PHOTO ENGAGEMENT: HOW PRESENTATION AND CONTENT OF IMAGES IMPACT THEIR ENGAGEMENT AND DIFFUSION

A Thesis
Presented to
The Academic Faculty

by

Saeideh Bakhshi

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Computer Science

Georgia Institute of Technology
December 2014

PHOTO ENGAGEMENT: HOW PRESENTATION AND CONTENT OF IMAGES IMPACT THEIR ENGAGEMENT AND DIFFUSION

Approved by:

Dr. Eric Gilbert, Committee Chair
School of Interactive Computing
Georgia Institute of Technology

Dr. Irfan Essa
School of Interactive Computing
Georgia Institute of Technology

Dr. Munmun De Choudhury
School of Interactive Computing
Georgia Institute of Technology

Dr. Polo Chau
School of Computational Science and
Engineering
Georgia Institute of Technology

Dr. David A. Shamma
HCI Research Group
Yahoo Labs

Date Approved: August 6 2014

To my parents, for their utmost love and support

ACKNOWLEDGEMENTS

During my five years of PhD journey at Georgia Tech, I have gone through so many ups and downs and experienced both good and bad times. Throughout the process, I was fortunate to have the unwavering support of my loved ones and to meet so many wonderful people who made my years of PhD more alive.

First and foremost, I would like to thank my advisor Prof. Eric Gilbert. When I first started working with Eric, I was already three years into graduate school, confused and undecided about continuing PhD. Eric believed in me and the work I wanted to pursue. He has been the ideal advisor, providing as much support as I needed and always being available to meet and discuss research, while letting me explore plenty of side-tracks. Not only he is a brilliant researcher with tons of great ideas, but also he is incredibly warm with my best interest at heart. He taught me ins and outs of research process. I cannot thank him enough for the freedom he gave me in pursuing topics I am interested in. He helped me gain independence and confidence in my research. If one day I become an advisor or a mentor to a student, he is the role model I want to look up to.

I'm also grateful for Dr. David A. Shamma's mentorship. He has shown me how to do research in a highly dynamic industry environment that is both useful to products as well as fundamental to research. I learned many research and technical skills from him during my internship at Yahoo Labs and I'm hoping to pick up more as I work with him at Yahoo Labs in the coming years. As a co-author and part of my PhD committee, David provided me with tons of feedback and insights. I would also like to thank my other committee members: Prof. Irfan Essa, Prof. Munmun De Choudhry and Prof. Polo Chau for their helpful comments and insights that helped me improve the work in this dissertation.

Next, I would like to thank my academic siblings at Comp.Social lab. Catherine Grevet, Tanushree Mitra, Chaya Hiruncharoenvate and CJ Hutto, thank you for sitting through my practice presentations, reading drafts, chatting in the lab, and being there for me. I couldn't

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

The type of media shared through social media channels has shifted from text content to include an increasingly large number of images. Visual traces resulting from people's online social behavior have the potential to reveal insights about our habits, activities and preferences. The role of social network-related factors have been well studied in previous research. Yet, few studies have sought to understand how user behavior in social networks is dependent on the image itself. The goal of my dissertation is to understand how people engage with image content, and I seek to uncover the role of presentation and image content on people's preferences.

To achieve this goal, I study the image sharing communities, Flickr, Instagram and Pinterest, using quantitative and qualitative methods. First, I show how colors – a fundamental property of an image – could impact the virality of an image on Pinterest. I consider three dimensions of color: hue, saturation and brightness and evaluate their role in the diffusion of the image on Pinterest, while controlling for social network reach and activity.

Next, I shift the focus from abstract colors to a higher-level presentation of images. I study the role of filters on the Flickr mobile application as proxies to visual computation. To understand how people use filters, I conduct an interview study with 15 Flickr mobile users about their filter use. I analyze Flickr mobile images to discover the role of filters in engaging users.

Presentation is not the only factor that makes an image interesting. To gain deeper insights in what makes an image more engaging in social image sharing sites, I study the images of people on the Instagram network. I compare images of people with those that do not have faces and find that images with human faces are more engaging. I also look at the role of age and gender of people in the image in engaging users.

Finally, I examine different content categories, with and without filters, and study the impact of content category on engagement. I use large-scale data from Flickr and interviews

with Flickr mobile users to draw insights into filter use and content engagement.

This dissertation takes a first step toward understanding content and presentation of images and how they impact one aspect of user behavior online. It provides several theoretical and design implications for effective design, creation and imposition of rules on image sharing communities. This dissertation opens up a new direction for future research in multimedia-mediated communication.

CHAPTER I

INTRODUCTION

What makes images engaging? Consider the two images taken from the Flickr social network in figure 1. The image in figure 1a was uploaded on November 30, 2008. It has been viewed 425 times, received 7 favorites and 36 comments and it is attributed to 7 tags. The image in figure 1b was uploaded four years later, on November 9, 2012 by a different user. It has received 3,942 views, 57 favorites and 2 comments. The image does not have any manual tags, but since it was uploaded through the Instagram app, it has Instagram machine tags. It was processed using a post-processing filter on Instagram. Aside from the influence of the user who uploaded the image and the efficiency of the tags associated with the images, what makes these images interesting? Why does one receive more comments and favorites than the other? Can we explain the engagement with these images only by the properties of the social network around them? How do aesthetics and color composition make a photo interesting? What is the role of color compositions in making a photo (e.g. figure 1a) an engaging image? Can filters make an image more interesting? Are photos of people more engaging that photos of landscape? Does the age of human subjects in the photo make the photo more appealing (See figure 1b)?

