Poster Abstract: Passive Activity Classification of Smart Homes through Wireless Packet Sniffing

Kwon Nung Choi
The University of Sydney
kcho7166@uni.sydney.edu.au

Thilini Dahanayaka
The University of Sydney
tdah5330@uni.sydney.edu.au

David Kennedy
The University of Sydney
dken5837@uni.sydney.edu.au

Kanchana Thilakarathna
The University of Sydney
kanchana.thilakarathna@sydney.edu.au

Suranga Seneviratne
The University of Sydney
suranga.seneviratne@sydney.edu.au

Salil S. Kanhere
The University of New South Wales
salil.kanhere@unsw.edu.au

Prasant Mohapatra
Univeristy of California, Davis
pmohapatra@ucdavis.edu

ABSTRACT

Network communications, despite being encrypted, leak crucial information via side channels. WiFi networks are more prone to such side-channel attacks since any attacker within the network's range can passively eavesdrop the channel. With the increasing number of smart home devices and sensors connecting to private WiFi networks, it is essential to understand the inadvertent information leakage through WiFi side-channels. Our work demonstrates how fine-granular information on the activities happening inside a house can be inferred by passively monitoring WiFi network traffic. In particular, we were able to correctly classify various user interactions with simple IoT devices such as smart bulbs or power sockets as well as advanced voice-based intelligent assistants.

1 INTRODUCTION

The Internet of Things (IoT) has revolutionised smart home technology and enriched the occupants' experience. For instance, more control over home security is enabled through smart cameras, smart locks, and motion sensors. Air quality sensors and environment monitors maintain user well-being, while smart meters manage energy efficiency. Such diversity in smart home IoTs has made voice assistants an increasing necessity in the smart home as they enable seamless control of the entire home, with active units projected to reach 555 million by 2024 [8]. A vast majority of voice assistants and off-the-shelf IoT devices for the home segment are WiFi based.

Although most home WiFi networks are secured, several works demonstrated the possibility of information leaks through side-channels. Early work by Chen et al. [4] showed how very detailed personal information such as illnesses, medications, family income, and investment details could be inferred simply by passively observing WiFi frames. Similarly, Li et al. [6] showed how a particular video streamed by a YouTube user can be identified through WiFi traffic fingerprinting. Proliferation of smart home IoT further increases such privacy threats by revealing to eavesdroppers not only our online, but also physical activities happening inside our homes. Existing work to date has primarily focused on identifying IoT sensors through hardware addresses [9] or detecting the presence of any interaction rather than identifying the interaction type [3]. To this end, our work aims to demonstrate the potential of side-channel information leakage outside the network (or home) for two purposes: *identifying IoT sensors and the device-human interactions*.

The possibility of device and activity fingerprinting poses multiple privacy and safety concerns. Device fingerprinting allows burglars to determine potential victims based on identified expensive devices, or equipped with specific connected door locks which are prone to electronic compromise. Vandals can also inflict significant harm by remotely manipulating the operation of these connected devices, e.g., a washing machine could cause flooding or smart oven could trigger a fire [5]. Activity fingerprinting enables inference of home occupancy based on the activity patterns of motion sensors, door locks or air pressure sensors throughout the day [7]. In particular, identifying activity preferences can also lead to user profiling such as gender and age.

In this paper, we consider a realistic *threat model* where an attacker is within close range of the target house (10m) so that encrypted WiFi frames can be passively sniffed without having authenticated access to the network, as shown in Fig. 1. We instrumented a smart home using eight IoT devices and passively collected WiFi data while performing a number of interactions with these devices. Next, we show that it is possible to identify sensor/device types and manufacturers by simply referring to online databases of MAC addresses and device names, and classify activities performed by developing a Convolutional Neural Network (CNN) based classifier.

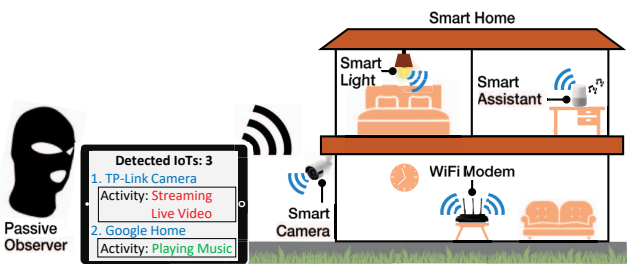

Figure 1: Threat Model. Scenario of the attacker scanning WiFi traffic from a home in close proximity, whereby the presence and activities of IoT devices are inferred.

978-1-7281-5497-8/20/$31.00 ©2020 IEEE
DOI 10.1109/IPSN48710.2020.00-13

No.	IoT Name	Total Hours	Total Packets
1	D-Link camera (8525LH)	20.19	51,219
2	D-Link camera (8300LH)	374.84	140,712
3	D-Link motion sensor	76.20	27,824
4	TP-Link smart plug	33.89	11,887
5	LIFX smart light	250.73	651,895
6	Netatmo environment monitor	119.63	287,665
7	AirBeam 2 air quality monitor	12.39	47,564
8	Amazon smart assistant (Echo)	39.37	5,312,349

Table 1: Summary of data collected from IoT devices.

