

Blur vs. Block: Investigating the Effectiveness of Privacy-Enhancing Obfuscation for Images

Yifang Li, Nishant Vishwamitra, Bart P. Knijnenburg, Hongxin Hu and Kelly Caine
Clemson University

{yifang2, nvishwa, bartk, hongxih, caine}@clemsun.edu

Abstract

Computer vision can lead to privacy issues such as unauthorized disclosure of private information and identity theft, but it may also be used to preserve user privacy. For example, using computer vision, we may be able to identify sensitive elements of an image and obfuscate those elements thereby protecting private information or identity. However, there is a lack of research investigating the effectiveness of applying obfuscation techniques to parts of images as a privacy-enhancing technology. In particular, we know very little about how well obfuscation works for human viewers or users' attitudes towards using these mechanisms. In this paper, we report results from an online experiment with 53 participants that investigates the effectiveness two exemplar obfuscation techniques: "blurring" and "blocking", and explores users' perceptions of these obfuscations in terms of image satisfaction, information sufficiency, enjoyment, and social presence. Results show that although "blocking" is more effective at de-identification compared to "blurring" or leaving the image "as is", users' attitudes towards "blocking" are the most negative, which creates a conflict between privacy protection and users' experience. Future work should explore alternative obfuscation techniques that could protect users' privacy and also provide a good viewing experience.

1. Introduction

Advances in computer vision have amplified already intense concerns about personal privacy. Whereas last century humans had to manually extract information from images, now computers can do so automatically [31]. Computer vision allows faster processing and analysis of images which means that any image that is captured digitally can provide more information than ever before, increasing potential privacy implications.

Privacy problems commonly arise in Online Social Networks (OSNs), where personal information about a user

such as relationships, affiliations and personal belongings can be learned from analyzing her/his photos [32, 58, 71]. In addition, deep learning has recently emerged as one of the most effective machine learning techniques for image segmentation and analysis [12, 46, 57, 66]. Using deep neural networks, attackers can extract objects from an image and also learn image semantics, such as the relationships between persons in an image, with a high degree of precision. Besides, using Optical Character Recognition (OCR) technology, attackers can also extract private textual information in an image, such as credit card numbers or social security numbers [49, 62]. Such computer vision-based techniques can be even used to perform "shoulder surfing attacks", which involve exposing private or sensitive information on a computer or smartphone screen to unauthorized individuals [49, 58, 71].

Aside from amplifying privacy consequences, computer vision may also be useful in preserving user privacy, because the mechanisms used to detect sensitive information can also be used to obfuscate that information in an image before posting it online.

Outside the computer vision community, there are two common approaches that protect image privacy: controlling the access to an image, and controlling the image content [8]. The first approach has been investigated extensively, especially in online image privacy research (e.g., [5, 29, 30, 59, 61]). For example, in the context of the OSN, Facebook, "Restrict Others" is a tool that provides a communication channel between the photo uploader and other stakeholders in the photo, so that the stakeholders can request that the owner sets access restrictions for certain viewers [5]. However, studies on this image access control mechanism unearthed negative effects, such as increased social tension between the owner and stakeholders [5], and an increased loss of information due to conflicting sharing requirements imposed by multiple users in an image, for example, if one stakeholder sets restrictions such that few people can access the image [61].

A second, less well-studied approach is to control the disclosure of image content [9]. In this approach, each im-

age element may be disclosed at a granular level. For example, instead of obfuscating the whole image, we may choose to obfuscate only part of the image content [10, 32, 67]. One benefit of this approach is that it may simultaneously preserve user privacy and reduce sharing loss. A common way to control image content is to blur a person’s faces in an image [4, 32, 37, 40, 42, 67]. In this way, even if the image is shared with unintended viewers, they may be unable to discern the identity of the person in the image [32].

Various elements in an image reveal personal information about stakeholders besides identity. People may consider some of these elements, such as objects in the image, the background setting and people within the image to be particularly sensitive [28]. For example, computer or phone screens may reveal sensitive information when captured in an image if not obscured [34], and Google Street View obscures faces and license plates by blurring them [22].

One widely used way to obfuscate sensitive information in an image is to smooth the image and remove the details using a Gaussian blur [27]. Indeed, blurring is the most commonly investigated and adopted obfuscation method in prior research across disciplines and has been applied to online photos, online videos, video surveillance and images published in newspapers or on television [22, 25, 40, 76]. There are different levels of blurring obfuscation. Increasing the blur level generally reduces the accuracy of identification by both machines and human viewers [25, 40, 42, 72, 73, 74].

In most of the aforementioned studies researchers used square face blurring, which has a number of drawbacks. First, since the blurred region is only the face, some body features can provide important hints for identification, such as clothes and gestures [1, 55]. Indeed, blurring the entire body is more effective than just blurring the face [11]. Second, square blurring may cover more than just the sensitive image content. For example, part of the adjacent background or objects may be blurred, causing loss of image quality without providing enhanced privacy. This may potentially reduce users’ satisfaction with an image.

