

Never Use Labels: Signal Strength-Based Bayesian Device-Free Localization in Changing Environments

Peter Hillyard and Neal Patwari

Abstract—Device-free localization (DFL) methods use measured changes in the received signal strength (RSS) between many pairs of RF nodes to provide location estimates of a person inside the wireless network. Fundamental challenges for RSS DFL methods include having a model of RSS measurements as a function of a person’s location, and maintaining an accurate model as the environment changes over time. Current methods rely on either labeled empty-area calibration or labeled fingerprints with a person at each location. Both need to be frequently recalibrated or retrained to stay current with changing environments. Other DFL methods only localize people in motion. In this paper, we address these challenges by, first, introducing a new mixture model for link RSS as a function of a person’s location, and second, providing the framework to update model parameters without ever being provided labeled data from either empty-area or known-location classes. We develop two new Bayesian localization methods based on our mixture model and experimentally validate our system at three test sites with seven days of measurements. We demonstrate that our methods localize a person with non-degrading performance in changing environments, and, in addition, reduce localization error by 11 – 51% compared to other DFL methods.



1 INTRODUCTION

WIRELESS sensor networks have opened up many opportunities for detecting breaches in physical property, for home automation, and for remotely monitoring the health and activity of home-bound patients. These systems depend on knowing the location and presence of people in an area of interest. Previous research has shown how a person, without an RF tag, can be localized through walls by processing received signal strength (RSS) measurements between many pairs of statically deployed RF nodes. This tag-less based localization technology is known as device-free localization (DFL) [1].

DFL in indoor environments presents many significant challenges. First, the links’ RSS are non-stationary in changing environments. Consequently, DFL methods that require an empty-area calibration [2], [3], [4], [5] or fingerprint training [6] will need frequent recalibration or retraining to adjust to a changing environment. In contrast, online calibration methods [7], [8] quickly adjust to a changing environment; yet, they can only locate people while they are in motion. For real-world installations of DFL, both DFL types may be unacceptable. For example, a home-bound patient may find it very inconvenient to leave their home to provide a DFL system an empty-area calibration period. They may find it equally inconvenient to retrain their system by labelling training data with true locations. Also, smart-home systems must estimate occupancy, not only motion, in order to control lighting, heating and cooling systems.

One contribution of this paper is that we develop and validate a localization system that addresses the draw-

backs of traditional empty room calibration, fingerprint training, and online calibration methods. The significance of our contribution is illustrated by stepping through the installation and unlabelled training process. In our system, which is built for a single occupant building, nodes are first installed around the building. After nodes have been deployed, the occupant walks around in the building during a system setup period. Alternatively, in the case of the home-bound patient, the node installer walks around the building while the home-bound patient remains stationary. During system setup, unlabelled training data is fed to the system and used to estimate model parameters. After system setup is complete, the system localizes the building occupant. Unlabelled training data continues to be fed to the system after system setup to adapt to non-stationary RSS measurements in changing environments, a method we call *continuous recalibration*. Using unlabelled training data from a person naturally moving inside the building provides a more convenient way of training the system and to locate both a stationary and moving person.

Other continuous recalibration methods with unlabelled training were presented in [9], [10]. However, these methods assumed that an empty-area calibration was performed before runtime, unlike our methods in which no empty-area calibration is assumed. In [9, §VI-D], an approach is proposed in which, instead of performing an empty-area calibration, it is estimated as a person walks inside the building, however, the localization performance when using this idea is not reported. This paper quantifies localization performance of a method in which a person moving inside the building provides unlabelled training data. Our experiments include intentional and unintentional changes to the background that allow us to evaluate our algorithms over time as the environment changes. Further, we note

- P. Hillyard is with Xandem Technology, Salt Lake City, USA. Contact e-mail: peter@xandem.com. N. Patwari is with the Dept. of Electrical and Computer Engineering, University of Utah, and Xandem Technology, Salt Lake City, USA.

that this method is complementary to the developments of [9], which performs filtering and smoothing to both image and coordinate estimates to refine the person's track over a period of time so that model parameters can be as accurate as possible. Our method primarily estimates areas where a person is very likely *not* located as a means to improve knowledge of link model parameters.

A second challenge with indoor DFL is that, because of multipath fading, it is difficult to model the effect of a person's location on the measured RSS of a link. With a fine resolution of fingerprint locations, fingerprint training can capture this relationship in a particular environment by measuring the relative frequency of RSS measurements at each location [11]. While fingerprinting can be very accurate in localizing people, the fingerprints must be frequently retrained to stay current in changing environments [12]. Alternative DFL methods like radio tomographic imaging (RTI) [2], Bayesian methods [13], and particle filters [5], [14], [15], [16] provide more flexibility for DFL because the relationship between measured RSS and a person's location is modeled *a priori*.

Fundamental to model-based DFL is that a link is *affected*, i.e., has significant measured changes in RSS, when a person is on the link line. The link line is the imaginary line segment connecting the link's nodes. In contrast, when the link is *unaffected*, i.e., has very little measured change in RSS, the person tends to be off of the link line. Model-based DFL methods have been built around the idea that the link is affected only when a person is inside an ellipse whose foci are the nodes of the link [2], [3], [5]. In reality, the link can be affected even when a person is far from the link line, or unaffected when a person is on the link line. These model inaccuracies confound Bayesian DFL methods.

As a second contribution of this paper, we develop a new mixture model where a link may be affected or unaffected no matter the person's position, but with probabilities that are a function of the person's distance from the link line. In our system, we learn RSS distribution parameters for both the affected and unaffected state of each link and we do so with unlabelled measurements. The weights in our mixture model are derived from a spatial model such that the affected RSS distribution is weighted more when a person is on the link line and weighted less the further the person is from the link line.

We incorporate our new mixture model in two Bayesian localization methods we develop which we refer to as maximum likelihood localization (MLL) and hidden Markov model localization (HMML). These two methods differ from other Bayesian localization methods [4], [5], [15] by incorporating randomness in the affected and unaffected state. MLL and HMML both compute the probability of observing the measured RSS given a person's location. Adding a temporal property to localization, HMML extends MLL by estimating the current location based on the previous location. In that both MLL and HMML operate on the same mixture and spatial models, and only differ in their temporal properties, we refer to them generally as model-based probabilistic localization (MPL). However, we differentiate between the localization method used in MPL as either MLL or HMML.

We experimentally validate HMML and MLL at three separate sites and with over 7 days of measured RSS data.

We demonstrate that MPL does not need an empty room calibration or fingerprint training period, that it adapts to changes in RSS due to a changing background, and that it is capable of localizing a stationary person. We compare HMML and MLL to an RTI method which uses empty room calibration [3], to an RTI method that uses online calibration [8], and to a Bayesian linear discriminant analysis method [6] which localizes with the help of a database of labelled fingerprint measurements. We show that HMML and MLL can match or decrease the localization error by 11 – 55% compared to these other DFL methods. Additionally, by reducing missed detection errors by orders of magnitude and reducing the false alarm rate by a factor of two to four, we show that we can track stationary targets despite changing environments.

2 RELATED WORK

The ability to locate a person indoors using sensors has changed the way we think about security, home automation and smart homes, and aging in place. Some of these sensing systems include: cameras that detect changes in pixel values caused by a person's presence [17]; pyroelectric sensors that detect and locate changes in thermal radiation due to a person's presence [18]; and vibration sensors to localize vibrations from a person walking [19]. Cameras and infrared sensors cannot sense through material opaque to visible light, and vibration sensors must be sensitive enough to detect vibrations on the inside of the home while ignoring ambient vibrations. Our MPL solution, like other RF solutions, is a more appropriate choice for whole-home, through-wall sensing since RF can sense through walls, smoke, and in any lighting condition.

RF sensing systems perform localization in a variety of ways. Ultra-wideband radios can be used in multi-static radar to measure the time-of-flight between pulse transmissions and received reflections caused by moving people [20]. The time-of-flight is proportional to the distance between the reflector and the transmitter and receiver which is used to localize a person. Ultra-wideband radios have also been used measure changes in the line-of-sight power and then perform tomography with those measurements [21]. A person's presence has also been shown to create significant changes in the amplitude of subcarriers in PHY layer measurements of commodity WiFi cards. These changes have been used in a fingerprint classification method to localize a person [22].

Our MPL method is complementary to methods that process RSS to perform DFL including particle filters [14], [15], fingerprint classification [6], [12], and RTI [2], [3]. These methods, however, either need a person to stand at several locations, to have the area completely vacant for a short period [2], [3], [10], [23], or to have the person continuously moving in order to perform localization [8], [24]. What sets our work apart is that MPL can perform DFL without fingerprinting, without a vacant area, and can localize stationary people.

DFL methods vary in how RSS measurements are used to estimate location. In fingerprint-based localization, a person stands at many locations in the area of interest while the statistics of the RSS distributions are recorded [6], [12].

