

Deep Adversarial Learning on Google Home devices

Andrea Ranieri¹, Davide Caputo², Luca Verderame², Alessio Merlo^{2*}, and Luca Caviglione¹

¹IMATI - National Research Council of Italy

{andrea.ranieri, luca.caviglione}@ge.imati.cnr.it

²DIBRIS - University of Genova, Italy

{davide.caputo, luca.verderame, alessio}@dibris.unige.it

Received: September 13, 2021; Accepted: November 1, 2021; Published: November 30, 2021

Abstract

Smart speakers and voice-based virtual assistants are core components for the success of the IoT paradigm. Unfortunately, they are vulnerable to various privacy threats exploiting machine learning to analyze the generated encrypted traffic. To cope with that, deep adversarial learning approaches can be used to build black-box countermeasures altering the network traffic (e.g., via packet padding) and its statistical information. This letter showcases the inadequacy of such countermeasures against machine learning attacks with a dedicated experimental campaign on a real network dataset. Results indicate the need for a major re-engineering to guarantee the suitable protection of commercially available smart speakers.

Keywords: Smart Speakers, IoT privacy, Deep Adversarial Learning, Machine Learning.

1 Introduction

The popularity of smart speakers (including voice-based virtual assistants) is rooted in their ability to control IoT nodes, network appliances, and other devices via natural speech. They can also be used to access multimedia contents and obtain various information, including news, weather forecasts, and traffic conditions. To implement such functionalities, the speaker exchanges information with a remote data center, leading to several security issues, e.g., device enumeration attacks, mass profiling, and privacy threats. An emerging trend in cybersecurity exploits machine learning techniques to obtain information from the encrypted traffic exchanged by the speaker with its ecosystem [25, 5].

Literature abounds in works investigating how statistical analysis of network flows produced by smart devices can be abused for reconnaissance or attack purposes. For instance, the traffic produced by home devices can be used to understand if a user is at home [10] as well as to model daily routines [1] or the sleep cycle [5]. In general, attacks leveraging machine learning proved to be effective, even when relying upon “poor” information. As an example, IoT nodes and connected devices can be identified by simply using the length of the produced protocol data units [16]. When HTTP-based interactions are present, it is possible to infer precise details, e.g., the status of a light bulb, as well as hijacking the conversation or physically endanger the target [3].

An emerging research trend explores the use of various artificial intelligence and machine learning techniques to classify the commands issued to smart speakers (see, e.g., [8] and the references therein). To this aim, attackers take advantage of traffic features not protected by the encryption, such as inter-packet time, throughput, the location of some endpoints, and the number of connections. Since the

Journal of Internet Services and Information Security (JISIS), volume: 11, number: 4 (November), pp. 33-43

DOI:10.22667/JISIS.2021.11.30.033

*Corresponding author: Department of Informatics, Bioengineering, Robotics, and System Engineering, Viale Causa 13, 16145, Genoa, Italy, Tel: +39-01033-52344

classification is typically accurate, the attacker can infer details like the number of devices controlled by the smart speaker, the presence of the user (even when the interaction is absent), and the “kind” of the issued commands [8, 1, 3, 10, 16]. Moreover, a relevant part of the traffic produced by smart speakers shares functional and technological traits with VoIP, meaning that it is also susceptible to attacks for disclosing the language of the talker or other sensitive behaviors [23].

Therefore, this paper focuses on investigating deep adversarial learning countermeasures against machine learning attacks targeting the traffic produced by smart speakers. To the best of our knowledge, this aspect has been mostly overlooked so far. The only notable exception is [17], which proposes a padding scheme to protect IoT and smart devices from statistical analysis. Instead, our work aims to showcase the limitations of traffic manipulation or morphing approaches, which often lead to flawed countermeasures [13].

To do so, we built a new dataset containing the network traffic of a typical smart home environment and an experimental testbed to evaluate the efficacy of deep adversarial learning techniques. The experimental activities exploit both theoretical approaches, i.e., the usage of Savitzky–Golay filters and Additive White Gaussian Noise (AWGN) on all the statistical features, and realistic ones, named “*Realistic Adversarial*”, that use constant padding and AWGN techniques that considers the constraints of the protocols and networks in use. The achieved results argue that smart speaker privacy needs a complete rethink.

The rest of the paper is structured as follows: Section 2 provides the background and the threat model, Section 3 presents the deep adversarial techniques used in this work. Section 4 describes the evaluation testbed and Section 5 presents the obtained results. Finally, Section 6 concludes the paper.