The photo-sharing applications are widely used on top of social networking platforms. Billions of photos are uploaded every day on Instagram, Flickr, Facebook and other social networking sites. The visual traces that are left as a result of these online social behaviors have the potential to reveal much about our habits, activities and preferences. Qualitative techniques developed by visual studies scholars can help us to understand the ways in which these images reflect social and cultural norms [116, 146, 174]. Computer vision has made it possible to automate the process of cataloging the content of these visual artifacts [?, 210], and data mining techniques have been designed to surface patterns of visual features across large collections of images [28]. These approaches each provide distinct benefits for working

(a) Photo received 425 views, 36 comments and 7 favorites on Flickr. Photo was originally uploaded to Flickr (Photo: ⓒ①◎ *krzychud1* http://flic.kr/p/5FBiNY.)

(b) Photo received 3,942 views, 2 comments and 57 favorites on Flickr. Photo was filtered on Instagram and shared on Flickr (Photo: ⓒ①◎ *lylevincent* http://flic.kr/p/ds4AS3.)

Figure 1: Example photos on Flickr with their engagement values.

with images, however they are rarely integrated. The research that investigates the image properties is often detached from the social meanings of the photograph and the research that focuses on the social behavior often ignores the content of images. My research bridges this gap by connecting social and visual research through studies of image engagement.

Images are viewed as having the ability to perform multiple *semiotic*, or meaning making roles [116]. These roles are based on the idea and the social meaning of the image. They include (i) representing ideas; (ii) mediating interactions between makers and viewers; (iii) providing genre specific cues for meaning making activities. The heart of my work is focused on understanding images through this social semiotic approach. Take, for example, the photograph presented in figure 2 which was posted on Flickr. First, and the most obvious, is idea of this image. Here, the photograph depicts a person. The social meaning of the photo is dependent on mutual understanding of how a person looks like. This is inherent in our visual system. Second, the photograph embodies interactions between the person who posted it and the people who view the post. The social meaning of the photograph also

Figure 2: Photo by: ©①◎ *throughpaintedeyes*
`http://flic.kr/p/91Xwqp`.

comes from its ability to create a moment of shared experience. Although the photo itself is central, evidence of this shared experience becomes visible when people interact with the photo. For example, on Flickr, some of these interactions take place by *viewing*, *favoriting* or *commenting* on the image. Third, the appearance of an image on the feed carries social meaning related to the visual mode in which it was created: candid snapshot or artfully arranged still life. Considering the photo in figure 2, we can see that the background is blurred and the focus of the camera is on the face of the person. This is a photography technique which makes the portrait appear with a depth of filed. On the other hand, the photo was filtered to remove the colors. Removal of colors and presenting the photo in black and white makes it look more artistic and perhaps helps to convey the message the photographer intended to communicate with the viewer. From this example, we can see that the analysis of images is complex, but also has a familiar rationale.

I was inspired by this simple, human perceivable, identification of cues from the image to understand the social meanings carried out by the image. I extract several pieces of information from images uploaded to social media sites and use them to model engagement
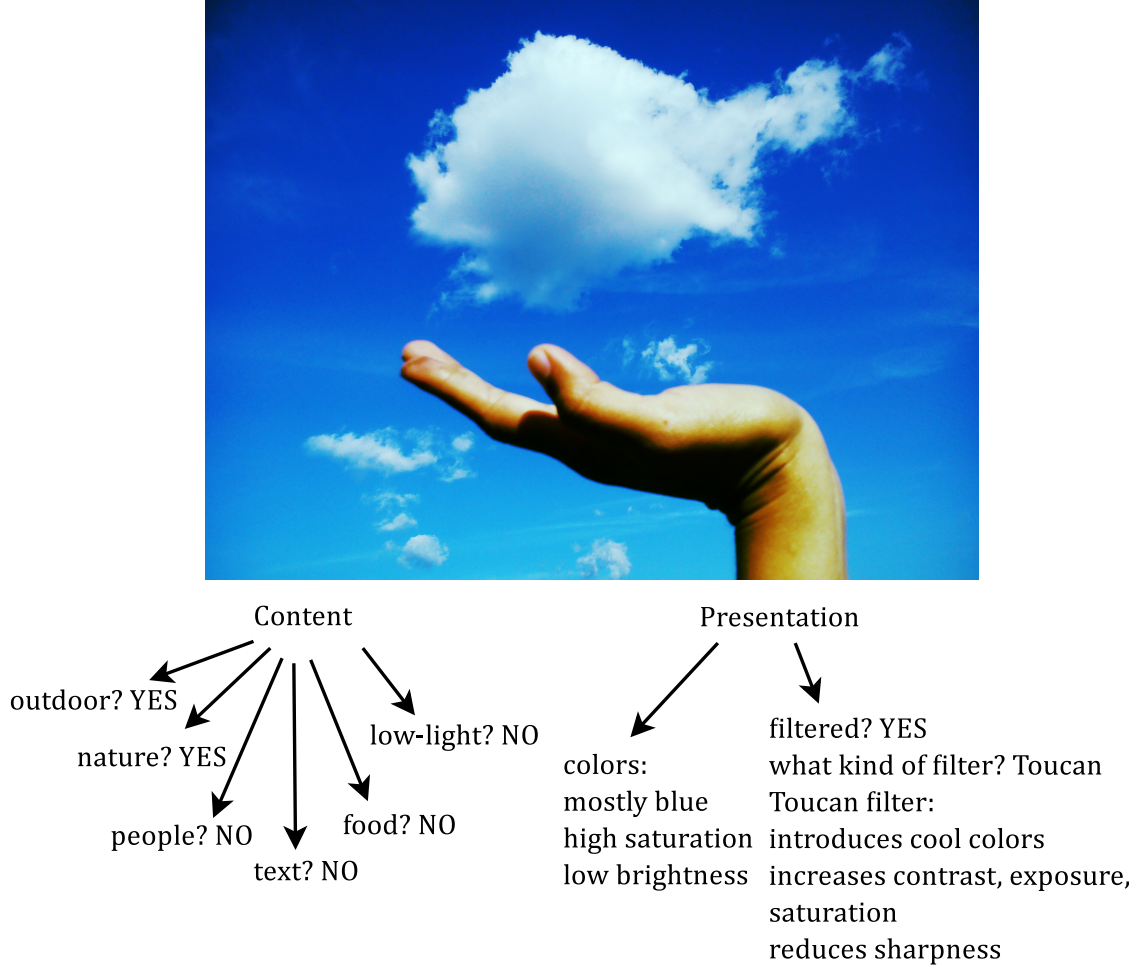
Figure 3: Example photo from Flickr (Photo: ©①◎ *situkangsampah* `http://flic.kr/p/fA98Sw`.) and the two dimensions of presentation and content.

