Ok Glass, Leave me Alone: Towards a Systematization of Privacy Enhancing Technologies for Wearable Computing

Katharina Krombholz[1], Adrian Dabrowski[1], Matthew Smith[2], and Edgar Weippl[1]

[1] SBA Research, Vienna, Austria
(firstletterfirstname)(lastname)@sba-research.org
[2] Usable Security and Privacy Group, University of Bonn, Germany
smith@cs.uni-bonn.de

Abstract. In the coming age of wearable computing, devices such as Google Glass will become as ubiquitous as smartphones. Their foreseeable deployment in public spaces will cause distinct implications on the privacy of people recorded by these devices. Particularly the discreet recording capabilities of such devices pose new challenges to consensual image disclosure. Therefore, new *Privacy Enhancing Technologies (PETs)* will be needed to help preserve our digital privacy. At the time of writing, no such PETs are available on the market to communicate privacy preferences towards Glass. In the scientific literature, a handful of approaches has been presented. However, none of them has been evaluated regarding their affordances and overall usefulness. In this paper, we provide the first systematization and qualitative evaluation of state of the art PETs that were designed to communicate privacy preferences towards (wearable) cameras, such as Google Glass. The purpose of this paper is to foster a broader discourse on how such technology should be designed in order to be fully privacy preserving and usable.

1 Introduction

Wearable computers with integrated cameras such as Google Glass might soon become as ubiquitous as smartphones. Due to their hands-free user interface and the discreet recording capabilities, collecting and sharing images and videos becomes easier than ever. In contrary to smartphones and other mobile devices, Google Glass literally remains in the wearer's face all the time. Consequentially, many bystanders view such wearables as invasive and fear substantial implications on their digital privacy. Since the paradigm shift to user-generated content on the Internet, the awareness for picture privacy has risen. The foreseeable deployment of wearable technology in public spaces is about to multiply the set of challenges related to non-consensual disclosure of graphical material on the Internet. In such situations, getting informed consent of all people recorded by such a device is unfeasible. Recently, attacks against Google Glass wearers in public have been reported in the media [1]. These scenarios highlight the high societal demand for *Privacy Enhancing Technologies (PETs)*. At the time of writing, no PETs are available on the market to communicate privacy preferences towards wearable cameras. In scientific literature, a handful of approaches has been published, however none

of them placed an emphasis on whether they are actually useful if deployed in particular situations where users are constrained in what artifacts they can carry or wear. The goal of this paper is to provide a first systematization of PETs that have been published in scholarly articles. To do so, we propose a collection of properties and criteria for categorization. The purpose of this paper is to start a discussion on both design and research directions in the fields of security, privacy and HCI in order to ensure that future PETs successfully counter privacy threats in the upcoming era of wearable computing. Additionally, our systematization provides a first suggestion how a standardized evaluation framework for (wearable) PETs could look like. In the course of our extensive literature review, we found approaches of how future PETs might look like. Most of them however, have substantial limitations such as that they address a narrow scenario, exclude particular user groups or cause further privacy challenges due to their privacy-violating functionality. These limitations may reduce the user's subjective satisfaction and introduce errors. Hence, they potentially have an impact on user experience.

2 Properties

The properties presented in this section haven been selected as a set we believe highlights important evaluation dimensions with respect to usability, in particular subjective user satisfaction, learnability, memorability, errors and efficiency.

User-Initiated: In order to mediate privacy-preferences, some PETs require a user to perform an action. Due to the unobtrusive recording capabilities of wearable cameras, users may be recorded by such a device without actually being aware of it. Therefore, user-triggered mediation hinders the consequent communication of a user's privacy preference. This is very likely to cause errors and misunderstandings which are very likely to have a negative impact on the overall user experience.

Location-Based: As determined by Denning et al. [4], privacy preferences can be determined by certain situations or locations. The definition of a privacy-sensitive space mostly varies from user to user, highly depending on their socio-cultural background. However, most users may require a device that works regardless of a certain location. Therefore, location-dependency may be a limiting factor. As a location-based approach usually requires the transmission of location information to other entities (e.g. a trusted server via a secure channel or another device in the surrounding), new privacy challenges are implied.

