Easing the Survey Burden: Quantitative Assessment of Low-Cost Signal Surveys for Indoor Positioning

Chao Gao
Computer Laboratory,
University of Cambridge, UK

Robert Harle
Computer Laboratory,
University of Cambridge, UK

Abstract—Indoor positioning through signal fingerprinting is a popular choice since it requires little or no additional infrastructure. However, the initial creation and subsequent maintenance of the signal maps remains a challenge since traditional manual surveying is not scalable. In this work we look at the use of *path surveys*, which attempt to construct the signal maps from a sparse set of fingerprints collected while a person walks through a space. As such, the survey points rarely provide a uniform coverage of the space of interest. We quantitatively evaluate path surveys with reference to a detailed manual survey using smartphone-grade equipment. We compare both the individual maps (generated using Gaussian Process regression) and their collective positioning performance. Our results are for both WiFi and Bluetooth Low Energy signals. We show that a path survey can provide maps of equivalent quality to a manual survey if a series of guidelines we provide are followed.

I. INTRODUCTION

Indoor location systems have a rich history of technologies and techniques [18], [16], [2]. However, no single system has yet emerged to provide ubiquitous indoor positioning, primarily due to the need to deploy local infrastructure of some sort. The most prevalent systems are those that are able to leverage existing infrastructure, and of those most are based on applying signal fingerprinting techniques to WiFi signals [4]. Devices are provided with a pre-existing signal survey (map) for the building, which they use to position themselves. The fingerprinting approach in general is a flexible one that can be applied to any spatially-varying but temporally stable signal. Furthermore, the fingerprint maps may have uses other than positioning (for example to assist the deployment of WiFi access points to ensure uniform coverage and minimal overlap).

The key issue with fingerprinting is the creation of the signal maps (often called the *survey* or *offline* process) and their maintenance over the long term. Early work in this area used manual surveys, requiring a surveyor to visit each point on a regular grid and measure the signal fingerprint there. This is a laborious and time-consuming process that is a deterrent to the creation and use of such systems.

An emerging approach is the use of what we term *path surveys*. Here the surveying device is somehow tracked as it moves continuously through the environment and records signal measurements. Signal maps are then generated from signals collected along these paths rather than from a regular

grid of points. The paths themselves may come from an alternative location system, or be jointly estimated based on the signal to be mapped using a variant of a Simultaneous Localisation and Mapping (SLAM) algorithm. These maps can be seen as approximations to the conventional manual survey. To date, however, there is no quantitative and rigorous analysis about how well such path surveys approximate their manual equivalents, and no guidelines about how the path survey should be performed.

In this paper, we explore these issues in detail. We make four primary contributions:

1) We describe the collection of a dataset containing a detailed manual survey and ground-truthed path surveys.
2) We consider Bluetooth Low Energy in addition to the usual WiFi signal.
3) We use an accurate ground truth to compare best-case path surveys to manual surveys in terms of map quality and positioning performance, which gives the upper bound of real-world path survey performance.
4) We provide guidelines for path survey based on our quantitative analysis; We demonstrate that by following our guidelines, a path survey can achieve good efficiency and accuracy.

II. RELATED WORK

Fingerprinting for indoor location began with the RADAR system [4]. Since then there have been a plethora of research and commercial systems using indoor fingerprinting (e.g. [12], [26], [15], [21], [14], [22]). These systems have generally used WiFi signals but cellular [19], DECT [17], and magnetic field strength [5] have also been used.

Manual surveys result in a set of survey points distributed throughout the space of interest. These points can be used as-is, forming a *raw map*. Positioning is then achieved through a k-NN approach to find the best matching k survey points, where k typically lies in the range 1–5 [12].

An alternative approach is to use the survey data as input to a regression algorithm to generate a continuous map. For a dense manual survey there is often little sensitivity to the sophistication of the regression algorithm since it is only being tasked with predicting the values a short distance from one or more inputs.

However, for a path survey the interpolation or regression stage is often crucial. Because the path only samples an

irregular subset of the space, there is a greater need to predict values further from the survey points. It is difficult to predict the variation of WiFi-like signals with distance in an indoor environment and so non-parametric regression techniques are favoured. Gaussian Process (GP) regression has emerged as the de-facto regression algorithm for WiFi data [9], [8]. Full details are available elsewhere ([20]): here we note only that the technique provides a normal *distribution* for the signal value at any given point in space, rather than a single value.

