



Elagroudy, P., Khamis, M. , Mathis, F., Irmischer, D., Bulling, A. and Schmidt, A. (2019) Can Privacy-Aware Lifelogs Alter Our Memories? In: 37th Annual ACM CHI Conference on Human Factors in Computing Systems (CHI '19 Extended Abstracts), Glasgow, UK, 04-09 May 2019, LBW0244. ISBN 9781450359719 (doi:[10.1145/3290607.3313052](https://doi.org/10.1145/3290607.3313052)).

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© The Authors 2019. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19), Glasgow, Scotland, UK, 04-09 May 2019, Article LBW0244. ISBN 9781450359719.
<https://doi.org/10.1145/3290607.3313052>.

<http://eprints.gla.ac.uk/181982/>

Deposited on: 15 March 2019

Can Privacy-Aware Lifelogs Alter Our Memories?

Passant Elagroudy
University of Stuttgart
Stuttgart, Germany
passant.el.agroudy@vis.uni-stuttgart.de

Florian Mathis
LMU Munich
Munich, Germany
Florian.Mathis@campus.lmu.de

Andreas Bulling
University of Stuttgart
Stuttgart, Germany
Andreas.Bulling@vis.uni-stuttgart.de

Mohamed Khamis
University of Glasgow
Glasgow, United Kingdom
Mohamed.Khamis@glasgow.ac.uk

Diana Irmischer
LMU Munich
Munich, Germany
d.irmscher@campus.lmu.de

Albrecht Schmidt
LMU Munich
Munich, Germany
albrecht.schmidt@ifi.lmu.de


ABSTRACT

The abundance of automatically-triggered lifelogging cameras is a privacy threat to bystanders. Countering this by deleting photos limits relevant memory cues and the informative content of lifelogs. An alternative is to obfuscate bystanders, but it is not clear how this impacts the lifelogger's recall of memories. We report on a study in which we compare viewing 1) unaltered photos, 2) photos with blurred people, and 3) a subset of the photos after deleting private ones, on memory recall. Findings show that obfuscated content helps users recall a lot of content, but it also results in recalling

CHI'19 Extended Abstracts, May 4–9, 2019, Glasgow, Scotland Uk

© 2019 Copyright held by the owner/author(s).

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI'19 Extended Abstracts), May 4–9, 2019, Glasgow, Scotland Uk*, <https://doi.org/10.1145/3290607.3313052>.



Sidebar 1: Original photo of three participants (top) vs. obfuscated version (bottom) where bodies are blurred (gaussian blur, radius = 40 px). We found that obfuscated lifelogs allow viewers to remember more details but with less accuracy.

“**Environmental lifelogging**” in this context is the recording of the environment using infrastructure cameras in the space capturing third-person views of the participants. For further examples, refer to [5].

Sidebar 2: Definition of “Environmental lifelogging”.

less accurate details, which can sometimes mislead the user. Our work informs the design of privacy-aware lifelogging systems that maximizes recall and steers discussion about ubiquitous technologies that could alter human memories.

KEYWORDS

Privacy; Filters; Obfuscation; Lifelogging; Blurring; Memory augmentation; Recall

ACM Reference Format:

Passant Elagroudy, Mohamed Khamis, Florian Mathis, Diana Irmscher, Andreas Bulling, and Albrecht Schmidt. 2019. Can Privacy-Aware Lifelogs Alter Our Memories?. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI’19 Extended Abstracts)*, May 4–9, 2019, Glasgow, Scotland Uk. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3290607.3313052>

INTRODUCTION AND BACKGROUND

Lifelogging can support reminiscence, reflection, and searching past moments [3]. Images promote more detail-rich recall compared to other types of data [9] as they contain rich contextual information [11]. Thus, pictorial lifelogs were extensively researched as a means to augment human memory (e.g., [1, 6, 15]). However, a key challenge to adopt a pictorial lifelogging service is properly addressing privacy concerns (e.g., [10]) as they expose the activities of individuals surrounding the wearer of the camera more than they expose the lifelogger themselves. This is particularly true for environmental lifelogging [5]. Previous work counteracted this problem by 1) deleting photos: this can be done post hoc or by preventing the capture of photos if privacy-sensitive situations are detected [16], or by 2) obfuscating the individuals, for example by inpainting them and replacing them by avatars [13]. The second approach is particularly promising as it protects the individuals’ privacy while maintaining good utility and user experience [13].

