

PrivacEye: Privacy-Preserving Head-Mounted Eye Tracking Using Egocentric Scene Image and Eye Movement Features

Julian Steil
Max Planck Institute for Informatics
Saarland Informatics Campus
jsteil@mpi-inf.mpg.de

Marion Koelle
University of Oldenburg
marion.koelle@uol.de

Wilko Heuten
OFFIS - Institute for IT
wilko.heuten@offis.de

Susanne Boll
University of Oldenburg
susanne.boll@uol.de

Andreas Bulling
University of Stuttgart
andreas.bulling@vis.uni-stuttgart.de

ABSTRACT

Eyewear devices, such as augmented reality displays, increasingly integrate eye tracking, but the first-person camera required to map a user’s gaze to the visual scene can pose a significant threat to user and bystander privacy. We present *PrivacEye*, a method to detect privacy-sensitive everyday situations and automatically enable and disable the eye tracker’s first-person camera using a mechanical shutter. To close the shutter in privacy-sensitive situations, the method uses a deep representation of the first-person video combined with rich features that encode users’ eye movements. To open the shutter without visual input, *PrivacEye* detects changes in users’ eye movements alone to gauge changes in the “privacy level” of the current situation. We evaluate our method on a first-person video dataset recorded in daily life situations of 17 participants, annotated by themselves for privacy sensitivity, and show that our method is effective in preserving privacy in this challenging setting.

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools; Ubiquitous and mobile computing; Ubiquitous and mobile devices; Human computer interaction (HCI);**

KEYWORDS

Privacy Protection; Egocentric Vision; Gaze Behaviour

ACM Reference Format:

Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2019. PrivacEye: Privacy-Preserving Head-Mounted Eye Tracking Using Egocentric Scene Image and Eye Movement Features. In *2019 Symposium on Eye Tracking Research and Applications (ETRA '19)*, June 25–28, 2019, Denver, CO, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3314111.3319913>

1 INTRODUCTION

Eyewear devices, such as head-mounted displays or augmented reality glasses, have recently emerged as a new research platform

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ETRA '19, June 25–28, 2019, Denver, CO, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6709-7/19/06...\$15.00
<https://doi.org/10.1145/3314111.3319913>

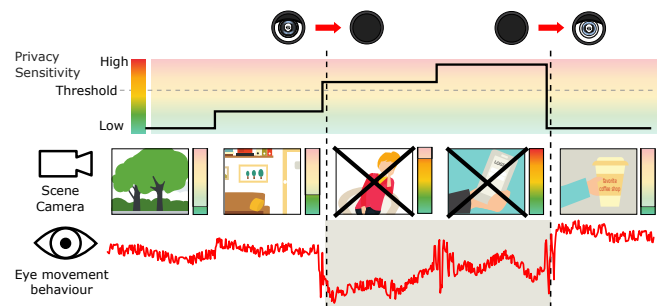


Figure 1: Our method uses a mechanical camera shutter (top) to preserve users’ and bystanders’ privacy with head-mounted eye trackers. Privacy-sensitive situations are detected by combining deep scene image and eye movement features (middle) while changes in eye movement behaviour alone trigger the reopening of the camera shutter (bottom).

in fields such as human-computer interaction, computer vision, or the behavioural and social sciences [Bulling and Kunze 2016]. An ever-increasing number of these devices integrate eye tracking to analyse attention allocation [Eriksen and Yeh 1985; Sugano et al. 2016], for computational user modelling [Fischer 2001; Itti and Koch 2001], or hands-free interaction [Hansen et al. 2003; Vertegaal et al. 2003]. Head-mounted eye tracking typically requires two cameras: An eye camera that records a close-up video of the eye and a high-resolution first-person (scene) camera to map gaze estimates to the real-world scene [Kassner et al. 2014]. The scene camera poses a serious privacy risk as it may record sensitive personal information, such as login credentials, banking information, or text messages, as well as infringe on the privacy of bystanders [Perez et al. 2017]. Privacy risks intensify with the unobtrusive integration of eye tracking in ordinary glasses frames [Tonsen et al. 2017].

In the area of first-person vision, prior work identified strategies of self-censorship [Koelle et al. 2017] that, however, are prone to (human) misinterpretations and forgetfulness, or the accidental neglect of social norms and legal regulations. In consequence, user experience and comfort are decreased and the user’s mental and emotional load increases, while sensitive personal information can still be accidentally disclosed. Other works therefore investigated alternative solutions, such as communicating a bystander’s privacy preferences using short-range wireless radio [Aditya et al. 2016],

visual markers [Schiff et al. 2007], or techniques to compromise recordings [Harvey 2012; Truong et al. 2005]. However, all of these methods require bystanders to take action themselves to protect their privacy. None of these works addressed the problem at its source, i.e. the scene camera, nor did they offer a means to protect the privacy of both the wearer and potential bystanders.

To address this limitation, we propose *PrivacEye*, the first method for privacy-preserving head-mounted eye tracking (see Figure 1). The key idea and core novelty of our method is to detect users' transitions into and out of privacy-sensitive everyday situations by leveraging both cameras available on these trackers. If a privacy-sensitive situation is detected, the scene camera is occluded by a physical shutter. Our design choice to use a non-spoofable physical shutter, which closes for some time and therefore provides feedback to bystanders, is substantiated by Koelle et al., who highlight an increased trustworthiness over LED lights on the camera or pure software solutions [Koelle et al. 2018]. While this approach is secure and visible to bystanders, it prohibits visual input from the scene. Thus, our method analyses changes in the users' eye movement behaviour alone to detect if they exit a privacy-sensitive situation and then reopens the camera shutter. A naive, vision-only system could reopen the shutter at regular intervals, e.g. every 30 seconds, to detect whether the current situation is still privacy-sensitive. However, this approach may negatively affect perceived reliability and increase mistrust in the system. Thus, our eye-tracking approach promises significant advantages over a purely interval-based approach in terms of user experience and perceived trustworthiness.

Our approach is motivated by prior work that demonstrates that eye movements are a rich source of information on a user's everyday activities [Bulling et al. 2011; Steil and Bulling 2015], social interactions and current environment [Bulling et al. 2013], or even a user's personality traits [Hoppe et al. 2018]. In addition, prior work showed that perceived privacy sensitivity is related to a user's location and activity [Hoyle et al. 2015]. We therefore hypothesize that *privacy sensitivity* transitively informs a user's *eye movements*. We are the first to confirm this transitivity, which results as a reasoned deduction from prior work.

The specific contributions of this work are three-fold: First, we present *PrivacEye*, the first method that combines the analysis of egocentric scene image features with eye movement analysis to enable context-specific, privacy-preserving de-activation and re-activation of a head-mounted eye tracker's scene camera. As such, we show a previously unconfirmed transitive relationship over the users' eye movements, their current activity and environment, as well as the perceived privacy sensitivity of the situation they are in. Second, we evaluate our method on a dataset of real-world mobile interactions and eye movement data, fully annotated with locations, activities, and privacy sensitivity levels of 17 participants. Third, we provide qualitative insights on the perceived social acceptability, trustworthiness, and desirability of *PrivacEye*, based on semi-structured interviews, using a fully functional prototype.

2 RELATED WORK

Research on eye tracking privacy is sparse. Thus, our work mostly relates to previous works on (1) privacy concerns with first-person cameras and (2) privacy enhancing methods for (wearable) cameras.

