2011 INTERNATIONAL CONFERENCE ON INDOOR POSITIONING AND INDOOR NAVIGATION (IPIN), 21-23 SEPTEMBER 2011, GUIMARÃES, PORTUGAL

Bayesian Radio Map Learning for Robust Indoor Positioning

Hui Wang*

*Sensor-Actuator-Systems Laboratory (ISAS),Karlsruhe Institute of Technology, Germany. Email: evanwang1981@gmail.com

***Abstract*—The uncertainty of radio propagation results in large errors in positioning systems based on the received signal strength (RSS). Especially in an indoor environment, the RSS distribution map, so called radio map, has a very complicated form due to numerous site-specific parameters. Therefore, modelling the radio map is a critical task for RSS based positioning systems. Researchers usually obtain an accurate radio map by measuring the RSS at a number of reference points. But in this way too many calibration efforts should be spent to guarantee a fine radio map accuracy. In this paper, a calibration-free radio map learning framework is proposed. In this framework, the system starts with a very simple and coarse radio map model, such as a radial model with default parameter values. A more accurate model is then obtained by learning the unlabelled online RSS data. The Expectation-Maximisation (EM) algorithm is used to calculate the posterior maximum likelihood (ML) of radio maps. Besides, we extend the standard EM algorithm by integrating expert knowledge of radio propagation. By applying the proposed algorithms in real-world data sets, we demonstrate that an accurate and robust radio map can be learned without requiring any calibration data.**

I. INTRODUCTION

Indoor positioning systems serve as the basis of a broad category of location-based applications such as tracking of assets and people, logistics, location-aware multimedia services and many others. So far, satellite navigation systems, such as the well-known Global Positioning System (GPS), still perform badly indoors because of attenuation and multipath propagation caused by buildings and walls. As an alternative, many researchers developed indoor positioning systems based on middle or short range wireless infrastructures, such as Wireless LAN (WLAN) [1], Zigbee [2] and Ultra-Wideband(UWB) [3].

Many of indoor positioning systems make use of received signal strength (RSS) to infer the location because all off-the-shelf devices support the RSS reporting. A standard RSS-based positioning approach consists of two steps: *offline* training and *online* localization/tracking. As shown in Fig. 1, in *offline* training step, a so-called radio map function, which represents the relationship between locations and RSS measurements, is learned from labelled data. Given the radio map function, the locations are estimated from online unlabelled RSS measurements in *online* localization/tracking step.

One of the biggest challenges of RSS-based positioning is how to obtain an accurate radio map function. As well-known, the indoor radio propagation is very complicated due to the reflection, refraction, diffraction and multi-path effect. Its complexity is highly related to the site parameters of a specific indoor environment such as the number and positions of walls, doors or furnitures. Thus, building a radio map function is a dilemma. A simple radio propagation model needs very few calibration samples to learn the parameters but often leads to a coarse accuracy, especially in complicated environments. On the other hand, if a complex model is used, many calibration samples are required, which results in a large installation effort, e.g. from hours to days. Besides, the radio map varies with time because either the parameters of transmitter or the parameters of environment change temporally. That means the calibration process has to be repeated periodically, bringing a large maintenance cost.

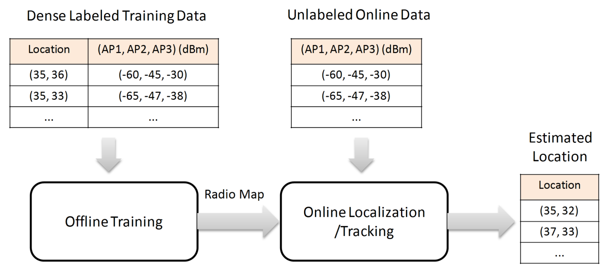

Fig. 1. Standard Procedure of RSS-based Positioning Approaches

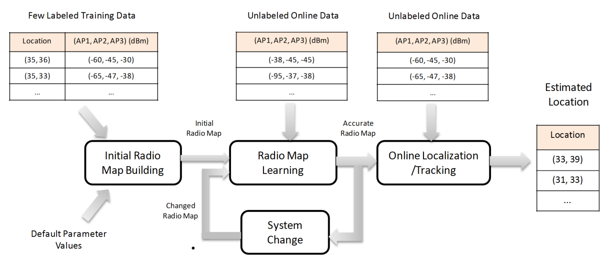

Fig. 2. Proposed Positioning Approaches Based on Radio Map Learning

In this paper, a radio map learning framework is proposed. As illustrated in Fig. 2, the key idea is to first build the radio map function using a very simple model. This model has very few parameters which can be learned from several labelled calibration samples or use default values. Then we generate many *pseudo* labelled data from the simple model and use these data to learn a complicated model. Of course, the complicated model is now not accurate. In the next learning step, a lot of online unlabelled data is used to improve the accuracy. Finally, we obtain an more accurate radio map function. If the environment changes, e.g., the furnitures are moved or the new access points (APs) are added, we just need to re-run the learning step.

