Indoor localization using ambient FM radio RSS fingerprinting: A 9-month study

Andrei Popleteev
University of Luxembourg,
Luxembourg
Email: contact@popleteev.com

***Abstract*—While indoor positioning systems aspire for higher accuracy, their coverage is typically limited to buildings with dedicated hardware. A possible alternative is offered by infrastructure-free positioning methods. In particular, several studies have demonstrated feasibility of indoor positioning using broadcast FM radio signals, which are available in most populated areas worldwide. However, previous work provides little information about long-term performance of FM-based indoor localization.**

This paper presents a longitudinal study of FM indoor positioning based on received signal strength (RSS) fingerprinting. We evaluate system's performance on a large dataset of real-world FM signals, systematically collected in several large-scale multi-floor testbeds over the course of 9 months. We also investigate the impact of different classifiers, training schedules and fingerprint sizes on localization accuracy. The results demonstrate that well-trained FM-based system can provide reliable indoor positioning even several months after deployment.

1. Introduction

Context-aware applications and mobile services need to know user location, wherever it is. While satellite-based systems like GPS and GLONASS are globally available in clear-sky areas, most of indoor positioning systems are constrained by the required in-building infrastructure — be it Bluetooth beacons [1], Wi-Fi networks [2], or ultrasound and ultrawideband transceivers [3], [4]. Inertial navigation, in turn, is infrastructure-independent, but still periodically needs an absolute positioning system to reset accumulated tracking errors [5].

Ambient signals of opportunity — such as commercial FM radio broadcasts — provide a promising alternative for infrastructure-free indoor positioning. High antenna elevations and transmission powers ensure good indoor availability over large areas. Due to the relatively long wavelength ($\lambda = 3$ m), FM radio waves are less attenuated by building materials than Wi-Fi or Bluetooth signals ($\lambda = 12.5$ cm). Moreover, small indoor objects reflect and scatter shorter Wi-Fi and Bluetooth waves, but are transparent for FM radio signals. Finally, FM receivers are readily available in many off-the-shelf smartphones.

A number of previous studies have investigated FM-based indoor positioning based on received signal strength (RSS) fingerprinting method [6], [7], [8], [9], [10], [11], [12]. These works showed that FM localization accuracy is comparable to that of Wi-Fi based solutions, and the two technologies can successfully complement each other. However, the limited duration of these studies provides little insight into long-term performance of FM-based indoor localization in changing environment conditions.

This paper aims to narrow the gap between short-term lab results and practical FM positioning. We present a long-term experimental evaluation of FM indoor localization performance in several scenarios, and investigate the impact of building population on FM signal stability. The results are based on an extensive dataset of FM signal samples, systematically collected in several large-scale multi-floor testbeds over the course of 9 months.

2. Background

FM-based indoor positioning systems typically employ the RSS fingerprinting method, originally created for Wi-Fi localization [13]. This approach leverages the fact that RSS fingerprints — vectors of RSS values from several stationary transmitters — depend on receiver location. Initially, the system is trained on a calibration dataset, which provides an empirical distribution of RSS values in the testbed. After calibration, the system combines its empirical knowledge about RSS distribution and machine learning methods to recognize locations by their fingerprints.

Despite the apparent similarity of localization methods, FM positioning is fundamentally different from Wi-Fi based approach. Firstly, FM radio waves are more than 20 times longer than Wi-Fi waves, and therefore do not interact with small indoor objects [14]. Secondly, in contrast to the short-range Wi-Fi signals that travel mainly indoors, FM signals propagate over kilometers of outdoor landscape, where they are exposed to environmental factors (such as road traffic or seasonal foliage). As a result, FM positioning performance cannot be directly inferred from existing Wi-Fi based results.

First studies of FM-based indoor localization also employed dedicated short-range FM transmitters [6]. However, the concept had been soon generalized to ambient FM radio stations [15]. A detailed investigation of both approaches was presented by Popleteev in his PhD thesis [8]. The author

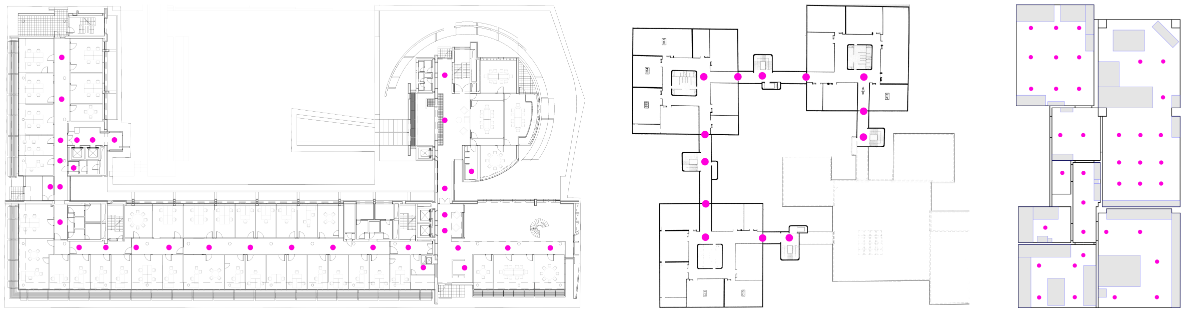

(a) Office Building (Floors 0 and 1)

