Real-World Robust Indoor Positioning in Smart Museums

Alessio Prete[1], Alessio Ferrato[1] and Giuseppe Sansonetti[1,*]

[1]*Department of Engineering, Roma Tre University, 00146 Rome, Italy*

Abstract

Indoor Positioning Systems (IPSs) based on Bluetooth Low Energy (BLE) are a promising solution for enabling location-aware services in smart cultural environments, where Global Navigation Satellite Systems are ineffective. However, real-world deployment remains challenging due to signal instability, infrastructure failures, and device heterogeneity. In this paper, we present an experimental study of BLE fingerprinting in a real museum environment using the BAR dataset collected at Palazzo Barberini in Rome. We evaluate machine learning approaches, focusing on Random Forest, under realistic conditions, including intra-device, cross-platform, and infrastructure-degradation scenarios. Results show near-perfect accuracy in controlled settings and strong robustness to beacon failures. However, performance drops significantly in cross-device scenarios due to hardware-induced domain shift. Despite this, Top-3 accuracy remains high, preserving user experience. Sensor fusion provides only marginal improvements, indicating a ceiling effect. These findings highlight that robustness and cross-device generalization, rather than raw accuracy, are the key challenges for real-world IPS deployment. This work contributes to social and cultural integration by enhancing the accessibility, reliability, and engagement of cultural spaces for diverse user groups, thereby fostering broader participation in cultural heritage experiences.

Keywords

Indoor positioning, Bluetooth Low Energy, Smart museums, Device heterogeneity, Machine learning, User experience

1. Introduction

The digital transformation of cultural heritage environments has led to the emergence of *smart museums*, where user experience is enhanced through context-aware services, such as automatic multimedia guides, personalized navigation, and crowd-flow analysis. In this scenario, Indoor Positioning Systems (IPSs) represent a fundamental enabling technology, providing the spatial awareness needed to deliver such services in environments where Global Navigation Satellite Systems are ineffective. Among the available solutions, Bluetooth Low Energy (BLE) has become a practical standard for museum deployments due to its low cost, limited infrastructural impact, and native compatibility with modern smartphones, supporting Bring Your Own Device (BYOD) paradigms. However, despite extensive research on BLE-based fingerprinting techniques, a significant gap remains between laboratory-level performance and real-world deployment. In practice, IPSs must operate under highly challenging conditions, including stochastic signal fluctuations, complex architectural layouts, and partial infrastructure failures. More importantly, a critical yet often underestimated issue is *device heterogeneity*: different smartphones perceive the same physical signal differently due to variations in antenna design, chipset calibration, and operating system constraints. This results in a domain shift that can severely degrade the performance of models trained on a specific device, thereby limiting their generalization capability in real-world scenarios.

In this paper, we present a study of BLE-based indoor positioning in a real museum environment, leveraging the BAR (Barberini Art Recognition) dataset collected at the Gallerie Nazionali di Arte Antica in Rome [1]. Unlike many prior works evaluated in controlled settings, our study focuses explicitly on robustness and cross-device generalization under realistic operational constraints. We adopt a machine learning approach based on Random Forest models and compare it against classical geometric

Joint Proceedings of the ACM UMAP Workshops 2026, UMAP 2026, June 8–11, 2026, Gothenburg, Sweden

*Corresponding author.

ale.prete1@stud.uniroma3.it (A. Prete); alessio.ferrato@uniroma3.it (A. Ferrato); gsansone@dia.uniroma3.it (G. Sansonetti)

0009-0006-6090-2905 (A. Prete); 0000-0001-5405-3018 (A. Ferrato); 0000-0003-4953-1390 (G. Sansonetti)

methods, such as k-Nearest Neighbors (k-NN), and state-of-the-art gradient boosting techniques, such as XGBoost. Through experimental trials, we analyze system performance across multiple scenarios, including intra-device evaluation, cross-platform transfer, and simulated infrastructure failures. The main contributions of this work are as follows:

- **Robust IPS in real-world environments:** We demonstrate that tree-based models can achieve near-perfect accuracy in complex museum settings, significantly outperforming traditional fingerprinting approaches.
- **Infrastructure resilience analysis:** We show that the proposed system exhibits gradual degradation, maintaining operational performance even under substantial beacon failures.
- **Cross-platform evaluation:** We provide a detailed analysis of device heterogeneity, highlighting a significant performance drop when transferring models across devices, and identifying domain shift as the primary limitation of BLE-based IPSs.
- **User-centric evaluation:** We demonstrate that, despite reduced Top-1 accuracy in cross-device scenarios, high Top-3 accuracy can preserve user experience, suggesting practical mitigation strategies.