Path surveying in an arbitrary environment is a natural extension of work in Pedestrian Dead Reckoning (PDR) [11]. PDR systems are based on the use of inertial sensors to estimate relative positions and headings. By combining building floor plans and particle filters, PDR can track a pedestrian to the metre level [23], [25], [24]. As such they are ideally suited to a dedicated surveyor performing a path survey.

Rai et al. proposed a similar approach for crowdsourcing [3]. In practice, however, consumers are resistant to permanently turn on the necessary sensors and processing due to high battery consumption. Turning it on for short periods might be acceptable, but it is typically difficult to set a start position on the floor plan when initialising the system from cold. In addition, PDR algorithms are not yet robust enough to deal with the arbitrary positions and movements of a smartphone seen when crowdsourcing.

Other systems have applied Simultaneous Localisation and Mapping (SLAM) techniques to constrain PDR drift. These systems use the spatially-varying signals to correct PDR error. Once a trustworthy trajectory is established, the observed signals are re-used to form a path survey. Most SLAM path survey approaches are based on graph optimisation (e.g. [13]), although particle filters are also used [7], as well as hybrids [6].

III. PATH SURVEYING TECHNIQUES

A path survey collects signal survey points along a continuous path through space rather than at discrete, evenly-distributed points. Like the signal it is measuring, the path may be dedicated to the survey task (where a surveyor moves with the sole purpose of building the maps); or opportunistic (where crowd-sourced measurements are collected from users moving with some other purpose in mind).

The primary advantage of a dedicated surveyor is that the movement can be constrained: the surveyor can be instructed to provide good spatial coverage; to provide occasional manual position fixes; and to move in a way that produces good position estimates. For example, some proposed schemes use some form of Pedestrian Dead Reckoning (PDR) algorithm applied to inertial sensors in the survey device. This can give high quality trajectory estimates if the user keeps the device in one position and only ever walks in the direction it points in (i.e. no side steps or back steps) at a constant speed. Applications such as WiFiSLAM, Google Indoor Maps[1] and IndoorAtlas[2] took this one step further and require the surveyor to move in

manually-specified straight line segments. The segments are then interpolated at the times signal measurements were made to produce a line of survey points. More advanced algorithms from the SLAM family may also be applied by having the surveyor perform particular actions (usually walking in loops).

Crowd-sourced data bears some similarity to the segment-walking technique just described: multiple users are expected to contribute short segments that can be combined over time to form a comprehensive survey. Since the users behave naturally, a higher density of measurements will be associated with more popular regions. Thus crowdsourced maps are naturally more detailed in high footfall areas such as corridors, at least in principle. In practice, crowdsourcing is challenging due to the lack of constraint in the trajectories and movements: patchy spatial coverage; devices being held in different ways; segments being non-linear with complex movements that are difficult to characterise. Even if a trajectory can be established, it is often difficult to anchor it to the floorplan, and thus to provide enough reliable survey points to make a comprehensive map. To date there has been no detailed study of the validity of crowd-sourced signal maps, although it shows promise.

The key requirement of the path survey is to recover the survey path taken by the surveyor. We consider two approaches:

- **A General Path Survey.** Here the surveyor moves freely around the environment and sensors are used to track them. Typical tracking algorithms will use PDR algorithms running on a consumer device with a building floor plan [23], [25].
- **Segment Survey.** This uses the approach popularised by Google Indoor Maps and IndoorAtlas. The surveyor manually marks a small number of points on a floor plan image and then walks in straight lines between them. Typically only a few segments along corridors are used.

The accuracy of the path derived by either method is affected by multiple factors. The most obvious one is the noisy sensors in the survey devices (typically smartphones). The algorithms (SLAM, particle filter etc.) used to correct PDR error also have impacts on the survey accuracy. The user's behaviour (e.g., whether the user walks along the specified segment accurately) can affect the segment survey accuracy. In this work we wish to bound the expected accuracy and so we use an external positioning system to simulate best-case tracking for path and segment surveys

IV. A GROUND-TRUTHED SURVEY DATASET

We have collected an extensive dataset in order to bound and assess the performance of path surveys in a real-world environment using consumer devices. A detailed manual survey of WiFi and BLE signals was taken in addition to a series of ground-truthed path survey walks.