(b) Campus Building (Floors 0 and 1)

(c) Apartment (Floor 3)

Figure 1. Testbed floorplans (color dots indicate the sampling locations — test points).

TABLE 1. EXPERIMENTAL TESTBEDS AND SAMPLING CHARACTERISTICS.

Testbed	Dimensions	Number of test points	Sampling period	Number of sessions	Samples collected
Office Building	100×50 m	$33 + 36 + 16$ (floors 1, 0, -2)	9 months (Feb–Oct 2016)	17	1445
Campus Building	80×80 m	$13 + 13$ (floors 1, 0)	9 months (Feb–Oct 2016)	16	416
Apartment	14×7 m	37 (floor 3)	3 months (Jan–Mar 2016)	6	222

analyzed fingerprint stability and receiver diversity issues, and compared RSS fingerprinting performance of FM, Wi-Fi and GSM signals [8], [10]. Another study, conducted in a slightly larger testbed (23×11 m vs 12×6 m), reported 1.3 m median error in same-dataset (leave-one-out) evaluation scenario [9]. In turn, [10] compared FM, GSM and Wi-Fi fingerprinting in next-day evaluation scenario, and found that FM and GSM accuracy decreases in larger testbeds. Later on, Chen et al. [11] reported 93% room recognition rate for same-dataset evaluation. However, they noticed that the accuracy decreased to 87% in "one-vs-many" cross-validation scenario based on three days in a larger testbed (shopping mall). Finally, Yoon et al. [12] proposed a method to facilitate system calibration using propagation modeling of FM radio signals from known transmitter sites; their method achieved 6 m average accuracy.

While previous studies focused on general feasibility of FM-based indoor positioning, the researchers also acknowledged that "localization performance is prone to degradation due to changing conditions in the environment" [7] and "temporal variations of the signal signatures can lead to noticeable degradation of localization accuracy" [11]. In particular, Matic et al. [7] reported accuracy degradation over the course of 7 months. Unfortunately, this period was only sparsely covered by only three datasets (collected in December, and followed by June and July of the next year). Chen et al. [11] continuously monitored FM RSS stability over 10 days, but only in one fixed location.

Overall, long-term performance of FM RSS-based indoor positioning in changing environments remains an open question, which is addressed by this paper.

3. Experiment Setup and Approach

3.1. Testbed setup

Due to the unusually long planned duration of the experiment, selecting appropriate testbeds was rather challenging. We looked for multi-floor buildings with varying human presence (sometimes empty, sometimes populated), yet accessible in off-hours. More importantly, all the test points (sampling locations) should have remained available for the whole multi-month duration of the study. Additionally, the look of the experimental equipment and its operator should not have disturbed people around. As a result, we have selected two university-owned buildings (Figure 1). One building featured research and administration offices, while the other one combined research labs and lecture halls. These testbeds are further referred to as "Offices" and "Campus". Additionally, a private apartment has been included for small-scale tests.

In each testbed, we defined a number of fixed test points, focusing primarily on wide coverage of the area. The test points were initially specified with regard to local landmarks (e.g. "in front of office 321") and later established with centimeter-level precision using laser rangefinders, in order to ensure consistent ground truth across the sessions [16], [17]. Main testbed characteristics are summarized in Table 1.

3.2. Data acquisition

Data acquisition has been performed within the scope of a long-term radio monitoring project [18], which employed software-defined radio (SDR) receivers [19] to acquire raw radio-frequency (RF) samples from multiple ambient radio sources (such as FM and TV stations, and cellular networks). At each test point, we recorded 2 s of full-band raw FM RF samples. (However, due to multi-band acquisition, tuning and storage delays, the process took on average 70 s per point.) Overall, we performed measurements approximately bi-weekly and collected 2083 radio samples across all the three testbeds (see Table 1).

