

Highlights

Detecting Hidden Webcams with Delay-Tolerant Similarity of Simultaneous Observation

Kevin Wu, Brent Lagesse

- Description of the streaming video attacker model with delays.
- A novel methodology that is able to detect hidden Wi-Fi cameras with a mobile phone.
- The ability to defend against a delayed-transmission attacker model.
- Evaluation of the methodology in a variety of environments and conditions.

Detecting Hidden Webcams with Delay-Tolerant Similarity of Simultaneous Observation

Kevin Wu^a, Brent Lagesse^b

*University of Washington Bothell
Box 358534
18115 Campus Way NE
Bothell, WA 98011-8246*

^a*kevinw9@uw.edu*

^b*lagesse@uw.edu*

Abstract

Small, low-cost, wireless cameras are becoming increasingly commonplace making surreptitious observation of people more difficult to detect. Previous work in detecting hidden cameras has only addressed limited environments in small spaces where the user has significant control of the environment. To address this problem in a less constrained scope of environments, we introduce the concept of similarity of simultaneous observation where the user utilizes a camera (Wi-Fi camera, camera on a mobile phone or laptop) to compare timing patterns of data transmitted by potentially hidden cameras and the timing patterns that are expected from the scene that the known camera is recording. To analyze the patterns, we applied several similarity measures and demonstrated an accuracy of over 87% and an F1 score of 0.88 using an efficient threshold-based classification. We used our data set to train a neural network and saw improved results with accuracy as high as 97% and an F1 score over 0.95 for both indoors and outdoors settings. We further extend this work against an attacker who is capable of delaying when the video is sent. With the new approach, we see increased F1 scores above .98 for the original data and delayed data. From these results, we conclude that similarity of simultaneous observation is a feasible method for detecting hidden wireless cameras that are streaming video of a user. Our work removes significant limitations that have been put on previous detection methods.

Keywords: Security, Privacy, Internet of Things, Streaming Video,

1. Introduction

Internet connected cameras have become a pervasive feature in the world. Most modern mobile phones contain at least one camera as do many laptops. Additionally cheap Wi-Fi connected cameras are easy to obtain and deploy. In addition to these devices, there are a variety of hidden cameras that are designed to evade visual detection. The cost of obtaining and deploying such devices continues to drop as retailers such as Amazon include Surveillance Camera and Hidden Camera shopping categories that include thousands of results. While Internet-connected cameras bring convenience to the owners, they also create security risks. Weak security mechanisms allow adversaries to exploit those IoT devices and have total control over such devices. In 2016, Mirai malware took advantage of the weak password settings of IoT devices and compromised 3.5 million devices, many of which were Wi-Fi cameras [1]. The infected devices were located globally, including most of the countries in Europe, Asia, and North and South America [2]. While one of the most widespread, the Mirai botnet is just one of many examples of cameras being compromised [3, 4, 5]. Furthermore, Wi-Fi cameras have been installed to spy on people in environments such as hotel rooms and AirBnB rentals [6, 7, 8, 9].

Given the ease of which cameras can collect information on people without them knowing it, there is very little that has been done to detect cameras that are spying on people. Previous work in detecting hidden cameras has generally relied on being indoors, having significant control of the environment, or performing significant manual inspection with custom hardware [9, 10, 11]. In this paper, we extend our work [12] in automatically detecting Wi-Fi cameras to mitigate the delayed-transmission attacker model. The approach works both indoors and outdoors in large or small areas and can be accomplished with common computing equipment such as a mobile phone or laptop.

To address this problem, we introduce *Similarity of Simultaneous Observation* to identify cameras that are streaming video of a user. This is accomplished by utilizing a known camera in the environment such as the camera on a mobile phone and recording the environment. Simultaneously, a networking interface enters into monitor mode and records nearby data

transmissions and logs the number of bytes transmitted in each time step by each wireless device. Next, we apply similarity measures between the data timing of the known recording and each network device. Note that due to similarities in the size of plaintext and its resulting ciphertext when encrypted, this approach works regardless of if the camera is using encryption or is on another wireless network that we do not have credentials to join. If the two transmissions are deemed similar enough, then we flag that device as potential webcam.

We have evaluated our approach using over 21 hours of recordings taken from indoors and outdoors environments with varying levels of motion, resolution, and relative angles of the cameras along with a variety of traffic sources that are not observing the user in order to demonstrate the robustness of this approach. Our experimental results show that we can achieve 100% recall and F1 scores of 0.965 with a simple neural network and F1 scores over 0.98 with an LSTM against a more advance attack model than our original work in [12].

Contributions. The major contributions of our work can be summarized as the following four items:

1. Description of a problem that has not previously been addressed in research literature in section 3.
2. A novel methodology that is able to detect hidden Wi-Fi cameras with a mobile phone in section 4.
3. The ability to defend against a delayed-transmission attacker model in section 5.3.9.
4. Evaluation of the methodology in a variety of environments and conditions in section 5.

While the focus of our work was on streaming Wi-Fi cameras, the techniques would apply to any streaming camera as long as the system could acquire the per-time step byte counts of the device transmitting the data (for example, at a router).

2. Background

Our preliminary work [10] was the first known research to demonstrate that it is feasible to detect hidden cameras that are streaming video of a user by causing a change in the physical environment and comparing the bandwidth usage of the devices that could potentially be recording the user.

In this work, the flash on a mobile phone was used to illuminate the room, thus changing the pixels recorded by a hidden camera. This would cause a spike in bandwidth usage. The mobile phone uses a network card set to promiscuous mode to capture the traffic and then calculates the correlation coefficient between a vector of byte counts per time step and a vector of when the phone was flashing or not. Liu et al. [13] and Cheng et al. [14] published similar research shortly after that also used probes to detect hidden Wi-Fi cameras. Unfortunately, the techniques described in this work require a disturbance in the environment to operate such as rapidly flashing the flash LED on the mobile phone. This is generally not an activity that a user would want to perform during a meeting. Furthermore, the techniques described in these papers became increasingly ineffective in larger spaces, so it is not suitable for detecting cameras in outdoor areas or large open spaces such as shopping malls.

The reason these techniques work is due to the inter-frame video compression algorithms commonly used by Wi-Fi cameras, mobile phones, and video streaming applications. The most common modern compression algorithm used by Wi-Fi cameras, H.264, was first introduced in [15]. One of the improvements of the H.264(MPEG-4 Part 10) is the ability to reduce the size of a video file, which requires less network bandwidth and storage space. The H.264 achieved this by removing unnecessary information, specifically, the unchanged pixels between frames. Instead, the algorithm only encodes the changing pixels with respect to reference frames. Thus, more movements occurring in the environment forced the Wi-Fi camera and the mobile phone to generate more data in network traffic and video frames. Our system is not exclusive to H.264 and should work with any compression technique where the size of encoding at a given time is a function of the scene it is observing.

3. Problem Statement

In this section, we introduce the problem that we address in our research. To the best of our knowledge, no previous research has directly addressed this problem. **Given an arbitrary space, is it feasible to detect whether or not somebody is streaming video of that space.**

3.1. System Model and Assumptions

We assume that the user is interested in detecting a camera that is streaming video of them in an environment with a significant number of wireless

networks and potentially wireless cameras. In this paper, we refer to a scene as the area of observation recorded by a given camera. It is not enough just to detect that a device on the network might be a camera, but also that the device is recording the scene in question. As a result, there may be dozens of networks, dozens of streaming devices, and hundreds or thousands of total devices within range of the user.

