CAPES.

References

Al-Maktary, O., Susanto, M., and Mardiana, M. (2025). Evaluation of rssi-based distance estimation with esp32 ble modules for indoor asset tracking. *ELKHA: Jurnal Teknik Elektro*, 17(2):172–179.

[1] `https://github.com/life-ufes/MedTracker`
[2] `https://github.com/ESPresense/ESPresense/tree/v3.3.4`
[3] `https://www.youtube.com/watch?v=loZT-bcbE_s`

Aziz, T. and Koo, I. (2025). A comprehensive review of indoor localization techniques and applications in various sectors. *Applied Sciences*, 15(3):1544.

Barach, P., Fisher, S. D., Adams, M. J., Burstein, G. R., Brophy, P. D., Kuo, D. Z., and Lipshultz, S. E. (2020). Disruption of healthcare: Will the covid pandemic worsen non-covid outcomes and disease outbreaks? *Progress in Pediatric Cardiology*, 59:101254.

ESPresense Project (2024). Espresense: Esp32-based ble presence detection.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830.

Sharma, A., Borah, S. B., and Moses, A. C. (2021). Responses to covid-19: The role of governance, healthcare infrastructure, and learning from past pandemics. *Journal of Business Research*, 122:597–607.

Zafari, F., Gkelias, A., and Leung, K. (2019). A survey of indoor localization systems and technologies. arxiv. *arXiv:1709.01015v3*.

Zhou, C., Yuan, J., Liu, H., and Qiu, J. (2017). Bluetooth indoor positioning based on RSSI and kalman filter. *Wireless Personal Communications*, 96(3):4115–4130.