2 Background and Attack Model

Figure 1 depicts a typical smart speaker ecosystem that provides the voice-activated user interface and acts as a hub for other IoT nodes and network appliances. In essence, the speaker collects, samples, and transmits voice commands to remote cloud services in charge of processing data to deliver back textual/binary representations as well as additional content, e.g., multimedia streams. The smart speaker can also provide feedback to the user, play content, retrieve data from third-party providers (e.g., music streaming services) or drive other nodes via local network or short-range links like IEEE 802.15.4. Even if the information gathered by commands exchanged locally between the speaker and various nodes can be used to threaten the privacy of the ecosystem [15]. This paper focuses on the attacks exploiting the network traffic exchanged between the speaker and its remote cloud.

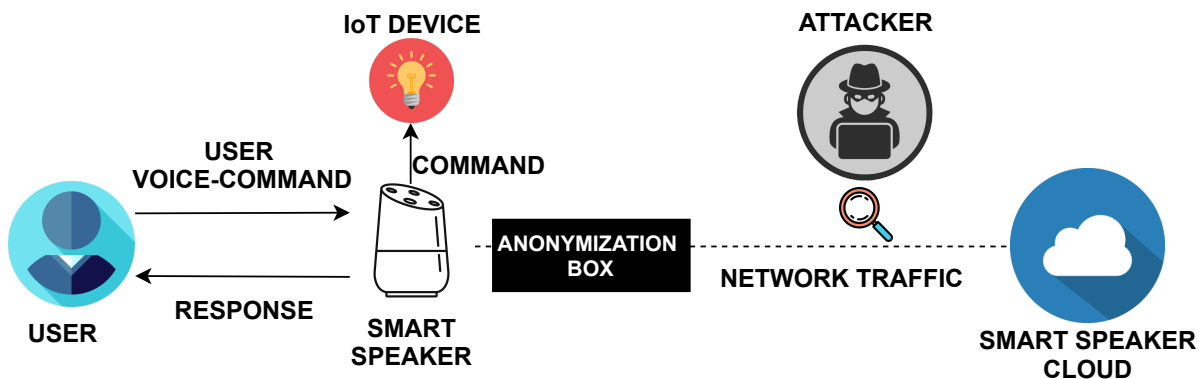


Figure 1: MITM threat model of a smart home ecosystem.

From a security perspective, the continuous exchange of data between the smart speaker and the cloud is a prime point of fragility. As depicted in Figure 1, an adversary (denoted as *attacker*) can mount MITM (Man-in-the-Middle) attacks [4] to gather network traffic even in the case of a communication encrypted with TLS/SSL [9].

Even if many commercial smart speakers implement countermeasures to protect the network traffic, the majority is still prone to a variety of privacy-breaking attacks targeting a composite set of features observable within the encrypted traffic flows [2, 25]. Specifically, we focus on an attacker willing to use machine learning or deep learning algorithms on encrypted traffic samples to infer “behavioral” information, e.g., the presence of the victim or the “type” of the requested information [8, 1, 3, 10, 16, 7]. Owing to the end-to-end encryption, the attacker can only observe and acquire the traffic produced by the smart speaker and cannot alter, manipulate or perform deep packet inspection operations. The attacker can then only rely on general statistics, e.g., the throughput, the size of protocol data units, IP addresses, the number of different endpoints, flags within the headers of the packets, or the behavior of the congestion control of the TCP [8].

The standard approach to mitigate machine learning attacks on the network traffic exploits the use of a middlebox (denoted as *anonymization box* in Figure 1) able to “sanitize” the network traffic by removing (or altering) the data that the attacker can exploit. For instance, the anonymization box can pad packets [17] or perform NAT-like operations to prevent profiling endpoints or probing [12]. All in all, since the anonymization box is outside the device, it cannot alter the protocol/communication architecture of the smart speaker ecosystem. Rather, it can only manipulate the traffic without disrupting the flow or penalizing the QoE perceived by the user, for instance, in terms of real-time guarantees (see, e.g., [11] and references therein). In the following, we will showcase the limits of such an approach, which appears to be unsuited to face modern machine learning-capable threats.

3 Deep Adversarial Learning Techniques

This paper investigates both theoretical and practical deep adversarial learning techniques to lower the classification accuracy. Ideally, we would like to reduce the classification accuracy to be as close as possible to a “coin toss” (e.g., 50% on a two-class classification problem). As in previous works [8, 21, 19], the attacker can only acquire encrypted traffic to compute statistical metrics and analyze them using machine learning techniques. Such a computation requires using a suitable number of packets grouped using either time spans of length Δt or bursts of a fixed size of N packets.

