accepted to appear in Proceedings of the 21st International Conference on Pervasive Computing and Communications (PerCom 2023) March 13-17, 2023, Atlanta, USA.

Scheduled Spatial Sensing against Adversarial WiFi Sensing

Steven M. Hernandez and Eyuphan Bulut
Department of Computer Science, Virginia Commonwealth University
401 West Main St. Richmond, VA 23284, USA
{hernandezsm, ebulut}@vcu.edu

***Abstract*—WiFi sensing aims to utilize the changes in the Channel State Information (CSI) of WiFi signals due to the reflections from objects in the environment for sensing purposes. It uses machine learning classification models to predict physical actions being performed in a given environment (e.g., human activities such as walking, running). Thanks to the existing WiFi infrastructure in most indoor areas, this device-free technology can be used to provide low-cost motion detection and activity recognition opportunities for smart-homes. However, as the WiFi signals can be sniffed by adversaries, it can also be utilized by malicious actors to learn private information about the residents. To address this issue, motivated by the fact that the accuracy of WiFi sensing systems is highly reliant on the location of transmitter and receiver devices, we propose a simple yet effective solution based on the utilization of spatially distributed transmitter antennas (connected to a single source device) which communicate to a receiver device. The legitimate or allowed receiver is provided the schedule of transmitter antennas; thus, it can leverage this information to more accurately recognize activities performed within the environment. On the other hand, an eavesdropper who is unaware of the transmission schedule will encode the CSI frames from all transmitter antennas as if they were transmitted by a single source and thus will fail to recognize the activities properly. Through experiments, we show the effectiveness of this approach considering different number of transmitter antennas as well as against different levels of eavesdroppers.**

***Index Terms*—WiFi sensing, security and privacy, human activity detection**

I. Introduction

The use of WiFi has recently been extended beyond communication purposes through the concept of *WiFi sensing*; thanks to the recent advances in deep learning as well as the tools [1]–[3] that made the access to the Channel State Information (CSI) of WiFi signals easy. WiFi sensing aims to leverage fine-grained WiFi signal variations caused by physical reflections from objects in the environment which can be used to perform sensing tasks. This is primarily achieved through the extraction of CSI over subcarriers in orthogonal frequency-division multiplexing (OFDM) systems [4]. CSI data represents how wireless signals propagate from the transmitter to the receiver through multiple paths. This data consists of a matrix of complex values representing the amplitude attenuation and phase shift of multi-path WiFi channels.

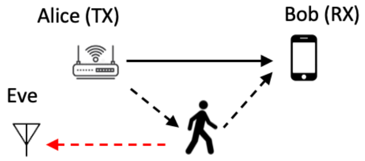

Fig. 1: A malicious eavesdropper (i.e., Eve) can obtain CSI data to perform adversarial WiFi sensing with a pretrained environment-independent ML model.

Several unique properties of WiFi sensing allow it to be a favorable solution compared to existing sensing systems. For example, WiFi sensing is device-free and thus it is not physically intrusive compared to wearable sensor-based systems, can be performed regardless of lighting conditions unlike camera-based sensing, and can go through walls contrary to both sensor-based and video-based solutions. As such, it has recently attracted a lot of attention by the research community and has been adopted in several applications (e.g., human activity/gesture recognition [5]–[7], health sensing [8]–[10]). These recent research efforts by academic community have also been supported by standardization efforts for next generation WiFi (e.g., 802.11 bf [11]) which consider sensing and communication together. Similarly, sensing as a side-service of WiFi is also becoming an industrial reality through new start-up companies (e.g., Origin Wireless [12]) and initial commercial products [13] for smart homes.

Despite this excitement for the use of WiFi for ubiquitous sensing in several different applications there is also an inevitable security and privacy risk of WiFi sensing for all of us. That is, an adversary (e.g., Eve in Fig. 1) sniffing the WiFi signals in the environment can use it for acquiring some private information about the users (e.g., if they are at home or not, or even which room they are in [25], their walking direction behind the wall [26]) and leverage this information for malicious purposes. Recognizing this risk, in some WiFi sensing studies [27], the machine learning model used in WiFi sensing is trained in a way such that only the

This work is supported in part by Commonwealth Cyber Initiative (CCI). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

TABLE I: Comparison of Existing Defense Methods

References	Method	Issues
[14], [15], [16]	Transmitter altered signals	By altering the signals, the data is no longer valid WiFi frames.
[17]	Signal strength variations	Reduces communication capacity of the network.
[18], [19]	External obfuscator node	Requires an additional device used solely for noisy transmissions.
[20]	Intelligent reflecting surface (IRS)	Requires hardware with low consumer usage.
[21], [22]	Omnidirectional jammer	Prevents all legitimate sensing and communication.
[23]	Directional jammer	Requires additional physically moving devices.
[24]	Alters signals to emulate activities	Requires specialized USRP equipment and non-standard WiFi frames.

allowed behaviors (e.g., falling of a senior) can be sensed properly while private activities (e.g., bathing) are prevented. However, such solutions provide only partial protection as it assumes that the trained model is the source of potential privacy leakage only. However, ambient WiFi signals can be sniffed by an eavesdropper and CSI data can be used for detection of activities using a pretrained environment-independent model [28]–[30] (i.e., a model trained using CSI data collected from different environment(s) but can perform accurate predictions in a totally new environment).

There are some recent efforts that aim to protect CSI signals from adversaries and thus invalidate their proper WiFi sensing capability. However, they are either more complicated as they use specialized hardware (e.g., using USRP [19], IRS [20]), and are not easy to implement in practice. Moreover, some of the solutions aim to totally avoid WiFi sensing even for legitimate devices thus are not desirable. Our goal is to allow legitimate WiFi sensing with allowed receiver (RX) devices but prevent illegitimate RX devices or eavesdroppers from performing adversarial WiFi sensing. To this end, we propose a WiFi sensing solution where multiple spatially distributed transmitter (TX) antennas are used to transmit WiFi packets to the RX.

The rest of the paper is organized as follows. In Section II, we provide a background on WiFi sensing and discuss the literature in particular in adversarial WiFi sensing and solutions to avoid it. In Section III, we present our system model together with the assumptions made and attacker and defense models. We then present our motivation for this work in Section IV and evaluate how our method can prevent eavesdroppers from performing WiFi sensing through experiments in Section V. Finally, we provide additional discussion about our method in Section VI and make our concluding remarks in Section VII.

II. Preliminaries

A. Background on WiFi Sensing

WiFi sensing uses the radio-frequency (RF) signals found throughout our homes and offices to detect and sense physical properties of the environment. These RF signals propagate over multiple unique physical paths (signal multipath) from the transmitter to the receiver. These multipaths cause slight variations in the signal due to the RF signals reflecting off of surfaces as well as propagating through objects such as walls, furniture, and people within the environment.

