Adversarial Localization against Wireless Cameras

Zhijing Li*, Zhujun Xiao+, Yanzi Zhu*, Irene Pattarachanyakul*,
Ben Y. Zhao+ and Haitao Zheng+
{zhijing,yanzi}@cs.ucsb.edu,ireneeypatt@hotmail.com,{zhujunxiao,ravenben,htzheng}@cs.uchicago.edu
*UC Santa Barbara, + University of Chicago

ABSTRACT

This paper identifies and empirically evaluates the effectiveness of adversarial localization attacks against wireless IoT devices, *e.g.,* wireless security cameras in the home. We use experiments in home and office settings to show that attackers can accurately pinpoint the location of WiFi cameras, using a small amount of stealthy, passive, exterior measurements coupled with unsupervised learning techniques. We also show that current defenses have minimal impact against these attacks, and are also easily circumvented via countermeasures. Thus significant work is needed to develop robust defenses against these attacks.

ACM Reference Format:
Zhijing Li*, Zhujun Xiao+, Yanzi Zhu*, Irene Pattarachanyakul*, Ben Y. Zhao+ and Haitao Zheng+. 2018. Adversarial Localization against Wireless Cameras. In *Proceedings of 19th International Workshop on Mobile Computing Systems & Applications (HotMobile '18).* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3177102.3177106

1 INTRODUCTION

Digital homes are becoming increasingly commonplace. Some of the most popular products today are for home security. Companies like Ring and Google Nest offer cheap, high quality wireless video devices for surveillance and intruder detection. They offer users a sense of security, especially when homeowners are away from home.

As these devices gain in popularity, their widespread deployment means any security vulnerability in their design will have large-scale impact across a large user population. For wireless home cameras, one key vulnerability comes in the form of signal leakage and remote localization by external attackers[1]. From outside the home, a burglar can use wireless measurements to detect the likely location of wireless security cameras, and can then plan their intrusion to avoid detection. Studies have shown that nearly 60% of home burglars will consider the presence of cameras or other video equipment when selecting targets [1]. Others show that burglars (even shoplifters) are adept at identifying and leveraging cameras' blind spots to evade detection [2, 3]. This type of attack is sometimes called *adversarial localization*[2], where an attacker applies localization techniques to locate third party wireless transmitters.

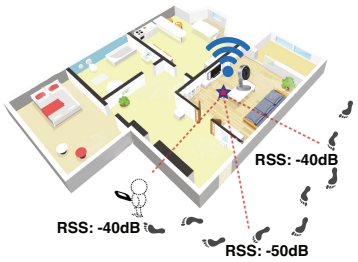

Figure 1: An example scenario of adversarial localization on behind-the-wall wireless camera.

The goal of this paper is to understand the practical effectiveness of adversarial localization in a realistic setting. We are particularly interested in *stealthy* attacks against wireless cameras, given the direct implications on home security systems. We consider the attack model shown in Figure 1. An adversary, a human or a robot, travels outside the home, passively collects received signal strength (RSS) data of a WiFi camera, then uses the RSS data to estimate the camera location. The key features of this attack model are stealthiness and simplicity. Being passive, the adversary does not need to communicate with the target; only requiring RSS measurements, the attack can be easily carried out by walking (normally) pass the home with a compact COTS WiFi receiver and applying existing RSS-based localization algorithms [5, 9]. Note that while more advanced localization algorithms (*e.g.* time, fingerprinting, AoA based) may lead to higher accuracy, they either require active communications with the target [15, 20, 28, 32][3] or bulky/specialized hardware with multiple antennas [31, 32]. These requirements largely reduce the attack stealthiness.

We perform numerous experiments using COTS video cameras in different home and office settings. In each case, the attacker and the camera are fully separated by walls. We show that despite significant noise in WiFi signals through the walls, localization attacks can be effective but the accuracy varies significantly across measurement instances. We propose unsupervised learning techniques an attacker can use to isolate high-quality measurement instances from noisy instances. For example, the attacker can determine on the fly whether the current measurement is sufficient to produce an accurate result, and use this information to terminate or continue the current effort, or plan a different route. Doing so dramatically improves efficacy of these attacks.

[1]Other extensions of this attack include jamming or even modifying the camera's wireless signal. In this paper, we focus on the attack to localize the camera rather than changing its signal.

