Privacy and Security in Augmentation Technologies

Mohamed Khamis and Florian Alt

Abstract In this chapter we present a privacy and security framework for designers of technologies that augment humans' cognitive and perceptive capabilities. The framework consists of several groups of questions, meant to guide designers during the different stages of the design process. The objective of our work is to support the need for considering implications of novel technologies with regard to privacy and security early in the design process rather than post-hoc. The framework is based on a thorough review of the technologies presented earlier on in this book as well as of prior research in the field of technology augmentation. From this review we derived several themes that are not only valuable pointers for future work but also serve as a basis for the subsequent framework. We point out the need to focus on the following aspects: data handling, awareness, user consent, and the design of the user interface.

1 Introduction

Novel technologies enter the market at a rapidly accelerating pace. A fundamental challenge is that designers of such technology usually need to think about the benefits it provides to the users first with the ultimate goal of developing a product that generates revenue. At the same time, this usually leads to that security and privacy are only a secondary design goal and are not considered but at a later stage of the design and development process, if at all.

Integrating security and privacy measures post-hoc is difficult for many reasons. This can even impact the success of augmentation technologies. For example, privacy concerns were among the main reasons Google Glass is no longer available as a

Mohamed Khamis
University of Glasgow, United Kingdom, e-mail: mohamed.khamis@glasgow.ac.uk

Florian Alt
Bundeswehr University, Munich, Germany e-mail: florian.alt@unibw.de

consumer product. There are even several documented cases where users of Google Glass were assaulted due to privacy concerns of bystanders [1]. Another example is head-mounted displays (HMDs). By being able to track users' moves, for example, it is possible to find out if users are slow to react to information displayed on HMDs. This is sensitive information a user might not want in the hand of third parties. At the same time, Mixed Reality companies such as Oculus deliberately decided not to focus on security measures such as data encryption [2], because in their view this would lead to unnecessary complexity, which would negatively influence the experience in the first place. An attack that causes the leakage of the user's sensitive data might have significant privacy implications. For example, the fact that data on OpenSim[1] is unencrypted enables attackers to steal or manipulate content and to impersonate users [3].

The aforementioned examples demonstrate the need to more carefully think about possible security and privacy implications as we design novel technologies early on in the design process. The objective of this chapter is to provide an overview of possible implications on security and privacy as we are designing technologies with the goal of augmenting human perception and cognition. In particular, we critically review state-of-the-art with the ideas presented in the previous chapters in consideration, and discuss how applications might put users' security and privacy at risk. Ultimately, we derive a framework meant to help designers of such novel technologies to consider privacy and security early on in the design process. Our chapter is complemented by a discussion of future research directions, including both challenges and opportunities, resulting from the aforementioned technologies becoming available.

2 Background

The following section introduces basic terms and concepts from privacy and security. In particular, we focus on the properties of secure systems, attack and threat modelling, and briefly motivate the need for user-centered security.

2.1 Properties of Secure Systems

The objective of security mechanisms is to preserve three properties of a system against misuse and interference: confidentiality, integrity, and availability [4].

Regarding *confidentiality*, two concepts can be distinguished: (1) *data confidentiality* refers to the need of preventing data to become available to unauthorized entities; (2) *privacy* means an individual's data should be used, disclosed and exchanged according to a set of rules that the user has consented to. An example of a

[1] OpenSim is an open-source platform for hosting virtual worlds. It was used for many years by Second Life and forms the basis of the US Military MOSES project.

confidentiality breach would be private user data (e.g., blood pressure readings) leaking from a server or an on-body sensor and being made accessible to an unauthorized entity, such as an insurance company.

Maintaining a system's *integrity* refers to ensuring that systems are consistently performing their function without intended or unintended unauthorized manipulation. A further aspect concerns the integrity of data, meaning that data should be altered only by authorized entities. An example of an integrity breach is when a malicious program (e.g., malware) manipulates the readings from a sensor or user data stored in a data base.

Finally, *availability* refers to ensuring that systems and data are promptly usable. Denial of service (DoS) attacks are among the most common ways to disrupt the availability of systems. Other examples include preventing users from making emergency calls.

2.2 Attack Modeling

The aforementioned properties of security systems may be compromised as results of an attack. For example, attackers try to exploit vulnerabilities to disclose data and, hence, breach confidentiality, alter data to breach integrity, or deny services or access to data to breach availability. In order to protect these properties – confidentiality, integrity and availability – a common approach is to model potential attacks. Subsequently, potential threats can be better understood and, ultimately, countermeasures be designed.

Firstly, it is necessary to understand the *causes of an attack*. Attacks can be a result of software and hardware vulnerabilities (e.g., backdoor, or network vulnerability); also, attackers often address humans as the weak link in secure systems (e.g., via social engineering attacks, such as phishing or deep fakes); or attackers often exploit unintended characteristics of a system, so-called side channels. In the latter form of attack, adversaries exploit information gained from a system's implementation rather than from a weakness of a system. Examples of such attacks include exploiting power consumption or network traffic to infer which type of data is being transmitted, or exploiting smudge [5] and heat traces [6] on interfaces to infer user input.

Secondly, two different *types of attacks* can be distinguished: active and passive attacks. During active attacks, adversaries try to actively alter system resources. In contrast, the goal of passive attacks is to learn about vulnerabilities of a system or to get access to confidential information without affecting system resources. As a result, the latter type of attacks is much more difficult to detect. An attempt to reset someone's account credentials is considered an active attack. Observing a user as they authenticate in public (i.e., shoulder surfing) is considered a passive attack.

Thirdly, designers need to be aware of *attackers' motivations*. This is not only useful to determine the likeliness of an attack but also the potential effort an attacker is willing to spend. Some attackers aim for profit (e.g., credit card fraud, ransomware, stealing and/or selling resources). Others aim at disrupting certain processes (e.g.,

hacking a political party's website), often with the intention to make a statement. Some individuals enjoy the challenge and perform attacks for "fun". Finally, so called white-hat hackers perform attacks to test systems (e.g., network security analyst performing penetration tests).

