

# Experience: Practical Indoor Localization for Malls

Yuming Hu<sup>1</sup>, Feng Qian<sup>1</sup>, Zhimeng Yin<sup>2</sup>, Zhenhua Li<sup>3</sup>, Zhe Ji<sup>4</sup>, Ye qiang Han<sup>4</sup>, Qiang Xu<sup>4</sup>, Wei Jiang<sup>5</sup>

<sup>1</sup>University of Minnesota, Twin Cities   <sup>2</sup>City University of Hong Kong   <sup>3</sup>Tsinghua University   <sup>4</sup>XYZ10 Technology  
<sup>5</sup>State Grid Corporation of China

## ABSTRACT

We report our experiences of developing, deploying, and evaluating MLoc, a smartphone-based indoor localization system for malls. MLoc uses Bluetooth Low Energy RSSI and geomagnetic field strength as fingerprints. We develop efficient approaches for large-scale, outsourced training data collection. We also design robust online algorithms for localizing and tracking users' positions in complex malls. Since 2018, MLoc has been deployed in 7 cities in China, and used by more than 1 million customers. We conduct extensive evaluations at 35 malls in 7 cities, covering 152K  $m^2$  mall areas with a total walking distance of 215 km (1,100 km training data). MLoc yields a median location tracking error of 2.4m. We further characterize the behaviors of MLoc's customers (472K users visiting 12 malls), and demonstrate that MLoc is a promising marketing platform through a promotion event. The e-coupons delivered through MLoc yield an overall conversion rate of 22%. To facilitate future research on mobile sensing and indoor localization, we have released a large dataset (43 GB at the time when this paper was published) that contains IMU, BLE, GMF readings, and the localization ground truth collected by trained testers from 37 shopping malls.

## CCS CONCEPTS

• **Information systems** → **Location based services**; *Sensor networks*; Global positioning systems; • **Networks** → **Location based services**.

## KEYWORDS

Indoor Localization; Bluetooth Low Energy; Geomagnetic Field.

### ACM Reference Format:

Yuming Hu, Feng Qian, Zhimeng Yin, Zhenhua Li, Zhe Ji, Ye qiang Han, Qiang Xu, Wei Jiang. 2022. Experience: Practical Indoor Localization for Malls. In *The 28th Annual International Conference On Mobile Computing And Networking (ACM MobiCom '22)*, October 24–28, 2022, Sydney, NSW, Australia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3495243.3517021>

## 1 INTRODUCTION

Indoor localization has been extensively researched in the past two decades. A Google Scholar search using “indoor localization” gives more than 40,000 results. In sharp contrast, large-scale deployment of indoor localization systems is far lagging behind. There are very

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*ACM MobiCom '22, October 24–28, 2022, Sydney, NSW, Australia*

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9181-8/22/10...\$15.00

<https://doi.org/10.1145/3495243.3517021>

limited news articles on commercial indoor localization deployments [8, 31, 32, 44, 45, 47], among which few offer deep insights into their deployment experiences.

This paper fills the above gap by reporting our experiences and findings of developing, deploying, and evaluating MLoc, a smartphone-based localization system for indoor malls (commercial complex buildings) typically with tens or hundreds of retail stores. MLoc helps customers find paths to stores (e.g., how to reach the nearest Starbucks) by providing accurate, easy-to-use localization and store-level navigation. Developing MLoc is very different from building an indoor localization prototype in the lab. It faces unique challenges, involves additional constraints, and requires judicious decisions considering numerous technical and non-technical factors, as elaborated below.

From the **infrastructural perspective**, an important decision is to select the appropriate physical signals as localization fingerprints. Given the plethora of research on WiFi-based localization [11, 12, 38, 56, 60] and the ubiquitous WiFi deployment in today's malls, we naturally sought to leverage WiFi as location fingerprints. However, we eventually rejected this design. We find that around half of our customers use iPhone, which does not offer public APIs for querying WiFi APs' RSSI. Even on Android devices, two unexpected factors render WiFi-based localization less feasible: many deeply customized Android systems have very low WiFi scanning frequency (e.g., every 20s), and commercial WiFi APs may periodically change MAC addresses for security consideration. MLoc instead uses the conventional Bluetooth Low Energy (BLE) RSSI and geomagnetic field (GMF) strength as the location fingerprints. BLE requires a light infrastructure consisting of cheap, small, battery-powered beacons, whereas GMF is infrastructure-free. We find their synergy can lead to an accuracy adequate for store-level navigation.

From the **data's perspective**, following the common wisdom in literature, MLoc adopts a landmark-based outsourcing approach (i.e., hiring human collectors to survey a few predefined landmarks) to collect BLE/GMF fingerprints and the ground truth location data. However, we note that the hired collectors are quite distinct from the knowledgeable collectors in academic research – they can easily miss certain landmarks, but meanwhile would like to move existing landmarks or even suggest adding new ones. To this end, we enhance the common approach by strategically restricting the landmark visiting paths; in addition, we respect collectors' on-site opinions by allowing them to improve the predefined landmarks (calculated based on imperfect floor plan information) through a dedicated, user-friendly GUI (§2.1).

From the **algorithmic perspective**, we choose not to build MLoc from scratch given the rich literature. Our main challenge thus lies in how to pick the suitable building blocks from existing works. We find that many sophisticated algorithms in the literature

aim at dealing with challenging cases in various indoor environments. In our domain of in-mall localization, surprisingly perhaps, we observe only several generic challenging areas (e.g., atrium, corridor dead ends, corridor connectors, and elevators) despite the malls' complex layouts, based on extensive field studies at 35 malls with different scales. Such an observation greatly simplifies our algorithm design. Encouragingly, we find that classical algorithms can be enhanced by simple yet strategic customizations such as fingerprint preprocessing, weight adjustment, and lightweight AI to tackle the challenging areas (§2.2–§2.5).

While we cannot claim the optimality of our solution, MLoc does achieve its design goal in a pragmatic sense. Since its debut in 2018, MLoc has been used by more than 1 million customers. Over more than one year, we conduct extensive evaluations at 35 malls in 7 cities in China, covering  $152K m^2$  localization areas.<sup>1</sup> Our evaluations consist of 4.3K paths with a total walking distance of 215 km (a path denotes a trajectory from the navigation source to the destination; its median length is around 45m). The underlying training data consists of 1,110 km walking, involving 21K BLE beacons and 244K landmarks. We find that MLoc yields a median location tracking error of 2.4m (10-th and 90-th percentile: 0.8m and 7.3m). Nevertheless, we do observe several limitations of MLoc, such as a long tail of high errors and imperfect floor detection. We find that many of these errors are attributed to non-algorithmic factors, such as beacon failures, poor device capability (e.g., some smartphone vendors throttle the BLE scanning frequency to save energy), and even poor Internet connectivity hindering the client-edge communication (§3).

Last but not least, we also make observations from the **business's perspective**. A key reason for the sharp disparity between the plethora of research and little commercial deployment of indoor localization is a lack of incentives and/or business models. To this end, we characterize the behaviors of MLoc's customers, based on 472K users' data collected from 12 malls over one year. The results suggest that MLoc is overall effective: 95% of the navigations ended at locations that are 20m within the destination storefronts. We also find that while MLoc can help many users find their destinations quickly, in at least 20% of the navigations, users spend more than 2 minutes using MLoc. Such an in-app session length during shopping time offers considerable business opportunities. For example, we demonstrate that MLoc is a promising marketing platform that can distribute targeted advertisements based on customers' real-time location. Through a sales event co-organized by MLoc and a large mall, we observe an ad conversion rate of 22%, significantly higher than those of online advertising [5] (§4).

**Ethical Concern.** All the analyses conducted in this paper comply with the agreement established between MLoc and its customers. No personally identifiable information (PII) was collected or used in this study. We never (and are unable to) correlate a user's location with his/her true identity.