A. Environment and Devices

Our tests were performed in a wing of the William Gates Building housing the Computer Laboratory at the University of Cambridge, UK. The building itself is a three-storey office

building constructed from brick, steel and concrete. Internal walls are constructed from plasterboard, wood and metal.

We gathered signal data using a custom application running on an Android smartphone. The application recorded WiFi and BLE Received Signal Strength (RSS) data. [3]

For accurate ground truth and to simulate an 'ideal' path survey, we used an established high-accuracy ultrasonic positioning system (the Bat system [1]). This system is capable of 3 cm accuracy in 3D 95% of the time. The active tag that the Bat system tracks was attached to the back of the smartphone. Time synchronisation was achieved using NTP (Network Time Protocol).

The test environment contained three WiFi APs; additional APs from the floor below and adjacent areas were also observed. The APs had been previously deployed by the building's IT staff to provide good communications coverage throughout the building (i.e. not explicitly with positioning in mind). As such they were representative of most office buildings. Seven BLE beacons were available in the area, each set to advertise a unique identifier at 12 Hz. These were deployed solely for the purpose of positioning. The beacons were attached to the ceiling (corridors) or on ledges above windows (offices). Figure 1 illustrates the test area and the positions of the WiFi APs and BLE beacons within it.
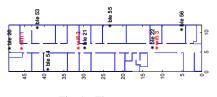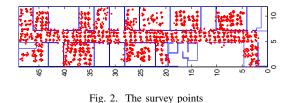
Fig. 1. The test area.

B. Manual Survey

We performed a manual survey in the test environment. Unlike previous works, we did not attempt to establish a set of positions on a regular grid since this was both very laborious and error-prone.[4] Instead, we used an approximate grid but measured the position of the device accurately at each survey point using the Bat system. In this way we achieved comprehensive spatial coverage and accurate survey positions. We aimed for an inter-point spacing of 1 m to provide a detailed survey—Figure 2 illustrates the survey points on a floor plan (rooms without survey points were not accessible during the survey).

During surveying the device was handheld as if being used to navigate. To examine body shadowing effects the surveyor rotated about a vertical axis through the phone to point in each of the four cardinal directions of the building.[5] 2.5 s of data

Fig. 2. The survey points

were collected in each orientation. This period allowed for tens of BLE beacon measurements, and up to two measurements of each WiFi AP per survey point per orientation.

C. Path Survey Walks

We performed four different survey walks (labelled W1, W2, W3, and W4) with the same device held by the surveyor as if navigating. The Bat system provided ground truth location. W1 was carefully chosen to represent a dedicated surveyor visiting every room and area covered by the manual survey; W2 and W3 involved walking the corridor multiple times (to simulate segment survey); W4 was free-form and visited many, but not all, of the surveyed rooms (see Figure 3). In total the walks contained 2,191 steps.
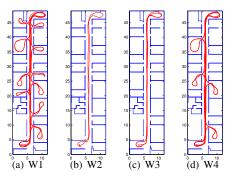
(a) W1   (b) W2   (c) W3   (d) W4

Fig. 3. Survey walk paths

D. GP Regressed Map

For every survey, we generated a GP map for each signal source[6]. Figure 4 shows sample WiFi and BLE RSS measurements taken during different surveys and the GP regressed maps (mean and variance at each point) that result from them. The variance grows quickly the further from survey points/paths we move, correctly encapsulating the growing signal propagation uncertainty.

V. POINT COMPARISON ON REGRESSED MAPS

When surveying for the purposes of wireless network deployment, the GP map itself is the end goal. Furthermore, we might intuitively expect that better map approximations allow for better online positioning. To meaningfully compare GP maps directly is a challenge since each position is associated with a normal *distribution* rather than a scalar value. Hence we first establish a metric to compare normal distributions.

---

[3]The Android WiFi subsystem was set to use 2.4 GHz only since this reduced the scan period from around 4 s to 1.5 s: a significant reduction that allowed us to better localise a scan spatially when the device was moving.

[4]It is difficult to accurately mark a 1 m grid across a large indoor space and even more difficult to ensure that a device is being held over a marked position when surveying

[5]These were approximately aligned with the magnetic cardinal directions. For simplicity we assume them to be the same herein.

