Viewer Experience of Obscuring Scene Elements in Photos to Enhance Privacy

Rakibul Hasan[1], **Eman Hassan**[1], **Yifang Li**[2],
Kelly Caine[2], **David J. Crandall**[1], **Roberto Hoyle**[3], **Apu Kapadia**[1]

[1] Indiana University, Bloomington, IN, USA, {rakhasan, emhassan, djcran, kapadia}@indiana.edu
[2] Clemson University, Clemson, SC, USA, {yifang2, caine}@g.clemson.edu
[3] Oberlin College, Oberlin, OH, USA, hoyle@oberlin.edu

ABSTRACT

With the rise of digital photography and social networking, people are sharing personal photos online at an unprecedented rate. In addition to their main subject matter, photographs often capture various incidental information that could harm people's privacy. While blurring and other image filters may help obscure private content, they also often affect the utility and aesthetics of the photos, which is important since images shared in social media are mainly for human consumption. Existing studies of privacy-enhancing image filters either primarily focus on obscuring faces, or do not systematically study how filters affect image utility. To understand the trade-offs when obscuring various sensitive aspects of images, we study eleven filters applied to obfuscate twenty different objects and attributes, and evaluate how effectively they protect privacy and preserve image quality for human viewers.

ACM Classification Keywords

K.4.1. Public Policy Issues: Privacy; H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

Privacy; Image Obfuscation; Image Filtering

INTRODUCTION

Sharing images has become popular on social media sites and applications such as Flickr, Instagram, Snapchat, Facebook, and WhatsApp — by one estimate, more than 1.8 billion photos are posted to popular social media systems each day.[1] Many of these images are shared despite the presence of private elements within the photo (e.g., an embarrassing facial expression or sensitive information visible on a computer screen), while other images may not be shared because of sensitive content that people prefer to keep hidden [30]. We seek to help people improve sharing decisions through image transforms that address privacy at an *element level*, rather than forcing people to take an all-or-nothing *image level* approach in sharing or withholding images online.

Transforming elements within an image present a trade-off between privacy and utility; these transformations need to be aggressive enough to remove or obscure private information, but not so aggressive that they destroy the value of sharing the image. For example, transformations such as blurring and pixelation can be used to redact portions of an image [5, 7, 14, 25, 30, 38]. However, much of this work does not consider the potential negative impact on image aesthetics or utility, and much of it focuses on obfuscating faces or bodies, not on various other scene elements that may also raise privacy concerns (e.g., monitors and financial documents). While some work considers how these transformations affect the user experience [13, 30] or studies particular transformations of objects [15], we believe a systematic study is needed on how well various transformations balance concealing private content with preserving image value for a human viewer.

Of course, what is considered private is highly individual and context-dependent, and further motivates why it is important to understand how various transforms obscure or reveal various elements. For example, in some cases, one may want to *conceal* the identity of a person in a photo but preserve some of their properties such as facial expression, gender, and skin color (e.g., documenting an angry protest while providing participants with anonymity). In other situations, one may want a filter that *reveals* the identity but conceals other properties (e.g., to obscure an embarrassing facial expression).

In this work we examine how obfuscating 'objects' affects various 'attributes' of those objects that a viewer can perceive. We present the findings of an experimental study conducted on Amazon Mechanical Turk[2] (N=570) on the effects of five different transforms (*masking*, *blurring*, *pixelation*, *edge detection*, and *silhouetting*) on both privacy and user experience (including visual aesthetics and satisfaction) for scenarios that previous studies have identified as important for privacy [1, 2, 9, 14, 18, 26, 46]. We find that it is possible to protect selected regions within an image while preserving utility and aesthetics. We also find that different filters work better at protecting different attributes within images, and provide quantitative information to guide future applications.

[1] https://www.dailydot.com/debug/mary-meeker-photo-report/

[2] https://www.mturk.com/

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

 ISBN 978-1-4503-5620-6/18/04…$15.00

DOI: https://doi.org/10.1145/3173574.3173621

RELATED WORK

Various approaches for protecting privacy in images by access control and information-content restriction have been proposed (Li et al. provide a review of such work [30]). PuPPIeS and P3 [16, 40] allow users to encrypt parts of images that contain sensitive information before sharing via social networking sites or storing in the cloud. POP [50] blurs, masks, and encrypts sensitive image regions before uploading to cloud servers, and allows users to search images in a privacy-preserving manner. These approaches focus on encryption techniques to cryptographically 'lock' sensitive regions of images and 'unlock' them for authorized users, but do not address the user experience of unauthorized users — e.g., what the transformed images look like aesthetically or what other information they reveal.

In many cases, images are meant for public consumption, and sensitive parts need to be redacted for privacy concerns with no intended authorized viewers of the redacted parts. For example, YouTube provides a feature that blurs faces in videos [6], while Google Street View obscures faces and license plates to avoid identity leakage [12]. In the context of remote collaboration via live video feed, Boyle et al. [7] studied how blurring and pixelating affect privacy and awareness, and Hudson et al. proposed techniques such as representing people's movement using dark pixels overlaid on a static image to reduce privacy risks while keeping information required for the collaboration [20]. Other work studies different forms of face de-identification [5, 14, 25, 38] for privacy protection. Our work fits within this general class of solutions that seeks to redact parts of images for wider consumption; these systems, however, focus primarily on faces and bodies, whereas we consider a host of other scenarios and objects. Furthermore, we aim to study which attributes are selectively revealed or hidden by different transforms. As mentioned earlier, in some cases, one may want to *reveal the identity* of a person, but not their facial expressions, whereas in other cases one may want to hide the identity but reveal other characteristics (e.g., gender and race [22, 45]).

Other work has considered threat models in which computer vision may be employed to attack privacy through techniques such as face detection [31], recognition [48], "hallucination" [21], completion [27], and attribute manipulation [44]. These computer vision advances introduce extra challenges for privacy protection techniques like obfuscations. The work of Brkic et al. [8] has shown that some obfuscation techniques can be defeated by neural network-based attacks. Hill et al. [17] built a Hidden Markov Model based system that can recover text from blurred and pixelated documents. McPherson et al. [36] used deep learning algorithms to correctly identify faces and recognize objects and handwritten digits even after they were blurred, pixelated, or encrypted with P3 [40]. Meanwhile, automated object detection and recognition [33, 41] and visual question answering [3, 34, 35] techniques can detect and recognize objects with a high degree of accuracy and answer questions about objects' attributes, such "what is written on this street sign?" or "what is the breed of this dog?" [3]. Our work addresses the related problem of what the transforms reveal to humans, and whether certain attributes can be meaningfully conveyed to humans despite the transformations. When computer vision-based adversaries are considered, our work still provides insight into what is revealed (or not) to humans when protections against computer vision algorithms are applied. In the end, these images must retain some utility for human viewers, while simultaneously addressing privacy concerns from the perspective of human viewers.

We believe more research is needed to study image obfuscation in social media, where humans are the primary consumers, and this warrants a human-centric evaluation. Gross et al. [13] demonstrated that for human faces, blurring and pixelating often either do not obscure enough details to provide adequate privacy, or obscure so much that they destroy the utility of the video. To understand these trade-offs, Li et al. [30] studied a set of obfuscations (blurring, pixelating, masking, avatar, inpainting, etc.) on both human faces and bodies, and evaluated person-identification accuracy and participants' perceptions towards obfuscated images. Our work extends theirs by studying a wider array of scene characteristics that may be considered sensitive, such as personal belongings, affiliations, computer monitor contents, and other private information [19, 49]. Directly related to our focus, Hassan et al. [15] study 'cartooning' transforms on images in order to conceal sensitive scene elements but still convey certain characteristics. Our work fills a gap by studying the effects of various transforms on different scene elements to understand the relative trade-offs in various application scenarios.

EXPERIMENT

We conducted an experiment to study the effectiveness of several image obfuscation methods (see Table 1) [3] designed to conceal objects and different properties of them, as well as how well these obfuscation methods retain image utility. We included twenty different scenarios in which we varied objects and their properties, as described in Table 2. Each of these scenarios had one of twelve different conditions, each using a different method and/or degree of obfuscation. Participants were randomly assigned to one of the twelve filter conditions (between subjects). Each participant was then presented with all 20 scenarios in random order (within subjects).

