

Indoor Quality-of-Position Visual Assessment using Crowdsourced Fingerprint Maps

CHRISTOS LAOUDIAS, University of Cyprus, Cyprus

ARTYOM NIKITIN, Skoltech, Russia

PANAGIOTIS KARRAS, Aarhus University, Denmark

MOUSTAFA YOUSSEF, Alexandria University, Egypt

DEMETRIOS ZEINALIPOUR-YAZTI, University of Cyprus, Cyprus

Internet-based Indoor Navigation (IIN) architectures organize signals collected by crowdsourcers in *Fingerprint Maps (FMs)*, in order to improve localization given that satellite-based technologies don't operate accurately in indoor spaces where people spend 80-90% of their time. In this paper, we study the *Quality-of-Position (QoP)* assessment problem, which aims to assess in an offline manner the localization accuracy that can be obtained by a user that aims to localize using a FM. Particularly, our proposed *ACCES* framework uses a generic interpolation method using *Gaussian Processes (GP)*, upon which a *navigability score* at any location is derived using the *Cramer-Rao Lower Bound (CRLB)*. We derive adaptations of *ACCES* for both Magnetic and Wi-Fi data and implement a complete visual assessment environment, which has been incorporated in the Anyplace open-source IIN. Our experimental evaluation of *ACCES* in Anyplace suggests the high qualitative and quantitative benefits of our propositions.

Additional Key Words and Phrases: Indoor Localization, Accuracy Estimation, Fingerprint Management

ACM Reference Format:

Christos Laoudias, Artyom Nikitin, Panagiotis Karras, Moustafa Youssef, and Demetrios Zeinalipour-Yazti. 2021. Indoor Quality-of-Position Visual Assessment using Crowdsourced Fingerprint Maps. *ACM Trans. Spatial Algorithms Syst.* 1, 1 (January 2021), 33 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Site survey tools (e.g., Ekahau.com, tamos.com, inssider.com) are typically used to generate radio coverage heatmaps of indoor spaces based on readings collected during measurement campaigns by mobile devices equipped with Wi-Fi sensors. Such heatmaps visualize the signal strength around the Wi-Fi Access Points (AP) deployed in an indoor space using a color map. Similar tools running on modern sensor-rich smartphones are available for surveying the magnitude of earth's magnetic field inside a building to deliver the corresponding magnetic signal heatmap. Looking at such heatmaps at different granularities, one can determine where sensor readings are deteriorated by cabling, electronics, appliances, or physical barriers and subsequently improve the situation by installing additional APs or by incentivizing users to contribute more data. Even though such heatmaps are useful, they only provide limited information about the location accuracy that one will experience

Authors' addresses: Christos Laoudias, University of Cyprus, Nicosia, 1678, Cyprus; Artyom Nikitin, Skoltech, Moscow, 143026, Russia; Panagiotis Karras, Aarhus University, Aarhus, 8200, Denmark; Moustafa Youssef, Alexandria University, Alexandria, 21526, Egypt; Demetrios Zeinalipour-Yazti, University of Cyprus, Department of Computer Science, 2109 Nicosia, University of Cyprus, Cyprus, +357-22-892755, dzeina@cs.ucy.ac.cy, Nicosia, 2109, Cyprus.

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.

© 2021 Association for Computing Machinery.

2374-0353/2021/1-ART \$15.00

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

using an indoor localization system. For example, the Wi-Fi APs that are deployed to serve users inside a large open-plan area will have relatively strong signals across this area; yet the positioning accuracy might be low, because the signals do not vary significantly to distinguish different locations effectively. This is because in open-plan areas the WiFi signals become weaker as the user walks away from a WiFi AP and the distance between them is increased; however, this distance can be several (even a few tens of) meters for the drop in the WiFi signal to be larger than the signal noise, which makes the signal values fluctuate when the user is not moving.

In other words, in open-plan areas two locations A and B that are several meters away may have only slightly different WiFi signal values (i.e., below the noise level) and a positioning algorithm that relies on these data can wrongly estimate that the user is at location B , while he/she is in fact at location A , thus leading to high positioning error. On the other hand, in an indoor space if there is a wall or obstacle separating locations A and B , then this will introduce additional signal attenuation so that the signal values at the two locations will differ significantly and the positioning algorithm will be able to estimate the correct location. Similarly, the magnetic field will be rather stable if there are no sources of magnetic disturbance (e.g., power cables, metal surfaces, elevators, escalators, etc.) leading to poor location accuracy. When there is only Earth's magnetic field, the magnetic data, i.e., the three components of the magnetic field, that are stored in a database will be very similar at locations A and B . Therefore, a positioning algorithm that compares the magnetic data observed when the user resides at location A with the data in the database may output the wrong location B . However, if there is a source of magnetic disturbance close to location A this will greatly affect the magnetic values. Consequently, the magnetic data at location A will differ significantly from the magnetic data at location B and the positioning algorithm will be able to estimate the correct user location.

Accurate and robust indoor location information could be used to extend the services of location-based social networks (e.g., [41]) to large indoor areas like shopping malls, train stations, and airports, for analyzing the presence of people using crowdsensing [46], or to identify stop-by behaviors [49]. At the same time, however, the underlying localization systems are expected to deliver not just bare location estimates, but also information about their reliability, i.e., how accurate these estimates are. For example, after a location request in any popular mobile navigation application it typically also displays a 'blue disc' centered at the user's location, whose radius reflects the uncertainty (i.e., level of accuracy) associated with his/her location. Such information about the computed location is highly desirable and greatly improves the user experience. Nevertheless, it is only available to the user *online* (i.e., after the location request) based on the type of localization technology (e.g., Wi-Fi, cellular, Bluetooth, RFID, or satellite-based localization among others).

From the perspective of the deployment staff though, aiming for a higher-accuracy localization service, it is important to know the anticipated location accuracy in different areas of a building *offline*, i.e., shortly after the data for the signal maps are collected and prior to any user-initiated location request. On one hand, a system operator who is aware that high location error is expected in some part of a building could take action to extend the localization infrastructure (e.g., install additional APs or beacons) and/or offer incentives to crowdsourcers for contributing more data in that area. On the other hand, a third-party application provider might choose not to offer a service (e.g., location-based advertisements and coupons or gaming) inside a certain venue until a minimum level of location accuracy is guaranteed.

Hybrid localization services [53] fuse multiple data sources to obtain both high accuracy and robustness in cases of data shortages. To exploit such services, an indoor localization service provider should be able to integrate new data sources and types. A common approach provides such integration capability by treating sensor measurements as *fingerprints*, each associated with a location. In the offline phase, such fingerprints are collected into a *Fingerprint Map (FM)* and stored in a database;

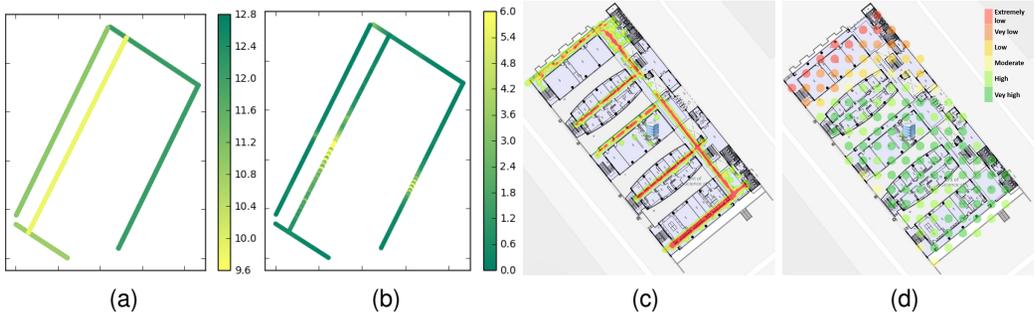


Fig. 1. Spatial density of collected fingerprints in an indoor environment (a) magnetic data; (c) Wi-Fi data. Resulting localization error in meters for magnetic data (b); resulting location accuracy for Wi-Fi data (d). The plots imply that there is a mismatch between fingerprint density and actual location accuracy.

subsequently, in the online phase, they are compared to readings from incoming user location requests. This *FM*-based approach treats all data sources equally, representing each fingerprint as a set of values associated with a location. For example, our in-house *Anyplace* Indoor Navigation Service [58] utilizes crowdsourced fingerprint data and achieved an average localization error of 1.96 *meter* by using Wi-Fi and inertial sensor data [35].

Recently, we introduced a generic framework that provides offline positioning accuracy assessment on arbitrary *FM*s. Our framework, titled *ACCES* (short for “ACCuracy ESTimation”) [37], achieves this by calculating a *navigability score* at any location of interest, while disregarding data origin. Our approach comprises three steps: First, we apply a black-box technique for interpolating arbitrary fingerprints based on a widely used statistical tool called *Gaussian Processes* (GP) [42], suitable for modeling smooth noisy data. This tool allows to: (i) predict sensor readings at chosen locations given the initial input data contained in the *FM*; and (ii) estimate the uncertainty of such predictions in the form of the variance of a Gaussian distribution. Then, given the predictive distribution calculated with such interpolation technique, we derive a lower bound for the uncertainty in the location estimate, i.e., the *localization error*, in the form of the *Cramer-Rao Lower Bound* (CRLB). The CRLB is used in estimation theory to derive lower bounds on the variance of an estimator. Obtaining such a lower bound on localization error is important for location and navigation service providers as it: (i) highlights theoretical accuracy limitations of a service; and (ii) provides insights on how real location accuracy could be improved. We utilize the derived CRLB as the *navigability score* for *FM* at any location and apply it to real-world fingerprint data.

Visualized in a user-friendly form such as a heatmap, this *navigability score* can be used by a location service architect or operator to either take actions towards improving the existing *FM* (e.g., by increasing the density of fingerprints) or to decide on exploiting other data sources at locations of low score. Figure 1 demonstrates that high fingerprint density does not guarantee high location accuracy, as there are other factors that may negatively impact accuracy including low signal variation and low dimension of the fingerprints (e.g., low number of Wi-Fi APs). For instance, Figure 1a shows the colored magnetic field fingerprint density map over a set of indoor corridors, where the green color indicates high density as opposed to low density (yellow color). Even though sensor readings are very dense (i.e., most of the regions are green), and each fingerprint contains more than one value (e.g., three components of a magnetic field vector), the location accuracy downgrades at certain locations. This is due to insufficient signal variation, as shown in Figure 1b where the yellow color

indicates high location error (i.e., low accuracy) and the green color low location error (i.e., high accuracy). Similarly for Wi-Fi data, the fingerprint density might seem sufficient and comparable to other parts of a building floor; see Figure 1c where the red color indicates high WiFi fingerprint density as opposed to low fingerprint density (yellow color). However, the resulting location accuracy may vary significantly from extremely low (red color) to very high (green color) accuracy, as shown in Figure 1d. This is due to the presence of less Wi-Fi APs covering the top wing, as opposed to the bottom wing.

This paper builds on our previous works in [37], in which we presented the design and preliminary results of our ACCES operator for indoor location accuracy estimation that is applicable to magnetic fingerprint data, and in [3], where we presented our Fingerprint Management Studio (FMS) that is a spatio-temporal platform for indoor signal data curation. In this paper we propose ACCES⁺ that incorporates several new improvements, especially a novel accuracy estimation operator coined ACCES-w, which is tailored to multi-dimensional Received Signal Strength (RSS) data from Wi-Fi routers. In particular, ACCES⁺ is a novel architecture for offline evaluation of Quality-of-Position (QoP) using multi-source fingerprint maps comprising two separate quality assessment operators, i.e., the ACCES operator for magnetic data and the ACCES-w operator for Wi-Fi data. As such, ACCES⁺ paves the way towards a deep understanding of where location data fails to provide the expected localization accuracy at the pre-deployment, rather than the post-deployment, stage. All our propositions are evaluated using real indoor localization data in a prototype architecture we have developed and fully integrated into our *Anyplace*¹ Indoor Navigation Service [58] with over 70K users and over 250K interactions. The overall contributions of our work are summarized as follows:

- We propose a generic interpolation technique for arbitrary fingerprint maps by *Gaussian Process Regression* (GPR), allowing to predict both measurements and their uncertainty at any location of a venue.
- Given the interpolation of fingerprints and corresponding uncertainties, we derive a theoretical upper bound on localization accuracy in the form of a CRLB, which can be utilized as a navigability score.
- We extend the notion of the ACCES operator, which is applicable to magnetic data, and introduce the ACCES-w operator that processes the ubiquitous Wi-Fi signals for the provision of reliable location accuracy estimates inside buildings.
- We augment the FMS capabilities to orchestrate the collection of location-dependent fingerprints and qualitatively assess the Wi-Fi coverage and data rates by implementing the ACCES-w operator as a plugin, while fully integrating FMS into the Anyplace service software stack to indicate areas where crowdsourcers should collect more data or new Wi-Fi APs should be installed to improve the localization accuracy.
- We evaluate our work with extensive experimentation and analysis on real data including both small scale magnetic datasets for micro-benchmarking, as well as large scale Wi-Fi datasets obtained through our award-winning Anyplace service that hosts thousands of buildings and associated localization data.

The rest of the paper is structured as follows. Section 2 overviews the related work on indoor localization methods and location accuracy estimation techniques. The proposed ACCES⁺ architecture for offline quality assessment of fingerprint maps is described in Section 3 while the domain-specific operators for Magnetic and Wi-Fi data are presented in Section 4. The prototype implementation of ACCES⁺ and integration with FMS into the Anyplace service is presented in Section 5. The performance evaluation results using real magnetic and Wi-Fi data are discussed in Section 6. Section 7 summarizes concluding remarks and directions for future work.

¹Available at: <https://anyplace.cs.ucy.ac.cy/>

2 BACKGROUND AND RELATED WORK

In the following, we outline the related work on indoor localization methods, as well as on techniques to estimate location accuracy with a fingerprint map.

2.1 Indoor Localization

The localization literature is very broad and diverse as it exploits several technologies [26, 27, 58]. Satellite positioning is ubiquitously available but has an expensive energy tag, may be negatively affected by the environment (cloudy days, forests, downtown areas), and is not available indoors. The localization research community has proposed numerous alternative solutions, including Infrared, Bluetooth [1, 5], Wireless LANs [7, 14, 56], Ambient Magnetic Field [17], Artificial Quasi-static Electromagnetic Field [8], Visual and Acoustic Analysis [13], Inertial Measurement Units (IMU) [16], Ultra-Wide-Band (UWB), and Sensor Networks, and their combinations in hybrid systems [34].

