



Passphrases Beat Thermal Attacks: Evaluating Text Input Characteristics Against Thermal Attacks on Laptops and Smartphones

Yasmeen Abdrabou, Reem Hatem, Yomna Abdelrahman, Amr Elmougy,
Mohamed Khamis

► To cite this version:

Yasmeen Abdrabou, Reem Hatem, Yomna Abdelrahman, Amr Elmougy, Mohamed Khamis. Passphrases Beat Thermal Attacks: Evaluating Text Input Characteristics Against Thermal Attacks on Laptops and Smartphones. 18th IFIP Conference on Human-Computer Interaction (INTERACT), Aug 2021, Bari, Italy. pp.712-721, 10.1007/978-3-030-85610-6_41 . hal-04215498

HAL Id: hal-04215498

<https://inria.hal.science/hal-04215498>

Submitted on 22 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

Passphrases Beat Thermal Attacks: Evaluating Text Input Characteristics Against Thermal Attacks on Laptops and Smartphones

Yasmeen Abdrabou¹, Reem Hatem², Yomna Abdelrahman¹, Amr Elmougy²
and Mohamed Khamis³

¹ Bundeswehr University Munich, Germany

² German University in Cairo, Egypt

³ University of Glasgow, United Kingdom
`yasmeen.essam@unibw.de`

Abstract. We investigate the effectiveness of thermal attacks against input of text with different characteristics; we study text entry on a smartphone touchscreen and a laptop keyboard. First, we ran a study (N=25) to collect a dataset of thermal images of short words, websites, complex strings (special characters, numbers, letters), passphrases and words with duplicate characters. Afterwards, 20 different participants visually inspected the thermal images to attempt to identify the text input. We found that long and complex strings are less vulnerable to thermal attacks, that visual inspection of thermal images reveals different parts of the entered text (36% on average and up to 82%) even if the attack is not fully successful, and that entering text on laptops is more vulnerable to thermal attacks than on smartphones. We conclude with three learned lessons and recommendations to resist thermal attacks.

Keywords: Thermal Imaging · Security · Privacy · Side-channel attack.

1 Introduction

Recent research has revealed that thermal cameras can be used to infer different types of passwords, such as text passwords [12] and PINs/Patterns [7, 8]. This presents a threat to authentication on safes [20], smartphones [7, 8], keyboards [8, 12], cash machines [16], digital door locks, and payment terminals [19]. Attacks that use thermal cameras are often referred to as “thermal attacks” [7]. They were shown to be feasible using both high-end and low-cost commercial cameras [7, 8], and were evaluated under threat models where attackers have access to automated image processing approaches [7, 12] and where attackers are non-experts that rely on visually inspecting the thermal images [8, 12, 20].

Despite the recent work on thermal attacks, it is unclear which characteristics of text strings impact their vulnerability to thermal attacks. To address this gap, this paper investigates how successful visual inspection of thermal images taken by an off-the-shelf thermal camera can retrieve text strings entered on laptops

and smartphones. We studied different input characteristics such as presence of duplicates and special characters and input length. We study attacks by non-expert attackers, which we define in this context as attackers who have access to a thermal camera and can visually inspect the recorded images, but do not have the skills to implement an automated approach as done in some prior research [7, 12]. Attacks by such non-experts are more likely as thermal cameras become significantly cheaper (e.g., at the time of writing this paper, a thermal camera can be bought for less than $< \text{£}200$ [1]). Our results show that employing thermal attacks on text input is effective; without training and using visual inspection, our participants were able to infer up to 90% of short text input and almost half short text inputs that contain duplicates. We found that longer text input, such as passphrases, and the inclusion of special characters in text input significantly reduce vulnerability. We conclude with learned lessons on securing text input against thermal attacks.

2 Background and Related Work

When a user touches a surface, heat is transferred from the user’s fingers to touched surfaces. This generates a temperature difference at the point of contact referred to as heat traces. Thermal cameras allow its users to see heat traces.