2.1 Privacy Concerns - First-Person Cameras

First-person cameras are well-suited for continuous and unobtrusive recordings, which causes them to be perceived as unsettling by bystanders [Denning et al. 2014]. Both users' and bystanders' privacy concerns and attitudes towards head-mounted devices with integrated cameras were found to be affected by context, situation, usage intentions [Koelle et al. 2015], and user group [Profita et al. 2016]. Hoyle et al. showed that the presence and the number of people in a picture, specific objects (e.g., computer displays, ATM cards, physical documents), location, and activity affected whether lifeloggers deemed an image "shareable" [Hoyle et al. 2014]. They also highlighted the need for automatic privacy-preserving mechanisms to detect those elements, as individual sharing decisions are likely to be context-dependent and subjective. Their results were partly confirmed by Price et al., who, however, found no significant differences in sharing when a screen was present [Price et al. 2017]. Chowdhury et al. found that whether lifelogging imagery is suitable for sharing is (in addition to content, scenario, and location) mainly determined by its sensitivity [Chowdhury et al. 2016]. Ferdous et al. proposed a set of guidelines that, among others, include semi-automatic procedures to determine the sensitivity of captured images according to user-provided preferences [Ferdous et al. 2017]. All highlight the privacy sensitivity of first-person recordings and the importance of protecting user and bystander privacy.

2.2 Enhancing Privacy of First-Person Cameras

To increase the privacy of first-person cameras for bystanders, researchers have suggested communicating their privacy preferences to nearby capture devices using wireless connections as well as mobile or wearable interfaces [Krombholz et al. 2015]. Others have suggested preventing unauthorised recordings by compromising the recorded imagery, e.g., using infra-red light signals [Harvey 2010; Yamada et al. 2013] or disturbing face recognition [Harvey 2012]. In contrast to our approach, these techniques all require the bystander to take action, which might be impractical due to costs and efforts [Denning et al. 2014].

A potential remedy are automatic, or semi-automatic approaches, such as *PlaceAvoider*, a technique that allows users to "blacklist" sensitive spaces, e.g., bedroom or bathroom [Templeman et al. 2014]. Similarly, *ScreenAvoider* allowed users to control the disclosure of images of computer screens showing potentially private content [Korayem et al. 2016]. Erickson et al. proposed a method to identify security risks, such as ATMs, keyboards, and credit cards, in images captured by first-person wearable devices [Erickson et al. 2014]. However, instead of assessing the whole scene in terms of privacy sensitivity, their systems only detected individual sensitive objects. Raval et al. presented *MarkIt*, a computer vision-based privacy marker framework that allowed users to use self-defined bounding boxes and hand-gestures to restrict visibility of content on two dimensional surfaces (e.g. white boards) or sensitive real-world objects [Raval et al. 2014]. *iPrivacy* automatically detects privacy-sensitive objects from social images users are willing to share using deep multi-task learning [Yu et al. 2017]. It warns the image owners what objects in the images need to be protected before sharing and recommends privacy settings.

While all of these methods improved privacy, they either only did so post-hoc, i.e. after images had already been captured, or they required active user input. In contrast, our approach aims to prevent potentially sensitive imagery from being recorded at all, automatically in the background, i.e. without engaging the user. Unlike current computer vision based approaches that work in image space, e.g. by masking objects or faces [Raval et al. 2014; Shu et al. 2016; Yamada et al. 2013], restricting access [Korayem et al. 2016], or deleting recorded images post-hoc [Templeman et al. 2014], we de-activate the camera completely using a mechanical shutter and also signal this to bystanders. Our approach is the first to employ eye movement analysis for camera re-activation that, unlike other sensing techniques (e.g., microphones, infra-red cameras), does not compromise the privacy of potential bystanders.

3 DESIGN RATIONALE

PrivacEye’s design rationale is based on user and bystander goals and expectations. In this section, we outline how PrivacEye’s design contributes to avoiding erroneous disclosure of sensitive information, so-called misclosures (User Goal 1), and social friction (User Goal 2), and detail on three resultant design requirements.

3.1 Goals and Expectations

Avoid Misclosure of Sensitive Data. A user wearing smart glasses with an integrated camera would typically do so to make use of a particular functionality, e.g., visual navigation. However, the device’s “always-on” characteristic causes it to capture more than originally intended. A navigation aid would require capturing certain landmarks for tracking and localisation. In addition, unintended imagery and potentially sensitive data is captured. Ideally, to prevent misclosures [Caine 2009], sensitive data should not be captured. However, requiring the user to constantly monitor her actions and environment for potential sensitive information (and then de-activate the camera manually) might increase the workload and cause stress. As users might be forgetful, misinterpret situations, or overlook privacy-sensitive items, automatic support from the system would be desirable from a user’s perspective.

Avoid Social Friction. The smart glasses recording capabilities may cause social friction if they do not provide a clear indication whether the camera is on or off: Bystanders might even perceive device usage as a privacy threat when the camera is turned off [Koelle et al. 2015, 2018]. In consequence, they feel uncomfortable around such devices [Bohn et al. 2005; Denning et al. 2014; Ens et al. 2015; Koelle et al. 2015]. Similarly, user experience is impaired when device users feel a need for justification as they could be accused of taking surreptitious pictures [Häkkinen et al. 2015; Koelle et al. 2018].

3.2 Design Requirements

As a consequence of these user goals there are three essential design requirements that PrivacEye addresses: (1) The user can make use of the camera-based functionality without the risk of misclosures or leakage of sensitive information. (2) The system pro-actively reacts to the presence or absence of potentially privacy-sensitive situations and objects. (3) The camera device communicates the recording status clearly to both user and bystander.



Figure 2: PrivacEye prototype with labelled components (B) and worn by a user with a USB-connected laptop in a backpack (A). Detection of privacy-sensitive situations using computer vision closes the camera shutter (C), which is reopened based on a change in the privacy detected level in a user’s eye movements (D).

4 PRIVACEYE PROTOTYPE

Our fully functional PrivacEye prototype, shown in Figure 2, is based on the PUPIL head-mounted eye tracker [Kassner et al. 2014] and features one 640×480 pixel camera (the so-called “eye camera”) that records the right eye from close proximity (30 fps), and a second camera (1280×720 pixels, 24 fps) to record a user’s environment (the so-called “scene camera”). The first-person camera is equipped with a fish eye lens with a 175° field of view and can be closed with a mechanical shutter. The shutter comprises a servo motor and a custom-made 3D-printed casing, including a mechanical lid to occlude the camera’s lens. The motor and the lid are operated via a micro controller, namely a Feather M0 Proto. Both cameras and the micro controller were connected to a laptop via USB. PrivacEye further consists of two main software components: (1) detection of privacy-sensitive situations to close the mechanical camera shutter and (2) detection of changes in user’s eye movements that are likely to indicate suitable points in time for reopening the camera shutter.

4.1 Detection of Privacy-Sensitive Situations

The approaches for detecting privacy-sensitive situations we evaluated are (1) *CNN-Direct*, (2) *SVM-Eye*, and (3) *SVM-Combined*.

