Can Privacy Be Satisfying? On Improving Viewer Satisfaction for Privacy-Enhanced Photos Using Aesthetic Transforms

Rakibul Hasan¹, Yifang Li², Eman Hassan¹, Kelly Caine², David J. Crandall¹, Roberto Hoyle³, Apu Kapadia¹

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

Pervasive photo sharing in online social media platforms can cause unintended privacy violations when elements of an image reveal sensitive information. Prior studies have identified image obfuscation methods (e.g., blurring) to enhance privacy, but many of these methods adversely affect viewers' satisfaction with the photo, which may cause people to avoid using them. In this paper, we study the novel hypothesis that it may be possible to restore viewers' satisfaction by 'boosting' or enhancing the aesthetics of an obscured image, thereby compensating for the negative effects of a privacy transform. Using a between-subjects online experiment, we studied the effects of three artistic transformations on images that had objects obscured using three popular obfuscation methods validated by prior research. Our findings suggest that using artistic transformations can mitigate some negative effects of obfuscation methods, but more exploration is needed to retain viewer satisfaction.

CCS CONCEPTS

• **Human-centered computing** → *Collaborative and social computing theory, concepts and paradigms*;

KEYWORDS

Privacy; Image Obfuscation; Image Filtering

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300597>

ACM Reference Format:

Rakibul Hasan, Yifang Li, Eman Hassan, Kelly Caine, David J. Crandall, Roberto Hoyle, Apu Kapadia. 2019. Can Privacy Be Satisfying? On Improving Viewer Satisfaction for Privacy-Enhanced Photos Using Aesthetic Transforms. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*, May 4–9, 2019, Glasgow, Scotland Uk. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3290605.3300597>

1 INTRODUCTION

Photo sharing provides a natural mechanism for people to express themselves and interact with one another [19], and online social media has dramatically increased the volume of photo sharing activity [7, 8]. Such sharing has, in turn, led to a rise in accidental privacy violations [46], for example by revealing embarrassing moments of photos. To address this risk, social media users engage in self-censoring measures ranging from restricting their sharing, to controlling access to their photos through privacy settings, to withdrawing from social media platforms altogether [40, 41, 47]. People also have to exercise control offline to avoid sharing co-owned photos with undesired audiences [42] or embarrassment by photos captured and shared by other people [37]. Unfortunately, these measures seem to be as inadequate as pervasive photo sharing on social media continues to generate major privacy concerns [37, 42].

To reduce the privacy risks of online photo sharing, recent studies have proposed using privacy-enhancing image obfuscations to obscure sensitive *regions* of photos [14, 15, 25], while trying to preserve the 'utility' (i.e., the viewer's experience) of the photos. These studies identify a set of obfuscations that can effectively obscure objects (or their properties) in a photo while minimizing the impact on the viewer's overall "satisfaction."¹ However, the set of such useful obfuscations is relatively small; most obfuscations reduce the

¹As described later, "satisfaction," "information sufficiency," and "aesthetics" are measured as dependent variables based on questions derived from prior work [14, 25].

perceived “information sufficiency” or aesthetics to the point of negatively impacting people’s satisfaction in viewing the image. Since one of the primary motivations for sharing photos is to convey information and seek acceptance, appreciation, and validation from peers [30, 35], preserving the utility (satisfaction) of obfuscated images is important if obfuscation methods are to be widely accepted.

At a high level, obfuscation imposes a trade-off that is easy to understand: mild obfuscations may not negatively affect viewer satisfaction but may also not remove private image content effectively, while aggressive obfuscations may preserve privacy but cause obvious visual changes that reduce viewer satisfaction. Prior work identified three useful variables in measuring the viewer experience — information sufficiency, satisfaction, and aesthetics — and measured how obfuscations affect each of these variables in isolation [14, 25]. Those studies, however, do not examine the inter-relationships among these dependent variables. Further, in addition to direct effects, obfuscations might have cascading effects on these variables (i.e. affecting one variable through another). Understanding these relationships would greatly benefit in designing novel obfuscation methods that can improve privacy without adversely impacting viewers’ experience.

Using data from our previous study [14], we conduct a path-model based analysis which suggests that the effects of the obfuscations on information sufficiency and visual aesthetics are much greater than the direct effects on satisfaction, but information sufficiency and visual aesthetics are significantly associated with satisfaction. This observation inspires our novel hypothesis that *it may be possible to compensate for the reduction in information sufficiency from obfuscations by increasing visual aesthetics, thus actually maintaining or improving overall satisfaction of the obfuscated image*. This approach is illustrated in Figure 1, where an object within the photo may be redacted with pixelation while the rest of the image is aesthetically ‘improved’ using an artistic transformation, resulting in satisfaction similar to the original image.

To test our hypothesis, we conduct a new online experiment with three obfuscation and three beautification transformations across a variety of photo types and scenarios. The experiment follows a between-subjects design that extends our previous experiment [14] by adding the beautification condition. Thus our design seeks to ascertain a causal relationship between the manipulation of aesthetics and the viewer’s overall satisfaction with various obfuscations. From the photo aesthetics literature [34], we know that colors and tones play an important role in aesthetics: pure and high saturation colors tend to be more appealing to viewers than dull colors, for example [5]. We pick three particular beautifications to represent different levels of abstraction: (1) a

Figure 1: An example illustrating how obfuscation and beautification change the utility aspects of an image: (a) an image without any alteration, (b) the image after a pixel obfuscation, and (c) the image after a pixel obfuscation applied to the food plate and a cartoon beautification on the other parts of the image.

low-level abstraction using color correction [10] to produce an effect similar to highly popular Instagram filters [31], (2) a ‘cartooning’ effect similar to a watercolor painting that moderately changes the appearance of the original image, and (3) a deep learning-based algorithm to render the photo in an bright, colorful style, inspired by the popular Prisma app [23], that produces a highly abstract and unrealistic version of the image. We refer to these three beautification transformations as ‘colors’, ‘cartoons’, and ‘abstract’, respectively.

Our results verify interactions among information content, aesthetics, and satisfaction, confirming that it is worthwhile to investigate whether satisfaction can be increased by increasing the other two variables. Although the gain in satisfaction was not statistically significant for our sample data using off-the-shelf artistic transforms, we hope our findings will inspire designing new transforms taking into account the negative effects of privacy obfuscations.

2 RELATED WORK

Prior work on mechanisms to reduce privacy risks in images mainly fall into two broad categories: controlling access, and limiting information content (see Li et al. [25] for a detailed review).

Access control mechanisms allow only authorized users to access a photo in its original form, while others either see a fully or partially scrambled version, or do not have any access at all. These methods are, however, not always effective in social media contexts, since managing access control using the ever-changing privacy settings of social media platforms requires substantial time and effort [27, 29]. Even if users can manage their own privacy by limiting sharing or using access control mechanisms, photos taken in public places may pose privacy treats to people who are not subjects of the photo. Finally, these methods defeat one of the primary motivations for using social networking platforms — reaching out and making new connections [35].

