International Journal of Development and Conflict
12(2022) 182–191

THE PREDICAMENT OF PRIVACY AND SIDE-CHANNEL ATTACKS

ALI AL MOUSA, MOHAMMAD AL QOMRI, MUHAMMAD IMAM
College of Computing and Mathematics
Computer Engineering Department
King Fahd University of Petroleum and Minerals
P.O. Box 5065, Dhahran 31261, Saudi Arabia

Attackers and their methodologies are getting creative. With the aid of growing technologies such as machine learning, meaningful information can be deduced from very little data. We show that attackers are able to utilize such promising technologies to perform side-channel attacks to obtain data that can identify individuals, their habits and interests, health information, and much more. Furthermore, we highlight that such data can be obtained in a legal and consensual context, making it very hard to determine if an individual's privacy has been breached or not.

Keywords: Privacy, Side-Channel Attacks, Machine Learning, Deep Learning, Invasion

Introduction

The topic of privacy has caught a lot of attention in the last several years. This is not a surprise as new technologies such as Machine Learning (ML) have made it easier for our privacy to be violated. Such technologies allow the processor of data to obtain so much information about the subject with minimal input when compared with the olden days. Given how capable these technologies are, it would not be a surprise to learn that they can be (and have been) misused to invade the privacy of individuals. Activists, governments, and independent organizations have identified and addressed this issue and are currently regulating how data should be collected, processed, and disposed of.

However, privacy is a non-trivial goal. It cannot be achieved solely by enforcing regulations and legislation. One of the fundamental challenges is that there is not a single and agreed-upon definition of privacy. Because of that, it can be challenging to assess whether an individual's privacy has been compromised or not, simply because we do not know where to draw the line.

The problem is even more prominent when the claimed privacy breach is based on physical attributes. For example, it can be reasonably asserted that a given person's privacy has been breached when the bank analyzed his purchasing data to infer that he is getting married, then used that data to offer a tailored mortgage proposal. On the other hand, it

*Corresponding author Email: @ kfupm.edu.sa

can be much harder to make that assertion when the bank teller sees the same person has started to wear a wedding ring and then offers the mortgage proposal accordingly.

Background

As technology evolves, attackers are getting more creative in their techniques and tactics. One of such is what is called side-channel attacks. In which attacks are not targeting the unexpected bugs in the algorithm. Rather, they are exploiting the implementation weaknesses of that algorithm. It can also include attacks on the physical attributes of the implementation. This type of attack varies greatly in terms of complexity. For example, an attacker could perform a side-channel attack simply by observing the reflections on the victim's glasses when he or she is typing a password on a mobile phone. In this case, we cannot blame the manufacturer of the glasses for this "vulnerability" in their product because it is not being exploited directly. In a more complex example, a capable attacker could perform a power analysis of the CPU to decode RSA keys [1]. By reading the high and low values of the CPU power consumption, the attack is able to translate it into bits.

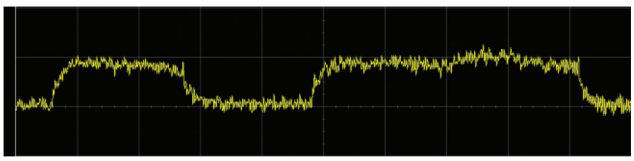

Figure 1. Power analysis on CPU to decode RSA keys [1]

Literature Review

With seemingly dreary data, researchers are able to deduce significant conclusions. Although such data is easily accessible in any given physical environment, it does not create any privacy concerns as comprehending such data is beyond any average human capacity. With that being said, and with the aid of machine learning (ML), researchers are able to consistently infer remarkable conclusions. We explore how such technologies are able to identify individuals, obtain personal information, protected health information (PHI), recovering typed text and acoustic conversations.

A. Identification and Personal Information

It is not a surprise to know that each person has a unique data footprint on the Internet. The way they view, interact with, and create data can be used to deduce personal information about them. As a matter of fact, such data can and have been used to incriminate an individual. This was done fifty-five years ago by Jan Svartvik when he published "The Evans Statements: A Case For Forensic Linguistics" [2]. In which he demonstrated that the grammatical style of an author could be forensically analyzed for identification purposes, marking the birth of a new science—forensic linguistics.