4.1.1 CNN-Direct. Inspired by prior work on predicting privacy-sensitive pictures posted in social networks [Orekondu et al. 2017], we used a pre-trained GoogleNet, a 22-layer deep convolutional neural network [Szegedy et al. 2015]. We adapted the original GoogleNet model for our specific prediction task by adding two additional fully connected (FC) layers. The first layer was used to reduce the feature dimensionality from 1024 to 68 and the second one, a Softmax layer, to calculate the prediction scores. Output of our model was a score for each first-person image indicating whether the situation visible in that image was privacy-sensitive or not. The cross-entropy loss was used to train the model. The full network architecture is included in the supplementary material.

4.1.2 SVM-Eye. Given that eye movements are independent from the scene camera’s shutter status, they can be used to (1) detect privacy-sensitive situations while the camera shutter is open and (2) detect changes in the subjective privacy level while the camera shutter is closed. The goal of this second component is to instead

detect changes in a user’s eye movements that are likely linked to changes in the privacy sensitivity of the current situation and thereby to keep the number of times the shutter is reopened as low as possible. To detect privacy-sensitive situations and changes, we trained SVM classifiers (kernel=rbf, C=1) with characteristic eye movement features, which we extracted using only the eye camera video data. We extracted a total of 52 eye movement features, covering fixations, saccades, blinks, and pupil diameter (see Table 2 in the supplementary material for a list and description of the features). Similar to [Bulling et al. 2011], each saccade is encoded as a character forming words of length n (wordbook). We extracted these features using a sliding window of 30 seconds (step size of 1 sec).

4.1.3 SVM-Combined. A third approach for the detection of privacy-sensitive situations is a hybrid method. We trained SVM classifiers using the extracted eye movement features (52) and combined them with CNN features (68) from the scene image, which we extracted from the first fully connected layer of our trained CNN model, creating 120 feature large samples. With the concatenation of eye movement and scene features, we are able to extend the information from the two previous approaches during recording phases where the camera shutter is open.

5 EXPERIMENTS

We evaluated the different approaches on their own and in combination in a realistic temporal sequential analysis trained in a person-specific (leave-one-recording-out) and person-independent (leave-one-person-out) manner. We assume that the camera shutter is open at start up. If no privacy-sensitive situation is detected, the camera shutter remains open and the current situation is rated “non-sensitive”, otherwise, the camera shutter is closed and the current situation is rated “privacy-sensitive”. Finally, we analysed error cases and investigated the performance of PrivacEye in different environments and activities.

5.1 Dataset

While an ever-increasing number of eye movement datasets have been published in recent years (see [Bulling et al. 2012, 2011; Hoppe et al. 2018; Steil and Bulling 2015; Sugano and Bulling 2015] for examples), none of them focused on privacy-related attributes. We therefore make resource to a previously recorded dataset [Steil et al. 2018]. The dataset of Steil et al. contains more than 90 hours of data recorded continuously from 20 participants (six females, aged 22-31) over more than four hours each. Participants were students with different backgrounds and subjects with normal or corrected-to-normal vision. During the recordings, participants roamed a university campus and performed their everyday activities, such as meeting people, eating, or working as they normally would on any day at the university. To obtain some data from multiple, and thus also “privacy-sensitive”, places on the university campus, participants were asked to not stay in one place for more than 30 minutes. Participants were further asked to stop the recording after about one and a half hours so that the laptop’s battery packs could be changed and the eye tracker re-calibrated. This yielded three recordings of about 1.5 hours per participant. Participants regularly interacted with a mobile phone provided to them and were also encouraged to use their own laptop, desktop computer, or music

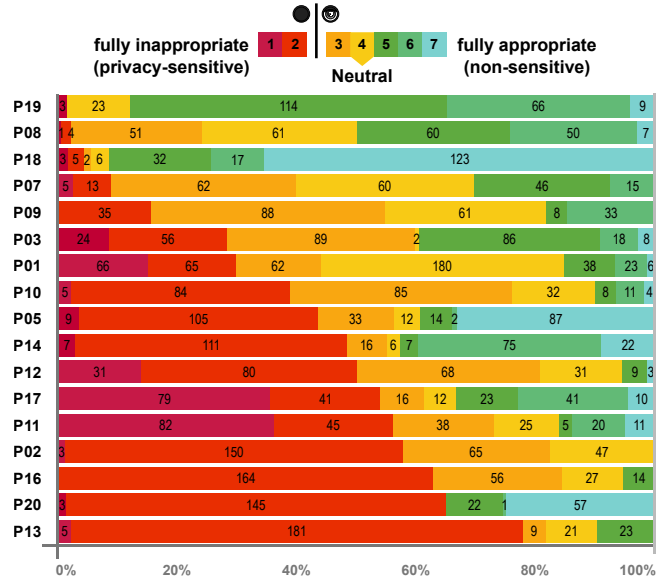


Figure 3: Privacy sensitivity levels rated on a 7-pt Likert scale from 1: fully inappropriate (i.e., privacy-sensitive) to 7: fully appropriate (i.e., non-sensitive). Distribution in labelled minutes/level per participant, sorted according to a “cut-off” between closed shutter (level 1-2) and open shutter (level 3-7). In practice, the “cut-off” level could be chosen according to individual ratings as measured by PAQ.

player if desired. The dataset thus covers a rich set of representative real-world situations, including sensitive environments and tasks. The data collection was performed with the same equipment as shown in Figure 2 excluding the camera shutter.

5.2 Data Annotation

The dataset was fully annotated by the participants themselves with continuous annotations of location, activity, scene content, and subjective privacy sensitivity level. 17 out of the 20 participants finished the annotation of their own recording resulting in about 70 hours of annotated video data. They again gave informed consent and completed a questionnaire on demographics, social media experience and sharing behaviour (based on Hoyle et al. [Hoyle et al. 2014]), general privacy attitudes, as well as other-contingent privacy [Baruh and Cemalcilar 2014] and respect for bystander privacy [Price et al. 2017]. General privacy attitudes were assessed using the *Privacy Attitudes Questionnaire* (PAQ), a modified Westin Scale [Westin 2003] as used by [Caine 2009; Price et al. 2017].

Annotations were performed using Advene [Aubert et al. 2012]. Participants were asked to annotate continuous video segments showing the same situation, environment, or activity. They could also introduce new segments in case a privacy-relevant feature in the scene changed, e.g., when a participant switched to a sensitive app on the mobile phone. Participants were asked to annotate each of these segments according to the annotation scheme (see supplementary material). Privacy sensitivity was rated on a 7-point Likert scale ranging from 1 (fully inappropriate) to 7 (fully appropriate). As we expected our participants to have difficulties understanding the concept of “privacy sensitivity”, we rephrased it

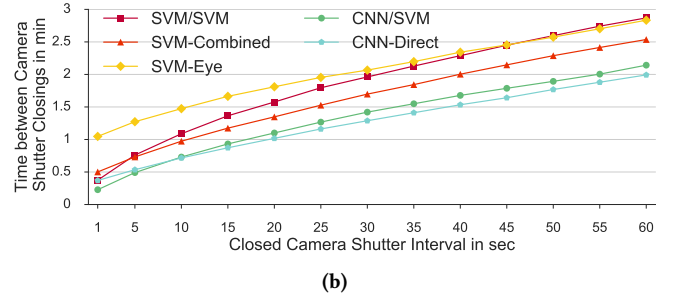
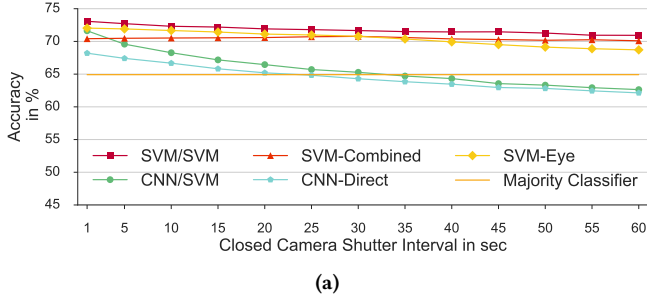


Figure 4: Person-specific leave-one-recording-out evaluation showing the achieved accuracy (a) and the time between camera shutter closings (b) across different closed camera shutter intervals.

for the annotation to “How appropriate is it that a camera is in the scene?”. Figure 3 visualises the labelled privacy sensitivity levels for each participant. Based on the latter distribution, we pooled ratings of 1 and 2 in the class “privacy-sensitive”, and all others in the class “non-sensitive”. A consumer system would provide the option to choose this “cut-off”. We will use these two classes for all evaluations and discussions that follow in order to show the effectiveness of our proof-of-concept system. The dataset is available at <https://www.mpi-inf.mpg.de/MPIIPrivacEye/>.

