“They see me scrollin”—Lessons Learned from Investigating Shoulder Surfing Behavior and Attack Mitigation Strategies

Alia Saad, Jonathan Liebers, Stefan Schneegass, and Uwe Gruenefeld

1 Introduction

People interact with an evergrowing number of mobile computing devices in everyday life. Nowadays, these devices have become ubiquitous and are commonly used in various places such as buses, trains, airports, coffee shops, and restaurants [3, 14]. As a result of the continuous growth, privacy and security challenges of these devices are becoming increasingly pressing. For example, smartphones hold sensitive information about users, including business records, financial interactions, personal details, and many more that should be kept hidden from others. Nevertheless, finding privacy-preserving solutions is not restricted to smartphones only. These solutions need to consider a variety of personal devices (e.g., smartwatches and tablets) as well as public or shared devices (e.g., ATMs and ticket machines).

All these devices are subject to various types of attacks. For instance, thermal attacks, where intruders use thermal cameras to analyze the heat traces of the entered authentication [1] or attacks that analyze the smudges on the screen for password reconstruction and gaining illegitimate access [49, 52]. However, smudge attacks are mainly focused on the authentication period, and thermal attacks require technical support and proper planning for a person to take a photo, feed it to a recognizer, and gain unauthorized access. On the other hand, observation attacks, commonly known as *shoulder surfing attacks*, are directly performed by humans and usually do not require additional hardware to be successfully completed. Despite a large body of work on these observation attacks, shoulder surfing remains a significant unresolved problem that requires more attention.

A. Saad (✉) · J. Liebers · S. Schneegass · U. Gruenefeld
University of Duisburg-Essen, Essen, Germany
e-mail: alia.saad@uni-due.de; jonathan.liebers@uni-due.de; stefan.schneegass@uni-due.de; uwe.gruenefeld@uni-due.de

N. Gerber et al. (eds.), *Human Factors in Privacy Research*,
https://doi.org/10.1007/978-3-031-28643-8_10

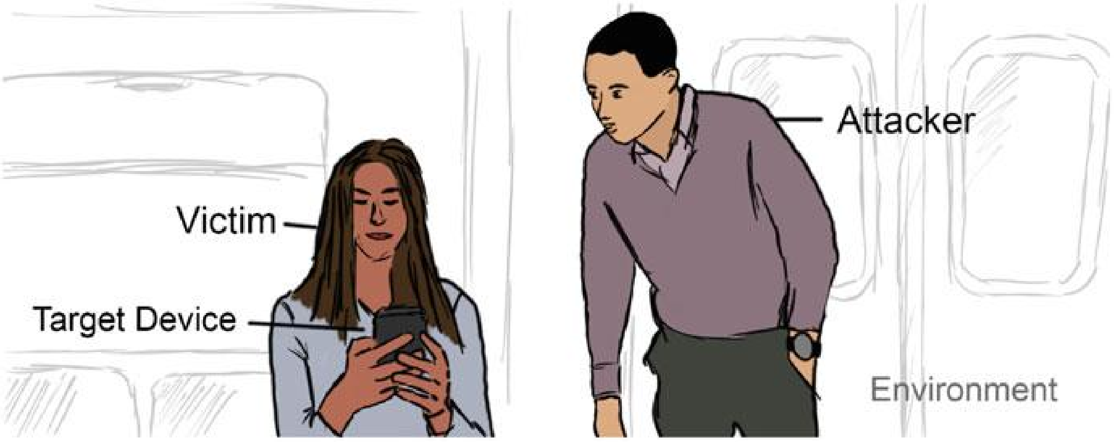

Fig. 1 Sketched example of a spontaneous shoulder surfing attack taking place during daily commute

Observation attacks are not limited to a specific device, location, or acquaintance level. Shoulder surfer can gaze at a person interacting with their personal phone or at someone's PIN, while they authenticate themselves after getting the phone out of the pocket. They do not need an extra device and can quickly memorize entered PINs or passwords. They could be standing in a train [46], or sitting next to the victim in an office [2] (see Fig. 1). The incident could occur between two closely tied people or with total strangers. Previous works confirm that observation attacks are widespread and highly likely to occur [14].

With this pervasiveness, nearly everyone is both *attacker* and *victim*. Albeit, recent studies showed that shoulder surfing incidents often take place opportunistically, and without malicious intent. To this end, we consider a person looking at the *user's* interaction as an *observer*, as we are not sure of their motives. Many researchers focused on understanding the occurrence of the observation attacks. However, regardless of the intentions of the observers, researchers also worked on various approaches to mitigate the risk of being observed, either by detection of the observer, or by providing novel solutions to prevent the looker from perceiving the content displayed.

Chapter Overview In the next section, we define the term shoulder surfing, describe different dimensions relevant for shoulder surfing attacks, and present key findings from previous research. Thereafter, we look at proposed strategies to mitigate shoulder surfing attacks. Here, we start by looking at threat models and algorithmic detection of shoulder surfers. Finally, we outline challenges and future research directions for shoulder surfing research.

2 Investigating the Phenomenon

In this section, we first define shoulder surfing to set the scope for this chapter. After that, we describe different methods with which researchers have investigated the phenomenon and discuss their advantages and disadvantages. Finally, we highlight the key findings from studies investigating shoulder surfing behavior.

2.1 *Defining Shoulder Surfing (Attacks)*

Observation attacks, commonly known as *shoulder surfing attacks*, are directly performed by humans and usually do not require additional technology to be successful. Farzand et al. [16] define shoulder surfing as observing someone's device screen without their consent. There are technology-based approaches to investigate observation attacks using machine vision, commonly referred to as recording attacks or video-based observation attacks (e.g., [30, 61]). Nonetheless, this chapter primarily focuses on shoulder surfing attacks performed by humans.

To be classified as shoulder surfing, it does not matter if the motivation to shoulder surf is simply curiosity or a deliberate attempt to steal information [9]. In fact, shoulder surfing mainly occurs in an opportunistic, non-malicious way [14]. Nonetheless, failing to prevent bystanders from observing sensitive information can lead to negative consequences such as financial loss, public exposure, and embarrassment [3]. An example of a shoulder surfing attack is shown in Fig. 1.

In the following, we provide an overview of different dimensions that help describe and classify shoulder surfing. The goal is not to present a complete overview of all dimensions relevant to shoulder surfing but rather to discuss different aspects that should be considered:

Motivation of Attack: Shoulder surfing attacks can be either *intentional* or *unintentional*, whereas unintentional means in an opportunistic, non-malicious way [9]. In most cases, shoulder surfing is unintentional and does not have serious consequences [14]. Nonetheless, it can evoke negative feelings for both parties and result in various coping strategies.

Attack Pattern: Shoulder surfing attacks can follow different attack patterns. Abdrabou et al. [2] found three different patterns: *continuous attacks*, *cautious attacks*, and *repeated attack*. While continuous attacks are characterized by bystanders looking at the target device for an extended period with few or no gaze shifts, cautious and repeated attacks alternate between observing the target device and looking away. For the latter two, the difference is the victim's behavior, who either looks up from the target device (from time to time) or shows high engagement. Friends, family, or colleagues at work may repeatedly observe their peers and thereby combine multiple partial observations to form a hypothesis of a target device's secret [37, 57].

Number of Attackers: In theory, a shoulder surfing attack can be performed by *multiple attackers*. While some research considers threat models with more than *one attacker* [24], many studies simplify this aspect and study 1:1 relationships between victim and attacker.

Relationship Between Victim and Attacker: Besides the number of attackers, the *type of relationship* (family, friend, colleague, stranger) is important as well. Muslukhov et al. [37] conducted surveys and interviews to investigate users' concerns about unauthorized access by insiders and strangers. They concluded that observing unlock attempts, memorizing it, and thus gaining unauthorized access by insiders are highly likely to occur. That is directly linked to insiders' ability to observe interactions closely and repeatedly. Farzand et al. [16] showed that the type of relationship impacts the choice of mitigation behavior. Moreover, depending on the relationship with the attacker, victims often do not want them to know they were caught.