[2]Our version of the adversarial localization differs from those proposed for wireless sensor networks [10, 13, 22, 30], which involve multiple sensors and focus on sensor routing mechanism design.

[3]Today's WiFi cards cannot report CSI/AoA in the monitor mode.

HotMobile '18, February 12–13, 2018, Tempe , AZ, USA

ACM ISBN 978-1-4503-5630-5/18/02…$15.00
https://doi.org/10.1145/3177102.3177106

WiFi Camera	TX Power (dBm)	Avg. Packet Size (bytes)	Packet Rate (pkt/s)
Yi Home Camera (720p, 2.4GHz)	16-18	650	60
Yi Home Camera2 (1080p, 2.4GHz)	16-18	645	75
Amcrest ProHD (1080p, 2.4GHz)	~19	1190	80
Samsung SmartCam (1080p, 5GHz)	~19	900	50

Table 1: WiFi cameras used in our experiments.

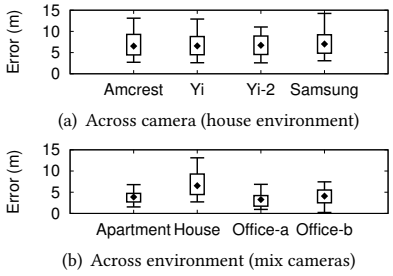

(a) Across camera (house environment)

(b) Across environment (mix cameras)

Figure 2: Localization performance in terms of quantile (5%, 25%, 50%, 75%, 95%).

We then study possible defenses against these attacks, which add noise to RSS signals in both temporal and spatial domains. We find that these defenses are moderately effective against the basic attack, but have minimum impact when the attacker uses advanced ML techniques. Further, a more sophisticated attacker deploys countermeasures in the form of additional stationary receivers, which help remove the impact of noise in the temporal domain. Overall, current defenses fall short against these localization attacks.

Our work seeks to bring attention to the practical dangers of widespread deployment of IoT devices like home video cameras. By taking a brief walk with today's COTS WiFi receivers, our simple attacks can already achieve 2.7 meters in median localization error and minimum variance in four common home/office environments. The efficacy will further improve as the adversary uses more advanced algorithms or hardware. We need to quickly develop effective defenses against these attacks to improve the security of wireless IoT devices in the home/office.

2 ADVERSARIAL LOCALIZATION

Our goal is to understand the effectiveness of adversarial localization on behind-the-wall wireless cameras. We performed actual attacks on popular COTS WiFi cameras in typical home and office scenarios. We now describe the attack model, our experiment configuration and dataset, and our findings.

2.1 Adversarial Model

We consider an adversary who physically moves outside the target house/apartment/office, seeking to localize WiFi cameras behind the walls (Figure 1). While moving, the adversary uses a standard WiFi receiver, *e.g.* a laptop or a smartphone, to *passively* sniff transmissions from nearby WiFi devices [25]. From the sniffed data, the adversary extracts non-payload information like MAC address, RSS and frame length. With these information, the adversary can identify WiFi cameras by directly matching their MAC addresses if she is knowledgeable enough or by analyzing traffic patterns, *e.g.* traffic volume and packet types [27]. After recognizing each target camera, the adversary builds a signal trace per target, *i.e.* a set of tuples (time, position, RSS) along a moving trajectory, and applies TX localization to estimate the target's position.

Environment	Adversary Behavior
Apartment	Adversary walks on the outdoor hallway.
House	Adversary walks on the lawn and sidewalk.
Office-a	Adversary walks on the indoor hallway.
Office-b	Adversary walks both inside & outside the building.

Table 2: Environments where we performed the attacks.

TX Localization. We consider RSS based passive TX localization, which does not require the adversary to communicate with the victim (for stealthiness). After de-noising RSS trace using window-based averaging, we apply the log-distance path loss model to estimate the camera location. We chose this method because it is simple, widely used, and has been shown to be robust against biased spatial coverage [17]. Thus our results provide a lower-bound on the accuracy of adversarial localization using COTS WiFi sniffers.

2.2 Measurements

We performed measurements on four WiFi security cameras with the highest ratings on Amazon (see Table 1). They use different WiFi chipsets and frequency bands (2.4GHz and 5GHz). We placed these cameras in rooms of resident houses, apartments and office buildings, and varied their locations (3 locations per room). In each experiment, the camera and the adversary were separated by walls (of different building materials). Table 2 summarizes the settings.

