Real-World NLOS Validation of Novel Low-Power BLE Devices for Indoor Localization

José L. López Ruiz[1(✉)], David Díaz Jiménez[2], Juan Francisco Gaitán-Guerrero[1], Joaquín Torres-Sospedra[3], and Macarena Espinilla Estévez[1]

[1] Department of Computer Science, Universidad de Jaén, Campus Las Lagunillas s/n, 23071 Jaén, Spain
{llopez,jgaitan,mestevez}@ujaen.es

[2] Department of Computer Systems, Universidad Politécnica de Madrid (UPM), Madrid, Spain
david.djimenez@upm.es

[3] Department of Computer Science, Universitat de València, Avda. Universitat s/n, 46100 Burjassot, Valencia, Spain
joaquin.torres@uv.es

Abstract. Indoor positioning remains an open challenge, as satellite signals suffer severe degradation from attenuation, diffraction, and multipath propagation in built environments. Although the literature has explored a wide range of wireless technologies, the limited configurability of commercial beacons restricts the performance of existing systems, particularly in Ambient Assisted Living settings. To address this gap, we propose an indoor positioning system based on Bluetooth Low Energy signal fingerprinting with custom-built beacons featuring configurable transmission power. The architecture follows an edge-fog-cloud model that integrates these beacons with an activity wristband, a mobile application, and a cloud back-end. Experiments were conducted in a real residential dwelling of 100 m^2 with 28 reference points, evaluating four machine learning classifiers: k-Nearest Neighbours, Support Vector Machine, Random Forest, and Multilayer Perceptron, under stratified cross-validation. All classifiers surpass 98 % accuracy, with Random Forest achieving the best and most consistent result at 99.4 %, thereby validating both the precision and the reliability of the proposed custom devices.

Keywords: Indoor Localization · Bluetooth Low Energy · Fingerprinting · Machine Learning · RSSI · Ambient Assisted Living

1 Introduction

Indoor Positioning Systems (IPS) determine the position of a person or object inside an enclosed space [8,18]. Unlike outdoor environments, where Global

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
M. J. del Jesus et al. (Eds.): CAEPIA 2026, LNAI 16999, pp. 396–407, 2027.
https://doi.org/10.1007/978-3-032-43170-7_32

Navigation Satellite Systems (GNSS) provide adequate coverage, indoor settings introduce Non-Line-of-Sight (NLOS) conditions caused by walls, metallic objects, and furniture that attenuate or block satellite signals [16,19]. Among the alternatives explored in the literature, including acoustic, inertial, and radio-frequency technologies, Bluetooth Low Energy (BLE) has gained the widest adoption thanks to its low power consumption, low deployment cost, and native availability on consumer mobile devices [11,19]. Overall system accuracy, however, depends not only on the positioning algorithms but equally on the configurability of the transmitting hardware.

In this context, fingerprinting has become the dominant paradigm for complex indoor environments [11,12]. Machine Learning (ML) and Deep Learning (DL) have substantially improved accuracy over classical geometric methods by capturing the non-linear and noisy nature of Received Signal Strength Indicator (RSSI) signals under multipath propagation [1,11,14].

The interest in BLE and ML-based IPS stems from its applicability across diverse domains. In Ambient Assisted Living (AAL) and healthcare, these systems enable continuous, non-invasive monitoring of patients and elderly individuals, supporting fall detection and context-aware services [7,16]. Industry 4.0 environments have also been targeted, where IPS facilitates real-time asset and personnel tracking [2,17,19], and public spaces such as airports and museums have seen DL models successfully estimate both position and movement direction with high accuracy [4,5,9].

However, commercially available beacons offer limited configurability, restricting the diversity of the resulting RSSI space. To address this shortcoming, custom BLE beacons with configurable transmission power were designed and prototyped specifically for AAL settings [3]. The central motivation of this paper is the experimental validation of those devices under real-world NLOS propagation conditions, with the obstructions and signal variability inherent in an occupied residential dwelling.