Victim–Attacker Relative Pose: To successfully shoulder surf, the content on the target device must be directly visible to the attacker (unless we reconstruct the screen content from visual reflections with machine learning [60]). Thus, the relative *pose between victim and attacker* is important, as the used term shoulder surfing illustrates. A sitting pose, for example, enables shoulder surfing more than a standing pose [46]. Furthermore, *viewing angle* and *distance* play an important role as well [6]. However, tilting the device away from the observer, a widely adopted defense strategy, provides limited protection from shoulder surfing attacks [25].

Type of Device: Different devices can be the target of a shoulder surfing attack, including but not limited to notebooks, tablets, smartphones, and smartwatches. However, shoulder surfing can also occur when using shared devices or accessing private information on public devices [9]. The main prerequisite for shoulder surfing is that a bystander can observe the user's screen. Hence, smartglasses are unaffected and can be used as a mitigation strategy [58].

Type of Content: Mainly, two different types exist: (1) authentication-based and (2) content-based shoulder surfing [18]. The primary focus of many shoulder surfing studies is to investigate secure password or PIN entry [8]. While authentication is, of course, important and prone to observational attacks, other types of content can also be observed. Moreover, content-based shoulder surfing is more frequently experienced than authentication-based shoulder surfing [18]. Previous work has examined different content types such as notifications, texts, photos, social media, and gaming [6, 46]. Nevertheless, while different types of content are affected by shoulder surfing, there are differences in their perceived sensitivity [17].

Type of Environment: Shoulder surfing can take place in different environments such as buses, trains, airports, coffee shops, and restaurants [3]. These environments can be classified in two different ways. One can either distinguish private, semi-public (work), or public contexts [45] or differentiate between personal and professional contexts [62]. Independent of the classification choice, the location cannot be neglected when studying shoulder surfing attacks as it influences victim and attacker behavior [48].

2.2 Research Methods

As outlined in the chapter "Empirical Research Methods in Usable Privacy and Security" , privacy and security research has applied various methods. In this section, we highlight the methods that were previously used to study shoulder surfing. In summary, we classify these methods into four categories: (1) surveys and interviews, (2) lab studies, (3) field/in-the-wild studies, and (4) studies in extended reality. The following subsection describes the different methods and highlights their advantages and disadvantages. Our goal is to provide an overview of the different methods to support researchers and practitioners (new to the field) in deciding which method to apply in their research.

Surveys and Interviews Surveys and interviews are helpful tools for privacy researchers to gather valuable insights into a broader population or specific user groups [36]. The difference between surveys and interviews is that in interviews, a researcher takes an active role and directs questions to the interviewee (cf., Lazar et al.[27, 28]), while in surveys, a set of predefined questions is presented to the participants. With surveys and interviews, it is possible to achieve various objectives. On one side, researchers can use them to gather evidence for shoulder surfing attacks in the real world and get insights into personal experiences with the phenomenon from both victims and attackers of shoulder surfing incidents (e.g., [14]). On the other side, they help to understand preliminary performance metrics of authentication techniques against observation attacks (e.g., robustness [4]) and can even be used to quantify which parameters of these techniques help to make them less observable (e.g., [54]). Different approaches to constructing surveys exist. Noticeable is the inclusion of video material to present recreations of shoulder surfing attacks to participants [4]. Aviv et al. [5] show that these videos embedded in surveys can achieve results comparable to user studies in the lab.

Compared to other research methods, surveys allow larger sample sizes as researchers can reach and recruit more participants. Nevertheless, sample sizes vary enormously for shoulder surfing research. Previous work has reported studies with more than 1000 participants ($n = 1173$) [4] to smaller numbers that remain in the hundreds (e.g., $n = 298$ [54] or $n = 174$ [14]). Compared with other research methods, surveys often report higher numbers of participants. Recently, crowdsourcing platforms have entered the stage of privacy research and provide researchers with access to different user groups (that can be specified concerning various dimensions) [23]. Nowadays, researchers can more easily recruit a diverse set of participants.

In addition to surveys, in-depth interviews can be a sensible next step that allows scientists to understand the reasons behind the observed data [14]. Nonetheless, interviews can also be applied as a standalone method. For interviews, the more active participation of a researcher asking questions can lead to more detailed responses [28]. Moreover, interviews allow the live demonstration of specific techniques under controlled conditions. For example, the interviewer can present different shoulder surfing mitigation strategies to participants during the interview [16].

Finally, there has been a recent study that explored shoulder surfing through a longitudinal investigation, meaning they performed a diary study with 23 participants over one month [18]. They found that content-based shoulder surfing takes place more frequently than authentication-based shoulder surfing.

While we presented different methods in this part, they all have in common that they rely on self-reporting. While self-reporting is frequently deployed in privacy research, it has a few noteworthy drawbacks. As researchers do not directly observe a phenomenon, factor, or effect, they rely on the subjective perception of the participant, which can include a recall bias [43]. Moreover, not every type of information can be gathered with self-reporting; however, asking indirect and anonymity-preserving questions can minimize social desirability bias [33, 53].

Lab Studies Scientists often conduct experiments to answer their research questions concerning shoulder surfing. In experiments, it is often necessary that researchers can observe a shoulder surfing situation taking place. Due to the challenges of researching the phenomenon during field or in-the-wild studies (see below), these studies are primarily carried out in the lab. Moreover, compared to surveys and interviews, recruiting participants is more difficult, and conducting the experiment is often more workload-intense. As a result, experiments generally report smaller sample sizes. Nevertheless, a lab study also has certain advantages, for example, compared to field or in-the-wild studies. The most significant benefit (compared to other study types) is the high degree of control over the experimental conditions. Moreover, a lab study allows gathering consent from all involved parties before the experiment.

When conducting a lab study to research different dimensions of a shoulder surfing attack (e.g., the resilience of authentication techniques against human shoulder surfers), a challenge is to replicate these attacks for the study [56]. In lab studies, participants often take over the role of the attacker (e.g., [46]). Nevertheless, it remains challenging to replicate realistic attacks, as often they are performed out of boredom in opportunistic moments [14]. Simply instructing participants to perform a shoulder surfing attack would broadly differ from the behavior observable during an actual attack. To overcome this challenge, researchers have designed studies that inform participants about the study's goals toward the end (e.g., [46]). These studies partially deceive participants by leaving out specific study details not to influence their behavior. However, it should be noted that deceiving participants in a user study can be problematic and not justified. Hence, it is strongly encouraged to balance ethical implications and knowledge gain and act cautiously when deceiving participants.

A different approach is to research factors and effects that are not related to the timing, occurrence, or behavior of shoulder surfing attacks but instead focus on aspects that can be researched with the research goal out in the open. For example, a previous study has investigated the effect viewing angle and distance have on the success of shoulder surfing attacks [6]. Here, a lab study can offer control to isolate research factors from others that would introduce too much complexity to the experiment.

In-the-Wild or Field Studies Researching the phenomenon of shoulder surfing with in-the-wild or field studies sheds more light on the contexts in which these attacks take place and could provide insights into the behavior of attackers and victims. However, performing these studies is very challenging and, thus, rarely conducted. One of these studies was a two-week in-the-wild study conducted by Schneegass et al. [48], where they investigated the likelihood of shoulder surfing attacks occurrence during unlock events. Nonetheless, shoulder surfing is socially unacceptable and privacy-invasive. Hence, observing these attacks requires consent, potentially biasing participants and making it very difficult to observe authentic interactions. Moreover, outside the lab, bystanders get involved quickly; when that happens, their consent is also necessary (e.g., when recording video for eye tracking). In the past, researchers have primarily relied on surveys and interviews to assess in-the-wild experiences [14], relying on self-assessment as the most frequent research method. To encompass both the benefits of a study in the lab (such as its associated high degree of control) and to enable researching more realistic (in situ) shoulder surfing scenarios, researchers have applied eXtended Reality as a study method.

Studies in Extended Realities Recently, eXtended Reality (XR) [42] entered human–computer interaction (HCI) as a means to conduct user studies that are not directly related to XR but use XR as a modality to conduct user studies instead (e.g., [31]). This is particularly the case for user studies that are taking place in virtual reality (VR) in a virtual environment (VE), whereas XR could implicate "augmented reality" (AR) or "mixed reality" (MR) as well. The trend of using XR as a research method got amplified with the ongoing Covid-19 pandemic as different frameworks appeared [19, 40].