Overall, our findings indicate that while accuracy in controlled settings is largely solved, robustness and cross-device generalization remain open challenges for the real-world deployment of indoor positioning systems. This shifts the focus from purely algorithmic improvements to adaptive, user-aware solutions for practical smart environments.

Although this work is primarily situated within the domains of indoor positioning and smart environments, it also addresses broader themes related to inclusive design, adaptive systems, and social and cultural integration. More specifically, the proposed approach provides a robust technological foundation for enabling personalized, accessible, and socially inclusive experiences within cultural spaces, thereby fostering wider participation among diverse, underrepresented, and marginalized user groups. Prior literature has extensively highlighted the role of museums as valuable environments for social and cultural participation, inclusion, and educational engagement (e.g., see [2, 3]). Moreover, recent studies have shown that socially aware and empathetic interactive technologies, including museum robots, can positively influence users' engagement and their perception of artificial agents in cultural contexts [4]. These findings further underscore the importance of designing robust, user-centered smart museum infrastructures that enable adaptive, emotionally engaging experiences.

2. Related Work

The research literature offers a rich body of contributions on approaches and technologies aimed at enhancing the accessibility and enjoyment of cultural heritage, while also fostering social and cultural inclusion. These efforts extend beyond traditional museum settings [5, 6, 7, 8] to include other cultural spaces [9, 10], with the objective of making cultural experiences more inclusive, participatory, and accessible to diverse audiences. For this purpose, indoor positioning can certainly also contribute. It has been extensively studied as a key enabling technology for context-aware and personalized services in smart environments. In cultural heritage scenarios, IPS solutions must satisfy constraints that are often less critical in other domains: low installation impact, limited maintenance cost, compatibility with visitors' personal devices, and respect for the aesthetics and preservation requirements of historical buildings [11]. Several technologies have been proposed for indoor localization, including Wi-Fi, Ultra-Wideband (UWB), computer vision, inertial sensing, and BLE [12, 13]. Wi-Fi fingerprinting has historically been one of the earliest and most widely used approaches [14], but recent restrictions imposed by mobile operating systems on scanning frequency and background access limit its suitability for real-time museum applications. UWB can achieve very high accuracy, but its deployment cost and limited availability on commodity smartphones reduce its applicability in BYOD settings. Vision-based systems may provide fine-grained localization, but they introduce privacy concerns and require favorable visual conditions. BLE, instead, offers a practical trade-off between cost, energy consumption, ease of installation, and smartphone compatibility, making it particularly suitable for smart museum deployments. Among BLE-based localization techniques, fingerprinting is one of the most widely adopted approaches. It relies on an offline phase, where radio signatures are collected at known positions, and an online phase, where new signal observations are matched against the radio map. Classical fingerprinting methods are often based on geometric similarity, with k-NN being a common baseline [15]. Although effective in controlled environments, these methods are sensitive to the stochastic nature of Received Signal Strength Indicator (RSSI) measurements, which are affected by multipath propagation, human-body absorption, occlusions, and temporal fluctuations. In complex historical buildings, where walls and room layouts strongly distort radio propagation, simple distance-based matching may fail to capture the non-linear structure of the signal space. To overcome these limitations, machine learning methods have increasingly been adopted for indoor positioning [16]. Tree-based models, such as Random Forest [17] and gradient boosting [18], are particularly attractive for RSSI fingerprinting because they can model non-linear decision boundaries, handle sparse feature spaces, and remain computationally efficient at inference time. Compared with deep learning approaches [19, 20], they also require less training data and provide better interpretability, which is relevant for real-world deployments where maintenance and debugging are important. While deep neural networks have shown promising results in large-scale indoor localization datasets, their computational cost, data requirements, and black-box nature can be problematic in museum scenarios, where calibration campaigns are expensive and mobile energy consumption is a practical constraint. A central open issue in BLE fingerprinting is device heterogeneity. RSSI values are not determined only by distance and environmental conditions, but also by the receiving device, whose antenna design, chipset, and operating system policies introduce systematic differences in the measured signal [21]. As a consequence, a radio map collected with one smartphone may not generalize to another. Prior work has addressed this issue through calibration functions, differential fingerprinting, and domain adaptation strategies [22]. However, linear calibration assumes a simplified relationship between devices, while differential signal-strength methods may amplify noise and discard useful absolute signal information. Therefore, cross-device generalization remains a major barrier to real-world adoption. This work builds on this line of research by focusing not only on accuracy but also on robustness under realistic deployment conditions. Unlike studies that evaluate IPS models primarily in ideal or single-device settings, we analyze BLE fingerprinting in a real museum environment, accounting for infrastructure degradation, Android-to-iOS transfer, and user-centric Top-k accuracy. In doing so, we position indoor positioning not merely as a classification problem, but as a reliability challenge for context-aware services in smart cultural spaces.