Measurements

In the experiment we asked five questions for each scenario:

What is the object (or property of the object) depicted in the image? This question varied slightly based on the scenario; in Table 2 we summarize the specific questions we used. Participants were asked to select from multiple-choice options consisting of the most common answers given in the pilot study, which had a free-form text box. Answers were marked either correct or incorrect. A green bounding box was overlaid surrounding the objects of interest to ease locating them only for this question, and later removed for subsequent questions (as described below) for the same scenario.

How confident do you feel that you correctly answered the previous question? This question used a 7-point Likert scale.

[3] We did not include actual photos that were used in the survey in Table 1 due to copyright issues. To obtain the photos please contact one of the authors.

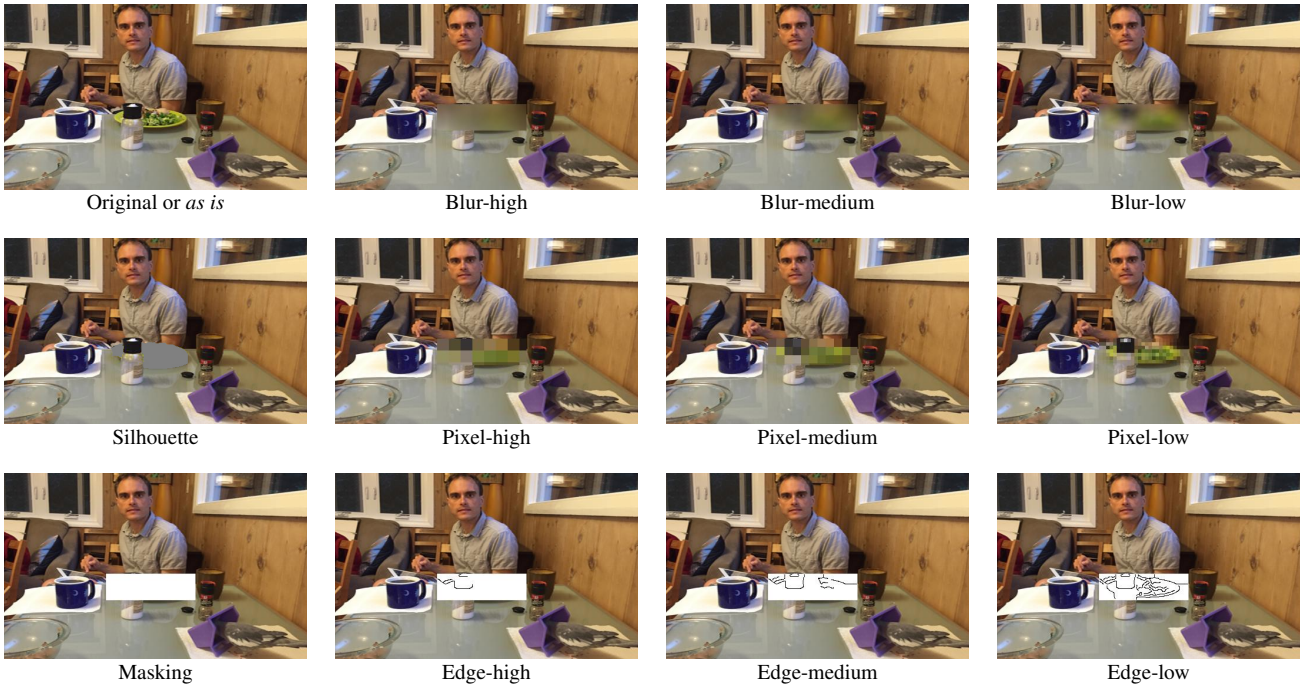

Table 1: Results of applying different filters to obscure food.

For the next three questions, we asked the participants whether they agreed or disagreed with the following statements.

The photo provides sufficient information. This item (also on a 7-point Likert scale) is adapted from the 'information quality scale' [43], which measures "the satisfaction of users who directly interact with the computer for a specific application" [11]. We adapted "Does the photo provide sufficient information," which loads onto the "content" factor and was strongly correlated with questions "is the system successful?" and "are you satisfied with the system?"

The photo is satisfying. We adapted this item from the validated 'image appeal scale' [10], which is the extent to which images are perceived as "appropriate and aligned to user expectations, satisfying, or interesting... and goes beyond aesthetics or the attractiveness." Specifically, this selected item measures the participants' overall 'satisfaction' with the image after the alterations, as also measured by Li et al. [30] when obscuring faces and bodies. A 7-point Likert was used.

This photo looks visually appealing. To frame this item, we asked participants to *"Imagine a friend of yours shares this photo on a social networking site, such as Facebook,"* and was also measured on a 7-point Likert scale.

Scene selection

Our scenarios are representative of the objects and properties about which privacy concerns were expressed in prior studies [1, 2, 9, 14, 18, 26, 46]. Through these scenarios we capture peoples' concerns related to privacy of information (e.g., leaking text from financial documents or computer screens), impression management based on appearance (e.g., facial expression, hair style), activities (e.g., using social media during work hours), and living conditions (e.g., messy room, eating habits).

Obfuscation Methods

Along with the *as is* (unaltered) condition as a control, we used five primary types of obfuscations: *Blurring*, *Pixelating*, *Edge* (i.e., line drawing), *Masking*, and *Silhouette*. This selection was informed by prior studies according to the appropriateness for the research questions we seek to answer. Earlier studies on *blurring* and *pixelating* were limited primarily to facial identity protection [28, 29], and found that these filters are well accepted by users but not effective when applied at a level that preserves photo utility [13, 39]. We thus wanted to determine their effectiveness to conceal other objects and properties. *Masking* has been found to be effective to protect identity but hides masked photo content completely [30]; we study its effect when applied on objects that are small or not the main subject matter of the photo. *Silhouette* is interesting because it preserves shape, which we hypothesized may be useful to retain an object's identity but remove finer details that might contain private information. On the other hand, *edge* preserves shape and some internal details and may be useful in cases where finer control is required.

While the *masking* and *silhouette* filters are binary, either completely obscuring the original object or not, the other three have continuous-valued filter parameters. Applying *blur* and *pixelating* filters with low parameter values generates output images that are similar to the originals, while increasing

Scenario	Question
Activity	What is the person inside the green rectangle doing?
Age	What is the age of the person inside the green rectangle?
Document class	What is the object inside the green rectangle?
Document text	What is the text inside the green rectangle?
Document type	What type of document (e.g. notebook, paper) is inside the green rectangle?
Dress	What type of clothing is the person inside the green rectangle wearing?
Ethnicity	What is the ethnicity of the person inside the green rectangle?
Expression	What is the facial expression of the person inside the green rectangle?
Food	What type of food is inside the green rectangle?
Gender	What is the gender of the person inside the green rectangle?
Hair	How long is the hair of the person inside the green rectangle?
Indoor	Was the following photo taken indoor or outside?
Indoor specific	What type of indoor place (e.g., library, concert hall) is shown in the following photo?
Laundry	What is the object inside the green rectangle?
Messy room	How well organized or messy is the place shown in the photo?
Monitor app.	What application is displayed on the computer monitor inside the green rectangle?
Monitor class	What is the object inside the green rectangle?
Monitor text	What is the text inside the green rectangle?
Outdoor	Was the following photo taken indoors or outside?
Outdoor specific	What type of outdoor place (e.g., field, street) is shown in the following photo?

Table 2: Scenarios and the recognition questions used in the survey.

values cause the filtered image regions to be more aggressively obscured. The *edge* filter parameter controls a threshold on edge strength, with higher values removing all but the strongest lines while lower values retain more detail.