The methods considered for this work are:

- smoothing of features through Savitzky–Golay filter;
- injecting Additive White Gaussian Noise (AWGN) into the features time series;
- applying a *Realistic Adversarial*, i.e., a targeted approach to feature degradation that also considers the constraints of the protocols and networks in use.

The first two approaches aim to show the theoretical performances that could be obtained by randomizing, without any constraints, the statistics of the packets and the features derived from them. In detail, the Savitzky–Golay filter [18] allows smoothing all features through a moving window. The length of the filter window determines the number of samples taken into consideration, and a polynomial of user-defined order subsequently approximates these samples.

The second adversarial technique uses Additive White Gaussian Noise (AWGN) – therefore with zero mean – whose variance is set proportionally to the variance of the original signal subject to adversarial.

The Realistic Adversarial, instead, represents an approach that takes into account which features can actually be distorted at the egress of the IoT device (therefore with external hardware such as an *anonymization box*) without compromising its operation (e.g., it is possible to add padding to the packets, but it is not possible to randomize the TCP window without disrupting the service completely). In this work, we exploited two Realistic Adversarial techniques: the *first* is a **constant padding** within the whole time series, which was simulated selecting the maximum value of the mean TCP packet length (*mean_len_pack*). The standard deviation of the TCP packet length (*std_len_pack*) was set to zero to match this operation. The *second* technique concerns the injection of **AWGN** in the following features:

- *std_ipt*: to simulate jitter while sending packets;
- *n_pack_tcp* and *n_pack_udp*: to simulate decoy connections and packets between endpoints;
- *n_pack_icmp*: to simulate decoy ping/traceroute packets between endpoints;
- *n_port_unique*:, to simulate decoy TCP/UDP packets addressed to random port numbers.

Moreover, to match the realistic scenario, we cannot modify the following features: *max_diff_time*, *n_ip_unique*, *mean_window* and *std_window*.

Feature Name	Adversarial Technique
Number of different IP address (<i>n_ip_unique</i>)	None
Number of different TCP/UDP ports (<i>n_port_unique</i>)	AWGN
Number of TCP packets (<i>n_pack_udp</i>)	AWGN
Number of UDP packets (<i>n_pack_udp</i>)	AWGN
Number of ICMP packets (<i>n_pack_icmp</i>)	AWGN
<i>Per-window</i> Inter packet time (<i>max_diff_time</i>)	None
Average of TCP Window (<i>mean_window</i>)	None
Standard deviation of TCP Window (<i>std_window</i>)	None
Average of IPT (<i>mean_ipt</i>)	None
Standard deviation of IPT (<i>std_ipt</i>)	AWGN
Average of packet length (<i>mean_len_pack</i>)	Const. Padding
Standard Deviation of packet length (<i>std_len_pack</i>)	Set to Zero

Table 1: Name and Acronym of statistical indicator used and the adversarial techniques applied to them.

4 Experimental Testbed

We developed an experimental testbed to prove the effectiveness of privacy threats and the inadequacy of countermeasures applied to smart speakers’ network traffic. Briefly, we collected a dataset with the network traffic of the IoT device during typical usage scenarios, e.g., during voice queries or media playback. Then, we used a set of deep adversarial learning techniques to alter the statistical information of outbound traffic. Finally, we applied a set of ML techniques to evaluate the corresponding degradation of the classification process.

The experimental activities were carried out on an Intel Core i7-3770 computer equipped with 16 GB of RAM and Ubuntu 16.04 LTS, and a Google Home Mini smart speaker. The traffic has been captured via an instrumented computer acting as an IEEE 802.11 access point running ad-hoc scripts for *tshark*¹. All the generated traffic traces have been anonymized and made available through Kaggle².

4.1 Dataset Definition

In this paper, we extend the dataset already published in our previous work [8]. Briefly, the original dataset consists of 9-days of network traffic that comprise: i) traffic with the microphone disabled (*Mic On-Off*), ii) microphone enabled in a quiet environment, and iii) microphone enabled with background noise (*Mic On-Noise*).

