

# CrossSense: Towards Cross-Site and Large-Scale WiFi Sensing

## ABSTRACT

WiFi has recently established itself as a powerful medium for sensing. By measuring how the wireless channel is affected by humans and activities and then correlating the measurement to pre-collected WiFi training samples, rich information ranging from gestures to gait patterns and even vital signs can be identified. However, WiFi sensing solutions currently have limited scale as they require training data to be meticulously collected in each deployment environment and as they are designed for small sets of activities. This paper presents CROSSSENSE, a novel system for scaling up WiFi sensing to new environments and larger problems. To reduce the cost of sensing model training data collection, CROSSSENSE employs machine learning to train, off-line, a *roaming model* to generate, from one set of measurements, synthetic training samples for each target environment. To scale up to a larger problem size, CROSSSENSE adopts a *mixture-of-experts* approach where multiple specialized sensing models, or *experts*, are used to capture the mapping from diverse WiFi inputs to the desired output. The experts are trained offline and at runtime the appropriate expert for a given input is automatically chosen. We evaluate CROSSSENSE by applying it to two representative WiFi sensing applications, gait identification and gesture recognition. We show that CROSSSENSE boosts the accuracy of state-of-the-art WiFi sensing techniques from 20% to over 80% and 90% for gait identification and gesture recognition respectively, delivering consistently good performance – particularly when the problem size is significantly greater than that current approaches can effectively handle.

## 1. INTRODUCTION

WiFi has emerged as a powerful medium for sensing information. By measuring how the wireless channel is affected by humans and their activities, tasks like gait identification [1–4], gesture recognition [5, 6], activity recognition [7, 8] and even vital sign monitoring [9, 10] could be possible. WiFi sensing is particularly attractive for smart spaces as it can be easily deployed using as few as two wireless routers, not requiring instrumenting the users, and being less privacy intrusive than other infrastructure-based solutions such video monitoring [11, 12].

While existing research has demonstrated the vast potential of WiFi as a sensing technique, currently there are two significant drawbacks that limit the uptake of WiFi sensing and the scale at which it can operate. Firstly, deploying WiFi sensing solutions requires a labour-intensive and time-

consuming process of collecting training measurements or fingerprints to characterize how wireless channel metrics, such as channel state information (CSI) or received signal strength indicator (RSSI), are affected by the target (*e.g., a gait or a gesture*). These training measurements must be carefully collected for each target subject or activity from each deployment site. While it may be feasible to collect such data from each occupant of a home, asking each employee or visitor to provide training measurements from each meeting room in a smart office setting is infeasible. Secondly, WiFi sensing solutions currently can handle only small sets of subjects. As shown in the paper, performance of state-of-the-art WiFi sensing techniques decreases quickly as the number of targets goes beyond a handful. This makes deployments to large-scale settings, such as enterprises or campuses, infeasible as these scenarios would require supporting hundreds of users and numerous activities.

We present CROSSSENSE, a system for scaling up WiFi sensing to new environments and larger sensing problems. By addressing both challenges simultaneously, CROSSSENSE not only increases the scale and uptake of WiFi sensing, but opens up new possibilities for innovative applications and services. As an example, one can offer personalized services in *every* room within a company building, hotel or conference center by collecting a training sample from the visitor *once* at the reception; and the system can now support many more users at a time. WiFi sensing could also be part of a multi-factor authentication protocol [13–16] to verify a user’s claimed identity across sites. These exciting applications can only be realized if the WiFi training measurements can be effectively translated and utilized across sites, and if sensing can target a larger problem size. CROSSSENSE has been designed to offer these capabilities.

To enable WiFi training measurements to be collected once and used across sites, CROSSSENSE integrates an artificial neural network based approach for roaming previously collected measurements from another site to automatically generate synthetic sensing model training samples for the new environment. The roaming model is first trained *off-line* and then used for any *unseen* subjects. To reduce the number of samples required for learning a roaming model for a new site or sensing task, we employ *transfer learning* [17] to leverage an existing roaming model. We exploit the fact that the WiFi signal properties that are abstracted by the beginning layers of the neural network are mostly independent of the optimization problem. We reuse these parts of the network across sites and domains, and, in the process, we speed up learning new roaming models considerably.

To enable WiFi sensing to target a larger problem size, we employ a *mixture-of-experts* based approach [18]. The central idea is that, instead of using a single monolithic model for a sensing task (e.g., gait recognition or gesture recognition), we use multiple models (*experts*) where each expert is specialized for a subset of application scenarios and input data. In this way, each model is used only when its predictions are effective. One major advantage of our approach is that new models can easily be added and are selected only when appropriate. The result is a new way of using machine learning for WiFi sensing, with a generalized framework for diverse inputs, application domains and environments.

We demonstrate the benefits of CROSSSENSE by applying it to WiFi-based gait identification and gesture recognition, and consider CSI and RSSI based metrics. We compare to five state-of-the-art WiFi sensing methods [1, 2, 6, 19, 20] and a wireless signal translation approach [21]. We perform an extensive evaluation across three sites, involving 100 users, 40 gestures, and over 1.2 million wireless samples. Experimental results show that CROSSSENSE delivers *the best and the most reliable* performance across evaluation scenarios and sensing tasks. It boosts the accuracy of state-of-the-art solutions by 4 folds, and enables WiFi sensing to work across sites and for significantly larger problem sizes.

The main contributions of this paper are:

- We present an automatic scheme to effectively leverage existing WiFi training measurements to build sensing models for new environments (Section 4);
- We apply, for the first time, transfer learning to effectively reuse the learned knowledge across different sites and problem domains for WiFi sensing (Section 4.3);
- Our work is the first to employ mixture-of-experts for WiFi sensing to scale up the size of the problems that can be supported, delivering significant performance improvement over state-of-the-art techniques. Our generic framework allows new models to be easily added to target a wider range of application contexts (Section 5);

## 2. MOTIVATION

CROSSSENSE addresses two significant drawbacks in current wireless sensing solutions, cross-site generalization and scale of the problems that can be supported. To further illustrate these limitations, in the following we consider gait identification as a representative example of WiFi sensing tasks and demonstrate how the performance of a state-of-the-art method, WIWHO [2], suffers when (a) the sensing model is built from training samples collected from *another* environment, and (b) targeting *more than a dozen* of users.

**Setup.** Our evaluation environments include a hall and a corridor. The two environments are of different sizes and thus have different multipath effects on the wireless signal. More details of our evaluation settings are given in Section 6.1. To collect training measurements, each user walked on a straight line and passed two WiFi devices 20 times. Dur-




Figure 1: The accuracy of WIWHO drops significantly if the model is not tuned for the target environment (a), and if the number of target users is more than a dozen (b).




Figure 2: The CSI patterns for the same person from two different environments are different.

ing testing, each user walked pass the evaluation scene on a straight line 10 times. The task of WIWHO is to map a testing measurement to one of the known users.

**Impact of training data.** The first experiment involves six users – the same number of users used to evaluate WIWHO in [2]. We expect WIWHO to work effectively on this user size. Figure 1a shows that a high accuracy (75% to 87%) can be achieved if WIWHO is built using examples obtained from the testing environment. The performance is comparable to the accuracy of 75% to 80% reported in [2]. However, the accuracy drops to below 25% when the sensing model is built using samples collected from another site.

**Environmental impact.** Figure 2a shows that the CSI amplitudes collected for the same person is visually different across sites. To quantify the differences, we calculate the dynamic time warping (DTW) distance between each pair of the 10 measurements collected for this person. This metric is often used to quantify the similarity between signal measurements [22, 23] – the further the DTW distance is, the less similar two measurements are. The cumulative distribution function (CDF) diagram in Figure 2b suggests that a larger number of samples have a closer DTW distance when the measurements are taken from the same environment over those taken from different environments. The average DTW distance is increased by 40% when the wireless measurements are taken from a different environment, indicating that the testing environment has a great impact on wireless channel metrics. We also found that using directional antennas does not eliminate the problem, which is in line with the finding in prior research [19]. Clearly, a sensing model tuned for the hall will be ineffective for the corridor.

**Impact of problem sizes.** In the second experiment, we collect the training samples from the target environment, but we consider more target users. Figure 1b shows that the accuracy drops quickly as the number of users increases. The accuracy drops from around 80% to below 60% and 30% when targeting more than 10 and 60 users, respectively. Later in

this paper, we demonstrate that the same issue is observed for gesture recognition – as the number of gestures increases, the accuracy of current sensing methods decreases greatly.

