An RSSI-based Wall Prediction Model for Residential Floor Map Construction

Xenofon Fafoutis, Evangelos Mellios,
Niall Twomey, Tom Diethe, Geoffrey Hilton and Robert Piechocki
Department of Electrical and Electronic Engineering
University of Bristol
Email: {xenofon.fafoutis, evangelos.mellios, niall.twomey, tom.diethe, geoff.hilton, r.j.piechocki}@bristol.ac.uk

***Abstract*—In residential environments, floor maps, often required by location-based services, cannot be trivially acquired. Researchers have addressed the problem of automatic floor map construction in indoor environments using various modalities, such as inertial sensors, Radio Frequency (RF) fingerprinting and video cameras. Considering that some of these techniques are unavailable or impractical to implement in residential environments, in this paper, we focus on using RF signals to predict the number of walls between a wearable device and an access point. Using both supervised and unsupervised learning techniques on two data sets; a system-level data set of Bluetooth packets, and measurements on the signal attenuation, we construct wall prediction models that yield up to 91% identification rate. As a proof-of-concept, we also use the wall prediction models to infer the floor plan of a smart home deployment in a real residential environment.**

***Keywords*-Wall prediction model; floor map construction; wearable systems; bluetooth low energy**

I. Introduction

The ubiquity of smart phones has triggered the development of various localisation services and location-based applications (e.g. [1]). Maps are often required by such applications to define the location in an absolute scale. While Global Positioning System (GPS) and map providers, such as Google and Microsoft, have simplified this problem in outdoors environments (see [2]); acquiring indoor floor plans is a challenging task. Occasionally, for few important public places, map providers also provide manually-constructed indoor floor maps.

In the most recent years, the problem of automatic indoor floor plan construction attracted attention in the research community with an substantial increase of the related literature. Such floor plan construction systems leverage various sources of information; primarily inertial sensors, but also cameras and Radio Frequency (RF) signals. Simultaneous Localisation and Mapping (SLAM) is a technique that is inherited from the robotics community [3] and used for tracking in unknown environments. SLAM is applied to the problem of indoor floor plan construction, exploiting the existence of inertial sensors, such as accelerometers and gyroscopes, in smart phones [4][5][6][7][8]. The core idea is to perform odometry to trace the movement of a user in an unknown environment while updating, in parallel, a map based on visited locations.

Inertial sensors do not provide absolute information. Hence, indoor floor map construction techniques often fuse other modalities, such as cameras [8], and RF signals [6], typically originating from WiFi local area networks. RF fingerprinting is the state-of-the-art technique for indoor localisation [9]. Upon the collection of a labelled data set, fingerprinting techniques identify correlations between locations and radio-signal patterns. In public places, the collection of fingerprints, inertial or other types of data can be crowed-sourced [2][5][6][8] through applications in the smart phones of passers-by.

In this paper, we consider the problem of automatic indoor floor map construction in a residential environment. In the context of smart residential environments for eHealth applications, for instance, floor map construction can assist location-based services that identify anomalies in the residents' daily routine [10], assist patients with moving difficulties [11] or prevent falls [12]. Envisioning the future of smart homes and the scalable adoption of the Internet of Things (IoT), users with little technical expertise will be expected to deploy IoT-enabling plug-and-play technologies. Without experts deploying the system, the floor map of the environment cannot be simply assumed by higher layers of the system (*i.e.* ambient intelligence algorithms). To make matters worse, users may potentially deploy their hardware in suboptimal locations.

The vast majority of the related work, assumes a public indoor environment, such as a university building, a shopping mall, or an office [4][5][6][8]. In contrast to public spaces, a residential environment has properties that constrain the applicability of some of the proposed solutions. WiFi fingerprinting, for instance, requires the collection of labelled fingerprints and, in the case of a residential environment, is typically limited to a single WiFi access point. Crowd-sourcing is not an option. Furthermore, the usefulness of smart phones (used in [4][5][6][7][8]) is also limited as, unlike public environments, it cannot be assumed that users will carry them continuously.