The fourth aspect to consider is the *resources and capabilities of the attacker*. On one hand, attacks can be performed by so-called script kiddies. On the other hand, adversaries may be well-organized cyber-terrorists or nation-sponsored hackers, having significant resources in terms of money and computing power.

2.3 Threat Modeling

An understanding of the attackers, their capabilities, their motivations, and their resources helps identifying the threats that a system can be exposed to. Subsequently, designers or system providers can decide, how to protect against possible attacks. It is important to realize that protecting against all possible attacks is usually unfeasible, since it requires considerable effort and resources. Hence, a more promising strategy is to prioritize which threats to protect against. For example, if an attack can be performed by adversaries with little to no technical background (e.g., most user-centered attacks, such as guessing or shoulder-surfing credentials), then defending against these potentially common attacks should be prioritized. Another common approach is to understand the potential consequences of an attack. For example, protecting against attacks that have mild consequences (e.g., embarrassment) can sometimes be prioritized over attacks that could impact health or lead to bankruptcy.

The following questions may guide designers and providers when designing appropriate security measures: Who are the most likely attackers? What are the capabilities of attackers? What are potential consequences of attacks? What is the weakest link in the system that attackers will likely exploit?

2.4 Human-Centered Security

The final part of this section is dedicated to human-centered security. After a brief motivation we summarize approaches to achieve human-centered security[2].

Researchers have long discussed the role of humans in secure systems. It has been argued that humans are often the weakest link in secure systems. At the same time, this is often a result of systems not being designed for the way in which people interact with computing systems. Take, for example, authentication mechanisms. To make people create strong passwords, policies today require users to chose passwords consisting of eight or more characters, containing uppercase letters, lowercase letters, symbols, and digits. At the same time, users are required to create such passwords for an average of 100 different accounts. The obvious consequence is that humans

[2] "Usable Privacy and Security" is another term that describes this field.

will reuse passwords or write them down. This, however, is less a result of the user's inability to remember such passwords, but of the poor design of this security mechanism for its use case. Text-based passwords emerged in the 1960s where people had on average access to one computer and authenticated a few times per day. Today, however, we interact fundamentally different with computers. Mobile, networked devices allow us to access sensitive information, everytime and everywhere. For example, we access the smart phone about 200 times per day [7], leading both to a significant overhead in authentication time and to the need to use more passwords that can be remembered. As a result, users will optimize for convenience especially because security is never their primary task but something that gets in their way as they are trying to do other things.

This has been recognized by the research community. The National Cyber Security Center in the UK has a team dedicated to 'people-centered Security', whose lead argues that 'security must work for people. If security doesn't work for people, it doesn't work" [8]. Similarly, researchers have acknowledged that 'users are not the enemy' [9] and identified the need to design secure systems that are usable by the average human.

In response to this, researchers came up with approaches to design such systems. Whitten and Tygar argue that "security software is usable if the people who are expected to use it are reliably made aware of the security tasks they need to perform; are able to figure out how to successfully perform those tasks; don't make dangerous errors; and are sufficiently comfortable with the interface to continue using it" [10]. Ka-Ping Yee suggests guidelines for usable and secure authorization [11]. For example, he suggests that the most straightforward way for users to perform tasks should be matched with the most secure option (e.g., default options are the most secure ones), and that users must maintain accurate awareness of their own authority to access resources, as well as being aware of and able to reduce others' authority to access own resources. Many of these recommendations correspond to Jakob Nielsen's usability heuristics for user interface design [12]. This similarity implies that there are many usability concepts that, when applied, would result in improved use of the security system, which in turn results in higher security.

Therefore, we conclude that the main aim of human-centered security, is to make privacy and security an integrated, natural, unburdened part of human-computer interaction. According to Cranor and her colleagues, the following are the core challenges of designing usable and human-centered security systems:

Security concepts are complicated for the average user. For the average user, understanding concepts such as encryption, HTTPS, etc. is hard and hence might result in poor security behavior [10].

Security is a secondary task. Humans never use a system with the aim to authenticate or download security updates. Hence, their motivation for secure behavior can be generally considered rather low [13].

Human capabilities are limited. For example, the average user has about 90 web accounts but not the cognitive abilities to memorize 90 unique passwords that abide to commonly used password policies [9].

Misaligned priorities. An organization's interest in protecting its data might lead them to requiring employees to use complicated security mechanisms. At the same time employee's want to get their work done as fast as possible. As a result, they may find workarounds to make security less cumbersome [9].

Habituation. People are used to dismissing warnings (e.g., clicking "next") which increases the likelihood the user will perform an insecure action [14].

2.5 Security and Privacy Implications of Human Augmentation

The proliferation of novel technologies always comes with new implications. This applies to any form of innovation – from the industrial revolution to AI-generated videos and content. In the last two decades, security and privacy implications were among the most discussed concerns of advances in computation as technology is becoming ubiquitous at an ever-increasing pace. This led to researchers investigating frameworks for designing privacy-aware ubiquitous systems [15]. Researchers have recently started investigating the privacy implications of particular augmentation technologies, such as eye tracking [16, 17], thermal imaging [6], and life logging [18, 19]. Similarly, recent work investigated the possibly malicious use of human augmentation technologies and proposed counter measures to address them. For example, researchers explored how to protect against thermal attacks that aim for retrieving passwords from heat traces left on touch surfaces [6], how to hide sensitive content that could be recorded by life loggers [20], and how to prevent the identification of a user through their eye tracking data [17].

While efforts to tackle privacy and security issues of individual technologies are a step in the right direction, we argue that there is a need for a high-level framework that would allow researchers and practitioners to address privacy and security issues from the beginning of the user-centered design process of human augmentation technologies. In other words, we need to address the issues proactively before they rise, rather than patching them up after the release of products. Without doing that, we risk that a) augmentation technologies are never picked up due to security or privacy concerns, and b) augmentation technologies are used maliciously.

3 Methodology

After introducing human-centered security and highlighting the importance of understanding the implications on user privacy, security and safety, now we explain our methodology in identifying said implications in the context of augmentation technologies.