The advantage of this radio map learning framework is to reduce the calibration effort without scarifying the accuracy. In the initialization step, we can use the mutual measurements of APs to train the model or directly take the default parameters. In the learning step, the online unlabelled data are automatically acquired after the system runs. No calibrations are necessary any more.

II. Problem Formulation

In the paper we use $\boldsymbol{x}$ to denote a location. $\boldsymbol{x}$ is either a 2-dimensional or a 3-dimensional vector, i.e., $\boldsymbol{x} = [x, y]^{\mathrm{T}}$ or $\boldsymbol{x} = [x, y, z]^{\mathrm{T}}$. $\mathcal{L} \subset \mathbb{R}^2$ or $\mathcal{L} \subset \mathbb{R}^3$ denotes an indoor location space, i.e., $\boldsymbol{x} \in \mathcal{L}$. A radio map function is defined by $R : \mathcal{L} \to \mathcal{P}$, where $\mathcal{P} \subset \mathbb{R}$ denotes a RSS space. Given an AP, a RSS measurement $s \in \mathcal{P}$ is modelled by

$$s = R(\boldsymbol{x}) + v, \tag{1}$$

where $v \backsim \mathcal{N}(0, \sigma_v^2)$ denotes the zero-mean Gaussian measurement noise. If I APs are available, we use s_i, R_i and v_i to denote the RSS measurement, radio map function and measurements for AP i, respectively.

Given radio map functions, the statistical theory can be used to solve the localization problem. Assuming that the RSS measurement from I APs are given by the vector $\boldsymbol{s} = [s_1, \cdots, s_I]^{\mathrm{T}}$, where $\{s_1, \cdots, s_I\}$ are the respective RSS values from AP_1 to AP_N. Then the posterior probability of being at the location $\boldsymbol{x}$ given the measurement $\boldsymbol{s}$ is expressed as $p(\boldsymbol{x}|\boldsymbol{s})$, which can be written using the Bayesian rule as:

$$p(\boldsymbol{x}|\boldsymbol{s}) = \frac{p(\boldsymbol{s}|\boldsymbol{x})p(\boldsymbol{x})}{\int p(\boldsymbol{s}|\boldsymbol{x})p(\boldsymbol{x})\, d\boldsymbol{x}} \tag{2}$$

Since all positions are equally probable, the prior probability $p(\boldsymbol{x})$ can be assumed as uniformly distributed. Additionally $\{s_1, \cdots, s_I\}$ can be regarded as independent for a given $\boldsymbol{x}$. Then (2) can be written as:

$$p(\boldsymbol{x}|\boldsymbol{s}) = \frac{\prod_{i=1}^{I} p(s_i|\boldsymbol{x})}{\int p(\boldsymbol{s}|\boldsymbol{x})\, d\boldsymbol{x}} \tag{3}$$

From (1), we know that the measurement noise is Gaussian, (3) becomes

$$p(\boldsymbol{x}|\boldsymbol{s}) = \frac{\prod_{i=1}^{I} \frac{1}{\sigma_{v,i}\sqrt{2\pi}} \exp\left(-\frac{(s_i - R_i(\boldsymbol{x}))^2}{2\sigma_{v,i}^2}\right)}{\int p(\boldsymbol{s}|\boldsymbol{x})\, d\boldsymbol{x}}, \tag{4}$$

where $\sigma_{v,i}$ is the standard deviation of measurement noise for AP i.

With Minimum Mean Variance Bayesian Estimator or Minimum Mean Square Error (MMSE) Estimator, an unbiased estimation is given by

$$\hat{\boldsymbol{x}} = \mathrm{E}[\boldsymbol{x}|\boldsymbol{s}] = \int \boldsymbol{x} \cdot p(\boldsymbol{x}|\boldsymbol{s})\, d\boldsymbol{x}, \tag{5}$$

.