The main contributions of this work are: (i) a proposed architecture for non-invasive localization in NLOS environments using custom BLE beacons and an activity wristband; (ii) a publicly available dataset collected in a real-world environment; and (iii) experimental validation of an IPS built on purpose-designed hardware.

The remainder of this paper is organised as follows. Section 2 reviews BLE- and fingerprinting-based IPS; Sect. 3 describes the edge-fog-cloud architecture, the RSSI preprocessing pipeline, and the classifiers; Sect. 4 details the use case, protocol, and results; and Sect. 5 concludes.

2 Related Work

2.1 BLE-Based Fingerprinting in Indoor Positioning Systems

Recent surveys cover several radio-frequency alternatives to GNSS for enclosed spaces: Wi-Fi, Ultra-Wideband (UWB), Radio-Frequency Identification (RFID), ZigBee and BLE [8,18,19]. Among them, BLE has progressively displaced Wi-Fi

for resource-constrained scenarios thanks to its advertising-mode operation [11]. Two methodological paradigms coexist: geometric methods, which invert propagation models from Time of Arrival (ToA), Time Difference of Arrival (TDoA), Angle of Arrival (AoA) or Round-Trip Time (RTT) measurements [6], and fingerprinting, which builds an empirical radio map at an offline phase and matches RSSI observations to it at runtime [12,15]. By treating multipath as part of the location signature rather than as noise, fingerprinting reports substantially better behaviour in cluttered indoor spaces [11,16].

2.2 Machine Learning and Deep Learning Classifiers for RSSI

k-Nearest Neighbours (kNN) is the historical baseline for fingerprint matching, and remains competitive when the radio map is sufficiently dense [11,17]. Beyond kNN, three classical classifiers are routinely benchmarked: Support Vector Machine (SVM), which draws non-linear boundaries in the high-dimensional RSSI space; Random Forest (RF), which combines noise resilience with an inherent feature-importance ranking [7,14]; and Multilayer Perceptron (MLP), which closes the gap with deeper architectures while remaining tractable on commodity hardware. DL architectures, Deep Neural Network (DNN) regressors [4], Convolutional Neural Networks (CNNs) on stacked RSSI matrices [5,9] and attention-based encoders [1], exploit the temporal correlations of RSSI sequences, yet their gains over well-tuned RF or SVM are modest at residential scale while computational cost grows significantly [11,19].

2.3 Configurable Hardware: Motivation and Gap

Despite this methodological progress, the BLE hardware deployed in the surveyed studies is largely off-the-shelf. Consumer-grade beacons (e.g. Estimote Proximity Beacon[1], Accent Systems iBKS 105[2], Minew E8[3]) expose only a limited subset of the BLE 5.x stack to the application layer: transmission power is restricted to a few preset levels and access to the three primary advertising channels (37, 38, 39) is mediated by closed firmware. Since RSSI fingerprinting relies on signal diversity, identical beacons transmitting at identical power generate a fingerprint space whose separability degrades rapidly with the number of reference points [13,19]; heterogeneous power and channel configurations, by contrast, improve discriminability between adjacent rooms.

Work on custom-built BLE devices remains, to the best of our knowledge, restricted to a handful of prototypes. A recent configurable BLE 5.1 beacon that exposes such hardware-level parameters to the application layer was characterised by Díaz Jiménez et al. under controlled propagation conditions [3], and provides the physical infrastructure of the present study, whose specifications will be detailed later in this paper.

[1] https://developer.estimote.com/how-beacons-work/.
[2] https://accent-systems.com/es/producto/ibks-105/.
[3] https://www.minew.com/product/e8-asset-tracking-tag/.

3 Methods

3.1 System Architecture

The system architecture follows a multi-tier edge-fog-cloud model built on four components:

(i) a set of custom BLE beacons with configurable transmission power that periodically broadcast advertising packets;
(ii) an activity wristband that captures both biometric measurements (heart rate, battery level) and the RSSI from selected beacons;
(iii) a mobile device running a cross-platform application developed in Flutter, which connects to the wristband via BLE, instructs it on which beacons to scan, and retrieves the resulting scan data; and
(iv) a NestJS-based Representational State Transfer (REST) Application Programming Interface (API) that acts as the back-end and persists the collected data in a PostgreSQL relational database.