Channel state information is a signal metric captured in communication systems which use orthogonal frequency-division multiplexing (i.e., 802.11), to allow data-symbols to be encoded in multiple subcarrier frequency allowing for higher symbol throughput as well as resilience to signal fading and shadowing caused by multipath interference in the channel. CSI is modeled using the following relation:

$$y^{(i)} = H^{(i)} x^{(i)} + \eta^{(i)} \tag{1}$$

where i is the subcarrier index, x is the transmitted signal, y is the received signal, η is a noise vector, and H is a complex vector containing the channel state information denoting the transformation change required from the input x to the output y. The complex CSI vector contains 64 subcarriers where 52 are data-subcarriers and 12 are null-subcarriers. The CSI value for each subcarrier is defined as a complex number with a real component ($H_r^{(i)}$) and an imaginary component ($H_{im}^{(i)}$). We can transform this raw CSI into amplitude:

$$A^{(i)} = \sqrt{\left(H_{im}^{(i)}\right)^2 + \left(H_r^{(i)}\right)^2}, \tag{2}$$

and phase:

$$\phi^{(i)} = atan2\left(H_{im}^{(i)}, H_r^{(i)}\right). \tag{3}$$

B. Related Work

With the growing number of studies (e.g., [25], [31]) showing various levels of activity information and location leakage through adversarial WiFi sensing systems, developing counter mechanisms has become a necessity. Thus, recently several studies have looked at this problem and proposed different solutions. Table I provides a summary of existing defense mechanisms against eavesdropping with WiFi sensing. In [14], an obfuscation based solution is proposed which captures ambient wireless signals and relays them back into the environment with randomized modifications. However, the proposed solution uses full-duplex radio which requires specialized and costly hardware. A similar approach without using full-duplex is studied in [23], but it uses a motorized component to change the orientation of the antenna and introduces randomized delay. In [17], a solution is proposed which varies signal strengths of the transmitters and a game-theoretical model is studied between the attacker and defender considering the trade-off between privacy and utility in the system. This can however reduce the communication capacity between the devices.

Jammer-based solutions [21], [22], introduce randomized signal noise to prevent proper sensing. However, these solutions hamper the communication, thus they may not be practical in most of the real-life scenarios. Instead, in [19], a selective obfuscating solution is proposed to avoid extraction of location information from CSI. The solution superimposes a duplicated copy of the signal on each frame which does not affect the reception but does hinder the location-relevant information. However, this is mainly for protection of location and not applicable to activity detection use cases.

In [24], a modification to the radio training system is proposed to change the transmitted symbols over time, space and frequency as if they are affected due to human activities in the environment. While this approach prevents eavesdroppers from distinguishing real and fake human gestures, due to the requirement of specialized hardware (e.g., USRPs), it incurs a high cost and will not be practical. Note that our work also differs from the studies (e.g., [16]) that look at solutions against malicious radiometric fingerprinting of devices. These studies focus on the device-specific fingerprinting which could be used for impersonation attacks.

III. System Model

A. Assumptions

In our proposed system, we assume multiple TX antennas that are spatially distributed in a target sensing area as illustrated in Fig. 2. There is a single source device ($\mathcal{D}$) that is equipped with an antenna switch to automatically select the transmitting antenna at a per-packet level. This ensures that low layer attributes (e.g., MAC address and sequence number) will not directly reveal antenna changes to the eavesdropper. Usage of TX antennas are determined by a predefined schedule model ($\mathcal{S}$) which is shared between $\mathcal{D}$ and any legitimate RX devices. When evaluating our proposed system, only a single TX antenna communicates at any given time, however, $\mathcal{S}$ may be extended to allow multiple antennas to communicate simultaneously.

B. Experiment Setup

In our experiments, we consider a home environment where ESP32 microcontrollers are used as both TX and RX devices using the 802.11n protocol and 2.4GHz frequency band for CSI collection. Five TX devices that are placed 70 centimeters apart in a room of size 2.8 meters × 3.2 meters and 1 RX device are used as illustrated in Fig. 2. The RX is placed in an adjacent room to emulate an attacker which does not have direct access to the targeted sensing area. Each of the five TXs transmit WiFi frames to the RX at 20Hz concurrently resulting in CSI samples arriving at the RX at an overall rate of 100Hz. During our experimental data collection we allow all TXs to transmit concurrently so that we can emulate different transmission schedules, however in a real world system, only selected TXs will transmit at any given point of time. We collect CSI data to a Raspberry Pi single-board computer for processing.

For our dataset[1], we perform the following 5 activities:

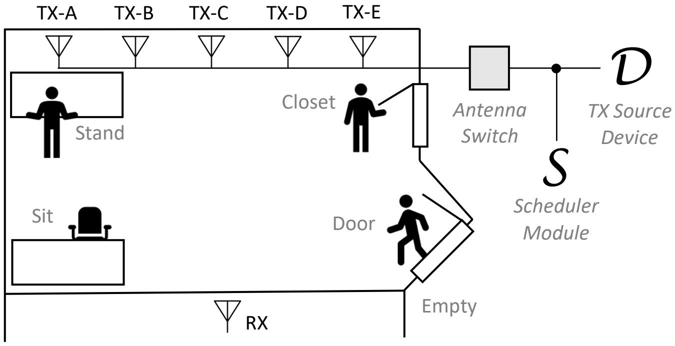

Fig. 2: Experimental setup with 5 TXs and 1 RX and 5 activities to be sensed. The scheduler ($\mathcal{S}$) decides which of the TXs that are wired connected to the same source device ($\mathcal{D}$) through an antenna switch needs to transmit.

- **Door:** Opening/closing main door
- **Sit:** Sitting and swiveling on a chair at a desk
- **Stand:** Standing at a desk and writing in a book
- **Closet:** Opening/closing closet door
- **Empty:** No movement within the room

These activities are performed in a round-robin fashion 6 distinct times. The first set of 3 repetitions are used for training our model while the final 3 are used for evaluation. Note that while the activities considered in this dataset are performed in diverse locations, WiFi sensing techniques are also applicable when multiple activities are performed at the same location [32], [33]. Additionally, while we use multiple ESP32s to act as TX antennas during our data collection phase, a similar system can also be achieved with a single WiFi device (e.g., one ESP32) and an antenna switch as illustrated in Fig. 2.

C. Tree-structured Parzen Estimator (TPE)

For our evaluations, we use the Tree-structured Parzen Estimator (TPE) [34] which is a hyperparameter optimization technique which selects some set of hyperparameters (θ) in an attempt to decrease some loss function $\mathcal{L}$ through the use of an expected improvement (EI) function:

$$EI(\theta) = \frac{p\big(\theta|\mathcal{L}'(\theta) > \mathcal{L}^*\big)}{p\big(\theta|\mathcal{L}'(\theta) \le \mathcal{L}^*\big)}, \tag{4}$$

where $\mathcal{L}^*$ is the average loss of the previously evaluated hyperparameter values for a given hyperparameter and $\mathcal{L}'(\theta)$ is formed based on previously observed hyperparameter values.

D. Attack Model

We consider a scenario where an attacker aims to sense the activities performed by an individual and subsequently localize the individual using the temporal CSI data obtained from the sniffed ambient WiFi signals. Fig. 1 illustrates our scenario where Alice transmits a signal to Bob. As the signals propagate through the environment, some of them reflect off of the human within the environment before continuing to propagate to Bob, thus allowing Bob to perform WiFi sensing. However, a malicious eavesdropper (Eve) can also receive the reflected

[1] Dataset and codebase for this project can be found in https://github.com/MoWiNG-Lab/AntiEave-WiFi-Sensing.

TABLE II: Scenarios considered during training and evaluation.

Scenario	Train on CSI from	Evaluate on CSI from	Section
Normal Sensing	Single TX	Single TX	Section IV
Naive Eve	Single TX	Multiple TX	Section V-A
Advanced Eve	Multiple TX	Multiple TX	Section V-B

signal which then allows Eve to perform WiFi sensing and thus Eve can achieve covert surveillance.