During testing, the likelihoods or Bayesian probability of measuring the observed RSS is computed for each fingerprint from which the estimated location is derived. With enough fingerprint locations, fingerprint-based localization captures the hard-to-predict RSS distribution of a link as a function of a person's location. However, this comes at an unsustainable cost of frequently retraining fingerprints to stay current with changing environments. MPL seeks to provide a highly accurate localization system, like that of fingerprinting, but by doing so with a model that provides more flexibility in changing environments.

In RTI methods, a spatial model is assumed which indicates where a person's presence will cause a change in RSS [2], [3]. An image of the most likely locations a person was is formed based on the change in RSS observed on the links in the network. A more recent version of RTI has been developed to learn the parameters of the spatial model for each link with unlabelled data [9]. Additionally, the empty room calibration measurements are continuously recalibrated as a person moves inside the area of interest to stay current in changing environments. Furthermore, the change in the RSS on a link is weighted as a function of the excess path length of the person's location and the link. MPL provides an alternative approach to localization by computing the probability of observing RSS measurements based on a person's location. Furthermore, our system performs continuous recalibration not only when a person is moving inside the area of interest but also when the area of interest is vacant. Continuous recalibration in both cases is necessary to avoid nuisance alarms when the area of interest is vacant.

Another variation of DFL is particle filtering. As in MPL, particle filters assume RSS distributions, one for when a person is on a link line, and one for when they are off of the link line [4], [5], [14], [15], [16]. Particles are then drawn from a Gaussian distributions and are said to be drawn from the affected RSS distribution when their excess path length to a link are less than some threshold. Otherwise they are drawn from the unaffected RSS distribution. MPL differs from this approach in that we do not place 0 or 1 weights to the affected and unaffected RSS distribution, but soft weights that are a function of the person's excess path length to a link. This approach inserts some uncertainty in the model to account for the reality that a link may not be affected even when a person is standing on the link line or that a link is affected when the person is far from the link line. As another point of differentiation, our methods, unlike [4], [5] actually implement and validate a method that tracks both a stationary target and a target in motion without using labelled calibration data.

3 METHODS

In this section, we describe the fundamental components of MPL generally and of MLL and HMML specifically. The components of MPL are shown in the block diagram in Fig. 1 and include: a one-time estimation of the weights of the mixture models that relate RSS to an occupied location; a lightweight online RTI method that runs in tandem with either MLL or HMML to provide a location of a moving person; a continuous recalibration block that continuously




Fig. 1. Block diagram of model-based probabilistic localization (MPL)

re-estimates the parameters of the links' affected and unaffected distributions; and a block where MLL or HMML is implemented. MLL and HMML compute the probabilities of observing the RSS measurements on the links given the person's location. We describe each of these components in more detail in the following sections.

3.1 Equipment and Measurements

In this paper, our wireless measurements are made using Texas Instruments CC2531 dongles that communicate using IEEE 802.15.4 in the 2.4 GHz ISM band. We deploy N nodes around the area of interest. The nodes are programmed to take turns transmitting a packet on a 802.15.4 channel during dedicated time slots using TDMA and a token-ring passing protocol [25]. This protocol is repeated on a predefined set of 802.15.4 channels.

As each node transmits on each channel, a separate node logs the RSS, also called the received power in decibel units, between each pairwise node. We denote the RSS measured on link $l = (i, j, c)$ formed by transmitting node i and receiving node j on channel c as r_l . The RSS is typically a discrete-valued measurement, and we denote its possible values as \mathcal{S}_r . \mathcal{S}_r also includes \emptyset , the event that there was a missed packet and as such RSS was not measured. We observe a vector $\mathbf{r} = [r_1, r_2, \dots, r_L]$ on L links.

3.2 Mixture Model

As in many model-based DFL methods, MPL adopts the idea that a link is either in an *affected* state or an *unaffected* state [13], [15], [16]. However, the novelty in this model is that, given the person's location, the state of the link is not known *a priori*. In contrast, some models state that a link deterministically is affected when a person is present in an ellipse whose foci are the node coordinates of the link and is unaffected when the person is outside of the ellipse [2], [3], [7], [15]. Another model states that a person's presence in a voxel intersected by the link line causes the link to be deterministically affected [26]. However, by virtue of the random nature of the multipath radio channel, any deterministic model for the state of a link as a function of person location is bound to be inaccurate. Further, if pixels are large (to keep computation time low), there may be, within the pixel, positions in which the person affects a link as well as positions in which they do not affect it. We develop a mixture model that places some uncertainty on the whether a link is affected or unaffected by a person's presence in a pixel.

Our mixture model follows the diagram shown in Fig. 2. First, assume that a person is present at one of $P + 1$ grid coordinates \mathbf{x}_k^{grid} for pixel $k \in \{0, \dots, P\}$. In one time step,




Fig. 2. Bayesian graphical model

a person can transition from one grid point to any other grid point that can be reached by a person moving at the maximum assumed velocity (we use 3 m/s unless otherwise stated). While a person is located in pixel k , a link l has a state $s_l \in \{a, u\}$ where a is the affected state and u is the unaffected state. Our link state probability, given the person's location, is denoted $p(s_l | k)$. The probability a link is affected is an exponentially decaying function of excess path length

$$p(s_l = a | k) = \beta_l \cdot e^{-\delta_{l,k}/\lambda_l} \quad (1)$$

where the excess path length of location k with respect to link l is $\delta_{l,k} = d(\mathbf{x}_k^{grid}, \mathbf{x}_l^{tx}) + d(\mathbf{x}_k^{grid}, \mathbf{x}_l^{rx}) - d(\mathbf{x}_l^{tx}, \mathbf{x}_l^{rx})$, where $d(\mathbf{x}, \mathbf{y})$ is the Euclidean norm between \mathbf{x} and \mathbf{y} , \mathbf{x}_l^{tx} is the coordinate of link l 's transmitter, \mathbf{x}_l^{rx} is the coordinate of link l 's receiver, and β_l and λ_l are parameters we will estimate (see Section 3.5). The probability that a link is unaffected by a person standing in pixel k is $p(s_l = u | k) = 1 - p_l(s_l = a | k)$.

If a link is affected, the RSS on link l , r_l , is generated from the conditional distribution $p(r_l | s_l = a)$. The probability of observing r_l in the affected state is weighted by $p(s_l = a | k)$, the probability that the link is affected given the person is standing in pixel k . If a link is unaffected, r_l is randomly generated from the conditional distribution $p(r_l | s_l = u)$. The probability of observing r_l in the unaffected state is weighted by $p(s_l = u | k)$. Via Bayes' Law, we can see that r_l given a person at position k is generated from the mixture model

$$p(r_l | k) = \sum_{s'=a,u} p(s_l = s' | k) \cdot p(r_l | s_l = s'). \quad (2)$$

Assuming link RSS measurements are independent, the likelihood that the person is in pixel k given \mathbf{r} is

$$p_k(\mathbf{r}) = \prod_{l=1}^L p(r_l | k). \quad (3)$$

However, this product may not be able to be represented by modern computers when L is large. To avoid these issues, we compute log probabilities first and then convert them back into probabilities as $p_k(\mathbf{r}) = \exp\{\sum_l \log p(r_l | k) - \psi\}$ where $\psi = \max_k \sum_l \log p(r_l | k)$. For MLL, the estimated location of the person is found from

$$\hat{\mathbf{x}}^{mll} = \arg \max_{0 \leq k \leq P} p_k(\mathbf{r}). \quad (4)$$

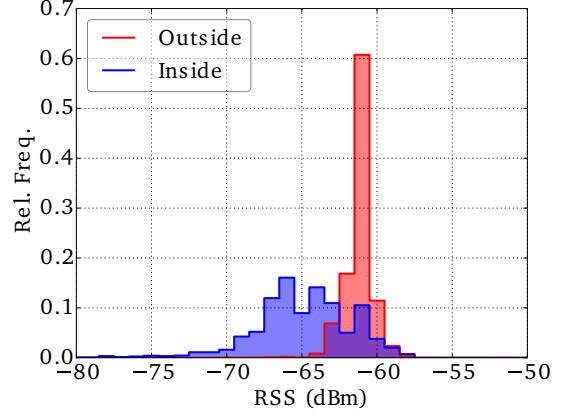


Fig. 3. Distribution of RSS on one link when a person is inside or outside an ellipse whose foci are the transmitter and receiver coordinates of a link. This link represents the behavior of the majority of links.

The HMML solves for the most likely location given a history of \mathbf{r} observations by inductively computing a forward probability vector, $\alpha_k[t]$, at time t for each grid coordinate $k = \{0, \dots, P\}$. The value of $\alpha_k[t]$ is the joint probability of current state k and all link RSS measurements \mathbf{r} through time t [27]. The HMML estimates the current location of the person as

$$\hat{\mathbf{x}}^{hmml}[t] = \arg \max_{0 \leq k \leq P} \alpha_k[t]. \quad (5)$$

The probability that the initial pixel of the person is k is denoted π_k . The forward algorithm initializes $\alpha_k[1] = \pi_k p_k(\mathbf{r}[1])$ where $\mathbf{r}[1]$ is the first measured RSS vector, and then computes $\alpha_k[t+1] = \left[\sum_{w=0}^P \alpha_w[t] p_{wk} \right] \cdot p_k(\mathbf{r}[t+1])$ for each $t > 1$ and for $0 \leq k \leq P$ where p_{wk} is the probability that a person transitions from pixel w to pixel k in one time step. For reference, the grid coordinates are evenly distributed. One coordinate represents the out-of-the-area coordinate which we denote $\mathbf{x}_P^{grid} = [\infty, \infty]$. For this out-of-the-area coordinate, we manually set $p_l(a | \mathbf{x}_P^{grid}) = 1 \times 10^{-3}$.