Both custom devices operate at the edge layer. The BLE beacons are compact form-factor devices designed and prototyped to overcome the limited configurability of commercial hardware [3]. They implement BLE 5.1 with the Zephyr Real-Time Operating System (RTOS) and broadcast advertising packets that include the device's Media Access Control (MAC) address and transmission power, with an adjustable power range from −20 to +8 dBm and configurable selection of advertising channels (37, 38, and 39). Their estimated battery life under continuous operation is three to four months. The activity wristband, also at the edge layer, is based on the Nordic Semiconductor nRF52833 System on Chip (SoC) (BLE 5.1, Zephyr RTOS) and integrates biometric sensors (photoplethysmography) and an Inertial Measurement Unit (IMU). It scans the configured beacons and exposes the acquired data to the upper tier through Generic Attribute Profile (GATT) services.

One tier up, at the fog layer, a mobile device runs the cross-platform Flutter application, which serves as the system orchestrator. It establishes a BLE connection with the wristband, issues GATT Write operations to specify which beacons to scan, and receives RSSI measurements and biometric data via GATT Notifications. The data are aggregated into batches and periodically forwarded to the cloud service; the mobile device performs no positioning computation.

At the top of the stack, the cloud layer hosts a server on which the REST API (implemented in NestJS) and the PostgreSQL relational database are deployed. This server receives the measurement batches from the mobile device, stores them in a structured form, and exposes the necessary endpoints for querying and exporting the data. The trained ML models for position estimation also reside at this layer.

Figure 1 summarises the distribution of the four components across the three tiers and the communication flows that connect them.

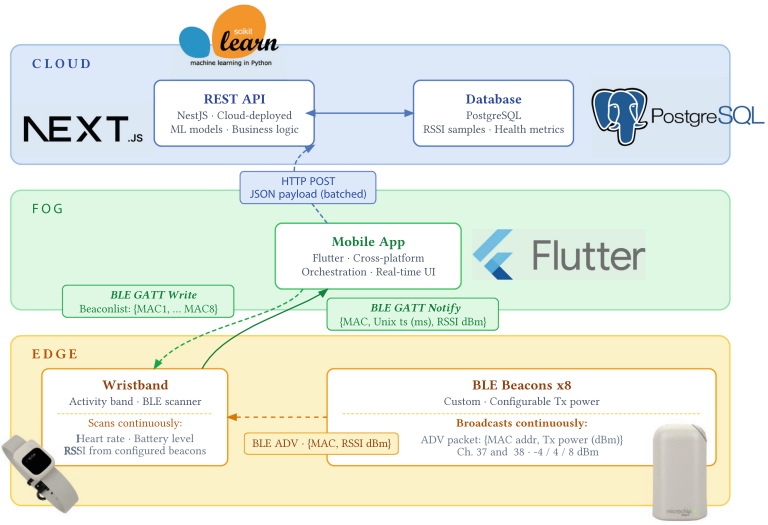


Fig. 1. Edge-fog-cloud architecture of the proposed positioning system, showing the components at each tier and their communication protocols.

3.2 Preprocessing Pipeline

Following an exploratory analysis of the dataset, a four-stage preprocessing pipeline was defined to improve the quality of the RSSI signals and adapt their structure to the requirements of the ML classifiers.

The first stage is outlier removal: anomalous RSSI samples identified during the exploratory analysis are discarded. Since these values correspond to spurious artefacts caused by constructive multipath interference or transient obstructions, row-wise deletion is preferred over imputation to avoid introducing distortions into the feature space.