We assume that the attacker knows the set of localized activities and has a pretrained ML model for these activities. We also assume that this model is obtained through the solutions in the literature that offer environment-independent ML models [28], [29] or generic models that are obtained through a federated learning process [30]. However, we consider CSI data generated from both a single TX and multiple TX devices for training the attacker's model. Additionally, we assume that the attacker has a device and a tool that can extract CSI data from sniffed signals, which can be easily achieved through recent low-cost off-the-shelf solutions [35]. Attacker then uses this extracted CSI for predictions using the pretrained model. Similar to the training scenario, we look at the predictions when attacker uses CSI data received from (i) a single TX and (ii) multiple TXs. These scenarios and the sections looking at the evaluation of each scenario are given in Table II.

E. Defense Model

To prevent eavesdroppers from sensing physical activities using WiFi sensing without also hindering allowed RX devices from sensing physical activities, we leverage a multi-TX setup as illustrated in Fig. 2. We define a scheduler $\mathcal{S}$ which pseudo-randomly decides which TX should transmit at a given time instance, t, such that $\mathcal{S}(t) \in \{1, 2, \ldots, |TX|\}$. Allowed RX devices are given access to $\mathcal{S}$ which ensures that they are able to accurately identify which TX is transmitting at any given time while the disallowed eavesdropper is unable to make this distinction. To further obfuscate the physical activity and reduce the sensing capability of the eavesdropper, we define specific probability values for each TX device to determine how often the TX is selected from our scheduler module ($\mathcal{S}$).

F. Allowed RX Emulation

To evaluate the proposed system, we begin by describing the CSI data as seen by the allowed RX. This RX can recognize which TX is transmitting at any given time (from scheduler information). In our evaluations, we begin with a 3-dimensional CSI tensor $\mathbb{H} \in \mathbb{R}^{|\mathcal{T}|\times|TX|\times|s|}$ where $|\mathcal{T}|$ is the number of time steps in our dataset, $|TX|$ is the number of transmitters, and $|s|$ is the number of subcarriers per CSI frame. We apply a transformation to $\mathbb{H}$ based on our scheduler model $\mathcal{S}$ as so:

$$\big(\mathbb{H} \oplus \mathcal{S}\big)[t, i, :] = \mathbb{H}[t, i, :] * \textit{soft\_equals}(i, \mathcal{S}(t)), \quad (5)$$

where $\mathbb{H}[t, i, :]$ is a tensor slice of all subcarriers for station i collected at time t and

$$\textit{soft\_equals}(a, b) = 1 - tanh(|a - b|\beta), \quad (6)$$

which has an output approaching 1 when $a = b$, and 0 when $a \neq b$ and when β is some large value (i.e., $\beta = 1e4$). Through this, the allowed RX receives a tensor $\big(\mathbb{H} \oplus \mathcal{S}\big) \in \mathbb{R}^{|\mathcal{T}|\times|TX|\times|s|}$, however for each $i \in \{1, 2, \ldots, |TX|\}$ which is not selected at time t, the values for $\big(\mathbb{H} \oplus \mathcal{S}\big)[t, i, :] = \mathbf{0}$ because the model being trained would not be able to witness the CSI for the i-th TX.

G. Disallowed (Eavesdropper) RX Emulation

Now that we have reviewed the CSI data as seen by an allowed RX, next we review the CSI data as seen by a disallowed RX (i.e., an eavesdropper). The only difference in the allowed RX versus the disallowed RX is that the disallowed RX is not able to directly identify the difference between which TX is actively transmitting at any given time. As such, we define:

$$\begin{aligned}\big(\mathbb{H} \breve{\oplus} \mathcal{S}\big)[t, :] &= \sum_{i=1}^{|TX|} \Big(\big(\mathbb{H} \oplus \mathcal{S}\big)[t, i, :]\Big) \\ &= \big(\mathbb{H} \oplus \mathcal{S}\big)[t, \mathcal{S}(t), :],\end{aligned} \quad (7)$$

where $\big(\mathbb{H} \breve{\oplus} \mathcal{S}\big) \in \mathbb{R}^{|\mathcal{T}|\times|s|}$. It is important to note that while $\big(\mathbb{H} \oplus \mathcal{S}\big)$ and $\big(\mathbb{H} \breve{\oplus} \mathcal{S}\big)$ have different tensor shapes, they both contain the same amount of CSI amplitude information, meaning that they both have the same number of non-zero entries within the tensors. However, $\big(\mathbb{H} \oplus \mathcal{S}\big)$ encodes slightly more information due to the structure itself which is derived due to the knowledge of the transmission schedule shared between the TXs and RX.

IV. Motivation

In our initial efforts to motivate the multi-TX based proposed solution, we begin by presenting our experiment results for a human activity detection and localization scenario. We train a Dense Neural Network (DNN) machine learning classifier model $\mathcal{M}_{\text{TX-}m}$ per TX-m with one input dense layer, two hidden dense layers and one dense output layer. We apply L_2 kernel regularization across each dense layer and apply a dropout layer between each dense layer to prevent the model from overfitting. Finally, we use Stochastic Gradient Descent (SGD) to optimize the loss function

$$\mathcal{L}(x, y) = -\frac{1}{|x|} \sum_{i=1}^{|x|} \sum_{c=1}^{|C|} y_{i,c} \log \mathcal{M}_{\text{TX-}m}(x_{i,c}), \quad (8)$$

where $\mathcal{M}_{\text{TX-}m}(x_{i,c})$ is the model prediction for input CSI $x_{i,c}$ and $y_{i,c}$ is the true class for the i-th CSI measurement. We apply a preprocessing step to transform the raw CSI through Principal Component Analysis (PCA), which is shown to be one of the most effective preprocessing methods for increasing prediction accuracy in WiFi sensing scenarios [36].

We begin by showing how the accuracy of a WiFi sensing system is affected by the different physical positions of the five

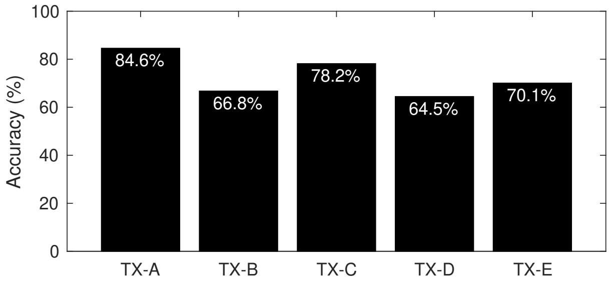

Fig. 3: Accuracy with ML models developed by CSI data coming from each TX.