In terms of data modeling, indoor localization algorithms can be categorized into: (i) *pure modeling*, where locations are estimated based on user-collected online measurements and a priori system information, e.g., positions of the Wi-Fi APs or Bluetooth beacons and an indoor signal propagation model that translates each signal measurement (e.g., timing, angle, or signal strength) to distance from the associated transmitter [15, 55]; (ii) *fingerprint-aided modeling*, where both user and AP locations are estimated based on user-collected online measurements and some pre-collected location measurements called *fingerprints* [9, 11, 18, 38, 48]; and (iii) *pure fingerprint-based*, which is based solely on the similarity of online measurements with pre-collected fingerprints [13, 17, 25, 39, 40, 50, 56].

An advantage of these pure fingerprint-based indoor localization algorithms is their applicability regardless of the underlying data sources. Several techniques collect sensor measurements, e.g., radio signals from Wi-Fi APs or Bluetooth beacons, magnetometer and light sensor readings, and store them in a database at high density; see [30] for an extensive survey of Wi-Fi fingerprint-based approaches. For instance, any commercial smartphone can collect Wi-Fi signal strength fingerprints from the surrounding Wi-Fi routers as part of the standard passive wireless network monitoring functionality. For indoor localization purposes, such fingerprints are collected at certain locations (x, y) , pin-pointed on a building floor map (e.g., every few meters) in an offline phase. Subsequently, these fingerprints are joined in an $N \times R$ matrix, the *Fingerprint Map (FM)*, where N is the number of unique locations and R is the number of Wi-Fi routers. Users can then compare the observed RSS measurements against the *FM* to find the best match, using pattern matching methods including deterministic algorithms such as k Nearest Neighbor (k NN) [4] and its popular Weighted k NN (Wk NN) variant [33] or probabilistic algorithms [43, 56]. This approach is also applicable to magnetic data, where a smartphone's magnetometer fingerprints consist of 3 values each, corresponding to the magnitude and direction (with respect to the measuring device's reference frame) of the earth's ambient magnetic field combined with the magnetic field from other sources, such as electronic devices.

2.2 Accuracy Estimation

The goal of estimating localization accuracy is to predict the error between the computed location and the actual location of a user given noisy measurements along with fingerprint and/or system information that are used by pure modeling techniques. In this paper, we focus on estimating the achievable accuracy offline given only a fingerprint map, in the context of pure fingerprinting. Here, we present the literature on the accuracy estimation problem.

Regarding *online* location accuracy estimation, i.e., while location requests are sent to the service, an approach for localization confidence estimation is presented in [12], which does not require

knowledge of the underlying localization algorithm. In the case of probabilistic fingerprint matching algorithms the estimated location error can be quantified by means of the covariance matrix of the expected user location given the observed Wi-Fi RSS fingerprint [28]. Authors in [32] propose the use of low-level features (i.e., RSS values from Wi-Fi APs) measured at different locations as well as their respective localization errors, to train different regression models that allows to predict the localization errors at new locations given new observed values of the low-level features at these locations. However, both methods can be applied to estimate accuracy only when using the localization service in order to have newly observed fingerprints. Along the same line, authors in [22] evaluate four FM construction methods based on Wi-Fi RSS data, i.e., point-by-point manual data collection, a walking survey, semi-supervised learning, and unsupervised learning, in terms of the location error attained by various fingerprint matching algorithms using online test data. The work in [54] focuses on estimating online the accuracy of raw sensor measurements (e.g., Wi-Fi or inertial sensors) rather than the accuracy of the computed locations.

With respect to *offline* accuracy estimation, i.e., prior to running the location service, algorithms based on the diversity of spatial measurements are proposed in [31]. One of the proposed approaches involves splitting the indoor environment into small clusters and merging adjacent clusters based on the similarity of RSS distributions. A cluster's final size indicates the localization accuracy that can be achieved at its area. While practical enough, this solution lacks a formal model to allow providing strong guarantees. Authors in [24] present an algorithm named TBNPD (Tolerance Based - Normal Probability Distribution) that calculates the correlation level between each pair of fingerprint entries forming the FM. This approach offers the possibility to assess the uniqueness of each fingerprint and evaluate the quality of FMs; however, it is applicable only to Wi-Fi RSS data. A solution by the mapping company *Here* computes a quality metric for the sufficiency of the collected fingerprints and another metric for the quality of the wireless infrastructure (e.g., Wi-Fi APs or Bluetooth beacons) to indicate areas where more data should be collected or more beacons should be installed [20]; however, no insight is provided about the expected location accuracy in problematic areas to assess the criticality of collecting more data and/or installing more beacons. The RMID (Radio Map Inherent Difficulty) parameter introduced in [44], attempts to quantify the 'difficulty' of a FM with respect to expected location error, i.e., tries to estimate how hard it would be to achieve good location accuracy with a given FM. However, this approach provides a single value for the quality of the entire FM that may span a large indoor environment, thus cannot deliver the fine-grain QoP information provided by our solution. Authors in [6] use the gradient of the FM to generate an expected location accuracy map of the environment assuming that fingerprint data are collected in every location. However, this may affect the applicability of this approach in the case that real data are sparsely collected.

To obtain a strict theoretical estimate, works such as [15, 19, 47] employ the CRLB, which provides a lower bound on the variance of estimators. In particular, [15] uses a CRLB to estimate the localization accuracy for several motion models and wireless sensor measurements (e.g., signal strength, time of arrival), and investigate whether particular policy requirements can be met; [19] uses CRLB to estimate the accuracy achieved by the *Signal Strength Difference* (SSD) on a signal strength propagation model. Likewise, [47] uses CRLB to estimate the localization accuracy for wireless data and optimize AP placement. Nevertheless, in all the above techniques the predicted accuracy depends on the particular measurement model. This fact constrains their generality and consequently limits their applicability if the model is unknown, as in the case of ambient magnetic field data. The first work to propose fingerprint-based localization using Gaussian Process Regression was [29]; however, it does not provide estimates on the accuracy of localization with CRLB as we do in this paper, but rather on the parameters of GPR. Authors in [45] propose an analytical model based on *proximity graphs* to determine the probability of correct localization in fingerprint-based

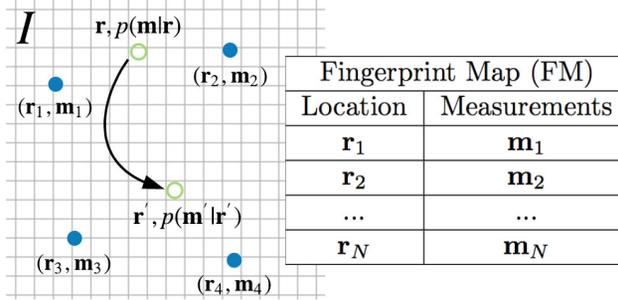


Fig. 2. System Model: (i) a user (empty circles) moving in an area with probability $p(\mathbf{m}|\mathbf{r})$ to measure \mathbf{m} at \mathbf{r} ; (ii) Fingerprint Map FM .

localization systems. However, this model is not geared towards predicting or bounding the actual localization error.

3 THE ACCES⁺ ARCHITECTURE

In this section we present ACCES⁺, which is a novel architecture for offline quality assessment of fingerprint maps comprising two separate quality assessment operators, i.e., the ACCES operator for magnetic data and the ACCES-w operator for Wi-Fi data. First, we formalize our system model assumptions, upon which a problem definition is provided. The main notation we use is presented in Table 1. First, we describe the internal calculation steps of ACCES⁺ pertaining to *fingerprint prediction* and *accuracy estimation*. Next, we discuss how these are applied to magnetic data for the ACCES operator and to Wi-Fi data for the ACCES-w operator.

3.1 System Model and Problem Formulation

We assume an indoor area I , with the d_r -dimensional location coordinates in this area denoted as $\mathbf{r} \in \mathbb{R}^{d_r}$. A fingerprint map FM of some data source for this area is a set of *fingerprints*, each represented as a pair of vectors: (i) a measurement vector (e.g., containing Wi-Fi RSS values or magnetometer readings) and (ii) an associated position vector where the measurement vector was collected, i.e., $FM = \{(\mathbf{r}_i, \mathbf{m}_i) : i = 1, \dots, N, \mathbf{r}_i \in \mathbb{R}^{d_r}, \mathbf{m}_i \in \mathbb{R}^{d_m}\}$, where N is the size of the FM and \mathbf{m}_i is the d_m -dimensional vector of measurements at location \mathbf{r}_i . Figure 2 depicts the system model including a user moving in area I with probability $p(\mathbf{m}|\mathbf{r})$ to observe measurement \mathbf{m} at location \mathbf{r} and a simple example of the FM .

We formulate the task of offline accuracy estimation for fingerprint-based localization in two steps: (i) given a fingerprint map FM , find the likelihood $p(\mathbf{m}|\mathbf{r})$ that \mathbf{m} will be measured at location \mathbf{r} ; (ii) given the likelihood $p(\mathbf{m}|\mathbf{r})$ find the smallest possible achievable *Root Mean Square Error* (RMSE) of the estimated position at some location \mathbf{r} . Such a bound on the RMSE at an arbitrary location can be used as a *navigability score*, which represents the associated QoP.

3.2 Outline of the Solution

To compute the navigability score for each indoor location given arbitrary fingerprint maps we consider two steps:

- **Fingerprint Prediction:** Given a set of fingerprints we interpolate them to any location on the floor plan map. The data interpolation method in our fingerprint prediction approach is based on the powerful GPR technique [42]. According to this technique, the interpolation output comprises the predicted values and their uncertainty in the form of a Gaussian distribution's

Table 1. Used Notation

Notation	Description
I	Indoor space
FM, N	Fingerprint Map and its size
L, M	Set of locations of interest and its size
\mathbf{r}, d_r	Coordinate-vector and its dimensionality
\mathbf{m}, d_m	Measurement vector and its dimensionality
θ, d_θ	Estimated parameter and its dimensionality
\mathbf{x}, d_x	Observed random-vector and its dimensionality
$p(\mathbf{x} \theta)$	Likelihood of observing \mathbf{x} given θ
$I(\theta)$	Fisher Information Matrix
$\mathbb{E}(\cdot), \mathbb{D}(\cdot)$	Expectation and Dispersion of a random variable
$ \cdot , (\cdot)^{-1}, tr(\cdot)$	Determinant, inverse and trace of a square matrix
$cov(\cdot)$	Covariance matrix of a random vector

mean and variance, respectively, while the Gaussian is a typical choice for the underlying distribution. This step allows to construct the picture of the possible signal (e.g., magnetic or Wi-Fi) distribution over the whole floor plan map along with the confidence in such distribution.

- **Accuracy Estimation:** Given such an interpolation of the fingerprints, we derive a bound on the achievable location accuracy, which we set as our ACCES⁺ navigability score. This step helps in finding the best accuracy (or equivalently the lowest localization error) that can be possibly achieved at each location given the confidence in the constructed signal distribution.

The first step allows us to perform FM interpolation disregarding the data source to obtain the possible data distribution over the indoor area along with a confidence in it, which capture the uncertainty of the collected measurements, spatial sparsity of the fingerprints and the spatial smoothness of the underlying true data. This approach is similar to the interpolation proposed in [29] and [2]. With the second step, a bound on the smallest possibly achievable localization error is calculated, given the knowledge and confidence about the collected and interpolated measurements. Similarly to [47], we assume that this bound reflects the localization error that a user will experience in practice when employing the location service. In the analysis of location systems, the CRLB is a bound on the localization error variance given the conditions in the target area including the number of signal sources and the noise in the signal measurements. This puts a bound on the accuracy that a user is expected to experience at a specific location. This should not be interpreted that accuracy will be X meters, but rather that accuracy cannot be better than X meters at that location. Yet, as it will be discussed in Section 3.4, this bound is a powerful tool to assess crowdsourced fingerprint maps and the adequacy of the underlying signal sources. In particular, the CRLB is employed to compute a navigability score that reflects the expected location accuracy as a qualitative metric (e.g., the relative location accuracy in an area of the indoor space will be better compared to another area where true data are sparsely collected), rather than as a quantitative metric (e.g., the absolute location accuracy in an area will be X meters as opposed to Y meters in the area with lower data density).

Algorithm 1 presents an outline of ACCES⁺ calculation at a set of locations given a FM . This algorithm implements the Fingerprint Prediction step based on the GPR technique and the Accuracy Estimation step based on the CRLB tool. In line 1, the semantics of the applied domain (e.g., magnetic, Wi-Fi, UWB, Bluetooth, etc.) are initialized. Consider for instance the Wi-Fi domain, where signals are associated with Wi-Fi APs that are identified by a unique MAC Address. Measurements are then compilations of these MAC addresses along with the corresponding RSS readings. On the other hand, magnetic data contains the three components of the earth's magnetic field without

Algorithm 1 - $ACCES^+(FM, L)$

Require: Fingerprint Map $FM = \{(\mathbf{r}_i, \mathbf{m}_i), i = 1, \dots, N\}$, locations of interest $L = \{\mathbf{r}_j, j = 1, \dots, M\}$.

- 1: InitializeDomain() ▷ Initialize the semantics of the applied domain
- 2: $score_j = 0, j = 1, \dots, M$ ▷ Navigability scores to return
- 3: $predictors = FingerprintPrediction(FM)$ ▷ Predictors for each component of a measurement vector
- 4: **for all** $\mathbf{r}_j \in L$ **do**
- 5: $score_j \leftarrow AccuracyEstimation(\mathbf{r}_j, predictors)$
- 6: **end for**
- 7: **return** $score$

any absolute reference points. As such, predictive frameworks like $ACCES^+$ are expected to be applicable to such highly diverse application domains. In line 2, the navigability scores for the locations in the area of interest are initialized to 0. In line 3, fingerprint prediction takes place for the locations that have no real measured data using the collected data points in the FM . To this end, the predictors for each component in measurement vector (i.e., fingerprint) are computed in the form of a Gaussian distribution's mean and variance; see Section 3.3 and Algorithm 2 for details. In lines 4-6, the predictors are employed for accuracy estimation using the CRLB tool to compute a score for each location; see Section 3.4 and Algorithm 3 for details. Finally, in line 7 the navigability scores are returned to visualize the quality of position across the area of interest. The topic of domain specialization is revisited in the next Section 4, where we describe the specifics for adapting the $ACCES^+$ architecture for both magnetic and Wi-Fi data. These adaptations are thoroughly assessed in the experimental evaluation of Section 6.

3.3 Fingerprint Prediction with GPR

The key objective in fingerprint prediction is to construct a picture of location dependent signals (e.g., magnetic intensity or Wi-Fi signal strength) in an indoor space, i.e., a signal map that reflects the signal variation at different locations that span the area of interest. The main idea is to leverage real signal measurements collected through crowdsourcing at specific locations and interpolate them to predict the multi-variate signal observations (i.e., fingerprints) at other locations, where no real data is available. In this way, the resulting FM , which contains the fingerprints and associated location coordinates, has the required granularity in the physical space to offer reliable accuracy estimation in the entire area by means of the proposed CRLB-based method described in Section 3.4.