The adversary uses a laptop (Macbook Air) to sniff WiFi traffic and a smartphone (Samsung Galaxy SIII) to track moving trajectory. The adversary applies dead-reckoning on smartphone sensor data, using the accelerometer readings to track walking distance and the orientation readings to track angles. We performed experiments to confirm that the trajectory error has negligible impact on localization performance.

Our experiments were carried out by six people with different heights, weights and walking behaviors. Since they lead to consistent results, we did not differentiate them in our following discussions. Overall, our experiments produced more than 1.2k walking traces (each of 25-60 meters long), mapping to more than 2.6 million (time, position, RSS) tuples.

2.3 Attack Effectiveness

We quantify the accuracy of adversarial localization by *absolute localization error*, *i.e.* the distance between the estimated transmitter location and the ground truth. Figure 2(a) plots the quantiles (5%, 25%, 50%, 75%, 95%) of the localization error for each camera across all the walking traces. We see that the localization performance is similar across the four cameras (who use different WiFi chipsets and carrier frequencies). The median error is around 4-5 meters, which is within the room level. Yet the variance is significant, and the error can reach 12 meters. This is as expected since previous works have shown that transmitter localization is highly sensitive to measurement coverage, environmental dynamics, noises and RF interference [17, 19].

Figure 2(b) then plots the localization error quantiles for different environments, mixing all the camera results. Interestingly,

the accuracy varies across the environments. The error is particularly large for "house" due to more complex fading profiles and longer distance between the camera and the adversary compared to the other environments. In this case, the median error rises to 6.8 meters, while the rest three environments remain 4 meters. But a consistent trend is that the variance of localization error is large.

3 MINIMIZING VARIANCE

Given the large variance in localization accuracy, the basic adversarial localization faces heavy uncertainty on its effectiveness. To reduce variance, the most intuitive method is to perform multiple rounds of measurements[4] and aggregate the raw data or the localization results. We tested this approach on our dataset and found that the improvement is limited and often saturates after three rounds. Furthermore, repeatedly wandering around the target home will easily raise red flags.

Instead, we propose a method to predict the localization accuracy or *fidelity* of any measurement instance, using unsupervised learning analysis. It allows the attacker to separate high-quality measurements from noises and carry out the attack more effectively.

3.1 Predicting Localization Accuracy

Given a measurement instance (*i.e.* a round of sniffing), the adversary first quantifies the accuracy of the localization result produced from the data (referred as *fidelity*), and only uses the localization result if the predicted accuracy is high.

To estimate fidelity, the adversary uses *unsupervised feature clustering* [19]. Feature clustering groups data instances together based on their *similarity* across a small group of key features. It assumes that a specific combination of data features tends to coexist in measurement instances that produce accurate localization results. If this assumption holds, such clusters will be easily identifiable, and will reveal the key features that strongly correlate with fidelity. Using these key features, the adversary can determine whether a measurement instance is of high fidelity or not. While our recent work has used this approach to prune crowdsourced cellular measurements [19], we seek to use it to analyze behind-the-wall WiFi signals, which have different propagation properties.

We applied this feature clustering approach on our dataset. The detailed procedure is similar to [19], thus omitted for brevity. We start from four groups of features extracted from the raw measurement data: *packet* features on traffic statistics, *spatial* features used by common spatial analysis [26], *RSS* features on RSS statistics, and *combined* features that capture the fitting error of the localization model and the joint distribution of packet/RSS and spatial properties. We also included the environment type as a feature.

The feature clustering results are shown in Figure 3. We obtained two key findings.

1. Natural Clusters on Localization Accuracy. Our results confirm the strong tie between feature clusters and localization accuracy. Figure 3(a) shows that the measurement instances in our dataset are divided into 3 clusters. Cluster A (49% of instances) produces fairly accurate localization results (2.5 meters and 5.5 meters for median and 95%-tile, respectively). Cluster B (35%) and C (16%) have relatively high localization errors (almost all > 5 meters).