These leveled filters might be effective in obscuring different types of information at different parameter values. For example, blurring with an aggressive filter value may be able to completely obscure an object such as a computer monitor, whereas blurring with a milder value might only obscure details (e.g., text on a monitor screen) but not the object itself. To study these effects, we included the *masking*, *silhouette*, and *blur* filters with 'high,' 'medium,' and 'low' levels of aggressiveness in our experiment. These values were determined through a smaller user study (more details are provided in the supplementary materials) in which we developed a tool and showed different images at different levels to participants. For each respondent, the filter levels were decreased until he or she was able to determine the identity of the object, and in turn (for the next level) detect lower-lever features in the image. The levels across all participants and images were averaged. The 'high' level was chosen as the average level for high-level details plus one standard deviation, and likewise 'low' was equal to the average for low-level details plus one standard deviation. The 'medium' level was chosen as the average for high-level details. These three filters with three levels each, along with the *masking*, *silhouette*, and *as-is* (no filtering), make up the twelve obfuscation methods in our study, and are summarized in Table 1.

Collecting Images

For each scenario we used different image sets, so that any image for one scenario would not reveal answers about any other scenario. Each set consisted of ten images collected from the internet. Using more than one image for each scenario allows us to incorporate some controlled variability and draw more useful conclusions from the study than for a single image. At the same time, we were careful to select images that had consistent image properties, such as brightness and object size, in each scenario. In particular, we tried to follow the following guidelines as closely as possible:

1. The quality, illumination, and size should be as consistent as possible across all images.

2. For any particular scenario, all five images should have a similar number of people and/or other objects with similar distribution and orientation.

3. For any particular scenario, the object of interest (e.g. face) should be of comparable size across all five images.

4. The object of interest should not be the focus of the image; e.g., when monitors are the object of interest, the monitor should not be in the center or 'too large' compared to other objects in the image. We are interested in cases where information is leaked through objects that are not the main subject matter and may go unnoticed.

5. It should not be possible to easily identify the object or property of interest from scene context, such as other objects or properties (e.g., computer monitors next to adjacent keyboards, type of indoor place like library from collection of bookshelves, food from the logo of the restaurant or other food in the vicinity, and so on).

From these images, we further sub-sampled the five images for each scenario which were most consistent with the guidelines. Finally, we scaled the images to be consistently sized. Our

pilot study did not reveal any systematic large differences in identification accuracy for any specific image in a scenario.

Organization of the Survey

The survey instrument was organized as follows:

1. Consent form.
2. Questions about which (if any) social media services the participant uses, how frequently they share images using those services, and four demographic questions.
3. Instructions on how to answer the survey questions along with a sample image either in *as is* condition, or a filter that was randomly selected and applied on a predefined region.
4. Twenty scenarios, presented in a random order (within subjects), each with five questions in a specified order. Each scenario presented one of five random images modified by one of the twelve obfuscation methods (between subjects — each participant was assigned to a single transform condition, e.g., *Blur-medium*, was selected at random and fixed for the participant).

The experiment was implemented in Qualtrics and is included as supplementary material.

Ethical Considerations

The study was approved by the Institutional Review Board at Indiana University.

Recruitment, Compensation, and Validation

The study was advertised on Amazon Mechanical Turk as an "Image Transformation Study." Participants were required to live and have resided in the United States for at least five years, in order to reduce cultural variability [24]. To ensure higher data quality [37], we restricted to MTurk workers with high reputation (above 95% approval rating on at least 1000 completed HITs). They were also required to be at least 18 years of age; studying photo obfuscation preferences and experiences of teenagers could be an interesting direction for future work. The average time to complete the survey was around 20 minutes, and respondents were compensated US$2.50 upon completion of the study. We paid all 725 respondents who completed the study, but eliminated participants from our sample if they failed any of the three attention-check questions, leaving 570 participants in our final sample.

Pilot Study

We first performed a pilot study with N=45 respondents, also administered via Amazon Mechanical Turk, but respondents were compensated $3.00. Data from this pilot study was used to estimate the sample size required to produce statistically significant findings through a power analysis. Moreover, the top five free-form text responses for the recognition questions were used as the multiple-choice options (instead of a text field) in the final study. We acknowledge the concern that providing a fixed number of choices can make picking the correct option easier than answering correctly in free-form text. In the pilot, however, we found that participants were already using contextual information present in the photos and for any particular question the number of different replies were less than ten. Furthermore, our experimental setup provides insights through the relative changes observed across conditions.

For each scenario, we used the most common response in the *as is* condition during the pilot study as the *correct answer* in the final study. The pilot also helped us to test for unforeseen variability within our images which might lead users to misidentify the objects of interest, but did not find any.

FINDINGS

Demographic Information

Among the 570 participants, 324 (56.8%) described themselves as male, and 172 (30.1%) were non-white. 197 (34.6%) were aged between 18–29 years, 303 (53.2%) between 30–49 years, 63 (11.1%) between 50–64 years, and 7 (1.2%) above 65 years. The highest level of education attained was high school for 203 (35.6%) participants, undergraduate degree for 306 (53.7%), Masters for 49 (8.68%), and Ph.D. or professional degree for 12 (2.08%). All participants reported using at least one social media service, and 293 (51.4%) reported sharing images using social media at least "a few times" a week.

Recognition Accuracy

In order to characterize how well filters obscure potentially private information we look at two metrics. First, we measured "accuracy" as the participants' ability to recognize objects and properties in transformed images by simply computing the fraction of correct responses. We analyzed the responses using Fisher's exact test, where we compared the accuracy of each filter with the accuracy of the *as is* condition, and present the results (recognition rate, p-value, and effect size) in Table 3. We applied the Bonferroni correction for these tests (i.e., for each row (filter) of the table, we corrected for 11 hypothesis tests against the baseline filter). Next, we also looked at the effect size to measuring the effectiveness of the filter over the as-is baseline. For Fisher's exact test, the effect size is the ratio of the odds of being correct in a treatment condition (i.e., filter) to the odds of being correct in the control condition (i.e., *as is*), so lower effect sizes correspond to lower odds of being correct when a filter is applied. In our case, this helps us determine how effectively the filter prevents recognition. We designate filters as *effective* when the recognition accuracy is less than 50% and the effect size is less than 0.05,[4] and *somewhat effective* when accuracy is less than 50% and effect size less than 0.1. For example, with 50% and 95.2% recognition accuracies for filtered and unfiltered conditions respectively, the odds ratio is 0.05, indicating that *effective* filters drastically lower the odds of recognition success.

In general, we observed that *blurring*, *pixelating*, and *edge* filters at low and medium levels are effective in protecting

[4]A threshold for recognition accuracy in the filtered condition is required to get a meaningful effect size. Otherwise when accuracy is 100% for the unfiltered condition, the odds ratio will be zero for any accuracy in the filtered condition and the filter will appear as *effective* even if it fails to prevent recognition (i.e., high recognition accuracy for the filtered object). Also we select these values to ensure that for effective filters, the recognition probability is less than 50% chance and without any filter the recognition probability is close to certainty (100%)

specific or minor details, such as text, but fail to obscure general properties such as whether an object is a document or monitor. These filters are almost always ineffective even in the strongest levels for scenarios that require obscuring the entire image. In contrast, *masking* is effective at obscuring objects (as well as almost all other scenarios) and *silhouette* is mostly effective for objects and attributes that cannot be recognized from shape (e.g., *ethnicity*). Below we describe the findings for each filter in more detail.

Blurring

We found that *blurring* at a low level is only effective in obscuring *activity*, *gender*, *document type*, *document text*, and *monitor text* (Table 3). In addition, mid-level blurring can prevent recognition of *monitor application* and *specific indoor environment*. On the other hand, a high level of *blurring* is effective in all scenarios except *expression*, *monitor class*, *general outdoor environment*, and *messy room*, and for *ethnicity*, *specific outdoor environment*, and *food*, it is only somewhat effective. In summary, *blurring* is not effective at protecting properties related to *Environment*, food, and laundry (specially at low and medium levels), but effective for other scenarios at high and medium levels.

Pixelating

High and medium levels of *pixelating* perform similarly to corresponding levels of *blurring* across all scenarios except for human attributes (e.g., *facial expression*, *dress*) where *pixleating* seems better than *blurring* (Table 3). On the other hand, a low level of *pixlating* is only effective for *activity*, *document text*, and *monitor text*, and performs worse than a low level of *blurring* in other cases. But a low level of *pixelating* preserves more information and generates more visually appealing photos compared to *blurring* (and other filters) as discussed in the section on Photo Utility, and so might be more desirable than other filters when effective.

