EDITORS: Florian Alt, florian.alt@unibw.de
Florian Schaub, fschaub@umich.edu

DEPARTMENT: USABLE SECURITY AND PRIVACY

Privacy-Aware Eye Tracking: Challenges and Future Directions

Céline Gressel, *University of Tübingen, 72074, Tübingen, Germany*
Rebekah Overdorf, *Université de Lausanne, 1015, Lausanne, Switzerland*
Inken Hagenstedt, *Apheris AI GmbH, 10999, Berlin, Germany*
Murat Karaboga, *Fraunhofer Institute for Systems and Innovations Research, 76135, Karlsruhe, Germany*
Helmut Lurtz, *Novartis Business Services GmbH, 83607, Holzkirchen, Germany*
Michael Raschke, *Blickshift GmbH, 70563, Stuttgart, Germany*
Andreas Bulling, *University of Stuttgart, 70569, Stuttgart, Germany*

What do you have to keep in mind when developing or using eye-tracking technologies regarding privacy? In this article we discuss the main ethical, technical, and legal categories of privacy, which is much more than just data protection. We additionally provide recommendations about how such technologies might mitigate privacy risks and in which cases the risks are higher than the benefits of the technology.

In addition to gaze direction, eye tracking gives access to other metrics, such as pupil dilation and micromovements, that provide rich information about a person's perceptual and cognitive processes.[1] While the availability of such information promises a range of exciting new applications, e.g., in health diagnostics, it also raises a number of ethical, legal, and privacy issues as eye tracking is increasingly deployed and used. Particularly in the areas of health and education, eye tracking will further incentivize companies and governments to collect large amounts of highly sensitive, personal data of both system users and nonusers. Large-scale collection of gaze data poses significant privacy risks, particularly if these data are centralized and analyzed automatically using machine learning (ML) and data mining methods. Despite the potential for significant harm, the eye-tracking community has only recently begun to consider such threats and few technical solutions for privacy-aware eye tracking exist (e.g., Steil et al.[7] and Steil et al.'s work[8]).

We argue that it is not sufficient to focus only on the technical challenges of future eye-tracking systems—research has to also consider the ethical, social, and legal aspects related to the privacy of eye-tracking technologies. Taking an interdiscplinary perspective, we identify key factors that have to be considered to strengthen privacy and illustrate them on two sample use cases that are currently still hypothetical but that we believe to have a high chance of being implemented in only a few years.

PERSPECTIVES ON PRIVACY

Our starting point is that privacy is more than protecting data—a broader perspective is needed to fully describe the privacy of a system or technology. Even when focusing only on data privacy, data do not exist on their own; they always have to be generated. While data are generated directly or indirectly by technology, they are always developed and applied by people. Hence, data represent the characteristics of these people and others. Consequently, privacy protection is necessary not

1536-1268 © 2023 IEEE
Digital Object Identifier 10.1109/MPRV.2022.3228660
Date of current version 27 March 2023.

only to protect the people who are represented by data but also the people who generate it.

Computer Science Perspective

In computer science, privacy can be viewed through the lens of system or information security. Consequently, developers typically focus on which data can be accessed and manipulated in a given system and on finding ways to protect against unwanted access. Thus, the foundation of a system analysis is a *threat model*. Which actor can technically do what, which actions are legitimate, and which are malicious? Malicious activities can then be obstructed by privacy enhancing technologies (PETs) for a particular threat model. These protections must be integrated into the system design[2] and not as a reactive add-on once the system has been designed or even developed.

While information security typically views privacy as protecting a user's data from a third party or a malicious actor, we must also consider privacy from the system itself. First, we must not only consider the data of a system, or even just the inferences that can be made with the data, but also the *goal* of the system itself. If the goal of the system is to manipulate and control the thought and actions of the user, it should be considered a violation of privacy. Technology deployers must ask if the system itself is counter to privacy ideals.

Second, the system may collect more data than its task requires (*counter to data minimization*) or use data beyond its initially intended purpose (*function creep*). These are especially relevant in the ML era.

Third, the balance between the privacy of users (and nonusers) and the utility of the system is decided by the system providers, so left unrestricted it will favor utility. The loss of privacy to individuals is an externality—a consequence that may not affect the system. This balance of utility and privacy must be considered not as something independent from the system, but as an essential part of its design.

