

ThermoSecure: Investigating the effectiveness of AI-driven thermal attacks on commonly used computer keyboards

NORAH ALOTAIBI, University of Glasgow, United Kingdom and Taif University, Saudi Arabia

JOHN WILLIAMSON, University of Glasgow, United Kingdom

MOHAMED KHAMIS, University of Glasgow, United Kingdom

Thermal cameras can reveal heat traces on user interfaces, such as keyboards. This can be exploited maliciously to infer sensitive input, such as passwords. While previous work considered thermal attacks that rely on visual inspection of simple image processing techniques, we show that attackers can perform more effective AI-driven attacks. We demonstrate this by presenting the development of ThermoSecure, and its evaluation in two user studies ($N=21$, $N=16$) which reveal novel insights about thermal attacks. We detail the implementation of ThermoSecure and make a dataset of 1,500 thermal images of keyboards with heat traces resulting from input publicly available. Our first study shows that ThermoSecure successfully attacks 6-symbol, 8-symbol, 12-symbol, and 16-symbol passwords with an average accuracy of 92%, 80%, 71%, and 55% respectively, and even higher accuracy when thermal images are taken within 30 seconds. We found that typing behavior significantly impacts vulnerability to thermal attacks, where hunt-and-peck typists are more vulnerable than fast typists (92% vs 83% thermal attack success if performed within 30 seconds). The second study showed that the keycaps material has a statistically significant effect on the effectiveness of thermal attacks: ABS keycaps retain the thermal trace of users presses for a longer period of time, making them more vulnerable to thermal attacks, with a 52% average attack accuracy compared to 14% for keyboards with PBT keycaps. Finally, we discuss how systems can leverage our results to protect from thermal attacks, and present 7 mitigation approaches that are based on our results and previous work.

CCS Concepts: • Human-centered computing → Human computer interaction (HCI); • Security and privacy → Authentication.

Additional Key Words and Phrases: Authentication, Deep learning, K-Mean Clustering, Mask RCNN

ACM Reference Format:

Norah Alotaibi, John Williamson, and Mohamed Khamis. 2022. ThermoSecure: Investigating the effectiveness of AI-driven thermal attacks on commonly used computer keyboards. *ACM Trans. Priv. Sec.* 1, 1, Article 1 (January 2022), 25 pages. <https://doi.org/10.1145/3563693>

1 INTRODUCTION

Thermal cameras are becoming ubiquitous and more affordable than before. Today, thermal cameras can be bought online for less than \$150. While they have many benefits, they present a new front for side-channel attacks. Namely, taking a thermal image of a user interface, such a keyboard or a touchscreen, reveals heat traces that can be used to determine the user's input. This input can range from day-to-day input on said devices, to sensitive input such as passwords, PINs, credit card numbers, and more. These types of attacks are referred to as thermal attacks [1, 3, 4].

Thermal cameras, unlike regular cameras, can reveal information without requiring the attacker to interact with the targeted victim, be present during the authentication attempt, or plant any

Authors' addresses: Norah Alotaibi, University of Glasgow, Glasgow, United Kingdom, n.alotaibi.3@research.gla.ac.uk, Taif University, Taif, Saudi Arabia, norah.t@tu.edu.sa; John Williamson, University of Glasgow, Glasgow, United Kingdom, johnh.williamson@glasgow.ac.uk; Mohamed Khamis, University of Glasgow, Glasgow, United Kingdom, mohamed.khamis@glasgow.ac.uk.

© 2022 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Privacy and Security*, <https://doi.org/10.1145/3563693>.

tool that can be linked to the attacker which could potentially exposing them. Such information includes heat residues left by the user during authentication, which can be retrieved using thermal cameras. Having acquired a thermal image of a keyboard or touchscreen after authentication, the attacker can then analyze the heat map and exploit it to uncover the entire password or pattern. Even without knowing the order of the keys, it is possible to significantly reduce the search space, which means fewer attempts are required to guess a password. Even if the order of the keys is unknown, it is possible to significantly reduce the search space, requiring fewer attempts to guess a password. Furthermore, our study and studies from prior work [1, 3, 4, 23] have shown that the order of entries can be leaked within a certain time frame.

Previous work has studied thermal attacks against ATMs [23], keyboards [3, 4, 13], smartphone touchscreens [1, 3], and laptop touchpads [3]. It was shown that thermal attacks can be successful if the resulting thermal images are visually inspected by non-experts [3, 4] or if they are analyzed using simple image processing techniques as done by Mowery et al. [23] and Abdelrahman et al. [1]. While previous work showed high success rates ranging from 72% to 100% [1] when thermal images of passwords are taken 30-60 seconds after user input in controlled conditions, we argue that higher accuracy under different contexts can be achieved by leveraging deep learning techniques. Machine learning is becoming increasingly accessible, making it more likely that attackers will employ it to improve their thermal attacks. Thus, there is a need to understand how successful thermal attacks can be if attackers employed more advanced methods for analyzing thermal images, and how user's behavior and input properties impact the success of thermal attacks.

We address this gap with a focus on thermal attacks against passwords entered on keyboards. Our work presents the implementation of ThermoSecure, a novel system that integrates deep learning to 1) determine the placement of keyboards in thermal images using Mask RCNNs, 2) determine which keys were pressed on the keyboard including accurate detection of keys that were pressed multiple times using K-mean clustering, 3) distinguish which keys were part of a username and which were part of a password entry, and 4) determine the order in which the keys were pressed to produce a list of the most likely user input using probability functions. We trained and evaluated our models using a dataset of 1500 thermal images taken in realistic conditions, which is made publicly available.

We then present the results of two user studies to assess the effectiveness of AI-driven thermal attacks against passwords of different properties, different input behaviors, and different keycap material types. First, in Study I (Section 6.1) we evaluate ThermoSecure in an empirical within-subjects user study in which 21 participants entered usernames and passwords on an external keyboard. Our participants entered passwords of different properties, and we took thermal images at 20, 30 and 60 seconds after entry. Our results reveal insights about 1) properties that make passwords more secure against thermal attacks and 2) typing behaviors that make input more secure against thermal attacks. For example, our analysis shows that hunt and peck typing is significantly more vulnerable to thermal attacks (92% thermal attack success if taken within 30 seconds) compared to fast typing (80%) and that this typing behavior can be determined in real time through keystroke dynamics. This creates avenues for future work on real time protection from thermal attacks by analyzing typing behavior. We also found that long passwords are significantly more resilient to thermal attacks; 100% of 6-symbol passwords are detected using ThermoSecure whereas 67% of 16-symbol passwords are detected within 20 seconds. Second, in Study II (Section 6.2) we investigate how some physical properties of external keyboards impact the success of thermal attacks through a follow up within-subjects user study in which 16 participants entered passwords on two keyboards: one that uses Acrylonitrile Butadiene Styrene (ABS) keycaps, and one that uses Polybutylene Terephthalate (PBT) keycaps. Our results indicate that Keycaps made of ABS were more vulnerable to thermal attacks than those made of PBT (52% and 14% attack success

respectively). We conclude with a discussion of how systems can protect users from thermal attacks by presenting 7 mitigation approaches that are based on our findings and previous work.

2 CONTRIBUTION STATEMENT:

In summary, this paper makes the following contributions:

- (1) We present the first dataset of thermal images showing keyboards with heat traces resulting from input. The dataset contains 1500 thermal images and can be used to further investigate and understand Thermal attacks¹.
- (2) We present the implementation of a novel model that:
 - (a) automates the detection of keyboards within the thermal image using a deep learning model: Mask RCNNs.
 - (b) employs K-mean clustering to detect multiple presses in the thermal image.
 - (c) distinguishes which key presses belong to the username and which belong to the password in an authentication attempt.
 - (d) offers a novel technique to infer the order of the pressed keys. This technique generates a combination of different passwords with different order probabilities.
- (3) Novel insights about the effectiveness of thermal attacks in realistic settings and how they are impacted by password properties, ages of heat traces, user typing behaviors, and the plastic material used for the keyboard's keycaps.

3 RELATED WORK

Previous work on thermal attacks is scarce. This is likely because the threat has only recently become feasible due to the falling prices of thermal cameras. Work in this area can be split into two categories based on the threat model: 1) thermal attacks in which the attacker utilized an automated approach to analyze the thermal images, and 2) attacks in which the attacker visually inspected the thermal image to determine the input.

3.1 Automated Thermal attacks

To the best our knowledge, the earliest published work on thermal attacks was by Mowery et al. [23]. In their work, they focused on ATM keypads and experimented with manual visual inspection of thermal images and an automated approach. They found that their automatic technique was more accurate than visually inspecting the thermal image, particularly as time went on. In comparison to visual inspection, which retrieved just 20-30% of codes after a minute, the automated technique retrieved roughly 50% of them.

Similarly, a study by Li et al. [19] investigated thermal attacks on ATM keypads. In addition to the thermal camera, an RGB camera was also used to help locate the keypad in the thermal image. The two cameras were aligned in advance for joint sequence analysis. In their model, the order of entries in the password was inferred by a frame by frame comparison of change in the temperature of keys. The main contribution of their work was developing a model estimation method of key touch time based on maximum likelihood, achieving an accuracy of 26.7% in attacking 6-digit PINs.

While the previously discussed work focused on ATM keypads, other researchers investigated different user interfaces. A study by Abdelrahman et al. [1] investigated the effectiveness of thermal attacks on user authentication on mobile devices. They developed the ThermalAnalyzer, which featured a recognition pipeline that reconstructs PINs and Android Lock Patterns by analyzing thermal images taken with an optris PI thermal camera. The approach used blob detection and the mean temperatures of the heat traces in the regions of interest to determine the input and order

¹<https://doi.org/10.5281/zenodo.7069957>

of entry. In their work, they achieved an overall accuracy of 78% when attacking PINs within 30 seconds of entry, and 38.89% when attacking patterns. It is notable that attack accuracy was 100% when attacking patterns that do not contain any overlapping input within 30 seconds.

