Reducing Privacy Risks in the Context of Sharing Photos Online

Rakibul Hasan
Indiana University
Bloomington, IN, USA
rakhasan@iu.edu

CHI '20 Extended Abstracts, April 25–30, 2020, Honolulu, HI, USA.
 ISBN 978-1-4503-6819-3/20/04 ...$15.00.
https://doi.org/10.1145/3334480.3375040.

Abstract
Through my doctoral research, I aim to reduce privacy risks in the context of photo-sharing online by developing tools and techniques that are both effective (in minimizing privacy risks) and usable. In solving this problem, I am taking a holistic socio-technical approach and proposing mechanisms to lessen the privacy risks of both the photo-sharers (or owners) and other people captured (sometimes unintentionally) in their photographs. More specifically, my goal is to i) design image transforms to *effectively obfuscate* privacy-sensitive content while *preserving utility* for human viewers, ii) develop techniques to automatically detect scene elements that may threaten privacy of people appearing photographs, and, iii) design behavioral interventions to persuade people to be more protective of their own and others' privacy. With my research, I hope to contribute to promoting privacy-protective behaviors and making online space more privacy-friendly.

Author Keywords
Privacy; Social media image sharing; Computer vision; Machine learning; Psychology; Decision making

CCS Concepts
•Security and privacy → **Social aspects of security and privacy(HCI);** *Privacy protections;* **Usability in security and privacy; •Human-centered computing** → *User stud-*

ies; Please use the 2012 Classifiers and see this link to embed them in the text: https://dl.acm.org/ccs/ccs_flat.cfm

Introduction and Motivation

Photos capture memorable life-events, and sharing them with others provides a natural mechanism for people to express themselves and interact with one another [20]. With the popularization of online social networks (OSNs) in the past few years, the volume of photo-sharing activity has dramatically increased [11, 28]. Such (re-)sharing has, in turn, led to a rise in (sometimes accidental) privacy violations, e.g., by revealing an individual's identity, location, activity, and so on, which may harm their impression management, subject them to surveillance or targeted advertising, and threaten their physical and property security [4, 45, 32, 1, 39]. As photo-sharing platforms are now becoming more popular than traditional social networks [19, 7] both the number of photos shared and the number of viewers of such photos will continue to increase, resulting in a higher rate of privacy violations.

Prior research has identified two strategies for reducing privacy risks – 1)limiting access to a shared item, and 2) modifying the item before sharing to remove sensitive content. Platforms like Facebook offer privacy settings that allow users to restrict unintended audiences from accessing a shared item. Such mechanisms, however, fall short in preventing re-sharing of photos in the same or other platforms that allow a larger audience to access the content. Furthermore, privacy violations may happen even when only the intended group of people can view the shared item, e.g., when some properties of the object or people captured in a photo are privacy-sensitive, rather than the subject matter of the photo or the identities of the individuals. Indeed, prior research has identified objects and properties that people consider as privacy-sensitive, even when viewed by only the intended audience [18, 36, 24, 9, 4, 2]. Access restriction mechanisms also fail to prevent privacy violations when multiple people co-own a photo and they have different sharing preferences (e.g., public vs. friend-only) [40]. Obscuring sensitive regions in a photo before sharing it might be a viable solution to this problem, and prior research has proposed image-filters to do that [16, 29, 48]. Since managing self-impression by sharing photos of high visual-aesthetic [35], and gathering information by observing other people's photos [43] are important motivations for using OSNs, an additional challenge is to retain enough utility (such as visual aesthetics and information content) for the *human* viewers in the obfuscated images. Significantly lowering photo utility may also discourage adopting these privacy-enhancing techniques (PETs). **Thus, one goal of my dissertation research is to design image obfuscation mechanisms that would effectively obscure sensitive content while preserving utility for human viewers.**

