Stay Cool! Understanding Thermal Attacks on Mobile-based User Authentication

Yomna Abdelrahman[1,*], **Mohamed Khamis**[2,*], **Stefan Schneegass**[1], **Florian Alt**[2]
[1]University of Stuttgart, HCI Group, Germany, {firstname.lastname}@vis.uni-stuttgart.de
[2]LMU Munich, Ubiquitous Interactive Systems Group, Germany, {firstname.lastname}@ifi.lmu.de
*** contributed equally**

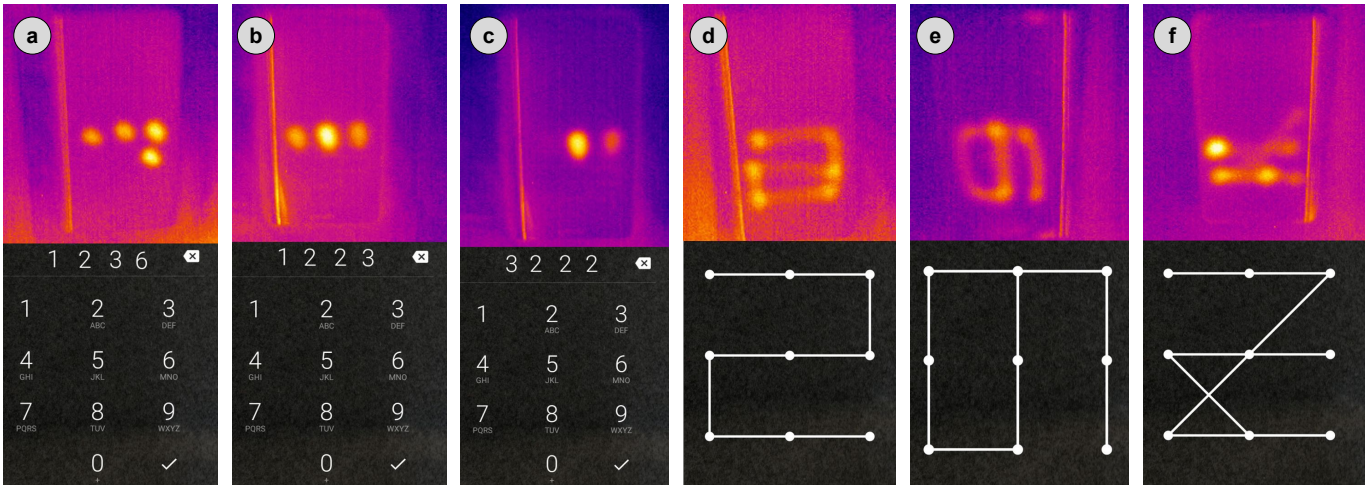

Figure 1: In this work we investigate thermal attacks against PINs and patterns on mobile devices. After entering PINs (a–c) or patterns (d–f) on a touch screen, a heat trace remains on the screen and can be made visible via thermal imaging.

ABSTRACT

PINs and patterns remain among the most widely used knowledge-based authentication schemes. As thermal cameras become ubiquitous and affordable, we foresee a new form of threat to user privacy on mobile devices. Thermal cameras allow performing thermal attacks, where heat traces, resulting from authentication, can be used to reconstruct passwords. In this work we investigate in details the viability of exploiting thermal imaging to infer PINs and patterns on mobile devices. We present a study (N=18) where we evaluated how properties of PINs and patterns influence their thermal attacks resistance. We found that thermal attacks are indeed viable on mobile devices; overlapping patterns significantly decrease successful thermal attack rate from 100% to 16.67%, while PINs remain vulnerable (>72% success rate) even with duplicate digits. We conclude by recommendations for users and designers of authentication schemes on how to resist thermal attacks.

ACM Classification Keywords

K.6.5 Security and Protection: Authentication

Author Keywords

Thermal Imaging; Mobile Authentication; TouchScreens.

CHI 2017, May 6-11, 2017, Denver, CO, USA.

ACM ISBN 978-1-4503-4655-9/17/05 ...$15.00.
http://dx.doi.org/10.1145/3025453.3025461

INTRODUCTION

The increasing amount of sensitive data available on personal mobile devices, such as personal photos, call logs, bank accounts, and emails underlines the need to secure them against various kinds of malicious attacks. For this reason, users protect access to their mobile devices using different authentication mechanisms, including patterns and PINs as well as biometric approaches, such as FaceUnlock or TouchID [23]. Privacy concerns are known to influence users' technology use decisions [39], suggesting that many users could abandon biometric schemes due to the associated privacy implications such as consequences of biometric information being disclosed. However, PINs and patterns remain among the most popular authentication mechanisms as of today [23, 49].

The usable security community lately focused on investigating different user-centered attacks, such as shoulder-surfing (e.g., [14, 18, 29]) and smudge attacks (e.g., [43, 51]). At the same time, a new threat emerged which received only little attention [4] so far from the research community, that is, thermal attacks on touch screens of mobile devices. The past years have witnessed portable thermal cameras becoming available on the mass market, in personal mobile devices such as CAT S60[1], or as attachable accessories for mobile devices (e.g., the FLIR One[2] or Seek thermal[3]). Falling hardware prices made these devices affordable. At the time of publishing

[1] `http://www.catphones.com/phones/s60-smartphone`
[2] `http://www.flir.de/flirone`
[3] `http://www.thermal.com`

this paper, a portable thermal camera with a temperature sensitivity of 0.05°C can be purchased for ∼400$. This creates an inherent need to understand the threat imposed by such devices that enable thermal attacks.

During a *thermal attack*, a thermal camera operating in the far infrared spectrum, captures the heat traces left on the surface of a mobile device after authentication. These traces are recovered and used to reconstruct the password. Unlike smudge attacks, thermal attacks can leak information about the order of entry for PINs and patterns (see Figure 1). Moreover, they can be performed after the victim had authenticated, alleviating the need for in-situ observation attacks (e.g., shoulder surfing attacks) that can be affected by hand occlusions.

While previous work utilized thermal conductance to recognize touch points for interaction [21, 32, 42], in this work we investigate its reliability to infer passwords from heat traces left on touch screens after authentication. Abdelrahman et al. [1] depicted a material space for on-surface heat trace recognition. Yet, their work did not cover touch screen's material. We investigate thermal attacks on Gorilla glass[4], the standard cover glass used for most touchscreens.

In this work, we explore how current authentication mechanisms are vulnerable to thermal attacks. We introduce an automated computer vision-based approach that analyzes the heat traces after an authentication process and extracts the potential PIN or pattern. Our implementation is open source, hence allowing further experimentation with thermal attacks[5]. We report on our findings from a user study where we investigated how properties of PINs and patterns influence the success of thermal attacks. In particular, we focus on the type of authentication scheme, the properties of the password and the time at which the attack was performed after authentication. We found that, although thermal images of PINs that contain duplicate digits do not leak the PIN to the naked eye (Figure 1), thermal attacks can yield 72% to 100% success rate when performed within the first 30 seconds after authentication. At the same time, thermal attack success decreases significantly in case of patterns (from 100% to 17% in the first 30 seconds) if the pattern includes one or more overlaps.

CONTRIBUTION STATEMENT

The contributions of this paper are as follows:

1. Assessment of the thermal contact conductance of state-of-the-art smartphone touch screens and how commercial thermal cameras can exploit them for thermal attacks.
2. An automated computer vision approach to analyze the thermal attack resistance of PINs and patterns by extracting them from heat traces.
3. An in-depth study of how properties of commonly used authentication schemes affect the success rate of thermal attacks.
4. A set of recommendations that help users and authentication scheme designers to overcome thermal attacks.

[4] `http://www.corning.com/gorillaglass/worldwide/en/products-with-gorilla.html`

[5] `https://github.com/Yomna-Abdelrahman/ThermalAttack.git`

RELATED WORK

Our work builds on two strands of prior work: (1) thermal imaging and (2) the different types of threats to user authentication on mobile devices.