Using VR to research the shoulder surfing phenomenon has several inherent benefits. First, a virtual environment allows a more believable recreation of a real-life situation, which would otherwise be hard to recreate in the lab (e.g., a bus stop or office environment with different people present [2]; see Fig. 2). In addition to the realistic recreated scenes, VR allows maintaining the consistency among study participants, avoiding external uncontrolled situations. With eye trackers embedded in the head-mounted displays (HMD), researchers are able to capture and analyze the gaze of the participants. Accordingly, they are able to profoundly understand the observation attacks cycles and expect what triggers the observers' attention. As VR is associated with a high degree of immersion, it allows placing the subject in a simulated, virtual environment, where they can experience the situation as intended

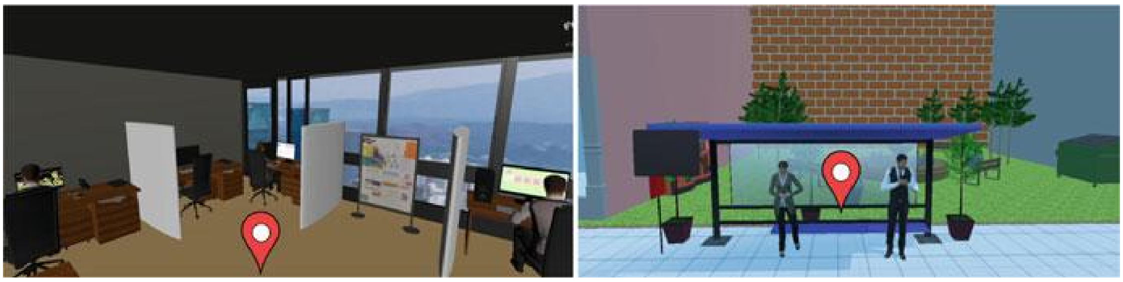

Fig. 2 Example taken from a previous paper that studied shoulder surfing in virtual reality [2]. The figure shows two virtual scenes that were used to investigate observing others' displays in an open office space (left) and a bus stop (right). The read markers indicate the participants' initial position

by the researchers. Here, the degree of presence can be assessed through the usage of presence questionnaires [50, 51, 59].

Potentially, such studies can also run outside the lab on HMDs owned by participants [40], and they were validated for usable security evaluations [35]. Additionally, user studies in XR allow fulfilling particular requirements specific for shoulder surfing studies. One is *privacy*, as conducting a user study in a real-world environment with real victims can be considered ethically challenging, whereas shoulder surfing a virtual avatar in a virtual environment (VE) is less likely an issue. Furthermore, conducting a user study in a VE allows for a very high degree of control since the environment is simulated by a computer, often exceeding the capability of control that an experimenter has over a real-world situation, even if it takes place in a lab. The high degree of control allows for replicability of such user studies between participants, as the experienced situation can be made to be precisely always the same.

2.3 Key Findings on Shoulder Surfing Behavior

With the growing number of studies investigating shoulder surfing events, we highlight the key findings on observers behaviors that we believe are of high relevance.

Observations Are Often More Random Than Planned In the survey by Eiband et al. [14], the main findings showed that despite the fact that observations are frequently conducted on an opportunistic basis, they go beyond exposing the authentication. Several participants reported negative feelings when other content such as personal photos or texts are exposed.

Victim–Attacker Pose Relationships Are Unalike In 2021, Saad et al. [46] explored the tendency of bystanders to shoulder surf in a scenario within an underground train. To that end, they varied the point of view of the attacker (standing vs. sitting) and the position of the victim (again standing vs. sitting) and used a

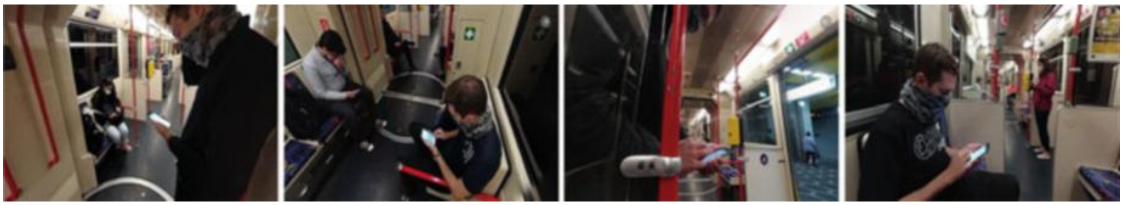

Fig. 3 User study conducted in virtual reality to investigate shoulder surfing attacks with prerecorded 360° videos [46]. Left to right: viewpoints of the participants with four different relative poses to the (virtual) victim: standing to standing, standing to sitting, sitting to standing, and sitting to sitting

360° camera to obtain a photorealistic recording of this setting, where several actors played either the role of the victim or became extras to simulate other people on the train. This recording then was played back to participants in a user study on an HMD that was equipped with an eye tracker in a lab study, and the point of view of the participants is seen in Fig. 3. Through the eye-tracking data, it was apparent that participants gazed at the object of interest, a smartphone held by the victim, and 11.16% of the time they were nearby.

VR Reflects Genuine Behavior… In 2022, Abdrabou et al. [2] conducted another project on the understanding of shoulder surfer behavior and the associated attack patterns. Here, they created a simulation in virtual reality with virtual, human-like avatars who were either located at a bus stop scene or within an office. The human participant of this study then was placed inside this VE through a VR HMD, which was again equipped with an eye tracker. The experimenters then recorded the participants’ gaze and their walking patterns in VR and found that participants looked at several objects of interest (e.g., smartphones in the bus stop scene or monitors in the office scene) 5.7 times on average, whereas the average eye contact duration was 1.61 s.

…but Immersion Is Needed. Also in 2022, Mathis et al. [34] considered the differences between non-immersive and immersive VR for shoulder surfing research and conducted a user study to explore the characteristics of both settings. They considered shoulder surfing attacks on automated teller machines, smartphone personal identification numbers (PIN), and smartphone pattern unlock mechanisms. They compare three scenarios, 2D video observations, 3D observations, and VR observations. The first scenario, 2D video observations, consists of the study participants watching a video of the shoulder surfing situation that they cannot influence on a traditional computer monitor, whereas in 3D observations, they could use the keyboard and mouse to walk around. These two conditions then were compared against each other and VR observations, where participants were wearing a VR headset and could freely move around and adjust their observation perspective. The authors found that VR observations lead to a significantly higher sense of presence and involvement and that VR observations also lead to the most accurate shoulder surfing observations.

There Is More than Smartphones There are other devices that are becoming more ubiquitous nowadays, smartwatches for instance. Recently, more studies are proposing authentication approaches for smartwatches, with resilience against shoulder surfing as a key metric for robustness [38, 39].

In conclusion, we can observe that there is an increasing number of publications that utilize XR, particularly VR, as a research method for shoulder surfing research. The high degree of immersion lets the participants of a user study easily take the role of the attacker, while such a lab study setting allows for an efficient resolution of the problematic aspects connected to ethics in this kind of research. Furthermore, VR allows the study to be exactly the same for each subject, as the computer-driven simulation creates an easily repeatable environment. Thereby, realistic scenarios can effectively be replicated in the lab.

3 Mitigating Shoulder Surfing Attacks

For the mitigation of shoulder surfing attacks, it is important to note that not every shoulder surfing incident is equally problematic. One important aspect to consider is the type of content visible. For content-based shoulder surfing, we need to understand what is considered sensitive content as it plays an important role in selecting a suitable mitigation strategy. To tackle this challenge, Farzand et al. [17] present a typology of perceived sensitivity that can help to understand the content sensitivity. Furthermore, one needs to take into account that the perception of shoulder surfing is different between cultures [47]. As a consequence, it also differs what is considered sensitive content.

In the following section, we look at research that aims to find solutions to mitigate shoulder surfing attacks. Therefore, we start by looking at different threat models against which researchers and practitioners can evaluate their mitigation strategies. After that, we briefly describe technical approaches to detect shoulder surfing and their current limitations. Finally, we present an overview of different mitigation strategies.

3.1 Threat Models

Threat models provide a systematic approach to investigate potential weaknesses to privacy and security [32]. For shoulder surfing, different threat models have been considered in the literature. Below, we provide a selection of these models and describe them briefly. It should be noted that also mixes of these are possible (e.g., a repeated attack that is technology-supported [7]):

Weak Attacks: A shoulder surfing attack is considered a weak attack if it is performed by a human observer without the help of any technology and with only limited practice [11].

