

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/266655817>

Faces engage us: photos with faces attract more likes and comments on Instagram

Article · April 2014

DOI: 10.1145/2556288.2557403

CITATIONS

168

READS

13,763

3 authors, including:



[David Ayman Shamma](#)

Centrum Wiskunde & Informatica

125 PUBLICATIONS 4,577 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



CozyMaps [View project](#)

Faces Engage Us: Photos with Faces Attract More Likes and Comments on Instagram

Saeideh Bakhshi
Georgia Tech
sbakhshi@cc.gatech.edu

David A. Shamma
Yahoo Labs
aymans@acm.org

Eric Gilbert
Georgia Tech
gilbert@cc.gatech.edu

ABSTRACT

Photos are becoming prominent means of communication online. Despite photos' pervasive presence in social media and online world, we know little about how people interact and engage with their content. Understanding how photo content might signify engagement, can impact both science and design, influencing production and distribution. One common type of photo content that is shared on social media, is the photos of people. From studies of offline behavior, we know that human faces are powerful channels of non-verbal communication. In this paper, we study this behavioral phenomena online. We ask how presence of a face, it's age and gender might impact social engagement on the photo. We use a corpus of 1 million Instagram images and organize our study around two social engagement feedback factors, *likes* and *comments*. Our results show that photos with faces are 38% more likely to receive likes and 32% more likely to receive comments, even after controlling for social network reach and activity. We find, however, that the number of faces, their age and gender do not have an effect. This work presents the first results on how photos with human faces relate to engagement on large scale image sharing communities. In addition to contributing to the research around online user behavior, our findings offer a new line of future work using visual analysis.

Author Keywords

mobile; photo; Instagram; faces; face detection; engagement; age; demographics; gender; content; image; social media; image sharing community

This material is based upon work supported in part by the Defense Advanced Research Projects Agency (DARPA) under Contract No. W911NF-12-1-0043. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or the U.S. Government. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

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 2014, April 26-May 1, 2014, Toronto, Ontario, Canada.
Copyright 2014 ACM 978-1-4503-2473-1/14/04..\$15.00.
<http://dx.doi.org/10.1145/2556288.2557403>

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Even as babies, humans love to look at faces; infants, barely minutes old, turn toward faces, sensing that they are important [39]. It is widely accepted in neuroscience that face perception is perhaps the most highly developed human visual skill [26]. Faces are also powerful channels of nonverbal communication [45]. We constantly monitor faces because they provide vital clues in an impressive variety of contexts: attraction, the complexity of emotions, identity, age, humor, and a person's regional and national background [24].

Many of the faces we see everyday now have an online presence. Photo sharing communities such as Instagram have made it possible to communicate with large groups of distributed people through an image—be it a picture of whats for dinner or a selfie—perhaps more easily than through words alone. As Kelsey puts it, “we are moving away from photography as a way of recording and storing the past, and instead turning photography into a social medium in its own right” [35].

Online photo sharing communities have grown at an impressive pace. At the time of this writing, Instagram users upload 55 million photos a day to the site¹. This presents a key research challenge for photo sharing communities like Instagram (cf. Flickr, Imgur, Tumblr): how do we discover the mechanisms by which users communicate around visual content and engage with such content. In other words, since engagement is vital to photo sharing communities, it is critical to understand what form of content drives engagement. While several research studies have focused on how users engage with textual content [6, 9, 11, 12, 31, 38], there are few studies on what makes visual content socially engaging online. To investigate this, we ask the following research questions in this paper, driven by social psychology work on face perception:

RQ1: Do photos with faces differ in online engagement compared to photos without them?

RQ2: If so, how do characteristics of the image subject, such as gender and age, affect engagement?

¹<http://instagram.com/press> (Accessed 9/2013)

We use Instagram to answer our research questions. Instagram has over 150 million active monthly users who collectively generate 1.2 billion likes per day. There are two social aspects of engagement here: the number of likes and the number of comments on the Instagram image². Using a corpus of 1 million images from the community, we find that on an average a photo that contains a face receives 38% more likes and 32% more comments compared to a photo that does not contain any faces (even after controlling for user activity levels and social network reach). Further, we find that the number of faces in the photo, their age and their gender do not impact engagement.

To our knowledge, our study is one of the first to show, systematically and at scale, how photos with faces drive online social engagement. In addition to contributing to the ongoing research conversation surrounding engagement, we believe that these findings create a path for future work, not only to uncover the impact of faces on other aspects of online user behavior, but also to use computer vision techniques to discover other antecedents of engagement. For example, we may be able to apply vision techniques to relate facial *expressions* of emotion to social behavior.

We begin with a review of related research on content and face perception, and a summary of the Instagram community. Next, we introduce the corpus we collected from Instagram and describe the statistical methods we used to isolate the effect of faces on likes and comments. Finally, we interpret our findings within the frame of existing work, both in theory and in practice.

LITERATURE REVIEW

In this section, we describe related work on role of content in engagement, face perception theories and HCI studies of faces. We then summarize the media community of our study, Instagram.

Role of Content on User Engagement

As Ellison and colleagues note, “the primary function of these [social network] sites is to consume and distribute personal content about the self” [22]. Sharing content can in turn ensure that users remain engaged and committed in the future visits [13, 47]. On the other hand, users have diverse motivations to share content on social network sites. For example, users may share useful content to appear knowledgeable or simply to help out [50]. Not only the content of posts, but also the emotional valence behind it can drive its usage. For example, in a recent study, researchers used New York Times articles to examine the relationship between the emotion evoked by content and its virality [9], finding that there is a direct relationship.