Despite the ubiquity of blurring and even with increased effectiveness of higher levels of blur that may protect image contents from human identification, we know that blurring is at least partially reversible using automatic re-identification. For example, the generative adversarial network for image super-resolution discussed in [41] helps recover the finer details; investigating the space spanned by the blurred image also facilitates re-identification [24]; artificial neural networks can be trained to recognize obscured faces, objects, and handwritten digits [47]; the attack algorithms discussed in [26, 65] are used to identify blurred and pixelated faces; the Custom Pictorial Structure technique used in [13] can re-identify the person by his/her pose; and similarly, faceless recognition can be accurate through an-

alyzing clothes and pose [51]. These techniques are fairly accurate; the accuracy of the top one guess of the identity of a person in a blurred image by a deep learning algorithm is 58%, and at five guesses, this rate increases to 86% [47].

In artificial neural networks designed to defeat image obfuscation, pixelating, blurring, and P3, which protects the privacy by splitting each JPEG into a public image and a secret image, were all found to be ineffective at protecting against attacks [47]. As an alternative method, block obfuscation, in which a solid block covers the person’s entire body, provides more privacy than blurring in the context of video assistive monitoring [19]. We similarly expect a higher objective effectiveness when using *blocking* compared to *blurring* for still images. Besides effectiveness, we are also interested in viewers’ experience of both the *blur* and *block* obfuscations.

Even if humans are not the “gold standard for extracting data from visual data” [47], it is still important to understand what obfuscation methods are effective for human viewers. From a social perspective, humans are the intended recipients, rather than machines, and while some may try to visually infer the identity of an obfuscated person, most will not go so far as to employ machine learning to achieve this goal. Considering human abilities alongside computer vision will therefore paint a more complete picture for our understanding of the overall effectiveness of obfuscation techniques.

OUR GOAL. Our goal is to investigate the effectiveness of *blurring* and *blocking* as privacy enhancing tools against human recognition, and to explore users’ perceptions of these obfuscations from the perspective of image satisfaction, information sufficiency, image enjoyment, and social presence.

Various elements in an image are considered sensitive, such as objects, people, or visible textual private information [28]. We select the *people* in the image as the element to be protected in this study because identity management is extremely important, especially in an online environment [2]. We expect that the findings about effectiveness and perhaps viewer experience may be applied to other image sensitive elements such as text, objects, and background setting.

We focus on two types of obfuscation: *blurring* and *blocking*. As mentioned in previous paragraphs, blurring the entire body is more effective than just blurring the face [11]. Therefore, in our study, we deployed entire body blurring. Moreover, to achieve both privacy and higher satisfaction, we made the blurred region follow the body shape rather than square. For *blocking*, we applied a gray square block covering the person’s body. We used a neutral gray across all images to avoid possible color bias in users’ perception of image satisfaction.

Besides investigating obfuscation effectiveness, we also




Figure 1. Experiment Interface.

explore users’ perceptions (in terms on satisfaction, information sufficiency, enjoyment, and social presence) of these obfuscations. In common usage scenarios, such as images shared via OSNs, the target audience is humans so it is important to understand whether these obfuscations may be detrimental to the viewers’ experience. To our knowledge, this is among the first research focusing on both effectiveness against human recognition and users’ perceptions of *blurring* and *blocking* as privacy-enhancing obfuscations, though see [44] for emerging work in this area.

2. Method

We conducted a within-subjects experiment with three obfuscation conditions: *as is* (no obfuscation is applied), *blurring*, and *blocking*. We investigated the effectiveness (defined here as the inability for humans to identify the target person) and users’ perceptions of these obfuscations.

For obfuscation effectiveness, we first measured the identification success by asking participants to “Please identify the person indicated by the orange arrow” with four choices (three photos including the target person and two distractors and “None of above”). Next, we measured participants’ identification confidence from 1 ‘completely unconfident’ to 7 ‘completely confident’ [53] (Figure 1).

For viewers’ experience, we measured image satisfac-



Figure 2. Examples of blurring and block obfuscations.

tion [16], image information sufficiency [56], image enjoyment [54], social presence [39], and obfuscation likability [50] by asking participants to rate the following five statements: “The photo is satisfying”, “The photo provides sufficient information”, “I enjoy the photo at this moment”, “There was a sense of human contact when I saw the photo”, and “I like the [blur, block, as is] obfuscation”. We used seven-point response scales from 1 ‘Strongly disagree’ to 7 ‘Strongly agree’ for each of these questions.

2.1. Participants

Fifty-three participants were recruited via Amazon Mechanical Turk. There were 21 men, and 30 women. Two participants preferred not to reveal their gender. Seventy-five percent of participants ranged in age from 24 to 44, and 74% were white.