5.3 Sequential Analysis

To evaluate PrivacEye, we applied the three proposed approaches separately as well as in combination in a realistic temporal sequential analysis, evaluating the system as a whole within person-specific (leave-one-recording-out) and person-independent (leave-one-person-out) cross validation schemes. Independent of CNN or SVM approaches, we first trained and then tested in a person-specific fashion. That is, we trained on two of the three recordings of each participant and tested on the remaining one – iteratively over all combinations and averaging the performance results in the end. For the leave-one-person-out cross validation, we trained on the data of 16 participants and tested on the remaining one. *SVM-Eye* is the only one of the three proposed approaches that allows PrivacEye to be functional when no scene imagery is available, i.e., when the shutter is closed. Additionally, it can be applied when the shutter is open thus serving both software components of PrivacEye. While the camera shutter is not closed, i.e., scene imagery is available, *CNN-Direct* or *SVM-Combined* can be applied. To provide a comprehensive picture, we then analysed the combinations *CNN-Direct* + *SVM-Eye* (*CNN/SVM*) and *SVM-Combined* + *SVM-Eye* (*SVM/SVM*). The first approach is applied when the camera shutter is open and *SVM-Eye* only when the shutter is closed. For the sake of completeness, we also evaluated *SVM-Combined* and *CNN-Direct* on the whole dataset. However, these two methods represent hypothetical best-case scenarios in which eye and scene features are always available. As this is in practice not possible, they have to be viewed as an upper-bound baseline. For evaluation purposes, we apply the proposed approaches within a step size of one second in a sequential manner. The previously predicted camera shutter position (open or close) decides which approach is applied for the prediction of the current state to achieve realistic results. We use $Accuracy = \frac{TP+TN}{TP+FP+TN+FN}$, where TP, FP, TN, and FN count sample-based true positives, false positives, true negatives, and false negatives, as performance indicator.

5.3.1 CNN-Direct. For training the CNN, which classifies a given scene image directly as privacy-sensitive or non-sensitive, we split the data from each participant into segments. Each change in environment, activity, or the annotated privacy sensitivity level starts a new segment. We used one random image per segment for training.

5.3.2 SVM-Eye and SVM-Combined. The SVM classifiers use only eye movement features (*SVM-Eye*) or the combination of eye movement and CNN features (*SVM-Combined*). We standardised the training data (zero mean, unit variance) for the person-specific and leave-one-person-out cross validation before training the classifiers, and used the same parameters for the test data.

5.4 Results

With potential usability implications in mind, we evaluate performance over a range of closed camera shutter intervals. If a privacy-sensitive situation is detected from the *CNN-Direct* or *SVM-Combined* approach, the camera shutter is kept closed for an interval between 1 and 60 seconds. If *SVM-Eye* is applied and no privacy change is detected, the shutter remains closed. In a practical application, users build more trust when the camera shutter remains closed, at least for a sufficient amount of time, to guarantee the protection of privacy-sensitive scene content when such a situation is detected [Koelle et al. 2018]. We also evaluated *CNN-Direct* and *SVM-Combined* on the whole recording as hypothetical best-case scenarios. However, comparing their performance against the combinations *SVM/SVM* and *CNN/SVM* illustrate the performance improvement using *SVM-Eye* when the camera shutter is closed.

5.4.1 Person-specific (leave-one-recording-out) evaluation. Figure 4a shows the person-specific accuracy performance of PrivacEye against increasing camera shutter closing time for two combinations *CNN/SVM* and *SVM/SVM*, and *SVM-Eye*, which can be applied independent of the camera shutter status. Besides *CNN-Direct* and *SVM-Combined*, the majority class classifier serves as a baseline, predicting the majority class from the training set. The results reveal that all trained approaches and combinations perform above the majority class classifier. However, we can see that *CNN-Direct* and its combination with *SVM-Eye* (*CNN/SVM*) perform below the other approaches and below the majority class classifier for longer closed camera shutter intervals. *SVM-Eye* and *SVM-Combined* perform quite robustly, around 70% accuracy, while *SVM-Eye* performs better for shorter intervals and *SVM-Combined* for longer intervals. The interplay approach *SVM/SVM*, which we would include in our prototype, exceeds 73% with a closed camera shutter interval of

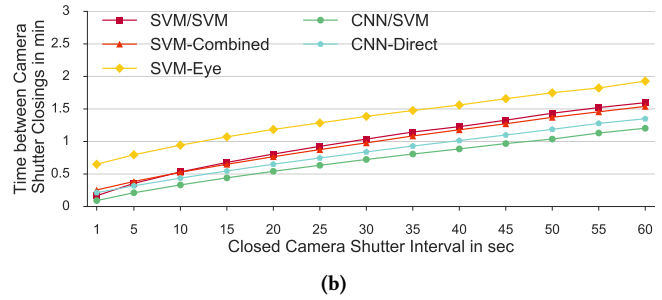
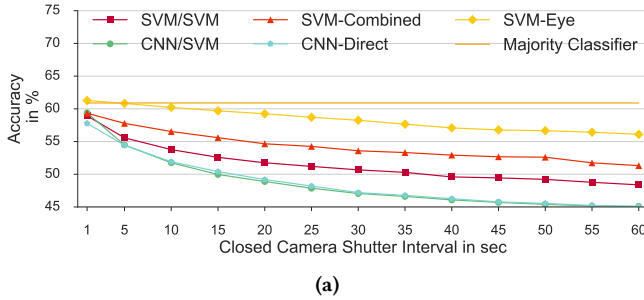


Figure 5: Person-independent leave-one-person-out evaluation showing the accuracy results (a) and the time between closing the camera shutter (b) across different closed camera shutter intervals.

one second and outperforms all other combinations in terms of accuracy in all other intervals. One reason for the performance improvement of *SVM/SVM* in comparison to its single components is that *SVM-Combined* performs better for the detection of privacy-sensitive situations when the camera shutter is open while *SVM-Eye* performs better for preserving privacy-sensitive situations so that the camera shutter remains closed. Another aim of our proposed approach is the reduction of opening and closing events during a recording to strengthen reliability and trustworthiness. A comparison of Figure 4a and Figure 4b renders a clear trade-off between accuracy performance and time between camera shutter closing instances. For very short camera shutter closing times the *SVM-Eye* approach, which only relies on eye movement features from the eye camera, shows the best performance, whereas for longer camera shutter closing times, the combination *SVM/SVM* shows better accuracy with a comparable amount of time between camera shutter closing instances. However, the current approaches are actually not able to reach the averaged ground truth of about 8.2 minutes between camera shutter closings.