We assume that the user has typical computing equipment available to them. For example, they possess a computer or a mobile phone and a network card that is capable of entering into monitor mode. We do not make explicit assumptions about whether the user is indoors or outdoors. We do not assume knowledge of the location of the Wi-Fi camera other than that it is within range of the wireless device that is in monitor mode. We do not assume that the user has credentials to join the network that the Wi-Fi camera is transmitting on.

3.2. Attacker Model and Assumptions

We make the following assumptions in this paper. This work focuses on currently publicized attacks such as those in hotels and off-the-shelf spy cameras. As a result, we assume the attacker lacks the motivation or technical skills to drastically reconfigure the camera. For example, the attacker may be an AirBnB owner or even somebody who has compromised a remote webcam by guessing the password. In this paper, we expand the attacker model beyond our previous work to include an attacker that has the ability to introduce delays in the video streaming as that caused misclassification in our previous work.

The work in this paper is designed to address 3 attacker models.

1. The attacker has placed a hidden camera.
2. The attacker has compromised a device with camera capabilities.
3. The user has deployed a device that is streaming video, but does not realize it.

3.3. Design Requirements

The purpose of our work is to help users detect that a device is streaming video of them. To this end, our work was approached with the following requirements:

- The system must work with common computing equipment that people tend to have with them most of the time.

- The system must work indoors or outdoors.
- The system must not require manipulation of the environment.
- The system must work even if the video is encrypted.

To the best of our knowledge, no known system or technique meets all of these requirements which has limited the effectiveness of camera detection techniques.

4. Methodology

We propose and evaluate the detection of Wi-Fi cameras passively by recording the environment. The detection mechanism analyzes timing characteristics that exist in the recorded video and the network traffic of the Wi-Fi camera.

The default behavior of Wi-Fi cameras is based on the video compression algorithm they use. H.264, a block-oriented, motion-compensation-based video compression standard, is utilized by many modern Wi-Fi cameras and streaming applications to transfer data efficiently. To reduce bandwidth usage, the standard only records motions between frames, in order to reduce storing overlapping information. Thus, a large amount of movement forces the Wi-Fi camera to generate and transfer large amounts of data, which creates peaks in network traffic.

The proposed framework has four major steps. The first step is to monitor the environment digitally by recording video and network traffic simultaneously. The recorded files contained timing characteristics that are essential to identify Wi-Fi camera. The second step is to extract a feature, specifically, the number of bytes per second, from both either the video file or the recorded network traffic file. This results in a vector of unsigned integers that represents each recording. The third step is to perform statistical analysis, calculating the Pearson correlation coefficient (CC), Dynamic Time Warping (DTW) distance, Kullback-Leibler divergence (KLD), and Jensen-Shannon divergence (JSD) on the bytes-per-time step vectors. The last step is to classify each vector as belonging to a spying camera or not. Descriptions of each steps and corresponding implementation are presented in the sections below. Figure 1 provides a visual overview of this process.

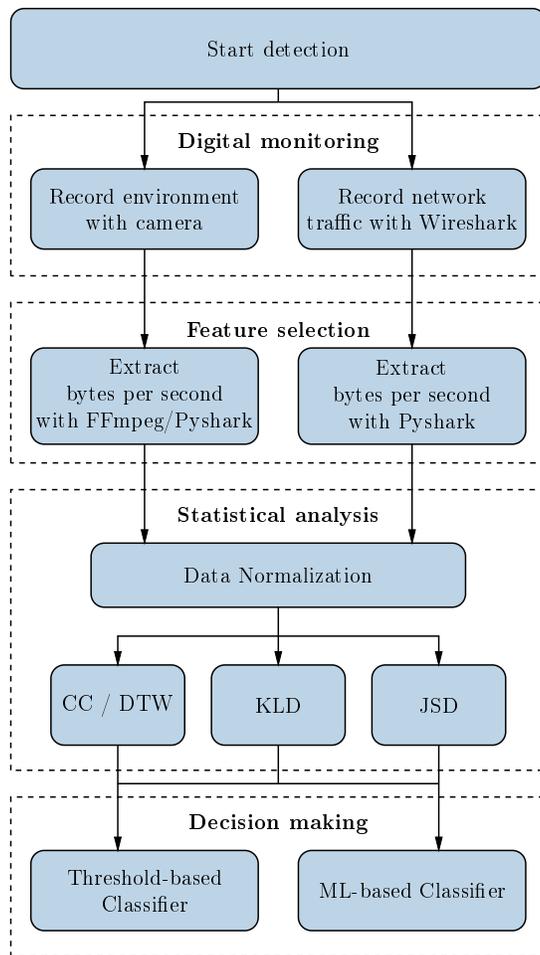


Figure 1: Flowchart of the two detectors.

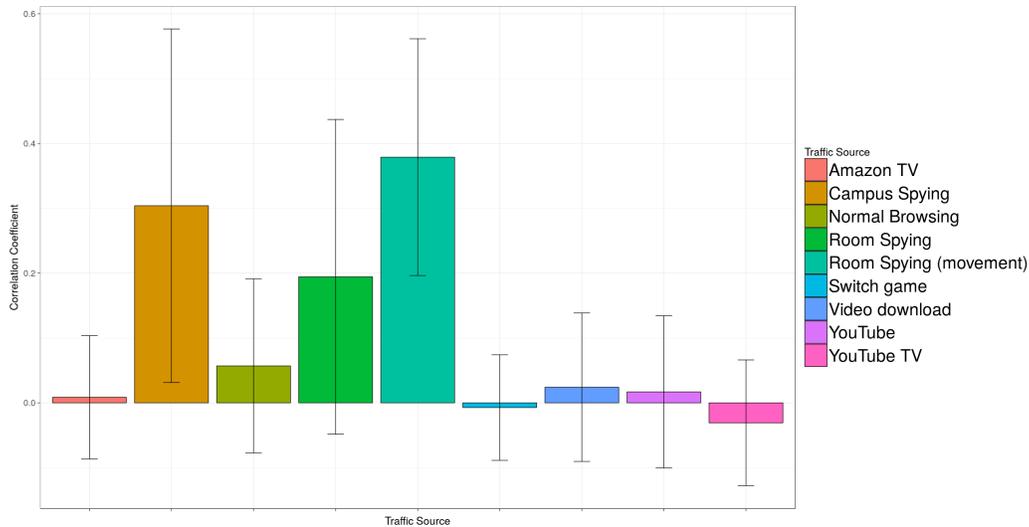


Figure 2: Correlation Coefficients for Various Traffic Sources (Error bars are one standard deviation above and below the mean)

4.1. Digital Monitoring

Digital monitoring is the first step in gathering data from the network traffic and the mobile phone. Network traffic is monitored while the mobile phone is recording the environment. In this step, the recording of the network traffic and the mobile phone are performed simultaneously.

4.1.1. Network Monitoring

In order to record the network traffic, a network sniffing tool is used with a network card in either promiscuous or monitor mode. Wireshark, an open source network sniffing tool supported in various platforms, is used to sniff the network traffic. In the experiments, Wireshark is used on a Macbook Pro with macOS High Sierra 10.13.4 to perform network monitoring. The version of the Wireshark software installed on the laptop is 2.4.2 and the Network Interface Card installed on the laptop is AirPort Extreme (0x14E4, 0x170) with firmware version of Broadcom BCM43xx 1.0 (7.77.37.29.1a7).