Trained Shoulder Surfers. Compared to weak attacks, trained shoulder surfers are more effective by training themselves. They often employ cognitive strategies that help to reach higher success rates [26]. Please note that trained shoulder surfers manage to be more effective without using recording devices.

Repeated Attacks: The repeated attacks threat model assumes that an attacker can repeatedly observe the target device of the victim. Moreover, this threat model often considers the attacker to be at close range—the attacker quite literally looks over the victims’ shoulder [7].

Insider Attacks: Quite similar to the repeated attacks threat model are the insider attacks. The main difference is that for this type of attack, family, friends, or colleagues perform them. They may repeatedly observe the victim, and by combining these partial observations, it is easier to form a hypothesis on the victim’s secret [57].

Multiple Attackers: The shoulder surfing attacks become more threatening when multiple attackers try to observe the target device. In this case, attackers can coordinate by either focusing on specific parts or organizing distraction and information stealing roles between attackers [24].

Technology-Supported Attacks: The probably strongest form of shoulder surfing attacks are technology-supported ones. In these cases, an attacker is recording the victim’s interactions, for example, when drawing money from an ATM [10]. With recent technology advances, camera-based sensors can be manufactured in very tiny proportions, allowing attackers to seamlessly integrate them in their clothing or accessories. When analyzing the recorded data with machine learning, breaches of privacy are possible even when the attacker is not direct line of sight because reflections on glasses are sufficient for reconstruction of screen content [60].

3.2 Algorithmic Detection of Attacks

To mitigate shoulder surfing attacks, they first need to be detected. In previous research, detecting shoulder surfing attacks is primarily achieved by focusing on the human attackers. Here, algorithmic approaches oftentimes rely on visual sensor data (i.e., monochrome and RGB cameras). As shoulder surfing is frequently researched for mobile devices, the built-in camera is a good source for visual information to detect attackers. For example, Ali et al. [3] investigated the use of the built-in camera on mobile devices to detect if an unauthorized person tries to gain access to the device. Here, to detect an observer, face detection is applied to the incoming video feed. Interestingly, popular operating systems such as Android come with real-time face detection capabilities that can be used for detecting

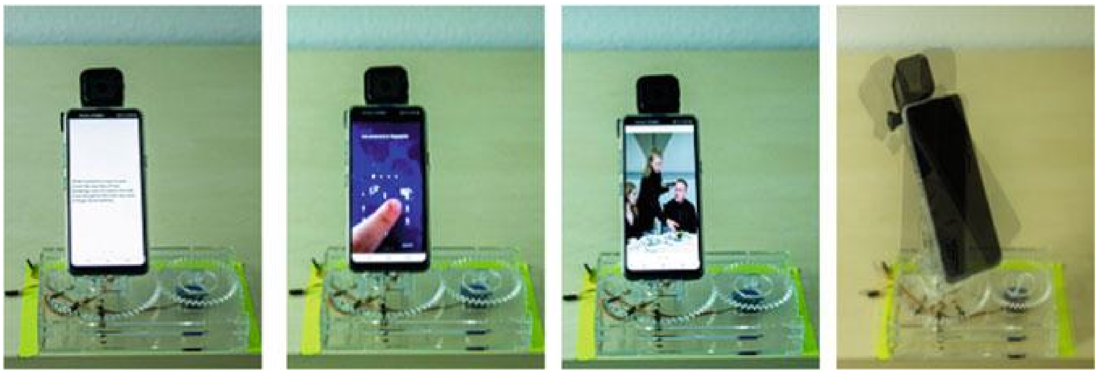

Fig. 4 Study apparatus to investigate the influence of distance and viewing angle on shoulder surfing success rate, figure taken from Bâce et al. [6] licensed under CC BY-NC-ND 4.0. The subfigures show examples of different content types on the phone display: (left to right): text, PIN, photo, and no content visible. The mechanical prototype visible rotated the smartphone between 0, 30, and 60°

shoulder surfers [7]. Nonetheless, not every detected face is necessarily a potential attacker as other factors play an important role as well, such as gaze direction and context, among others. In a recent study, different angles and distances have been investigated to understand which of them are most critical as they provide a good position for shoulder surfing [6]. The threat model was also based on evaluating people's perception on the displayed content that varied between visual, textual, and authentication, as seen in Fig. 4.

Nevertheless, visual detection of potential shoulder surfing also comes with a few downsides. First, they require the camera to be active and to record the scene. This scene likely involves the users of the device as well and, thereby, introduces another privacy risk. Furthermore, not only the privacy of a user may be violated, but also that of bystanders (as it continuously records the scene). Another issue is that the continuous recording and processing of the video feed drains the battery more quickly [7]. Hence, researchers have explored other options as well. For example, Lian et al. [29] used "multiple sensors, i.e., video camera module, ultrasonic distance module, light sensor module, to detect screen peeping, user distance and environmental lightness." Here, future studies should compare the different sensor technologies and develop adaptive strategies that take the context into consideration. For example, when a user is logged in to their wireless network at home and no other Bluetooth signatures are around, continuous monitoring via the built-in camera to detect shoulder surfing may not be necessary.

3.3 Prevention Strategies

Oftentimes, a detection algorithm proposed by researchers goes hand in hand with an implementation of a mitigation strategy (cf. [44]). In the following, we discuss two different strategy types into which proposed systems can be classified.

On one side, there are strategies that try to be generalizable toward every kind of content, and on the other side, there are strategies that focus on mitigating attacks against specific types of contents. These two strategies are in line with how we categorize shoulder surfing attacks into authentication-based and content-based shoulder surfing. Here, it is important to note that while authentication-based shoulder surfing is perceived as more problematic, content-based shoulder surfing is occurring more frequently [18].

Strategies Independent of Content Often times, researchers propose systems that mitigate shoulder surfing attacks independent of the content shown by the target device. Different systems have been proposed that try to create awareness for an actively ongoing shoulder surfing attack. For example, Ali et al. [3] proposed a system that informs users whether text on the screen could be read by an attacker. To better understand, in which way users want to be alerted, researchers have conducted a user study to compare four different methods: vibro-tactile, front LED, on-screen icons, and video feedback, finding that vibro-tactile feedback works best, as seen in Fig. 5. Their findings showed that vibration feedback allowed for a faster response time, in comparison to the other three methods [44]. Moreover, it has been examined how additional parameters such as distance and orientation can benefit victims in applying appropriate actions [62].

While awareness-based systems leave it to the user to decide on how they want to react, researchers have proposed different strategies that help users in their actions [9] or automatically react to shoulder surfing attacks [29]. Here, users can either move or hide information presented on the screen by performing explicit interactions [9] or information is automatically masked [9, 29] (e.g., with the help of eye tracking [41]). Lian et al. [29] found that with limited brightness or contrast, only the user could read the screen, while others have trouble reading it [29].

Furthermore, different strategies have been proposed that do not rely upon detecting a shoulder surfer at first, but rather are applied constantly. For example, Chen et al. [12] developed Hide Screen, which utilizes human vision characteristics to preserve privacy. Simplified, the approach allows changing the readability of information based on the viewing angle. Instead of hiding the information from an attacker, Watanabe et al. [55] suggest adding additional information that is designed to throw an attacker off. They suggest showing multiple cursors on the screen

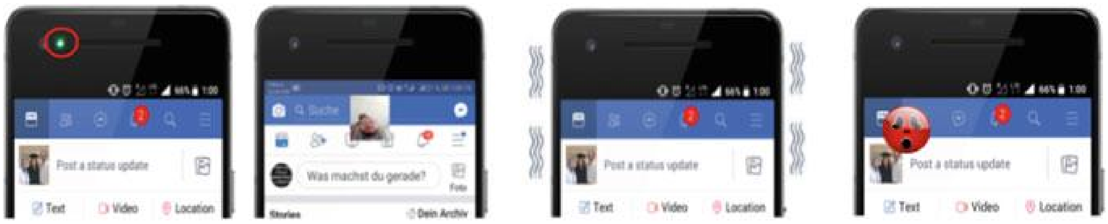

Fig. 5 Different feedback conditions to communicate a shoulder surfing incident investigated in previous work [44]. The different feedback conditions are (from left to right): (1) front LED, (2) video preview, (3) vibro-tactile, and (4) on-screen icon. The authors found that vibro-tactile feedback results in the lowest reaction time

and, thereby, effectively hiding the real cursor for an observer. Finally, it has been proposed to extend an observable screen with a second screen that is not observable and can be used to show private information. For example, Winkler et al. [58] are using smartglasses to show private information that would have otherwise be shown on the smartphone display.