Apart from the radio samples, each measurement session includes metadata about the environment, such as human presence in the building (manually specified by the operator). To avoid radio signal shadowing by the operator's body, the receiver was raised above the head level. Moreover, the operator has always faced the next test point along the predefined path, so the receiver orientation was consistent throughout the study. All floors of a specific testbed were sampled within the same day (one floor — one session) and in the same order (top to bottom).

It should be noted, that while our radio measurements have been performed using a professional SDR receiver, the results are directly generalizable to FM RSS fingerprints acquired by low-cost USB radio tuners or FM-enabled smartphones [8], [11].

3.3. Data processing

Collected raw RF samples were preprocessed offline using GnuRadio toolkit [20]. Each 2 s long sample was split to 0.2 s chunks, resulting in 10 full-band RSS fingerprints per point. Active channels with ongoing broadcasts were detected by the presence of 19 kHz FM stereo pilot.

In addition to fingerprint normalization and k-nearest neighbor (kNN) classifiers used in the previous studies [7], [8], [9], [10], [11], we also investigated a more advanced fingerprint standardization procedure, as well as random forest and support vector machine (SVM) classifiers [21].

Fingerprint standardization is a statistical method which involves independent centering and scaling of RSS values in an active FM channel to zero-mean unit-variance sequence:

$$x'_{ch} = \frac{x_{ch} - \langle x_{ch} \rangle}{\sigma[x_{ch}]}$$

where x_{ch} is a sequence of RSS values in channel ch, while $\langle x_{ch} \rangle$ and $\sigma[x_{ch}]$ are its mean and standard deviation. This approach equalizes the impact of stronger (high-RSS) and weaker (low-RSS) stations. Moreover, in contrast to simple min/max normalization, statistical standardization is robust to occasional outliers.

4. Experimental Results

4.1. Classifier selection

In the first experiment, we compare localization performance of different machine learning approaches in order to choose the most appropriate method for further analysis. In addition to the kNN classifier with Euclidean and Manhattan metrics (commonly used in previous studies), we also tested SVM and random forest algorithms.

For evaluation, we use a leave-one-session-out approach, where one measurement session is selected for testing, while the rest of data are used to train the system. After looping through all the sessions, the total ratio of correctly recognized test points provides an estimate of the localization accuracy of the system. In contrast to other performance metrics, such as median error distance, classification accuracy does not directly depend on the distance between test points and thus is more appropriate for our diverse testbeds.

TABLE 2. LOCALIZATION ACCURACY OF DIFFERENT CLASSIFIERS, WITH (WITHOUT) STANDARDIZATION, IN PERCENT.

Testbed	kNN (Eucl)	kNN (Manh)	R.Forest	SVM
Apartment (fl. 3)	52.0 (47.7)	51.3 (51.6)	52.2	**57.4**
Campus (fl. 1)	91.6 (89.6)	**95.4** (92.0)	93.7	92.0
Campus (fl. 0)	93.8 (87.2)	94.0 (87.6)	89.9	**94.3**
Offices (fl. 1)	85.4 (84.7)	86.5 (86.2)	87.6	**89.6**
Offices (fl. 0)	74.7 (69.2)	74.3 (71.0)	77.9	**79.3**
Offices (fl. -2)	51.0 (54.0)	51.7 (52.6)	54.4	**60.4**
Average	74.7 (72.1)	75.5 (73.5)	75.9	**78.8**

The results are presented in Table 2. On average, fingerprint standardization improves kNN accuracy by several percents. The other two classifiers, in turn, are indifferent to the procedure — either by design (random forest) or because they always perform standardization internally (SVM).

Among the tested classifiers, kNN with Euclidean metric showed the lowest average localization accuracy (74.7%), followed by kNN with Manhattan metric (75.5%) and random forest (75.9%). The highest average accuracy of 78.8% was demonstrated by the SVM; this classifier is used in all further experiments.

4.2. General performance evaluation

In this section we evaluate general performance of the system in two extreme scenarios: with minimal and with maximal training. In both cases we leverage the leave-one-session-out approach described in the previous section: measurement sessions are taken one by one and compared to the rest of data. In the one scenario, the system is trained on multiple sessions and tested on the selected single session ("maximal training"); this provides an optimistic performance estimate of a well-trained system. In turn, in the "minimal training" scenario the system is trained on a single session and tested on the rest of data; this provides a lower-bound estimate of system performance.

As the evaluation procedure iterates through all the sessions, we calculate localization error distances for every fingerprint in each test. In the final step, cumulative distribution function (CDF) plot of all the error distances provides an overview of general system performance.

The results are shown in Figure 2.

In the small-scale Apartment testbed, only 57.4% of test locations were correctly recognized despite the maximal training (and only 33.7% with minimal training). However, due to the spatial correlation of RSS values in densely placed test points, 90% of position estimates were still within 4.1 m from the ground truth (5.2 m with minimal training).