**Dataset release.** To facilitate future research on mobile sensing and indoor localization, we have released a large dataset (43 GB at the time when this paper was published) that contains IMU,

<sup>1</sup>MLoc only serves malls' shared areas (e.g., corridors, atrium, stairs, rest areas). It does not cover areas inside stores (i.e., gross leasable areas, GLA) due to many stores' privacy policies.



Figure 1: BLE Beacon on ceiling.

BLE, GMF readings, and the localization ground truth collected by trained testers from 37 shopping malls. The dataset (including its detailed data format) can be downloaded at:

<https://kumius.github.io/MLoc/>

## 2 THE MLOC SYSTEM DESIGN

Not surprisingly, MLoc consists of two phases: *offline training*, where (fingerprint, location) pairs are collected to build a localization model, and *online inference*, where a user's smartphone collects fingerprints, uploads them to the edge, and obtains the location and/or navigation guidance in real time.

**Fingerprint Selection.** A wide range of physical signals (e.g., acoustics [43, 62], radio signals [10, 11, 14, 24, 25, 49, 57], visible light [26, 59], magnetic field [42, 63], and camera images [54]) can be leveraged as localization fingerprints. MLoc leverages two sources as fingerprints: signal strength of infrastructural Bluetooth Low Energy (BLE) beacons and geomagnetic field (GMF) strength in X, Y, Z dimensions. We make this design decision due to three main reasons.

- Both fingerprints are accessible on commodity devices (Android and iOS smartphones) through standard SDKs. Recall from §1 that we reject WiFi-based localization due to several practical difficulties: a lack of APIs on iOS, low scanning frequencies on deeply customized Android devices, and MAC address randomization observed on commodity APs.
- The management teams of all the shopping malls we approached only allow us to deploy small-sized, battery-powered hardware. Installation of additional power and networking cables is forbidden due to aesthetic considerations.
- The judicious combination of BLE and GMF can yield the desired accuracy (Figure 6, §3). We surveyed more than 4K stores in 12 malls, and find that only 3% of the storefronts have a width less than 5m. Therefore, a localization accuracy of  $\sim 5m$  is sufficient for store-level localization.

### 2.1 Offline Training Data Collection

**BLE Beacon Deployment.** Due to aesthetic considerations, we adopt small-sized, battery-powered beacons. Each beacon costs around 5 USD and has a small form factor of  $59 mm \times 59 mm \times 18 mm$ . It is equipped with a 2,400 mAh battery. The default broadcast interval is 200 ms unless otherwise stated.

According to our agreements with the malls, we are only allowed to deploy the beacons in shared areas (as opposed to the gross leasable areas, GLA) in a mall. We mount the beacons on the ceiling



Figure 2: An atrium (beacons in red circles).




Figure 3: Landmarks and N-shaped path.

of the corridor (Figure 1) or on the surrounding edges in atrium (open space) areas (Figure 2). The typical distance between two beacons is from 10 to 15m, and the typical beacon density is from 2 to 5 per  $1000m^2$ , depending on the specific layout of the area. For the areas where localization is challenging (e.g., atrium, corridor, and elevators, see the remainder of this section), we reduce the inter-beacon distance to 6m to ensure good localization accuracy. Following the above guidelines, we deployed MLoc in 35 shopping malls across 7 cities in China.

Once deployed, MLoc incurs small maintenance overheads, which mostly come from replacing fallen beacons due to glue failure. This is handled by shopping malls' management teams. In the long run, another major type of maintenance is battery replacement, which is detailed in §3. Our deployment experience suggests the following:

■ **Guideline 1.** For store-level localization, a desirable BLE beacon density is every 5 to 15 meters. Compared to corridors, more dense beacons are needed in atrium areas where BLE fingerprints are more likely to cause confusion.

**BLE/GMF Training Data Collection.** MLoc adopts an outsourcing approach (i.e., hiring paid human workers) for collecting BLE/GMF fingerprints and the ground truth location data. This is a labor-intensive task: it takes on average 20 seconds for the collector to survey a location. A typical mall in a large city has an area of 50K to 100K  $m^2$ , with MLoc's coverage (all shared areas in the mall such as entrances, corridors, atrium, stairs, rest areas) accounting for  $\sim 10\%$  of the total area. It thus takes 28 to 56 hours for a human collector to survey, say, all  $1m \times 1m$  grids in a mall.

To reduce the data collection overhead, we only ask collectors to visit a small subset of locations called *landmarks*. Given the large scale of data collection, it is infeasible to ask collectors to identify




Figure 4: Trajectory correction.

the landmarks. Instead, as shown in Figure 3, a collector performs the job using a custom mobile app developed by us, which uses the floor plan to automatically generate the vertices of the stores' bounding boxes as landmarks. The collector will need to visit each landmark and collect the corresponding fingerprint. Each mall only has 500 to 1,000 landmarks, thus reducing the collector's workload by 10 $\times$  compared to the "exhaustive traversal" approach.

We notice two major issues in the pilot deployment. First, many collectors do not visit the landmarks in an efficient manner; they oftentimes miss landmarks and need to go back, thus increasing the travel distance. Second, many collectors suggest that there are more visually recognizable landmarks such as pillars and doors than those automatically generated based on the floor plan, which may not contain complete and detailed environmental information. We thus improve the data collection design accordingly as follows. First, the data collection app generates not only the landmarks, but also a suggested path for the collector. We use an N-shape path for each corridor (Figure 3), and a circular path for each atrium. Second, the app also allows collectors to add or modify the landmarks through a GUI. In this way, more visually recognizable landmarks could be identified in a "crowd-sourced" manner and incorporated into future collections. After these improvements, the vast majority of the collectors reported much better data collection experiences.

**Data Processing.** The data collected by human workers is processed in three steps.

(i) *Validation.* We use several heuristics (e.g., footstep counting, dead reckoning, see §2.3) to validate the worker's walking distance and direction, to ensure that the worker has indeed visited all the landmarks. A cross-collector fingerprint check is also performed to identify outliers.

(ii) *Dead-reckoning Trajectory Correction.* When the collector is walking from Landmark A to B, the collection app uses dead reckoning (§2.3) to track the collector's trajectory. Meanwhile, fingerprints are also collected as the collector is walking, e.g., at location  $p_1$ ,  $p_2$ , and  $p_3$  illustrated in Figure 4. These fingerprints will be used in the interpolation step to be described next. Due to the dead reckoning errors, the tracked trajectory  $A \rightarrow B'$  may differ from the actual trajectory  $A \rightarrow B$ . We thus calculate a transformation  $T$ , which consists of a vector rotation followed by linear scaling, that transforms vector  $\vec{AB'}$  to  $\vec{AB}$ . We then apply  $T$  to all the on-trajectory locations. As shown in Figure 4,  $\{p_i\}$  is transformed to  $\{q_i\}$  that is closer to the actual trajectory.

(iii) *Grid Interpolation.* We divide the entire mall's localization areas into equal-sized grids ( $1m \times 1m$  for BLE and  $0.1m \times 0.1m$  for GMF). Next, using the locations with (corrected) fingerprints as input vertices (e.g.,  $A$ ,  $q_1$ ,  $q_2$ ,  $q_3$ , and  $B$  in Figure 4), we apply

Delaunay Triangulation [27] to generate a mesh of triangles. For each grid, its fingerprint is then calculated by performing linear interpolation of the three vertices of the triangle that the center of the grid belongs to. This results in a grid map of fingerprints that will be used for online localization and tracking.

■ **Guideline 2.** *In complex malls, it is feasible to survey a small number of landmarks and use them to generate a fingerprint grid map. To reduce the data collection overhead, not only the landmarks but also their visiting paths should be pre-generated. It is yet beneficial to allow on-site collectors to improve landmark selections through user-friendly GUIs.*

## 2.2 Online Initial Positioning

We now describe how MLoc provides online localization and tracking services to end users. There are two stages involved: *initial positioning* (this subsection) that finds a user’s initial location, and *location tracking* (§2.3) that tracks and refines the location as the user is walking.