**Lessons learned.** This example shows that current approaches have to collect the training samples of each target subject or activity from each deployment environment and only work for a small problem size. We want to have a technique that can utilize the WiFi training samples collected in one site for each new environment. If we can achieve this, we can then build sensing models for new environments using the same set of training samples at a low cost. We would also like the sensing method to work effectively on a larger problem size to support more subjects and activities. CROSSSENSE has been designed to address both issues.

### 3. OVERVIEW OF CROSSSENSE

The focus of our work is on enabling WiFi sensing (1) to work efficiently across sites by using a single set of training measurements, and (2) to scale to large problem sizes.

To enable cross-site sensing, CROSSSENSE employs machine learning to learn a function to *roam* the WiFi training measurement collected from one environment to another. The roaming model takes in WiFi measurements collected from the data collection site and generates synthetic training instances as if the synthesized data were collected from the target environment. The model is trained *off-line* on a set of examples collected from the data collection site and each deployment site. The learned model is used to translate training measurements for any *unseen* subjects or activities. Crucially, the roaming model needs to capture how the environment affects the WiFi signal from a small set of examples, so that we can apply the model to synthesize a larger set of *unseen* measurements. This is detailed at Section 4.

To scale WiFi sensing to larger problem sizes, CROSSSENSE uses multiple models to map the *on-site* WiFi measurement to the desired output. It then uses a classifier to select which expert should be used for the input measurement. Depending on the characteristics of the input signal, different models can be dynamically selected. This is different from all prior WiFi sensing studies that employ a *one-size-fits-all* model. We argue that a single model is unlikely to precisely capture the diverse inputs for a large problem size. This is described in more details at Section 5.

### 4. ROAMING WIFI MEASUREMENTS

CROSSSENSE uses an artificial neural network (ANN) for roaming WiFi sensing measurements across environments. The motivation for using ANN is twofold. Firstly, ANN can capture both linear and non-linear relationships, making it well suited for generalizing across diverse environments. For example, open halls tend to have different multipath effects than small corridors, requiring a flexible way to model diverse types of relationships. Secondly, ANNs are well-suited for transfer learning and can be used to reduce training costs




Figure 3: Training the roaming model.




Figure 4: Our roaming model is a multi-layer neural network. It takes a WiFi measurement and produces a synthetic sample for the target environment.

for new environments (Section 4.3). As part of our evaluation we also consider other techniques and conclude that ANNs have the best overall performance (Section 7.7).

The input to the roaming model are numerical feature vectors extracted from previously collected WiFi training measurements, and the output is the *expected* feature values of those measurements in the target environment. The set of features used is described in Section 5.1. Building and using such a model follows the 3-step process for supervised machine learning: (i) generate training data (ii) learn a model (iii) use the model. These steps are described as follows.

#### 4.1 Building a Roaming Model

Figure 3 illustrates the process for building a roaming model. We first collect some WiFi signal measurements from the deployment and the data collection sites. We then extract important quantifiable properties, or *features*, of the raw signal data. Finally, we apply a learning algorithm over the collected data to construct the roaming model.

Figure 4 depicts the structure of our roaming model, which is a fully connected, feed-forward ANN with 7 hidden layers. The number of nodes of the input and the output layers is determined by the dimensionality of the feature set (see Table 5.1). Keeping the network structure simple is essential for learning an effective model from a relatively small training dataset. We also evaluate various neural network structures. This is discussed in Section 7.7.

##### 4.1.1 Generate Roaming Model Training Data

In this work, roaming model training data are collected from some users. The subjects (e.g., users or gestures) involved in training do not need to be the same as the ones to be targeted during deployment. This is because the goal is to learn the wireless representation translation for a particular task (e.g., gait identification or gesture recognition) rather than for a specific subject. Specifically, the roaming model learns from training samples how to map a WiFi measure-

ment from one environment to another. Once the model has learned the mapping, it can then apply the learned knowledge to measurements of unseen subjects. In our experiments, training data consists of gait and gesture data collected from the data collection site and each deployment site. The initial training involves 50 and 10 users respectively to learn a roaming model for gait identification and gesture recognition for *the first* deployment site. Using *transfer learning* [17], the number of roaming model training samples can be reduced by 4x when targeting additional sites without sacrificing performance (see Section 7.5).

**Noise control.** Unlike past work that asks users to repeat an activity a fixed number of times, we use statistical reasoning to determine the times on a per user basis. Specifically, we calculate the confidence range of the measurements based on a 95% confidence interval for four metrics: the maximum, average, minimum, and standard deviation values of each user’s measurements; we make sure that the confidence range of each metric is within 10% of its mean. For example, in our experiments, a user on average needs to walk 10 times and repeat a gesture 5 times to gather clean data. Walking pass wireless receivers and performing a gesture take less than 5 and 2 seconds respectively.

**Training instance pairing.** Recall that we need to learn, from the roaming model training data, how to map a subject’s WiFi measurement from one site to the other. The mapping relation must be captured by the training examples. As such, a roaming model training instance consists of measurements gathered from the data collection site ( $m_f$ ) and the deployment site ( $m_d$ ) for a training subject or activity. Because we have taken multiple measurements from both sites, a sample from one site can be mapped to any of the samples gathered from the other site. To determine the mapping, we calculate the Pearson product-moment correlation coefficient [24] between each possible mapping from  $m_f$  to  $m_d$ . To form a training example, we choose the pair,  $(m_f, m_d)$ , which has the strongest correlation coefficient. Furthermore, instead of presenting the raw measurement values to the learning algorithm, we transform the WiFi signal data into a fixed length feature vector. As detailed in Section 5, we use a range of features to capture diverse behaviors. Therefore, we generate, from the raw signal measurements of each site, a training dataset for each feature set.

#### 4.1.2 Learning a Roaming Model

After gathering sufficient training data, we build the roaming model by employing a learning algorithm. The algorithm takes in feature values of the data gathered from (1) the data collection site (*input*) and (2) the deployment site (*output*). It produces a roaming model for each feature set. Training is a *one-off* cost performed off-line, and the learned model can be used for any *new, unseen* subjects without incurring extra training. Therefore, the training cost can be amortized when the number of activities to be targeted is greater than that involved in training the roaming model, which is likely

to be the case for large-scale deployments.

Our roaming model is trained using Back-propagation with Stochastic Gradient Descent (SGD). For a set of training examples  $X_1 \dots X_n$ , SGD attempts to find the network parameters  $\Theta$  that minimize the output of a loss function:

$$\Theta = \arg \min_{\Theta} \frac{1}{n} \sum_{i=1}^n \ell(X_i, \Theta) \quad (1)$$

where loss function  $\ell(x, \Theta)$  computes the logarithmic difference between the model output and expected values.

**Training cost.** The total training time is comprised of three parts: gathering the raw wireless channel metrics, processing the raw data, and then building a model. Gathering the raw data consumes most of the total training time, in this work it took about 4 hours to collect gait and gesture data from 50 users from one environment. In comparison processing the raw data and learning the model took a negligible amount of time, less than 20 minutes on a multi-core PC.

## 4.2 Deployment

Once we have built and trained the roaming models – one model per site, per feature set, we can use them to quickly translate the sensing model training measurements for each target subject from one site to another.

## 4.3 Transfer Learning

Prior work in machine learning has shown that ANN models trained on similar inputs for different tasks often share useful commonalities. The idea is that the information learned at the early layers of a neural network (i.e., those closer to the input layer) are used to capture the input and are mostly independent of the optimization goal. The later the network layers are, the more specialized the layers become.

We hypothesize that this would also be the case for wireless signal. That is, we can utilize a trained roaming model between sites A and B to speed up the process for learning a model between sites A and C. This is because the first few layers of our ANN model are likely to focus on abstracting the input WiFi channel metrics and are largely independent to the model output. Since we use the same network structure, transfer learning is as simple as copying the learned weights of the learned model to initialize the new network. Then we train the model as normal but with fewer instances. In Section 7.5, we show that transfer learning greatly reduces the training cost across deployment sites and sensing tasks.

## 5. SCALING TO LARGE PROBLEM SIZES

CROSSSENSE uses multiple distinct expert models to scale WiFi sensing to larger problem sizes. The expert models are built from the WiFi training measurements of the target subjects such as users or gestures. When targeting a new environment, we use the synthetic training samples produced by our roaming model to tune the expert models for the new environment. During deployment, an expert selector decides which model to use, based on the on-site signal measurement. This process is illustrated at Figure 5.




Figure 5: Overview of our mixture-of-expert approach.

