






Microbenchmarking Documents in the Air

Dragoş Ştefan Niculescu^(✉) , Laura Ruse^{}, and Costin Carabaş^{}

University Politehnica of Bucharest, Bucharest, Romania
suotiuqibom.dragos@0sg.net

Abstract. *Documents in the Air* is a middleware system that allows placing and retrieving virtual objects or documents at different indoor locations **without requiring a positioning system**. It consists of an Android application and an intranet or cloud server, and only makes use of existing WiFi or BLE infrastructure to produce location specific signatures. We evaluate the performance of the system and study signature behavior with respect to: WiFi network characteristics, dissimilarity and real distances, and various collection methods. Using our own measurements, as well as publicly available data from several buildings, we show that the document retrieval process is accurate under conditions of signature impairment, signature aging, reduced AP density, and heterogeneous devices.

Keywords: WiFi fingerprints · Location centric services · Performance evaluation

1 Introduction

One of the main driving forces of the IoT is the desire to connect everything to the network, and the hope to access and control everyday life's processes. With cluttered environments in homes, and institutions, one of the challenges is the management of the Internet connected objects. Physical object databases come with significant challenges in management of devices, topologies, inter-operation, security, privacy, portability, and context awareness. The concept of context is central for pervasive computing and IoT, so that ongoing research still requires extensive surveying effort and building of taxonomies [6, 15, 18, 28, 32, 33]. Context is heavily overloaded concept, and while position is certainly a context, obtaining it, especially indoors requires nontrivial effort. Indoor location will no doubt play an important role in the quest for context awareness, and research in locating and tracking devices, which has ramped up significantly in the last decade [16, 20, 21, 23], has shown that obtaining indoor location is costly in several ways: necessary infrastructure (specialized measurement hardware), low accuracy (WiFi, BLE), effort (training and maintaining location databases), battery consumption (GPS, WiFi, 4G methods), erosion of privacy (Google tracking).

Work supported in part by the Romanian National Authority for Scientific Research and Innovation, UEFISCDI project PN-III-P2-2.1-PED-2019-5413.

Documents in the Air(*AirDocs*) is a recently proposed [26] middleware system that avoids the use of Cartesian location, relying instead on context specific signatures to allow placing and retrieving virtual documents at different indoor locations. It relies on WiFi/Bluetooth infrastructure existing in most homes and institutions, requiring a single additional server visible in the intranet, and an application that can be installed on any mobile device. Collecting this information in a **signature of the context**, which includes the WiFi fingerprint and other context specific information, enables retrieval of information based on signatures. *AirDocs* enables many applications that involve natural placing and retrieving of documents at locations, but without actually requiring a location system. In this article, we experiment with a dissimilarity measure as a proxy for Euclidean distance, that allows for several operations in signal space to enable placing and retrieving virtual objects using **only** context derived from wireless fingerprints. We benchmark the behavior of the signatures in signal space, and Cartesian locations are never used in the operation of the system, except for reporting purposes. The contributions of the papers are:

- implement a proof of concept for the middleware, which includes a simple Android application that places and retrieves signatures, and a server that implements searches in signature space.
- propose mapping between Euclidean distance and signature dissimilarity
- benchmark the main part of signatures, namely the WiFi fingerprints, with respect to creation, classification, collision behavior, aging
- show that object retrieval is robust with respect to: infrastructure density, time variation, methods of fingerprint collection, and device variability.

2 System Architecture

The *AirDocs* architecture is shown in Fig. 1: the middleware provides an API for scanning for Wi-Fi APs, cellular networks, Bluetooth Low Energy (BLE) devices, GPS information, and sound, in order to build signatures. Also, it includes methods for sending documents to the server along with the associated signature, and for retrieving documents from the server for a recently collected signature. This middleware can then be used by actual applications in order to store and retrieve documents depending on their specific.

The server is responsible with storing documents and their associated signatures, and also with identifying the appropriate document for a certain signature. It does this by comparing the collected signature with other signatures stored in the database, by using a (dis)similarity function. The most similar signature is identified and the associated document is retrieved and displayed in the application.

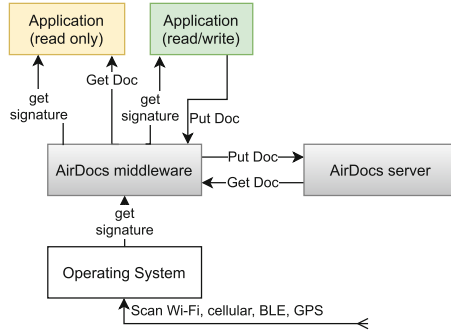


Fig. 1. AirDocs architecture: the middleware module on the mobile collects signatures from the environment. The server indexes document databases and answers queries based on signatures.

The unique rich signatures obtained can then be used to manage a document collection without mapping them to geographical locations, but in fact obtaining an association between a document and its location in the building, but not physical position. The data-structure obtained provides many functionalities of a location indexed database. The middleware will offer three main primitives to the applications:

- $S = CreateSgn()$ collects a location specific signature from the phone sensors (WiFi, BLE, 4G, sound, etc.) and creates a multidimensional signature that is unique from any other location signature in the building;
- $Put(S, document)$ stores a document on the server associated with the signature S ; The signature is created by a phone, but the indexing of the signatures and the document storage happen on a server in the intranet (or Internet).
- $Get(S)$ - a phone harvests its current signature, and asks the server for a list of documents that have “close” signatures, meaning documents that have been stored at nearby locations. The server **searches in the signature space, and real geographical coordinates are never needed.**

On the server, documents are indexed by their signatures, based on their similarities between them. For a given signature query, the server may: 1. retrieve the document with the lowest dissimilarity with the query, or 2. retrieve all documents with dissimilarities below a threshold. Since real physical positions are not known, the database of signatures on the server needs to be organized using clustering and labeling methods.

2.1 Usage Scenarios

The system is akin to augmented reality with the users having the illusion of the virtual objects being spread in the physical environment, visible only at certain locations. Leaving an object or a document “in the air” allows for a natural way to use it as a virtual wireless post-it for museums explanations, navigation in complex buildings like universities, airports and malls, advertising, lab door announcements, restaurant menus, office pin-boards, refrigerator post-its, general reminders, and notices around the house and office. Many of these applications would usually require location, and a building-wide indoor coordinate system, but if the *AirDocs* service is available, the functionality of placing and retrieving documents can be used right away requiring only the installation of the Android application.