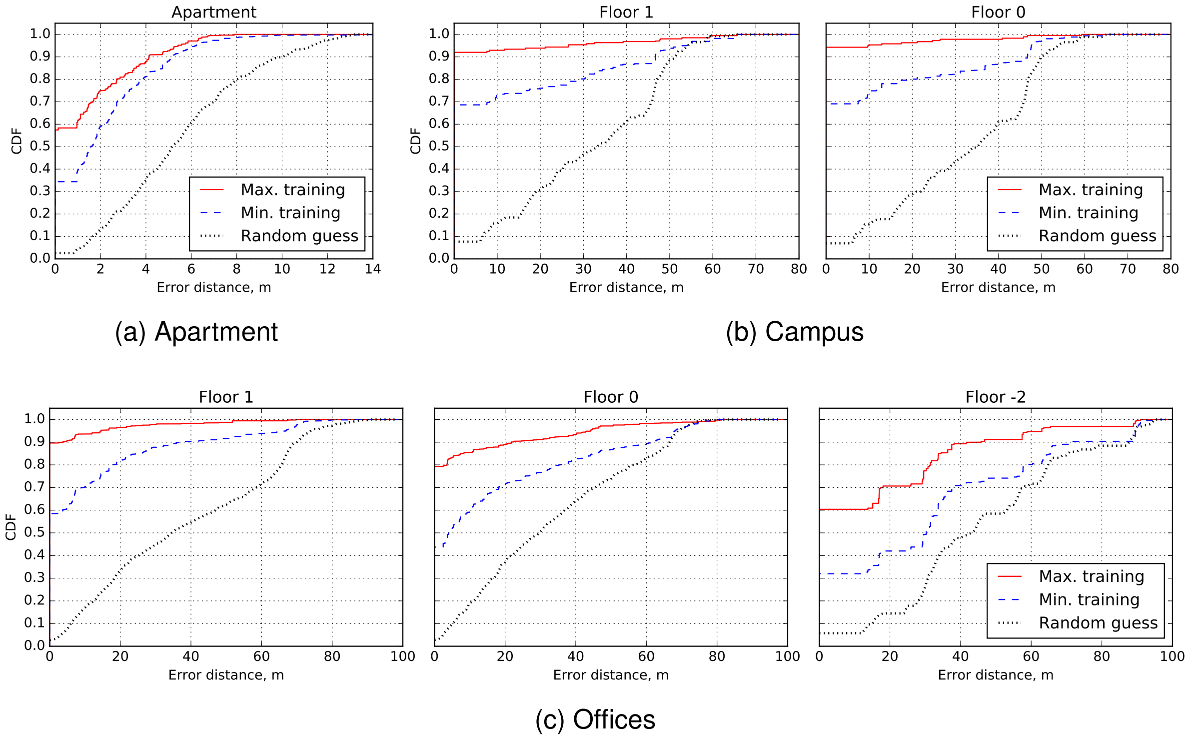

Figure 2. General localization performance.

In the large-scale testbeds, with larger distances between test points, the system showed higher localization accuracy. In particular, even with the minimal one-session training the system recognized 68.6% and 69.1% of test points in the Campus testbed (on floors 1 and 0, respectively). With maximal training the results increased to 92.0% and 94.3% (see Figure 2b).

The Offices testbed, in turn, demonstrated a rather surprising dependence of localization accuracy on the floor number (see Figure 2c). With maximal training, the accuracy varied from 89.6% to 79.3% and 60.4% on floors 1, 0 and -2, respectively (58.5%, 43.7%, and 32.0% with minimal training).

A closer analysis explains this dependency by the different number of active FM channels available at each floor (Figure 3). Indeed, as Floor 1 is higher above the ground, it is less shadowed by nearby buildings; on this floor we detected on average 24 active FM stations. In contrast, the underground parking on Floor -2 proved to be a particularly harsh environment for RSS fingerprinting: not only there were only 6 active stations, but the floor was populated with large metallic objects (cars) which changed their positions between the sessions.

4.3. Performance over time

Having evaluated the general performance of the system, we now proceed with an in-depth analysis of how FM positioning performance changes over time. Since displaying even aggregated CDFs for the high number of test combinations is not feasible, in this section we focus on a single-number metric, namely, the localization accuracy (the percentage of correctly recognized fingerprints/test points).