2.2. Stimuli

IMAGE CREATION. We selected 14 targets (11 males, 3 females) who were the people to be identified in the experiment. Their races include White, Asian, African American, and Hispanic/Latino Americans [63]. We applied three obfuscation conditions (*as is*, *blurring*, and *blocking*) on each target image (Figure 2). All images have the same quality and composition, with three background people, and a similar campus scene background.

ID PHOTOS. We collected one ID photo consisting of a full body image from each target, and three full body images of similar looking people (Figure 1).

2.3. Procedure

Participants first gave informed consent, and then began the online experiment. Before the data collection began, participants were provided with training to familiarize them with the tasks they would perform. Each participant viewed one image in each of the three conditions (*as is*, *blurring*, and *blocking*), and the order of the images was randomized. For each image, they were asked to identify the target person by selecting one of four options (three full body images, and “None of above”) (Figure 1). The target was presented among the four options in 79% of the test rounds, and absent in the remaining 21% (selected at random). After identification, participants rated their confidence. Next, they rated

		As is	Blurring	Blocking
Identification Success	All cases	81%	68%	19%
	Target present	82%	80%	7%
	Target absent	79%	44%	70%

Table 1. Successful identification by obfuscation type.

the four statements regarding their satisfaction, information sufficiency, enjoyment, social presence, and the likability of the image, before moving to the next round.

3. Results

We present the results in terms of obfuscation effectiveness, which contains identification success and confidence, and users' perception of the three obfuscation conditions¹ in terms of image satisfaction, information sufficiency, enjoyment, social presence, and likability.

3.1. Obfuscation Effectiveness

IDENTIFICATION SUCCESS. We used a signal detection analysis approach to classify the identification results into four categories: hit (the target is present, and the response is correct), miss (the target is present, but the response is wrong, including selecting the wrong person, and "None of above"), correct rejection (the target is absent, and the response is "None of above"), and false alarm (the target is absent, but participant does not select "None of above").

The Tukey post hoc tests on logistic mixed-effects models showed that the overall identification success rate (hit + correct rejection) of *as is* (81%) is slightly higher than *blurring* (68%), but the difference is not significant ($p = .19$); and the success of both is much higher than *blocking* (19%) (both $ps < .001$). When the target is present, the success rates (hit) of *as is* (82%) and *blurring* (80%) are almost identical ($p = .98$), while the difference between these two and *blocking* (7%) becomes larger (both $ps < .001$). On the contrary, for "target absent" cases, the success rates (correct rejection) of *as is* (79%) and *blocking* (70%) are similar ($p = .78$). The rate of *blurring* decreases to 44%. There is a difference between *blurring* and *as is* ($p < .05$), while between *blurring* and *blocking*, the difference is not significant ($p = .36$) (Table 1).

IDENTIFICATION CONFIDENCE. In Figure 3, the confidence of "total correct" shows the confidence of both "hit" and "correct rejection"; the confidence of "total




Figure 3. Mean values of the identification confidence, error bars represent the standard error of the mean.

wrong" represents both "miss" and "false alarm". The patterns of "hit", "miss", and "correct rejection", are similar: the confidence for *as is* is slightly higher than for *blurring* (scores are all around five to six, corresponding to confident) though not significant ($ps > .05$). In "hit", the confidence of *as is* is much higher than *blocking* (between two and three, which corresponds to not confident; $p < .001$).

For "false alarm", the confidence ratings of all three obfuscation conditions are below four, which corresponds to not confident, although the confidence for *as is* ($M = 3.67$, $SE = 0.21$) and *blurring* ($M = 3.40$, $SE = 0.19$) is slightly higher than for *blocking* ($M = 2.33$, $SE = 0.32$) (both $ps > .05$).

3.2. Users' Perception

Image satisfaction, enjoyment, social presence, and likability scores are presented in Figure 4 and Table 2. An analysis of the linear mixed effect model on image satisfaction scores yielded significant variation among three obfuscation conditions, $\chi^2(2) = 60.55$, $p < .0001$, indicating the obfuscations affected satisfaction differently. The Tukey post hoc test showed that each obfuscation type is significantly different from the other at $p < .001$, with *as is* being the most satisfying, then *blurring*, and then *blocking*.

Similar patterns occur for image information sufficiency ($\chi^2(2) = 131.15$, $p < .0001$), image enjoyment ($\chi^2(2) = 75.10$, $p < .0001$), social presence ($\chi^2(2) = 77.63$, $p < .0001$), and likability ($\chi^2(2) = 60.10$, $p < .0001$). From the post hoc test results, each condition differs from both others significantly (all $ps < .001$), except for likability: the likability of *as is* and *blurring* are significantly higher than *blocking* (both $ps < .001$), but there is no significant difference between *as is* and *blurring*.