Due to the inevitable influence of noise to measuring systems and the difference in sensors across devices, fingerprints should not be considered as deterministic, but rather as probabilistic entities. Thus, fingerprint values should be defined with a *probability distribution* in the form of a *likelihood function*, which shows how probable is to measure a particular value at a location. The likelihood function is usually constructed by considering the measured fingerprint values \mathbf{m}_i (along with the noise) and *interpolating* them to any arbitrary location with coordinates \mathbf{r} based on the distances $\|\mathbf{r} - \mathbf{r}_i\|$. In our approach we use the *predictive distribution* of a GPR estimator to define the likelihood function.

Algorithm Description: One can think of a fingerprint as a sample from a noisy function f of some physical quantity (e.g., Earth's magnetic field or Wi-Fi electromagnetic signal), where noise is assumed to be Gaussian. In order to estimate the underlying function, one can apply a *regression* technique to the measurement data, so as to obtain a *predictor* that can be used to estimate unknown values at arbitrary locations. Remarkably, with GPR, such a predictor also provides, along with

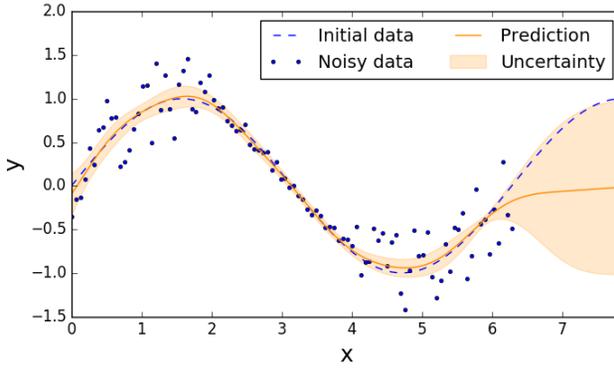


Fig. 3. Example application of *GPR* on the *sin* function. Noisy samples with smaller noise for $2 \leq x \leq 4$ and no data for $x \geq 6$, *GPR* prediction and its uncertainty. In the interval of smaller noise, uncertainty decreases, whereas in the region of data unavailability it increases dramatically.

the predicted value, an uncertainty estimate per se as the Gaussian distribution’s variance. This uncertainty estimate captures: (i) the spatial sparsity of the measurements (i.e., the more sparse measurements are, the larger is the uncertainty); (ii) the variability of the measurements (i.e., less complex functions are predicted more accurately); and (iii) the noise in measurements (i.e., larger noise leads to worse accuracy).

Given a map of scalar fingerprints *FM* where measurements consist of single values (i.e., $d_m = 1$), such as magnetic field magnitude, the output of the *GPR* algorithm is given by:

$$GPRScalar(FM) \rightarrow f(\mathbf{r}) \sim \mathcal{N}(\mu(\mathbf{r}), \sigma^2(\mathbf{r})), \quad (1)$$

where \mathcal{N} denotes a Gaussian random variable, and $\mu(\mathbf{r})$, $\sigma(\mathbf{r})^2$ denote its mean and variance, respectively. The mean and variance of the Gaussian random variable in eq. 1 come from the predictive distribution of the *GPR* estimator, while learning is achieved through standard techniques; see [42] for techniques applicable to *GPR*. The likelihood function value at \mathbf{r} is given by the *probability density* of the *GPR*’s predictive distribution, which is Gaussian.

Figure 3 depicts the output of *GPR* applied on noisy scalar sinusoidal data with the uncertainty shown as the standard deviation values. In this example application of *GPR*, the noise-free *Initial data* are generated by the underlying sinusoidal function and are actually not available. Only the *Noisy data*, i.e., initial data points perturbed by variable random noise, can be observed and measured. The objective is that the *Prediction* curve produced by *GPR* is as close as possible to the *Initial data*. When $x \geq 6$, the *Uncertainty* (shaded area) increases significantly because there are no measured data to help in the prediction. In this case, the *Prediction* curve can merely follow the trend of the previously measured data and therefore deviates considerably from the *Initial data*.

Multidimensional data: As *GPR* cannot be directly applied to multidimensional (or vector) data (i.e., $d_m \geq 2$), but only to scalar data (i.e., $d_m = 1$), we need to assume that measurement channels are independent, i.e., magnetic field components vary independently over an area, or separate Wi-Fi APs do not interfere with each other [23]. Thus, *GPR* can be applied separately to each component $k = 1, \dots, d_m$ of fingerprint measurements to obtain d_m predictors, each allowing to estimate the parameters of a Gaussian distribution for one component of a vector-function $f(\mathbf{r})$ separately (see Algorithm 2). Combining these parameters into a single vector, one obtains the following *GPR* output

from a map FM of vector fingerprints:

$$GPR(FM) \rightarrow \mathbf{f}(\mathbf{r}) \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{r}), \text{diag}(\boldsymbol{\sigma}(\mathbf{r}))^2), \quad (2)$$

where $\text{diag}(\mathbf{x})$ denotes a diagonal matrix with a diagonal equal to the vector \mathbf{x} and \mathcal{N} denotes a Gaussian random vector. The likelihood function value at \mathbf{r} is the respective probability density function of the *multivariate* Gaussian distribution.

Choosing parameters: Although GPR is a non-parametric technique and can be used as a black box, some parameters, dependent on the origin of the fingerprints and localization scenarios, must be selected in order to obtain sensible error estimate values.

First, we should choose a *Kernel Function* or a *kernel*, which, in the case of GPR, denotes how spatially near measurements are *correlated* or, in other words, influenced by each other. We opt for the popular *Radial Basis Function* (RBF) kernel, which has the following form:

$$K^{RBF}(\mathbf{r}_i, \mathbf{r}_j) = \exp\left(-\frac{\|\mathbf{r}_i - \mathbf{r}_j\|^2}{2\ell^2}\right), \quad \mathbf{r}_i, \mathbf{r}_j \in \mathbb{R}^{d_r}, \quad (3)$$

where ℓ is a scaling factor. RBF kernels are good to model arbitrary non-periodic and smooth functions, which is the case for physically-driven data, e.g., magnetic and electromagnetic.

Secondly, we should consider the uncertainty due to noise and limited sensor accuracy, which can be directly incorporated in the kernel as follows:

$$K(\mathbf{r}_i, \mathbf{r}_j) = K^{RBF}(\mathbf{r}_i, \mathbf{r}_j) + \sigma_k^2 \delta_{ijk}, \quad (4)$$

where σ_k^2 is a noise level for measurement k , and δ_{ijk} is a multi-index Kronecker delta. Note that no more hyper-parameters are required, other than the Kernel hyper-parameters ℓ and σ_k that can be optimized using Maximum Likelihood Estimation (MLE) or its variations [21].

Reducing complexity: A drawback of GPR is its computational complexity. Regression takes $\mathcal{O}(N^3)$ operations, where N is the number of data samples. Therefore, the GPR calculation may be very costly for large fingerprint maps. However, such a calculation is rarely necessary to be carried out in an online manner as the FM changes slowly and in cases of change, the GPR calculation can be carried out in an offline manner. Clearly, there is a trade-off between the accuracy of the GPR model, which affects the subsequent location accuracy estimation with CRLB, and the efficiency of the proposed solution in terms of time to compute the predictors with GPR, especially when new collected data are contributed to the FM and the predictors need to be re-evaluated. To overcome this limitation, we propose to split the indoor area I into parts, with each part containing a moderate amount of measurements, i.e., split each building into floors, floors into rooms, corridors, etc. and calculate GPR only locally. This seems to be a natural solution, even though the fingerprints on each area edge affect each other. This can be justified by the fact that distant fingerprints are not interrelated. Moreover, the collected fingerprints contained within the walls of a room are more relevant for predicting the signal strength values at other locations inside the room, than fingerprints collected outside.

3.4 Accuracy Estimation with CRLB

The key objective in our accuracy estimation is to derive a qualitative indication of the relative position accuracy that a user is expected to experience at an area in an indoor space compared to another area. As this QoP measure can be computed offline, i.e., before the fingerprint-based localization service starts receiving location requests, it offers valuable knowledge to indoor location service operators for taking corrective actions towards improving the location accuracy where needed. The main idea is to exploit the real and predicted signal fingerprints in the FM and employ the CRLB

Algorithm 2 - *FingerprintPrediction(FM)***Require:** Fingerprint Map $FM = \{(\mathbf{r}_i, \mathbf{m}_i), i = 1, \dots, N\}$.

```

1:  $predictors_k = \emptyset, k = 1, \dots, d_m$   $\triangleright$  Predictors for each component of a measurement vector
2: for  $k = 1 \dots d_m$  do
3:    $FM_i^k = \emptyset, i = 1, \dots, N$   $\triangleright$  A sub-map of  $FM$ , containing  $k^{th}$  component of each fingerprint
4:   for all  $(\mathbf{r}_i, \mathbf{m}_i) \in FM$  do
5:      $FM_i^k \leftarrow (\mathbf{r}_i, \mathbf{m}_i^k)$   $\triangleright$  Constructing the sub-map
6:   end for
7:    $predictor = GPRScalar(FM^k)$   $\triangleright$  Obtaining a scalar GPR predictor from the sub-map
8:    $predictors_k \leftarrow predictor$ 
9: end for
10: return predictors

```

tool to estimate location accuracy solely on the characteristics of the signals in the FM , e.g., sparsity of real data at an area implies worse accuracy compared to another area where the collected data are more dense or higher signal variation suggests that the location accuracy will be better than the case that the signals are relatively stable.

In estimation theory, the CRLB indicates that the covariance matrix of the estimation error vector is greater than or equal (i.e., lower bound) than the inverse of the *Fisher Information Matrix* (FIM). In our context, the FIM contains information about the number of signal sources and associated noise in the signal measurements at a given location \mathbf{r} , while the estimator can be any fingerprinting algorithm that uses RSS measurements to estimate the unknown user location. The covariance matrix of the estimation error vector, is a 2-D matrix that conveniently reflects the expected localization error at a given location [15] and the CRLB can be viewed as a lower bound on the accuracy that a user will experience in practice. Therefore, the CRLB implies that accuracy at a specific location cannot be better than the computed bound. As discussed in Section 3.2, due to the assumptions in CRLB computation, this bound should be interpreted as a qualitative (i.e., relative accuracy), rather than a quantitative (i.e., absolute accuracy) metric. In particular, the CRLB becomes larger (i.e., lower accuracy should be expected) when the number of collected data is smaller (i.e., lower fingerprint density), or fewer signal data sources are available, or the spatial gradient of the measurements is low at different locations or the noise disturbing the measurements is high.

For instance, the density of the fingerprint map is reflected in the CRLB and consequently the proposed ACCES score. For the GPR model, as shown in Figure 3 (right-hand side), the predictive variance becomes equal to the prior distribution, thus reflecting the navigability score as desired, i.e., i) the predictive mean converges to a constant value asymptotically and ii) the predictive variance increases and then converges to constant value asymptotically. This means that the predictive distribution tends to become less and less indistinguishable with each small step while moving to the right due to the missing data, so our ACCES navigability score will decrease as expected (seen from equations 5 and 6).

Given the interpolated fingerprint values obtained by GPR, we can derive the CRLB on the best possibly achievable accuracy of any localization algorithm. We emphasize that our interest is neither in the CRLB's absolute value, nor in the possibility of some localization algorithms outperforming this bound; we are rather curious about the CRLB's behavior that reflects the relative performance of an arbitrary localization algorithm in different parts of the building or even the same floor. Algorithm 3 provides an outline of the CRLB calculation for accuracy estimation.

Algorithm 3 - *AccuracyEstimation*(\mathbf{r} , *predictors*)**Require:** Location \mathbf{r} , fingerprint *predictors*

- 1: $\mathcal{I} = \text{evalFIM}(\mathbf{r}, \text{predictors})$ ▷ Calculate Fisher Information Matrix as in eq. 12
- 2: $\text{crlb} = \text{sqrt}(\text{tr}(\text{inv}(\mathcal{I})))$
- 3: **return** *crlb*

Algorithm Description: The CRLB computation requires the intermediate calculation (line 1 of Algorithm 3) of the FIM, which shows how much information an observable random variable (or vector) carries about some deterministic parameter it depends on, e.g., how much sensor measurements at an unknown location could shed a light on the location coordinates.

Now, we discuss an analytical representation of the FIM, inferred from the GPR interpolation, with the derivations available in the appendix. Let $\boldsymbol{\theta} \in \mathbb{R}^{d_\theta}$ denote a vector-parameter that is being estimated and $\mathbf{x} \in \mathbb{R}^{d_x}$ denote a random vector, the distribution of which is dependent on $\boldsymbol{\theta}$. Then, the FIM is

$$\mathcal{I}(\boldsymbol{\theta}) = [\mathcal{I}(\boldsymbol{\theta})_{ij}]_{i,j=1}^{d_\theta, d_\theta} = -\mathbb{E} \left(\frac{\partial^2 \log p(\mathbf{x}|\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \right), \quad (5)$$

where θ_i denotes the i -th component of $\boldsymbol{\theta}$, $\mathcal{I}(\boldsymbol{\theta})_{ij}$ denotes an element of the matrix $\mathcal{I}(\boldsymbol{\theta})$ in the i -th row and the j -th column, $p(\mathbf{x}|\boldsymbol{\theta})$ is a likelihood showing the probability of sampling \mathbf{x} from its distribution given the parameters $\boldsymbol{\theta}$, and mathematical expectation \mathbb{E} is taken over \mathbf{x} . In the case of GPR, the likelihood is given by the probability density function of a Gaussian distribution.

The CRLB is given by the following bound on the RMSE of the location estimation

$$\text{RMSE} \geq \sqrt{\text{tr}(\mathcal{I}^{-1}(\boldsymbol{\theta}))} = \text{crlb}, \quad (6)$$

where we let *crlb* to be such bound that represents the ACCES⁺ navigability score. Low navigability score means high QoP and vice versa.

Numerical Calculation: As it is shown in eq. 12 in the appendix, computation of the FIM involves evaluation of a *Hessian* matrix, which raises a concern about its computational overhead. One approach is to derive it analytically from the expression for the predictive distribution of GPR. The main drawback of this approach is that the analytical form depends on the particular choice of kernel and is therefore not generalizable. Another approach is a numerical approximation based on finite difference schemes. With this approach, it is possible to evaluate a Hessian matrix for any values of $\boldsymbol{\theta}$ with sufficient precision. We stick to the latter approach, and for the numerical calculation we utilize the *Numdifftools* package for *Python*, assuming the conditions for differentiability are satisfied. Last, we let $\boldsymbol{\theta} \triangleq \mathbf{r}$ and $\mathbf{x} \triangleq \mathbf{m}$ and, thus, obtain the desired CRLB for fingerprint-based localization accuracy estimation. Specifically, the corresponding probability density functions $p(\mathbf{m}|\mathbf{r})$ are required to compute the CRLB and these are always available through the GPR technique in the Fingerprint Prediction step, which interpolates the collected data in the *FM* to output a Gaussian probability density function (i.e., a predictor with a mean and a variance computed from the collected data as the parameters). These Gaussian distributions (predictors) are then used for the numerical calculation of the Hessian matrix that is required for the computation of the FIM to obtain the CRLB for location accuracy estimation.