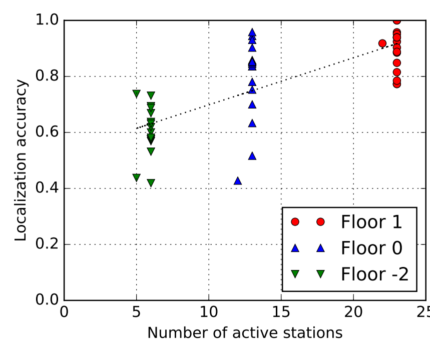

Figure 3. Localization accuracy on different floors of the Offices testbed (max training). Each point represents one test session.

First of all, we perform an exhaustive pairwise evaluation, where the system is trained on one session and tested on another one — for all the possible combinations. As shown in Figure 4, sessions close in time provide higher localization accuracy, as indicated by the more saturated color of the corresponding cells (near the main diagonals). This confirms that the accuracy of recently calibration system is likely to be high. However, the figure does not show any

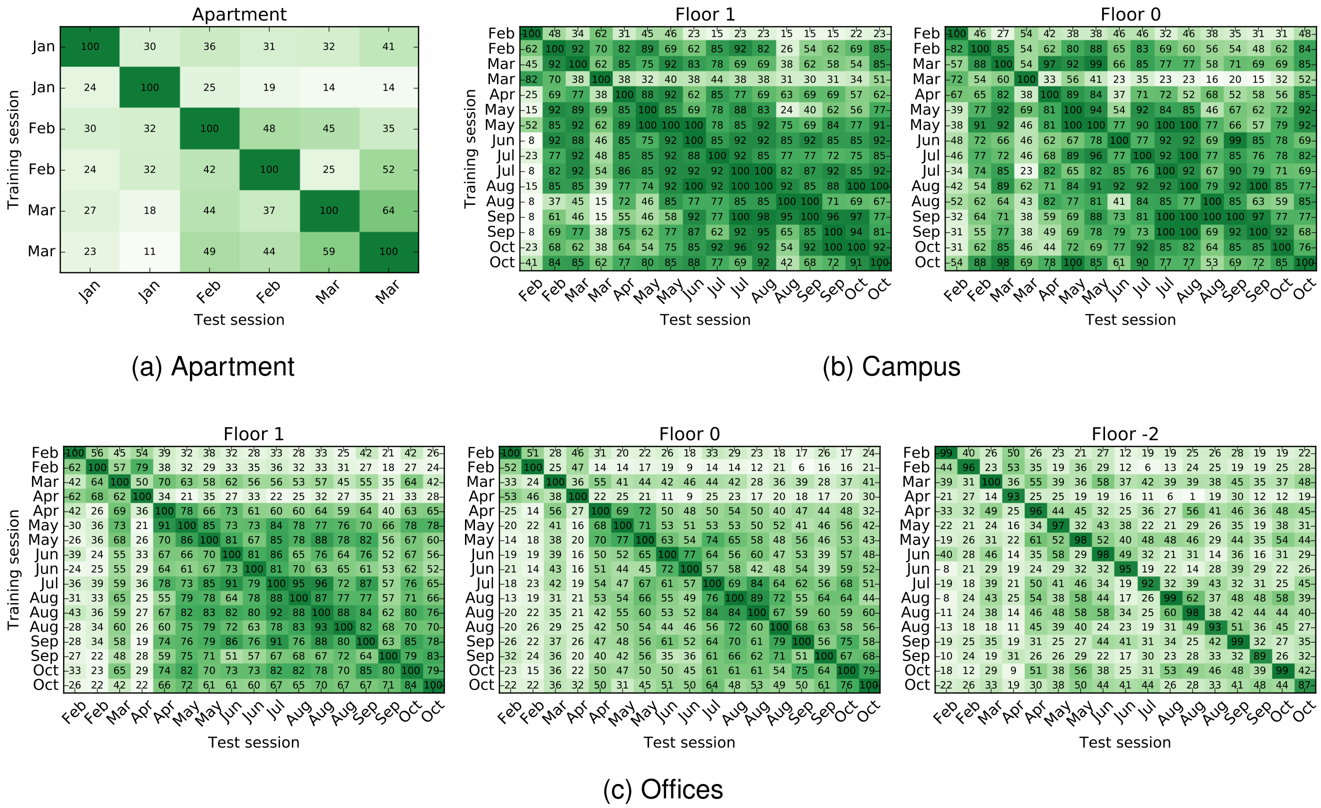

Figure 4. Positioning performance over time, with pairwise evaluation. Numbers in the cells are localization accuracy in percents. Random-guess levels are 3% for Apartment, 8% for Campus, and 3%, 3%, and 6% for floors 1, 0 and -2 of the Offices testbed.

consistent gradient of performance degradation. Instead, as evidenced by eventual blank cells with low accuracy, the system seems to be more sensitive to specific environment conditions of the given session pair than to the age of calibration data.

