EFFECTIVENESS AND USERS' EXPERIENCE OF FACE BLURRING AS A PRIVACY PROTECTION FOR SHARING PHOTOS VIA ONLINE SOCIAL NETWORKS

Yifang Li, Nishant Vishwamitra, Hongxin Hu, Bart P. Knijnenburg & Kelly Caine
School of Computing, Clemson University

Photo sharing on online social networks (OSNs) can cause privacy issues. Face blurring is one strategy to increase privacy while still allowing users to share photos. To explore the potential blurring has as a privacy-enhancing technology for OSN photos, we conducted an online experiment with 47 participants to evaluate the effectiveness of face blurring compared to the original photo (as-is), and users' experience (satisfaction, information sufficiency, enjoyment, social presence, and filter likeability). Users' experience ratings for face blurring were positive, indicating blurring may be an acceptable way to modify photos from the users' perspective. However, from a privacy-enhancement perspective, while face blurring may be useful in some situations, such as those where the person in the photo is unknown to the viewer, in other cases, such as in an OSN where the person in the image is known to the viewer, face blurring does not provide privacy protection.

INTRODUCTION

Hundreds of millions of Online Social Network (OSN) users present themselves, communicate, and share thoughts and pictures every day (Ellison, Steinfield, & Lampe, 2007). By 2016, 69% of U.S. adults used at least one OSN, and 68% of U.S. adults use Facebook (Pew Research Center, 2016). Facebook users alone generate four million likes every minute and have uploaded more than 250 billion photos (Smith, 2016).

The data shared on OSNs sometimes includes sensitive details, which generates privacy issues such as unintentional facial recognition, inference attacks, location leakage, identity theft, phishing, profiling risk, or fake product sale (Fire, Goldschmidt, & Elovici, 2014; Kumar, Gupta, Rai, & Sinha, 2013). All of these can be viewed as cybersecurity issues with the potential to lead to cybercrime (Wall, 2004). Cybersecurity, which is defined as "technologies and processes constructed to protect computers, computer hardware, software, networks and data from unauthorized access, vulnerabilities supplied through Internet by cyber criminals, terrorist groups and hackers" (Goutam, 2015, p. 1), is very important for not only organizations and governments, but also individuals and families (Goutam, 2015). Users have a role to play in cybersecurity, including maintaining privacy of sensitive personal information.

Most OSN users have some degree of privacy protection awareness. For example, when they upload a photo, users may choose settings so that the photo to be only visible to their friends. Nevertheless, users often still do not achieve the privacy they desire (Lewis, Kaufman, & Christakis, 2008). For example, users accept friend requests of mutual friends of their Facebook friends, even though they do not know those people (Boshmaf, Muslukhov, Beznosov, & Ripeanu, 2011). Consequently, information intended only for friends may be revealed to strangers.

One proposed solution to the issue of unintentional sharing is to reduce privacy leakage by increasing users' ability to control *who* can access information via an OSN. For example, on Facebook, users can select a subset of friends to share their photos or posts with; or if they are unsatisfied with a photo, they can hide the photo from the network altogether. This approach, and other similar approaches such as un-tagging or deleting the photo, address privacy concerns by preventing unwanted others from viewing their photos (Besmer & Lipford, 2010; Strater & Lipford, 2008). This approach of controlling photo recipients has been studied extensively (e.g., Besmer & Lipford, 2010; Squicciarini, Shehab, & Paci, 2009; Thomas, Grier, & Nicol, 2010), but it has some drawbacks, including causing social tension between the photo uploaders and other stakeholders (Besmer & Lipford, 2010), and increasing sharing loss when one stakeholder sets the visibility as "only me" (Thomas et al., 2010).

In this paper, we focus on another approach, controlling photo *content* disclosure. In this approach, part of the photo is hidden, for example by blurring, pixelating or distorting a person's face to avoid identification. Blurring is the most commonly used and widely studied approach to controlling photo content disclosure (e.g., Ilia, Polakis, Athanasopoulos, Maggi, & Ioannidis, 2015; Li, Li, & Gao, 2016).