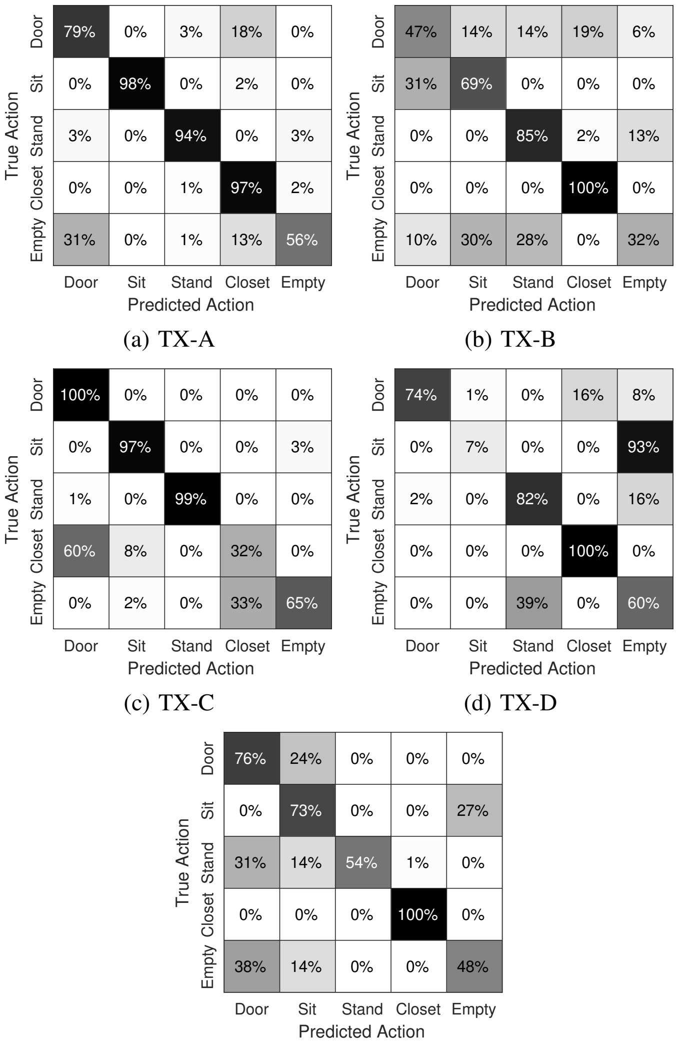

	Door	Sit	Stand	Closet	Empty
Door	79%	0%	3%	18%	0%
Sit	0%	98%	0%	2%	0%
Stand	3%	0%	94%	0%	3%
Closet	0%	0%	1%	97%	2%
Empty	31%	0%	1%	13%	56%

(a) TX-A

	Door	Sit	Stand	Closet	Empty
Door	47%	14%	14%	19%	6%
Sit	31%	69%	0%	0%	0%
Stand	0%	0%	85%	2%	13%
Closet	0%	0%	0%	100%	0%
Empty	10%	30%	28%	0%	32%

(b) TX-B

	Door	Sit	Stand	Closet	Empty
Door	100%	0%	0%	0%	0%
Sit	0%	97%	0%	0%	3%
Stand	1%	0%	99%	0%	0%
Closet	60%	8%	0%	32%	0%
Empty	0%	2%	0%	33%	65%

(c) TX-C

	Door	Sit	Stand	Closet	Empty
Door	74%	1%	0%	16%	8%
Sit	0%	7%	0%	0%	93%
Stand	2%	0%	82%	0%	16%
Closet	0%	0%	0%	100%	0%
Empty	0%	0%	39%	0%	60%

(d) TX-D

	Door	Sit	Stand	Closet	Empty
Door	76%	24%	0%	0%	0%
Sit	0%	73%	0%	0%	27%
Stand	31%	14%	54%	1%	0%
Closet	0%	0%	0%	100%	0%
Empty	38%	14%	0%	0%	48%

(e) TX-E

Fig. 4: Confusion matrix for each model in Fig. 3.

TXs relative to the RX as well as relative to the actions being performed. To this end, we train an ML model on training data captured by each TX and then evaluate the models on the testing data from the same TX. The accuracy for the models trained at each TX is shown in Fig. 3. We can see that TX-A achieves the highest accuracy at 84.59% and TX-D achieves the lowest accuracy at 64.47%. This demonstrates that the accuracy possible from each TX varies due to the unique physical positions of the TXs within the environment.

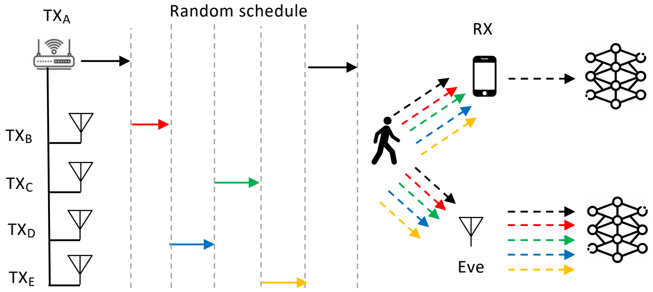

Fig. 5: Multiple TX antennas are used to transmit the WiFi signals at different times based on a predefined schedule known by a legitimate RX device, which then can filter the necessary CSI data for use in the prediction model, while eavesdropper uses all CSI and obtains inaccurate results.

The confusion matrices in Fig. 4 show which classes of actions are accurately predicted and which classes are commonly predicted incorrectly per TX. From this, we can see that each TX is better at distinguishing different sets of activities due to the spatially distributed nature of the TXs in the environment as well as the unique physical locations where each physical activity is performed. For example, TX-A achieves high classification accuracy on classes *sit, stand, closet*, TX-C achieves high classification accuracy on classes *door, sit, stand* and TX-B, TX-D, TX-E can each distinguish the *closet* action with high accuracy. This means that each of the TXs has unique strengths as well as unique weaknesses in our experiment scenario. In the next section, we will evaluate how we can leverage these differences due to TX positioning against a malicious eavesdropper.

V. Evaluation

We demonstrated how CSI captured from a single TX can be used to predict the localized physical activity of humans in an environment. However, achieving high accuracy in the previous scenario not only means that legitimate RXs can sense actions being performed, but it also means that malicious eavesdroppers can also covertly perform surveillance on the human target by sniffing these same signals.

To obfuscate the physical actions being performed in the environment, we allow the TXs to transmit one at a time on a random schedule every 50ms as illustrated in Fig. 5. This random schedule is emulated during our evaluations using the data collected and described in the previous section, but in a real-world deployment, we can assume that each TX adheres to the random schedule.

We study two scenarios: (i) a naive attacker that is not aware of the multiple TX antennas thus uses a sensing model trained with CSI data from one TX, and (ii) a more intelligent and advanced attacker which trains a sensing model using CSI

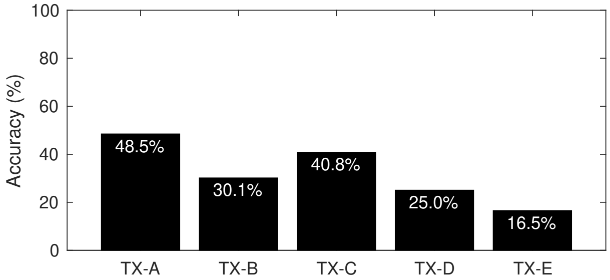

Fig. 6: Accuracy of eavesdropper's model trained on a single TX CSI data and used in obfuscated CSI data from all 5 TXs on a random schedule.