and diffusion. I control for the well known factors influencing engagement such as network structure and user activity. My research leverages the computational tools made available by computer vision to uncover the role of images in shaping interactions around online multimedia content. I focus on two dimensions of information in the image: the presentation and the content (See figure 3).

The first dimension of information that O look at is the presentation of a photo. At the pixel level, I look at colors. I pull from psychological studies of color to investigate the role of colors on photo engagement. Colors are shown to evoke emotions. I ask how these emotions affect our online behavior? Do people like some photos simply because the colors are nice? Filters are another visual computation tools that process the presentation of a

photo. They are applied by users, so there is a user dimension to this type of information in the image. Are filtered photos more likely to be engaging than non-filtered ones?

The second dimension of information available in images is the content of the image. Was the photo taken outdoors? Was it taken under low-light? Does it include natural landscapes? How about people? What is the age and gender of people in the photo? These are some of the content categories we can extract from images using visual computation. I construct models that take these information from images along with the network structure and activity and find the role of each in engaging the viewer. Further, I ask what presentation works the best with each type of content. Are photos of people more engaging when they are filtered? How about natural landscape photos? Answering these research questions is a step towards understanding the underlying behavior that depends on images.

My research directly connects to studies in cognitive psychology where researchers evaluate the role of emotionally evocative content on offline behavior. I extend this area by adding findings from the online world. My work shows how certain properties of images such as colors, post processing and content category might impact the reaction and behavior of people in an online setting. This work can be expanded by deeper analysis of content, for example emotion in the faces.

My work also contributes to multimedia-mediated communication and human computer interaction in many ways. I show here how image as one type of multimedia might impact various aspects of user interaction with the content. The presentation and content of images might directly affect the communication among users as well. I shed light on effectiveness of communication through images and how what is inside an image can convey meanings to the viewer and attract their attention.

My approach in this dissertation is a combination of *(i)* empirical work, measurements and experiments, *(ii)* qualitative studies on users of image sharing communities, *(iii)* explanatory modeling and analysis of user behavior and *(iv)* drawing implications for design of effective social platforms based on images and image components. The goal of this dissertation is to understand impact of various components of image media on users interactions and their engagement with image content. I conduct the research at multiple levels looking

at how different building blocks of images: colors, filters (as proxies to visual computation), content of images such as faces, natural scenes and objects, impact the behavior that affects interaction with content.

1.1 Problem Statement and Research Questions

In this dissertation, I wish to provide nuance into how image presentation and content affects its virality and engagement on social network sites. I pull from the areas of cognitive psychology, HCI and Multimedia-Mediated Communication in order to ground my analysis and methodology.

I look at this area as a social computing researcher, with focus on understanding practices of users around visual content. The overarching question here is: How does image content and presentation reflect on user interactions around the image? I show that microspocic and macroscopic properties of images (such as colors and post processing) as well the subjects in the image are of substantial importance in shaping its engagement. I then set to quantify role of each by developing new models of engagement and diffusion that take into account the properties of the image itself. This dissertation is divided into four research questions:

RQ1: What is the role of colors in image diffusion in the network? To answer RQ1, I conduct a quantitative study of the social networking site Pinterest. I collect images and their meta-data from Pinterest, and build a statistical model to find the relationship between the number of repins and the dominant color in the image. The results are reported in Chapter 3.