[6]We adopt the open source library provided by [20] to train the survey data and generate GP maps.
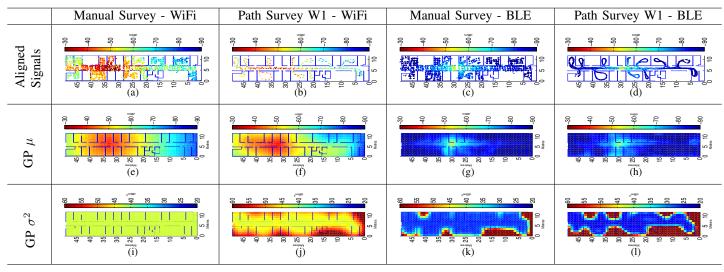
Fig. 4. Sample signal surveys and the resulting GP maps

Comparison of normal distributions. For two normal distributions, $\mathcal{N}(\mu_1, \sigma_1^2)$ and $\mathcal{N}(\mu_2, \sigma_2^2)$, we form the distribution of the differences, which has the form $\mathcal{N}(\mu_1 - \mu_2, \sigma_1^2 + \sigma_2^2)$. Since we are only interested in the magnitude of the difference, we convert this normal distribution to a folded normal distribution, the CDF of which is given by:

$$F(X; \mu, \sigma) = \frac{1}{2} \left(\mathrm{erf} \left(\frac{x + \mu}{\sqrt{2}\sigma} \right) + \mathrm{erf} \left(\frac{x - \mu}{\sqrt{2}\sigma} \right) \right) \quad (1)$$

for $\mu = \mu_1 - \mu_2$ and $\sigma^2 = \sigma_1^2 + \sigma_2^2$. We then estimate the RSS difference corresponding to the 90% confidence interval, RSS_{90}. Hence 90% of the time the two distributions would agree to within RSS_{90}.

Comparison of GP maps. We evaluate RSS_{90} at each grid point position covered by the manual survey since these points are associated with the lowest uncertainty in the manual survey-derived GP map. The distribution of RSS_{90} acts as a measure of the agreement between the two maps. Two very similar maps would show a strong concentration of RSS_{90} values within the expected measurement noise, whilst differing maps would not.

We visualise the RSS_{90} by heatmap (Figure 5) and CDF (Figure 6). The colour (value) c of a point on the RSS_{90} map represents that at this point, the predicted RSS values by a path survey GP map and a manual survey GP map agree to within c dBm 90% of the time. Similarly, the interpretation of a point (x, y) on the CDF is the proportion of grid points (y) where we expect the two maps to agree to within x dBm 90% of the time.

Figure 6(b) shows that for the BLE signal, the GP maps of W1 and W4 (which covered the space more comprehensively than W2 and W3) are very similar to the manual survey. This is because these survey paths cover the space comprehensively. Conversely W2 and W3 (which only visit a subset of the rooms in the manual survey and hence require prediction further from

the survey path) exhibit good results only within a metre or so of the survey path (Figure 5(f) and 5(g)).

However, this phenomenon is less obvious for the WiFi signal. One possible reason is that the WiFi signal variance over space is lower than that of the BLE signal. As can be seen in Figure 4, the WiFi signal attenuates more slowly than BLE over space. This is confirmed by the GP regression results—as is common we use a *squared exponential* (SE) covariance function (or "Gaussian kernel") to encode the correlation between nearby measurements:

$$k_{SE}(r) = \sigma_f^2 exp \left(-\frac{r^2}{2l^2} \right) \quad (2)$$

where r is the distance between two spatial locations, σ_f^2 is the signal variance, and l is the *characteristic length-scale* that determines how the correlation weakens with distance [20], [9]. The signal variance σ_f^2 and the characteristic length-scale l are usually learnt from the training data. For the W1 WiFi and BLE maps in Figure 4, the learnt values of the signal variance σ_f^2 are 44.08 and 193.72 respectively, and the learnt values for the characteristic length-scale l are 4.96 m and 2.50 m respectively. These values show that WiFi signals in distant locations correlate better than their BLE counterparts. This means it is easier for GP regression to predict the signal strength distribution in more distant locations for WiFi than BLE.