Controlling photo content, such as obscuring part of a photo to protect privacy, has been widely studied and adopted by many existing applications. For example, people’s faces and vehicle license plates are blurred in Google Street View [13]. Blurring and pixelation are commonly used obfuscations in existing research and applications (e.g., [13, 17, 25]), but are often ineffective [14, 25]. From the viewer experience perspective, more effective obfuscations, such as silhouette, masking, and point-light, are unable to retain enough information, destroy visual aesthetics, and cause dissatisfaction for viewers [14, 25]. These utility variables affect photo owners’ willingness to adopt and photo viewers’ willingness to view an obscured photo, since a major motivation to use online social networks is information gathering [44]. Moreover, users tend to upload aesthetically-pleasing photos to their social networks in order to manage the impression that they leave on others [39]. Research has shown that more visually pleasing images are shared with larger crowds [19] by their owners, and also re-shared more frequently by viewers [45]. We thus need image transformations that are able to obscure parts of a photo without destroying important content or negatively affecting visual aesthetics. The balancing of privacy-utility trade-offs are hoped to inspire widespread adoption of these transformations, both as standalone privacy enhancing techniques, as well as complementing other techniques. For example, in the context of privacy conflicts around sharing photos with multiple owners, Such et al. found that an ‘all-or-none’ approach to resolve the conflicts is dominating [42]. While the privacy-concerned co-owner(s) of a photo might restrain from sharing it (i.e. adopting a self-regulation strategy), privacy violation might occur if any of the other co-owner(s) chose to share it with others. Sharing a transformed version of the photo will complement the privacy preserving behavior of the concerned owner(s) and preserve their privacy, at the same time will allow other co-owner(s) to enjoy the photo sharing activity with negligible reduction in utilities.

Of course, it is challenging to measure subjective properties of an image such as creativity, aesthetics and memorability. Isola et al. [18] define sets of images features such as colors, scene semantics, objects statistics, etc., and correlate them with image memorability. In the context of art analysis, Elgammal and Saleh [9] develop a computational model to estimate the creativity of a work of art based on historical context. Aydin et al. [1] analyze how perceived image aesthetics are related to lower-level image attributes, such as sharpness, depth, clarity, tone, and colorfulness, and develop a metric to quantify these values. Deep learning techniques that attempt to estimate aesthetic attribute values have also been introduced, as in Lu et al. [28] and Kao et al. [21] using the AVA dataset [33]. While useful for quantifying image aesthetics, these systems do not attempt to measure viewer

satisfaction or the effect of image transformations in the context of preserving privacy, as we do here.

A wide variety of techniques has been proposed for adding artistic effects to images, as well as for removing noise and correcting other defects. These effects are popular on social photo sharing platforms such as Instagram and Snapchat. In the academic literature, there is much work in computer vision, image processing, and computational photography on problems like creating panoramas [11], generating High Dynamic Range (HDR) images [20], and performing color correction and enhancement [10]. Recent work in deep learning-based image style transfer has created beautifications that mimic the style of particular artists [12, 48]. Other systems also attempt to generate images with specific artistic effects like cartooning [16, 24] or water-coloring [2]. We draw on several of these beautification techniques, studying them here in the context of how they affect perceived visual aesthetics and viewer satisfaction for obfuscated images.

3 METHOD

In earlier work, we studied the privacy-protecting and utility-preserving qualities of four obfuscation methods (i.e., image filters) [14]. There we applied these filters on people and other objects to obfuscate properties or attributes (such as the age of a person, the organization of a room) that were identified as privacy sensitive in prior work. We experimented with 20 attributes, and analyzed how effective each of the filters was in obscuring these attributes and how they affected the utility variables (i.e., information content, aesthetics, and viewers’ satisfaction). In this work we conducted additional analysis of that data using path models to study the inter-dependencies of the utility variables. The next two subsections describe the procedure and results of this analysis. We then provide details of our new experiment, which was inspired by the results of the path model analysis.

Path Model Analysis

We constructed separate path models using data from our previous experiment [14] for each of the 20 attributes (e.g., activity, gender, document class, document type). In these path models the exogenous variable was obfuscation type (such as blur and pixel) and the endogenous variables were information sufficiency, photo aesthetics, and photo satisfaction. We excluded data about identification accuracy and confidence from our model since we focus on the utility variables. For each attribute, we began with the initial model shown in Figure 2, and then trimmed insignificant effects.

In this graph representation the vertices represent variables, and arrows represent relationships between the variables. Further, the blue rectangular vertices are the exogenous variables (e.g., different obfuscations), and the orange

Figure 2: Initial path model.

ellipses are the dependent variables measured (e.g., information content). The directional edges in this graph express causal relationships, where changing the variable denoted by the starting vertex of an edge ‘causes’ a change in another variable denoted by the finishing vertex of the same edge (e.g., changing the obfuscation method causes a change in ‘satisfaction’). This model also captures indirect causal effects, such as obfuscation methods’ effects on ‘satisfaction’ through ‘information content’. The causal effects between the endogenous variables are speculative, but we describe the rationale for this particular arrangement of the vertices and the directions of the edges, e.g., why we think ‘information content’ causally affects ‘satisfaction’ and not the other way around.

An important motivation for people to use online social networks is to gather information, such as by observing other people’s photos [44]. This means that for high satisfaction, viewers need to be able to see important content (‘sufficient information’) in the photo. Aesthetics also contributes to satisfaction; in fact, users often edit their photos before sharing to improve aesthetics and to help control the impressions conveyed to others [39]. Our initial path model thus assumes there are causal relationships (and our experiment seeks to test such causality) from information sufficiency and aesthetics to satisfaction. For example, increasing aesthetics or information sufficiency may improve satisfaction when viewing the photo. From a photo composition perspective, we expect that displaying sufficient information improves photo aesthetics. Finally, based on previous work that shows that obfuscations affect information content sufficiency, photo aesthetics, and satisfaction in most scenarios, our initial model includes causal arrows from transformations to each dependent variable [14].

Path Model Results

As expected, our findings generally indicate that obfuscations have negative effects on information sufficiency, while their effects on aesthetics vary based on attribute types. For example, in the laundry scenario (Fig. 3), applying obfuscation has no effect on image aesthetics ($\chi^2(11) = 12.87$, $p = 0.30$). A possible explanation is that laundry is typically not

an appealing or important visual element, so obscuring it does not affect the aesthetics of the overall photo. Additionally, in half of the scenarios (age, document type and text, dress, ethnicity, expression, food, hair, indoor general and specific, and messy room), there is no direct effect on photo satisfaction by different types of obfuscations, although there are indirect effects mediated by information sufficiency and aesthetics.