Much research attention has gone into investigating what makes content in an online community interesting to its members. In a series of studies conducted on Usenet newsgroups, researchers investigated properties that influenced the likelihood of reply, a measure of the community’s

engagement. Explicit requests, personal testimonials relating one’s connection to the group, and staying on-topic increased the probability of receiving a reply. Newcomers to a group were less likely to receive a reply than veterans [6]. Following up on this work, Burke et al. studied the role of self-disclosing introductions, mentions of the poster’s age, and an acknowledgment that this is the poster’s first post; these factors were found to increase reply probability [11]. In another study, Burke and Kraut, studied the effect of the politeness of a post, finding that politeness leads to more replies in certain types of groups, while in other types of groups, rudeness actually increases replies [12]. On Twitter, researchers have used retweeting as a measure of community interest/engagement, and have investigated the features that predict retweeting. Suh et al. found that the presence of URLs and hashtags in tweets predicted more retweeting, as did a richer connection with the community [44]. In a recent work, Gilbert et. al. studied Pinterest- a social networking site based on images- and found that which properties of an image makes the content more interesting to users [23]. The properties used in this work are based on meta data and not the content of images.

As far as we know, however, we have no such similar line of work on how image content can affect different aspects of online behavior, such as engagement, diffusion or link formation. In our work, we intend to provide an understanding of image engagement by looking at the photo content.

Face perception

One of the common types of photo content shared on social networking sites is the photos of people or the ones with human faces in them. Through daily experience, we know that human faces are readily distinguishable. People tend to find faces in unexpected scenes and photographs even where faces do not exist. For example, the 1976 Viking 1 prob photographed a shadowed region on Mars’ northern planes that resembled a face. While higher resolution imagery has shown the region to actually be a mesa, the face on Mars remains a pop icon and the source of many books, TV shows, and films [40].

Faces have long been a source of scientific interest in a wide range of disciplines. In recent years, this breadth of interests, approaches and expertise has led directly to rapid advances in our understanding of many different aspects of how we perceive and process faces [10]. The human brain has evolved to recognize faces within hours after birth [33]. Human infants only minutes old attend particularly to face-like stimuli relative to equally complicated non-face stimuli [10, 30]. We prefer to look at faces from that early age and thereafter, often opting to spend more time looking at faces than any other type of object [53]. By the age of two months, infants begin to differentiate specific visual features of the face [39] and process facial expressions [20]. Our brains have a specific region, Fusiform Face Area (FFA), that is specialized for facial recognition [34, 41]. Faces are important for social cognition as well, not only because we are able to recognize them earlier than other objects, but also because they display our feelings about past, current and future events through expressions [17,

²We use the terms image and photo interchangeably throughout this paper to refer to the images on Instagram.

21]. This can be highly important to very practical concerns: faces, particularly attractive ones, are found to be effective in improving consumer responses to advertisements [7].

Our research examines the presence of this phenomena on-line, by analyzing the effect of having faces in engaging users on Instagram.

Faces and HCI research

In HCI research, there is a great deal of work exploring the benefits of using face icons and faces in interfaces [36, 43, 45, 49]. Walker et al. [49] studied how having faces and facial expressions for a computer application affects users' performance and productivity. They compared subjects' responses to an interview survey under three conditions: questions spoken by a synthesized face with neutral expressions, spoken by a face with stern expressions, or text only. Subjects who responded to the spoken face made more effort to answer the questions by spending more time, writing more comments and making fewer mistakes. They reported that having a face is engaging and takes more effort and attention from the user.

Takeuchi et al. [46] compared users' impressions of an agent which helped them to win a card game. The agent was represented either as an arrow or a face. They showed that users respond differently to systems having a face than to those without. The arrow was recognized as useful and reliable, while the face was rated as fun and entertaining. They conclude that a face in an interface captures more attention and people try to interpret the meaning behind the expression.

Studies on embodied interfaces showed similar results. Agents are visual digital representations of a computer interface often in the form of human-like faces [15]. In a review study of embodied agents [18], authors reported that adding an embodied agent to an interface made the experience more engaging.

Role of Age and Gender

Age and gender have been studied extensively as factors affecting social media use [8, 16, 27, 28]. Recent data³ shows that women form a majority of Facebook and Twitter users, as well as dominating Pinterest; however, men are the majority of users on Google+ and LinkedIn. In a recent study Gilbert et. al. [23] found that females are more likely to receive repins and fewer followers than males on Pinterest. Moreover, Pew Internet Research [4] ran a survey to give marketers a clearer picture of who they can expect to reach on Instagram. According to the source, 28% of U.S. internet users aged 18 to 29 snapped and posted photos on the network in December 2012. 14% of those aged 30 to 49 did the same, and very few users older than 50 participated in any way on Instagram.

Inspired by previous research on disparities in internet usage and social network audience, we used age and gender variables to investigate whether they affect the number of likes and comments on photos.

Instagram

³<http://mashable.com/2012/07/04/men-women-social-media/> (Accessed 9/2013)

Instagram is a social network site designed around photo- and video-sharing. It enables users to take photos and videos with their mobile devices, apply digital filters to them and share them on variety of social networking services, such as Facebook, Twitter, Tumblr and Flickr [2], all of which are social media sites in their own right. Instagram has rapidly gained popularity with over 100 million active users as of April 2012 [19]. The total number of photographs uploaded recently exceeded one billion [1, 3].

Instagram accounts are public by default, unless users elect to create a private account; there is no tier privacy photo by photo. To add a photo, users can take a photo from inside the app. It is also possible to choose a photo from an existing album on the mobile device to share with Instagram followers. Instagram users can apply filters on their photos. An Instagram filter is a digital layer that when added to a photo, gives it the appearance of stylistic editing. Some filters enhance the colors in a photo, while others dull the light to a soft glow for an aged, vintage appearance.