High sensitivity to session conditions can be addressed by additional training. In order to test this, we replicate a practical scenario of incremental training, where the system is trained on several consecutive sessions and then repeatedly tested over the following months.

Detailed results of incremental training are shown in Figure 5. Here, each row represents an evaluation history of the system trained on N first sessions and tested on the remaining ones (in contrast to Figure 4, where each row corresponds to a single training session). From top to bottom rows, the system is provided with more and more calibration data, which improves the localization accuracy.

This improvement is evident in Figure 6, which shows row-wise averages of Figure 5 — that is, the average accuracy of the system trained on N first sessions. As the figure shows, the results vary among the testbeds. In the campus building, one additional session almost doubled the accuracy (see Figure 6b). The Apartment and Offices testbeds, in turn, required longer training (around five sessions), but also managed to double the initial results.

Long-term observations of the Campus and Offices testbeds also show that after some training, system performance saturates at a certain level (plateau-like intervals in Figure 6b and 6c). On Campus testbed, the system achieves almost perfect accuracy, while for the Offices building the saturation levels vary by the floor — an effect which is explained by the different FM reception quality on these floors (cf. Figure 3).

4.4. RSS stability and human presence

Robust localization performance requires high stability of received signals. However, human presence in the building may dynamically change spatial patterns of RSS distribution due to wave scattering and reflection by human body [8], [22]. This section investigates the impact of human presence on ambient FM signals.

In order to characterize signal stability during a measurement session s, we first center RSS values of each active FM channel ch and each test point pt in that session, by subtracting their all-time average across all sessions:

$$x'_{s,pt,ch} = x_{s,pt,ch} - \langle x_{s,pt,ch} \rangle \Big|_s$$

These centered RSS values are then averaged across all the test points and active channels of the session. This 'centered