Strategies Focused on Specific Types of Content Because not every type of content requires the same level of protection, many proposed strategies that are highly dependent on the type of content that they protect. In particular, authentication approaches need high protection against shoulder surfing attacks. Hence, researchers have suggested a variety of authentication techniques that are more resilient against observational attacks.

Bianchi et al. [7] proposed to use a composition of non-visual cues (i.e., audio and haptic cues) to enter a password. As a result, an observational attack cannot rely on visual information only to decipher the password. Furthermore, others have suggested to use gaze as an input modality in combination with graphical passwords [10]. Thereby, an attacker would need to observe the eye gaze of the victim additionally to the phone screen, making it very challenging to reconstruct the password. Another strategy is to extend the input surface for the authentication scheme toward the backside of the smartphone, which is more difficult to observe [13].

Besides authentication approaches, researchers have focused on other types of content. For example, Eiband et al. [15] have investigated how text can be presented in a way that is readable to the user but unreadable to an observer. In essence, they propose to display text in the user's own handwriting. While this does not prevent an attacker from reading the text, it significantly slows them down.

4 Challenges and Future Research Directions

In the following, we present challenges and research directions concerning the methodology of researching shoulder surfing and the phenomenon itself. These are particularly related to the methodology of shoulder surfing research and the attacker's behavior.

Research Methods to Investigate Shoulder Surfing While conducting research on shoulder surfing in the wild, several challenges regarding the methodology became apparent. First of all, a central element is an ethical dilemma associated with the necessity of obtaining the shoulder surfer's consent. When researchers ethically design an experiment on shoulder surfing that involves participants, participants usually have to get into the role of either the victim or attacker. However, shoulder surfing usually is an interaction that is very affective by its nature [14], hence instructing participants on the roles that they should get into highly inflects their behavior, and thus, results elicited from the study. Consequently, there is a dichotomy between asking for consent and subjects' unchanged behavior that needs to be weighed individually for each study, taking the objectives of the study into account.

Another argument on shoulder surfing studies is to *simultaneously* consider both roles of the attacker and the victim. Considering only the role of the observer and not the victim could leave out vital parts of the shoulder surfing incident, such as the occlusion of the phone display by the victim [6].

Virtual Reality for User Studies To overcome some of the challenges related to this ethical dichotomy, several research projects utilized virtual reality to simulate the shoulder surfing interaction with virtual avatars [2, 34, 46]. Although it is not necessary to obtain consent from a virtual avatar that has the role of the victim, it, however, still is necessary to obtain consent from a participant that gets into the role of the attacker. Furthermore, virtual reality allows for a simulation of the environment; hence, the interaction can be explored in different settings that would be hard to replicate in a physical lab.

However, virtual reality is also only a limited solution, as there are certain aspects impacted by the simulation of the environment. For example, today's head-mounted displays can influence people's behavior such as their movement [20] or also their social comfort distance that is less in virtual reality than in reality [22]. They can, however, help in recreating scenarios from the real-world by simulating them in a lab, as conducting field studies or in-the-wild experiments is particularly challenging due to the ethical aspects, particularly, when uninvolved third parties become part of the investigation. The same applies to other methodologies such as the usage of recording videos outside the lab, the so-called "lifelogs", as using cameras impacts the protection of private information of both the wearer and potential bystanders [21].

Identifying Sensitive Content In general, two types of shoulder surfing are distinguished: authentication-based and content-based shoulder surfing. While authentication-based shoulder surfing is inherently problematic as it exposes sensitive information (e.g., PIN or password), it is more complicated for content-based shoulder surfing that happens more frequently [18]. Privacy is an individual concept. Hence, what one person considers sensitive information may not be considered sensitive by someone else. This makes it very difficult to have an overall solution that equally protects all users. As a consequence, we need to investigate what content is considered sensitive (e.g., [17]). Furthermore, we need to examine different factors that can influence the perception of what is considered sensitive content such as cultural differences [47].

Understanding the Attacks and Behavior Another open research direction is to create an understanding of the shoulder surfing interaction itself, by, for instance, creating models of it. Here, Abdrabou et al. have created one of the first works in creating a model of attack patterns [2]. Their study took place in virtual reality; hence, creating a model-based understanding of the phenomenon, in reality, is still an open research opportunity nowadays. It is therefore necessary to conduct further studies to determine more substance to derive models about behavior within more contexts of the interaction. This includes, but is not limited to, in-the-wild studies as well as long-term studies to understand, whether the behavior changes over time.

Additionally, recent studies focus on password attacks but do not have a strong focus on understanding shoulder surfing behavior in general [8]. However, when considering only the attacks on passwords, such as android pattern locks, models were already created that predict the grade of observability [54]. This also opens up the opportunity to further explore the type of content that is particularly attracting shoulder surfing attacks, which partly has been covered by recent studies [2, 46].

5 Conclusion

In this chapter, we presented lessons learned from research on the shoulder surfing phenomenon and attack mitigation strategies. We started with a definition of shoulder surfing and an introduction of different types of attacks. After that, we present different research methods that have been applied in the past and discussed key findings related to shoulder surfing behavior. Next, we gave an overview of different threat models for shoulder surfing and discussed algorithmic detection of these attacks and different mitigation strategies. We concluded the chapter with an outlook on persistent challenges and future research directions. We believe that this book chapter offers a great starting point for new researchers and practitioners in the field. Moreover, we see great potential for eXtended Reality to overcome the limitations that field and in-the-wild studies introduce.

References

1. Abdelrahman, Y., Khamis, M., Schneegass, S., & Alt, F. (2017). Stay cool! Understanding thermal attacks on mobile-based user authentication. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (pp. 3751–3763).
2. Abdrabou, Y., Rivu, S. R., Ammar, T., Liebers, J., Saad, A., Liebers, C., Gruenefeld, U., Knierim, P., Khamis, M., Makela, V., Schneegass, S., & Alt, F. (2022). Understanding shoulder surfer behavior and attack patterns using virtual reality. In P. Bottoni & E. Panizzi, (Eds.), *Proceedings of the 2022 International Conference on Advanced Visual Interfaces* (pp. 1–9). ACM.
3. Ali, M. E., Anwar, A., Ahmed, I., Hashem, T., Kulik, L., & Tanin, E. (2014). Protecting mobile users from visual privacy attacks. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, UbiComp '14 Adjunct (pp. 1–4). Association for Computing Machinery.
4. Aviv, A. J., Davin, J. T., Wolf, F., & Kuber, R. (2017). Towards baselines for shoulder surfing on mobile authentication. In *Proceedings of the 33rd Annual Computer Security Applications Conference* (pp. 486–498).
5. Aviv, A. J., Wolf, F., & Kuber, R. (2018). Comparing video based shoulder surfing with live simulation. In *Proceedings of the 34th Annual Computer Security Applications Conference*, ACSAC '18 (pp. 453–466). Association for Computing Machinery.