Thermal Imaging

Thermal cameras capture the thermal map of a scene. They operate in the far-infrared spectrum with wavelengths between 7.5 and 13 μm. There are multiple differences between properties of thermal imaging and those of visible light.

The first thermal property is heat radiation. Compared to visible light, heat radiation has different reflection properties that depend on the surface [1]. Thermal reflections were exploited in previous work to enable body-worn and hand-held devices to detect mid-air gestures [42].

The second unique property is that thermal imaging is independent of light and coloring conditions, which allows thermal cameras to be used for face and expression recognition [30, 31]. Thermal cameras can provide information about the sensed body's temperature, which can be used to infer the physiological and cognitive state of users in a contact free manner [41] by, for example, evaluating their stress levels [28].

A third unique property is that thermal imaging is capable of detecting input that has been performed in the past. When a user touches a point on a surface, heat is transferred from the user to the surface, generating heat traces that slowly fade away. These traces can be detected using thermal imaging. Heat traces have been utilized for input [21, 32, 42] and to authenticate users based on their thermal hand print [11].

In this paper, we investigate the use of thermal imaging to infer passwords entered on mobile devices, exploiting the fact that heat traces take time to fade away. We investigate the thermal properties of state-of-the-art touch screens and study the impact of password properties on the heat trace and, thus, the successful retrieval of passwords via thermal imaging.

Threats to Authentication on Mobile Devices

Mobile devices, such as tablets and smart phones, store and allow access to a plethora of private content. Prior work investigated a number of threat models that put the user's private data at risk.

Shoulder Surfing Attacks

One of the most discussed threats are shoulder surfing attacks, in which an observer attempts to eavesdrop a user to uncover private information, among which are login credentials [18]. Different approaches have been introduced to mitigate shoulder surfing attacks, ranging from adding random cues [6, 7, 8, 46], splitting the attackers' attention by requiring them to observe multiple cues [14, 29], and disguising the user input [15, 22]. Despite focusing on login credentials, research also investigated methods to protect users from shoulder surfing text messages [19] and pictures [50]. Most of the schemes that counter shoulder surfing address a threat model where the attacker can clearly observe the password entry once. Other threat models cover multiple observation attacks [24, 36, 29, 52] or video attacks [14, 46].

Smudge Attacks

Another type of attack that has been addressed by previous work is smudge attack, in which an attacker exploits the oily residues left on the touch screen after interaction to uncover the password [5]. Smudge attacks perform particularly well against patterns, as smudges give hints on where the pattern started. However they can hardly provide any useful information about the order of PIN entries. Approaches to mitigate smudge attacks include graphically transforming the visual cue on which the password is entered [43, 51], introducing a random element that leads to different smudges at every authentication attempt [51], or using multiple fingers to increase the complexity of the pattern [35]. Threat models that consider smudge attacks assume that the attacker has access to the mobile device, in addition to clearly visible smudge traces and optimal lighting conditions to see the smudges clearly.

Thermal Image Attacks

Thermal image attacks exploit properties of thermal imaging. Namely, heat traces are transferred from the user's hands to the touch screen during authentication. These traces fade away slowly [32], allowing thermal cameras to perceive which parts of the display have been touched even after the user had already entered the password. Similar to shoulder surfing, thermal attacks leak information about the order of entered PINs and patterns [5]. In contrast to shoulder surfing, however, thermal attacks can be performed after the user had left the device. This gives attackers an advantage as they no longer need to observe the user while authenticating, which makes the attack more subtle and eliminates hand occlusions. Although thermal images can be distorted by interaction, a user who performs limited interactions or leaves the device after authentication is still vulnerable to thermal attacks.

Mowery et al. investigated the effectiveness of thermal attacks on ATMs with plastic keypads [34]. They found that thermal attacks are feasible even after the user authenticated. While Mowery et al. investigated thermal attacks on plastic keypads of ATMs, little work was done regarding thermal attacks on mobile devices and other touch screens devices. In a preliminary study, Andriotis et al. [4] were able to observe heat traces resulting from entering a pattern for 3 seconds after authentication. This allowed them to retrieve parts of the pattern.

In our work, we perform an in-depth analysis of how well thermal attacks perform on PINs and patterns on mobile device touch screens with respect to different password properties. We also consider duplicate digits in PINs, and overlaps in patterns. To do this, we implemented ThermalAnalyzer, which automatically retrieves passwords from heat traces. ThermalAnalyzer shows that thermal attacks can be successful even if they take place 30 seconds after authentication (i.e. 10 times longer compared to previous work [4]).

UNDERSTANDING THERMAL ATTACKS

Our work relies on the phenomenon of heat transfer from one object to another. Heat transfers from users' hands to surfaces they interact with, leaving traces behind that can be analyzed. This relies on the surface's material property know as thermal contact conductance [12], which refers to the conductivity of heat between two objects (surfaces) that are in contact.

According to the blackbody model [27], any object above absolute zero (e.g., surrounding objects in our environment) emits thermal radiation. This radiation is absorbed, reflected, and transmitted. However, for fully opaque surfaces the transmitted portion is discarded [20]. This limits the portions of interest to the reflected and absorbed radiation. Hence, thermal radiation could be presented as in $Thermal\ reflectivity + Thermal\ absorptivity = 1$.

As soon as an object contacts a surface, thermal radiation is transmitted and absorbed by the surface, causing a temperature change. This leads to heat traces accumulating on the surface. To compute the transferred heat and identify whether or not it is detectable by commercial thermal cameras, we measured the temperature at the contact point ($T_{contact}$). We used a well-established model by Ray [40] to compute the temperature at the contact point of the two bodies. In our scenario, the two bodies are: the human skin (i.e. the user's finger), and the mobile device's touchscreen (i.e. a plate of Gorilla glass).

$$T_{contact} = \frac{b_{skin}T_{skin} + b_{gorilla\ glass}T_{gorilla\ glass}}{b_{skin} + b_{gorilla\ glass}} \quad (1)$$

$$b = \sqrt{K.P.C} \quad (2)$$

$T_{contact}$ depends on the temperature of the contact points (T_{skin} and $T_{gorillaglass}$) as well as their *thermal penetration coefficient (b)*. It is the amount of thermal energy penetrated and absorbed by the surface. The b is defined in Equation 2. It is composed of the product of thermal conductivity (K), thermal density (P), and specific heat capacity (C) [38]. The b of human skin and the gorilla glass for short contact are 1000 $JS^{-1/2}m^{-2}K^{-1}$ [38] and 1385 $JS^{-1/2}m^{-2}K^{-1}$ [44] [6] respectively.

Additionally, the detection of temperature changes at the contact point depends on the camera's sensitivity. The change in temperature must be higher than the camera's temperature sensitivity to be distinguishable by the camera. For example, if the touch screen's glass has a temperature $T_{gorillaglass}$ of 23°C and the user's hand temperature T_{user} is 30°C, then $T_{contact}$ would be 25.9°C according to Equation 1. This results in a temperature difference of 2.9°C ($T_{contact} - T_{gorillaglass}$). Hence, a thermal camera with thermal sensitivity $\leq$ 2.9°C would be able to recover the order in which a PIN/pattern entry was performed by utilizing the heat trace decays. In our work, the thermal camera has a thermal sensitivity of 0.04°C, allowing different hand temperatures to be sensed.