Besides utilizing access control mechanisms provided by OSNs, to reduce privacy risks, people engage in self-censoring behaviors ranging from restricting their sharing to withdrawing themselves from OSNs altogether [37, 46, 38]. People also exercise control offline to avoid sharing co-owned photos with undesired audiences [40] or photos captured and shared by other people [30]. On one hand, abandoning OSNs prevents people from receiving the social benefits they offer [43, 13, 22, 50], on the other hand, it does not provide adequate protection from privacy risks. Photos taken in public places inadvertently capture *bystanders* (i.e., people who were captured unintentionally and are not essential to the photo) and when these photos are shared online can pose privacy risks for the bystanders. Technologies have been proposed that allow bystanders to communicate their privacy preferences with the photographers [6, 34, 3, 33, 49] (e.g., using a smartphone app that

broadcasts preferences using Bluetooth). Unfortunately, this approach promotes the notion of 'by-default-public', since it requires the bystanders, who are the victims of privacy violations, to be proactive to keep their data private. Moreover, broadcasting privacy preferences publicly (e.g., using visual markers [6] or hand gestures [34]) in itself might be a privacy violation. Finally, most of these tools require a bystander's device to share sensitive data (such as facial features [3, 49] and location [33]) so that the photographer's device can identify them and apply the intended privacy policy. **The second goal of my doctoral research is to detect *bystanders* in photos automatically using computer vision and machine learning**. This approach has the potential to enforce a privacy-by-default policy, in which identities and other sensitive attributes of bystanders can be protected (e.g., by obscuring them) without requiring explicit action or sharing any sensitive data.

Table 1: Privacy filters (Blur, Edge, and Pixel had three levels of strength, resulting in 11 filters in total)

- Blur
- Edge
- Pixel
- Silhouette
- Mask

Most often than not, there is a trade-off between the usability of a system (e.g., online social platforms) and ensuring security and privacy in that system. In the context of photo-sharing online, using privacy-enhancing tools (e.g., image filters) may hamper the users' primary goal of the sharing act (e.g., by reducing utility for the viewers and hence discouraging social interactions). How can we persuade OSN users to adopt privacy-enhancing tools, especially when other peoples' (e.g., bystanders) privacy, rather than their own, is at stake? **The third goal of my dissertation research is to design effective textual and visual interventions to persuade people to adopt PETs and be more protective of their own and others' privacy.**

Research Projects and Current Progress

Obfuscating sensitive contents in images

The goal of this project is to identify and/or design image filters that would *effectively* obscure privacy-sensitive attributes of people (e.g., facial expression) and other objects (e.g., text on an electronic screen) while preserving enough utility (e.g., visual aesthetics) for the viewers. In a study, we applied eleven filters (Table 1) on twenty privacy-sensitive attributes (Table 2) and showed the obfuscated images to the participants. We measured five dependent variables – if the obfuscated content could still be identified, identification confident of the participants, and three utility variables: perceived information content [31], perceived visual aesthetic, and overall satisfaction as a viewer [10]. While we identified many filters that were effective in obscuring the intended attribute, our findings also identified a clear trade-off between privacy and utility – the more effective a filter was, the more aggressively it reduced utility for the viewers. These findings were published at CHI'2018 [14].

The previous experiment [14] revealed a negative association between the effectiveness of a filter to protect the privacy and how much it preserves utility. To design filters that can balance both privacy and utility, however, it is important to understand interactions among the utility variables themselves. We conducted path model analyses on the data from previous study [14] to investigate how the three utility variables affect one another and to quantify the direct and indirect impact of the filters on 'viewer satisfaction' [10], the dependent variable we want to ultimately maximize. We found that the indirect (negative) effect of the filters on satisfaction via information-content and visual-aesthetic overwhelmed the (negative) direct impact. This led to the idea that, if information-content or visual-aesthetic of images with privacy-enhancing filters applied can be improved, that may lead to an eventual increase in viewers' satisfaction. Since obfuscating image regions with filters will inevitably reduce information-content, we focused on improving the visual aesthetics of the remaining parts of the images using artistic image transforms. We selected three artistic

Table 2: Privacy sensitive attributes

- Activity
- Age
- Dress
- Ethnicity
- Facial expression
- Gender
- Hair
- Computer monitor
- Computer application
- Text on monitor
- Document class
- Text on a document
- Document type
- Indoor
- Specific indoor space
- Outdoor
- Specific outdoor space
- Laundry
- Food

transforms that are well-studied in computer vision literature [12, 42, 47], and applied them on the privacy-enhanced images to beautify them. Using these images as stimuli, we conducted a new study similar to the prior one [14] and measured the same three utility variables. We found that, in some cases, the visual-aesthetic improved, but the transforms failed to overcome the destructive effects of the filters and improve satisfaction. These results were published at CHI'2019 [15].