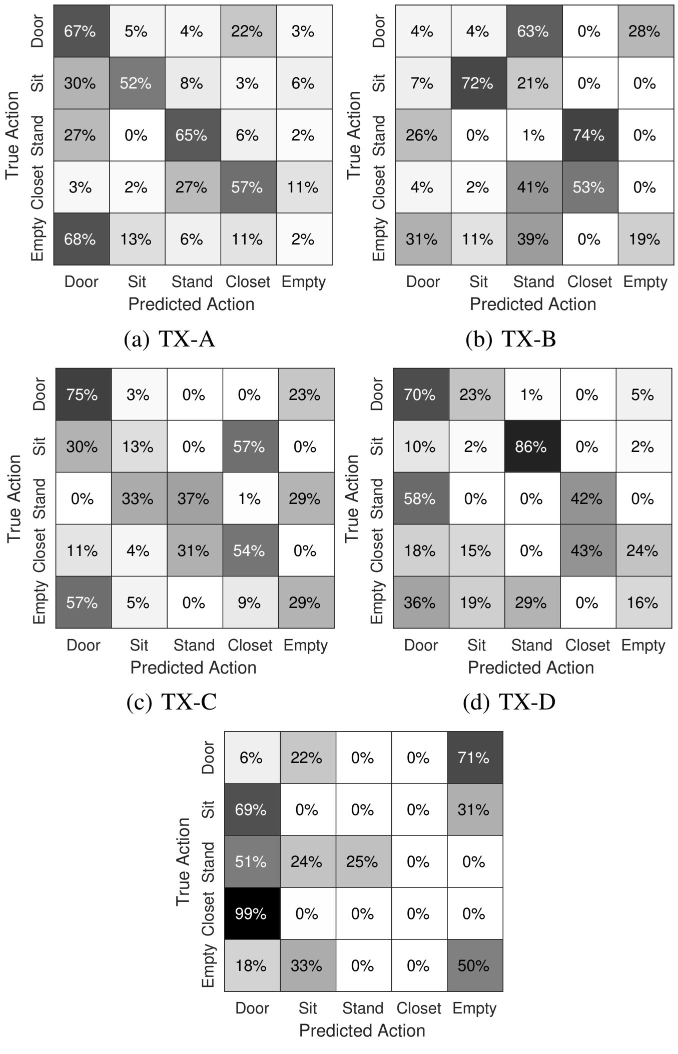

(a) TX-A

(b) TX-B

(c) TX-C

(d) TX-D

(e) TX-E

Fig. 7: Confusion matrix for each model in Fig. 6.

from multiple TXs which are generated based on the scheduler. In the latter, however, we still assume that the attacker does not know which packet comes from which TX device.

Note that it is not trivial for an advanced attacker to generate a pretrained environment independent model using CSI data from multiple TX locations (as done in the single source with a single antenna scenario [28]–[30]). This is because different spatial distribution of TX devices with respect to a receiver device can generate different results. However, to explore the extent to which an attacker can achieve sensing, we assume that the attacker is able to acquire CSI data from the same spatial distribution of TX devices as in the environment of interest along with the corresponding labels for each activity.

A. Naive Attacker

We begin by evaluating the naive attacker which considers that there is only one TX in the environment communicating with an RX device to generate the necessary signaling for WiFi sensing. As such, the model that is trained by this naive attacker will likely be confused by the CSI data coming from multiple TX antennas located in unique physical positions.

In Fig. 6, we can see the accuracy of the eavesdropper model when trained on CSI from a single TX and then applied to our obfuscation scenario where 5 TXs transmit on a random schedule. Since the eavesdropper does not know the random order of the transmitting devices, the eavesdropper must assume the use of all incoming CSI frames. We can see that the accuracy for each of the TX models has decreased significantly by as much as 53.5% for TX-E and a decrease of accuracy more than 35% for all other TXs (compared to the results in Fig. 3). Overall, this suggests that increasing the number of TXs even beyond five will allow for an ever lower accuracy for the naive attacker.

Fig. 7 shows the confusion matrix for each of the eavesdropper models in this same scenario. These figures show that our random scheduling method causes the eavesdropper model to randomly and incorrectly guess the current action being performed in the environment. Unlike Fig. 4 where each TX was able to achieve greater than 80% accuracy for more than one class, with our random scheduling method, the eavesdropper is unable to predict any of the individual classes with an accuracy greater than 80% for any of the TX models. The class that is most accurately predicted for the eavesdropper would be the *door* class using the model trained at TX-C. However, because the remaining predictions are so poor, it is not reasonable for an eavesdropper to believe that these predictions are correct. For example, while *door* is correctly predicted 75% of the time, *empty* is incorrectly predicted to be the *door* class 57% of the time and similarly, *sit* is incorrectly predicted to be the *door* class 30% of the time. Thus, because the accuracy is so poor for most of the action classes, any accurately predicted class cannot be distinguished from incorrectly predicted classes by the eavesdropper thus rendering the predictions useless.

Another fascinating observation is that the TX-E model incorrectly predicts *sit, stand, closet* classes most often to be the *door* class, yet the *door* action is rarely ever predicted correctly. This suggests that our method can be used to deceive the eavesdropper such that the eavesdropper will have a high propensity for predicting one given class while also concealing

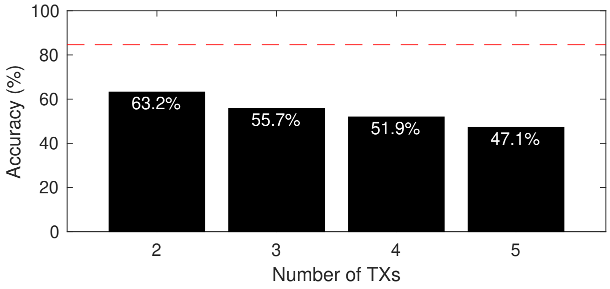

Fig. 8: Accuracy of an eavesdropper with different number of TXs communicating on a random schedule. Red dashed line shows the accuracy if random scheduling was not used.

the action when it actually occurs in the environment.

Now that we have evaluated the eavesdropper model when 5 TXs are used in our random schedule, we next look at the accuracy of our system when different numbers of TXs are used during evaluation. In Fig. 8, we evaluate the accuracy for a model trained at TX-A and then evaluate when multiple TXs are used in the random schedule including TX-A and some number of other TXs. The red dashed line shows the accuracy (84.59%) of the model when the random schedule was not applied. We can see that the accuracy decreases as more TXs are added to our random schedule, however, even when the number of TXs is 2 (i.e., TX-A and one other TX), the accuracy is 63.2% which is far lower than the 84.59% that could be achieved without the random schedule.

B. Advanced Attacker

In the previous scenario, we assumed that the eavesdropper naively trains a model using CSI collected from a single TX and then applies this model in a randomly scheduled multi-TX setting. However, a more advanced eavesdropper may train their model using CSI collected from all TXs as they actively communicate in the environment. As such, in this section, we begin by evaluating the advanced attacker in a random station schedule scenario. After this, our goal is to identify a transmission schedule which reduces the ability of the eavesdropper to perform sensing.

1) Random Schedule: In order to test the accuracy of models generated by Eve using the multi-TX data, we initially consider a random schedule of TXs in the system. Eve trains a model based on the data from all TXs using this random schedule, then the model is also used for predictions again using the CSI data from all TXs involved. In Table III, we review two forms of random TX scheduling, namely: periodic and non-periodic. In the periodic case, we create a pseudo-random schedule of size w which is repeated across the entire dataset. For the non-periodic case, we create a pseudo-random schedule across all timesteps within our dataset without actively ensuring periodicity. From this, we can observe that the periodic case allows Eve to achieve an accuracy of $+29.62\%$ greater than the non-periodic case. This demonstrates that any repeating patterns in the transmission schedule will actually improve the accuracy of Eve compared to a single-TX system (i.e., 86.20% is greater than all accuracy values in Fig. 3). This also demonstrates that an advanced attacker can achieve greater prediction accuracy (i.e., 56.58%) compared to a naive attacker (i.e., 48.5% with TX-A in Fig. 6) but still less accuracy than if only a single TX was used in the environment (i.e., 64.5% worst-case with TX-D in Fig. 3).

TABLE III: Eavesdropper accuracy with periodic and non-periodic random schedulers ($N = 50$ each).