Assuming the use of wearable sensors instead of smart phones, in this paper, we identify and propose an alternative use of RF signals in the problem of indoor map construction. Upon a wireless propagation campaign in a residential environment, we observed that the nature of this environment (*i.e.* small rooms, thick walls) gives unique properties to the signal attenuation. As shown in Section II, most of the signal attenuation occurs because of propagation through walls, rather than because of the distance within the same room or the body shadow. As a result, the Received Signal Strength

Indicator (RSSI) at the receiver contains information about the number of walls that the signal traversed through. In smart homes with multiple receivers an RSSI-based wall prediction model can be used, eventually together with other modalities (*e.g.* cameras or inertial sensors), for floor map construction.

In fact, our previous measurement campaigns using the Sensor Platform for Healthcare in Residential Environments (SPHERE) infrastructure [13] suggest that a Bluetooth Low Energy (BLE) wearable sensor requires multiple receiver units for full-house coverage. Therefore, with regard to the infrastructure, the integration of an RSSI-based wall prediction model for floor map construction comes at no additional overhead. Furthermore, it introduces no overheads to the energy-constrained wearable sensor.

In addition to indoor floor map construction, which is the focus of this paper, an RSSI-based wall prediction model can be potentially used by IoT service providers to identify cases where the users have deployed access points in sub-optimal locations (*e.g.* two access points in the same room). Furthermore, it can give spatial context to other sensing modalities (*e.g.* identify whether sensors are deployed in different rooms).

It should be noted that, contrary to related work that uses inertial sensors, a wall prediction model cannot estimate the dimensions of the rooms. Instead, it identifies wall relationships between them. Therefore, this work complements the state-of-the-art with an additional modality. In this paper we are interested in identifying the potential of the wall prediction model itself. Thus, the use of our platforms' accelerometers is out of our scope and considered future work.

In short, the contribution of this work can be summarised as follows. First, we identify that in certain types of residential environments the RSSI contains information on the number of walls between the transmitter and the receiver. We compare the accuracy of various learning techniques on building an RSSI-based wall prediction model, considering both supervised and unsupervised solutions. We validate the robustness of this model to the locations of the nodes within the rooms, and to different platforms, by training it with measurements on the channel attenuation of different links. Lastly, we apply the predictions to the case of the SPHERE infrastructure [13] and evaluate its accuracy on constructing the floor map of the SPHERE house.

The remainder of the paper is structured as follows. Section II provides insight on the signal propagation patterns in a residential environment and on how they can be leveraged for determining the number of walls in a link. In Section III, we use different learning techniques to create models that estimate the number of walls. Section IV evaluates the performance of the considered algorithms using both system-level RSSIs and channel gain measurements. Section V applies the wall prediction model to the problem of residential floor map construction. Lastly, Section VI concludes the paper.

II. Motivating Experiment

In this section, we present results of a measurement campaign in a residential environment, using the SPHERE infrastructure [13], in an attempt to demonstrate the RF signal attenuation patterns and provide intuition on how they can be leveraged for predicting the existence of a wall. In this measurement campaign, we use the body-centric sensing hardware of [13] that includes a BLE-enabled wearable sensor. This set of experiment considers various links within a residential environment. In each scenario, a person was rotating 360^o (in 4 steps of 90^o), holding the wearable device in front of his body.

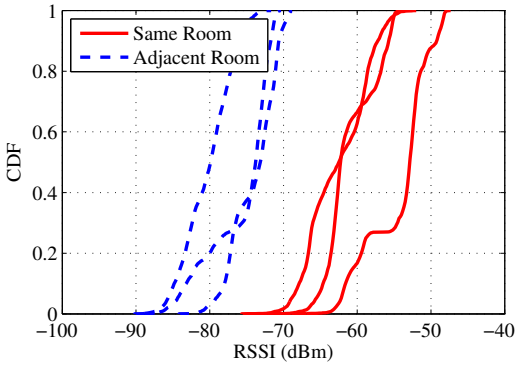

Fig. 1. CDFs of the RSSI of all the packets received for 3 links, as measured by a receiver in the same and adjacent room.