Perhaps one of the most prevalent examples of such practice is the Unabomber case. From 1978 to 1995, a terrorist by the name of Theodore Kaczynski had launched a campaign of bombing attacks targeting universities and airline companies, hence the name

Unabomber [3]. Accompanied by the bombs, he also sent notes and letters. In addition, he had also published a 35,000-word manifesto titled "Industrial Society and its Future" [4]. By analyzing the notes, letters, and manifesto, an FBI agent by the name of Roger Shuy was able to correctly ascertain demographic information about the terrorist, such as his age, origins, residency, education level, specialization, and religious beliefs [5].

In a slightly different yet similar case, literature such as [6] has illustrated the possibility of identifying an individual based on their typing patterns, also known as keystroke dynamics. By having a baseline of the legit user's keystroke timing patterns, a ML model will be able to differentiate a malicious user from another with decent accuracy. The inception of this field was marked in the early 1980s when the National Science Foundation and the National Burau of Standards in the United States of America published studies stating that data in typing patter can be considered a unique biometric to identify an individual [7]. With that being said, early evidence suggested that the use of such data for identification has been linked back to the World War II era. A methodology known as "First of the Sender" was used by the soldiers to identify telegraph operators on each end according to their typing rhythm [8].

Other researchers, such as [9], were able to deduce demographic information by analyzing the voice. By utilizing deep neural networks (DNN), [9] have classified the age of speakers from their voice, with an accuracy exceeding 77%. On a similar note, [10] [11] [12] have demonstrated that the ethnicity of the speaker can also be reliably identified by analyzing his or her voice, with accuracy going as high as 97%.

B. Recovering Protected Health Information

Researchers from Universidad Politécnica de Madrid, MIT, and Johns Hopkins University published a study to diagnose Parkinson's disease by utilizing speech recognition techniques [13]. While currently there are no biomarkers to diagnose such disease definitively, physicians typically combine two to three symptoms for the diagnosis. Such symptoms include but are not limited to tremors, stiffness, and balance problems. With that being said, the researchers at [13] were able to diagnose Parkinson's patients reliably with an accuracy of 87%. On a similar note, researchers in [14] have managed to achieve an outstanding 99% accuracy for the same diagnosis. The researchers in [14] did not need more than a ten-second audio clip of subjects saying "Aaaah" to achieve such high accuracy.

Another study by [15] was conducted to evaluate whether there is a measurable correlation between the quality of the voice and obesity. [15] concluded that there is indeed a statistically significant difference in the maximum phonation time between a population with different body mass index (BMI). Many reasons could cause this measurable difference, but perhaps the most likely factors are the increased tissue bulk in the chest and neck, as suggested by [16].

More complex models and algorithms have also been developed by other literature to tackle even more challenging problems. For example, [17] used a convolutional neural network (CNN) to diagnose respiratory-related diseases. With accuracy as high as 89%, [17] is able to reliably diagnose bronchitis, bronchiolitis, and pertussis based on the acoustic features of the patients' coughs.

C. Recovering Typed Text

The mentioned literature demonstrates that acoustic side-channel attacks are a growing field with promising potential. To add to that, [18] demonstrated that by analyzing the acoustic features of a dot-matrix printer sound, an attacker could recover up to 95% of the printed text. Furthermore, [18] noted that 60% of doctors in Germany are using such printers to print patients' information. In addition, 30% of banks in Germany are also using the same type of printers to print bank statements, transcripts, and other sensitive information.

Similar methodologies have been proposed to utilize acoustic side-channel attacks to recover typed text from keyboards. Using standard ML techniques combined with speech recognition, [19] [20] [21] were able to recover the typed text by extracting features from the sound of each keystroke emitted from the keyboard. Figure 2 shows a sample signal spectrum of "q" and "w" keystrokes. The extracted text could include sensitive information about the person or login credentials. This approach has proven to be reliable as it reaches an accuracy of up to 96%.