Closest to our work is the work by Kaczmarek et al. [13], which focused on thermal attacks against keyboards. They collected thermal images of keyboards from fixed camera locations and orientations. They asked participants to locate, draw the regions where the heat traces are located and label the keys. Their system, called ThermaNator, used a blob detection technique to separate the thermal traces from the background. The authors show that the key-presses that constitute the password can be recovered within 30 seconds, however they do not estimate the order of key-presses. Instead, the authors suggest using a dictionary attack to generate possible passwords that use the detected keys.

3.2 Thermal Attacks by Visual Inspection

While the previously discussed works presented automated methods for analyzing the thermal images, other works explored how well visually inspecting thermal images can reveal input.

A study by Wodo and Hanzlik [31] presented several scenarios that simulated thermal attacks against passwords on computer keyboards, cash machines, digital doors locks and payment terminals. In these scenarios most of user's passwords were successfully retrieved within a time frame that varies between 30 seconds and 40 seconds. No further details about the accuracy of the attacks were reported. Another study by Abd Rabou et al [3] investigated thermal attacks against touch gestures and taps on smartphone touchscreens and laptop touchpads. In this study, two sets of participants were recruited: the first entered passwords while input was recorded using a Flir C2 Compact thermal camera after 4 seconds of entry, and the second set inspected the recorded thermal images with the aim to determine the passwords. The authors found that touch gestures are more vulnerable compared to tapping on the touchscreens/touchpads (60.65% vs 23.61% success rate), and that touchscreens are more vulnerable than touchpads (87.04% vs 56.02%). One main insight from that work is that thermal attacks are feasible using affordable cameras even when the attackers are not trained, suggesting that thermal attacks can potentially become ubiquitous. This is supported by further evidence from recent work by Bekaert et al. [8], which showed that users' typical behavior makes them frequently at risk of thermal attacks due to, for example, their choice of authentication methods, and leaving devices unattended.

3.3 How our work advances state of the art

Our work advances state of the art in three directions 1) our work features improved attack accuracy, 2) our evaluation uses a more realistic threat model, and 3) our work reveals novel insights regarding how password properties, plastic material of keyboards, and user input behavior impact vulnerability to thermal attacks.

While most previous work on thermal attacks considered smaller search spaces (e.g., 10 keys on an ATM keypad or a mobile touchscreen [1, 3, 4, 23, 31]), our work focuses on thermal attacks against keyboards, which are relatively understudied. In contrast to the work by Kaczmarek et al. [13], our model is able to determine the order of entries by comparing the temperatures of the heat traces at the different keys and uses the information obtained from the thermal image to probabilistically decode multiple hypothesized passwords ranked by probability of correctness. Our approach also distinguishes username and password key-presses by comparing every pair of consecutive keys to find the ones with the highest temperature difference (i.e., the last key of the username and the first key of the password), which can indicate a possible shift that corresponds to typing something new (i.e., password). Furthermore, our deep learning approach achieves high accuracy: up to 100%, 93%, 82% and 67% for 6-symbol, 8-symbol, 12-symbol and 16-symbol passwords respectively. This

is significantly higher than the values reported in Sections 3.1 and 3.2. We also reveal valuable insights on the impact of typing behavior and keycaps' material on the success of thermal attacks.

Furthermore, all previous work assumed a fixed setup where the distance between the user interface and the camera is fixed [13]. However, it cannot be assumed that these methods will always work in a real life scenarios where the attacker will need to configure the method for every image captured. In contrast, a) our study involves thermal images from different angles and distances from the keyboard and b) our model is able to distinguish the entry of usernames and passwords.

Taking this into consideration, understanding how successful thermal attacks can be if they benefit from all the aforementioned information is essential to assess the attacks' impact on security.

4 THREAT MODEL

In our threat model, the attacker possesses a thermal camera and uses it to take a thermal image of the surface of the keyboard after the user has authenticated by entering a username and a password. Our experiment considers cases where the thermal image is taken 20, 30 and 60 seconds after authentication. This could happen in situations where the user logs in and then shortly leaves their workstation e.g., to take a break. The attacker can then take a thermal image after the user has left their workstation unattended. Another possible scenario is where the attacker inconspicuously uses a thermal camera while the victim is still using their computer provided they do not occlude the keyboard. This can be done either by using a small add-on thermal cameras, such as FLIR one², that can be attached to smartphones, and using a smartphone that comes with an integrated thermal camera³), or through a mounted thermal camera used for security⁴. It does not matter whether the user has logged out before leaving their workstation without using their keyboard or if the device automatically logs them out due to inactivity – this does not impact our threat model. We assume the user does not interact with the keyboard after log in. This is a reasonable assumption as users may spend complete interaction sessions using their mouse or consuming content displayed on the screen.

Unlike previously studied threat models where the user entered a password only [1, 3, 4, 13], our model assumes the user has entered both a username and a password. This means that the attacker in our threat model needs to distinguish heat traces resulting from interaction before authentication (i.e., entering the username) from those resulting from authentication (i.e., entering the password).

5 THERMOSECURE: CONCEPT AND IMPLEMENTATION

In the following, we present our implementation of ThermoSecure, a method that retrieves input on keyboards through thermal imaging. The main objectives of ThermoSecure were shaped through comparative review of how previous studies approached the topic of thermal attacks [1, 3, 4, 13, 23], and addressing the gaps by said approaches. While our threat model and some of the approaches used to process the thermal images were inspired by previous work, ThermoSecure advances state of the art through the following:

- (1) ThermoSecure incorporates an object detection technique based on Mask RCNN to ensure that the placement of the keyboard in the thermal image does not reduce the effectiveness of the attack.
- (2) While previous approaches for thermal attacks demonstrated some success, none of the previously proposed methods outputs the correct entry every time [1, 13]. Instead of producing a

²<https://www.flir.com/flir-one/>

³<https://www.catphones.com/en-gb/features/integrated-thermal-imaging/>

⁴<https://www.flir.co.uk/browse/security/thermal-security-cameras/>

single output, ThermoSecure uses the information obtained from the thermal image to probabilistically decode multiple hypothesised passwords ranked by probability of correctness.

- (3) While previous work assumed the user does not provide any input apart from the password, ThermoSecure is evaluated under a threat model where the user provides both a username and a password. ThermoSecure's ability to distinguish heat traces resulting from password entry and those resulting from interactions prior to authentication demonstrates that the attack is feasible in scenarios that are more realistic compared to previously studied threat models [1, 13].

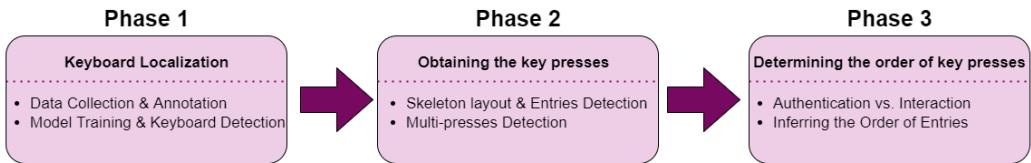


Fig. 1. ThermoSecure architecture

Figure 1 summarizes the architecture of the model used in ThermoSecure. The next subsections detail the different phases that form ThermoSecure. The first phase is concerned with the detection of the keyboard within the thermal image. The coordinates obtained from this phase are then used in the second phase to apply a skeleton layout to the detected keyboard to extract the pressed keys. Having defined which keys form the password, the final phase produces the order of the key presses to retrieve the full password.

5.1 Phase 1: Keyboard Localization

In this phase, our goal is to detect the keyboard within the thermal image and obtain the bounding box coordinates surrounding it, such that the thermal image of the keyboard can be rectified. We apply a deep learning approach to extract the keyboard corners and warp the thermal image to standardized coordinates. This process is detailed in the following sections.

5.1.1 Data collection/Annotation of the images. First, we captured 1500 thermal images of a standard ISO QWERTY keyboard with ABS keycaps using an optris PI 450i (see samples in Figure 3). In each thermal image, we pressed random keys to create random heat traces on different parts of the keyboard. This was done to avoid biasing the model towards detecting keyboards that have specific patterns of heat traces or those with no traces at all. The direction of the camera and the distance from the keyboard were randomly changed in each image to evaluate ThermoSecure, and in particular phase 1, against challenging and realistic scenarios. The thermal images were annotated using Labelme package⁵. The annotation process includes labeling, setting the coordinates of the keyboards in the thermal image and saving these information in JSON files (see example in Figure 2).

The thermal images along with their corresponding annotation files were then randomly divided to three sets:

- (1) Train Set: 1300 images that were used to train the model.
- (2) Validation Set: 170 images that were used during the training to tune the parameters of the model.
- (3) Test Set: 30 thermal images that were used to test how well the model performs on unseen (new) data.

⁵<https://github.com/wkentaro/labelme>

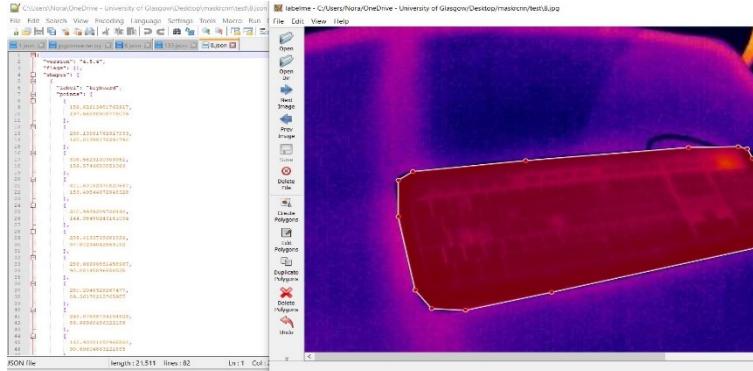


Fig. 2. An example of the labeling process in which each point on the keyboard edges was annotated.