4.1.2. Video Recording

To retrieve data from the environment that is monitored by the Wi-Fi camera, video recording is performed from the back camera of the mobile

phone. The video recordings on the mobile phone also use a video compression algorithm to shrink the size of the video file. Mobile phones used H.264 to encode the video. This paper uses a Motorola-Z, with the OS version Android 8.0.0, to perform the experiments. The videos were recorded as either 720p or 1080p depending on the experiment, and are all in the length of one minute. The videos are encoded as MP4 files with audio support.

4.2. Process for Features

After the recording is completed, features are extracted from the recorded files to form data streams between IP addresses (if in promiscuous mode) or MAC addresses (if in monitor mode). Two data streams are further extracted from the recorded network traffic and the video file. While the recorded video is encoded as a MP4 file and the recorded network traffic is saved as a PCAP file, it is necessary to extract the same feature from the recorded files to perform statistical analysis. Bytes-per-time step, a shared feature in both MP4 and PCAP files, is extracted from the recordings. Experimentally we determined that 1 second time steps provided a good trade-off between timing differences of the devices and the amount of data that the device needed to send.

4.3. Perform Similarity Analysis

Initially, we utilized the techniques used by [10] to detect cyber-physical correlations; however, relying solely on Pearson's correlation coefficient resulted in an unacceptable number of false positives in some of our environments. As shown in figure 2, the correlation coefficient did result in visually different results; however, the standard deviations were so large that it was not useful as a classifier by itself. To counter this problem, we utilized several additional distance measures. In the case of comparing recorded videos with streaming network traffic, the correlation coefficient had so little predictive power that we did not include its results in the evaluation.

After the byte-per-second streams are extracted, we further conduct statistical analysis to calculate the relationship between the two data streams. Before performing any statistical analysis, data normalization is applied. In this project, Correlation Coefficient (CC), Dynamic Time Warping (DTW), Jensen-Shannon divergence (JSD), Kullback-Leibler divergence (KLD), Cramer distance (CD), Energy distance (ED), and Wasserstein distance (WD) are selected to measure the relationships between the two data streams. These

tests were selected because they capture a wide variety of ways that two distributions can be similar. CC is a statistical measure to calculate the correlation between two variables and was examined due to its effectiveness in our previous work [10], and DTW was examined since it is used to measure similarity between two temporal sequences. CC and DTW had the least predictive power, so we then considered other statistical measures, but we mention them here for informational purposes. KLD calculates the differences between two normally distributed data samples and JSD measures the similarity between two probability distributions. KLD was chosen because of the distances between the averages of spying and non-spying traffic while JSD was chosen because of the separation of the standard deviation of spying and non-spying traffic. WD measures the underlying difference in geometries of two probability distributions. CD is very similar to WD except that it also provides unbiased sample gradients. ED is also similar to CD except that it is rotation invariant. WD, CD, and ED were selected because they rely on the distance caused by the Cumulative Distribution Function which makes them more robust to minor timing mismatches caused by processing delay in the hidden camera.

4.3.1. Data normalization

Data normalization is performed to standardize the range of the variables in byte-per-second streams. This pre-processing step eliminates the effect of particular outliers and prevents certain objective algorithms from failing. This study utilized feature scaling to perform data normalization. Feature scaling re-scales all values in the data stream into the range between 0 and 1.

4.4. Decision Making

The results of the similarity analysis are used to decide whether the network stream is a Wi-Fi camera that is spying on the scene. We examined two methods for classification. One is a threshold-based approach where we identified values that most effectively differentiated between spying and non-spying devices. The second is a machine learning based classifier where we trained a neural network to differentiate between spying and non-spying devices.

4.4.1. Threshold-based approach

The threshold selection was conducted based on the number of tests. Each collected result is further compared with the proposed threshold to determine

the strength of the relationships. The threshold values are selected based on the corresponding F1 score. For each measure, we computed the F1 scores for various threshold values and selected the one with the highest F1 score.

4.4.2. Machine-learning-based approach

After studying the threshold-based approach, we observed that when the system produced errors, it was usually not for all of the metrics. Only in 24% of our errors did we observe that all of our metrics were incorrect. As a result, we decided to combine the metrics using supervised machine learning. We examined a variety of machine learning algorithms and were able to achieve significantly improved results by training a neural network.

5. Evaluation Procedure

In this section we evaluate the effectiveness of our approach to detecting hidden cameras in a variety of environments. The goal of our evaluation is to understand under which circumstances the approach is effective. We have evaluated the approach by analyzing both the network output of a Wi-Fi camera and a recording taken (but not transmitted) on a mobile phone. We have collected data under a variety of conditions as described in table 1 by varying the relative angle between the devices, motion in the space, resolution of the cameras, and whether the environment is indoors or outdoors. Through these experiments we demonstrate that our work is effective in environments that prior work [10] was not effective.

5.1. Detectors

We selected two likely options that a user would have to detect a streaming camera. The first of these is to use a Wi-Fi camera and the second is to use the camera on a mobile phone or laptop. Two Wi-Fi cameras are more likely to have stronger correlations between their network outputs due to the similarity of hardware; however, a user is more likely to carry a mobile phone than a Wi-Fi camera, so we examined both options.

5.2. Environmental Setup

The baseline of environment for our experiments is an 80 square meter room with lights on and with two individuals moving in space. For reference, the results in [10] began to significantly degrade when the device was further than 2 meters from the spying camera. For our outdoor testing, we recorded

a 250 square meter courtyard during the evening of a sunny day with one individual walking around in the space. We also performed some experiments on a university campus with a scene that was approximately 3000 square meters (results pertaining to this environment are labeled "campus").

As seen in Table 1, the testing environment of the experiments is an 80 square meter room with illumination. The window size of the recordings (network traffic recording and video recording) is 60 seconds. Different angles between the hidden Wi-Fi camera and the detectors are also being considered. Testing angles included 0 degree, 90 degrees, and 180 degrees. The video compression algorithm of the Wi-Fi camera is H.264 with 720p resolution, and the video compression algorithm of the mobile phone is H.264 with both 720p and 1080p as resolutions.

Table 1: PARAMETER SETTINGS OF THE EXPERIMENT.

Parameters settings	Parameters Tested
Wi-Fi camera	DCS-936L
Video compression	H.264/MPEG-4
Mobile phone	Google Nexus 6P
OS platform	Android 8.0.0
Video resolution	720p and 1080p
Room size	80 square meters
Courtyard size	250 square meters
Illumination level of the room	Bright
Testing angles	0, 90, and 180 degrees
Window of recording	60 seconds

5.2.1. Additional environments

In addition to the original test data set, we have also added a new environment. We collected data from two different cameras in a 10-square meter room with the light on. An Android-based Nexus 6P and the camera from MacBook Pro are used to perform data collection. We have collected data from different angles, including 0 degrees, 90 degrees, and 180 degrees. In total of 366 data samples have been collected. Testing Parameters of the new experiment is shown in Table 2 below.

5.2.2. Parameter setting

For this research, we used an Android-based Nexus 6P and a D-Link Wi-Fi camera (DCS-936L) to perform data collection. Network data was

Table 2: PARAMETER SETTINGS FROM NEW EXPERIMENTS.

Parameters settings	Parameters Tested
Video compression	H.264/MPEG-4/MOV
Mobile phone	Google Nexus 6P
OS platform	Android 8.0.0
Video resolution	720p and 1080p
Laptop camera	MacBook Pro
Video resolution	720p
Room size	10 square meters
Illumination level of the room	Bright
Testing angles	0, 90, and 180 degrees
Window of recording	60 seconds

encrypted via WPA2. Unless otherwise noted, the parameters in Table 1 were used for our experiments.