4 DOMAIN-SPECIFIC OPERATORS

In this section we describe how the ACCES⁺ architecture incorporates two separate domain-specific operators, namely the ACCES operator for magnetic data and the ACCES-w operator for Wi-Fi data.

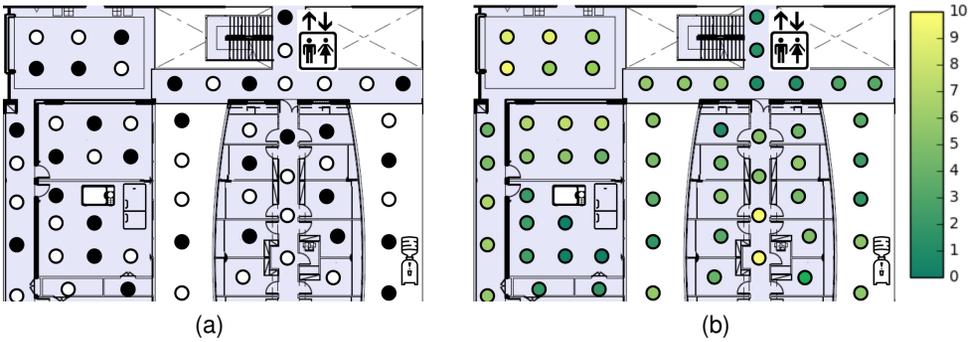


Fig. 4. Example application of ACCES⁺ on magnetic data. (a) Black circles denote locations of the collected magnetic fingerprints, white circles denote locations, where measurements are not collected, but predicted using GPR; (b) navigability scores calculated at a set of locations inside indoor environment with the fingerprint data comprising magnetometer readings.

Figure 4 exemplifies the application of the proposed ACCES⁺ architecture in the case of magnetic data where the ACCES operator is invoked. An indoor environment with the ambient magnetic field serving as the source of fingerprint data is assumed in Figure 4a. Black circles represent locations that constitute a FM collected during offline phase. White circles represent locations, where magnetometer measurements were not collected, but are predicted using GPR to obtain a picture of how the magnetic field could be distributed over the whole indoor area. Figure 4b depicts the navigability scores calculated for a set of locations (of both collected and predicted measurements) inside an indoor environment during the offline phase. Lower navigability scores (dark green circles) values imply higher QoP, i.e., better accuracy during the operation of the location service, whereas higher scores (light yellow circles) imply lower QoP, i.e., the user will experience worse accuracy. We make the following observations: (i) in the room with the microwave and fridge additional magnetic perturbations occur, and accuracy is expected to be higher than in another area; (ii) in other parts of the building close to the escalator, elevator and water dispenser, where magnetic disturbances occur, accuracy is expected to be relatively good; and (iii) far from magnetic disturbance, like in the top left part of the building, accuracy is expected to be worse, which also applies to locations across the corridors where magnetic field does not vary significantly.

4.1 The ACCES Operator for Magnetic Data

Typically, the magnetic signal is quite stable and does not vary significantly in the geographical area covered even by large buildings. However, the magnetic signal can be disturbed indoors due to various sources including metal surfaces, power cables, electrical appliances, operating equipment, elevators, escalators, etc. Such disturbances are prevalent in modern indoor environments, which make it possible to distinguish between distant locations by using magnetic fingerprints. For the ACCES operator, first we must select the values of ℓ and σ_k to be used for fingerprint prediction in eq. 3 and eq. 4, respectively. We may pick ℓ as the smallest achievable motion magnitude in the particular localization setting, which could be the smallest wheeled robot movement, or human's foot size; the noise level σ_k can be expressed as the sensitivity of the sensors, which is 150 to 600 nT for typical smartphone magnetometer readings. However, such sensitivity may underestimate the true noise level and, thus, it is preferable to obtain the noise level from empirical data, e.g., collect multiple measurements for each location \mathbf{r} and calculate the standard deviation. Besides, it is not

necessary to infer parameter values directly; they can rather be derived using a *parameter estimation* technique to get the best approximation of the data with GPR. Thus, one may consider choosing ℓ and σ_k manually when the fingerprint map is sparse and some prior information is available, or automatically using parameter estimation methods when the map is dense. In our experiments with magnetic data, we employed the automatic optimization approach because of the high density of the magnetic fingerprints (see Figure 1a). In particular, we provide reasonable bounds for the parameters and the outputs of the optimization are optimal parameter values (within those bounds) for different *FM*s.

4.2 The ACCES-w Operator for Wi-Fi Data

As part of the functionality in Wi-Fi networks (i.e., IEEE 802.11 standard) the Wi-Fi APs transmit periodically beacon packets that contain hardware specific information (e.g., MAC address of the AP) and connection options (e.g., encryption protocol). These packets are received by Wi-Fi enabled mobile devices that scan sequentially the frequency channels to discover available Wi-Fi APs. Upon reception, the RSS value pertaining to each AP is measured by the mobile device as an indicator of the signal quality to the corresponding AP. The resolution of the RSS readings reported by modern smartphones is in the range $[0.5 \dots 1.0]$ dBm.

Typically, in open spaces the RSS values follow a log-distance attenuation model and become weaker as the distance from the Wi-Fi AP increases. In indoor environments, the presence of walls, doors, equipment, furniture, and people moving introduce additional signal attenuation. Moreover, due to the multi-path signal propagation (e.g., reflections on walls) the RSS values at a specific location are not stable and may fluctuate significantly. This behavior is demonstrated in Figure 5, where the histograms of the RSS samples collected at two distinct locations from the same Wi-Fi AP over a period of several minutes are plotted. These histograms suggest that when enough data is collected at a single point in space, then the σ_k parameter could be set equal to the sample standard deviation in the RSS histogram. However, in crowdsourced location systems the crowdsourcers rarely stand still at the same location for enough time to collect multiple measurements (Wi-Fi readings are reported by modern smartphone every few seconds) and typically record fingerprint data while walking. Therefore, it is not feasible to build reliable RSS histograms. In such a case, the σ_k parameter is selected manually and setting it equal to the sensitivity of the Wi-Fi adapter, i.e., $\sigma_k = 1.0$ dBm in eq. 4, works well in practice. This is because the RSS values pertaining to the same Wi-Fi AP are blended with the respective RSS values from neighboring locations with the selection of the ℓ parameter in eq. 3, as described next.

Wi-Fi RSS and its inherent noise characteristics do allow the determination of transmitter-receiver distance with uncertainty, which leads to room-level localization accuracy. It is important to notice that this level of accuracy is in many application scenarios [36] (e.g., smart factories, hospitals, and ships) adequate and useful to a wide range of applications. In case additional location-related measurements (e.g., timing, angle of arrival) are available together with RSS data, then the CRLB can be properly adapted to consider all data sources, thus improving the accuracy estimation. We observe that the RSS values at location *A* with distance d_A from the Wi-Fi AP are in the range $[-58 \dots -80]$ dBm, fluctuate around the mean value -69 dBm, and can be approximated by a Gaussian distribution, as shown in Figure 5a. Similar behavior is observed at a nearby location *B* with distance $d_B > d_A$ from the Wi-Fi AP, where the RSS values are in the range $[-60 \dots -93]$ dBm and the mean value of the Gaussian bell (-80 dBm) is shifted to the right, i.e., the signal is weaker; see Figure 5b. Note that the same behavior could occur if locations *A* and *B* were at the same distance from the Wi-Fi AP, but location *B* was behind a wall or obstacle. Fingerprint matching algorithms are able to capture this behavior and exploit the RSS values from multiple Wi-Fi APs that are readily available on commercial smartphones to distinguish indoor locations. Note that the signals at a location in a new

environment most likely will not match any of the Gaussian distributions shown in Figure 5. In a new environment, however, a different *FM* will be available (e.g., through data crowdsourcing) and the associated RSS data can be used to obtain the corresponding signal distributions at each location in the new environment. In case the RSS histogram can be better approximated by a different distribution, then the associated probability density function (e.g., skewed, multimodal, non-Gaussian, etc.) can be used with a different modelling technique [52], instead of the GPR that relies on the normality assumption, to improve the fingerprint prediction accuracy. In this way, the proposed method for accuracy estimation, as well as any fingerprint-based localization algorithm, are applicable to a different indoor area.

Obviously, there is an overlap in the two histograms that reflects the spatial correlation of the RSS values at nearby locations. The closer two locations are in the physical space (assuming no wall or obstacle is between them) the larger the spatial correlation of RSS values and vice versa. Two locations that are far from each other (e.g., tens of meters) are not expected to have highly correlated RSS values. In addition, two neighboring locations that are within 2-3 meters but are separated by a wall or obstacle, which causes a large shift in the mean RSS value, will not have highly correlated RSS values. In this sense, there is a natural way to select the ℓ parameter of the exponential kernel in eq. 3. In large open spaces where the Wi-Fi signal decays smoothly as the distance from an AP increases without any abrupt changes, ℓ can be set in the range [5 . . . 10]meter; whereas for locations that are far from each other or are separated by walls or obstacles smaller values should be used to impose low spatial correlation in the RSS values. Note that when fingerprint prediction with GPR is performed separately for separate rooms, as discussed in Section 3.3 (Reducing complexity), then it is not necessary to select a small value for ℓ to consider the impact of walls to the correlation of RSS values that lie at different sides of a wall. However, when obstacles are present, then this is recommended because the RSS values in the fingerprints, which are separated by an obstacle that lies inside a room, will be less correlated (due to the additional signal attenuation); yet, these fingerprints will be considered together for spatial regression even if rooms are treated separately.

In practice, for simplicity a single value can be used for the ℓ parameter. Similar to magnetic data, ℓ can be selected automatically using parameter estimation algorithms when the Wi-Fi *FM* is dense. Alternatively, ℓ can be selected empirically when the Wi-Fi *FM* is sparse, as in most practical application scenarios. We selected $\ell = 3$ meters that yields adequate fingerprint estimation performance in our experimental evaluation.

5 PROTOTYPE IMPLEMENTATION

5.1 Anyplace Service Architecture & Software Stack

Our Anyplace system [57] follows a service-oriented architecture design that allows to plug-n-play additional modules, either for extending system capabilities - by implementing new features - or for enhancing user-experience, by improving existing functionalities (e.g., map-matching and sophisticated data fusion to increase localization accuracy). The public Anyplace service has to this date supported more than 250K real user interactions, with many more users using its standalone installations.

The Anyplace native Android application is composed of the *Navigator* and the *Logger* that can benefit from Wi-Fi fingerprinting [35, 40] available under this platform. The *Logger* application enables users to record Wi-Fi readings from nearby Wi-Fi APs and upload them to our *Server* through a Web 2.0 API (in JSON). It is used by volunteers for contributing Wi-Fi data and for crowdsourcing the *FMs* of buildings [10] (i.e., four directional fingerprinting in multiple rounds to remove noise). The *Navigator* allows users to see their current location on top of the floorplan map and navigate between Points-of-Interest (POI) inside the building with high accuracy (i.e., 1.96 *meter* at the

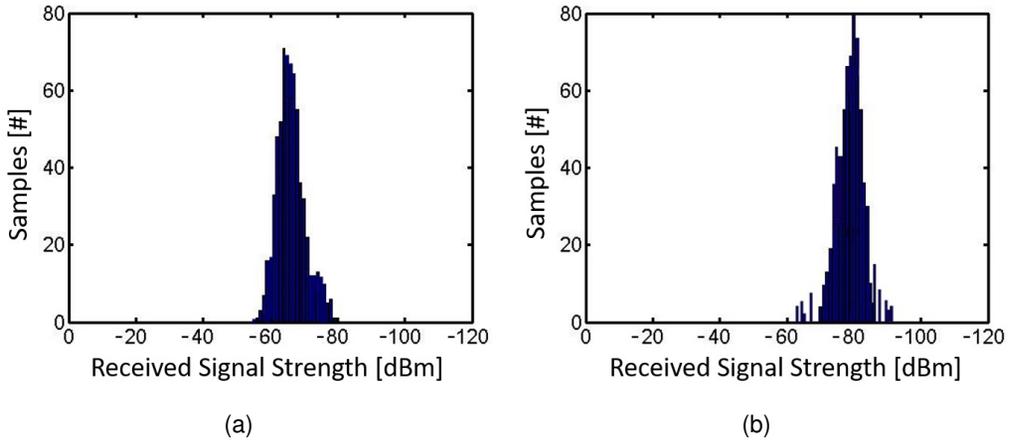


Fig. 5. Histograms of Wi-Fi RSS values pertaining to the same Wi-Fi AP and two distinct locations: (a) Location A at distance d_A from the Wi-Fi AP; (b) Location B at distance $d_B > d_A$ from the Wi-Fi AP.

Microsoft Indoor Localization Competition at ACM/IEEE IPSN'14 [35]). When the *Navigator* is launched, the building map and the associated POIs are automatically loaded by using the rough user location provided by the Google Geolocation API (see Figure 6). Then, the application downloads the *FM* of the relevant floor (subsequently the complete building) and displays the user location on top of the map. Moreover, users may search for POIs and get navigation directions from their current location. The *Navigator* also uses the on-board smartphone sensors (i.e., accelerometer, gyroscope and digital compass), which are seamlessly integrated in our tracking module to smooth the Wi-Fi locations and enhance the navigation experience.

The localization function $loc()$ of Anyplace comprises of the following phases: in the first offline phase, it records the Wi-Fi RSS fingerprints at certain locations (x,y) pin-pointed on a building floor map (e.g., every few meters). In the second offline phase, the Wi-Fi fingerprints are merged into the *FM*, which is essentially a $N \times M$ matrix where N is the number of unique fingerprints and M the total number of APs. Finally, in the online phase, a user can compare its currently observed RSS fingerprint against the *FM* to find the best match, either on the server side or in-situ at the smartphone device after downloading the whole *FM*, by using known algorithms such as KNN or WKNN [33].