The initial positioning includes three steps: (1) Mall (building) Identification, which is done by examining the received BLE beacon IDs; (2) Floor Estimation, which is described in §2.4; and (3) Location Estimation, which we detail next.

The initial positioning stage only uses BLE fingerprints because GMF is much noisier (Figure 6). MLoc applies the  $k$ -nearest neighbors (kNN) to find the  $k$  grids  $g_1, \dots, g_k$  whose BLE fingerprints  $f_1, \dots, f_k$  are closest to (in L2 distance) the fingerprint  $f$  collected at the current location. We set  $k=20$  based on our controlled experiments in multiple malls. We exclude from fingerprints the BLE beacons whose signal strengths are weaker than a threshold (empirically set to -95dB), because weaker signals are less sensitive to distance changes and introduce more noises. In addition, we observe that the fingerprint readings across different smartphone brands/models oftentimes exhibit disparities (*i.e.*, Model A’s RSSI reading is always slightly higher than Model B). To overcome this issue, MLoc adopts a simple yet effective method: it normalizes the BLE fingerprints by subtracting from each RSSI reading the average RSSI across all the samples of all the beacons collected by the same device. Given that a device can sense a large number of beacons with diverse RSSI readings, their average RSSI provides a good per-device “baseline”. The normalization is applied to both the training and testing fingerprints. After the above normalization, MLoc estimates the user’s location (grid)  $g$  as a weighted sum of  $g_1, \dots, g_k$ , *i.e.*,  $g = \sum_{i=1}^k w_i g_i$  where  $w_i$  is inversely proportional to the L2 distance between (normalized)  $f_i$  and  $f$ , and  $\sum_{i=1}^k w_i = 1$ .

Through extensive field experiments, we identify two types of challenging areas. The first is illustrated in Figure 5(a) where the user is near the connection point between a narrow corridor and a wide corridor (or an open area). In this case, there are more nearby grids in the wide corridor than in the narrow corridor. The weighted sum of the  $k$ -nearest neighbors will thus be biased towards the wide corridor. The second case is shown in Figure 5(b) where the user is at the dead end of a corridor. In this case, the weighted sum of the nearest neighbors will be shifted to the open end. Note that these two types are generic and representative, observed in almost all the malls (35) we have studied. For identifying the dead end areas, we




Figure 5: Two types of challenging areas identified in malls.

manually mark them based on the mall’s (known) floor map. The incurred overhead is small because this is a one-time effort.

Both cases in Figure 5 are attributed to the non-uniform distribution of the grids imposed by the mall’s layout. We thus augment the localization algorithm using the floor layout information. To address the case in Figure 5(a), we assign an additional weight  $w'_i$  to each grid  $g_i$  to account for the nearby floor layout. We set  $w'_i$  to be inversely proportional to the width of the corridor or the diameter of the open space, which is computed offline based on the floor plan. The localization result is thus  $g = \sum_{i=1}^k w_i w'_i g_i$ . To handle the “dead end” scenario in Figure 5(b), we use a small  $k$  (the number of nearest neighbors) to minimize the impact of floor layout when 60% of the neighbors are in predefined dead end areas. Since the number of dead ends in each mall is limited, we manually set  $k$  (*e.g.*,  $k=10$ ) for each affected area.

■ **Guideline 3.** *Simple algorithms (kNN) can give an initial location estimation with reasonable accuracy. Preprocessing is important: fingerprints should be normalized to account for the device heterogeneity; floor layout information can also be leveraged to improve the accuracy. Most of the above computation can be done offline to minimize the runtime overhead.*

## 2.3 Online Location Tracking

After the initial positioning, MLoc will keep tracking the user’s movement and updating his/her location in real time.

**Dead Reckoning (DR)** is a widely used navigation technique that calculates an object’s position using its base position and measurements of the object’s speed and heading [23]. In MLoc, the base position is provided by the initial positioning stage (§2.2), and DR is performed on a per-footstep basis. Our DR solution is assembled from four building blocks developed in the literature. (1) We use the technique in [37] for *footstep detection*. (2) We use the algorithm in [53] to calculate the *footstep length* through the empirical formula  $\gamma \sqrt{acc_{max} - acc_{min}}$  where  $acc_{min}$  and  $acc_{max}$  are the minimum and maximum acceleration in Z axis over a detection window, respectively. Initially  $\gamma$  is empirically set to 0.4. (3) We use the technique in [38] to *correct the footstep length estimation* ( $\gamma$ ) based on the floor plan (so we know the walking distance) and the number of footsteps. (4) MLoc further leverages prior work [35, 39] for *heading estimation* using the accelerometer, gyroscope, magnetometer, and compass.

**Particle-filtering based Tracking.** MLoc employs Particle Filtering (PF [16]) as the basis for location tracking. PF uses a set of particles to model the posterior distribution of a stochastic process (localizing the user in our case), given the noisy observations (initial

positioning and error-prone sensor reading in our case). We next describe how MLoc instantiates the generic PF framework.

- Step ❶: MLoc performs the initial positioning. It initializes the observation  $o$ , the currently tracked location, to the initial positioning result. MLoc then generates  $n$  particles  $p_1, \dots, p_n$  following a 2D Gaussian distribution centering at  $o$ .  $n$  is set to 2K to 5K depending on the floor size. Next, go to ❸.
- Step ❷: MLoc performs weighted sampling (with replacement) of  $\eta \cdot n$  particles from the existing pool of particles, and generates  $(1 - \eta)n$  new particles forming a 2D Gaussian at  $o$ .  $\eta$  is set to 95%. This is called the *resampling* step in PF.
- Step ❸: Now MLoc starts tracking the user footstep by footstep. In each footstep, DR produces a vector  $\vec{v}$  indicating the user's footstep direction and distance. MLoc updates the observation  $o \leftarrow o + \vec{v}$ , as well as each particle  $p_i \leftarrow p_i + \vec{v}$ .
- Step ❹: For  $o$  and each  $p_i$ , MLoc maintains a trajectory of the most recent  $m$  footsteps' fingerprints (BLE and GMF features), denoted as  $tj(o) = \{f_1^o, \dots, f_m^o\}$  and  $tj(p_i) = \{f_1^{p_i}, \dots, f_m^{p_i}\}$  respectively.  $tj(o)$  is collected on the user's device, and  $tj(p_i)$  is retrieved from the grid map (§2.1) stored on the edge using  $p_i$ 's trajectory. MLoc applies Dynamic Time Warping (DTW) [36] to calculate the similarity between  $tj(o)$  and  $tj(p_i)$ , and uses the similarity value as the weight of  $p_i$ .
- Step ❺: The particles' weights are normalized to make their sum be 1. The observation  $o$  (i.e., the estimated location) is then updated as the (weighted) centroid of all the particles.
- Step ❻: Go to ❷ to process the next footstep. Ideally, as more and more footsteps are observed, the centroid of the particles will converge to the user's true location. However, due to errors of initial positioning and/or DR, the particles may diverge instead of converging. If the mean weight of the top 10% particles (ranked by their weights) is less than 0.25, we reset the entire PF algorithm by returning to ❶.

In practice, we find that there is no need to invoke ❷, ❹, and ❺ for every footstep, given that a typical footstep length is only 50 to 70 cm. Instead, performing the above steps every 4 footsteps yields almost no localization accuracy loss. In addition, the trajectory window length  $m$  needs to be judiciously chosen. A large  $m$  possibly benefits the tracking accuracy by giving more fingerprints for DTW matching, but at the cost of accumulating more DR errors. We design a mechanism that gradually probes  $m$ . Upon initial positioning,  $m$  is (re)set to  $m_{min}$ . We then increase  $m$  in every invocation of ❹ until it reaches  $m_{max}$ .  $m_{min}$  and  $m_{max}$  are empirically set to 4 and 12, respectively.