For example, consider the path model for *dress* (Figure 3), which has a good model fit ($\chi^2(11) = 13.02$, $p = 0.29$, $CFI = 0.998$, $TLI = 0.994$, $RMSEA = 0.018$). Overall, there is a difference in information sufficiency between different transformation conditions ($\chi^2(11) = 106.36$, $p < 0.001$) compared to the baseline condition; most of the transformations (blur-medium, blur-high, pixel-medium, pixel-high, edge-low, edge-medium, edge-high, masking, and silhouette) decrease information sufficiency (all $p < 0.01$) while blur-low and pixel-low do not have any effect. On the other hand, obfuscations also have a generally negative effect on aesthetics ($\chi^2(11) = 31.79$, $p < 0.001$). Photos on which blur-low, edge-low, and edge-medium have been applied have lower aesthetics compared with the original photos (all $p < 0.05$). Meanwhile, information sufficiency appears to have a significant effect on aesthetics ($p < 0.001$), with a one-point difference in information sufficiency associated with a 0.491-point difference in aesthetics ($SE = 0.037$). Furthermore, aesthetics appears to positively affect satisfaction ($p < 0.001$), and information sufficiency also appears to have a direct effect on satisfaction ($p < 0.001$).

More generally, in all scenarios, controlling for manipulations, information sufficiency has a highly significant positive association with aesthetics (all $p < 0.001$). Additionally, aesthetics (all $p < 0.001$) and information sufficiency (all $p < 0.001$) have a direct positive association with satisfaction. These results indicate that increasing either information sufficiency or aesthetics may boost image satisfaction, and beautification on the remaining (non-obfuscated) part of the image could make up for the viewers’ satisfaction lost through obfuscation. To test this causal effect, we conducted a new online experiment; the design and methodology of this experiment is described in the following sections.

Experimental Design

In this experiment, we presented participants with photos that had been manipulated using various combinations of privacy-enhancing obfuscation and beautification transformations, and collected their ratings on utility variables. The privacy-enhancing obfuscations were applied on specific regions of a photo (to obscure attributes of people/objects) and the beautification transforms were applied on the rest of the photo. With 3 obfuscations and 3 beautification transforms, our study had 13 between-subjects experimental conditions

Figure 3: Example path model diagrams.

Privacy enhancing obfuscations	Beautification transformations
Masking	Abstract
Pixelation	Cartoon
Edge	Color

Table 1: Obfuscations and transformations used in this study. Each obfuscation was combined with each transformation, resulting in nine conditions. In addition, we included 3 obfuscation-only conditions, as well as a condition with the original, unaltered image, totaling 13 experimental conditions.

(3 obfuscations + 3 obfuscations $\times$ 3 beautifications + 1 unfiltered) (see Table 1). The baseline (i.e. unfiltered) condition included images without any alteration. The other conditions had only an obfuscation or an obfuscation combined with a beautification. Participants were randomly assigned to one of these conditions (between subjects), but each participant viewed images for all six object attributes (described below). Similar to our prior study [14], each participant viewed an image for each attribute and then answered five questions corresponding to the five dependent variables that we measured, as described below.

Participants

For our previous study [14], the number of participants per condition was calculated using a power analysis based on data from a pilot study. We planned a similar number of conditions and analysis for this new study, hence we used the same number of participants (48) for each condition. With thirteen conditions, we needed at least $13 \times 48 = 624$ participants in total. We advertised our experiment on Amazon

Mechanical Turk² and hosted it on Qualtrics³, restricting participation to MTurk workers with a high reputation (above 95% approval rating on at least 1000 completed HITs) to ensure data quality [32]. We also required workers to be at least 18 years old and living in the United States for at least five years to help control the cultural variability [22]. We included three attention check questions to maintain data quality [26]. After removing the responses from participants who provided wrong answers for one or more attention checks, we were left with 653 responses (out of a total of 780) that we used for analysis. Each participant was paid \$1.50, whether or not we used their response. The study was approved by Indiana University’s ethics board.

Selecting Attributes

From the set of twenty privacy-sensitive attributes used in our earlier experiment [14], we selected six to include in this study (see Table 2). We chose these six attributes to balance the size of the private image regions, since the sizes of obfuscated regions may otherwise vary dramatically depending on the size of the object to be obfuscated and/or the attribute itself. For example, we did not include any scenarios where the whole image needed to be obfuscated (e.g., hiding whether a photo was taken indoors or outdoors), since we wanted to study our hypothesis in the context of object obfuscations.

Image dataset

We used the same image set we previously used in [14], which allowed us to isolate and measure the effects of beautifications on the filtered images in this experiment. The

²<https://www.mturk.com>

³<https://www.qualtrics.com>

Attribute	Question
Document class	What is the object inside the green rectangle?
Dress	What type of clothing is the person inside the green rectangle wearing?
Gender	What is the gender of the person inside the green rectangle?
Laundry	What is the object inside the green rectangle?
Computer app.	What application is displayed on the computer monitor inside the green rectangle?
Monitor text	What is the text inside the green rectangle?

Table 2: The six attributes and corresponding detection questions used in the survey.

dataset contains sets of five images for each attribute, all collected from online sources. Care was taken to ensure that all images in each set were consistent with each other in terms of the number of objects and people, the shapes and sizes of these objects, the overall image quality and brightness, and the effort required to infer a certain attribute.

Privacy-enhancing Transformations and Artistic Transformations

We identified three main obfuscations: masking, pixelation, and edge. Previous work applied each of these transformations with three strength levels (high, medium, and low) and found that the ‘high’ level was most effective at obscuring sensitive attributes [14], so we use only that level here.

We chose three different beautifications that abstract scene content to different degrees. Our most conservative transformation, which we call *No-abstraction* or ‘Colors’, applies the color correction technique of Finlayson et al. [10], which modifies colors but does not affect the semantic content of the image. *Mid-abstraction* or ‘Cartoons’ applies a simple technique for “cartooning” the image, by applying bilateral filter-based blurring (Tomasi and Manduchi [43]), detecting edges from image gradients and highlighting them in black, and performing luminance quantization to 8 levels. This beautification abstracts some image content, since the blurring reduces resolution and the luminance quantization and added edges create an artistic effect. Finally, *Max-abstraction* or ‘Abstract’ applies deep-learning based artistic style transfer [48] for Henri Matisse’s famous painting *Woman with a hat*. Using artistic transforms that abstract photo content to compensate for lost information (due to the application of privacy obfuscations) might seem counter intuitive; we hypothesize that, since the abstraction happens at the global level, local information loss due to obfuscations may be less noticeable. Further, the abstracted form of the photos may help viewers absorb the high level story of a photo more easily, thus creating a sense of complete information.