Despite the popularity of Instagram, there is little scholarly work on it. In a recent piece, Hochman et al. [29] analyzed colors in photos uploaded from two different cities of New York and Tokyo and found differences across the two locations. For instance, hues of pictures in New York were mostly blue-gray, while those in Tokyo were characterized by dominant red-yellow tones.

METHODS

We take a quantitative approach in this paper to investigate the relationship between faces and engagement. While engagement can be quantified in various ways, we use two essential aspects of content on Instagram that can signal for engagement: likes and comments. The number of likes signals for the extent to which the content is interesting to users and the number of comments quantifies the level of discussion on the social network. In this section, we describe the data we collected from Instagram and how we detected faces and their age and gender; followed by clarifying our statistical methods and analysis process.

Face detection

Face detection and recognition from images or video is a popular topic in vision research and it has received lots of attention [48, 51]. A general statement of the problem of machine recognition of faces is usually formulated as follows: given still or video images of a scene, identify or verify one or more persons in the scene using a stored database of faces or facial features. The solution to the problem involves segmentation of faces (face detection) from cluttered scenes and extraction of features from the face.

While the current state of the art in face detection and recognition is highly accurate [32], we did not have access to an implementation that can work for large scale image analysis. We therefore used a publicly available face detection API developed by Face++ [5]. We only use the detection modules, as the goal of this paper is to find relationship between existence of faces and the social engagement. Face++ provides a set of compact, powerful, and cross-platform vision services,

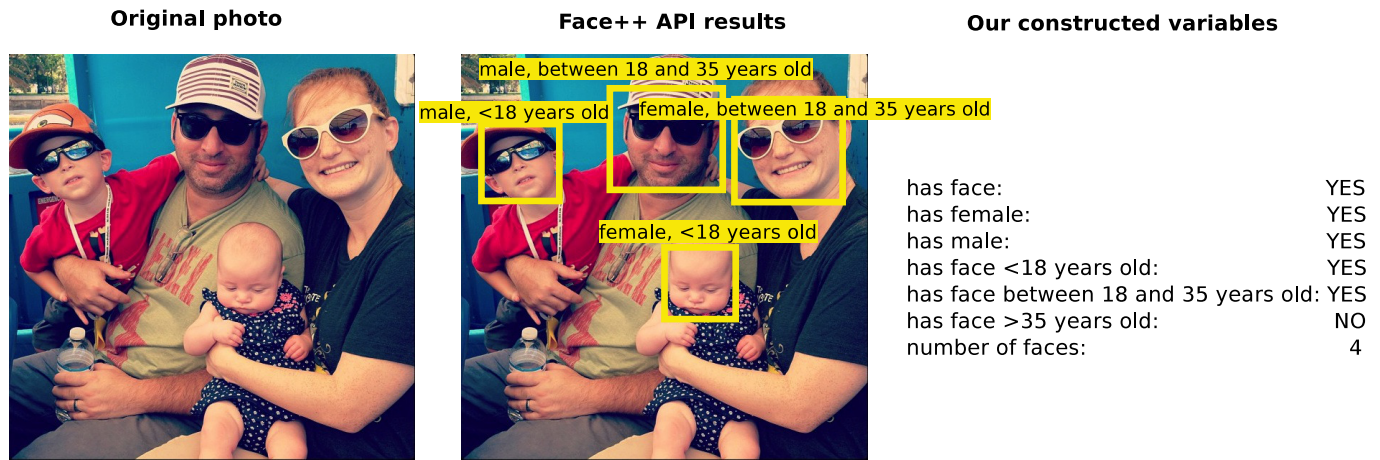


Figure 1. Example Face++ face detection and how we construct our variables. The photo used in this example is a photo under Creative Commons license from Flickr.

which enabled us to use cutting-edge vision techniques. The API does not provide us with an estimation of accuracy, so we turn to a crowd-sourced validation method to confirm the accuracy of our face detector described later in the validation section.

Face++ provides us with an API that accepts the URL of an Instagram image and returns information about detected faces. This information includes the position of the face in the image, as well as the detected gender and age range of all faces. We then reduce the dimensionality of data by converting the results into a binary space, where we mark only when there is a face in an image. We also identify whether any of the faces in the image belong to certain age ranges. The three age ranges we consider in this paper are (1) children and teens-younger than 18, (2) young adults- faces with age between 18 and 35, and (3) older adults- older than 35. To evaluate the role of gender, we construct another binary feature which determines whether at least one female or one male face is in the image. Figure 1 shows an example Face++ detection and how we construct our variables.

Data

Our goal is to obtain a random sample of photos from Instagram. Even though Instagram provides us with a publicly available API, gathering a random subset of photos is a challenging task. We can either search for photos by location or query on the list of most recent popular photos. We opted to start with a set of 2,000 popular Instagram photos, collected on November 2012. We then used snowball sampling [25] to collect the users and their followers as well as a random set of their photos. Our dataset consists of 23 million Instagram photos and over 3 million Instagram users. To soften biases due to snowball sampling, we randomly selected 1.1 million photos from this data set. The snowball sampling method was necessary because Instagram does not provide any mechanism by which to monitor the global stream of photos. Figure 2 shows a detailed flowchart of data collection, evaluation and analysis processes.

Response variables (dependent measures)

Our goal in this paper is to investigate the role of photos in predicting user engagement on Instagram. We chose number of likes and number of comments as two features that represent fundamental aspects of engagement on the site. An overview of each of these variables is provided in Table 1.

Likes: Number of likes is a measure of engagement for the photo. It quantifies the number of distinct users who liked the photo. Like is a strong social signal on Instagram that demonstrates the extent to which users liked the photo.

Comments: Number of comments is another measure of engagement, or as Yew and Shamma [52] note, a measure of explicit action on the content. The number of comments is the number of distinct comments posted on the photo. The number of comments determines the extent to which users discussed the photo and hence it can be considered as measure of discussion.