A radio map function can be approximated by various models. Of course, no model is perfect. The true radio map is the summation of an approximated radio map and the corresponding model error, which is expressed as

$$R(\boldsymbol{x}) = \tilde{R}(\boldsymbol{x}, \boldsymbol{\theta}) + \Delta\tilde{R}(\boldsymbol{x}), \tag{6}$$

where $\tilde{R}$ denotes the approximated radio map like radial model R^{RM} or multi-wall model R^{MWM}. $\boldsymbol{\theta}$ represents the parameter of $\tilde{R}$. $\Delta\tilde{R}(\boldsymbol{x})$ is the model error.

In order to generate an accurate radio map with minimum calibration effort, we propose a novel radio map learning framework. This framework is illustrated by Fig. 2. The key idea is to start with a simple model and then build a complicated model by using unlabelled online RSS data. Here we use $R_{\mathrm{s}}(\boldsymbol{x}, \boldsymbol{\theta}_{\mathrm{s}})$ to denote a simple initial model such as R^{RM} or R^{MWM}, whose parameter vector $\boldsymbol{\theta}_{\mathrm{s}}$ has a low dimension $n_{\boldsymbol{\theta}_{\mathrm{s}}}$. $R_{\mathrm{c}}(\boldsymbol{x}, \boldsymbol{\theta}_{\mathrm{c}})$ denotes a complicated model such as mass fingerprinting model R^{FP}, which has a high dimensional parameter $\boldsymbol{\theta}_{\mathrm{c}}$. The dimension of $\boldsymbol{\theta}_{\mathrm{c}}$ is $n_{\boldsymbol{\theta}_{\mathrm{c}}}$, satisfying $n_{\boldsymbol{\theta}_{\mathrm{c}}} \gg n_{\boldsymbol{\theta}_{\mathrm{s}}}$.

Initially, we need to know the value of parameter $\boldsymbol{\theta}_{\mathrm{s}}$ to build R_{s}. We can use default parameters. Or in the scenarios where the positions of APs are known, we can take mutual measurements among APs as calibration data to train the parameter. No matter which approach is applied, no manual calibration is needed. Besides, by starting with a simple model, we naturally integrate the knowledge regarding the physics of radio propagation into the learning framework.

After obtaining R_{s}, we can generate as many labelled RSS data as we want. These *pseudo* labelled data are used to initialize parameter $\boldsymbol{\theta}_{\mathrm{c}}^0$ of a complicated model $R_{\mathrm{c}}(\boldsymbol{\theta}_{\mathrm{c}}^0)$. Of course, $\boldsymbol{\theta}_{\mathrm{c}}^0$ is not accurate and needs to be corrected. In the next step, a more accurate radio map $R_{\mathrm{c}}(\boldsymbol{\theta}_{\mathrm{c}})$ will be learned from unlabelled online data. If the system or the environment changes, we just re-learn the new radio map from the old one.

The core of proposed learning framework is to learn the parameter $\boldsymbol{\theta}_{\mathrm{c}}$ from RSS measurements, given an initial parameter $\boldsymbol{\theta}_{\mathrm{c}}^0$. This is equivalent to a parameter estimation or system identification problem.

III. Expectation-Maximisation (EM) Algorithm for Radio Map Estimation

A. EM algorithm

In our learning framework, a radio map is parameterized by $\boldsymbol{\theta}$. The true value of $\boldsymbol{\theta}$ is unknown due to the shortness of satisfactory calibration. So we start with an inaccurate radio map function $R\left(\boldsymbol{x}, \boldsymbol{\theta}^0\right)$, where $\boldsymbol{\theta}^0$ is the initial parameter. When the positioning system runs, the mobile terminal receives RSS measurements continuously. We denote the RSS vector from I APs at time k as $\boldsymbol{s}_k = [s_{1,k}, s_{2,k}, \cdots, s_{I,k}]^T$. The series of RSS observations is denoted as $\boldsymbol{S}_K = [\boldsymbol{s}_1, \boldsymbol{s}_2, \cdots, \boldsymbol{s}_K]$. The $\boldsymbol{\theta}$ can be estimated by maximizing the parameter likelihood function

$$\mathcal{L}(\boldsymbol{\theta}) = \log p(\boldsymbol{S}|\boldsymbol{\theta}). \tag{7}$$