Since the WiFi coverage is often associated with administrative control of the physical space, the placement of the document server may be in the intranet. In this case, the server can be accessible on a standard port once the mobile joins the building WiFi, as any other service that runs inside a home or institution. Applications using the *AirDocs* API would detect local servers using a name discovery (e.g. DNS-SD, zeroconf, or UPnP) and interact with them in an application specific fashion.

As we will show in Sect. 4.1, the resolution of the system is currently around 2 m¹. This allows for placing documents at higher density, but at query time they will be reported as being in the same place. As the performance of the system will improve by enriching signatures (Sect. 6.1, more applications would become possible. In a shopping mall for example, signatures are based on all visible and fixed WiFi infrastructure, and the documents are placed for announcement or for advertisement purposes. The system would be read-only for shoppers, and a document could contain a floor map, a sale, a coupon, or any other digital object that could be useful to the shoppers. A museum would have a similar setup in that documents placing and contents are curated by the institutions and the visitors would only read them when the appropriate locations. An university, on the other hand, may allow limited posting by the students in certain areas, or for a maximum document size, or limited lifetime imposed by weekly cleanups. Conferences could place maps, programs, and other pointers at hotel entrances and in appropriate presentation rooms.

Full-blown applications based on this API will have to consider actual institution specific virtual object types (pdf, gif, URLs, multimedia), document policies (their maximum size and life length), security (who can create documents), visibility (certain user-groups might see different sets of documents), scalability, and an appropriate GUI (simple browser refresh, or full augmented reality) to facilitate production and consumption of spatial data that are all application specific. If *AirDocs* is run in the intranet on a standardized port, a common smartphone app would cover many usage cases, and the user would not need to

¹ Limited by the Android WiFi throttling procedure taking about 3–4 s for a full scan of both 2.4 GHz and 5 GHz WiFi bands.

install a new one when attending a conference, visiting an exhibit, or browsing a brick and mortar store.

The *AirDocs* API can be used even without ownership of the physical space or of the WiFi infrastructure, with a server in the cloud, as long as the collected signatures are stable (see Sect. 6, whitelists). Since connectivity to the visible APs is not needed, signatures can be collected anywhere and stored on an application specific, or group specific private server that allows *Put*-ting and *Get*-ting documents to create a private document collection embedded in any WiFi rich physical space.

3 Dissimilarity of Signatures

In this section, we explore the behavior of signatures collected from the WiFi infrastructure, called fingerprints, which are most readily available, and present in ever increasing AP (access point) densities. Since the system does not use any Cartesian location and distances, we first look at the mapping between signature dissimilarity and Cartesian distance.

As far as indoor positioning based on radio fingerprints is concerned, accuracy is predictable [3], and can be quantified in terms of *meters per dBm* of RSSI (received signal strength indicator) measurement accuracy. The organizers of several editions of Microsoft Indoor Location Competition mention that for the infrastructure-free mode, the expected indoor positioning accuracy was around 2m in most years, but was as high as 4m in tough environments [16]. Some improvements implemented for these wining methods could be employed by *AirDocs*, but not the ones that require extensive setup, training, and calibration. We next propose a dissimilarity measure that employs methods that have been validated by other researchers, and exhibit a monotonic behavior with real distance.

3.1 Dissimilarity Measure

Generally, positioning using fingerprints uses some function of distance in signal space, with Euclidean used in the RADAR paper [1], and many others tested in the literature. Caso et al. [4] tests Minkowski, cosine, Pearson correlation, and Shepard, finding that Euclidean and Pearson correlation provide the best results. In other studies, Mahalanobis is found to have the best performance, but for our setup it cannot be applied, since Android only gives one RSSI per 3s reading, therefore a covariation matrix between RSSI distributions of different APs cannot be obtained without extensive waiting. Torres-Sospedra et al. [34] explore many others distances ad dissimilarities used in the literature, and found Sørensen (BrayCurtis coefficient) to perform best. In addition, we adopted some other improvements proposed in [34]: zero-to-one normalized representation (Eq. 1) of a RSSI value x_i in dBm:

$$X_i = \text{normalized}(x_i) = \alpha \left(1 - \frac{x_i}{\min}\right)^e \quad (1)$$

We chose the scale value α so that the range of x_i observed values -99 dBm .. -30 dBm get mapped to X_i in the interval $[0,1]$. The purpose of this normalization is double: it maps negative power reading in dBm to positive values that are needed by some similarity measures, but also discounts more differences between low power readings. The latter means that differences in stronger signals are penalized, for example a -90 dBm to -85 dBm difference is less important than a -40 dBm to -35 dBm difference, as RSSI readings are known to be much noisier at low power values.

The Bray-Curtis dissimilarity (Sørensen distance) relies on APs common between the two fingerprints:

$$BCurtis(X, Y) = \frac{\sum_{i=1}^c |X_i - Y_i|}{\sum_{i=1}^c (X_i + Y_i)} \tag{2}$$

Most dissimilarity measures consider only common APs between two signatures, but as stated by Beder and Klepal [2], non common APs are a critical factor in reducing false similarities.




Fig. 2. Braycurtis dissimilarity vs real distance for all pairs of points on a square floor with 46 signatures.

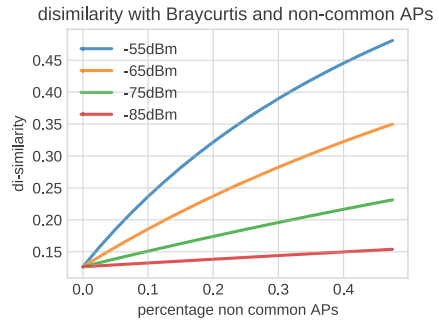


Fig. 3. Braycurtis dissimilarity (Eq. (2)) when a high percentage of non common APs received at different powers. Low power APs have sporadic reception, but have low effect on dissimilarity.

Device Variability. Another issue to be considered when evaluating dissimilarity is the harvesting of signatures with different devices, which might have different receiving characteristics. These are caused by different antenna patterns, and different RF chains. In order to mitigate the RSSI differences introduced by device heterogeneity, we use robust fingerprints that include differences of RSSIs corresponding to pairs of APs from the same WiFi fingerprint. The method is

derived from [7] and [17], and has been evaluated in [9] to be effective in dealing with device heterogeneity.