Thermal attacks exploit this phenomenon by interpreting the user’s input on a user interface based on the heat traces [7, 16]. Thermal attacks have an advantage over other types of side-channel attacks, such as shoulder surfing [10] in that they allow determining the order of the entered input; for example the last touched digit in a PIN often has the warmest heat trace. Another advantage for attackers is that thermal attacks are performed after the user had left the device. This gives an advantage over shoulder surfing as attackers no longer need to observe the user while interacting, which makes the attack more covert.

Previous work demonstrated the effectiveness of thermal attacks on PINs on different devices such as safes [20], smartphones [7, 8], keyboards [8, 12], cash machines [16], digital door locks, and payment terminals [19]. High-end thermal cameras as well as low-cost commercial cameras have been used in different research showing the real threat to users’ privacy [7, 8, 16]. Both automated approaches [7, 12] and visual inspection attacks by non-expert human attackers [8, 20] were highly successful. Most of the research has investigated thermal attacks on PINs or patterns which are usually short entries. An exception is one work that compared the success of thermal attacks against strong and weak passwords [12] which are relatively longer than PINs and patterns. In their research, Kaczmarek et al. [12] studied attacking weak and strong passwords entered on external keyboards using a commercial thermal camera. They found that a non-expert attacker can recover key presses up to 30 seconds after entry.

This paper complements previous work by reporting on an in-depth analysis of how well thermal attacks perform on user text input. We study different input characteristics that were not studied before, such as text length, presence of duplicates, and presence of special characters, and investigate their effect on

the thermal attack’s success. We collect thermal images of inputs entered on a smartphone’s touchscreen and a laptop’s keyboard. In our threat models, non-expert attackers visually inspect thermal images taken by an off the shelf thermal camera. Our results shed light on how to protect against thermal attacks.

3 Evaluation

In this work, we investigate different text input characteristics (length, duplicate letters and special characters, upper and lower case letters) and their influence on the attack success rate (**RQ**). To answer our research question, we ran a lab study consisting of two phases (1) Thermal Images Collection and (2) Thermal Attacks. We implemented a simple website interface for users to enter their input. The interface shows the input text to be typed and a text area for input entry.

3.1 Threat Model

In our threat model, the attacker waits until the victim leaves the laptop or a smartphone unattended. This could be the case when the user texts someone, googles something or visits a website and then leaves the device. To ensure optimal but realistic conditions, we assume the user’s only interaction was to enter the text. The attacker then captures the thermal image and visually inspects.

3.2 Phase 1: Thermal Images Collection

To collect thermal images to be used in the analysis of thermal attacks, we recorded thermal images of a laptop and a smartphone after participants entered text with different characteristics. We studied five categories of text:

Category 1: Short average length of 4 ± 1 characters and no duplicates.

Category 2: Duplicates short words with average length 4 ± 1 characters and include duplicates.

Category 3: Websites RootDomains (domain name and Top level Domain) [5] with average length 9 ± 2 . The websites were chosen from the 50 most visited websites [6] and included one duplicate letter.

Category 4: Complex medium length inputs with special characters of an average length of 8 characters and no duplicates. We generated those using a password generator [4].

Category 5: Passphrase a passphrase is a password in the form of a phrase e.g., aMooseSendsAgoal. The idea behind passphrases is that they are long yet easy to remember by users [18]. The passphrases were chosen from a passphrase generator [3] with an average length of 20 ± 2 and includes three duplicate letters.

Participants, Apparatus, and Setup We recruited 25 participants (16 males, 9 females) using a mailing list. Their ages varied between 18 to 23 ($M = 21.32$, $SD = 1.14$). Only one participant was left-handed and none of the participants had experience with thermal cameras. We used a Flir C2 Camera [2] thermal




Fig. 1. The thermal images collection setup showing the input device, and the thermal camera placed on a tripod placed 30 cm away from the device.

camera with resolution ($80 \text{ px} \times 60 \text{ px}$). The camera was mounted on a 25 cm high tripod placed 30 cm away from the device (see Figure 1). Input was provided on a Lenovo Tango Phab 2 Pro smartphone with a gorilla glass screen ($1440 \text{ px} \times 2560 \text{ px}$) pixels, and a MacBook Air Laptop ($1440 \text{ px} \times 900 \text{ px}$).