The KNN and WKNN fingerprint methods used in Anyplace have been reported to achieve higher accuracy in real-life deployments than the standard NN fingerprint matching method, which returns the location coordinates associated with the fingerprint in the *FM* that has the lowest distance from the observed fingerprint during localization. This is due to the fact that the noise in the signal values (e.g., the fluctuations in the RSS values at the same location shown in Figure 5) can cause the NN method to return another location with a similar fingerprint that can be several meters away. In contrast, the KNN and WKNN methods consider K locations with the lowest distances in the signal space and are able to smooth out the effect of noise to some extent and thus achieve higher localization accuracy. In addition, fingerprint methods have several advantages compared to traditional positioning methods, i.e., *pure modeling* methods [15, 55], described in Section 2.1. These methods require a priori system information including the positions of the Wi-Fi APs or Bluetooth beacons and a signal propagation model that translates each signal reading (e.g., timing, angle, or signal strength) to a distance measurement from the associated transmitter, while the user location is estimated at the intersection or overlapping area of the circles defined by multiple distance measurements. However,

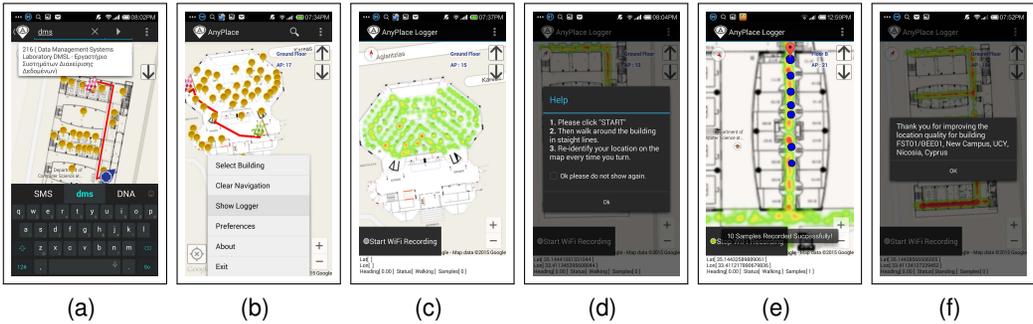


Fig. 6. **Anyplace Logger and Navigator.** This is a native Android application that enables fine-grain indoor localization (up to 1.96m accuracy [35]) through the use of RSS fingerprints contributed by the crowd.

in indoor application scenarios the transmitter locations may not be available or are hard to obtain (e.g., a large mall typically has its Wi-Fi infrastructure and individual stores deploy additional Wi-Fi APs without keeping any record of the installation location). Moreover, it is very challenging to derive an accurate signal propagation model for indoor environments where signal reflections and multi-path propagation, as opposed to line-of-sight channels, prevail due to walls and obstacles. This severely degrades the accuracy of traditional positioning methods indoors. In contrast, fingerprint methods such as those employed in Anyplace, do not require knowledge of the transmitter locations, while they can model reliably data sources such as ambient magnetic and light intensity that are not emitted by specific transmitters and cannot be modeled by traditional positioning methods. In addition, signal variations due to reflections and multi-path conditions are captured in the data collected and stored in the *FM*; thus fingerprint methods employ powerful data-driven pattern matching techniques to estimate user location and consequently outperform traditional positioning methods in terms of location accuracy.

5.2 Fingerprint Management Studio and ACCES⁺ integration

The FMS is a signal management studio fully integrated in Anyplace, which provides a spatio-temporal platform to manage the collection of location-dependent sensor readings (i.e., fingerprints) in indoor environments, estimate the expected localization accuracy based on the collected fingerprints, and assess Wi-Fi coverage and data rates. FMS comprises the following components: (i) *CSM (Crowd Signal Map)*, which is a map-based visual management environment to orchestrate the crowdsourcing effort of indoor signals for ethical benefit; (ii) *ACCES-w (Accuracy Estimation for Wi-Fi data)*, which enables the qualitative assessment of location accuracy before deploying the localization service; and (iii) *WS (Wi-Fi Surveying)*, which enables the qualitative assessment of Wi-Fi coverage using the signals collected by crowdsourcers. In the following, we outline the technical details of the CSM and WS components and describe how they interplay with the newly developed ACCES-w component.

CSM (Crowd Signal Map): is a map-based visual management environment to orchestrate the crowdsourcing effort of indoor signals. We provide multi-granular visual analytic structures (heatmap and coverage maps) that enable the management of data in space and time exposing both performance and high resolution. Figure 7 shows one such example where the stencil bar on the left provides quick access to all functionality of FMS. Particularly, the user observes through the heatmap where fingerprints have been sampled (red shows high density). It also shows with purple squares where the

Wi-Fi signal strength is between -90dBm and -100dBm (i.e., intermittent connectivity). A histogram time selector at the bottom enables the user to focus on different time ranges.

ACCES-w (Accuracy Estimation for Wi-Fi data) is our novel operator for the offline QoP assessment based on Wi-Fi data collected at arbitrary locations, as discussed in Section 4.2.

WS (Wi-Fi Surveying): surveying formally refers to the process of planning and designing a wireless network so that Wi-Fi APs are positioned optimally, with respect to wireless coverage, data rates, network capacity, roaming capability and Quality of Service (QoS). Given that an AP network is already in place when an IIN architecture is deployed, the aim of FMS is to provide retrospective visual analytics to network architects for subsequent decision support, e.g., where to install new APs and stronger antennas or where to change the SSID broadcast channels to avoid collisions at the MAC layer. The WS component shows the reception quality of Wi-Fi in five classes, based on the RSS indicator of the MAC addresses. It also records in a database the prefixes (3-Byte to 5-Byte) of the 6-Byte MAC addresses for various manufacturers (we use Wireshark.org OUI database). AP inference is carried out with a weighted threshold algorithm: if all signals for an AP MAC address are stronger than a given threshold (e.g., -60dBm), we derive the AP location using the centroid of the AP signal strengths. Otherwise, the AP location is derived based on the strongest signal.

Notably, the ACCES-w component expands the FMS workflows by interplaying with CSM and WS to optimize certain scenarios with particular focus on improving the QoP of the localization system and consequently the overall indoor location experience. Indicative scenarios include the following novel FMS workflows where the system operator is able to:

- visualize the fingerprint density through the CSM fingerprint heatmap and compare it side-by-side with the expected accuracy through the ACCES-w heatmap to identify where additional effort is needed to collect more Wi-Fi data and direct crowdsourcers to sparsely-sampled areas.
- visualize the approximate Wi-Fi AP location and number of detected APs in each location through the WS component and easily correlate with the expected accuracy provided the ACCES-w heatmap to get hints on where to install additional APs for increasing the dimension of the Wi-Fi fingerprints and thus their ability to discriminate among distant locations.

6 PERFORMANCE EVALUATION

This section presents an experimental evaluation of the ACCES⁺ architecture.

6.1 Series-1: ACCES-w Real-Life Usage

In this series we evaluate the performance of the proposed ACCES⁺ framework and focus in particular to the assessment of the ACCES-w operator using real-life Wi-Fi data from the *Anyplace* indoor navigation service.

6.1.1 Experimental Wi-Fi Data. We consider the following Wi-Fi datasets that are readily available through the Anyplace platform.

University of Cyprus Campus: The CSUCY dataset was collected at the Department of Computer Science (CS), University of Cyprus. In particular, it consists of 5,000 reference fingerprints taken from ~266 Wi-Fi APs installed in the four floors of the CS and neighboring buildings. We collected our data by walking over a path that consists of ~2,900 locations. The density of the collected fingerprints at floor 0, which covers around 2,500 m^2 , is shown in Figure 8a.

Mall of Cyprus: The MALL dataset was collected at the Mall of Cyprus. In particular, it consists of 800 reference fingerprints taken from ~279 Wi-Fi APs installed in the two floors of the mall and



Fig. 7. The Fingerprints Management Studio (FMS) provides a spatio-temporal platform to: (i) manage the collection of location-dependent *fingerprints* in indoor environments; (ii) estimate the localization accuracy based on the collected fingerprints; and (iii) assess Wi-Fi coverage and data rates.

neighboring buildings. The fingerprint density at floor 1, which covers around $18,500\text{ m}^2$, is shown in Figure 8b.

6.1.2 Methodology. The objective of this experimental series is to assess qualitatively the performance of the ACCES-w operator under real-life conditions, thus demonstrating the benefits for the location service providers to improve the delivered localization accuracy. In particular, we perform two experiments namely

- *E1* – In this scenario, we selectively remove from the *FM* a subset of the Wi-Fi fingerprints collected in the corresponding locations inside a specific area of a building floor, which corresponds to the case where no data have been collected in that area.
- *E2* – In this scenario, we consider all the fingerprints in the *FM* (as opposed to experiment *E1*), but we tamper the fingerprints by removing the RSS measurements from a subset of randomly selected Wi-Fi APs, which corresponds to the case that the fingerprints in an area contain measurements only from a few Wi-Fi APs.

6.1.3 Experimental Evaluation. Next we present the details and discuss our findings with respect to the experiments *E1* and *E2*.

E1 – The goal of this scenario is to show the accuracy degradation in that area, due to the absence of data, through the heatmap visualization powered by ACCES-w. From the system operator’s viewpoint, this provides useful information on where to direct crowdsourcers for collecting data. We consider floor 0 of CSUCY dataset to demonstrate the behavior of the ACCES-w operator in case real data are not available for a specific section of the floor. We emulate this by manually removing from the original *FM* the subset of fingerprints collected at the corresponding locations in that section. Figure 9 portrays the density of the collected fingerprints as a heatmap produced by the CSM component overlaid with the accuracy heatmap computed by ACCES-w. The color bar

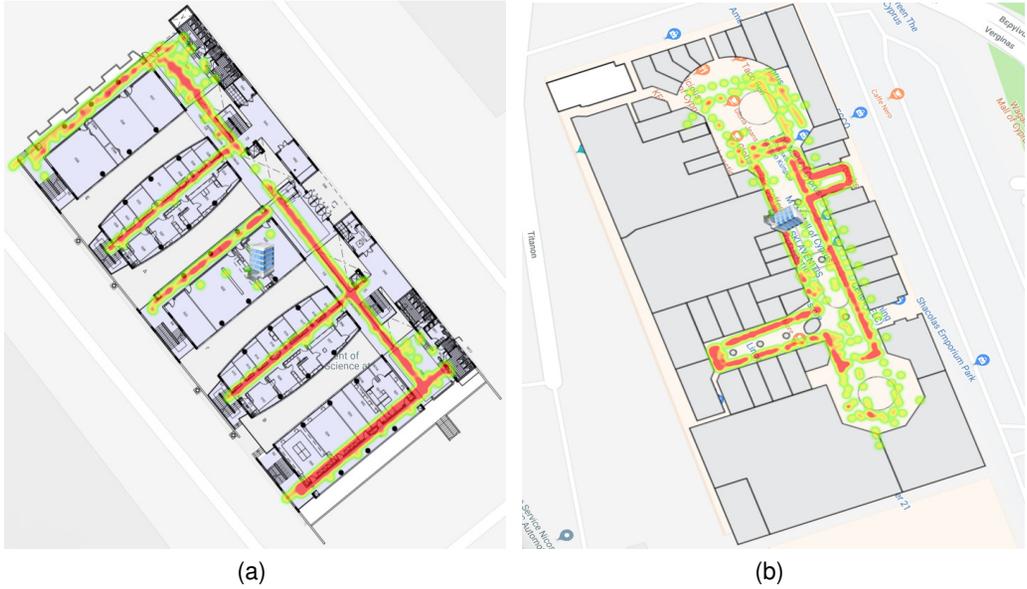


Fig. 8. *Wi-Fi datasets*. The heatmaps indicate the density of the collected crowdsourced signal strength fingerprints: (a) University of Cyprus Campus and (b) Mall of Cyprus.

indicates the QoP expected at various locations across the floor. In the case that all fingerprints in the original *FM* are considered, we can make the following interesting observations; see Figure 9a. First, in those locations where real data have been collected, the accuracy score indicates that the users will experience very high QoP (i.e., low location error). Second, in the locations where users have not crowdsourced real data but are surrounded by data collection points, high or moderate QoP will be delivered. This shows the ability of the GPR function within ACCES-w to reliably interpolate real data and produce artificial, yet realistic, fingerprint data.

Figure 9b depicts the location accuracy heatmap after removing the fingerprints collected across the bottom wing of the floor (marked with blue dashed line). The color of the locations in that part of the floor turn from dark or light green to orange/red indicating that significantly worse QoP (i.e., higher location error) will be experienced by users traversing that area. This is due to the absence of real data, as predicted by the ACCES-w operator. In this case, the system operator can direct volunteers to that area for collecting Wi-Fi data and eventually reaching the level of QoP shown in Figure 9a. Even though this might seem like an obvious conclusion for the operator when Wi-Fi data are missing, the proposed solution offers additional insights. Firstly, with the visualization feature the operator can quickly identify the indoor spaces (e.g., rooms, corridors, etc.) that suffer from missing or low-density data as the estimated QoP is conveniently displayed on top of the indoor floor plan map. In contrast, without the proposed solution, the operator can only see the locations of the collected data in the *FM* and it is not easy to tell in which area new data should be collected by simply looking at location coordinates with no underlying floor plan map. Secondly, when crowdsourcers are directed to the target area for collecting additional data, the QoP in that area will start improving as new data comes in and the operator can easily see when sufficient location accuracy is reached. Thus, unnecessary crowdsourcing effort in that area is avoided and crowdsourcers can be utilized more efficiently, e.g., direct them to other areas of low QoP.

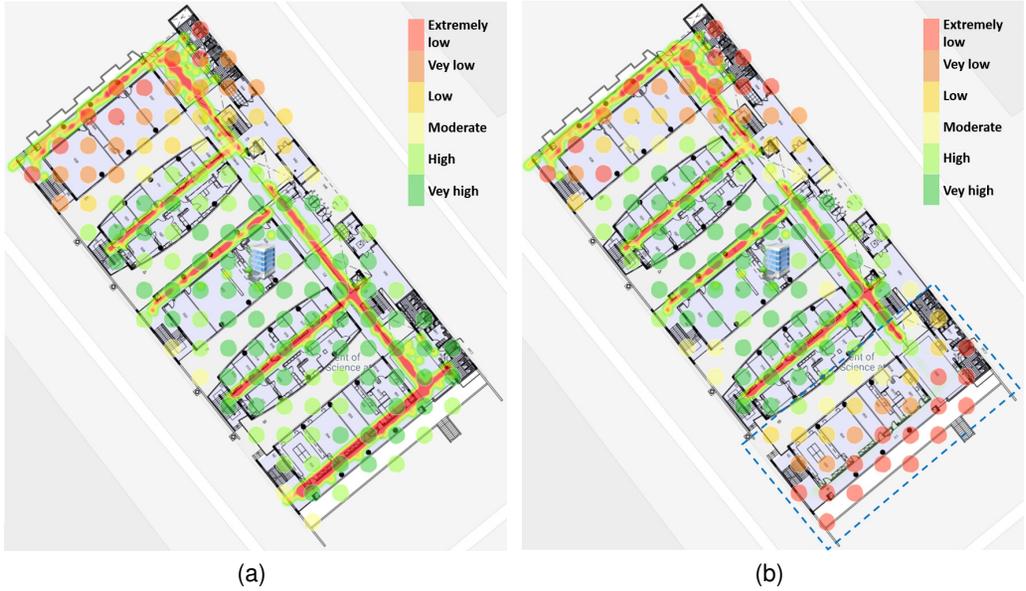


Fig. 9. (E1) Performance of the ACCES-w operator when real data are not available for a specific section of the floor: (a) All fingerprints in the original *FM*; and (b) Fingerprints removed from the bottom wing of the floor.