For a WiFi fingerprint, **only** for common APs $X = \{X_i\}$, the device invariant fingerprint is represented by differentiating power for consecutive APs:

$$X_{invariant} = \{(X_1 - X_2), (X_2 - X_3), \dots, (X_{n-1} - X_n), (X_n - X_1)\} \quad (3)$$

In *AirDocs* we use Bray-Curtis based dissimilarity function with all the mentioned features, and compare its monotonicity against real distance on data collected in an eight-story building, dataset 1a (described in Sect. 4.1). In Fig. 2, we consider all pairs of points on a floor that are 15 m or closer. We can see that dissimilarity increases linearly with actual distance, and beyond the 13 m mark, which is the length of a corridor, we begin to see values of 1.0 in dissimilarity, and a higher deviation due to more wall attenuation.

For APs missing between the two signatures, we consider them visible at -100 dBm (-99 dBm is the minimum observed value in datasets in this article), so that they contribute to the dissimilarity. To quantify the contribution of non common APs, we increase the percentage of non common APs between compared signatures to understand how they affect dissimilarity. As will be detailed later in the paper (Sect. 4.1), these are typical ratios of non common APs between spots that can be meters away from each other, and a dissimilarity of 0.25 will be later used as a threshold for selecting nearby documents. In Fig. 3, we see that weak APs that are not common do not affect the dissimilarity much, as it is common for mobile phones not to reliably pick them. In contrast, strong missing APs are a clear indicator of a faraway spot. In addition to that, we simply set the dissimilarity to 1.0 if the fraction of common APs is lower than a threshold (25%).

4 Dissimilarity Microbenchmarks

We evaluate the behaviour of the dissimilarity function and of the retrieval process on three different setups: one that we produced in our building, which will be made available as supporting material [27], and two which are publicly available [19], and [22]. In all the following sections, we use the real position of the documents only for reporting purposes, that is, to quantify the distance at which documents would be found, but **the real position is not used in the calibration of the system, indexing of the database, or anywhere in the search and retrieval process**, as it will not be available or necessary for the users of the system.

4.1 Building 1 Results

Datasets 1a and 1b were collected in our own office building with methodologies described in Sect. 3. The documents are spread across a square corridor (Fig. 4)




Fig. 4. Dataset 1a contains 8 floors with similar collection trajectory, and dataset 1b one floor.




Fig. 5. Searching the document with a different device at the same point, or at points closeby. Increasing the dissimilarity threshold guarantees that the document is found even if searching some distance from it.

consisting of true position and signature. For dataset 1a, eight floors were collected with collection points 2 steps apart (approximately 1.2m), holding two Android devices (Google Pixel 4A and Redmi Note 8) at waist level close to the body. For each collection point 4 directions were collected, rotating 90° after each scan. Dataset 1b contains the ground floor, sampled along the same trajectory, but at one step (0.6 m) resolution, and with the devices held at face level, away from the body. The building has an infrastructure WiFi, and a measurement point receives a median of 32 APs (minimum 20 APs, 95% = 49 APs). The collection methodology is described in more detail in Sect. 6.

We use documents placed in the environment (Fig. 4) to validate the searching process in several ways. First we search with a different device at the same physical point where the document was placed, and by using a large enough threshold we can guarantee that the document is found with high probability. In Fig. 5 we see that using a dissimilarity threshold of 0.22 guarantees that the document would be found in 99.5% of the cases. When searching at a distance from where the document was placed, an increasingly higher value of the threshold is necessary. As shown in Fig. 5, for dataset 1b, a search tolerance of 0–1.8 m would require a dissimilarity threshold of 0.22–0.3.

An alternate way of looking at the problem is to compute for each query signature the distance to the closest document in the database. This can be zero when a document collected with the other device is found at that same location, or non-zero for nearby documents. In Fig. 6, we see that in 72% of the cases the closest document (in signal space) returned is the one taken at the query spot, and in the rest of the cases is a spot 0.6 m away. The closer the response document, the more discriminate the signature is, and fewer faraway documents




Fig. 6. Distance to the closest point: *Put* and *Get* devices are different. In 72% of the cases, the same point (collected with the other device) is returned. Sample documents are 0.6 m apart.

| dataset 1b | | |
|------------|-----------|-----------|
| dissim | 50% of | 95% of |
| thresh | documents | documents |
| 0.20 | 0.63 m | 2.54 m |
| 0.25 | 1.21 m | 3.13 m |
| 0.30 | 1.52 m | 3.81 m |
| 0.35 | 1.88 m | 4.45 m |

Fig. 7. Document queries based on dissimilarity thresholds in signature space will also return some further away documents.

will be returned. For the rest of the experiments in the paper, we will use this measure of distance to closest document reported as a performance for all the datasets we explore.

As shown in Fig. 2, the relation between dissimilarity and real distance will make any decision based on a threshold in signal space to return several documents that have similar signatures, but are far from the query point. In Fig. 7, increasing the dissimilarity threshold returns, in addition to the desired document, other documents that have similar signatures. Whether these are considered as false positives depends on the actual application *AirDocs* is used for.

Simulated Impairments. Knowing that devices used by the system may vary in their antenna characteristics, we stress test the searching process to better understand higher diversity in devices.

We add Gaussian $\mathcal{N}(\mu, \sigma)$ noise to the database taken with the Pixel smartphone. The Redmi device queries for the closest point in the signal space from the database. As mentioned in Sect. 6, Fig. 19, the Pixel device already receives a $\mathcal{N}(-4.4 \text{ dBm}, 3.5 \text{ dBm})$ lower power. To explore a wider range of impairments, we add noise with $\mu = -5 \text{ dBm}$, and σ increased from 1 dBm to 8 dBm, and compute through cross-validation the all the obtained distances. In Fig. 8, we plot the 50%, 90% and 99% percentiles of the CDF for the obtained closest distances. The system degrades gracefully, even for a standard deviation of 8 dBm, with the resulted median distance to the closest point increasing from 0.6 m to 0.9 m.

Then, we alter the database with a fixed deviation of $\sigma = 3 \text{ dBm}$, and a varying offset $\mu = -1 \text{ dBm} \dots -10 \text{ dBm}$. Due to the differentiating feature of

the proposed dissimilarity used (Sect. 3.1), the median error obtained remains constant at 0.6m, and the 95% at 1.8m throughout the entire interval studied.




Fig. 8. Adding Gaussian (-5 dBm, stddev) noise to RSSI measurements, dataset 1b.