Study Procedure After arriving at our lab, participants signed a consent form and provided their demographics data. Participants were asked not to exert physical activity before arriving in the lab including taking the stairs lest it impacts their body temperature. Participants were asked to remove any metal objects in their hand as well as gloves. To let the participants familiarize themselves with the keyboards, they were asked to enter random text 3 times, then the experiment started. Participants provided text input based on a predefined list which covered the 5 aforementioned categories. Our web interface displayed the text input to be entered. Participants were asked to wait one minute in-between each two text entries to make sure heat trackers had decayed, and the thermal image was taken 4 seconds after entry. These decisions were inspired by prior work on thermal attacks [8]. Each participant entered a total of 40 entries (4 per category, half on the laptop and the other half on the smartphone). The room temperature was kept the same throughout the study 24°C .

3.3 Phase 2: Thermal Attacks

The aim of this phase was to measure the success of thermal attacks by visually inspecting the thermal images collected in the previous phase.

Study Design and Measures This phase followed a within-subjects design and included two independent variables: the input category and the input device. We measured: 1) Levenshtein distance and 2) successful attack rate. The Levenshtein distance is a metric for measuring the difference between two strings and was used to measure how close the attacker’s guess is to the original text. The *success rate* is the percentage of attacks that fully recover the full entered input in the correct order from visually inspecting the thermal image. These metrics are recommended by literature [8, 9, 11, 13–15, 21].

Participants and Procedure We invited 20 participants (9 males, 11 females) to our study. Their ages varied from 15 to 29 ($M = 21.85$, $SD = 3.297$). None of them had participated in phase 1 nor had experience with thermal cameras.

After arriving at the lab, participants were asked to fill in the consent and demographics forms, then were explained the study. We created a questionnaire that included the thermal images collected from the previous phase, and 3 text fields to allow participants to provide up to 3 guesses. Each participant inspected 20 images, 5 for each input device and two for each input category. The images were randomly chosen from our collected dataset and the order was counter-balanced using a Latin square.

4 Results

4.1 Levenshtein Distance

Figure 2 (left) shows the mean Levenshtein distance per input category and input device. The Levenshtein distance is “The smallest number of insertions, deletions, and substitutions required to change one string or tree into another” [17]. Overall, participants were more successful in guessing the correct characters and positions on the smartphone than on the laptop (shorter Levenshtein distance). This was also proven statistically: a repeated measures ANOVA showed statistical significant effect of device (laptop ($M = 7.85$; $SD = 1.18$) and smartphone ($M = 6.70$; $SD = .71$)) on Levenshtein distance ($F_{1,19} = 46.6875$, $P < .001$).

Overall the Levenshtein distance was shorter on laptops for short and duplicate inputs than for complex, websites, and passphrase inputs. This was also reflected in a repeated measures ANOVA where it showed a statistically significant effect of the input category on the Levenshtein distance ($F_{4,72} = 264.855$, $P < .001$). Pairwise comparison showed statistical significance between all input categories, short ($M = .97$; $SD = 1.33$), duplicates ($M = 2.71$; $SD = 2.47$), complex ($M = 7.85$; $SD = .24$), website ($M = 6.5$; $SD = 1.83$) and passphrases ($M = 15.72$; $SD = .34$), all $P < .001$.

For the laptop, the Levenshtein distance is shorter for short and duplicate inputs. This was also reflected statistically, where a repeated-measures ANOVA revealed a significant main effect of the input category on the Levenshtein distance ($F_{4,72} = 264.855$, $P < .001$). Pairwise comparison with Bonferroni correction showed statistical significance between all input categories, short input ($M = .28$; $SD = .61$), duplicates ($M = 2.71$; $SD = 2.47$), complex ($M = 7.53$; $SD = .72$), websites ($M = 10.18$; $SD = 3.5$), passphrases ($M = 18.71$; $SD = 1.18$), all $P < .05$.