5.4.2 Person-independent (leave-one-person-out) evaluation. The more challenging task, which assumes that privacy-sensitivity could generalise over multiple participants, is given in the person-independent leave-one-person-out cross validation of Figure 5a. Similar to the person-specific evaluation, *CNN-Direct* and *CNN/SVM* perform worse than the other approaches. Here, *SVM-Eye* outperforms *SVM-Combined* and *SVM/SVM*. However, none of the approaches are able to outperform the majority classifier. These results show that eye movement features generalise better over multiple participants to detect privacy-sensitive situations than scene image information. Comparing the number of minutes between camera shutter closing events of person-specific and leave-one-person-out in Figure 4b and Figure 5b, the person-specific approach outperforms the person-independent leave-one-person-out evaluation scheme for each approach. This shows that privacy sensitivity does not fully generalise, and consumer systems would require a person-specific calibration and online learning.

5.5 Error Case Analysis

For PrivacEye, it is not only important to detect the privacy-sensitive situations (TP), but equally important to detect non-sensitive situations (TN), which are relevant to grant a good user experience. Our results suggest that the combination *SVM/SVM* performs best for the person-specific case. For this setting we carry out a detailed error case analysis of our system for the participants'

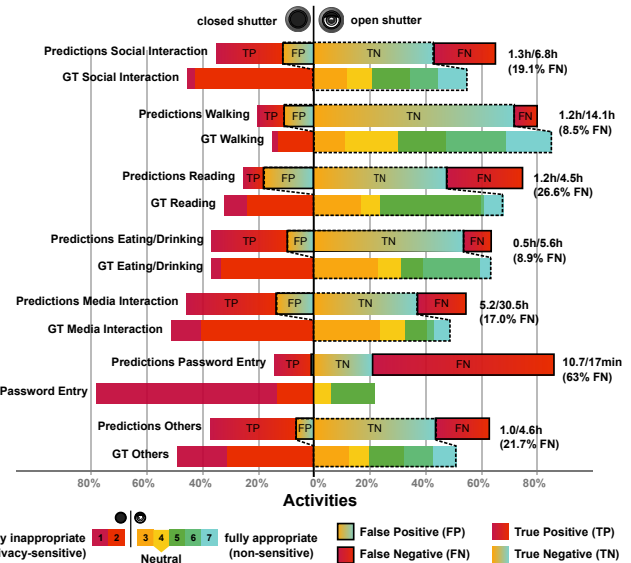


Figure 6: Error case analysis for different and activities showing the “cut-off” between closed shutter (left, *privacy-sensitive*) and open shutter (right, *non-sensitive*) with PrivacEye prediction and the corresponding ground truth (GT). False positives (FP) are *non-sensitive* but protected (closed shutter), false negatives (FN) are *privacy-sensitive* but unprotected (open shutter).

different activities. For the activities outlined in Figure 6, PrivacEye works best while eating/drinking and in media interactions. Also, the results are promising for detecting social interactions. The performance for password entry, however, is still limited. Although the results show that it is possible to detect password entry, the amount of true negatives (TN) is high compared to other activities. This is likely caused by the dataset’s under-representation of this activity, which characteristically lasts only a few seconds. Future work might be able to eliminate this by specifically training for password and PIN entry, which will enable the classifier to better distinguish between PIN entry and, e.g., reading. In the supplementary material we provide an in-depth error case analysis to further investigate error cases in different environments.

6 USER FEEDBACK

Collecting initial subjective feedback during early stages of system development allows us to put research concepts in a broader context

and helps to shape hypotheses for future quantitative user studies. In this section, we report on a set of semi-structured one-to-one interviews on the use of head-worn augmented reality displays in general, and our interaction design and prototype in particular. To obtain the user feedback, we recruited 12 new and distinct participants (six females), aged 21 to 31 years ($M=24$, $SD=3$) from the local student population. They were enrolled in seven highly diverse majors, ranging from computer science and biology to special needs education. We decided to recruit students, given that we believe they and their peers are potential users of a future implementation of our prototype. We acknowledge that this sample, consisting of rather well educated young adults (with six of them having obtained a Bachelor's degree), is not representative for the general population. Interviews lasted about half an hour and participants received a 5 Euro Amazon voucher. We provide a detailed interview protocol as part of the supplementary material. The semi-structured interviews were audio recorded and transcribed for later analysis. Subsequently, qualitative analysis was performed following inductive category development [Mayring 2014]. Key motives and reoccurring themes were extracted and are presented in this section, where we link back to PrivacEye's design and discuss implications for future work.

6.1 User Views on Transparency

Making it transparent (using the 3D-printed shutter), whether the camera was turned on or off, was valued by all participants. Seven participants found the integrated shutter increased perceived safety in contrast to current smart glasses; only few participants stated that they made no difference between the shutter and other visual feedback mechanisms, e.g. LEDs ($n=2$). Several participants noted that the physical coverage increased trustworthiness because it made the system more robust against hackers (*concerns:hacking*, $n=3$) than LEDs. Concluding, the usage of physical occlusion could increase perceived safety and, thus, could be considered an option for future designs. Participants even noted that the usage of the shutter as reassuring as pasting up a laptop camera (*laptop comparison*, $n=4$), which is common practice.

6.2 User Views on Trustworthiness

In contrast, participants also expressed technology scepticism, particularly that the system might secretly record audio (*concerns:audio*, $n=5$) or malfunction (*concerns:malfunction*, $n=4$). With the increasing power of deep neural networks malfunctions, system failures, or inaccuracies will be addressable in the future, interaction designers will have to address this fear of "being invisibly audio-recorded". A lack of knowledge about eye tracking on both the user's and the bystander's side might even back this misconception. Therefore, future systems using eye tracking for context recognition will have to clearly communicate their *modus operandi*.

6.3 Perceived Privacy of Eye Tracking

The majority of participants claimed to have no privacy concerns about smart glasses with integrated eye tracking functionality: "I do see no threat to my privacy or the like from tracking my eye movements; this [the eye tracking] would rather be something which could offer a certain comfort." (P11) Only two participants expressed concerns about their privacy, e.g., due to fearing eye-based emotion

recognition (P3). One was uncodeable. This underlines our assumption that eye tracking promises privacy-preserving and socially acceptable sensing in head-mounted augmented reality devices and, thus, should be further explored.

6.4 Desired Level of Control

Participants were encouraged to elaborate on whether the recording status should be user-controlled or system-controlled. P10 notes: "I'd prefer if it was automatic, because if it is not automatic, then the wearer can forget to do that [de-activating the camera]. Or maybe he will say 'Oh, I do not want to do that' and then [...] that leads to a conflict. So better is automatic, to avoid questions." Four other participants also preferred the camera to be solely controlled by the system (*control:automatic*, $n=4$). Their preference is motivated by user forgetfulness ($n=5$), and potential non-compliance of users (in the bystander use case, $n=1$). Only two participants expressed a preference for sole (*control>manual*) control, due to an expected lack of system reliability, and technical feasibility. Two responses were uncodable. All other participants requested to implement manual confirmation of camera de-activation/re-activation or manual operation as alternative modes (*control:mixed*, $n=4$), i.e., they like to feel in control. To meet these user expectations, future interaction designs would have to find an adequate mix of user control and automatic support through the system; for example, by enabling users to explicitly record sensitive information (e.g. in cases of emergency) or label seemingly non-sensitive situations "confidential".