Another issue to explore is AP density, which has been increasing in recent years, even if sometimes the APs are virtual, being emitted from the same physical card. We eliminate fractions of APs from the query (Redmi) before querying the database (Pixel), and summarize the results in Fig. 9. The performance degrades gracefully and maintains performance even with one third of the existing APs. Please see the Sect. 6 for further discussion about AP density.




Fig. 9. Distance to closest point obtained when removing a fraction of APs (dataset 1b).

4.2 Building 2 Results

We use a dataset published by Indoor Location Competition, 2020 edition [22]. It contains data from two shopping mall buildings, that have been sampled at walking speed with the collector facing the necessary direction to complete the desired path. The collector stopped and marked certain points in the measurement trips, but WiFi and BLE beacon collection went on continuously. We post-processed the data to interpolate linearly the position of the collector at the time of each WiFi scan based on the timestamped positions logged during collection. The first floor of the first site is shown in Fig. 10, together with the resulted sampled points.

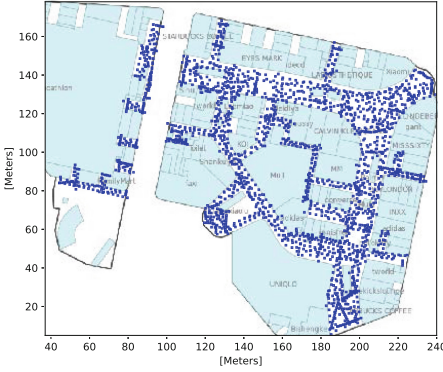


Fig. 10. Dataset 2 floor topology

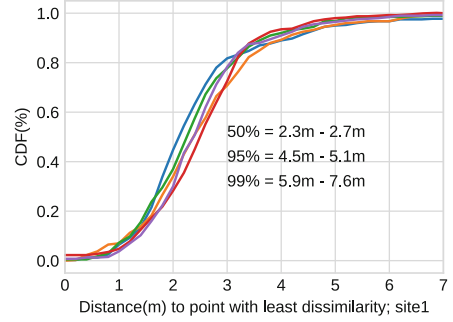


Fig. 11. Dataset 2 closest point

The first site has a high density of APs with a total of 4053 unique MAC addresses for the entire building, with 1452–2524 APs per floor, and 130–401 APs collected on average for each point. Despite this density, collecting at high speed and with the collector facing only one direction, yields a closest point that is higher than for datasets 1a and 1b (Fig. 11). Dissimilarities of these closest points appear invariant to density. *AirDocs* is usable even with these relaxed collection methods, but the operating circle around the user would be larger, which could be appropriate for malls, with larger spaces and higher mobility patterns.




Fig. 12. UJI 3rd floor topology, 25 complete sets of WiFi signatures taken at the same points over a span of 11 months

On the server side, the implication of high density and variability is that $q = 9$ offers 99% retrieval rate and $q = 20$ is necessary to achieve 100% document retrieval rate, where we consider documents all the WiFi sampled points. This result shows that high AP density is not always a blessing, since variability in the strongest AP received increases the search time on the server. This situation can be partially mitigated by identification and merging of virtual APs.

4.3 Building 3 Results

The third dataset used is based on data from the UJI repository at [19], from which we select one particular floor that has been measured repeatedly over a period of almost one year. The topology is depicted in Fig. 12 and consists of

24 points visited sequentially by the collector, and at the end returning in the opposite sense over the same points after a time of about 10 min. The reference points have rather high distances between them 1.8 m, respectively 4.2 m on each axis, but at each location 6 samples were collected with Android mobile devices. We post-processed the data so we merged the two readings 10 min apart, assuming that the collectors probably faced opposite direction during the return trip, and also averaged all the collected values resulting in one signal strength per AP per point.

Given that the dataset has a lower spatial resolution, we only used points 3–20 as queries for cross-validation since corner points only have one reasonable option to be returned, besides the point itself. We first verify how real distance and dissimilarity are linked for this topology by measuring dissimilarity between points taken in the same session (Fig. 13a, blue dots and line fit). This partially resembles the behavior for dataset 1 (Fig. 2), with the collection lattice binning of possible distances. Then for each collection point, we computed all dissimilarities to its past and future measurements, and collected all the values in the gray boxplot. The boxplot is manually placed horizontally based on its median value and the line fit to give an estimation of the error resulted from fingerprint aging.




Fig. 13. (a) Blue: dissimilarity vs distance for measurements taken at the same time; Boxplot: dissimilarities between scans of the same point at different times, cross validated across all times and all sets. (b) UJI 3rd floor number of APs per point in time. (Color figure online)

To understand better the behavior of fingerprints in time, we look at a timeline of the number of APs available for a point at different moments. In Fig. 13b, each point represents a measurement set, and the actual number of APs is averaged across all 24 points, with the standard deviation indicated at each point. As the collecting device does not pick up all APs at all times, even if the instantaneous number of APs is relatively stable, the total number of APs that are historically available is substantially higher. This mismatch behaves like an aging effect, but can be countered either by taking more measurements, using more collector orientations, or using a different stances for the device.

To evaluate how retrieving of documents works with old databases, we cross validate fingerprints against entire database of 24 points taken at future or past

time, and compute the closest point returned. In Fig. 14, we see that in 67% of the cases, the same location would result, and the rest distributed between the two available candidates, at 1.8 m and 4.2 m. The dissimilarities associated to these closest points are higher than the ones resulted in dataset 1, due to both fingerprint aging, and the actual spatial resolution of the sample points.

5 App and Server Implementation

The *AirDocs* middleware consists of an Android module that collects signatures, and a server that stores documents indexed by signatures and responds to queries based on signatures.

The Android module performs WiFi scans every 3 s, which is a limitation imposed by the operating system, as mentioned by other works as well [14]. The operating system provides a list of APs and the associated RSSI in dBm, which constitutes the WiFi fingerprint, the main part of the context signature.

The module also scans cellular networks, BLE devices, obtains the GPS location when possible and records the background sound. All this is performed in the 3 s frame and all information is included in the signature.

From Android version 8, an additional throttling mechanism was included in the operating system, in order to limit the frequency of scans and reduce power consumption. In Android 8, a background application can only perform one scan in 30 min. However, this is mitigated by using a foreground service, which is not limited by the throttling mechanism. In Android 9, each foreground application can scan 4 times in 2 min. This problem can not be circumvented and it limits the scanning capabilities of *AirDocs* middleware on this version of Android. From Android 10 and above, the throttling mechanism can be disabled from Settings, so it does not affect the scanning procedure of our middleware.