The de-identification effectiveness of blurring has been studied both in photo and video-based media. Blurring and/or pixelation is less effective when the intensity of the low (Boyle, Edwards, & Greenberg, 2000; Lander, Bruce, & Hill, 2001). In three experiments exploring the effect of pixelation on unfamiliar face identification results indicated that even in their highest pixelation degree condition (a horizontal image resolution of 8 pixels per face), the successful identification rate is as high as 48% (Bindemann, Attard, Leach, & Johnston, 2013). Moreover, identification accuracy increases when identifying familiar faces (Demanet, Dhont, Notebaert, Pattyn, & Vandierendonck, 2007), and blurred faces are easier to identify than pixelated faces (Harmon & Julesz, 1973). Two studies evaluated the effectiveness of blurring and pixelating on the identification success of familiar faces. Both results showed that although the photos were degraded using these filters, people still could identify people correctly (Demanet et al., 2007; Lander et al., 2001). Notably, a lower degree of blurring or pixilation led to higher successful identification rates (Demanet et al., 2007). However, all of these studies focus on video surveillance, or photos in newspaper/TV, rather than

identification on OSNs. Furthermore, all evaluated strict identification rates, which may not be a good indicator of the potential usefulness of these filters as a possible privacy-enhancing technology. Even though the overall identification rate of a moderate amount of blurring is around 44% (Boyle et al., 2000; Harmon & Julesz, 1973; Lander et al., 2001), this could result in plausible deniability about the identity of a person in an OSN, and therefore may be an effective approach to controlling photo content disclosure. Finally, users' experience of blurring is unknown. To our knowledge, there is no study focusing on both effectiveness and users' experience of blurring as a privacy-enhancing technology.

The goal of our study is to evaluate the effectiveness of blurring as a privacy-enhancing technology and to understand users' experience (satisfaction, perceived information sufficiency, photo enjoyment, social presence, and filter likeability) of facial blurring as a privacy enhancement.

METHOD

We conducted a within-subjects experiment testing 14 levels of privacy filter, but focus our analysis in this paper on two levels because of sample size limitations: *as is* (no filter is applied) and *face blurring* (see (Li, Vishwamitra, Knijnenburg, Hu, & Caine, 2017a; Li, Vishwamitra, Knijnenburg, Hu, & Caine, 2017b) for information about other privacy filters). The dependent variables were privacy filter effectiveness, and users' experience.

Privacy Filter Effectiveness. We measured the identification success by asking participants to "identify the person indicated by the orange arrow" among four options (three ID photos and "None of above"), and their identification confidence from 1 'completely unconfident' to 7 'completely confident' (Phillips, McAuliff, Kovera, & Cutler, 1999).

Users' Experience. Next, we measured four aspects of the privacy filter in the photo: photo satisfaction (Cyr, Head, Larios, & Pan, 2009), information sufficiency (Seddon & Kiew, 1996), photo enjoyment (Redden, 2008), social presence (Kumar & Benbasat, 2006), and privacy filter likeability (Murray & Häubl, 2010), using a 7-point scale, from 1 'Strongly disagree' to 7 'Strongly agree'.

Participants

We recruited 60 participants via Amazon Mechanical Turk. Fifty-seven percent of the participants were aged 25-34, and 23% were aged 35-44. There were 27 women and 20 men, and the majority of them were white (79%). Ninety-eight percent reported using OSNs.

Stimuli

Photos with Privacy Filters. We used Photoshop to create photos of 14 different targets, who were the people participants were asked to identify. Targets were selected from our campus and included white, African American, Asian, and Hispanic and Latino American persons approximating race and ethnicity in the United States (United States Census Bureau, 2010). Eleven targets were male, and three were female. Targets were randomly assigned to the 14 trials for each participant. For each target, we created an *as is* and a *face blurring* version. To make sure all the photos were similar quality, each photo had three random background people and a similar background (campus building etc.; see Figure 1).

ID Photos. We also collected one ID photo of each target, as well as three ID photos of similar-looking people, in terms of hair style, skin color, and height (see Figure 1).