Edge

Similar to *blurring* and *pixelating*, the *edge* filter becomes more effective as the filter parameter becomes more aggressive. However, unlike the other two, a high level *edge* filter is at least somewhat effective for all the scenarios related to *document* and *computer monitor* (Table 3). *Edge* is also effective at obscuring *food* at both high and medium levels, and effective for *laundry* even at a low level. In short, *edge* seems to be more effective than *blur* and *pixelate* when the object to be obfuscated has irregular shape and/or internal texture that produces *noise-like curves* when the filter is applied.

Silhouette and Masking

Silhouette and *masking* filters are similar in the sense that they completely remove or replace the filtered region. But since *silhouette* preserves shape information, we expected it to be effective for objects and properties that cannot be recognized by boundaries. We found *silhouette* effective for all scenarios except *hair*, *monitor class*, and *food* (Table 3), which we expected, but also for *age*, which was surprising because high levels of *blurring* and *pixelating* are effective in this case. We believe this is because a person's body shape and posture can be strong cues for *age*, and *silhouette* reveals information about these. Another interesting finding is that *masking* fails to protect *facial expression* despite being effective in all other cases. This is because our definition of effectiveness required an effect size (ratio of recognition accuracies in filtered to unfiltered conditions) less than 0.05, but the accuracy on the *as-is* condition was already so low that *masking* did not create enough *additional* confusion to be considered effective. Note that we did not test environmental and text related scenarios with *silhouette* because it is not clear how this transform would be applied in these cases different from masking.

Recognition Confidence

In general, the mean confidence value for the *as is* condition was the highest for all scenarios as expected.[5] Next we analyzed participants' confidence levels separately for correct and incorrect answers. For incorrect identification, there were no significant differences in confidence levels for any filter across any scenario. When identified correctly, generally the mean confidence levels were higher than when identified incorrectly. Moreover, we found that confidence levels vary significantly for different filters for most of the scenarios, but interestingly, for difficult and/or confusing scenarios (such as *document text*, *expression*, and *hair*), we did not find any significant difference in confidence for any filter. This indicates the inherent ambiguity involved in identifying these properties of images and participants were not very confident about their (correct) identification.

Photo Utility

We next analyzed how well filters preserved the perceived *utility* of images, using the three questions from the survey on whether an image "provides sufficient information," "is satisfying," and "looks visually appealing."

Information Sufficiency

Viewers' perceptions of information sufficiency in an image is affected by obfuscation [30]. By performing an overall Kruskal-Wallis test for all conditions in each scenario, we found significant variations, but the actual H-statistic values differ for different scenarios, and for any particular scenario not all obfuscation methods have a significant effect on information sufficiency, as shown in Table 4. The least effective filters in terms of recognition accuracy also have the least damaging effect on information sufficiency. Conversely, filtered images that obscure sensitive information also tend to lack sufficient information even at low levels (such as *edge* for activity and text). Unsurprisingly, the highest level of these filters and *masking* significantly destroy information content in most of the scenarios. We examine the relationship between information sufficiency and recognition accuracy in more depth in the next section.

To allow us to draw more general conclusions, we categorized the scenarios into five groups: Human (activity, age, dress, expression, ethnicity, gender, hair), Monitor (monitor class, monitor application, monitor text), Document (document class, document type, document text), Environment (indoor, indoor

[5]For *specific indoor environment*, *monitor text*, and *specific outdoor environment*, the highest values were for *masking*, *pixelate-medium*, and *pixelate-low* respectively, although the differences with the *as is* condition were not statistically significant.

		Original	*Masking*	*Blur*			*Pixelate*			*Edge*			*Silhouette*
				high	medium	low	high	medium	low	high	medium	low	
Human	Activity	97%	$8\%^{***}_{H}$	$15\%^{***}_{H}$	$6\%^{***}_{H}$	$16\%^{***}_{H}$	$14\%^{***}_{H}$	$11\%^{***}_{H}$	$48\%^{***}_{H}$	$14\%^{***}_{H}$	$35\%^{***}_{H}$	$64\%^{***}_{N}$	$10\%^{***}_{H}$
	Age	88%	$17\%^{***}_{H}$	$26\%^{***}_{H}$	$36\%^{***}_{M}$	$62\%^{**}_{N}$	$42\%^{***}_{M}$	$42\%^{***}_{M}$	$85\%^{*}_{N}$	$37\%^{***}_{M}$	$35\%^{***}_{M}$	$50\%^{***}_{N}$	$50\%^{***}_{N}$
	Dress	100%	$28\%^{***}_{H}$	$44\%^{***}_{H}$	$51\%^{***}_{N}$	$81\%^{**}_{N}$	$44\%^{***}_{H}$	$38\%^{***}_{H}$	$93\%^{*}_{N}$	$27\%^{***}_{H}$	$44\%^{***}_{H}$	$44\%^{***}_{H}$	$48\%^{***}_{H}$
	Ethnicity	88%	$17\%^{***}_{H}$	$28\%^{***}_{M}$	$59\%^{**}_{N}$	$65\%^{*}_{N}$	$48\%^{***}_{N}$	$52\%^{***}_{N}$	$78\%^{*}_{N}$	$8\%^{***}_{H}$	$24\%^{***}_{H}$	$46\%^{***}_{N}$	$25\%^{***}_{H}$
	Expression	72%	$22\%^{***}_{N}$	$28\%^{***}_{N}$	$21\%^{***}_{N}$	$13\%^{***}_{M}$	$19\%^{***}_{M}$	$11\%^{***}_{M}$	$38\%^{**}_{N}$	$31\%^{***}_{N}$	$13\%^{***}_{M}$	$16\%^{***}_{M}$	$18\%^{***}_{M}$
	Gender	100%	$20\%^{***}_{H}$	$24\%^{***}_{H}$	$25\%^{***}_{H}$	$46\%^{***}_{H}$	$28\%^{***}_{H}$	$33\%^{***}_{H}$	$72\%^{***}_{N}$	$37\%^{***}_{H}$	$26\%^{***}_{H}$	$52\%^{***}_{N}$	$43\%^{***}_{H}$
	Hair	52%	$2\%^{***}_{H}$	$20\%^{**}_{N}$	$21\%^{**}_{N}$	$20\%^{**}_{N}$	$12\%^{***}_{N}$	$7\%^{***}_{M}$	$34\%^{*}_{N}$	$14\%^{***}_{N}$	$22\%^{**}_{N}$	$12\%^{***}_{N}$	$33\%^{*}_{N}$
Document	Document class	86%	$13\%^{***}_{H}$	$17\%^{***}_{H}$	$42\%^{***}_{N}$	$27\%^{***}_{M}$	$26\%^{***}_{M}$	$40\%^{***}_{N}$	$65\%^{*}_{N}$	$27\%^{***}_{M}$	$24\%^{***}_{M}$	$48\%^{***}_{N}$	$10\%^{***}_{H}$
	Document type	97%	$8\%^{***}_{H}$	$6\%^{***}_{H}$	$23\%^{***}_{H}$	$25\%^{***}_{H}$	$12\%^{***}_{H}$	$14\%^{***}_{H}$	$72\%^{**}_{N}$	$33\%^{***}_{H}$	$26\%^{***}_{H}$	$54\%^{***}_{N}$	$45\%^{***}_{H}$
	Document text	91%	$0\%^{***}_{H}$	$0\%^{***}_{H}$	$2\%^{***}_{H}$	$4\%^{***}_{H}$	$0\%^{***}_{H}$	$2\%^{***}_{H}$	$21\%^{***}_{H}$	$16\%^{***}_{H}$	$11\%^{***}_{H}$	$16\%^{***}_{H}$	—
Monitor	Monitor class	100%	$42\%^{***}_{H}$	$68\%^{***}_{N}$	$57\%^{***}_{N}$	$88\%^{*}_{N}$	$60\%^{***}_{N}$	$71\%^{***}_{N}$	$89\%^{*}_{N}$	$45\%^{***}_{H}$	$55\%^{***}_{N}$	$64\%^{***}_{N}$	$81\%^{**}_{N}$
	Monitor app.	88%	$15\%^{***}_{H}$	$22\%^{***}_{H}$	$14\%^{***}_{H}$	$34\%^{***}_{M}$	$35\%^{***}_{M}$	$45\%^{***}_{N}$	$74\%^{*}_{N}$	$14\%^{***}_{H}$	$28\%^{***}_{M}$	$34\%^{***}_{M}$	$9\%^{***}_{H}$
	Monitor text	94%	$0\%^{***}_{H}$	$0\%^{***}_{H}$	$0\%^{***}_{H}$	$0\%^{***}_{H}$	$0\%^{***}_{H}$	$2\%^{***}_{H}$	$0\%^{***}_{H}$	$0\%^{***}_{H}$	$0\%^{***}_{H}$	$0\%^{***}_{H}$	—
Environment	Indoor general	97%	$6\%^{***}_{H}$	$35\%^{***}_{H}$	$57\%^{***}_{N}$	$83\%^{*}_{N}$	$42\%^{***}_{H}$	$59\%^{***}_{N}$	$91\%^{*}_{N}$	$50\%^{***}_{N}$	$73\%^{**}_{N}$	$76\%^{**}_{N}$	—
	Indoor specific	94%	$2\%^{***}_{H}$	$13\%^{***}_{H}$	$31\%^{***}_{H}$	$65\%^{**}_{N}$	$23\%^{***}_{H}$	$52\%^{***}_{N}$	$91\%^{*}_{N}$	$16\%^{***}_{H}$	$55\%^{***}_{N}$	$72\%^{**}_{N}$	—
	Outdoor general	100%	$6\%^{***}_{H}$	$66\%^{***}_{N}$	$91\%^{*}_{N}$	$97\%^{*}_{N}$	$67\%^{***}_{N}$	$90\%^{*}_{N}$	$93\%^{*}_{N}$	$58\%^{***}_{N}$	$95\%^{*}_{N}$	$96\%^{*}_{N}$	—
	Outdoor specific	80%	$4\%^{***}_{H}$	$24\%^{***}_{M}$	$42\%^{***}_{N}$	$72\%^{*}_{N}$	$16\%^{***}_{H}$	$23\%^{***}_{M}$	$80\%^{*}_{N}$	$20\%^{***}_{M}$	$48\%^{**}_{N}$	$58\%^{*}_{N}$	—
	Messy room	58%	$0\%^{***}_{H}$	$17\%^{***}_{N}$	$23\%^{**}_{N}$	$27\%^{**}_{N}$	$30\%^{**}_{N}$	$21\%^{***}_{N}$	$57\%^{*}_{N}$	$18\%^{***}_{N}$	$51\%^{*}_{N}$	$46\%^{*}_{N}$	—
Other	Laundry	94%	$17\%^{***}_{H}$	$26\%^{***}_{H}$	$48\%^{***}_{M}$	$48\%^{***}_{M}$	$21\%^{***}_{H}$	$40\%^{***}_{H}$	$57\%^{***}_{N}$	$31\%^{***}_{H}$	$28\%^{***}_{H}$	$44\%^{***}_{H}$	$39\%^{***}_{H}$
	Food	91%	$22\%^{***}_{H}$	$37\%^{***}_{M}$	$48\%^{***}_{M}$	$41\%^{***}_{M}$	$37\%^{***}_{M}$	$42\%^{***}_{M}$	$76\%^{*}_{N}$	$33\%^{***}_{H}$	$20\%^{***}_{H}$	$50\%^{***}_{N}$	$57\%^{***}_{N}$