The *AirDocs* server communicates with the *AirDocs* smartphone module and is responsible for two scenarios. First, it receives a new document together with an associated signature and stores them in a database associated with the building, institution, or group. Second, it responds to searches based on signatures, identifying documents with low dissimilarities. The most similar signatures and their associated documents are sent to the mobile and then displayed in the application.

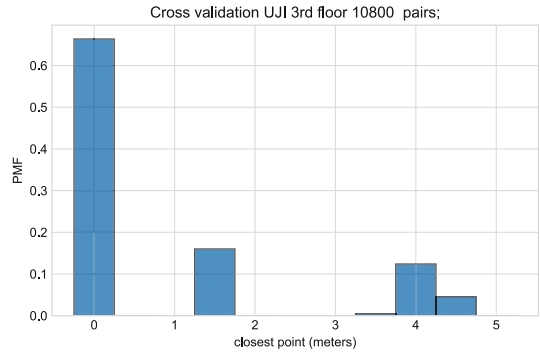


Fig. 14. UJI 3rd floor dissimilarity of the closest point across time

With a naive linear database implementation, insertions would take $O(1)$, and searches $O(n)$, where n is the number of documents stored on the server. *AirDocs* stores a spatial database, in the sense that documents are tied to location, but the actual locations of documents is not known. Therefore, most of the methods to index spatial information, such as *M-trees* and *R*trees*, or others used optimize databases of position labeled fingerprints are not applicable.

However, dissimilarities between documents impose some structure in the signal space. A natural spatial clusterization should produce hierarchical structure into buildings, floors, rooms. This would allow for searching in a tree like structure that uses spatial relationships of inclusion and adjacency.

We use a multi-labeling system, where clusters are headed by the strongest AP observed, which imposes locality to the search process. When receiving a document, the server assigns q labels, where q is the number of the strongest APs visible in the WiFi fingerprint. For every query, the server then searches only in q lists associated with the strongest APs. Due to instabilities of indoor WiFi, the strongest AP might not always be the same and the question is *how to determine the minimum value of q required to retrieve any document?* If we had perfectly regular signal propagation, and $q = 1$, APs in a plane would divide space among them in Voronoi regions. The strongest AP would be the Voronoi seeds, and the documents points in the Voronoi cells, based on their strongest RSSI value. In reality, due to irregular indoor propagation, these regions are not Voronoi shaped, and do not have definite borders, but are partially overlapping. The shape and the amount of overlap depend on the building, AP density, and the collection method.

The search complexity of this data-structure is therefore $O(q * n_V)$, where n_V is the number of documents inside a Voronoi region. The areas of Voronoi regions depend on the density of APs (as Voronoi vertices), but because of varying RSSI values for the strongest APs, the size of the region will be larger in practice. This search process is trivial to parallelize since a fingerprint can be searched in parallel in each of the q lists. In Sect. 4.1 we will estimate typical values for q , and for the size of the regions.




Fig. 15. Diameter of Voronoi-like regions. Multi label based searching will only query documents inside a diameter of 11m(median) for dataset 1b.

For dataset 1b, we now characterize the data structures necessary on the server, given that they are a function of AP density and of indoor propagation specifics. As detailed in Sect. 5, searches on the server are performed in q lists of documents, corresponding the q strongest APs for a fingerprint. We tested increasing values of q to determine the minimum value that allows retrieval of all documents. We considered all 85 measured points in this dataset as potential documents, since the resolution is of 0.6m. If for $q = 1$ the associated region is akin to a Voronoi region, for increasing q , this region is enlarged, and we measure the size of the region by measuring the longest distance between two points in such a region, calling it a diameter. In Fig. 15 we plot diameters for values of $q = 1..4$. For $q = 4$ all documents are retrieved successfully for dataset 1b, querying regions 11m wide on average.

6 Discussion

In previous sections we validated the fact that WiFi signatures can discriminate between close-by points in a variety of operating conditions, but there are still a number of open questions remaining:

Tunables: while the target is for the system to work out of the box for both users and sysadmins, there are still a few values that need to be calibrated: The dissimilarity threshold that governs the area around the user is invariant on the density of APs but the actual values that correspond to a given radius in meters tend to be different depending on the collection density. For the server part, the q value that governs the efficiency of the document search depends on AP density and on collection method ($q = 4$ for dataset 1, and $q = 20$ for dataset 2).

AP density affects the performance of the system in more than one way: on one hand many APs means more ways to discriminate between close locations, on the other hand unstable AP picked up by Android scanning introduces additional noise in the dissimilarity. More study is needed to understand whether certain APs contribute positively or negatively to the signatures.

Whitelists: in many setups there will be temporary APs, or APs that change location. If these are a small fraction of the total, their effect will not be visible, as shown in Sect. 4.1. However, for low AP density it is indicated that the system only use APs in a whitelist with MAC addresses that belong to the infrastructure. Also, most modern APs create virtual SSIDs, so the same physical card would broadcast under MAC addresses differing by one byte (Cisco), therefore a whitelist would be beneficial in unifying these readings. For searching on the server q value is also affected by virtual APs since the physically strongest AP might appear with several MACs, thus artificially increasing the search complexity.

Better harvesting: We target an app that is usable on most Android phones, and decided to rely on a default scanning procedure that takes 3–4 s, only getting one RSSI reading per AP. But using monitor mode on a laptop would allow receiving 10 beacons/second from most APs, which could allow

using the entire sampled distribution of received power, enabling richer signatures and better dissimilarity measures (Mahalanobis). Unfortunately, the use of laptops would decrease the accessibility of the project, but could be used for anchor documents or other high quality signatures.

Bad spots are those where document resolution is weaker, the WiFi fingerprint is not discriminate enough, or measurement is insufficient for retrieval of the document within a reasonable radius. More study is needed on how to identify these situations when needed, and either alert the user to take extra measurements, or prompt the sysadmin to improve the density of APs (physical, not virtual).

Curating documents: since documents are not placed on a map, a method to manage document collections by the server administration is needed. The proxy of distance used provides good clustering properties, in that documents beyond a certain distance, on different floors, or not having enough common APs have their dissimilarity set to 1. This allows for some organization of documents on buildings and floors, but management of documents in signal space is needed to perform periodic cleanup (because of institution policy for example), retrieving of lost/non-accessible documents, refreshing of fingerprints with changes in WiFi infrastructure, or addition of maps if they are available.