As such, solutions for privacy in eye tracking must go beyond the technical. Ethical, legal, and social considerations are necessary to distinguish legitimate and malicious activities, both from outside of the eye-tracking system and from within. They also help define which parts of the system can or should be trustworthy to enforce the use of PETs. The technical measures to prevent attacks, as well as the collection, storage, *and use* of data, can be legally enforced and/or required by society such that all users and nonusers are safe.

Socio-Ethical and Legal Perspective

In scientific debates, privacy has been understood as an asset worth protecting since the end of the 19th century. In 1967, Alan Westin defined the essential properties of privacy. He focused on the control of information flows. Accordingly, *privacy is "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others."*[10] Consequently, privacy fulfills four specific functions.

1) Privacy fulfills the need not to be manipulated or dominated by others.
2) Privacy allows one to relieve stress, be free from social pressures and expectations, and "be oneself."
3) Privacy allows space and time to reflect the experiences of everyday life to enable an authentic, self-determined life.
4) Privacy defines the boundaries of interpersonal mental closeness or distance to others. This enables the individual to differentiate between the addressees of personal information.

Privacy is thus a space in which actions and states are controlled solely by the individual(s). Although this appears to be a purely spatial phenomenon, privacy is not limited to local dimensions. Rather, privacy removes specific actions, situations, and other states of mind or body from the control of others. Shifting the perspective to the private as the inner world of the person, it becomes evident that a boundary worth protecting is crossed by quantifying, evaluating, storing, and further processing data that inform about the inner world of a person. This boundary must only be crossed if well-defined conditions apply.

Numerous arguments exist about why privacy should be protected, usually, to protect the individual. According to Rössler, privacy enables an autonomous life.[5] Nissenbaum extended this and defined privacy as contextual integrity. If, for example, in a health context, confidential information about the patient's state of health falls into the wrong hands, this is not only a violation of individual privacy but a contextual collapse that fundamentally damages trust in the doctor–patient relationship. Thus, any data that relate to individuals must always be contextualized. In general, privacy is considered protected when individuals and groups can decide to whom they want to disclose their personal data, which includes finding a balance between the possibility of controlling one's own data and being part of communicative communities. These insights have to be carefully applied to eye tracking.

The protection of these different privacy dimensions is reflected today in data protection law, such as the General Data Protection Regulation (GDPR) in

Europe. The main goal of developing these laws was to protect the individual from the feared data supremacy of the state. However, with the emergence of private databases, data protection laws were extended to apply both to private and state data processors. Modern data protection laws pursue the goal of striking a balance between the processing interests of data controllers and the protection of data subjects. One result of this balance is that it is not always necessary to obtain the consent of data subjects because other reasons, such as legitimate interests, the performance of a contract, or the protection of life justify lawful processing. However, it is important in these further cases that processing may not take place to the detriment of an individual and, in case of doubt, consent should always be obtained.

PRIVACY-AWARE EYE TRACKING

Analyzing both of these perspectives in the context of eye tracking is challenging given that privacy-aware eye tracking is not yet a widely researched field. Consequently, our analysis could not build on a large number of papers in this area and we instead had to follow a methodology that is inspired by grounded theory and methods of constant comparison, contrasting, and dimensionalizing.[9] Based on our technical expertise, we first identified two future ideal-typical, yet realistic case studies. From these case studies we then derived concrete categories. It is important to note that in practice, situations never occur as unambiguously as they are described here but rather as mixed forms. Accordingly, our considerations should always be seen in the specific context in which eye tracking is being applied.

CASE STUDIES

We selected the following two case studies because eye tracking is likely to be readily integrated in cars and upcoming smart glasses, and hence, given the significant and imminent privacy risks that eye tracking will pose in these settings.

Case Study 1: Eye Tracking of Drivers

When car drivers become tired their ability to concentrate diminishes, which can lead to accidents. Similarly, external distractions, such as a phone call or operating the entertainment system, can cause drivers to miss traffic lights, signs, or obstacles on the road. The combined analysis of the gaze direction, fixation duration, and pupil dilation using eye tracking integrated into the car dashboard could help reduce these sources of accidents. Future cars could include pre-trained gaze analysis and fatigue detection systems to predict the driver's ability to concentrate. If these systems predict that the driver's concentration has waned or that the driver did not perceive traffic signs or obstacles, it could warn them or control certain aspects of the vehicle, e.g., limit the maximum speed of the vehicle. The system could further automatically adapt itself to the driver but we assume here all personal information is stored only locally in the vehicle. We further assume the analysis system is installed locally in the car, thus data are generated, processed, and stored in aggregated statistics but no data leave the vehicle. Such a system could be installed in private cars but also in car-sharing cars or taxis.