Fig. 1 plots the cumulative distribution functions (CDFs) of the RSSI of all the packets received for 3 different wireless links[1], as measured by a receiver in the same room and in a room that is separated by one wall. Observe that there is a dynamic range in the RSSI that is mainly caused by the body shadow of the user (approximately 20 dB in the same room and 15 dB in the adjacent room). Despite the effect of the body shadow, and despite the differences of the links (user and receiver unit locations), observe that the CDFs of the RSSI of the packets measured by the receivers, in the same and adjacent room, do not overlap significantly; indicating that the RSSI contains information that can be used for predicting the existence of a wall.

III. Wall Prediction Model

An RSSI-based wall prediction model is, essentially, equivalent to a function that transforms the RSSI, P_{RSSI}, to a number of walls, $f : \mathbb{R} \rightarrow \mathbb{N}$. Given the nature of wave propagation, f is defined by a sequence of RSSI thresholds, $t_1, t_2, \ldots t_K$, that identify the limits of a room.

$$f(P_{RSSI}) = \begin{cases} 0 & : P_{RSSI} < t_1 \\ 1 & : t_1 \leq P_{RSSI} < t_2 \\ \ldots & \\ K-1 & : t_{K-1} \leq P_{RSSI} < t_K \\ K & : t_K \leq P_{RSSI} \end{cases} \quad (1)$$

K depends on the sensitivity threshold of the receiver, the transmission power of the transmitter and the antennas used; and defines an upper limit on the number of walls that can be

[1] Additional details on the measurements and the links in Section IV-A.

traversed. As a matter of fact, the RSSI thresholds are dictated by the link budget formula.

Since a wall cannot amplify the radio signal, the RSSI thresholds are also dictated by the following constraint.

$$t_k < t_{k-1} : k \in [1, K] \tag{2}$$

In this paper, we use three algorithms to generate the RSSI thresholds, $t_1, t_2, \dots t_K$, namely Supervised Medians, Support Vector Machines (SVMs) and K-Means.

The first algorithm, Supervised Medians, assumes a set of labelled data. Let $\tilde{M}_k$ be the median RSSI of the measurements that correspond to k walls. We define the threshold t_k as the average of the corresponding medians.

$$t_k = \frac{\tilde{M}_{k-1} + \tilde{M}_k}{2} : k \in [1, K] \tag{3}$$

The SVM [14] is a supervised classifier that leverages training data to identify a hyperplane that separates them into two classes. As with Supervised Medians, it requires a set of labelled data input. The RSSI thresholds are derived from multiple SVM classifiers, which are based on the radial basis function (RBF) kernel (configured with $c = 2$ and $\sigma = 1$), and used in a variation of the one-versus-all fashion. Assuming M_k is a set of RSSI measurements that correspond to k walls, $\forall\ k \in [1, K]$, each threshold t_k is estimated by an SVM classifier that is trained with data from sets $\bigcup_{i\in[0,k-1]} M_i$ and $\bigcup_{j\in[k,K]} M_j$. Before training, the input data are centered to their mean and scaled to have a unit standard deviation. We did not observe any significant variance of the results based on the configuration parameters c and σ.

K-Means [15] is an algorithm that clusters data into K groups. Starting with K arbitrary centroids, the algorithm iterates between assigning observations to the closest centroid and updating the centroids to the mean of the clustered observations. Contrary to the previous methods, K-Means is an unsupervised algorithm that operates on unlabelled data. This is of particular value in our application of interest. Unlike the two previous methods that require the predicted RSSI thresholds to be generalisable to an unknown new residential environment, K-Means constructs new thresholds for each new environment. In this application, K derives from the upper limit on the amount of walls. The output of K-Means is $K + 1$ centroids. Using (2), we correspond each centroid to a number of walls by sorting the centroids in descending order. Assuming $C_0, C_1 \dots C_K$ is the sorted sequence of centroids, the RSSI thresholds are defined as follows.

$$t_k = \frac{C_{k-1} + C_k}{2} : k \in [1, K] \tag{4}$$

Next, we evaluate the performance of the three aforementioned models on predicting the number of walls.