Although previous works (e.g., [7, 13]) investigated the impact of obfuscation on privacy protection and user experience, it remains unclear *how obfuscation impacts the recall of memories*. Closing this gap is crucial as recalling memories is one of the main motivations behind lifelogging. Therefore, we report on a two-stage pilot user study (N=12) in which we compare the impact of obfuscation using body blurring vs. the deletion on the recalled memories from pictorial lifelogs (we will hereafter refer to them as lifelogs for simplicity). Participants first took part in an eventful interaction session and then returned after 4–5 days to recall memories when viewing 1) 20 unaltered photos (baseline), 2) obfuscated versions of the 20 photos where persons are blurred, and 3) five of the 20 original photos after deleting private ones.

Our results suggest that blurred photos enabled users to remember more details than deleted photos. However, blurred photos degraded the accuracy of the recalled details in comparison to deletion of the photos in some cases. This implies that privacy-preserving obfuscated photos using

C1 (baseline) Participants received 20 original photos.

C2 (obfuscation) Participants received an obfuscated version of photos in (C1), where all persons were blurred (see sidebar 1). We used *body blurring* because of positive results in prior work [2, 8, 12] and its wide adoption in research and industry (e.g. Google Maps).

C3 (deletion) Participants received only a subset of five original photos from (C1), mimicking deletion for privacy protection. We equally sampled them from C1 (*baseline*) dataset across time to avoid biases in deleting particular participants.

Sidebar 3: The three conditions used to evaluate the impact of the privacy-protection method on information recall from lifelogs.

Personal questions (e.g., “Every participant introduced themselves at the beginning by giving a short talk. Can you remember the hair color of Anna?”)

Procedural questions (e.g., “Please tell us the procedure of the first session as specific as possible.”)

Game questions (e.g., “How many times did the players in your group have to swap their playing pieces?”).

Sidebar 4: Memory questions administered in session 2

body blurring are good for remembering more about a forgotten topic but not for recalling accurate details. As such, our work sheds light on a trade-off between protecting the privacy in lifelogs and undermining the lifelogs’ potential as a memory prosthetic. Additionally, it is a first milestone to spark a discussion about potential use cases for altering memories via ubiquitous technologies.

METHODOLOGY

Our pilot study was composed of *two lab sessions*. The *first session* (Session 1) was to create an environmental lifelog of all participants in a common eventful interaction to evaluate their recall in the next session. The *second session* (Session 2) was to compare the impact of applying the privacy-protection methods to the common lifelogs on the recall quality of the events.


Session 1: Building the lifelog in a controlled event

Twelve participants were invited to Session 1 that lasted approximately 90 minutes. The session was photographed using two cameras in the room from different angles to create an environmental lifelog (third-person view). Participants were informed about the recording, their consent and demographic data was collected. We introduced the participants to the domain of lifelogging and obfuscation of photos creating an environment of open discussions. Afterwards, participants were split into random teams of four and played a locally popular board game. The game’s goal is to move your playing pieces as fast as possible into the safe zone using your dice score. We modified the game rules and asked participants to physically swap seats upon getting certain dice values. The objective was to make the lifelogs recorded during gameplay more dynamic and to reduce potential bias and confusion from having most of the lifelogs seemingly static.

Dataset collection. In the *first session*, we collected over 450 photos. We used fixed temporal sampling to select the presented photos (memory cues). We sampled at five-minute intervals during the introduction and discussions. However, we reduced the interval to three minutes during the game part as it lasted for a shorter period of time (about 20 minutes). Each participant appeared at least once in their experimental lifelog dataset.