Case Study 2: Smart Glasses

Smart glasses project information into the field of view of the wearer using specialized optics. An integrated processing unit controls this projection and can be connected to the network via a mobile phone. Gaze direction is used in an increasing number of these glasses to adapt the visual content in real time, e.g., for gaze-contingent rendering. Gaze direction in 3D can further be linked with the users' orientation and location as well as objects in their environment. If eye movements and additional information are analyzed together, a detailed user model can be generated that can be used, for example, to anticipate user behavior or infer personality traits.[4] Eye-based classification of psychological and cognitive characteristics could further indicate eye or brain diseases or other psychological illnesses, even at an early stage. For the purpose of our discussion, we assume gaze data are stored on a cloud server, potentially for extended periods of time. We additionally assume gaze data are collected from many users to periodically retrain computational models that are then shipped to all smart glass users. We further assume that gaze data are analyzed jointly with additional data about location and objects in the environment, and that the smart glasses manufacturer observed increased model accuracy when predicting the personality and behavioral traits of the wearer. Therefore, the manufacturer predicts these properties only for the purpose of user experience and not for medical purposes.

PRIVACY CATEGORIES

Both use cases involve different privacy threats and have different data protection requirements. The violation and protection depend on several factors. The privacy-relevant aspects are not only about protecting users from

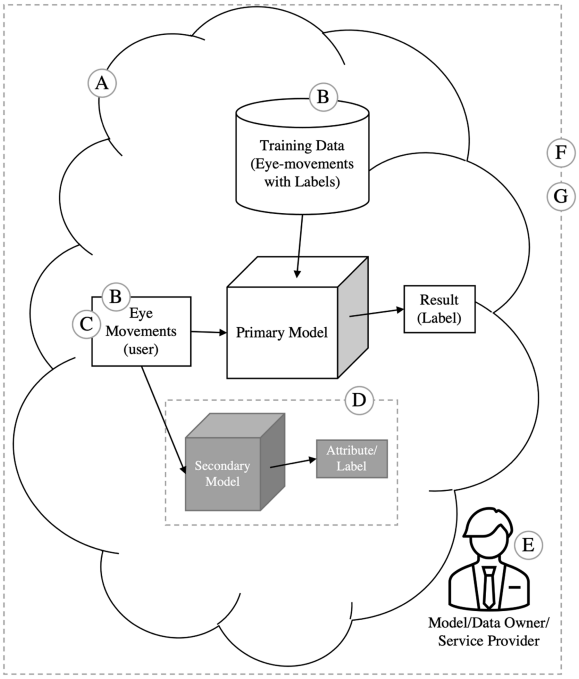

FIGURE 1. Diagram of the standard flow of data in a modern eye-tracking system. A primary (white) model is trained on an existing dataset. An entity deploys in a context and a user's eye movements are recorded and sent to the model to obtain a result. Secondarily (gray), the data may be fed to additional models trained on other tasks.

third-party malicious attacks or malfunctions but also of questioning the consequences for the privacy of the intended use of eye tracking. Different aspects of privacy relate to eye tracking and numerous design decisions during the development of new eye-tracking technologies affect privacy (see Figure 1):

Ⓐ Context
Ⓑ Identifiability
Ⓒ Type of data and derived information
Ⓓ Purpose of data processing and further processing
Ⓔ Data processor
Ⓕ Transparency
Ⓖ IT Security

Context

Privacy consequences of data processing must always be considered in the context of a particular application. The context can be classified in terms of its degree of privacy, e.g., regarding the number of people affected or their impact on different vulnerable populations. There are private spaces (e.g., a restroom) that differ fundamentally from semi-private spaces (e.g., doctors' offices or private cars) and public spaces (e.g., marketplaces or taxis). Private spaces are characterized by few individuals interacting in a manageable and foreseeable way. Whereas in public spaces, we can always expect to encounter strangers observing us, (semi-) private spaces are a likely setting for intimate activities. The threat to privacy posed by eye tracking is especially problematic in public spaces because the increased amount of external stimuli allows for better predictions about personality traits. For example, during a bus ride, a person's gaze can react to political advertisements or people of different genders or skin colors. Furthermore, while the eye-tracking context can change often and rapidly in public (smart glasses, e.g., can be worn while going to work where the person is first at home, then on the bus, afterward walking along a street and then entering the office), privacy settings are usually static and not adjusted to each new context, let alone automatically.