THREAT MODEL

In our threat model, the attacker (i.e., a person who wants to access a device without permission) waits for the victim to complete the authentication process and to leave the mobile device. This could be the case when the user quickly checks his latest messages before getting something to drink from the coffee machine, while leaving the device on his/her desk. To ensure optimal conditions for the attacker in our threat model, the user does not interact with the device but merely

[6] This value was confirmed by lab measurements by the Institute of Applied Optics in our university

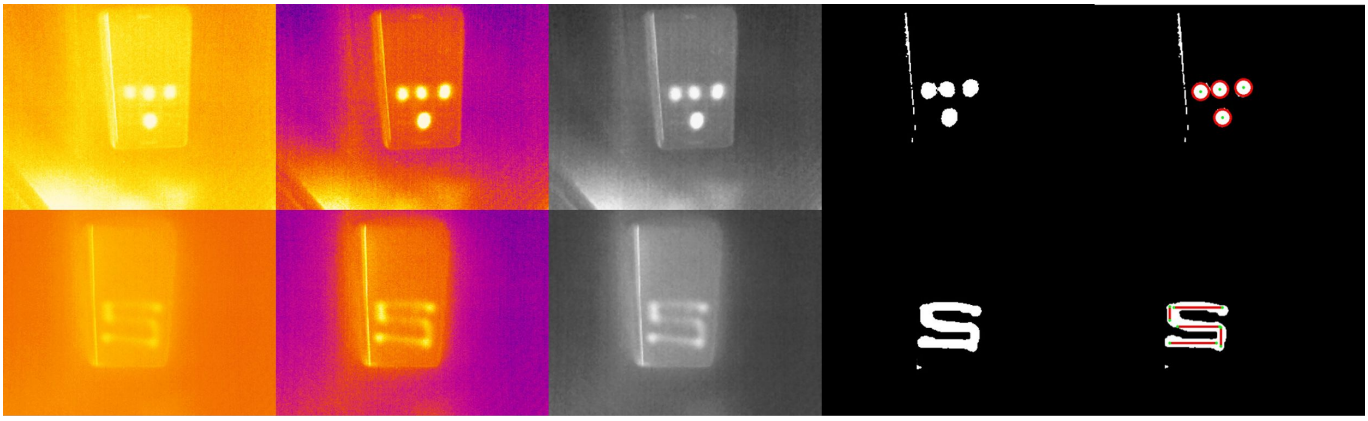

Figure 2: The figure illustrates the recognition pipeline of PINs (top) and patterns (bottom).

authenticates (e.g., to check an update from a notification or a widget) then leaves the device idle. The attacker then uses a thermal camera (e.g., integrated into a smart phone) to take a thermal image of the device's touchscreen. The attacker then analyzes the thermal image in a manner similar to our analysis presented in the following section to identify the PIN/pattern. Similar to previously discussed threat models [24, 36, 29, 52], the attacker exploits a chance where the device is unattended to login and access the user's private information.

THERMALANALYZER

In the following, we describe the design and implementation of the ThermalAnalyzer. The ThermalAnalyzer consists of a thermal camera capturing an image and a recognition pipeline used to extract the PINs and patterns.

Recognition Pipeline

The recognition pipeline consists of six steps, performed to extract a PIN or patterns from an image (Figure 2). The steps are performed using OpenCV[7] and include frame extraction, pre-processing, noise and background removal, and thresholding. The final step is feature extraction to deduce the touch points' location and temperature information.

Frame Extraction and Camera Configuration

We captured the thermal image through the Optris thermal camera API[8]. Using the interprocessor communication, we capture the frames in a 16-bit color format along with the encoding of the temperature information. We configured the camera using its API to capture temperature values between 19°C to 32°C. This was done to achieve higher contrast of colors that represent different temperature values as depicted in Figure 2. For each captured frame a pre-processing procedure is performed. This included noise filtering, background subtraction, and thresholding.

Noise Filtering

We adapted the noise filtering process used by [1, 32, 42] by applying a 5×5 px median filter, converting the image to grey scale and reapplying the filter for enhanced noise reduction.

[7] `http://opencv.org/`
[8] `http://www.optris.com/software`

Background Subtraction

We built a semi-static background model for background subtraction. A static model is preferred in our case as we want the heat trace detected to last over the frames and not to be adsorbed by a dynamic background model. Yet, on the other hand a dynamic model is required to tolerate slight temperature difference of the device along the operation. Hence, we built a semi-static background model, where the update is controlled by the learning rate (α) parameter, which is a value that controls the rate of background model updates. An α value of 0.001 showed the best result in preliminary tests. As a result, the latest heat trace stemming from password entry stays in the foreground, whereas heat traces from slight changes from the environment temperature are merged with the background.

Thresholding

To segment the regions that are relevant to identifying heat traces (Figure 2), we used Otsu's thresholding method [37]. The frame is classified into two set of pixels with minimum overlap between them based on a dynamically computed threshold by Otsu's algorithm. Then, we applied an additional morphological closing operation to highlight the boundaries of the thresholded foreground and reduce the background.

Feature Extraction

Our features are classified into (1) circular fitted traces for PIN detection and (2) line fitted traces for patterns detection.

The heat trace is detectable via extracting the contours from the binary images, where the image is scanned to detect arrays of contours. Similar to the work of Sahami et al. [42], we used a *circular fitted* contour detection to identify the PIN entries. The contour center is computed as the spatial moment of the extracted contour. Using the same approach for detecting the circular fitted heat traces, we used the Hough Transform [17] for extracting *line fitted* contour detection to identify the entered patterns, as depicted in Figure 2.

PIN and Pattern Sequence Detection

At this point in the processing pipeline, the PIN or pattern entry has been extracted from the captured frames but with no information about the sequence of entry. To infer the sequence for the PINs, we utilized a pre-set frame with the keypad to

identify the PIN location using squares. The squares represent the entire set of regions of interest (ROI). Mowery et al. [34] reported that representing the ROI with the mean temperature yields best performance for recovering the order of the entry sequence. Hence, we compute the mean temperature for each ROI, and sort them based on their weights.

To identify duplicate entries, we compute the overall average temperature of each digit. Thereby, the background temperature is subtracted. Hence, a digit that was never pressed would have a temperature value of almost zero. Consequently, duplicate entries (i.e., the digit that was touched multiple times) have a value that exceeds the overall average. The number of duplicates can be inferred from the relative temperature values of the overall number of detected presses. In summary, given a four-digits PIN, there would be four detection scenarios:

- Four different heat traces: This means there are no duplicates. Hence, ordering the traces based on their temperature in a descending manner would infer the sequence.
- Three different heat traces: The heat trace that has a temperature of $T_{contact}$ is the last entry in the PIN, as it will maintain the $T_{contact}$ value. This leaves only 3 possibilities for the remaining sequence, which are sufficient for the attacker to try without being locked out. This approach, however, will work with recently captured frames as the heat trace, i.e. $T_{contact}$, decays over time.
- Two different heat traces: According to the relative ratios of the weights, the number of repetitions of each digit is identified. Normalizing the weights would then show the last touched digit. Once the last digit is identified, the attacker can tell whether it is the duplicated one (i.e. the other duplicate is either in position 1, 2 or 3, while the remaining digits are ordered according to their heat traces), or the last digit is a non-duplicated digit, hence the attacker has only 3 possibilities to try without being locked out.
- One heat trace: This means that the PIN consists of the same number repeated 4 times.

One of the former three conditions could be experienced due to heat trace decay. In that case, we identified the missing digit to be unidentified and set it to be the beginning of the PIN (e.g., if 3 traces where detected with no evidence of duplicates, the first digit is labeled unknown and the remaining three are sorted by their temperature weights).

The same approach was followed for the patterns, where the extracted lines are analyzed and ordered by their mean temperature. Additionally, the temperature of the tips of the extracted lines were compared to identify the direction. Our algorithm does not account for a specific pattern length, hence we present the available heat trace to be the regenerated pattern.

For a more conservative analysis, ThermalAnalyzer is not optimized for detecting patterns of specific length (max of 9). This is because in our threat model, and most likely in a real scenario, the attacker does not know the pattern length. This means that in cases where the ThermalAnalyzer generated a guess that is of length n instead of 9, the heat traces of the remaining 9–n had already decayed by the time of the attack.

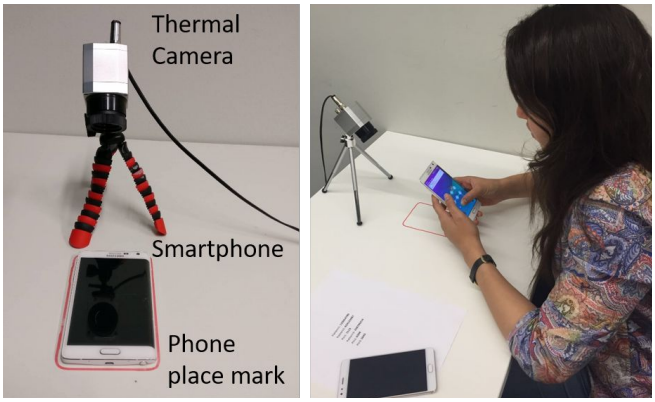

Figure 3: Setup with the thermal camera capturing the phone's screen.