Thus, to mimic the normal use of smart speakers by users, we extended the available data with three different classes of queries for the smart speaker, i.e., media, travel, and utility³. To do so, we executed three different rounds of measurements that last three days each. In essence, a synthetic talker has been created by using various voice records representing a wide range of speakers (e.g., male and female or with different accents or talking speeds) and it has been used to issue commands to the smart speaker. In the first round, we focused on retrieving the network traffic generated to playback multimedia content. For example, we captured traffic when the synthetic talker asked questions like “*What’s the latest news?*” or “*Play some music*”. For the second round, we performed queries related to travels, thus accounting for the interaction with services providing traffic indications or weather forecasts. In this case, we asked questions like “*How is the weather today?*”. Lastly, we performed general queries belonging to the utility category, like “*What’s on my agenda today?*” and “*What time is it?*”.

For each query, we collected 30 seconds of inbound and outbound network traffic to have a proper tradeoff between accuracy and size of the data. The newly collected data contains 2,500k packets for media queries, 400k packets for travel queries, and 310k for utility queries.

4.2 Deep Adversarial Techniques Setup

We implemented the three deep adversarial learning techniques presented in Sect. 3. In detail, we set the Savitzky–Golay filter moving window to 51, and we applied different polynomials degrees, i.e., $\psi = \{1, 3, 5, 7, 9\}$. For the Additive White Gaussian Noise (AWGN) techniques, we multiplied the original variance using a constant $v = [0.2, 2.0]$ for the *Mic On-Off* and *Mic On-Noise* scenario and $v = \{2, 8, 16, 32, 64\}$ for the *utility/media/travel* scenario.

Table 1 lists the set of network features computed for the experimental activity of this paper and the corresponding employed deep adversarial technique.

4.3 ML techniques

To implement the attacker, we used machine learning algorithms provided by the *scikit-learn* and *Fast.ai* libraries. We considered the most popular techniques commonly used in the literature, i.e., AdaBoost (AB), Decision Tree (DT), k-Nearest Neighbors (kNN), Random Forest (RF) and Neural Networks (NN) [8, 14, 24]. In our experiment, we assume that the attacker is able to collect 3 days of traffic for each scenario, similar to [8]. The classifiers were neither pre-trained with the original data nor fed with previously trained models.

¹<https://www.wireshark.org/docs/man-pages/tshark.html>

²<https://www.kaggle.com/smartspeaker/google-home-pcap>

³<https://voicebot.ai/2019/03/12/smart-speaker-owners-agree-that-questions-music-and-weather-are-killer-apps-what-comes-next/>

5 Experimental Results

In this section, we show and discuss the numerical results obtained. First, we show the performance and accuracy of machine learning algorithms used to infer the query category. Next, we discuss the efficacy of the deep adversarial learning techniques to protect the network traffic.

5.1 Query Classification

Figure 2 shows the classification accuracy for ML techniques trained on the original data (not subjected to any adversarial technique) of the *Utility/Media/Travel* scenario. With sampling and feature generation intervals higher than 500 packets, all the considered algorithms achieve a classification accuracy higher than 90%, being the kNN algorithm the less precise. The accuracy of the neural network model (i.e., NN) also tends to drop slightly as the interval increases because the number of samples decreases too much for the training “from scratch” of a neural network (at 1500 packets, the number of samples is just above 300 units for this scenario). Such behavior is also consistent with the classification of the *Mic On-Off* and *Mic On-Noise* scenarios whose accuracy is higher than 80% (as described in [8]).