A structural difficulty relevant for all contexts is miniaturization that increasingly complicates the users' ability to know whether eye tracking is active. This increased miniaturization also means that eye tracking is invisibly penetrating the intimate areas of private life. Thus, both developers and users should be aware that people who have not consciously decided to use eye tracking may nonetheless be tracked.

We must not only be concerned with how many people are impacted but also by their vulnerability, e.g., people with low education levels, disabilities, or mental disorders. At best, vulnerable groups are partially covered, such as children via regulations for parental responsibility. Developers and operators of eye-tracking systems should therefore be aware of which group could be exposed to the technology and take measures to mitigate these risks early on in the development. Developers should also be aware of people appropriating technologies in their own ways and counteracting the intended use in the intended context. E.g., smart glasses intended to navigate can be used as a diary that stores images and the situational felt emotions together.

Identifiability

The European data protection regulation is only applicable if personal data are processed, leading to two options: either the processed data are personal data, and therefore fully protected, or not. Under GDPR (Art. 4 Nr. 1) personal data are any information relating to an identified or identifiable natural person ("data subject"). Whether a person is identifiable is a difficult technical question with many works showing how statistical or ML-based approaches can deanonymize data under different conditions and threat models.

There have been some technical solutions presented to aid in masking the identity of a user, e.g., using differential privacy which can reduce privacy loss while simultaneously allowing classification of the primary task.[7] However, such solutions do impact utility and do not apply if identity is paramount to the goal of the system. Here, a case-by-case, conscientious weighing of the advantages and disadvantages is necessary.

For eye tracking, the first question is how precisely gaze data itself can uniquely identify a person. This could qualify the processed data as personal according to the GDPR. Even by only using diameter measurements allows to identify a person with an accuracy of at least 60%. Furthermore, eye-tracking systems might process additional data, such as sound or video of the surroundings. Even if the controller could not identify a person via their gaze data alone, the enrichment with additional data, e.g., from voice control of the smart glasses, enables identification. This becomes a critical issue if eye tracking has been previously legally assessed as not processing personal data because the system was assumed to process different data types that by themselves do not identify a natural person. In this case, it would not fall under almost any data protection rules. When data then become personal data, they are processed without the safeguard of the GDPR. Thus, from a legal point of view, it is important to consider whether eye tracking can be used to identify a natural person. In terms of the first use case, it has to be evaluated whether the eye-tracking data itself, voice data (e.g., voice control), drivers profile (e.g., for seat settings or log-in for car-sharing app), break, and acceleration data can identify a person.

Predicted Information

Ethically, even aggregated data not directly identifying a natural person may be problematic. Models built on aggregated data can classify people based on their attributes or actions, leading to discrimination or other harms. In our first use case, eye tracking could be used to determine which drivers are more likely to be distracted and sets insurance rates based on driver attributes and eye-tacking data of other drivers. Because modern deep neural methods may infer gender and use it as a feature even if they are not explicitly instructed to consider gender as an attribute, it may classify women as more likely to be distracted while driving because women may statistically be more likely to drive alone with their small children than their male partners. Similar biases have been found in existing systems that do not account for the fact that, on average, women and men still traverse the world differently.[3]

Type of Data and Derived Information

Our use cases demonstrate that eye-tracking data are very specific but diverse. To check the driver's ability to steer a car, data about gaze (direction, duration, saccades, pupil size, etc.), the relative head position, and facial muscle activity around the eye area are analyzed. These data may, e.g., be used to infer the state of the physical or mental health of the driver. Personal data concerning health are specifically protected and processing is only possible if certain requirements are fulfilled, e.g., explicit consent, vital interests of the data subject, sustainable public interest, for preventive or occupational medicine, or necessary for reasons of public interest in the area of public health. One key difference from other personal signals of the human body, such as voice, is that eye movements are rarely under conscious control, which must be considered in the legal evaluation (e.g., in the balancing of interests).