¹We tested 14 conditions in all, but focus on three exemplars here due to sample size limitations. See [45] for more information.




Figure 4. Mean values of the image satisfaction, image information sufficiency, image enjoyment, social presence, and obfuscation likability, error bars represent the standard error of the mean.

4. Discussion

Overall, it is easiest for participants to identify the people in images where no obfuscation is applied (*as is*): the identification rate for this condition is 80%. However, the identification rate for *blurring*, the most commonly used obfuscation method in previous academic research and industry practice [32, 42, 72], is also quite high, at 68%. When the target is present in the photo, the success rate further increases to 80%, which is almost the same as *as is*. If *blurring* is applied to the face only, as is often the case (e.g., [25, 32, 42]), we can infer that the identification success rate would probably be even higher.

Hence, *blocking* appears to be a better privacy-enhancing technology, because it removes all hints about targets, thus prevents them from being identified. Note that when the target is absent, the success rate increases to 70%. This is because without any clues, most participants tend to select “None of above” when both targets are present or absent, resulting in wrong responses in cases where the target is present, and more correct responses in those cases where the target is indeed absent.

In cases where identifications were correct (“hit” and “correct rejection”), the confidence ratings for *blurring* and *blocking* are lower than for *as is*. This indicates that even if participants are able to successfully identify the targets in images, they are less certain of their answer for obfuscated images versus non-obfuscated images. This means that while blocking and even blurring may provide some level of *plausible deniability* which is a desirable characteristic for enhancing privacy [69]. Note that participants were still rather confident for *blurring* cases, so *blocking* is again the better obfuscation method, as it results in lower levels of identification confidence.

Note that “miss” and “false alarm” cases mean that people mistakenly identify a different person as the target in an

	As is	Blurring	Blocking
Image satisfaction	4.83 (0.20)	3.74 (0.21)	2.51 (0.19)
Image Information Sufficiency	5.45 (0.18)	4.06 (0.20)	2.00 (0.15)
Image Enjoyment	4.81 (0.20)	3.79 (0.19)	2.42 (0.19)
Social Presence	5.02 (0.21)	4.08 (0.20)	2.51 (0.18)
Likability	4.98 (0.27)	4.19 (0.24)	2.15 (0.21)

Table 2. Image satisfaction, image information sufficiency, image enjoyment, social presence, and likability means for *as is*, *blurring*, and *blocking*. Standard errors of the means appear in parentheses below the means.

image. Such mis-identifications, while perhaps contributing positively to plausible deniability, may have a negative effect on the mis-identified person. For example, when either the image’s background setting or the behavior of the target itself is inappropriate the mis-identified person could suffer the social consequences. The low identification confidence of *blocking* eliminates this threat.

From the perspective of obfuscation effectiveness against human recognition, *blocking* performs the best. However, considering image satisfaction, information sufficiency, enjoyment, and social presence, we found that users’ experience of the *blocking* was far worse than the *blurring* condition. A person is an important component of a group photo, and for our stimuli we added obfuscation on the person who is in the middle of the photo. Hence, both *blurring* and *blocking* lead to a reduction in aesthetic and integrity of the photo [43], which results in lower satisfaction compared to the original image [23]. Note that users’ satisfaction with *blocking* is lower than with *blurring*, arguably because the square mask covers more content.

In some cases, part of the adjacent people in the image may also be covered, which explains the lower information sufficiency of *blocking* compared to *as is* and *blurring*. Although in *blurring* the target’s whole body is blurred, some traits are still detectable, such as height and skin color [1, 55], which yield to higher information sufficiency than *blocking*.

The integrity and aesthetic of the image may influence enjoyment as well, so it is not surprising that *blocking* tends to be the least enjoyable.

Since viewers are less likely to easily identify the person, and because it is hard to perceive the person’s facial

expression, the sense of human contact appears to be lower between not only viewers and people in the image, but also the people’s interaction within the image. This arguably reduces social presence for the *blurring* condition and even more pronouncedly so for the *blocking* condition. If applying these obfuscations on OSNs, lack of human contact in a photo would decrease users’ participation and reduce their motivations of using the medium [21].

For obfuscation likability, there is no major difference between *as is* and *blurring*, as the ratings indicate participants like both. However, they dislike *blocking*, with a rating as low as 2.15 out of 7.

Notably, the ratings across all five measurements for *as is* are not as high as we expected, as none of them are greater than 6 (Figure 4). This may indicate an overall lack of quality of the images we created. On the other hand, this shows that we did not have ceiling effect. The ratings of satisfaction, information sufficiency, enjoyment, social presence, and likability may be higher than current ratings if we apply obfuscation on an image with better quality.

To summarize: *blocking* is much more effective in de-identification comparing to *blurring* and *as is* for both human recognition and, as others have found (*cf.* [19]), machine identification. However, users attitude towards *blocking* is not as good as it is towards *blurring* and *as is*, which creates a conflict between privacy protection and users’ experience. Alternative obfuscation methods that provide both privacy protection and a good user experience should be explored in future studies.