3.3 Conditional RSS Distributions

In this work, we have adopted the idea that a link is either in an affected or unaffected state. We perform our own experiments to support this claim. In this experiment, a person walks around a room at known times and at known coordinates while we record RSS measurements. We show in Fig. 3 one link's distribution of RSS when a person is far from the link line and when the person is on or near the link line. While the RSS does not assume these distributions in all cases, we use this link to represent the behavior of the majority of links. The affected and unaffected RSS conditional distributions can be modeled as skew-Laplace [15], or Ricean [28], but we sacrifice model accuracy by using a Gaussian model for simplicity. A normal distribution for RSS in decibel units has also been adopted in [6], [16]. The mean and variance of the unaffected distribution we denote as $\mu_{l,u}$ and $\sigma_{l,u}^2$ where the subscript u specifies unaffected. The mean and variance of the affected distribution we denote as $\mu_{l,a}$ and $\sigma_{l,a}^2$ where the subscript a specifies affected.

For link l , we estimate the mean and variance of both distributions using RSS measurements when there is ev-

idence that the link is unaffected. We describe in Section 3.7 how we decide when RSS is measured when the link is unaffected, but for now, we create a FIFO buffer of length B for link l . When a link is unaffected, we add r_l to the buffer. When a measurement is added to the buffer, we compute both the sample mean and the sample variance of the buffer which we save as $\hat{\mu}_{l,u}$ and $\hat{\sigma}_{l,u}^2$ respectively. When there are no changes to objects in the background environment, we anticipate $\hat{\mu}_{l,u}$ to be about the same as $\mu_{l,u}$. So, we only perform the update $\mu_{l,u} \leftarrow \hat{\mu}_{l,u}$ and $\sigma_{l,u}^2 \leftarrow \hat{\sigma}_{l,u}^2$ when $|\mu_{l,u} - \hat{\mu}_{l,u}| > 1$ dBm.

From Fig. 3, we also observe that the mean of the unaffected histograms is a few dBm greater than the affected histogram's mean. Also, the variance of the affected histogram is larger than the unaffected variance. In our model, we use these observations to also estimate the mean and variance of the affected distribution by $\mu_{l,a} \leftarrow \hat{\mu}_{l,u} - \Delta$ and $\sigma_{l,a}^2 \leftarrow \eta \hat{\sigma}_{l,u}^2$. We have found that $\Delta = 3$ dBm and $\eta = 2.5$ are appropriate parameters to use for indoor settings. We also note that to estimate $\hat{\sigma}_{l,u}^2$, we use the maximum of the sample variance of the buffer and a minimum constant $\omega^2 > 0$. Due to quantization of RSS, the sample variance may be zero even though the true real-valued received power would have had a positive variance. We impose a minimum variance of $\omega^2 > 0$ to avoid numerical instability. We have found that $\omega = 0.75$ is an appropriate value for this application.

In reality, the effects of multipath fading would mean each link would have a unique Δ and η instead of the fixed value we use. Although fixed in this paper, the values of Δ and η for all links captures the general RSS response, which is a drop in RSS and an increase in variance, when a person is nearby the link line. The choice to estimate each link's Δ and η from noisy target location estimates would add another layer of high-dimension estimation which we chose not to explore in this paper.

When the mean and variance of a link's unaffected and affected distributions have been re-estimated, we recompute their RSS mass functions as

$$p(r_l | u) = \begin{cases} \epsilon, & r_l = \emptyset \\ \max \left\{ \epsilon, \frac{1}{\gamma} \mathcal{N}(r_l; \mu_{l,u}, \sigma_{l,u}^2) \right\}, & r_l \neq \emptyset \end{cases} \quad (6)$$

and

$$p(r_l | a) = \begin{cases} \epsilon, & r_l = \emptyset \\ \max \left\{ \epsilon, \frac{1}{\gamma} \mathcal{N}(r_l; \mu_{l,a}, \sigma_{l,a}^2) \right\}, & r_l \neq \emptyset \end{cases} \quad (7)$$

where γ is constant such that the pdf (probability distribution function) sums to one, and $\epsilon > 0$ is a small-valued lower bound on the probability value away from zero. The use of the minimum probability ϵ is due to the fact that, in practice, we may observe values far from the mean more often than described by equations (6) and (7) because temporal fading does not always fit the log-normal distribution [29]. Using a small value ϵ conveys the model uncertainty and avoids numerical issues with very low probabilities in the likelihood computations. We use 1×10^{-5} in this work.

3.4 Spatial Model

In Section 3.2, we cited previous work that modeled the affected and unaffected state of the link as deterministic.




Fig. 4. The absolute difference between the mean RSS during empty room and the mean RSS when a person occupied each square. Red squares represent an absolute difference of 2 dBm or less. Blue squares represent an absolute difference greater than 2 dBm. The white squares were never occupied. The nodes and link line are shown in orange.

Our choice for using a decaying elliptical model instead is based on an experiment that we perform in an empty classroom. In this experiment, a person moves inside many $1.22m^2$ areas for 30 seconds each, during which time RSS for many links are measured and recorded. An additional 30 seconds of RSS is recorded when the person is not in the area of interest. In post-processing, we find the mean RSS for each link and for each location the person occupied, including when the person stood outside the area of interest. In Fig. 4, we show the absolute difference in mean RSS when a person occupies each $1.22m^2$ area and the mean RSS when the area of interest in vacant for link l . We threshold the image so that the absolute differences that are greater than 2 dBm are shown in blue, and smaller differences are shown in red. White squares are never occupied.

What Fig. 4 shows is which areas experience a change in mean RSS when a person occupies the area. In Fig. 4, we observe that, in general, areas near the link line tend to result in a decrease in RSS. Areas that are further away tend to experience small differences in RSS. However, we also observe that some locations show no measured change in RSS even when the person is on the link line. Additionally, when a person is very far from the link line, the link's RSS can significantly change in mean. A simple elliptical model does not capture the uncertainties due to multipath fading. The spatial model in equation (1) creates some uncertainty in our mixture model so that there is always a small but significant probability that an RSS measurement was drawn from either the affected or unaffected conditional distribution. When the person is near the link line, the affected distribution is weighted more heavily than the unaffected distribution. The choice of β_l and λ_l gives us some control over how the weights in the model are selected so that we can adjust to the different fading characteristics of each link.

3.5 Estimating Spatial Model Parameters

In this section we describe how we estimate the mixture model parameters β_l and λ_l for each link l . To accomplish this, our goal is to estimate β_l and λ_l such that our mixture model closely matches the distribution of RSS measurements as a function of the excess path length of the person's location and link l . This estimation process refers to the KRTI and parameter estimation block seen in Fig. 1. KRTI is an online DFL method that does not require an empty room calibration period [8]. We choose

KRTI because of its relatively low computational complexity and its highly accurate localization capability. KRTI updates a long and short term RSS histogram with every new RSS measurement. The difference between these two histograms is computed using the kernel distance. The differences from all of the links are then used to form an image and estimate the location of a person. During a training period a person walks inside the area of interest. During this time, KRTI provides an estimated location, $\hat{\mathbf{x}}^{krti}$, for each r_l . We store all $\langle r_l, \delta_l^{krti} \rangle$ tuples where δ_l^{krti} is the excess path length between $\hat{\mathbf{x}}^{krti}$ and link l .

After the training period is complete, we first estimate the mean and variance of the unaffected distribution using, respectively, the median and median absolute deviation (MAD) of the RSS during the training period. We use these statistics to estimate the mean and variance of the unaffected RSS distribution. The median and MAD ignore RSS measurements that fall far from the true unaffected mean and robustly estimate the mean and variance. Additionally, we multiply the MAD by 1.48 and square the value to make it an estimate of the variance for Gaussian data [30]. Once the unaffected mean and variance have been estimated, we apply the same shift to the mean and scale to the variance to get the affected mean and variance as described in Section 3.3.

Instead of basing the affected distribution parameters on the unaffected distribution parameters, an algorithm could discriminate between times when a link is in either the affected or unaffected state and then directly estimate the parameters for the affected and unaffected state. However, we choose to not pursue this approach because it would require a long training period to ensure that there were a sufficient number of RSS measurements from each state to estimate the distribution parameters. Additionally, a person may not be physically capable of reaching locations where they would be affecting a link, and therefore, there would be no RSS measurements to estimate the affected distribution parameters.