COLLECTING THERMAL IMAGES

Despite the variety of authentication schemes that were introduced in the past years, personal identification numbers (PIN) is one of the most commonly used schemes [49]. Moreover, as Android devices dominate the market, there is an increasing adoption of pattern, which is an Android graphical password scheme where users draw a line pattern that connects dots displayed in a 3×3 grid [45].

In this study we analyze thermal images of a smart phone screen after a user entered a password. These images are evaluated using ThermalAnalyzer. We particularly focus on understanding how (1) different authentication schemes, (2) the time between password entry and attack, and (3) password properties influence the feasibility of thermal attacks.

Design

The study uses a repeated measures design, where all participants were exposed to all conditions. We studied the effect of three independent variables on the success of thermal attacks: (1) the password type: whether the used scheme is a PIN or a pattern, (2) the age of the heat trace: we analyzed heat traces 0, 15, 30, 45 and 60 seconds after authentication to investigate for how long they remain exploitable by an attacker, and (3) the properties of the PINs and patterns.

In case of PINs, the property we studied is the number of duplicates in the PIN. On one hand duplicates distort the heat traces, making the entry order less distinguishable. On the other hand the presence of duplicates reduces the password space, which means that less information from the thermal attack would be sufficient to uncover the password.

We studied the influence of having No-duplicates, 1-duplicate, and 2-duplicates (e.g., 1236, 1223, and 3222 respectively). Examples are shown in Figures 1a, 1b, and 1c respectively. In case of patterns, we investigated the effect of the number of overlaps [48] in the pattern. An overlap occurs when the user's finger passes through a node that is already selected. We expect that overlaps can distort the heat traces enough to make it infeasible to reconstruct the entered pattern. We studied the influence of having one, two, or no overlaps in the patterns (cf. Figures 1e, 1f, and 1d respectively).

Apparatus

Our setup included two Samsung Galaxy Note Edge smart phones, a thermal camera (Optris PI450[9]), and a GoPro Hero3 RGB camera, both mounted on a tripod. One smart phone was used for practicing the passwords and the other one for the actual input. The thermal camera has an optical resolution of 382×288 pixels and a frame rate of 80 Hz. It is able to measure temperatures between -20°C and 900°C, and operates with a thermal sensitivity of 0.04°C represented by the noise equivalent temperature difference (NETD)[10]. The wavelengths captured by the camera are in the spectral range between 7.5μm and 13μm. The lens provides a 80°×58° field of view. The thermal camera uses USB as power source as well as to transfer data. It provides temperature information in the form of 16-bit color values encoding the temperature information.

To ensure that heat traces are recorded at the intended times, we used a place mark at a distance of 80 cm in front of the camera (cf. Figure 3) to indicate the optimal position of the smart phone to record the heat trails with the thermal camera while minimizing thermal reflection. Additionally, we recorded the whole study using an RGB video camera. The RGB video feed was later used to determine the time at which the users fingers were no longer in contact with the screen.

Participant and Procedure

We recruited 18 participants (10 female and 8 male) with an average age of 28.3 years ($SD = 4.7$) using university mailing lists. All participants were students in different majors. Two participants were left handed. None of the participants had any previous experience with thermal cameras.

After participants arrived in the lab, we first asked them to sign a consent form and explained the purpose of the study. Next, we handed a set of PINs and patterns printed on cards as well as the two smart phones to the participants. To avoid errors and pauses during entry, we asked the participants to familiarize themselves beforehand with the passwords by entering them multiple times on the practice smart phone first. We instructed the participants to enter the password, then immediately place the study smart phone on the place mark on the table in front of them (cf. Figure 3). We waited for three minutes between each entry, to ensure full heat trace decay of the previous entry. Each participant entered three passwords of each type (i.e. 18 passwords). The order was counter balanced using a Latin-square.

The study took approximately 40 minutes. We video recorded the study for post-hoc analysis of the input times. Throughout the experiment we recorded the temperature of the participant's dominant hand (i.e., the hand used to enter the password), in addition to the phone's temperature. The experiment was conducted in a maintained room temperature of 24°C.

To analyze the thermal attacks, we considered two approaches: (1) visually inspecting heat traces and (2) using our computer vision approach ThermalAnalyzer. The analysis was done by one of the authors who was not aware and was never exposed to the list of entered passwords. Additionally, the feed was analyzed from the thermal cameras. Using the ThermalAnalyzer earlier, the author reported the regenerated PINs and patterns in a csv file defining all possible combinations.

[9] `http://www.optris.com/thermal-imager-pi400`

[10] NETD refers to the electronic noise that is interpreted as a temperature difference of an object.

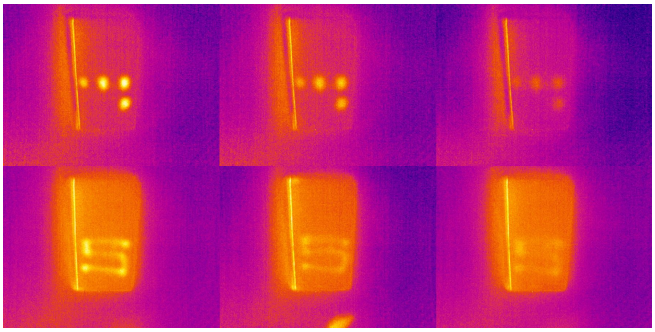

Figure 4: The figure shows the heat traces resulting from entering PINs (top) and Patterns (bottom) as they decay from 0 to 30 to 60 seconds.

RESULTS

To evaluate the success of thermal attacks against PINs and patterns, we measured

1. The success rate: the percentage of cases in which the thermal attack successfully revealed the entire password correctly.

2. The Levenshtein distance: the distance between the generated guesses and the correct password.

The success rate and Levenshtein distance were used in previous work to reflect how successful attacks are (success rate) and how close the guess is to the original password (Levenshtein distance) [16, 29, 47].

We visually inspected a sample of thermal images from the data (3 participants). However we could not visually recover the entire order of the PIN, nor the direction of the patterns. This is apparent in Figure 1a, where identifying the order of 3 and 6 is challenging to the naked eye. Additionally, the starting points of the pattern is not visually deducible (see Figure 1e). Hence, we only considered the PINs and patterns from the ThermalAnalyzer. We investigated the effect of three independent variables: (1) authentication scheme, (2) age of the heat trace and (3) password properties. The tasks performed during the study typically require between 26% to 44% CPU usage.

Statistical Analysis

As we have three independent variables, we analyzed the data using a three-factor repeated measures ANOVA (with Greenhouse-Geisser correction if sphericity was violated). This was followed by post-hoc pairwise comparisons using Bonferroni-corrected t-tests.

Figures 5 and 7 show the success rate per age of the heat traces and password property. Additionally, Figures 6 and 8 show the Levenshtein distances per age of the heat traces and password property. The results show that thermal attacks are more successful against PINs than against patterns.

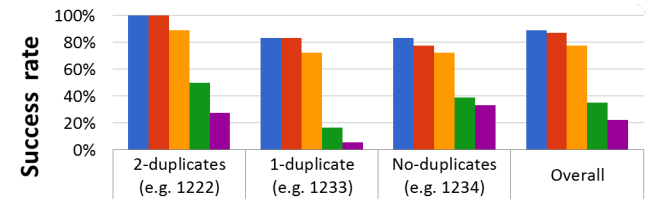

	2-duplicates (e.g. 1222)	1-duplicate (e.g. 1233)	No-duplicates (e.g. 1234)	Overall
Immediate	100%	83%	83%	89%
15 seconds	100%	83%	78%	87%
30 seconds	89%	72%	72%	78%
45 seconds	50%	17%	39%	35%
60 seconds	28%	6%	33%	22%

Figure 5: The successful thermal attack rate against PINs is significantly higher when the thermal image is taken in the first 30 seconds. Thermal attacks perform well against PINs with duplicate digits despite the noise introduced by touching the same digit multiple times.