7 DISCUSSION

We discuss PrivacEye in light of the aforementioned design and user requirements and results of the technical evaluation.

7.1 Privacy Preserving Device Behaviour

Design Requirements 1 and *2* demand privacy-preserving device behaviour. With PrivacEye, we have presented a computer vision routine that analyses all imagery obtained from the scene camera, combined with eye movement features with regard to privacy sensitivity and, in case a situation requires protection, the ability to de-activate the scene camera and close the system's camera shutter. This approach prevents both accidental mislosure and malicious procurement (e.g. hacking) of sensitive data, as has been positively highlighted by our interview participants. However, closing the shutter comes at the cost of having the scene camera unavailable for sensing after it has been de-activated. PrivacEye solves this problem by using a second eye camera that allows us, in contrast to prior work, to locate all required sensing hardware on the user's side. With PrivacEye we have provided proof-of-concept that context-dependent re-activation of a first-person scene camera is feasible using only eye movement data. Future work will be able to build upon these findings and further explore eye tracking as a sensor for privacy-enhancing technologies. Furthermore, our results provide first prove that there is indeed a transitive relationship over privacy sensitivity and a user's eye movements.

7.2 Defining Privacy Sensitivity

Prior work indicates that the presence of a camera may be perceived appropriate or inappropriate depending on social context,

location, or activity [Hoyle et al. 2015, 2014; Price et al. 2017]. However, related work does, to the best of our knowledge, not provide any insights on eye tracking data in this context. For this reason, we run a dedicated data collection and ground truth annotation. Designing a practicable data collection experiment requires the overall time spent by a participant for data recording and annotation to be reduced to a reasonable amount. Hence, we made use of an already collected data set, and re-invited the participants only for the annotation task. While the pre-existing data set provided a rich diversity of privacy-sensitive locations and objects, including smart phone interaction, and realistically depicts everyday student life, it is most likely not applicable to other contexts, e.g., industrial work or medical scenarios.

For PrivacEye, we rely on a 17-participant-large, ground truth annotated dataset with highly realistic training data. Thus, the collected training data cannot be fully generalised, e.g., to other regions or age groups. On the plus side, however, this data already demonstrates that in a future real-world application, sensitivity ratings may vary largely between otherwise similar participants. This might also be affected by their (supposedly) highly individual definition of “privacy”. Consequently, a future consumer system should be pre-trained and then adapted online, based on personalised retraining after user feedback. In addition, users should be enabled to select their individual “cut-off”, i.e., the level from which a recording is blocked, which was set to “2” for PrivacEye. Future users of consumer devices might choose more rigorous or relaxed “cut-off” levels depending on their personal preference. Initial user feedback also indicated that an interaction design that combines automatic, software-controlled de- and re-activation, with conscious control of the camera by the user, could be beneficial.

7.3 Eye Tracking for Privacy-Enhancement

Eye tracking is advantageous for bystander privacy given that it only senses users and their eye movements. In contrast to, e.g., microphones or infra-red sensing, it senses a bystander and/or an environment only indirectly via the user’s eye motion or reflections. Furthermore, eye tracking allows for implicit interaction and is non-invasive, and we expect it to become integrated into commercially available smart glasses in the near future. On the other hand, as noted by Liebling and Preibusch [Liebling and Preibusch 2014; Preibusch 2014], eye tracking data is a scare resource, which can be used to identify user attributes like age, gender, health, or user’s current task. For this reason, the collection and use of eye tracking data could be perceived as a potential threat to user privacy. However, our interviews showed that eye tracking was not perceived as problematic by a large majority of our participants. Nevertheless, eye tracking data must be protected by appropriate privacy policies and data hygiene.

To use our proposed hardware prototype in a real-world scenario, data sampling and analysis need to run on a mobile phone. The CNN feature extraction is currently the biggest computational bottleneck, but could be implemented in hardware to allow for real-time operation (c.f., Qualcomm’s Snapdragon 845). Further, we believe that a consumer system should provide an accuracy >90% which could be achieved using additional sensors such as GPS or inertial tracking. However, presenting the first approach for automatic de-

and re-activation of a first-person camera that achieves ~73% with competitive performance to *ScreenAvoider* (54.2 - 77.7%) [Korayem et al. 2014] and *iPrivacy* (~75%) [Yu et al. 2017], which are restricted to scene content protection and post-hoc privacy protection, we provide a solid basis for follow up work. We note that a generalized person-independent model for privacy sensitivity protection is desirable. For this work only the participants themselves labelled their own data. Aggregated labels of multiple annotators would result in a more consistent and generalizable “consensus” model and improve test accuracy, but would dilute the measure of perceived privacy sensitivity, which is highly subjective [Price et al. 2017]. Specifically, similar activities and environments were judged differently by the individual participants, as seen in Figure 3. The availability of this information is a core contribution of our dataset.

7.4 Communicating Privacy Protection

The interaction design of PrivacEye tackles *Design Requirement 3* using a non-transparent shutter. Ens et al. [Ens et al. 2015] reported that the majority of their participants expected to feel more comfortable around a wearable camera device if it clearly indicated to be turned on or off. Hence, our proposed interaction design aims to improve a bystander’s awareness of the recording status by employing an *eye metaphor*. Our prototype implements the “eye lid” as a retractable shutter made from non-transparent material: open when the camera is active, closed when the camera is inactive. Thus, the metaphor mimics “being watched” by the camera. The “eye lid” shutter ensures that bystanders can comprehend the recording status without prior knowledge, as eye metaphors have been widely employed for interaction design, e.g., to distinguish visibility or information disclosure [Motti and Caine 2016; Pousman et al. 2004; Schlegel et al. 2011] or to signal user attention [Chan and Minamizawa 2017]. Furthermore, in contrast to visual status indicators, such as point lights (LEDs), physical occlusion is non-spoofable (c.f., [Denning et al. 2014; Portnoff et al. 2015]). This concept has been highly appreciated during our interviews, which is why we would recommend adopting it for future hardware designs.

8 CONCLUSION

In this work, we have proposed PrivacEye, a method that combines first-person computer vision with eye movement analysis to enable context-specific, privacy-preserving de-activation and re-activation of a head-mounted eye tracker’s scene camera. We have evaluated our method quantitatively on a 17-participant dataset of fully annotated everyday behaviour as well as qualitatively, by collecting subjective user feedback from 12 potential future users. To the best of our knowledge, our method is the first of its kind and prevents potentially sensitive imagery from being recorded at all, without the need for active user input. As such, we believe the method opens up a new and promising direction for future work in head-mounted eye tracking, the importance of which will only increase with further miniaturisation and integration of eye tracking in head-worn devices or even in normal glasses frames.

ACKNOWLEDGMENTS

This work was funded, in part, by a JST CREST research grant under Grant No.: JPMJCR14E1, Japan.