## 2.4 Online Floor Detection

Floor detection is triggered during the initial positioning stage, as well as when an irregular change of the accelerometer's Z-axis reading is detected [55].<sup>2</sup> MLoc adopts a simple floor detection algorithm by default: performing a majority vote of the floors associated with the 5-strongest BLE beacon signals captured over a 5-second window. We find that this simple method works well in most places

<sup>2</sup>MLoc does not use barometer for floor change detection, because barometer is only available on limited smartphones models.




Figure 6: Localization and tracking accuracy.

in a mall with two exceptions. First, most users take elevators or escalators instead of climbing stairs. We find that detecting a user entering/leaving an elevator or an escalator is often difficult, leaving floor detection untriggered. To address this issue, instead of developing sophisticated activity detection algorithms, we choose a simple design: when the user is approximated to be near an elevator or an escalator (conservatively determined based on the strongest BLE beacons), MLoc invokes floor detection periodically (every 2 seconds). Periodical polling incurs additional overheads than the event-triggered mode. However, given that the floor detection algorithm is very lightweight, we believe our selective polling design strikes a desired balance between accuracy and system overhead.

The second issue relates to the atrium area where a user may see strong BLE beacons from multiple floors due to the atrium's tall open space. This leads to degraded floor detection accuracy. To address this issue, when the user is found to be around the (pre-designated) atrium area, MLoc uses a different floor detection method based on deep neural networks (DNN). For each mall, we train a DNN model, which uses perceived BLE RSSIs to predict a one-hot vector of floors. The simple DNN model consists of 6 autoencoder-based layers for feature extraction [50] and 6 fully connected layers for classification. Compared to the basic detection algorithm, DNN can significantly increase the detection accuracy despite its higher overhead (Figure 11 in §3).

■ **Guideline 4.** *Despite the complex floor layouts, there are only several generic types of challenging areas (e.g., atrium, corridor dead ends, corridor connectors, and elevators) based on our extensive field studies. They can be tackled by classical algorithms (e.g., kNN and PF) enhanced by simple yet strategic customizations (e.g., weight adjustment and lightweight AI).*

## 2.5 Operational Model

MLoc employs the edge computing paradigm. An edge server is deployed in each mall. When the user launches the app, MLoc's centralized gateway server identifies which mall the user is at; then the remaining localization and tracking tasks are handed over to the local edge server. All the computation tasks are performed on the edge. The encrypted client-edge communication is over the Internet (in-mall WiFi or cellular).

MLoc's thin-client approach facilitates cross-platform development, but brings a side effect. Poor Internet connectivity in some malls can delay the client-edge communication, thus increasing localization errors, as complained by some users and verified by us. We will add the "offline mode" to MLoc.




Figure 7: Temporal stability.

### 3 LARGE-SCALE EVALUATION

This section presents large-scale evaluations via *trained testers*, while measurements on real users are described in §4.

**Evaluation Methodology.** MLoc has been deployed since 9/2018 with improvements being made over the past three years. We have conducted principled, large-scale evaluations by hiring tens of trained testers. The evaluation methodology is as follows. The testers adopt a specially designed testing app (a modified version of MLoc) and are asked to walk along multiple store-to-store paths to mimic normal user behaviors. Each path consists of a series of landmarks (§2.1). When the tester passes a landmark, he/she will tap a button so the testing app learns the ground truth. Meanwhile, the testing app runs MLoc’s localization algorithms as described in §2.

The evaluations (*i.e.*, the testing data collection) were conducted from 11/2019 to 01/2021 (with some additional benchmarks conducted in 2018), at 35 malls in 7 cities in China: Hangzhou, Shanghai, Wuxi, Wuhan, Guangzhou, Tianjin, and Shenyang (up to 2200 km apart). Each building has an average of 5.0 floors, and each floor has an average area of 49K  $m^2$ . Recall from §2.1 that the beacons are only deployed in shared areas of malls (*e.g.*, corridors, atrium, and rest areas), which account for around 10% of a floor’s area. The median path length is around 45 meters. The evaluations cover 152K  $m^2$  localization area with a total number of 4.3K paths and a total walking distance of 215 km. The training data collection consists of 22.3K paths (1,100 km of walking) involving 21K beacons and 244K landmarks. The ratio between the training and testing data size is roughly 5:1. The mobile devices used in our evaluations include Huawei, Xiaomi, Oppo, and iPhone. Unless otherwise mentioned, the evaluations (*i.e.*, testing) were conducted shortly (1 week to 1 month) after the training data was collected, on a per-mall basis. By default, the beacons have a broadcast interval of 200ms.

**Positioning Accuracy.** Figure 6 plots the overall positioning accuracy across all the tested landmarks. The four curves correspond to the errors in the {initial positioning, tracking} stage using {BLE only, BLE+GMF}. As expected, initial positioning gives a low accuracy (median error 4.1m for BLE and 5.0m for BLE+GMF), which is significantly improved in the tracking stage (median error 2.4m for BLE+GMF and 3.5m for BLE).<sup>3</sup> Meanwhile, we observe that for 3% of the landmarks, the location tracking error is higher than 10m. This is caused by various factors such as failed BLE beacons, fingerprint noises, erroneous floor detection, and low smartphone scanning frequency. The standard deviation of the median tracking

<sup>3</sup>Unless otherwise mentioned, MLoc uses BLE for initial positioning and BLE+GMF for location tracking, according to the design in §2.

Table 1: Beacon failure rate.

Time	% Failure
12/19	0.3%
06/20	1.5%
01/21	4.9%

error across the 35 malls is 0.3m (min: 1.9m, max: 3.0m). Smaller malls tend to have lower errors.

**Temporal Stability of Fingerprints.** To assess the temporal stability of fingerprints, we conduct a separate long-term experiment in three malls in Hangzhou. The training data was collected in 12/2019. Since then, no maintenance such as training data update and beacon replacement was performed. Then we launch two test campaigns involving the same paths, one in 01/2020, and the other in 01/2021. As shown in Figure 7, after one year, the median localization error increases from 2.4m to 4.1m. This is attributed to two factors: (1) the change of the physical environment such as store renovations and various events held in atrium areas, and (2) the failure of BLE beacons, as to be quantified next.

Table 1 shows the ratio of failed BLE beacons (*i.e.*, the fraction of beacons that stop broadcasting), based on our on-site surveys conducted in the 12 malls in 7 cities. As shown, over one year, the failure rate was non-trivial. We find that the beacon failures were not only due to hardware issues, but also because many beacons fell off the ceilings. It is thus important to ensure the secure attachment of the beacons.

■ **Finding 1.** MLoc achieves a median tracking error of 2.4m (10-th and 90-th percentile: 0.8m and 7.3m), more than adequate for store-level navigation. After one-year usage without replacing failed beacons (5%) or updating the training data, the error remains at an acceptable level (median 4.6m). We do observe a long tail of errors caused by various factors, many of which are non-algorithmic, such as beacon failures, poor client capability (Figure 14), and poor Internet connectivity (§2.5).

**BLE vs. GMF Features.** Figure 6 shows that compared to BLE only, further using GMF can improve the tracking accuracy. However, in the initial positioning stage, BLE+GMF underperforms BLE only. To better understand the interplay between the BLE and GMF features, we select 100 landmarks located on a single floor in a mall in Hangzhou. The landmarks are numbered from 1 to 100 such that two numerically close IDs imply that the two landmarks are physically close. Figure 8 plots the heatmaps of normalized fingerprint differences (in Manhattan distance) across all landmark pairs, for BLE, GMF, and BLE+GMF features. As shown, BLE features can easily distinguish two far-apart landmarks, but they are not good at differentiating nearby locations. In contrast, GMF features are not sensitive to distance: both nearby and far-apart landmarks may have different GMF fingerprints. Based on our measurement, the resolution<sup>4</sup> of GMF features is 5.4× better than BLE. Also, GMF fingerprints can be frequently sampled (every 20ms or shorter) on smartphones, whereas scanning BLE beacons takes a much longer

<sup>4</sup>The resolution of a grid  $g$  is calculated as the minimum physical distance to another grid  $g'$  such that  $\|f_g - f_{g'}\| < 10\% \cdot \|f_g\|$  where  $f_g$  and  $f_{g'}$  are the fingerprints. We calculate all grids’ resolutions then take their average.