3. Methodology

We address indoor positioning as a multi-class classification problem based on BLE fingerprinting. The proposed approach is designed to be robust to signal instability, missing data, and real-world deployment constraints.

3.1. Fingerprint Construction

Raw BLE scans are inherently asynchronous, as beacons transmit independently. To obtain stable representations, we adopt a sliding window strategy with a window size of three seconds and a step of one second. Within each window, multiple RSSI samples are collected for each beacon. To mitigate the impact of noise and outliers caused by multipath propagation and transient interference, we aggregate RSSI values using the median rather than the mean. This choice provides a more robust estimate of the signal distribution without requiring additional filtering. Missing values are handled through a deterministic imputation strategy: non-detected beacons are assigned a fixed value of -110 dBm, representing signal absence below the receiver sensitivity threshold. This allows the model to distinguish between weak and non-observable signals. Each fingerprint is therefore represented as a fixed-length vector containing RSSI values for all known beacons.

3.2. Model Selection

We adopt a Random Forest classifier as the primary model, due to its ability to capture non-linear relationships and its robustness to noisy and sparse data. Hyperparameters are optimized via grid search with three-fold cross-validation, exploring the number of trees, maximum depth, and minimum split size. For comparison, we also implement three alternative models: k-NN, a geometric baseline based on the Manhattan distance; Logistic Regression, a linear probe to assess the separability of the feature space; and XGBoost, a state-of-the-art gradient boosting approach. These models allow us to evaluate the impact of non-linearity, model complexity, and learning paradigms on localization performance.

3.3. Inference Pipeline

During operation, incoming BLE packets are continuously buffered within a sliding window. At the end of each window, a fingerprint is generated using median aggregation and missing value imputation, and the trained model predicts the user's location. This design ensures low-latency inference while maintaining robustness against transient signal fluctuations.

4. Experimental Setup

This section presents the experimental framework used to evaluate the proposed localization approach under realistic conditions, including the dataset, devices, evaluation scenarios, and validation protocol adopted to assess performance, robustness, and generalization capabilities.

4.1. Environment and Dataset

Experiments are conducted on the BAR dataset, collected at the Gallerie Nazionali di Arte Antica, Palazzo Barberini, in Rome [1]. This environment represents a challenging real-world scenario characterized by thick walls, irregular layouts, and uncontrolled visitor flows, all of which significantly affect radio propagation. The infrastructure consists of BLE beacons configured with standard iBeacon parameters [23]. Their spatial distribution is non-uniform, reflecting both architectural constraints and exhibition design.

4.2. Devices and Data Collection

To evaluate device heterogeneity, data are collected using two representative platforms: Android (Samsung Galaxy S21) and iOS (iPhone 16). The dataset is organized into multiple acquisition sessions, separated in time and device, enabling both temporal and cross-platform evaluation.

4.3. Evaluation Scenarios

We define three evaluation scenarios to assess system performance under realistic conditions:

- **Intra-device:** training and testing on the same device, representing upper-bound performance.
- **Cross-platform:** training on one device and testing on another, measuring domain shift.
- **Robustness stress test:** progressive removal of beacon signals to simulate infrastructure failures.

4.4. Validation Protocol

To ensure statistical reliability, we adopt a stratified 10-fold cross-validation protocol. This allows us to estimate both average performance and variance across different data splits. Additionally, paired statistical tests are used to verify whether performance differences between models are significant. To evaluate resilience to infrastructure failures, we apply an anomaly-injection procedure that simulates beacon outages by forcing selected features to the signal-absence value of -110 dBm. This enables a systematic analysis of performance degradation as sensor loss increases. In addition to standard

Table 1
Intra-device performance comparison under 10-fold cross-validation.