5. Limitations and Future Work

We note a number of limitations of this work. First, we only focus on *blurring* and *blocking* in this paper. We chose to focus on these obfuscations because the represented extreme ends of a spectrum of possible obfuscations. Other possible obfuscations along this spectrum include pixelating [6, 17, 20, 25, 35, 37, 40, 68], silhouette [18, 38, 52, 75], morphing [33, 36], and inpainting [14, 15, 38, 52, 60, 64, 70, 75, 77, 78]. Now that we know the effectiveness and viewers’ experience about the extreme ends of the spectrum, future work should explore other points along the spectrum.

Second, we know that the level or intensity of *blurring* affects obfuscation effectiveness against human recognition [25, 40]. We adopted a lower-intensity blurring (the lowest blurring level in [40]) which may have led to a higher identification success over stronger blur options. Future work should investigate the effect of more intense blurs, or other blur modifications that may make blurring more effective against human recognition. From a machine vision perspective, with the increasing blurring intensity, the identification accuracy decreases [3, 25].

Third, all targets in the images were unknown to the

participants, which does not match a likely usage scenario, where targets may include both unfamiliar and familiar people. When people are familiar with each other, we expect a higher recognition rate [7, 17]. Fourth, the targets in our images and the participants in our study included a number of different races. This may influence identification success, since people tend to be better able to identify faces of others who match their race [48]. Last, we used *blurring* along the body shape rather than square *blurring*, which may increase the identification success, because some information may be gleaned from the target’s body shape.

In future studies, we plan to explore alternative obfuscation methods that are both effective and satisfying. We also plan to study how these obfuscation methods can be applied to other privacy-sensitive image elements such as a diabetic bracelet or a beer can, or privacy sensitive background settings such as a seedy bar or a hospital.

6. Conclusion

In this paper, we evaluated two obfuscations (*blurring* and *blocking*) in terms of their user experience and effectiveness against human recognition. Results suggest that *blurring* (the most commonly used obfuscation) is much less effective at preserving privacy than *blocking*. However, from the perspective of satisfaction, information sufficiency, photo enjoyment, social presence, and likability, *blocking* is less desirable than for *blurring* and *as is*. Hence, there is a need to develop and evaluate alternative obfuscation methods that both protect privacy and provide good user experience.

7. Acknowledgement

This research was supported by the National Science Foundation under grant no.1527421. We thank Cheng Guo for great assistance with the Qualtrics experiment creation, and all colleagues from Hatlab for suggestions that improved this study. We would also like to thank Mahdi Nasrullah Al-Ameen, Byron Lowens, Yangyang He, Darcia Wilkinson, Pratitee Sinha, Mary-Jo May, Brian Justice, Ellie Ebrahimi, Geoffry Sheehan, Brady Bannister, Jason Alexander, Levin Czenkusch, Freddie Pope and Matthew Ellsworth who served as models for our stimuli. Finally, we are grateful to the people who participated in this study.

References

- [1] P. Agrawal. *De-Identification for Privacy Protection in Surveillance Videos*. PhD thesis, International Institute of Information Technology Hyderabad, India, 2010. 2, 5
- [2] S. Ahern, D. Eckles, N. S. Good, S. King, M. Naaman, and R. Nair. Over-exposed?: privacy patterns and considerations in online and mobile photo sharing. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 357–366. ACM, 2007. 2