VI. POSITIONING COMPARISON ON GP REGRESSED MAPS

Positioning using GP maps requires the use of multiple maps (or equivalently, signals from multiple sources). It is therefore feasible that good positioning results could be achieved even when one or more GP maps are very poor approximations to the true map. To see this, consider a GP map with high variances associated with most of the evaluation points (which is an indication that the GP map is not of high quality). A sensible positioning algorithm would derive very little information from such a map—in effect, it should be
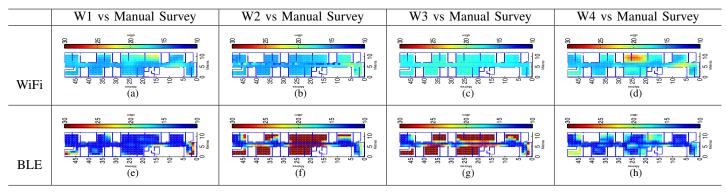
W1 vs Manual Survey	W2 vs Manual Survey	W3 vs Manual Survey	W4 vs Manual Survey
WiFi (a)	(b)	(c)	(d)
BLE (e)	(f)	(g)	(h)

Fig. 5. RSS_{90} maps from point comparison of various GP maps with the manual survey GP map.

(a) WiFi

(b) BLE

Fig. 6. RSS_{90} distribution (CDF).

Fig. 7. Corridor directionality for BLE. It is similar in the case of WiFi so we omit it here for brevity.
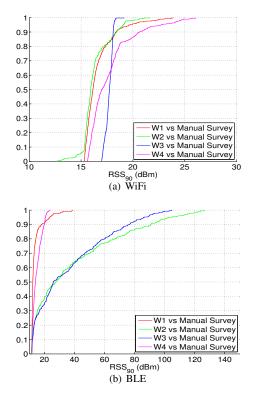
discarded. Provided there are enough low-variance evaluations for other GP maps, a good positioning results is still possible. In this section we study the results of positioning using maps from path and manual surveys.

A. Positioning Method

We adopt the signal strength-based location estimation method described in [9] to evaluate the positioning performance. We first use the raw data to generate GP maps and then use a standard Bayesian localisation algorithm to incorporate the variance estimate as well as the mean at each GP map location. To estimate a person's location, x, conditioned on observation of a set of signal strength measurements (fingerprints), z, we used:
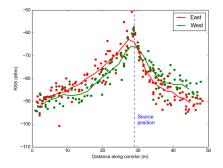
$$p(x|z) \propto p(x)p(z|x) \tag{3}$$

where $p(x)$ is the prior probability of the person being at x; and $p(z|x)$ is the likelihood of observing a set of signal strength measurements at location x. $p(z|x)$ is directly given by the GP regression.

We divided the environment into a grid of square cells of length 1 m or less, For each step at each cell we computed $p(x|z)$ and selected the cell with the maximum as the person's location. Signal strength measurements lower than -90 dBm were discarded. Note that the positioning algorithm did *not* assume prior information—i.e. our results are for one-shot positioning. This is achieved by setting $p(x)$ to a uniform distribution over the whole space. A recursive tracking scheme that would seed each positioning calculation with the posterior from the last would be expected to perform better, but could mask a bad positioning result.

B. Map Directionality And Spatial Coverage of Survey Path

Previous works have highlighted an orientation dependency in RSS measurements on handheld consumer devices. The dependency is primarily caused by body shadowing: the attenuation of a signal that must pass through the human body as it travels directly from source to receiver. The extent of this dependency is important for a path survey, since it may demand paths be traversed in both directions. In addition, if the

	Manual Survey	Path Survey W1
East-facing WiFi	1460	77
West-facing WiFi	1452	70
East-facing BLE	1374	703
West-facing BLE	1407	626

orientation plays an important role in the positioning accuracy, than more complicated orientation-based positioning algorithm is preferred.

We expect directionality to be strongest along vectors that pass through the source, and it is thus most easily observed along the straight corridor in our data. Figure 7 shows the East-facing and West-facing RSS values observed at the survey points in the corridor, plotted against the distance along the corridor (the distance increases in the westerly direction). The general trends are shown by the two smoothed lines. In both cases there is an offset between the trends that corresponds to higher RSS when approaching a source compared to receding from it.