Figure 8: Heatmaps of pairwise fingerprint differences for (a) BLE, (b) GMF, (c) BLE+GMF.




Figure 9: BBI vs. localization error.




Figure 10: Beacon spacing vs. localization error.

time (Figure 15). Despite these advantages, GMF fingerprints are noisier (Figure 8(b)) and more temporally unstable than BLE. We find that, for example, symmetric areas in a mall such as multiple wings may share similar GMF fingerprints.

■ **Finding 2.** *BLE and GMF are complementary. BLE is accurate, having low resolutions, and slow to scan; GMF is noisy, having high resolutions, and fast to collect. MLoc thus only uses BLE in the initial positioning to combat noises, and further uses GMF during the tracking stage to refine the location as more fingerprint samples accumulate.*

**BLE Beacon Broadcast Interval (BBI)** is a key parameter in MLoc. A small BBI helps improve the localization accuracy by allowing more frequent device-side fingerprint collection, but at the cost of a shorter beacon battery life. In our initial deployment in late 2018, we use a BBI of 500ms. In 2019, we reduce the BBI to 200ms. To justify this change, Figure 9 compares the localization accuracy of the two configurations, under best-effort fair comparisons (similar training-testing intervals and testing paths). As shown, reducing the BBI slightly improves the median localization error from 3.0m to 2.6m, but the tail improvement is much more significant: the 90-percentile error decreases from 10.6m to 7.1m. We find a major reason to be client-specific: smartphones with long BLE scanning intervals work poorly with long BBIs, as to be shown in Figure 16. Regarding the other side of the tradeoff, Table 2 shows a beacon’s expected battery life at different BBIs, based on our in-lab power modeling. At a 200ms BBI, a beacon can last for 1.75 years, which is short but still acceptable. Considering the results in Figure 9 and Table 2, we changed the BBI to 200ms in 2019.

■ **Finding 3.** *Increasing the BBI brings small improvements in most cases, but can significantly improve the accuracy for the long tail (~20% cases), benefiting devices with long BLE scanning intervals (Figure 16). For short BBIs (e.g., <300ms), the battery replacement overhead is non-trivial.*

Table 2: BBI vs. beacon battery life (2400 mAh).

BBI (ms)	Bat. Life (yr)
100	0.92
200	1.75
500	4.50
1000	7.75

**BLE Beacon Spacing** (i.e., the distance from a beacon to its nearest beacon) also affects the localization accuracy. To illustrate this, we study two malls, one in Tianjin (denoted as S1) and the other in Hangzhou (denoted as S2). Both malls have 5 floors, similar floor layouts (e.g., large atrium), similar floor areas (33.5K  $m^2$  for S1 and 30.8K  $m^2$  for S2), and the same BBI (200ms). The above factors foster fair comparisons (in a best-effort manner) between the two sites. MLoc has been deployed in both malls since 2019.

The key difference between S1 and S2 is their average BLE beacon spacing: around 14m for S1 and around 18m for S2. Figure 10 plots their distributions of localization errors. The results indicate that there is no noticeable difference in terms of the median localization error (~0.5m) between S1 and S2. However, similar to Figure 9, a smaller beacon spacing helps improve the tail accuracy: at 14m average beacon spacing (S1), the 90-percentile localization error is 6.7m, compared to 10.5m as for 18m average beacon spacing (S2).

**Floor Detection.** Recall from §2.4 that MLoc employs two floor detection algorithms: a majority vote of the top-5 strongest BLE beacons, and the DNN-based approach. Figure 11 compares their accuracy for atrium and non-atrium areas (the error bars are across different malls). We only include the results of the initial positioning stage, which is more challenging compared to the tracking stage due to a lack of observed BLE beacons. As shown, in non-atrium areas, both top-5 and DNN algorithms work well, with a median




Figure 11: Floor detection accuracy.




Figure 12: Walking distance estimation error.

accuracy of higher than 97%. However, in atrium areas, the accuracy of the top-5 method drops drastically to 61%, while the DNN-based approach remains robust, with a median accuracy of 96%. This justifies our design decision of using DNN-based floor detection in atrium areas (§2.4).

Despite the seemingly good floor detection results, we are a bit surprised to see that among the negative feedback from users, there are much more complaints about incorrect floor detection than inaccurate localization on the same floor. Users can tolerate large same-floor localization errors (e.g., up to 20m), but they can hardly accept *any* floor detection error, which will lead to a totally unexpected floor map.

■ **Finding 4.** *A lightweight DNN-based floor detection method (a 12-layer network adopted by MLoc) can yield good accuracy in the challenging atrium areas. However, customers have an extremely low tolerance to floor detection errors, compared to their large tolerance to same-floor localization errors. Future systems should aim at 100% accuracy for floor detection.*

**Dead Reckoning (DR).** Recall that each path walked by a tester consists of multiple landmarks with their location ground truth known. For each *segment* (i.e., a pair of consecutive landmarks)  $A \rightarrow B$ , we compare  $|AB|$  with the walking distance calculated by DR, and compare the direction of  $\vec{AB}$  with the average heading calculated by DR. In other words, our DR evaluation is performed on a per-segment basis, because obtaining the per-footstep ground truth is difficult.

Figure 12 plots the estimation error of walking distance across all paths' segments. We plot the first segment of a path and the remaining segments separately. The latter bears a higher accuracy than the former, because as the user walks, DR improves the step length estimation based on the floor plan and the number of footsteps (§2.3). Overall, the 25th, 50th, and 75th-percentile of the estimation error are 6%, 12%, and 19%, respectively. Figure 13 plots the




Figure 13: Heading estimation error.




Figure 14: Localization error across smartphone brands.

heading estimation error distribution, with the 25th, 50th, and 75th-percentile being 3.9°, 9.1°, and 16.4°, respectively. The results are not as impressive as those reported by the research papers [38, 39] due to noises from multiple sources (imperfect sensors, different handheld positions of phones, busy malls). Nevertheless, the accuracy is adequate for store-to-store navigation, and many errors can be corrected by particle-filtering (§2.3). We therefore do not pursue more sophisticated DR approaches.

**Impact of Smartphone Brands.** Figure 14 plots the localization errors for different smartphones (measured in 2021, matching Figure 7). We observe considerable differences across the devices. iPhone, which accounts for ~25% of the total testing devices, owns the best accuracy of 3.1m (median value). Android devices belong to a more fragmented ecosystem and thus exhibit diverse performance. Popular smartphones such as Huawei and Xiaomi (account for ~50% of the devices) bear not only higher errors, but also larger variations. The cross-device homogeneity is attributed to the devices' different software (e.g., OS and driver) and hardware (BLE radio and sensor chips). Vendors such as Huawei are known to deeply customize Android for energy saving purpose [2, 3]. This may cause degraded localization accuracy.

We showcase a major root cause of the cross-device accuracy difference: the Successful Scanning Interval (SSI), defined as the interval between two consecutive BLE scans each capturing at least one beacon. As plotted in Figure 15, there are considerable statistical differences among the SSIs measured on different devices, and the SSI is largely correlated with localization errors reported in Figure 14. In extreme cases, a Huawei phone cannot see even a single beacon for more than 10 seconds. This inevitably leads to high localization errors. Figure 16 further shows that, as the beacon broadcast interval (BBI) increases, the SSI and its variation increase correspondingly, in particular for the Huawei phones with throttled scanning activities. This explains our findings in Figure 9 where




Figure 15: SSI (Successful Scanning Interval) across brands.