Figure 1. Experiment interface (*face blurring* condition)

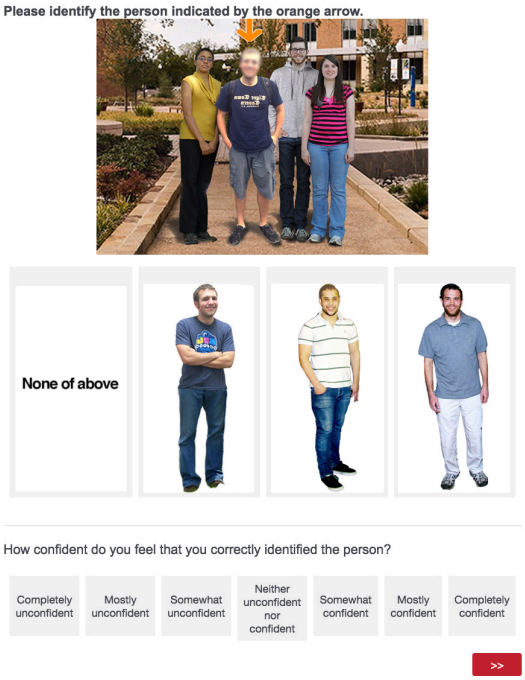

Procedure

In the experiment, after consenting, participants answered six demographic questions and two social network familiarity questions. They received training to learn about the tasks they would perform later, and then began the test. Each participant viewed two group photos, one with the *face blurring* filter on a person's face, and one without a privacy filter (*as is*). The order of these two photos was randomized to avoid order effects. Next, we asked participants to identify the person in the image (the "target") by choosing one of four options, consisting of three persons' ID photos and a "None of above" option. In most cases, the target was among the four options; but in 21% of the trials, the target was NOT present. Following the identification task, we asked participants to rate their confidence (see Figure 1). Afterwards, participants were asked to provide ratings for five statements (satisfaction, information sufficiency, enjoyment, social presence, and likeability) about their perceptions of the privacy filter. At the conclusion of the study, we collected qualitative feedback by asking participants for their thoughts about the *as is* condition and the *face blurring* filter.

RESULTS

Sixty participants completed the experiment. We excluded the data of 13 participants because they failed more than one attention check question, resulting in a final sample size of 47.

Filter Effectiveness

Identification Success. We categorized the identification results using a signal detection approach (Swets, 1964): hit (the target is present, and the response is correct), miss (the target is present, but the response is the wrong person, or "None of above"), correct rejection (the target is absent, and the response is "None of above"), and false alarm (the target is absent, but the response is not "None of above"). For all cases, the overall identification success of *as is* (74%) is higher than *face blurring* (49%) with $p < .05$, (see Table 1). When the target was present, the difference is smaller and not significant (*as is* 76%; *face blurring* 69%). When the target was absent, the difference is larger (*as is* 69%; *face blurring* 24%) ($p < .05$).

Identification Confidence. In Figure 2, "total correct" represents the confidence of both hit and correct rejection; "total wrong" represents both miss and false alarm. For hit, correct rejection, false alarm, and total correct, the differences between *as is* and *face blurring* are small and not significant (*p*s > .05). For total wrong, the confidence rating of *as is* is higher than for the *face blurring* condition ($p < .01$). Overall, the confidence ratings are all above or equal to 4, no matter the filter condition and identification success. This means that participants were generally confident about their selections.

Table 1. Successful identification by filter

		Privacy Filters	
		As is	Face blurring
Identification Success	All cases	74%	49%
	Target present	76%	69%
	Target absent	69%	24%

Figure 2. Mean values of the identification confidence, error bars represent the standard error of the mean

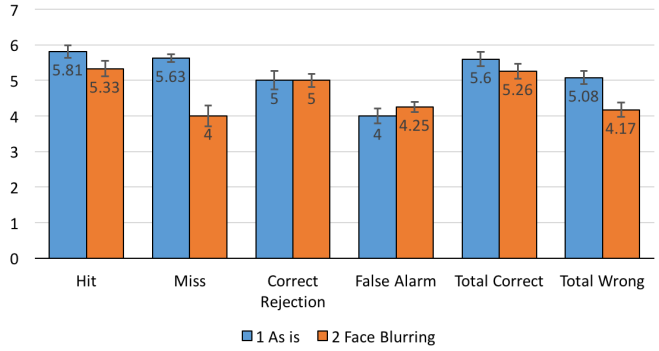

Users' Experience

We conducted five dependent t-tests to compare the differences between *as is* and *face blurring* for our five experience measurements (photo satisfaction, photo information sufficiency, photo enjoyment, social presence, and privacy filter likeability; see Table 2 and Figure 3).