Type	Avg. Accuracy (Std. Dev.)
Non-Periodic	56.58% (±9.90%)
Periodic	86.20% (±1.75%)

2) Probabilistic Schedule: In our previous experiments, we observed that each TX can achieve different levels of accuracy. For example, TX-A achieves the greatest accuracy in Fig. 3 at 84.6% while TX-D only achieves the lowest accuracy of 64.5%. We propose that we can leverage this knowledge to determine a schedule by setting pseudo-random probabilities uniquely per-station. Since different environments will have different TXs which achieve the best and worst sensing accuracy values, thus, we propose a learning approach to determine these pseudo-random per-station probabilities. Specifically, we use TPE to determine the optimal hyperparameter values for the probability of each station.

Towards this, when selecting the per-station probabilities, we define $0 \leq m \leq \frac{100}{|TX|}$, the minimum probability that all TXs are selected. In our experiments, since we have 5 TXs, the maximum value for m is 20%. The order in which hyperparameters are selected is important to ensure that the entire search space is explored by TPE. We find that if we use TPE to select a station probability in order for TX-A, TX-B, ..., TX-E, then TX-E will inevitably result in only very low probability values being explored due to it being the last selected probability value. As such, we instead select the probabilities of each TX in a random order for each TPE trial, thus allowing the full search space to be explored.

In Fig. 9, we illustrate the results of TPE when $N = 100$ TPE trials are performed and the minimum per-station probability $m = 5\%$. In this figure, we can see that as the probability for TX-A increases, the eavesdropper accuracy increases as well, thus TPE is able to recognize that low-values are more useful for our experiment environment. TX-C shows a similar upward trend while TX-B, TX-D, and TX-E show a negative trend as probability increases for each station. This is understandable considering that these TXs achieve the lowest accuracy values in Fig. 3 when evaluated on their own. While high probability values for TX-E appear to achieve the lowest accuracy for Eve, the achievable accuracy distribution range is wide, demonstrating that high probability values for TX-E do not always translate to the same low accuracy for Eve. This may be due to other TXs like TX-A or TX-C being selected along with TX-E in those trials.

Next we consider how applying different minimum probability values for m affects accuracy of Eve. By applying a minimum probability for all TXs, we can ensure that we

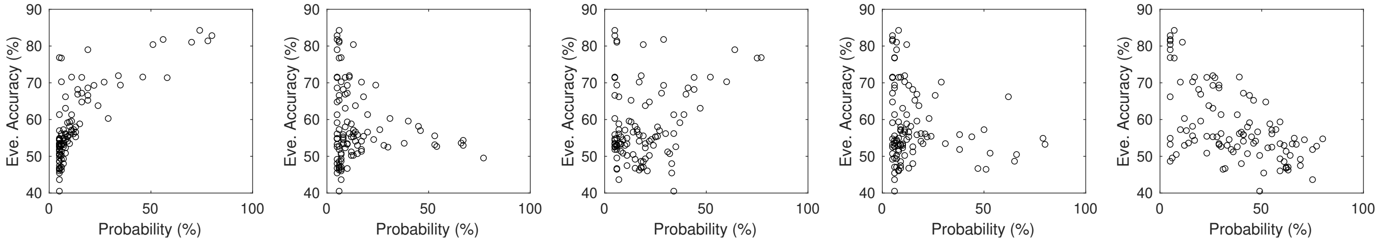

(a) TX-A (b) TX-B (c) TX-C (d) TX-D (e) TX-E

Fig. 9: Eavesdropper accuracy for different per-station probabilities when using TPE ($N = 100$, minimum per-station probability: 5%).

TABLE IV: Average accuracy ($N = 25$ each) for different per-station probabilities.

Station Probabilities					Eavesdropper Accuracy			Allowed-RX
TX-A	TX-B	TX-C	TX-D	TX-E	TPE Accuracy	Avg. Accuracy ($N = 25$)	Difference	Avg. Accuracy ($N = 25$)
4%	34%	23%	39%	0%	40.26%	40.60%	(+0.34%)	86.12%
6%	29%	25%	36%	4%	40.06%	40.94%	(+0.87%)	87.61%
6%	26%	21%	40%	7%	39.89%	40.85%	(+0.96%)	88.19%

leverage all of the available hardware which is deployed in the environment. Since our experimental design uses five TXs, when $m = 20$, each TX is selected equally by $\mathcal{S}$, however, with $m = 0$, it is possible that some TXs are unused for communication and sensing. In Fig. 10, we show the mean and standard deviation of $N = 100$ TPE trials when $m \in \{0, 5, 10, 15, 20\}$. The average accuracy of Eve decreases slightly as m decreases from 62.38% when $m = 20$ down to 59.02% when $m = 0$. However, the standard deviation increases from 2.43% when $m = 20$ up to 13.33% when $m = 0$. This is because with low values of m, there are more possibilities for better Eve accuracy as well as lower Eve accuracy values. This demonstrates that allowing some stations to be selected with a minimum probability $m < \frac{100\%}{|TX|}$ ensures that we can further decrease the achievable accuracy of even an advanced attacker.

The three best performing station probability values found through TPE are shown in Table IV along with the accuracy achieved during TPE optimization. We can see that TX-D is given the highest probability values. This is a reasonable choice considering that TX-D achieves the lowest accuracy (i.e., in Fig. 3) when evaluated alone. When TPE is used to optimize the per-station probability values, only a single training repetition is performed. As such, it is possible that the accuracy achieved is artificially low. To ensure that the accuracy values found through TPE are legitimate, we repeat the experiment with the same per-station probability hyper-parameter values over $N = 25$ repetitions and calculate the average and the difference from the TPE accuracy. We can see that for most of the best-selected per-station probabilities, the TPE accuracy and the average after 25 repetitions is within 1%. This demonstrates that the per-station probabilities selected by TPE are generalizable and not due to random chance. From this, we can observe that by using unique selection probabilities for each TX allows us to reduce the expected accuracy of the eavesdropper from 56.58% (i.e., non-periodic in Table III) down to approximately 40% accuracy.

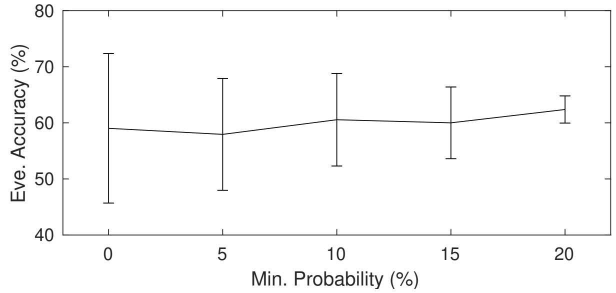

Fig. 10: Effect of minimum per-station probability on eavesdropper accuracy ($N = 100$ each).

Now that we have demonstrated that these per-station probability values can successfully decrease the accuracy of the advanced attacker, next we look at how these random station probabilities affect any legitimate WiFi sensing RX device. A legitimate RX knows the exact random schedule of the TXs while the eavesdropper does not and as such, our allowed TX can actually leverage the CSI coming from more than one TX when making predictions. In Table IV, we identified the station probabilities which achieve lowest accuracy for Eve through TPE. Using these same station probabilities, we trained an allowed RX model by replacing eq. (7) with eq. (5). By doing this, we encode some additional structure in the CSI tensor without including any additional CSI amplitude data. Through this, we find that all station probabilities achieve between 86.12% and 88.19% accuracy for the allowed RX. In fact, these accuracy values are similar and even greater than the best single TX in Fig. 3 (i.e., TX-A with 84.6% accuracy). As such, we can say that while applying the pseudo-random schedule reduces the effectiveness of disallowed eavesdropper devices in performing sensing, the same system does not affect and may even improve the performance of allowed sensing device. Note that these accuracy values for Eve are based on the assumption that the eavesdropper can obtain training CSI data from the TX devices in the environment, which could be challenging. Any missing information during such training process (e.g., wrong labels, missing CSI from some time frames or from some TXs temporarily) will potentially lower the accuracy even further.