E2: The goal of this scenario is to demonstrate through the ACCES-w heatmap the accuracy degradation in those parts of the floor where the tampered fingerprints contain RSS values only from a few APs. In this case, the system operator receives valuable insight on where additional Wi-Fi APs could be deployed to improve the location accuracy. We consider the data in the Anyplace service that correspond to floor 1 of the MALL dataset and investigate the behavior of the ACCES-w operator when the collected fingerprints in some parts of the floor contain RSS measurements from a few Wi-Fi APs. In this case, the fingerprints have limited discriminative power, i.e., they are not capable to distinguish between locations that might be far from each other because of the low dimension of the fingerprints. The location accuracy heatmap that corresponds to the original *FM*, where all Wi-Fi APs are considered, is illustrated in Figure 10a. In total, there are 91 Wi-Fi APs heard on this floor, while the number of APs sampled in the fingerprints across different locations ranges from 1–23 APs.

We emulate this scenario by manually removing the RSS values pertaining to a subset of Wi-Fi APs across all fingerprints in the original *FM*. Note that when an AP covers the floor partially, which is the typical case in large indoor spaces, then disregarding that AP will affect only the fingerprints within its coverage area. Figure 10b depicts the resulting location accuracy heatmap when eight APs (marked with red color) have been removed. We see that in several parts of the floor, where previously better QoP was anticipated, location accuracy will degrade significantly due to the lower number of Wi-Fi APs in the fingerprints. For instance, two relatively large areas that are now suffering worse QoP are indicated with blue dashed line. In this case, the system operator can easily identify the problematic areas (i.e., those covered by a few Wi-Fi APs) and decide to strategically deploy additional APs to increase the dimension of the fingerprints, thus improving the QoP.

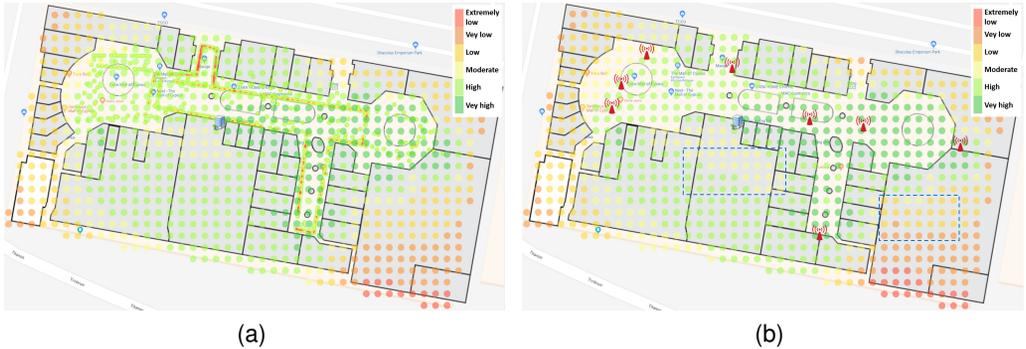


Fig. 10. (E2) Performance of the ACCES-w operator when the fingerprints contain RSS measurements from less Wi-Fi APs: (a) Fingerprints with measurements from all Wi-Fi APs; and (b) Measurements from 8 Wi-Fi APs (red color) are removed.

6.2 Series-2: ACCES Microbenchmarking

In this series we focus on microbenchmarking of the ACCES operator using magnetic data. Since both operators behave similarly in respect to the metrics described next, we only present the results for magnetic data.

6.2.1 Real Magnetic Data. We use datasets covering magnetometer readings to show the applicability of our analysis to fingerprint data of any modeling complexity. This dataset, obtained from the UJIIndoorLoc-Mag database [51], consists of 270 continuous training and 11 testing samples with each sample comprising a set of discrete samples collected along 8 corridors of a 260 m^2 lab, with a total of 40,159 discrete samples, each containing readings from the magnetometer, the accelerometer and the orientation sensor taken by an Android smartphone.

Training samples are split into two groups: “lines” and “curves”. Samples from the first group represent a sequence of measurements taken along a single corridor in two opposing directions with 5 passes for each line, whereas samples from the latter group are taken along each possible pair of connected corridors with 5 passes for each curve. Testing samples represent a sequence of measurements obtained over complex routes consisting of several corridors.

Dataset Construction: We consider this dataset consisting of 1-D data of fingerprints taken only over corridors (similarly to the approach proposed in [1]); thus, we consider the venue being split into the sub-areas (corridors), which we study separately. For the quality assessment of a FM , we are interested only in the 3 values of magnetometer readings, thus, we discard accelerometer and orientation data. To simplify the localization scenarios we use only “lines” group of samples both for the FM construction and testing phases. Moreover, we consider only passes in one direction, thus, we leave only 5 sequences of readings for each corridor. The FM data used for the interpolation step consists of the 4 combined sequences of fingerprints per corridor with the 5th used for navigability score calculation purposes during the testing step. In particular, for the sake of RMSE calculation, each of five attempt sequences are one by one chosen as testing data and the rest four as FM data, thus, producing five FM s and five testing sequences.

6.2.2 Methodology. We conjecture that changes in the observed RMSE along locations would be reflected in corresponding changes in the calculated score. We assess our hypothesis as follows: (i) given the fingerprint map FM , we evaluate the navigability score and the RMSE of location

estimation using the *WkNN* algorithm at a subset of locations $L = \{r_j : j = 1, \dots, M\}$; and (ii) we calculate the *relative similarity* between the respective values of RMSE and the navigability score. Since the navigability score may be defined not only as a lower-bound on RMSE, but in arbitrary manner carrying any unit measure and physical meaning, it is not possible to directly compare real localization error with it. Therefore, we construct the relative similarity so that it captures not difference in the values, but rather difference in relative changes (over space) of the values. We assume L to comprise not 2-D but 1-D locations (i.e., $d_r = 1$) in order to simplify the calculations, and moreover, without loss of generality, we assume they are sorted by coordinate values.

Metrics: Now, we describe how we calculate such relative similarity. Let $X = X_j$ and $Y = Y_j$ denote sets of values associated to the 1-D locations $r_j \in L$. Then, we define a similarity measure between these sets as follows: (i) to capture the behavior of X and Y , we calculate their *Difference Quotients* (DQ), $DQ(X_j)$ and $DQ(Y_j)$ for X and Y at each location $r_j \in L$; (ii) to evaluate the distance between the respective DQs we apply sequences similarity measure algorithm, which we choose to be *Dynamic Time Warping* (DTW); and (iii) we normalize the obtained distance value to the $[-1; 1]$ domain, where 1 corresponds to identical behavior, 0 to dissimilar, and -1 to opposite behavior.

The navigability score is a qualitative metric of the location accuracy that a user is expected to experience in practice. In other words, it indicates whether locations in an area of the indoor environment will have better accuracy or not than the locations in another area, i.e., how location accuracy changes over space. The DQ is a discrete approximation of a function's derivative, showing how the function grows or decreases (i.e., the measure for the average rate of change); thus, it is suitable for a qualitative description of a function's behavior. In our setting it is calculated as follows:

$$DQ(X_j) = \frac{X_j - X_{j-1}}{r_j - r_{j-1}}, j = 2, \dots, M \quad (7)$$

The boundary case for X_1 can be treated differently and is not of significant importance, we let $DQ(X_1) = DQ(X_2)$, while $DQ(Y_j)$ is calculated similarly.

DTW allows for the comparison of temporal sequences that may vary in speed, i.e., data values may be shifted and/or stretched relatively. DTW accommodates such deformations, as a *window size* parameter allows, and finds the dislocations that minimize a distance function. For example, consider the sequence $X_j = [0, 0.1, 0.2, 0.3, 0.4, 0.5]$ and its element-wise square $Y_j = X_j^2$; the sum of element-wise absolute differences $\sum |X_j - Y_j|$ yields 0.95, whereas DTW with window size 2 gives value of 0.72, and, with window size 6, value 0.55.

The DTW result is a pair of optimally warped DQ-sequences $DQ(X)'$ and $DQ(Y)'$, on which the optimal DTW distance function is obtained. Given these M' warped sequences, we compute the Relative DQ-Similarity $DQRelSim$ as follows:

$$\begin{aligned} DQRelSim(X, Y) &= \\ &= -\frac{1}{M'} \sum_{j=1}^{M'} \frac{|DQ(X_j)' - DQ(Y_j)'|}{\max(|DQ(X_j)'|, |DQ(Y_j)'|)} + 1. \end{aligned} \quad (8)$$

If $DQ(X_j)' = DQ(Y_j)'$, then $DQRelSim(X, Y) = 1$, and if $DQ(X_j)' = -DQ(Y_j)'$, then $DQRelSim(X, Y) = -1$. Moreover, if, say, X is constant, and, thus, $DQ(X_j)$ and $DQ(X_j)'$ are all zero, whereas Y is varying, then, from the construction of the normalization step, $DQRelSim(X, Y) = 0$. Note that the time warped DQ-sequences $DQ(X)$ and $DQ(Y)$ may be of different length. Therefore, common similarity measures, which typically assume sequences of the same sample size, such as the correlation coefficient are not directly applicable in our case. As it is shown in the experimental section, two sequences X and Y (i.e., RMSE and the navigability score values) are considered to behave similarly if $DQRelSim(X, Y) \geq 0.1$.

Algorithms: The RMSE can be interpreted as the achieved position error in meters [15]. It is a spatial average in the sense that it provides the expected position error at a specific location based on a number of estimated locations computed by a location estimator (i.e., localization algorithm) using noisy measurements. We consider the RMSE to be the true localization error computed at specific test locations using the associated test data, i.e., one test fingerprint at each test location. In this sense, a sequence of true position errors is obtained for a number of subsequent test locations and the derivative of RMSE, i.e., the corresponding DQ, is calculated using eq. 7. For the real indoor localization scenario we utilize the $WkNN$ localization algorithm and obtain the **RMSE** using the ground truth location from the dataset. Given the fingerprint map FM and a currently observed fingerprint \mathbf{m} , $WkNN$ calculates the Euclidean distance against it and all of the fingerprints, i.e., $d_i = \|\mathbf{m} - \mathbf{m}_i\|, \forall (\mathbf{r}_i, \mathbf{m}_i) \in FM$. Then, weights $w_i \propto d_i^{-1}$ are assigned to each fingerprint and the k most significant are chosen. The user's location is calculated using a convex combination of those k fingerprints' locations. In our experiments we set $k = 3$.

FSSI is a basic approach for FM assessment based on the map's spatial sparsity. It measures the spatial sparsity at a location \mathbf{r} via a simple navigability score called *Fingerprint Spatial Sparsity Indicator* (FSSI), defined as the area of a circle centered at \mathbf{r} and touching its nearest fingerprint (see Figure 11):

$$FSSI(\mathbf{r}) = \min_{i \in \{1, \dots, N\}} \|\mathbf{r} - \mathbf{r}_i\|, \quad (9)$$

where \mathbf{r}_i denotes the i -th fingerprint in FM . Presumably, accuracy should be lower (equivalently, RMSE should be higher) at locations with high FSSI values, as less information is obtained with fewer collected fingerprints. This allows to capture the spatial sparsity of fingerprints and is evaluated as in equation 9. Larger FSSI values correspond to smaller numbers of fingerprints per area, thus, potentially, worse accuracy. Yet, this approach will be shown to be inadequate, as, even with high FM density, the measured values may vary a little over space, leading to accuracy degradation.

Note that FSSI can adequately reflect the density of the fingerprint data in an indoor space. Thus, it is sufficient for identifying low-accuracy areas due to missing or low-density data. However, it fails to capture other intrinsic characteristics of the signals including the variation of the signals or the number of data sources (e.g., number of Wi-Fi APs) that affects the dimension of the fingerprints. For this reason, it might be misleading if the fingerprint density is high, but the signals do not vary significantly (e.g., Wi-Fi signal strength in open areas) or a few Wi-Fi APs are available. In these cases, FSSI would wrongly estimate high location accuracy simply because of the high fingerprint density, whereas in reality location accuracy will be degraded.

ACCES is the proposed navigability score, where GPR-based interpolation is done via the open-source *scikit-learn* library for Python with kernel parameters optimized using the in-built *Broyden-Fletcher-Goldfarb-Shanno* algorithm. As opposed to FSSI, the ACCES algorithm takes into consideration not only the density of the fingerprint data, but also the level of signal variation and the number of data sources to derive a more accurate and reliable navigability score. Therefore, the ACCES score follows closely the true localization error RMSE and can better estimate the location accuracy that the user will experience in practice.

The comparison of RMSE against the ACCES and FSSI values is performed via the $DQRelSim$ metric described above. For the given metric, we use a DTW window size equal to 20% of the sequence length, which appeared to be provide a reasonable performance-utility trade-off in our experiments.

To evaluate quantitatively the performance of ACCES we compared ACCES and FSSI against the RMSE values based on relative similarity in three experiments, namely:

- $E3$ – “cut” scenario, where a continuous sequence of measurements is removed from FM , which corresponds to unavailability of a part of the venue during the FM construction. In this

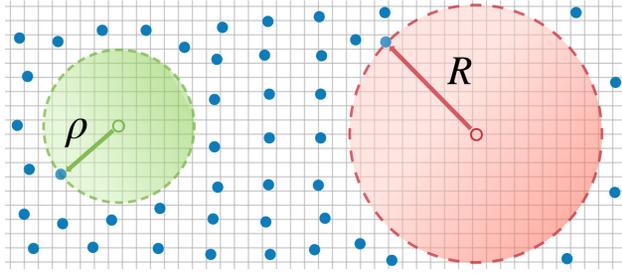


Fig. 11. *FSSI*: $FSSI_{green} = 4\pi\rho^2$, $FSSI_{red} = 4\pi R^2$. Higher *FSSI* values correspond to locations of lower *FM* density.

scenario, the estimated location accuracy is expected to degrade significantly in the part of the venue that was not available during the *FM* construction, i.e., no real data were collected in that part, because fingerprints cannot be predicted reliably without real data.