Table 3: Recognition accuracy for different filters across different scenarios. Recognition accuracies are shown as percentages, while subscripts and colors indicate whether each filter is effective (H), somewhat effective (M), or not effective (N) in preventing recognition, and asterices indicate significance: * is significant at $p < .05$, ** is significant at $p < .001$, and *** is significant at $p < 0.001$, after Bonferroni correction.

specific, outdoor, outdoor specific, messy room), and Other (laundry, food). Figure 1 presents mean responses for the information sufficiency question (in terms of the 7-point Likert scale) for each filter and scenario group. We noticed that for scenarios where only small portions of images are obfuscated, all filters have comparable mean values (Figure 1). For Human properties, *pixel-low* has the highest mean value among all filters, followed by *blur-low*. For Document, all levels of *blur* and *pixel* (except *pixel-low*) along with *masking* have lower values than the average value of the scale (3.5), while *silhouette* and *edge-low* have values close to *as is*.This is probably due to the fact that documents have rigid shapes which are better preserved by *silhouette* and *edge* filters compared to others. For monitor attributes, *pixel-low* and *silhouette* retain more information compared to others, while for environment scenarios where we obfuscate the whole image, we observe large differences in mean values of *pixel-low* and *edge-low* compared with others. In summary, the weakest filters (e.g. *pixelate-low*) preserve the most information, and information content is inversely proportional to the filter strength, conforming to prior studies [30], and is proportional to the area of the filtered region.

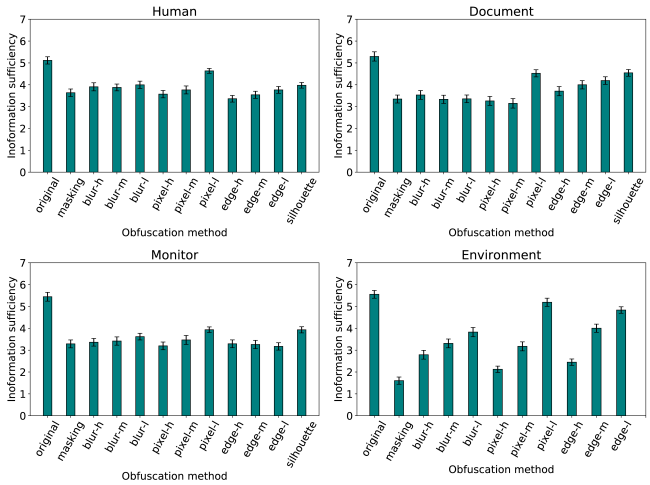

Figure 1: Information sufficiency across scenario groups and filters, in terms of mean values and standard error.

Photo Satisfaction and Visual Aesthetics

We observe that less aggressive (and thus often less effective) filters such as *blur low* and *pixelate low* generate images that are more satisfactory and visually appealing. However, satisfaction and aesthetics also depend on the size of the obfuscated region. For full image obfuscation (such as indoor/outdoor environment) and full human body obfuscation (such as dress and ethnicity), both satisfaction and aesthetics are hampered.