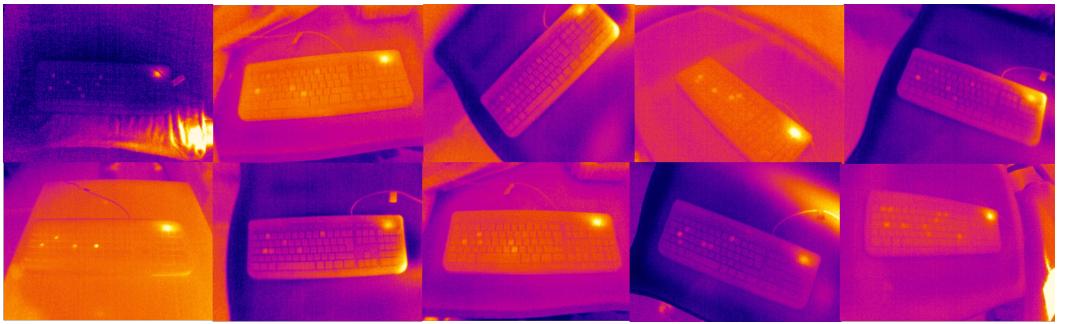


Fig. 3. Samples of the collected thermal images used for training the model. We took thermal images of a standard ISO QWERTY keyboard with ABS keycaps from random angles and distances. In each image, heat traces were produced on the keyboard by pressing random keys.

5.1.2 Pre-processing/Data augmentation. To make the thermal traces in the image more prominent, we need to reduce high-frequency noise and standardize the contrast. Noise can arise as an effect of different factors such as having an insufficient light levels during image acquisition and the interference in the electronic circuits inside the camera that increase the thermal energy of heat inside the sensors.

As illustrated in Figure 4, in addition to applying Contrast Limited Adaptive Histogram Equalization (CLAHE) [32] to the image, the noise filtering process in a previous study [1] was adapted as follows:

- (1) Applying a 5x5 median filter.
- (2) Converting the images from RGB color to grayscale.
- (3) Reapplying the median filter in (1) for enhanced noise reduction.

Furthermore, data augmentation was applied to increase the training data without capturing new images. We did this by applying random rotation, padding, and horizontal flipping (see Figure 5).

5.1.3 Training. Our model uses Mask RCNN [12] for object detection. Mask RCNN is a framework for Image Segmentation tasks that uses the ResNet architectures (101, 50) to extract features from the images. In order to achieve optimal results with less training time on our relatively small dataset

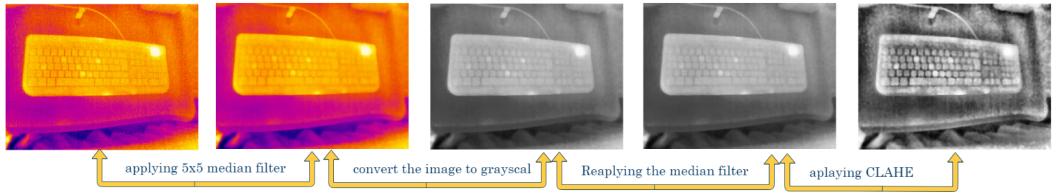


Fig. 4. Sample thermal image undergoing the pre-processing stage for denoising. We removed the noise by applying a 5x5 median filter, converting the image to grayscale, reapplying the same filter again, and then enhancing it using Contrast limited adaptive histogram equalization (CLAHE) [32].

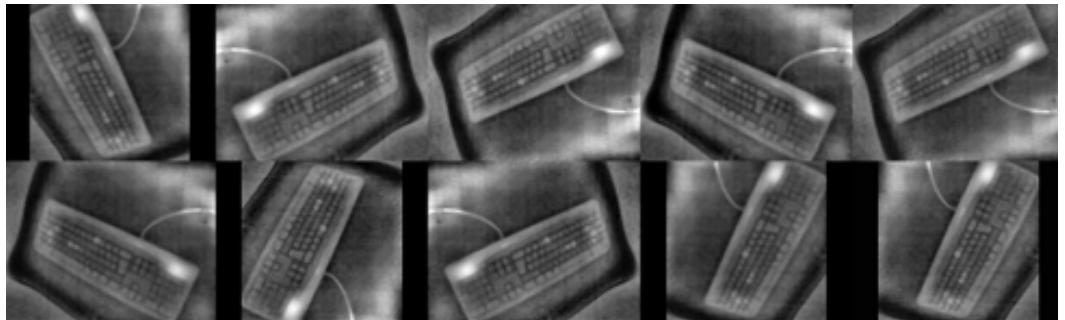


Fig. 5. An example of an augmented thermal image

we follow the transfer learning concept. A COCO⁶ model which is pre-trained on classifying and detecting more than 800K objects was used as a feature extractor. By choosing the optimal hyper-parameters from several trials, the model is fine-tuned on our custom dataset.

5.1.4 Detection. The model loads the best trained weights to detect an object that belongs to a class in a given image. In our case, the class is a keyboard. This returns a dictionary that includes: class name (i.e., a keyboard or an undefined object), bounding box coordinates of the detected object, mask information that include the coordinates of every pixel within the bounding box, and a prediction score of how confident the model is that the detected object is from a certain class. The bounding box coordinates of the detected keyboard are then saved to be used in the next phase.

5.2 Phase 2: Obtaining the key presses

The bounding box coordinates obtained from the previous phase were used to approximate the location of the keys using the following method:

- (1) The Rotated Bounding Box (Rbbox) was calculated to be used instead of the bounding box (BBox) obtained from the Mask RCNN Network (see Figure 6).
- (2) The Rbbox is then used to locate the four corners of the detected Keyboard.
- (3) We assume the target keyboard has the typical layout for most keyboards: 8 regions of interest, 6 rows of keys and 2 areas without keys at the top and bottom of the keyboard.
- (4) The coordinates that represent the start and the end of each region of interest is calculated by adding or subtracting displacement values unique to every image. The direction and the amount of displacement needed in the 2D axis is determined based on the layout of the

⁶<https://cocodataset.org/>

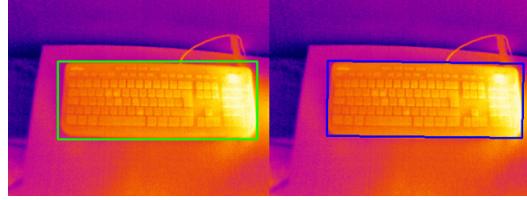


Fig. 6. The BBOX is calculated with no regards to the rotation of the object (left). Thus, we used the RBBox as shown in the right image instead. The RBBox mainly depends on the mask coordinates that include all of the data points that were detected as part of the keyboard. The RBBox was calculated by estimating the corner points of the mask.

keyboard; in particular on the height, width and the rotation of the keyboard which was obtained using the RBBox.

- (5) The start and end of each region, the region length and the displacement values are used to estimate the box coordinates of the keys (see examples in Figure 7).

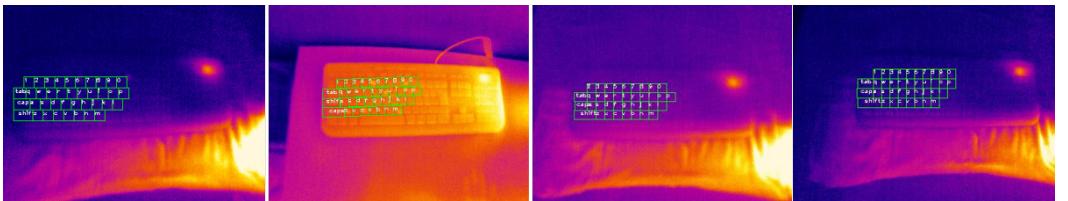


Fig. 7. Example of Keys that were located using the method described in Section 5.2.

5.2.1 Blob Detection. The last described step of the second phase provided us with box coordinates for each key. We then use this information along with blob detection to obtain the entries of the password: a blob is a group of connected pixels in a binary image, while blob detection is the process of detecting regions that have different properties such as shape or brightness. To perform blob detection the following steps are necessary:

- (1) We apply a threshold to convert the image to a binary image, i.e., each pixel value greater than 200 was set to zero while preserving the pixel values of the background.
- (2) We then apply a morphological closing to fill (i.e., close) any small holes/black points in the foreground objects.
- (3) The final step is finding and extracting the contours of the blobs.

After performing blob detection, we use the box coordinates of the keys to find all the keys that have contour points within their bounding box as shown in Figure 8.

5.2.2 K-mean clustering. K-mean clustering is a machine learning algorithm that is used to cluster data points into different clusters based on how similar they are. In our model, K-mean clustering was employed to analyze the temperature values for each of the detected keys in order to detect multiple clusters (presses) if any. The mean of an idle (unpressed) key was used to remove the temperatures of data points with a regular temperature. To choose the optimal value of k, we used the elbow method. Namely, we ran k-means several times and incremented k every iteration. We then selected the value at the “elbow” start of the linear increase: K=2 (see Figure 9).

We ran the algorithm again with the appropriate k value and assigned each point to its cluster. For example the letter N now is divided into N1 and N2 with different temperature values (mean,

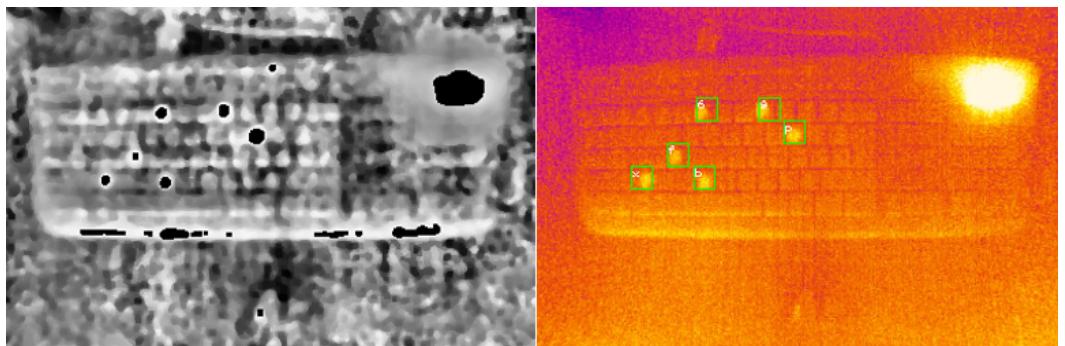


Fig. 8. The figure shows an example of detected blobs that represent pressed keys (left) and their bounding boxes (right).