- *E4* – “flat” scenario, where a continuous sequence of measurements from *FM* is made constant, which corresponds to insufficient magnetic field variation in the corridor of the building due to the lack of sources of magnetic disturbance. In this scenario, the estimated location accuracy is expected to decrease considerably in that corridor because the magnetic signal has no variation, thus the signal stability will inevitably introduce ambiguity in the estimated user location and accuracy in practice will be worse.
- *E5* – “sparse” scenario, where fingerprints are removed from the *FM* uniformly, which relates to variability of a sensor reading rate during fingerprint collection. In this scenario, the estimated location accuracy is expected to become worse because of the lower density of collected data, which leads to increased uncertainty in the fingerprint prediction.

6.2.3 Experimental Evaluation. To assess the behavior of ACCES and naïve *FSSI* approach, we evaluate them over test locations along the 8 corridors in the “cut”, “flat” and “sparse” scenarios, and compare them against the respective RMSE values using the *DQRelSim* metric. The graphs in Figure 12 and Figure 13 present the results pertaining to a subset of the corridors for brevity; the *DQRelSim* values in all 8 corridors and the corresponding graphs for additional corridors can be found in [37].

E3 – Figure 12a shows the distribution of the magnetic field magnitude for the complete and “cut” *FM* data, when a sufficiently large interval of fingerprints is removed. Figure 12b shows the respective distributions of the real localization error RMSE and navigability scores *FSSI* and ACCES (normalized to the maximum value). Both *FSSI* and ACCES show reasonable behavior patterns and yield good similarity values *DQRelSim* to RMSE that are greater than 0.2.

E4 – Figure 12c shows the distribution over the corridor of the magnetic field magnitude for the initial and “flat” *FM* data, when a sufficiently large interval of the measurements is made constant. Figure 12d shows the respective distributions of the real localization error RMSE and the navigability scores *FSSI* and ACCES. Remarkably, ACCES shows a correspondence to the RMSE, especially in the interval of measurements’ stagnation and yields *DQRelSim* values greater than 0.2, which indicates good similarity; on the other hand, *FSSI* does not detect any potentially problematic areas, keeping a nearly constant value level over the corridors and yielding nearly zero *DQRelSim* values, which indicates no similarity.

E5 – With respect to the “sparse” scenario, Figure 13 shows the ACCES values (Figure 13a,c) and RMSE values (Figure 13b,d) for 3 cases, i.e., (i) complete *FM*; (ii) half of fingerprints removed

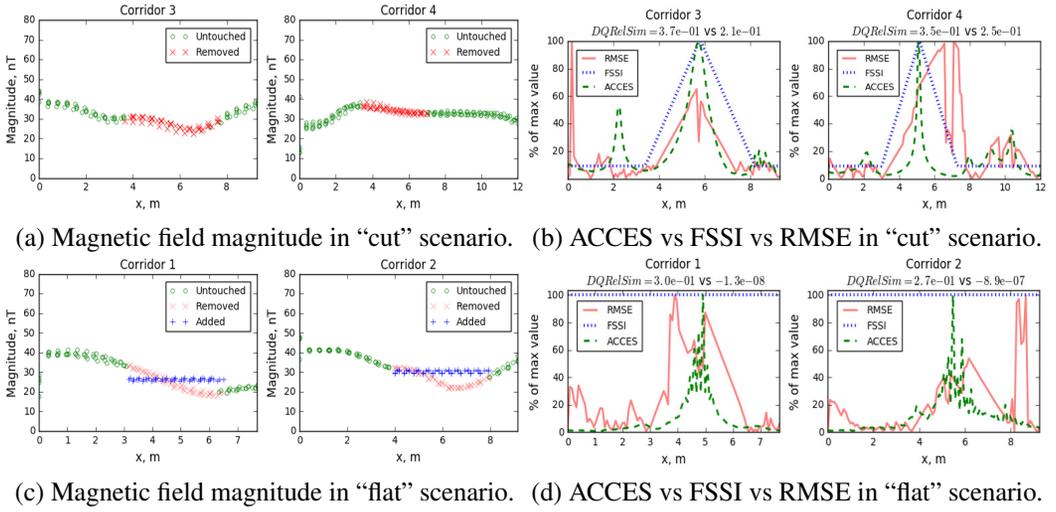


Fig. 12. (E3, E4) Magnetic field magnitude and ACCES vs FSSI vs RMSE for 4 corridors; $DQRelSim(RMSE, ACCES)$, $DQRelSim(RMSE, FSSI)$ values on top.

uniformly from the FM ; and (iii) 75% of fingerprints removed uniformly from the FM . Overall, we note that both ACCES and RMSE values indeed grow as the number of fingerprints decreases because collecting fewer data leads to accuracy degradation.

7 CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we study the *Quality-of-Position (QoP)* assessment problem, which aims to assess the accuracy of FMs in an offline manner. Particularly, the proposed ACCES method operates in two phases: the *interpolation* phase, in which an arbitrary fingerprint map is used as input for *Gaussian Process Regression*, yielding the fingerprint prediction at any location in the form of a *likelihood* function; and the *accuracy estimation* phase, in which the likelihood function is used for the calculation of *Cramer-Rao Lower Bound* on localization error, which is considered as a navigability score, qualitatively describing the real localization performance and consequently the QoP. We have derived adaptations of ACCES for both Magnetic and Wi-Fi data (ACCES-w) and implemented them in a complete visual assessment interface, coined Fingerprint Management Studio (FMS), which has been incorporated in the Anyplace open-source IIN. Our experimental study reveals that the QoP visualization and the underlying navigability scores correspond well to the real localization errors. Additionally, the FMS user interface allows indoor location service providers to fine tune the localization hardware infrastructure and to reduce in this way the deployment costs.

In the future, we aim to apply our framework to better cope with mobility assets (e.g., Wi-Fi mesh topology that changes over time) in emergency response scenarios (e.g., inside ships, trains and airplanes). Another major future challenge is to fuse multi-modal location-dependent sensing data, coming from highly diverse, low-cost and error-prone smartphone sensors, in order to push indoor accuracy to the limit. As such, we aim to derive extensions that will expand to other localization technologies as well (e.g., BLE). Finally, given that the technology road-map is towards indoor GIS integration where indoorgml.net, geojson.org or any other standard may appear in the future and become fully inter-operable, we aim to align the data formats in Anyplace and ACCES with these developments.

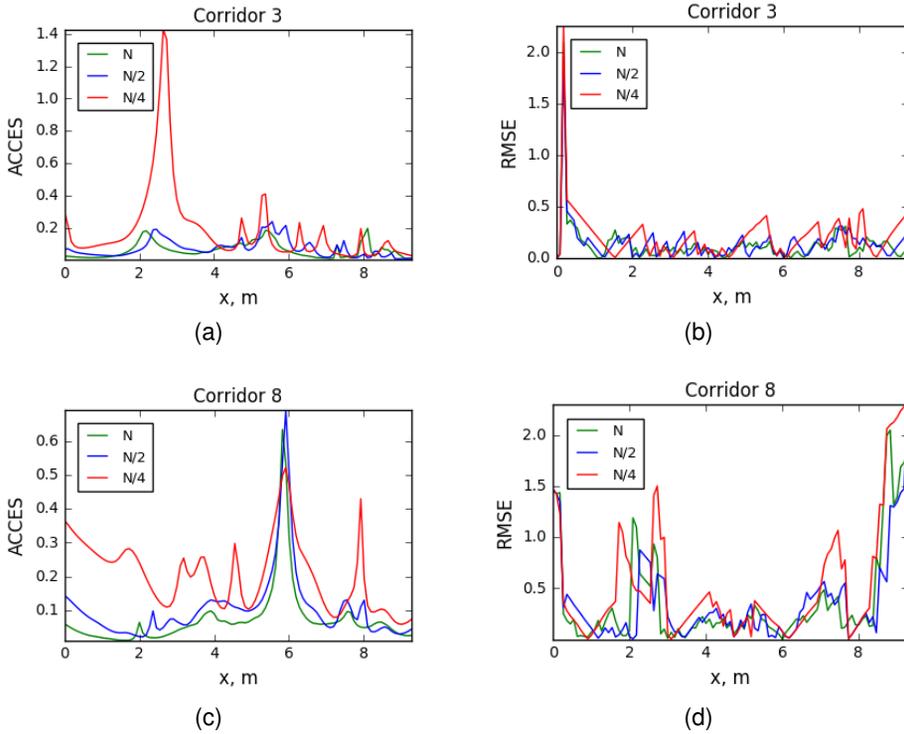


Fig. 13. (E5) ACCES and RMSE values for 2 corridors in the “sparse” scenario.

ACKNOWLEDGMENTS

We would like to thank the researchers at the Data Management Systems Laboratory of the University of Cyprus and the Anyplace Community at wide that have contributed to this open-source project over the years. This work was supported in part by Cyprus’ Research Promotion Foundation RESTART 2016-2020 programme, under project EnterCY INTEGRATED/0609/0020; EU’s Horizon 2020 research and innovation programme project LASH FIRE under grant agreement No 814975; and EU’s Marie Skłodowska-Curie Actions (MSCA) H2020-MSCA-RISE-2018 project ENDORSE, under grant agreement No 823887. The work of the first author has been supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 739551 (KIOS CoE) and from the Republic of Cyprus through the Directorate General for European Programmes, Coordination and Development.

REFERENCES

- [1] Dragan Ahmetovic, Cole Gleason, Kris M. Kitani, Hironobu Takagi, and Chieko Asakawa. 2016. NavCog: Turn-by-Turn Smartphone Navigation Assistant for People with Visual Impairments or Blindness. In *Proceedings of the 13th Web for All Conference (W4A’16)*. Association for Computing Machinery, New York, NY, USA, Article Article 9, 2 pages. <https://doi.org/10.1145/2899475.2899509>
- [2] N. Akai and K. Ozaki. 2015. Gaussian processes for magnetic map-based localization in large-scale indoor environments. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE/RSJ, New York, NY, USA, 4459–4464. <https://doi.org/10.1109/IROS.2015.7354010>
- [3] M. Angelidou, C. Costa, A. Nikitin, and D. Zeinalipour-Yazti. 2018. FMS: Managing Crowdsourced Indoor Signals with the Fingerprint Management Studio. In *2018 19th IEEE International Conference on Mobile Data Management*

- (MDM). IEEE, New York, NY, USA, 288–289. <https://doi.org/10.1109/MDM.2018.00054>
- [4] P. Bahl and V. 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, New York, NY, USA, 775–784 vol.2. <https://doi.org/10.1109/INFCOM.2000.832252>
 - [5] Artur Baniukevic, Christian S. Jensen, and Hua Lu. 2013. Hybrid Indoor Positioning with Wi-Fi and Bluetooth: Architecture and Performance. In *Proceedings of the 14th IEEE International Conference on Mobile Data Management*. IEEE, New York, NY, USA, 207–216.
 - [6] C. Beder, A. McGibney, and M. Klepal. 2011. Predicting the expected accuracy for fingerprinting based WiFi localisation systems. In *2011 International Conference on Indoor Positioning and Indoor Navigation*. IEEE, New York, NY, USA, 1–6. <https://doi.org/10.1109/IPIN.2011.6071939>
 - [7] J. Biswas and M. Veloso. 2010. WiFi localization and navigation for autonomous indoor mobile robots. In *2010 IEEE International Conference on Robotics and Automation*. IEEE, New York, NY, USA, 4379–4384. <https://doi.org/10.1109/ROBOT.2010.5509842>
 - [8] J. Blankenbach and A. Norrdine. 2010. Position estimation using artificial generated magnetic fields. In *2010 International Conference on Indoor Positioning and Indoor Navigation*. IEEE, New York, NY, USA, 1–5. <https://doi.org/10.1109/IPIN.2010.5646739>
 - [9] L. Bruno, M. Khider, and P. Robertson. 2013. On-line training of the path-loss model in Bayesian WLAN indoor positioning. In *International Conference on Indoor Positioning and Indoor Navigation*. IEEE, New York, NY, USA, 1–9. <https://doi.org/10.1109/IPIN.2013.6817897>
 - [10] Georgios Chatzimilioudis, Andreas Konstantinidis, Christos Laoudias, and Demetrios Zeinalipour-Yazti. 2012. Crowdsourcing with Smartphones. *IEEE Internet Computing* 16, 5 (2012), 36–44.
 - [11] 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 (MobiCom'10)*. Association for Computing Machinery, New York, NY, USA, 173–184. <https://doi.org/10.1145/1859995.1860016>
 - [12] Rizzanne Elbakly and Moustafa Youssef. 2016. CONE: Zero-Calibration Accurate Confidence Estimation for Indoor Localization Systems. *CoRR* abs/1610.02274 (2016). arXiv:1610.02274 <http://arxiv.org/abs/1610.02274>
 - [13] Carlos E Galván-Tejada, Juan Pablo García-Vázquez, Jorge I Galván-Tejada, J Rubén Delgado-Contreras, and Ramon F Brena. 2015. Infrastructure-less indoor localization using the microphone, magnetometer and light sensor of a smartphone. *Sensors* 15, 8 (2015), 20355–20372.
 - [14] Kyriakos Georgiou, Timotheos Constambeys, Christos Laoudias, Lambros Petrou, Georgios Chatzimilioudis, and Demetrios Zeinalipour-Yazti. 2015. Anyplace: A Crowdsourced Indoor Information Service. In *The 16th IEEE International Conference on Mobile Data Management*, Vol. 1. IEEE, New York, NY, USA, 291–294. <https://doi.org/10.1109/MDM.2015.80>
 - [15] Fredrik Gustafsson and Fredrik Gunnarsson. 2005. Mobile positioning using wireless networks: possibilities and fundamental limitations based on available wireless network measurements. *IEEE Signal processing magazine* 22, 4 (2005), 41–53.
 - [16] Robert Harle. 2013. A Survey of Indoor Inertial Positioning Systems for Pedestrians. *IEEE Communications Surveys and Tutorials* 15, 3 (2013), 1281–1293.
 - [17] Janne Haverinen and Anssi Kemppainen. 2009. Global indoor self-localization based on the ambient magnetic field. *Robotics and Autonomous Systems* 57, 10 (2009), 1028–1035.
 - [18] Suining He, Tianyang Hu, and S.-H. Gary Chan. 2015. Contour-Based Trilateration for Indoor Fingerprinting Localization. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys'15)*. Association for Computing Machinery, New York, NY, USA, 225–238. <https://doi.org/10.1145/2809695.2809703>
 - [19] A. K. M. M. Hossain and W. Soh. 2010. Cramer-Rao Bound Analysis of Localization Using Signal Strength Difference as Location Fingerprint. In *2010 Proceedings IEEE INFOCOM*. IEEE, New York, NY, USA, 1–9. <https://doi.org/10.1109/INFCOM.2010.5462020>
 - [20] Pavel Ivanov and Lauri Aarne Johannes Wirola. 2020. Metric for evaluating indoor positioning systems, U.S. patent number 10,555,132, February 4, 2020. (2020).
 - [21] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An introduction to statistical learning*. Vol. 112. Springer, Berlin, Heidelberg.
 - [22] Suk Hoon Jung, Byeong-Cheol Moon, and Dongsoo Han. 2017. Performance evaluation of radio map construction methods for Wi-Fi positioning systems. *IEEE Transactions on Intelligent Transportation Systems* 18, 4 (2017), 880–889.
 - [23] K. Kaemarungsi and P. Krishnamurthy. 2004. Properties of indoor received signal strength for WLAN location fingerprinting. In *The First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services, 2004. MOBIQUITOUS 2004*. IEEE, New York, NY, USA, 14–23. <https://doi.org/10.1109/MOBIQ.2004.1331706>