Our experiment used 13 versions of each image as shown in Table 3: unaltered, obscured (3 versions), and obscured and beautified (3x3 versions), resulting in the thirteen experimental conditions. For privacy-enhancing transforms, we used

the same transformation size, position, and other parameters reported in [14]. For the obscured and beautified versions, we first applied the obfuscations on the specific image regions, and then the artistic transform to the rest of the image using one of the three beautifications. The obscured areas were not beautified to hold the degree of privacy constant when comparing the obscured version with the obscured and beautified version; otherwise a higher satisfaction score could be attributed to lower privacy through first obscuring and then beautifying a sensitive object.

Measurements

For each attribute, we asked five questions from two perspectives: obfuscation effectiveness and utility to the viewer. We note that all these questions and response options were adapted from our previous study [14].

- (1) **Identification.** Participants first saw an image with a green bounding box overlaid on an object of interest. They were asked to identify the object in the box by answering a multiple-choice question, “What is the object (or property of the object) depicted in the image?” The specific questions were slightly different based on the attribute, as shown in Table 2. For this question, we provided a list of options (including “Cannot tell”) to select from as an answer. The green bounding boxes surrounding the objects/attributes of interest were shown only in this question and not for the following ones.
- (2) **Identification Confidence.** Participants answered “How confident do you feel that you correctly answered the previous question?” on a seven-point Likert scale from 1 ‘Completely unconfident’ to 7 ‘Completely confident’ [36].
- (3) **Information Content Sufficiency.** We asked participants to rate their agreement with “The photo provides sufficient information,” on a 7-point Likert from ‘Strongly disagree’ to ‘Strongly agree.’ This item was adapted from the ‘information quality scale’ [6], which measures “the satisfaction of users who directly interact with the computer for a specific application.” Our

Table 3: Results of applying different obfuscations and beautifications.

item loads onto the “content” factor and is strongly correlated with “is the system successful?” [6]

- (4) **Visual Aesthetics.** To measure photo aesthetics, we used “This photo looks visually appealing” from the image appeal scale [4], again on a 7-point Likert scale.
- (5) **Satisfaction.** Similarly, “The photo is satisfying” was adapted from the image appeal scale [4], which has also been used when measuring satisfaction of face and body obfuscation [25]. This item measures participants’ overall satisfaction with the photo and again was rated on a 7-point Likert scale.

Procedure

The experiment flowed as follows:

- (1) Consent form detailing the experiment, estimated time to finish, and compensation.
- (2) Questions about social media usage and frequency of image sharing activities, along with demographics.
- (3) Instructions on how to respond to the survey questions with a sample image and questions.

- (4) Six blocks of questions corresponding to the six attributes, in random order. Each block showed the five questions corresponding to the five measurements for each attribute. One of the five photos for each attribute was randomly selected to be presented to the participant with the assigned condition (‘unaltered’, ‘obfuscated’, or ‘obfuscated plus beautified’).

Data Analysis Procedure

We used non-parametric versions for all of our statistical tests as our data do not meet the assumptions of parametric tests, such as normality and equal variance of errors. For each dependent variable (information content, visual aesthetics, satisfaction), we first conducted an overall Kruskal-Wallis test across all conditions to see if there was any significant difference in the measured variables among the conditions. We followed this with a Dunn’s post hoc test with Bonferroni correction, where we compared between specific pairs. For each dependent variable, we selected the pairs to compare as follows: for each of the three obfuscation conditions (masking, pixel, edge) was compared with the three corresponding *obfuscation plus beautification* conditions. Therefore, for each of the three obfuscations, we had three pairwise tests, for a total of nine. This set of pairwise tests allowed us to study whether combining beautification transforms with privacy

obfuscations increases the utility of photos. Next, we conducted additional pairwise tests to see how combinations of privacy obfuscations and beautification transforms preserve utility when compared with the original (i.e. unaltered) photos. To do this, for each of the three obfuscations, we picked one beautification transform that performed best (i.e. highest mean value of the measured variable) when combined with it, yielding three *obfuscation plus beautification* conditions. Then these *obfuscation plus beautification* conditions were compared with the *unfiltered* condition. This resulted in three additional comparisons, or twelve in total. We present results of these pairwise tests in the supplementary document, where, in addition to the test statistics, we report the Pearson's product moment correlation (r) [3].

As an example of the process, for the *dress* attribute and the information content dependent variable, we first conducted an overall Kruskal-Wallis test for any difference in information content across the experimental conditions. If the p-value was not significant, we did not conduct any follow-up. If the p-value was significant ($p < 0.05$), then there were significant differences involving at least two different conditions. To find the pairs of conditions having differences, we followed up with Dunn's post hoc test for pairs of *only obfuscation* and *obfuscation plus beautification*. For example, for the *masking* obfuscation, we compared the *masking* condition with each of *masking + abstract*, *masking + cartoon*, and *masking + color* applied on the *dress* attribute. Also, if for example *masking + abstract* retained more information among these three *obfuscation plus beautification* conditions, we compared it to the *unfiltered* condition for the same measured variable (i.e., information content). This setting allowed us to test the effects of beautifications on obfuscated images, and also study the behavior of obfuscation-beautification combinations compared with the *unfiltered* condition.

4 FINDINGS

We now present the results of our experiment.

Demographic Characteristics of the Participants

Out of 653 participants, 436 (66.7%) identified themselves as male and 216 (33%) as female. Our participants were typically under 49 years of age, with 351 (53.7%) between 18 and 29 years, 250 (38.3%) between 30 and 49 years, 54 (6.7%) between 50 and 64 years, and eight (1.2%) participants 65 years or older. Three hundred and thirty five (51.3%) participants were white, 152 (23.2%) were Asian, and 43 (6.5%) were black or African American. For the highest level of education, 320 (49%) participants reported an undergraduate degree, 172 (26.3%) high school, 145 (22.2%) a Master's degree, and 16 (2.4%) a professional degree. All participants reported having at least one social network account, while 512 (78.3%)

reported sharing photos online with frequency ranging from several times a day to a few times a week, and only 25 (3%) participants reported never posting photos online.

Effects of Transformations on Information Content

For all attributes, perceived information content was the highest for the *unfiltered* condition (Table 4). The *abstract* transform, when combined with privacy obfuscations resulted in the lowest information content for most of the attributes (Table 4). Surprisingly, the *color* transform, which alters the image content the least, reduced more information than the *cartoon* transform, which, when combined with *edge* and *pixelation* privacy obfuscations, actually increased perceived information content for most of the attributes (Table 4). We conducted an overall Kruskal-Wallis test and detected significant differences in perceived information content among different *obfuscated*, *obfuscated plus beautified*, and *unfiltered* conditions (*document*: $\chi^2(11) = 54.75$, *dress*: $\chi^2(11) = 55.55$, *gender*: $\chi^2(11) = 57.55$, *computer application*: $\chi^2(11) = 81.00$, *monitor text*: $\chi^2(11) = 45.26$, *laundry*: $\chi^2(11) = 39.07$, all $p < 0.01$).