RQ2: Why do people use filters on images and how does it impact engagement? To explore RQ2, I conduct two studies. The first one is a qualitative study where I interview 15 Flickr mobile app users to understand their motivations behind filtering their photos. Then I perform a large-scale analysis of Flickr mobile images to understand how filters impact image engagement.

RQ3: How do human faces in images impact engagement? To explore RQ3, I use a large dataset from Instagram images. I apply computer vision algorithms to detect faces,

and their age and gender in the image. I build a model of image engagement as a function of several control variables and face attributes.

RQ4: How does the content of an image impact its engagement, and how do content categories interact with filters? Inspired by findings of RQ2 and RQ3, I design a more general study where I investigate the role of different content categories on engagement. I also consider each category in presence and absence of filters. I use a mixed-method approach to perform this study: a large-scale analysis of Flickr mobile images complemented by interviews with 15 Flickr mobile users.

1.2   Overview of this Dissertation

This dissertation is organized as follows. In chapter 2, I summarize the theoretical and practical foundations of this work based on previous research literature. In Chapter 3, I present a quantitative study of colors on Pinterest. In this chapter, I show how colors impact diffusion and adoption of images on Pinterest. In Chapter 4, I present a mixed-method study of filters on Flickr. I perform a qualitative study, interviewing Flickr mobile app users to find how they use filters on their photos. I then conduct a quantitative study to investigate the role of filters on photo engagement. In Chapter 5, I describe a quantitative study of Instagram, where I explore the differences in engagement between photos of people and other types of photos. I also look at the role of age and gender of faces in the photos. In chapter 6, I compare engagement value of different types of photo content on Flickr. I also measure the interaction between filters and content types to understand what photos are more engaging with and without filters. Finally, in Chapter 7, I conclude with reflections on this dissertation and future research directions.

CHAPTER II

RELATED WORK

This thesis builds on diverse streams of related work. I group the related work into two categories: the fundamental theoretical framework that I use to motivate my work and the practical research that is related to my research.

2.1 Theoretical Framework

In this section, I review the relevant literature supporting theoretical aspects of the current dissertation. I start with the previous work on multimedia perception and it's effectiveness, I then go through the previous work in multimedia-mediated communication. Next, I summarize the findings of psychology and behavioral sciences in perceptions of colors, human faces and natural scenes.

2.1.1 Multimedia, Effectiveness and Perception

The term *multimedia* is often misleading because it does not indicate the existence of multiple media; rather, as Marmolin note: "multimedia implies multiple senses used in processing a stimulus, or multiple modalities or channels used in transmitting a message" [134]. Hoogeveen treats multimedia as a property of a system or object wherein "multiple perceptual representation media, such as speech, music, text, graphic, still, animation and video, are used in an integrated manner" [97].

Regardless of which definition of multimedia one adopts, prior research is inconclusive about the direction of measurable effects of multimedia functionality on online users' memory as well as their perceptions of the interface. Understanding people's perceptions of web-based information systems, or social media sites is a complex phenomenon, requiring a series of systematic studies [70]. However, previous research in communication and psychology might help us to gain intuitions on the relative cognitive effects of multimedia modalities.

Previous research in psychology suggests that each individual modality (e.g., text, photo, video, etc.) contains unique characteristics, and people encode this modality-specific content when they process information [206]. That is why information is considered more memorable when presented in certain modalities than in others. For example, scholars testing the *separate-streams* hypothesis have shown that auditory information is recalled better than visual information [161].

Some researchers have used the concept of *redundancy* to argue that information presented in multiple modalities has a stronger chance of getting through to receivers than information presented in a single modality [98]. This is because repetition of content, in different forms, contributes to cognitive rehearsal, thereby enhancing its likelihood of storage in memory. There is some support for this notion in broadcast news research which shows that the addition of images that are redundant with audio serves to enhance memory for content [170].

Two theories shed some light on memory processes pertaining to multiple modalities of presentation. The *Dual-Coding* theory [159] assumes that there are two cognitive subsystems; one specialized in processing verbal stimuli, and the other specialized in non-verbal or image stimulus that operate independently as far as encoding into memory is concerned. Therefore, this theory would predict that delivering information in two modalities (instead of one) acts as a double dose thereby enhancing the storage potential of that information. The general theoretical perspective of multiple modalities resulting in cognitively superior outcomes (in comparison to individual modalities) has a long history of support in psychology, ever since a study showed that participants in bi-sensory modality had higher recall than those presented to appeal to a single sense [36].

On the other hand, the *Limited Capacity Information Processing* theory [120] and the *Multiple Resource* theory [217] argue that media messages, delivered simultaneously in a number of modalities, are cognitively complex and serve to overload the processing system. In this formulation, recognition measures of memory indicate how much information was encoded whereas recall measures how much information was stored and is available for retrieval [119]. According to this perspective, the addition of multimedia to text-only systems

should result in superior recognition memory but inferior recall memory.

In this dissertation, the goal is to cover this gap, by analyzing the content of images and discovering their relationship with users' perception. As part of this dissertation, I aim to evaluate the multiple modalities presented in the form of text on image as well and compare it with the single modality of the image only.