5.2.3. Collected data

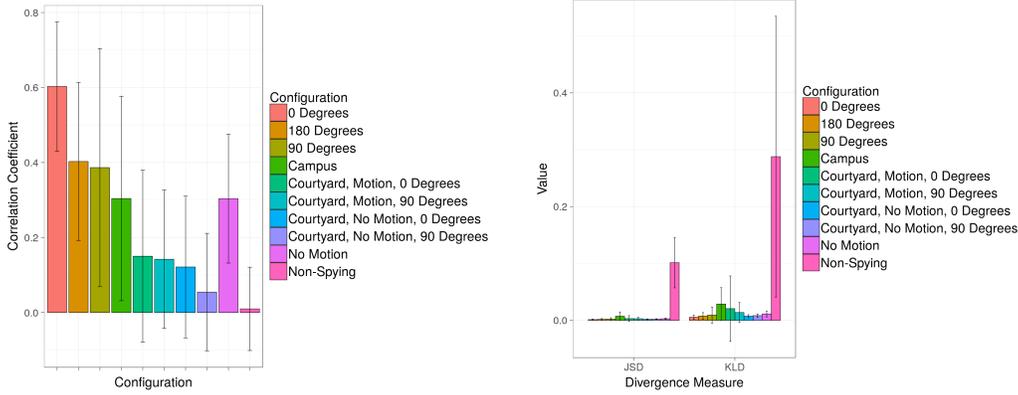
In this research, we have collected in total 830 data samples from the indoors room using the Wi-Fi camera, mobile phone. We collected 217 samples of traffic from outdoors. We collected 260 samples of non-spying traffic.

There is a mix of videos that capture motion and no motion. The Wi-Fi camera recorded at 720p and observed the scene relative to the spying camera at angles of 0, 90, and 180 degrees. The recorded video from the mobile phone included similar data except we also recorded additional data at 1080p.

We collected videos with both the Wi-Fi camera and the mobile phone of the outdoors courtyard. The videos were collected with and without motion. The videos that were collected without motion were done at a time when nobody was using the courtyard. For the motion videos, several people walked around in the courtyard while the cameras were recording. The camera and phone were both used to record the courtyard at 0 and 90 degrees relative to the spying camera. We also collected data from an outdoors portion of a university campus.

For non-spying camera traffics, we collected in total 260 data samples of network traffic from Skype, YouTube, YouTube TV, Amazon TV, Switch gaming, Normal browsing, and Video downloading. Those non-spying camera traffics are used in this paper to not only produce true positives, but also avoid false positives. We mostly focused on video-related traffic patterns,

but also included non-video data for diversity.



(a) Correlation Coefficient between Wi-Fi Camera and Spy Camera (b) JSD and KLD between Wi-Fi Camera and Spy Camera

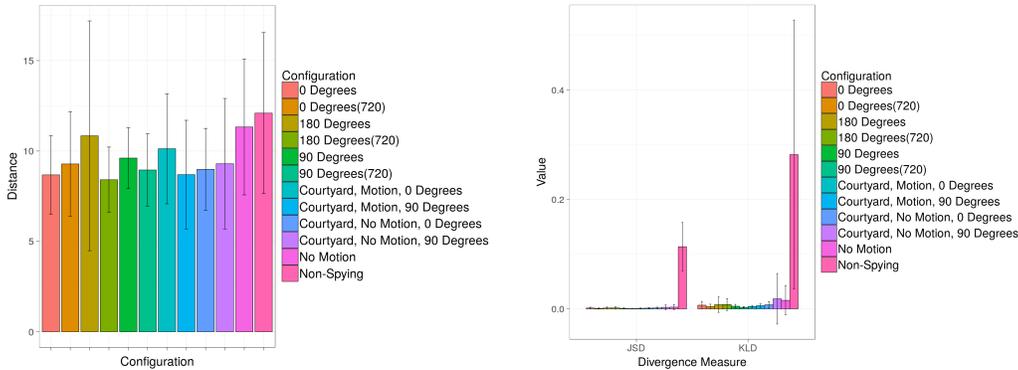
Figure 3: Wi-Fi Camera Detector

5.3. Results

In this section we present the results of the analysis of the data we collected. These results show that the correlation coefficient measurement used in [10] does not hold for larger outdoors spaces. They also show the added difficulty of measuring similarity between different types of devices. From these results, utilize additional distant measures and train a neural network to assist with classification.

5.3.1. Correlation Coefficient

Since previous work had relied on Pearson’s correlation coefficient, we first examined it as a similarity measure. These results can be seen in figure 2. Note that while all of the situations in which there was a spying camera on average are different than the non-spying traffic, the standard deviations caused a significant overlap between spying and non-spying traffic, so we concluded that we would be unable to use only correlation coefficients for classification. Likewise, we demonstrate in figure 3a that the difference between non-spying traffic and spy cameras degrades even further when we consider results from the outdoors scenario.



(a) DTW between Mobile Phone and Spy Camera (b) JSD and KLD between Mobile Phone and Spy Camera

Figure 4: Mobile Phone Detector

5.3.2. Similarity Measures

Next, we considered other measures for determining the similarity and differences between our recorded stream and the spy camera. We examined JSD and KSD as divergence measures and showed that they provided significantly different results in spying vs non-spying traffic. In figures 3b and 4b we see that for both the camera and the mobile phone, JSD has the most distance between one standard deviation above the mean for the spying video and one standard deviation below the mean for the non-spying video. Likewise, KLD provides the largest gap between the mean of the spying video and the non-spying video.

In our experiments between the Wi-Fi camera and the mobile phone, we noticed that there was a significant difference between the data usage of encoding on the phone and the traffic patterns of the Wi-Fi camera. We attribute this to the low power hardware used in the Wi-Fi camera as we noticed that there were often times of significant movement where the Wi-Fi camera did not transmit any data at all and then spiked in traffic shortly after the movement. This pattern caused the correlation coefficient to become almost useless, so we examined DTW as a distance measure. DTW distance was only a weak predictor of whether or not a device was a spy camera as seen in figure 4a.

5.3.3. Threshold-based Classifiers

After we analyzed similarity measures as suitable for determining the distance between spying and non-spying traffic, we analyzed our results to identify optimal thresholds for classification. The advantage of threshold classification is that it has a very low computational cost, so it has value as a classifier for low power devices. From this analysis, we identified the best thresholds for each measure based on F1 score as shown in table 3. Note that these are not necessarily always going to be the optimal threshold, but they do provide us with an understanding of an approximate starting point for a threshold-based classifier.

Table 3: CLASSIFICATION THRESHOLDS

CC	0.21
DTW	12.51
KLD	0.021
JSD	0.005

The results of the threshold-based classifiers can be found in table 4. As expected from the analysis of distance between the means and standard deviations, KLD and JSD greatly outperformed DTW with the mobile phone detector.

Table 4: THRESHOLD-BASED CLASSIFIERS.