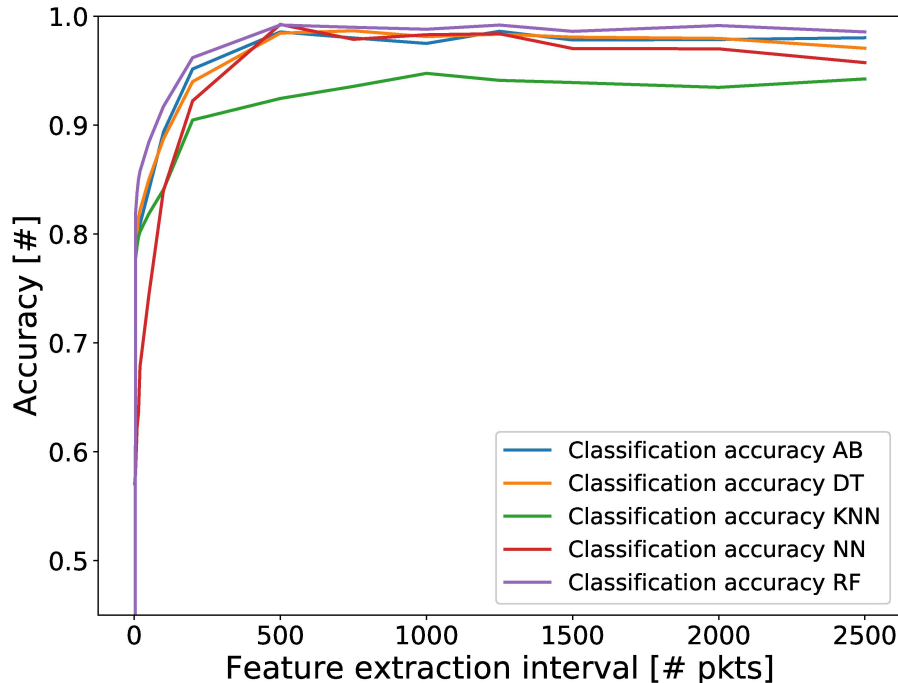


Figure 2: *Utility/Media/Travel* scenario: classification accuracy of the different machine learning/deep learning models using unmodified original features.

5.2 Mitigation Results

The set of images in Figure 3 - 8 and Figure 9 - 11 depict the accuracy of a classifier trained on data subjected to the three deep adversarial techniques.

In detail, Figure 3 to 5 show the accuracy obtained in the *Mic On-Off* classification scenario, Figure 6 to 8 refer to the *Mic On-Noise* scenario, and Figure 9 - 11 show the classification performance of a classifier trained on the traffic features from the *Utility/Media/Travel* scenario.

In the first two scenarios, the trend is similar: the theoretically most effective adversarial technique is

the polynomial smoothing as it deprives the signal of most of its high-frequency components. On the contrary, AWGN injection is less effective as it leaves much of the information untouched, especially with low values of ν (i.e. variance multiplier values). Figure 5 and 8, conversely, show the implementation of a realistic adversarial technique which is therefore not able to degrade all the features simultaneously. As can be seen from the accuracy levels (for some sampling intervals > 1250 packets, higher than 90%), the neural network model is sufficiently “intelligent” to correctly predict the class using just the few remaining features and ignoring all the others. In this case, higher AWGN intensity does not appear to affect the prediction in any way.

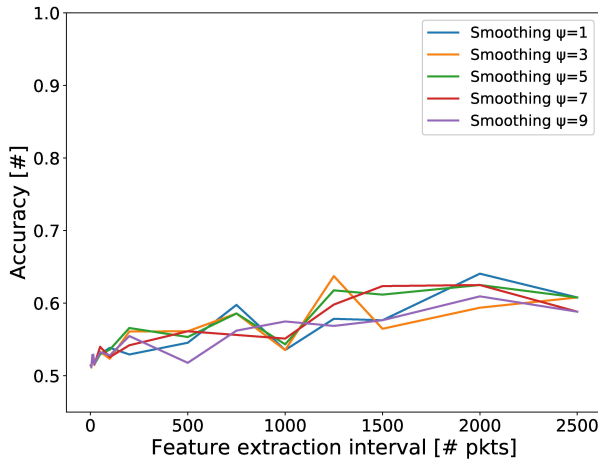


Figure 3: Mic On-Off Adversarial with smoothing (*Savitzky-Golay*).

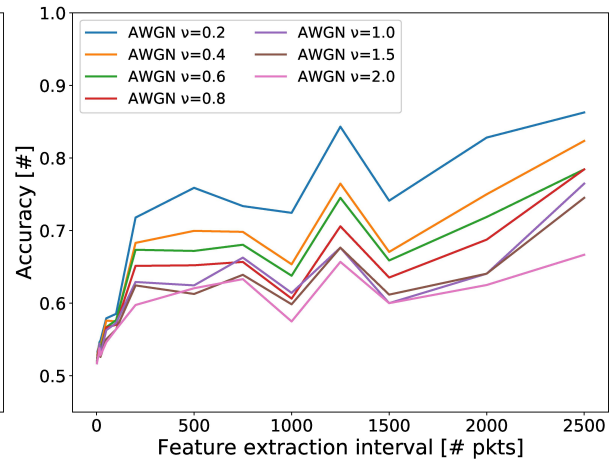


Figure 4: Mic On-Off Adversarial with noise (AWGN).




Figure 5: Mic On-Off Realistic adversarial technique.