#### Statistical Features

*Time domain*: min/max, min/max 10th/90th, mean, std. deviation, etc.

*Frequency domain*: energy, domain-frequency ratio, FFT peaks;

#### Compression Features

3-level discrete wavelet transform on the signal data;

#### Spectrogram Features

normalized energy at the main frequency, the maximum, minimum, average, and variance for the movement speed, etc.

#### Transformed Features

The signal shape; the autocorrelation and the partial autocorrelation functions.

Table 1: Wireless signal features used in this work

Naive Bayes (NB)	Random Forest (RF)
Support Vector Machine (SVM) w/ the RBF kernel	SVM with a linear kernel
K-Nearest Neighbor (KNN)	Adaboost

Table 2: Classification techniques used in the work

## 5.1 Wireless Signal Features

Table 1 gives the set of features considered in this work. We use a range of features that were found to be useful in prior work of WiFi sensing, but new features can be added. These features are grouped into three categories: (1) statistics of the wireless channel [2, 5, 20]; (2) compression features [25, 26] for which we apply the Discrete Wavelet Transform [27] to the wireless channel metrics to sample the time and frequency information; (3) spectrogram features such as the normalized energy at the main frequency [1]; and (4) transformed features for which we convert the raw wireless channel measurements to a different form where discriminative features can be easily found [28].

Supervised learning typically requires feature values lie in a certain range. Therefore, we normalize the value for each of our features to the range between 0 and 1. We record the maximum and minimum value of each feature found at the training phase, and use these values to scale features extracted from a new signal measurement during deployment.

## 5.2 Expert Models

To construct the expert models, we consider six classification techniques listed in Table 2. For this work, a classifier takes in a feature vector of the input measurement to predict the target’s label that the input corresponds to.

Each classifier can be used together with one of the feature sets given in Table 1. To decide which feature-model combinations should be used as an expert, we perform cross-validation on the *translated* training measurements. We break the dataset into two parts. We train a feature-model combination on one dataset (*expert model training* set), and test the trained model on another dataset (*expert model testing* set). We keep a combination if it either gives the highest accuracy

ID	Expert: (Feature, Model)	ID	Expert: (Feature, Model)
1	(Transformed, SVM-Linear)	2	(Spectrogram, SVM-RBF)
3	(Statistics, RF)	4	(Transformed, RF)
5	(Transformed, SVM-RBF)	6	(Statistics, SVM-RBF)
7	(Statistics, NB)	8	(Transformed, Adaboost)
9	(Compression, SVM-RBF)	10	(Compression, KNN)

Table 3: Chosen expert models, sorted by capabilities.




Figure 6: Percentages of samples that can be correctly classified by an expert on our expert model testing dataset.

or can correctly classify testing cases where others fail. As a result, we keep the 10 combinations given in Table 3.

Expert models are tuned for each environment using the *translated* WiFi measurements produced by our roaming model. Learning in this context is to find model parameters that can closely correlate the feature values of the *translated* measurements to the corresponding subject or activity.

Figure 6 shows the accuracy of each selected expert model on the *expert model testing set* for gait identification. While no individual expert can correctly classify over 50% of the test cases, collectively, they can successfully identify all.

We note that expert models are automatically learned from training data, treating the problem domain and data patterns as black boxes; new applications and data patterns would be learned similarly, potentially causing the addition of new models. Like the roaming model, expert training and selection are also performed *off-line* and are an one-off cost.

## 5.3 Expert Selector

After learning individual expert models, we need to have a mechanism to determine which expert to use for a given on-site WiFi measurement. Our expert selector is a KNN classifier where  $k$  is set to 3. We use the *translated* WiFi training measurements of the target subjects or activities as *fingerprints* for choosing expert models. To determine which expert models to use, we find which of the three fingerprints are most closest to the input collected during deployment. We then use the experts that are found, during training, to be effective for the closest neighboring fingerprints. Neighbor evaluation is performed by applying the DTW algorithm [29] to the de-noised wireless channel metrics. We choose to use DTW because it is a well-established method for measuring the similarity between two temporal sequences [22].

We use the Principal Component Analysis (PCA) based method in [8] to reduce the noises and dimensionality of the raw measurement. We then choose the first 30 principal components given by PCA, which account for over 95% of the data variances. We record the PCA coefficients and use them to transform the input collected during deployment. The KNN algorithm then works on the de-noised input.

Modern WiFi devices typically have multiple transmit and receive antennas and a wireless channel is typically divided into multiple subcarriers. Therefore, the wireless channel metric can be measured independently from each subcarrier. For this reason, we first use KNN to select three experts from each subcarrier measurements, and then use majority voting to choose the final expert model.

Intuitively, learning an expert selector is easier than learning a one-size-fits-all classifier. This is because there are less classification labels to learn – as the number of expert models is typically smaller than the target problem size. It is also because there are often multiple experts that can correctly classify a given input and the expert selector only needs to pick one of them. We also note that one advantage of using KNN is that the distance used for expert selection can also be used to measure the prediction confidence, which essentially provides a degree of soundness guarantee. For example, if an on-site measurement is too far from any fingerprint, this might mean that either we do not have a fingerprint for the target, or the environment has changed significantly where a retraining of the roaming model is required.

## 5.4 Computational Cost

The computational overhead of mixture-of-experts mainly comes from training the expert models. Expert model training is a one-off cost unless new subjects or activities are targeted. It takes less than 15 minutes to train in parallel our 10 expert models on a multi-core server. This can be further accelerated using multiple servers to provide a near real-time model update. The overhead of selecting and using the expert is negligible, as only one model is used at a time. In our case, the expert selector takes less than 1 second to run, and classification takes less than 0.5 second by running the most expensive expert model, Adaboost, on a PC.

## 6. EXPERIMENTAL SETUP

We thoroughly evaluate CROSSSENSE using over one million wireless measurements collected from 100 users across three sites. All our experiments were approved by our IRB.

### 6.1 Evaluation Scenarios and Setup

**Application scenarios.** CROSSSENSE is evaluated on two WiFi sensing applications: gait identification and gesture recognition, and two wireless channel metrics: CSI and RSSI.

**Evaluation environments.** We test CROSSSENSE in three indoor environments of different sizes. Figure 7a shows the layouts and wireless setups of the three scenes. The first is a spacious hall entrance which mimics the reception area of a building. The second is a smaller, narrow corridor, and the third is a typical indoor environment with furniture include desks, chairs, bookshelves and appliances (see Figure 7b). The CDF diagram in Figure 8a suggests that the multipath effects of our evaluation environments can have a great impact on the wireless channel metrics. A larger number of gait samples (for the same person) have a further distance if

the measurements are collected from two different environments. Figure 8b projects 60 measurements (20 per site) for the user seen in Figure 2a onto a 2-dimensional space using PCA. As can be seen from the diagram, measurements from the same environment are more similar to each other; and as a result, we can group measurements into three clusters based on where the data come from.

**Wireless setup.** As an example, Figure 7b depicts the wireless setup of scene-3. To collect CSI, we used a mini PC with an Intel 5300 NIC as the receiver and a TP-Link WD7500 WiFi router as the transmitter (sender). The sender and the receiver are equipped with six and three antennas respectively to acquire and record wireless channel measurements from 30 channels for each antenna. We run an open source CSI measurement tool [30] on the mini PC to obtain CSI measurements in a 5GHz WiFi environment. To measure the CSI, the sender pings the receiver at a rate of 1,000 packets per second, which is a standard sample rate used in past work [31, 32]. We use a XiaoMI note2 smartphone as the receiver to collect RSSI measurements. We varied the distance between the sender and the receiver in different sites (ranging from 0.4 to 2 meters). Like WIAG [19], we collect gesture data from five positions. This is illustrated in Figure 7c where each data collection point gives the user’s absolute position in centimeters (using the sender as the origin) as well as the orientation with respect to the receiver. The work of WIAG shows that a change in position of up to 31 cm or a change in orientation of up to 45° does not have a significant impact on gesture recognition accuracy [19]. We use the gesture position location method of WIAG to determine the position and orientation of a user.

**Participants.** We recruited 100 volunteers (52 females) to participate in our experiments. As can be seen from Figure 9, our subjects have different heights, weights and somatotype, as indicated by various body mass index (BMI) values, and thus represent wide user groups.

**Gait data.** To collect CSI data for gait identification, each volunteer walked pass the two wireless devices at their natural speeds. We collected 20 wireless channel measurements per user per site, which yield in total 6,000 samples (100 users  $\times$  20 measurements per user  $\times$  3 sites).