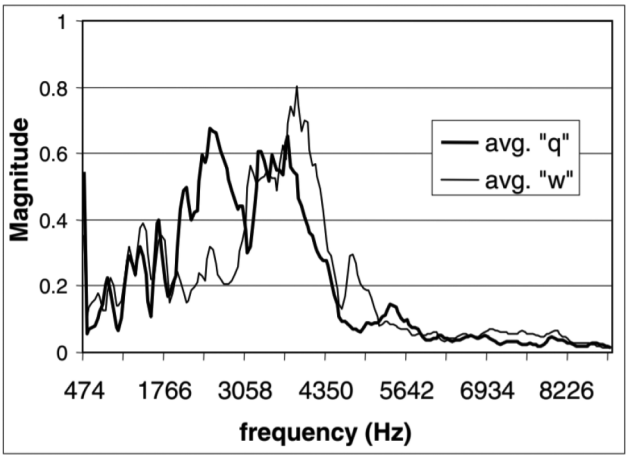

Figure 2. Average signal spectrum of "q" and "w" [21]

Other literature is targeting the same information but with a different methodology. Instead of relying on the sound signature of each keystroke, [22] have suggested using the vibration captured by an accelerometer. Because each keystroke generates a unique vibration based on its physical position, a ML model will be able to analyze the data and retrieve what keys have been pressed. By doing so, [22] have reached accuracies as high as 80%.

On a more advanced approach, [23] was able to retrieve and recover keyboard strokes remotely by utilizing a laser microphone. Depending on the quality of the laser, the distance of the attack can be massive. The only requirement for this attack is a line of sight on the victim's machine. From there, the laser will be able to pick up the vibrations caused by the keyboard from the laptop's screen. What makes this attack so powerful is the fact that the victim and the attacker could be in two different buildings and separated by a great distance. In addition, the attack leaves little to no measurable footprints, making it extremely hard to detect and prevent.

An even more impressive method to recover typed text was proposed by [27], called WiKey. As the name suggests, WiKey exploits Wi-Fi signals to recognize and identify keystrokes. As each keystroke is performed by a unique hand formation and direction, a unique pattern of Channel State Information (CSI) time series is generated. By analyzing the CSI value of each keystroke, WiKey is able to recover typed sentences with accuracies exceeding 97%

Another unusual approach covered by [29] shows how a user can type his PIN password into any touchable device. The mechanism is simple, once a user touches the screen to insert the number, the haptic feedback can emit from the device traveling through the user's finger and moving all the way to his muscles and bones and ending up being detected by a wearable device. Such detection is possible due to the fact most wearable accessories can have an embedded accelerometer that can sense the vibration and locate where the touches were on the screen of the device. [29] also highlighted another aspect where an attacker could use it to his advantage. Such as utilizing a microphone that can be used to retrieve the sound of clicking noise produced by the touchable device. Thus, resulting in revealing the PIN code for that specific user by utilizing ML to reconstruct the original PIN code. This approach is very similar to what was presented by the [22] and [24] in their papers.

D. Recovering Conversations

Other techniques utilizing ML have also been used to retrieve information that would otherwise be very difficult for an average human to do. An example of such utilization is demonstrated by [24], where a ML model is developed and trained to retrieve and recover audio conversation. By watching a video of an individual speaking (without sound), the model is able to read the individual's lips to determine what words have been spoken, with accuracies reaching 92%.

Another approach with a similar goal is demonstrated by [25], where they demonstrated that audio conversation could be recovered and reconstructed by observing the vibrations from nearby objects. An example of such objects is a glass of water, a bag of chips, or a box of tissues. By using a high-speed camera, the attacker is able to measure the vibration on the object caused by the speaker, even if the speaker is not visible. What makes this attack so powerful is the fact that there is no need for a line of sight with the victim. The victim could be speaking in a different room, and the attack is possible. Unlike other attacks discussed previously, this one is entirely passive, making it nearly impossible to detect.

Similarly, [26] have used remote mechanisms to detect and "hear" human conversations. Instead of relying on vibration, [26] utilizes Wi-Fi signals to detect and analyze micromovement generated by the lips during conversations by utilizing a ML model. Figure 3 shows a sample of lips micromovement recognized by Wi-Fi. As a result, [26] were able to reach an accuracy of 91%.

Another paper presented by [30] show's how an attacker can utilize a ML model to reconstruct the actual text and exact words that have been said. Such an approach allows injecting any video without a sound and retrieving all the sensitive information. Thus, it was possible to break the video down into pictures and then analyze it frame by frame.