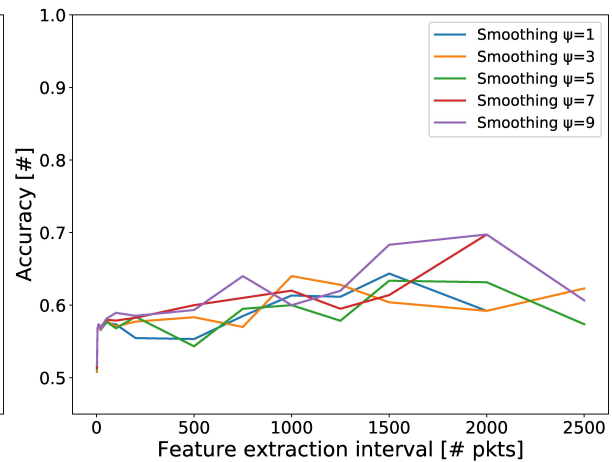


Figure 6: Mic On-Noise Adversarial with smoothing (*Savitzky-Golay*).

The analysis of the queries scenario in Figure 9 - 11 confirms that the polynomial smoothing techniques are the most efficient, bringing the accuracy of the classifier to around 40%, very close to the theoretical level of a 33% random choice of a problem with three classes (cf. Figure 9).

On the contrary, Figure 10 highlights the need for a substantially higher ν value for the AWGN in order to sufficiently degrade the features for this scenario. Indeed, unlike the *Mic On-Off* and *Mic On-*




Figure 7: figure
Mic On-Noise Adversarial
with noise (AWGN).

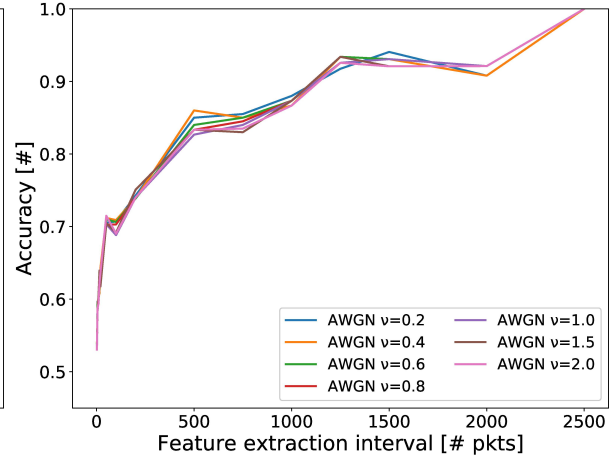


Figure 8: figure
Mic On-Noise Realistic
adversarial technique.

Noise scenarios (with $v = 2.0$), this scenario requires a variance multiplier $v = 64$ to bring the accuracy level of the ML classifiers to values close to the 33%.

Finally, Figure 11 shows the accuracy obtained from the same neural network architecture trained on degraded features with realistic adversarial. As in Figure 5 and 8, the impossibility of degrading all the features at once, leaving some of them intact, preserves enough information for the neural network to learn how to correctly identify the three classes of the problem, even with an accuracy of 95% considering an interval of 1000 packets, independently from the magnitude of the injected AWGN.

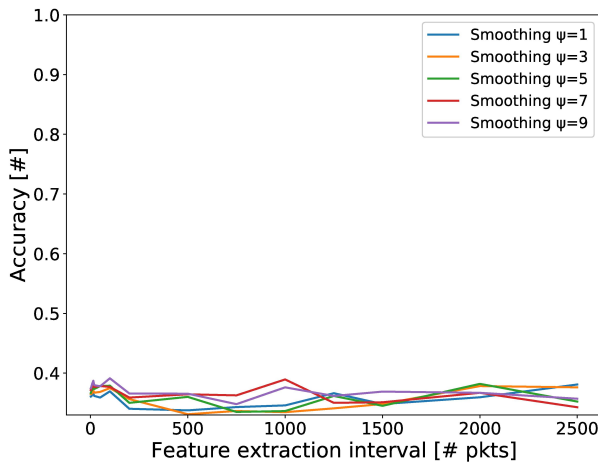


Figure 9: *Utility / Media / Travel* Adversarial
with smoothing.




Figure 10: *Utility / Media / Travel* Adversarial
with noise.

6 Conclusions

In this paper, we have empirically demonstrated how adversarial learning countermeasures applied to the virtual assistant’s outbound traffic are ineffective against machine learning attacks, thus leading to




Figure 11: *Utility / Media / Travel* Realistic adversarial

serious concerns for the privacy of users in smart home environments.

The results indicate the need for a major HW/SW redesign of virtual assistant platforms to ensure adequate protection of commercially available smart speakers, such as the introduction of trusted computing platforms [6, 22] to anonymize the outbound traffic or the usage of misbehavior detection [20].