Figure 16: BBI vs. SSI across smartphone brands.

reducing the BBI helps improve the tail localization errors, which indeed mostly occur on high-SSI phones as verified by us.

■ **Finding 5.** *Different smartphone brands/models exhibit considerable differences in localization accuracy. Some vendors throttle the BLE scanning frequency to save energy. This may severely impact the performance of localization/tracking applications. Smartphone OS vendors should consider making the scanning frequency (among a wide range) configurable.*

**The Energy Consumption** of MLoc is measured by controlled experiments. We examine the battery drain while running MLoc. We use several fully charged iPhones and Android phones to perform continuous store-to-store navigation for 1 hour, with WiFi, BLE, screen turned on and the display brightness level set to 50%. We use the Medium Power mode in Table 3 (described next) for the experiment. In 1 hour, the battery life drops by 7% to 11%, and the traffic usage is less than 12 MB. In contrast, keeping the phones idle for one hour with the same screen brightness consumes 4% to 6% of the battery energy. The results indicate low resource footprints of MLoc.

In addition to the above coarse-grained measurement, we also perform a closer examination of MLoc’s power consumption by logging the battery usage events [1] of an OPPO R15 smartphone under three configurations of MLoc: low power, medium power (the default configuration), and high power. As shown in Table 3, IMU and GMF sensors are enabled in all three modes as they incur very low power consumption. BLE sensing is disabled in the low power mode while enabled in the medium and high power mode. The localization frequency increases from 1Hz to 10Hz in the high power mode. The comparison baseline corresponds to the scenario where the screen (at 50% display brightness level), Bluetooth, WiFi, and GPS are turned on, but MLoc is not running. The average baseline power consumption on OPPO R15 is 1032mW. As shown in the rightmost column in Table 3, compared to the baseline, the average power consumption increases by 5%, 12%, and 16%, in the




Figure 17: Distribution of floor changes (0 = same floor).

Table 3: Power usage of MLoc on OPPO R15 smartphone.

Power Mode	IMU+GMF Sensing	BLE Sensing	Localization Frequency	Avg. Device Power
Low Power	50 Hz	Off	1 Hz	+5%
Med. Power	50 Hz	1 Hz	1 Hz	+12%
High Power	50 Hz	1 Hz	10 Hz	+16%

three power modes respectively. The above results confirm that MLoc’s incurred energy overhead is acceptable, at least for its low and medium power configurations.

## 4 USER BEHAVIORS

This section analyzes the user behaviors of MLoc’s customers in the real world. Note that this is different from the evaluations using data collected by trained testers in §3. We collected MLoc’s customer data at 12 malls in 7 cities (Hangzhou, Shanghai, Wuxi, Wuhan, Guangzhou, Tianjin, and Shenyang) from 04/2020 to 04/2021. The data includes 472K unique users (48% are female). The user devices are dominated by iPhone (46%), Huawei (31%), Xiaomi (19%), and Oppo (3%), but we do see a long tail of 8 other smartphone brands.

MLoc is offered to the end customers as both a standalone localization app and a mini-program of WeChat (a multi-purpose instant messaging, video/audio chat, social media, and payment app). It is also released as a library that is embedded into shopping malls’ mobile apps. We find that more than 95% of the customers use the mini-program or shopping malls’ mobile apps with MLoc embedded. This is in part because of the advertisements (both on-site and online) conducted by the malls to promote the two types of apps. The advertisements are usually effective. For example, in our one-month study of using MLoc as a marketing platform (described shortly), online advertisement attracted 46% of new users (as tracked by the scanned QR codes).

While we are unable to evaluate the localization accuracy due to a lack of ground truth from the customers’ data, MLoc records customers’ last tracked locations in their navigations. Using this information, we find that 95% of the navigations ended at locations that are 20m within the destination storefront. This indicates the overall effectiveness of MLoc.

The MLoc app also logs basic information regarding the app usage. Figure 17 plots the distribution of floor changes from the beginning to the end of the navigation (or user closing our app). We find that 27% of the navigations involve floor changes. 31% of the floor changes are going downstairs – higher than what we




Figure 18: The usage of MLoc.

expected. Figure 18 plots the distributions of duration and users' walking distance in each navigation, whose median values are 73 seconds and 81 footsteps (about 50m), respectively (25-th percentile: 44 seconds and 35 footsteps; 75-th percentile: 108 seconds and 189 footsteps). Note that our measured duration and footsteps are the *lower bounds* of the actual walking time/distance, because a user may close our app early during navigation. Figure 19 plots the distribution of the number of daily navigation sessions across all (user, day) when MLoc is used at least once. In 28% of the cases, a customer uses MLoc more than once within a day.

■ **Finding 6.** While MLoc can help many users find their destinations quickly, in at least 20% of the navigations, users spend more than 2 minutes using MLoc or walk for more than 230 steps (~150m). Customers use MLoc for same-floor, upstairs, and downstairs navigations.

We also investigate the destination stores users go to. The location types ordered by popularity include food and drink (53%), clothing (15%), in-mall playground (5%), electronics (4%), mall exit & subway entrance (3%), and others (20%). Through our informal chat with customers during field studies, we find that people need MLoc not only because today's malls are large, but also due to many other reasons: signs being too small to recognize in distance; highly symmetric floor layout confusing users (they may also confuse localization algorithms); and young people who get used to (outdoor) navigation apps being inexperienced in reading maps, *etc.*

**MLoc as a Marketing Platform.** Indoor localization systems are promising marketing platforms. To demonstrate this, from 11/2020 to 12/2020, we organized a sales event by working with a large mall in Wuxi (in eastern China). In this event, customers can retrieve from the MLoc app e-coupons that offer discounts in tens of the stores in the mall. The e-coupons are automatically pushed to the user when he/she is physically near the store (*i.e.*, when the store shows up on the in-app navigation UI). Overall, we observe high engagement and conversion rate, as reported below.

During the one-month event, MLoc recorded about 11K navigation sessions from 7K customers. Among them, 73% of the navigation sessions from the 3K participating customers (PCs) involve at least one coupon retrieval. 56% of the PCs are in their 30s or older. We also track the origin of the PCs. 43% of the PCs were existing MLoc users; 26% learned this promotion event through online ads; 19% were attracted by on-site ads so they scanned the QR code and installed MLoc; 10% found out about the event from social media; and 2% learned it from other sources.




Figure 19: Distribution of the number of daily navigation sessions per user.

It is encouraging to see that within the PCs (who have retrieved at least one coupon), 22% of them actually used the coupon(s) in stores. On average, each PC retrieved 2.3 coupons and used 1.1 coupons. Also note that the coupons have a limited supply. 87% of the prepared coupons were retrieved by the PCs. Among the retrieved coupons, 41% of them were used in stores. The above results indicate the overall success of this sales event.

■ **Finding 7.** Localization systems and their served malls can mutually benefit each other: MLoc gains many new users through this sales event (57% of the PCs were new); meanwhile, the event achieved a conversion rate of 22%, significantly higher than typical conversion rates for pay-per-click (PPC) online advertising (2% to 3% [5]).

## 5 RELATED WORK AND CONCLUSION

**Commercial Indoor Localization Systems.** There are several news and blog articles [8, 31, 32, 44, 45, 47] on commercial indoor localization systems deployed in museums, airports, railway stations, *etc.*, using either BLE or WiFi as fingerprints. None of them offered detailed insights or experiences as we did. Their deployment scales were also smaller.