Photo Satisfaction. On average, participants were more satisfied with the *as is* condition than with the *face blurring* condition: $t(46) = 3.74$, $p < .001$.

Photo Information Sufficiency. Participants felt they received more information from the photo without a filter than with the *face blurring* filter: $t(46) = 5.61$, $p < .001$.

Photo Enjoyment. The photos without the privacy filter were perceived as more enjoyable, than those with the *face blurring* filter: $t(46) = 3.79$, $p < .001$.

Social Presence. Participants perceived more human contact when they saw the photo without a filter, than the photo with the *face blurring* filter: $t(46) = 2.57$, $p < .05$.

Likeability. There is no difference of the likeability between *as is* and *face blurring*, $t(46) = 1.24$, $p = .22$, although the mean of *as is* is slightly higher than *face blurring*.

Table 2. Photo satisfaction, information sufficiency, photo enjoyment, social presence, and likeability means for *as is* vs. *face blurring*.

	Privacy Filters			
	As is	Face blurring	t	df
Satisfaction	5.0(0.23)	4.15(0.23)	3.74***	46
Information Sufficiency	5.43(0.22)	4.30(0.22)	5.61***	46
Photo Enjoyment	4.77(0.22)	4.0(0.23)	3.79***	46
Social Presence	5.06(0.22)	4.60(0.22)	2.57*	46
Likeability	4.83(0.28)	4.40(0.28)	1.24	46

Note. * = p < .05, *** = p < .001. Standard Error of the Mean appear in parentheses below means.

Figure 3. Mean values of the photo satisfaction, information sufficiency, photo enjoyment, social presence, and filter likeability, error bars represent the standard error of the mean.

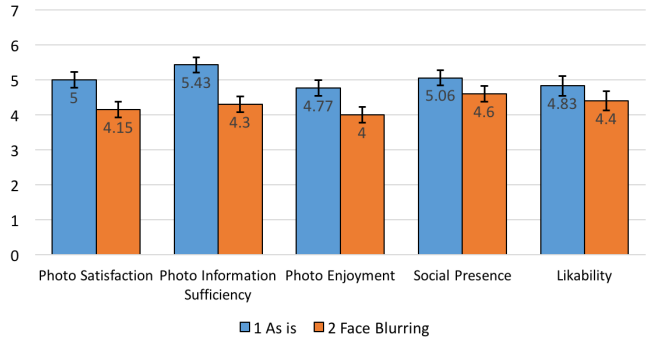

DISCUSSION

Overall, it is easier for viewers to identify people in photos when their faces are not blurred, indicating that facial blurring may provide enhanced privacy compared to photos where no filter is applied. This result suggests that applying *face blurring* filters may improve users' identity privacy to some extent.

However, in cases when a target is present, *face blurring* becomes ineffective with a very high identification success of almost 70%, effectively the same as photos with no filter applied. One possible reason *face blurring* is ineffective when a target was present is that, even though the target's face was blurred, some other hints about identity remained, such as skin color, body shape, or height. Thus, participants could still compare the target with the three ID photos, and identify successfully using these clues. This indicates that *face blurring* may not be effective as a means of privacy-enhancement for OSNs because many of these clues would be available.

Photo-sharing on OSNs is one of the most common identity disclosure elements that may reveal information about users (Stutzman, 2006). This may result in photos being used

for unintended purposes, or revealing unintended details. For example, athletics officials in several colleges have investigated and restricted photo-sharing online because some photos showed bad behavior of student athletes, and their faces were easy to identify, which damaged their reputation (Wolverton, 2006); employers used photos from OSNs that were associated with their employees to evaluate their behaviors (The Washington Times, 2006).

The "target present" case in our study represents a situation similar to this, where the viewer may have a list of candidates of who might be in a blurred image; when targets are familiar or limited, identification rate will be higher (Demanet et al., 2007). Anytime a viewer can imagine or narrow down to a set of potential people who may be in an image (using other clues about context such as body shape, clothing and/or other people in the photo) and therefore may be able to distinguish among that set. Even when the target's face is blurred, viewers may use some other hints to match a blurred face to an identity (Agrawal, 2010; Saini, Atrey, Mehrotra, & Kankanhalli, 2014). This suggests that *face blurring* may not be a good choice for protecting users' privacy for OSNs.