References

- [1] A. Acar, H. Fereidooni, T. Abera, A. K. Sikder, M. Miettinen, H. Aksu, M. Conti, A.-R. Sadeghi, and S. Uluagac. Peek-a-boo: I see your smart home activities, even encrypted! In *Proc. of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec'20)*, Linz, Austria, pages 207–218. ACM, July 2020.
- [2] E. Alepis and C. Patsakis. Monkey says, monkey does: Security and privacy on voice assistants. *IEEE Access*, 5:17841–17851, August 2017.
- [3] Y. Amar, H. Haddadi, R. Mortier, A. Brown, J. Colley, and A. Crabtree. An analysis of home iot network traffic and behaviour. *arXiv preprint arXiv:1803.05368*, March 2018.
- [4] S. Andy, B. Rahardjo, and B. Hanindhito. Attack scenarios and security analysis of mqtt communication protocol in iot system. In *Proc. of the 4th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI'17)*, Yogyakarta, Indonesia, pages 1–6. IEEE, September 2017.
- [5] N. Apthorpe, D. Reisman, and N. Feamster. A smart home is no castle: Privacy vulnerabilities of encrypted iot traffic. *arXiv preprint arXiv:1705.06805*, May 2017.
- [6] A. Armando, A. Merlo, and L. Verderame. Trusted host-based card emulation. In *Proc. of the 2015 International Conference on High Performance Computing & Simulation (HPCS'15)*, Amsterdam, Netherlands, pages 221–228. IEEE, July 2015.
- [7] D. Caputo, L. Verderame, A. Merlo, A. Ranieri, and L. Caviglione. Are you (google) home? detecting users' presence through traffic analysis of smart speakers. In *Proc. of the 4th Italian Conference on Cyber Security (ITASEC'20)*, Ancona, Italy, pages 105–118. CEUR-WS.org, February 2020.
- [8] D. Caputo, L. Verderame, A. Ranieri, A. Merlo, and L. Caviglione. Fine-hearing google home: why silence will not protect your privacy. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 11(1):35–53, March 2020.
- [9] M. Conti, N. Dragoni, and V. Lesyk. A survey of man in the middle attacks. *IEEE Communications Surveys & Tutorials*, 18(3):2027–2051, March 2016.

- [10] B. Copos, K. Levitt, M. Bishop, and J. Rowe. Is anybody home? inferring activity from smart home network traffic. In *Proc. of the 37th IEEE Symposium on Security and Privacy Workshops (SPW'16)*, San Jose, California, USA, pages 245–251. IEEE, May 2016.
 - [11] J. Fan, C. Guan, K. Ren, Y. Cui, and C. Qiao. Spabox: Safeguarding privacy during deep packet inspection at a middlebox. *IEEE/ACM Transactions on Networking*, 25(6):3753–3766, October 2017.
 - [12] M. Gregorczyk, P. Żórawski, P. Nowakowski, K. Cabaj, and W. Mazurczyk. Sniffing detection based on network traffic probing and machine learning. *IEEE Access*, 8:149255–149269, August 2020.
 - [13] A. Houmansadr, C. Brubaker, and V. Shmatikov. The parrot is dead: Observing unobservable network communications. In *Proc. of the 34th IEEE Symposium on Security and Privacy (S&P'13)*, San Francisco, California, USA, pages 65–79. IEEE, May 2013.
 - [14] Z. Li, R. Yuan, and X. Guan. Accurate classification of the internet traffic based on the svm method. In *Proc. of the 2007 IEEE International Conference on Communications (ICC'07)*, Glasgow, UK, pages 1373–1378. IEEE, June 2007.
 - [15] B. Nour, K. Sharif, F. Li, and Y. Wang. Security and privacy challenges in information-centric wireless internet of things networks. *IEEE Security & Privacy*, 18(2):35–45, August 2019.
 - [16] A. J. Pinheiro, J. d. M. Bezerra, C. A. Burgardt, and D. R. Campelo. Identifying iot devices and events based on packet length from encrypted traffic. *Computer Communications*, 144:8–17, August 2019.
 - [17] A. J. Pinheiro, P. F. de Araujo-Filho, J. d. M. Bezerra, and D. R. Campelo. Adaptive packet padding approach for smart home networks: A trade-off between privacy and performance. *IEEE Internet of Things Journal*, 8(5):3930–3938, September 2021.
 - [18] A. Savitzky and M. J. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8):1627–1639, July 1964.
 - [19] M. R. Shahid, G. Blanc, Z. Zhang, and H. Debar. IoT Devices Recognition Through Network Traffic Analysis. In *Proc. of the 2018 IEEE International Conference on Big Data (Big Data'18)*, Seattle, Washington, USA, pages 5187–5192. IEEE, December 2018.
 - [20] V. Sharma, I. You, K. Yim, I.-R. Chen, and J.-H. Cho. Briot: Behavior rule specification-based misbehavior detection for iot-embedded cyber-physical systems. *IEEE Access*, 7:118556–118580, May 2019.
 - [21] D. Su, J. Liu, S. Zhu, X. Wang, and W. Wang. "are you home alone?" "yes" disclosing security and privacy vulnerabilities in alexa skills. *arXiv preprint arXiv:2010.10788*, October 2020.
 - [22] J. Winter. Trusted computing building blocks for embedded linux-based arm trustzone platforms. In *Proc of the 3rd ACM workshop on Scalable trusted computing (STC'08)*, Alexandria, Virginia, USA, pages 21–30. ACM, October 2008.
 - [23] C. V. Wright, L. Ballard, F. Monrose, and G. M. Masson. Language identification of encrypted voip traffic: Alejandra y roberto or alice and bob? In *Proc. of the 16th USENIX Security Symposium (USENIX Security'07)*, Austin, Texas, USA, pages 43–54, August 2007.
 - [24] A.-m. Yang, S.-y. Jiang, and H. Deng. A p2p network traffic classification method using svm. In *Proc. of the 9th International Conference for Young Computer Scientists (ICYCS'08)*, Hunan, China, pages 398–403. IEEE, November 2008.
 - [25] Y. Yang, L. Wu, G. Yin, L. Li, and H. Zhao. A survey on security and privacy issues in internet-of-things. *IEEE Internet of Things Journal*, 4(5):1250–1258, October 2017.
-