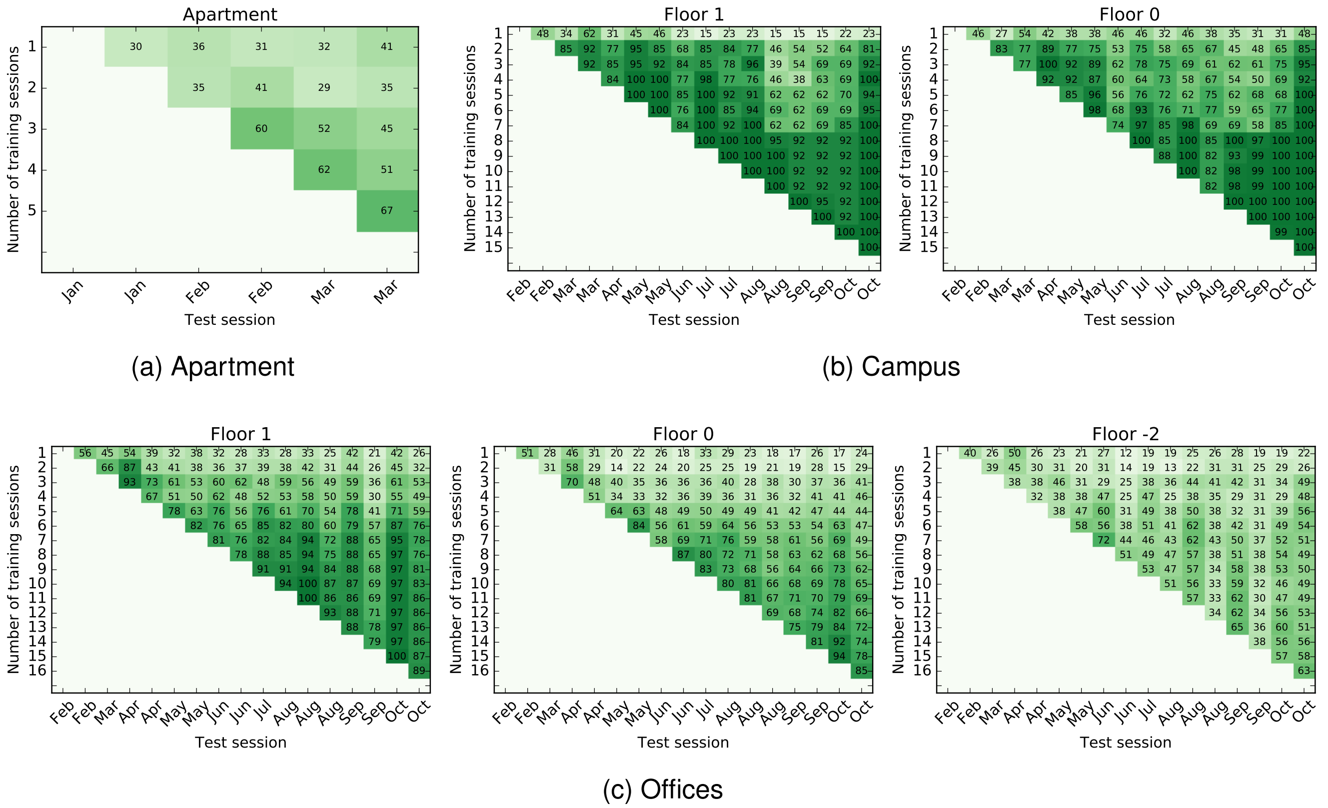

Figure 5. Positioning performance over time, with incremental training. Numbers in the cells are localization accuracy in percents. Random-guess levels are 3% for Apartment, 8% for Campus, and 3%, 3%, and 6% for floors 1, 0 and -2 of the Offices testbed.

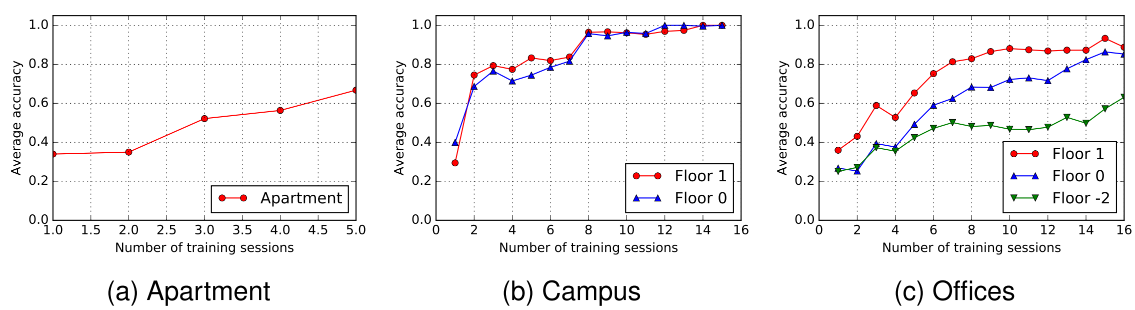

Figure 6. The effect of incremental training.

average' represents the *systematic bias* of RSS values in the session:

$$X_s = \langle x_{s,pt,ch} \rangle \Big|_{pt,ch}$$

In turn, fast *random fluctuations* of RSS are characterized in a similar way, by averaging standard deviations of x' values:

$$\sigma_s = \langle \sigma[x'_{s,pt,ch}] \rangle \Big|_{pt,ch}$$

The impact of human presence on FM RSS stability is shown in Figure 7. As expected by the propagation theory, RSS fluctuations in a populated building are significantly higher than in an empty one. A notable exception is the underground parking on Floor -2: although there was very little human activity, the level of RSS noise has always been relatively high due to strong attenuation by the ground and only few detectable FM channels.

Figure 7 also shows slight systematic RSS differences between sessions. This can be explained by slow large-scale environment changes, from internal restructuring (building new walls) to precipitation-caused changes of soil conductivity around the building. In absolute terms, however, both random and systematic RSS variations were typically less than 1 dB, which should not have any significant effect on the localization accuracy.

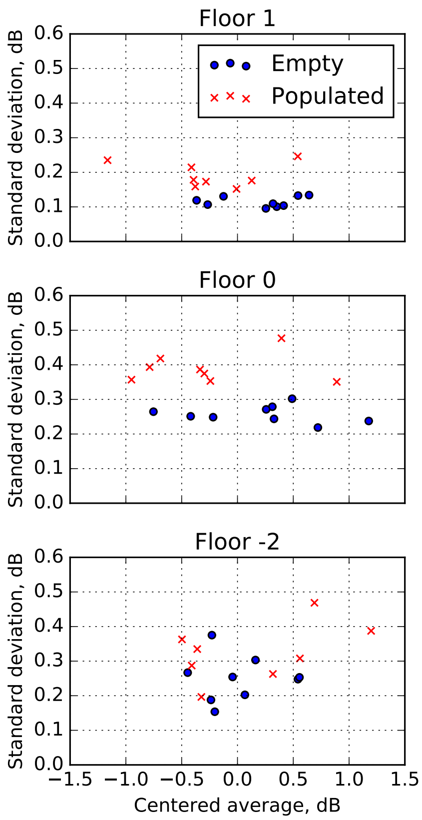

Figure 7. Impact of human presence on FM RSS statistics (Offices testbed). Each point represents one measurement session.