On the smartphone, the same behavior was noticed where the Levenshtein distance was also shorter for short and duplicates input categories. A repeated-measures ANOVA revealed a significant main effect of the input category on the Levenshtein distance ($F_{4,72} = 409.246$, $P < .001$). Pairwise comparison with Bonferroni correction showed statistical significance between all input categories, short input ($M = .28$; $SD = .6$), duplicates ($M = 2.55$; $SD = 1.73$), complex




Fig. 2. The Levenshtein Distance between the attackers’ guesses and the original input (left) – shorter bars represent smaller differences between the guessed and the actual input – and the attack success rate (right). Passphrases and complex input are significantly more secure than short inputs, words with duplicated and websites. Attacks are more successful on the laptop than on the smartphone.

($M = 10.25$; $SD = 3.43$), websites ($M = 7.5$; $SD = .71$), passphrases ($M = 18.72$; $SD = 1.16$), all $P < .001$. This means that guesses against short input are significantly closer to the correct input than those against all other categories.

4.2 Success Rate

Figure 2 (right) shows the overall success rate of all input categories. As seen, the success rate is higher on laptops than on the smartphone. This was also statistically significant as found by a repeated-measures ANOVA where the input interface (laptop ($M = 30$; $SD = 11.23$) and smartphone ($M = 22.50$; $SD = 15.52$)) have a significant effect on the success rate of the attack $F_{1,19} = 7.703$, $P = .012$. The figure also shows that short and duplicate inputs are the most vulnerable.

Success rate in case of attacks against input on the laptop was the highest for short inputs. A repeated measures ANOVA showed a statistically significant effect of the input category on the success rate of attacks against input on the laptop ($F_{4,72} = 56.410$, $P < .001$). Pairwise comparison with Bonferroni-correction showed statistical significance for the success rate between all pairs of the following input categories: short ($M = 90$; $SD = 20.51$), duplicates ($M = 50$; $SD = 37.28$), website ($M = 7.5$; $SD = 24.47$), $P < .001$. There was 0% success rate against complex and passphrases inputs.

The same patterns were also found for attacks against input on smartphone, where the attack success rate was higher for short, duplicates and website inputs. A repeated-measures ANOVA revealed a significant main effect of the input category on the attack success rate against input on the smartphone ($F_{4,72} = 33.921$, $P < .001$). Pairwise comparison showed statistical significance for all pairs of input categories: short input ($M = 80$; $SD = 29.91$), duplicates ($M = 34.21$; $SD = 41$), websites ($M = 10$; $SD = 26.15$), $P < .001$. Attacks against passphrases and complex inputs were never successful (0% success).

The results suggest that attackers can guess many parts of all input categories, but are less likely to guess the entire string. This is particularly clear in attacks against laptops, as the success rate against laptops is higher than against smartphones, but on average the guesses seem to be closer to the original text

on smartphones than on laptops. Even a partial reveal of input poses a risk as this can reveal personal info to attackers e.g., search history, text messages, etc.

5 Discussion and Future Work

Our results show the possibility of using a low-cost thermal camera to conduct thermal attacks against text input by visually inspecting the thermal images. The attack success rate highly depends on many aspects, the length of the input, the characters entered, and the victim’s hand and device temperature.

Lesson 1: Passphrases and Complex Entries are Less Vulnerable to Thermal Attacks We found that attacks were more successful on short words and words with duplicates rather than websites, long sentences and complex combinations of letters. This is likely because the longer length of the input results in longer entry time, which means that by the time the attacker takes a thermal image, the heat traces resulting from entering the first characters will have decayed. In the case of duplicates, the overlapping heat traces made identifying the input more challenging. This can be seen in the successful attack percentage of 53% for duplicates and 90% for short input. For medium input length represented in websites and complex combinations of letters, we found that attacks were significantly more successful against website inputs than against complex input. In case of complex entries, the need to press the shift button first and then select the uppercase, special characters or even a number adds more heat traces that distract the attacker. In case of passwords, this means passwords with special characters or duplicates are less likely to be vulnerable to thermal attacks than those without. This also applies when users choose a passphrase for their password as its length alone sufficiently reduces the thermal attack risk.