Purpose of Data Processing and Further Processing

Purpose is a critical factor in assessing the privacy implications of some technology. Data can only be collected for specified, explicit, and legitimate purposes and cannot be further processed in a manner that is incompatible with those purposes. There are only few exceptions to this principle, e.g., scientific or historical research purposes. Moreover, the purpose determines the legal basis for data processing. Since gaze data are considered sensitive according to Art. 9 GDPR, explicit consent will likely be the common legal basis for eye tracking in practice. Situations, such as carpooling, present a challenge, as passengers may be unaware of the presence of eye tracking and, therefore, cannot consent.

In the car use case, the primary purpose of eye tracking is to enhance the safety of the driver, the passengers, and others on the road. Following this goal is always associated with a price, e.g., loss of freedom or autonomy and invasion of privacy. But evidence shows that people are willing to pay this price and even pay more generally a higher price to ensure security than for other purposes, such as entertainment. It can, therefore, be assumed that people are more likely to give their consent to eye tracking in the car than for purely entertainment purposes. This means that privacy issues in connection with the purpose of the application have direct consequences for the user acceptance of eye tracking. Commercial purposes (marketing/advertising), as could be attributed to smart glasses, are not regarded as an accepted legitimation for curtailing privacy, at least in western cultures. According to the privacy paradox, however,

there can be a gap between people's behavior in specific contexts and what they would consider legitimate in a broader valuation of privacy.[6] There is a social, culturally dependent question of which goods (freedom, security, privacy) are important and desirable. If the curtailment of one good for the benefit of another systematically excludes certain groups, this can lead to significant social problems.

Shift of Purpose

Users actively appropriate technologies in their own unique ways and may not use them as intended. Developers should be aware that the actual use of eye-tracking technology may differ from the intended purpose. The legal judgement requires a separate legal basis for most processing for purposes other than the initially intended one. Fulfilling these legal requirements is particularly challenging if gaze data are processed within ML systems that create detailed user models and make far-reaching prediction about individuals and groups. The task of processors is to do everything to technically prevent any misuse and illegitimate further processing of the data. It is important to think in terms of those affected by a data processing operation to distinguish between illegitimate and legitimate processing operations.

Data Processor

Only if users know who is processing their data they can make an informed decision about disclosing it. Whoever can regulate an individual's privacy options has great power over that person. It makes a difference whether the users themselves initially select privacy settings, whether settings are preset but can be changed (privacy-by-default and privacy-by-design), or whether certain settings are immutable and predetermined by the service provider. To assess the additional risks posed by data processors to privacy, we must consider different dimensions.

First, the location of the data controller determines the legal, social, and ethical background for data processing. An eye-tracking application run by an European company within Europe falls under the stricter legal regulation of the GDPR than in the United States. Even if the data stay on the eye-tracking device, there is a high risk in the U.S. that it could be disclosed as evidence during court proceedings—and thereby used for other purposes. Therefore, the GDPR (Art. 44) sets specific rules for third-country data transfers to cover the differences in controller locations.

Second, the relationship between the data controller and the data subject can influence the intensity and probability of privacy infringement. If eye-tracking systems are deployed in the cars of professional drivers, they might fear negative consequences from their employer if they do not consent to data processing. As such, the GDPR (Art. 88) regulates working environments specifically.

Third, the controller's technical and financial capabilities—whether it is a private person, (big) company, or state—influence the extent of the processing.

Transparency

According to the transparency principle [Art. 5(1)(a) GDPR], personal data shall only be processed transparently. This key requirement of transparency has multiple dimensions as follows.

Particular Use (Recognizability of the System)

Transparency forms the legal basis of consent, and any device with eye-tracking functionality must inform the user about the data usage. In the case of smart glasses, users must put them on willingly, although this does not unequivocally mean that they are aware of or consent to their eye movements being tracked. Steil et al.[8] proposed automatically closing a shutter around the scene camera of smart glasses in privacy-sensitive situations. In the car use case, it is even more critical to reflect on whether all (co)drivers are aware that the vehicle integrates such a system.

Types, Amount, Storage, and Analysis of the Collected Data

The transparency principle requires that the data subject be informed (Art. 12 ff. GDPR). Therefore, the basic elements of the processing must be communicated: the identity and contact details of the controller, the purposes of the processing, the recipients, etc.

Transfer of Data to Third Parties and Possible Secondary use of the Data

Fulfilling the legal requirement of transparency is more difficult for others, like guests in a car. They may not be familiar with how the information is encoded (e.g., colors of light signals).