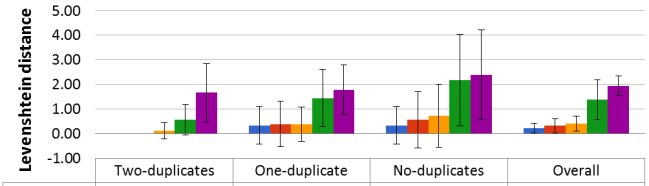

	Two-duplicates	One-duplicate	No-duplicates	Overall
immediate	0.00	0.33	0.33	0.22
15 sec	0.00	0.39	0.56	0.31
30 sec	0.11	0.39	0.72	0.41
45 sec	0.56	1.44	2.17	1.39
60 sec	1.67	1.78	2.39	1.94

Figure 6: The mean Levenshtein distances and standard deviation between the guessed PINs and the correct PIN.

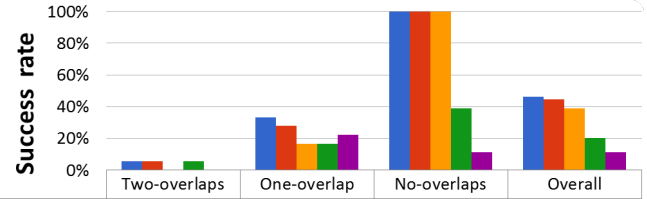

	Two-overlaps	One-overlap	No-overlaps	Overall
immediate	5.56%	33.33%	100.00%	46.30%
15 sec	5.56%	27.78%	100.00%	44.44%
30 sec	0.00%	16.67%	100.00%	38.89%
45 sec	5.56%	16.67%	38.89%	20.37%
60 sec	0.00%	22.22%	11.11%	11.11%

Figure 7: The successful thermal attack rate against patterns is significantly higher when the analyzed thermal image is taken in the first 30 seconds after authentication. Furthermore, thermal attacks are significantly less successful against patterns with overlaps.

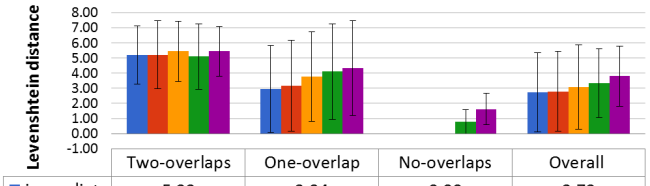

	Two-overlaps	One-overlap	No-overlaps	Overall
immediate	5.22	2.94	0.00	2.72
15 sec	5.22	3.17	0.00	2.80
30 sec	5.44	3.78	0.00	3.07
45 sec	5.11	4.11	0.78	3.33
60 sec	5.44	4.33	1.61	3.80

Figure 8: The mean Levenshtein distances and standard deviation for the guessed patterns and the correct pattern.

Authentication Scheme: PINs vs Patterns

Overall, thermal attacks were more successful for PINs (M = 0.62, SD = 0.31) than for patterns (M = 0.32, SD = 0.16). Similarly, the Levenshtein is shorter for PINs (M = 0.856, SD = 0.127) than for patterns (M = 3.14, SD = 0.28). We found a significant main effect of password type on the Levenshtein distance between the guess and the entered password $F_{1,17} = 91.923$, $p < 0.001$. Post-hoc analysis showed significant differences ($p < 0.001$) between passwords of type PIN (M = 0.856, SD = 0.127) compared to those of type pattern (M = 3.14, SD = 0.28). This means guesses against PINs are generally closer to the original password compared to those against patterns.

Age of Heat Traces

PINs

Looking at the age of the heat trace, the results show that the earlier the heat attack is performed, the higher the success rate and the lower the Levenshtein distance are (cf., Table 1). The results of the ANOVA revealed a significant main effect of the heat trace's age on the Levenshtein distance between the correct password and the guess $F_{1.79,30.45} = 41.7$, $p < 0.001$. Post-hoc analysis using Bonferroni corrected t-tests showed statistically significant differences between 60 seconds and all other durations ($p < 0.001$) as well as between 45 seconds and all other durations ($p < 0.001$). This shows that thermal attacks against PINs that take place within the first 30 seconds after authentication result in guesses that are significantly closer to the correct password compared to those done after 30 seconds. This is also reflected in the success rate as shown in Figure 5. Overall, this suggests that thermal attacks are very effective against PINs when performed within 30 seconds after authentication.

Patterns

Similar to the results of the PINs, the results for the patterns show that the older the traces are, the less likely a thermal attack is successful and the higher the Levenshtein distances are (cf., Table 1). We found a significant main effect of the heat trace's age on the Levenshtein distance between the correct pattern and the guess $F_{2.228,38.876} = 13.295$, $p < 0.001$. Post-hoc analysis using Bonferroni corrected t-tests showed significant differences ($p < 0.05$) between 60 seconds and all other durations.

This shows that thermal attacks against patterns that take place 60 seconds after authentication result in guesses that are significantly farther away from the correct password, compared to those done within the first 45 seconds. This is also reflected in the success rate shown in Figure 7. Overall, this suggests that thermal attacks are very effective against patterns when performed within 45 seconds after authentication.

Age	PIN		Pattern	
	Levenshtein	Success Rate	Levenshtein	Success Rate
immediate	M=0.222, SD=0.76	M=0.89, SD=0.08	M=0.222, SD=0.76	M=0.46, SD=0.40
15 seconds trace	M=0.222, SD=0.76	M=0.87, SD=0.09	M=0.315, SD=0.139	M=0.44, SD=0.40
30 seconds trace	M=0.407, SD=0.134	M=0.78, SD=0.08	M=0.407, SD=0.134	(M = 1, SD = 0.39)0.44
45 seconds trace	M=1.39, SD=0.2	M=0.35, SD=0.14	M=1.39, SD=0.2	M=0.20, SD=0.14
60 seconds trace	M=1.94, SD=0.23	M=0.22, SD=0.12	M=3.8, SD=0.32	M=0.11, SD=0.09

Table 1: The success rate and Levenshtein distances for different ages of the heat trace.

Hand and Screen Temperature

We found that the difference in temperature (D_t) between the hand and screen influences the success of a thermal attack. The higher D_t, the more successful is a thermal attack, since more thermal energy is transferred to the screen (cf. Equation 1). Using Pearson's product-moment correlation, we found that the correlation between D_t and the successful thermal attack rates increases from 0.55 (at 0 seconds) to 0.85 (at 60 seconds). This means that there is a strong correlation between D_t and the success of an attack and that D_t is particularly important for attacks happening some time after authentication.

Password Properties

PINs Duplicates

We found a significant main effect of number of duplicate digits on resistance to thermal attacks $F_{2,34} = 13.23$, $p < 0.01$. Post-hoc analysis revealed statistically significant differences ($p < 0.05$) between No-duplicates (M = 1.23, SD = 0.25) and 2-duplicates (M = 0.47, SD = 0.08) and between 1-duplicate (M = 0.87, SD = 0.15) and 2-duplicates (M = 0.47, SD = 0.08). This means that the more duplicates a PIN has, the closer the guess is to the correct PIN.

This shows that although the presence of duplicate digits makes it harder to determine the order of the detected touches, the approach is able to determine if a digit is repeated two or three times. As a result, the security added by overwritten heat traces in case of duplicate PINs is outweighed by the significantly reduced password space.

Patterns Overlaps

We found a significant main effect of the number of overlaps on the distance between the correct pattern and the guess $F_{1.441,24.503} = 28.563$, $p < 0.001$. Post-hoc analysis showed significant differences between two pairs ($p < 0.001$): patterns with no overlaps (M = 0.48, SD = 0.08) compared to those with one overlap (M = 3.67, SD = 0.68), and between patterns with no overlaps (M = 0.478, SD = 0.08) compared to those with two-overlaps (M = 5.29, SD = 0.43). No significant differences were found between the third pair ($p > 0.05$).