Metrics	F1 score	Accuracy	Error	Precision	Recall(TP)
Wi-Fi camera-based detection model					
CC	77.642	77.005	22.994	81.159	74.418
KLD	88.643	87.165	12.834	84.384	93.355
JSD	83.208	84.841	15.158	76.497	91.208
Mobile phone-based detection model					
DTW	78.947	72.173	27.826	67.415	95.238
KLD	89.185	87.304	12.695	83.611	95.555
JSD	88.656	86.782	13.217	83.661	94.285

5.3.4. Machine Learning Classifiers

We examined the false positives that resulted from each of the different threshold measures and noted that only 24% of the time did all of the measures simultaneously produce a false positive. Table 5 provides a breakdown of the false positives. We hypothesized that we could utilize the lack of agreement between the similarity measures to improve our results via machine learning. We did not examine a majority vote system because that would have only eliminated the false positives in 56% of our samples.

We examined many standard classifiers to attempt to improve above the threshold classification method. Of these, we achieved the best performance

Table 5: FALSE POSITIVE COUNT

False Positives	Wi-Fi Camera	Mobile Phone
Total Samples	61	135
3	3.28%	33.33%
2	80.33%	10.37%
1	16.39%	56.30%

Table 6: F1 SCORES FOR PORTABILITY BETWEEN INDOORS AND OUTDOORS TRAINING FOR WI-FI CAMERA DETECTOR

		Tested		
		Indoors	Outdoors	Both
Trained	Indoors	96.55	62.50	67.24
	Outdoors	81.11	92.31	83.67
	Both	83.02	84.21	85.71

with a neural network. We performed grid search with 10-fold cross validation. For this study, a in total of 768 combinations of hyper-parameters are tested. We performed the grid search separately for both the Wi-Fi camera detector and the mobile phone detector and they both produced very similar models. The Wi-Fi camera detector’s selected model had L-BFGS as the solver and the Logistic activation function. It also had three hidden layers with 13 neurons in each of them. The only difference with the mobile phone detector was that each layer had 14 neurons.

5.3.5. Best Classifiers

Based on the results from sections 5.3.3 and 5.3.4, we selected the best threshold-based and machine-learning-based classifiers for the two detection models. The selected best classifiers are presented in Table 8 below.

Table 8 presents the best classifiers for the two detection models. As seen in the table, neural network models outperformed threshold-based classifiers both in terms of the F1 score and an accuracy rate achieving above 94%. Moreover, both of the neural network models had a 100% recall rate, so

Table 7: F1 SCORES FOR PORTABILITY BETWEEN INDOORS AND OUTDOORS TRAINING FOR MOBILE PHONE DETECTOR

		Tested		
		Indoors	Outdoors	Both
Trained	Indoors	96.55	73.68	78.79
	Outdoors	82.62	95.23	66.67
	Both	72.72	82.35	89.15

Table 8: THE SELECTED BEST CLASSIFIERS.

Classifiers	F1 score	Accuracy	Error	Precision	Recall
Wi-Fi camera-based detection model					
Threshold-based: KLD	88.643	87.165	12.834	84.384	93.355
Neural Network Indoors	96.551	97.436	2.564	93.333	100.000
Neural Network Outdoors	92.307	94.118	5.882	85.714	100.000
Mobile phone-based detection model					
Threshold-based: KLD	88.814	87.453	12.546	81.846	97.080
Neural Network Indoors	96.550	96.078	3.922	93.333	100.000
Neural Network Outdoors	95.238	96.774	3.226	90.909	100.000

scoring measurements that focus more heavily on True Positives would result in even better scores.

5.3.6. Convergence Time

While all of the tests described in this paper were run on 60 seconds of observation, we also examined the convergence rate of detection. We randomly selected 1 spying camera device and 69 non-spying camera devices then analyzed our results at each time step. Figure 5 shows that our results when averaged over 40 trials. Generally the spying camera is identified within a 10 seconds, and the rest of the time is spent weeding out the false positives. We see that the F1 score exceeds 0.90 within 20 seconds.

5.3.7. Model Portability

In this portion of the evaluation we examined the portability of the models between indoors and outdoors spaces. Figures 6 and 7 present a matrix

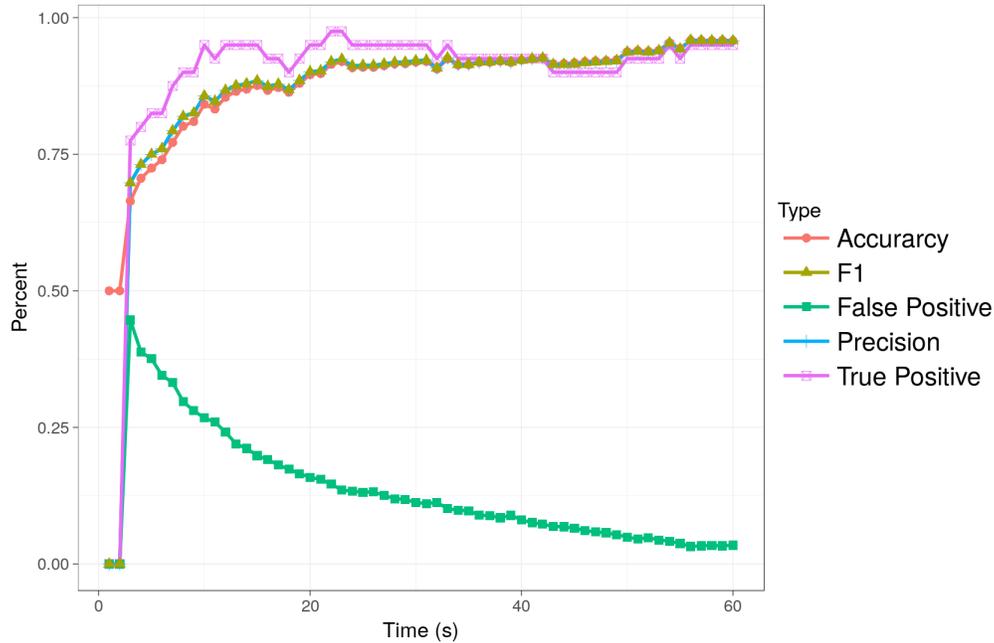


Figure 5: Time to Convergence

summary of the results by showing the F1 scores for our models when the data is partitioned into Indoors, Outdoors, and Both and then the model is trained and tested on samples from each set. From these results, we see that, as one would expect, the best results are achieved when the model is trained only with the class of data that it will be used to test with. We also note that training with the outdoor data provided much better results for non-outdoors testing than occurred with indoor training data. In general we conclude that it is best to use separate models for drastically different types of space, but even if you use a combined model, there will still be value to the results.

5.3.8. Performance of the ANN

We further examined the performance of the selected model in the new environment. The best ANN model was selected and trained with the old dataset, then made predictions on the new dataset we collected. The results of the predictions on the new dataset shown in Table 9. As shown in the table, the performance of the model suffered as the accuracy rate drops to

81.94%. The overall performance also drops while numbers of false positive and false negative increases. It is stated that different environment settings did affected the magnitudes of the changes pixels in the H.264 encoding thus degrades the performance of the model.

In order to boost the performance of the model, we introduced several new features into the ANN model. Cramer distance, Energy distance, and Wasserstein distance are further implemented into the new ANN model. Performance of the new ANN model is also presented in Table 9. Compared with the previous model, the model with new features is able to boost the recall rate to 93.37%. The new model trained with the new features can classify more true positives than the previous model. However, the accuracy rate and overall F1 score stays around the same ranged between 82% to 84%. It is also noticed that the Precision rate decreased by 4%, which indicates the new model classified slightly more false positives.