Since the maximisation of $\mathcal{L}(\boldsymbol{\theta})$ is usually hard to be tracked directly, the Expectation-Maximisation (EM) algorithm [4] is widely used to solve this problem. The EM algorithm is a powerful method for finding the maximum likelihood solutions for models with latent variables. It consists of two major steps: an expectation (E) step, which computes an expectation of the likelihood by including the latent variables as if they were observed, and a maximization (M) step, which computes the maximum likelihood estimates of the parameters by maximizing the expected likelihood found on the E step. The parameters found on the M step are then used to begin another E step, and the process is repeated. The estimated parameters can be proven to converge to local maxima after a number of steps [5].

For our radio map learning problem, the EM algorithm includes the following two steps.

E-Step: calculate

$$Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right) = \int_{\mathcal{L}^K} p\left(\boldsymbol{X}_K | \boldsymbol{S}_K, \boldsymbol{\theta}^t\right) \log p(\boldsymbol{X}_K, \boldsymbol{S}_K | \boldsymbol{\theta}) d\boldsymbol{X}_K; \tag{8}$$

M-Step: find

$$\boldsymbol{\theta}^{t+1} = \arg\max_{\boldsymbol{\theta}} Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right), \tag{9}$$

where $\boldsymbol{X}_K = [\boldsymbol{x}_1, \boldsymbol{x}_2, \cdots, \boldsymbol{x}_K]$ denotes the hidden locations. Using Bayes rules, the joint probability $p\left(\boldsymbol{X}_K, \boldsymbol{S}_K | \boldsymbol{\theta}\right)$ can be expressed by

$$p\left(\boldsymbol{X}_K, \boldsymbol{S}_K | \boldsymbol{\theta}\right) = p\left(\boldsymbol{S}_K | \boldsymbol{X}_K, \boldsymbol{\theta}\right) p\left(\boldsymbol{X}_K | \boldsymbol{\theta}\right). \tag{10}$$

Because $\boldsymbol{X}_K$ and $\boldsymbol{\theta}$ are independent, $p\left(\boldsymbol{X}_K | \boldsymbol{\theta}\right) = p\left(\boldsymbol{X}_K\right)$. And under Markov assumption, RSS measurement $\boldsymbol{s}_k$ is only dependent on its location $\boldsymbol{x}_k$. So (10) can be further expressed by

$$p\left(\boldsymbol{X}_K, \boldsymbol{S}_K | \boldsymbol{\theta}\right) = \prod_{k=1}^{K} p\left(\boldsymbol{s}_k | \boldsymbol{x}_k, \boldsymbol{\theta}\right) p\left(\boldsymbol{X}_K\right). \tag{11}$$

After a number of calculations, the function $Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right)$ becomes

$$Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right) = \sum_{k=1}^{K} \int_{\mathcal{L}} p\left(\boldsymbol{x}_k | \boldsymbol{S}_K, \boldsymbol{\theta}^t\right) \log p\left(\boldsymbol{s}_k | \boldsymbol{x}_k, \boldsymbol{\theta}\right) d\boldsymbol{x}_k + E,$$

where E is a constant, which is calculated by

$$E = \int_{\mathcal{L}^K} p\left(\boldsymbol{X}_K | \boldsymbol{S}_K, \boldsymbol{\theta}^t\right) \log p\left(\boldsymbol{X}_K\right) d\boldsymbol{X}_K. \tag{12}$$

If the measurement noise is Gaussian, we can derive

$$p\left(\boldsymbol{s}_k | \boldsymbol{x}_k, \boldsymbol{\theta}\right) = \prod_{i=1}^{I} \frac{1}{\sigma_{v,i}\sqrt{2\pi}} \exp\left(-\frac{(s_{i,k} - R_i(\boldsymbol{x}_k, \boldsymbol{\theta}))^2}{2\sigma_{v,i}^2}\right). \tag{13}$$

In M step, the maximisation of $Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right)$ is tractable by letting

$$\frac{dQ\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right)}{d\boldsymbol{\theta}} \stackrel{!}{=} 0. \tag{14}$$