This shows that although patterns can be successfully uncovered using thermal attacks up to 30 seconds after authentication (100% success rate), the presence of overlaps significantly increases its resistance against thermal attacks.

DISCUSSION

The results of our study and a review of prior work reveal that surfaces with specific properties can be used to detect on-surface interaction using thermal imaging. On this basis, in the previous sections, we presented the results from collecting and analyzing thermal traces of authentication processes, which we summarize and discuss grouped by the most important observations in the following.

We particularly focused on PINs and patterns since they are currently the most common knowledge-based authentication schemes [23, 49]. However, other authentication schemes may be vulnerable to thermal attacks as well. We expect attacks on graphical passwords that rely on cued-recall [2, 3, 43] to be similarly effective compared to the patterns we investigated.

PINs are typically easy to uncover using observation attacks (De Luca et al. report 95% successful attack rate for PINs [15]). Our results indicate that PINs poorly resist thermal attacks as well, with overall success rates ranging from 78% to 100% when attacks are performed within the first 30 seconds after authentication (Figure 5). Although smudge attacks against PINs can uncover which digits were entered, hence significantly reducing the password space, thermal attacks can additionally uncover the order in which the digits were entered.

Without overlaps, patterns of maximum length are uncovered in 100% of the cases when thermal attacks are performed within 30 seconds after authentication (Figure 7). However, just adding one overlap significantly increases resistance to thermal attacks, as it influence the direction detection and the order of the performed patterns. Overlaping patterns did not have the same effect as duplicate PINs, as they also influence the detected direction and the order of the performed pattern. Hence we recommend including an overlap movement in patterns to increase resistance against thermal attacks.

In contrast to overlaps, knight moves do not distort the heat trace of pattern points but only the path at intersections. Hence, knight moves are similarly ineffective in making thermal attacks more difficult as they are against smudge attacks.

Moreover, unlike smudge attacks, *thermal attacks* do not require finding an optimal angle at which the traces are visible. Thermal attacks were shown to be tolerant to viewing angle/distance, as reported by Mowery et al. [34]. Mowery et al. evaluated different distances (30–70 cm) and did not observe changes in the detection. In our setup the camera was placed at 80 cm above the phone, hence we expect minimal to no influence of the distance on the results.

In contrast to observation attacks, thermal images are taken after authentication, hence the attack is less obvious to the victim and is not influenced by authentication speed. Additionally, the operation of thermal imaging allows seamless attacks, as it operates in a light invariant manner, where lighting conditions do not influence the capturing of thermal information [34].

Using a thermal camera with high temperature sensitivity and an automated computer vision approach to detect the traces, outperformed the results reported by prior related work [34]. Our approach unveils PINs/Patterns with high success after 30 seconds while previous work was successful up to only 3 seconds after authentication. While a higher sensitivity camera might have led to better results in manual analysis, we believe the main enhancement to come from the automated computer vision approach which allowed detection of heat traces despite being invisible to manual visual inspection.

RECOMMENDATIONS TO RESIST THERMAL ATTACKS

There are ways to resist thermal attacks. We present three categories: (1) based on the results of our study, we are able to guide users in selecting PINs/patterns that are resistant to thermal attacks, (2) based on a literature review, we recommend schemes that are theoretically unaffected by thermal attacks, and (3) we present novel approaches that distort the heat traces, reducing the chances for successful thermal attacks.

Selection of PINs and Patterns

Our results indicate that adding a single overlap in an authentication pattern significantly increases the resistance to thermal attacks. When it comes to PINs, although duplicates distort the heat traces thermal attacks rely on, also other factors contribute to the ease/difficulty of uncovering duplicate PINs.

We recommend to increase the resistance of PINs against thermal attacks by increasing the number of digits in the PIN. The longer the PIN the longer it takes the user to enter it, which would in turn decrease the intensity of heat traces of the first digits by the time the user authenticates.

Thermal Attack Resilient Schemes

Many authentication schemes have been proposed to resist different types of attacks. We are not aware of systems that were built with the main aim of resisting thermal attacks on touch screens. However, some existing knowledge-based schemes do resist them by design.

One group of authentication schemes resilient against thermal attacks rely on one or more modalities other than touch input. For example, biometrics schemes (for example, [10, 13, 25, 26]) rely on data collected by sensors, such as accelerometers, to identify the user. Since they do not use the touch screen for dedicated input, they are not vulnerable to thermal attacks.

Similarly, authentication schemes that combine touch input with another modalites increase the resilience towards thermal attacks. PhoneLock [6], SpinLock [7], TimeLock [8], and ColorLock [8] augment PIN entry by using auditory and haptic cues the user needs to respond to when authenticating. These cues are randomized to counter shoulder-surfing attacks. Other examples utilized eye movements. For example Liu et al. [33] and Bulling et al. [9] used gaze input to authenticate. Similarly, Khamis et al. [29] introduced GazeTouchPass which combines gaze gestures and touch-input. Depending on the authentication scheme, the use of thermal cameras can still help the attacker to reveal the part of the input made on the touch screen. Being untied to the touchscreen, thermal attacks against these schemes would fail to uncover the PIN.

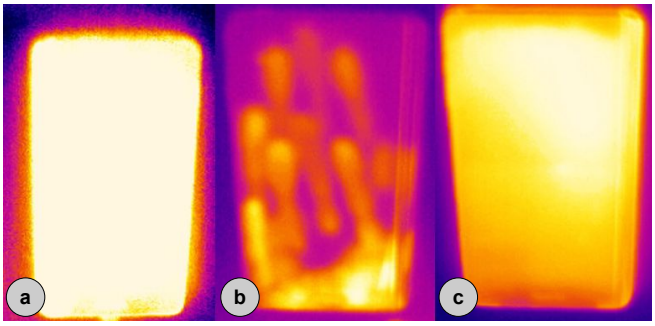

Figure 9: In addition to solutions based on our results and on prior work, we additionally propose the following approaches for resisting thermal attacks after entering the password: (a) using a white flash light for 3 second, (b) users swaps their hand randomly on the screen to distort the password's heat trace and (c) forcing maximum CPU usage for 3 seconds.

Moreover, novel authentication schemes designed to resist smudge attacks also increase the resilience towards thermal attacks since smudge attacks exploit a similar weakness in touch-based input. For example, SmudgeSafe [43] complicates smudge attacks against graphical passwords by randomly transforming the underlying image, causing the smudges to be different at every login attempt. Von Zezschwitz et al. [51] proposed three token-based graphical password schemes, two of which were significantly more secure against smudge attacks compared to patterns. The schemes rely on randomly positioned dragable objects. Hence thermal attacks are not expected to perform better than smudge attacks against these.

Physical Protection Measures

While novel authentication schemes increase the resistance towards thermal attacks, increasing the security of current PIN and pattern input against thermal attacks is still an important aspect. Placing the hand on the display might remove all thermal traces on the screen as shown in Figure 9b. However, there are different procedures that decrease the success rate of thermal attacks without involving the user. For example, increasing the brightness of the display to the maximum for a few seconds heats up the display temperature and, thus, reduces the time thermal traces are visible, as depicted in Figure 9a. Similarly, running computationally heavy processes on the phone quickly heats the phone up, resulting in a similar effect as shown in Figure 9c.

LIMITATIONS

In this work we explore and understand the effect of PIN/Pattern properties on its vulnerability to thermal attacks. Hence, our threat model assumes a best case scenario for the attacker. In this case, the user unlocks the phone (e.g., to check notifications or a calendar entry on the home screen) without further interaction. We acknowledge that in real-world situations, users are likely to interact after unlocking the phone, thus creating further traces. More sophisticated approaches (e.g., using deep learning) in the future could separate authentication and interaction patterns by, for example, utilizing the trace's age to only consider the oldest trace.

We did not measure CPU usage during the study, although it might influence the success of the thermal attacks. This could be considered in future work to investigate the influence of the CPU usage on the success rate of thermal attacks.