To help designers of technology that augments our perception and cognition mitigate potential privacy and security issues, we set out by obtaining an *understanding of potential privacy and security concerns*. Therefore, we carefully reviewed both

the technology presented in the previous chapters alongside prior work in this field. Selected projects considered for our analysis are listed in Table 1. For each project we (1) extract the core concept, (2) identify which data is collected and the consequences of storing and sharing this data, (3) discuss which information could be derived if somebody had access to the collected data, (4) discuss new attack channels that are now feasible due to the use of this concept, (5) discuss how the technology can be used maliciously, and (6) derive possible implications on users' security and privacy with a focus on data manipulation and physical harm. This part primarily evolves around implications on the user of a technology, implications from the surroundings of the user of a technology (e.g., new attack vectors), implications of third parties on the user (e.g., implementation of privacy and data confidentiality by a company), and implications on those people around the user.

In a second step, these implications were used as a basis to *identify themes that require further research*. To do this, we performed a data walkthrough and discussed what it would take to address the concerns revealed in the first step.

Third, we derived a *framework* that provides designers a structured approach of considering potential security and privacy implications during the design process. In particular, the framework evolves around questions regarding data collection, data storage, user control over the data, and the user interface design.

Although unanticipated, our review identified not only potential security and privacy concerns (Section 4), but also opportunities for novel security mechanisms (Section 5). More specifically, we discuss how perception and cognition enhancing technologies can be leveraged to build novel authentication mechanisms that are both secure and usable. With this discussion we hope to provide fertile ground for future research between people working in the field of technology augmentation and usable security and privacy.

Chapter	**Project**	**Description**
Chapter 7	RainSense	A system supporting users to develop a sense of precipitation through thermal output; the system receives weather information via Bluetooth and subsequently provides thermal feedback by means of a Peltier element.
Chapter 7	Solo	A system enabling users to focus on sounds in a given direction by means of pointing at a sound source. It filters out surrounding noise, enabling users to perceive individual sound signals.
Chapter 7	Clairbuoyant	A system enhancing the sense of direction of open water swimmers. It uses augmented swimming goggles for providing visual directional cues.
Chapter 6	Insertables for non-medical purposes	A survey with 115 participants to understand what devices they are putting in their body, what they use these devices for, their motivations for doing so, and how they identify themselves.

Table 1 Selected projects related to human augmentation.

4 Privacy, Security, and Safety Implications of Technology-Augmented Perception and Cognition

The past chapters have provided forward-looking concepts and systems that demonstrate a great potential of augmenting the perception and cognition of humans. Without doubt, these developments will bring tangible benefits to users. However, similar to any technological innovation, there are downsides that, if not accounted for, could have significant negative impacts on humans' lives.

In the following, we will highlight a number of issues related to privacy, security and safety that researchers and practitioners should account for when designing systems to augment human perception and cognition.

4.1 Understanding Consequences of Data Sharing

The abundance of sensors users will be carrying in the future allows myriads of personal information to be collected, almost anytime and anywhere. This information can be used for many applications, providing benefits for the user. At the same time, the collected data are usually transferred to servers for analysis. This potentially affects users' privacy: non-trusted companies that manage the collected data could potentially sell the data to third parties.

For example, Chapter 6 demonstrated how biometric data collected through insertables can be transmitted to external devices or even remote servers for analysis. Interviews reveal that early adopters of insertables hope to see them used to read their internal blood oxygen level and interpret data from their eyes. At the same time, there could be severe privacy implications in case such data is shared with third parties. For example, insurance companies could use the data about a user's biometrics to decide whether or not to insure them.

Another example is the work on task resumption in Chapter 4 where the idea is to help users recover from being interrupted during tasks, as it has been shown that such interruptions strongly influence productivity. At the same time, information on interruptions is sensitive: for example, such interruptions may be caused by a colleague to have a private chat. As a result, knowledge on the type of interruption might influence on personal evaluation of even payment if known to employers. For example, if an employer learns that an employee is being interrupted a lot by matters that are not related to their work (e.g., personal messaging notifications), they might use this information against them.

The particular challenge we see here is that it is often difficult to infer which consequences it might have if data is available to third parties . Designers and developers need to exercise great care and develop an understanding of what it means to collect, store, and process certain types of data.

4.2 Loss of Data Control

The increasing number, ubiquity and capability of sensors makes it significantly harder for the user to control what information is being collected and shared with other parties. This could make informed consent of data collection and sharing almost impossible. The reason is that the amount of data that is collected is so massive that humans are cognitively incapable of keeping up with it and are left alone understanding, or even realizing, that data is being collected. In the following we discuss different forms of loss of control over data.

For example, Chapter 6 describes a trend to insert RFID and NFC tags into people's bodies. These chips carry information that could potentially be sensitive (e.g., credentials to access a building). At the same time, users do not have full control on when this data can be accessed. If an attacker approaches one of those users with an NFC reader, they can extract the information without the user's consent. Similarly, an attacker could read off information on a user's inserted RFID if the user is unconscious, sleeping, or inattentive.

Insertables (Chapter 6) introduce a further threat: communicating biometric data to surrounding devices or to remote servers may introduce a new attack channel: if data is sent in an unencrypted way, man-in-the-middle attacks can result in attackers intercepting sensitive information about the user as they are transferred wirelessly. This can happen without the user's control.

In contrast to an attacker actively extracting data, the user could disclose private data through unexpected means when using augmenting technologies. For example, in Chapter 7, Poguntke and Kiss introduced a technology amplifying a users' sense of hearing where users could perform a mid-air gesture to specify a direction for which they want to increase their sense of hearing, for example, a group of people in close vicinity. The design of this gesture in itself might be problematic, since it reveals a user's interest.

Another challenge arises in situations where information can be inferred from data. For example, the data collected to support task resumption strategies discussed in Chapter 5 does not only reveal which content the user accessed and interacted with, but it could also reveal when the user worked on what. While a user might have consciously consented to sharing data on the content they accessed, they might not be aware of additional information it allows to be derived nor what it could be used for.

In summary, it is unrealistic to expect that users are aware of the consequences as novel technologies take away control over their data. Designers need to carefully think how to enable control over users' data. Furthermore, as users are asked to provide consent, it needs to be clearly communicated to the user what potential consequences are.