Next, we conducted Dunn's post-hoc pairwise tests with Bonferroni correction to detect any significant differences in information content (see supplementary document). For all attributes, pairwise Dunn's tests comparing the *only obfuscation* and *obfuscation plus beautification* conditions revealed no significant difference in information, meaning that combining *beautification* with *obfuscation* does not reduce any more information. When compared with the *unfiltered* condition, we found that for *gender*, *computer application*, and *monitor text*, all *obfuscation plus beautification* transforms resulted in significant reduction in information content with medium to high effect sizes ($.45 \leq r \leq .75$, all $p < .01$). For *document* and *dress*, except for *edge + cartoon* and *pixelation + cartoon* respectively, all other *obfuscation plus beautification* transforms significantly reduce information content ($.43 \leq r \leq .6$, all $p < .05$). Finally, for *laundry*, only the *pixelation + cartoon* transform results in reduction in information with medium effect size ($r = .42$, $p < .05$).

Overall, despite being a source of additional abstraction, the beautification transforms do not cause any significant additional reduction in information content to an obscured image.

Effects of Transformations on Visual Aesthetics

Overall Kruskal-Wallis tests indicated that there are significant differences in perceived image aesthetics across conditions for all attributes except *monitor text* (*gender*: $\chi^2(11) = 21.26$, *dress*: $\chi^2(11) = 37.75$, *document*: $\chi^2(11) = 21.54$, *computer application*: $\chi^2(11) = 23.88$, *laundry*: $\chi^2(11) = 24.36$, $p < 0.01$ for *dress*, $p < 0.05$ for other attributes). For the *document* and *dress* attributes, the *unfiltered* condition has

Condition	Document		Dress		Gender		Laundry		Computer app		Monitor text	
	mean	std	mean	std	mean	std	mean	std	mean	std	mean	std
Unfiltered	5.34	1.24	5.56	1.36	5.48	1.16	5.18	1.37	5.42	1.49	5.26	1.55
Edge	4.76	1.55	4.13	2.30	3.41	2.07	3.37	1.96	3.17	2.08	3.43	2.23
Edge + Abstract	4.04	1.87	3.49	2.00	3.40	2.13	3.49	1.94	3.00	2.02	3.36	2.13
Edge + Cartoon	4.64	1.64	4.22	1.97	3.94	1.90	4.24	1.64	3.74	1.75	3.86	1.74
Edge + Color	4.18	1.99	4.12	2.00	3.35	2.11	3.37	1.95	3.47	2.27	3.55	2.20
Pixelation	3.90	2.04	4.84	1.70	3.90	2.04	3.82	1.86	3.78	1.81	3.57	2.33
Pixelation + Abstract	3.76	1.62	4.25	1.75	3.57	1.70	3.96	1.75	3.90	1.88	3.33	1.91
Pixelation + Cartoon	4.08	1.61	4.98	1.52	3.67	1.93	4.08	1.61	3.92	1.70	3.17	1.72
Pixelation + Color	3.66	1.91	4.26	1.81	3.64	1.88	3.60	1.82	3.43	1.80	3.19	2.15
Masking	3.98	1.97	4.24	2.11	3.58	2.17	3.90	1.94	3.34	2.07	3.42	2.01
Masking + Abstract	3.38	1.89	3.27	1.76	2.88	1.62	4.06	1.83	2.77	1.77	3.17	1.86
Masking + Cartoon	3.74	1.88	4.43	1.91	3.61	1.95	4.13	2.05	3.22	1.98	3.41	2.01
Masking + Color	3.38	1.74	4.02	1.85	2.98	1.81	3.80	1.78	2.38	1.60	3.00	1.96

Table 4: Means and standard deviations of information content scores for different attributes.

the highest scores for visual aesthetics, and combining aesthetic transforms reduced scores compared to applying only privacy obfuscations (Table 5). For *document*, the reductions were not significant for any beautification transform (all $p > .05$) but for *dress*, combining *cartoon* with *masking* significantly lowered visual aesthetics ($z = 3.02$, $r = .42$, $p < .05$) compared to the condition when only the *masking* obfuscation was applied. On the other hand, for *laundry*, *edge* obfuscation combined with the *cartoon* transform produced significantly more visually appealing photos compared to both when only *edge* obfuscation was used ($z = 3.01$, $r = .42$, $p = .03$) and the *unfiltered* condition ($z = 3.03$, $r = .43$, $p < .05$). For *gender* and *computer app*, no *obfuscation plus beautification* transform significantly increased aesthetics over *only obfuscation* conditions (see supplementary material).

Overall, except for the cartoon transform (in one case), the beautification transforms did not significantly increase the visual aesthetics of obscured photos.

Effects of Transformations on Viewers' Satisfaction

Except for the *computer application* and *monitor text* attributes, photo satisfaction had the highest scores in the *unfiltered* condition for all other attributes (Table 6). Kruskal-Wallis tests across all conditions detected significant differences in satisfaction scores for *document* ($\chi^2(11) = 22.38$, $p < .05$), *dress* ($\chi^2(11) = 47.1$, $p < .0001$), and *computer app* ($\chi^2(11) = 25.96$, $p < .01$). For the other three attributes, none of the conditions had significantly different satisfaction scores than others. Hence we conducted pairwise tests only for *document*, *dress*, and *computer app* attributes (see

supplementary material). We did not find any statistically significant increase in satisfaction when comparing *only obfuscation* with *obfuscation plus beautification* (see supplementary material). Finally, comparing the *obfuscation plus beautification* conditions with the *unfiltered* condition, we found that for *dress*, *masking + cartoon* significantly lowered satisfaction ($z = 3.6$, $r = 0.51$, $p < .001$). All other results were non-significant ($p > 0.05$).

Overall, we did not find statistically significant evidence that our selected artistic transforms increase viewers' satisfaction compared to obscured photos.