The second stage is per-second vectorisation. Raw BLE packets arrive asynchronously and without uniform temporal ordering; the vectorisation step groups them into one-second intervals, takes the maximum RSSI per beacon in each interval, and constructs a fingerprint-time matrix. Time instants with no detection for a given beacon are encoded with a sentinel value below the global minimum, so that signal absence acts as an informative feature rather than missing data.

The third stage applies a sliding temporal window to smooth signal fluctuations without sacrificing real-time responsiveness. Different window sizes and aggregation operators (mean and maximum) are evaluated; the selection procedure is detailed in Sect. 4.2.

The fourth and final stage is class balancing via the Synthetic Minority Oversampling Technique (SMOTE). Some reference points accumulate fewer samples than others, which may bias classifiers towards majority classes. SMOTE synthesises new fingerprints for under-represented classes through interpolation

between real samples, and is applied exclusively to the training set to prevent any information leakage into the validation set.

Formally, the output of the pipeline is a labelled dataset $\{(\mathbf{x}_i, y_i)\}$, where each feature vector $\mathbf{x}_i \in \mathbb{R}^8$ contains the maximum RSSI recorded from each of the eight deployed beacons within one temporal window, and the class label $y_i \in \{1, \ldots, 28\}$ identifies the reference point at which the fingerprint was collected. Formally, this is a supervised multi-class classification problem: the goal is to learn a classifier $f : \mathbb{R}^8 \to \{1, \ldots, 28\}$ that maps an eight-dimensional RSSI observation to one of the 28 reference points while minimising the expected misclassification rate over unseen fingerprints.

3.3 Classifier Selection

The classifier selection covers a representative range of classical supervised algorithmic families. In BLE-based fingerprinting, the strong empirical dependency between received signal strength and proximity to a reference point means that classical classifiers consistently match or outperform deeper architectures at residential scale, while incurring substantially lower computational cost [11,19]. These are therefore the most widely adopted family in the area. Four families are included for systematic comparison [11,14]:

- kNN: the historical reference algorithm in fingerprinting-based IPS, providing a robust baseline for comparison [11]. The `ball_tree` index structure is adopted, reducing nearest-neighbour search complexity from $\mathcal{O}(N)$ to $\mathcal{O}(\log N)$, which is particularly advantageous in the high-dimensional RSSI feature space and benefits real-time deployment.
- SVM: included for its ability to construct non-linear decision boundaries in the high-dimensional RSSI feature space [14].
- RF: selected for its resilience to multipath fading noise and its provision of feature importance scores, which allow the contribution of each beacon to the prediction to be quantified.
- MLP: representative of shallow neural networks, introducing the perspective of neural models in the face of RSSI signal variability encountered indoors.

Because consecutive fingerprints are classified independently, these models do not exploit the sequential structure of human motion; this is a deliberate trade-off favouring low-latency, low-power operation, with sequence models such as Long Short-Term Memory (LSTM) networks or Transformers left to future work.

4 Experiments and Results

4.1 Use Case

The use case was set in a real residential dwelling of approximately $100\,\mathrm{m}^2$ divided into several rooms separated by walls, thus reproducing real-world NLOS

propagation conditions, namely attenuation, multipath, and interference, representative of AAL environments. Eight custom BLE beacons were deployed at strategic elevated positions to maximise coverage and ensure that the signal path from each beacon traverses at least one physical barrier, and 28 reference points were defined, distributed uniformly across all rooms. The resulting dataset is publicly available to the research community [10]. The spatial distribution of beacons and reference points is shown in Fig. 2.