- [24] L. Kanaris, A. Kokkinis, A. Liotta, and S. Stavrou. 2017. Quality of fingerprint radiomaps for positioning systems. In *2017 24th International Conference on Telecommunications (ICT)*. IEEE, New York, NY, USA, 1–5. <https://doi.org/10.1109/ICT.2017.7998235>
- [25] Yungeun Kim, Seokjun Lee, Yohan Chon, Rhan Ha, and Hojung Cha. 2017. Scalable and Consistent Radio Map Management Using Participatory Sensing. *Pervasive Mob. Comput.* 40, C (Sept. 2017), 397–413. <https://doi.org/10.1016/j.pmcj.2017.04.001>
- [26] Andreas Konstantinidis, Georgios Chatzimilioudis, Demetrios Zeinalipour-Yazti, Paschalis Mpeis, Nikos Pelekis, and Yannis Theodoridis. 2015. Privacy-Preserving Indoor Localization on Smartphones. *IEEE Transactions on Knowledge and Data Engineering* 27, 11 (Nov 2015), 3042–3055. <https://doi.org/10.1109/TKDE.2015.2441724>
- [27] Andreas Konstantinidis, Panagiotis Irakleous, Zacharias Georgiou, Demetrios Zeinalipour-Yazti, and Panos K. Chrysanthis. 2018. IoT Data Prefetching in Indoor Navigation SOAs. *ACM Trans. Internet Technol.* 19, 1, Article Article 10 (Nov. 2018), 21 pages. <https://doi.org/10.1145/3177777>
- [28] P. Kontkanen, P. Myllymaki, T. Roos, H. Tirri, K. Valtonen, and H. Wetzig. 2004. Topics in probabilistic location estimation in wireless networks. In *2004 IEEE 15th International Symposium on Personal, Indoor and Mobile Radio Communications (IEEE Cat. No.04TH8754)*, Vol. 2. IEEE, New York, NY, USA, 1052–1056 Vol.2. <https://doi.org/10.1109/PIMRC.2004.1373859>
- [29] Sudhir Kumar, Rajesh M Hegde, and Niki Trigoni. 2016. Gaussian Process Regression for Fingerprinting based Localization. *Ad Hoc Networks* 51 (2016), 1–10.
- [30] Christos Laoudias, Adriano Moreira, Sunwoo Kim, Sangwoo Lee, Lauri Wirola, and Carlo Fischione. 2018. A survey of enabling technologies for network localization, tracking, and navigation. *IEEE Communications Surveys & Tutorials* 20, 4 (2018), 3607–3644.
- [31] Hendrik Lemelson, Mikkel Baun Kjærgaard, Rene Hansen, and Thomas King. 2009. Error estimation for indoor 802.11 location fingerprinting. In *International Symposium on Location-and Context-Awareness*. Springer, Berlin, Heidelberg, 138–155. https://doi.org/10.1007/978-3-642-01721-6_9
- [32] Filip Lemic, Vlado Handziski, and Jeroen Famaey. 2019. Toward Regression-Based Estimation of Localization Errors in Fingerprinting-Based Localization. In *2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*. IEEE, New York, NY, USA, 1–5. <https://doi.org/10.1109/VTCSpring.2019.8746700>
- [33] Binghao Li, James Salter, Andrew G. Dempster, and Chris Rizos. 2006. Indoor positioning techniques based on wireless LAN. In *1st International Conference on Wireless Broadband and Ultra Wideband Communications*. IEEE, New York, NY, USA, 13–16.
- [34] Hui Liu, H. Darabi, P. Banerjee, and Jing Liu. 2007. Survey of Wireless Indoor Positioning Techniques and Systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* 37, 6 (2007), 1067–1080.
- [35] 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 (IPSN'15)*. Association for Computing Machinery, New York, NY, USA, 178–189. <https://doi.org/10.1145/2737095.2737726>
- [36] Paschalis Mpeis, Thierry Roussel, Manish Kumar, Constantinos Costa, Christos Laoudias, Dennis Capot-Ray, and Demetrios Zeinalipour-Yazti. 2020. The Anyplace 4.0 IoT Localization Architecture. In *21st IEEE International Conference on Mobile Data Management (MDM)*. IEEE, New York, NY, USA, 218–225. <https://doi.org/10.1109/MDM48529.2020.00045>
- [37] Artyom Nikitin, Christos Laoudias, Georgios Chatzimilioudis, Panagiotis Karras, and Demetrios Zeinalipour-Yazti. 2017. Indoor localization accuracy estimation from fingerprint data. In *18th IEEE International Conference on Mobile Data Management (MDM)*. IEEE, New York, NY, USA, 196–205.
- [38] H. Nurminen, J. Talvitie, S. Ali-L  y  t  , P. M  ij  ller, E. Lohan, R. Pich  l, and M. Renfors. 2012. Statistical path loss parameter estimation and positioning using RSS measurements in indoor wireless networks. In *2012 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, New York, NY, USA, 1–9. <https://doi.org/10.1109/IPIN.2012.6418856>
- [39] F. Pereira, C. Theis, A. Moreira, and M. Ricardo. 2011. Evaluating location fingerprinting methods for underground GSM networks deployed over Leaky Feeder. In *2011 International Conference on Indoor Positioning and Indoor Navigation*. IEEE, New York, NY, USA, 1–6. <https://doi.org/10.1109/IPIN.2011.6071940>
- [40] Lambros Petrou, George Larkou, Christos Laoudias, Demetrios Zeinalipour-Yazti, and Christos G. Panayiotou. 2014. Demonstration Abstract: Crowdsourced Indoor Localization and Navigation with Anyplace. In *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks (IPSN'14)*. IEEE Press, New York, NY, USA, 331–332.
- [41] Sanjay Purushotham and C.-C. Jay Kuo. 2016. Personalized group recommender systems for location-and event-based social networks. *ACM Transactions on Spatial Algorithms and Systems (TSAS)* 2, 4 (2016), 16.

- [42] Carl Edward Rasmussen and Christopher K. I. Williams. 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, Cambridge, MA, USA.
- [43] Teemu Roos, Petri Myllymaki, Henry Tirri, Pauli Misikangas, and Juha Sievanen. 2002. A probabilistic approach to WLAN user location estimation. *International Journal of Wireless Information Networks* 9, 3 (2002), 155–164.
- [44] E. Sansano, R. Montoliu, and J. Torres-Sospedra. 2017. A novel methodology to estimate a measurement of the inherent difficulty of an indoor localization radio map. In *2017 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, New York, NY, USA, 1–8. <https://doi.org/10.1109/IPIN.2017.8115939>
- [45] E. Sansano, R. Montoliu, and J. Torres-Sospedra. 2017. A novel methodology to estimate a measurement of the inherent difficulty of an indoor localization radio map. In *2017 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, New York, NY, USA, 1–8. <https://doi.org/10.1109/IPIN.2017.8115939>
- [46] Georgios Skoumas, Dieter Pfoser, Anastasios Kyriillidis, and Timos Sellis. 2016. Location estimation using crowdsourced spatial relations. *ACM Transactions on Spatial Algorithms and Systems (TSAS)* 2, 2 (2016), 5.
- [47] Maja Stella, Mladen Russo, and Dinko Begusic. 2012. RF localization in indoor environment. *Radioengineering* 21, 2 (2012), 557–567.
- [48] Sukhoon Jung, Choon-oh Lee, and Dongsoo Han. 2011. Wi-Fi fingerprint-based approaches following log-distance path loss model for indoor positioning. In *2011 IEEE MTT-S International Microwave Workshop Series on Intelligent Radio for Future Personal Terminals*. IEEE, New York, NY, USA, 1–2. <https://doi.org/10.1109/IMWS2.2011.6027190>
- [49] Shan-Yun Teng, Wei-Shinn Ku, and Kun-Ta Chuang. 2017. Toward Mining Stop-by Behaviors in Indoor Space. *ACM Transactions on Spatial Algorithms and Systems (TSAS)* 3, 2 (2017), 7.
- [50] Joaquín Torres-Sospedra, Raúl Montoliu, Sergio Trilles, Oscar Belmonte, and Joaquín Huerta. 2015. Comprehensive analysis of distance and similarity measures for Wi-Fi fingerprinting indoor positioning systems. *Expert Systems with Applications* 42, 23 (2015), 9263–9278.
- [51] J. Torres-Sospedra, D. Rambla, R. Montoliu, O. Belmonte, and J. Huerta. 2015. UJIIndoorLoc-Mag: A new database for magnetic field-based localization problems. In *2015 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, New York, NY, USA, 1–10. <https://doi.org/10.1109/IPIN.2015.7346763>
- [52] Bo Wang and Jian Qing Shi. 2014. Generalized Gaussian Process Regression Model for Non-Gaussian Functional Data. *Journal of the American Statistical Association* 109, 507 (2014), 1123–1133. <https://doi.org/10.1080/01621459.2014.889021>
- [53] He Wang, Souvik Sen, Ahmed Elgohary, Moustafa Farid, Moustafa Youssef, and Romit Roy Choudhury. 2012. No Need to War-Drive: Unsupervised Indoor Localization. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services (MobiSys '12)*. Association for Computing Machinery, New York, NY, USA, 197–210. <https://doi.org/10.1145/2307636.2307655>
- [54] Hongkai Wen, Zhuoling Xiao, Andrew Markham, and Niki Trigoni. 2015. Accuracy estimation for sensor systems. *IEEE Transactions on Mobile Computing* 14, 7 (2015), 1330–1343.
- [55] Xinrong Li, K. Pahlavan, M. Latva-aho, and M. Ylianttila. 2000. Comparison of indoor geolocation methods in DSSS and OFDM wireless LAN systems. In *Vehicular Technology Conference Fall 2000. IEEE VTS Fall VTC2000. 52nd Vehicular Technology Conference (Cat. No.00CH37152)*, Vol. 6. IEEE, New York, NY, USA, 3015–3020 vol.6. <https://doi.org/10.1109/VETEFCF.2000.886867>
- [56] Moustafa Youssef and Ashok Agrawala. 2005. The Horus WLAN Location Determination System. In *Proceedings of the 3rd International Conference on Mobile Systems, Applications, and Services (MobiSys '05)*. Association for Computing Machinery, New York, NY, USA, 205–218. <https://doi.org/10.1145/1067170.1067193>
- [57] Demetrios Zeinalipour-Yazti and Christos Laoudias. 2017. The anatomy of the anyplace indoor navigation service. *SIGSPATIAL Special* 9, 2 (2017), 3–10.
- [58] Demetrios Zeinalipour-Yazti, Christos Laoudias, Kyriakos Georgiou, and Georgios Chatzimilioudis. 2017. Internet-Based Indoor Navigation Services. *IEEE Internet Computing* 21, 4 (2017), 54–63.

A APPENDIX

In this section we derive CRLB for the predictive distribution of the GPR. As we discussed, in the case of GPR, the predictive distribution, i.e., the distribution of noisy function values for some parameter value, is Gaussian

$$\mathbf{x}|\theta \sim \mathcal{N}(\boldsymbol{\mu}(\theta), \Sigma(\theta)),$$

where \mathbf{x} is the predicted value for the parameter value θ , $\boldsymbol{\mu}(\theta)$ and $\Sigma(\theta)$ are the mean vector and the covariance matrix, respectively. We assume that the covariance matrix is diagonal. Thus, the

log-likelihood function for the predicted value \mathbf{x} given the parameter $\boldsymbol{\theta}$ is:

$$\begin{aligned} \log p(\mathbf{x}|\boldsymbol{\theta}) &\sim -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \Sigma^{-1}(\boldsymbol{\theta})(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta})) - \\ &\quad - \frac{1}{2} \log |\Sigma(\boldsymbol{\theta})|, \end{aligned} \quad (10)$$

where $|\cdot|$ denotes a *determinant* of a square matrix. Plugging (10) into (5) results in the following expression for the FIM:

$$\begin{aligned} \mathcal{I}(\boldsymbol{\theta}) &= \frac{1}{2} \cdot \mathbb{E} \left[\frac{\partial^2}{\partial \theta_i \partial \theta_j} \left[(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \Sigma^{-1}(\boldsymbol{\theta})(\mathbf{x} - \boldsymbol{\mu}(\boldsymbol{\theta})) + \right. \right. \\ &\quad \left. \left. + \log |\Sigma(\boldsymbol{\theta})| \right] \right]. \end{aligned} \quad (11)$$

Subsequently, it can be transformed to

$$\begin{aligned} \mathcal{I}(\boldsymbol{\theta}) &= \frac{1}{2} \sum_{k=1}^{d_x} \left[(\sigma_k^2 + \mu_k^2) H(\sigma_k^{-2}) + H(\mu_k^2 \sigma_k^{-2}) - \right. \\ &\quad \left. - 2\mu_k H(\mu_k \sigma_k^{-2}) + 2H(\log \sigma_k) \right], \end{aligned} \quad (12)$$

where we omit the arguments of the mean and variance for the sake of readability, $H(\cdot)$ is a *Hessian Matrix*, μ_k is the k -th component of $\boldsymbol{\mu}$ and σ_k is the k -th diagonal element of Σ . The CRLB is then given by the inequality:

$$\text{cov}(\hat{\boldsymbol{\theta}}) \geq \mathcal{I}^{-1}(\boldsymbol{\theta}),$$

where $\hat{\boldsymbol{\theta}}$ is any unbiased estimator of $\boldsymbol{\theta}$, and $\text{cov}(\cdot)$ its covariance matrix. Now the Root Mean Square Error of the estimate is given by the square root of the *trace* of the covariance matrix, thus:

$$\text{RMSE} \geq \sqrt{\text{tr}(\mathcal{I}^{-1}(\boldsymbol{\theta}))}. \quad (13)$$

For the case of the scalar parameter θ , i.e., $d_\theta = 1$ (which is considered in our experimental evaluation), the FIM is a one-element matrix, and, thus, the derived formulas are directly applicable with a Hessian matrix being just a second-order partial derivative.