5 DISCUSSION

Our hypothesis was that by applying beautification techniques to an image in which privacy-sensitive content has been obfuscated, we can increase both perceived information content (possibly by providing a high-level story) and visual aesthetics, and thus recover some or all of the viewer satisfaction that would otherwise have been lost to the privacy transform. Our results show that the beautifications we experimented with did *not* significantly increase satisfaction. Certain combinations of obfuscation and beautification transforms (e.g., when using the *cartoon* transform), however, appeared to increase some or all of the three dependent variables (information content, visual aesthetics, and satisfaction). These combinations could be studied with more statistical power in the future, or with modifications that attempt to increase aesthetics and satisfaction. It is interesting to see that the *cartoon* transform boosted information sufficiency despite being a form of abstraction. This supports our speculation that presenting photo content at a high level

Condition	Document		Dress		Gender		Laundry		Computer app		Monitor text	
	mean	std	mean	std	mean	std	mean	std	mean	std	mean	std
Unfiltered	4.92	1.65	5.00	1.29	4.24	1.59	3.62	1.99	3.48	1.55	3.50	1.64
Edge	4.54	1.62	4.20	1.75	3.87	1.78	3.65	1.99	3.54	2.03	3.67	1.93
Edge + Abstract	3.80	1.66	3.38	1.57	3.38	1.93	3.76	1.79	3.20	1.82	3.47	1.67
Edge + Cartoon	4.62	1.71	4.38	1.81	4.34	1.79	4.72	1.75	4.08	1.84	4.20	1.74
Edge + Color	4.04	1.98	3.80	1.77	3.47	1.82	3.61	1.99	3.43	1.88	3.53	1.97
Pixelation	4.57	1.70	4.61	1.59	4.31	1.59	3.92	1.65	3.71	1.80	3.90	1.98
Pixelation + Abstract	4.53	1.47	4.33	1.65	4.27	1.63	4.35	1.60	4.37	1.62	4.16	1.60
Pixelation + Cartoon	4.31	1.68	4.44	1.77	4.06	1.62	3.96	1.71	3.58	1.77	3.73	1.77
Pixelation + Color	4.26	1.93	4.09	1.79	3.92	1.80	3.57	1.86	3.26	1.77	3.51	1.72
Masking	4.76	1.92	4.96	1.73	4.26	1.87	4.34	1.94	3.64	1.99	3.78	1.93
Masking + Abstract	4.46	1.74	4.23	1.64	4.50	1.68	4.42	1.53	3.81	1.66	3.83	1.58
Masking + Cartoon	3.93	1.83	4.00	1.65	3.78	1.63	3.93	1.80	3.22	1.90	3.61	1.75
Masking + Color	4.48	1.76	4.48	1.47	4.06	1.65	3.92	1.76	3.06	1.63	3.34	1.56

Table 5: Means and standard deviations of visual aesthetics scores for different attributes.

condition	Document		Dress		Gender		Laundry		Computer app		Monitor text	
	mean	std	mean	std	mean	std	mean	std	mean	std	mean	std
Unfiltered	4.84	1.42	5.32	1.45	4.62	1.46	4.18	1.87	3.88	1.70	4.10	1.82
Edge	4.43	1.77	4.20	1.98	3.69	1.97	3.87	1.90	3.50	1.94	3.65	2.06
Edge + Abstract	3.80	1.82	3.31	1.73	3.36	1.91	3.67	1.75	3.22	1.74	3.47	1.96
Edge + Cartoon	4.74	1.60	4.42	1.89	4.14	1.68	4.52	1.71	4.08	1.70	3.90	1.59
Edge + Color	4.06	1.92	3.67	1.95	3.51	1.95	3.65	1.97	3.45	1.81	3.51	2.00
Pixelation	4.45	1.79	4.71	1.65	4.24	1.74	4.29	1.85	3.86	2.00	3.94	2.07
Pixelation + Abstract	4.27	1.66	4.25	1.61	4.02	1.67	4.20	1.50	4.16	1.77	3.78	1.64
Pixelation + Cartoon	4.31	1.49	4.65	1.54	3.94	1.73	3.75	1.60	3.69	1.75	3.58	1.60
Pixelation + Color	3.92	1.81	4.15	1.81	3.89	2.01	3.64	1.89	3.15	1.83	3.34	1.89
Masking	4.14	1.82	4.58	1.91	4.06	1.85	4.36	1.80	3.50	1.98	3.56	1.99
Masking + Abstract	4.02	1.74	3.83	1.59	3.96	1.69	4.19	1.54	3.50	1.68	3.56	1.67
Masking + Cartoon	3.70	1.88	4.02	1.95	3.70	1.92	4.04	1.89	3.28	1.90	3.54	1.86
Masking + Color	3.92	1.93	4.00	1.65	3.68	1.78	3.70	1.74	2.90	1.66	3.24	1.62

Table 6: Means and standard deviations of photo satisfaction scores for different attributes.

might increase overall information absorption. Also it may be that viewers found the ‘beautified’ versions more interesting and derived more information from the transformed photo. For example, an ordinary object may appear more interesting following the cartoon transformation.

We found that the *abstract* transform appeared to increase aesthetics in some cases, but lowered information content without increasing viewers’ satisfaction; we expected a greater increase in perceived visual aesthetics since this is the most artistic transform among the three. One possible explanation is that the reduction in information content by the *abstract*

transform might negatively affect the other two variables, since our results from the path model analysis show that information content is associated with both of those variables. Finally, we found the *color* transform did not increase any of the measured variables. We expected *color* to have a lesser effect on both lowering information content and increasing visual aesthetics compared to the other two transforms. It might be the case that the negative effect of the loss of information on visual aesthetics and satisfaction was greater than the increase, if any, in the latter two variables.

Although we did not have sufficient statistical power to ascertain the difference in satisfaction between the obfuscated and beautified conditions, our findings still suggest the validity of an approach where a combination could increase satisfaction, with the cartoon filter being the most promising. Overall, we believe future work should explore other possible beautification transforms to study the novel privacy vs. satisfaction trade-off. It may be particularly promising to study obfuscating transforms that are themselves aesthetically pleasing or ‘fun’ instead of beautifying the rest of the image – as people grow accustomed to filters and effects (such as ‘stickers’) in photo-sharing applications, it will be increasingly acceptable to apply such obfuscations and transforms in general. By understanding and quantifying the effects of obfuscation on privacy and satisfaction, as well as the effects of beautification on satisfaction, we may be able to design the ‘correct’ combination of transformations for sensitive and non-sensitive image regions in order to both improve privacy and retain (or improve) satisfaction for the viewer. Indeed, improving privacy could be ‘fun’ too, both for the person transforming the photo and the viewer.

Limitations

We note several limitations of our study, which could be addressed in future work. We purposely restricted our pool of MTurk participants to users in the United States of at least 18 years of age to control for cultural differences. Although MTurk participants resemble US population fairly well and better than other web panels [38], our findings may not generalize for other age groups. Further, photo sharing behaviors as well as perceptions of privacy and aesthetics differ across cultures, and explicitly studying these differences in the context of beautification and obfuscation transformations would be interesting for future work. Moreover, we used the same pool of photos as past work to allow for direct comparison with published results, but these photos were collected from web sources. Participant views of aesthetics and satisfaction on these images may not reflect how they would feel about transformations applied to their own images. Follow-up studies could request users to subject their own photos to transformations, and compare outcomes on those photos versus the web images we consider here. Our selections of obfuscation and beautification transformations were made based on past work, and they were designed precisely for the same purposes as ours – to obfuscate objects and increase photo aesthetics. There are many other possible combinations of such transformations, and studying a larger set may reveal techniques that are more effective at balancing privacy, aesthetics, and satisfaction. Finally, we did not consider other obfuscation techniques (such as Snapchat filters and Apple Memoji) that can add or replace information instead of just obscuring (e.g., a smiley face replacing the original

emotion of a person). While the popularity of these features indicates their effectiveness in retaining and/or increasing viewers’ satisfaction, it would be interesting to study their effectiveness in protecting privacy.