IV. Model Evaluation

The evaluation of the RSSI-based wall prediction model is twofold. We experiment with training the model with both system-level and channel-level measurements. For the system-level training, we use the SPHERE infrastructure [13], which is based on BLE for wireless communications, to collect a data set of RSSIs for various links in the SPHERE house (shown in Fig. 2). This data set is split into two; one is used for training the models and the other is used for testing them. To identify the resilience of the model to different access point and user locations, and identify the importance of the particular platform, we also train the model using data from a separate channel-level measurement campaign. This data set is, essentially, measurements of the channel gain using continuous waves. Considering that in the end platform only RSSI values will be available, the model is tested using the system-level measurements.

A. Training with System-Level Measurements

For the collection of the system-level RSSI measurements, we use the prototype SPHERE infrastructure [13] that is deployed to a residential property in the city of Bristol, UK. The SPHERE infrastructure incorporates a BLE-enabled wearable sensor that broadcasts to an infrastructure. The wearable sensor employs a low profile patch antenna (dimensions: $17.8 \times 18.5 \times 1.3$ mm, efficiency: 55%, maximum directivity at 2440 MHz: 6.8 dBi) [16]. The system achieves high energy-efficiency by using BLE in connectionless mode (*i.e.* broadcasting mode). To provide reliable full-house coverage, three BLE receivers are deployed in the SPHERE house. Each BLE receiver incorporates two radios that employ two orthogonally polarised dipole antennas, working in parallel. For details on the receiver antennas we refer the reader to [17]. We found empirically that the system can support up to 2-wall links ($K = 2$).

Fig. 2 depicts the floor plan of the SPHERE house, denoting the locations of the access points (APx) and the user with the wearable sensor (Wx) that are considered in this paper. For the system-level measurements, the three access points are deployed in locations AP3, AP5 and AP6. The RSSI measurement campaign considers 5 user locations (W1, W3, W4, W6 and W8); that is 15 links in total. In fact, the links shown in Fig. 1 are part of this data set. As explained in Section II, in each case, the user was rotating 360^o (in 4 steps of 90^o), holding the wearable device in front of his body.

Prior to training the various models, we process the raw RSSI data as follows. First, we average the RSSI values of the packets that are received in both orthogonally polarised antennas. As a user changes his orientation to the access point, the antenna that is aligned with transmitting antenna on the wearable yields a higher RSSI and vice versa. Therefore, the average of the two values decreases the dynamic range of the RSSI values. The raw RSSI data are also passed through a low pass filter, *i.e.* a moving average with a window of w. The moving average improves the performance of all the considered training algorithms, primarily by filtering the RSSI variations caused by fast-fading. For the algorithms that require labelled data, Supervised Medians and SVM, labelling is conducted as follows: AP3-W3, AP3-W1, AP5-W6 and

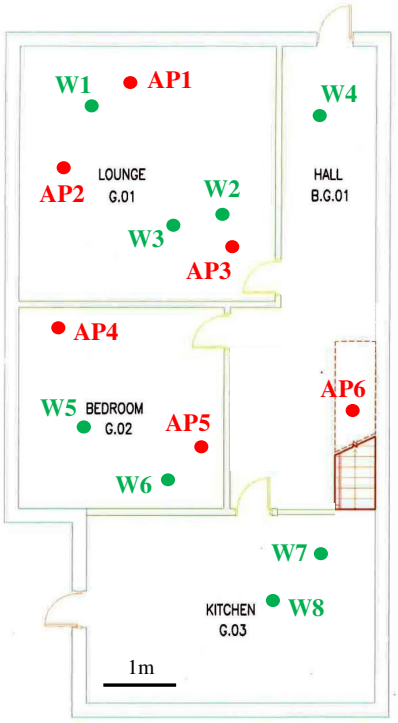

Fig. 2. SPHERE house floor plan.

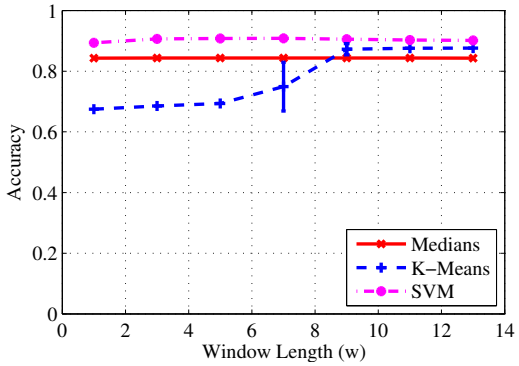

Fig. 3. Wall identification rate with 50% of the system-level measurements for training and 50% for testing (unless plotted, $\sigma < 0.006$).