Predictor variables

In this paper, we use two major control variables to adjust for the impact of social network reach and a user’s activity.

Control: user’s followers count. An Instagram photo is posted by an Instagram user. The nature of relationship on Instagram is follower/following. Users form a social network based on “follow” relationships. When A follows B, B’s photos will show up in A’s photo-stream. The number of followers signals the social network reach. The more number of followers, the more people can see the photo and there is presumably a higher chance of receiving likes and comments.

Control: user’s photo count. Photo count is the feature we use to quantify a user’s activity on the site. It represents the number of photos on a user’s profile. The larger values of photo count show the user has shared more content on the site; in other words the user is more active.

As we discussed in related work section, faces are found to be effective stimuli [10, 30] in attracting people’s attention. We use a binary variable as our predictor to account for presence of a face in the photo.

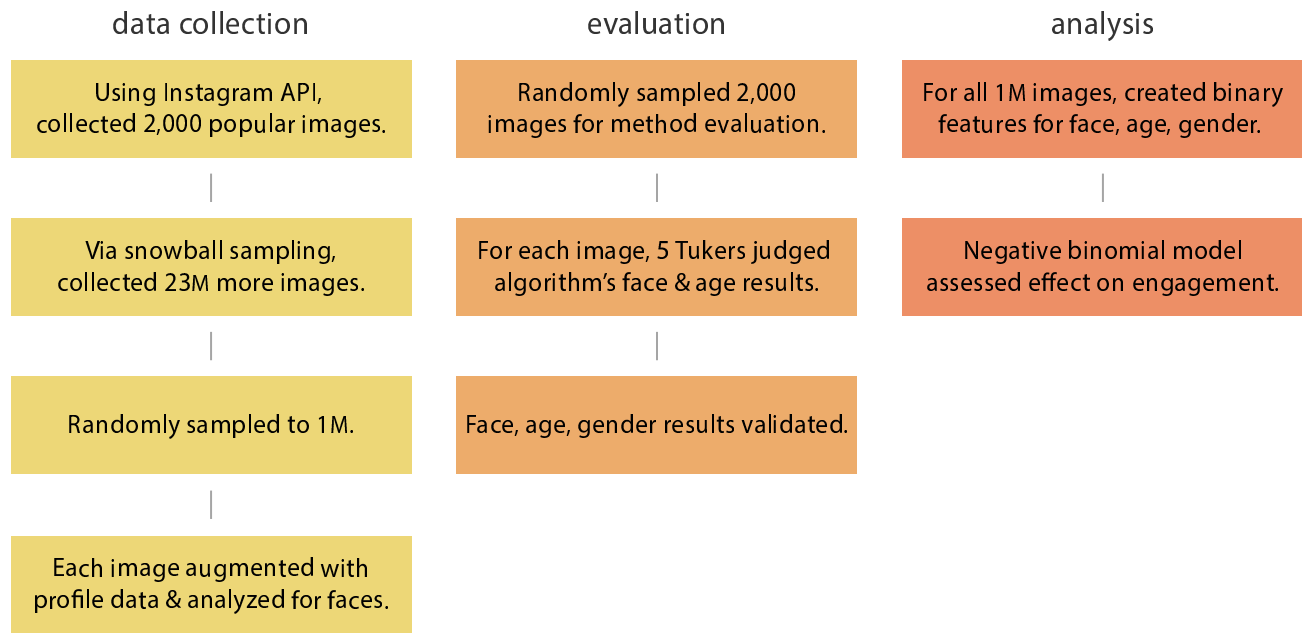


Figure 2. An overview of the steps taken in this paper for collection, evaluation and analysis of the data.

Has face. For each Instagram photo, we determine whether at least one human face exists in the photo. This is a binary feature; when it is set to 1, there is at least one face in the image, otherwise it is set to 0.

Other than presence of faces, we consider variables identifying age and gender of them. Our age and gender variables are derived using face detection method.

Has children and teens- has face < 18 years old. We use a binary feature to determine whether the photo has any faces in the age group <18 years old. The variable is set to 1 when at least one of the identified faces in the image appears to be younger than 18 years old, and set to 0 otherwise.

Has young adults- has face > 18 and < 35 years old. This is another age feature that is set to 1 when at least one of the identifies faces in the image appears to be between 18 and 35 years old, and it is set to 0 otherwise.

Has older adults- Has face > 35 years old. Our final age feature is to identify presence of older adults in the image. If at least one of the faces in the image appears to be older than 35 years old, this variable is set to 1, 0 otherwise.

Has female face. This feature is a binary feature reflecting whether there is a female face in the photo. When the variable is set to 1, the image has at least one female face, and it is set to 0 otherwise.

Has male face. This feature is a binary feature reflecting whether there is a male face in the photo. When the variable is set to 1, the image has at least one male face, and it is set to 0 otherwise.

The distribution and short summary of each of these features is provided in Table 1.

Statistical methods

Next, we present statistical methods we used to model our two dependent variables, number of likes and number of comments. Both dependent variables are *count* variables. We model the number of likes and the number of comments using negative binomial regression, on two classes of independent variables: the control variables (followers count and photos count) and our variables of interest (related to existence of a face, age group of the face and gender of the face). Negative binomial regression is well-suited for *over-dispersed* distributions of *count* dependent variable [14]. We use negative binomial regression instead of Poisson regression since the variance of the dependent variable is larger than the mean for both likes and comments. We use over-dispersion to test whether Poisson or negative binomial regression should be used. This test was suggested by Cameron and Trivedi [14], and involves a simple least-squares regression to test the statistical significance of the over-dispersion coefficient.