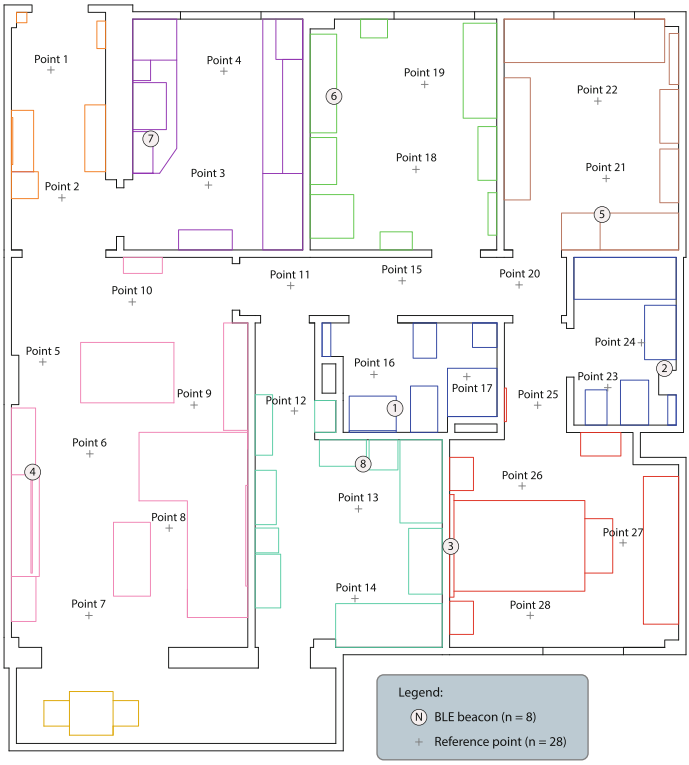


Fig. 2. Floor plan of the experimental dwelling ($\approx$100 m^2) with the positions of the 28 reference points and the 8 BLE beacons.

Table 1 details the configuration of each beacon. Transmission power was deliberately made heterogeneous to diversify the RSSI feature space and increase discriminability between adjacent rooms; this configuration was determined from prior anechoic chamber measurements [3].

RSSI fingerprints were acquired at each reference point using the mobile application, yielding a total of 31 978 raw samples—each corresponding to a single Beacon–RSSI detection event received by the wristband—collected over approximately 67 min. After per-second vectorisation, these reduce to 3 878 one-second time-instance vectors; following the 10-sample sliding window with the

Table 1. BLE beacons deployed in the experimental dwelling: identifier, MAC address, transmission power, and advertising channels.

Beacon	MAC Address	Tx Power (dBm)	Adv. channels
B1	C5:C1:22:55:5A:69	4	37, 38
B2	C6:8C:6E:50:1C:66	−4	37, 38
B3	CA:FA:83:43:9D:6A	8	37, 38
B4	CB:F4:EA:1D:3F:31	4	37, 38
B5	DB:36:66:4B:BD:C4	4	37, 38
B6	EC:0F:79:C9:CB:8A	4	37, 38
B7	F2:34:2D:BE:43:1E	4	37, 38
B8	F2:7C:81:CB:BB:5B	4	37, 38

maximum operator, 3 869 eight-dimensional RSSI feature vectors are obtained for model training and evaluation.

4.2 Experimental Setup

After applying the preprocessing pipeline described in Sect. 3.2, sliding temporal windows of 5, 10, and 15 samples were evaluated with mean and maximum aggregation operators, and the resulting distributions were estimated via kernel density estimation. The 5-sample window proved insufficient to filter background noise, while the 15-sample window introduced excessive latency for real-time localization. A window of 10 samples with the maximum operator was therefore selected: in BLE propagation, the maximum value typically corresponds to the packet that has undergone the fewest reflections, thereby providing a more faithful proximity estimate at low computational cost. For evaluation, a grouped k-fold cross-validation scheme with $k = 7$ was adopted (an $\approx$85 %/15 % train/validation ratio). Folds are formed at the level of temporal groups: all overlapping windows from the same contiguous acquisition segment are kept in the same fold, so no validation fingerprint shares its temporal support with a training one. This removes the optimistic bias that a random sample-wise split would introduce through autocorrelation between adjacent windows, while preserving the per-class (reference-point) proportion within each fold and providing a robustness indicator: high fold-to-fold variability would point to signal instability at a specific reference point.