The solution of above equations depends on the form of radio map function $R(\boldsymbol{x}, \boldsymbol{\theta})$. A widely used type of radio map model has a form

$$R(\boldsymbol{x}, \boldsymbol{\theta}) = \sum_{m=1}^{M} \phi^m(\boldsymbol{x}) r^m, \tag{15}$$

where $\boldsymbol{\theta} = [r^1, r^2, \cdots, r^M]^{\mathrm{T}}$ is parameter vector and $\phi(\boldsymbol{x})$ represents a basis function. The numerical solution can be obtained by discretizing $\boldsymbol{x}_k$. In a very special case where the radio map itself is discrete, i.e., fingerprint model, the basis function is actually a dirac function, i.e.,

$$\phi^i(\boldsymbol{x}_k) = \delta(\boldsymbol{x}_k - \boldsymbol{x}^i). \tag{16}$$

Then we get the analytical solution parameter r_n^m, which is

$$r_n^m = \frac{\sum_{k=1}^{K} p\left(\boldsymbol{x}_k = \boldsymbol{x}^m | \boldsymbol{S}_K, \boldsymbol{\theta}^t\right) s_{n,k}}{\sum_{k=1}^{K} p\left(\boldsymbol{x}_k = \boldsymbol{x}^m | \boldsymbol{S}_K, \boldsymbol{\theta}^t\right)}. \tag{17}$$

(17) indicates that given enough observations, the accurate radio map can be estimated by statistically averaging the observations. The weight $p\left(\boldsymbol{x}_k = \boldsymbol{x}^m | \boldsymbol{S}_K, \boldsymbol{\theta}^t\right)$ is the posterior probability given the current radio map and all RSS measurements, which can be calculated using different algorithms under different assumptions.

- If we assume the current location is only dependent on the current RSS measurement, i.e., no movement is considered, we can get
$$p\left(\boldsymbol{x}_k = \boldsymbol{x}^m | \boldsymbol{S}_K, \boldsymbol{\theta}^t\right) = p\left(\boldsymbol{x}_k = \boldsymbol{x}^m | \boldsymbol{s}_k, \boldsymbol{\theta}^t\right). \tag{18}$$
This posterior density can be calculated by Bayesian localization algorithm in (2) to (4).
- If we assume the current location is dependent on the previous measurements, i.e.,
$$p\left(\boldsymbol{x}_k = \boldsymbol{x}^m | \boldsymbol{S}_K, \boldsymbol{\theta}^t\right) = p\left(\boldsymbol{x}_k = \boldsymbol{x}^m | \boldsymbol{S}_k, \boldsymbol{\theta}^t\right), \tag{19}$$
we can calculate the posterior density by Bayesian filtering.
- Finally, if the current location is assumed to depend on all the RSS measurements, the problem becomes a Bayesian smoothing problem. Various algorithm can solve this problem, like [6]

B. Extensions

1) EM with Constraints: The main problem of EM is that it can only converge to local maxima of likelihood function. Especially for high dimensional parameters, there will be more local stationary points. Therefore, it is very important to choose an initial point for EM estimation. Here we extend the standard EM algorithm by using prior knowledge parameters as the constraints. The benefit of this extension is twofold. On one hand, parameter constrains can limit the search space for EM , hence reducing the possibilities of local maxima. On the other hand, parameter constrains integrate expert knowledge about radio map. The way of adding constraints to EM algorithm is to deploy a new cost function instead of the one in (8),which is given by

$$Q^{'}\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right) = Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right) + \log p\left(\boldsymbol{\theta}\right), \tag{20}$$

where all the information or assumptions about parameter $\boldsymbol{\theta}$ is expressed by the prior density $p\left(\boldsymbol{\theta}\right)$. Different types of constraints can be deployed. Here we discuss some of them as examples.

2) Constraint for Independent Parameters: Sometimes, we know the range for specific parameters. For example, we know that despite that it is inaccurate to predict the RSS value at some point using radial model, the radial model still provides some constraints. If RSS by radial model is -40 dBm, true RSS can not be -90 dBm.

