Contextual Privacy Ethics and Wearable Devices

Roberto Hoyle
School of Informatics
Indiana University
rjhoyle@indiana.edu

David Crandall
School of Informatics
Indiana University
djcran@indiana.edu

Qatrunnada Ismail
School of Informatics
Indiana University
qismail@indiana.edu

Apu Kapadia
School of Informatics
Indiana University
kapadia@indiana.edu

Luke Stark
Department of Sociology
Dartmouth College
luke.stark@dartmouth.edu

Denise Anthony
Department of Sociology
Dartmouth College
denise.anthony@dartmouth.edu

Abstract
The increasing ubiquity of digital wearable devices, including those which collect images, has prompted a variety of novel questions regarding both the ethics of privacy preservation in new social contexts, and the ethics of doing research within these contexts, that are not being addressed by device manufacturers. The use of these devices as intended generates a large amounts of data, and the curation of the data to filter out potentially sensitive information is not only burdensome and difficult, but fraught for consumers, designers, and scholars. Based on empirical studies of lifelogging "in the wild" and users' expressed privacy preferences around the images collected via these lifelogging practices, we explore several novel ethical questions for designers and policymakers around privacy research in its social context, wearables, and pervasive data collection.

Introduction
Over the past five years, wearable devices have become a growing sector of the social computing market and research area. Lifelogging, or the collection of masses of everyday data about one's own daily activities through means such as ubiquitous photography and pervasive metadata collection, is just one part of the expansion of these personal devices. One of the purposes behind the collection of this everyday pervasive data is for sharing.

Fitbits,[1] for example, have sharing options in which people can share their steps and other data with their friends. Internet of Things (IoT) devices, such as Withings smart weight scales,[2] are able to post a person's daily readings, motivating them to stick to their diet. Narrative clips[3] record a picture every 30 seconds with the intention of posting a stream of images that show what a person did during the day.

These novel technologies present new challenges around the ethics of personal privacy, both for users and for researchers. As this visual data collection becomes more and more widespread, its dissemination and impact across diverse social contexts and among a variety of social actors has begun to produce novel and unpredictable effects on the shared human experience of privacy. The problem across many scenarios involving wearable computing is not one of protecting the data by default, but of understanding how social and contextual norms might help users curate the large amount of generated data in order to protect the particular pieces that may need protection – a problem, in other words, of privacy in context. This question of curation, however, itself suggests new questions regarding the ethical responsibilities for designers, researchers, and users.

Our prior work [4] has shown that people prefer **in-situ controls** for protecting images over curation at the end of the day. This is most likely because the vast amount of data that gets collected by wearable devices is banal and mundane; most of us do not experience many sensitive moments during the day. It is much harder to sift through all the data at the end of the day and recall when the sensitive moments were. Participants found it much easier to exert physical control over the device, either before a sensitive moment happened, or just afterwards by deleting data that had just been collected. We suggest that these kinds of studies in which we look at people's actual behavior are vital for designing systems that will protect people's privacy by understanding the contextual social aspects of their behavior. These studies allow us to adapt insights regarding the social dynamics of privacy around visual data into ethical guidance around the potential public uses and impacts of novel digital photographic technologies. Specifically, we draw on the following themes to note emergent privacy ethics problems:

- **Curation Assistance** – Wearable computing devices need to incorporate tools to assist people in contextual curation - but who sets the defaults? How can researchers help shape these norms responsibly?
- **Concern for Bystanders** – Lifeloggers show concern about the privacy of others captured in lifelogging images, so-called bystanders - but how to balance the ethical responsibilities in these contextual encounters?
- **Content and Context of Data** – Both the context and the content of the collected data contains information that can guide privacy decisions - how can this information be incorporated into system design, and what special risks do privacy researchers face in handling this data?
- **Heuristics** – What common social elements can we recognize to use as rules assisting people's privacy around wearable devices?

[1] http://www.fitbit.com
[2] http://www.withings.com
[3] http://getnarrative.com

Curation Assistance

Wearable devices that are currently on the market do not assist people well in the contextual aspects of data curation. Some, like Google Glass [3] do not allow for physical *in situ* control; they cannot simply be put into a pocket. A person may not want to put a Glass in a pocket if it is their prescription eyewear, nor may they be able to turn off collection from an implanted device. Narrative Clip allows for physical control, but the end-of-day curation tools are very inefficient, specially when one takes into account the data collection rate of the device. In our study, we found that users overwhelmingly preferred to exert in-context physical control over the device, such as turning it off, putting it upside down, storing it in a pocket, etc. over going over the set of data at the end of the day and filtering for sensitivity. Most devices do not have usable *in situ* controls on them, despite their popularity. Narrative has no "off" switch, for example, and wearable camera devices do not tend to have a mechanism for physically blocking the camera to assist with bystander concerns. Given the disconnect between design affordances and user behavior, how can companies be convinced to change the default design affordances of their products? How can user hacks or modifications of these devices be supported by both legal and design communities? And what responsibilities do researchers have regarding advocacy for such design changes?