On the other hand, *face blurring* may be somewhat effective for situations where the viewer does not have a list of candidate for who might be in the blurred image. For example, a blurred photo of a non-public person in a national newspaper would be less likely to be recognized because of the huge number of potential people the blurred person could be and the lack of available contextual information (e.g., mutual friends).

Participants were generally confident about their identification when viewing face-blurred photos (mean scores are equal or above 4; see Figure 2). In cases where responses were correct, the confidence ratings of *as is* and *face blurring* are similar, both above five. However, in the "miss" condition, the confidence of *face blurring* is four, which is much lower than *as is* (5.63) suggesting viewers were neither confident nor unconfident when viewing images with *face blurring*, but somewhat confident when viewing *as is* images. In "false alarm", the mean confidence for *face blurring* is similar to "miss" both hovering around four. Based on the confidence ratings in the "miss" and "false alarm" conditions, we could argue that the potential threat of viewers mistakenly identifying the wrong person is somewhat mitigated by the fact that they are not very confident when mis-identifying.

Unsurprisingly, from the satisfaction, information sufficiency, enjoyment, and social presence aspects, participants rated the *as is* condition significantly higher than the *face blurring*. A person's face is an important component of a group photo (Li, Gallagher, Loui, & Chen, 2010), and face blurring reduces the integrity and aesthetic, making it less satisfying and less enjoyable. In addition, the amount of information is limited compared to the original photo. Viewers are less likely to identify the person and distinguish their facial expression because of lack of details (Wang & Tang, 2005). The sense of human contact decreases between not only viewers and the people in the photo, but also the interaction within the group photo. Indeed, from the qualitative feedback, several participants said that blurring might ruin the photo, and they preferred to see the original photos. However, participants' attitudes on these four aspects are generally positive (above or equal to 4). Notably, the ratings for the *as is* condition are not as high as we expected, in that none of them reached 6, which means that maybe the overall quality of our photos is not particularly high. The positive side-effect of this is that we avoid a ceiling effect. This also supports the idea that the *face blurring* filter is acceptable compared to original photos. We can infer that when applying *face blurring*, users' experience will not be negatively impacted to a problematic extent.

Filter likeability is an exception. There is no significant difference on likeability between the *as is* and the *face blurring* condition, and both means are above 4, indicating participants liked both conditions. This is consistent with the above discussion that *face blurring* filter may be an acceptable way to treat some photos. *Face blurring* is as likeable and only slightly less attractive in terms of other elements of users' experience and does provide some amount of privacy enhancement. Given this, it make sense that *face blurring* is a commonly used filter in existing literature (Ilia et al., 2015; A. Li et al., 2016) and in practice (e.g., in Google's street view, people's faces are blurred). However, *face blurring* may be ineffective as a privacy filter when a viewer can determine a set of potential targets. One implication of this is that designers of future photo privacy technology should consider other privacy filters that are more effective across situations.

LIMITATIONS AND FUTURE WORK

We only focus on *face blurring* in this paper because it is one of the most widely used privacy filters and represents an extreme end of a spectrum of possible filters, such as pixelating (Boyle et al., 2000; Demanet et al., 2007; Lander et al., 2001; Vishwamitra, Li, Wang, Hu, Caine & Ahn, 2017), morphing (Jana, Narayanan, & Shmatikov, 2013), and blocking (Li, Vishwamitra, Knijnenburg, Hu, & Caine, 2017a). In our study, all the targets in the photos were unknown to the participants limiting our understanding about the potential for identification success when viewers see familiar faces (Demanet et al., 2007). In the future, we plan to study this by using familiar people (e.g., famous people or participants' OSN friends). In addition, the cross-racial face recognition may influence identification success and confidence. People are more accurate when identifying faces of people of their own race (Meissner & Brigham, 2001). We plan to do a detailed analysis of this in a follow-up study. Finally, we know from prior work in computer vision that even though *face blurring* is effective in decreasing the identification rate among human observers, it is ineffective against machines (Ledig et al., 2016). Therefore, if users are concerned about being identified using automated facial recognition technologies, this technique may be insufficient.