- [3] T. Ahonen, E. Rahtu, V. Ojansivu, and J. Heikkilä. Recognition of blurred faces using local phase quantization. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, pages 1–4. IEEE, 2008. 6
- [4] A. Besmer and H. Lipford. Tagged photos: concerns, perceptions, and protections. In *CHI'09 Extended Abstracts on Human Factors in Computing Systems*, pages 4585–4590. ACM, 2009. 2
- [5] A. Besmer and H. Richter Lipford. Moving beyond untagging: photo privacy in a tagged world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1563–1572. ACM, 2010. 1
- [6] M. Boyle, C. Edwards, and S. Greenberg. The effects of filtered video on awareness and privacy. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, pages 1–10. ACM, 2000. 6
- [7] V. Bruce, Z. Henderson, C. Newman, and A. M. Burton. Matching identities of familiar and unfamiliar faces caught on cctv images. *Journal of Experimental Psychology: Applied*, 7(3):207, 2001. 6
- [8] K. E. Caine. *Exploring everyday privacy behaviors and misclosures*. PhD thesis, Georgia Institute of Technology, 2009. 1
- [9] K. E. Caine, W. A. Rogers, and A. D. Fisk. Privacy perceptions of an aware home with visual sensing devices. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 49, pages 1856–1858. SAGE Publications, 2005. 1
- [10] K. E. Caine, C. Y. Zimmerman, Z. Schall-Zimmerman, W. R. Hazlewood, L. J. Camp, K. H. Connelly, L. L. Huber, and K. Shankar. Digiswitch: A device to allow older adults to monitor and direct the collection and transmission of health information collected at home. *Journal of medical systems*, 35(5):1181–1195, 2011. 2
- [11] D. Chen, Y. Chang, R. Yan, and J. Yang. Protecting personal identification in video. In *Protecting Privacy in Video Surveillance*, pages 115–128. Springer, 2009. 2
- [12] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014. 1
- [13] D. S. Cheng, M. Cristani, M. Stoppa, L. Bazzani, and V. Murino. Custom pictorial structures for re-identification. In *Bmvc*, volume 2, page 6, 2011. 2
- [14] S.-C. Cheung, M. Venkatesh, J. Paruchuri, J. Zhao, and T. Nguyen. Protecting and managing privacy information in video surveillance systems. In *Protecting Privacy in Video Surveillance*, pages 11–33. Springer, 2009. 6
- [15] A. Criminisi, P. Perez, and K. Toyama. Object removal by exemplar-based inpainting. In *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*, volume 2, pages II–II. IEEE, 2003. 6
- [16] D. Cyr, M. Head, H. Larios, and B. Pan. Exploring human images in website design: a multi-method approach. *MIS quarterly*, pages 539–566, 2009. 3
- [17] J. Demanet, K. Dhont, L. Notebaert, S. Pattyn, and A. Vandierendonck. Pixelating familiar people in the media: Should masking be taken at face value? *Psychologica belgica*, 47(4), 2007. 6
- [18] G. Demiris, D. P. Oliver, J. Giger, M. Skubic, and M. Rantz. Older adults’ privacy considerations for vision based recognition methods of eldercare applications. *Technology and Health Care*, 17(1):41–48, 2009. 6
- [19] A. Edgcomb and F. Vahid. Privacy perception and fall detection accuracy for in-home video assistive monitoring with privacy enhancements. *ACM SIGHIT Record*, 2(2):6–15, 2012. 2, 6
- [20] A. Erdélyi, T. Barát, P. Valet, T. Winkler, and B. Rinner. Adaptive cartooning for privacy protection in camera networks. In *Advanced Video and Signal Based Surveillance (AVSS), 2014 11th IEEE International Conference on*, pages 44–49. IEEE, 2014. 6
- [21] A. J. Flanagan and M. J. Metzger. Internet use in the contemporary media environment. *Human communication research*, 27(1):153–181, 2001. 6
- [22] A. Frome, G. Cheung, A. Abdulkader, M. Zennaro, B. Wu, A. Bissacco, H. Adam, H. Neven, and L. Vincent. Large-scale privacy protection in google street view. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 2373–2380. IEEE, 2009. 2
- [23] B. Geng, L. Yang, C. Xu, X.-S. Hua, and S. Li. The role of attractiveness in web image search. In *Proceedings of the 19th ACM international conference on Multimedia*, pages 63–72. ACM, 2011. 5
- [24] R. Gopalan, S. Taheri, P. Turaga, and R. Chellappa. A blur-robust descriptor with applications to face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 34(6):1220–1226, 2012. 2
- [25] R. Gross, E. Airoldi, B. Malin, and L. Sweeney. Integrating utility into face de-identification. In *International Workshop on Privacy Enhancing Technologies*, pages 227–242. Springer, 2005. 2, 5, 6
- [26] R. Gross, L. Sweeney, F. De la Torre, and S. Baker. Model-based face de-identification. In *Computer Vision and Pattern Recognition Workshop, 2006. CVPRW'06. Conference on*, pages 161–161. IEEE, 2006. 2
- [27] S. Hill, Z. Zhou, L. Saul, and H. Shacham. On the (in) effectiveness of mosaicing and blurring as tools for document redaction. *Proceedings on Privacy Enhancing Technologies*, 2016(4):403–417, 2016. 2
- [28] R. Hoyle, R. Templeman, S. Armes, D. Anthony, D. Crandall, and A. Kapadia. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 571–582. ACM, 2014. 2
- [29] H. Hu, G.-J. Ahn, and J. Jorgensen. Detecting and resolving privacy conflicts for collaborative data sharing in online social networks. In *Proceedings of the 27th Annual Computer Security Applications Conference*, pages 103–112. ACM, 2011. 1
- [30] H. Hu, G.-J. Ahn, Z. Zhao, and D. Yang. Game theoretic analysis of multiparty access control in online social networks. In *Proceedings of the 19th ACM symposium on Access control models and technologies*, pages 93–102. ACM, 2014. 1