Detecting sensitive content in photos

A large body of prior research is dedicated on reducing the privacy risks of the image-owners and their friends, e.g., by detecting sensitive objects in photos [41, 21] and providing an overall 'privacy-sensitivity' ratings [27] using computer vision and machine learning. My research is complementary to those efforts, where I attempt to detect *bystanders* in images automatically so that measures can be taken to protect their privacy (e.g., by obscuring their identity). This is a challenging problem because subjects and bystanders in a photo may not always have very distinctive visual characteristics. This categorization also depends on the interactions among people appearing in the photo as well as the context and the environment in which the photo was taken. We approach this problem by conducting a user study to investigate how humans conceptualize and classify *bystander* and *subject* in a photo.[1] Study participants labeled people in a set of 5000 photos taken in the wild [23] as subject or bystander and provided reasons for particular labeling. They also rated each person in terms of several visual characteristics (Table 3). Correlation and regression analyses on data obtained from this study justified our intuition that

[1]Merriam Webster defines 'Bystander' as "one who is present but not taking part in a situation or event: a chance spectator," but this leaves much open to context and socio-cultural norms.

Figure 1: Factor loadings of the *high-level* visual characteristics and other features obtained from annotation data across the two extracted factors. The numeric values of the loadings are displayed within braces with the legend.

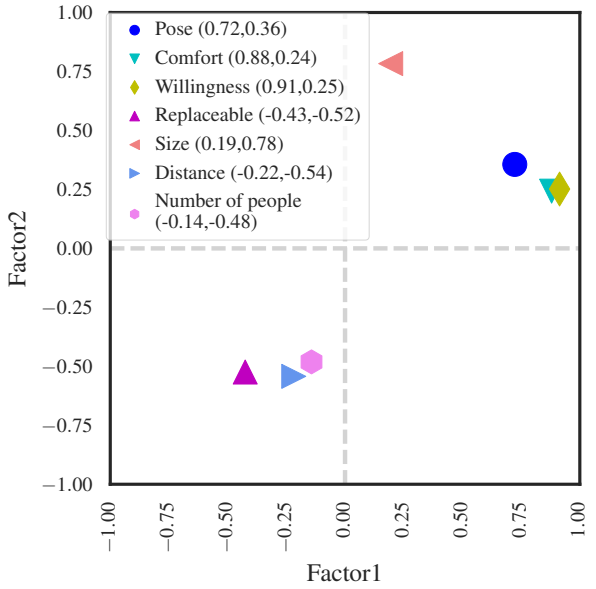

these visual characteristics are relevant for the classification task. From exploratory factor analysis, we identified two underlying constructs (factors) that humans use to classify a person as a subject or bystander – i) visual appearance, and ii) how prominent a person is in the photo. The factor loadings on these two factors are shown in Figure 1.

Our eventual goal is to classify bystanders and subjects automatically, and we experimented with two different approaches to doing it. In the first approach, we trained deep

neural networks with various features extracted from image data, such as features for distinguishing a person from other objects [17]), features related to body orientation [8], and features related to facial expressions [25]. In the second approach, we used the aforementioned features to first predict the high-level, intuitive visual characteristics (Table 3). Then these predicted values were used to train the classifier model. We evaluated both models by doing 10-fold cross-validations. The average classification accuracy obtained from the first approach was 76%, whereas the second approach yielded an accuracy of 85% (also see Figure 2). This improvement suggests that the high-level features contain information more pertinent to the classification of subject/bystander and with less noise compared to the lower-level features they were derived from. This again justifies our selection of the intuitive, high-level features, and helps to interpret the classification model parameters more easily. These findings are currently being reviewed for publication.

Designing interventions to encourage privacy-protective behavior

To identify and/or design persuasive interventions, it is important to understand how people conceptualize privacy in the context of image sharing and how do they make decisions to share images. In a recent study (to appear at S&P'2020 [5]), we asked participants how likely they are to share photo memes containing people unknown to them in three experimental conditions. In two of these conditions, we incorporated textual priming by asking them first to i) imagine themselves as the subjects of the photos, and ii) consider the privacy preference of the subjects in the photos. These memes were also rated for valence (i.e., if they portray the subjects positively or negatively) by 400 people in a separate experiment. We found that, in all conditions, people were less likely to share negative photos than

Figure 2: Receiver operating characteristic (ROC) plots for classifier models using *predicted* values of the *high-level* visual characteristics.