**Gesture data.** We consider 40 gestures presented in prior work [5, 6, 19, 20, 33–36]. The list of gestures is given in Table 4. These include primitive gestures like hand moves down, combinations of primitive gestures such as up-down (the hand first moves up and then moves down), and more complex patterns like putting up the cigarette [36]. Our evaluation targets a significantly larger number of gestures than any of the prior work seen to date, which typically was evaluated using a handful of gestures. To collect gesture data, each volunteer repeatedly perform each gesture 20 times in each site. This results in 1,200,000 measurements (100 users  $\times$  40 gestures  $\times$  10 times per gesture  $\times$  2 channel metrics  $\times$  5 positions  $\times$  3 evaluation sites).

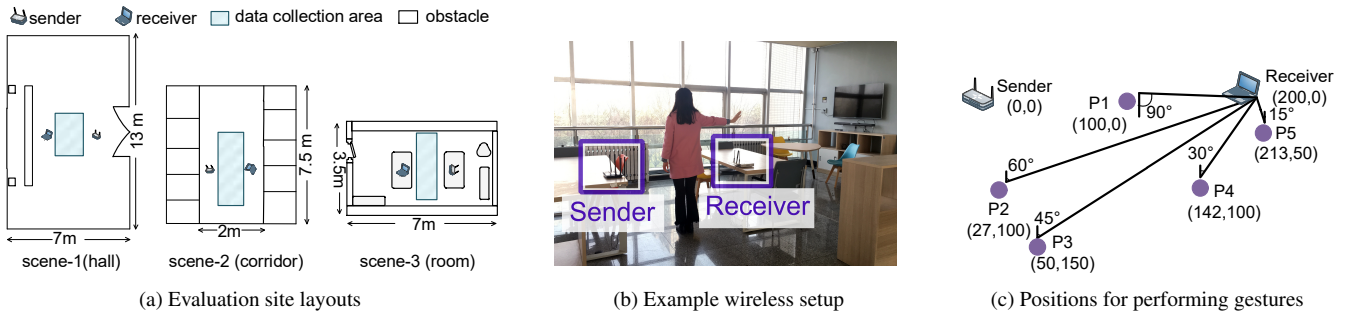


Figure 7: Experiment setup of evaluation site layouts (a), wireless setup (b) and gesture collection positions (c).

ID	Gesture	ID	Gesture	ID	Gesture	ID	Gesture	ID	Gesture	ID	Gesture	ID	Gesture
1	Flick	7	Pull	13	Hand moves down	19	Swipe right	25	Up-down	31	Cooking	37	Brushing teeth
2	Double flick	8	Throw	14	Hand moves up	20	Swipe left	26	Down-up	32	Studying	38	Putting up the cigarette
3	Punch	9	Circle	15	Kick	21	Infinity	27	Sit down	33	Taking a bath	39	Putting down the cigarette
4	Punch x 2	10	Dodge	16	Bowling	22	Left-right	28	Zoom In	34	Washing dishes	40	Inhaling/exhaling smoke
5	Level	11	Drag	17	Right	23	Right-left	29	Zoom Out	35	Turn on washing machines		
6	Push	12	Strike	18	Left	24	Head left	30	Eating	36	Play video games		

Table 4: Gestures used in our evaluation. These include primitive gestures and combinations of primitive gestures.




Figure 8: DTW distances of gait measurements for the same person increase if the samples are collected from different environments (a). The measurements of one user can be grouped into three clusters on a 2-d PCA space, each corresponds to where the data are gathered (b).




Figure 9: Heights, weights and BMIs of our participants. Our participants represent a wide range of user groups.

**Naming conventions.** We use *collection\_site - testing\_site* to denote a cross-site scenario. For example, *s1-s2* means that we use the WiFi measurements collected in scene-1 and to generate synthetic training examples to train a sensing model for scene-2, and then test the trained model in scene-2. Moreover, *cross-site transfer learning* is denoted as  $(s_i - s_j, s_k)$ , which is a pair of (i) an existing roaming model from sites  $i$  to  $j$ , and (ii) the target model for the new site  $k$ . For instance,  $(s_1 - s_2, s_3)$  means that we first learn a roaming model,  $m_1$ , to translate WiFi training samples from scene-1 to scene-2, and then transfer the learned model,  $m_1$ , to build a new roaming model,  $m_2$ , to translate measurements from scene-1 to scene-3.

## 6.2 Evaluation Methodologies

**Model Evaluation.** We use *cross-validation* to evaluate our approach. The expert models are trained on the *synthetic* WiFi training measurements of the targets in a cross-site sensing scenario. We report the *geometric mean* accuracy across evaluation scenarios using the *top-1* score, i.e., if the model output (the one with the highest probability) matches the expected answer. There is little changes in results when using *top-2* and *top-3* scores. Performance variances are shown using min-max bars.

**Comparisons.** We compare CROSSSENSE against four state-of-the-art CSI-based sensing methods. These include WiWho [2] and WiFiU [1] for gait identification, and WiAG [19] and WiG [6] for gesture recognition. In addition to CSI, we also apply CROSSSENSE to RSSI-based gesture recognition and compared it with TELEPATHICPHONE [20]. We compare our roaming method to FITLOC [21], a WiFi signal translation approach for RSSI-based localization.

**Implementation.** CROSSSENSE<sup>1</sup> is implemented using scikit-learn [37]. It is tested on a PC with a dual-core 2.4GHz Corei5 CPU and 16GB of RAM running Ubuntu 16.04.

## 7. EXPERIMENTAL RESULTS

Highlights of our evaluation are as follows:

- CROSSSENSE improves the classification accuracy of state-of-the-arts from around 20% to over 80% (Section 7.1).
- CROSSSENSE delivers consistently good performance regardless of the problem size (Section 7.2).
- Our roaming scheme allows existing WiFi sensing methods to work effectively across sites using a single set

<sup>1</sup>Code and data of CROSSSENSE are available at: [url redacted for double-blind review].




Figure 10: Overall performance per problem size across site pairs and cross-validation. CROSSSENSE significantly outperforms all competitive approaches by delivering the best and the most robust performance.




Figure 11: Accuracy distributions for CROSSSENSE per training\_data\_collection-testing-site pair. The thick black line shows where 50% of the data lie. CROSSSENSE delivers consistently good performance across evaluation scenarios.




Figure 12: How often (as percentages) an expert model in Table 3 correctly classifies an input across sites (s1, s2, s3).

of WiFi training measurements (Section 7.3). Transfer learning reduces the number of samples needed for learning the roaming model by 4x (Section 7.5).

- We thoroughly evaluate CROSSSENSE and provide detailed analysis on its working mechanisms.

## 7.1 Overall Performance

Figure 10 compares CROSSSENSE against alternative approaches across training-testing-site pairs (i.e., using the WiFi training measurements collected from one site to build sensing models for another site). In all experiments, we make sure that the training and the testing datasets are different. As expected, existing approaches give disappointing results – less than 20% for most of the cases. This is because they cannot translate and re-use training measurements to build sensing models for a new environment. Further, the precision of existing approaches drops quickly as the number of target users or gestures increases. By contrast, CROSSSENSE achieves the best accuracy for all evaluation scenarios, delivering an average accuracy of above 80% (and over 90% for CSI-based gesture recognition). CROSSSENSE also gives less than 5% reduction in accuracy as the problem size increases from the smallest setting to the largest one. Compared to the over 2x drop in accuracy given by other schemes, CROSSSENSE thus delivers the most reliable performance.




Figure 13: CSI-based sensing performance with different numbers of target activities. CROSSSENSE provides consistently good performance regardless of the problem size.

Figure 11 shows that CROSSSENSE delivers consistent performance across site pairs. While CROSSSENSE only use one set of WiFi training measurements, its performance is comparable to or even better than other approaches when their sensing models are trained using samples collected from the target environment. As such, CROSSSENSE allows scaling WiFi sensing to new environments and larger problem sizes with higher precision but less cost.

Figure 12 summarizes how often an expert gives a correct classification across evaluation sites when targeting 100 users and 40 gestures. No single model correctly classifies more than 60% of the test cases, with model capability changing from one environment to the other. The results emphasize that an universal model is unlikely to deliver good performance for large-scale and cross-site sensing problems.

Finally, we note that a *theoretically perfect* expert selector for our collection of expert models would give an average accuracy of 94.5% and 98.5% respectively for CSI-based gait identification and gesture recognition. The results suggest our chosen expert models are highly efficient. If we can improve the accuracy of our expert selector, we can then further improve the accuracy of CROSSSENSE.