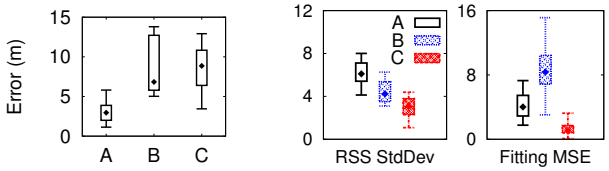

(a) Clustering result (b) Top features

Figure 3: Feature clustering performance.

Figure 3(b) shows that two key features, *RSS standard deviation* and *fitting mean squared error (MSE)*, can be used to distinguish cluster A from the others. Measurement instances in cluster A display large RSS standard deviation but low model-fitting MSE. This is because RSS standard deviation increases as the adversary moves closer to the transmitter (camera). In this case, RSS measurements become more reliable in the presence of environmental artifacts. The next key feature is model-fitting MSE. Intuitively, this value should be small to have desirable localization accuracy. The exception is cluster C which has low values and yet bad localization results. This is because cluster C mainly consists of measurements that are far away from the transmitter and the RSS readings are *flat* across the trajectory. This type of data leads to a good model-fit, but is unsuitable for localization (since they do not capture the distance-RSS relationship).

2. Consistency across Environments/Cameras. We also observe that the clustering results and feature properties are *consistent* across all environments, cameras types and locations. This aligns with that of [19] on outdoor cellular localization. Such consistency greatly facilitates the attack. An adversary can use offline, local measurements to build clusters and determine the key features used to identify high quality measurement instances. It can then derive the fidelity level instantaneously during the actual attack.

Based on this observation, we configure the feature thresholds to identify a high-fidelity measurement instance (*i.e.* those in cluster A) as RSS standard deviation above 5 and fitting MSE lower than 6. These feature thresholds do not change when we only use data from any individual environment or camera.

3.2 Advanced Attack Performance

We now show the localization performance when the adversarial uses fidelity estimation to *prune* the data, *i.e.* only using a measurement instance if its fidelity value is sufficient. As baselines, we also include results of the basic attack and *data combining* where we aggregate three rounds of measurements since the performance saturates at this point.

Gain of Pruning. We first look at the "house" environment, which has the *worst* localization performance of the four environments. Figure 4(a) shows the localization performance for basic attack, after combining, and after pruning. Consistently across all four cameras, pruning offers significant performance improvement, reducing the median error from more than 6.8 meters to 4.2 meters, and bounding the error by 6 meters (compared to 15 meters in the basic attack). Next, Figure 4(b) plots results of the four environments while mixing results of all four cameras. Again pruning

[4]We assume that the attacker walks in the same area across measurement rounds, since his moving space within the camera's WiFi signal coverage is often limited, *e.g.* sidewalks, hallways etc.

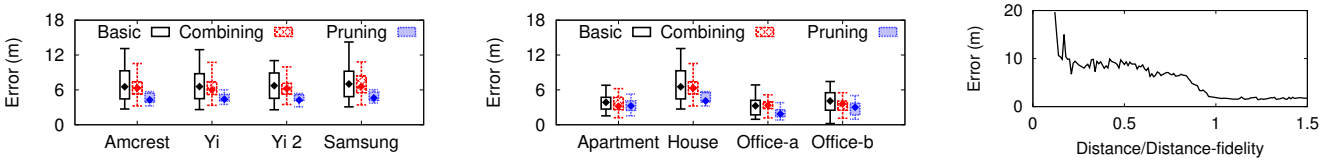

(a) Across camera (house environment) (b) Across environment (mix cameras) (c) Normalized walking distance vs. error

Figure 4: Localization performance of advanced localization attacks.

leads to significant improvement in localization accuracy, reducing variance from 6.4 to 1.8. This shows that identifying useful data is much more effective than blindly adding data.

Cost of Pruning. While effectively reducing the variance of localization accuracy, pruning comes at the cost of often requiring extra rounds of measurements until a high-fidelity instance is reached. Since we carried out 15 rounds of measurements per environment/camera/camera location configuration, we were able to compute the number of extra rounds required. Overall, 64% of measurement instances were determined as high fidelity by our pruning tool. For 17% of cases, no useable instance was found after 15 rounds of measurements, meaning that the attacker was unable to localize the camera with reasonable accuracy. For the rest, it takes an average of 3.25 rounds to reach a high-fidelity instance.