CONCLUSION

We evaluated the effectiveness and users' experience towards the *face blurring* privacy filter compared to the original photo (*as is* condition) to see the potential of applying such filters on OSNs. The results show that while *face blurring* may be useful as a privacy-enhancement in some situations, such as those where the person in the photo is unknown to the viewer; in other cases, such as an OSN where the person in the image is

known to the viewer, *face blurring* does not provide privacy protection. However, from the perspectives of satisfaction, information sufficiency, photo enjoyment, social presence, and likeability, although the ratings for *face blurring* were not as high as *as is*, the participants' attitudes were all positive, which means *face blurring* may be acceptable.

ACKNOWLEDGMENTS

This research was supported by the National Science Foundation under grant no. 1527421. We thank the participants for their willingness to be in the study and our colleagues from the HATlab for suggestions that improved this study.

REFERENCES

Agrawal, P. (2010). *De-Identification for Privacy Protection in Surveillance Videos*. (Masters thesis, International Institute of Information Technology Hyderabad, India).

Besmer, A., & Lipford, H. R. (2010). Moving beyond untagging: photo privacy in a tagged world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1563-1572). ACM.

Bindemann, M., Attard, J., Leach, A., & Johnston, R. A. (2013). The Effect of Image Pixelation on Unfamiliar-Face Matching. *Applied Cognitive Psychology, 27*(6), 707-717.

Boshmaf, Y., Muslukhov, I., Beznosov, K., & Ripeanu, M. (2011). The socialbot network: when bots socialize for fame and money. In *Proceedings of the 27th annual computer security applications conference* (pp. 93-102). ACM.

Boyle, M., Edwards, C., & Greenberg, S. (2000). The effects of filtered video on awareness and privacy. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work* (pp. 1-10). ACM.

Cyr, D., Head, M., Larios, H., & Pan, B. (2009). Exploring human images in website design: a multi-method approach. *MIS quarterly*, 539-566.

Demanet, J., Dhont, K., Notebaert, L., Pattyn, S., & Vandierendonck, A. (2007). Pixelating Familiar People in the Media: Should Masking Be Taken at Face Value? *Psychologica belgica, 47*(4).

Ellison, N. B., W Steinfield, C. W., & Lampe, C. (2007). The benefits of Facebook "friends:" Social capital and college students' use of online social network sites. *Journal of Computer-Mediated Communication, 12*(4), 1143-1168.

Fire, M., Goldschmidt, R., & Elovici, Y. (2014). Online social networks: threats and solutions. *IEEE Communications Surveys & Tutorials, 16*(4), 2019-2036.

Goutam, R. K. (2015). Importance of Cyber Security. *International Journal of Computer Applications, 111*(7).

Harmon, L. D., & Julesz, B. (1973). Masking in visual recognition: Effects of two-dimensional filtered noise. *Science, 180*(4091), 1194-1197.

Ilia, P., Polakis, I., Athanasopoulos, E., Maggi, F., & Ioannidis, S. (2015). Face/off: Preventing privacy leakage from photos in social networks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (pp. 781-792). ACM.

Jana, S., Narayanan, A., & Shmatikov, V. (2013). A Scanner Darkly: Protecting user privacy from perceptual applications. In *Security and Privacy (SP), 2013 IEEE Symposium on* (pp. 349-363). IEEE.

Kumar, A., Gupta, S. K., Rai, A. K., & Sinha, S. (2013). Social networking sites and their security issues. *International Journal of Scientific and Research Publications, 3*(4), 1-5.

Kumar, N., & Benbasat, I. (2006). Research note: the influence of recommendations and consumer reviews on evaluations of websites. *Information Systems Research, 17*(4), 425-439.

Lander, K., Bruce, V., & Hill, H. (2001). Evaluating the effectiveness of pixelation and blurring on masking the identity of familiar faces. *Applied Cognitive Psychology, 15*(1), 101-116.

Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., . . . & Shi, W. (2016). Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*.

Lewis, K., Kaufman, J., & Christakis, N. (2008). The taste for privacy: An analysis of college student privacy settings in an online social network. *Journal of Computer-Mediated Communication, 14*(1), 79-100.