Figure 14: Cross-site sensing with and without using our roaming mechanism. Roaming training measurements help existing methods to achieve comparable precision as if the training samples were collected from each deployment site.

## 7.2 Changing Problem Sizes

Figure 13 shows how the accuracy of CSI-based sensing changes of as the problem size increases. To isolate the issue, we use the target’s WiFi training measurements collected *from the deployment site* to build sensing models. We pick a given number of randomly chosen users and gestures as sensing targets, and repeat this process until all users and gestures are tested for at least once.

Using training measurements collected from the target environment does help existing approaches on a small problem size to deliver an accuracy of around 80%. However, performance of the competitive schemes degrades as the problem size increases. Their accuracy can drop by over 20% when we target more than 50 users or 10 gestures. When moving from the smallest setting to the largest one, their accuracy can drop by more than 50%. By contrast, CROSSSENSE delivers not only the best accuracy for all testing scenarios but also consistent performance across test cases. The drop in accuracy for CROSSSENSE is much smaller (less than 2% on average) when the problem size increases from the smallest to the largest one.

## 7.3 Roaming WiFi Training Measurements

We now examine if the synthetic training samples generated by our roaming model are useful. To isolate the issue, we test all models on a *small-scale* problem of six target users and gestures at a time. Our roaming model is trained on data collected from users and gestures that are not used in testing. The roaming model is trained on 2,000 randomly chosen samples (1,000 measurements per site), and we repeat this process to ensure that each user and gesture is tested at least once.

Figure 14 compares the accuracy of each method with and without using our roaming scheme. Directly using training samples of the target subjects collected from another site is a poor choice, giving only around 20% accuracy, which is only slightly better than making a random guess from the six possible options in this experiment. Our roaming model gives a clear boost to cross-site WiFi sensing, improving the sensing accuracy by over 3.5x with an accuracy of above 75% (and over 80% for most of the cases). This level of performance is in comparable to the accuracy when the model is built using training samples collected from the target en-




Figure 15: Comparison of the CROSSSENSE’s roaming model with the one used by FITLOC. The existing signal transferring method is ill-suited for cross-site WiFi sensing.

vironment. We want to highlight that our roaming scheme is particularly useful when the sensing targets (e.g., visitors) constantly change, because the roaming model only needs to be trained once but can apply to unseen sensing targets.

## 7.4 CrossSense versus FitLoc

Wireless signal transferring has been exploited in prior work in indoor localization to translate wireless fingerprints between different areas within an environment. The transferring method employed by FITLOC [21] is the current state-of-the-art. This experiment compares our roaming model against FITLOC’s transferring method. We use the translated (or synthetic) WiFi training examples collected from one site to build our expert models for another site. We then test the trained expert models on unseen WiFi measurements collected from the testing environment. Note that we use the same set of features to train the expert models, albeit that the translated samples given by FITLOC and CROSSSENSE, from the same raw WiFi training measurements, are different. The experiment targets 100 users and 40 gestures.

Results in Figure 15 show that the method used by FITLOC is ill-suited for gait and gesture classification, leading to a low classification accuracy – less than 10%. This is because the translated measurements produced by FITLOC have lost discriminative information and do not provide distinguishing features to separate sensing targets. By contrast, our roaming model can capture the essential characteristics of the wireless representations, and thus leads to better classification performance.

## 7.5 Transfer Learning

This experiment applies transferring learning (TL) to learn a roaming model for both a new site and a different sensing task. The accuracy is calculated by applying WIIWHO to six target users and WIAG to six gestures using cross-validation.

**TL for a new site.** Sub-figures a and b of Figure 16 show the results for using TL to learn a roaming model for an additional site. Figure 16a gives the mean accuracy across site permutations when the model is trained using 500 samples. Using TL, we achieve a comparable accuracy of over 80% and 90% respectively for gait identification and gesture recognition, but using only one quarter of the samples that would be required without TL. Figure 16b shows the average performance when the new roaming model is trained with




Figure 16: Cross-site (a & b) and cross-task (c & d) transfer learning (TL). The accuracy is calculated by applying W1WHO to six users and WIAG to six gestures. TL significantly reduces the number of samples required to learn a new roaming model.




Figure 17: Internal neural states for two different roaming models when processing the same input. The activations in each layer increasingly diverge the lower down the network.

different numbers of examples for site permutation ( $s_1 - s_2, s_3$ ). TL reduces the number of samples needed to reach a certain level of accuracy by approximately 4.5x.

**TL across tasks.** Figures 16 c and d show the resulting accuracy when the roaming model is built from a model that was previously designed for another sensing task. Figure 16c gives the accuracy for different cross-site configurations. Figure 16d shows the accuracy achieved when different numbers of training samples are used to build a roaming model to transfer sensing model training samples from scene-1 to scene-2. The use of TL across problem domains can match the performance given by direct training, but requires nearly 4x fewer training examples.

**Visualize network states.** In an attempt to explain TL, Figure 17 depicts the internal state of the two roaming models – each translates a CSI measurement from scene-1 to scene-2 and scene-3, respectively. We use a heatmap to visualize the intensity of each neuron activation of a network layer. From the top to bottom, we begin with a specific input feature vector. As information flows through the network, the layers become progressively more specialized to the target deployment environment. The activations from the top to the bottom layers become increasingly diverge. The mean variance of activations across the two models increases by three folds from 0.70 at the first hidden layer to 2.25 at the output layer. TL allows us to reuse the early layers of the network to speedup learning when targeting a new site or task. We also see that the output feature vectors are significantly different from the input. This reinforces our claim that a WiFi training measurement collected from one site must be translated before it can be used in another site.




Figure 18: SNR variances due to environmental changes (a) the resulting accuracy using synthetic samples produced by WIAG and CROSSSENSE (b).

## 7.6 Adapt to Environmental Changes

Changes in the environment could have an impact on fine-grained sensing like gesture recognition. Prior work requires to recollect training data to update the sensing model if this happens. WIAG tackles the problem by using training samples collected from one location to create virtual samples for other target locations. Here we compare our roaming model against the virtual sample generation scheme of WIAG.

In the experiment, we added and removed a few stationary objects such as chairs and tables near the wireless devices. We repeated this process 10 times for each environment. To adapt to the change, CROSSSENSE first collects 10 new samples from each target location from the changed environment (50 new samples for the five gesture locations in Figure 7c); it then uses the new samples to update the roaming model to generate *new* synthetic samples to refresh the sensing model. Our *initial* roaming model is trained using 2,000 samples, which is a *one-off* cost. For WIAG, we use 500 *new* training samples to update its gesture translation model whenever the environment has changed. To isolate the impact of problem sizes, we use WIAG’s KNN classifier to recognize six target gestures. Our roaming model is trained on gestures that are not used in testing.

The violin diagram in Figure 18a shows the distribution of the signal-to-noise ratio (SNR) variance after the environment changed. This is calculated from the participants’ positions. Figure 18b shows that our roaming model uses 10x fewer new samples to refresh the sensing model, but leads to a better accuracy. This is because our ANN-based model together with TL can better capture the environmental impact to the wireless signal; and consequently, it generates synthetic training samples with a higher quality. To match the level of accuracy given by our roaming model, WIAG would require to collect at least 1,200 new samples

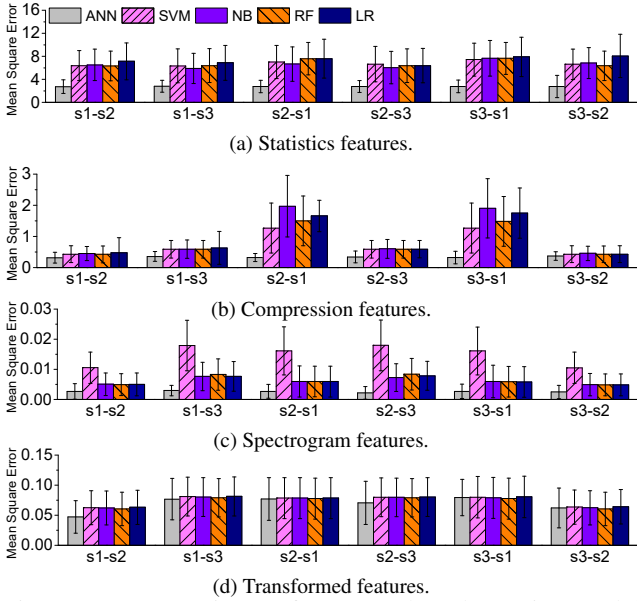


Figure 19: Comparisons of our ANN-based roaming model with other modeling techniques across feature sets for gait identification. Our model has the lowest mean square error.

every time in our evaluation environments. Therefore, while CROSSSENSE incurs slightly higher setup overhead (to build the first roaming model), it can adapt to the environment faster in a constantly changing environment.