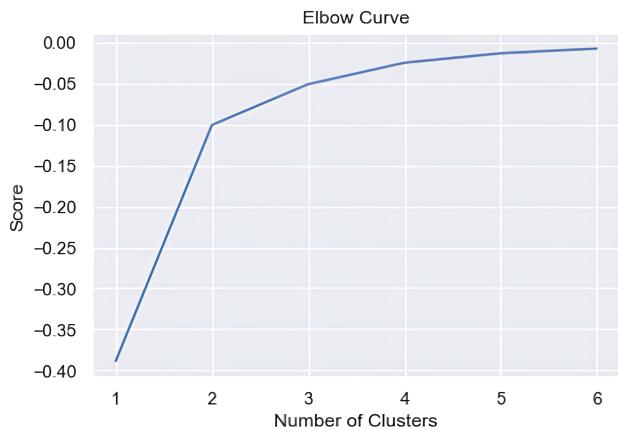


Fig. 9. We selected the value of K at the “elbow” start of the linear increase. In that case we selected K=2.

max, min) and (x,y) coordinates. The (x,y) values obtained from the blob detection were then used to visualize the different clusters.

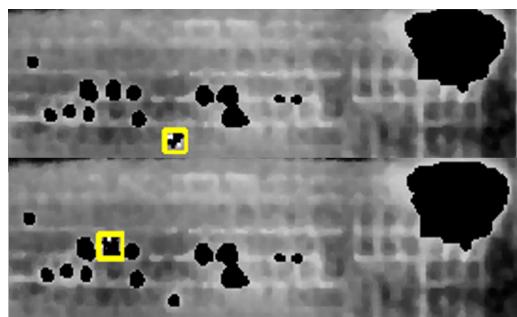


Fig. 10. The figure shows examples of different clusters (Multiple-presses).

State\Key	6	f	b	x	9	p
1	0.029392	0.02839	0.027722	0.027555	0.027054	0.02672
2	0.029392	0.063392	0.062592	0.062392	0.061792	0.061392
3	0.029392	0.063392	0.104892	0.104642	0.103892	0.103392
4	0.029392	0.063392	0.104892	0.159342	0.158352	0.157692
5	0.029392	0.063392	0.104892	0.159342	0.240342	0.239342
6	0.029392	0.063392	0.104892	0.159342	0.240342	0.400342

Table 1. The password in this example is generated by calculating the probabilities of each key in different states.

5.3 Phase 3: Determining the order of key presses

Using the coordinates of the detected password keys obtained from the previous phase, the temperature data for each thermal image is used to extract the mean, minimum and maximum temperatures for each detected key. Each key's (mean, max, min) triple is averaged, then used to arrange the keys in the appropriate order to obtain the correct password.

5.3.1 Distinguishing Authentication and Interaction. As explained in our threat model, the attacker needs to distinguish which heat traces result from entering the password, and which of those result from entering the username. The temperature values of the obtained keys were analyzed to find a temperature transition threshold. This can be done by comparing the transition value between each consecutive keys; the highest transition is then assumed to be the one between the last entry in the username and the first entry in the password.

5.3.2 Determining the order of entries. After removing the keys which are part of the username, the temperature data for each key is used as a transition probability to move between different states (keys). Using the transition probability and the state probability we calculated the probability of each key at a different state as shown in Figures 11 and ??.

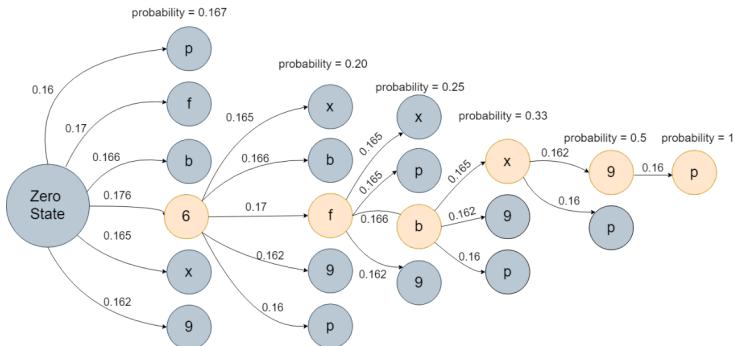


Fig. 11. Example of a possible path to determine the password. For example, the highest probability for the first key is 6, followed by f, b, x, 9 and p.

To infer the password order, we developed a symbol-level model that uses the thermal probabilities vector for each key at a given state to generate the most likely passwords by predicting one symbol at a time (see Figure 12). The model will generate a probability distribution over all of the possible symbols in the password sequence, having done that we will need a decoding algorithm to convert

the probabilistic output into a readable form. To do this we used the beam search decoding algorithm. The beam search decoder considers the top-k sequences with the highest probabilities. Thus, the beam search decoder generates k-different output sequences, where we will have a high chance of obtaining the correct sequence in the top-k sequences (see Sample output 1).

```

1 ('Password : ', '6fbx9p', 'with a probability: ' , 0.40740188103110087)
2 ('Password : ', '6fxx9p', 'with a probability: ' , 0.4071952200350172)
3 ('Password : ', 'ffbx9p', 'with a probability: ' , 0.40670664210825924)
4 ('Password : ', '6fb99p', 'with a probability: ' , 0.40659163972726375)
5 ('Password : ', '6f9x9p', 'with a probability: ' , 0.4065814008654436)
6 ('Password : ', 'ffxx9p', 'with a probability: ' , 0.40649998111217556)
7 ('Password : ', '6fx99p', 'with a probability: ' , 0.40638497873118007)
8 ('Password : ', '6bbx9p', 'with a probability: ' , 0.4059259564615335)

```

Listing 1. Sample output from ThermoSecure outlining the possible passwords and the probability of each to be the correct password.

6 USER STUDIES

We conducted two user studies to address the following research questions:

- RQ1 Does the length of a password affect the feasibility of a thermal attack using ThermoSecure?
- RQ2 Does the age of the heat trace has an impact of the effectiveness of a thermal attack?
- RQ3 Does the way people type makes them more vulnerable to thermal attacks?
- RQ4 Does the type of keycaps have an impact on the feasibility of thermal attacks?

We address RQs 1-3 in user study I (Section 6.1) and RQ4 in user study II (Section 6.2).

Participation in both studies was voluntary, and participants were not compensated for their participation. Due to the situation with Coronavirus (COVID-19) at the time, participants were chosen through convenience and snowball sampling. The studies were carried out in accordance with local health and safety regulations as laid out by the Saudi Ministry of Health⁷, which included the use of masks and disinfection of all touched equipment upon the completion of the study tasks by each participant. We also followed University of Glasgow's guidelines for conducting user studies during COVID-19. Both studies took place in different locations and at different times, including participants' homes and offices in the cities of Jeddah, Mecca, and Taif in Saudi Arabia).

6.1 Study I: Effect of Password Length, Age of the Heat Trace, and Typing Behavior

This study aims to evaluate ThermoSecure's effectiveness in retrieving passwords of different lengths when entered on keyboards, and how this is impacted by the age of the heat trace and user's typing behavior. To this end, we first collected a dataset of a) new thermal images of keyboards, taken after participants have authenticated by entering a username and a password, paired with b) usage logs of the keystroke dynamics while authenticating. Apart from evaluating ThermoSecure, the analysis can reveal whether certain users groups or typing behaviors are more vulnerable to thermal attacks, and general insights about factors that impact the effectiveness of thermal attacks. The first part of the study follows a within-subject design and aims to investigate the impact of two independent variables on the success of thermal attacks:

- (1) **Password Length:** We covered the following conditions: 6-symbol passwords (resembling weak passwords), 8-symbol passwords (oftentimes used as the minimum length for strong passwords), and 12-16 symbol passwords. For the last one, we used passphrases as users are unlikely to memorize passwords of that length if they consist of random characters [15].

⁷<https://www.moh.gov.sa/en/Pages/Default.aspx>

All the symbols used to formulate the passwords are alphanumeric passwords consisting of digits and uppercase/lowercase characters. Moreover, some of these passwords have special characters that must be entered while holding down the shift key. For easier referencing, we will refer to these passwords hereafter as **short**, **medium** and **long passwords**.

- (2) **Heat trace age:** We took thermal images of the keyboard at **20**, **30** and **60** seconds after authentication.