- [31] T. Huang. Computer vision: Evolution and promise. *CERN EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH-REPORTS-CERN*, pages 21–26, 1996. 1
- [32] P. Iliia, I. Polakis, E. Athanasopoulos, F. Maggi, and S. Ioanidis. Face/off: Preventing privacy leakage from photos in social networks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 781–792. ACM, 2015. 1, 2, 5
- [33] S. Jana, A. Narayanan, and V. Shmatikov. A scanner darkly: Protecting user privacy from perceptual applications. In *Security and Privacy (SP), 2013 IEEE Symposium on*, pages 349–363. IEEE, 2013. 6
- [34] M. Korayem, R. Templeman, D. Chen, D. Crandall, and A. Kapadia. Enhancing lifelogging privacy by detecting screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4309–4314. ACM, 2016. 2
- [35] P. Korshunov, C. Araimo, F. De Simone, C. Velardo, J.-L. Dugelay, and T. Ebrahimi. Subjective study of privacy filters in video surveillance. In *Multimedia Signal Processing (MMSP), 2012 IEEE 14th International Workshop on*, pages 378–382. Ieee, 2012. 6
- [36] P. Korshunov and T. Ebrahimi. Using face morphing to protect privacy. In *Advanced Video and Signal Based Surveillance (AVSS), 2013 10th IEEE International Conference on*, pages 208–213. IEEE, 2013. 6
- [37] P. Korshunov, A. Melle, J.-L. Dugelay, and T. Ebrahimi. Framework for objective evaluation of privacy filters. In *SPIE Optical Engineering+ Applications*, pages 88560T–88560T. International Society for Optics and Photonics, 2013. 2, 6
- [38] T. Koshimizu, T. Toriyama, and N. Babaguchi. Factors on the sense of privacy in video surveillance. In *Proceedings of the 3rd ACM workshop on Continuous archival and retrieval of personal experiences*, pages 35–44. ACM, 2006. 6
- [39] N. Kumar and I. Benbasat. Research note: the influence of recommendations and consumer reviews on evaluations of websites. *Information Systems Research*, 17(4):425–439, 2006. 3
- [40] K. Lander, V. Bruce, and H. Hill. Evaluating the effectiveness of pixelation and blurring on masking the identity of familiar faces. *Applied Cognitive Psychology*, 15(1):101–116, 2001. 2, 6
- [41] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016. 2
- [42] A. Li, Q. Li, and W. Gao. Privacycamera: Cooperative privacy-aware photographing with mobile phones. In *Sensing, Communication, and Networking (SECON), 2016 13th Annual IEEE International Conference on*, pages 1–9. IEEE, 2016. 2, 5
- [43] C. Li, A. C. Loui, and T. Chen. Towards aesthetics: A photo quality assessment and photo selection system. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 827–830. ACM, 2010. 5
- [44] Y. Li, N. Vishwamitra, H. Hu, B. Knijnenburg, and K. Caine. Effectiveness and users’ experience of face blurring as a privacy protection for sharing photos via online social networks. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 61. SAGE Publications, 2017. 3
- [45] Y. Li, N. Vishwamitra, B. Knijnenburg, H. Hu, and K. Caine. Effectiveness and users’ experience of obfuscation as a privacy-enhancing technology for sharing photos. Manuscript submitted for publication, 2017. 4
- [46] J. Mao, W. Xu, Y. Yang, J. Wang, and A. L. Yuille. Explain images with multimodal recurrent neural networks. *arXiv preprint arXiv:1410.1090*, 2014. 1
- [47] R. McPherson, R. Shokri, and V. Shmatikov. Defeating image obfuscation with deep learning. *arXiv preprint arXiv:1609.00408*, 2016. 2
- [48] C. A. Meissner and J. C. Brigham. Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review., 2001. 6
- [49] M. Mitchell, A.-I. Wang, and P. Reiher. Cashtags: Prevent leaking sensitive information through screen display. In *Proceedings of the USENIX Security Symposium*, 2015. 1
- [50] K. B. Murray and G. Häubl. Freedom of choice, ease of use, and the formation of interface preferences. 2010. 3
- [51] S. J. Oh, R. Benenson, M. Fritz, and B. Schiele. Faceless person recognition: Privacy implications in social media. In *European Conference on Computer Vision*, pages 19–35. Springer, 2016. 2
- [52] J. R. Padilla-López, A. A. Chaaoui, F. Gu, and F. Flórez-Reuelta. Visual privacy by context: proposal and evaluation of a level-based visualisation scheme. *Sensors*, 15(6):12959–12982, 2015. 6
- [53] M. R. Phillips, B. D. McAuliff, M. B. Kovera, and B. L. Cutler. Double-blind photoarray administration as a safeguard against investigator bias. *Journal of Applied Psychology*, 84(6):940, 1999. 3
- [54] J. P. Redden. Reducing satiation: The role of categorization level. *Journal of Consumer Research*, 34(5):624–634, 2008. 3
- [55] M. Saini, P. K. Atrey, S. Mehrotra, and M. Kankanalli. W3-privacy: understanding what, when, and where inference channels in multi-camera surveillance video. *Multimedia Tools and Applications*, 68(1):135–158, 2014. 2, 5
- [56] P. Seddon and M.-Y. Kiew. A partial test and development of delone and mclean’s model of is success. *Australasian Journal of Information Systems*, 4(1), 1996. 3
- [57] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1
- [58] A. C. Squicciarini, D. Lin, S. Sundareswaran, and J. Wede. Privacy policy inference of user-uploaded images on content sharing sites. *IEEE transactions on knowledge and data engineering*, 27(1):193–206, 2015. 1
- [59] A. C. Squicciarini, M. Shehab, and F. Paci. Collective privacy management in social networks. In *Proceedings of the 18th international conference on World wide web*, pages 521–530. ACM, 2009. 1