## 7.7 Roaming Model Analysis

We also compare our chosen roaming model against alternative modeling techniques and neural network structures.

**Alternative modeling techniques.** Figure 19 compares our ANN-based roaming model against four alternative regression models: support vector machines (SVM), Naïve Bayes (NB), random forests (RF) and linear regression (LR), for each of the four feature sets listed in Table 1. We consider CSI-based gait identification and use the same process and data to train all models. The results are given in Figure 19. Since a measurement from one site can be mapped to multiple samples collected from the other site, we calculate the mean square error for each possible mapping and show the range of errors on the min-max bars. The results suggest that the ANN model gives the lowest error across feature sets. Using an ANN also allows us to easily employ TL to reduce the cost for learning the roaming model.

**Impact of neural layers and training samples.** Figure 20a shows the classification accuracy when the roaming model is constructed with different numbers of hidden layers. The accuracy is calculated by applying WIWHO and WIAG to six users and six gestures respectively using cross-validation, and the roaming model is trained using 2,000 examples without TL. We observe that using 7 and 10 hidden layers gives the best performance for gait identification and gesture recognition respectively. In this work, we choose to use a unify model structure with 7 hidden layers as it would require less

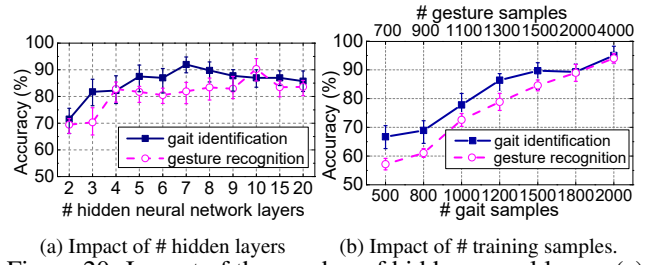


Figure 20: Impact of the number of hidden neural layers (a) and training samples of the roaming model (b).




Figure 21: CSI-based gait identification accuracy when using different numbers of expert models to target 100 users.




Figure 22: Using samples collected in the deployment environment to improve accuracy.

examples to train compared to the one with 10 hidden layers. Figure 20b shows a steady improvement in sensing model accuracy when using more examples to train the roaming model. We need more training data to build the gesture roaming model because gesture recognition needs to capture the subtle change of wireless signals at a finer-grained level. The figure also suggests that one can continuously improve the sensing model by providing more training samples over time (see also Section 7.9).

## 7.8 Impact of Expert Sizes

Figure 21 shows how the number of expert models affects the gait identification accuracy on each feature set. In addition to the six models listed in Table 2, we also consider four additional techniques: ANNs, logical regression, stochastic gradient descent (SGD) and linear discriminant analysis (LDA). The accuracy is calculated on a per feature set basis using 10-fold cross-validation on WiFi measurements. We target 100 users, test all possible model mixtures, and collect the WiFi measurements from the target environment.

As we increase the number of expert models, there is an increase in accuracy, confirming that a single model is insufficient. At the same time, we observe that the accuracy reaches a plateau when using six classifiers. We choose to use six classification techniques in the work because there is little gain in accuracy after that point to justify the increased cost in expert model training.

## 7.9 Continuous Learning

Figure 22 shows that the sensing performance can be continuously improved using samples collected from the deployment environment. The results are obtained by applying our mixture-of-experts method to 100 users and 40 gestures. The accuracy of the initial sensing model (learned using synthetic samples) can be improved by 4% for gait identification and 2% for gesture identification as we adding more samples obtained from the deployment site to the training dataset.

## 8. DISCUSSIONS

Naturally there is room for further work and improvements. We discuss a few points here.

**Training cost.** To further reduce the cost of learning the roaming model, one can employ a continual learning strategy, i.e., starting with a modest accurate model (that requires less samples to learn) and then using the data collected in target environment to continuously improve the model over time (see Section 7.9). Active learning [38] can also be employed to only focus on collecting samples that are likely to improve the accuracy of the model.

**Regression-oriented problems.** CROSSSENSE is evaluated on classification-based sensing problems. For regression-based problems, e.g., tracking and localization [39], we would need to replace our current expert models with regression techniques. This may also require employing new features, but our methods for training measurement roaming, feature engineering and expert model tuning remain applicable.

**Deep learning based feature representation.** Recently, deep neural networks [40] are shown to be powerful in extracting feature representations. However, they require many more training examples than that we use, which would incur prohibitive overhead because of user involvement. On the other hand, deep learning can be used within a continual- and transfer-learning framework to replace some of our expert models or obtain signal embeddings [41] as features.

**Other sensing platforms.** We believe CROSSSENSE can be applied to radio frequency (RF) based sensing [42, 43] and the extension is our future work. It would also be interesting to see if CROSSSENSE can be combined with specialized sensing hardware [44–46] to provide higher precision.

## 9. RELATED WORK

Recent years have witnessed a growing interest in exploiting wireless signals for tasks like gait identification [1–4], gesture recognition [6, 15, 19, 34, 42, 47–49, 49–59], localization [23, 39, 60–64], health and risk assessment [10, 35, 46, 65–71], activity detection [7, 8, 42, 72–80], and emotion recognition [81]. Indeed, wireless sensing has now moved from an a research niche [82] to a mainstream activity.

Wireless sensing maps a wireless measurement to an output. Prior work requires collect the targets’ training data from each deployment environment. WIAG [19] tackles the problem through generating virtual training samples for dif-

ferent locations in a single environment. While promising, WIAG only targets gesture recognition and does not address the issue of cross-site sensing. CROSSSENSE builds upon WIAG, utilizing its location estimation method to offer a generic framework for cross-site sensing.

Another drawback of prior work is that the proposed sensing method typically only work on a small problem size. After reviewing 70 research papers in wireless sensing from top-tier conferences between 2013 and 2018: SIGCOMM, MobiCom, MobiSys, Ubicomp, CHI, NSDI, and IPSN, we found that previous systems typically target a handful of subjects or activities – the largest one targets 50 users in a single environment [1]. While this may be a compromise of limited academic resources, we show that a single sensing model performs poorly as the problem becomes larger. We address this by combing and selecting multiple sensing models.

Transfer learning is shown to be useful in activities recognition [83], localization [84], crowdsourced mobile activity learning [85], and human activity recognition using sensor data [86]. CROSSSENSE builds on these past foundations to enable cross-site and cross-domain WiFi sensing.

Our mixture-of-experts approach is a form of ensemble learning [87]. Ensemble learning has been used in *video-based* gesture recognition [88] and gait identification [89], face recognition [90], information filtering [91], and other optimization tasks [92, 93]. However, no work so far has employed ensemble learning for WiFi sensing and this work is the first to do so. The goal of CROSSSENSE is not to advance ensemble learning; instead, it uses the technique together with signal transferring methods to design a general learning framework for wireless sensing.

## 10. CONCLUSION

This paper has presented CROSSSENSE, a novel framework to enable WiFi sensing to work effectively across deployment sites and to target a large number of sensing targets. To reduce the cost and human involvement for cross-site sensing, CROSSSENSE employs a machine learning model to generate, from a single set of WiFi training measurements, synthetic training samples for each deployment environment. To enable WiFi sensing to scale up to a larger problem size, CROSSSENSE integrates a mixture-of-experts based sensing approach. It determines at runtime, the best sensing model out of a collection of models (experts), as there is no one-size-fits-all universal best model. Such an approach provides a mechanism to gracefully add additional new sensing models to target a wider range of sensing scenarios and tasks. We demonstrate the effectiveness of CROSSSENSE by applying it to gait identification and gesture recognition and thoroughly evaluate it using over 1.2 million WiFi measurements. Compared to prior work, CROSSSENSE delivers the best and the most reliable performance across evaluation scenarios, and can work efficiently on problem sizes that are significantly bigger than that the current approaches can effectively handle.