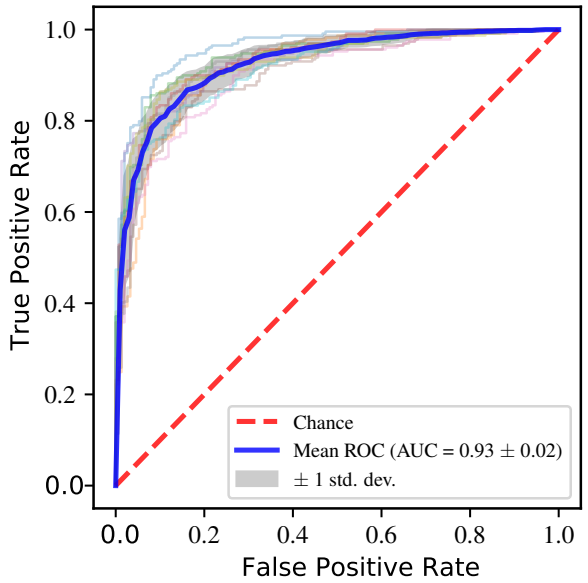

Table 3: Measured visual characteristics of people in photos.

- If the person appears to be *comfortable* being in the photo.
- If the person appears to be *aware of* being photographed.
- If the person appears to be *willing* to be in the photo.
- If the person appears to be *posing* for the photo.

the positive ones. When we compared sharing likelihood across conditions, to our surprise, we found that the likelihood was significantly higher in the two experimental conditions than the control condition. To explain this counter-intuitive result, we conducted a follow-up study where in addition to the previous questions, we asked the participants to explain why they had made a certain sharing decision. Analysis of the qualitative responses (after coding) revealed that, indeed, participants considered the privacy implications of their sharing decision more often when primed to think about others' privacy (condition 3) compared to the other two conditions. But they decided to share anyway because they perceived no privacy risks. One possible reason for thinking in this way could be reactance, or the tendency for seemingly unnecessary rules or prompts to elicit the opposite effect as intended [26]. Past research has also shown that when people explicitly reject the values of the primes, priming might encourage the opposite action [44]. For example, if participants do not value others' privacy, but are reminded of others' privacy, this could cause them to share more rather than less. Yet another possible reason might be that, when primed to think about privacy, participants might have decided exclusively based on whether there was any privacy risk associated with their sharing decision. But in the other conditions they might have considered many other aspects (e.g., photo quality, whether the audience will like it, and if the photo is entertaining enough and so on) including privacy, thus found more reasons to *not share*. Currently, we are conducting more experiments to understand this complex and multifaceted problem and identify effective interventions.

Future work

I am planning and designing new research experiments for all three of the research projects that I am involved in.

New obfuscation method. Currently, I am designing an experiment where I propose to re-purpose masks, emojis, and clip-arts, which are already popular in photo-sharing platforms, as utility-preserving obfuscations.

Bystander Detection. To further improve bystander-detection accuracy, I am continuing to experiment with other features, such as activity and interaction among people in a photo.

Designing new behavioral interventions. To better understand how people make photo-sharing decisions so that we can design appropriate and effective interventions to influence those decisions, I am designing an experiment in lab-settings. My plan is to collect data about eye-movement, pupil dilation, and heart rate while participants view and share photos. I believe these data will improve our understanding of i) on what objects people focus in a photo, ii) the effects of photo-content on emotional arousal, and iii) the association between emotional states and photo sharing behavior.

Conclusion

Online social networks provide immensely helpful services toward society, but also pose great privacy and security risks toward the consumers. With my research, I aim to promote safe and healthy interaction in the online space by understanding the human decision-making process and building tools and technologies that encourage and facilitate privacy-preserving, respectful, and beneficial behavior.

Acknowledgement

This material is based upon work supported in part by the National Science Foundation under grants CNS-1408730, IIS-1253549, IIS-1527421, and CNS-1814513.