Session 2: Evaluating the recall

Session 2 took place four to five days after session 1 to ensure a realistic decay of information in the memory [11], lasting for approximately 90 minutes. We re-invited the participants to individually review the lifelogs from session 1. To reduce potential bias due to learning effects about the lifelog’s content, we therefore opted for a between-subjects experiment design. We covered one independent variable, *the privacy-protection method*, with three conditions (see Sidebar 3) and measured their impact on information recall.



Sidebar 5: Overview of the correctness level for each condition (C1-C3). (A) Obfuscation of persons was associated with the highest number of semi-correct overall answers, while deletion of photos was associated with the highest number of wrong answers. (B) However, reviewing the answers using obfuscated photos (C2) resulted in the lowest ratio of correct answers and the highest ratio of semi-correct and wrong answers.

Each team of four was assigned to a condition. Our aim was to measure participants' recall and if they perceived the photos as helpful for remembering details from session 1. Thus, participants filled in a questionnaire where they answered 30 questions about details that happened in Session 1 (see sample questions in Sidebar 4). For each condition, we asked participants to answer the questionnaire: 1) before viewing the memory cues (pre-questionnaire) and 2) after viewing them, i.e., photos of the respective condition (post-questionnaire). This was done to account for prior knowledge of the answers and to identify any improvements resulting from having seen the memory cues. Participants were allowed to navigate through the photos as long as they wanted. They were also allowed to improve their answers to the pre-questionnaire when filling the post-questionnaire. On 5-point Likert scales, participants estimated the helpfulness of the photos in aiding recall during the *post-questionnaire*.

Participants and Recruitment


We recruited 12 participants (3 females) via university mailing lists aged between 20 and 32 years old (mean=24.2 years, SD=3.72). Participants were compensated with an e-shop voucher worth 20€. To motivate participants to put an effort in their guesses, we also arranged a raffle for an additional voucher, where participants who performed the best in session 2 had the highest chance in winning.

RESULTS: HOW DID THE PRIVACY-PROTECTION METHOD IMPACT THE RECALL?

Inspired by the methodology of Le et al. [11], two researchers labeled each answer to the memory questionnaire's questions as: 1) *correct* when the answer is consistent with the expected one, 2) *semi-correct* when parts of the answer are correct (e.g., remembering the outfit of a participant instead of his name) and 3) *wrong* when the answer is incorrect. We also computed a custom metric, which we refer to as the *Recall Correctness Score (RCS)*. The correctness labels were scored as follows: two points for correct answers, one point for semi-correct answers and zero points for wrong answers. The RCS is the summation of the weights per participant.

Correctness of recalled memories

We compared the distribution of correctness labels across the conditions accounting for all answers in the *post-questionnaire* (see Sidebar 5A). Our preliminary results suggest that participants in C2 (*obfuscation*) had more *semi-correct* answers compared to C3 (*deletion*) (34.17% and 22.5% respectively). C3 (*deletion*) resulted in more *wrong* answers compared to C2 (*obfuscation*) (33.33% and 23.33% respectively). The trend persisted if we only evaluate the answers that were reviewed and updated. All conditions had a comparable ratio of *correct* answers. We compared the average RCS scores between the *pre-questionnaire* (without photos) and the *post-questionnaire* (with photos) for each condition (see Sidebar 6). The RCS increased after reviewing the photos in all conditions by 29.46% for C1 (*baseline*), followed by 14.66% in C3 (*deletion*), then 12.6% in C2 (*obfuscation*).



Sidebar 6: The figure shows a comparison between the trends of the *Recall Correctness Score (RCS)* before viewing photos, and after viewing photos. The RCS is directly proportional to the quantity and the quality of the photos used in reviewing the answers.