## 11. REFERENCES

- [1] Wei Wang, Alex X. Liu, and Muhammad Shahzad. Gait recognition using wifi signals. In *UbiComp*, 2016.
- [2] Yunze Zeng, Parth H. Pathak, and Prasant Mohapatra. Wiwho: Wifi-based person identification in smart spaces. In *IPSN*, 2016.
- [3] Feng Hong, Xiang Wang, Yanni Yang, Yuan Zong, Yuliang Zhang, and Zhongwen Guo. Wfid: Passive device-free human identification using wifi signal. In *MOBIQUITOUS*, 2016.
- [4] Jin Zhang, Bo Wei, Wen Hu, and Salil S. Kanhere. Wifi-id: Human identification using wifi signal. In *DCOSS*, 2016.
- [5] Heba Abdelnasser, Moustafa Youssef, and Khaled A. Harras. Wigest: A ubiquitous wifi-based gesture recognition system. In *INFOCOM*, 2015.
- [6] Wenfeng He, Kaishun Wu, Yongpan Zou, and Zhong Ming. Wig: Wifi-based gesture recognition system. In *ICCCN*, 2015.
- [7] Stephan Sigg, Markus Scholz, Shuyu Shi, Yusheng Ji, and Michael Beigl. Rf-sensing of activities from non-cooperative subjects in device-free recognition systems using ambient and local signals. *IEEE Transactions on Mobile Computing*, 2014.
- [8] Wei Wang, Alex X Liu, Muhammad Shahzad, Kang Ling, and Sanglu Lu. Understanding and modeling of wifi signal based human activity recognition. In *MobiCom*, 2015.
- [9] Heba Abdelnasser, Khaled A. Harras, and Moustafa Youssef. Ubibreathe: A ubiquitous non-invasive wifi-based breathing estimator. In *MobiHoc*, 2015.
- [10] Hao Wang, Daqing Zhang, Junyi Ma, Yasha Wang, Yuxiang Wang, Dan Wu, Tao Gu, and Bing Xie. Human respiration detection with commodity wifi devices: do user location and body orientation matter? In *UbiComp*, 2016.
- [11] Chen-Yu Hsu, Yuchen Liu, Zachary Kabelac, Rumen Hristov, Dina Katabi, and Christine Liu. Extracting gait velocity and stride length from surrounding radio signals. In *CHI*, 2017.
- [12] Jin Kun, Fang Si, Peng Chunyi, Teng Zhiyang, Mao XuFei, Zhang Lan, and Li Xiangyang. Vivisnoop: Someone is snooping your typing without seeing it! In *CNS*, 2017.
- [13] Zhiping Jiang, Jizhong Zhao, Xiang-Yang Li, Jinsong Han, and Wei Xi. Rejecting the attack: Source authentication for wi-fi management frames using csi information. In *INFOCOM*, 2013.
- [14] Hongbo Liu, Yan Wang, Jian Liu, Jie Yang, and Yingying Chen. Practical user authentication leveraging channel state information (csi). In *ASIA CCS*, 2014.
- [15] Muhammad Shahzad, Alex X Liu, and Arjmand Samuel. Behavior based human authentication on touch screen devices using gestures and signatures. *IEEE Transactions on Mobile Computing*, 2017.
- [16] H. T. T. Truong, Xiang Gao, B. Shrestha, N. Saxena, N. Asokan, and P. Nurmi. Comparing and fusing different sensor modalities for relay attack resistance in zero-interaction authentication. In *PerCom*, 2014.
- [17] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *NIPS*, 2014.
- [18] Michael I. Jordan and Robert A. Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural Comput.*, 1994.
- [19] Aditya Virmani and Muhammad Shahzad. Position and orientation agnostic gesture recognition using wifi. In *MobiSys*, 2017.
- [20] Stephan Sigg, Ulf Blanke, and Gerhard Troster. The telepathic phone: Frictionless activity recognition from wifi-rssi. In *PerCom*, 2014.
- [21] L. Chang, X. Chen, Y. Wang, D. Fang, J. Wang, T. Xing, and Z. Tang. Fitloc: Fine-grained and low-cost device-free localization for multiple targets over various areas. *IEEE/ACM Transactions on Networking*, 2017.
- [22] Jue Wang and Dina Katabi. Dude, where's my card?: Rfid positioning that works with multipath and non-line of sight. In *SIGCOMM*, 2013.
- [23] Ju Wang, Hongbo Jiang, Jie Xiong, Kyle Jamieson, Xiaojiang Chen, Dingyi Fang, and Binbin Xie. Lifis: low human-effort, device-free localization with fine-grained subcarrier information. In *MobiCom*, 2016.
- [24] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., 2006.
- [25] Kamran Ali, Alex X. Liu, Wei Wang, and Muhammad Shahzad. Keystroke recognition using wifi signals. In *MobiCom*, 2015.
- [26] Mengyuan Li, Yan Meng, Junyi Liu, Haojin Zhu, Xiaohui Liang, Yao Liu, and Na Ruan. When csi meets public wifi: Inferring your mobile phone password via wifi signals. In *CCS*, 2016.
- [27] Guanhua Wang, Yongpan Zou, Zimu Zhou, Kaishun Wu, and Lionel Ni. We can hear you with wi-fi! In *MobiCom*, 2014.
- [28] Anthony Bagnall, Jason Lines, Jon Hills, and Aaron Bostrom. Time-series classification with cote: The collective of transformation-based ensembles. *IEEE Transactions on Knowledge and Data Engineering*, 2015.
- [29] Stan Salvador and Philip Chan. Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis*, 11(5), 2007.
- [30] Dan Halperin. Linux-80211n-csitool. <http://dhalperi.github.io/linux-80211n-csitool/>.
- [31] Tong Xin, Bin Guo, Zhu Wang, Mingyang Li, Zhiwen

- Yu, and Xingshe Zhou. Freesense: Indoor human identification with wi-fi signals. In *GLOBECOM*, 2016.
- [32] Kun Qian, Chenshu Wu, Zimu Zhou, Yue Zheng, Zheng Yang, and Yunhao Liu. Inferring motion direction using commodity wi-fi for interactive exergames. In *CHI*, 2017.
- [33] Qifan Pu, Sidhant Gupta, Shyamnath Gollakota, and Shwetak Patel. Whole-home gesture recognition using wireless signals. In *MobiCom*, 2013.
- [34] Bryce Kellogg, Vamsi Talla, and Shyamnath Gollakota. Bringing gesture recognition to all devices. In *NSDI*, 2014.
- [35] Yan Wang, Jian Liu, Yingying Chen, Marco Gruteser, Jie Yang, and Hongbo Liu. E-eyes: device-free location-oriented activity identification using fine-grained wifi signatures. In *MobiCom*, 2014.
- [36] Xiaolong Zheng, Jiliang Wang, Longfei Shangguan, Zimu Zhou, and Yunhao Liu. Smokey: Ubiquitous smoking detection with commercial wifi infrastructures. In *INFOCOM*, 2016.
- [37] Fabian Pedregosa et al. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 2011.
- [38] David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. *JAIR*, 1996.
- [39] Jiang Xiao, Kaishun Wu, Youwen Yi, Lu Wang, and Lionel M Ni. Pilot: Passive device-free indoor localization using channel state information. In *ICDCS*, 2013.
- [40] Shuochao Yao, Shaohan Hu, Yiran Zhao, Aston Zhang, and Tarek Abdelzaher. Deepsense: A unified deep learning framework for time-series mobile sensing data processing. In *WWW*, 2017.
- [41] Daniel Garcia-Romero, David Snyder, Gregory Sell, Daniel Povey, and Alan McCree. Speaker diarization using deep neural network embeddings. In *ICASSP*, 2017.
- [42] Fadel Adib, Zachary Kabelac, and Dina Katabi. Multi-person localization via rf body reflections. In *NSDI*, 2015.
- [43] Ju Wang, Jie Xiong, Xiaojiang Chen, Hongbo Jiang, Rajesh Krishna Balan, and Dingyi Fang. Tagscan: Simultaneous target imaging and material identification with commodity rfid devices. In *MobiCom*, 2017.
- [44] Bastien Lyonnnet, Cornel Ioana, and Moeness G Amin. Human gait classification using microdoppler time-frequency signal representations. In *Radarconf*, 2010.
- [45] Fadel Adib, Chen-Yu Hsu, Hongzi Mao, Dina Katabi, and Frédo Durand. Capturing the human figure through a wall. *TOG*, 34(6), 2015.
- [46] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C. Miller. Smart homes that monitor breathing and heart rate. In *CHI*, 2015.
- [47] Fadel Adib and Dina Katabi. See through walls with wifi! In *SIGCOMM*, 2013.
- [48] Sriram Nandha Premnath, Suman Jana, Jessica Croft, Prarthana Lakshmane Gowda, Mike Clark, Sneha Kumar Kasera, Neal Patwari, and Srikanth V. Krishnamurthy. Secret key extraction from wireless signal strength in real environments. In *TMC*, 2013.
- [49] Sheng Tan and Jie Yang. Wifinger: leveraging commodity wifi for fine-grained finger gesture recognition. In *MobiHoc*, 2016.
- [50] Ouyang Zhang and Kannan Srinivasan. Mudra: User-friendly fine-grained gesture recognition using wifi signals. In *CoNEXT*, 2016.
- [51] Jiguang Lv, Wu Yang, Dapeng Man, Xiaojiang Du, Miao Yu, and Mohsen Guizani. Wii: Device-free passive identity identification via wifi signals. In *GLOBECOM*, 2017.
- [52] Jiacheng Shang and Jie Wu. A robust sign language recognition system with sparsely labeled instances using wi-fi signals. In *MobiHoc*, 2017.
- [53] Jiacheng Shang and Jie Wu. A robust sign language recognition system with multiple wi-fi devices. In *MobiArch*, 2017.
- [54] Rajalakshmi Nandakumar, Bryce Kellogg, and Shyamnath Gollakota. Wi-fi gesture recognition on existing devices. *Eprint Arxiv*, 3(2), 2014.
- [55] Longfei Shangguan, Zimu Zhou, and Kyle Jamieson. Enabling gesture-based interactions with objects. In *MobiSys*, 2017.
- [56] Cong Shi, Jian Liu, Hongbo Liu, and Yingying Chen. Smart user authentication through actuation of daily activities leveraging wifi-enabled iot. In *MobiHoc*, 2017.
- [57] Pedro Melgarejo, Xinyu Zhang, Parameswaran Ramanathan, and David Chu. Leveraging directional antenna capabilities for fine-grained gesture recognition. In *UbiComp*, 2014.
- [58] J. Wang, L. Zhang, Q. Gao, M. Pan, and H. Wang. Device-free wireless sensing in complex scenarios using spatial structural information. *TWC*, 2018.
- [59] Lei Yang, Qiongzhen Lin, Xiangyang Li, Tianci Liu, and Yunhao Liu. See through walls with cots rfid system! In *Mobicom*, pages 487–499, 2015.
- [60] Zimu Zhou, Zheng Yang, Chenshu Wu, and Longfei Shangguan. Towards omnidirectional passive human detection. In *INFOCOM*, 2013.
- [61] Moustafa Youssef. Cosdeo 2016 keynote: A decade later ??? challenges: Device-free passive localization for wireless environments. In *Percom Workshops*, pages 1–2, 2016.
- [62] Heba Aly and Moustafa Youssef. An analysis of device-free and device-based wifi-localization systems. In *IJACI*, 2015.