The regression coefficients β allow us to understand the effect of an independent variable on the number of likes and comments (note that to be able to compare coefficients, we z-score all numerical variables before performing regression). For the variables with heavy tail distribution, such as followers count and photos count, we log transformed the variables before performing regression. We use Chi-squared statistics to find the statistical significance of our regression models, computing the reduction in deviance as compared to a null model.


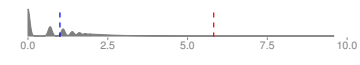


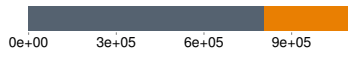
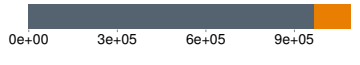


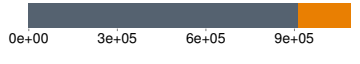

Type	Variable	Description	Distribution
Engagement	likes*	Number of likes on each photo.	
	comments*	Number of comments on each photo.	
Audience & Activity	followers*	Number of users who follow the photo's owner.	
	photos*	Number of photos shared by photo's owner.	
Faces	has face	1 if the photo contains a face, 0 otherwise.	
	has face < 18 years old	1 if there is at least one face younger than 18, 0 otherwise.	
	has face ∈ [18, 35] years old	1 if there is at least one face with age between 18 and 35, 0 otherwise.	
	has face > 35 years old	1 if there is at least one face older than 35, 0 otherwise.	
	has female face	1 if there is at least one female face in the photo, 0 otherwise.	
	has male face	1 if there is at least one male face in the photo, 0 otherwise.	

Table 1. Distributions of quantitative and binary variables used in this paper. Variables marked with “*” are log-transformed. The red and blue lines identify mean and median of the distribution, respectively. Orange refers to 1’s in the bar graphs. The engagement variables are our dependent measures. Audience and activity variables are used as controls, and faces variables are the focal point of this study.

FACE DETECTION VALIDATION

As we mentioned in the previous section, we use Face++ API to detect faces in Instagram photos. Even though the currently used face detection mechanisms are high in accuracy (over 95%), we undertake an additional evaluation step to validate and confirm the accuracy of our methods. For this purpose, we crowd-source a random sample of photos from our dataset to Mechanical Turkers in order to verify the results of API.

The validation process is as follows: we select a random sample of 2,000 images from our dataset. We then create tasks on Mechanical Turk, where the image is shown to five different Turkers. We ask questions regarding the faces they see in the image. They answer the questions about each image by identifying how many human faces they see in the image, how many of them are female and how many are male. We then ask Turkers to categorize the faces into different age groups and report the number of people in each age group. We specifically ask Turkers to only report the human faces and avoid reporting the people they see in the picture if the faces are not visible.

We take the majority votes on each image and the results of Mechanical Turkers evaluations are in agreement with API output 97% ± 0.75% of the time. 93% ± 1.11% of the Turker evaluations are in agreement in detecting faces with age range under 18, 96% ± 0.86% in age range agreement between 18 and 35, and 99% ± 0.44% in detecting ages over 35. Overall the evaluation of our face detecting API shows high accuracy.

The results of the human validation are summarized in Table 2.

RESULTS

We use negative binomial regression to model the number of likes and comments on photos. The results of the regression are presented in Table 3. We use the Chi-squared Test to find the significance of the regression model, by computing the reduction in deviance from a null model. For our likes model, we find reduction in deviance of $\chi^2 = (5.2M - 1.2M)$, or 76%, on 8 degrees of freedom. The test rejected the null hypothesis of a null model ($p < 10^{-15}$); hence, the model is well-suited to characterize the effects of the described variables.

For our comments model, the reduction in deviance is $\chi^2 = (1.79M - 1.1M)$, or 38%, on 8 degrees of freedom. The test rejected the null hypothesis of a null model ($p < 10^{-15}$). The model for comments is also well-suited to characterize the effects of the independent variables.

We test coefficients of all independent variables for the null hypothesis of a zero-valued coefficient (two-sided) and find that the test rejects the null hypothesis ($p < 10^{-5}$) in all cases.

Effect of control variables

We use number of followers and number of photos as our control variables. As expected, the followers count has a large positive effect on the number of likes and comments. This

Validation test	Accuracy	Margin of error
has face	97%	0.75%
has female face	96%	0.86%
has male face	96%	0.86%
has face < 18 years old	93%	1.11%
has face between 18 and 35 years old	96%	0.86%
has face > 35 years old	99%	0.44%

Table 2. Results of Mechanical Turk evaluation for our face detection approach. Margin of Error is computed for 95% confidence. Our face detector works correctly $97\% \pm 0.75\%$ of the time.

means the higher the number of followers, the more likely it is for the photo to receive likes and comments. The higher number of followers guarantees a larger audience and so the photo is expected to be seen by more number of people, increasing the likelihood of receiving likes and comments.

On the other hand the number of photos shared by user shows a negative effect on both likes and comments. The number of photos is an indicator of activity on Instagram. As we can see in our results (Table 3) the higher activity (number of photos), the lower chances of receiving likes and comments. This might also be interpreted another way: the more photos a user has, the lower probability any single one has of being liked or commented on.

Effect of faces

All other predictors in our model come from the face detection results. We are interested in quantifying the effect of faces and their comparative importance on social engagement. We use a binary variable that reflects the existence of a face in the image. We can see in Table 3 that number of likes and comments are significantly higher when there is at least one face in the image ($\beta_{likes} = 0.32$, $\beta_{comments} = 0.28$, $p < 10^{-15}$). This means that photos with faces are 38% more likely ($IRR = 0.38$) to receive likes and 32% more likely ($IRR = 0.32$) to receive comments⁴.