In this case, parameters are assumed to be independent, i.e.,

$$\log p(\boldsymbol{\theta}) = \sum_{m=1}^{M} \log p(\theta_m). \tag{21}$$

For the sake of simplicity, the parameter range is usually expressed by Gaussian form. Then we further deduce that

$$\log p(\theta_m) = -\frac{(\theta_m - f)^2}{2\sigma_{\theta_m}}. \tag{22}$$

For example, we want to use another model $f(\boldsymbol{\theta}^m)$ to constrain the linear interpolation model in (15). We get

$$\log p(r^m) = -\frac{(r^m - f(\boldsymbol{x}^m))^2}{2\sigma_{r^m}} - \log(\sqrt{2\pi}\sigma_{r^m}). \tag{23}$$

Now in M-step, we need to maximize the new function as

$$\boldsymbol{\theta}^{t+1} = \arg\max_{\boldsymbol{\theta}} \left[Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right) - \sum_{m=1}^{M} \frac{(r^m - f(\boldsymbol{x}^m))^2}{2\sigma_{r^m}} \right]. \tag{24}$$

3) Constraint for Correlated Parameters: Sometimes, the parameters are not independent. For example, when we think of interpolation model, the RSS values at different calibration points are correlated. The RSS value at one calibration point can somehow be predicted by its neighbours. This correlation can be expressed by the following equation

$$\log p(\boldsymbol{\theta}) = \log p(\theta_m | \boldsymbol{\theta}_{n\neq m}) + \log p(\boldsymbol{\theta}_{n\neq m}). \tag{25}$$

Because in the above equation the second term does not include θ_m, so the maximisation step can be further expressed by

$$\boldsymbol{\theta}_{\boldsymbol{m}}^{t+1} = \arg\max_{\theta_m} \left[Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^t\right) + \log p(\theta_m | \boldsymbol{\theta}_{n\neq m}) \right]. \tag{26}$$

IV. Experimental Results

A. Experiment Setup

We evaluate our algorithm in a typical office building depicted in Fig. 3,which has an area of $50m$x$80m$. There are 14 WiFi access points installed . Usually only 3 to 5 access points can be reached depending on different locations in the building. We take almost equally distributed 138 reference points. The RSS values at these points are measured offline and the mean of measurements are used to build an accurate yet complicated fingerprinting radio map model. To evaluate our learning algorithm, we use a radial model and default parameter as the starting point to derive coarse radio maps. Then we use the radial model to generate RSS values at 138 reference points as initial parameters. In the next step we walked randomly in the building and recorded 155 online unlabelled RSS vectors. By feeding these online data into our learning algorithm described in previous sections we got learned radio maps.

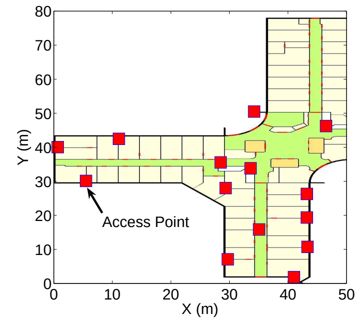

Fig. 3. Test environment

TABLE I
LOCALIZATION RESULTS

	Fingerprinting	Radial Model	EM Learning
Mean of Loc. Error (m)	1.4	12.4	2.27
STD of Loc. Error (m)	0.61	6.35	1.18
Max of Loc. Error (m)	2.9	24.51	4.89
Mean of RM Error	-	28.47	5.36

B. Results

We compared the mean of localization errors, standard deviation of localization errors ,the max localization error and mean of radio map errors for accurate radio maps by mass fingerprinting, coarse initial radio maps by radial model and the learned radio maps. Table I shows the results. We clearly see that the EM learning algorithm improves the localization performance remarkably even without any calibration.

V. Conclusions and Future Work

In this paper, we demonstrated that an improved localization performance can be achieved by learning the unlabelled online data, even without calibration. A better result is expected if some level of calibration is incorporated. We also briefly discussed the effect of EM extensions. This framework can be applied for all RSS-based positioning systems. In the future we will learning other models rather than only fingerprinting model and also analyse the effect of some chosen parameters.

References

[1] P. Bahl and V. N. Padmanabhan, "RADAR: an building RF-based user location and tracking system," in *Proceedings of IEEE INFOCOM 2000*, 2000.

[2] S. Tadakamadla, "Indoor local positioning system for ZigBee, based on RSSI," Master's thesis, Mid Sweden University, 2006.

[3] A. Fujii, H. Sekiguchi, M. Asai, S. Kurashima, H. Ochiai, and R. Kohno, "Impulse radio UWB positioning system," in *Proceedings of IEEE Radio and Wireless Symposium, 2007*, 2007.

[4] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society*, vol. B 39(1), pp. pp. 1–38, 1977.

[5] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.

[6] M. Briers, A. Doucet, and S. Maskell, "Smoothing algorithms for sate-space models," *IEEE Transactions on Signal Processing*, vol. 33, pp. 2951–2970, 2004.