After the training period is complete and the affected and unaffected RSS distribution parameters are estimated, we turn to the RSS, excess path length tuples previously mentioned. An example of the tuples for one of the links is shown in Fig. 5. We next divide r_l into bins according to excess path length δ_l^{krti} . We choose to bin all tuples $\langle r_l, \delta_l^{krti} \rangle$ into groups such that their excess path lengths are equal. The possible ordered bin values are in the set $\{\delta_l^{krti}(0), \dots, \delta_l^{krti}(M-1)\}$ where M is the total number of bins. The RSS measurements for one group of these tuples are seen in the grey box in Fig. 5 and the histogram of these RSS measurements is shown in Fig. 6. We denote the histogram of the RSS measurements whose excess path length is $\delta_l^{krti}(m)$ as $\mathbf{h}_{l,m}$ where index m indexes in the set of all excess path lengths. We wish to find $b_{l,m}$ such that the mixture model $p(r_l | b_{l,m}) = b_{l,m} \cdot p(r_l | a) + (1 - b_{l,m}) \cdot p(r_l | u)$ most closely matches $\mathbf{h}_{l,m}$. To do this, we perform

$$b_{l,m}^* = \arg \min_{b_{l,m} \in \mathcal{S}_b} \|p(r_l | b_{l,m}) - \mathbf{h}_{l,m}\| \quad (8)$$

where \mathcal{S}_b is a set of equally-spaced real valued numbers between 1×10^{-5} and 1 and $\|\cdot\|$ is the ℓ^2 -norm. An example optimal mixture model is shown in Fig. 6.

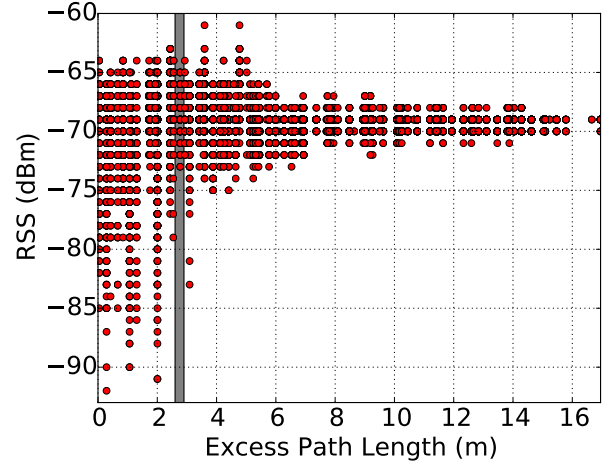


Fig. 5. Measured RSS as a function of excess path length, computed using the estimated location from the KRTI block seen in Fig. 1. One group of RSS measurements with the same excess path length is shown in the gray box.

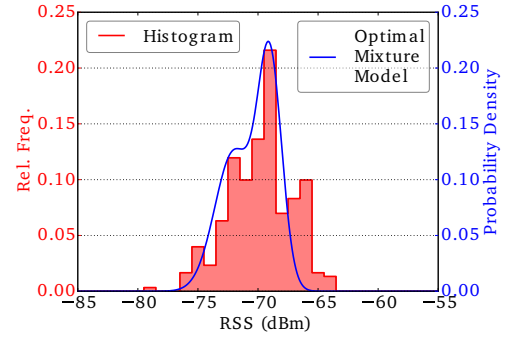


Fig. 6. Histogram of RSS within the gray box of Fig. 5. The mixture model is found using equation 8 for this histogram and is overlaid.

By performing this process for all excess path length bins, we get the tuples $\langle \delta_l^{krti}(m), b_{l,m}^* \rangle$. We plot these tuples for a link in Fig. 7. The relationship between $b_{l,m}^*$ and $\delta_l^{krti}(m)$ follows our spatial exponential decay function in equation (1). We estimate β_l and λ_l from a nonlinear least squares solution. The estimation includes constraining $0 < \beta_l < 1$ to keep the conditional probabilities between 0

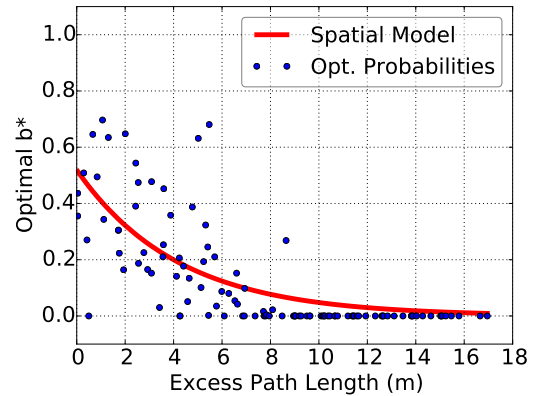


Fig. 7. Optimal probabilities b^* as a function of excess path length. The estimated spatial model is overlaid.

and 1 and $\lambda_l > 0$ so that a link is always more likely to be affected by a person on or near the link line over a person far from the link line. After β_l and λ_l are estimated, we do not re-estimate them in this paper and consequently no longer run KRTI. However, re-estimation can be optionally performed at any later time as desired.

3.6 Initial and Transition Probabilities for HMML

A hidden Markov model includes a transition matrix which defines the probabilities of transitioning from one state to another in one time step. In our model, we incorporate the physical constraints of walking inside a building, like walking speed and fixed barriers, into our transition probabilities. To do this, we first label each grid coordinate as either an entrance-exit or as a non-entrance-exit coordinate. Entrance-exits are locations in the area of interest where a person can enter or exit the area of interest. Second, for each grid coordinate, the grid coordinates that are ≤ 0.75 m away are labelled as neighbors. For entrance-exit states, we include the out-of-area grid coordinate as a neighbor since the only way to leave the area of interest is via an entrance-exit. For the out-of-area grid coordinate, we label the entrance-exit states as neighbors. However, a grid coordinate cannot be a neighbor if a person must travel through a wall to get to that grid coordinate. Third, we assume that a person is more likely to stay at the current grid coordinate than to transition to another.

For transition probabilities, the probability of remaining in the same grid coordinate after one time step is set to 0.9 for all states. For all non-neighbor grid coordinate, we assign a probability 10^{-200} . We found that 10^{-200} was the closest value to 0 we could use without encountering numerical representation issues when computing the forward probabilities. This value also gave some, but very little probability, of a target reaching any other state in one time step. For all neighbor grid coordinate, we assign a equal probability so that the sum of probabilities of transitioning from the current grid coordinate to any other grid coordinate equals 1. We note that wall and entrance-exit information is extra information required to create these transition probabilities. Consequently, we will show in Section 5.4 how HMML's localization performance is affected if we ignore wall and entrance-exit information.

A hidden Markov model also includes the probability π_k that the Markov chain starts in grid coordinate \mathbf{x}_k^{grid} . We assume that when the system turns on, the person is located outside of the area of interest with probability 0.95. All other initial state probabilities are assigned $0.05/P$.

3.7 Continuous Recalibration

An important element of MPL is that it does not use an empty room calibration period, nor a fingerprint training period, to estimate the mean and variance of the affected and unaffected RSS distributions. Instead, MPL uses unlabelled training data when a person is moving inside a building, a feature we believe adds convenience in deploying a DFL system. Furthermore, MPL is capable of adapting to non-stationary RSS distributions. We enable these features of MPL by running a light weight companion localization method called VRTI [7]. Using online calibration, VRTI

localizes motion by computing the sample variance of a buffer of RSS for each link. The sample variance for each link is used to form an image of the motion, from which we estimate a person's location. We denote the location estimate from VRTI as $\hat{\mathbf{x}}^{vrti}$.

The purpose of running VRTI in tandem with MPL is that VRTI can localize a moving person in spite of a changing environment. With VRTI's location estimate, we not only know where the moving person is located but also where they are *not* located. Here, we are assuming our system is used in a home with a single occupant. If no person is near a link, we can safely update that link's RSS unaffected distribution parameters. We say that $\hat{\mathbf{x}}^{vrti}$ is far from a link if the excess path length of $\hat{\mathbf{x}}^{vrti}$ with respect to link l , which we denote $\delta_l^{vrti} = d(\hat{\mathbf{x}}^{vrti}, \mathbf{x}_l^{tx}) + d(\hat{\mathbf{x}}^{vrti}, \mathbf{x}_l^{rx}) - d(\mathbf{x}_l^{tx}, \mathbf{x}_l^{rx})$, is greater than $\delta_l^{max}/2$ where δ_l^{max} is the maximum excess path length of any coordinate in \mathbf{x}_k^{grid} for $k \in \{0, \dots, P-1\}$ with respect to link l . When $\delta_l^{vrti} > \delta_l^{max}/2$, we add r_l to the B -length FIFO buffer referred to in Section 3.3.

In as much as VRTI is unable to distinguish between a stationary person and when the area of interest is vacant, the RSS distribution parameters won't be re-estimated when the person is stationary or when a person is outside of the area of interest. However, it is important to update the RSS distributions when the area of interest is vacant. To re-estimate the RSS distribution parameters when the area of interest is vacant, we add r_l to the B -length buffer, if it has not been added already, when HMML or MLL says the area of interest is empty, i.e. when $k = P$ is the solution to $\arg \max_{0 \leq k \leq P} \alpha_k$ or $\arg \max_{0 \leq k \leq P} p_k(\mathbf{r})$. The mean and variance of the buffer is then used to periodically re-estimate the distribution parameters for both a link's unaffected and affected state as described in Section 3.3. With both the location estimate of VRTI and HMML or MLL, we are able to perform continuous recalibration without an empty room calibration period. We found that $B = 15$ was an appropriate buffer length for our application.