4.3 Novel Attack Channels

Many technologies we reviewed add novel output channels and, hence, enable novel attacks via these channels. A particular challenge stems from the fact that many of these channels might initially not have been designed to convey sensitive data and designers thus need to think about novel ways to protect these channels.

For example, Chapter 4 presented how visual, auditory, haptic, and tangible output modalities can be leveraged to provide cues for task resumption. This includes, for example, large screens able to convey personal information. Traditionally, large screens were not meant to convey this type of information. In particular, if used in shared work spaces or public, there is a need to develop means for protecting users' privacy.

Similarly, the technologies presented in Chapter 7 offer novel means to present information. Examples of output channels that have recently became mainstream include head-mounted displays (HMDs) and smart watches. The visual feedback provided through their screens makes them subject to shoulder surfing [21]. Note, that not only visual output is subject to eavesdropping. For example, Solo (Chapter 7) allows for listening to conversations at a distance and these conversations could well be eavesdropped by bystanders of a person wearing the Solo system.

From this we learn that designers need to carefully consider the output channel used for technologies augmenting human senses. Of particular interest here is the context of use. If such a technology is used in a public or shared space, means need to be provided to secure the technology against attacks from bystanders.

4.4 The Cyborg Stalker

An interesting aspect is that technologies augmenting users' perception and cognition could be exploited by the augmented users themselves in malicious ways. This means that designers should not only think of how to protect their users' privacy, but also try to understand how users can misuse their innovations.

For example, an attacker could use the Solo system (Chapter 7) to eavesdrop on conversations of people in public. Lifelogging cameras are known to cause privacy issues when bystanders are in the field of view [22, 23] or screens that are showing sensitive content (e.g., desktop screen showing emails) [20]. Some HMDs are now augmented with thermal cameras [24], which could be used to infer the emotions of surrounding bystanders [25], or even to perform thermal attacks to extract users' inputs (e.g., passwords) on screens and keyboards [6].

This points at interesting directions for future work. To mitigate malicious use, technologies could consider their current context and accordingly enable or disable certain features. For example, a lifelogging camera can automatically blur people's faces and black out recorded screens.

4.5 Impact of Data Manipulation

As we become more dependent on the technologies we use, their impact on our lives increases. Consequently, any manipulation of system variables, configurations, or any other aspects impacting system integrity can likewise influence us.

One particular challenge we foresee as augmentation technologies are becoming ubiquitous is fake data. For example, RainSense (see Chapter 7) is meant to make users aware of weather conditions and potentially notify them about that it might rain during the day. Access to the data on which the prediction for rain is based, adversaries could misuse the system to trigger the prediction of rain and ultimately make the user purchase accessories (such as an umbrella) or clothes even though they do not need them.

This is just one example highlighting the need to carefully think about how data can be collected, transmitted and be stored in a secure manner, such that adversaries cannot simply temper with data.

4.6 Physical Harm

We identified a number of ways in which augmentation technologies could have severe consequences and physically harm the user, if not carefully designed. In the following we sketch cases in which this can happen directly or indirectly.

Examples, where the system could indirectly lead to physical harm is the wrong display of information. For example, Poguntke and Kiss proposed using visual feedback to convey a swimmer's orientation to herself. Any potential tampering with the shown orientation could result in misleading the user. This might not only have social and economic implications (e.g., losing a competition), but in extreme cases it could even have safety implications (e.g., leading the swimmer to a dangerous area in open water).

Another example of indirect harm could be the increased value of human body parts as augmentation technology, for example, in the form of insertables, are added. As humans start to use expensive insertables, a new type of crime could arise from illegally harvesting insertables from their users. Or, a human with an insertable used for authentication (e.g., unlock a security door) could be physically harmed to gain access to their authentication token. Note that similar attacks occur in the context of biometric features (finger, iris, etc.).

This amplifies the need to ensure that these technological advancements are safe and do not subject users to hazardous situations. The previous section highlighted how manipulating the data that is perceived by the user could have fatal consequences. Indeed, an HMD that is communicating with other servers could be prone to man-in-the-middle attacks, in which an attacker can overlay a virtual bridge over a real cliff, hence subjecting the user's life to danger [26]. Designers should keep similar issues in mind when designing such technologies.

4.7 Summary

To summarize, while technology augmentations can bring many benefits to users, there is a number of privacy, security, and safety issues that should be considered and further investigated: 1) the augmentation itself could become the very reason behind the leakage of private data; 2) obtaining informed consent before data is exchanged or used is becoming increasingly challenging; 3) adversaries will come up with novel attacks for each new, exposed channel; 4) designers should keep in mind that the augmentations can be potentially exploited by its users (e.g., to spy on others); and 5) designers should be aware of the amplified impact of data manipulation which could lead to physical harm.

5 Directions for Future Research

Our review of related work revealed a need for further research in different areas. Particular areas we identified are supporting informed consent, understanding and mitigating the effects of data manipulation, and exploring the influence of augmentation on malicious use.

Of particular interest in this context is the European Union's General Data Protection Regulation (GDPR)[3]. One of the core principles is the need for a lawful basis for processing data. In particular, explicit informed consent from users is required in case data is made publicly available. One specific challenge here is also that it needs to be possible to revoke this consent at any time. Furthermore, there is a need to inform users about the extent of data collection and provide them an overview of which data is stored and how it is being processed upon request. Finally, providers need to ensure that data can be erased within 30 days. While not being specific to augmenting technologies, these requirements are nevertheless highly relevant, since they strongly influence the way in which technologies, underlying architectures, and user interfaces need to be built.

5.1 Informed Consent

Achieving informed consent is becoming increasingly challenging in the ubiquitous computing age. Augmentation technologies raise similar concerns. Some of the augmentations, such as Insertables (Chapter 6), require collecting data. There are two main aspects where it is becoming increasingly challenging to ensure the user gives full, informed consent in the context of augmentation technologies:

[3] https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules/eu-data-protection-rules_en

5.1.1 Informed Consent in Data Usage

More and more data is being collected about the user. At the same time, the amount of information that can be inferred from this data is not only increasing but also becoming more complex and less obvious for the regular user. For example, it is not immediately clear to users that their eye movements can reveal their mental states [27] or even their political temperament [28].