Model	Mean accuracy	Standard deviation
k-NN	87.07%	±1.33%
Random Forest	**99.78%**	±0.16%
XGBoost	99.41%	±0.32%

Table 2
Artwork-level accuracy under simulated beacon failures.

Removed beacons	Random Forest	k-NN	Logistic Regression	Proximity
0	**89.77%**	85.70%	55.64%	64.18%
2	**89.65%**	85.23%	53.30%	63.07%
4	**89.88%**	84.65%	50.91%	61.37%
10	**84.77%**	81.73%	48.01%	60.91%
20	**81.70%**	79.18%	37.08%	54.01%

Top-1 accuracy, we evaluate Top-*k* accuracy, specifically Top-3, to better reflect the user experience. In practical applications, suggesting a small set of likely nearby locations can mitigate the impact of uncertainty in signal-based positioning.

5. Experimental Results

In this section, we present the experimental results and discuss the key findings with a focus on robustness and real-world applicability.

5.1. Overall Performance

We first evaluate the proposed approach in an intra-device scenario, representing ideal operating conditions. The Random Forest model achieves near-perfect performance, with an average accuracy of up to 99.78%, significantly outperforming both the geometric baseline and the proximity-based approach (see Table 1). Compared to k-NN, the proposed model reduces the residual error by an order of magnitude. This confirms that indoor positioning in complex environments is not merely a geometric-matching problem but rather a nonlinear pattern recognition task. The improvement is statistically significant and consistent across cross-validation folds, with very low variance, indicating high model stability. Interestingly, while XGBoost also achieves strong performance, Random Forest slightly outperforms it. This suggests that, in highly noisy RSSI environments, bagging-based ensemble methods may generalize better than boosting approaches, which are more prone to overfitting noise.

5.2. Robustness to Infrastructure Failures

A key requirement for real-world deployment is resilience to partial infrastructure failure. To assess this, we simulate beacon outages through anomaly injection. The results show a clear *gradual degradation* behavior (see Table 2). The model maintains high accuracy even when multiple beacons are removed, remaining above 80% accuracy under severe degradation conditions. Small failures, such as removing a few beacons, have a very low impact and, in some cases, even improve performance by removing noisy features. This behavior highlights an important advantage of ensemble models: the decision process is distributed across multiple features and trees, providing intrinsic redundancy. In contrast, k-NN exhibits a sharper performance decline, as its distance-based formulation is directly affected by the loss of feature dimensions.

Table 3
User-centric Top-1 and Top-3 accuracy comparison.

Scenario	Top-1	Top-3	Gain
Android intra-device	88.80%	97.92%	+9.12
iOS intra-device	66.37%	93.71%	+27.34

5.3. Cross-Platform Generalization

The most critical experiment concerns cross-platform evaluation. When the model is trained on Android data and tested on iOS, accuracy drops dramatically from near-perfect values to approximately 60% (see Table 3). This result reveals a fundamental limitation of RSSI-based fingerprinting: a strong domain shift induced by device heterogeneity. Different smartphones perceive the same physical signal differently due to variations in antenna design, hardware calibration, and operating system constraints. As a result, the feature distributions shift, leading to misaligned decision boundaries. Error analysis shows that this degradation is not random. Instead, misclassifications are structured and tend to occur between spatially adjacent locations. This indicates that the underlying topological structure of the environment is preserved but shifted in the feature space. These findings suggest that improving cross-device generalization requires explicit domain-adaptation strategies rather than simply increasing model complexity.

5.4. User-Centric Evaluation

From a practical perspective, strict Top-1 accuracy may not fully capture the quality of the user experience. When evaluating Top-3 accuracy, the system achieves values above 90% even in challenging conditions (see Table 3). This result is particularly relevant for museum applications, where nearby artworks may have similar signal fingerprints. In such cases, presenting a small set of likely candidates can effectively mitigate positioning uncertainty while preserving usability. We also observe a pronounced discrepancy between Top-1 and Top-3 performance on iOS devices, highlighting a phenomenon we refer to as the *iOS paradox*: while exact identification is difficult due to signal instability, the correct location is often included among the top-ranked predictions. This reinforces the importance of user-centric evaluation metrics in real-world systems.