- [60] Y. Tanaka, A. Kodate, Y. Ichifuji, and N. Sonehara. Relationship between willingness to share photos and preferred level of photo blurring for privacy protection. In *Proceedings of the ASE BigData & SocialInformatics 2015*, page 33. ACM, 2015. 6
- [61] K. Thomas, C. Grier, and D. M. Nicol. unfriendly: Multi-party privacy risks in social networks. In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 236–252. Springer, 2010. 1
- [62] L. Tran, D. Kong, H. Jin, and J. Liu. Privacy-cnh: A framework to detect photo privacy with convolutional neural network using hierarchical features. In *AAAI*, pages 1317–1323, 2016. 1
- [63] U.S. Census Bureau. American factfinder - race results, 2010. Available at https://factfinder.census.gov/faces/tableservices/jsf/pages/productview.xhtml?pid=DEC_10_DP_DPDP1&src=pt. 3
- [64] M. V. Venkatesh, S.-c. S. Cheung, and J. Zhao. Efficient object-based video inpainting. *Pattern Recognition Letters*, 30(2):168–179, 2009. 6
- [65] D. Venkatraman. Why blurring sensitive information is a bad idea, 2007. 2
- [66] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):652–663, 2017. 1
- [67] N. Vishwamitra, Y. Li, K. Wang, H. Hu, K. Caine, and G.-J. Ahn. Towards pii-based multiparty access control for photo sharing in online social networks. In *Proceedings of the 22nd ACM on Symposium on Access Control Models and Technologies*. ACM, 2017. 2
- [68] E. von Zezschwitz, A. De Luca, and H. Hussmann. Filter selection and evaluation. 2015. 6
- [69] D. Walton. Plausible deniability and evasion of burden of proof. *Argumentation*, 10(1):47–58, 1996. 5
- [70] J. Wickramasuriya, M. Datt, S. Mehrotra, and N. Venkatasubramanian. Privacy protecting data collection in media spaces. In *Proceedings of the 12th annual ACM international conference on Multimedia*, pages 48–55. ACM, 2004. 6
- [71] K. Xu, Y. Guo, L. Guo, Y. Fang, and X. Li. My privacy my decision: Control of photo sharing on online social networks. *IEEE Transactions on Dependable and Secure Computing*, 2015. 1
- [72] C. Yan, M. Tien, and J. Wu. Interactive background blurring. In *Proceedings of the 17th ACM international conference on Multimedia*, pages 817–820. ACM, 2009. 2, 5
- [73] YouTube Creator Blog. Blur moving objects in your video with the new custom blurring tool on youtube, February 2016. Available at <https://youtube-creators.googleblog.com/2016/02/blur-moving-objects-in-your-video-with.html>. 2
- [74] YouTube Official Blog. Face blurring: when footage requires anonymity, July 2012. Available at <https://youtube.googleblog.com/2012/07/face-blurring-when-footage-requires.html>. 2
- [75] X. Yu, K. Chinomi, T. Koshimizu, N. Nitta, Y. Ito, and N. Babaguchi. Privacy protecting visual processing for secure video surveillance. In *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, pages 1672–1675. IEEE, 2008. 6
- [76] L. Zhang, T. Jung, C. Liu, X. Ding, X.-Y. Li, and Y. Liu. Pop: Privacy-preserving outsourced photo sharing and searching for mobile devices. In *Distributed Computing Systems (ICDCS), 2015 IEEE 35th International Conference on*, pages 308–317. IEEE, 2015. 2
- [77] L. Zhang, K. Liu, X.-Y. Li, P. Feng, C. Liu, and Y. Liu. Enable portrait privacy protection in photo capturing and sharing. *arXiv preprint arXiv:1410.6582*, 2014. 6
- [78] W. Zhang, S. Cheung, and M. Chen. Hiding privacy information in video surveillance system. In *Image Processing, 2005. ICIP 2005. IEEE International Conference on*, volume 3, pages II–868. IEEE, 2005. 6