As an example of how we perform continuous recalibration, we show in Fig. 8 the measured RSS on a link before and after a couch nearby is moved 15 cm. After the couch is moved at 2550 s, the unaffected RSS increases by 6 dBm. We also show μ_u for the link during this time period as it is re-estimated. After a few minutes, our unaffected mean RSS estimate adjusts to the increase in RSS due to the changing environment.

3.8 Method Summary

There are several components to our methods and so we summarize the major points here so that the reader can gain a high level understanding of the algorithm.

The algorithm begins with an unsupervised training period that is performed once. While a person walks around the whole space for several minutes, a history of $\hat{\mathbf{x}}^{krti}$ location estimates and RSS measurements are saved (see Section 3.5). After the training period is complete, the history of $\hat{\mathbf{x}}^{krti}$ location estimates and RSS measurements are used to estimate β_l and λ_l for each link (see Section 3.5). The estimated β_l and λ_l values are then used to compute $p(s_l = a | k)$ and $p(s_l = u | k)$ using (1) in Section 3.2. Though not part of the unsupervised training, the transition




Fig. 8. Measured RSS on a link before and after a couch is moved at 2550 s. The unaffected RSS mean changes by 6 dBm. The red line shows how our continuous recalibration eventually adjusts the estimate of the unaffected mean after environment changes.

probabilities used in HMML are also computed just once (see Section 3.6).

After the unsupervised training has been completed, one of the MPL methods is used to perform location estimation. For every new measurement \mathbf{r} , the likelihoods are updated using (2) and (3) in Section 3.2. The likelihoods are then used to update the MLL or HMML location estimate (see Section 3.2). In addition to updating the MPL location estimate, a new \mathbf{r} measurement is used to get a new VRTI estimate $\hat{\mathbf{x}}^{vrti}$ (see Section 3.7). Both $\hat{\mathbf{x}}^{vrti}$ and $\hat{\mathbf{x}}^{mpl}$ location estimates are then used to update $\mu_{l,u}, \mu_{l,a}, \sigma_{l,u}^2$ and $\sigma_{l,a}^2$ which are used in (6) and (7) described in Section 3.3.

3.9 Baseline DFL Methods

To demonstrate the advantages of MPL, we compare HMML and MLL with other DFL methods that lack the same advantages. One of these methods is attenuation-based RTI which we refer to as RTI [3]. RTI requires an empty room calibration where the mean RSS for each link is computed and stored. The absolute difference between \mathbf{r} and the mean RSS is computed and stored as \mathbf{y}^{rti} , which is in turn used to compute an image and estimate the person's location. The second method is kernel-based RTI which we refer to as KRTI [8]. KRTI continuously updates a long and short-term RSS histogram. The kernel distance between these histograms are then computed and stored as \mathbf{y}^{krti} , from which the image and the person's location is estimated. For both RTI and KRTI, we use an elliptical model for the weight matrix \mathbf{W} [2]. A regularized-least squares solution is then used to estimate the image \mathbf{z} using the linear relationship $\mathbf{y} = \mathbf{W}\mathbf{z} + \tilde{\mathbf{n}}$ where $\tilde{\mathbf{n}}$ is the noise. The pixels in the image, \mathbf{z} , for both RTI and KRTI map to the same grid coordinates \mathbf{x}_k^{grid} for $k = \{0, \dots, P-1\}$ mentioned in Section 3.2. We use the pixel with the greatest value as the location estimate, which we denote $\hat{\mathbf{x}}^{rti}$ for RTI and $\hat{\mathbf{x}}^{krti}$ for KRTI. However, when the image maximum falls below a threshold, we set the location estimate as the out-of-the-area pixel \mathbf{x}_P^{grid} . We note that MPL also uses KRTI to estimate spatial parameters, but the baseline KRTI mentioned in this section is separate from KRTI used in MPL. After this point, we distinguish between the two when needed.

The last method is a linear discriminant analysis classifier, which we refer to as LDA, that requires RSS fingerprints

at many locations [6]. During fingerprinting, a person moves inside of a small area around a known location. The mean RSS of all L links is recorded for fingerprint location index k' and stored as $\boldsymbol{\mu}_{k'}^{lda}$ where $k' = \{0, \dots, K'\}$, $K' + 1$ is the total number of fingerprints. The covariance of the RSS over all fingerprint locations, $\hat{\Sigma}$, is then estimated using Ledoit-Wolf shrinkage as $\hat{\Sigma} = (1 - \nu)\Sigma' + \nu\rho I$ where $\Sigma' = \sum_{k'=1}^{K'} \sum_{t \in class_{k'}} (\mathbf{r}[t] - \boldsymbol{\mu}_{k'}^{lda})(\mathbf{r}[t] - \boldsymbol{\mu}_{k'}^{lda})^T / (T - K')$ and T is the number of RSS measurement vectors measured during fingerprinting. Ledoit-Wolf shrinkage is a traditional way to estimate a covariance matrix when the number of samples used for estimation is small but the number of variables to estimate is high. We find ourselves in this situation since the number of measurements we record at each fingerprint tends to be small. Finally, we find the k' that maximizes $\mathbf{r}^T \hat{\Sigma}^{-1} \boldsymbol{\mu}_{k'}^{lda} - 0.5 \boldsymbol{\mu}_{k'}^{lda T} \hat{\Sigma}^{-1} \boldsymbol{\mu}_{k'}^{lda}$ which gives us our location estimate $\hat{\mathbf{x}}^{lda}$.

4 EXPERIMENTATION

In this section, we describe the three test sites we used to evaluate the localization performance. We also describe the localization metric used for the evaluation.

4.1 Test Sites

In our evaluation, we perform experiments at three different test sites. At each site, we first collect a training data set which we use to perform supervised fingerprint training for LDA and to perform unsupervised estimation of β_l and λ_l for each link in MPL. Additional testing data sets are then performed. Both training and testing data sets include the known location of the person moving through the area. Experimentation at each test site was performed differently, so we describe each test site individually. The floorplans for each site are shown in Fig. 9.

There are differences in the number of frequency channels measured for each site. We programmed the nodes for two of the sites to use four frequency channels. The data for the third site was collected by other experimenters and they chose to measure eight frequency channels. We chose to measure on four to have a higher sampling rate. In our algorithms, all channels measured are used in the DFL methods.

4.1.1 Classroom

Our first test site, which we refer to as site CR, is an empty classroom. We deploy twenty nodes, which measure on four channels, on the inside perimeter of the classroom such that a majority of the links are line of sight. It takes the protocol 0.24 s to get an RSS measurement for each link and channel.

A total of 30 s of RSS are collected for fingerprints at 100 locations spaced 0.61m apart. At each fingerprint location, the person moved inside a 0.61m² before moving to the next fingerprint location. The total duration of the training experiment was 55 min.

During the testing experiments, the room was vacant for the first minute. A person then entered the room and continuously moved to each fingerprint location at least once. A total of twelve test experiments were performed. No objects inside the room were intentionally moved at any time during the training and testing data sets.

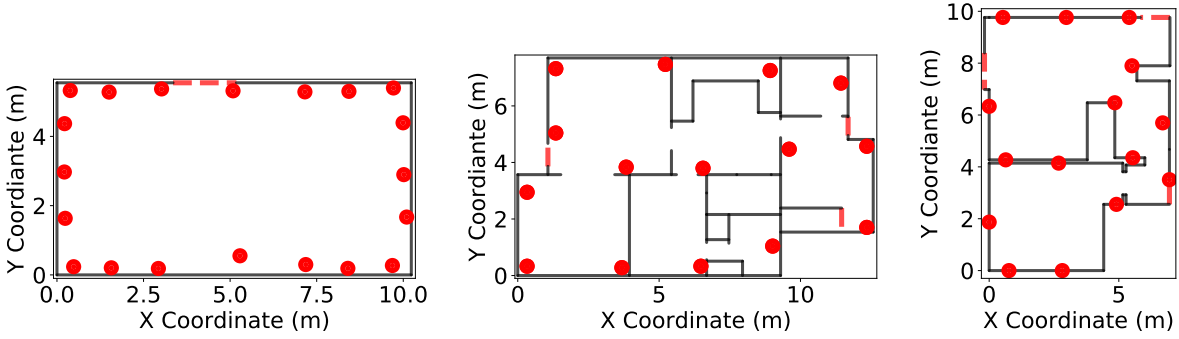


Fig. 9. Experiment locations where the walls or barriers are shown in black and the nodes are shown as red circles. The red dashed lines indicate where entrance-exits are located for (left) class room, (middle) first floor, and (right) basement.