Author Biography



Andrea Ranieri is a researcher at CNR-IMATI where he works with Deep Learning mainly in the field of Computer Vision and Time Series Forecasting. Andrea has a degree in Computer Engineering, a Ph.D. in Space Science and Engineering and an extremely varied technological background ranging from networking and distributed systems to robotics and perception. He has participated in many Italian and European research projects (FP6 INTERMEDIA-NoE, PRIN Italian project MARIS, FP7 MORPH, FP7 CADDY and others) and has co-authored many scientific publications in conference proceedings and international journals. His research interests range from the simple classification of images and time series, to unsupervised pre-learning tasks for transfer learning to supervised tasks.



Davide Caputo is a third-year Ph.D. student in Computer Science at the University of Genova, Genova, Italy. He obtained both his BSc and MSc in Computer Engineering at the University of Genova, and now he is working under the supervision of Prof. Alessio Merlo and Dott. Luca Verderame. His research topic focuses on Mobile Security and IoT Security.



Luca Verderame obtained his Ph.D. in Electronic, Information, Robotics, and Telecommunication Engineering at the University of Genoa (Italy) in 2016, where he worked on mobile security. He is currently working as a post-doc research fellow at the Computer Security Laboratory (CSEC Lab), and he is also the CEO and Co-founder of Talos, a cybersecurity startup and university spin-off. His research interests mainly cover information security applied, in particular, to mobile and IoT environments.



Alessio Merlo is an Associate Professor in Computer Engineering in the Department of Informatics, Bioengineering, Robotics and System Engineering Department (DIB-RIS) at the University of Genoa, and a member of the Computer Security Laboratory (CSEC Lab). His main research field is Mobile Security, with a specific interest on Android security, automated static and dynamic analysis of Android apps, mobile authentication and mobile malware.

More information can be found at: http://csec.it/people/alessio_merlo/.



Luca Caviglione received the Ph.D. degree in electronics and computer engineering from the University of Genoa, Genoa, Italy. He is a Research Scientist with the Institute for Applied Mathematics and Information Technologies of the National Research Council of Italy, Genoa. His research interests include optimization of large-scale computing frameworks, wireless and heterogeneous communication architectures, and network security. He is an author or co-author of more than 140 academic publications and several patents in the field of p2p and energy-aware computing. He has been involved in many research projects funded by the European Space Agency, the European Union, and the Italian Ministry of Research. He is a Work Group Leader of the Italian IPv6 Task Force, a contract professor in the field of networking/security and a professional engineer. He is a member of the technical program committee of many international conferences. From 2016 he is an Associate Editor of the International Journal of Computers and Applications, Taylor&Francis