We also check the effect of number of faces on engagement and find that while existence of a face positively correlates with the number of likes and comments, the number of faces does not particularly change this effect. Regardless of whether it is a group photo or a single person's photo, the fact that a face is in the image significantly impacts the number of likes and comments. It does not matter how many faces are in the image. We did not include the number of faces in the final model to avoid co-linearity of the predictor variables.

Effects of age and gender

To test whether the demographic of users [4] biases toward photos with younger face groups, we considered using three binary variables each identifying the age of a face. Table 3 shows that the age group of the faces are generally not strong predictors for the number of likes. In case of number of comments the photos with adult age groups negatively affects the number of comments. This could be related to lower presence of older age groups in the social network of Instagram.

⁴We use IRR to refer to *Incidence Rate Ratio*. We compute IRR for a categorical independent variables x as the ratio of amount of change in the dependent variable (outcome) for x relative to a reference level of x .

Gender of the faces in the image does not show any strong effect on the image's engagement. Table 3 shows that the β coefficients for gender variables are of negligible size compared to some other features such as existence of a face.

DISCUSSION

Using Instagram as our research context, we set out to investigate how photos with faces relate to engagement, as measured as the number of likes and comments, compared to those without. We considered presence of a face in a photo, its gender and age as predictors, controlling for social network reach and activity. From this we asked two research questions: are photos with faces more engaging than those without and if so how do the characteristics of a face in a photo affect engagement?

As expected, we find that among the factors we measured, the number of followers is the main driver of engagement for both likes and comments on the photo. The number of followers is a proxy for the size of a user's audience. Having a larger audience increases the likelihood of a like or comment, a common sense fact realized in our models. Furthermore, we find that activity level is negatively correlated with likes and comments. The more photos a user posts, the less likely it is that her photos receive likes and comments. As we mentioned earlier, this most likely represents the intuition that the more photos a user posts, the less likely any one of them is to be highly liked or commented.

Faces engage us

The major finding of this paper is that the existence of a face in a photo significantly affects its social engagement. This effect is substantial, increasing the chances of receiving likes by 38% and comments by 32%. We also find that number of faces in the image does not have significant impact on engagement. Having a photo with a face, regardless of how many faces are in the photo, increases the likelihood of receiving likes and comments. Our findings connect to the findings from offline studies in psychology, marketing and social behavior, as well as qualitative studies of HCI, confirming that people engage more with photos of faces.

Age and gender do not impact engagement

Our results show that the age and gender of faces in the photo does not seem to drive or hinder its engagement values. This is a surprising finding, given the bias of demographics using the site and the general belief that photos of kids or female faces may get more attention. For comments, we see in the results that there is a small negative effect of older adult photos. Since the comments are mostly related to the extent to which a photo is discussed, the lower number of comments on this type of photos can be related to the lower demographics of older adults on Instagram. Future work can look at effect of similar factors on other photo sharing communities such as Pinterest with biased gender demographics.

Implications and future work

This work raises many fundamental questions about the nature of social interaction around multimedia content. We believe that this is an initial step and there is a rich landscape

Variable	β	Std.Err	p	Variable	β	Std.Err	p
number of followers	1.32	0.00	$< 10^{-15}$	number of followers	0.97	0.00	$< 10^{-15}$
number of photos	-0.21	0.00	$< 10^{-15}$	number of photos	-0.12	0.00	$< 10^{-15}$
has face	0.32	0.01	$< 10^{-15}$	has face	0.28	0.00	$< 10^{-15}$
has face <18 years old	0.02	0.01	$< 10^{-15}$	has face <18 years old	-0.01	0.01	$< 10^{-15}$
has face >18 and <25	-0.03	0.01	$< 10^{-15}$	has face >18 and <25	-0.07	0.01	$< 10^{-15}$
has face >25 years old	-0.03	0.01	$< 10^{-15}$	has face >25 years old	-0.04	0.01	$< 10^{-15}$
has female face	-0.04	0.01	$< 10^{-8}$	has female face	-0.01	0.01	$< 10^{-4}$
has male face	-0.02	0.01	$< 10^{-3}$	has male face	-0.02	0.01	$< 10^{-6}$
(Intercept)	3.47	0.00	$< 10^{-15}$	(Intercept)	3.47	0.00	$< 10^{-15}$
Null deviance			5208940	Null deviance			1790136
Residual deviance			1227787	Residual deviance			1105145

Table 3. Results of negative binomial regression with number of likes (left) and number of comments (right) as dependent variable.

of research directions and open questions in this area. Future work can look at other visual characteristics of multimedia and study their impact on online behavior. Here we find that faces might have an impact on engagement, but faces are just one visual feature. Other signals can be gathered from people in photos, including facial expression, gaze direction, as well as, body posture and movement. Although facial expressions reliably signal the so-called basic emotions such as fear or happiness, human viewers are also surprisingly adept at making reliable judgments about social information from impoverished stimuli, such as faint changes in facial expressions [30]. Emotional expressions in faces are known to activate several areas of the brain [24]. Future work can look at emotional expressions of faces and explore the effects on user behavior. For example, are we more likely to comment on wry smiles or broad grins?

Our quantitative results illuminate what is the response to the photos with faces, but not why users behave this way or what kind of connections they make with such photos. Additional work, particularly using qualitative methods, is needed to answer these questions. Some of the most compelling questions concern the person in the photo; for example, are users engaging with faces as generic objects or are they connecting with the face as a person they know.

As our work is based on quantitative studies and observational data, we cannot make any strong causal claims. We find that photos with faces have higher chances of being liked and commented on, but we don't know if faces are the exact cause of this. More experimental work needs to corroborate these findings. Further, the statistical methods we used examine only a small segment of behavior on the site.