Face-Recognition-Based: In order to correlate the user of a PET with a user in an image or video taken by a wearable camera, facial recognition could be an efficient method of choice as state of the art algorithms provide sufficient accuracy to correctly identify individuals. Certainly, facial recognition requires the transmission of privacy-sensitive data to an associated service and therefore poses significant challenges to preserve the user's privacy.

Visual-Marker-Based: As widely used in augmented reality applications, person-tracking can efficiently be performed using visual markers. For PETs, this means that a user has to carry or wear one or multiple visual markers in order to communicate privacy policies towards wearable cameras. Markers can be designed in an obtrusive or unobtrusive way. In certain situations or cultural groups, subtle markers could be preferred over invasive ones and vice versa.

Gesture-Based: In augmented reality applications, gestures are another feasible approach to track individuals in videos. This approach is mostly limited to long or full-shot videos as it is obviously difficult to perform gesture-recognition on single images or smaller image selections. Furthermore, gestures have to be actively performed and therefore require a user who is aware of being filmed. It furthermore poses distinct accessibility challenges for people with physical disabilities and elderly users.

Signal-Emission-Based: As cameras are sensitive to a certain light spectrum, signal-emitting jamming could be used in PETs. Signal-emission based approaches are mostly expensive as they require a dedicated technical artifact which makes them significantly obtrusive.

Physical-Artifact-Required and/or Dedicated-Device-Required: This is an umbrella property for all approaches that require a dedicated physical artifact in order for the PET to function (e.g. an electronic device or visual markers). From the user's perspective, tangible interfaces as provided by such artifacts offer advantages and disadvantages. The main disadvantage is that the user has to carry or wear the artifact at any time. In some situations, this is unfeasible (e.g. in spaces where digital artifacts cannot be operated due to environmental constraints). In contrary, physical artifacts are often easier to understand for the user and give them a sense of control.

Requires-Trusted-Third-Party-Service: If communication with an external service is required, an Internet connection is indispensable. This might be unfeasible in some scenarios, where users are limited in what devices they can carry (e.g. at a beach)

Smartphone-Based: Smartphone-based approaches are easy to deploy since most individuals in today's society carry one with them all the time. However, situations where smartphones are not applicable but preserving an individual's privacy is required. An example of such a situation would be sunbathing and wearing only a bikini.

Visibility: Some PETs require physical artifacts in order to function. When deployed, some of them are highly visible to bystanders and therefore instantly disclose a certain privacy preference to nearby individuals. In some cultures or even particular situations, a more subtle and unobtrusive technology could be preferred by the user. However, others may want to use an obtrusive PET in order to disclose their privacy preference or policy openly when recorded by a wearer of Google Glass or a similar device. Low visibility however may constrain the communication of a cognitive model towards the user. It also implies a lack of feedback options.

Accessibility: As digital privacy affects all user groups likewise, a PET should work regardless of disabilities or other physical or cognitive conditions, such as low motor control or visual impairments.

Anonymity: PETs should not imply further privacy violations due to their functionality. Approaches that use facial recognition or location-tracking potentially violate the privacy of their users. Presumably, users of PETs are highly concerned about their privacy and potentially perceive privacy-violating PETs as paradoxical.

Impacts-User-Behavior: Some PETs heavily impact the user's behavior, as they either require a high effort in preparation or require the user to perform an action. PETs that require a user to be aware of being filmed also potentially influence the user's behavior.

Requires-Devices-To-Comply: This property indicates whether a PET can be deployed only if (wearable) cameras are updated accordingly (software and/or hardware).

3 Systematization of Privacy Enhancing Technologies

Table 1 presents our systematization in which we indicate if a certain PET has a certain attribute or not. If a PET could be configured in a way to evoke a certain property, we assume a best-case working scenario supposing that a poor implementation would make any concept potentially unusable. Obviously, some properties are disadvantageous for the overall user experience. For this systematization, we refrain from introducing a rating scheme and prefer a non-judgmental presentation. The reason for this is that neither we nor the authors of the respective PETs conducted user studies that would confirm these assumptions. While some of the PETs presented in this section have been particularly designed for the mobile/wearable computing domain, others were designed to preserve picture privacy in general. For the purpose of fostering a fruitful discourse however, we discuss some potentially impacting factors in a qualitative way.