To prevent overfitting the model, we examined the performance of the model by reversed the training dataset with the testing dataset. By training the model with the new dataset and made predictions on the old dataset, we can observed the sustainability of the model’s performance in different environments. As shown in table 9, The model performs about the same as the previous model. Reversing the train-test dataset gave us slightly better results as all of the performance metrics improved slightly. The accuracy rate increased by 3% as well as the overall F1 score. The persistence of the performance has proved that the model could serve as a uniformed detection classifier that can make predictions in other environments with promising results.

Table 9: PERFORMANCE COMPARISON OF THE NEW MODELS.

Classifiers	F1 score	Accuracy	Error	Precision	Recall
ANN Indoors	96.55	96.08	3.92	93.33	100.00
ANN: New environment	83.30	81.94	18.06	80.94	85.82
ANN: New features	84.08	82.32	17.68	76.47	93.37
ANN: Reversed train-test	87.34	85.47	14.53	80.54	95.39

5.3.9. Video Delays

The adversaries might perform delay viewing attacks on the hidden camera internet stream to counter the detection. Delay viewing attack is carried

out by the intentionally delay the internet stream for several seconds to differentiate it from normal internet stream of a Wi-Fi camera. Addressing the possibility of the delay viewing attack on the hidden Wi-Fi camera, we compare and evaluate the performance of the old ANN model and the new ANN model. The results of the performance of the old model is presented in Figure 6.

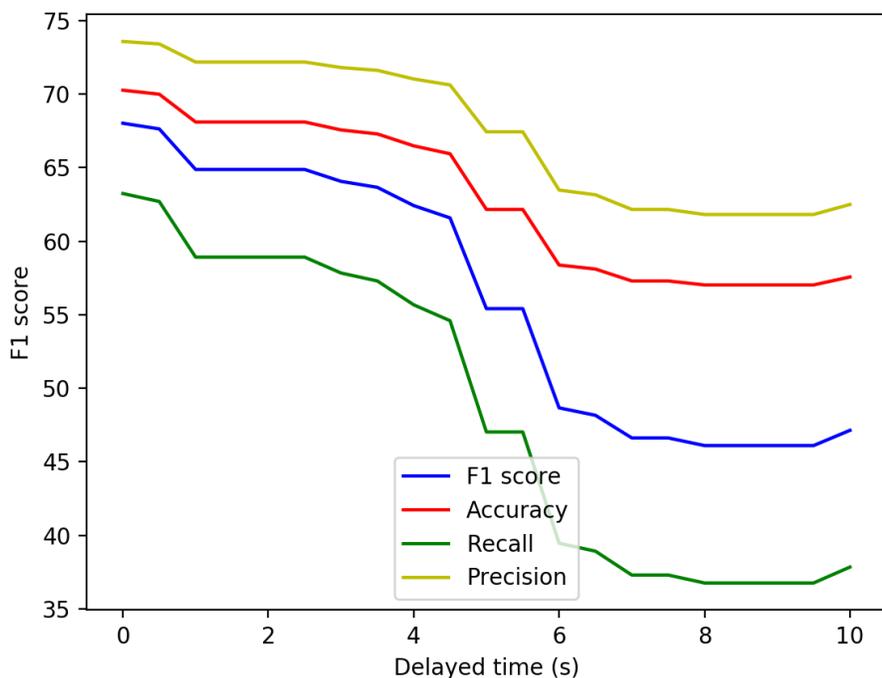


Figure 6: Effect of Delay Attack on Original Model

As shown in Figure 6, the old model did not perform well against the delay viewing attack. The accuracy rate stays around 67% within 5 seconds of delay then drops dramatically below 60% for longer delay. The precision rate is always higher than the recall rate indicates that the model predicts more false negatives than false positives. This has shown that the model was falsing predicting hidden camera stream as other internet streams. The performance of the model stays mediocre until the delay reached over 5 seconds. Recall rate even drops below 40% if the delay is longer than 6 seconds. It has shown

that the old model did not have the ability to detect hidden cameras under the delay viewing attack.

The performance of the new model under delay viewing attack is shown in Figure 7. As shown from the figure, the new model performs better than the old model while it maintained an accuracy rate around 83% within 5 seconds of delay then slightly decreased to 70% for longer delay. The new model performs opposite as the older model by having a higher recall rate than the precision rate. It is showing that the model predicts more false positive than false negative. The new model is more feasible than the older model since it had less false negatives, indicates that it can predict more hidden cameras rather labeled them as other internet streams. The recall rate maintained higher than precision rate and was higher than 80% until delay reach around 7.8 seconds. It is showing that the model had the ability to predict sufficient amount of true positives until the delay is longer than 8 seconds. The performance of the model drops dramatically over 8 seconds of delay and the accuracy rate maintained over 70% for 10 seconds delay. Although the performance of the new model degrades, it still gave us promising results for any delay attack shorter than 8 seconds. As a result, the new model served as a better classifier than the older model to counter delay viewing attack.

5.3.10. LSTM-based Approach

In order to achieve better predicting results, we implemented an LSTM model. The LSTM model is a well-known artificial recurrent neural network which predicts time series data. The selected LSTM model had 50 neurons in the first hidden layer, with 1 neuron in the output layer with the sigmoid function. It also utilized an Adam optimizer and the selected batch size is set to 48. A stepped function is applied after the output layer. A threshold of 0.00001 is applied within the stepped function. The LSTM model is trained with the old dataset and make predictions based on the new dataset. The performance of the LSTM model is shown in Table 10. As shown in the table, the LSTM model performs well on predicting time-series data. The accuracy rate had achieved 99.72% and the overall F1 score reached 99.76%. Moreover, the model is able to eliminate all the false positives while the recall rate achieved 100%. It shows that the model is perfectly fit for detecting all hidden cameras without mispredict any hidden camera stream.

To prevent overfitting the model, we also examined the performance of the LSTM model by reversing the training dataset with the testing dataset. As shown in table 10, The model performs about the same before reversing

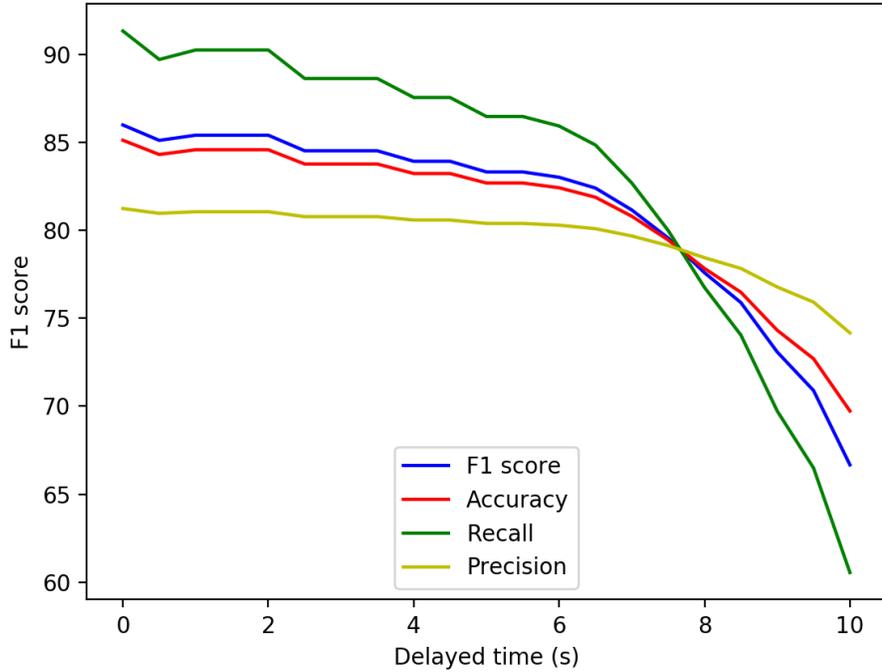


Figure 7: Effect of Delay Attack on Model with Improved Metrics

the train-test dataset. Both the accuracy rate and the overall F1 score drops around 1%, and the recall rate drops around 2.5%. It indicates that reversing the train-test dataset made the model allowed more false negatives. Besides the slight downgrade of the evaluation metrics, the overall performance of the model stays excellent as the accuracy rate of 98.61% and overall F1 score of 98.66%. It indeed showed that the LSTM model is perfectly fit for detecting hidden Wi-Fi cameras, moreover, it is a uniformed classifier that can predict hidden cameras in different environment settings.