2.1.2 Multimedia-Mediated Communication

Theory and research that directly address the effectiveness of multimedia in web based communication are relatively scant. Some researchers propose that the value added by introducing multimedia to existing interfaces is merely perceptual. That is, multimedia enhancements will serve to generate positive impressions about a website because of their sheer presence and not because of their greater cognitive utility. Hoogeveen calls this the *self-fulfilling prophecy* hypothesis [97]. A recent experimental study offers limited support for this notion by showing that individuals exposed to a multimedia news site expressed greater likelihood of re-visiting the site than their counterparts exposed to a text-only version of the same site even through the two groups did not differ in the amount of information learned from the site [27].

Some have theorized that adding multimedia components to a compute program will foster positive attitudes toward the program because their dynamic presence facilitates greater involvement and engagement with the system [133]. Biocca argues that the senses are channels to the mind, and since multimedia appeals to a variety of senses, it might generate more immersion with the interface [30]. This is akin to the *begin there* effect documented by Reeves and Nass, whereby formal features of media, such as those used for building multimedia websites, have the potential to create the illusion, at least momentarily, of being transported to the world portrayed in media [171]. Reeves and Nass argue that such perceptions can be generated with relatively low-tech interfaces involving "simple textual and pictorial material shown on garden-variety technology" [171]. Similar ideas characterize the notion of *para-social interaction* in the television research literature [128] and the concept of *social presence* in technology studies [187].

Steuer argues that vividness is the key [192]. He conceptualizes the vividness of an interface in terms of its sensory breadth (i.e., the number of senses engaged by it) and sensory depth (the resolution within each perceptual channel). He operationalized vividness in terms of modality, and ranks text as being low in vividness and moving images with voice as being high in vividness. However, in his experiment that involved presentation of an educational tutorial in a variety of modalities (ranging in vividness from low to high), participants in his text-only condition showed the most positive attitudes toward the system and the most positive effect toward the computerized tutor. Similarly, another experiment shows that attitudes toward a web page with text and graphics were more positive than attitudes toward an identical page that featured animation in addition to text and graphics [13].

On the other hand, research with instructional technology and computer games has shown that adding multimedia leads to positive attitudes such as greater confidence [151], motivation [222], and enthusiasm [74] among users. While the perceptual consequences of adding multimedia functionality is mostly in favor of new media; it is not clear whether users appreciate the extra cost and effort involved in designing, producing, and placing multimedia enhancements in text-only sites, and whether this appreciation leads to more positive attitudes toward the sites and more positive evaluation of content on those sites. Sundar suggested in another study that pictures and audio are powerful psychological cues, hindering memory for story content and leading to negative evaluations of the websites, while improving memory for advertisement [196].

In several early works, researchers have explored the impact of video on the processes of communication and collaboration [48, 61, 157, 183]. Several theories have been proposed which describe how different communication media, including video-mediated communication systems, can be characterized in terms of the facilities they offer, and how these lead to differing impacts upon users. For example, Short et al. [187] describe how the *social presence* of communication media vary and how these in turn affect how users of media interact. They claim that the communication media that provides a wide range of visual cues engenders a strong sense of social presence among users. Daft and Lengel outlined how

communication media differs in *richness*, with some rich media such as face-to-face communication able to support rapid exchange of multiple types of information whilst *leaner* media such as email provide fewer cues and/or slower feedback [54]. Most of the studies within these theoretical frameworks have compared communication media which offer different facilities to users.

Although the literature contains very mixed results on the benefits of richer media such as photo and video-mediated communications systems compared to text and audio-only systems, in terms of task performance, there is generally a reported preference among users for these richer media [11, 183, 199, 215].

Neurologically all images are by nature gestalt, made up of fragments of visual experience processed modularly and then coordinated through perceptual process into what Walter Lippmann called *pictures in our heads* [131]. They are stories, always implying more than their parts, always in process and actively seeking meaning. Because vision developed before verbal language, images are a natural part of our primal sense of being and represent the deepest recesses of ourselves. Although verbal language represents an evolutionary advance in that it allows us to abstract thought from experience, it, too, is of necessity grounded in perception [22]. Birdwhistell estimates that 65 to 70 percent of the social meaning within a conversation is carried by the visual [31]; Mehrabian estimates this to be as high as 93 percent [138].

My dissertation takes the first step towards understanding multimedia contents in the social settings.

2.1.3 Color as Affective Stimulus

Given the ubiquity of color in our lives, it's not surprising that a great deal of research has been conducted over the past century on it. Scientists widely recognize colors as source of impact on our emotions and feelings [89, 111, 121, 132]. Color is believed to affect the degree of felt arousal [143, 214]. This view of arousal has been predominant in both psychology [26] and marketing [156, 177]. Apter and his colleagues [14, 189] used a two dimensional theory to explain arousal and pleasure in color research. According to their theory, there are two

dimensions of arousal: one goes from boredom to excitement, the other goes from tension to relaxation. Both excitement and tension can produce equivalently high states of arousal, but the former would be pleasant while the latter would be unpleasant [214]. Using the two-dimensional view of arousal, they discovered a link between red and felt excitement, and blue and felt relaxation (also see [85, 200]).