Other techniques that have been used are presented by [31]; the researcher shows that utilizing the time reaction for responding to a specific question can expose the identity

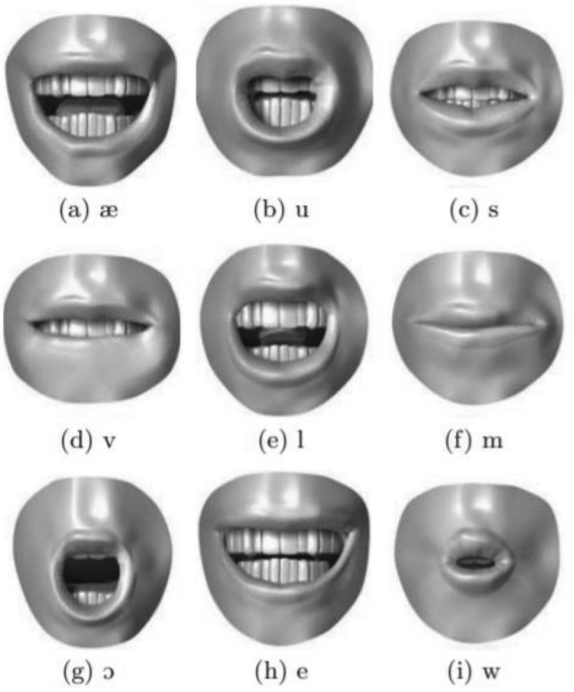

Figure 3. Vowels and consonants recognized by Wi-Fi [26]

of a lier to prove that he's not the person who claims. Doing so has proven to help detect and separate a truth-teller from a lier. Thus, the mix of reaction time and the unexpected questions led to success in the accuracy rate. The research pointed out that the giving ML model can detect lair if used as mentioned with a success rate of 90%, where the metrics were the speed and the accuracy. Some of the questions used were regarding the actual personal information such as age, date of birth, and zodiac sign. Thus, selecting a random zodiac sign with the wrong date exposed the claims of a person. Therefore, planning for such a lie requires initial preparation by the claimer to prove his point. So, knowing that the researcher believes that he can rely on these metrics to be solid evidence for his approach. Also, [30] mentioned that responding to personal information when engaged in typing the information took more time than usual.

Table 1. Summary of Literature Review

Input	Retrieved Information	Technique	References
	The age	You can utilize the sound to estimate how old the person is based on his vocals	[9]
	The weight	By analyzing the person's voice, you can estimate his BMI based on the maximum phonation time	[15] [16]
Acoustics	Diagnosis of respiratory diseases	Using CNNs to analyze the acoustic features of the patients'coughs.	[17]

Input	Retrieved Information	Technique	References
	Text of printed transcripts	Using ML model to analyze the sound signature of dot-matrix printers	[18]
	Text typed on a keyboard	Because each keystroke has a unique sound signature, a ML model can be used to map them and recover the typed text	[19] [20] [21]
Appearance	What was said in a video of a person talking without a sound.	Using ML to read the lip movement and translate it into text	[24] [30]
	Diagnosis of movement diseases and disorders	By analyzing the walking rhythm and style of the patients	[13] [14]
	Text typed on a keyboard	Because each keystroke has a unique vibration signature, a ML model can be used to map them and recover the typed text	[22]
Behaviors	Text typed on a keyboard	By pointing a laser in a specific location to read the vibration emitted by the target device	[23]
	What is being said in a different room (without a line of sight)	Utilizing Wi-Fi signals to detect and analyze micromovement generated by the lips during conversations with the aid of a ML model	[26][27]
	Lie detection	Using a ML model to record and analyze the reaction times and other speaking attributes	[31]

Analysis and Discussion

A. The Philosophical Predicament

In the literature review section, we have demonstrated the current state-of-the-art capabilities of ML technologies to derive significant information from seemingly dreary data. Most of that data is retrieved from legal and consensual means. For example, you wouldn't feel that your privacy has been violated if a person hears the sound of keystrokes while you're typing. However, how would you feel if the same data is fed into a ML model to determine what you have typed? Similarly, when your voice pattern is heard in public and when it's being analyzed. Although the data is obtained from the same environment, the capabilities of the data processor make this act uncomfortable for the subject. This also raises the philosophical question of do you own your data? For example, as you're walking across a street, other people may infer information or judgment about you just by the way you walk, sound, and appear. In this case, would you have the right to prevent people from making such conclusions about you? How is that different when the same is done by a capable human (e.g., physician) or a trained ML model? Given that, the topic of privacy and the usage of data in the future can be more complex than what we have anticipated.