REFERENCES

[1] Alessandro Acquisti, Ralph Gross, and Frederic D Stutzman. 2014. Face recognition and privacy in the age of augmented reality. *Journal of Privacy and Confidentiality* 6, 2 (2014), 1.

[2] Anne Adams, Sally Jo Cunningham, Masood Masoodian, and University of Waikato. 2007. *Sharing, privacy and trust issues for photo collections*. Technical Report. https://hdl.handle.net/10289/59

[3] 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 *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '16)*. ACM, New York, NY, USA, 235–248. DOI: http://dx.doi.org/10.1145/2906388.2906412

[4] Shane Ahern, Dean Eckles, Nathaniel S Good, Simon King, Mor Naaman, and Rahul Nair. 2007. Over-exposed?: Privacy Patterns and Considerations in Online and Mobile Photo Sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 357–366. DOI: http://dx.doi.org/10.1145/1240624.1240683

[5] Mary Jean Amon, Rakibul Hasan, Kurt Hugenberg, Bennett I Bertenthal, and Apu Kapadia. 2020. Influencing Photo Sharing Decisions on Social Media: A Case of Paradoxical Findings. In *the Proceedings of the IEEE Symposium on Security & Privacy (SP '20), To appear*.

[6] Cheng Bo, Guobin Shen, Jie Liu, Xiang-Yang Li, YongGuang Zhang, and Feng Zhao. 2014. Privacy.Tag: Privacy Concern Expressed and Respected. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems (SenSys '14)*. ACM, New York, NY, USA, 163–176. DOI: http://dx.doi.org/10.1145/2668332.2668339

[7] By Dawn C. Chmielewski. 2018. YouTube, Instagram And Snapchat All More Popular Than Facebook Among Teens, Pew Reports. (2018). https://deadline.com/2018/05/youtube-instagram-snapchat-more-popular-facebook-american-teens-pe

[8] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2018. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. *arXiv preprint arXiv:1812.08008* (2018).

[9] Eun Kyoung Choe, Sunny Consolvo, Jaeyeon Jung, Beverly Harrison, and Julie A Kientz. 2011. Living in a Glass House: A Survey of Private Moments in the Home. In *Proceedings of the 13th International Conference on Ubiquitous Computing (UbiComp '11)*. ACM, New York, NY, USA, 41–44. DOI: http://dx.doi.org/10.1145/2030112.2030118

[10] Dianne Cyr, Milena Head, Hector Larios, and Bing Pan. 2009. Exploring human images in website design: a multi-method approach. *MIS quarterly* (2009), 539–566.

[11] Jim Edwards. 2014. PLANET SELFIE: We're Now Posting A Staggering 1.8 Billion Photos Every Day. http://www.businessinsider.com/were-now-posting-a-staggering-18-billion-photos-to-social-media-every-day-2014-5. (2014).

[12] Graham D Finlayson, Michal Mackiewicz, and Anya Hurlbert. 2015. Color correction using root-polynomial regression. *IEEE Transactions on Image Processing* 24, 5 (2015), 1460–1470.

[13] Sehee Han, Jinyoung Min, and Heeseok Lee. 2015. Antecedents of social presence and gratification of social connection needs in SNS: A study of Twitter users and their mobile and non-mobile usage. *International Journal of Information Management* 35, 4 (2015), 459–471. DOI:http://dx.doi.org/https://doi.org/10.1016/j.ijinfomgt.2015.04.004

[14] 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 *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 47:1–47:13. DOI: http://dx.doi.org/10.1145/3173574.3173621

[15] Rakibul Hasan, Yifang Li, Eman Hassan, Kelly Caine, David J. Crandall, Roberto Hoyle, and Apu Kapadia. 2019. Can privacy be satisfying? On improving viewer satisfaction for privacy-enhanced photos using aesthetic transforms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, Vol. 14. ACM, 25. DOI: http://dx.doi.org/10.1145/3290605.3300597

[16] Jianping He, Bin Liu, Deguang Kong, Xuan Bao, Na Wang, Hongxia Jin, and George Kesidis. 2014. PuPPIeS: Transformation-Supported Personalized Privacy Preserving Partial Image Sharing. In *IEEE International Conference on Dependable Systems and Networks*. IEEE Computer Society, Atlanta, Georgia USA.