6. Bâce, M., Saad, A., Khamis, M., Schneegass, S., & Bulling, A. (2022). PrivacyScout: Assessing vulnerability to shoulder surfing on mobile devices. *Proceedings on Privacy Enhancing Technologies, 1*, 21.
7. Bianchi, A., Oakley, I., Kostakos, V., & Kwon, D. S. (2010). 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 (pp. 197–200). Association for Computing Machinery.
8. Bošnjak, L., & Brumen, B. (2020). Shoulder surfing experiments: A systematic literature review. *Computers & Security, 99*, 102023.
9. Brudy, F., Ledo, D., Greenberg, S., & Butz, A. (2014). Is anyone looking? Mitigating shoulder surfing on public displays through awareness and protection. In *Proceedings of The International Symposium on Pervasive Displays*, PerDis '14 (pp. 1–6). Association for Computing Machinery.
10. Bulling, A., Alt, F., & Schmidt, A. (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 (pp. 3011–3020). Association for Computing Machinery.
11. Chakraborty, N., & Mondal, S. (2014). An improved methodology towards providing immunity against weak shoulder surfing attack. In A. Prakash & R. Shyamasundar (Eds.), *Information Systems Security* (pp. 298–317). Springer International Publishing.
12. (Daniel) Chen, C.-Y., Lin, B.-Y., Wang, J., & Shin, K. G. (2019). Keep others from peeking at your mobile device screen! In *The 25th Annual International Conference on Mobile Computing and Networking*, MobiCom '19. Association for Computing Machinery.
13. De Luca, A., Harbach, M., von Zezschwitz, E., Maurer, M.-E., Slawik, B. E., Hussmann, H., & Smith, M. (2014). Now you see me, now you don't: Protecting smartphone authentication from shoulder surfers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14 (pp. 2937–2946). Association for Computing Machinery.
14. Eiband, M., Khamis, M., Von Zezschwitz, E., Hussmann, H., & Alt, F. (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* (pp. 4254–4265).
15. Eiband, M., von Zezschwitz, E., Buschek, D., & Hußmann, H. (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 (pp. 2041–2048). Association for Computing Machinery.
16. Farzand, H., Bhardwaj, K., Marky, K., & Khamis, M. (2021). The interplay between personal relationships & shoulder surfing mitigation. In *Mensch Und Computer 2021*, MuC '21 (pp. 338–343). Association for Computing Machinery.
17. Farzand, H., Marky, K., & Khamis, M. (2022). "I hate when people do this; there's a lot of sensitive content for me": A typology of perceived privacy-sensitive content in shoulder surfing scenarios. In *Proceedings of the Eighteenth USENIX Conference on Usable Privacy and Security*. USENIX Association.
18. Farzand, H., Marky, K., & Khamis, M. (2022). Shoulder surfing through the social lens: A longitudinal investigation & insights from an exploratory diary study. In *2022 European Symposium on Usable Security* (pp. 85–97).
19. Gruenefeld, U., Auda, J., Mathis, F., Schneegass, S., Khamis, M., Gugenheimer, J., & Mayer, S. (2022). VRception: Rapid prototyping of cross-reality systems in virtual reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22. Association for Computing Machinery.
20. Hollman, J. H., Brey, R. H., Robb, R. A., Bang, T. J., & Kaufman, K. R. (2006). Spatiotemporal gait deviations in a virtual reality environment. *Gait & Posture, 23*(4), 441–444.
21. Hoyle, R., Templeman, R., Anthony, D., Crandall, D., & Kapadia, A. (2015). Sensitive lifelogs: A privacy analysis of photos from wearable cameras. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15 (pp. 1645–1648). Association for Computing Machinery.

22. Iachini, T., Coello, Y., Frassinetti, F., Senese, V. P., Galante, F., & Ruggiero, G. (2016). Peripersonal and interpersonal space in virtual and real environments: Effects of gender and age. *Journal of Environmental Psychology, 45*, 154–164.
23. Jin, H., Shen, H., Jain, M., Kumar, S., & Hong, J. I. (2021). Lean privacy review: Collecting users' privacy concerns of data practices at a low cost. *ACM Transactions on Computer-Human Interaction, 28*(5), 1–55.
24. Khamis, M., Bandelow, L., Schick, S., Casadevall, D., Bulling, A., & Alt, F. (2017). They are all after you: Investigating the viability of a threat model that involves multiple shoulder surfers. In *Proceedings of the 16th International Conference on Mobile and Ubiquitous Multimedia*, MUM '17 (pp. 31–35). Association for Computing Machinery.
25. Khan, H., Hengartner, U., & Vogel, D. (2018). Evaluating attack and defense strategies for smartphone PIN shoulder surfing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18 (pp. 1–10). Association for Computing Machinery.
26. Kwon, T., Shin, S., & Na, S. (2014). Covert attentional shoulder surfing: Human adversaries are more powerful than expected. *IEEE Transactions on Systems, Man, and Cybernetics: Systems, 44*(6), 716–727.
27. Lazar, J., Feng, J. H., & Hochheiser, H. (2017). Chapter 5: Surveys. In J. Lazar, J. H. Feng, & H. Hochheiser (Eds.), *Research methods in human computer interaction* (2nd ed., pp. 105–133). Morgan Kaufmann.
28. Lazar, J., Feng, J. H., & Hochheiser, H. (2017). Chapter 8: Interviews and focus groups. In J. Lazar, J. H. Feng, & H. Hochheiser (Eds.), *Research methods in human computer interaction* (2nd ed., pp. 187–228). Morgan Kaufmann.
29. Lian, S., Hu, W., Song, X., & Liu, Z. (2013). Smart privacy-preserving screen based on multiple sensor fusion. *IEEE Transactions on Consumer Electronics, 59*(1), 136–143.
30. Maggi, F., Volpatto, A., Gasparini, S., Boracchi, G., & Zanero, S. (2011). Poster: Fast, automatic iPhone shoulder surfing. In *Proceedings of the 18th ACM Conference on Computer and Communications Security* (pp. 805–808).
31. Mäkelä, V., Radiah, R., Alsherif, S., Khamis, M., Xiao, C., Borchert, L., Schmidt, A., & Alt, F. (2020). Virtual field studies: Conducting studies on public displays in virtual reality. In R. Bernhaupt, F. F. Mueller, D. Verweij, J. Andres, J. McGrenere, A. Cockburn, I. Avellino, A. Goguey, P. Bjørn, S. Zhao, B. P. Samson, & R. Kocielnik (Eds.), *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–15). ACM.
32. Marback, A., Do, H., He, K., Kondamarri, S., & Xu, D. (2013). A threat model-based approach to security testing. *Software: Practice and Experience, 43*(2), 241–258.
33. Marques, D., Guerreiro, T., & Carriço, L. (2014). Measuring snooping behavior with surveys: It's how you ask it. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '14 (pp. 2479–2484). Association for Computing Machinery.
34. Mathis, F., O'Hagan, J., Khamis, M., & Vaniea, K. (2022). Virtual reality observations: Using virtual reality to augment lab-based shoulder surfing research. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (pp. 291–300). IEEE.
35. Mathis, F., Vaniea, K., & Khamis, M. (2021). RepliCueAuth: Validating the use of a lab-based virtual reality setup for evaluating authentication systems. In Y. Kitamura, A. Quigley, K. Isbister, T. Igarashi, P. Bjørn, & S. Drucker (Eds.), *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1–18). ACM.
36. Müller, H., Sedley, A., & Ferrall-Nunge, E. (2014). Survey research in HCI. In *Ways of knowing in HCI* (pp. 229–266). Springer.
37. Muslukhov, I., Boshmaf, Y., Kuo, C., Lester, J., & Beznosov, K. (2013). Know your enemy: The risk of unauthorized access in smartphones by insiders. In *Proceedings of the 15th International Conference on Human-Computer Interaction with Mobile Devices and Services* (pp. 271–280).
38. Nagatomo, M., Watanabe, K., Aburada, K., Okazaki, N., & Park, M. (2019). Proposal and evaluation of authentication method having shoulder-surfing resistance for smartwatches using shift rule. In *International Conference on Network-Based Information Systems* (pp. 560–569). Springer.