Minimizing Measurement Distance. The attacker can continuously monitor the fidelity level of the current measurement. After a sufficient level of fidelity is reached, he can terminate the measurement (for stealthiness and efficiency). Figure 4(c) plots the localization accuracy with the moving distance. We normalize the x-axis by the distance where the required fidelity level is achieved (Distance-fidelity). We see that our fidelity estimation can accurately identify the "stopping point". Across all the high-fidelity instances (766), the min, median, max of Distance-fidelity are 6.1, 23.9, and 53.2 meters, respectively. Similarly, the attacker can use fidelity to terminate "unsuccessful" measurements.

Together, our experiments on the advanced attacks show that adversarial localization is highly effective against WiFi cameras (behind the walls). By walking around the building/house briefly, the attacker can estimate the locations of indoor WiFi cameras with room-level accuracy (median of 2.7 meters), and identify unreliable measurements.

4 PRACTICAL DEFENSES

In term of defense, ideally, one can reconfigure WiFi hardware or install signal reflectors to prevent signal coverage beyond the wall/room [8, 14, 29]. Yet in practice, such approach incurs high cost, limits the camera's connection to the AP, and is often infeasible to deploy. Instead, we consider a suite of practical defenses that add noise to RSS signals in both the temporal and spatial domains, by modifying the camera WiFi transmit power or adding extra "cameras".

- *Temporal Obfuscation* – The camera adjusts its WiFi transmit power randomly over time, creating random noise on RSS values observed by the adversary. This idea originated from the power adjustment scheme in [12].
- *Spatial Obfuscation* – Extra WiFi transmitters, *e.g.* cameras, are placed away from the victim (in neighboring rooms). They coordinate with the victim to produce transmissions that are well mixed into the victim's transmissions so that the attacker cannot separate them during sniffing[5]. The design was inspired by recent works on anti-sensing using full duplex radios [24] and device cooperation [23].

Next, we describe both defenses, our experiments to validate them, and potential countermeasures by the adversary.

Key Findings. Current defenses fall short against adversarial localization attacks. While moderately effective against the basic attack, they have minimum impact when the attacker uses advanced ML techniques like pruning. The attacker can also deploy countermeasures in the form of additional stationary receivers, which help remove the impact of noise in the temporal domain.

4.1 Defense 1: Temporal Obfuscation

When the victim camera changes its transmit power randomly, the attacker will observe noisy RSS values that could degrade the localization results. The power variation needs to be large enough to create RSS distortion, and yet stealthy enough to avoid being predicted or removed.

Randomization Configuration. Three factors matter when configuring power randomization: *variation range*, *pattern*, and *frequency*. The range depends on the hardware capability, *e.g.* the dynamic range of power amplifier, and the required transmission range of the camera to reach its AP at a specific rate. The choice of the randomization pattern is non-trivial since there are many options. We consider *distribution-based*, *e.g.* Beta, Uniform, Exponential, Gaussian, and *pseudo-random number generators*. Among them, Beta provides the opportunity to maximize variance while *Uniform* (a special case of Beta) maximizes the entropy.

We first used trace-driven emulation to narrow down our choices before implementing them on our testbeds. Leveraging our RSS datasets, we emulate transmit power randomization by adding the instantaneous offset in transmit power (in dB) directly onto the attacker's RSS value (in dB). While making an ideal assumption that there's no packet loss and interference, this effort allows us to study the choice of variation range, pattern and frequency using repeatable experiments. For all our experiments, the average and maximum transmit power levels remain the same.

Our results show that the variation range needs to be larger than 10dB to make any visible change in localization results. Also the variation frequency needs to be carefully chosen. Rapid changes

[5]They use the same MAC address of the victim, and synchronize their packet sequence numbers for WiFi transmissions.

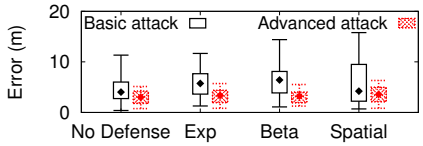

Figure 5: Localization performance of no defense, temporal obfuscation (Exponential (Exp), Beta distribution), and spatial obfuscation.