VI. Discussion

A. *Effect on Communication*

WiFi sensing combines RF sensing into preexisting pervasive communications systems (i.e., WiFi). As such, it is important that a scheme which decreases the sensing ability of a system does not also decrease the communication ability of the system. For example, a signal jammer may be an efficient method for adding random signal noise into WiFi sensing measurements, however it also hinders the ability for legitimate WiFi devices to communicate while jamming is in progress. Our proposed method achieves the following regarding both sensing and communication:

1) Sensing is still possible and even improved for legitimate RXs through the use of multiple TXs.
2) Sensing is falsified and obscured for illegitimate eavesdropper RXs.
3) Communication packets are captured like normal for legitimate RXs.
4) Communication packets are captured like normal for eavesdropper RXs.

Notice, that our method does not worry about the content of the communication and even allows both legitimate and eavesdropper RXs to still capture the packet data. If the data in the packets must be hidden from eavesdroppers, then the data can easily be encrypted before transmission, however this is unrelated to the privacy concerns discussed in this paper.

B. *Generalizability to New Environments*

In this work, we demonstrated that we can confuse an eavesdropper device by transmitting over multiple TX antennas following a pseudo-random schedule rather than transmitting over just a single TX antenna. Due to the placement of these TX antennas and the physical locations of the activities being sensed, we showed that different TX antennas are better for recognizing different sets of activities. As such, the best per-station probabilities selected in this experimental environment will not necessarily be applicable to new environments, which may also have more or fewer TX antennas in the setup. As such, the proposed system is structured such that:

1) Per-station probabilities are learned through the TPE using real-world CSI data collected in the environment. Thus, the probability values can be selected automatically for each new environment.
2) The allowed RX (i.e., eq. (5)) and disallowed RX (i.e., eq. (7)) are designed as differentiable functions which allows for a machine learning model-based optimization of station probabilities. Thus, more complex station scheduling can be performed in new environments.

Furthermore, towards machine learning model-based station scheduling, it has been shown in [36] that even low level WiFi sensing devices such as the ESP32 used in this study can leverage machine learning models directly on-board. This means that such a system is possible even with low cost equipment, thus improving the scalability of such a system.

C. *Future Work*

In this work, we evaluated the effect of five TX antennas that are positioned at a constant distance apart, however different distributions of TX antennas will have different effects within each unique environment. As such, more work can be done in understanding how different TX antenna positions and different number of antennas affect the proposed system. Furthermore, because each environment has unique activities to be obfuscated, it may be possible to automatically determine optimal placement of these TX antennas through metrics such as the sensing-signal-to-noise-ratio (SSNR) [37] or through wireless sensing signal simulators [38].

In our experiments, the eavesdropper uses CSI exclusively to recognize and localize the activities performed. However, metrics such as the received signal strength indicator (RSSI), angle of arrival (AoA), or other signal metrics may also be available for the eavesdropper and may reveal the physical locations of the TXs [39]. While this work focuses directly on obfuscating physical activities, additional work into obfuscating antenna locations (e.g., [40]) will directly benefit our proposed system.

VII. Conclusion

In this work, we proposed a defense mechanism against adversarial WiFi sensing through the use of multiple spatially distributed TX antennas connected to the same source device. These antennas are utilized to transmit data based on a pseudo-random schedule which is known to legitimate RX devices but hidden from eavesdroppers. Legitimate RX devices can filter the received data per TX device based on the schedule used and use a specific ML model for predictions, while the eavesdropper uses the CSI data from all TX devices and uses them as input into its prediction model. Through various experiments, we showed that accuracy of the eavesdropper model is much lower than the accuracy of the legitimate RX model thanks to the obfuscation generated through the spatially distributed TX antennas. The accuracy of the eavesdropper also reduces as the number of TX devices increases. Additionally, we demonstrated that setting a per-station probability for our pseudo-random scheduler allows for a further decrease in the accuracy of an eavesdropper. We proposed a Tree-structured Parzen Estimator (TPE) approach to identify optimal per-station probability values which ensure that the system can be automatically adaptable in new environments. Finally, we also showed that accuracy for legitimate WiFi sensing RX devices can even be improved through the use of CSI from multiple TXs. As such, the proposed system is able to allow legitimate sensing to occur while reducing the feasibility of illegitimate eavesdropper-based sensing from occurring.

References

[1] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, “Tool Release: Gathering 802.11n Traces with Channel State Information,” *ACM SIGCOMM CCR*, vol. 41, no. 1, p. 53, Jan. 2011.

[2] Nexmon: The C-based Firmware Patching Framework, 2019. [Online]. Available: https://nexmon.org

[3] Atheros CSI Tool, 2019. [Online]. Available: https://wands.sg/research/wifi/AtherosCSI/

[4] Y. Ma, G. Zhou, and S. Wang, “WiFi Sensing with Channel State Information: A Survey,” *ACM Computing Surveys (CSUR)*, vol. 52, no. 3, p. 46, 2019.

[5] W. Wang, A. X. Liu, M. Shahzad, K. Ling, and S. Lu, “Understanding and Modeling of WiFi Signal Based Human Activity Recognition,” in *Proceedings of the 21st ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2015, pp. 65–76.

[6] J. Zhang, Z. Tang, M. Li, D. Fang, P. Nurmi, and Z. Wang, “CrossSense: towards cross-site and large-scale WiFi sensing,” in *Proceedings of the 24th ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2018, pp. 305–320.

[7] Y. Zheng, Y. Zhang, K. Qian, G. Zhang, Y. Liu, C. Wu, and Z. Yang, “Zero-Effort Cross-Domain Gesture Recognition with Wi-Fi,” in *Proceedings of the 17th ACM Annual International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2019, pp. 313–325.

[8] X. Wang, C. Yang, and S. Mao, “PhaseBeat: Exploiting CSI phase data for vital sign monitoring with commodity WiFi devices,” in *Proceedings of the 37th IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2017, pp. 1230–1239.

[9] Z. Chen, L. Zhang, C. Jiang, Z. Cao, and W. Cui, “WiFi CSI based passive human activity recognition using attention based BLSTM,” *IEEE Transactions on Mobile Computing*, vol. 18, no. 11, pp. 2714–2724, 2018.

[10] M. Huang, J. Liu, Y. Gu, Y. Zhang, F. Ren, X. Wang, and J. Li, “Your WiFi Knows You Fall: A Channel Data-Driven Device-Free Fall Sensing System,” in *Proceedings of IEEE International Conference on Communications (ICC), Shanghai, China, May 20-24*, 2019, pp. 1–6.

[11] F. Restuccia, “IEEE 802.11 bf: Toward Ubiquitous Wi-Fi Sensing,” *arXiv preprint arXiv:2103.14918*, 2021.

[12] Origin Wireless, “Wireless AI for a Smarter World,” 2020. [Online]. Available: https://www.originwirelessai.com/

[13] Linksys-aware, “The First Mesh WiFi Motion Sensing Technology,” 2020. [Online]. Available: https://www.linksys.com/us/linksys-aware/

[14] Y. Qiao, O. Zhang, W. Zhou, K. Srinivasan, and A. Arora, “PhyCloak: Obfuscating Sensing from Communication Signals,” in *Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, 2016, pp. 685–699.