AP6-W4 are 0-wall links; AP3-W6, AP3-W4, AP5-W1, AP5-W3, AP5-W4, AP5-W8 and AP6-W8 are 1-wall links; and AP3-W8 is a 2-wall link.

Fig. 3 shows the average wall identification rate of the three considered algorithms averaged over 100 repetitions and for various values of the window w. In each repetition, the model is reconstructed using a new randomly selected 50% of the data set. Then, the model is tested on the other half of the data set. We observe that the window of the rolling average primarily affects the performance of K-Means that maximises the wall identification rate at a window of 9 samples (87%). Less affected by the rolling average, SVM yields the best performance (90%). Lastly, the performance of Supervised Medians is at 84%.

As indicated in the previous section, K-Means has particular practical value. Contrary to supervised algorithms, which have practical value only if their model can be generalised to an unknown scenario, K-Means constructs a unique model for each particular scenario, in an unsupervised manner. Moreover, the results, presented in this section, suggest that its performance is comparable to the supervised alternatives.

B. Training with Channel-Level Measurements

In this section, we are interested in generalising the results of the previous section to different links and platforms. In particular, we train the models using raw data from a channel measurement campaign that is protocol-independent; *i.e.* uses continuous waves generated by a Vector Network Analyser (VNA). We then test the performance of these models on the system-level RSSI measurements of Section IV-A.

In the channel-level measurements, we measure the channel gain in the 2.4 GHz band using a VNA. It should be clarified that, in addition to the path loss, the measurements include the transmit and receive antenna gains. All channel measurements were performed with the same patch antenna mounted on the wrist of a human subject, who was sitting on a stool, which was mounted on a turntable and rotated in azimuth through 360° with the help of motors (with a step of 30°). Each measurement was conducted twice; once with the arm parallel to the body and once with the arm orthogonal to the body. On the receiver end, the same orthogonally polarised monopole antennas were deployed in locations AP1, AP2 and AP4, as shown in Fig. 2. The user locations considered are W2, W5 and W7; a total of 12 links considering the 2 arm positions.

Before training the models, the raw channel gain data is transformed into RSSI. First, we account for the transmission power used by the end platform, by adding 4 dBm to the channel gain data set. Indeed, the addition of the transmission power to the channel gain estimates the received signal strength. We then account for the error of the RSSI measurements reported by the BLE platform. To measure the error in a controlled environment, we connect the BLE platform directly to a spectrum analyser, using low-loss wires and 30dB attenuators. Fig. 4 reports the results, which we fit with a quadratic function.

Similarly to the system-level data set, we then average the values that correspond to the two orthogonally polarised antennas and pass the data through a low pass filter (with a window of w). For the algorithms that require labelled data, Supervised Medians and SVM, labelling is conducted as follows: AP1-W2, AP2-W2 and AP4-W5 are 0-wall links; AP1-W5, AP2-W5, AP4-W2 and AP4-W7 are 1-wall links; and, AP1-W7 and AP2-W7 are 2-wall links.

Fig. 5 shows the average wall identification rate of the three considered algorithms averaged over 100 repetitions and for various values of the window w. In each repetition, the model is reconstructed using a subset of randomly selected channel-level measurements. Then, the model is tested on the system-level data set. We observe that all three algorithms yield

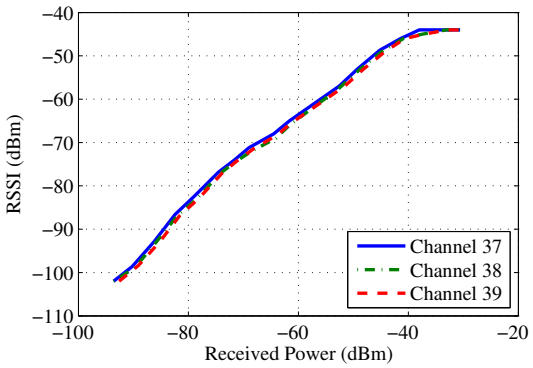

Fig. 4. Mapping the RSSI of nRF51822 to received signal power.