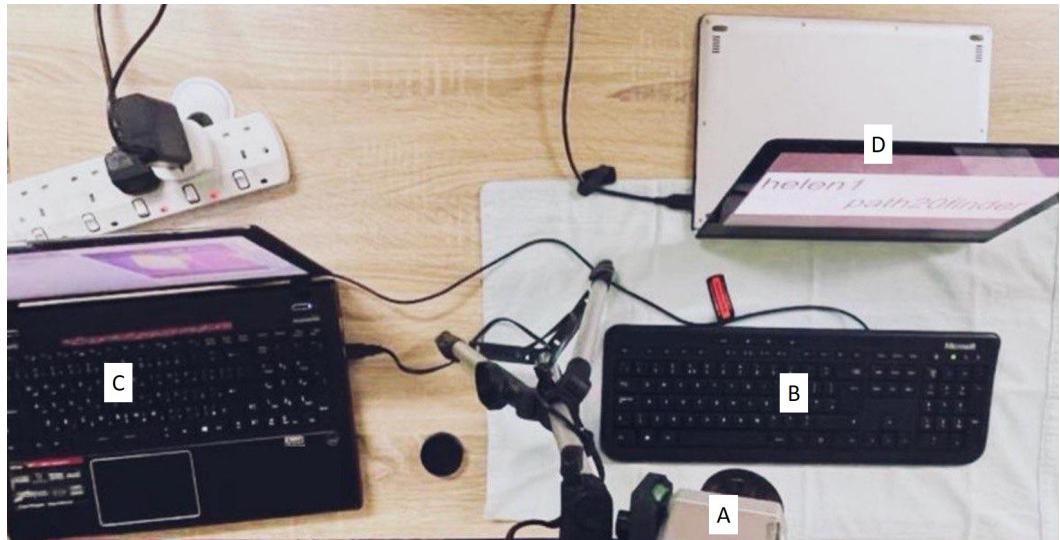


Fig. 12. Snapshot of a session from the study: The thermal camera (A) was placed on a tripod where the height and the distance from the keyboard (B) were different at every session. The participant was shown the content they should enter on laptop1 (D), while laptop2 (C) was used by the experimenter to store the thermal images and log the keystroke dynamics.

6.1.1 Apparatus. Figure 12 shows the setup of our experiment. Participants typed on a Microsoft Wired ISO Keyboard 600 (QWERTY) with ABS keycaps, while thermal images were taken using an Optris pi 450 (764px × 480px, 80 Hz, 40mK NETD, -20°C to 100°C). The thermal camera produced a 16-bit color video from which we took snapshots in the timestamps indicated above. The Optris API generated a csv data file that shows the temperature at each pixel in the thermal image. The thermal camera was mounted on a tripod to ensure the entire keyboard is captured. Two tripods⁸ (Figures 12 and 15) were used in the user studies, with the camera placed at random distances ranging from 50cm to 121cm and 60cm to 90cm. To simulate realistic scenarios and evaluate our keyboard localization model, the height of the tripod and the distance from the keyboard were randomly changed every time the participant authenticated.

Both the keyboard and thermal camera were connected to a laptop that stored the thermal data and logged the keystrokes. A second laptop was used to show the participant which username/password they should enter next.

6.1.2 Participant and Procedure. 21 participants (12 female and 9 male) were invited to participate in the study. Due to the Covid-19 pandemic, the materials presented to the participants were restricted to electronic materials only. Each participant was presented with an information sheet, a

⁸<https://amzn.eu/d/6io8vY1>

consent form, and a task sheet. Before reading the task sheet the participants were asked to read and sign the consent form digitally. The task sheet provided the participants with instructions to complete the task. The participants were asked to complete 4 tasks. In each task, the participant entered a username followed by a password. The username and password to be entered were displayed to the participant on a screen (see Figure 12). Participants spent 1.00 to 1.30 min on average in each task and were asked to wait 4 to 5 min between tasks to ensure that the heat traces from the previous task have faded away. While performing the tasks, the keyboard was video recorded by a thermal camera. Approval from our ethics committee for this experiment was received prior to conducting it.

6.1.3 Collected Data. We collected two types of data: 1) thermal images and 2) behavioral typing data. The behavioral typing metrics are essential to investigate if typing behavior has any impact on susceptibility to thermal attacks. Linking this data with the thermal images can be very useful for further analysis by the Human-Computer Interaction and Security research communities. We collected:

- (1) **Thermal images** of the heat traces left after completing each task. Thermal images were taken 20, 30 and 60 seconds after authentication.
- (2) **Behavioral typing data:** We collected the following for each key entry: a) Key Press time, b) Key-release time, c) Key-press duration, d) Latency (time between releasing the key and pressing the following key), and e) Flight time (time between pressing two consecutive keys).

6.1.4 Measuring the Accuracy of Attacks. By analyzing the thermal images using ThermoSecure (see Section 5 for the detailed steps/phases), we were able to obtain a set of passwords. Inspired by prior work in thermal attacks [1, 3, 4], we calculated the similarity between the password produced by ThermoSecure and the actual password using the Levenshtein Distance [18], which is commonly used by the user-centered security community to measure accuracy of attacks against passwords [20, 21, 28]. The Levenshtein distance is a metric that measures the minimum number of modification (addition or deletion) needed to map a word to a different one. We then convert the Levenshtein distance to a percentage to allow comparing guesses against passwords of different length. The result of the mapping is the accuracy of the guess against the actual password as a percentage. Furthermore, the behavioral key metrics data along with the notes collected during the study were used to objectively assign each participant to the appropriate typing behavior type. Figure 13 shows a summary of how the age of the heat trace and password length affect attack accuracy.

Distinguishing interaction and authentication: The flight time and the latency values revealed that there is a notable time difference between entering the username and entering the password which led us to further investigate the temperature data for each key to seek a temperature transition threshold. We used this to distinguish between entries that were part of a username and entries that were password. This was done by comparing the transition value between the alphanumeric keys within the password and the username with the transition value between the last key of the username and first key of the password.

6.1.5 Results. As we have two independent variables, we analyzed the data using a two-way repeated measures ANOVA. We applied Greenhouse-Geisser correction due to the violation of Mauchly's test of sphericity. The analysis revealed that there was no statistically significant two-way interaction between the password length and the age of the heat trace, $F(2.7, 54.5) = 2.645, p = 0.1$. But significant main effects were found as we report next: **Effect of Password Length on Attack Accuracy:** Observing the mean accuracy values of the predicted passwords suggests that short and medium password are more prone to be cracked using thermal attacks (92% and 80%

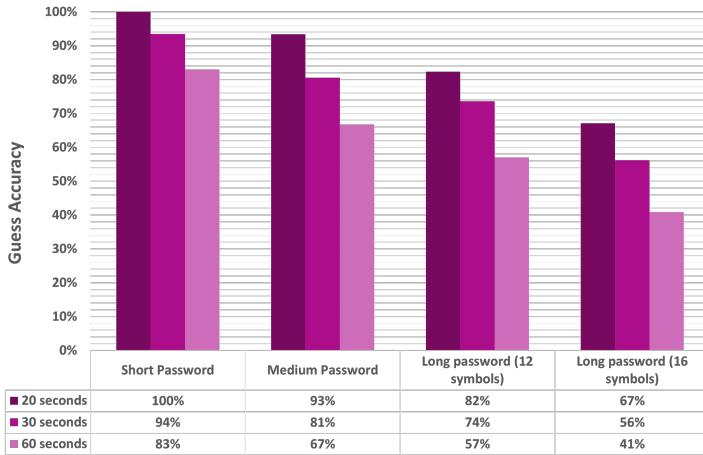


Fig. 13. Impact of the heat trace age and the length of the password on the accuracy of the guess. The figure shows how accuracy of attacks drop as more passwords become longer, and also as the time period between entry and taking the thermal image increases.

respectively) compared to long passwords (12-16 symbols) where the average accuracy was 63% (71% for 12-symbol passwords and 55% for 16-symbols passwords).

Furthermore, a repeated-measures ANOVA revealed a significant main effect of length of password on accuracy of the guess $F(2.338, 46.77) = 275.118, p < 0.05$. Post-hoc pairwise comparisons with Bonferroni correction showed significant differences ($p < 0.05$) between all the different pairs of password length (see values in Table 2 and Figure 13).

Effect of Age of Heat traces on Attack Accuracy: The collective values of the accuracy of the guess based on the age of the heat traces suggests that we have a better chance in retrieving the password with 76-86% accuracy if the thermal images were taken within the first 30s. The accuracy of the guess then drops to 62% when we reach the 60 seconds mark. A repeated-measures ANOVA was run to determine the effect of time (age of heat trace) on the accuracy of the guess and revealed a significant main effect of the age of heat trace on accuracy of the guess $F(1.524, 30.49) = 300.146, p < 0.05$. Post-hoc pairwise comparisons with Bonferroni correction showed significant differences ($p < 0.05$) between all the different pairs of the heat trace ages (20s, 30s and 60s).

Effect of Typing behavior on Attack Accuracy In addition to the two within-subjects variables above, we additionally investigated whether the **typing behavior** impacts the success of thermal attacks. Note that we did not control for typing behavior when recruiting participants but rather analyzed it posthoc. After analyzing the typing behavior of our participants by observing the way they typed and the behavioral typing metrics, we found that our participants' typing behavior can be classified into two types:

- (1) **Fast typists:** these typists lightly touch the keyboard's keys. The average key press duration is less than 200ms and the latency is less than 1000ms.
- (2) **Hunt-and-peck typists:** these typists spend more time looking for the key to press (>1000ms latency) and their key presses are also longer (>200ms).