Faces and their presence connect to psychological studies of human behavior, and emphasize the importance of engaging our unconscious perceptual biases—instantiated in this work as face perception. Future work can investigate the relationship between face perception theories and other aspects of online user behavior. For example are faces effective when it comes to spreading the content on the social network? Are photos or topics, accompanied by human faces more/less persuasive in terms of delivering the content?

The context in which faces appear also invites interesting questions about individual and group behavior. Are photos of friends group more/less popular than the family ones? What about selfies and people's reaction to self portraits? It is worth studying the cultural impacts on photo sharing, say for example are group photos more engaging in collectivism cultures rather than individualistic ones?

Camera-phones and mobile photo capture has changed how we perceive photo-work in the academic community. This work takes one of the first steps into understanding modern photo capture and consumption through the study of Instagram. That said, Instagram is one online ecosystem and it has been claimed that perception and semantics in social media sites is a construction of the community on that site [42]. For example, Instagram is a people-centric site and the influence of faces might be different in a product-centric site such as Pinterest. On the other hand a community such as Instagram, which is strongly based on social connections might react differently to faces than a professional photography community such as Flickr.

The practical implication of social engagement in online photo sharing lies strongly in search and recommendation. Knowing photos with faces increase engagement suggests one could increase their search ranking to keep people on site and active. Our results highlight the importance of effective methods that take advantage of presence of faces in photos for personalization of site content. Additionally, while we have seen face finding applications for social media sites [37], these tools have been designed for the utility of retrieval and not for conversation and comments.

For designers, the present findings may shed light on how to filter, prioritize and highlight photos from the global image stream, especially ones that have just been submitted and therefore haven't had time to accumulate very many likes and comments.

CONCLUSION

Faces are shown to be powerful visual tool used in human non verbal communication. With the widespread use of image sharing communities, most of which are on top of social

platforms, a key challenge in research community is to understand the role of the image content in online user behavior. In this paper, we took a first step toward uncovering an important feature of some of images, the human faces. We find that photos with faces are 38% more likely to be liked and 32% more likely to be commented on. Our results, however, show that number of faces, their age and gender do not have significant impact. In addition to speaking to the ongoing studies in online user behavior and social engagement, our findings open a new thread of future work, suggesting research in visual analysis.

REFERENCES

- 150,000,000 photos. <http://blog.instagram.com/post/8758450298/150-000-000-photos>, August 2011.
- Here's how to use instagram. <http://www.businessinsider.com/instagram-2010-11?op=1>, May 2011.
- Testing, testing... <http://blog.instagram.com/post/8758396471/testing-testing>, July 2011.
- A demographic portrait of users of various social media services. <http://pewinternet.org/Reports/2013/Social-media-users/Social-Networking-Site-Users/Demo-portrait.aspx>, 2013.
- Face++. <http://en.faceplusplus.com/>, 2013.
- Arguello, J., Butler, B. S., Joyce, E., Kraut, R., Ling, K. S., Rosé, C., and Wang, X. Talk to me: foundations for successful individual-group interactions in online communities. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM (2006), 959–968.
- Baker, M. J., and Churchill J., G. A. The impact of physically attractive models on advertising evaluations. *Journal of marketing research* (1977), 538–555.
- Beckwith, L., Kissinger, C., Burnett, M., Wiedenbeck, S., Lawrance, J., Blackwell, A., and Cook, C. Tinkering and gender in end-user programmers' debugging. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM (2006), 231–240.
- Berger, J., and Milkman, K. L. What makes online content viral? *Journal of marketing research* 49, 2 (2012), 192–205.
- Bruce, V., and Young, A. *In the eye of the beholder: the science of face perception*. Oxford University Press, 1998.
- Burke, M., Joyce, E., Kim, T., Anand, V., and Kraut, R. Introductions and requests: Rhetorical strategies that elicit response in online communities. In *Communities and Technologies 2007*. Springer, 2007, 21–39.
- Burke, M., and Kraut, R. Mind your ps and qs: the impact of politeness and rudeness in online communities. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, ACM (2008), 281–284.
- Burke, M., Marlow, C., and Lento, T. Feed me: motivating newcomer contribution in social network sites. In *Proceedings of the SIGCHI conference on human factors in computing systems*, acm (2009), 945–954.
- Cameron, C. A., and Trivedi, P. K. *Regression analysis of count data (econometric society monographs)*. Cambridge university press, September 1998.
- Cassell, J., et al. Nudge nudge wink wink: elements of face-to-face conversation for embodied conversational agents. *Embodied Conversational Agents* (2000), 1–27.
- Cunha, E., Magno, G., Almeida, V., Gonçalves, M. A., and Benevenuto, F. A gender based study of tagging behavior in twitter. In *Proceedings of the 23rd ACM conference on Hypertext and social media*, ACM (2012), 323–324.
- Darwin, C. *The expression of the emotions in man and animals*. Oxford University Press, 1998.
- Dehn, D. M., and Van Mulken, S. The impact of animated interface agents: a review of empirical research. *International journal of human-computer studies* 52, 1 (2000), 1–22.
- Desmarais, C. Facebook's instagram says it has 90 million monthly active users, February 2013.
- Diener, E., Fraser, S. C., Beaman, A. L., and Kelem, R. T. Effects of deindividuation variables on stealing among halloween trick-or-treaters. *Journal of personality and social psychology* 33, 2 (1976), 178.
- Ekman, P., and Friesen, W. *Pictures of facial affect*. consulting psychologists press, 1975.
- Ellison, N., Steinfield, C., and Lampe, C. Connection strategies: social capital implications of facebook-enabled communication practices. *New media & society* 13, 6 (2011), 873–892.
- Gilbert, E., Bakhshi, S., Chang, S., and Terveen, L. I need to try this?: a statistical overview of pinterest. In *Proceedings of the SIGCHI conference on human factors in computing systems*, ACM (2013), 2427–2436.
- Goldman, A. I., and Sripada, C. S. Simulationist models of face-based emotion recognition. *Cognition* 94, 3 (2005), 193–213.
- Goodman, L. A. Snowball sampling. *The Annals of Mathematical Statistics* 32, 1 (1961), 148–170.
- Haxby, J., Hoffman, E., and Gobbini, M. The distributed human neural system for face perception. *Trends in cognitive sciences* 4, 6 (2000), 223–233.
- Herring, S. C. Gender and power in on-line communication. *The handbook of language and gender* (2008), 202.