4.1.2 First Floor

The second test site, which we refer to as site 3F, is the furnished first floor of a home. Thirty nodes are deployed, which measure on eight channels, on the inside perimeter of the house. It takes the protocol 0.82 s to get an RSS measurement for each link and channel. A pair of nodes are attached to a tall stand such that the nodes are 0.3 and 1.3 m above the floor.

RSS Fingerprints are collected at thirty-two locations in the house. During the testing experiments, the house was vacant for the first 50 s. A person then entered the house and moved to the first fingerprint location, standing there for 50 s. After the 50 s elapsed, the person moved to the next fingerprint location where the process continued. The same procedure was followed for the test experiments except that the person stands at each location for 20 s.

After the training experiment and each testing experiment, an intentional change to the house was made. For example, a couch was moved, a washer lid was shut, or a sink was filled with water. These intentional changes were performed to simulate the passage of time in a typical house where objects are moved, added, or removed from the area of interest. We note that this training and testing data set was originally created and used in [12].

4.1.3 Basement Living

Our last test site, which we refer to as site BL, is a furnished basement. We deploy fifteen nodes, which measure on four channels, on the inside perimeter of the area of interest. It takes the protocol 0.53 s to get an RSS measurement for each link and channel.

During the training experiment, a person continuously moved around the basement at known locations at known times. For fingerprinting, we create several reference locations that serve as the fingerprint location since the person was moving for the duration of the training period.

During the testing experiments, the basement was vacant for the first minute. A person then entered and continuously walked around the basement. However, the person also reclined on a bed, and sat in an armchair, on a couch, and in a chair for 2 min each at different points during the experiment. This was done to show how DFL methods with online calibration lose track of a stationary person. RSS measurements were recorded for 24 h for seven days. We divide each day into an individual experiment. Each day, a person's ground truth location was recorded for 14 min for

performance evaluation. During the seven days, the person performed normal day-to-day tasks and activities, including moving furniture and adding, moving or removing other household items in the area.

4.2 Localization Accuracy

The DFL methods we evaluate produce a location estimate for each time t . These methods can also indicate that the area of interest is vacant. To evaluate each DFL method, we compute the localization error at time t as

$$e[t] = d(\hat{\mathbf{x}}[t], \mathbf{x}^{true}[t]) \quad (9)$$

where \mathbf{x}^{true} is the true location coordinate, $\hat{\mathbf{x}}[t]$ is the estimated location coordinate from one of the localization methods, and $d(\hat{\mathbf{x}}[t], \mathbf{x}^{true}[t])$ is the Euclidean distance between the true and estimated location. We then compute the median of $e[t]$ for all t and call it the median Euclidean error e^{med} .

5 RESULTS

In this section, we discuss the localization performance of MLL and HMML and our baseline DFL methods, RTI, KRTI, and LDA. We show how MLL and HMML outperform all of the baseline methods at three different sites, how MLL and HMML robustly localize moving and stationary people, and how MLL and HMML adapt to changing environments. We then make intentional modifications to MLL and HMML and show how their localization performance is affected.

5.1 DFL Method Comparison

For each of the three test sites, we show the median error, e^{med} , for MLL, HMML, and the baseline DFL methods in Fig. 10. For all sites, MLL and HMML outperform or match the baseline methods in localization accuracy. At sites 3F and BL, where there were considerable changes to the environment, MLL and HMML reduce e^{med} by 51% or greater when compared to RTI and LDA. Since the environment changes often at site 3F, RTI and LDA's empty room and fingerprint calibration methods become outdated, resulting in poor performance over time. When the environment does not change, like in the CR experiments, RTI and LDA's localization performance closely matches MLL, HMML and KRTI. However, an unchanging environment, like at site CR, is not likely to exist in most applications.

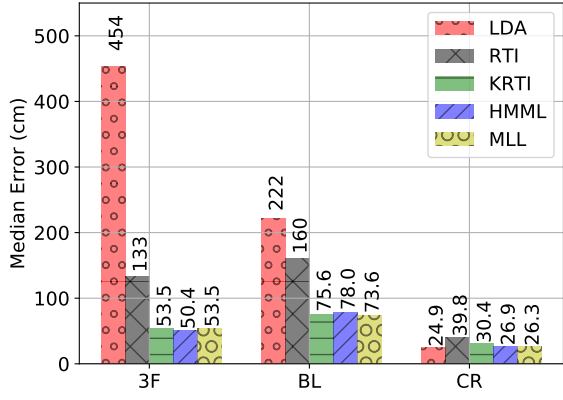


Fig. 10. Median Euclidean error, e^{med} , for LDA, RTI, KRTI, MLL and HMML at three different experiment sites. Only site CR maintained a relatively unchanged environment during the course of the training and testing experiments.

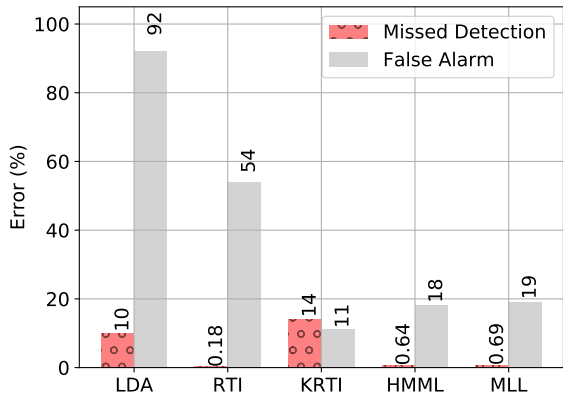


Fig. 11. The missed detection and false alarm percentages incurred by each DFL method in the BL experiment. Both the HMML and MLL achieve less than 1% missed detection error which is two orders of magnitude lower than what was achieved by KRTI and LDA. HMML and MLL were also able to achieve a false alarm rate that was two or more times as low as LDA and RTI.

In contrast, we observe that MLL reduces the median error by up to 13% and HMML by up to 11% compared to KRTI at the three sites. While this reduction in error seems small, it is important to recognize that it becomes more challenging to make significant reductions in error when the errors are already considerably low given the size of the site areas, the number of nodes deployed, and the spatial diversity of the nodes. Additionally, MLL and HMML have the advantage of localizing stationary people, a feature that is missing in KRTI and other online DFL methods.

5.2 Tracking Stationary People Evaluation

To evaluate each method’s ability to localize a stationary target, we show the missed detection and false alarm errors incurred by each method for the BL site in Fig. 11. We show just the results for the BL site since we deliberately included times when the person was stationary during the test.

HMML and MLL both achieve lower missed detection rates than LDA and KRTI by two orders of magnitude. A lower missed detection rate is desirable because it means that the method is better able to localize a stationary target. Interestingly, RTI achieved the lowest missed detection rate.




Fig. 12. Median error for each of the sixteen test experiments performed at 3F. Intentional changes to the environment were made after each experiment. HMML and KRTI adjusts to these changes.. RTI gradually suffers in localization performance as the empty room calibration measurements diverge from measurements made when the system turned on.

However, it also suffered from a 54% false alarm rate which is more than two times as great as MLL and HMML’s rate. Overall, MLL and HMML achieve the best missed detection rate without sacrificing on their false alarm rate.

We note that the false alarm rates for all five DFL methods are very high. One reason for this is because there were few samples when the person was not in the area of interest. As such, falsely detecting even a few of those samples as presence raised the false alarm rate. Had we collected more samples during which the space was vacant, we anticipate that the false alarm percentages would reduce significantly. A second reason why the false alarm rates are high in general for each DFL method is because we chose to weight the risk of losing track of a stationary person higher than falsely detecting a person’s presence. We also desired to have a low localization error. Both of these factors influenced parameter selection and naturally increased the false alarm rate. However, we felt these design choices would be useful in elder care applications when it is more important to keep track of where a person is inside a space while allowing for some false alarms.

5.3 Continuous Recalibration Evaluation

Another important feature of MLL and HMML is that it can robustly localize a person even in changing environments. We show an example of localization performance in changing environments in Fig. 12. The median error is shown for RTI, KRTI and HMML for each of the fifteen test experiments at site 3F. Again, MLL’s performance is similar to HMML, and we therefore did not want to clutter the figure by including MLL’s results. We exclude LDA’s results since its large errors make it difficult to see the differences between HMML and KRTI. HMML, KRTI, and RTI perform equally as well for the first experiment since RTI’s empty room calibration measurements are current. With each successive experiment, intentional changes are made to the environment. As a result, RTI’s localization error increases, even doubling by experiment five. Without frequent empty room calibration, RTI is unable to provide




Fig. 13. Median error for each of the seven test experiments performed at BL. Changes to the environment after each test were the result of day-to-day living. HMML adjusts to these changes with continuous recalibration. RTI and LDA gradually suffers in localization performance as the empty room and fingerprint calibration measurements become out dated.

a reliable location estimate. On the other hand, HMML, as well as KRTI, robustly localizes the person in spite of a changing environment. These same observations can be seen in Fig. 13 which shows the median error for each of the seven test experiments at site BL.

5.4 MPL Feature Evaluation

In this section, we intentionally modify parts of MLL and HMML to see how localization is affected. We make the following four modifications to HMML and MLL.