B. Possible Use and Misuse Cases

We have demonstrated the capabilities of ML models and established that significant conclusions could be deducted from dreary data that is acquired from consensual means. This begs the question of how can such capabilities be used and misused?

Governments and anti-terrorism agencies may utilize the discussed techniques to combat terrorism. Without a search warrant or any other legal paperwork, such agencies may get further insights about their suspect with significant efficiency in time and resources. A relevant example can be traced way back to 1945. During that time, a Swedish newsreel containing Hitler walking slowly was leaked from German censorship. While the videotape had only a few seconds of Hitler's movement, it was thoroughly analyzed to conclude that he had stage 1.5 Parkinson's disease [28]. By using much higher fidelity videotapes and other forms of data, the same principle can be stretched to deduce an immense amount of information about terrorists and war criminals.

On the other hand, and as expected by any technology, the same principle can be exploited to breach the privacy of individuals. Employees in the workplace are naturally producing and exposing a significant number of data about themselves, ranging from personal habits and demographic data to political views. When that data is combined with data that is obtained from side-channel attacks, we can draw significant conclusions about the employee and his personal life. Most organizations hold the right to log and inspect the user's activities in full detail. From CCTV footage to computer logs, a lot of details can be found when we feed such data into a ML model. Given that, it has become increasingly more challenging to draw the line between what is rightful for the organization and the employee's privacy.

Conclusion And Future Work

In conclusion, attackers are getting creative with their methodologies, and technologies are more capable than ever. By analyzing seemingly dreary data, so much can be inferred about an individual, such as demographic, health, and behavioral information. Further, this data can be obtained via legal and consensual means, making it very hard to draw a line for privacy violations. Therefore, and in future work, we believe that in order to solve this predicament methodically, we have to approach it from a philosophical point of view.

REFERENCES

[1] **Lerman, Liran, Gianluca Bontempi, and Olivier Markowitch.(2014).** "Power analysis attack: an approach based on machine learning." International Journal of Applied Cryptography 3.2

[2] **Svartvik, Jan. (1968).** The Evans Statements. University of Goteburg.

[3] **O'Keeffe, Anne, and Michael McCarthy, eds. (2010).** The Routledge handbook of corpus linguistics. Routledge.

[4] **Kaczynski, (1995).** Theodore John. "Industrial society and its future."

[5] **Leonard, Robert A., Juliane ER Ford, and Tanya Karoli Christensen. (2016).** "Forensic linguistics: Applying the science of linguistics to issues of the law." Hofstra L. Rev. 45.

[6] **Peacock, Alen, Xian Ke, and Matthew Wilkerson. (2004).** "Typing patterns: A key to user identification." IEEE Security & Privacy 2.5.