Hyperparameter tuning was performed via randomised search over the configuration space rather than exhaustive grid search ($n_{\text{iter}} = 10$ candidate configurations per classifier). Given the large number of hyperparameters in models such as RF and MLP, exhaustive search incurs an exponential computational cost; randomised search explores the same space within a fixed evaluation budget, reducing training time without significantly compromising the quality of the final model. Moreover, randomised search is a standard and competitive alternative to exhaustive grid search under a fixed budget, and the common

$n_{\text{iter}} = 10$ setting keeps the comparison computationally homogeneous across the four classifier families. It is worth noting that the optimal RF configuration includes `class_weight=balanced`, which already incorporates an internal class imbalance correction through sample re-weighting in each tree; the effect of the SMOTE applied at the preprocessing stage is therefore marginal for this model. The same pipeline is nonetheless kept for all classifiers to ensure a homogeneous comparison. The five reported metrics are: Accuracy, Precision, Recall, and F1-Score (the latter three computed as weighted averages across the 28 classes, weighting each class by its support in the validation fold), and the standard deviation of accuracy across folds. The system performs symbolic, room-level localization, so it is assessed with per-class classification metrics rather than a metric coordinate error: each reference point is treated as a discrete class, and a Euclidean distance in metres is not defined under this formulation.

4.3 Results

Table 2 reports, for each classifier, the best hyperparameter combination identified by randomised search along with the corresponding cross-validation accuracy.

Table 2. Best hyperparameter combination and cross-validation accuracy for each classifier.

Algorithm	Accuracy	Best hyperparameters
RF	99.403 %	`n_estimators`=364, `max_depth`=30, `max_features`=log2, `min_samples_split`=3, `min_samples_leaf`=2, `bootstrap`=True, `class_weight`=balanced
kNN	99.080 %	`n_neighbors`=7, `weights`=distance, `metric`=manhattan, `p`=2, `algorithm`=ball_tree, `leaf_size`=43
SVM	98.459 %	`C`=5, `kernel`=rbf, `gamma`=scale, `shrinking`=True
MLP	98.235 %	`hidden_layer_sizes`=(100,), `activation`=relu, `solver`=lbfgs, `alpha`=0.00765, `learning_rate`=invscaling, `learning_rate_init`=0.00668, `batch_size`=32, `early_stopping`=False

All classifiers exceed 98 % accuracy, confirming both the quality of the acquired dataset and the effectiveness of the preprocessing pipeline. RF leads the ranking at 99.403 %, followed by kNN (99.080 %), SVM (98.459 %), and MLP (98.235 %).

Table 3 presents the validation metrics obtained on data unseen during training. Classifiers are ordered by decreasing accuracy.

Table 3. Validation metrics of the classifiers on unseen RSSI fingerprints.

Algorithm	Accuracy	Std. Dev.	Precision	Recall	F1
RF	99.403 %	0.3584 %	0.9944	0.9940	0.9940
kNN	99.080 %	0.4531 %	0.9914	0.9908	0.9908
SVM	98.459 %	0.6536 %	0.9856	0.9846	0.9846
MLP	98.235 %	0.3652 %	0.9831	0.9823	0.9823

The figures in Table 3 are telling. The fact that all models exceed 98 % accuracy under real-world NLOS conditions, with no post-hoc signal correction, suggests that the custom beacons generate a sufficiently discriminative and reproducible RSSI fingerprint space: hardware quality is thus validated by the system's own performance. The consistency between Precision, Recall, and F1, all within [0.98, 0.99] across all classifiers, reinforces this reading, since an imbalance between these metrics would indicate saturation in particular rooms or localised signal instability. RF achieves the lowest standard deviation (0.358 %), confirming that its performance is also the most stable across folds and that the model has generalised robust radio-frequency patterns; a particularly relevant property for real-world deployments. SVM, by contrast, exhibits the highest fold-to-fold variability (0.654 %), suggesting greater sensitivity to the signal fluctuations inherent in NLOS environments. Considering accuracy, stability, and RF's ability to quantify each beacon's contribution jointly, this classifier is the most suitable choice for the proposed deployment.