Several empirical projects have studied the role of color in affective marketing. One stream of research examines the specific colors used in magazine ads [123, 180]. The second stream of research has investigated the efficiency of colors compared with black and white ads [140, 190]. The third stream has focused on the effects of specific colors on consumer responses [24, 53]. This work suggests, for example, that red backgrounds elicit greater feelings of arousal than blue ones, whereas products presented against blue backgrounds are liked more than products presented against red ones [24, 142]. Increases in arousal are at first pleasurable and exciting, but after a certain point, more will decrease pleasure and increase tension [26].

Color theorists believe that color also influences cognition and behavior through learned associations [65]. Color can affect cognitive task performance, with red and blue activating different motivations and consequently affecting different types of tasks [139]. In another piece, researchers studied the relationship of color and emotion among college students by asking them about the feelings colors brought to mind [111]. They found that the principled hues showed the most positive emotional responses, followed by intermediate hues and achromatic colors.

In addition to hues, studies have examined saturation and brightness as stimuli. More deeply saturated colors can be more exciting and cause surprising behavior [207, 8]. Other research supports the idea that higher levels of chroma are more widely liked [84, 85, 136, 186]. Another study concludes from experiments that hue, chroma and value are linked to consumers' feelings and shopping attitudes. Higher levels of saturation seem to increase excitement, while increases in brightness lead to feelings of relaxation [82].

Prior psychological studies argue that colors are associated with certain abstract concepts. For example, blue is associated with wealth, trust and security; gray is associated

13

with strength, exclusivity, and success; and, orange connotes cheapness [118]. Marketing psychologists suggest that a sustained color impression is made on a subject within 90 seconds and that color accounts for 60% of the acceptance or rejection of an object, place, individual or circumstance [141]. Because color impressions are made quickly and are long lasting, decisions regarding choice of color can be highly important to marketing success [32].

Somewhat anecdotally, color choice has been suggested to impact the market performance of products and brands. For example, the use of a subdued blue for Nabisco's Honeycomb Graham snacks was blamed for its sluggish sales, leading the company to redesign and re-launch the brand [169]. McDonald's use of saturated red for its interiors even led customers to complain of headaches [212]. The poor sales performance of McCormuck led to changes in its advertising and packaging, from drab olive green to deeper forest green [145].

Green is seen as cool, fresh, clear and pleasing, but when illuminated on skin tones it becomes repulsive [52]. While red is associated with excitement, yellow is perceived as cheerful, purple as dignified and stately, and blue connotes comfort and security [216]. Moreover, red is known as a dominant and dynamic color, but this can have both positive and negative effects. Green can be used as a symbol of refreshment, restfulness, quietness, naturalness as well as tiredness and guilt [56, 132]. Purple was mostly associated with children and laughing, a positive association reported by Kaya et al. [111].

Psychological studies suggest that white has a calming effect, producing the least amount of tension [106, 158]. The implication is that higher value, lighter colors should be more relaxing than lower value, darker colors. The well-known design Faber Birren claimed that blue, red, grey, orange and yellow color preferences are nearly identical for both sexes and exist beyond cultural boundaries [32].

This dissertation builds on the previous research by looking for the first time at diffusion as a function of image features. I connect the anecdotal, experimental and theoretical studies of color with the spread of online content.

2.1.4 Face Perception

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 [166].

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 [38]. The human brain has evolved to recognize faces within hours after birth [109]. Human infants only minutes old attend particularly to face-like stimuli relative to equally complicated non-face stimuli [38, 103]. 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 [223].

By the age of two months, infants begin to differentiate specific visual features of the face [148] and process facial expressions [60]. Our brains have a specific region, Fusiform Face Area (FFA), which is specialized for facial recognition [184, 110]. Faces are important for social cognition, 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 [55, 64]. This can be highly important to very practical concerns: faces, particularly attractive ones, are found to be effective in improving consumer responses to advertisements [18].

In HCI research, there is a great deal of work exploring the benefits of using face icons and faces in interfaces [122, 191, 197, 213]. Walker et al. studied how having faces and facial expressions for a computer application affects users' performance and productivity [213]. 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. compared users' impressions of an agent which helped them to win a card game [198]. 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 [45]. In a review study of embodied agents, authors reported that adding an embodied agent to an interface made the experience more engaging [58].

2.1.5 Perception of Nature

Research has found that viewing nature can help reduce stress or help wellness [204]. Although perception of nature is multi-sensory and involves responses to sounds and smells as well as visual content, research to date has focused on influences of viewing nature. The intuitively-based belief that visual exposure to trees, water and other nature tends to produce restoration or recovery from stress dates as far back as the earliest large cities, such as ancient Rome. In the U.S. in the 19th century, intuitively-based arguments about healthful effects of viewing nature were influential in establishing urban pastoral parks, such as New York's Central Park, and later in preserving wilderness for public use [155].

Historically, a theme running through these beliefs is the notion that if individuals are stressed or unhappy, views of most natural settings will have positive influences, whereas views of urban or built settings will tend to impede recuperation, especially if they lack nature content such as vegetation and water. A body of research has tested this old belief that visual contacts with nature have restorative or stress-reducing influences [96, 201, 203, 204]. This research suggests that many nature scenes or elements foster stress recovery because

they elicit positive feelings, reduce negatively toned emotions such as fear, anger, and sadness, effectively hold attention/interest, and accordingly might block or reduce stressful thoughts. Research also indicates that views dominated by nature content, in contrast to built or urban scenes lacking nature, foster more rapid and complete restoration in terms of another critical component of stress, the physiological.