Out of our 21 participants, 11 were fast typists and 10 were hunt-and-peck typists. There were two outliers in the data, as assessed by inspection of a boxplot, which were removed: (P20 (77% guess accuracy, Hunt-and-Peck) and P6 (60% guess accuracy, Fast typist)). Thus, we treated the typing behavior as a between-subject factor with the two typing methods as its conditions. An

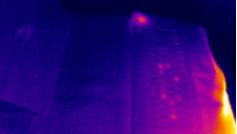
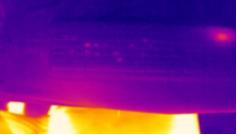
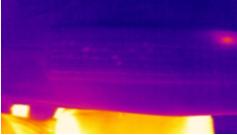
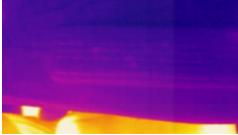
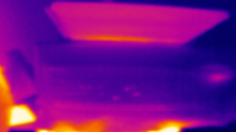
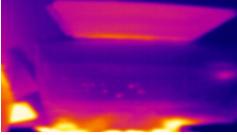
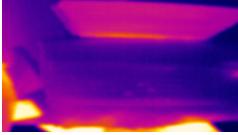
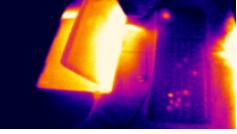
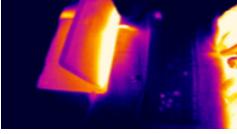
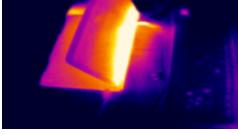
	20 seconds after authentication	30 seconds after authentication	60 seconds after authentication	Average accuracy
Long Passwords (16 symbols)				55%
	Guess Accuracy : 67%	Guess Accuracy : 56%	Guess Accuracy : 41%	
Long Passwords (12 symbols)				71%
	Guess Accuracy : 82%	Guess Accuracy : 74%	Guess Accuracy : 57%	
Medium Passwords				80%
	Guess Accuracy : 93%	Guess Accuracy : 81%	Guess Accuracy : 67%	
Short Passwords				92%
	Guess Accuracy : 100%	Guess Accuracy : 94%	Guess Accuracy : 83%	
Average accuracy	86%	76%	62%	

Table 2. Accuracy of guess results from Study I alongside a sample thermal images from different angles and distances to the keyboard. Shorter passwords are more vulnerable to thermal attacks and the sooner the thermal image is taken the more effective it is. Still, even 60 seconds after authentication up to 6 characters in a 16-character password are determined in the correct position and order through ThermoSecure.

independent-samples t-test was run to determine if there were differences in the attack accuracy between the two typing behaviors. Accurate guess scores for each type of Typing behavior were normally distributed, as assessed by Shapiro-Wilk's test ($p > .05$). and there was homogeneity of variances, as assessed by Levene's test for equality of variances ($p = .075$). There was a statistically significant difference in average attack accuracy between the two typing behaviors which indicates that an attacker can obtain more accurate password guesses if the typist was a hunt and peck typist ($83\% \pm 0.01$) rather than a Fast typist ($68\% \pm 0.017$), $t(17) = -22.7$, $p < 0.05$.

Figure 14 shows how some user groups, specifically hunt and peck typists, are more vulnerable to thermal attacks than others. Figure 3 shows a detailed sample from participants P1 (hunt-and-peck typist) and P12 (fast typist).

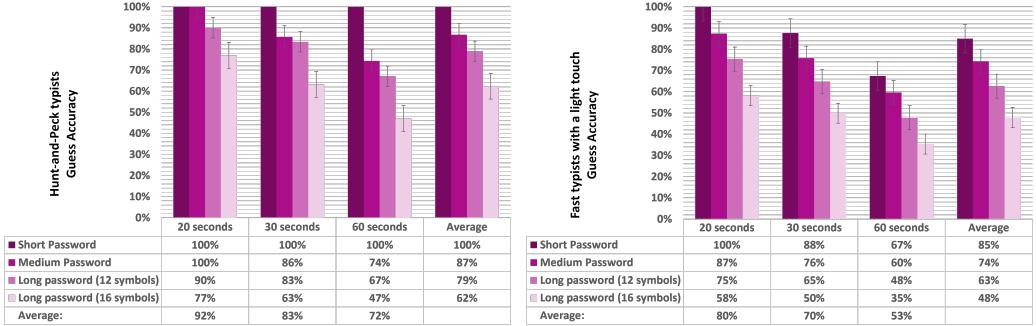


Fig. 14. Hunt-and-Peck typists (left) vs Fast typists (right). Attack accuracy is significantly higher for Hunt-and-peck typists, which means they are more vulnerable to thermal attacks. On the other hand, fast typists lightly touch the keys in comparison, resulting in less heat traces due to the short contact time.

While previous work suggested that typing behavior can impact the susceptibility to thermal attacks [13] by visually inspecting typing behavior, our work is the first to determine whether the user's typing behavior makes them vulnerable to thermal attacks by using behavioral typing metrics that can be determined in real time. This means that our findings allows future systems to preemptively determine whether the user is likely to be vulnerable to thermal attacks by monitoring their typing behavior.

6.2 Study II: Effect of Type of Keycaps

This study aims to evaluate the impact of the most commonly used materials for keyboard keycaps on ThermoSecure's effectiveness in retrieving passwords (RQ4). Other keycap styles are available, but they are much less common and more expensive. Rubber and brass keycaps, for example, can be difficult to find and purchase. This study also follows a within-subject design and has one independent variable: the **Keycap Type**, which had two conditions:

- Keycaps made of PBT plastic (Polybutylene Terephthalate).
- Keycaps made of ABS plastic (Acrylonitrile Butadiene Styrene).

6.2.1 Apparatus. Figure 15 shows the setup of our experiment. Participants typed on a Razer Huntsman Tournament Edition Gaming Keyboard (PBT keycaps) and a Razer Cynosa Lite Gaming Keyboard (ABS keycaps), while thermal images were taken using the same thermal camera and setup mentioned in Section 6.1.1. The second laptop, however, was used to randomize the order of the predefined passwords list and show the participant which password to enter next.

6.2.2 Participants and Procedure. A total of 16 participants (10 female and 6 male) were invited to take part in the study. The same documents in Section 6.1.2 were presented. Participants were

	Short Passwords			Medium Passwords			Passphrases (12 symbols)			Passphrases (16 symbols)		
	20	30	60	20	30	60	20	30	60	20	30	60
P1	100%	100%	100%	100%	86%	67%	91%	86%	80%	76%	62%	56%
P12	100%	91%	67%	86%	77%	67%	74%	67%	59%	62%	56%	43%

Table 3. Sample of results from participants P1 and P12. P1 was a hunt-and-peck typist – their short password was detected with 100% accuracy even 60 seconds after entry. In general, As the password length increased and more time passed, the accuracy decreased. But the decrease in accuracy was sharper for P12 as they were a fast typer.



Fig. 15. An example of a session from Study II: In every session, the thermal camera (F) was mounted on a tripod at varying heights and distances from the keyboards. (A) a Razer Huntsman Tournament Edition gaming keyboard (**PBT** keycaps). (B) a Razer Cynosa Lite Gaming Keyboard (**ABS** keycaps). The content participants should type was shown to the participant on Laptop1 (D), while the experimenter recorded the thermal images using Laptop2 (C) which was connected to the thermal camera.

instructed to complete three tasks on each keyboard (6 tasks in total). Each task required the participant to enter a password with different properties. To reduce learning effects, the passwords to be entered were shown to the participant in a random order. Participants spent 30 to 60 seconds on average in each task. Similar to Study I, participants of this study were asked to wait 4 to 5 minutes between tasks to ensure that the heat traces from the previous task had faded away. While the participants were performing the tasks, the thermal camera recorded the keyboards.

6.2.3 Collected Data. In this study, thermal images of the heat traces left after completing each task were collected. We captured thermal images 20, 30, and 60 seconds after authentication.

6.2.4 Measuring the Accuracy of Attacks. Using ThermoSecure (as described in Section 5), we inferred the passwords using the information gathered from the heat traces left after each task on each keyboard. As done for Study I (see 6.1.4), we measured the Levenshtein Distance between ThermoSecure’s generated password and the actual password to estimate the accuracy. The distance was then converted to a percentage to account for passwords of different lengths.

6.2.5 Results. As we need to evaluate the attack’s feasibility against two different keycap types (i.e., ABS vs. PBT), the data was analyzed using a Paired Sample T-test. Prior to the test, the data was tested for outliers and was checked for normal distribution. Boxplot inspection revealed that there were no outliers. The differences between PBT keycaps and ABS keycaps in term of the accuracy of the guess were normally distributed, as assessed by Shapiro-Wilk’s test ($p = 0.162$). the paired sample T-Test revealed that passwords entered on an ABS keyboard ($52\% \pm 16\%$) are more vulnerable to thermal attacks regardless of the length of the password or the age of the heat trace compared to the PBT keyboard ($14\% \pm 19\%$) which are more resilient to thermal attacks $t(15)= 8.124$ $d=2.03$ $p < 0.0005$. The detailed results are shown in Table 4.

	20 seconds after authentication		30 seconds after authentication		60 seconds after authentication	
	ABS	PBT	ABS	PBT	ABS	PBT
Long Passwords						
	61%	26%	50%	13%	37%	5%
Medium Passwords						
	69%	25%	60%	21%	45%	8%
Average guess accuracy	63%	20%	54%	16%	39%	6%

Table 4. Accuracy of guess results from Study II against passwords of different lengths when thermal images are taken 20, 30 and 60 seconds after authentication on keyboards with ABS (left) and PBT (right) keycaps. The results show that the type of the keycaps has a significant impact on the success of the thermal attacks (52% against ABS and 14% against PBT). We observed that regardless of the length of the password, gathering any information from a keyboard with PBT keycaps 20 seconds or longer after password entry is very difficult.

7 DISCUSSION AND FUTURE WORK

In the following, we discuss which factors impact the success of thermal attacks, and discuss ways to mitigate them in light of our results and previous work.

7.1 Factors that impact the success of thermal attacks

Our findings indicate that a variety of factors can explain why thermal attacks are more successful in some cases but may fail in others. These factors can be classified into two types: 1) factors related to the **input** such as the password length, and user typing behavior, and 2) factors related to the **interface** such as the material out of which the keycaps are made and their thermal conductivity, which in turn impact how fast heat traces decay off the keys. Our findings explain how these factors influence the feasibility of thermal attacks.

7.1.1 Input factors. The length of the password has a significant main effect on the accuracy of the guessed password. In particular, Short and medium passwords are significantly more vulnerable to thermal attacks (up to 100% attack success).

Main finding 1 : Increasing the length of the password significantly increases the resistance to thermal attacks.

This is in line with previous work on thermal attacks against different types of passwords (e.g., graphical passwords [1, 3]).