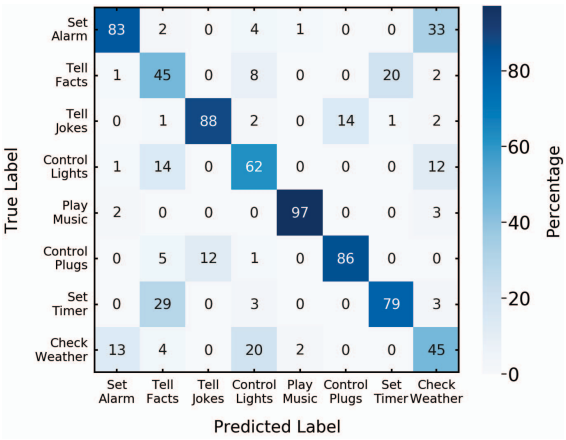

Figure 2: Confusion matrix of Amazon Echo activities.

2 DATASET & METHODOLOGY

Dataset - We collected packet traces sent to and from seven *simple* and one *complex* IoT devices when idle or active, using a Raspberry Pi 3 running Kali Linux using `tcpdump` with the wireless interface in monitor mode. Table 1 summarizes the total amount and duration of data collected for each device. *Simple devices* are those that have single purpose and limited scope of activity such as a smart light or motion sensor. *Complex devices* have a wider range of use cases and supported activities like smart assistants. As illustrated in Fig. 2, we selected eight common activities of a smart assistant for our investigation. In order to create large enough data samples and to overcome bias over changes in voice command, we utilized machine generated voice commands to interact with the assistant, collecting 1000 instances per activity over two weeks.

Methodology - We focused on two tasks: *identifying devices and activities*. **Devices** - We identified the *manufacturer* by looking up the device MAC address on the IEEE database [2], while *device type* was obtained by querying the same MAC address from a crowd sourced database Fingerbank [1]. **Activities** - For simple devices, differences in traffic rates, packet sizes and frequencies between idle and active state were used for inference using thresholding rules. For complex devices, we developed a CNN classifier and split the training, validation and test set to 80%/10%/10%, respectively. We selected packet sizes from each instance as the input vector and normalized it, assigning outgoing packets from sensors as positive and incoming as negative. The input vector size was set as 300 based on the mean number of packets per activity instance, with shorter instances zero-padded at the end. As music had significantly longer packets than others, we excluded it from the mean calculation.

3 RESULTS

We were able to correctly identify all 8 IoT devices in terms of their device type and manufacturer. For simple devices, large differences observed in the packet size between idle and active state made classification straight forward, apart from the motion sensor which had identical packet size and patterns even when active.

Interactions with Home-Assistants. For classifying home assistant activities, we obtained mean classification accuracy of 73.2%. Out of the activities investigated, classifying music performs best at 97%, as seen in Fig. 2. This is unsurprising given its large frame size and frequency, which is anomalous to other activities considered. Asking for jokes and the weather perform the poorest at 45%, mostly being misclassified as setting the timer or alarm, respectively. Although less in magnitude, the converse is also true for predicting classes for setting the alarm or timer. It appears that activities requiring a simple online query such as asking for facts or the weather portray similar traffic patterns to activities that use offline apps to carry out the task, such as setting the alarm.

4 CONCLUSION & FUTURE WORK

To summarize, our work examined the efficacy of classifying IoT devices and their activities under a realistic threat scenario. We developed a novel CNN classifier that only relies on the time and direction of packet transmissions to infer activities performed through smart assistants. We found that activities like streaming music that have distinctly large frame size and frequency are easier to classify, while activities involving simple Internet queries are more difficult due to having patterns similar to local query reliant tasks that do not require communication over the Internet, like setting an alarm.

In future work, we plan to increase the generalizability of the classifier by collecting data from multiple locations, under different networks and in different time frames. We will also investigate the limits of inference capabilities under considerable packet loss by increasing the distance between the observer and target location.

ACKNOWLEDGMENTS

This project is partially funded by Data61 and the Defence Science and Technology Group (DSTG) CRP through the Next Generation Technology Fund.

REFERENCES

[1] 2019. Fingerbank. https://fingerbank.org/. Accessed: 2019-12-23.
[2] 2019. IEEE Registration Authority. https://regauth.standards.ieee.org/standards-ra-web/pub/view.html#registries. Accessed: 2019-12-23.
[3] Noah Apthorpe, Danny Yuxing Huang, Dillon Reisman, Arvind Narayanan, and Nick Feamster. 2019. Keeping the smart home private with smart (er) iot traffic shaping. *Proceedings on Privacy Enhancing Technologies* 2019, 3 (2019), 128–148.
[4] Shuo Chen, Rui Wang, XiaoFeng Wang, and Kehuan Zhang. 2010. Side-channel leaks in web applications: A reality today, a challenge tomorrow. In *2010 IEEE Symposium on Security and Privacy*. IEEE.
[5] Tamara Denning, Tadayoshi Kohno, and Henry M Levy. 2013. Computer security and the modern home. *Commun. ACM* 56, 1 (2013), 94–103.
[6] Ying Li et al. 2018. Deep Content: Unveiling video streaming content from encrypted WiFi traffic. In *2018 IEEE 17th International Symposium on Network Computing and Applications (NCA)*. IEEE.
[7] Shwetak N Patel, Matthew S Reynolds, and Gregory D Abowd. 2008. Detecting human movement by differential air pressure sensing in HVAC system ductwork: An exploration in infrastructure mediated sensing. In *International Conference on Pervasive Computing*. Springer, 1–18.
[8] Juniper Research. 2020. Smart Home Statistics. https://www.juniperresearch.com/resources/infographics/smart-home-statistics. Accessed: 2020-01-24.
[9] Danny Yuxing Huang, Noah Apthorpe, Gunes Acar, Frank Li, and Nick Feamster. 2019. IoT Inspector: Crowdsourcing Labeled Network Traffic from Smart Home Devices at Scale. *arXiv preprint arXiv:1909.09848* (2019).