In laboratory research, visual exposure to everyday nature has produced significant positive changes in physiological measures such as blood pressure and muscle tension [205]. Also, a study of unstressed individuals found that slides of nature sustained attention much more effectively through a lengthy viewing session and produced more positive feeling states that did built scenes [202]. In the same study, recordings of brain electrical activity in the alpha frequency range suggested that individuals were more wakefully relaxed during the nature exposures. In summary, these studies indicate that for stressed individuals, restorative influences of viewing nature involve, among other responses, a broad shift in feelings towards a more positively-toned feeling state, positive changes in activity levels in different physiological systems, and that these changes are accompanied by moderately high levels of sustained attention.

2.1.6 Filters and Visual Effects

Recent work on creativity found that digital artifacts that are special are often self-made, such as presentations, animations and photo montages [164]. These results have shown that crafting and making with digital media can make these media more special or cherished, and in fact, being self-made or augmented appears to be one of the main reasons people cherish their digital possessions [80, 153, 175].

Image editing, also known as photo retouching, is the process of altering images. Prior to widespread use of mobile camera phones, filters were mostly used by photographers to correct problems, apply effects or simply to make a photo look different. Even before the age of digital photography, people used to make changes in the dark rooms. They could adjust the contrast, use different papers and different tones on those papers to make their photos look better [182].

Digital images are stored in the form of picture elements or pixels. Process of editing can change the pixels to enhance the image in many ways. The pixels can be changed in a group or individually. Camera or computer image editing programs often offer basic automatic image enhancement features that correct color hue and brightness imbalances as well as other image editing features, such as red eye removal, sharpness adjustments, zoom features and automatic cropping. These tools are complex and more catered to experts rather than end users. Filters, on the other hand, are tools that give users the opportunity to enhance their photos, without the need to go through professional software. Most photo filters manipulate colors, saturation, light exposure or simulate a change in focus. There are different use cases introduced by each filter. Filters can age a photo, make colors more vibrant, or give photos a cooler color temperature. Some filters overlay masks in the image or add borders or frames [137]. Instagram and Flickr are among the most popular mobile photo-sharing platforms that provide users with a variety of filter choices.

In this dissertation, I look at filtering and visual post processing practices on photo sharing communities and the impacts of these visual effects in engaging users.

2.2   *Empirical Literature*

In this section, I summarize the literature on online engagement, social network diffusion and behavior around online images.

2.2.1   Information Diffusion in Social Networks

Several recent empirical and theoretical research papers have addressed diffusion in various social networks. Examples include studies on Facebook diffusion trees [195], diffusion of gestures between friends on Second Life [20], diffusion of health behavior [46] and adoption of mobile phone applications over the Yahoo! Messenger network [15]. In most of these studies, network structure is considered the main driving cause of influence and diffusion. For instance, a study by Kwak et al. compared three network measures of influence (influence is largely regarded as the ability to cause diffusion): number of followers, Page-Rank, and number of retweets [117]. Cha et al. also considered an additional measure, number of mentions on Twitter, to quantify diffusion [47].

Diffusion of information also known as viral marketing, specially when it is about the product and its adoption over the social network. Primarily in social sciences there is a long history of the research on the influence of social networks on innovation and product diffusion. For example, Brown and Reingen interviewed the families of students being instructed by three piano teachers, in order to find out the network of referrals [37]. They found that strong ties, those between family or friends, were more likely to be activated for information flow and were also more influential than weak ties between acquaintances [83]. Similar observations were also made by DeBruyn and Lilien in the context of electronic referrals [57]. They found that characteristics of the social tie influenced recipients behavior but had different effects at different stages of decision making process: tie strength facilitates awareness, perceptual affinity triggers recipients interest, and demographic similarity had a negative influence on each stage of the decision-making process.

Social networks can be composed by using various information, i.e. geographic similarity, age, similar interests and so on. Yang and Allenby showed that the geographically defined network of consumers is more useful than the demographic network for explaining consumer behavior in purchasing Japanese cars [220]. A study by Hill et al. found that adding network information, specifically whether a potential customer was already talking to an existing customer, was predictive of the chances of adoption of a new phone service option [92]. For the customers linked to a prior customer the adoption rate of was 35 times greater than the baseline.

Factors that influence customers willingness to actively share the information with others via word of mouth have also been studied. Frenzen and Nakamoto surveyed a group of people and found that the stronger the moral hazard presented by the information, the stronger the ties must be to foster information propagation [73]. Also, the network structure and information characteristics interact when individuals form decisions about transmitting information. Bowman and Narayandas found that self-reported loyal customers were more likely to talk to others about the products when they were dissatisfied, but interestingly not more likely when they were satisfied [34]. Richardson and Domingos used Epinions trusted reviewer network to construct an algorithm to maximize viral marketing efficiency assuming

that individuals probability of purchasing a product depends on the opinions on the trusted peers in their network [172]. Kempe, et al. have followed up on Richardson and Domingos challenge of maximizing viral information spread by evaluating several algorithms given various models of adoption we discuss next [49].