- [63] Xiaohua Tian, Mei Wang, Wenxin Li, Binyao Jiang, Dong Xu, Xinbing Wang, and Jun Xu. Improve accuracy of fingerprinting localization with temporal correlation of the rss. *IEEE Transactions on Mobile Computing*, 2018.
- [64] Xiaohua Tian, Ruofei Shen, Duowen Liu, Yutian Wen, and Xinbing Wang. Performance analysis of rss fingerprinting based indoor localization. *IEEE Transactions on Mobile Computing*, 2017.
- [65] Neal Patwari, Joey Wilson, Sai Ananthanarayanan, Sneha K. Kasera, and Dwayne R. Westenskow. Monitoring breathing via signal strength in wireless networks. In *TMC*, 2011.
- [66] Chunmei Han, Kaishun Wu, Yuxi Wang, and Lionel M. Ni. Wifall: Device-free fall detection by wireless networks. In *INFOCOM*, 2014.
- [67] Fadel Adib, Zachary Kabelac, Hongzi Mao, Dina Katabi, and Robert C. Miller. Demo: real-time breath monitoring using wireless signals. *Mobicom*, 2014.
- [68] Xuefeng Liu, Jiannong Cao, Shaojie Tang, and Jiaqi Wen. Wi-sleep: Contactless sleep monitoring via wifi signals. In *RTSS*, 2015.
- [69] Sameera Palipana, David Rojas, Piyush Agrawal, and Dirk Pesch. Falldefi: Ubiquitous fall detection using commodity wi-fi devices. In *UbiComp*, 2018.
- [70] B. Mager, N. Patwari, and M. Bocca. Fall detection using rf sensor networks. In *PIMRC*, 2013.
- [71] Neal Patwari, Lara Brewer, Quinn Tate, Ossi Kaltiokallio, and Maurizio Bocca. Breathfinding: A wireless network that monitors and locates breathing in a home. *IEEE Journal of Selected Topics in Signal Processing*, 2013.
- [72] Wei Xi, Jizhong Zhao, Xiang-Yang Li, Kun Zhao, Shaojie Tang, Xue Liu, and Zhiping Jiang. Electronic frog eye: Counting crowd using wifi. In *INFOCOM*, 2014.
- [73] Bo Chen, Vivek Yenamandra, and Kannan Srinivasan. Tracking keystrokes using wireless signals. In *MobiSys*, 2015.
- [74] Kamran Ali, Alex X Liu, Wei Wang, and Muhammad Shahzad. Recognizing keystrokes using wifi devices. *J-SAC*, 35(5), 2017.
- [75] Dan Wu, Daqing Zhang, Chenren Xu, Hao Wang, and Xiang Li. Device-free wifi human sensing: From pattern-based to model-based approaches. *IEEE Communications Magazine*, 55(10), 2017.
- [76] Liang Wang, Tao Gu, Xianping Tao, and Jian Lu. Toward a wearable rfid system for real-time activity recognition using radio patterns. *IEEE Transactions on Mobile Computing*, 2016.
- [77] Stephan Sigg, Shuyu Shi, Felix Buesching, Yusheng Ji, and Lars Wolf. Leveraging rf-channel fluctuation for activity recognition: Active and passive systems, continuous and rssi-based signal features. In *MoMM*, 2013.
- [78] Dan Wu, Daqing Zhang, Chenren Xu, Yasha Wang, and Hao Wang. Widir: walking direction estimation using wireless signals. In *UbiComp*, 2016.
- [79] Yunze Zeng, Parth H. Pathak, Chao Xu, and Prasant Mohapatra. Your ap knows how you move: fine-grained device motion recognition through wifi. In *HotWireless*, 2014.
- [80] Simone Di Domenico, Mauro De Sanctis, Ernestina Cianca, and Giuseppe Bianchi. A trained-once crowd counting method using differential wifi channel state information. In *WPA*, 2016.
- [81] Mingmin Zhao, Fadel Adib, and Dina Katabi. Emotion recognition using wireless signals. In *Mobicom*, 2017.
- [82] Kristen Woyach, Daniele Puccinelli, and Martin Haenggi. Sensorless sensing in wireless networks: Implementation and measurements. In *WiOpt*, 2006.
- [83] Diane Cook, Kyle D Feuz, and Narayanan C Krishnan. Transfer learning for activity recognition: A survey. *KAIS*, 2013.
- [84] Sinno Jialin Pan, James T Kwok, and Qiang Yang. Transfer learning via dimensionality reduction. In *AAAI*, 2008.
- [85] Zhongtang Zhao, Yiqiang Chen, Junfa Liu, Zhiqi Shen, and Mingjie Liu. Cross-people mobile-phone based activity recognition. In *IJCAI*, 2011.
- [86] Jindong Wang, Yiqiang Chen, Lisha Hu, Xiaohui Peng, and Philip S Yu. Stratified transfer learning for cross-domain activity recognition. In *PerCom*, 2018.
- [87] Cha Zhang and Yunqian Ma. *Ensemble machine learning: methods and applications*. Springer, 2012.
- [88] Di Wu, Fan Zhu, and Ling Shao. One shot learning gesture recognition from rgb-d images. In *CVPRW*, 2012.
- [89] Y. Guan, C. T. Li, and F. Roli. On reducing the effect of covariate factors in gait recognition: A classifier ensemble method. *TPAMI*, 37(7), 2015.
- [90] Srinivas Gutta and Harry Wechsler. Face recognition using hybrid classifier systems. In *IJCNN*, volume 2, 1996.
- [91] Hichem Snoussi and Cédric Richard. Ensemble learning online filtering in wireless sensor networks. In *ICCS*, 2006.
- [92] Vicent Sanz Marco, Ben Taylor, Barry Porter, and Zheng Wang. Improving spark application throughput via memory aware task co-location: a mixture of experts approach. In *Middleware*, 2017.
- [93] Murali Krishna Emani and Michael O'Boyle. Celebrating diversity: A mixture of experts approach for runtime mapping in dynamic environments. In *PLDI*, 2015.