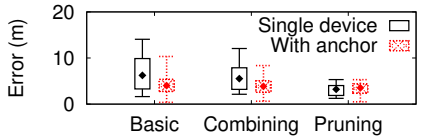

Figure 6: Attack performance w/ and w/o anchor. The victim uses power randomization (Beta).

can be easily mitigated via de-noising, and slow changes lead to little impact within each measurement round. Overall a frequency of one per 200 packets is desirable assuming the adversary moves at a walking speed. Finally, we narrowed down to the Beta (α=β=0.5) and Exponential distributions since they lead to the heaviest degradation in localization accuracy while holding large entropy values.

Testbed Experiments. While today's WiFi APs support power adaptation, this feature is not yet available on WiFi cameras. Thus we implement power randomization using USRP-GNU radios. Using the IEEE 802.11 implementation [7], we configure each USRP N210 node (XCVR2450 daughterboard) to mimic the camera operations in Table 1. We modify the automatic gain adaptation to enable 0-20dB gain variation at a granularity of 0.5dB. For each experiment, we configure the USRP transmitter to operate on three modes sequentially: fixed power (with 10dB gain), Beta randomization, and Exponential randomization. Overall, our dataset includes 0.9K walking traces in all four environments, with 4.2 million (time, position, RSS) tuples.

Results. We compare the efficacy of both basic and advanced attacks with and without the defense. We mix the results across all experiments since they are consistent. Figure 5 shows that power randomization is moderately effective against the basic attack, but has minimum impact when the attacker uses pruning. Here we assume the adversary is aware of the use of power randomization, and can adjust the feature thresholds properly. We verified that the clustering results and key features under power randomization are consistent with those under fixed power. The exact feature thresholds depend on the randomization pattern and range, which can be easily estimated by sniffing signals at a fixed location over a short period of time (<8 minutes shown by our experiments).

On the other hand, we do observe that power randomization increases the cost of data pruning. Now only 32% of measurement instances can offer high fidelity results (compared to 63% under fixed power). And in average it now takes 12 measurement rounds to obtain a high-fidelity localization outcome (compared to 3 rounds under fixed power). Therefore, temporal obfuscation by power randomization can increase the cost of localization attacks. However, later in §4.3 we show that such advantage can be easily diminished by attackers deploying additional stationary receivers.

Impact on Cameras. By lowering the camera's transmit power, power randomization could lead to undesirable packet losses and throughput drop. While we did not observe this artifact in our experiments, it can become a limiting factor when the link between the camera and the AP is already weak.

4.2 Spatial Obfuscation

In this defense, we place another camera (USRP with fixed power) in the neighboring room of the victim camera (USRP with fixed power), and repeat our experiments. The distance between the two is between 3-22 meters (with median of 8 meters). To avoid implementing coordination between the two cameras, we simply mix their RSS traces as a single one. Note that given the tight coordination requirement, this defense is much harder to implement in practice compared to power randomization.

Results. Figure 5 shows the defense is moderately effective against the basic attack. It raises the variance of localization error beyond that of power randomization. However, the attacker can overcome the defense by applying feature clustering to identify high-quality measurement instances, using the same feature thresholds of single camera scenarios. With pruning, the measurement overhead is slightly worse than that under single camera case: 47% of the measurement instances are now useable (compared to 63%); it takes an average of 3.67 rounds of measurements to reach a high-fidelity result.

Another interesting finding is that the location of the second camera is an important factor but hard to optimize. It should be away from the victim to create large noise on RSS, and yet close to the attacker (of unknown location) to produce strong WiFi signals that will be used in localization. Overall, this defense is difficult to optimize and costly to implement.

4.3 Countermeasure by Adversaries

Defenses via power randomization can also be countered. In theory, the RSS contribution from power randomization can be removed by using a stationary WiFi sniffer (referred to as anchor) to measure RSS continuously in the attack area [16, 18]. Subtracting the RSS change seen by the anchor from its own RSS, the attacker can remove the randomization contribution. Of course, this assumes that the anchor and the receiver face similar channel conditions.

To validate this countermeasure, in each of our previous USRP measurements, we also deployed three stationary anchors to collect RSS. We find that the performance depends heavily on the anchor choice. A good anchor will capture the majority of the victim's packets to build a comprehensive RSS trace for power subtraction. Thus in each experiment, we picked the anchor with the lowest packet loss rate.