Because temporally adjacent windows are confined to the same fold (Sect. 4.2), this accuracy cannot stem from autocorrelation leakage; the residual optimism is due to the single-dwelling scope and room-level granularity examined in Sect. 5. In efficiency terms, the four models are lightweight at inference, kNN uses a ball-tree index and RF, SVM and the shallow MLP incur negligible per-sample cost, suiting the cloud-tier deployment of Sect. 3.1.

5 Conclusions and Future Work

Indoor localization in AAL environments presents demands that GNSS cannot meet. This work approaches the problem from the hardware perspective: two custom, low-power BLE devices have been experimentally validated as the physical infrastructure of a fingerprinting-based IPS.

To this end, an edge-fog-cloud architecture was designed integrating beacons with configurable power, an activity wristband, a mobile device, and a cloud back-end. On a dataset of 31 978 RSSI samples collected at 28 reference points in a real residential dwelling of 100 m^2, a preprocessing pipeline was applied (outlier removal, per-second vectorisation, 10-sample maximum temporal window, and SMOTE), and four classifiers (kNN, SVM, RF, and MLP) were compared under stratified 7-fold cross-validation. All surpass 98 % accuracy on unseen data, confirming both the pipeline's effectiveness and the quality of the signals produced

by the custom hardware. RF stands out as the most balanced option (99.403 %), combining low noise sensitivity, computational efficiency, and the ability to quantify each beacon's importance.

The present validation has two intended boundaries. First, it covers a single single-floor dwelling of 100 m^2; larger, multi-floor or more heavily obstructed layouts, where beacon density, fingerprint dimensionality and data volume interact differently, remain to be assessed and motivate the scalability line below. Second, the radio map relies on 28 reference points, that is, room-level granularity: this density is a deliberate design choice that bounds the one-off offline calibration effort while matching the symbolic localization required by AAL services.

Regarding future work, three lines of continuation are identified: (i) Extending the validation to larger-area environments to evaluate the system's scalability; (ii) Integrating the AoA feature of BLE 5.1 and exploring sequential architectures such as LSTM networks or Transformers to exploit the temporal correlations present in RSSI sequences and (iii) Adopting Federated Learning to retrain the models without location data leaving the user's device, thereby strengthening data privacy.

Acknowledgments. This work was partially supported by grant PID2024-156412OB-I00 funded by MICIU/AEI/10.13039/501100011033 and by "ERDF/EU" and grant M.2 PDC_000756 funded by Consejería de Universidad, Investigación e Innovación and by ERDF Andalusia Program 2021-2027. J. Torres-Sospedra gratefully acknowledge funding from Generalitat Valenciana (CIDEXG/2023/17, Conselleria d'Educació, Universitats i Ocupació).

Disclosure of Interests. The authors declare that they have no relevant conflicts of interest with respect to the content of this article.

References

1. Ai, H., Sun, X., Tao, J., Liu, M., Li, S.: DRVAT: Exploring RSSI series representation and attention model for indoor positioning. Int. J. Intell. Syst. **37**(7), 4065–4091 (2021). https://doi.org/10.1002/int.22712
2. Aranda, F.J., Parralejo, F., Álvarez, F.J., Torres-Sospedra, J.: Multi-slot BLE raw database for accurate positioning in mixed indoor/outdoor environments. Data **5**(3), 67 (2020). https://doi.org/10.3390/data5030067
3. Díaz-Jiménez, D., López Ruiz, J.L., Cuevas-Martínez, J.C., Torres-Sospedra, J., Navarro, E.A., Espinilla Estévez, M.: BLE-based custom devices for indoor positioning in ambient assisted living systems: Design and prototyping. Sensors **25**(20), 6499 (2025). https://doi.org/10.3390/s25206499
4. Echizenya, K., Kondo, K.: Indoor position and movement direction estimation system using DNN on BLE beacon RSSI fingerprints. Adv. Sci. Technol. Eng. Syst. J. **7**(3), 129–138 (2022). https://doi.org/10.25046/aj070315
5. Ghozali, R.P., Kusuma, G.P.: Indoor positioning system using regression-based fingerprint method. Int. J. Adv. Comput. Sci. Appl. **10**(8), 231–239 (2019). https://doi.org/10.14569/IJACSA.2019.0100829