6 CONCLUSIONS

We explored the novel question of whether a viewer’s satisfaction of a photo with obfuscated elements can be improved. While one might expect there to be a strict privacy-satisfaction trade-off, where applying obfuscations to improve privacy degrades the viewing experience, we hypothesize that ‘beautification’ transforms can be applied to the *rest* of the image to compensate for or counteract the loss in satisfaction, in order to create an image that *both* preserves privacy and viewer satisfaction.

As a first step, we experimented with three off-the-shelf beautification transforms and extended prior work on obfuscation transforms to evaluate combinations of obfuscation and beautification. While we did not find statistically significant support for our hypothesis that these transforms boost viewers’ satisfaction, we hope the gain in information content and visual aesthetics will inspire the exploration of new transforms that take into account the negative effects of privacy obfuscations, as well as obfuscating transforms that are themselves aesthetically pleasing (e.g., a sticker obfuscating a face but also making the image more fun to look at). We believe this line of work is particularly salient with the popularity of photo sharing and adding photo effects and stickers, and hope it inspires further exploration of how such transforms can be used not only for entertainment but to simultaneously afford more privacy.

7 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 are grateful to the participants of our online study.

REFERENCES

- [1] T. O. Aydin, A. Smolic, and M. Gross. 2015. Automated Aesthetic Analysis of Photographic Images. *IEEE Transactions on Visualization and Computer Graphics* 21, 1 (Jan 2015), 31–42. <https://doi.org/10.1109/TVCG.2014.2325047>
- [2] Adrien Bousseau, Fabrice Neyret, Joëlle Thollot, and David Salesin. 2007. Video Watercolorization Using Bidirectional Texture Advection. In *ACM SIGGRAPH 2007 Papers (SIGGRAPH '07)*, Vol. 26. ACM, New York, NY, USA, Article 104. <https://doi.org/10.1145/1275808.1276507>
- [3] Jacob Cohen. 1992. A power primer. *Psychological bulletin* 112, 1 (1992), 155.
- [4] Dianne Cyr, Milena Head, Hector Larios, and Bing Pan. 2009. Exploring Human Images in Website Design: A Multi-Method Approach. *MIS Quarterly* 33, 3 (2009), 539–566. <http://www.jstor.org/stable/20650308>
- [5] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z. Wang. 2006. Studying Aesthetics in Photographic Images Using a Computational Approach.

- In *Computer Vision – ECCV 2006*, Aleš Leonardis, Horst Bischof, and Axel Pinz (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 288–301.
- [6] William J. Doll and Gholamreza Torkzadeh. 1988. The Measurement of End-User Computing Satisfaction. *MIS Quarterly* 12, 2 (1988), 259–274. <http://misq.org/the-measurement-of-end-user-computing-satisfaction.html>
 - [7] Maeve Duggan. 2013. Photo and video sharing grow online. (2013). <http://www.pewinternet.org/2013/10/28/photo-and-video-sharing-grow-online/>.
 - [8] Jim Edwards. 2014. PLANET SELFIE: We're Now Posting A Staggering 1.8 Billion Photos Every Day. <http://www.businessinsider.com/were-now-posting-a-staggering-18-billion-photos-to-social-media-every-day-2014-5>. (2014). Accessed April 19, 2018.
 - [9] Ahmed M. Elgammal and Babak Saleh. 2015. Quantifying Creativity in Art Networks. *CoRR* abs/1506.00711 (2015). arXiv:1506.00711 <http://arxiv.org/abs/1506.00711>
 - [10] Graham D Finlayson, Michal Mackiewicz, and Anya Hurlbert. 2015. Color correction using root-polynomial regression. *IEEE Transactions on Image Processing* 24, 5 (2015), 1460–1470.
 - [11] J. Gao, S. J. Kim, and M. S. Brown. 2011. Constructing image panoramas using dual-homography warping. In *CVPR 2011*. 49–56. <https://doi.org/10.1109/CVPR.2011.5995433>
 - [12] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. 2015. A Neural Algorithm of Artistic Style. *CoRR* abs/1508.06576 (2015). arXiv:1508.06576 <http://arxiv.org/abs/1508.06576>
 - [13] Google Street View. 2018. Image Acceptance and Privacy Policies. (2018). Retrieved March 07, 2018 from <https://www.google.com/streetview/privacy/>.
 - [14] Rakibul Hasan, Eman Hassan, Yifang Li, Kelly Caine, David J. Crandall, Roberto Hoyle, and Apu Kapadia. 2018. Viewer Experience of Obscuring Scene Elements in Photos to Enhance Privacy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 47, 13 pages. <https://doi.org/10.1145/3173574.3173621>
 - [15] Eman T. Hassan, Rakibul Hasan, Patrick Shaffer, David Crandall, and Apu Kapadia. 2017. Cartooning for Enhanced Privacy in Lifelogging and Streaming Videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop on Computer Vision Challenges and Opportunities for Privacy and Security (CV-COPS)*. 29–38. <https://doi.org/10.1109/CVPRW.2017.175>
 - [16] James Hays and Irfan Essa. 2004. Image and video based painterly animation. In *International Symposium on Non-photorealistic Animation and Rendering*. ACM, ACM, New York, NY, USA, 113–120. <http://doi.acm.org/10.1145/987657.987676>
 - [17] Panagiotis Ilia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. 2015. Face/Off: Preventing Privacy Leakage From Photos in Social Networks. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*. ACM, New York, NY, USA, 781–792. <https://doi.org/10.1145/2810103.2813603>
 - [18] Phillip Isola, Jianxiong Xiao, Devi Parikh, Antonio Torralba, and Aude Oliva. 2014. What makes a photograph memorable? *IEEE transactions on pattern analysis and machine intelligence* 36, 7 (2014), 1469–1482.
 - [19] Sanjay Kairam, Joseph 'Jofish' Kaye, John Alexis Guerra-Gomez, and David A. Shamma. 2016. Snap Decisions?: How Users, Content, and Aesthetics Interact to Shape Photo Sharing Behaviors. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 113–124. <https://doi.org/10.1145/2858036.2858451>
 - [20] Nima Khademi Kalantari and Ravi Ramamoorthi. 2017. Deep high dynamic range imaging of dynamic scenes. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 144.
 - [21] Y. Kao, R. He, and K. Huang. 2017. Deep Aesthetic Quality Assessment With Semantic Information. *IEEE Transactions on Image Processing* 26, 3 (March 2017), 1482–1495. <https://doi.org/10.1109/TIP.2017.2651399>
 - [22] RM Khan and MA Khan. 2007. Academic sojourners, culture shock and intercultural adaptation: A trend analysis. *Studies About Languages* 10 (2007), 38–46.
 - [23] Hope King. 2016. *Wildly popular Prisma app just made a major breakthrough*. CNN. <http://money.cnn.com/2016/08/23/technology/prisma-offline/index.html>
 - [24] Jan Eric Kyprianidis and Jürgen Döllner. 2008. Image Abstraction by Structure Adaptive Filtering. In *Proc. EG UK Theory and Practice of Computer Graphics*. Eurographics Association, Manchester, UK, 51–58.
 - [25] Yifang Li, Nishant Vishwamitra, Bart P. Knijnenburg, Hongxin Hu, and Kelly Caine. 2017. Effectiveness and Users' Experience of Obfuscation As a Privacy-Enhancing Technology for Sharing Photos. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 67 (Dec. 2017), 24 pages. <https://doi.org/10.1145/3134702>
 - [26] 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. <https://doi.org/10.1002/meet.14504901100>
 - [27] Yabing Liu, Krishna P. Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing Facebook Privacy Settings: User Expectations vs. Reality. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (IMC '11)*. ACM, New York, NY, USA, 61–70. <https://doi.org/10.1145/2068816.2068823>
 - [28] Xin Lu, Zhe Lin, Xiaohui Shen, Radomir Mech, and James Z Wang. 2015. Deep multi-patch aggregation network for image style, aesthetics, and quality estimation. In *Proceedings of the IEEE International Conference on Computer Vision*. 990–998.
 - [29] Michelle Madejski, Maritza Johnson, and Steven M Bellovin. 2011. The failure of online social network privacy settings. *Department of Computer Science, Columbia University, Tech. Rep. CUCS-010-11* (2011).
 - [30] Aqdas Malik, Amandeep Dhir, and Marko Nieminen. 2016. Uses and Gratifications of digital photo sharing on Facebook. *Telematics and Informatics* 33, 1 (2016), 129 – 138. <https://doi.org/10.1016/j.tele.2015.06.009>
 - [31] Farhad Manjoo. 2017. *Why Instagram Is Becoming Facebook's Next Facebook*. <https://www.nytimes.com/2017/04/26/technology/why-instagram-is-becoming-facebook-next-facebook.html>
 - [32] Adam W Meade and S Bartholomew Craig. 2012. Identifying careless responses in survey data. *Psychological methods* 17, 3 (Sept 2012), 437–455. <https://doi.org/10.1037/a0028085>
 - [33] N Murray, L Marchesotti, and F Perronnin. 2012. AVA: A large-scale database for aesthetic visual analysis. *Proceedings / CVPR, IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 2408–2415. <https://doi.org/10.1109/CVPR.2012.6247954>
 - [34] M. Nishiyama, T. Okabe, I. Sato, and Y. Sato. 2011. Aesthetic quality classification of photographs based on color harmony. In *CVPR 2011*. 33–40. <https://doi.org/10.1109/CVPR.2011.5995539>
 - [35] Anne Oeldorf-Hirsch and S. Shyam Sundar. 2016. Social and Technological Motivations for Online Photo Sharing. *Journal of Broadcasting & Electronic Media* 60, 4 (2016), 624–642. <https://doi.org/10.1080/08838151.2016.1234478> arXiv:https://doi.org/10.1080/08838151.2016.1234478
 - [36] Mark R Phillips, Bradley D McAuliff, Margaret Bull Kovera, and Brian L Cutler. 1999. Double-blind photoarray administration as a safeguard against investigator bias. *Journal of Applied Psychology* 84, 6 (1999), 940.
 - [37] Yasmeen Rashidi, Tousif Ahmed, Felicia Patel, Emily Fath, Apu Kapadia, Christena Nippert-Eng, and Norman Makoto Su. 2018. "You don't