Currently, consent for collecting, using and sharing the user's data is obtained by asking users to read and accept long and complicated terms and conditions that lay out the privacy policies. This presents a research opportunity: we should move away from the use of text-based privacy policies that suffer from low usability to novel approaches that are more understandable and support informed consent. For example, Kelly et al. [29] proposed designing a "Nutrition label" for privacy. In their work, they suggest standardizing privacy related facts by presenting them in a concise form similar to how food products are labeled with nutrition facts that summarize the amount of calories, vitamins, fats, etc. in a product. In their proposal, they suggest emphasizing "what" information is collected, "how" it can be used, and "who" may the user's information be shared with. Another example is how many tech companies are starting to use alternatives to text-based privacy policies. Apple, Google, and Facebook use videos to illustrate which data is collected, how it is used, and with whom it is shared. The privacy policy pages are structured in a more intuitive way compared to the traditional text-based privacy policies.

While these improvements are all in the right direction, augmentation technologies require us to design ways to communicate an unprecedented amount of information to ensure informed consent. Therefore, there is a need to design methods to effectively and efficiently communicate privacy related information when using augmentation technologies. This presents another research opportunity: a promising starting point is to explore different mediums (e.g., videos, virtual reality, narrated stories, games) to communicate privacy policies of augmentation technologies to the users. As for example in Solo, augmentation technologies may also create an increasing number of situations, in which data of people who are not the main users of a technology are collected willingly or unwillingly. This creates another opportunity for future research, i.e. how can consent be obtained from non-users?

5.1.2 Informed Consent in User Actions

There is a growing trend to build systems that work in the background without the user's explicit input. Examples include systems that respond to the user's behavior (e.g., gaze behavior, body movements, etc.) or systems that respond to implanted RFIDs (e.g., see Chapter 6), the user's face, or fingerprint. While these technologies make interactive systems seamlessly integrated around us and reduce the cognitive load required to control them, they present new challenges.

As interaction becomes more passive rather than active, it becomes more likely that users perform unintended, or even unauthorized, inputs. This has negative implications on security. For example, contact-free bank cards have become more common recently as they are more usable. They are faster because they do not require the user to enter a PIN as long as the payment is below a certain threshold. This usability improvement also means that payments can be done without the user's permission if the card is stolen. Similarly, there have been reported cases where a smartphone's fingerprint sensor unlocks a user's phone upon sensing their fingerprint, even if the user did that unconsciously (e.g., while asleep [30]) or forcibly (e.g., police unlocking a dead person's phone [31]). A similar issue that is heavily studied is the Midas touch problem in gaze interfaces [32, 33, 34], in which systems interpret regular eye movements as gaze input and result in unintended input. While fingerprint unlock without consent and Midas touch are two distinct problems, they are both forms of unintentional input, which suggests that some of the solutions proposed to address one of them might be promising for the other.

This presents a research opportunity: Future work in augmentation technologies should strive balance between high usability of systems based on implicit interaction and a high level of security. One approach is to detect intention during implicit interactions. Similar to how gaze interfaces require users to perform special gaze movements [35] or to dwell at targets before selecting them [33], implicit systems such as fingerprint sensors and readable inserables should similarly adopt measures to confirm the user's intention. For example, fingerprint sensors can be augmented with other sensors to estimate if the user is intending to authenticate (e.g, by exploiting gaze direction). Systems that read input from insertables should rely on more than just the presence of the insertable in close proximity and instead involve a user action to make it less likely that user's input is interpreted unintentionally (e.g., contracting certain body muscles).

Therefore, important questions that will drive future research are: How can we make sure that the user's implicit action is intended? How can we ensure consent without greatly compromising usability and responsiveness of systems?

5.2 When Users are Evil – Augmentations can be Used Maliciously

In the past, using a technology maliciously required resources (e.g., expensive equipment) and special skills (e.g., programming knowledge). As hardware becomes cheaper, people with limited technical knowledge can use technologies maliciously to spy or cause harm to others. A recent example is the work of Abdelrahman et al. [6], which showed that thermal cameras can be used to infer passwords entered on mobile devices. Although they used an algorithm to map the thermal images to PIN and pattern entries, many of the thermal images reveal the user's password through visual inspection by a non-expert attacker. Similarly, the ubiquity of mobile devices means that users are accessing sensitive information in public areas, which in turn means that attackers can potentially gain access to sensitive information by

merely shoulder-surfing others [21]. This is another attack that does not require any technical expertise.

Human augmentation raises similar concerns. Augmented users will carry cameras and sensors that may allow them to spy on others (e.g., listening to distant conversations as highlighted in Chapter 7), infer sensitive information about them without their consent (e.g., revealing their PINs using thermal imaging [6]), or even harm others (e.g., robotic limbs [36] could unintentionally hit bystanders).

An interesting direction for future research is investigating applications of augmentation technologies that are potentially malicious with the goal of identifying privacy and security issues before a technology becomes ubiquitous.

6 Towards a Security and Privacy Framework for Technology Augmentation

We complement our exploration of security and privacy implications in the context of human augmentation with presenting a framework. This framework is meant to make the designers of such technologies think of suitable approaches that mitigate issues related to privacy and security.

We envision the framework to be used in the following way. Questions are grouped based on different phases of the design process. In particular, questions relate to how data is being handled (this is relevant as designers create a system architecture), user consent, and control over data (these aspects mainly relate to the user interface). Designers can now answer each question in the context of their design and reflect on whether or not it has been addressed meaningfully.

In the following, we present the questions, alongside a brief explanation for why we think this is relevant. Table 2 summarizes the questions of the framework and demonstrates how it could be applied by designers and developers who aim to make a commercial product out of one of the concepts presented in this book. Note, that in the context of the framework we not only refer to users but more generally to *stakeholders*, since our exploration revealed that also other groups of people, such as bystanders, are affected by augmentation technologies.