		Masking	*Blur*			*Pixel*			*Edge*			*Silhouette*
			High	Medium	Low	High	Medium	Low	High	Medium	Low	
		P I S V	P I S V	P I S V	P I S V	P I S V	P I S V	P I S V	P I S V	P I S V	P I S V	P I S V
Human	Activity	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✗✗✓✓	✓✗✓✓
	Age	✓✓✓✓	✓✓✓✓	✗✓✓✓	✗✓✓✓	✗✗✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✗	✗✓✓✓
	lDress	✓✗✓✓	✓✓✓✓	✗✓✗✓	✗✓✗✗	✓✗✗✓	✓✗✓✓	✗✓✓✓	✓✗✗✗	✓✗✗✗	✓✗✗✗	✓✓✓✓
	Ethnicity	✓✗✗✗	✗✗✗✗	✗✓✓✗	✗✓✗✗	✗✗✓✓	✗✓✓✓	✗✓✓✓	✓✗✗✗	✓✗✗✗	✗✗✗✗	✓✗✗✓
	Expression	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✗	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓
	Gender	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✗✓	✓✗✗✓	✓✗✓✓	✗✓✓✓	✓✗✗✓	✓✗✗✓	✗✗✓✓	✓✗✓✓
	Hair	✓✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✗✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✗✗	✗✓✓✓
Document	Document class	✓✗✓✓	✓✗✓✓	✗✗✓✓	✗✓✓✓	✗✗✓✓	✗✗✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✗✓✓✓	✓✓✓✓
	Document type	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✗✓	✓✗✓✓	✓✗✓✓	✗✓✓✓	✓✗✓✓	✓✗✓✓	✗✓✓✓	✓✓✓✓
	Document text	✓✗✓✓	✓✗✓✓	✓✗✗✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✗✗	—
Monitor	Monitor class	✓✗✓✓	✗✓✓✓	✗✗✓✓	✗✓✓✓	✗✗✓✓	✗✗✓✓	✗✓✓✓	✓✗✓✓	✗✗✓✓	✗✗✓✓	✗✓✗✓
	Monitor application	✓✗✓✓	✓✗✓✓	✓✗✓✓	✗✗✓✓	✗✗✓✓	✗✗✓✓	✗✓✓✓	✓✗✓✓	✗✗✓✓	✗✗✓✓	✓✗✓✓
	Monitor text	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✗✓	—
Environment	Indoor general	✓✗✗✗	✓✗✗✗	✗✗✗✗	✗✗✗✗	✓✗✗✗	✗✗✗✗	✗✓✓✗	✗✗✗✗	✗✓✓✓	✗✓✓✓	—
	Indoor specific	✓✗✗✗	✓✗✗✗	✓✗✗✗	✗✗✗✗	✓✗✗✗	✗✗✗✓	✗✓✓✓	✓✗✗✗	✗✗✗✗	✗✓✓✓	—
	Outdoor general	✓✗✗✗	✗✗✗✗	✗✗✗✗	✗✓✗✗	✗✗✗✗	✗✗✗✗	✗✓✓✓	✗✗✗✗	✗✓✓✗	✗✓✓✗	—
	Outdoor specific	✓✗✗✗	✗✗✗✗	✗✗✗✗	✗✓✗✗	✓✗✗✗	✗✗✗✗	✗✓✓✓	✗✗✗✗	✗✗✗✓	✗✓✗✗	—
	Messy room	✓✗✗✗	✗✗✓✓	✗✗✓✓	✗✗✓✓	✗✗✗✗	✗✗✓✓	✗✓✓✓	✗✗✗✓	✗✗✓✓	✗✓✓✓	—
Other	Laundry	✓✗✓✓	✓✗✓✓	✗✓✓✓	✗✗✓✓	✓✓✓✓	✓✗✓✓	✗✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓	✓✗✓✓
	Food	✓✗✗✗	✗✗✓✓	✗✗✓✗	✗✗✓✗	✗✗✗✗	✗✗✗✗	✗✓✓✓	✓✗✗✗	✓✗✗✗	✗✗✗✗	✗✗✓✓

Table 4: Privacy and utility trade-offs. For each filter, a green checkmark or red cross indicates whether that filter 1) protects privacy (i.e. recognition accuracy $< 50\%$ and odds-ratio < 0.05) (P), 2) provides sufficient information (I), 3) creates a satisfactory image (S), and 4) creates a visually appealing image (V).

Interestingly, while satisfaction and visual appeal are highly correlated (0.66 correlation), information sufficiency is much less correlated with both of these (correlations 0.44 and 0.25), suggesting that reduced information is not necessarily always accompanied by lower satisfaction (as we discuss in the next section). We also observe similar mean values across filters for these two measures both for individual scenarios and grouped scenarios, so we only include the plot of photo satisfaction of grouped scenarios, in Figure 2. Similarly, we study the relationship of recognition accuracy with only visual aesthetics in detail in the next section.

Privacy-Utility Trade-off

Figure 3 visualizes the trade-off between obscuring sensitive information and retaining image utility, using scatter plots of information sufficiency (y-axis) versus recognition accuracy (x-axis) for grouped scenarios. We see a roughly linear, positive correlation between detection accuracy and information sufficiency. This suggests that viewers of an image perceive it to be lacking information when they fail to recognize objects or properties of interest in the image. For groups *Human*, *Document*, and *Monitor*, we see clusters of filters in the left region of the plots. We find that *blur-high* for *Human*, and *silhouette* for all categories except *Environment* might strike the best balance between privacy and perceived information sufficiency. For *Environment,* where the whole image is obfuscated, the points form a diagonal line, indicating a clear trade-off between privacy protection and information content of images. In this case, a medium level of *blur* and *pixelate* provides a reasonable balance between recognition accuracy and the amount of information retained in obfuscated photos.

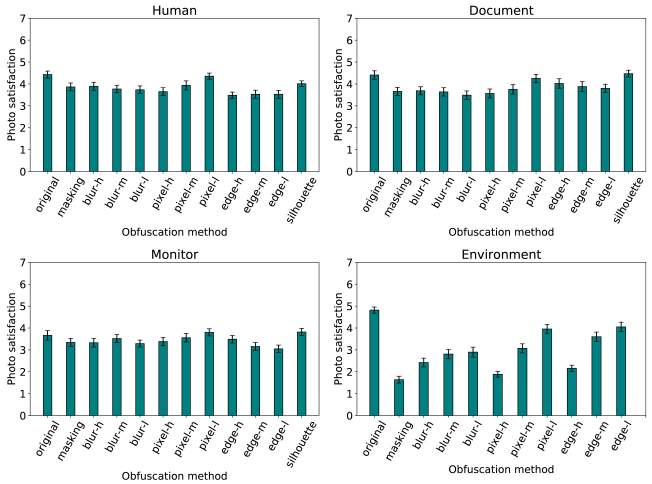

Figure 2: Photo satisfaction across scenario groups and filters, in terms of mean values and standard error.

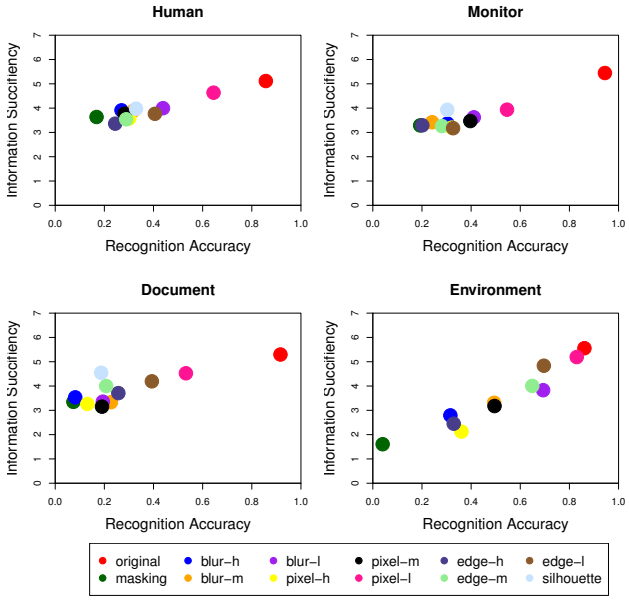

Figure 3: Trade-off between protecting against information leaks and information sufficiency across filters, in terms of recognition accuracy (x-axis) and mean information sufficiency (y-axis). Note that *Silhouette* was not studied for any property related to *Environment*.

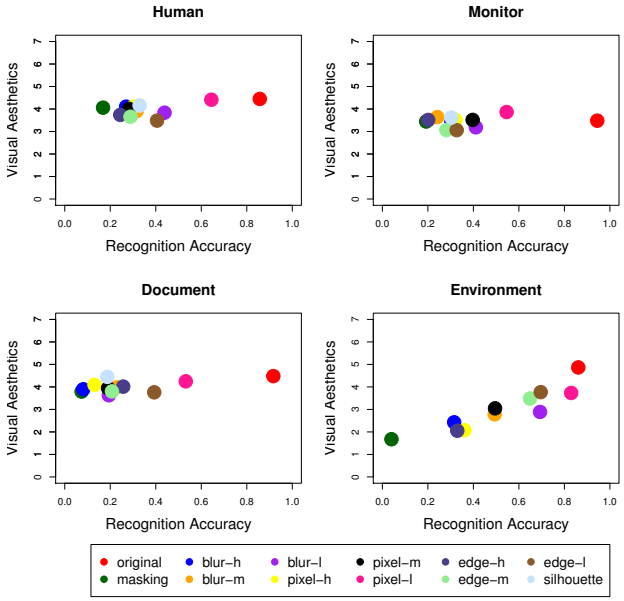

Figure 4: Trade-off between protecting against information leaks and aesthetics across filters, in terms of recognition accuracy (x-axis) and mean visual aesthetics (y-axis). Note that *Silhouette* was not studied for any property related to *Environment*.