We had a stationary setup with the thermal camera at fixed distance. Initial studies have shown robustness of thermal attacks against viewing distance. However, exploring different viewing angles ranging from 0 to 180 degrees from the phone with combination of different distance, might enhance the understanding and practicality of thermal attacks.

FUTURE WORK

In future work, our results can be used to generalize thermal attacks on devices with touch screens. We could consider wider scenarios including tablets and shared public touch screen devices (e.g., IKEA self check-out, where users enter their PIN code on a touch screen without any further interaction on the screen, making their PIN vulnerable to thermal attacks). Additionally, it would be interesting to extend our ThermalAnalyzer to include neural networks and learning mechanisms for a better detection of PINs and patterns. We analyzed one image for each point in time. We expect that with more sophisticated thermal recording stream of images or video and the use of the trace history can further increase the success rate of attacks.

Our findings are based on the thermal camera sensitivity. The use of a thermal camera with higher thermal sensitivity would allow the heat traces to be detected even after 60 seconds. We considered PINs of length four. However the approach can be replicated as much as required to infer longer PINs.

We analyzed the thermal contact conductance to identify the applicability of our proposed attack. Additionally, computing and analyzing the heat transfer coefficient of the touch screen would provide more detailed information about the decay rate and the age of the traces. If users know these two thermal properties, they would be able to identify the possibility of a thermal attack on their devices.

Another direction could be the analysis of different screen protectors. We tested some materials, yet a visual analysis revealed that many other materials perform worse as heat traces are shown with more intensity compared to the gorilla glass. A detailed investigation of different screen protector could yield further insights.

CONCLUSION

We investigated the viability of thermal attacks on state-of-the-art touch screens and authentication schemes of mobile devices. To analyze the thermal images we implemented the ThermalAnalyzer, which was capable of uncovering 72%–100% of PINs in the first 30 seconds, and 100% of patterns that do not have overlaps. We additionally found that pattern overlaps significantly increase resistance to thermal attacks. Our work validates that thermal attacks are indeed a threat to mobile devices and should be considered by users and authentication scheme designers alike. We also furnish several solutions to protect from thermal attacks that are based on our results, previous work, and approaches to distort heat traces.

ACKNOWLEDGMENT

This work was partly conducted within the Amplify project which received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 683008) and received funding from the German Research Foundation within the SimTech Cluster of Excellence (EXC 310/2).

REFERENCES

1. Yomna Abdelrahman, Alireza Sahami Shirazi, Niels Henze, and Albrecht Schmidt. 2015. Investigation of Material Properties for Thermal Imaging-Based Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 15–18. DOI: http://dx.doi.org/10.1145/2702123.2702290
2. Florian Alt, Mateusz Mikusz, Stefan Schneegass, and Andreas Bulling. 2016. Long-term Memorability of Cued-Recall Graphical Passwords with Saliency Masks. In *Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia (MUM'16)*. ACM, New York, NY, USA. DOI: http://dx.doi.org/10.1145/3012709.3012727
3. Florian Alt, Stefan Schneegass, Alireza Sahami Shirazi, Mariam Hassib, and Andreas Bulling. 2015. Graphical Passwords in the Wild: Understanding How Users Choose Pictures and Passwords in Image-based Authentication Schemes. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '15)*. ACM, New York, NY, USA, 316–322. DOI:http://dx.doi.org/10.1145/2785830.2785882
4. Panagiotis Andriotis, Theo Tryfonas, George Oikonomou, and Can Yildiz. 2013. A Pilot Study on the Security of Pattern Screen-lock Methods and Soft Side Channel Attacks. In *Proceedings of the Sixth ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '13)*. ACM, New York, NY, USA, 1–6. DOI: http://dx.doi.org/10.1145/2462096.2462098
5. 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 (WOOT'10)*. USENIX, Berkeley, CA, USA, 1–7. http://dl.acm.org/citation.cfm?id=1925004.1925009
6. Andrea Bianchi, Ian Oakley, Vassilis Kostakos, and Dong Soo Kwon. 2011. The Phone Lock: Audio and Haptic Shoulder-surfing Resistant PIN Entry Methods for Mobile Devices. In *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '11)*. ACM, New York, NY, USA, 197–200. DOI: http://dx.doi.org/10.1145/1935701.1935740
7. Andrea Bianchi, Ian Oakley, and DongSoo Kwon. 2011. Spinlock: A Single-Cue Haptic and Audio PIN Input

Technique for Authentication. In *Haptic and Audio Interaction Design*, EricW. Cooper, VictorV. Kryssanov, Hitoshi Ogawa, and Stephen Brewster (Eds.). Lecture Notes in Computer Science, Vol. 6851. Springer Berlin Heidelberg, 81–90.

8. Andrea Bianchi, Ian Oakley, and Dong Soo Kwon. 2012. Counting clicks and beeps: Exploring numerosity based haptic and audio PIN entry. *Interacting with Computers* 24, 5 (2012), 409 – 422.

9. Andreas Bulling, Florian Alt, and Albrecht Schmidt. 2012. Increasing the Security of Gaze-based Cued-recall Graphical Passwords Using Saliency Masks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 3011–3020. DOI: **http://dx.doi.org/10.1145/2207676.2208712**

10. Daniel Buschek, Alexander De Luca, and Florian Alt. 2015. Improving Accuracy, Applicability and Usability of Keystroke Biometrics on Mobile Touchscreen Devices. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1393–1402. DOI: **http://dx.doi.org/10.1145/2702123.2702252**

11. Kun Woo Cho, Feng Lin, Chen Song, Xiaowei Xu, Fuxing Gu, and Wenyao Xu. 2016. Thermal handprint analysis for forensic identification using Heat-Earth Mover's Distance. In *2016 IEEE International Conference on Identity, Security and Behavior Analysis (ISBA)*. IEEE, 1–8.

12. MG Cooper, BB Mikic, and MM Yovanovich. 1969. Thermal contact conductance. *International Journal of heat and mass transfer* 12, 3 (1969).

13. Alexander De Luca, Alina Hang, Emanuel von Zezschwitz, and Heinrich Hussmann. 2015. I Feel Like I'M Taking Selfies All Day!: Towards Understanding Biometric Authentication on Smartphones. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1411–1414. DOI: **http://dx.doi.org/10.1145/2702123.2702141**

14. Alexander De Luca, Marian Harbach, Emanuel von Zezschwitz, Max-Emanuel Maurer, Bernhard Ewald Slawik, Heinrich Hussmann, and Matthew Smith. 2014. Now You See Me, Now You Don't: Protecting Smartphone Authentication from Shoulder Surfers. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2937–2946. DOI: **http://dx.doi.org/10.1145/2556288.2557097**

15. Alexander De Luca, Emanuel von Zezschwitz, Ngo Dieu Huong Nguyen, Max-Emanuel Maurer, Elisa Rubegni, Marcello Paolo Scipioni, and Marc Langheinrich. 2013a. Back-of-device Authentication on Smartphones. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2389–2398. DOI: **http://dx.doi.org/10.1145/2470654.2481330**

16. Alexander De Luca, Emanuel von Zezschwitz, Laurent Pichler, and Heinrich Hussmann. 2013b. Using Fake Cursors to Secure On-screen Password Entry. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2399–2402. DOI: **http://dx.doi.org/10.1145/2470654.2481331**

17. Richard O Duda and Peter E Hart. 1972. Use of the Hough transformation to detect lines and curves in pictures. *Commun. ACM* 15, 1 (1972), 11–15.

18. Malin Eiband, Mohamed Khamis, Emanuel von Zezschwitz, Heinrich Hussmann, and Florian Alt. 2017. Understanding Shoulder Surfing in the Wild: Stories from Users and Observers. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 11.

19. Malin Eiband, Emanuel von Zezschwitz, Daniel Buschek, and Heinrich Hußmann. 2016. My Scrawl Hides It All: Protecting Text Messages Against Shoulder Surfing With Handwritten Fonts. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. ACM, New York, NY, USA, 2041–2048. DOI: **http://dx.doi.org/10.1145/2851581.2892511**

20. Robert C Folweiler and William J Mallio. 1964. *Thermal Radiation Characteristics of Transparent Semi-Transparent and Translucent Materials under Non-isothermal Conditions*. Technical Report. DTIC Document.