5. Conclusion

We presented a longitudinal study of FM indoor positioning system based on RSS fingerprinting. The experiments are based on a large dataset of radio samples, collected bi-weekly over a 9-month period in several testbeds with different environment conditions.

Our long-term evaluation demonstrates that while instantaneous FM localization performance varies due to environment dynamics, these variations can be substantially reduced by additional training. For instance, after only two training sessions, the system correctly recognized on average 70% of fingerprints over the following 8 months (Campus testbed). After six more training sessions, the average localization accuracy surpassed 96%.

Overall, the results show that with sufficient training and good radio reception, ambient FM signals can provide robust indoor localization for weeks and months after calibration.

Acknowledgments

This work was supported by the National Research Fund of Luxembourg (project INDOORS, grant number C14/IS/8311593).

References

[1] R. Faragher and R. Harle, “Location fingerprinting with Bluetooth Low Energy beacons,” *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 11, pp. 2418–2428, 2015.

[2] Z. Yang, Z. Zhou, and L. Yunhao, “From RSSI to CSI: Indoor localization via channel response,” *ACM Computing Surveys*, vol. 46, no. 2, p. 25, Nov. 2013.

[3] “Ubisense: Precise real-time location,” http://www.ubisense.net.

[4] P. Lazik, N. Rajagopal, O. Shih, B. Sinopoli, and A. Rowe, “ALPS: A bluetooth and ultrasound platform for mapping and localization,” in *Proceedings of SenSys-2015*. ACM, 2015, pp. 73–84.

[5] R. Harle, “A survey of indoor inertial positioning systems for pedestrians,” *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 3, pp. 1281–1293, 2013.

[6] A. Papliatseyeu, N. Kotilainen, O. Mayora-Ibarra, and V. Osmani, “FINDR: Low-cost indoor positioning using FM radio,” in *Proceedings of MobilWare-2009*, Apr. 2009, pp. 15–26.

[7] A. Matic, A. Popleteev, V. Osmani, and O. Mayora, “FM radio for indoor localization with spontaneous recalibration,” *Journal of Pervasive and Mobile Computing*, vol. 6, pp. 642–656, 2010.

[8] A. Popleteev, “Indoor positioning using FM radio signals,” Ph.D. dissertation, University of Trento, 2011.

[9] V. Moghtadaiee, A. Dempster, and S. Lim, “Indoor localization using FM radio signals: A fingerprinting approach,” in *Proceedings of IPIN-2011*, 2011.

[10] A. Popleteev, V. Osmani, and O. Mayora, “Investigation of indoor localization with ambient FM radio stations,” in *Proceedings of PerCom-2012*. IEEE, 2012, pp. 171–179.

[11] Y. Chen, D. Lymberopoulos, J. Liu, and B. Priyantha, “FM-based indoor localization,” in *Proceedings of MobiSys 2012*, 2012.

[12] S. Yoon, K. Lee, and I. Rhee, “FM-based indoor localization via automatic fingerprint DB construction and matching,” in *Proceeding of MobySys-2013*. ACM, 2013, pp. 207–220.

[13] P. Bahl and V. N. Padmanabhan, “RADAR: An in-building RF-based user location and tracking system,” in *Proceedings of INFOCOM-2000*, vol. 2, 26-30 March 2000, pp. 775–784.

[14] L. W. Barclay, Ed., *Propagation of Radiowaves*, 2nd ed. The Institution of Engineering and Technology, 2002.

[15] A. Matic, A. Papliatseyeu, V. Osmani, and O. Mayora-Ibarra, “Tuning to your position: FM-radio based indoor localization with spontaneous recalibration,” in *Proceedings of PerCom-2010*. IEEE, Apr. 2010, pp. 153–161.

[16] A. Popleteev, “HIPS: Human-based indoor positioning system,” in *Proceedings of IPIN-2016*. IEEE, 4–7 Oct. 2016.

[17] ——, “Poster: Impact of ground truth errors on Wi-Fi localization accuracy,” in *Proceedings of MobiSys-2017*, 2017.

[18] ——, “Indoor positioning using ambient radio signals: Data acquisition platform for a long-term study,” in *Proceedings of WPNC-2016*, 2016.

[19] Ettus Research, “USRP B-210 USB software defined radio (SDR),” Online: https://www.ettus.com/product/details/UB210-KIT.

[20] “GNU radio toolkit,” http://gnuradio.org.

[21] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[22] K. Kaemarungsi, “Distribution of WLAN received signal strength indication for indoor location determination,” in *Proceedings of ISWPC-2006*, 2006, pp. 1–6.