To address the possibility of delay-viewing attack on the hidden Wi-Fi camera, we further evaluate the performance of the LSTM model on delayed dataset. We have evaluated the performance of the model on two different scenarios: delay viewing with 5 seconds and delay viewing with 30 seconds. As shown in table 10, The LSTM model still performs well for both of the delaying scenarios. Both of the accuracy rate and the overall F1 score stayed

between 97% and 98% while the error rate increased slightly by 2%. It has shown that the LSTM model had the ability to counter delay-viewing attack even with 30 seconds of delay.

Table 10: PERFORMANCE OF THE LSTM MODELS.

Classifiers	F1 score	Accuracy	Error	Precision	Recall
LSTM model	99.76	99.72	0.28	99.53	100.00
LSTM: Reversed train-test	98.66	98.61	1.39	99.90	97.44
LSTM: Delay 5 sec.	98.91	98.07	1.93	100.00	97.85
LSTM: Delay 30 sec.	98.89	97.80	2.20	97.80	100.00

6. Discussion

The results we obtained in this study demonstrate that there are 4 main points of concern for determining how accurately one can detect hidden cameras using the passive approaches described in this paper. These include the changes in the physical world that can be observed by the devices, the fidelity of the camera, the network transmission, and the background traffic from other devices. In other words, to theoretically predict your results, you need to answer the following questions: i) What is happening in the physical world? ii) How is it being recorded? iii) How is it being transmitted? iv) How is it different from other transmissions?

6.1. Scene Change

Scene change describes the scene that the cameras are recording. To demonstrate this point, consider two cameras that are facing each other with a television in between them. The camera facing the front of the television would record significant change whereas the one facing the back would record no change. The primary variables that can affect detection are the relative placement of the cameras which affects the portion of overlap of the recorded scene, and the magnitude of the movement in the overlap of the recorded scene. The placement of the cameras affect the detection since their location affects the number of pixels that are simultaneously altered due to a change in the scene between shared between two recordings. The magnitude of the movement in the scene affects the detection since no movement or constant movement will be easy to confuse with periodic network traffic that has a similar transmission frequency to the I-Frame transmission frequency for the codec or for near constant bitrate traffic, respectively.

6.2. Camera Fidelity

Camera fidelity describes the quality of the recording made by the camera. To demonstrate this point, consider an extreme case where the camera only records a single pixel that is either black or white vs a camera with 1920x1080 resolution. The higher resolution camera would be able to pick up subtle changes whereas the 1 pixel camera would not be able to do so. The primary variables that can affect detection are the resolution of the camera, the video codec, and the optics of the camera. The resolution affects the number of pixels that a change in the scene affects; normalization can mask this in some cases, but not when a particular movement fails to register a change in lower resolution cameras. The video codec and its associated parameters can affect how many pixels are reported as changed especially depending on the compression technique. The optics of the camera can affect how sensitive a camera is to change and whether or not minor changes are detected.

6.3. Network Transmission

Network transmission describes how the data is disseminated by the camera. To demonstrate this point, consider a camera that is streaming over TCP and a camera that is streaming over UDP. Congestion in the network could cause the TCP camera to back off and modify its transmission speed whereas the UDP camera would transmit as fast as data was available, so the exact same scene could appear on the network with different bandwidth consumption. The primary variables that can affect detection are transmission delays, differing protocols, and the differing parameters used even when the protocols are the same. The delay can be due to processing delay because of low-power computing hardware, a phenomenon we experienced in our experiments, or due to customization by the attacker to try to evade detection. As mentioned before, different protocols for transmitting data can affect the timing and quantity of data transmitted. Furthermore, some protocols that adapt to bandwidth availability can cause issues if they adapt during the middle of bandwidth sampling since it would throw off our normalization process. Similarly, each network transmission protocol can be configured with different parameters that could result in different timings or bandwidth usage patterns.

6.4. Background Traffic

Background traffic describes the network traffic that is being transmitted by devices other than the spy camera. Since the usefulness of detecting spy

cameras depends on being able to differentiate between the spy camera and other network devices, devices that have transmission patterns similar to the timing of movement in the recorded scene will result in false positives as mentioned in section 6.1.

6.5. Limitations

If an attacker switches from an interframe compression algorithm such as H.264 to an intraframe or constant bit rate compression algorithm then our technique will be ineffective at detecting that camera; however, this switch comes with a cost of increased bandwidth usage. While many cameras still support MJPEG our experience has been that the cameras we have evaluated default to H.264 and some of them no longer include MJPEG support. We could extend our approach to also record using MJPEG and look for correlations since JPEG will compress each frame differently based on the colors in the scene.

An attacker could also modify the software running on the webcam to inject additional signals into the data transmission rates that are not as expensive as a CBR codec would be. The topic of how an attacker can optimize this injection is a topic for future work. We would need to augment our system with approaches such as network anomaly or protocol detection to be able to detect such an attack.

Additionally, we are limited to streaming cameras with this approach. As future work we are examining improved techniques for detecting cameras that are not streaming data. Currently, this approach would need to be used as one technique in an anti-spying toolkit.

7. Related Work

Related research has focused on identifying services, applications, websites, and connected devices with various detecting mechanisms. Since network traffic contained critical information regarding communicating entities and ongoing communications, most of the research concentrated on detecting targets by utilizing the data embedded within network traffic. Some studies introduced in perform timing analysis is also related to our work.

7.1. Network traffic analysis

Geer et al. [16] demonstrate that network traffic analysis is a powerful tool to identify targets regarding of the network traffic volume that is generated

daily. Their research included several features of the network traffic, such as frequency, volume, and timing, that are favorable for the attackers to identify particular patterns. Moreover, encryption over network traffic does not prevent adversaries from studying those features. The findings allowed adversaries to identify certain behavior and services from the network traffic. Coull et al. [17] researched network traffic analysis for Apple iMessage. The study looked into the volume of the encrypted network traffic that is being transferred and found that adversaries can successfully learn the victim's actions, language used, and the length of the messages with 96% of accuracy.

Siby et al. [18] focused on an IoT-rich environment and privacy concerns. They discovered existing wireless infrastructure by analyzing the numbers of Frames, mFrames, cFrames, and dFrames; network traffic volume; and send-to-received ratio passively identify IoT devices. Gong et al. [19] studied the feasibility of Dynamic Time Warping (DTW) on network traffic patterns. The study showed that website fingerprinting is applicable, even with noisy network traffic, by applying DTW with traffic analysis.