Improvements after viewing the photos

We compared the distribution of the correctness labels across the conditions using the *changed answers* only from the *post-questionnaire* (see Sidebar 5B). A changed answer corresponds to adding extra information to the answer or completely changing it irrespective of its correctness label. C1 (*baseline*) and C3 (*deletion*) had a similar ratio of reviewed *correct* answers (46% and 44.83% of the reviewed answers respectively). However, C2 (*obfuscation*) resulted in an inferior ratio of *correct* answers compared to other conditions (31.58%). Nevertheless, C2 (*obfuscation*) resulted in the highest *semi-correct* reviewed answers (36.84%) compared to C1 (*baseline*) (32%) and C3 (*deletion*) (27.59%). C2 (*obfuscation*) also resulted in the highest ratio of *wrong* answers (31.58%) while C1 (*baseline*) had the least ratio (22%). This suggests that changes made after viewing C2 (*obfuscation*) photos are the least correct. However, it encouraged semi-correct and wrong changes.

LESSON LEARNT: AMBIGUOUS LIFELOGS MIGHT DISTORT MEMORIES

In contrast to C2 (*obfuscation*), participants updated their answers to the *post-questionnaire* in the C1 (*baseline*) and C3 (*deletion*) conditions when they remembered the information on their own or saw it in one of the photos, leading to correct answers. However, the number of correct answers in C2 (*obfuscation*) drops compared to the other conditions because it involves the possibility of guessing or wrongful cueing due to participants seeing distorted photos. We believe this is also magnified because of the *retrieval induced forgetting* phenomenon [4], where remembering an unintended piece of information leads to inhibiting the recall of another requested one. This also aligns with Schachter's work [14] about two common memory sins: *mis-attribution*, i.e., the tendency to confuse the source of a memory with another, and *the suggestibility*, i.e., the tendency to mix false suggestions made by others with the original memory. Thus, our pilot study results suggest that blurring is good for remembering more about a forgotten topic, but not for recalling accurate details of memories.

Conclusion and Future Work. In this work we reported on a pilot study where we investigated the impact of privacy-aware obfuscation in lifelogging on recall of memories. Participants viewed unaltered photos, photos with blurred people, and a subset of the photos after deleting private ones. We found that blurring improves quantity of recalled details, but does not improve the accuracy of recall. Although blurring for obfuscation is commonly used in research [7, 13] and industry (e.g., Google Street View), we cannot make claims about the impact of other techniques on memory recall. We plan to compare in the impact of obfuscation in explicit photo capturing contexts such as taking selfies. We also plan to further investigate: 1) mechanisms to utilize this effect in supporting useful memory alterations in cases like blocking memories for patients with post traumatic stress disorder or reviving memories for patients with temporary memory loss and 2) protecting users against malicious

ACKNOWLEDGMENTS



This work is partially conducted within the Amplify project (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 683008).

use-cases such as implanting fake memories. Exploring such issues brings lifelogging one step closer to real life applications as a viable alternative for externalizing memory prosthetics.