39. Park, M., Aburada, K., & Okazaki, N. (2021). Proposal and evaluation of a gesture authentication method with peep resistance for smartwatches. In *2021 Ninth International Symposium on Computing and Networking Workshops (CANDARW)* (pp. 359–364). IEEE.
40. Radiah, R., Mäkelä, V., Prange, S., Rodriguez, S. D., Piening, R., Zhou, Y., Köhle, K., Pfeuffer, K., Abdelrahman, Y., Hoppe, M., Schmidt, A., & Alt, F. (2021). Remote VR studies: A framework for running virtual reality studies remotely via participant-owned HMDs. *ACM Transactions on Computer-Human Interaction, 28*(6), 1–36.
41. Ragozin, K., Pai, Y. S., Augereau, O., Kise, K., Kerdels, J., & Kunze, K. (2019). Private reader: Using eye tracking to improve reading privacy in public spaces. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI ’19. Association for Computing Machinery.
42. Rauschnabel, P. A., Felix, R., Hinsch, C., Shahab, H., & Alt, F. (2022). What is XR? Towards a framework for augmented and virtual reality. *Computers in Human Behavior, 133*, 107289.
43. Robins, R. W., Fraley, R. C., & Krueger, R. F. (2009). *Handbook of research methods in personality psychology*. Guilford Press.
44. Saad, A., Chukwu, M., & Schneegass, S. (2018). Communicating shoulder surfing attacks to users. In *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia*, MUM 2018 (pp. 147–152). Association for Computing Machinery.
45. Saad, A., Gruenefeld, U., Mecke, L., Koelle, M., Alt, F., & Schneegass, S. (2022). Mask removal isn’t always convenient in public!—The impact of the Covid-19 pandemic on device usage and user authentication. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI EA ’22. Association for Computing Machinery.
46. Saad, A., Liebers, J., Gruenefeld, U., Alt, F., & Schneegass, S. (2021). Understanding bystanders’ tendency to shoulder surf smartphones using 360-degree videos in virtual reality. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction* (pp. 1–8). ACM.
47. Saleh, M., Khamis, M., & Sturm, C. (2019). What about my privacy, Habibi?. In D. Lamas, F. Loizides, L. Nacke, H. Petrie, M. Winckler, & P. Zaphiris (Eds.), *Human-computer interaction —INTERACT 2019* (pp. 67–87). Springer International Publishing.
48. Schneegass, S., Saad, A., Heger, R., Delgado, S., Poguntke, R., & Alt, F. (2022). An investigation of shoulder surfing attacks on touch-based unlock events. In *Proceedings of the 24th International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI ’22. Association for Computing Machinery. To Appear.
49. Schneegass, S., Steimle, F., Bulling, A., Alt, F., & Schmidt, A. (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 (pp. 775–786). Association for Computing Machinery.
50. Schubert, T. W. (2003). The sense of presence in virtual environments: A three-component scale measuring spatial presence, involvement, and realness. *Zeitschrift für Medienpsychologie, 15*(2), 69–71.
51. Schwind, V., Knierim, P., Haas, N., & Henze, N. (2019). Using presence questionnaires in virtual reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, volume 2019 of CHI ’19* (pp. 1–12). Association for Computing Machinery.
52. Shin, H., Sim, S., Kwon, H., Hwang, S., & Lee, Y. (2022). A new smart smudge attack using CNN. *International Journal of Information Security, 21*(1), 25–36.
53. Tourangeau, R., & Yan, T. (2007). Sensitive questions in surveys. *Psychological Bulletin, 133*(5), 859.
54. von Zezschwitz, E., De Luca, A., Janssen, P., & Hussmann, H. (2015). 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 (pp. 2339–2342). Association for Computing Machinery.
55. Watanabe, K., Higuchi, F., Inami, M., & Igarashi, T. (2012). CursorCamouflage: Multiple dummy cursors as a defense against shoulder surfing. In *SIGGRAPH Asia 2012 Emerging Technologies*, SA ’12 (pp. 1–2). Association for Computing Machinery.

56. Wiese, O., & Roth, V. (2015). Pitfalls of shoulder surfing studies. In *In NDSS Workshop on Usable Security 2015 (USEC'15)* (pp. 1–6). Internet Society.
57. Wiese, O., & Roth, V. (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 (pp. 453–464). Association for Computing Machinery.
58. Winkler, C., Gugenheimer, J., De Luca, A., Haas, G., Speidel, P., Dobbelstein, D., & Rukzio, E. (2015). Glass unlock: Enhancing security of smartphone unlocking through leveraging a private near-eye display. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15 (pp. 1407–1410). Association for Computing Machinery.
59. Witmer, B. G., & Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments, 7*(3), 225–240.
60. Xu, Y., Heinly, J., White, A. M., Monrose, F., & Frahm, J.-M. (2013). Seeing double: Reconstructing obscured typed input from repeated compromising reflections. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, CCS '13 (pp. 1063–1074). Association for Computing Machinery.
61. Ye, G., Tang, Z., Fang, D., Chen, X., Wolff, W., Aviv, A. J., & Wang, Z. (2018). A video-based attack for Android pattern lock. *ACM Transactions on Privacy and Security, 21*(4), 1–31.
62. Zhou, H., Ferreira, V., Alves, T., Hawkey, K., & Reilly, D. (2015). Somebody is peeking! A proximity and privacy aware tablet interface. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '15 (pp. 1971–1976). Association for Computing Machinery.

Privacy Research on the Pulse of Time: COVID-19 Contact-Tracing Apps

Eva Gerlitz and Maximilian Häring

1 Introduction

In 2020, COVID-19 hit the World, and with it came the desire for a well-functioning and a fast-working possibility to trace contacts of those people who tested positive for the virus, a method called *contact tracing*.

> **Definition: Contact Tracing**
> Following the Cambridge Dictionary, contact tracing is "the process of finding any other people that an infected person has met or had close contact with, usually in order to control the spread of an infectious disease" [8]. Similar definitions are used elsewhere, e.g., by the World Health Organization (WHO) [11] and the European Centre for Disease Prevention and Control (ECDC) [9].

Early on, digital contact tracing was seen as a tool to interrupt chains of infection. This led to a discussion about apps to automatically trace and store with whom a user had been in contact with and, as a result, would warn those who might have become infected. Digital contact tracing was even advertised as a "key" in fighting the pandemic [12]. It has several advantages compared to a manual approach done

Authors are listed in alphabetical order and contributed equally

E. Gerlitz
Fraunhofer FKIE, Bonn, Germany
e-mail: gerlitz@cs.uni-bonn.de

M. Häring (✉)
Universität Bonn, Institut für Informatik, Bonn, Germany
e-mail: haering@cs.uni-bonn.de

N. Gerber et al. (eds.), *Human Factors in Privacy Research*,
https://doi.org/10.1007/978-3-031-28643-8_11

by health workers, e.g., that it enables to warn more people who else could not have been notified due to incomplete memory or knowledge about contacts of an infected person. Digital contact tracing also supports the authorities by notifying contacts of positive tested persons: Instead of calling each person one by one, the information can be transferred immediately to all persons at once.

Most of the digital contact tracing approaches were realized through smartphone apps. The idea of using apps that help fight a disease was not new in 2020. In Africa, e.g., an app supported contact tracing personnel in faster submitting the information to help combat Ebola in 2019 [33].

One of the first COVID-19 focusing apps was launched in February 2020 by the Chinese government. It was specifically designed to warn its users about a contact with someone who is infected with the virus [7]. Many other governments followed, and a lot of those contact tracing apps (CTA) based their tracing on Bluetooth or the users' location. As of March 2021, the MIT Technology Review lists 49 contact tracing apps in 48 countries from around the World [1] and an overview from Google lists 60 apps that make use of their provided framework [27].

Depending on how automated tracing is implemented, it is necessary to capture and store sensitive information about the user, such as where the user has been, who they were in contact with, and their health status. All of this entails the potential of mission creep and surveillance. Based on the possibility of misuse, a lot of public discussions in 2020 revolved around the architecture of such tracing apps. Many experts and organizations worldwide made a strong case for apps that should technically prevent abuse [10].

Researchers from the University of Oxford estimated what percentage of the population would need to install a contact tracing app for it to be effective, depending on further measures that were taken throughout the country. Their results indicate that adoption of 60% could stop the pandemic, but already smaller installation numbers would reduce the number of infections and deaths [17]. In public discussions, this number of 60% was often misreported to be the threshold that needs to be achieved in order to fight COVID-19 [24].

Taken together, these requirements (privacy preserving and the need to reach a large part of the population) were able to influence political decisions, e.g., in Germany [6], where the government switched to a more privacy-preserving app after another one had already been planned.

The concerns for misuse of the captured data were, in fact, not unfounded: Later, in at least one case in Germany, data of a private app that was used to check-in into restaurants and that stored the data centrally were used by the criminal investigation department to find witnesses of an accident. This happened even though it is illegal to use these data for law enforcement purposes, according to the Infection Protection Act for reasons of data protection [22].

In Singapore, data captured through the widespread contact tracing app "Trace-Together," about which it was claimed after its release that the data would only be accessed if a user tests positive for COVID-19, were used in a murder investigation [32].