REFERENCES

- Paarijaat Aditya, Rijurekha Sen, Peter Druschel, Seong Joon Oh, Rodrigo Benenson, Mario Fritz, Bernt Schiele, Bobby Bhattacharjee, and Tong Tong Wu. 2016. I-pic: A platform for Privacy-compliant Image Capture. In *Annual International Conference on Mobile Systems, Applications, and Services (MobiSys)*. ACM, 235–248. <https://doi.org/10.1145/2906388.2906412>
- Olivier Aubert, Yannick Prié, and Daniel Schmitt. 2012. Advence As a Tailorable Hypervideo Authoring Tool: A Case Study. In *Proceedings of the 2012 ACM Symposium on Document Engineering (DocEng '12)*. ACM, New York, NY, USA, 79–82. <https://doi.org/10.1145/2361354.2361370>
- Lemi Baruh and Zeynep Cemalcilar. 2014. It is more than personal: Development and validation of a multidimensional privacy orientation scale. *Personality and Individual Differences* 70 (2014), 165–170. <https://doi.org/DOI:10.1016/j.paid.2014.06.042>
- Jürgen Bohn, Vlad Coroamă, Marc Langheinrich, Friedemann Mattern, and Michael Rohs. 2005. Social, economic, and ethical implications of ambient intelligence and ubiquitous computing. In *Ambient Intelligence*. Springer, 5–29. https://doi.org/10.1007/3-540-27139-2_2
- Andreas Bulling and Kai Kunze. 2016. Eyewear Computers for Human-Computer Interaction. *ACM Interactions* 23, 3 (2016), 70–73. <https://doi.org/10.1145/2912886>
- Andreas Bulling, Jamie A. Ward, and Hans Gellersen. 2012. Multimodal Recognition of Reading Activity in Transit Using Body-Worn Sensors. *ACM Transactions on Applied Perception* 9, 1 (2012), 2:1–2:21. <https://doi.org/10.1145/2134203.2134205>
- Andreas Bulling, Jamie A. Ward, Hans Gellersen, and Gerhard Tröster. 2011. Eye Movement Analysis for Activity Recognition Using Electrooculography. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 4 (April 2011), 741–753. <https://doi.org/10.1109/TPAMI.2010.86>
- Andreas Bulling, Christian Weichel, and Hans Gellersen. 2013. EyeContext: Recognition of High-level Contextual Cues from Human Visual Behaviour. In *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. 305–308. <https://doi.org/10.1145/2470654.2470697>
- Kelly Caine. 2009. *Exploring everyday privacy behaviors and misclosures*. Georgia Institute of Technology.
- Liwei Chan and Kouta Minamizawa. 2017. FrontFace: Facilitating Communication Between HMD Users and Outsiders Using Front-facing-screen HMDs. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, USA, Article 22, 5 pages. <https://doi.org/10.1145/3098279.3098548>
- Soumyadeb Chowdhury, Md Sadek Ferdous, and Joemon M Jose. 2016. Exploring Lifelog Sharing and Privacy. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16)*. ACM, New York, NY, USA, 553–558. <https://doi.org/10.1145/2968219.2968320>
- Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In situ with Bystanders of Augmented Reality Glasses: Perspectives on Recording and Privacy-mediating Technologies. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. ACM, 2377–2386. <https://doi.org/10.1145/2556288.2557352>
- Barrett Ens, Tovi Grossman, Fraser Anderson, Justin Matejka, and George Fitzmaurice. 2015. Candid interaction: Revealing hidden mobile and wearable computing activities. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. ACM, 467–476. <https://doi.org/10.1145/2807442.2807449>
- Zackory Erickson, Jared Compiano, and Richard Shin. 2014. Neural Networks for Improving Wearable Device Security. (2014).
- Charles W Eriksen and Yei-yu Yeh. 1985. Allocation of attention in the visual field. *Journal of Experimental Psychology: Human Perception and Performance* 11, 5 (1985), 583. <https://doi.org/10.1037/0096-1523.11.5.583>
- Md Sadek Ferdous, Soumyadeb Chowdhury, and Joemon M Jose. 2017. Analysing privacy in visual lifelogging. *Pervasive and Mobile Computing* (2017). <https://doi.org/10.1016/j.pmcj.2017.03.003>
- Gerhard Fischer. 2001. User modeling in human-computer interaction. *User modeling and user-adapted interaction* 11, 1-2 (2001), 65–86. <https://doi.org/10.1023/A:1011145532042>
- Jonna Häkikilä, Farnaz Vahabpour, Ashley Colley, Jani Väyrynen, and Timo Koskela. 2015. Design Probes Study on User Perceptions of a Smart Glasses Concept. In *Proceedings of the 14th International Conference on Mobile and Ubiquitous Multimedia (MUM '15)*. ACM, New York, NY, USA, 223–233. <https://doi.org/10.1145/2836041.2836064>
- John Paulin Hansen, Anders Sewerin Johansen, Dan Witzner Hansen, Kenji Itoh, and Satoru Mashino. 2003. Command without a click: Dwell time typing by mouse and gaze selections. In *Proceedings of Human-Computer Interaction-INTERACT*. 121–128.
- Adam Harvey. 2010. Camoflash-anti-paparazzi clutch. (2010). <http://ahprojects.com/projects/camoflash/> accessed 13/09/2017.
- Adam Harvey. 2012. CVDazzle: Camouflage from Computer Vision. *Technical report* (2012).
- Sabrina Hoppe, Tobias Loetscher, Stephanie Morey, and Andreas Bulling. 2018. Eye Movements During Everyday Behavior Predict Personality Traits. *Frontiers in Human Neuroscience* 12 (2018). <https://doi.org/10.3389/fnhum.2018.00105>
- Roberto Hoyle, Robert Templeman, Denise Anthony, David Crandall, and Apu Kapadia. 2015. Sensitive lifelogs: A privacy analysis of photos from wearable cameras. In *Proceedings of the 33rd Annual ACM conference on human factors in computing systems*. ACM, 1645–1648. <https://doi.org/10.1145/2702123.2702183>
- Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy Behaviors of Lifeloggers using Wearable Cameras. In *International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. ACM, 571–582. <https://doi.org/10.1145/2632048.2632079>
- Laurent Itti and Christof Koch. 2001. Computational modelling of visual attention. *Nature reviews neuroscience* 2, 3 (2001), 194. <https://doi.org/10.1038/35058500>
- Moritz Kassner, William Patera, and Andreas Bulling. 2014. Pupil: an open source platform for pervasive eye tracking and mobile gaze-based interaction. In *Adj. Proc. UbiComp*. 1151–1160. <http://dx.doi.org/10.1145/2638728.2641695>
- Marion Koelle, Wilko Heuten, and Susanne Boll. 2017. Are You Hiding It?: Usage Habits of Lifelogging Camera Wearers. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, USA, Article 80, 8 pages. <https://doi.org/10.1145/3098279.3122123>
- Marion Koelle, Matthias Kranz, and Andreas Möller. 2015. Don't look at me that way!: Understanding User Attitudes Towards Data Glasses Usage. In *Proceedings of the 17th international conference on human-computer interaction with mobile devices and services*. ACM, 362–372. <https://doi.org/10.1145/2785830.2785842>
- Marion Koelle, Katrin Wolf, and Susanne Boll. 2018. Beyond LED Status Lights-Design Requirements of Privacy Notices for Body-worn Cameras. In *Proceedings of the Twelfth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 177–187. <https://doi.org/10.1145/3173225.3173234>
- Mohammed Korayem, Robert Templeman, Dennis Chen, David Crandall, and Apu Kapadia. 2014. Screenavoider: Protecting computer screens from ubiquitous cameras. *arXiv preprint arXiv:1412.0008* (2014).
- Mohammed Korayem, Robert Templeman, Dennis Chen, David Crandall, and Apu Kapadia. 2016. Enhancing lifelogging privacy by detecting screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 4309–4314. <https://doi.org/10.1145/2702123.2702183>
- Katharina Krombholz, Adrian Dabrowski, Matthew Smith, and Edgar Weippl. 2015. Ok glass, leave me alone: towards a systematization of privacy enhancing technologies for wearable computing. In *International Conference on Financial Cryptography and Data Security*. Springer, 274–280. https://doi.org/10.1007/978-3-662-48051-9_20
- Daniel J Liebling and Sören Preibusch. 2014. Privacy considerations for a pervasive eye tracking world. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 1169–1177. <https://doi.org/10.1145/2638728.2641688>
- Philipp Mayring. 2014. *Qualitative content analysis: theoretical foundation, basic procedures and software solution*. 143 pages.
- Vivian Genaro Motti and Kelly Caine. 2016. Towards a Visual Vocabulary for Privacy Concepts. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 60. SAGE Publications Sage CA: Los Angeles, CA, 1078–1082. <https://doi.org/10.1177/1541931213601249>
- Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. 2017. Towards a Visual Privacy Advisor: Understanding and Predicting Privacy Risks in Images. In *International Conference on Computer Vision (ICCV 2017)*. Venice, Italy. <https://doi.org/10.1109/ICCV.2017.398>
- Alfredo J Perez, Sherali Zeadally, and Scott Griffith. 2017. Bystanders' Privacy. *IT Professional* 19, 3 (2017), 61–65. <https://doi.org/10.1109/MITP.2017.42>
- Rebecca S Portnoff, Linda N Lee, Serge Egelman, Pratyush Mishra, Derek Leung, and David Wagner. 2015. Somebody's watching me?: Assessing the effectiveness of Webcam indicator lights. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1649–1658. <https://doi.org/10.1145/2702123.2702611>
- Zachary Pousman, Giovanni Iachello, Rachel Fithian, Jehan Moghazy, and John Stasko. 2004. Design iterations for a location-aware event planner. *Personal and Ubiquitous Computing* 8, 2 (2004), 117–125. <https://doi.org/10.1007/s00779-004-0266-y>
- Sören Preibusch. 2014. Eye-tracking. Privacy interfaces for the next ubiquitous modality. In *2014 W3C Workshop on Privacy and User-Centric Controls*. <https://www.w3.org/2014/privacysw/pp/Preibusch.pdf>
- Blaine A. Price, Avelie Stuart, Gul Calikli, Ciaran McCormick, Vikram Mehta, Luke Hutton, Arosha K. Bandara, Mark Levine, and Bashar Nuseibeh. 2017. Logging You, Logging Me: A Replicable Study of Privacy and Sharing Behaviour in Groups of Visual Lifeloggers. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 2, Article 22 (June 2017), 18 pages. <https://doi.org/10.1145/3090087>
- Halley Profita, Reem Albaghli, Leah Findlater, Paul Jaeger, and Shaun K Kane. 2016. The AT Effect: How Disability Affects the Perceived Social Acceptability of Head-Mounted Display Use. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 4884–4895. <https://doi.org/10.1145/2858036.2858130>
- Nisarg Raval, Animesh Srivastava, Kiron Lebeck, Landon Cox, and Ashwin Machanavajjhala. 2014. MarkIt: Privacy markers for protecting visual secrets. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 1289–1295. <https://doi.org/10.1145/2638728.2641707>
- Jeremy Schiff, Marci Meingast, Deirdre K Mulligan, Shankar Sastry, and Ken Goldberg. 2007. Respectful cameras: Detecting visual markers in real-time to address privacy