6.1 Data Handling

What data is collected? The purpose of this question is to make designer reflect on the need to collect certain types of data. Whereas some data might be essential for the functionality of the technology, others might serve secondary purposes, such as identifying potential usability issues or help improve a technology or service.

How is data collected? One important question is how data is collected. While some services might collect data from accessing sensors, other data might be

Part I: Data Handling	**Example: Solo**
What data is collected?	audio data, head pose / visual attention
How is data collected?	microphone / eye tracker or camera to detect head pose
How is data being transmitted?	Over a (secure/insecure) network
Where is data being processed?	In the cloud / on the device
Where is data being stored?	In the cloud / on the device
Part II: Awareness & Consent	
How to communicate to stakeholders that data is being collected?	Visual / audio feedback can be provided to the user and other stakeholders whose privacy and security are impacted.
How to communicate to stakeholders that data is being shared?	Visual / audio feedback can be used to communicating this information to the user and stakeholders.
Do stakeholders understand what happens to their data?	Investigate how to make the functionality of the system clear to stakeholders.
How to obtain informed consent from stakeholders?	While obtaining informed consent from the user can be straight forward, obtaining it from stakeholders, such as bystanders, is more difficult. One way is to make stakeholders aware by, for example, broadcasting notifications to nearby smartphones and smartwatches.
Part III: Control Over The Data	
Which types of control do stakeholders have over their data?	Investigate how to allow both the augmented user and stakeholders to access and be able to request deletion of the data that was recorded.
How can users request their data being deleted?	The augmented user should be provided with a mechanism (e.g., an app) that allows deletion of collected data. One way to ensure the same for stakeholders, is to provide each of them with a link from which they can review, and if desired, delete their data that was collected by the device.

Table 2 Framework for privacy and security assessment of augmentation technologies. The right column describes how the framework could be applied by designers and developers who are aiming to make a commercial product out of the Solo prototype.

collected from third party services. This is relevant, since it ultimately means that others might have access to sensitive data as well.

How is data being transmitted? With modeling and machine learning becoming increasingly important and powerful tools, the processing power of devices is often not sufficient. As a result there is a need to transmit data to a server that then executes performance-heavy tasks. At the same time, this poses the risk of data leaking during transmission. Hence, designers need to ensure that data is transferred in a secure manner and neither be intercepted nor manipulated (e.g., by encrypting them).

Where is data being processed? Closely related to the question above, one question is where data is being processed. From a user perspective it is clearly desireable to process data on their personal device. As this is not possible (e.g., due to limitations in computing power), designers need to think careful where data is being processed and by whom.

Where is data stored? Data storage is another important aspect. In particular in cases where lots of data is being collected, storage on the device itself might not be possible (think about fine-grained behavioral data or image/video data). In this case, a designer needs to think carefully, where data is being stored and who can access it.

6.2 Awareness & Consent

In response to the Lederer's pitfalls in designing for privacy, the following part of the framework postulates that designers carefully think about how users can be made aware of what happens to their data and how consent can be obtained.

How to communicate to stakeholders that data is being collected? Most fundamentally, stakeholders should be made aware at any point in time, which data is collected about them.

How to communicate to stakeholders that data is being shared? In particular in cases where data is shared with third parties, e.g., for processing, there is a need to communicate this to stakeholders.

Do stakeholders understand what happens to their data? One fundamental aspect of protecting stakeholders' privacy is comprehension. In particular, designer need to make sure that stakeholders are fully aware of what happens to their data. Important questions to ask here are whether the stakeholder understands what kind of information is shared, with whom it is shared, through which medium it is conveyed, where and how it is processed and where it is being stored.

How to obtain informed consent from stakeholders? According to the GDPR, informed consent must be obtained from people once data is made available to third parties. Here, designers must make sure to provide suitable means for (a) obtaining informed consent and (b) enabling stakeholders to revoke this consent at any time.

6.3 Control over data

Finally, designers need to take into consideration how users could be provided control over their data. Means of control need to be realized as part of the user interface of a technology.

Which types of control do stakeholders have over their data? Beyond requesting mere deletion of their data, designers might want to think about providing stakeholders an opportunity to only delete parts of the data. This might be useful specifically for data that are not essential for the functionality of a technology. Enabling stakeholders to do so might also be beneficial for the providers of technology, since rather than completely opting out of a service or technology, stakeholders might only disable the collection of or delete parts of the data stored about them.

How can stakeholders request their data being deleted? A core principle in GDPR is the opportunity to have their data be deleted within 30 days. Designers need to think about a way how stakeholders can make this request (in particular, if they are not the users) and how the system architecture can be designed in a way to do this with minimal effort.

7 Discussion

In the following sections we reflect on the framework, in particular discussing aspects that require further investigation.

7.1 Required Expertise

Our framework is generally targeted at the *designers* of systems augmenting humans. At the same time, employing the framework is likely to require different types of expertise, in particular such that is not available in traditional design processes. For example, to properly design secure transmission or storage of data, experts in network or data security may be required. This need has been recognized by the community [37] and is also backed by prior work, showing that software developer often either do not have the required expertise for building secure systems or do not see the need for it [38].

Another example that becomes particularly apparent with the GDPR is the need for experts, overseeing that data is handled in a privacy-preserving way, e.g., data protection officers or even lawyers.

7.2 Interplay With Commercial Interest

Another aspect that would be interesting to explore is how our framework interplays with commercial interests. This is particularly important in times, where data is an important currency. Many business models today are based on access and control over user data. As an ever-increasing amount of sensitive data is being used, companies may need to rethink their business models in such a way that these comply with the suggestions put forth by this framework. For example, companies may want to return to traditional pay-per-use or subscription-based business models.

7.3 Need for End User Involvement

Many of the aspects identified in our framework can be addressed by experts, e.g., implementing encryption to ensure secure data transmission. At the same time, there are several aspects that may require the design of novel approaches and, subsequently, the involvement of end users to test these approaches. This is particularly true for aspects that concern the user interface of human augmentation technologies. For example, there is no standard way of communicating to people that data about them is being collected. This strongly depends on the technology and its output modalities. If the device has a display, it could be used to design an appropriate visualization to communicate to users that data is being collected. In other cases, this need might even require adding additional output technologies, such as an LED, that were previously not part of the product. In order to find out how to best design such novel approaches, designers may need to involve end users and conduct user studies to find out how to optimally design for a certain aspect.