Figure 6 compares the localization error of the basic and advanced attacks with and without the use of anchors, when the victim uses power randomization (Beta). We see that the use of anchors largely improves localization accuracy under the basic attack. When the attacker uses data pruning, the accuracy improvement becomes negligible, but the pruning overhead reduces significantly. 52% of measurement traces are useable (compared to 32%) and the

average number of measurement rounds reduces from 12 to 3.8. Thus the countermeasure is effective.

5 RELATED WORK

Indoor WiFi TX Localization. Existing works can be classified into *fingerprinting*, *time*, *AoA*, and *RSS*. We use RSS based methods because they do not require active communication with the target, and can be easily carried out via compact, COTS WiFi receivers/sniffers. Our goal in this work is not to minimize localization error (which we leave to future work), but to demonstrate the practicality and effectiveness of adversarial localization attacks.

TX Location Privacy. Researchers have developed multiple mechanisms to protect TX location privacy. Device anonymization [12] prevents the adversary from recognizing a device by its ID, but traffic and signal analysis [27] can still recognize devices like cameras. Others use antenna arrays [14, 29], directional antenna [6] or signal reflectors [8] to limit transmission coverage, but suffer from high cost and limited effectiveness due to environment constraints. Similar to our work, [4, 11] applied power randomization to prevent untrusted users from localizing themselves, but only drew conclusions from simulations. [24] designs full-duplex obfuscator to jam or obfuscate signals, but the high hardware cost prevents its adoption by IoT devices. [23] creates ghost locations from synchronized transmissions of devices in close proximity, but only targets fingerprinting-based localization. Our work is not to invent new defenses, but to show that existing defenses are inadequate. Finally, our work differs from adversarial localization in wireless sensor networks [10, 13, 22], which design routing protocols across a group of sensors to avoid localization of source nodes.

6 OPEN CHALLENGES

Our work highlights the dangers of reliance on IoT devices in a security context. Using commodity devices, attackers today can perform highly effective localization attacks against WiFi cameras inside the home/office. Significant work is needed to develop robust defenses against these attacks.

Adversarial localization for mobile devices. While this work targets stationary WiFi cameras, an open question is how adversarial localization attacks apply to other IoT devices in the current (and future) market. An example is the commercial mobile surveillance robots that survey a home while streaming videos. The camera movements and rotations certainly introduce new challenges in launching adversarial localization attacks. To locate and track a moving WiFi camera, the attacker will likely need more information beyond RSSI values.

Identifying and tracking multiple devices. Each home/room will likely deploy multiple (WiFi-based) IoT devices. Thus the attacker needs to separate and identify various target devices before localization. A straightforward solution is to use MAC addresses as the unique identifiers for each device. But some of today's devices already employ MAC address randomization to prevent being tracked constantly. Although existing works have proposed methods to reverse engineer MAC randomization via traffic analysis [21], its impact on adversarial localization remains unclear and needs investigation. Aside from MAC address based identification, the attacker can also perform traffic analysis to separate IoT devices if they have different traffic patterns. There are interesting follow-up works for our study.

ACKNOWLEDGEMENT

We thank the anonymous reviewers for their feedback. This work was supported by NSF grants AST-1443956 and CNS 1705042.