Explainability and Comprehensibility of Software Results

Even if people consent to eye tracking, one cannot assume that they understand what additional conclusions can be drawn from the generated data. The fact that privacy-relevant data are discretely generated poses a major problem from an ethical and legal perspective. Users should be informed as to which measurements form the basis of the data generation and which can be derived from the generated data.

Regarding the first three dimensions, legal obligations exist for data controllers, while the fourth dimension is currently covered only rudimentary. Therefore, a sensitive design that goes beyond the legal requirements is necessary if privacy is to be ensured. The first step to more transparency could be a privacy dashboard showing the possible insights of the gathered data. This information reduces the privacy risks for the data subject, so it is more likely that the application is considered legal by Data Protection Authorities. Developers should always strive to design eye-tracking systems that break down the complexity of processing to a level that can be understood by the general user.

IT Security

Suitable PETs should be selected and implemented during the design of the technology. To determine which technologies are suitable, the threat model should be established. What can be realistically controlled by a malicious party? The answer depends on the use case.

In the first use case, we assume the data are recorded, processed, and stored only at the local device, i.e., the car. Thus, an attacker would need to break into the car's software and open up a channel back outside. As cars need an Internet connection for automatic emergency calls, the car software manufacturer needs to ensure this channel cannot be abused. Technical measures for this are sandboxing or access control, which should be explained by the car software manufacturer. Better would be an independent review of the security-critical parts by domain experts.

In the second use case, we assume that data are recorded at the local device but processed by a service provider online. Thus, an attacker would need to intercept the Internet connection, compromise the device, or compromise the service provider. While the service provider may be trustworthy, it is clear that the data need to be protected while in transit. A simple, yet effective way to implement this protection is with standard TLS protocols that are also used elsewhere for secure Internet connections. The service provider should inform users about this protection mechanism. This protection can be verified by users via packet inspection and is strengthened by further allowing users to check open-sourced software components.

In both cases, if we assume the service providers are malicious, privacy protection becomes challenging. Privacy-related issues also arise regarding the storage of data. In the car use case, where we assume that no data are stored beyond the period of use, there is a low risk of the data being accessed or further processed in a way that is not appropriate for the users. In the smart glasses case, the presence of ML indicates that the data will both be stored so that the system can learn and persist in the model itself. The location of the data and model storage plays an important role. Assuming that the users themselves want to determine who can have access to their data, this is made much easier for them if they can locate the spatial location of the data storage and, as in the car example, even have access control over that location and, if necessary, the ability to manually turn off the device. The existence of a model trained on user data complicates this picture, however, and removing samples from a trained model is not technically feasible.

CONCLUSION AND OUTLOOK

Eye tracking has significant potential to become a key component in a range of future applications but, at the same time, this prospect also raises a number of serious and far-reaching privacy questions. In this position article we argued that eye-tracking technology design must ensure that each person can not only truly decide independently what data they disclose about themselves—to whom, how, and for what purpose, but also control what information can be inferred about them and their communities. To achieve this, we call for the eye-tracking community to develop new tools and frameworks for ethically developing and deploying eye-tracking technology, including the use of privacy-by-design principles from the beginning of development through deployment and fully local processing of data.

ACKNOWLEDGMENTS

The work of author Céline Gressel was supported by German Federal Ministry for Education and Research under Grant 16SV8185. The work of author Murat Karaboga was supported by German Federal Ministry for Education and Research under Grants 16KIS1374K and 16KIS1372K. The work of author Michael Raschke was supported by German Federal Ministry for Education and Research under Grant 16SV8105). The work of author Andreas Bulling was supported by European Research Council (ERC) under Grant 801708.

REFERENCES

1. A. Bulling and D. Roggen, "Recognition of visual memory recall processes using eye movement analysis," in *Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput.*, 2011, pp. 455–464.