21. 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 (PerDis '15)*. ACM, New York, NY, USA, 225–231. DOI: **http://dx.doi.org/10.1145/2757710.2757713**

22. Jan Gugenheimer, Alexander De Luca, Hayato Hess, Stefan Karg, Dennis Wolf, and Enrico Rukzio. 2015. ColorSnakes: Using Colored Decoys to Secure Authentication in Sensitive Contexts. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '15)*. ACM, New York, NY, USA, 274–283. DOI: **http://dx.doi.org/10.1145/2785830.2785834**

23. Marian Harbach, Alexander De Luca, and Serge Egelman. 2016. The Anatomy of Smartphone Unlocking: A Field Study of Android Lock Screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 4806–4817. DOI: **http://dx.doi.org/10.1145/2858036.2858267**

24. Eiji Hayashi, Sauvik Das, Shahriyar Amini, Jason Hong, and Ian Oakley. 2013. CASA: Context-aware Scalable Authentication. In *Proceedings of the Ninth Symposium on Usable Privacy and Security (SOUPS '13)*. ACM, New York, NY, USA, Article 3, 10 pages. DOI: **http://dx.doi.org/10.1145/2501604.2501607**

25. Christian Holz and Frank R. Bentley. 2016. On-Demand Biometrics: Fast Cross-Device Authentication. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3761–3766. DOI: **http://dx.doi.org/10.1145/2858036.2858139**

26. Christian Holz and Marius Knaust. 2015. Biometric Touch Sensing: Seamlessly Augmenting Each Touch with Continuous Authentication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15)*. ACM, New York, NY, USA, 303–312. DOI: **http://dx.doi.org/10.1145/2807442.2807458**

27. John R Howell, M Pinar Menguc, and Robert Siegel. 2010. *Thermal radiation heat transfer*. CRC press.

28. H. Kataoka, H. Kano, H. Yoshida, A. Saijo, M. Yasuda, and M. Osumi. 1998. Development of a skin temperature measuring system for non-contact stress evaluation. In *Proceedings of the Conference on Engineering in Medicine and Biology Society*. 940–943.

29. Mohamed Khamis, Florian Alt, Mariam Hassib, Emanuel von Zezschwitz, Regina Hasholzner, and Andreas Bulling. 2016. 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 (CHI EA '16)*. ACM, New York, NY, USA, 2156–2164. DOI: **http://dx.doi.org/10.1145/2851581.2892314**

30. Masood Mehmood Khan, Michael Ingleby, and Robert D. Ward. 2006. Automated Facial Expression Classification and affect interpretation using infrared measurement of facial skin temperature variations. *ACM Transactions on Autonomous Adaptive Systems* 1, 1 (2006), 91–113.

31. Seong G Kong, Jingu Heo, Faysal Boughorbel, Yue Zheng, Besma R Abidi, Andreas Koschan, Mingzhong Yi, and Mongi A Abidi. 2007. Multiscale fusion of visible and thermal IR images for illumination-invariant face recognition. *International Journal of Computer Vision* 71, 2 (2007), 215–233.

32. Eric Larson, Gabe Cohn, Sidhant Gupta, Xiaofeng Ren, Beverly Harrison, Dieter Fox, and Shwetak Patel. 2011. HeatWave: Thermal Imaging for Surface User Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 2565–2574. DOI: **http://dx.doi.org/10.1145/1978942.1979317**

33. Dachuan Liu, Bo Dong, Xing Gao, and Haining Wang. Exploiting Eye Tracking for Smartphone Authentication. In *Proc. of ACNS '15*. 20.

34. 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 (WOOT'11)*. USENIX Association, Berkeley, CA, USA.

35. Ian Oakley and Andrea Bianchi. 2012. Multi-touch Passwords for Mobile Device Access. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, 611–612. DOI:**http://dx.doi.org/10.1145/2370216.2370329**

36. Ian Oakley and Andrea Bianchi. 2014. Keeping Secrets from Friends. *Archives of Design Research* 27 (2014), 49–62. **http://www.dbpia.co.kr/Article/NODE02465396**

37. Nobuyuki Otsu. 1975. A threshold selection method from gray-level histograms. *Automatica* 11, 285-296 (1975), 23–27.

38. KC Parsons. 1992. Contact between human skin & hot surfaces equivalent contact temperature. In *Proc. ICEE*.

39. Alexander P. Pons and Peter Polak. 2008. Understanding User Perspectives on Biometric Technology. *Commun. ACM* 51, 9 (Sept. 2008), 115–118. DOI: **http://dx.doi.org/10.1145/1378727.1389971**

40. RD Ray. 1984. The theory and practice of safe handling temperatures. *Applied ergonomics* 15, 1 (1984).

41. E F J Ring and K Ammer. 2012. Infrared thermal imaging in medicine. *Physiological Measurement* 33, 3 (2012), R33. **http://stacks.iop.org/0967-3334/33/i=3/a=R33**

42. 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 (CHI '14)*. ACM, New York, NY, USA, 3483–3492. DOI: **http://dx.doi.org/10.1145/2556288.2557208**

43. Stefan Schneegass, Frank Steimle, 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 (UbiComp '14)*. ACM, New York, NY, USA, 775–786. DOI: **http://dx.doi.org/10.1145/2632048.2636090**

44. Corning Display Technologies. 2006. Glass Material Information. **http://www.sydor.com/wp-content/uploads/Corning-EAGLE-XG-Display-Glass.pdf**. (2006). Accessed September 19, 2016.

45. Sebastian Uellenbeck, Markus Dürmuth, Christopher Wolf, and Thorsten Holz. 2013. Quantifying the Security of Graphical Passwords: The Case of Android Unlock Patterns. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (CCS '13)*. ACM, New York, NY, USA, 161–172. DOI: **http://dx.doi.org/10.1145/2508859.2516700**

46. Emanuel von Zezschwitz, Alexander De Luca, Bruno Brunkow, and Heinrich Hussmann. 2015a. SwiPIN: Fast and Secure PIN-Entry on Smartphones. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1403–1406. DOI: **http://dx.doi.org/10.1145/2702123.2702212**

47. Emanuel von Zezschwitz, Alexander De Luca, and Heinrich Hussmann. 2013. *Survival of the Shortest: A Retrospective Analysis of Influencing Factors on Password Composition*. Springer Berlin Heidelberg, Berlin, Heidelberg, 460–467. DOI: **http://dx.doi.org/10.1007/978-3-642-40477-1_28**

48. Emanuel von Zezschwitz, Alexander De Luca, Philipp Janssen, and Heinrich Hussmann. 2015b. Easy to Draw, but Hard to Trace?: On the Observability of Grid-based (Un)Lock Patterns. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2339–2342. DOI: **http://dx.doi.org/10.1145/2702123.2702202**

49. Emanuel von Zezschwitz, Paul Dunphy, and Alexander De Luca. 2013. Patterns in the Wild: A Field Study of the Usability of Pattern and Pin-based Authentication on Mobile Devices. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '13)*. ACM, New York, NY, USA, 261–270. DOI: **http://dx.doi.org/10.1145/2493190.2493231**

50. Emanuel von Zezschwitz, Sigrid Ebbinghaus, Heinrich Hussmann, and Alexander De Luca. 2016. You Can'T Watch This!: Privacy-Respectful Photo Browsing on Smartphones. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 4320–4324. DOI: **http://dx.doi.org/10.1145/2858036.2858120**

51. 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 (IUI '13)*. ACM, New York, NY, USA, 277–286. DOI: **http://dx.doi.org/10.1145/2449396.2449432**

52. Oliver Wiese and Volker Roth. 2016. See You Next Time: A Model for Modern Shoulder Surfers. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '16)*. ACM, New York, NY, USA, 453–464. DOI:**http://dx.doi.org/10.1145/2935334.2935388**