Signature Collection Methods. Since mobile phones have downsized antennas, the collector’s body orientation with respect to the building, and the relative position of the phone with respect to the collector are factors affecting the sampling of the signal strength. We explore both aspects of collecting, by having the collector gather one sample in four consecutive directions, 90° apart. The collector holds the phone either near to the body, at hip level, or at the face level, arm length away from the body. Data is collected simultaneously with two different phones - Google Pixel 4A, and Redmi Note 8, both running Android 10.




Fig. 16. Maximum power difference between same AP measured in: 2 directions at 90°; 3 directions at -90°, 0°, 90°; all 4 cardinal directions




Fig. 17. APs gained when harvesting several samples with collector body rotating around 90°.

In Fig. 16 we can see that the maximum power difference for the same AP ranges from 2.5 dBm when rotating 90° , 7 dBm when rotating 180° , and 9 dBm when considering all 4 directions. Boxplots² represent distribution of results gathered across spots and directions in the entire building (dataset 1a, described in Sect. 4.1). For comparison, Beder and Klepal [3] mention 2 m (best case) to 15 m (worst case) of positioning error per dBm RSSI measurement accuracy.

The effect of collecting data facing different directions can be seen in Fig. 17 as we count the number of APs that are cumulatively collected with 1–4 poses 90° apart. When rotating, Redmi keeps 75% of the APs, and drops/gains another 25%, when compared to the number of APs seen before rotation. For Pixel, the same figures are 87% and 13% respectively. This statistics are gathered for all consecutive 90° rotations, for all points in the dataset. By comparing all same AP readings in all the spots and in all directions (a total of 13000 measurements), we found that the Pixel - Redmi power difference is -4.3 dBm with a deviation of 3.6 dBm (Gaussian shaped). The interpretation of these numbers is that Redmi has a more directional antenna, and sees more of a power difference when rotating, but number of APs gathered is lower since it only gathers in the preferred direction.

To better quantify all these factors, we search for a given fingerprint taken with Pixel among all fingerprints taken with Redmi, considering the dissimilarity measure presented in Sect. 3.1. We cross validate using all the points on the floor (dataset 1b, described in Sect. 4.1), and report the real distance to the fingerprint with the minimum dissimilarity. In Fig. 18 boxplots summarize distances to the closest point in signal space for several combinations of directions used during collection (*Put*) and testing (*Get*). The best case is when we use all 4 directions collected, and the worse when selecting a random directions out of the 4. An acceptable performance is obtained using any two directions, including random. Since one measurement takes 3s in Android, we conclude that measuring at least




Fig. 18. All 4 dir: average of four samples measured in 4 directions 90° apart; dir 0: the direction along the walk; dir 0,2: experimenter uses direction 0, and opposite 180° ; dir 1,3: two orientations sideways, also 180° apart; dir 0,1: two directions 90° apart; dir 0–2: three directions sweeping from left to right, 90° apart; rand: one random direction; rand 2: average of two random directions (out of 4 possible). Measuring in two random directions either 90 or 180° apart is invariant on actual orientation of sampling and reduces measurement time.

² Boxplots in this article indicate middle quartiles (25%–75%), median, and whiskers at 5% and 95%. Outliers are shown outside the whiskers.

two different angles provides a good balance between the collection latency and the accuracy of the results. King et al. [11] also reports that measuring in two directions is enough for positioning purposes.

Phone Position Relative to the Body. We collected two datasets: 1a) at waist level and close to the body, and 1b) at face level and arm length away from the body. In Fig. 19, we summarize the findings comparing these two collection methods. Collecting away from the body, and at a higher pose brings several advantages: more APs are gathered by either phone by 18%-31%; There is less of a difference in number of APs gathered as the effect of rotating about; Less of a difference in power gathered as the effect of rotating about, therefore less variability; The consistent power difference between devices is reduced.

All these factors show that collecting high and away from the body is beneficial, as it reduces measurement variability for the three factors that affect dissimilarity: absolute power, number of APs, and device difference. These recommendations have an impact on the way the user collects the fingerprint whenever documents are placed or queried.

6.1 Future Work

One method to obtain increased resolution for the dissimilarity of the signatures is to use of additional sensors besides WiFi. BLE infrastructures are not as prevalent as WiFi, but all the issues explored in this paper for WiFi apply directly when beacons are available (datasets 1 and 2 also have BLE information but their density is not operational).

| direc- tions | Pixel | | Redmi | |
|------------------------------|-----------------------------|------|----------|------|
| | Waist | Face | Waist | Face |
| | Median # of APs | | | |
| 1 | 41 | 52 | 29 | 38 |
| 2 | 48 | 60 | 39 | 46 |
| 4 | 54 | 64 | 45 | 54 |
| | Power difference[dBm] | | | |
| 2 | 2.0 | 3.0 | 3.0 | 3.0 |
| 3 | 6.0 | 6.0 | 7.0 | 6.0 |
| all | 9.0 | 8.0 | 10.0 | 8.0 |
| | APs after rotation | | | |
| com- mon gain/ loss | 64% | 87% | 64% | 82% |
| | 36% | 13% | 33% | 18% |
| | RSS power diff. Pixel-Redmi | | | |
| | Waist | | Face | |
| mean | -4.4 dBm | | -1.8 dBm | |
| stddev | 3.5 dBm | | 3.8 dBm | |

Fig. 19. Improvements obtained with signal harvesting stance: Collecting high and away from the body improves: power, number of APs collected, and reduces effect of device variation.




Fig. 20. Location specific signatures built using signals received to or from the smartphone.

Figure 20 shows several possible sources of data to enrich signatures making them more discriminate with respect to location. 4G/5G has a rather low positioning accuracy, but is available in all smartphones, and could be used to speed up the searching structures in the server. Sound reflections (as used in project EchoTag [36]) are another source of enriching the signature that does not require deploying of additional infrastructure. Basically, any context information that is *stable, available, and easily collectable* by the phone can become part of the signature.

Contact tracing [24,25,35] has recently seen a surge of interest, and has similar requirements with our system: no additional infrastructure, and simple operation with existing smartphones. *AirDocs* explores the same idea of proximity based on dissimilarity, and can be used as support for a contact tracing app, since the 1 m–4 m proximity detection is within range of current health advisories.

Finally, as part of future work, we plan to open-source the client app (in public Application Stores) and the server, as well as publish all the measured data on *Zenodo* [27].