5.5. Impact of Sensor Fusion

We further investigate whether integrating inertial sensor data improves performance. The results show only a marginal gain of 0.07 percentage points, though it is statistically significant. This limited improvement can be explained by a *ceiling effect*: in environments with highly distinctive radio signatures, the additional information provided by inertial sensors does not substantially reduce uncertainty. Moreover, the increased computational and energy costs associated with high-frequency sensor processing make this approach less attractive for mobile deployments.

5.6. Discussion

Overall, the experimental results highlight a key insight: while high accuracy can be achieved in controlled conditions, real-world deployment introduces challenges that are not captured by standard evaluation protocols. In particular, robustness and cross-device generalization emerge as the primary limitations of BLE-based indoor positioning. The proposed Random Forest approach effectively addresses signal noise and infrastructure variability, but remains sensitive to domain shifts caused by hardware heterogeneity. These findings suggest a shift in research focus: from maximizing accuracy under ideal conditions to designing adaptive, robust systems that operate reliably across devices and environments.

6. Limitations

Although the results are promising, the study has three main limitations. First, the evaluation is based on a single museum environment. Palazzo Barberini is a complex and representative historical building, but results may differ in modern open-space museums, where radio propagation has different characteristics. Second, while data were collected under real operating conditions, visitor density was not explicitly controlled or quantified; therefore, the isolated effect of crowding cannot be measured. Third, the robustness test simulates complete beacon failure by forcing signals into radio silence, but does not model partial malfunctions, such as battery-related power fluctuations or sensor physical displacement.

7. Conclusion and Future Work

In this paper, we investigated BLE-based indoor positioning in a real-world museum environment, focusing on robustness and cross-device generalization. Using the BAR dataset collected at Palazzo Barberini, we demonstrated that a Random Forest-based approach can achieve near-perfect accuracy under controlled conditions, significantly outperforming classical geometric methods and competing with state-of-the-art models. Beyond accuracy, our results highlight two key findings. First, the proposed approach exhibits strong robustness to infrastructure degradation, maintaining reliable performance even under partial beacon failures. Second, and more critically, we show that device heterogeneity introduces a substantial domain shift, leading to a significant drop in performance when models are transferred across platforms. This confirms that cross-device generalization remains a major unresolved challenge for real-world deployment of RSSI-based indoor positioning systems. From a user-centric perspective, we observe that high Top-k accuracy can mitigate the impact of positioning errors, preserving usability even in degraded scenarios. Additionally, we find that sensor fusion provides only marginal benefits in environments with highly distinctive radio signatures, suggesting that simpler BLE-only solutions may offer a better trade-off between accuracy and efficiency.

Overall, our study indicates that achieving high accuracy in controlled settings is no longer the primary bottleneck for indoor positioning. Instead, future research should focus on adaptive techniques that handle domain shift and ensure consistent performance across heterogeneous devices and dynamic environments. Promising directions include unsupervised domain adaptation for cross-device calibration, collaborative fingerprinting through crowdsourcing, and hybrid integration with visual positioning systems for fine-grained artwork-level localization.

We hope this work can contribute to social and cultural integration by making cultural spaces more accessible, reliable, and engaging for a wide range of users, thereby supporting broader participation in cultural heritage experiences.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT and Grammarly for the execution of the following tasks: Paraphrase and reword, Grammar and spelling check. After using these tools, the authors reviewed and edited the content as needed. The authors take full responsibility for the content of the publication.

References

[1] A. Ferrato, F. Gasparetti, C. Limongelli, S. Mastandrea, G. Sansonetti, J. Torres-Sospedra, Cross-platform smartphone positioning at museums, in: 2025 International Conference on Indoor Positioning and Indoor Navigation (IPIN), IEEE, 2025, pp. 1–6.

[2] A. I. González-Herrera, A. B. Díaz-Herrera, P. Hernández-Dionis, D. Pérez-Jorge, Educational and accessible museums and cultural spaces, Humanit. Soc. Sci. Commun. 10 (2023) 67.

[3] A. Ferrato, C. Gena, C. Limongelli, G. Sansonetti, Multimodal llm question generation for children's art engagement via museum social robots, in: Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct '25, ACM, New York, NY, USA, 2025, pp. 144–150.