Our survey data contained orientation information, allowing us to generate both:

- **Directional Maps**. For each signal source, divide corresponding signal strength sample dataset into several subsets according to the surveyor's direction when taking each sample. Then, create a map for each cardinal direction. When positioning, use only the map corresponding to the direction of movement. Here we only use the east-facing and west-facing datasets because these correspond to the major direction of the building. The number of samples in each dataset for a single signal source is shown in Table I.
- **Omnidirectional Map**. For each signal source, create a single map using all the survey data *after* discarding the orientation labels. The directionality should then manifest as a larger variance in the values collected at each survey point.

We produce these two kinds of map using various datasets and evaluate their positioning performance using different datasets as fingerprints (input). We adopt the notation AA_{BB} to indicate that we evaluated the maps generated from AA using the data in BB as input. For example, MS_{W1} evaluates the manual survey maps using the W1 data. Figure 8 shows the positioning results separated into those in and out of the corridor (where we expect the most significant directionality).

We see that the directional maps outperform the omnidirectional maps for the manual survey (MS_{W1}) as expected given the observations made earlier. However, the omnidirectional map generated from W1 achieves better results than the directional map when evaluated using the manual survey data ($W1_{MS}$). We attribute this result to the lower volume of survey points used in the regression for a path survey (Table I). Separating the low number of points into directional sets further sparsifies the regression input, making the directional

maps of particularly poor quality. The overall $W1_{MS}$ positioning accuracy is notably worse than the MS_{W1} accuracy. However, we see in Figures 8(c) and 8(f) that the higher-error positions are associated with distances further from the W1 path. Within approximately 2 m of the path, the path survey maps performed only slightly worse than the manual survey maps.

In practice, a given space is not utilised uniformly: people walk in the centre of corridors, around furniture, and so on. In recognition of this, we also show the results of evaluating W1 maps using W4 input data ($W1_{W4}$). Since W4 is a subset of W1, this simulates the likely real-world situation where the survey path follows the most likely paths in the space. We observe that the positioning performance $W1_{W4}$ is similar or even better than manual survey map positioning. We attribute this to the fact we are matching data collected in a similar manner (i.e. while moving).

C. Segment Survey Evaluation

A segment survey will typically aim to generate the maps based on a few uni-directional walks along corridors. From our results so far, we expect the positioning performance for such maps to be good only near the corridor. We verify this using walks W2 (walking up and down the corridor once) and W3 (walking up and down the corridor twice), generating omnidirectional maps from each. Results are shown in Figure 9, where we see that W1 maps significantly outperform W2 and W3 maps for both WiFi and BLE *outside* of the corridor as expected.

Within the corridor, the performance of the W2 and W3 maps is much closer to that of the W1 maps, although it is notable that there is a greater discrepancy associated with the WiFi maps. This illustrates the importance of sample density. W1 has more samples in the corridor area than W3, and W3 has more than W2, so WiFi positioning performance of W1 is better than W2 and W3. But the sample rate of BLE is much higher than that of WiFi (10 Hz vs 0.5 Hz) on the android phone we use, which mitigates the difference in the sample density of the three surveys. So the difference in the BLE positioning performance is less significant than the difference in the WiFi positioning performance.

We conclude that segment surveys that only cover the corridor area of the environment will not give good positioning performance beyond the corridors. Comprehensive spatial coverage is necessary when performing a path survey.

VII. OBTAINING ACCURATE REAL-WORLD TRAJECTORIES

Our analysis used an external positioning system to reconstruct the path survey trajectory very accurately. In practice this will be unavailable and the trajectory must be derived from other techniques such as PDR. Such techniques are likely to be associated with higher trajectory error, although the gap is fast closing due to the application of SLAM-based techniques. These require the trajectory is repeated in whole or part to provide *loop closures*. We have previously demonstrated that
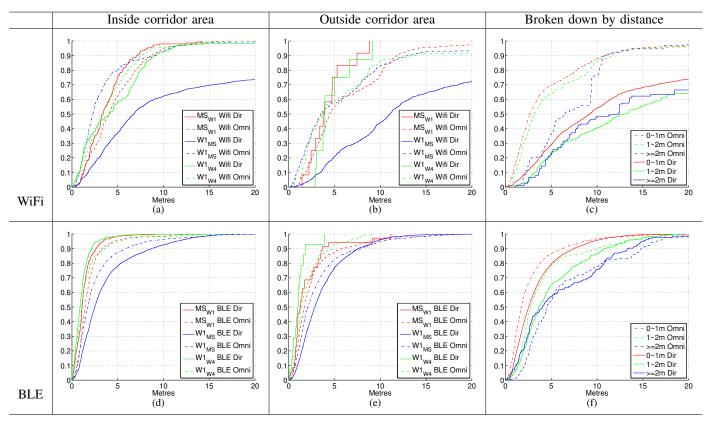
Fig. 8. Positioning results for directional maps. AA_{BB} implies a map generated from AA data tested using the data from BB. Directional maps are labelled 'Dir', non-directional maps 'Omni'. The rightmost column shows the results of $W1_{MS}$ broken down by distance of (manual survey) test point from the W1 survey path.