28. Herring, S. C., Kouper, I., Scheidt, L. A., and Wright, E. L. Women and children last: The discursive construction of weblogs. In *Into the blogosphere: Rhetoric, community, and culture of weblogs*, Citeseer (2004).
29. Hochman, N., and Schwartz, R. Visualizing instagram: Tracing cultural visual rhythms. In *Proceedings of the workshop on social media visualization (socmedvis) in conjunction with the sixth international AAAI conference on weblogs and social media (ICWSM-12)* (2012).
30. Iacoboni, M., Koski, L., Brass, M., Bekkering, H., Woods, R., Dubeau, M., Mazziotta, J., and Rizzolatti, G. Reafferent copies of imitated actions in the right superior temporal cortex. *Proceedings of the national academy of sciences* 98, 24 (2001), 13995–13999.
31. Jamali, S., and Rangwala, H. Digging digg: comment mining, popularity prediction, and social network analysis. In *International Conference on Web Information Systems and Mining, 2009. WISM 2009*, iee (2009), 32–38.
32. Jenkins, R., and Burton, A. M. 100% accuracy in automatic face recognition. *Science* 319, 5862 (2008), 435.
33. Johnson, M. H., Dziurawiec, S., Ellis, H., and Morton, J. Newborns' preferential tracking of face-like stimuli and its subsequent decline. *cognition* 40, 1 (1991), 1–19.
34. Kanwisher, N., McDermott, J., and Chun, M. M. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience* 17, 11 (1997), 4302–4311.
35. Kelsey, R. E., and Stimson, B. *The meaning of photography*. Clark Art Institute, 2008.
36. Laurel, B. Interface agents: Metaphors with character. *Human values and the design of computer technology* (1997), 207–219.
37. Lei, Y.-H., Chen, Y.-Y., Iida, L., Chen, B.-C., Su, H.-H., and Hsu, W. H. Photo search by face positions and facial attributes on touch devices. In *Proceedings of the 19th ACM international conference on Multimedia, MM '11*, ACM (New York, NY, USA, 2011), 651–654.
38. Millen, D. R., and Patterson, J. F. Stimulating social engagement in a community network. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*, ACM (2002), 306–313.
39. Morton, J., and Johnson, M. H. Conspic and conlern: a two-process theory of infant face recognition. *Psychological review* 98, 2 (1991), 164.
40. Phillips, T. Unmasking the face on mars. new high-resolution images and 3d altimetry from nasa's mars global surveyor spacecraft reveal the face on mars for what it really is: a mesa.
http://science.nasa.gov/science-news/science-at-nasa/2001/ast24may_1/, May 2001.
41. Sergent, J., Ohta, S., and Mdonald, B. Functional neuroanatomy of face and object processing: a position emission tomography study. *Brain* 115, 1 (1992), 15–36.
42. Shamma, D. A., Shaw, R., Shafton, P. L., and Liu, Y. Watch what i watch: using community activity to understand content. In *MIR '07: Proceedings of the international workshop on Workshop on multimedia information retrieval* (2007), 275–284.
43. Sproull, L., Subramani, M., Kiesler, S., Walker, J. H., and Waters, K. When the interface is a face. *Human-Computer Interaction* 11, 2 (1996), 97–124.
44. Suh, B., Hong, L., Pirolli, P., and Chi, E. H. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *2010 IEEE Second International Conference on Social Computing (SocialCom)* (2010), 177–184.
45. Takeuchi, A., and Nagao, K. Communicative facial displays as a new conversational modality. In *Proceedings of the INTERSECT'93 and CHI'93 conference on human factors in computing systems*, ACM (1993), 187–193.
46. Takeuchi, a., and naito, t. Situated facial displays: towards social interaction. In *Proceedings of the SIGCHI conference on human factors in computing systems*, acm press/addison-wesley publishing co. (1995), 450–455.
47. Vasalou, A., Joinson, A., and Courvoisier, D. Cultural differences, experience with social networks and the nature of true commitment in facebook. *International journal of human-computer studies* 68, 10 (2010), 719–728.
48. Viola, p., and jones, m. j. Robust real-time face detection. *International Journal of Computer Vision* 57, 2 (2004), 137–154.
49. Walker, J. H., Sproull, L., and Subramani, R. Using a human face in an interface. In *Proceedings of the SIGCHI conference on human factors in computing systems*, ACM (1994), 85–91.
50. Wojnicki, A., and Godes, D. Word-of-mouth as self-enhancement. *HBS marketing research paper*, 06-01 (2008).
51. Wright, J., Yang, A., Ganesh, A., Sastry, S., and Ma, Y. Robust face recognition via sparse representation. *Pattern analysis and machine intelligence, IEEE Transactions on* 31, 2 (2009), 210–227.
52. Yew, J., and Shamma, D. A. Know your data: Understanding implicit usage versus explicit action in video content classification. *IS&T/SPIE Electronic Imaging (January 2011)* (2011).
53. Young, A., Mcweeny, K., Hay, D., and Ellis, A. Matching familiar and unfamiliar faces on identity and expression. *Psychological research* 48, 2 (1986), 63–68.