Photo satisfaction and visual aesthetics were closely correlated (0.66 correlation), so we only discuss visual aesthetics. Figure 4 compares recognition accuracy and visual aesthetics. We see that for *Environment* the filters are distributed diagonally, meaning that there is a clear trade-off between privacy and visual aesthetics. But for other scene categories there are filters that both protect privacy and leave the filtered image visually appealing: such as *silhouette* for all categories; *pixelate-high* and *blur-medium* for *Monitor*; *blur-high*, *pixelate-high*, and surprisingly, *masking* for *Human*.

DISCUSSION

We now examine the implications, future work possibilities, and limitations of our study.

Privacy vs. Utility

In general, our findings are in line with earlier work [7, 30]: stronger filters increase perceived privacy and decrease perceived information content, satisfaction, and aesthetics. This is especially true for scenarios with specific answers (e.g., dress and gender) or when the whole photo is filtered. However, when a filtered object is small and/or not an integral part of the scene but nevertheless potentially privacy sensitive (e.g., a document), perceived information content and visual aesthetics remain high. This indicates that enhancing the privacy of images does not always result in a reduced user experience. For example, at one extreme, silhouetting objects provides complete privacy for all object attributes other than the type of object, but results in high scores for visual aesthetics. We also found that weaker filters and levels (such as *blur-low*) have little effect on obscuring people, monitors, and documents across a range of situations, again confirming prior findings. This demonstrates that all filters are not equivalent, and different solutions may be more appropriate for different user needs and content types.

Effectiveness of Filters Throughout Categories

The effectiveness of obscuring information for the leveled filters is highly correlated with the specificity of the information that the filter is intended to obscure. At their most aggressive levels, these filters can prevent leaking major details (such as the photo environment or gender of a person), but at medium

and low levels are effective only in protecting minor details and specific information (e.g., text or age). On the other hand, since *silhouette* preserves the shape of object boundaries but redacts everything else, we expect it to fail to protect information leakage only when the information can be inferred from the shape of the boundary (such as *food* and *monitor class*). For objects with rigid boundaries, *silhouette* is as effective as *masking*, which is the most effective filter we found.

For subjective and difficult scenarios (as indicated by low recognition accuracy in the *as-is* condition in Table 3) such as *facial expression*, *age*, *messy room*, and *hair length*, all filters seem to be less effective than scenarios with straightforward answers (such as *text*). But note that effect size is a relative measure with respect to the *as-is* condition, so that low baseline accuracy worsens the effect size, meaning applying any filter does not add much confusion.

Edge Detection Side Effects

The *edge* filter behaves differently than the other leveled filters: while *blur* and *pixelate* produce an image very similar to the original at their lowest levels, *edge* always produces a binary "edge map" (as shown in Table 1). At its most aggressive, *edge* shows only the most prominent lines, and as the parameter is decreased, progressively weaker lines are revealed. Intuitively, the edge map contains more information for lower values of the parameter, but in some cases, detection accuracy actually *decreased* for lower parameter values (e.g., gender, hair). We speculate that in some scenarios, such as *document type* and *monitor type*, applying the edge filter at a high level leaves the obfuscated region with lines that amplify prominent rectangular objects that are distinctive of these objects. Meanwhile, the abundance of distracting edges at lower filter values makes it more difficult to correctly identify objects. In effect, *edge* applied with a low parameter actually increases noise, and can make it harder to infer information when the filtered regions have too much detail.

Implications and Practical Applications

We expect that our work will shed light on how to transform elements within an image to preserve privacy. Our work shows, as one might expect, that there is no 'one size fits all' filter for obscuring scene elements. Depending on the application, different objects can be obscured with custom filters, and our work makes the first step at trying to characterize how different filters applied at varying levels affect what is concealed and revealed about objects. These findings may improve user acceptability and privacy protection applications such as transforming real-time video streams [15] by selecting the transformation type in an object-dependent way. Our work also offers insight for mobile applications such as VizWiz [4], which allow people with visual impairments to take photos of their environment and ask questions about it to crowd workers, social media friends, or automated applications. While these applications have tremendous potential to help people with visual impairments, there are also severe privacy risks, since users do not necessarily know what their photos contain. Our findings provide a way to transform images so that only the image elements required to answer a particular question are retained. Finally, photos shared via social media can be privacy sensitive for their owner and/or bystanders, and our findings can be integrated into privacy preserving image sharing frameworks such as PuPPIeS [16], and combined with proposed methods to automatically detect sensitive contents in photos [47].

Human vs. Computer Viewers

As discussed in the Introduction, this work does not consider computer-vision based attacks. While certain types of transforms can be defeated by computer vision better than humans, other transforms (such as those applied to CAPTCHAs) defeat computer vision algorithms but not humans. Our work considers human viewers of images and our findings can be interpreted in conjunction with research on computer vision based attacks, based on the application and adversary model, in particular considering whether or not information needs to be revealed to human viewers and the impact of the transforms on human experience. Future work can further study the trade-offs of computer-vision based adversaries.

Limitations

We speculate that filters covering only background or foreground elements or of different sizes may exhibit different results. We made attempts to control for this by making sure that the main, centered, foreground object was not the one that was filtered, and that the filtered area was not so big so as to occlude most of the image or so small that it was hard to spot. However, we have not systematically studied how the size or location of obscured regions within a photo affects how they are perceived; this is a worthwhile direction for future work.

Another limitation is that we only compare the performance of each filter against an *as is* baseline, as opposed to other myriad possible comparisons. Due to the number of conditions in our study, we struck a balance between the resources needed and the number of insights that could be drawn with sufficient statistical significance.

Finally, this study was conducted on Amazon Mechanical Turk, whose user demographics are not representative of the general U.S. population and are known to be more privacy conscious [23, 42]. Nevertheless, we attempt to measure the information loss through objective measures. Other studies have shown that such crowd platforms are a reasonable choice for studying user experience [32].

CONCLUSIONS

Our work sheds light on the effects of applying various types of image transforms to scene elements in an image. In particular we studied the relative trade-offs between privacy (revealing and concealing selective attributes of objects) and utility (the visual aesthetics and user satisfaction of the image) of five different image transforms and show that while in some cases a clear privacy vs. utility trade-off is realized, in other scenarios a high degree of privacy can be attained while retaining utility. Our work also contributes significantly to the existing literature by examining these trade-offs for a range of objects and their attributes, whereas previous work had focused largely on obscuring people and faces. We hope our work spurs further

research on studying the relative trade-offs of image transformations for enhanced privacy without (significantly) degrading the user experience of the viewers.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation (under grants CNS-1408730, IIS-1253549, and IIS-1527421), Google, and NVidia. We would like to thank Dr. Brad Luen, Dept. of Statistics, Indiana University and Dr. JangDong Seo of the Indiana Statistical Consulting Center for their valuable suggestions regarding the statistical analysis of our experiment data. We also thank the volunteer participants of our lab study. Finally, we are grateful to the people who participated in the online study.