Li, A., Li, Q., & Gao, W. (2016). PrivacyCamera: Cooperative Privacy-Aware Photographing with Mobile Phones. In *Sensing, Communication, and Networking (SECON), 2016 13th Annual IEEE International Conference on* (pp. 1-9). IEEE.

Li, C., Gallagher, A., Loui, A. C., & Chen, T. (2010). Aesthetic quality assessment of consumer photos with faces. In *Image Processing (ICIP), 2010 17th IEEE International Conference on* (pp. 3221-3224). IEEE.

Li, Y., Vishwamitra, N., Knijnenburg, B., Hu, H., & Caine, K. (2017a). Blur vs. Block: Investigating the effectiveness of privacy-enhancing obfuscation for images. In *The First International Workshop on The Bright and Dark Sides of Computer Vision: Challenges and Opportunities for Privacy and Security (CV-COPS 2017)*.

Li, Y., Vishwamitra, N., Knijnenburg, B., Hu, H., & Caine, K. (2017b). *Effectiveness and users' experience of obfuscation as a privacy-enhancing technology for sharing photos.* Manuscript submitted for publication.

Meissner, C. A., & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review.

Murray, K. B., & Häubl, G. (2010). Freedom of choice, ease of use, and the formation of interface preferences.

Pew Research Center. (2016). *Social Media Fact Sheet*. Retrieved from: http://www.pewinternet.org/fact-sheet/social-media/

Phillips, M. R., McAuliff, B. D., Kovera, M. B., & Cutler, B. L. (1999). Double-blind photoarray administration as a safeguard against investigator bias. *Journal of Applied Psychology, 84*(6), 940.

Redden, J. P. (2008). Reducing satiation: The role of categorization level. *Journal of Consumer Research, 34*(5), 624-634.

Saini, M., Atrey, P. K., Mehrotra, S., & Kankanhalli, M. (2014). W3-privacy: understanding what, when, and where inference channels in multi-camera surveillance video. *Multimedia Tools and Applications, 68*(1), 135-158.

Seddon, P., & Kiew, M.-Y. (1996). A partial test and development of DeLone and McLean's model of IS success. *Australasian Journal of Information Systems, 4*(1).

Smith, K. (2016). *Marketing: 47 Facebook Statistics for 2016*. Retrieved from https://www.brandwatch.com/blog/47-facebook-statistics-2016/

Squicciarini, A. C., Shehab, M., & Paci, F. (2009). Collective privacy management in social networks. In *Proceedings of the 18th international conference on World wide web* (pp. 521-530). ACM.

Strater, K., & Lipford, H. R. (2008). Strategies and struggles with privacy in an online social networking community. In *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction-Volume 1* (pp. 111-119). British Computer Society.

Stutzman, F. (2006). An evaluation of identity-sharing behavior in social network communities. *Journal of the International Digital Media and Arts Association, 3*(1), 10-18.

Swets, J. A. (1964). Signal Detection and Recognition in Human Observers: Contemporary Readings.

The Washington Times. (2006). Face It: 'Book' No Secret to Employers. Retrieved from http://www.washingtontimes.com/news/2006/jul/17/20060717-124952-1800r/

Thomas, K., Grier, C., & Nicol, D. M. (2010). unFriendly: Multi-party privacy risks in social networks. In *International Symposium on Privacy Enhancing Technologies* Symposium (pp. 236-252). Springer Berlin Heidelberg.

United States Census Bureau. (2010). *American FactFinder - Race Results*. Retrieved from https://factfinder.census.gov/faces/tableservices/jsf/pages/productview.xhtml?pid=DEC_10_DP_DPDP1&src=pt

Vishwamitra, N., Li, Y., Wang, K., Hu, H., Caine, K., & Ahn, G. J. (2017). Towards PII-based multiparty access control for photo sharing in online social networks. In P*roceedings of the 22nd ACM on Symposium on Access Control Models and Technologies* (pp. 156-166). ACM.

Wall, D. (2004). What are Cybercrimes? *Criminal Justice Matters, 58*(1), 20-21.

Wang, X., & Tang, X. (2005). Hallucinating face by eigentransformation. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reivews), 35*(3), 425-434.

Wolverton, B. (2006). Hazing Photos Spur Debates on Complicity of Coaches. Retrieved from http://www.chronicle.com/article/Hazing-Photos-Spur-Debates-on/8324