Therefore, to protect from thermal attacks, we recommend users to use long passwords and/or complex entries. The fact that these two types of input are resilient to thermal attacks is positive because users are recommended to use them to protect against other types of attacks (e.g., offline dictionary attacks).

Lesson 2: Text Input on Laptop Keyboards are more Vulnerable to Thermal Attacks than on Smartphones Touchscreens We found that thermal attacks against text input were more successful on laptop keyboards than on smartphone touchscreens. Smartphones’ screen is directly attached to the processing unit which means that it may heat up due to CPU usage. This may distort heat traces. On the other hand, laptop’s processing unit is rather affecting a relatively smaller area of the keyboard, leaving heat traces on the rest of the keyboard unaffected and potentially more visible (see Figure 3). This correlates with the literature where thermal attacks on passwords entered on external keyboards were 80% successful [12]. Another reason for the ineffectiveness of thermal attacks on smartphone keyboards is that the soft keys are too close to each other, which makes it more likely for the attacker to mix nearby keys.

In a study on thermal attacks against graphical passwords by Abdrabou et al. [8], they found that entering graphical passwords on laptops using their touchpads is more secure against visual inspection thermal attacks compared to




Fig. 3. Thermal images for (Left) Laptop keyboard with visible heat traces. (Right) Touch screen smartphone with less visible heat traces due to device temperature.

when entering them on smartphone touchscreens. This attributed to the fact that entering graphical passwords on touchpads requires first navigating to the initial point, which distorts the heat traces. This does not contradict our results. In a nutshell, previous work [8] shows that entering *graphical passwords* is more secure on touchpads of laptops than on smartphones, whereas we show that entering *text* is less secure on laptop keyboards compared to smartphones.

Lesson 3: Victim and Device Temperature Affects the Success of the Thermal Attack We found a statistically significant difference between the temperatures of user’s hand the device temperature. This means having colder or warmer hand temperature affects the thermal attack success. This suggests that holding hot or cold objects shortly before typing impacts the user’s vulnerability to thermal attacks. For future work, we suggest running an uncontrolled study where thermal images are collected from interfaces after users have interacted in everyday conditions e.g., after exercising, holding hot/cold drinks, etc.

6 Conclusion

In this work, we investigated the possibility of attacking user text input on laptop keyboards and touchscreens of smartphones using a thermal camera. We collected a dataset of thermal images of a smartphone’s touchscreen and a laptop’s keyboard after participants entered different text input categories varying in length and complexity (e.g. some including special characters). In a second study, 20 participants visually inspected the thermal images to infer the input text. We found a significant effect of Input characteristics on vulnerability to thermal attacks. We also found that attacks are more successful on laptop keyboards than smartphone touchscreens due to the high temperature of the smartphone device. We showed that even if the success rate is not high in the case of long and complex input, the Levenshtein distance showed that attackers were still able to detect different parts of the input text but not enough to make the attacks successful. We concluded with three main learned lessons.

Acknowledgments

This work was supported by the Royal Society of Edinburgh (RSE award number 65040), and the EPSRC (EP/V008870/1).