- concerns. In *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*. IEEE, 971–978. https://doi.org/10.1007/978-1-84882-301-3_5
- Roman Schlegel, Apu Kapadia, and Adam J Lee. 2011. Eyeing your exposure: quantifying and controlling information sharing for improved privacy. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*. ACM, 14. <https://doi.org/10.1145/2078827.2078846>
- Jiayu Shu, Rui Zheng, and Pan Hui. 2016. Cardea: Context-aware visual privacy protection from pervasive cameras. *arXiv preprint arXiv:1610.00889* (2016).
- Julian Steil and Andreas Bulling. 2015. Discovery of Everyday Human Activities From Long-term Visual Behaviour Using Topic Models. In *Proc. ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. 75–85. <https://doi.org/10.1145/2750858.2807520>
- Julian Steil, Philipp Müller, Yusuke Sugano, and Andreas Bulling. 2018. Forecasting user attention during everyday mobile interactions using device-integrated and wearable sensors. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services*. ACM, 1. <https://doi.org/10.1145/3229434.3229439>
- Yusuke Sugano and Andreas Bulling. 2015. Self-Calibrating Head-Mounted Eye Trackers Using Egocentric Visual Saliency. In *Proc. of the 28th ACM Symposium on User Interface Software and Technology (UIST 2015)*. 363–372. <https://doi.org/10.1145/2807442.2807445>
- Yusuke Sugano, Xucong Zhang, and Andreas Bulling. 2016. Aggregaze: Collective estimation of audience attention on public displays. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, 821–831. <https://doi.org/10.1145/2984511.2984536>
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going Deeper with Convolutions. In *Computer Vision and Pattern Recognition (CVPR)*. <http://arxiv.org/abs/1409.4842>
- Robert Templeman, Mohammed Korayem, David J Crandall, and Apu Kapadia. 2014. PlaceAvoider: Steering First-Person Cameras away from Sensitive Spaces. In *NDSS*. <https://doi.org/10.14722/ndss.2014.23014>
- Marc Tonsen, Julian Steil, Yusuke Sugano, and Andreas Bulling. 2017. InvisibleEye: Mobile Eye Tracking Using Multiple Low-Resolution Cameras and Learning-Based Gaze Estimation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 1, 3 (2017), 106:1–106:21. <https://doi.org/10.1145/3130971>
- Khai Truong, Shwetak Patel, Jay Summet, and Gregory Abowd. 2005. Preventing camera recording by designing a capture-resistant environment. *UbiComp 2005: Ubiquitous Computing* (2005), 903–903. https://doi.org/10.1007/11551201_5
- Roel Vertegaal et al. 2003. Attentive user interfaces. *Commun. ACM* 46, 3 (2003), 30–33. <https://doi.org/10.1145/636772.636794>
- Alan F Westin. 2003. Social and political dimensions of privacy. *Journal of social issues* 59, 2 (2003), 431–453. <https://doi.org/10.1111/1540-4560.00072>
- Takayuki Yamada, Seiichi Gohshi, and Isao Echizen. 2013. Privacy visor: Method based on light absorbing and reflecting properties for preventing face image detection. In *Systems, Man, and Cybernetics (SMC), 2013 IEEE International Conference on*. IEEE, 1572–1577. <https://doi.org/10.1109/SMC.2013.271>
- Jun Yu, Baopeng Zhang, Zhengzhong Kuang, Dan Lin, and Jianping Fan. 2017. iPrivacy: image privacy protection by identifying sensitive objects via deep multi-task learning. *IEEE Transactions on Information Forensics and Security* 12, 5 (2017), 1005–1016. <https://doi.org/10.1109/TIFS.2016.2636090>