2. A. Cavoukian et al., "Privacy by design: The 7 foundational principles," *Inform. Privacy Commissioner Ontario*, 2009.
3. L. Gauvin et al., "Gender gaps in urban mobility," *Humanities Social Sci. Commun.*, vol. 7, no. 1, pp. 1–13, 2020.
4. S. Hoppe, T. Loetscher, S. Morey, and A. Bulling, "Eye movements during everyday behavior predict personality traits," *Front. Hum. Neurosci.*, vol. 12, 2018, Art. no. 105.
5. B. Rössler, *The Value of Privacy*. Hoboken, NJ, USA: Wiley, 2018.
6. D. J. Solove, "The myth of the privacy paradox," *George Washington Law Rev.*, vol. 89, 2021, Art. no. 1.
7. J. Steil, I. Hagestedt, M. X. Huang, and A. Bulling, "Privacy-aware eye tracking using differential privacy," in *Proc. ACM Int. Symp. Eye Tracking Res. Appl.*, 2019, pp. 1–9.
8. J. Steil, M. Koelle, W. Heuten, S. Boll, and A. Bulling, "PrivacEye: Privacy-preserving head-mounted eye tracking using egocentric scene image and eye movement features," in *Proc. ACM Int. Symp. Eye Tracking Res. Appl.*, 2019, pp. 1–10.
9. A. Strauss and J. M. Corbin, *Grounded Theory in Practice*. Newbury Park, CA, USA: Sage, 1997.
10. A. F. Westin, "Privacy and freedom," *Washington Lee Law Rev.*, vol. 25, no. 1, 1968, Art. no. 168.

CÉLINE GRESSEL is a sociologist of technology at the International Center for Ethics in the Sciences and Humanities (IZEW), 72074, Tübingen, Germany. She investigates social and ethical aspects, the use of virtual and augmented realities for a healthy life, as well as the question of the conditions for a successful integration of these aspects in technology development projects. She studied sociology, psychology, and education. Contact her at celine.gressel@izew.uni-tuebingen.de.

REBEKAH OVERDORF is an assistant professor at the School of Law, Criminal Justice and Public Administration, University of Lausanne, 1015, Lausanne, Switzerland, working in security, privacy, and digital forensics, more specifically anonymity networks (traffic analysis vulnerabilities of Tor hidden services), stylometry (authorship attribution based on writing style), cybercriminal networks, and social media-based attacks with particular expertise and interest in using data science and machine learning methods to study security. She was a postdoctoral researcher at EPFL and a postdoctoral researcher at the Computer Security and Industrial Cryptography (COSIC), KU Leuven, Leuven, Belgium. Overdorf received her Ph.D. degree in computer science from the College of Computing, Drexel University. Contact her at rebekah.overdorf@unil.ch.

INKEN HAGESTEDT is a privacy expert at Apheris AI GmbH, 10999, Berlin, Germany, working on collaborative data ecosystems. Her research interests include privacy-preserving computation and sharing of biomedical data, including epigenetic and eye tracking data. Hagestedt received her Ph.D. degree from the CISPA Helmholtz Center for Information Security, Saarbrücken, Germany. Contact her at i.hagestedt@apheris.com.

MURAT KARABOGA is a senior researcher and has been at the Competence Center Emerging Technologies, Fraunhofer Institute for Systems and Innovations Research, 76135, Karlsruhe, Germany, since January 2014. His work focuses on policy analysis and the analysis of governance and actor structures in the domain of information and communication technologies. Contact him at murat.karaboga@isi.fraunhofer.de.

MICHAEL RASCHKE is co-founder and managing director of Blickshift GmbH, 70563, Stuttgart, Germany, and an expert for a visualization-based eye movement analysis. Since 2009, he has been working on new methods and techniques for the analysis of perceptual and cognitive processes at the Institute for Visualization and Interactive Systems at the University of Stuttgart. Contact him at michael.raschke@blickshift.de.

HELMUT LURTZ is a lawyer in the field of data protection law and is currently at the Novartis Business Services GmbH, 83607, Holzkirchen, Germany. He was a research assistant at the University of Kassel, where he especially focused on employee data protection law in Germany. Within the scope of his project activities, he has been working in interdisciplinary teams, especially on the implementation of human-focused cyber-physical systems in the Industrial-IoT (MyCPS) and diagnosis/e-assistance systems based on a VR/AR-system for ophthalmic diseases (IDeA). Contact him at helmut.lurtz@novartis.com.

ANDREAS BULLING is full professor of computer science at the University of Stuttgart, 70569, Stuttgart, Germany, where he directs the research group "Human-Computer Interaction and Cognitive Systems." He was a Feodor Lynen and Marie Curie research fellow at the University of Cambridge, Cambridge, U.K., and a senior researcher at the Max Planck Institute for Informatics, Germany. His research interests include computer vision, machine learning, and human–computer interaction. Bulling received his Ph.D. degree in information technology and electrical engineering from ETH Zurich, Zurich, Switzerland. Contact him at andreas.bulling@vis.uni-stuttgart.de.