[15] M. Cominelli, F. Kosterhon, F. Gringoli, R. L. Cigno, and A. Asadi, “An Experimental Study of CSI Management to Preserve Location Privacy,” in *Proceedings of the 14th ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, ser. WiNTECH’20, New York, NY, USA, Sep. 2020, pp. 64–71.

[16] L. F. Abanto-Leon, A. Bäuml, G. H. Sim, M. Hollick, and A. Asadi, “Stay Connected, Leave no Trace: Enhancing Security and Privacy in WiFi via Obfuscating Radiometric Fingerprints,” vol. 4, no. 3, 2020, pp. 1–31.

[17] P. M. Wijewardena, A. Bhaskara, S. K. Kasera, S. A. Mahmud, and N. Patwari, “A Plug-n-Play Game Theoretic Framework For Defending Against Radio Window Atacks,” in *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2020, pp. 284–294.

[18] M. Cominelli, F. Gringoli, and R. L. Cigno, “Non Intrusive Wi-Fi CSI Obfuscation Against Active Localization Attacks,” in *Proceedings of the 16th IEEE Annual Conference on Wireless On-demand Network Systems and Services Conference (WONS)*, Klosters, Switzerland, Mar. 2021, pp. 1–8.

[19] M. Cominelli, F. Gringoli, and R. L. Cigno, “AntiSense: Standard-compliant CSI obfuscation against unauthorized Wi-Fi sensing,” *Elsevier Computer Communications*, vol. 185, pp. 92–103, 2022.

[20] P. Staat, S. Mulzer, S. Roth, V. Moonsamy, A. Sezgin, and C. Paar, “IRShield: A Countermeasure Against Adversarial Physical-Layer Wireless Sensing,” *arXiv preprint arXiv:2112.01967*, 2021.

[21] J. Zhang, M. Li, Z. Tang, X. Gong, W. Wang, D. Fang, and Z. Wang, “Defeat Your Enemy Hiding behind Public WiFi: WiGuard Can Protect Your Sensitive Information from CSI-Based Attack,” *Applied Sciences*, vol. 8, no. 4, p. 515, 2018.

[22] Q. Wang, “A Novel Anti-Eavesdropping Scheme in Wireless Networks: Fri-UJ,” in *Proceedings of the International Conference on Embedded Wireless Systems and Networks*. Junction Publishing, 2019, pp. 316–317.

[23] Y. Yao, Y. Li, X. Liu, Z. Chi, W. Wang, T. Xie, and T. Zhu, “Aegis: An Interference-Negligible RF Sensing Shield,” in *Proceedings of IEEE International Conference on Computer Communications (INFOCOM)*, 2018, pp. 1718–1726.

[24] S. A. Mahmud, N. Patwari, and S. K. Kasera, “How to Get Away with MoRTr: MIMO Beam Altering for Radio Window Privacy,” in *Proceedings of the 18th IEEE International Conference on Mobile Ad Hoc and Smart Systems (MASS)*, 2021, pp. 315–323.

[25] Y. Zhu, Z. Xiao, Y. Chen, Z. Li, M. Liu, B. Y. Zhao, and H. Zheng, “Adversarial WiFi Sensing,” *arXiv preprint arXiv:1810.10109*, 2018.

[26] S. M. Hernandez and E. Bulut, “Adversarial Occupancy Monitoring using One-Sided Through-Wall Wi-Fi Sensing,” in *Proceedings of IEEE International Conference on Communications (ICC)*, 2021, pp. 1–6.

[27] S. Zhou, W. Zhang, D. Peng, Y. Liu, X. Liao, and H. Jiang, “Adversarial WiFi Sensing for Privacy Preservation of Human Behaviors,” *IEEE Communications Letters*, vol. 24, no. 2, pp. 259–263, 2019.

[28] W. Jiang, C. Miao, F. Ma, S. Yao, Y. Wang, Y. Yuan, H. Xue, C. Song, X. Ma, D. Koutsonikolas, W. Xu, and L. Su, “Towards Environment Independent Device Free Human Activity Recognition,” in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2018, pp. 289–304.

[29] Y. Ma, S. Arshad, S. Muniraju, E. Torkildson, E. Rantala, K. Doppler, and G. Zhou, “Location and Person Independent Activity Recognition with WiFi, Deep Neural Networks and Reinforcement Learning,” *ACM Transactions on Internet Things*, vol. 2, no. 1, pp. 3:1–3:25, 2020.

[30] S. M. Hernandez and E. Bulut, “WiFederated: Scalable WiFi Sensing Using Edge-Based Federated Learning,” *IEEE Internet Things Journal*, vol. 9, no. 14, pp. 12 628–12 640, 2022.

[31] Y. Zhu, Z. Xiao, Y. Chen, Z. Li, M. Liu, B. Y. Zhao, and H. Zheng, “Et Tu Alexa? When Commodity WiFi Devices Turn into Adversarial Motion Sensors,” in *Proceedings of the 27th Annual Network and Distributed System Security Symposium, (NDSS), San Diego, California, USA, February 23-26*, 2020.

[32] S. M. Hernandez, M. Touhiduzzaman, P. E. Pidcoe, and E. Bulut, “Wi-PT: Wireless Sensing based Low-cost Physical Rehabilitation Tracking,” in *Proceedings of IEEE International Conference on E-health Networking, Application & Services (HealthCom)*, 2022, pp. 113–118.

[33] B. Tan, Q. Chen, K. Chetty, K. Woodbridge, W. Li, and R. Piechocki, “Exploiting WiFi Channel State Information for Residential Healthcare Informatics,” *IEEE Communications Magazine*, vol. 56, no. 5, pp. 130–137, May 2018.

[34] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, “Algorithms for Hyper-Parameter Optimization,” *Advances in Neural Information Processing Systems*, vol. 24, 2011.

[35] S. M. Hernandez and E. Bulut, “Lightweight and Standalone IoT based WiFi Sensing for Active Repositioning and Mobility,” in *Proceedings of the 21st IEEE International Symposium on “A World of Wireless, Mobile and Multimedia Networks” (WoWMoM)*, Cork, Ireland, Jun. 2020, pp. 277–286.

[36] S. M. Hernandez and E. Bulut, “WiFi Sensing on the Edge: Signal Processing Techniques and Challenges for Real-World Systems,” *IEEE Communications Surveys & Tutorials*, 2022.

[37] X. Wang, K. Niu, J. Xiong, B. Qian, Z. Yao, T. Lou, and D. Zhang, “Placement Matters: Understanding the Effects of Device Placement for WiFi Sensing,” *in Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 6, no. 1, pp. 1–25, 2022.

[38] S. Vishwakarma, W. Li, C. Tang, K. Woodbridge, R. Adve, and K. Chetty, “SimHumalator: An Open-Source End-to-End Radar Simulator for Human Activity Recognition,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 37, no. 3, pp. 6–22, 2021.

[39] D. Konings, N. Faulkner, F. Alam, F. Noble, and E. Lai, “Do RSSI values reliably map to RSS in a localization system?” in *Proceedings of 2nd IEEE Workshop on Recent Trends in Telecommunications Research (RTTR)*, 2017, pp. 1–5.

[40] R. Ayyalasomayajula, A. Arun, W. Sun, and D. Bharadia, “Users are Closer than they Appear: Protecting User Location from WiFi APs,” *arXiv preprint arXiv:2211.10014*, 2022.