7 Related Work

Several systems were proposed to achieve positioning based solely on the existing wireless infrastructure using propagation properties, but many of them require extensive training and updating to maintain a positioning service [10,39]. In contrast, *AirDocs* proposes management of documents in a context aware fashion, but not linked to geographic locations which are natural contexts. Association of documents with locations has been explored—web documents are being geo-tagged and geo-referenced [29], and in the database community there are efforts to formalize searches for objects distributed in space [37].

Use of fingerprints for positioning has began in 2000 with the seminal paper by Padmanabhan [1], but has since developed into a rich research area in which several engineering approaches are possible. For a taxonomy, see [12], which describes choices of types of measurements, estimation methods, radio maps, collection methods, types of collectors. Collection effort, also called training, or war driving, is the main disadvantage for fingerprint based location, and some researchers have proposed the use of monitors [13] to minimize the training process. Yang et al. [38] tries to reduce training effort using an informal site survey by untrained users. Signatures are recorded with gait measurements and mapped to real space using MDS (multi dimensional scaling) and ground truth points obtained by GPS or manually. Google uses undisclosed methods to crowdsource data from all the users, and offers sparsely available indoor positioning, but does not have a public API, and comes with serious privacy concerns.

The EchoTag project [36] uses the microphone and speaker of the mobile phone to create a sound signature specific to the location. We plan to explore this direction with the purpose of creating an even richer signature for *AirDocs*.

Augmented reality is an emerging technology that “supplements the real world with virtual (computer-generated) objects that appear to coexist in the

same space as the real world” [5]. [30] mentions projects spawned from MIT Media Lab’s project sixth sense, that achieves a form of augmented reality by requiring the user to carry a projector and a camera to recognize hand gestures. *AirDocs* is an enabler of augmented reality in the sense that documents are embedded in physical space, but without requiring positioning, head mounted displays, or instrumentation of the environment.

Dousse et al. [8] develop a purely fingerprint-based place learning method. Its core is a density-based clustering algorithm that works directly on the raw WiFi fingerprints. They also study the behavior of fingerprints with respect to space and time, but their focus on learning about stationary *places* by using 60s sampling, manually labeled sets, and an unspecified spatial resolution of these places. Also, locations are visited for more than 5 min, in contrast with *AirDocs*, which aims for a more fluid user experience.

Proximity based on fingerprint comparison has been explored both for the purpose of privacy implication [31], and contact tracing [35]. *AirDocs* exploits more the resolution available in the dissimilarity - distance function, and can be used as a primitive for both problems.

8 Conclusion

We benchmarked *AirDocs*, a system that makes use of signatures composed of stable information about the location, that is easily collectable by smartphones. Documents are managed spatially, but without the use of a location system, which usually requires extra infrastructure, training, or crowd-sourcing of measurements. We explore the use of WiFi fingerprints as the main component of a location dependent signature, and define a measure of dissimilarity that is mostly monotonic with real distance. We show that typical WiFi deployments enable reliable retrieval of documents in areas with radius 0.6 m–2 m (median values), and characterize the behavior of dissimilarity with respect to: impairments and differences between measuring devices, collection methods, density of APs, and signature aging.

References

1. Bahl, P., Padmanabhan, V.N.: Radar: an in-building RF-based user location and tracking system. In: Proceedings IEEE INFOCOM 2000, vol. 2, pp. 775–784. IEEE (2000)
2. Beder, C., Klepal, M.: Fingerprinting based localisation revisited a rigorous approach for comparing RSSI measurements coping with missed access points and differing antenna attenuations. In: 2012 International Conference on Indoor Positioning and Indoor Navigation, IPIN 2012 (2012). <https://doi.org/10.1109/IPIN.2012.6418940>
3. Beder, C., McGibney, A., Klepal, M.: Predicting the expected accuracy for fingerprinting based wifi localisation systems. In: 2011 International Conference on Indoor Positioning and Indoor Navigation, pp. 1–6 (2011). <https://doi.org/10.1109/IPIN.2011.6071939>

4. Caso, G., De Nardis, L., Di Benedetto, M.G.: A mixed approach to similarity metric selection in affinity propagation-based wifi fingerprinting indoor positioning. *Sensors* **15**(11), 27692–27720 (2015). <https://doi.org/10.3390/s151127692>. <https://www.mdpi.com/1424-8220/15/11/27692>
5. Chatzopoulos, D., Bermejo, C., Huang, Z., Hui, P.: Mobile augmented reality survey: from where we are to where we go. *IEEE Access* **5**, 6917–6950 (2017)
6. de Matos, E., et al.: Context information sharing for the internet of things: a survey. *Comput. Netw.* **166**, 106988 (2020). <https://doi.org/10.1016/j.comnet.2019.106988>. <https://www.sciencedirect.com/science/article/pii/S1389128619310400>
7. Dong, F., Chen, Y., Liu, J., Ning, Q., Piao, S.: A calibration-free localization solution for handling signal strength variance. In: Fuller, R., Koutsoukos, X.D. (eds.) MELT 2009. LNCS, vol. 5801, pp. 79–90. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04385-7_6
8. Dousse, O., Eberle, J., Mertens, M.: Place learning via direct wifi fingerprint clustering. In: 2012 IEEE 13th International Conference on Mobile Data Management, pp. 282–287. IEEE (2012)
9. Fang, S.H., Wang, C.H., Chiou, S.M., Lin, P.: Calibration-free approaches for robust wi-fi positioning against device diversity: a performance comparison. In: 2012 IEEE 75th Vehicular Technology Conference (VTC Spring), pp. 1–5 (2012). <https://doi.org/10.1109/VETECS.2012.6240088>
10. He, S., Chan, S.H.G.: Wi-fi fingerprint-based indoor positioning: recent advances and comparisons. *IEEE Commun. Surv. Tutor.* **18**(1), 466–490 (2016). <https://doi.org/10.1109/COMST.2015.2464084>
11. King, T., Kopf, S., Haenselmann, T., Lubberger, C., Effelsberg, W.: Compass: a probabilistic indoor positioning system based on 802.11 and digital compasses. In: Proceedings of the 1st International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization, WiNTECH 2006, pp. 34–40. Association for Computing Machinery, New York (2006). <https://doi.org/10.1145/1160987.1160995>
12. Kjærgaard, M.B.: A taxonomy for radio location fingerprinting. In: Hightower, J., Schiele, B., Strang, T. (eds.) LoCA 2007. LNCS, vol. 4718, pp. 139–156. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75160-1_9
13. Lim, H., Kung, L.C., Hou, J.C., Luo, H.: Zero-configuration indoor localization over IEEE 802.11 wireless infrastructure. *Wirel. Netw.* **16**(2), 405–420 (2010)
14. Liu, H.H., Liu, C.: Implementation of wi-fi signal sampling on an android smartphone for indoor positioning systems. *Sensors* **18**(1) (2018). <https://doi.org/10.3390/s18010003>. <https://www.mdpi.com/1424-8220/18/1/3>
15. Lohiya, R., Thakkar, A.: Application domains, evaluation data sets, and research challenges of IoT: a systematic review. *IEEE Internet Things J.* **8**(11), 8774–8798 (2021). <https://doi.org/10.1109/JIOT.2020.3048439>
16. Lymberopoulos, D., Liu, J.: The microsoft indoor localization competition: experiences and lessons learned. *IEEE Signal Process. Mag.* **34**(5), 125–140 (2017). <https://doi.org/10.1109/MSP.2017.2713817>
17. Mahtab Hossain, A., Jin, Y., Soh, W.S., Van, H.N.: SSD: a robust RF location fingerprint addressing mobile devices’ heterogeneity. *IEEE Trans. Mob. Comput.* **12**(1), 65–77 (2013). <https://doi.org/10.1109/TMC.2011.243>
18. Makris, P., Skoutas, D.N., Skianis, C.: A survey on context-aware mobile and wireless networking: on networking and computing environments’ integration. *IEEE Commun. Surv. Tutor.* **15**(1), 362–386 (2012)