7.2. Timing analysis

Feghhi et al. [20] researched the effectiveness of timing-based attacks against encrypted network traffic and were able to infer web pages more than 87% of the time. Other studies have demonstrated that performing timing analysis reveals victim nodes within anonymizing systems [21, 22].

Apthorpe et al. [23] performed experiments on IoT smart home devices. They discovered that the network traffic of those devices often revealed potential information about user interactions. Based on the sending/receiving rates of the streams, they were able to map live traffic to user behaviors. This research indicates that the network streams of IoT devices have certain attributes that are controllable by the users. We expect to adapt their findings to build a novel IoT sensor detection method based on certain movement interactions. A timing analysis on a low-latency network has also been discussed [21, 22]. Both studies have pointed out that the timing characteristics of network traffic tend to be remained. We intend to extend their findings to perform statistical analysis on the timing characteristics of Wi-Fi cameras.

8. Conclusion

This paper has proposed and evaluated a novel method, *Similarity of Simultaneous Observation*, for detecting streaming Wi-Fi cameras. This

method, as with the most effective prior research [10], works with common computing equipment and still works even if the attacker is using encryption or is on a different Wi-Fi network. Unlike prior work, this method works both indoors and outdoors without requiring any manipulation of the environment.

To validate the feasibility of this approach, we first analyzed the significance of the difference of several computationally efficient similarity measurements. Then, we examined the effectiveness of using those similarity measurements as a threshold-based classifier. Next, we applied machine learning to further improve our classification results. As a result, we demonstrated a threshold-based similarity measure that achieved an F1 score of 0.886 and a neural network model that achieved an F1 score of 0.966 with 100% recall across all of our scenarios.

Next we introduced a signal delay attacker model. The attacker can delay the streaming of the video. This attack drastically reduces the effectiveness of the original Similarity of Simultaneous Observation algorithm with very little delay necessary. To combat this, we introduce additional similarity metrics and use them to train an LSTM model. The new model not only defeats the delayed streaming attacker, but it also improves the overall performance of the system over our previous work.

From these results, we conclude that Similarity of Simultaneous Observation is an effective approach to detecting hidden streaming cameras in a variety of environments where previous work has failed. We have identified that there are some environments in which the technique performs better than others, but even in the most difficult environments our work is valuable.

9. Human Subjects and Ethical Considerations

The experiments described in this paper were reviewed by our IRB and were determined to be exempt from a full IRB review since any humans that were incidentally captured by our cameras were in public locations and the techniques rely only on the bytes per time step of the recorded video, not the content of the video.

References

- [1] B. Herzberg, D. Bekerman, I. Zeifman, Breaking down mirai: An IoT DDoS botnet analysis.

URL <https://www.incapsula.com/blog/malware-analysis-mirai-ddos-botnet.html>

- [2] S. Bobby, F5 labs hunt for IoT vol 3, accessed 2018-01-02.
URL <https://www.cbronline.com/whitepapers/f5-labs-hunt-iot-vol-3/>
- [3] Y. M. Pa Pa, S. Suzuki, K. Yoshioka, T. Matsumoto, T. Kasama, C. Rossow, IoTPOT: Analysing the Rise of IoT Compromises, Usenix Workshop on Offensive Technology (WOOT), 2015.
- [4] B. Krebs, Hacked Cameras, DVRs Powered Today's Massive Internet Outage, accessed 2017-11-04.
URL <https://krebsonsecurity.com/2016/10/hacked-cameras-dvrs-powered-todays-massive-internet-outage/>
- [5] S. Fogie, Abusing and Misusing Wireless Cameras, 2007, accessed 2017-11-04.
URL <http://www.informit.com/articles/article.aspx?p=1016099>
- [6] H. Coffey, How to spot a hidden camera in your airbnb, 2017, accessed 2017-12-05.
URL <https://www.independent.co.uk/travel/news-and-advice/airbnb-hidden-cameras-how-to-spot-online-holiday-rentals-apartments-secret-surveillance-a8092661.html>
- [7] Yvonne Edith Maria Schumacher vs Airbnb, Inc., a foreign corporation, and Fariah Hassim and Jamil Jiva, accessed 2017-12-05.
URL https://cdn2.vox-cdn.com/uploads/chorus_asset/file/5398067/1-main.0.pdf
- [8] J. Steinberg, These Devices May Be Spying On You (Even In Your Own Home), accessed 2017-11-04.
URL <https://www.forbes.com/sites/josephsteinberg/2014/01/27/these-devices-may-be-spying-on-you-even-in-your-own-home/>
- [9] P. Polstra, Am I Being Spied On? Low-tech Ways Of Detecting High-tech Surveillance (2014).

- [10] B. Lagesse, K. Wu, J. Shorb, Z. Zhu, Detecting Spies in IoT Systems using Cyber-Physical Correlation, IEEE Workshop on Mobile and Pervasive Internet of Things (2018).
- [11] M. Roessler, How to find hidden cameras (2002).
URL https://archive.org/stream/How_to_Find_Hidden_Cameras/How_to_Find_Hidden_Cameras_djvu.txt
- [12] K. Wu, B. Lagesse, Do You See What I See? Detecting Hidden Streaming Cameras Through Similarity of Simultaneous Observation, IEEE Pervasive Computing and Communications (2019).
- [13] T. Liu, Z. Liu, J. Huang, R. Tan, Z. Tan, Detecting Wireless Spy Cameras Via Stimulating and Probing, in: Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services, 2018.
- [14] Y. Cheng, X. Ji, T. Lu, W. Xu, DeWiCam: Detecting Hidden Wireless Cameras via Smartphones, in: Proceedings of the 2018 on Asia Conference on Computer and Communications Security, ACM, 2018.
- [15] H.264:advanced video coding for generic audiovisual services, accessed 2018-05-19.
URL <https://www.itu.int/rec/T-REC-H.264-200305-S>
- [16] K. Geers, Core illumination: Traffic analysis in cyberspace, in: 9th International Conference on Cyber Conflict (CyCon), 2017. doi:10.23919/CYCON.2017.8240328.
- [17] S. E. Coull, K. P. Dyer, Traffic analysis of encrypted messaging services: Apple iMessage and beyond, in: ACM SIGCOMM Computer Communication Review, Vol. 44, 2014.
- [18] S. Siby, R. R. Maiti, N. O. Tippenhauer, Iotscanner: Detecting privacy threats in iot neighborhoods, in: Proceedings of the 3rd ACM International Workshop on IoT Privacy, Trust, and Security, IoTPTS '17, ACM, New York, NY, USA, 2017.
- [19] X. Gong, N. Borisov, N. Kiyavash, N. Schear, Website detection using remote traffic analysis, in: S. Fischer-Hübner, M. Wright (Eds.), Privacy

Enhancing Technologies, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

- [20] S. Feghhi, D. J. Leith, Time and place: robustness of a traffic analysis attack against web traffic, in: 14th IEEE Annual Consumer Communications Networking Conference (CCNC), 2017. doi:10.1109/CCNC.2017.8067722.
- [21] S. J. Murdoch, G. Danezis, Low-cost traffic analysis of tor, in: 2005 IEEE Symposium on Security and Privacy, 2005. doi:10.1109/SP.2005.12.
- [22] V. Shmatikov, M.-H. Wang, Timing analysis in low-latency mix networks: attacks and defenses, Springer-Verlag, 2006.
- [23] N. Apthorpe, D. Reisman, N. Feamster, A smart home is no castle: Privacy vulnerabilities of encrypted iot traffic, Workshop on Data and Algorithmic Transparency (2016).