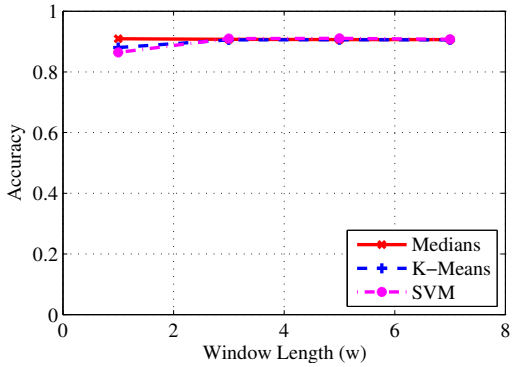

Fig. 5. Wall identification rate with the channel-level measurements for training and the system-level measurements for testing ($\sigma < 0.006$).

similar performance (91%) and are marginally affected by the window of the rolling average filter. Interestingly, for some cases, the performance of a model trained with channel gain measurements performs better than the respective experiment of Section IV-A. This phenomenon is explained by the fact that the channel-level measurements is a larger data set that captures better the effect of multi-path propagation and body shadowing with finer rotations of 30° instead of 90°. The results further demonstrate that the wall prediction models are resilient to the exact locations of the access points and the exact location of the user.

Lastly, the results presented in this section, suggest that the adaptation of the wall prediction models to another BLE radio or another wireless communications protocol is possible, assuming the deviation of the reported RSSI from a known point of reference is measured and considered in the calculations.

V. Floor Plan Construction

As a use case scenario, we apply the RSSI-based wall prediction model of Section III to the problem of residential floor plan construction. The algorithm operates as follows. A transmitted packet is received by multiple access points, either correct or with channel errors. Assuming a number of access points (N), we start by assigning to the access points an arbitrary label from 1 to N. A given vector of RSSI values, sorted by the label of the respective access point, is passed through (1) and translated into a vector that counts the walls between the user and the respective access point. We consider two types of such vectors. The vectors that contain a zero are measurements where the user is estimated to be in the same room with an access point. We use these vectors to identify the wall relationships between access points. In particular, we calculate the number of walls between access point i and j, R_{ij}, as the average absolute difference between the wall predictions for i and j on all the vectors that contain at least one 0-wall prediction. To construct the floor plan, the R_{ij} matrix is used as a set of constraints. Each access point, i, is sequentially placed on a hexagon grid in the first position that satisfies all the wall relationships, R_{ij}, of the access points placed in previous iterations, *i.e.* $j < i - 1$. Wall prediction vectors that do not contain a 0-wall prediction are cases where the user is estimated to be in room that does not have an access point. A new room label is assigned to each unique vector like that (labelling continues from $N + 1$). Such rooms are placed in the hexagon grid in the first position that satisfies all the established wall relationship constraints.

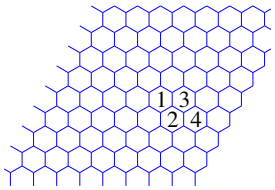

Fig. 6. Floor map prediction for the SPHERE house.

The performance of the different wall prediction models on floor map construction is evaluated next. In particular, we use the data set described in Section IV-A, which employs three access points ($N = 3$). It should be noted here that in deployments with fewer access points per room, certain rooms may demonstrate the same wall relationship profiles and cause confusion. Nevertheless, in such cases, full-house wireless coverage becomes a more important challenge [13]. In line with the previous experiments, a different set of samples is used for training and for testing. Fig. 6 demonstrates the outcome of the floor plan construction algorithm for the case of the SPHERE house (Fig. 2) in the case of a correct prediction. Arbitrarily, the algorithm assigns the following labels: 1 to the lounge (AP3), 2 to the bedroom (AP5) and 3 to the hall (AP6). Label 4 is, then, assigned to the kitchen.