7.4 Influence Beyond the Design Phase

Considering our framework may have consequence beyond the design process. For example, the way in which a company decides to implement ways for users to have control over their data or to request their data being deleted, may require thinking about a support infrastructure or even create the need to hire people that deal with such requests. This might ultimately influence also the business model.

8 Conclusion

In this chapter we introduced a framework for privacy and security, meant to guides the design of technologies augmenting humans' perception and cognition. The framework was derived from work presented in the chapters of the book.

We first identified the implication of novel technologies on privacy, security and safety. We found the understanding of consequences of data sharing, control of the data, novel attack channels, the opportunity to leverage such technologies for malicious purposes, data manipulation and safety to be critical aspects.

The analysis also revealed a need for more research. In particular, future work needs to obtain a better understanding of what can be learned from the stakeholders' data obtained from technologies meant to augment humans' perception and cognition; researchers need to think about how informed consent could be obtained from the stakeholders, both regarding the use of their data as well as regarding user actions; and researchers could focus on understanding how augmentations could be used maliciously and how to mitigate such cases.

Finally, our framework provides three sets of questions that guide the design of secure and privacy-preserving designs of novelx augmentation technologies. Specifically, designers need to carefully consider how they handle data (collecting, processing, storing, sharing), how stakeholders could be made aware of the implications of using the technology and how they could provide consent, and how they can be provided control over their data.

For the future we intend to evaluate the framework with researchers working on augmentation technologies. In particular, we plan to interview them on how the framework helped them in creating concepts, developing a system architecture, and designing their user interfaces.

References

1. K. Wolf, A. Schmidt, A. Bexheti, M. Langheinrich, IEEE Pervasive Computing **13**(3), 8 (2014). DOI 10.1109/MPRV.2014.53
2. A. Ng. VR systems Oculus Rift, HTC Vive may be vulnerable to hacks. https://www.cnet.com/news/hack-a-vr-system-lead-a-player-astray-yes-say-researchers/ (2018). Last accessed 29. April 2019
3. M. Korolov. Army reveals OpenSim's top security risks. https://www.hypergridbusiness.com/2016/10/army-reveals-opensims-top-security-risks/ (2016). Last accessed 29. April 2019
4. J. Andress, *The basics of information security: understanding the fundamentals of InfoSec in theory and practice* (Syngress, 2014)
5. S. Schneegass, F. Steimle, A. Bulling, F. Alt, A. Schmidt, in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (ACM, New York, NY, USA, 2014), UbiComp '14, pp. 775–786. DOI 10.1145/2632048.2636090. URL http://www.florian-alt.org/unibw/wp-content/publications/schneegass2014ubicomp.pdf. Schneegass2014ubicomp
6. Y. Abdelrahman, M. Khamis, S. Schneegass, F. Alt, in *Proceedings of the 35th Annual ACM Conference on Human Factors in Computing Systems* (ACM, New York, NY, USA, 2017), CHI '17. DOI 10.1145/3025453.3025461. URL http://dx.doi.org/10.1145/3025453.3025461
7. M. Harbach, E. Von Zezschwitz, A. Fichtner, A. De Luca, M. Smith, in *Proceedings of the Tenth USENIX Conference on Usable Privacy and Security* (USENIX Associa-

tion, Berkeley, CA, USA, 2014), SOUPS'14, pp. 213–230. URL http://dl.acm.org/citation.cfm?id=3235838.3235857
8. E. W. Cyberuk 2017: People - the strongest link. Webpage (2017). URL https://www.ncsc.gov.uk/blog-post/cyberuk-2017-people-strongest-link. Retrieved April 25, 2019
9. A. Adams, M.A. Sasse, Commun. ACM **42**(12), 40 (1999). DOI 10.1145/322796.322806. URL http://doi.acm.org/10.1145/322796.322806
10. A. Whitten, J.D. Tygar, in *Proceedings of the 8th Conference on USENIX Security Symposium - Volume 8* (USENIX Association, Berkeley, CA, USA, 1999), SSYM'99, pp. 14–14. URL http://dl.acm.org/citation.cfm?id=1251421.1251435
11. K.P. Yee, in *Security and Usability: Designing Secure Systems that People Can Use*, ed. by S. Garfinkel, L. Cranor (O'Reilly Media, Champaign, IL 61820, USA, 2005), chap. 13, pp. 253–280
12. J. Nielsen, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (ACM, New York, NY, USA, 1994), CHI '94, pp. 152–158. DOI 10.1145/191666.191729. URL http://doi.acm.org/10.1145/191666.191729
13. R. West, Commun. ACM **51**(4), 34âĂŞ40 (2008). DOI 10.1145/1330311.1330320. URL https://doi.org/10.1145/1330311.1330320
14. J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, L.F. Cranor, in *Proceedings of the 18th Conference on USENIX Security Symposium* (USENIX Association, USA, 2009), SSYM'09, pp. 399–416
15. M. Langheinrich, in *Proceedings of the 3rd International Conference on Ubiquitous Computing* (Springer-Verlag, Berlin, Heidelberg, 2001), UbiComp '01, pp. 273–291. URL http://dl.acm.org/citation.cfm?id=647987.741336
16. C. Katsini, Y. Abdrabou, G. Raptis, M. Khamis, F. Alt, in *Proceedings of the 38th Annual ACM Conference on Human Factors in Computing Systems* (ACM, New York, NY, USA, 2020), CHI '20. DOI 10.1145/3313831.3376840. URL http://dx.doi.org/10.1145/3313831.3376840
17. J. Steil, I. Hagestedt, M.X. Huang, A. Bulling, in *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (ACM, New York, NY, USA, 2019), ETRA '19, pp. 27:1–27:9. DOI 10.1145/3314111.3319915. URL http://doi.acm.org/10.1145/3314111.3319915
18. P. Elagroudy, M. Khamis, F. Mathis, D. Irmscher, A. Bulling, A. Schmidt, in *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (ACM, New York, NY, USA, 2019), CHI EA '19, pp. LBW0244:1–LBW0244:6. DOI 10.1145/3290607.3313052. URL http://doi.acm.org/10.1145/3290607.3313052
19. J. Steil, M. Koelle, W. Heuten, S. Boll, A. Bulling, in *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications* (ACM, New York, NY, USA, 2019), ETRA '19, pp. 26:1–26:10. DOI 10.1145/3314111.3319913. URL http://doi.acm.org/10.1145/3314111.3319913
20. M. Korayem, R. Templeman, D. Chen, D. Crandall, A. Kapadia, in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (ACM, New York, NY, USA, 2016), CHI '16, pp. 4309–4314. DOI 10.1145/2858036.2858417. URL http://doi.acm.org/10.1145/2858036.2858417
21. M. Eiband, M. Khamis, E. von Zezschwitz, H. Hussmann, F. Alt, in *Proceedings of the 35th Annual ACM Conference on Human Factors in Computing Systems* (ACM, New York, NY, USA, 2017), CHI '17. DOI 10.1145/3025453.3025636. URL http://dx.doi.org/10.1145/3025453.3025636
22. Y. Li, N. Vishwamitra, B.P. Knijnenburg, H. Hu, K. Caine, Proc. ACM Hum.-Comput. Interact. **1**(CSCW), 67:1 (2017). DOI 10.1145/3134702. URL http://doi.acm.org/10.1145/3134702
23. E. Thomaz, A. Parnami, J. Bidwell, I. Essa, G.D. Abowd, in *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (ACM, New York, NY, USA, 2013), UbiComp '13, pp. 739–748. DOI 10.1145/2493432.2493509. URL http://doi.acm.org/10.1145/2493432.2493509