- want to be the next meme": College Students' Workarounds to Manage Privacy in the Era of Pervasive Photography. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 143–157. <https://www.usenix.org/conference/soups2018/presentation/rashidi>
- [38] E M Redmiles, S Kross, and M L Mazurek. 2019. How Well Do My Results Generalize? Comparing Security and Privacy Survey Results from MTurk, Web, and Telephone Samples. In *2019 IEEE Symposium on Security and Privacy (SP)*, Vol. 00. IEEE, 227–244. <https://doi.org/10.1109/SP.2019.00014>
- [39] Andra Siibak. 2009. Constructing the self through the photo selection-visual impression management on social networking websites. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace* 3, 1 (2009).
- [40] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. 2013. The Post That Wasn't: Exploring Self-censorship on Facebook. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. ACM, New York, NY, USA, 793–802. <https://doi.org/10.1145/2441776.2441865>
- [41] Manya Sleeper, Justin Cranshaw, Patrick Gage Kelley, Blase Ur, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2013. "I Read My Twitter the Next Morning and Was Astonished": A Conversational Perspective on Twitter Regrets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 3277–3286. <https://doi.org/10.1145/2470654.2466448>
- [42] Jose M. Such, Joel Porter, Sören Preibusch, and Adam Joinson. 2017. Photo Privacy Conflicts in Social Media: A Large-scale Empirical Study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 3821–3832. <https://doi.org/10.1145/3025453.3025668>
- [43] C. Tomasi and R. Manduchi. 1998. Bilateral filtering for gray and color images. In *Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271)*. 839–846. <https://doi.org/10.1109/ICCV.1998.710815>
- [44] Leman Pinar Tosun. 2012. Motives for Facebook use and expressing "true self" on the Internet. *Computers in Human Behavior* 28, 4 (2012), 1510 – 1517. <https://doi.org/10.1016/j.chb.2012.03.018>
- [45] Luam Catao Totti, Felipe Almeida Costa, Sandra Avila, Eduardo Valle, Wagner Meira, Jr., and Virgilio Almeida. 2014. The Impact of Visual Attributes on Online Image Diffusion. In *Proceedings of the 2014 ACM Conference on Web Science (WebSci '14)*. ACM, New York, NY, USA, 42–51. <https://doi.org/10.1145/2615569.2615700>
- [46] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. 2011. "I Regretted the Minute I Pressed Share": A Qualitative Study of Regrets on Facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS '11)*. ACM, New York, NY, USA, Article 10, 16 pages. <https://doi.org/10.1145/2078827.2078841>
- [47] Pamela Wisniewski, Heather Lipford, and David Wilson. 2012. Fighting for My Space: Coping Mechanisms for sns Boundary Regulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 609–618. <https://doi.org/10.1145/2207676.2207761>
- [48] Hang Zhang and Kristin J. Dana. 2017. Multi-style Generative Network for Real-time Transfer. *CoRR* abs/1703.06953 (2017). [arXiv:1703.06953](http://arxiv.org/abs/1703.06953)