Fig. 7 shows the accuracy of the three considered models on predicting the floor map of the SPHERE house over 10000 attempts. Similarly to Section IV, in every attempt, the model is reconstructed on a different randomly selected training set and the map is constructed with a random subset of the remaining test set. The test set is also passed through a low pass filter (*i.e.* rolling average with a window of w). The accuracy, shown in Fig. 7, is based on a binary evaluation, where any prediction different to Fig. 6 is considered incorrect. Alternatively, a score-based evaluation may be more appropriate in scenarios

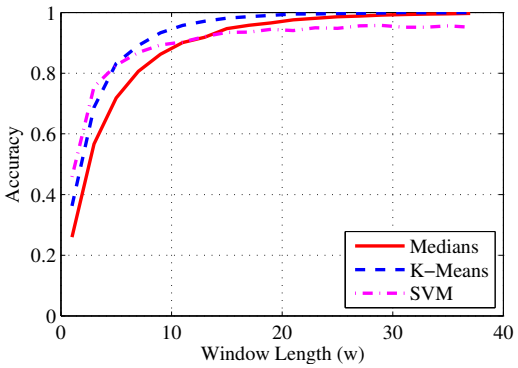

Fig. 7. Floor map construction accuracy using the system-level measurements for training the wall prediction model.

where wall identification models are combined with other modalities for floor map construction. We observe that all three algorithms improve their accuracy with a higher w with K-Means converging earlier at predicting the correct floor plan (shown in Fig. 6). The SVM converges at a 95% recognition rate, which indicates that the high mean wall identification rate shown in Fig. 3 is not evenly distributed in all classes.

VI. Conclusions

In this paper, we construct a RSSI-based wall prediction model that leverages the propagation patterns of RF signals in residential environments with small rooms and thick walls, to estimate the number of walls between a user (*i.e.* wearable sensor unit) and a receiver unit. We construct wall prediction models using three different learning algorithms, namely Supervised Medians, SVMs and K-Means. K-Means, which has particular practical value due to the fact that it generates unique models for each separate case in an unsupervised manner, yields a wall identification rate of 87%. For comparison, the supervised SVM yields the best wall identification performance (90%). To explore the resilience of the model to different wireless protocols and different links, we use data from a channel gain measurement campaign to construct the models. Testing these models on the system RSSI measurements yields similar results, *i.e.* a maximum performance of 91%. Lastly, we apply the wall prediction model to predict the floor plan of a residential environment.

Given the promising results on a proof-of-concept scenario, our future work follows two directions. The first is to generalise and validate the concept on more residential floor plans, including 2-storey houses. Indeed, the plan of our research team for the following years is to deploy the SPHERE technologies [13] to the houses of 100 volunteers in the region of Bristol, UK. Such a deployment will give us the opportunity to test the concepts presented in this paper on a larger scale. The second direction is to fuse the wall prediction models with other sensing modalities in a smart home. This would enable assigning contextual labels to the rooms instead of unique numbers. For instance, combining the prediction of an activity recogniser (*e.g.* watching TV) with localisation information (*e.g.* room 1), would allow us to assign a contextual label to the respective room (*e.g.* lounge).

Acknowledgment

This work was performed under the SPHERE IRC funded by the UK Engineering and Physical Sciences Research Council (EPSRC), Grant EP/K031910/1.

References

[1] A. Leonidis, G. Baryannis, X. Fafoutis, M. Korozi, N. Gazoni, M. Dimitriou, M. Koutsogiannaki, A. Boutsika, M. Papadakis, H. Papagiannakis, G. Tesseris, E. Voskakis, A. Bikakis, and G. Antoniou, "Alertme: A semantics-based context-aware notification system," in *Proc. 33rd Int. Annu. IEEE Comput. Soft. and Appl. Conf. (COMPSAC)*, vol. 2, July 2009, pp. 200–205.

[2] U. Blanke, R. Guldener, S. Feese, and G. Troester, "Crowdsourced Pedestrian Map Construction for Short-Term City-Scale Events," in *Proc. First Int. Conf. IoT Urban Sp.* ICST, Oct. 2014, pp. 25–31.

[3] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.