19. Mendoza-Silva, G.M., Richter, P., Torres-Sospedra, J., Lohan, E.S., Huerta, J.: Long-term wi-fi fingerprinting dataset and supporting material. Zenodo (2017). <https://doi.org/10.5281/zenodo.1066041>
20. Microsoft: Microsoft indoor localization competition (2017). <https://www.kaarta.com/kaarta-wins-microsoft-indoor-localization-competition-for-second-consecutive-year/>
21. Microsoft: Microsoft indoor localization competition (2018). <https://www.microsoft.com/en-us/research/event/microsoft-indoor-localization-competition-ipsn-2018/>
22. Microsoft: Indoor location competition dataset (2020). <https://github.com/location-competition/indoor-location-competition-20>. gitHub repository
23. Microsoft: Microsoft indoor localization competition (2020). <https://github.com/location-competition/indoor-location-competition-20>
24. Nguyen, C., et al.: A comprehensive survey of enabling and emerging technologies for social distancing-Part I: fundamentals and enabling technologies. *IEEE Access* **8**, 153479–153507 (2020)
25. Nguyen, C., et al.: A comprehensive survey of enabling and emerging technologies for social distancing-Part II: emerging technologies and open issues. *IEEE Access* **8**, 154209–154236 (2020)
26. Niculescu, D.S., Ruse, L., Carabas, C.: Documents in the air. In: 2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops), pp. 1–4 (2022). <https://doi.org/10.1109/PerComWorkshops53856.2022.9767318>
27. Niculescu, D., Ruse, L., Carabas, C.: Documents in the air dataset and supporting materials. Zenodo (2021). <https://doi.org/10.5281/zenodo.4999707>
28. Perera, C., Zaslavsky, A., Christen, P., Georgakopoulos, D.: Context aware computing for the internet of things: a survey. *IEEE Commun. Surv. Tutor.* **16**(1), 414–454 (2013)
29. Poesse, I., Uhlig, S., Kaafar, M.A., Donnet, B., Gueye, B.: IP geolocation databases: unreliable? *ACM SIGCOMM Comput. Commun. Rev.* **41**(2), 53–56 (2011)
30. Sanchez-Vives, M.V., Slater, M.: From presence to consciousness through virtual reality. *Nat. Rev. Neurosci.* **6**(4), 332–339 (2005)
31. Sapiezynski, P., Stopczynski, A., Wind, D.K., Leskovec, J., Lehmann, S.: Inferring person-to-person proximity using wifi signals. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* **1**(2) (2017). <https://doi.org/10.1145/3090089>
32. Sezer, O.B., Dogdu, E., Ozbayoglu, A.M.: Context-aware computing, learning, and big data in internet of things: a survey. *IEEE Internet Things J.* **5**(1), 1–27 (2018). <https://doi.org/10.1109/JIOT.2017.2773600>
33. Shit, R.C., Sharma, S., Puthal, D., Zomaya, A.Y.: Location of things (LoT): a review and taxonomy of sensors localization in IoT infrastructure. *IEEE Commun. Surv. Tutor.* **20**(3), 2028–2061 (2018)
34. Torres-Sospedra, J., Montoliu, R., Trilles, S., Óscar Belmonte, Huerta, J.: Comprehensive analysis of distance and similarity measures for wi-fi fingerprinting indoor positioning systems. *Expert Syst. Appl.* **42**(23), 9263–9278 (2015). <https://doi.org/10.1016/j.eswa.2015.08.013>. <https://www.sciencedirect.com/science/article/pii/S0957417415005527>
35. Trivedi, A., Zakaria, C., Balan, R., Becker, A., Corey, G., Shenoy, P.: Wifitrace: network-based contact tracing for infectious diseases using passive wifi sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* **5**(1) (2021). <https://doi.org/10.1145/3448084>

36. Tung, Y.C., Shin, K.G.: Echotag: accurate infrastructure-free indoor location tagging with smartphones. In: Proceedings of the 21st Annual International Conference on Mobile Computing and Networking, pp. 525–536 (2015)
37. Wu, D., Cong, G., Jensen, C.S.: A framework for efficient spatial web object retrieval. *VLDB J.* **21**(6), 797–822 (2012)
38. Yang, Z., Wu, C., Liu, Y.: Locating in fingerprint space: wireless indoor localization with little human intervention. In: Proceedings of the 18th Annual International Conference on Mobile Computing and Networking, pp. 269–280 (2012)
39. Zhu, X., Qu, W., Qiu, T., Zhao, L., Atiquzzaman, M., Wu, D.O.: Indoor intelligent fingerprint-based localization: principles, approaches and challenges. *IEEE Commun. Surv. Tutor.* **22**(4), 2634–2657 (2020). <https://doi.org/10.1109/COMST.2020.3014304>