[17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[18] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy Behaviors of Lifeloggers Using Wearable Cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. ACM, New York, NY, USA, 571–582. DOI: http://dx.doi.org/10.1145/2632048.2632079

[19] Isobel Asher Hamilton. 2019. Instagram has avoided Facebook's trust problem, beating its parent as app of choice for Generation Z. (2019). https://www.businessinsider.com/nstagram-is-more-popular-among-generation-z-than-facebook-2019-3/

[20] Sanjay Kairam, Joseph 'Jofish' Kaye, John Alexis Guerra-Gomez, and David A Shamma. 2016. Snap Decisions?: How Users, Content, and Aesthetics Interact to Shape Photo Sharing Behaviors. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 113–124. DOI: http://dx.doi.org/10.1145/2858036.2858451

[21] 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 (CHI '16)*. ACM, New York, NY, USA, 4309–4314. DOI: http://dx.doi.org/10.1145/2858036.2858417

[22] Yi-Cheng Ku, Rui Chen, and Han Zhang. 2013. Why do users continue using social networking sites? An exploratory study of members in the United States and Taiwan. *Information & Management* 50, 7 (2013), 571–581. DOI:http://dx.doi.org/https://doi.org/10.1016/j.im.2013.07.011

[23] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Malloci, Tom Duerig, and Vittorio Ferrari. 2018. The Open Images Dataset V4: Unified image classification, object detection, and visual relationship detection at scale. (2018).

[24] H Lee and A Kobsa. 2017. Privacy preference modeling and prediction in a simulated campuswide IoT environment. In *2017 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 276–285. DOI:http://dx.doi.org/10.1109/PERCOM.2017.7917874

[25] Shan Li, Weihong Deng, and JunPing Du. 2017. Reliable Crowdsourcing and Deep Locality-Preserving Learning for Expression Recognition in the Wild. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2584–2593.

[26] Anca M Miron and Jack W Brehm. 2006. Reactance theory-40 years later. *Zeitschrift für Sozialpsychologie* 37, 1 (2006), 9–18.

[27] Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. 2017. Towards a Visual Privacy Advisor: Understanding and Predicting Privacy Risks in Images. In *The IEEE International Conference on Computer Vision (ICCV)*.

[28] Pew Research Center. 2013. Photo and Video Sharing Grow Online. *Pew Research Center* (2013). http://www.pewinternet.org/2013/10/28/photo-and-video-sharing-grow-online/

[29] Moo-Ryong Ra, Ramesh Govindan, and Antonio Ortega. 2013. P3: Toward Privacy-preserving Photo Sharing. In *USENIX Conference on Networked Systems Design and Implementation (nsdi'13)*. USENIX Association, Berkeley, CA, USA, 515–528.

[30] Yasmeen Rashidi, Tousif Ahmed, Felicia Patel, Emily Fath, Apu Kapadia, Christena Nippert-Eng, and Norman Makoto Su. 2018. "You don't want to be the next meme": College Students' Workarounds to Manage Privacy in the Era of Pervasive Photography. In *Fourteenth Symposium on Usable Privacy and Security ({SOUPS} 2018)*. {USENIX} Association, Baltimore, MD, 143–157. https://www.usenix.org/conference/soups2018/presentation/rashidi

[31] Peter Seddon and Min-Yen Kiew. 1996. A Partial Test and Development of Delone and Mclean's Model of IS Success. *Australasian Journal of Information Systems* 4, 1 (1996). DOI:http://dx.doi.org/10.3127/ajis.v4i1.379

[32] Ryan Shaw. 2006. Recognition markets and visual privacy. *UnBlinking: New Perspectives on Visual Privacy in the 21st Century* (2006).

[33] Jiayu Shu, Rui Zheng, and Pan Hui. 2016. Cardea: Context-aware visual privacy protection from pervasive cameras. *arXiv preprint arXiv:1610.00889* (2016).

[34] Jiayu Shu, Rui Zheng, and Pan Hui. 2017. Your Privacy Is in Your Hand: Interactive Visual Privacy Control with Tags and Gestures. In *Communication Systems and Networks*, Nishanth Sastry and Sandip Chakraborty (Eds.). Springer International Publishing, Cham, 24–43.

[35] Andra Siibak. 2009. Constructing the self through the photo selection-visual impression management on social networking websites. *Cyberpsychology: Journal of psychosocial research on cyberspace* 3, 1 (2009).