REFERENCES

[1] 2013. *Burglars Confess: Why Your Home is a Target.* http://www.alarm.org/HomeSafety/BurglarsSpillAboutSecuritySystems.aspx.
[2] 2016. *Minimize Blind Spots.* http://www.vivint.com/neighborhood/tech-neighbor/minimize-blind-spots-for-your-security-camera-system/.
[3] 2016. *Security Camera Blind Spots: How to Find and Avoid Them.* https://reolink.com/find-and-avoid-security-camera-blind-spots/.
[4] F. Anjum, S. Pandey, and P. Agrawal. 2005. Secure Localization in Sensor Networks Using Transmission Range Variation. In *Proc. of MASS.*
[5] P. Bahl and V.N. Padmanabhan. 2000. RADAR: An in-building RF-based User Location and Tracking System. In *Proc. of INFOCOM.*
[6] K. Bauer and et al. 2009. The Directional Attack on Wireless Localization. In *Proc. of GLOBECOM.*
[7] B. Bloessl, M. Segata, C. Sommer, and F. Dressler. 2013. An IEEE 802.11 a/g/p OFDM Receiver for GNU Radio. In *SRIF.*
[8] J. Chan, C. Zheng, and X. Zhou. 2015. 3d Printing Your Wireless Coverage. In *Proc. of HotWireless.*
[9] K. Chintalapudi, A.P. Iyer, and V.N. Padmanabhan. 2010. Indoor localization without the pain. In *Proc. of MobiCom.*
[10] N. Dutta, A. Saxena, and S. Chellappan. 2010. Defending Wireless Sensor Networks against Adversarial Localization. In *Proc. of MDM.*
[11] R. El-Badry, A. Sultan, and M. Youssef. 2010. Hyberloc: Providing Physical Layer Location Privacy in Hybrid Sensor Networks. In *Proc. of ICC.*
[12] T. Jiang, H.J. Wang, and Y. Hu. 2007. Preserving Location Privacy in Wireless LANs. In *Proc. of MobiSys.*
[13] P. Kamat, Y. Zhang, W. Trappe, and C. Ozturk. 2005. Enhancing Source-Location Privacy in Sensor Network Routing. In *Proc. of ICDCS.*
[14] Y.S. Kim, P. Tague, H. Lee, and H. Kim. 2012. Carving Secure Wi-Fi Zones with Defensive Jamming. In *Proc. of Asia CCS.*
[15] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti. 2015. Spotfi: Decimeter Level Localization Using Wifi. In *Proc. of SIGCOMM.*
[16] J.H. Lee and R.M. Buehrer. 2009. Location Estimation using Differential RSS with Spatially Correlated Shadowing. In *Proc. of GLOBECOM.*
[17] L. Li and et al. 2014. Experiencing and Handling the Diversity in Data Density and Environmental Locality in an Indoor Positioning Service. In *Proc. of MobiCom.*
[18] X. Li, Y. Chen, J. Yang, and X. Zheng. 2011. Designing Localization Algorithms Robust to Signal Strength Attacks. In *Proc. of INFOCOM.*
[19] Z. Li, A. Nika, X. Zhang, Y. Zhu, Y. Yao, B. Zhao, and H. Zheng. 2017. Identifying Value in Crowdsourced Wireless Signal Measurements. In *Proc. of WWW.*
[20] A.T. Mariakakis, S. Sen, J. Lee, and K. Kim. 2014. Sail: Single Access Point-based Indoor Localization. In *Proc. of MobiSys.*
[21] J. Martin and et al. 2017. A Study of MAC Address Randomization in Mobile Devices and When it Fails. *arXiv preprint arXiv:1703.02874* (2017).
[22] K. Mehta, D. Liu, and M. Wright. 2012. Protecting Location Privacy in Sensor Networks against a Global Eavesdropper. *IEEE Transactions on Mobile Computing* 11, 2 (2012).
[23] S. Oh, T. Vu, M. Gruteser, and S. Banerjee. 2012. Phantom: Physical Layer Cooperation for Location Privacy Protection. In *Proc. of INFOCOM.*
[24] Y. Qiao and et al. 2016. PhyCloak: Obfuscating Sensing from Communication Signals. In *Proc. of NSDI.*
[25] C. Sanders. 2011. *Practical Packet Analysis: Using Wireshark to Solve Real-World Network Problems.* No Starch Press.
[26] S. Shekhar and S. Chawla. 2003. *Spatial Databases: A Tour. Introduction to Spatial Data Mining.* Pearson.
[27] S. Siby, R.R. Maiti, and N.O. Tippenhauer. 2017. IoTScanner: Detecting Privacy Threats in IoT Neighborhoods. In *IoTPTS.*
[28] J. Wang and et al. 2016. Lifs: Low Human-effort, Device-free Localization with Fine-grained Subcarrier Information. In *Proc. of MobiCom.*
[29] T. Wang and Y. Yang. 2011. Location Privacy Protection from RSS Localization System using Antenna Pattern Synthesis. In *Proc. of INFOCOM.*
[30] Y. Xi, L. Schwiebert, and W. Shi. 2006. Preserving Source Location Privacy in Monitoring-based Wireless Sensor Networks. In *Proc. of IPDPS.*
[31] J. Xiong and K. Jamieson. 2013. ArrayTrack: A Fine-Grained Indoor Location System.. In *Proc. of NSDI.*
[32] M. Youssef and et al. 2006. Pinpoint: An Asynchronous Time-based Location Determination System. In *Proc. of MobiCom.*