a handheld smartphone can produce accurate trajectories (sub-metre for any given point) using such techniques [10].

VIII. CONCLUSIONS AND FURTHER WORK

In this paper we have quantitatively evaluated the signal maps produced by Gaussian Processes regression applied to a grid-like manual survey and a more opportunistic path survey. We used a high-accuracy tracking system during the path survey to bound the performance by simulating best-case inputs. Our results motivate the following guidelines for path surveying:

1) The survey path should visit each room and pass within 1∼2 m of any given point where positioning may be subsequently required;

2) The surveyor should repeat some parts of the path to increase the signal sampling density. This is especially necessary for WiFi positioning because the WiFi sample rate on a modern smartphones is typically as low as 0.5∼1.0 Hz; and

3) For best results, directionality should be incorporated into the generated maps. This clearly requires traversal of the path in both directions. However, traversing once in either direction does not give a sufficient sample density if the signal sample rate is low (< 2 Hz). In such cases the path must be traversed multiple times in

multiple directions *or* omnidirectional maps should be used.

If these guidelines are followed and an accurate trajectory is computed for a device, we found that a path survey can produce similar results to a manual survey but at a much lower cost. In our experiments the manual survey took more than 3 hours (despite using an accurate external positioning system to speed up the process) but each of the W1–W4 took less than 10 minutes (Table II).

TABLE II
SURVEY DURATIONS (MINS)

W1	W2	W3	W4	Manual
9.6	2.1	3.3	7.5	206

REFERENCES

[1] M. Addlesee, R. Curwen, S. Hodges, J. Newman, P. Steggles, A. Ward, and A. Hopper. Implementing a sentient computing system. *IEEE Computer*, 34(8), August 2001.
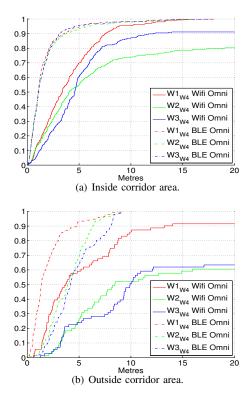
(a) Inside corridor area.

(b) Outside corridor area.

Fig. 9. The positioning CDF results for W1, W2 and W3 maps evaluated against fingerprints generated from W4 data.

[2] K. Al Nuaimi and H. Kamel. A survey of indoor positioning systems and algorithms. In *Innovations in Information Technology (IIT), 2011 International Conference on*, pages 185–190. IEEE, 2011.

[3] A. Anshul Rai, K. K. Chintalapudi, P. Venkat, and R. Sen. Zee : Zero-effort crowdsourcing for indoor localization. In *Proceedings of The 18th Annual International Conference on Mobile Computing and Networking (MobiCom)*, august 2012.

[4] P. Bahl, V. N. Padmanabhan, and A. Balachandran. Enhancements to the radar user location and tracking system. Technical report, 2000.

[5] J. Chung, M. Donahoe, C. Schmandt, I.-J. Kim, P. Razavai, and M. Wiseman. Indoor location sensing using geo-magnetism. In *Proceedings of the 9th International Conference on Mobile Systems, Applications, and Services*, MobiSys '11, pages 141–154, New York, NY, USA, 2011. ACM.

[6] R. Faragher and R. Harle. Smartslam - an efficient smartphone indoor positioning system exploiting machine learning and opportunistic sensing. In *Proceedings of the 26th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2013, Nashville, Tennessee*, September 2013.