[4] C. Gena, F. Manini, A. Lieto, A. Lillo, F. Vernero, Can empathy affect the attribution of mental states to robots?, in: E. André, M. Chetouani, D. Vaufreydaz, G. M. Lucas, T. Schultz, L. Morency, A. Vinciarelli (Eds.), Proceedings of the 25th International Conference on Multimodal Interaction, ICMI 2023, Paris, France, October 9-13, 2023, ACM, 2023, pp. 94–103.

[5] M. Mezzini, C. Limongelli, G. Sansonetti, C. De Medio, Tracking museum visitors through convolutional object detectors, in: Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '20 Adjunct, ACM, New York, NY, USA, 2020, pp. 352–355.

[6] A. Ferrato, C. Limongelli, M. Mezzini, G. Sansonetti, The meta4rs proposal: Museum emotion and tracking analysis for recommender systems, in: Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization, UMAP '22 Adjunct, ACM, New York, NY, USA, 2022, pp. 406–409.

[7] A. Ferrato, C. Limongelli, M. Mezzini, G. Sansonetti, A deep learning-based approach to model museum visitors, in: CEUR Workshop Proceedings, volume 3124, 2022, pp. 217–221.

[8] A. Ferrato, C. Limongelli, M. Mezzini, G. Sansonetti, Using deep learning for collecting data about museum visitor behavior, Applied Sciences 12 (2022).

[9] I. Amendola, F. Cena, L. Console, A. Crevola, C. Gena, A. Goy, S. Modeo, M. Perrero, I. Torre, A. Toso, Ubiquito: A multi-device adaptive guide, in: S. A. Brewster, M. D. Dunlop (Eds.), Mobile Human-Computer Interaction - Mobile HCI 2004, 6th International Symposium, Glasgow, UK, September 13-16, 2004, Proceedings, LNCS, Springer, 2004, pp. 409–414.

[10] C. Gena, I. Torre, The importance of adaptivity to provide onboard services: A preliminary evaluation of an adaptive tourist information service onboard vehicles, Appl. Artif. Intell. 18 (2004) 549–580.

[11] A. Chianese, F. Piccialli, Designing a smart museum: When cultural heritage joins iot, in: Proc. NGMAST 2014, IEEE, 2014, pp. 300–306.

[12] Y. Gu, A. Lo, I. Niemegeers, A survey of indoor positioning systems for wireless personal networks, IEEE Communications Surveys & Tutorials 11 (2009) 13–32.

[13] Y. Sartayeva, H. C. Chan, Y. H. Ho, P. H. Chong, A survey of indoor positioning systems based on a six-layer model, Computer Networks 237 (2023) 110042.

[14] P. Bahl, V. N. Padmanabhan, Radar: An in-building rf-based user location and tracking system, in: Proceedings IEEE INFOCOM 2000, volume 2, IEEE, 2000, pp. 775–784.

[15] J. Torres-Sospedra, R. Montoliu, A. Martínez-Usó, J. P. Avariento, T. J. Arnau, M. Benedito-Bordonau, J. Huerta, Ujiindoorloc: A new multi-building and multi-floor database for wlan fingerprint-based indoor localization problems, in: IPIN 2014, IEEE, 2014, pp. 261–270.

[16] A. Sonny, A. Kumar, L. R. Cenkeramaddi, A survey of application of machine learning in wireless indoor positioning systems, arXiv preprint arXiv:2403.04333 (2024).

[17] L. Breiman, Random forests, Machine learning 45 (2001) 5–32.

[18] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 785–794.

[19] I. Goodfellow, Y. Bengio, A. Courville, Y. Bengio, Deep learning, MIT Press Cambridge, 2016.

[20] X. Wang, L. Gao, S. Mao, S. Pandey, Csi-based fingerprinting for indoor localization: A deep learning approach, IEEE Transactions on Vehicular Technology 66 (2016) 763–776.

[21] F. Lemic, A. Behboodi, V. Handziski, A. Wolisz, Experimental decomposition of the performance of fingerprinting-based localization algorithms, in: 2014 International Conference on Indoor Positioning and Indoor Navigation (IPIN), IEEE, 2014, pp. 355–364.

[22] A. M. Hossain, Y. Jin, W.-S. Soh, H. N. Van, Ssd: A robust rf location fingerprint addressing mobile devices' heterogeneity, IEEE Transactions on Mobile Computing 12 (2011) 65–77.

[23] Apple Inc., Getting Started with iBeacon, Apple Developer, 2014.