**Prior Deployment Experiences.** A recent study [15] reported a large-scale deployment of a BLE beacon system for goods delivery. It detects arrivals and departures of couriers at merchants, thus operating at a much coarser granularity than MLoc from the localization perspective. Some other works also conducted real-world deployment of localization systems [19, 33, 61], albeit at much smaller scales. A prior study [34] describes the results and lessons learned from the 2014 Microsoft Indoor Localization Competition; the authors compared indoor localization solutions from 22 teams around the world. We take into account their experiences when developing MLoc. More recently, Microsoft sponsored and co-organized another indoor localization competition, and released a large dataset (with ground truth) used by the participating teams [4]. To our knowledge, there is no detailed report of the competition or any characterization study of the dataset. In our future work, we plan to evaluate MLoc on their datasets and use the results to improve our localization algorithms.

**Indoor Localization and Sensing.** Researchers have developed numerous techniques and models for indoor localization, such as fingerprinting [7, 13, 20, 22, 28, 41, 46, 54], RSSI propagation models [7, 10, 11, 21, 48], wireless models [6, 24, 35, 49, 57, 58] such as angle of arrival (AoA) [57] and Time of flight (ToF) [35, 49], dead reckoning [40], crowdsourcing [56], sensor fusion [18], and

image or light based localization [17, 26, 29, 30]. Many other studies develop sensing techniques [9, 51, 52] for, e.g., identifying human behaviors [52]. MLoc leverages robust, time-tested algorithms from the literature, and customizes them for in-mall localization.

To conclude, our experiences suggest that high localization accuracy is only one of the multiple objectives of MLoc, which needs to carefully balance the tradeoffs among accuracy, human labor, infrastructure complexity, usability, and maintenance overhead, to name a few. Beacons' limited battery life may also pose an obstacle towards their large-scale deployment and maintenance in the wild. We hope our insights can boost future efforts on transforming the two-decade research on indoor localization into commercial products.

## ACKNOWLEDGMENTS

We thank the anonymous reviewers and our shepherd for their insightful comments. This paper was supported in part by NSF Award 1915122 and 2128489. Zhimeng Yin's research was supported by City University of Hong Kong 9610491 and NSF China 62102332.

## REFERENCES

- [1] [n.d.]. BatteryManager. <https://developer.android.com/reference/android/os/BatteryManager>.
- [2] [n.d.]. BLE devices issues with Huawei Phone. <https://stackoverflow.com/questions/51203419/bluetooth-low-energy-ble-devices-connection-issue-with-huawei-phone>.
- [3] [n.d.]. Don't kill my app (on Huawei devices). <https://stackoverflow.com/questions/51203419/bluetooth-low-energy-ble-devices-connection-issue-with-huawei-phone>.
- [4] [n.d.]. Microsoft Indoor Location Competition. <https://www.microsoft.com/en-us/research/publication/indoor-location-competition-2-0-dataset/>.
- [5] [n.d.]. What Is a Good Conversion Rate? <https://www.webfx.com/blog/marketing/what-is-a-good-conversion-rate/>.
- [6] Koshan Ayyalasoamayajula, Aditya Arun, Chenfeng Wu, Sanatan Sharma, Abhishek Rajkumar Sethi, Deepak Vasisht, and Dinesh Bharadia. 2020. Deep learning based wireless localization for indoor navigation. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–14.
- [7] Paramvir Bahl and Venkata N Padmanabhan. 2000. RADAR: An in-building RF-based user location and tracking system. In *Proceedings IEEE INFOCOM 2000. Conference on computer communications. Nineteenth annual joint conference of the IEEE computer and communications societies (Cat. No. 00CH37064)*, Vol. 2. IEEE, 775–784.
- [8] Airport benchmarking. 2016. Airlines and airports are beaconizing. <https://www.airportbenchmarking.com/airlines-and-airports-are-beaconizing/>.
- [9] Andreas Biri, Neal Jackson, Lothar Thiele, Pat Pannuto, and Prabal Dutta. 2020. SociTrack: infrastructure-free interaction tracking through mobile sensor networks. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–14.
- [10] Dongyao Chen, Kang G Shin, Yurong Jiang, and Kyu-Han Kim. 2017. Locating and tracking ble beacons with smartphones. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*. 263–275.
- [11] Yin Chen, Dimitrios Lymberopoulos, Jie Liu, and Bodhi Priyantha. 2012. FM-based indoor localization. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*. 169–182.
- [12] Krishna Chintalapudi, Anand Padmanabha Iyer, and Venkata N Padmanabhan. 2010. Indoor localization without the pain. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*. 173–184.
- [13] Jaewoo Chung, Matt Donahoe, Chris Schmandt, Ig-Jae Kim, Pedram Razavai, and Micaela Wiseman. 2011. Indoor location sensing using geo-magnetism. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*. 141–154.
- [14] Patrick Dickinson, Gregorz Cielniak, Olivier Szymanezyk, and Mike Mannion. 2016. Indoor positioning of shoppers using a network of Bluetooth Low Energy beacons. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 1–8.
- [15] Yi Ding, Ling Liu, Yu Yang, Yunhuai Liu, Desheng Zhang, and Tian He. 2021. From Conception to Retirement: a Lifetime Story of a 3-Year-Old Wireless Beacon System in the Wild. In *18th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 21)*. 859–872.
- [16] Petar M Djuric, Jayesh H Kotecha, Jianqui Zhang, Yufei Huang, Tadesse Ghirmai, Mónica F Bugallo, and Joaquin Miguez. 2003. Particle filtering. *IEEE signal processing magazine* 20, 5 (2003), 19–38.
- [17] Jiang Dong, Yu Xiao, Marius Noreikis, Zhonghong Ou, and Antti Ylä-Jääski. 2015. iMoon: Using smartphones for image-based indoor navigation. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*. 85–97.
- [18] Wan Du, Panrong Tong, and Mo Li. 2020. UniLoc: A unified mobile localization framework exploiting scheme diversity. *IEEE Transactions on Mobile Computing* (2020).
- [19] Rizanne Elbaky, Moustafa Elhamshary, and Moustafa Youssef. 2018. HyRise: A robust and ubiquitous multi-sensor fusion-based floor localization system. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 1–23.
- [20] Ramsey Faragher and Robert Harle. 2015. Location fingerprinting with bluetooth low energy beacons. *IEEE journal on Selected Areas in Communications* 33, 11 (2015), 2418–2428.
- [21] Chen Feng, Wain Sy Anthea Au, Shahrokh Valaee, and Zhenhui Tan. 2011. Received-signal-strength-based indoor positioning using compressive sensing. *IEEE Transactions on mobile computing* 11, 12 (2011), 1983–1993.
- [22] AKM Mahtab Hossain and Wee-Seng Soh. 2010. Cramer-Rao bound analysis of localization using signal strength difference as location fingerprint. In *2010 Proceedings IEEE INFOCOM*. IEEE, 1–9.
- [23] Wonho Kang and Youngnam Han. 2014. SmartPDR: Smartphone-based pedestrian dead reckoning for indoor localization. *IEEE Sensors journal* 15, 5 (2014), 2906–2916.
- [24] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. Spotfi: Decimeter level localization using wifi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*. 269–282.
- [25] Swaran Kumar, Stephanie Gil, Dina Katabi, and Daniela Rus. 2014. Accurate indoor localization with zero start-up cost. In *Proceedings of the 20th annual international conference on Mobile computing and networking*. 483–494.
- [26] Ye-Sheng Kuo, Pat Pannuto, Ko-Jen Hsiao, and Prabal Dutta. 2014. Luxapose: Indoor positioning with mobile phones and visible light. In *Proceedings of the 20th annual international conference on Mobile computing and networking*. 447–458.
- [27] Der-Tsai Lee and Bruce J Schachter. 1980. Two algorithms for constructing a Delaunay triangulation. *International Journal of Computer & Information Sciences* 9, 3 (1980), 219–242.
- [28] Binghao Li, Thomas Gallagher, Andrew G Dempster, and Chris Rizos. 2012. How feasible is the use of magnetic field alone for indoor positioning?. In *2012 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 1–9.
- [29] Lingkun Li, Pengjin Xie, and Jiliang Wang. 2018. Rainbowlight: Towards low cost ambient light positioning with mobile phones. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. 445–457.
- [30] Song Liu and Tian He. 2017. Smartlight: Light-weight 3d indoor localization using a single led lamp. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. 1–14.
- [31] Locatify. 2020. BLE Beacon Museum Guide with real-time user location. <https://locatify.com/blog/case-studies/eldheimar-museum/>.
- [32] Locix. 2019. Locix Releases Ultra-Precise WiFi-Based Local Positioning System. <https://www.businesswire.com/news/home/20191015005251/en/Locix-Relases-Ultra-Precise-WiFi-Based-Local-Positioning-System>.
- [33] Chengwen Luo, Hande Hong, Mun Choon Chan, Jianqiang Li, Xinglin Zhang, and Zhong Ming. 2017. MPiLoc: Self-calibrating multi-floor indoor localization exploiting participatory sensing. *IEEE Transactions on Mobile Computing* 17, 1 (2017), 141–154.
- [34] Dimitrios Lymberopoulos, Jie Liu, Xue Yang, Romit Roy Choudhury, Vlado Handziski, and Souvik Sen. 2015. A realistic evaluation and comparison of indoor location technologies: Experiences and lessons learned. In *Proceedings of the 14th international conference on information processing in sensor networks*. 178–189.
- [35] Alex T Mariakakis, Souvik Sen, Jeongkeun Lee, and Kyu-Han Kim. 2014. Sail: Single access point-based indoor localization. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*. 315–328.
- [36] Meinard Müller. 2007. Dynamic time warping. *Information retrieval for music and motion* (2007), 69–84.
- [37] Jan Racko, Peter Brida, Arto Perttula, Jussi Parviainen, and Jussi Collin. 2016. Pedestrian dead reckoning with particle filter for handheld smartphone. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 1–7.
- [38] Anshul Rai, Krishna Kant Chintalapudi, Venkata N. Padmanabhan, and Rijurekha Sen. 2012. Zee: zero-effort crowdsourcing for indoor localization. In *MobiCom*. ACM, 293–304.
- [39] Nirupam Roy, He Wang, and Romit Roy Choudhury. 2014. I am a smartphone and i can tell my user's walking direction. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*. 329–342.
- [40] Jochen Seitz, Jasper Jahn, Javier Gutiérrez Boronat, Thorsten Vaupel, Steffen Meyer, and Jörn Thielecke. 2010. A hidden markov model for urban navigation