[7] R. Faragher, C. Sarno, and M. Newman. Opportunistic radio slam for indoor navigation using smartphone sensors. In *Position Location and Navigation Symposium (PLANS), 2012 IEEE/ION*, pages 120 –128, April 2012.

[8] B. Ferris, D. Fox, and N. Lawrence. Wifi-slam using gaussian process latent variable models. In *Proceedings of the 20th international joint conference on Artifical intelligence*, IJCAI'07, pages 2480–2485, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.

[9] B. Ferris, D. Haehnel, and D. Fox. Gaussian processes for signal strength-based location estimation. In *In proc. of robotics science and systems*. Citeseer, 2006.

[10] C. Gao and R. Harle. Sequence-based magnetic loop closures for automated signal surveying. In *Indoor Positioning and Indoor Navigation (IPIN), 2015 International Conference on*, pages 1–12. IEEE, 2015.

[11] R. Harle. A survey of indoor inertial positioning systems for pedestrians. *Communications Surveys Tutorials, IEEE*, 15(3):1281–1293, Third 2013.

[12] V. Honkavirta, T. Perala, S. Ali-Loytty, and R. Piche. A comparative survey of wlan location fingerprinting methods. In *Positioning, Navigation and Communication, 2009. WPNC 2009. 6th Workshop on*, pages 243–251, March 2009.

[13] J. Huang, D. Millman, M. Quigley, D. Stavens, S. Thrun, and A. Aggarwal. Efficient, generalized indoor wifi graphslam. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 1038–1043, May.

[14] E. Jedari, Z. Wu, R. Rashidzadeh, and M. Saif. Wi-fi based indoor location positioning employing random forest classifier. In *Indoor Positioning and Indoor Navigation (IPIN), 2015 International Conference on*, pages 1–5. IEEE, 2015.

[15] T. King, S. Kopf, T. Haenselmann, C. Lubberger, and W. Effelsberg. Compass: A probabilistic indoor positioning system based on 802.11 and digital compasses. In *Proceedings of the 1st international workshop on Wireless network testbeds, experimental evaluation & characterization*, WiNTECH '06, pages 34–40, New York, NY, USA, 2006. ACM.

[16] H. Koyuncu and S. H. Yang. A survey of indoor positioning and object locating systems. *IJCSNS International Journal of Computer Science and Network Security*, 10(5):121–128, 2010.

[17] M. Kranz, C. Fischer, and A. Schmidt. A comparative study of dect and wlan signals for indoor localization. In *Pervasive Computing and Communications (PerCom), 2010 IEEE International Conference on*, pages 235–243. IEEE, 2010.

[18] H. Liu, H. Darabi, P. Banerjee, and J. Liu. Survey of wireless indoor positioning techniques and systems. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 37(6):1067–1080, 2007.

[19] V. Otsason, A. Varshavsky, A. LaMarca, and E. De Lara. Accurate gsm indoor localization. In *UbiComp 2005: Ubiquitous Computing*, pages 141–158. Springer, 2005.

[20] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.

[21] M. Schüssel and F. Pregizer. Coverage gaps in fingerprinting based indoor positioning: The use of hybrid gaussian processes. In *Indoor Positioning and Indoor Navigation (IPIN), 2015 International Conference on*, pages 1–9. IEEE, 2015.

[22] P. Wilk, J. Karciarz, and J. Swiatek. Indoor radio map maintenance by automatic annotation of crowdsourced wi-fi fingerprints. In *Indoor Positioning and Indoor Navigation (IPIN), 2015 International Conference on*, pages 1–8. IEEE, 2015.

[23] O. Woodman and R. Harle. Pedestrian localisation for indoor environments. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 114–123. ACM, 2008.

[24] O. Woodman and R. Harle. RF-Based Initialisation for Inertial Pedestrian Tracking. *Pervasive Computing 7th International Conference Pervasive 2009 Nara Japan May 1114 2009 Proceedings*, 5538:238–255, 2009.

[25] O. J. Woodman. *Pedestrian localisation for indoor environments*. Phd thesis, Computer Laboratory, University of Cambridge, September 2010.

[26] M. Youssef and A. Agrawala. The horus wlan location determination system. In *Proceedings of the 3rd international conference on Mobile systems, applications, and services*, MobiSys '05, pages 205–218, New York, NY, USA, 2005. ACM.