- First, we fix λ and β for all links instead of estimating them. The values for λ and β are set to achieve the lowest localization error. We call this modification FIXED. We perform this modification for both HMML and MLL.
- Second, we use the true location \mathbf{x}^{true} instead of \mathbf{x}^{krti} to estimate λ and β in the spatial model parameter estimation block in Fig. 1. We call this modification TRUE and make the modification for both MLL and HMML.
- Third, we ignore wall and entrance-exit information when creating the transition probabilities. We call this modification NO WALL but only apply this to HMML since MLL does not use transition probabilities.

The unmodified MLL and HMML we call BASELINE. Only one modification is made to MLL and HMML at a time. For each modification, we perform localization using the data from each site and show the results for MLL in Fig. 14 and for HMML in Fig. 15.

With the FIXED modification, we eliminate the KRTI and parameter estimation blocks as seen in Fig.1. In their place, λ and β are tuned by the user and, consequently, MLL and HMML are ready to run when the system starts. There is no need for any calibration. As seen in Fig. 15, setting all spatial parameters to be the same value for all links increases the median error for HMML by 7 cm at site 3F, reduces the error by 1 cm at site CR, and reduces the error by 9 cm at site BL

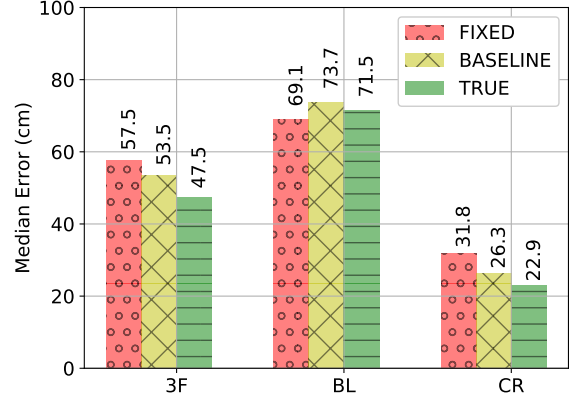


Fig. 14. Median error for MLL at three different experiment sites. In FIXED, we use the same λ and β parameters for all links. In TRUE, the true location \mathbf{x}^{true} is used instead of \mathbf{x}^{krti} to estimate λ and β . BASELINE indicates no modification.

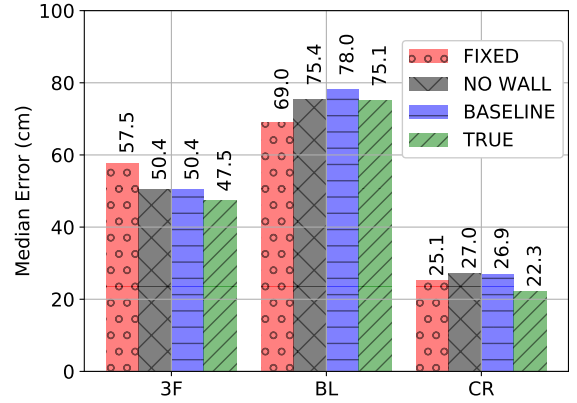


Fig. 15. Median error for a modified HMML at three different experiment sites. In FIXED, we use the same λ and β parameters for all links. In NO WALL, we ignore wall and entrance-exit information when creating the transition probabilities. BASELINE is HMML without modifications. In TRUE, the true location \mathbf{x}^{true} is used instead of \mathbf{x}^{krti} to estimate λ and β .

when compared to BASELINE. From Fig. 14, we observe that setting all spatial parameters to be the same value for all links increases the median error for MLE by 4 cm at site 3F, 5 cm at site CR, and reduces the median error by 4 cm at site BL when compared to BASELINE.

Our first observation is that there are cases when performing a calibration for MLL and HMML actually lead to poorer performance than if a fixed set of parameters were applied. We suspect that there are environments where the estimated location provided by the KRTI block of MPL has a high error which leads to less accurate estimated system parameters. The convenience of the unsupervised calibration may be worth a slight loss in localization accuracy.

With the TRUE modification, we require the user to provide labelled RSS data with their true location during the training phase. We wish to see if having labelled training data would improve the estimation of the spatial parameters, and, in turn, improve localization. From Fig. 15 and Fig. 14, we observe that at all three sites, supplying labelled RSS data matches or decreases localization error for MLL and HMML when compared to BASELINE. The improvement to localization performance can be as great

as 6 cm. The user may decide that the performance gains are too small to warrant having to provide labelled RSS data. However, the trade off is the additional computational overhead needed to run KRTI for the location estimates.

Finally, with the NO WALL modification, we eliminate the need for extra wall information to be entered into HMML prior to operation. From Fig. 15, we observe that the localization error increases at most by 0.1 cm and decreases by 2.5 cm when we ignore wall information. This suggests that there is really no significant loss or gain by including wall information into the transition probabilities. As a result, we can save in some overhead cost in computing the transition probabilities without risking loss in localization accuracy.

5.5 Complexity and Feature Trade Offs

We have shown how MML and HMML achieve a lower localization error than other DFL methods and can do so in a changing environment and without an empty room or fingerprint calibration period. Other DFL methods do not share all of these same properties. The trade off, however, for using MLL and HMML over other DFL methods is the greater memory and computational complexity required to run them. In Table 1, we compare the properties of MLL, HMML and other DFL methods, their calibration requirements, and their memory and computational complexity.

The table shows that DFL systems that require labelled empty-area calibration or labelled fingerprints like RTI and LDA are not able to achieve constant localization performance over the long term unless they are recalibrated or retrained. Needing labelled training can be a major inconvenience, e.g., for a home-bound patient who has limited mobility. DFL systems, like KRTI, that only locate motion does not address the need for occupancy estimation for home automation and in-area monitoring. MLL and HMML do not have either of these drawbacks.

However, the trade off when using MLL or HMML is their relatively higher memory and computational complexity. We observe that HMML has an extra P^2 memory factor which is used to store the transition probabilities and an extra P^2 term in computational complexity which is needed to compute the forward algorithm. The greater memory and computational cost of HMML were used to add a temporal component to localization. We set out to see if the temporal properties of HMML would provide greater localization accuracy than MLL, but we did not observe those gains. HMML is therefore at a disadvantage when P is increased when compared to any of the DFL methods we compared.

The alternative to HMML is MLL, which does not include the extra P^2 memory and P^2 computation like for HMML. We also saw in the previous sections that MLL often performed localization just as well as HMML. Since MLL ignores the temporal component that HMML embraces, it reduces the computational and memory cost by a nontrivial amount. One question to be asked though is, why doesn't HMML benefit from the addition of transition probabilities? We note that both MLL and HMML use the likelihood probabilities in equation (3) which turn out to be values either very near 0 or very near 1 since L is large for all of our experiments. Therefore, the transition probabilities

play an insignificant role when inductively computing the joint probabilities α_k . The localization results of MLL and HMML demonstrate that, at least in the experiments we performed, there is no clear advantage for including temporal properties into the localization problem by using HMML. In general, estimating with temporal properties is helpful since it smooths unlikely jumps in the location estimate. However, in the experiments we performed, the measurement dimension was so high that it resulted in likelihoods that overwhelmed any contribution the transition probabilities could make in HMML. Had this not been the case, and had we been able to sample more frequently, HMML may have been able to reduce the number of big jumps in the location estimate.

6 CONCLUSION

In this paper, we have presented a new signal strength-based Bayesian device-free localization system called model-based localization (MPL) that can localize stationary people, does not require an empty room calibration period, and achieves constant localization performance in changing environments. We developed a new mixture model where the probability of a person occupying a location is a function of signal strength measurements from a wireless sensor network. Our mixture model allowed for uncertainty in the state of the link as a function of the person's location. We developed two realizations of MPL including MLL and HMML which compute the probabilities of a person's location based on the RSS measurements observed. We also developed a method to continuously recalibrate our model to a changing environment.

To validate the performance of MLL and HMML, we performed a series of experiments at three different sites and compute the localization error of MLL, HMML and three other DFL methods. We demonstrated that MLL and HMML outperform the baseline methods in terms of localization accuracy, that MLL and HMML are capable of localizing a stationary person when other baseline methods cannot, and that MLL and HMML achieves constant localization performance even when the environment changes. For assisted living and home automation applications, MPL offers an important advantage of constant localization performance and tracking stationary people without significant costs in computational complexity, memory usage, or convenience.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation under grant #1407949 and the Army Research Office under grant #69215CS.

REFERENCES

- [1] M. Youssef, M. Mah, and A. Agrawala, "Challenges: Device-free passive localization for wireless environments," in *Proceedings of the 13th Annu. ACM Int. Conf. on Mobile Computing and Networking*, ser. MobiCom '07, 2007, pp. 222–229.
- [2] J. Wilson and N. Patwari, "Radio tomographic imaging with wireless networks," *IEEE Trans. on Mobile Computing*, vol. 9, no. 5, pp. 621–632, May 2010.