[7] **L. F. Coppenrath and Associates. (2001).** Biometric Solutions By Classification.http://www.lfca.net/Reference%20Documents/Biom etric%20Solutions%20By%20Classification.pdf.
[8] **Banerjee, Salil P., and Damon L. Woodard. (2012).“**Biometric authentication and identification using keystroke dynamics: A survey.” Journal of Pattern Recognition Research 7.1.
[9] **Büyük, Osman, and Mustafa Levent Arslan. (20182).** “Combination of Long- Term and Short-Term Features for Age Identification from Voice.” Advances in Electrical and Computer Engineering 18.2.
[10] **Torres-Saillant, Silvio. (2003).** “Inventing the race: Latinos and the ethnoracial pentagon.” Latino Studies 1.1.
[11] **Hawkins, Francine Dove. (1992).** Speaker ethnic identification: The roles of speech sample, fundamental frequency, speaker and listener variations. Diss. University of Maryland, College Park.
[12] **Foreman, Christina Gayle. (2000).** “Identification of African-American English from prosodic cues.” Texas Linguistic Forum. Vol. 43. Austin; University of Texas.
[13] **Moro-Velazquez, Laureano, et al.(2018).** “Analysis of speaker recognition methodologies and the influence of kinetic changes to automatically detect Parkinson’s Disease.” Applied Soft Computing 62.
[14] **Singh, Sanjana, and Wenyao Xu. (2020).** “Robust detection of Parkinson’s disease using harvested smartphone voice data: a telemedicine approach.” Telemedicine and e-Health 26.3.
[15] **Souza, Lourdes Bernadete Rocha de, and Marquiony Marques dos Santos. (2018).** “Body mass index and acoustic voice parameters: is there a relationship?” Brazilian journal of otorhinolaryngology 84.
[16] **Celebi S, Yelken K, Develioglu ON, Topak M, Celik O, Ipek HD, Kulekci M. (2013).** Acoustic, perceptual and aerodynamic voice evaluation in an obese population. The Journal of Laryngology & Otology.
[17] **Bales, Charles, et al. (2020).** “Can machine learning be used to recognize and diagnose coughs?.” 2020 International Conference on e-Health and Bioengineering (EHB). IEEE.
[18] **Halevi, Tzipora, and Nitesh Saxena. (2015).** “Keyboard acoustic side channel attacks: exploring realistic and security-sensitive scenarios.” International Journal of Information Security 14.5.
[19] **Li Zhuang, Feng Zhou, and J. Doug Tygar. (2005).** Keyboard acoustic emanations revisited. In Proc. 12th ACM Conference on Computer and Communication Security (CCS 2005), pages 373–382. ACM Press, (2005).
[20] **Yigael Berger, Avishai Wool, and Arie Yeredor (2006).** Dictionary attacks using keyboard acoustic emanations. In Proc. 13th ACM Conference on Computer and Communication Security (CCS 2006), pages 245–254. ACM.
[21] **Dmitri Asonov and Rakesh Agrawal. (2004).** Keyboard acoustic emanations. In Proc. 2004 IEEE Symposium on Security and Privacy (Oakland 2004), pages 3–11.
[22] **Marquardt, Philip, et al. (2011).** “(sp) iphone: Decoding vibrations from nearby keyboards using mobile phone accelerometers.” Proceedings of the 18th ACM conference on Computer and communications security. (2011).
[23] **Barisani, Andrea, and Daniele Bianco. (2009).** “Sniffing keystrokes with lasers/voltmeters: Side channel attacks using optical sampling of mechanical energy and power line leakage.” Black Hat Technical Security Conference: USA. (2009).
[24] **Chung, Joon Son, and Andrew Zisserman. (2016).** “Lip reading in the wild.” Asian conference on computer vision. Springer, Cham, (2016).
[25] **Davis, Abe, Michael Rubinstein, Neal Wadhwa, Gautham J. Mysore, Fredo Durand, and William T. Freeman.(2014).** “The visual microphone: Passive recovery of sound from video.”

[26] **Wang, Guanhua, et al.(2016).** "We can hear you with Wi-Fi!." IEEE Transactions on Mobile Computing 15.11

[27] **Kamran Ali, Alex X. Liu, Wei Wang, and Muhammad Shahzad. (2015).** Keystroke Recognition Using Wi-Fi Signals. In Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom' 15). Association for Computing Machinery, New York, NY, USA, 90–102. DOI:https://doi.org/10.1145/2789168.2790109.

[28] **Bhattacharyya, Kalyan B. (2015).** "Adolf Hitler and his parkinsonism." Annals of Indian Academy of Neurology 18, No. 4.

[29] **Ling, Caijin, et al. (2016).** "You cannot sense my pins: A side-channel attack deterrent solution based on haptic feedback on touch-enabled devices." 2016 IEEE Global Communications Conference (GLOBECOM). IEEE,

[30] **Chung, Joon Son, and Andrew Zisserman. (2018).** "Learning to lip read words by watching videos." Computer Vision and Image Understanding 173

[31] **Monaro, Merylin, et al. (2021).** "The detection of faked identity using unexpected questions and choice reaction times." Psychological Research 85.6