Apart from the properties of the input, another input factor is the user's typing behavior. We investigated the possible link between input behavior and thermal attacks, and found that user's

typing behavior significantly impacts the quality of information obtained from the thermal traces. As a result, there are typing behaviors that make the users more vulnerable to thermal attacks.

Main finding 2: Users who are Hunt-and-Peck typists are particularly vulnerable to thermal attacks.

Compared to the most relevant prior work [13], ThermoSecure is successful in inferring both the keys used to input the passwords and their order of entry in the majority of cases. This was done by relying on the information gained from the thermal images only without the need to integrate any other form of attacks such as dictionary attacks as done in [13].

7.1.2 Interface factors. By performing a frame-by-frame comparison at different time intervals, we find that the age of the heat trace has a significant impact on the quality of the information that can be used in performing a thermal attack.

Considering the thermodynamic nature of the surface (i.e., keyboard) the heat traces will decay over time across all the keys. Thermal images taken 20 seconds after entry provide more thermal information, which makes it easier to reveal both the entries of the password and the order of entry compared to thermal images that were taken later. Although these thermal images revealed less information about the authentication process, we can still uncover most of the password entries with accuracy 40% to 80% after 60 seconds). This significantly reduces the password space, thereby making it easier for the attacker to guess the rest of the password or use other means to uncover a smaller portion of the password. These other means could be guessing attacks, dictionary attacks (as done in [13]), smudge attacks [7] (in case of touchscreen interfaces), video-based attacks [34] or even a second thermal attack.

Main finding 3: The age of the heat trace has a significant impact on the attack's success. The more time passes, the less likely the attack will succeed.

Furthermore, Study II investigated how thermal attacks fare against two models of computer keyboards with keycaps constructed of two different plastic materials in the (i.e., ABS vs PBT). PBT keycaps were found to be more resilient to thermal attacks as heat traces fade faster due the material's lower thermal conductivity [33].

Main finding 4 :There is a link between keyboard materials and the feasibility of thermal attacks. Keyboards using ABS keycaps are more vulnerable to thermal attacks than those that use PBT keycaps.

7.2 Increasing password length

Main finding 1 indicates that longer passwords are more secure against thermal attacks. The straight forward consequence of this is to recommend that users create longer passwords (**Mitigation approach 1**). The longer the password, the more likely heat traces of the first entries to decay by the time the thermal image is taken. Additionally, the longer the password, the less pronounced the differences in temperatures at different keys will be, which makes it harder for attackers to infer the correct order of entries.

However, there are human factor challenges in creating long passwords; it is unreasonable to expect users to create and memorize different long passwords for their many accounts [5].

While many platforms recommend increasing password lengths to improve security, increasing the complexity of the password instead (e.g., by including special characters and mixing upper- and lower-case characters) can yield better results in terms of usability and security against offline attacks [27]. For this reason, we recommend using passphrases as they were shown to be more memorable [15] and can help users create longer passwords. That being said, there are mitigation strategies that could potentially be more usable, which we discuss next.

7.3 Estimating vulnerability based on user properties

Main finding 2 indicates that typing behavior can be a predictor of how vulnerable a user can be to thermal attacks. We were able to classify users into hunt-and-peck and fast typists by objectively analyzing their typing behavior. This means that future systems can leverage the typing behavior to do the same in real time, and consequently take measures to improve security against thermal attacks (**Mitigation approach 2**). For example, these users may be required to use longer passwords, or they may be asked to provide random input after entering their passwords, so that the heat traces can be distorted.

In a similar vein, a promising direction for future work is to explore if certain user groups are inherently vulnerable to thermal attacks. For example, gender [16] and age [30] impact the body and hand heat temperature. A future study could investigate if these demographic factors impact vulnerability as systems can then deploy different authentication methods for these user groups.

7.4 Thermal conductivity of surfaces

Main finding 4 indicates that PBT keycaps are more resistant to thermal attacks due to their lower thermal conductivity. Mowery et al. [23] has shown that metallic keypads are also not vulnerable to thermal attacks as they reflect the hands temperature. Previous work in human-computer interaction investigated the thermal conductivity of surfaces to exploit them for interaction [2, 24]. The idea is that if a surface can maintain enough heat traces to be captured by a thermal camera, the thermal camera can then channel this information to other systems to make said surface interactive. This concept has already been used in patents and research applications [11, 17]. These developments have several implications on thermal attacks: 1) there is ongoing work on assessing the feasibility of determining input through heat traces on different surfaces, which means that more work is needed to understand if other surfaces are vulnerable to thermal attacks, 2) as these applications and patented ideas make their ways into consumer products, the risk of thermal attacks becomes greater, and 3) another mitigation method is to use interactive surfaces that are less vulnerable to thermal attacks (**Mitigation approach 3**). For example, metal keypads should be used instead of plastic ones [23], and PBT keyboards instead of ABS ones.

7.5 Inducing heat to “erase” heat traces

Some keyboards come with backlighting to improve their aesthetics and usage in dark environments. The heat generated by backlighting in the keyboards can potentially be used to balance the temperature across the keyboard, thereby erasing the heat traces resulting from interaction (**Mitigation approach 4**). As shown in Figure 16, our pilot tests with the keyboards we tried show that this approach is slightly effective, but no significant results were found. Future work can explore how to customize backlighting to increase the decay rate of heat traces.

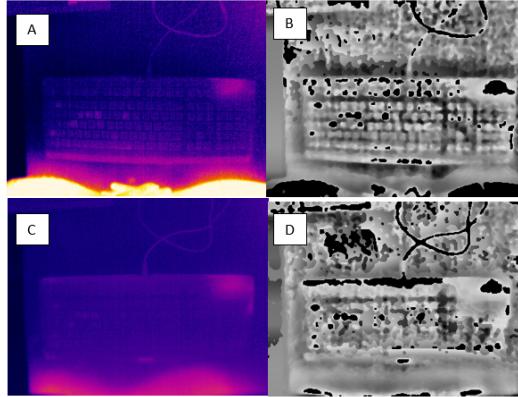


Fig. 16. An illustration of the effect of increasing the keyboard's backlight to accelerate the decay of heat traces. The thermal image **A** was captured after authenticating with a medium password (Backlight=0). **B**) The heat traces revealed by segmenting **A**. **C**) A thermal image of the same password as in **A**, but with the backlight turned up to 100. **D**) After adjusting the backlighting on the keyboard, segmenting **C** revealed that there were fewer heat traces left.

7.6 Hiding heat traces in the thermal camera's view

An alternative approach to mitigate thermal imaging attacks is by detecting interfaces in the feed of the thermal camera, and obfuscating it to prevent users from performing thermal attacks. There has been recent preliminary work in this direction [6]. However, more research is needed to offer protection against thermal attacks without significantly impacting the utility of the thermal camera.

7.7 Alternative authentication methods

“The continued domination of passwords over all other methods of end-user authentication is a major embarrassment to security researchers” [9]. While there had been predictions that passwords will cease to exist, they are pervasively integrated into systems we use in our day-to-day lives, making them difficult to replace. Nevertheless, there has been a lot of efforts in developing alternative methods that are both more secure and more usable. Some of these methods are also promising for protecting against thermal attacks.

Biometric authentication is a possible alternative authentication method that is not vulnerable to thermal attacks at the time of writing this article (**Mitigation approach 5**). Examples include physiological biometrics such as fingerprint and facial recognition, but also behavioral biometric methods that rely on keystroke dynamics [10, 22] or gaze behavior [14]. Authentication schemes that rely on eye gaze are also resistant to thermal attacks as they do not result in heat traces. See [14] for an overview of the use of gaze for implicit and explicit authentication (**Mitigation approach 6**). Finally, many of the methods that were proposed in the literature to resist smudge attacks [7], are also resilient to thermal attacks by design (**Mitigation approach 7**). Examples include smudgesafe [26], and the work of von Zezschwitz et al. [29].

While the aforementioned authentication methods are expected to be resilient to thermal attacks, they come with their own disadvantages. For example, biometric passwords are difficult to change, and the collection of biometric and gaze data by third parties has privacy implications [14]. Some of the methods proposed to resist smudge attacks employ graphical passwords, which might be not suitable for integration with existing backends, and usually suffer from low password space and/or vulnerability to shoulder surfing [25].

8 CONCLUSION

In this work, we examined the feasibility of thermal attacks on commonly used computer keyboards. We presented ThermoSecure, a system that analyzes thermal images to estimate user input. We also presented the first publicly available dataset of 1500 thermal images of keyboards. Through two user studies, we found that ThermoSecure reveals the vast majority of passwords within 20 seconds (86%) and slightly less in 30 seconds (76%). Accuracy drops significantly after 60 seconds (62%). Accuracy also decreases as passwords become longer, and as users type fast rather than using the hunt-and-peck approach. We found through a second study that PBT keyboards are significantly more secure against thermal attacks compared to ABS keyboards. We concluded with recommendations for mitigating thermal attacks and directions for future work in this area.

ACKNOWLEDGMENTS

This work was supported by the Royal Society of Edinburgh (RSE award number 65040), the EPSRC (EP/V008870/1), and the PETRAS National Centre of Excellence for IoT Systems Cybersecurity, which is also funded by the EPSRC (EP/S035362/1).