Concern for Bystanders

One of the findings that our prior work discovered is that lifeloggers have significant concern for the privacy of bystanders – what Irving Goffmann termed "civil inattention," or purposefully ignoring useful information about others for the benefit of social cohesion. This concern might be leveraged in a privacy by design system by doing recognition of data captured from bystanders, and not sharing it. A wearable camera system could, for example, run a face detection algorithm on an image, and if there were no faces on it, could mark the image as low-sensitive, while trying to match faces that are detected against a known "white list" of people whose images could be shared by default. However, ethical problems remain: should bystanders be informed, either in situ or after the fact, that their image has been recorded? Do lifeloggers have either ethical or legal obligations to provide visual data to authorities such as law enforcement? And will a world of wired and connected devices create situations in which different individual devices conflict around desired pre-set technical defaults? How can we build models of "civil inattention" into the design of these devices? And when we conduct research on "civil" inattention, how do we protect the bystanders involved?

Content and Context of Data

Wearable cameras have usually used only contextual information and metadata in determining whether an image is sensitive or not because this information is relatively easy to determine based on prior behavior. A system can be imagined that has a "work", "home", and "bar" setting that are automatically detected based on a person's location, and with pre-determined sharing profiles that are tailored to the person for these settings. This ignores, however, the content of the data, which can be vital in determining whether information should or should not be shared [7, 1]. In pictures, for example, the content of the images often contains vital information about its sensitivity. Content is unfortunately difficult to detect, but there are some simple things that seem to indicate sensitivity that may be easier to detect. The presence of computer screens, for example, seems to be indicative of sensitivity. Similarly, the more people present in an image, the less sensitive an image seems to be [4]. How might

these contextual insights be worked into the technical defaults underpinning wearable devices? And how much should developers and academic researchers limit their own data collection in the service of privacy?

Heuristics

People are better able to manage their privacy when their options are presented in simpler, contextually appropriate ways, rather than with all of the options available to them to choose from [2]. To facilitate this, the concept of heuristics is extremely important. These simple rules can be used to indicate whether data is normal or abnormal, and can be used to assist in curation. Location sharing systems have started looking into this question [5, 8] by generating rules from a person's location traces. Recently, computer vision mechanisms have been developed that can recognize rooms, marking some of them as private areas that should not be photographed [9, 6]. This kind of work needs to be extended to consider the ethical norms behind these human behaviors, and explore what dynamic rules can be created from awareness of what a person normally shares in their everyday social context.

Conclusion

Given that the volume of data produced by ubiquitous cameras and other forms of pervasive tracking, including those used in lifelogging, is often so large as to be difficult to curate manually, it is vital that privacy scholars develop and recommend contextually-sensitive best practices both for research methods, technical privacy protection and regulatory engagement around data collected from wearable devices. As these devices become more automated, they get better at collecting data about us. As such, policymakers and ordinary citizens need to grapple with the privacy ethics implications of visual data's collection, dissemination, and use. While privacy-protecting tools and algorithms are an important part of the sociotechnical response to increases in visual data, they should not and cannot be separated from robust ethical and normative responses by policymakers, designers, and the public.

Acknowledgements

This material is based upon work supported by the National Science Foundation under grants CNS-1016603, CNS-1252697, CNS-1408730, and IIS-1253549, the Office of the Vice Provost for Research at Indiana University Bloomington through the Faculty Research Support Program, and a Google Faculty Research Award. Ismail is funded by the College of Computer and Information Sciences in King Saud University, Saudi Arabia. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[1] Anthony, D., Henderson, T., and Kotz, D. Privacy in location aware computing environments. *IEEE Pervasive 6*, 4 (Oct–Dec 2007), 64–72.

[2] Egelman, S., Oates, A., and Krishnamurthi, S. Oops, i did it again: Mitigating repeated access control errors on facebook. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, ACM (New York, NY, USA, 2011), 2295–2304.

[3] *Google Glass*. http://glass.google.com.

[4] Hoyle, R., Templeman, R., Armes, S., Anthony, D., Crandall, D., and Kapadia, A. Privacy behaviors of lifeloggers using wearable cameras. In *ACM Int'l Joint Conference on Pervasive and Ubiquitous Computing* (2014), 571–582.

[5] Kelley, P. G., Hankes Drielsma, P., Sadeh, N., and

Cranor, L. User-controllable learning of security and privacy policies. In *the 1st ACM workshop*, ACM Press (New York, New York, USA, 2008), 11–8.

[6] Korayem, M., Templeman, R., Chen, D., Crandall, D., and Kapadia, A. Screenavoider: Protecting computer screens from ubiquitous cameras. In *CoRR arXiv Technical Report arXiv:1412.0008* (2014).

[7] Prasad, A., Sorber, J., Stablein, T., Anthony, D., and Kotz, D. Understanding sharing preferences and behavior for mHealth devices. In *WPES '12: Proceedings of the 2012 ACM workshop on Privacy in the electronic society*, ACM Request Permissions (New York, New York, USA, Oct. 2012), 117–128.

[8] Ravichandran, R., Benisch, M., Kelley, P. G., and Sadeh, N. M. Capturing social networking privacy preferences.

[9] Templeman, R., Korayem, M., Crandall, D., and Kapadia, A. PlaceAvoider: Steering First-Person Cameras away from Sensitive Spaces. In *21st Annual Network and Distributed System Security Symposium (NDSS)* (2014).