REFERENCES

- [1] Emma Berry, Narinder Kapur, Lyndsay Williams, Steve Hodges, Peter Watson, Gavin Smyth, James Srinivasan, Reg Smith, Barbara Wilson, and Ken Wood. 2007. The use of a wearable camera, SenseCam, as a pictorial diary to improve autobiographical memory in a patient with limbic encephalitis: A preliminary report. *Neuropsychological Rehabilitation* 17, 4-5 (2007), 582–601. <https://doi.org/10.1080/09602010601029780> arXiv:<https://doi.org/10.1080/09602010601029780>
- [2] Andrew Besmer and Heather Lipford. 2009. Tagged Photos: Concerns, Perceptions, and Protections. In *Proc. CHI EA'09*. ACM, New York, NY, USA, 4585–4590. <https://doi.org/10.1145/1520340.1520704>
- [3] Eun Kyoung Choe, Nicole B. Lee, Bongshin Lee, Wanda Pratt, and Julie A. Kientz. 2014. Understanding Quantified-selves' Practices in Collecting and Exploring Personal Data. In *Proc. CHI'14*. ACM, New York, NY, USA, 1143–1152. <https://doi.org/10.1145/2556288.2557372>
- [4] Michael A. Ciranni and Arthur P. Shimamura. 1999. Retrieval-induced forgetting in episodic memory. *Journal of experimental psychology. Learning, memory, and cognition* 25 6 (1999), 1403–14.
- [5] Sarah Clinch, Paul Metzger, and Nigel Davies. 2014. Lifelogging for 'Observer' View Memories: An Infrastructure Approach. In *Adjunct Proc. UbiComp'14*. ACM, New York, NY, USA, 1397–1404. <https://doi.org/10.1145/2638728.2641721>
- [6] Cathal Gurrin, Alan F. Smeaton, and Aiden R. Doherty. 2014. LifeLogging: Personal Big Data. *Foundations and Trends in Information Retrieval* 8, 1 (2014), 1–125. <https://doi.org/10.1561/15000000033>
- [7] Rakibul Hasan, Eman Hassan, Yifang Li, Kelly Caine, David J. Crandall, Roberto Hoyle, and Apu Kapadia. 2018. Viewer Experience of Obscuring Scene Elements in Photos to Enhance Privacy. In *Proc. CHI'18*. ACM, New York, NY, USA, Article 47, 13 pages. <https://doi.org/10.1145/3173574.3173621>
- [8] Panagiotis Ilia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. 2015. Face/Off: Preventing Privacy Leakage From Photos in Social Networks. In *Proc. CCS'15*. ACM, 12. <https://doi.org/10.1145/2810103.2813603>
- [9] Vaiva Kalnikaite, Abigail Sellen, Steve Whittaker, and David Kirk. 2010. Now Let Me See Where I Was: Understanding How Lifelogs Mediate Memory. In *Proc. CHI'10*. ACM, New York, NY, USA, 2045–2054. <https://doi.org/10.1145/1753326.1753638>
- [10] Marion Koelle, Matthias Kranz, and Andreas Möller. 2015. Don't Look at Me That Way!: Understanding User Attitudes Towards Data Glasses Usage. In *Proc. MobileHCI'15*. ACM, New York, NY, USA, 11. <https://doi.org/10.1145/2785830.2785842>
- [11] Huy Viet Le, Sarah Clinch, Corina Sas, Tilman Dingler, Niels Henze, and Nigel Davies. 2016. Impact of Video Summary Viewing on Episodic Memory Recall: Design Guidelines for Video Summarizations. In *Proc. CHI'16*. ACM, New York, NY, USA, 4793–4805. <https://doi.org/10.1145/2858036.2858413>
- [12] Yifang Li, Nishant Vishwamitra, Bart P. Knijnenburg, Hongxin Hu, and Kelly Caine. 2017. Blur vs. Block: Investigating the Effectiveness of Privacy-Enhancing Obfuscation for Images. In *Proc. CVPRW'17*. <https://doi.org/10.1109/CVPRW.2017.176>
- [13] Yifang Li, Nishant Vishwamitra, Bart P. Knijnenburg, Hongxin Hu, and Kelly Caine. 2017. Effectiveness and Users' Experience of Obfuscation As a Privacy-Enhancing Technology for Sharing Photos. *PACM HCI* 1, CSCW, Article 67 (Dec. 2017), 24 pages. <https://doi.org/10.1145/3134702>
- [14] Daniel L. Schachter. 2002. *The seven sins of memory*. Boston: Houghton Mifflin.
- [15] Abigail J. Sellen, Andrew Fogg, Mike Aitken, Steve Hodges, Carsten Rother, and Ken Wood. 2007. Do Life-logging Technologies Support Memory for the Past?: An Experimental Study Using Sensecam. In *Proc. CHI'07*. ACM, New York, NY, USA, 81–90. <https://doi.org/10.1145/1240624.1240636>
- [16] Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2018. PrivacEye: Privacy-Preserving First-Person Vision Using Image Features and Eye Movement Analysis. (2018). arXiv:1801.04457 <http://arxiv.org/abs/1801.04457>