[4] M. Alzantot and M. Youssef, "CrowdInside: Automatic Construction of Indoor Floorplans," in *Proc. 20th Int. Conf. Adv. Geogr. Inf. Syst. - SIGSPATIAL '12*. ACM Press, Nov. 2012, p. 99.

[5] H. Shin, Y. Chon, and H. Cha, "Unsupervised Construction of an Indoor Floor Plan Using a Smartphone," *IEEE Trans. Syst. Man, Cybern. Part C (Applications Rev.*, vol. 42, no. 6, pp. 889–898, Nov. 2012.

[6] Y. Jiang, Y. Xiang, X. Pan, K. Li, Q. Lv, R. P. Dick, L. Shang, and M. Hannigan, "Hallway based automatic indoor floorplan construction using room fingerprints," in *Proc. 2013 ACM Int. Jt. Conf. Pervasive ubiquitous Comput. - UbiComp '13*. ACM Press, Sep. 2013, p. 315.

[7] M. Hardegger, D. Roggen, and G. Tröster, "3D ActionSLAM: wearable person tracking in multi-floor environments," *Pers. Ubiquitous Comput.*, vol. 19, no. 1, pp. 123–141, Sep. 2014.

[8] R. Gao, M. Zhao, T. Ye, F. Ye, Y. Wang, K. Bian, T. Wang, and X. Li, "Jigsaw: indoor floor plan reconstruction via mobile crowdsensing," in *Proc. 20th Annu. Int. Conf. Mob. Comput. Netw. - MobiCom '14*. ACM Press, Sep. 2014, pp. 249–260.

[9] R. Faragher and R. Harle, "An Analysis of the Accuracy of Bluetooth Low Energy for Indoor Positioning Applications," in *Proc. 27th Int. Tech. Meet. Satell. Div. Inst. Navig. (ION GNSS+)*, 2014, pp. 201–210.

[10] M. D'Souza, M. Ros, and M. Karunanithi, "An Indoor Localisation and Motion Monitoring System to Determine Behavioural Activity in Dementia Afflicted Patients in Aged Care," *Electron. J. Heal. Informatics*, vol. 7, no. 2, p. 14, Jan. 2012.

[11] S. Mazilu, M. Hardegger, Z. Zhu, D. Roggen, G. Troster, M. Plotnik, and J. Hausdorff, "Online detection of freezing of gait with smartphones and machine learning techniques," in *Proc. 6th Int. Conf. Pervasive Comput. Technol. Healthc.*, 2012, pp. 123–130.

[12] F. Sposaro and G. Tyson, "iFall: an Android application for fall monitoring and response." *Conf. Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. IEEE Eng. Med. Biol. Soc. Annu. Conf.*, vol. 2009, pp. 6119–6122, Jan. 2009.

[13] P. Woznowski, X. Fafoutis, T. Song, S. Hannuna, M. Camplani, L. Tao, A. Paiement, E. Mellios, M. Haghighi, N. Zhu, G. Hilton, D. Damen, T. Burghardt, M. Mirmehdi, R. Piechocki, D. Kaleshi, and I. Craddock, "A Multi-modal Sensor Infrastructure for Healthcare in a Residential Environment," in *Proc. Int. Conf. Communications Workshops (ICCW)*, Jun. 2015, pp. 271–277.

[14] N. Cristianini and J. Shawe-Taylor, *An introduction to support vector machines and other kernel-based learning methods*. Cambridge university press, 2000.

[15] J. MacQueen, "Some methods for classification and analysis of multivariate observations," in *Proc. 5th Berkeley Symp. Math. Stat. Probab. Vol. 1 Stat.* The Regents of the University of California, 1967.

[16] M. W. Abdullah, X. Fafoutis, E. Mellios, M. Klemm, and G. Hilton, "Investigation into Off-Body Links for Wrist Mounted Antennas in Bluetooth Systems," in *Proc. Loughborough Antennas and Propagation Conf. (LAPC)*, 2015.

[17] E. Mellios, A. Goulianos, S. Dumanli, G. Hilton, R. Piechocki, and I. Craddock, "Off-body Channel Measurements at 2.4 GHz and 868 MHz in an Indoor Environment," in *Proc. 9th Int. Conf. on Body Area Networks (BODYNETS)*, 2014.