- based on fingerprinting and pedestrian dead reckoning. In *2010 13th International Conference on Information Fusion*. IEEE, 1–8.
- [41] Yuanchao Shu, Cheng Bo, Guobin Shen, Chunshui Zhao, Liqun Li, and Feng Zhao. 2015. Magical: Indoor localization using pervasive magnetic field and opportunistic WiFi sensing. *IEEE Journal on Selected Areas in Communications* 33, 7 (2015), 1443–1457.
- [42] Yuanchao Shu, Kang G Shin, Tian He, and Jiming Chen. 2015. Last-mile navigation using smartphones. In *Proceedings of the 21st annual international conference on mobile computing and networking*. 512–524.
- [43] Stephen P Tarzia, Peter A Dinda, Robert P Dick, and Gokhan Memik. 2011. Indoor localization without infrastructure using the acoustic background spectrum. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*. 155–168.
- [44] Fablian Technologies. 2018. Eddystone beacon installation at Indian Railway stations by Google. <https://www.fabliantechnologies.com/eddytone-beacon-installation-at-indian-railway-stations-by-google/>.
- [45] THINKPROXI. 2017. FAMOUS BEALE STREET IMPLEMENTED BEACON TECHNOLOGY. <https://thinkproxi.com/thinkproxi-announces-famous-beale-street-implemented-beacon-technology/>.
- [46] Xiaohua Tian, Ruofei Shen, Duowen Liu, Yutian Wen, and Xinbing Wang. 2016. Performance analysis of RSS fingerprinting based indoor localization. *IEEE Transactions on Mobile Computing* 16, 10 (2016), 2847–2861.
- [47] Future travel experience. 2017. Gatwick's beacon installation provides partners with blue dot navigation and augmented reality wayfinding. <https://www.futuretravelexperience.com/2017/05/gatwick-airports-beacon-installation-enables-blue-dot-navigation/>.
- [48] David Tse and Pramod Viswanath. 2005. *Fundamentals of wireless communication*. Cambridge university press.
- [49] Deepak Vasisht, Swarun Kumar, and Dina Katabi. 2016. Decimeter-level localization with a single WiFi access point. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*. 165–178.
- [50] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research* 11, 12 (2010).
- [51] Jue Wang, Deepak Vasisht, and Dina Katabi. 2014. RF-IDraw: Virtual touch screen in the air using RF signals. *ACM SIGCOMM Computer Communication Review* 44, 4 (2014), 235–246.
- [52] Yan Wang, Jian Liu, Yingying Chen, Marco Gruteser, Jie Yang, and Hongbo Liu. 2014. E-eyes: device-free location-oriented activity identification using fine-grained wifi signatures. In *Proceedings of the 20th annual international conference on Mobile computing and networking*. 617–628.
- [53] Harvey Weinberg. 2002. Using the ADXL202 in pedometer and personal navigation applications. *Analog Devices AN-602 application note 2*, 2 (2002), 1–6.
- [54] Martin Werner, Moritz Kessel, and Chadly Marouane. 2011. Indoor positioning using smartphone camera. In *2011 International Conference on Indoor Positioning and Indoor Navigation*. IEEE, 1–6.
- [55] Oliver Woodman and Robert Harle. 2008. Pedestrian localisation for indoor environments. In *Proceedings of the 10th international conference on Ubiquitous computing*. 114–123.
- [56] Chenshu Wu, Zheng Yang, Yunhao Liu, and Wei Xi. 2012. WILL: Wireless indoor localization without site survey. *IEEE Transactions on Parallel and Distributed Systems* 24, 4 (2012), 839–848.
- [57] Jie Xiong and Kyle Jamieson. 2013. Arraytrack: A fine-grained indoor location system. In *Presented as part of the 10th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 13)*. 71–84.
- [58] Jie Xiong, Karthikeyan Sundaresan, and Kyle Jamieson. 2015. Tonetrack: Leveraging frequency-agile radios for time-based indoor wireless localization. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. 537–549.
- [59] Zhice Yang, Zeyu Wang, Jiansong Zhang, Chenyu Huang, and Qian Zhang. 2015. Wearables can afford: Light-weight indoor positioning with visible light. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*. 317–330.
- [60] Zheng Yang, Chenshu Wu, and Yunhao Liu. 2012. Locating in fingerprint space: wireless indoor localization with little human intervention. In *Proceedings of the 18th annual international conference on Mobile computing and networking*. 269–280.
- [61] Zuwei Yin, Chenshu Wu, Zheng Yang, and Yunhao Liu. 2017. Peer-to-peer indoor navigation using smartphones. *IEEE Journal on Selected Areas in Communications* 35, 5 (2017), 1141–1153.
- [62] Sangki Yun, Yi-Chao Chen, Huihuang Zheng, Lili Qiu, and Wenguang Mao. 2017. Strata: Fine-grained acoustic-based device-free tracking. In *Proceedings of the 15th annual international conference on mobile systems, applications, and services*. 15–28.
- [63] Chi Zhang, Kalyan P Subbu, Jun Luo, and Jianxin Wu. 2014. GROPING: Geomagnetism and crowdsensing powered indoor navigation. *IEEE Transactions on Mobile Computing* 14, 2 (2014), 387–400.