6. González Díaz, N., Zola, E., Martín-Escalona, I.: Assessing the impact of coupling RTT and RSSI measurements in fingerprinting Wi-Fi indoor positioning. In: Proceedings of the 26th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems, pp. 19–26. Association for Computing Machinery (2023). https://doi.org/10.1145/3616388.3617528
7. Han, T.T., Nguyen Dinh, P., Nguyen Duc, T., Nguyen Long, V., Dinh Tan, H.: An efficient method for BLE indoor localization using signal fingerprint. EAI Endorsed Trans. Ind. Netw. Intell. Syst. **12**(1) (2024). https://doi.org/10.4108/eetinis.v12i1.6571
8. Kordi, K.A., Roslee, M., Alias, M.Y., Alhammadi, A., Waseem, A., Osman, A.F.: Survey of indoor localization based on deep learning. Comput. Mater. Continua **79**(2), 3261–3298 (2024). https://doi.org/10.32604/cmc.2024.044890
9. Liu, L., Zhao, Q., Miki, S., Tokunaga, J., Ebara, H.: Indoor fingerprinting positioning system using deep learning with data augmentation. Sensors Mater. **34**(8), 3047 (2022). https://doi.org/10.18494/SAM3912
10. López Ruiz, J.L., Díaz Jiménez, D., Espinilla, M.: UJALoc-BLE-NLOS: Real-World BLE Dataset for Indoor Localization under NLOS Conditions (2026), https://zenodo.org/records/19076975, Accessed 25 June 2026
11. Martín-Frechina, S., Dura, E., Miralles, I., Torres-Sospedra, J.: From fingerprinting to advanced machine learning: A systematic review of Wi-Fi and BLE-based indoor positioning systems. Sensors **25**(22), 6946 (2025). https://doi.org/10.3390/s25226946
12. Nabati, M., Ghorashi, S.A.: A real-time fingerprint-based indoor positioning using deep learning and preceding states. Expert Syst. Appl. **213**(118889), 118889 (2023). https://doi.org/10.1016/j.eswa.2022.118889
13. Qian, J., Komuro, N., Kim, W.S., Yoo, Y.: Enhanced position estimation via RSSI offset correction in BLE fingerprinting-based indoor positioning. Future Internet **17**(10), 440 (2025). https://doi.org/10.3390/fi17100440
14. Samavedam, A., Hallisey, V.: Indoor localization: BLE, machine learning, and Kalman filtering for resource-constrained devices. J. Student Res. **14**(1) (2025). https://doi.org/10.47611/jsrhs.v14i1.8702
15. Singh, N., Choe, S., Punmiya, R.: Machine learning based indoor localization using Wi-Fi RSSI fingerprints: An overview. IEEE Access **9**, 127150–127174 (2021). https://doi.org/10.1109/ACCESS.2021.3111083
16. Sonny, A., Kumar, A., Cenkeramaddi, L.R.: A survey of application of machine learning in wireless indoor positioning systems (2024). https://doi.org/10.48550/arXiv.2403.04333
17. Tahat, A., Smadi, M.N., Syouf, M.: A practical evaluation of ML algorithms for a tag-based BLE indoor positioning system. Int. J. Online Biomed. Eng. **16**(8), 39–54 (2020). https://doi.org/10.3991/ijoe.v16i08.14795
18. Zafari, F., Gkelias, A., Leung, K.K.: A survey of indoor localization systems and technologies. IEEE Commun. Surv. Tutorials **21**(3), 2568–2599 (2019). https://doi.org/10.1109/COMST.2019.2911558
19. Zholamanov, B., et al.: RSSI fingerprint-based indoor localization solutions using machine learning algorithms: A comprehensive review. Smart Cities **8**(5), 153 (2025). https://doi.org/10.3390/smartcities8050153