24. Y. Abdelrahman, P. Wozniak, P. Knierim, N. Henze, A. Schmidt, in *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia* (ACM, New York, NY, USA, 2018), MUM 2018, pp. 245–252. DOI 10.1145/3282894.3282920. URL http://doi.acm.org/10.1145/3282894.3282920
25. Y. Abdelrahman, E. Velloso, T. Dingler, A. Schmidt, F. Vetere, Proc. ACM Interact. Mob. Wearable Ubiquitous Technol. **1**(3), 33:1 (2017). DOI 10.1145/3130898. URL http://doi.acm.org/10.1145/3130898
26. F. Mathis, M. Khamis, in *Proceedings of the CHI 2019 Workshop on Challenges Using Head-Mounted Displays in Shared and Social Spaces* (2019), SHMD '19
27. P. Majaranta, A. Bulling, *Eye Tracking and Eye-Based Human–Computer Interaction* (Springer London, London, 2014), pp. 39–65. DOI 10.1007/978-1-4471-6392-3_3. URL http://dx.doi.org/10.1007/978-1-4471-6392-3_3
28. M.D. Dodd, J.R. Hibbing, K.B. Smith, Attention, Perception, & Psychophysics **73**(1), 24 (2011). DOI 10.3758/s13414-010-0001-x. URL https://doi.org/10.3758/s13414-010-0001-x
29. P.G. Kelley, J. Bresee, L.F. Cranor, R.W. Reeder, in *Proceedings of the 5th Symposium on Usable Privacy and Security* (ACM, New York, NY, USA, 2009), SOUPS '09, pp. 4:1–4:12. DOI 10.1145/1572532.1572538. URL http://doi.acm.org/10.1145/1572532.1572538
30. TheGuardian. Qatar airways plane forced to land after wife discovers husband's affair midflight. Webpage (2017). URL https://www.theguardian.com/world/2017/nov/08/qatar-airways-plane-forced-to-land-after-wife-discovers-husbands-affair-midflight. Retrieved April 19, 2019
31. BBC. Police 'visit funeral home to unlock dead man's phone'. Webpage (2018). URL https://www.bbc.co.uk/news/technology-43865109. Retrieved April 19, 2019
32. H. Drewes, M. Khamis, F. Alt, in *Proceedings of the 18th International Conference on Mobile and Ubiquitous Multimedia* (ACM, New York, NY, USA, 2019), MUM '19. DOI 10.1145/3365610.3365626. URL https://doi.org/10.1145/3365610.3365626
33. R.J.K. Jacob, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (ACM, New York, NY, USA, 1990), CHI '90, pp. 11–18. DOI 10.1145/97243.97246. URL http://doi.acm.org/10.1145/97243.97246
34. M. Khamis, C. Oechsner, F. Alt, A. Bulling, in *Proceedings of the 2018 International Conference on Advanced Visual Interfaces* (ACM, New York, NY, USA, 2018), AVI '18. DOI 10.1145/3206505.3206522. URL https://doi.org/10.1145/3206505.3206522
35. M. Vidal, A. Bulling, H. Gellersen, in *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (ACM, New York, NY, USA, 2013), UbiComp '13, pp. 439–448. DOI 10.1145/2493432.2493477. URL http://doi.acm.org/10.1145/2493432.2493477
36. M. Al-Sada, T. Höglund, M. Khamis, J. Urbani, T. Nakajima, in *Proceedings of the 10th Augmented Human International Conference 2019* (ACM, New York, NY, USA, 2019), AH2019, pp. 37:1–37:9. DOI 10.1145/3311823.3311850. URL http://doi.acm.org/10.1145/3311823.3311850
37. F. Alt, E. von Zezschwitz, Journal of Interactive Media (icom) **18**(3) (2019). DOI 10.1515/icom-2019-0019. URL http://florian-alt.org/unibw/wp-content/publications/alt2019icom.pdf. Alt2019icom
38. A. Naiakshina, A. Danilova, E. Gerlitz, E. von Zezschwitz, M. Smith, in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Association for Computing Machinery, New York, NY, USA, 2019), CHI âĂŹ19. DOI 10.1145/3290605.3300370. URL https://doi.org/10.1145/3290605.3300370