References

1. Affordable thermal camera on amazon. Webpage, <https://www.amazon.co.uk/dp/B07CMDZGV/>, accessed 13 April 2021
2. Flir c2. Webpage (2021), <http://www.flir.eu/instruments/c2/>, accessed 13 April 2021
3. Make me a password. Webpage (2021), <https://makemeapassword.ligos.net/>, accessed 13 April 2021
4. Online password generator. Webpage (2021), <https://passwordsgenerator.net>, accessed 13 April 2021
5. Rootdomains. Webpage (2021), <https://moz.com/learn/seo/domain>, accessed 13 April 2021
6. Top 50 most visited websites. Webpage (2021), <https://www.alexa.com/topsites>, accessed 13 April 2021
7. Abdelrahman, Y., Khamis, M., Schneegass, S., Alt, F.: Stay Cool! Understanding Thermal Attacks on Mobile-Based User Authentication, p. 3751–3763. Association for Computing Machinery, New York, NY, USA (2017), <https://doi.org/10.1145/3025453.3025461>
8. Abdrabou, Y., Abdelrahman, Y., Ayman, A., Elmougy, A., Khamis, M.: Are thermal attacks ubiquitous? when non-expert attackers use off the shelf thermal cameras. In: Proceedings of the International Conference on Advanced Visual Interfaces. AVI '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3399715.3399819>, <https://doi.org/10.1145/3399715.3399819>
9. De Luca, A., von Zezschwitz, E., Pichler, L., Hussmann, H.: Using Fake Cursors to Secure On-Screen Password Entry, p. 2399–2402. Association for Computing Machinery, New York, NY, USA (2013), <https://doi.org/10.1145/2470654.2481331>
10. Eiband, M., Khamis, M., von Zezschwitz, E., Hussmann, H., Alt, F.: Understanding Shoulder Surfing in the Wild: Stories from Users and Observers, p. 4254–4265. Association for Computing Machinery, New York, NY, USA (2017), <https://doi.org/10.1145/3025453.3025636>
11. George, C., Khamis, M., Buschek, D., Hussmann, H.: Investigating the third dimension for authentication in immersive virtual reality and in the real world. In: 2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR). pp. 277–285 (2019). <https://doi.org/10.1109/VR.2019.8797862>
12. Kaczmarek, T., Ozturk, E., Tsudik, G.: Thermanator: Thermal residue-based post factum attacks on keyboard data entry. In: Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security. p. 586–593. Asia CCS '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3321705.3329846>, <https://doi.org/10.1145/3321705.3329846>
13. Katsini, C., Abdrabou, Y., Raptis, G.E., Khamis, M., Alt, F.: The role of eye gaze in security and privacy applications: Survey and future hci research directions. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. p. 1–21. CHI '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3313831.3376840>, <https://doi.org/10.1145/3313831.3376840>
14. Khamis, M., Alt, F., Hassib, M., von Zezschwitz, E., Hasholzner, R., Bulling, A.: Gazetouchpass: Multimodal authentication using gaze and touch on mobile devices. In: Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors

- in *Computing Systems*. p. 2156–2164. CHI EA '16, Association for Computing Machinery, New York, NY, USA (2016). <https://doi.org/10.1145/2851581.2892314>, <https://doi.org/10.1145/2851581.2892314>
15. Mathis, F., Williamson, J.H., Vaniea, K., Khamis, M.: Fast and secure authentication in virtual reality using coordinated 3d manipulation and pointing. *ACM Trans. Comput.-Hum. Interact.* **28**(1) (Jan 2021). <https://doi.org/10.1145/3428121>, <https://doi.org/10.1145/3428121>
 16. Mowery, K., Meiklejohn, S., Savage, S.: Heat of the moment: Characterizing the efficacy of thermal camera-based attacks. In: *Proceedings of the 5th USENIX Conference on Offensive Technologies*. p. 6. WOOT'11, USENIX Association, USA (2011)
 17. Navarro, G.: A guided tour to approximate string matching. *ACM Comput. Surv.* **33**(1), 31–88 (Mar 2001). <https://doi.org/10.1145/375360.375365>, <https://doi.org/10.1145/375360.375365>
 18. Porter, S.N.: A password extension for improved human factors. *Computers & Security* **1**(1), 54–56 (1982)
 19. Wodo, W., Hanzlik, L.: Thermal imaging attacks on keypad security systems. In: *SECRYPT*. pp. 458–464 (2016)
 20. Zalewski, M.: Cracking safes with thermal imaging. ser. <http://lcamtuf.coredump.cx/tsafe> (2005)
 21. von Zezschwitz, E., De Luca, A., Hussmann, H.: Survival of the shortest: A retrospective analysis of influencing factors on password composition. In: Kotzé, P., Marsden, G., Lindgaard, G., Wesson, J., Winckler, M. (eds.) *Human-Computer Interaction – INTERACT 2013*. pp. 460–467. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)