TABLE 1

Features offered by DFL methods along with their calibration requirements and memory and computational complexity. For reference, L is the number of RSS measurements, P is the number of grid coordinates, and R is the number of bins used for histograms. The costs for continuously running VRTI for HMML and MLL are included in their complexity.

DFL Method	Initial Calibration				Constant perf. in changing env.'s	Localize		Complexity	
	Labelled Empty	Labelled Fingerprint	Unlabelled Occupied	None		Motion	Stationary Person	Memory	Computation
LDA		X				X	X	$LP + P$	$LP + P$
KRTI				X	X	X		$LP + 2LR$	$LR + LP$
RTI	X					X	X	LP	$L + LP$
HMML			X		X	X	X	$2LP + 2LR + P + P^2$	$L + 2LP + P^2$
MLL			X		X	X	X	$2LP + 2LR$	$L + 2LP$

- [3] O. Kaltiokallio, M. Bocca, and N. Patwari, "Enhancing the accuracy of radio tomographic imaging using channel diversity," in *2012 IEEE 9th Int. Conf. on Mobile Ad-Hoc and Sensor Sys. (MASS 2012)*, Oct 2012, pp. 254–262.
- [4] S. Savazzi, S. Sigg, M. Nicoli, V. Rampa, S. Kianoush, and U. Spagnolini, "Device-free radio vision for assisted living: Leveraging wireless channel quality information for human sensing," *IEEE Signal Processing Magazine*, vol. 33, no. 2, pp. 45–58, March 2016.
- [5] S. Savazzi, M. Nicoli, F. Carminati, and M. Riva, "A bayesian approach to device-free localization: Modeling and experimental assessment," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 1, pp. 16–29, Feb 2014.
- [6] C. Xu, B. Firner, Y. Zhang, R. Howard, J. Li, and X. Lin, "Improving RF-based device-free passive localization in cluttered indoor environments through probabilistic classification methods," in *2012 ACM/IEEE 11th Int. Conf. on Inform. Process. in Sensor Networks (IPSN)*, April 2012, pp. 209–220.
- [7] J. Wilson and N. Patwari, "See-through walls: Motion tracking using variance-based radio tomography networks," *IEEE Trans. on Mobile Computing*, vol. 10, no. 5, pp. 612–621, May 2011.
- [8] Y. Zhao, N. Patwari, J. M. Phillips, and S. Venkatasubramanian, "Radio tomographic imaging and tracking of stationary and moving people via kernel distance," in *2013 ACM/IEEE Int. Conf. on Inform. Process. in Sensor Networks (IPSN)*, April 2013, pp. 229–240.
- [9] O. Kaltiokallio, R. Jntti, and N. Patwari, "Arti: An adaptive radio tomographic imaging system," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 8, pp. 7302–7316, Aug 2017.
- [10] M. Bocca, O. Kaltiokallio, and N. Patwari, *Radio Tomographic Imaging for Ambient Assisted Living*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 108–130.
- [11] M. Seifeldin, A. Saeed, A. E. Kosba, A. El-Keyi, and M. Youssef, "Nuzzer: A large-scale device-free passive localization system for wireless environments," *IEEE Trans. on Mobile Computing*, vol. 12, no. 7, pp. 1321–1334, July 2013.
- [12] B. Mager, P. Lundrigan, and N. Patwari, "Fingerprint-based device-free localization performance in changing environments," *IEEE J. on Select. Areas in Commun.*, vol. 33, no. 11, pp. 2429–2438, Nov 2015.
- [13] P. Hillyard, A. Luong, and N. Patwari, "Highly reliable signal strength-based boundary crossing localization in outdoor time-varying environments," in *2016 15th ACM/IEEE Int. Conf. on Inform. Process. in Sensor Networks (IPSN)*, April 2016, pp. 1–12.
- [14] X. Chen, A. Edelstein, Y. Li, M. Coates, M. Rabbat, and A. Men, "Sequential Monte Carlo for simultaneous passive device-free tracking and sensor localization using received signal strength measurements," in *Proceedings of the 10th ACM/IEEE Int. Conf. on Inform. Process. in Sensor Networks*, April 2011, pp. 342–353.
- [15] J. Wilson and N. Patwari, "A fade-level skew-Laplace signal strength model for device-free localization with wireless networks," *IEEE Trans. on Mobile Computing*, vol. 11, no. 6, pp. 947–958, June 2012.
- [16] Y. Zheng and A. Men, "Through-wall tracking with radio tomography networks using foreground detection," in *2012 IEEE Wireless Commun. and Networking Conf. (WCNC)*, April 2012, pp. 3278–3283.
- [17] D. Anderson, J. M. Keller, M. Skubic, X. Chen, and Z. He, "Recognizing falls from silhouettes," in *2006 Int. Conf. of the IEEE Engineering in Medicine and Biology Soc.*, Aug 2006, pp. 6388–6391.
- [18] Q. Hao, D. J. Brady, B. D. Guenther, J. B. Burchett, M. Shankar, and S. Feller, "Human tracking with wireless distributed pyroelectric sensors," *IEEE Sensors J.*, vol. 6, no. 6, pp. 1683–1696, Dec 2006.
- [19] M. Mirshekari, S. Pan, P. Zhang, and H. Y. Noh, "Characterizing wave propagation to improve indoor step-level person localization using floor vibration," in *SPIE Smart Structures and Materials+ Nondestructive Evaluation and Health Monitoring*. International Society for Optics and Photonics, 2016, pp. 980 305–980 305.
- [20] M. McCracken and N. Patwari, "Hidden Markov estimation of bistatic range from cluttered ultra-wideband impulse responses," *IEEE Trans. on Mobile Computing*, vol. 13, no. 7, pp. 1509–1521, July 2014.
- [21] B. Beck, X. Ma, and R. Baxley, "Ultrawideband tomographic imaging in uncalibrated networks," *IEEE Trans. on Wireless Commun.*, vol. 15, no. 9, pp. 6474–6486, Sept 2016.
- [22] H. Abdel-Nasser, R. Samir, I. Sabek, and M. Youssef, "Monophy: Mono-stream-based device-free WLAN localization via physical layer information," in *2013 IEEE Wireless Commun. and Networking Conf. (WCNC)*, April 2013, pp. 4546–4551.
- [23] A. Saeed, A. E. Kosba, and M. Youssef, "Ichnaea: A low-overhead robust WLAN device-free passive localization system," *IEEE J. of Select. Topics in Signal Process.*, vol. 8, no. 1, pp. 5–15, Feb 2014.
- [24] A. Edelstein and M. Rabbat, "Background subtraction for online calibration of baseline RSS in RF sensing networks," *IEEE Trans on Mobile Computing*, vol. 12, no. 12, pp. 2386–2398, Dec 2013.
- [25] M. Bocca, multi-SPIN communication protocol. [Online]. Available: <https://sites.google.com/site/boccamaurizio/home/software-data>
- [26] M. A. Kanso and M. G. Rabbat, "Compressed RF tomography for wireless sensor networks: Centralized and decentralized approaches," in *Proceedings of the 5th IEEE Int. Conf. on Distributed Computing in Sensor Sys.*, ser. DCOSS '09. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 173–186.
- [27] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
- [28] M. Ghaddar, L. Talbi, T. A. Denidni, and A. Sebak, "A conducting cylinder for modeling human body presence in indoor propagation channel," *IEEE Trans. on Antennas and Propagation*, vol. 55, no. 11, pp. 3099–3103, Nov 2007.
- [29] H. Hashemi, "A study of temporal and spatial variations of the indoor radio propagation channel," in *PIMRC-94*, vol. 1, Sep 1994, pp. 127–134.
- [30] P. J. Rousseeuw and C. Croux, "Alternatives to the median absolute deviation," *J. of the Amer. Statistical Assoc.*, vol. 88, no. 424, pp. 1273–1283, 1993.



Peter Hillyard Peter Hillyard received the B.S., M.E., and Ph.D. degrees from the Department of Electrical and Computer Engineering, University of Utah. He completed a postdoctoral research position in the Sensing and Processing Across Networks (SPAN) Lab at the University of Utah in wireless RF-based respiratory monitoring technology and now works at Xandem Technology as a Senior Engineer. His research interests include signal processing, estimation and detection theory, and pattern classification.



Neal Patwari is a Professor in the Dept. of Electrical and Computer Engineering at the University of Utah, with an adjunct appointment in the School of Computing. He directs the Sensing and Processing Across Networks (SPAN) Lab, which performs research at the intersection of statistical signal processing and wireless networking, for improving wireless sensor networking and for RF sensing, in which the radio interface is the sensor. His research perspective was shaped by his BS and MS in EE at Virginia

Tech, his research work at Motorola Labs in Plantation, Florida, and his Ph.D. in EE at the University of Michigan. He received the NSF CAREER Award in 2008, the 2009 IEEE Signal Processing Society Best Magazine Paper Award, and the 2011 U. of Utah Early Career Teaching Award. He has co-authored papers with best paper awards at IEEE SenseApp 2012 and at the ACM/IEEE IPSN 2014 conference. Neal has served on technical program committees for IPSN, MobiCom, SECON, and SenSys.