So, privacy has been a big topic in the development and the public discussions centered around contact tracing apps. But how big of a role does and did privacy actually play in the mind of potential users when they needed to decide whether or not to install a contact tracing app? And what can privacy research learn from that?

This chapter is a starting point for every reader interested in these questions. In this chapter, we:

- Give a brief outline of the tracing technologies and their implications for the users' data and therefore privacy.
- Look at scientific studies with end users and how their privacy concerns impacted their decision to install a contact tracing app.
- Set the study results in the context of the used methodology (e.g., the time the study was conducted or who was asked).

After reading this chapter, the reader will have an overview of the general privacy discussion on contact tracing apps in the context of COVID-19 and hints on where to find further information.

2 Tracing Technologies

This section gives a brief overview of technical possibilities to automatically warn people who had been in contact with someone who later tests positive for COVID-19. Worldwide, different versions of contact tracing apps were proposed, discussed, and rolled out. The task of apps in this context ranged from simply informing users about their contact and asking them to start a voluntary quarantine (e.g., in Germany [25]) to functioning as access control (e.g., in China [23]).

Obviously, it is (currently) not feasible to technically directly track whether a person met another person; therefore, many solutions use the personal smartphone as a proxy. The apps capture whether a device was in proximity to another device, therefore also called proximity tracing. For simplicity, we assume in the following that people always carry their smartphones with them, and we will use the ideas of "Who met whom" and "Which device encountered which device" interchangeably.

The following two sections detail the steps of such a digital contact tracing: The tracing itself and the details of when and how a user is informed about meeting someone who tested positive. Our goal is to give enough detail about the essential technology for the reader to have a general overview and can follow the debates around the different apps, their approaches, and possible implications for the users. Please note that this is not a complete list of technologies.

2.1 Proximity Tracing

For a contact tracing app to work, first and foremost, it must be logged who was in contact with whom. There are different approaches to accomplish this and different ways to categorize them: Huan et al. [18], for example, used a categorization where approaches are separated based on the data collection method: *cell phone base station data*, *location history*, and *Bluetooth proximity data*. Another possible taxonomy could be built based upon the interaction and setup needed (e.g., device-to-device communication directly via Bluetooth), indirect via participation tracking (e.g., at an event through QR codes [15]), or the not-so-common usage of already existing data (e.g., cell phone base station data).

To understand a lot of the research focusing on privacy in the contact tracing context, one has to look at the storage location of the logged contact data and the usage of Bluetooth Low Energy (LE). It works as follows: devices broadcast IDs via Bluetooth LE. The received IDs are stored together with the sent ones, and some information is added/derived, such as a distance and time metric. Those stored IDs are later matched with a list of IDs representing infected persons. If a device keeps the gathered IDs stored locally and compares them locally to a public list of IDs representing an infected person, the approach is called *decentral*. On the other hand, *central* means that the devices upload at least the seen and gathered IDs to a central entity/server.

Both approaches have their disadvantages, but the threat model differs. In the centralized approach, parties hosting or having access to the service (e.g., the government) could gain access to the data [28]. In this case, the third party could, for example, learn about the users' social graph. Compared to this, in the decentralized approach, an attacker needs to be in close vicinity to gain knowledge, as explained by Baumgärtner et al. [5].

Independent of how the approaches are categorized, tracing was discussed in many different ways, and for further research in this area, we suggest further literature and projects (e.g., [4, 5, 14, 26, 29]).

2.2 Risk Calculation and Informing Those at Risk

For efficient contact tracing, it is not only necessary to trace contacts, but also to inform those who had been in close contact with infected people (and possibly also give advice or instructions on how they should behave). This can be divided into the following three problem spaces:

Medical Basis for Risk Calculation The fundamental question is who should be informed and under what circumstances. For this, requirements from epidemiologists and virologists need to be implemented, concerning, for example, the distance and time after which an infection becomes more likely.

Technical Implementation of Risk Calculation There are different possibilities for where the actual risk calculation can occur. Research and politics in the EU favored mainly the previously outlined decentralized approach. In this approach, the assessment of whether the user is at risk is calculated on the phones directly. In the centralized approach, this calculation happens on a central server. Independent of the approach is the fact that the risk calculation can only be an estimation of what actually happened. False positives and true negatives have to be balanced. On either side, it can result in a negative effect on the adoption and effectiveness of the app.

How to Inform Those at Risk In the decentralized approach, no central entity knows the contacts of an infected person and therefore cannot inform them. Each device itself is "responsible" to inform its user. In a centralized setting, the server knows who is at risk. Therefore, even out-of-band contact, e.g., via phone, is possible depending on what data are available.

3 Privacy and Contact Tracing Apps—User Studies

The previous sections concerned technical circumstances to give the reader an overview of the situation. This section now focuses on the end user, thus the person owning a smartphone, and who is the potential user of an app. We give insight into what the studied participants think about contact tracing apps in terms of privacy, and how privacy considerations impact the willingness to use such apps.

For this, we conducted a literature review. In 2020, the topic of contact tracing apps was highly relevant and design decisions needed to be made urgently, so many researchers around the world examined the effect of different app properties and their general acceptance in the public population: The ACM Digital Library [2], for example, as of September 2022, lists around 32K publications published since 2020 when searching for "contact tracing."

We thus specified our search term such that the terms "contact," "trac*," and "priv*" had to be found in either the title or the abstract. Our full search comprised the databases ACM Digital Library [2], IEEE Xplore [19], and Web of Science [38]. We also analyzed the Google Scholar top twenty security conferences and journals if their names included "privacy" and the A* and A CORE-ranked privacy conferences and journals. Only those that were not already included in the previous database search underwent a manual title search. This included the Symposium On Usable Privacy and Security (SOUPS) and the International Conference on Security and Privacy for Communication Networks (SecureComm).

After this search, we ended up with 245 papers. We manually reviewed all abstracts and only picked those that fit our requirements. Articles were excluded if they matched the following criteria:

- Not related to contact tracing technology to combat COVID-19.
- No user study was conducted. (This included all studies that looked at user feedback from the App stores of Apple or Google.)
- The user study did not look at sentiments of users concerning the privacy aspects of contact tracing apps.

We ended up with 13 papers that are covered in this chapter. Table 1 gives a brief overview of the included studies.

It must be noted that because of the urgency and its possible high relevance to ongoing discussions, many studies were not only published in a peer-reviewed conference or journal but faster published, e.g., by uploading on arXiv. Those are not necessarily of bad quality but have to be read more carefully than work that was

Table 1 Brief overview of the presented studies. If a specific contact tracing app was investigated, this information is included in brackets

Authors	Country	*n*	Purpose of the app (CT = Contact tracing)	Used standardized questionnaire?
Huang et al. [18]	USA	44	CT, Home quarantine, Epidemiological investigation support system, Information tracking of dine-in customers, E-permit service	No
Häring et al. [16]	Germany	744	CT (CWA)	No
Utz et al. [37]	Germany, USA, China	1003, 1003, 1019	CT, Symptom Check, Quarantine Enf., Information, Health Certificate	IUIPC, 2004
Redmiles et al. [28]	USA	1000	CT, Information	No
Xie et al. [39]	Ireland	286	CT (COVID Tracker)	Westin's privacy segmentation index (PSI), privacy attitude questionnaire (PAQ)
Trestian et al. [36]	Ireland	258	CT (COVID Tracker)	Westin's privacy segmentation index (PSI)
Lu et al. [21]	USA	291	CT (identifying and notifying close contacts) + monitoring symptoms	No
Dooley et al. [13]	USA	7,010,271 impressions	CT	–
Zampedri et al. [40]	Belgium	15	CT	No
Sharma et al. [30]	27 different countries	261	CT, information, self-assessment	No
Trestian et al. [35]	Ireland	1001	CT (COVID Tracker)	Westin's privacy segmentation index (PSI)
Jamieson et al. [20]	USA	290	CT	UTAUT
Aji et al. [3]	Malaysia	505	CT (MySejahtera)	No

already peer-reviewed. For this reason, we only include peer-reviewed work in this chapter but would like to point out that many (in our sample of papers 9) of those cite such publications. Also, we want to point out to the reader that the studies were not conducted in the same setting: Some asked about a hypothetical app, others studied an existing app, and others an app that was about to be published. Additionally, the design of the presented apps differed, making the comparison additionally hard.