The **Privacy Makeup** and hair-style approach as presented by Harvey et al. [5] exploits the weaknesses of commonly used face detection systems. To inhibit the feature response of face detection algorithms, significantly invasive distortions are created with camouflage makeup. This approach is time consuming in preparation and visually dominant. It therefore hinders everyday social interaction and can provoke unwanted reactions. It is only feasible when facial recognition algorithms are used or the makeup is applied in a way that its wearer is unrecognizable. For unexperienced users, it is hard to apply the makeup correctly. The **Respectful Cameras** approach as presented in [9] uses colored hats and scarfs as visual markers. Depending on whether an individual prefers to be made irrecognizable or not, the corresponding artifact is chosen and worn in front of a camera. The Picture-Privacy-Policy framework (**P3F**) as presented in [3] uses a similar approach, however the privacy policies used in this scheme are more complex and fine-grained. The visual markers of the respectful cameras approach [9] is based on a binary privacy policy and obtrusive markers. The P3F use not only ded-

Table 1. Systematization of PETs

	User-Initiated	Location-Based	Face-Recognition-Based	Visual-Marker-Based	Gesture-Based	Signal-Emission-Based	Dedicated-Device-Required	Physical-Artifact-Required	Communicates-With-Server	Internet-Connection-Required	Smartphone-Based	Visibility	Accessibility	Anonymity	Impacts-User-Behavior	Requires-Devices-To-Comply
Privacy Makeup [5]				•								•		•	•	
Respectful Cameras [9]				•				•				•	•	•		•
P3F [3]				•				•								•
OfflineTags [7]	•			•				•				•	•	•	•	•
Privacy Visor [11, 12]						•	•					•	•	•		
SnapMe [6]		•	•						•	•	•					•
FaceBlock [13]		•	•						•	•	•					•
BlindSpot [8]		•				•	•						•	•		•
Place Avoider [10]		•											•	•		•
Privacy Gestures [2]	•				•							•		•	•	•

icated accessories but aims at providing a clothing pattern database with fashionable clothing patterns that are then used as visual markers. A large-scale deployment as presented in the paper, however, would require all cameras or picture publishing platforms to use the P3F software to detect the visual markers and to deduct the privacy policies from them. Another visual-marker based approach is **Offlinetags** [7]. Offlinetags uses four different symbols readable by the open-source Offlinetags software. These symbols can simply be printed on a piece of paper and then presented to a camera. In contrary to the other visual-marker-based approaches presented in this section, the obtrusive markers must be presented actively towards a camera. Yamada et al. [11, 12] presented the **Privacy Visor**, i.e. glasses with infrared light sources that are visible to most camera sensors but invisible to the human eye. The goggles approach requires a constant power supply and infrared LEDs that can keep up with the ambient light. As most portable devices come with GPS sensors, location-based technologies such as the **SnapMe** privacy watchdog [6] or **Blind Spot** [8] are feasible to mediate privacy preferences. These approaches are based on correlated location information of a camera and its bystanders. Additionally to the location-reference, SnapMe proposes the use of facial recognition to identify individuals in pictures. In comparison to SnapMe, the Blind Spot approach is based on fixed cameras and intended for CCTV-like surveillance systems and thus limited to a specific location. **FaceBlock** [13] is based on biometric features on images taken by a (wearable) camera. Similar to the other facial-recognition-based approaches in this section, the FaceBlock system implies further privacy challenges, as privacy-sensitive biometric information is processed and transferred to a (trusted) server. Both FaceBlock and SnapMe provide a smartphone app where users can configure their privacy-settings. The **PlaceAvoider** [10] approach is not only intended to protect the privacy of bystander but also of the wearer of a wearable camera. Similar to BlindSpot, it provides blacklisting of privacy-sensitive spaces like bathrooms and bedrooms. Similar to other location-based approaches, it requires a predefined location and might therefore not be applicable in all desired situations.

Barhm et al. [2] presented a gesture-based method (**Privacy Gestures**) to communicate privacy preferences. Individuals perform defined gestures when recorded by a camera. Even though no additional artifact is required, its feasibility is limited to situations where an individual is aware of being recorded.