REFERENCES

- [1] Yomna Abdelrahman, Mohamed Khamis, Stefan Schneegass, and Florian Alt. 2017. Stay Cool! Understanding Thermal Attacks on Mobile-Based User Authentication. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 3751–3763. <https://doi.org/10.1145/3025453.3025461>
- [2] Yomna Abdelrahman, Alireza Sahami Shirazi, Niels Henze, and Albrecht Schmidt. 2015. *Investigation of Material Properties for Thermal Imaging-Based Interaction*. Association for Computing Machinery, New York, NY, USA, 15–18. <https://doi.org/10.1145/2702123.2702290>
- [3] Yasmeen Abdrabou, Yomna Abdelrahman, Ahmed Ayman, Amr Elmougy, and Mohamed Khamis. 2020. Are Thermal Attacks Ubiquitous? When Non-Expert Attackers Use Off the Shelf Thermal Cameras. In *Proceedings of the International Conference on Advanced Visual Interfaces* (Salerno, Italy) (AVI '20). Association for Computing Machinery, New York, NY, USA, Article 47, 5 pages. <https://doi.org/10.1145/3399715.3399819>
- [4] Yasmeen Abdrabou, Reem Hateri, Yomna Abdelrahman, Amr Elmougy, and Mohamed Khamis. 2021. Passphrases Beat Thermal Attacks: Evaluating Text Input Characteristics Against Thermal Attacks on Laptops and Smartphones. In *Human-Computer Interaction – INTERACT 2021*, Carmelo Ardito, Rosa Lanzilotti, Alessio Malizia, Helen Petrie, Antonio Piccinno, Giuseppe Desolda, and Kori Inkpen (Eds.). Springer International Publishing, Cham, 712–721.
- [5] Anne Adams and Martina Angela Sasse. 1999. Users Are Not the Enemy. *Commun. ACM* 42, 12 (Dec. 1999), 40–46. <https://doi.org/10.1145/322796.322806>
- [6] Norah Alotaibi, Md Shafiqul Islam, Karola Marky, and Mohamed Khamis. 2022. Advanced Techniques for Preventing Thermal Imaging Attacks. In *27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) (IUI '22 Companion). Association for Computing Machinery, New York, NY, USA, 18–21. <https://doi.org/10.1145/3490100.3516472>
- [7] Adam J. Aviv, Katherine Gibson, Evan Mossop, Matt Blaze, and Jonathan M. Smith. 2010. Smudge Attacks on Smartphone Touch Screens. In *Proceedings of the 4th USENIX Conference on Offensive Technologies* (Washington, DC) (WOOT'10). USENIX Association, USA, 1–7.
- [8] Paul Bekaert, Norah Alotaibi, Florian Mathis, Nina Gerber, Aiden Rafferty, Mohamed Khamis, and Karola Marky. 2022. Are Thermal Attacks a Realistic Threat? Investigating the Preconditions of Thermal Attacks in Users' Daily Lives.. In *Proceedings of the 12th Nordic Conference on Human-Computer Interaction* (NordiCHI '22). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3546155.3546706>
- [9] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. 2012. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. In *2012 IEEE Symposium on Security and Privacy*. 553–567. <https://doi.org/10.1109/SP.2012.44>
- [10] Alexander De Luca, Alina Hang, Frederik Brudy, Christian Lindner, and Heinrich Hussmann. 2012. *Touch Me Once and I Know It's You! Implicit Authentication Based on Touch Screen Patterns*. Association for Computing Machinery, New York, NY, USA, 987–996. <https://doi.org/10.1145/2207676.2208544>
- [11] Markus Funk, Stefan Schneegass, Michael Behringer, Niels Henze, and Albrecht Schmidt. 2015. An Interactive Curtain for Media Usage in the Shower. In *Proceedings of the 4th International Symposium on Pervasive Displays*

- (Saarbruecken, Germany) (*PerDis '15*). Association for Computing Machinery, New York, NY, USA, 225–231. <https://doi.org/10.1145/2757710.2757713>
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*. 2961–2969.
- [13] Tyler Kaczmarek, Ercan Ozturk, and Gene Tsudik. 2019. Theremanator: Thermal Residue-Based Post Factum Attacks on Keyboard Data Entry. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security* (Auckland, New Zealand) (*Asia CCS '19*). Association for Computing Machinery, New York, NY, USA, 586–593. <https://doi.org/10.1145/3321705.3329846>
- [14] Christina Katsini, Yasmeen Abdabou, George E. Raptis, Mohamed Khamis, and Florian Alt. 2020. *The Role of Eye Gaze in Security and Privacy Applications: Survey and Future HCI Research Directions*. Association for Computing Machinery, New York, NY, USA, 1–21. <https://doi.org/10.1145/3313831.3376840>
- [15] Mark Keith, Benjamin Shao, and Paul John Steinbart. 2007. The usability of passphrases for authentication: An empirical field study. *International Journal of Human-Computer Studies* 65, 1 (2007), 17–28. <https://doi.org/10.1016/j.ijhcs.2006.08.005> Information security in the knowledge economy.
- [16] Han Kim, Clark Richardson, Jeanette Roberts, Lisa Gren, and Joseph L Lyon. 1998. Cold hands, warm heart. *The Lancet* 351, 9114 (1998), 1492.
- [17] Daniel Kurz. 2020. Method and device for detecting a touch between a first object and a second object. US Patent 10,877,605.
- [18] Vladimir I Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, Vol. 10. 707–710.
- [19] Duo Li, Xiao-Ping Zhang, Menghan Hu, Guangtao Zhai, and Xiaokang Yang. 2018. Physical password breaking via thermal sequence analysis. *IEEE Transactions on Information Forensics and Security* 14, 5 (2018), 1142–1154.
- [20] Florian Mathis, Kami Vaniea, and Mohamed Khamis. 2021. *RepliCueAuth: Validating the Use of a Lab-Based Virtual Reality Setup for Evaluating Authentication Systems*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3411764.3445478>
- [21] Florian Mathis, John H. Williamson, Kami Vaniea, and Mohamed Khamis. 2021. Fast and Secure Authentication in Virtual Reality Using Coordinated 3D Manipulation and Pointing. *ACM Trans. Comput.-Hum. Interact.* 28, 1, Article 6 (Jan. 2021), 44 pages. <https://doi.org/10.1145/3428121>
- [22] Fabian Monrose and Aviel Rubin. 1997. Authentication via Keystroke Dynamics. In *Proceedings of the 4th ACM Conference on Computer and Communications Security* (Zurich, Switzerland) (*CCS '97*). Association for Computing Machinery, New York, NY, USA, 48–56. <https://doi.org/10.1145/266420.266434>
- [23] Keaton Mowery, Sarah Meiklejohn, and Stefan Savage. 2011. Heat of the moment: Characterizing the efficacy of thermal camera-based attacks. In *Proceedings of the 5th USENIX conference on Offensive technologies*. 6–6.
- [24] Alireza Sahami Shirazi, Yomna Abdelrahman, Niels Henze, Stefan Schneegass, Mohammadreza Khalilbeigi, and Albrecht Schmidt. 2014. Exploiting Thermal Reflection for Interactive Systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) (*CHI '14*). Association for Computing Machinery, New York, NY, USA, 3483–3492. <https://doi.org/10.1145/2556288.2557208>
- [25] Florian Schaub, Marcel Walch, Bastian Königs, and Michael Weber. 2013. Exploring the Design Space of Graphical Passwords on Smartphones. In *Proceedings of the Ninth Symposium on Usable Privacy and Security* (Newcastle, United Kingdom) (*SOUPS '13*). Association for Computing Machinery, New York, NY, USA, Article 11, 14 pages. <https://doi.org/10.1145/2501604.2501615>
- [26] Stefan Schneegass, Frank Steinle, Andreas Bulling, Florian Alt, and Albrecht Schmidt. 2014. SmudgeSafe: Geometric Image Transformations for Smudge-Resistant User Authentication. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Seattle, Washington) (*UbiComp '14*). Association for Computing Machinery, New York, NY, USA, 775–786. <https://doi.org/10.1145/2632048.2636090>
- [27] Richard Shay, Saranga Komanduri, Adam L. Duritz, Phillip (Seyoung) Huh, Michelle L. Mazurek, Sean M. Segreti, Blase Ur, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2016. Designing Password Policies for Strength and Usability. *ACM Trans. Inf. Syst. Secur.* 18, 4, Article 13 (May 2016), 34 pages. <https://doi.org/10.1145/2891411>
- [28] Emanuel von Zezschwitz, Alexander De Luca, Bruno Brunkow, and Heinrich Hussmann. 2015. *SwiPIN: Fast and Secure PIN-Entry on Smartphones*. Association for Computing Machinery, New York, NY, USA, 1403–1406. <https://doi.org/10.1145/2702123.2702212>
- [29] Emanuel von Zezschwitz, Anton Koslow, Alexander De Luca, and Heinrich Hussmann. 2013. Making Graphic-Based Authentication Secure against Smudge Attacks. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces* (Santa Monica, California, USA) (*IUI '13*). Association for Computing Machinery, New York, NY, USA, 277–286. <https://doi.org/10.1145/2449396.2449432>
- [30] Jill Waalen and Joel N. Buxbaum. 2011. Is Older Colder or Colder Older? The Association of Age With Body Temperature in 18,630 Individuals. *The Journals of Gerontology: Series A* 66A, 5 (02 2011), 487–492. <https://doi.org/10.1093/gerona/glr011>

glr001 arXiv:<https://academic.oup.com/biomedgerontology/article-pdf/66A/5/487/1529621/glr001.pdf>

- [31] Wojciech Wodo and Lucjan Hanzlik. 2016. Thermal Imaging Attacks on Keypad Security Systems.. In *SECRYPT*. 458–464.
- [32] Garima Yadav, Saurabh Maheshwari, and Anjali Agarwal. 2014. Contrast limited adaptive histogram equalization based enhancement for real time video system. In *2014 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*. 2392–2397. <https://doi.org/10.1109/ICACCI.2014.6968381>
- [33] Yong Yang. 2007. *Thermal Conductivity*. Springer New York, New York, NY, 155–163. https://doi.org/10.1007/978-0-387-69002-5_10
- [34] Guixin Ye, Zhanyong Tang, Dingyi Fang, Xiaojiang Chen, Willy Wolff, Adam J. Aviv, and Zheng Wang. 2018. A Video-Based Attack for Android Pattern Lock. *ACM Trans. Priv. Secur.* 21, 4, Article 19 (July 2018), 31 pages. <https://doi.org/10.1145/3230740>