REFERENCES

1. 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
2. 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
3. Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*. 2425–2433.
4. Jeffrey P. Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C. Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samual White, and Tom Yeh. 2010. VizWiz: Nearly Real-time Answers to Visual Questions. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, New York, NY, USA, 333–342. DOI: http://dx.doi.org/10.1145/1866029.1866080
5. Dmitri Bitouk, Neeraj Kumar, Samreen Dhillon, Peter Belhumeur, and Shree K. Nayar. 2008. Face Swapping: Automatically Replacing Faces in Photographs. *ACM Trans. Graph.* 27, 3, Article 39 (Aug. 2008), 8 pages. DOI:http://dx.doi.org/10.1145/1360612.1360638
6. YouTube Official Blog. 2012. Face blurring: when footage requires anonymity. Blog. (18 July 2012). Retrieved April 13, 2017 from https://youtube.googleblog.com/2012/07/face-blurring-when-footage-requires.html.
7. Michael Boyle, Christopher Edwards, and Saul Greenberg. 2000. The Effects of Filtered Video on Awareness and Privacy. In *ACM Conference on Computer Supported Cooperative Work (CSCW '00)*. ACM, New York, NY, USA, 1–10. DOI: http://dx.doi.org/10.1145/358916.358935
8. Karla Brkic, Ivan Sikiric, Tomislav Hrkac, and Zoran Kalafatic. 2017. I Know That Person: Generative Full Body and Face De-Identification of People in Images. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE, 1319–1328.
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. William J Doll and Gholamreza Torkzadeh. 1988. The measurement of end-user computing satisfaction. *MIS quarterly* (1988), 259–274.
12. Arturo Flores and Serge Belongie. 2010. Removing pedestrians from google street view images. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*. IEEE, 53–58.
13. Ralph Gross, Edoardo Airoldi, Bradley Malin, and Latanya Sweeney. 2006. Integrating Utility into Face De-identification. In *International Conference on Privacy Enhancing Technologies (PET'05)*. Springer-Verlag, Berlin, Heidelberg, 227–242. DOI: http://dx.doi.org/10.1007/11767831_15
14. Ralph Gross, Latanya Sweeney, Jeffrey Cohn, Fernando Torre, and Simon Baker. 2009. *Protecting Privacy in Video Surveillance*. Springer London, London, Chapter Face De-identification, 129–146. DOI: http://dx.doi.org/10.1007/978-1-84882-301-3_8
15. Eman T Hassan, Rakibul Hasan, Patrick Shaffer, David Crandall, and Apu Kapadia. 2017. Cartooning for Enhanced Privacy in Lifelogging and Streaming Videos. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE, 1333–1342.
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. Steven Hill, Zhimin Zhou, Lawrence Saul, and Hovav Shacham. 2016. On the (in) effectiveness of mosaicing and blurring as tools for document redaction. *Proceedings on Privacy Enhancing Technologies* 2016, 4 (2016), 403–417.
18. Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014a. Privacy Behaviors of Lifeloggers using Wearable Cameras. In *Proceedings of the ACM International Joint*

Conference on Pervasive and Ubiquitous Computing (UbiComp). 571–582. DOI: http://dx.doi.org/10.1145/2632048.2632079

19. Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014b. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 571–582.
20. Scott E. Hudson and Ian Smith. 1996. Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (CSCW '96)*. ACM, New York, NY, USA, 248–257. DOI: http://dx.doi.org/10.1145/240080.240295
21. Junjun Jiang, Ruimin Hu, Zhongyuan Wang, and Zhen Han. 2014. Face super-resolution via multilayer locality-constrained iterative neighbor embedding and intermediate dictionary learning. *IEEE Transactions on Image Processing* 23, 10 (2014), 4220–4231.
22. A. Jourabloo, X. Yin, and X. Liu. 2015. Attribute preserved face de-identification. In *2015 International Conference on Biometrics (ICB)*. 278–285. DOI: http://dx.doi.org/10.1109/ICB.2015.7139096
23. Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. 2014. Privacy Attitudes of Mechanical Turk Workers and the U.S. Public. In *Symposium On Usable Privacy and Security (SOUPS '14)*. 37–49.
24. RM Khan and MA Khan. 2007. Academic sojourners, culture shock and intercultural adaptation: A trend analysis. *Studies About Languages* 10 (2007), 38–46.
25. P. Korshunov and T. Ebrahimi. 2013. Using face morphing to protect privacy. In *Advanced Video and Signal Based Surveillance (AVSS), 2013 10th IEEE International Conference on*. IEEE Computer Society, Krakow, Poland, 208–213. DOI: http://dx.doi.org/10.1109/AVSS.2013.6636641
26. 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
27. Yijun Li, Sifei Liu, Jimei Yang, and Ming-Hsuan Yang. 2017a. Generative Face Completion. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
28. Yifang Li, Nishant Vishwamitra, Hongxin Hu, Bart P. Knijnenburg, and Kelly Caine. 2017b. Effectiveness and Users' Experience of Face Blurring as a Privacy Protection for Sharing Photos via Online Social Networks. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 61, 1 (2017), 803–807. DOI: http://dx.doi.org/10.1177/1541931213601694
29. Yifang Li, Nishant Vishwamitra, Bart P Knijnenburg, Hongxin Hu, and Kelly Caine. 2017c. Blur vs. Block: Investigating the effectiveness of privacy-enhancing obfuscation for images. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*. IEEE, 1343–1351.
30. Yifang Li, Nishant Vishwamitra, Bart P. Knijnenburg, Hongxin Hu, and Kelly Caine. 2018. Effectiveness and Users' Experience of Obfuscation as a Privacy-Enhancing Technology for Sharing Photos. *Proceedings of the ACM: Human Computer Interaction (PACM)* (2018).
31. Shengcai Liao, Anil K Jain, and Stan Z Li. 2016. A fast and accurate unconstrained face detector. *IEEE transactions on pattern analysis and machine intelligence* 38, 2 (2016), 211–223.
32. Di Liu, Randolph G. Bias, Matthew Lease, and Rebecca Kuipers. 2012. Crowdsourcing for usability testing. *Proceedings of the American Society for Information Science and Technology* 49, 1 (2012), 1–10. DOI: http://dx.doi.org/10.1002/meet.14504901100
33. Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*. Springer, 21–37.
34. Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical question-image co-attention for visual question answering. In *Advances In Neural Information Processing Systems*. 289–297.
35. Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. 2015. Ask your neurons: A neural-based approach to answering questions about images. In *Proceedings of the IEEE international conference on computer vision*. 1–9.
36. Richard McPherson, Reza Shokri, and Vitaly Shmatikov. 2016. Defeating Image Obfuscation with Deep Learning. *CoRR* abs/1609.00408 (2016). http://arxiv.org/abs/1609.00408
37. Adam W Meade and S Bartholomew Craig. 2012. Identifying careless responses in survey data. *Psychological methods* 17, 3 (sep 2012), 437–455. DOI: http://dx.doi.org/10.1037/a0028085
38. Y. Nakashima, T. Koyama, N. Yokoya, and N. Babaguchi. 2015. Facial expression preserving privacy protection using image melding. In *Multimedia and Expo (ICME), 2015 IEEE International Conference on*. IEEE Computer Society, Torino,Italy, 1–6. DOI: http://dx.doi.org/10.1109/ICME.2015.7177394
39. Carman Neustaedter, Saul Greenberg, and Michael Boyle. 2006. Blur Filtration Fails to Preserve Privacy for Home-based Video Conferencing. *ACM Trans. Comput.-Hum. Interact.* 13, 1 (March 2006), 1–36. DOI: http://dx.doi.org/10.1145/1143518.1143519
40. 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.

41. Joseph Redmon and Ali Farhadi. 2017. YOLO9000: Better, Faster, Stronger. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
42. Joel Ross, Lilly Irani, M. Six Silberman, Andrew Zaldivar, and Bill Tomlinson. 2010. Who Are the Crowdworkers?: Shifting Demographics in Mechanical Turk. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems (CHI EA '10)*. ACM, New York, NY, USA, 2863–2872.
43. 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
44. Wei Shen and Rujie Liu. 2016. Learning residual images for face attribute manipulation. *arXiv preprint arXiv:1612.05363* (2016).
45. T. Sim and L. Zhang. 2015. Controllable Face Privacy. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, Vol. 04. 1–8. DOI: http://dx.doi.org/10.1109/FG.2015.7285018
46. 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
47. Lam Tran, Deguang Kong, Hongxia Jin, and Ji Liu. 2016. Privacy-CNH: A Framework to Detect Photo Privacy with Convolutional Neural Network Using Hierarchical Features. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*. AAAI Press, 1317–1323. http://dl.acm.org/citation.cfm?id=3015812.3016006
48. Luan Tran, Xi Yin, and Xiaoming Liu. 2017. Disentangled Representation Learning GAN for Pose-Invariant Face Recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
49. Nishant Vishwamitra, Yifang Li, Kevin Wang, Hongxin Hu, Kelly Caine, and Gail-Joon Ahn. 2017. 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, 155–166.
50. 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