[36] Samarth Singhal, Carman Neustaedter, Thecla Schiphorst, Anthony Tang, Abhisekh Patra, and Rui Pan. 2016. You Are Being Watched: Bystanders' Perspective on the Use of Camera Devices in Public Spaces. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)*. ACM, New York, NY, USA, 3197–3203. DOI: http://dx.doi.org/10.1145/2851581.2892522

[37] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. 2013a. The Post That Wasn'T: Exploring Self-censorship on Facebook. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. ACM, New York, NY, USA, 793–802. DOI: http://dx.doi.org/10.1145/2441776.2441865

[38] Manya Sleeper, Justin Cranshaw, Patrick Gage Kelley, Blase Ur, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2013b. "I Read My Twitter the Next Morning and Was Astonished": A Conversational Perspective on Twitter Regrets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 3277–3286. DOI: http://dx.doi.org/10.1145/2470654.2466448

[39] Michelle Starr. 2014. Facial recognition app matches strangers to online profiles. (2014). https://www.cnet.com/news/facial-recognition-app-matches-strangers-to-online-profiles/

[40] Jose M Such, Joel Porter, Sören Preibusch, and Adam Joinson. 2017. Photo Privacy Conflicts in Social Media: A Large-scale Empirical Study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 3821–3832. DOI: http://dx.doi.org/10.1145/3025453.3025668

[41] Robert Templeman, Mohammed Korayem, David Cr, and Apu Kapadia. 2014. PlaceAvoider: Steering first-person cameras away from sensitive spaces. In *In NDSS*.

[42] C Tomasi and R Manduchi. 1998. Bilateral filtering for gray and color images. In *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*. 839–846. DOI: http://dx.doi.org/10.1109/ICCV.1998.710815

[43] Leman Pınar Tosun. 2012. Motives for Facebook use and expressing "true self" on the Internet. *Computers in Human Behavior* 28, 4 (2012), 1510–1517. DOI: http://dx.doi.org/https://doi.org/10.1016/j.chb.2012.03.018

[44] Stephanie Trudeau, Sara Sinclair, and Sean W Smith. 2009. The Effects of Introspection on Creating Privacy Policy. In *Proceedings of the 8th ACM Workshop on Privacy in the Electronic Society (WPES '09)*. ACM, New York, NY, USA, 1–10. DOI: http://dx.doi.org/10.1145/1655188.1655190

[45] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. 2011. "I Regretted the Minute I Pressed Share": A Qualitative Study of Regrets on Facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS '11)*. ACM, New York, NY, USA, 10:1–10:16. DOI: http://dx.doi.org/10.1145/2078827.2078841

[46] Pamela Wisniewski, Heather Lipford, and David Wilson. 2012. Fighting for My Space: Coping Mechanisms for sns Boundary Regulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 609–618. DOI: http://dx.doi.org/10.1145/2207676.2207761

[47] Hang Zhang and Kristin J Dana. 2017. Multi-style Generative Network for Real-time Transfer. *CoRR* abs/1703.0 (2017). http://arxiv.org/abs/1703.06953

[48] L Zhang, T Jung, C Liu, X Ding, X Y Li, and Y Liu. 2015. POP: Privacy-Preserving Outsourced Photo Sharing and Searching for Mobile Devices. In *Distributed Computing Systems (ICDCS), 2015 IEEE 35th International Conference on*. IEEE Computer Society, Columbus,Ohio,USA, 308–317. DOI: http://dx.doi.org/10.1109/ICDCS.2015.39

[49] Lan Zhang, Kebin Liu, Xiang-Yang Li, Cihang Liu, Xuan Ding, and Yunhao Liu. 2016. Privacy-friendly Photo Capturing and Sharing System. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 524–534. DOI: http://dx.doi.org/10.1145/2971648.2971662

[50] Yin Zhang, Leo Shing-Tung Tang, and Louis Leung. 2011. Gratifications, Collective Self-Esteem, Online Emotional Openness, and Traitlike Communication Apprehension as Predictors of Facebook Uses. *Cyberpsychology, Behavior, and Social Networking* 14, 12 (2011), 733–739. DOI: http://dx.doi.org/10.1089/cyber.2010.0042