4 Conclusion and Work in Progress

As wearables such as Google Glass are very likely to capture private information of individuals recorded by these devices, PETs have become necessary to preserve our digital privacy. In this work, we provide an evaluation of PETs to communicate privacy preferences towards wearable cameras. At the time of writing, no such technology is available on the market. In this work, we assembled and systematized PETs that were mostly published at distinguished scientific conferences. We found that most of them are limited to certain pre-defined scenarios or exclude specific user groups: Smartphone-based approaches exclude smartphone abstainers who might refrain from using smartphones for privacy reasons. Some PETs require the collection and transmission of private information such as the location or biometric features and therefore imply further privacy challenges. The purpose of this work is to initiate a discourse between designers as well as security and usability experts and researchers and to pro-

vide the fundamentals for establishing a standard benchmark to evaluate conceptual PETs. Based on the results presented in this paper, we are currently conducting a comprehensive user study with qualitative interviews in the field. Our preliminary results suggest that privacy-aware potential users highly desire fully privacy-preserving tools. Furthermore, we found that some particularly unobtrusive PETs are hard for the user to understand and to use as PETs with low visibility do not sufficiently communicate cognitive models to the user.

Acknowledgements

We would like to thank the reviewers for their insightful comments. The research was funded by COMET K1, FFG - Austrian Research Promotion Agency and netidee.

References

1. Google Glass targeted as symbol by anti-tech crowd. `http://edition.cnn.com/2014/04/14/tech/mobile/google-glass-attack/`, accessed 10/7/2014.
2. M. S. Barhm, N. Qwasmi, F. Z. Qureshi, and K. El-Khatib. Negotiating privacy preferences in video surveillance systems. In *Modern Approaches in Applied Intelligence*, pages 511–521. Springer, 2011.
3. A. Dabrowski, E. R. Weippl, and I. Echizen. Framework based on privacy policy hiding for preventing unauthorized face image processing. In *Systems, Man, and Cybernetics (SMC), 2013 IEEE International Conference on*, pages 455–461. IEEE, 2013.
4. T. Denning, Z. Dehlawi, and T. Kohno. In situ with bystanders of augmented reality glasses: perspectives on recording and privacy-mediating technologies. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 2377–2386. ACM, 2014.
5. A. Harvey. CV Dazzle, 2010-2012. `http://cvdazzle.com/`, accessed 10/4/2014.
6. B. Henne, C. Szongott, and M. Smith. Snapme if you can: privacy threats of other peoples' geo-tagged media and what we can do about it. In *Proceedings of the sixth ACM conference on Security and privacy in wireless and mobile networks*, pages 95–106. ACM, 2013.
7. F. Pallas, M.-R. Ulbricht, L. Jaume-Palasí, and U. Höppner. Offlinetags: A novel privacy approach to online photo sharing. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '14, pages 2179–2184, New York, NY, USA, 2014. ACM.
8. S. N. Patel, J. W. Summet, and K. N. Truong. Blindspot: Creating capture-resistant spaces. In *Protecting Privacy in Video Surveillance*, pages 185–201. Springer, 2009.
9. J. Schiff, M. Meingast, D. K. Mulligan, S. Sastry, and K. Goldberg. Respectful cameras: Detecting visual markers in real-time to address privacy concerns. In *Protecting Privacy in Video Surveillance*, pages 65–89. Springer, 2009.
10. R. Templeman, M. Korayem, D. Crandall, and A. Kapadia. Placeavoider: Steering first-person cameras away from sensitive spaces. In *Network and Distributed System Security Symposium (NDSS)*, 2014.
11. T. Yamada, S. Gohshi, and I. Echizen. Use of invisible noise signals to prevent privacy invasion through face recognition from camera images. In *Proceedings of the 20th ACM international conference on Multimedia*, MM '12, pages 1315–1316, New York, NY, USA, 2012. ACM.
12. T. Yamada, S. Gohshi, and I. Echizen. Privacy Visor: Method for Preventing Face Image Detection by Using Differences in Human and Device Sensitivity. 2013. unpublished, under review for CMS 2013.
13. R. Yus, P. Pappachan, P. K. Das, E. Mena, A. Joshi, and T. Finin. Demo: Faceblock: privacy-aware pictures for google glass. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*, pages 366–366. ACM, 2014.