Other than properties of the social network, one reason that certain content may be highly shared is because it has inherent value or contains useful information. Discount coupons or articles about good restaurants help people save money and eat better. Consumers may share such practically useful content for altruistic reasons (e.g., to help others) or for self-enhancement purposes (e.g., to appear knowledgeable) [218]. Practically useful content also has social exchange value [95], and people may share it to generate reciprocity [71].

Emotional aspects of content may also impact whether it is shared. People report discussing many of their emotional experiences with others, and customers report greater word-of-mouth at the extremes of satisfaction (i.e., highly satisfied or highly dissatisfied [12]). People may share emotionally charged content to make sense of their experiences, reduce dissonance, or deepen social connections [72, 147, 162, 173].

But while this suggests that emotionally evocative content may be particularly viral, which is more likely to be highly shared, positive or negative content? While there is a lay belief that people are more likely to pass along negative news [77], this prediction has never actually been tested. Further, the study on which this idea is based actually focused on understanding what types of news people encounter, not what they transmit [78]. Consequently, researchers have noted that "more rigorous research into the relative probabilities of transmission of positive and negative information would be valuable to both academics and managers," [78], yet little empirical work has examined this issue.

Whereas the focus of most of these studies is largely network structure, content can of course also be the reason behind diffusion. Users may share useful content to appear knowledgeable or simply to help out [218]. The emotional valence behind content can also drive its sharability. For example, Jamali et al. used a Digg dataset to predict the popularity of stories, where sentiment emerges as a major predictor [105]. In another study,

Berger et al. used New York Times articles to examine the relationship between the emotion evoked by content and virality, finding that stories that inspire awe in readers get shared the most [25].

2.2.2 User Engagement

User engagement is the quality of the user experience that emphasizes the positive aspects of the interaction, and in particular the phenomena associated with being captivated by a web application, and so being motivated to use it [125]. According to classification that Lehmann et al. provided, there are three groups of engagement: the self-reported engagement, the cognitive engagement and the engagement obtained through online behavior metrics [126].

The self-reported engagement data is usually obtained by conducting interviews, or collecting responses to questionnaires and interviews (e.g. [152, 178]). Some of these studies are carried out in the lab settings or crowd-sourced through online resources. The drawback of such methods is their subjectivity [188] to the sample of users who participated in the study.

The cognitive engagement method uses task-based models (e.g dual task [181]) and physiological measures to evaluate the cognitive eye tracking [63], heart rate monitoring and mouse tracking [99]. Although this measures are more objective, they are not scalable to the web scale. In contrast, the web-analytics community has been studying user engagement through online behavior metrics that assess users depth of engagement with a site. For instance, Peterson and Carrabis describe engagement metrics that indicate whether or not users consume content slowly and methodically, return to a site, or subscribe to feeds [163]. Widely used metrics include click-through rates, number of page views, time spend on a site, how often users return to a site, number of users, and so on. Only online behavior metrics are able to collect data from millions of users. Although these metrics cannot explicitly explain why users engage with a service, they act as proxy for online user engagement: the higher and the more frequent the usage, the more engaged the user [125].

2.2.3 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" [67]. Sharing content can in turn ensure that user remain engaged and committed in the future [39, 209]. Users have diverse motivations to share content on social network sites (SNSs). For example, users may share useful content to appear knowledgeable or simply to help out [218]. Not only the content of posts, but also the emotional valence behind it can drive its usage. In a study on Digg, the authors predicted the popularity of stories, and used sentiment as a major predictor [105]. In another study, researchers used New York Times articles to examine the relationship between the emotion evoked by content and its virality, finding that that there is a direct relationship [25].

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 interest. 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 [16]. 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 [40]. In another study, Burke and Kraut, studied the effect of the politeness of a post, with the interesting finding that politeness leads to more replies in certain types of groups, while in other types of groups, rudeness actually increases replies [41].

On Twitter, researchers have used retweeting as a measure of community interest, and have investigated what features will predict retweeting. Suh et al. found that the presence of URLs and hash-tags in tweets predicted more retweeting, as did a richer connection with the community [194]. In a recent study, Hutto et al. studied 507 twitter users over 15 months and found that a variety of predictors influenced follower growth, including message content, social behavior and network structure [102].

2.2.4   Images and Social Behavior

Widespread adoption of camera phones is resulting in increasingly large visual datasets representing a range of technology mediated social practices. While techniques for analysis of textual data are well established, systematic interpretation of images is challenging for many reasons. Visual cues and signifiers are mainly subjective and can take on varied significance when seen in combination with other forms of expression. Traditional qualitative approaches to image analysis are typically designed to highlight nuanced social and cultural cues that contribute to an images meaning [146, 174, 193]. However, these techniques can also be time consuming to perform and can produce subjective results that are difficult to generalize.

Automated processes for image recognition using computer vision techniques [28, **?**, 210] are effective for annotating content in very large collections of images but are not designed to support interpretivist techniques for understanding the meaning images carry beyond literal depiction. Previous literature shed light on those aspects of images that can help with social connections. Lin and Faste [129] highlight the potential of images to promote social connections in the online space. Contextual interviews focusing on users' photo sharing, organizing, and viewing behaviors indicated that people are socially motivated by photographs, are selective in what they view, and use photographic narratives to correspond with others and to browse information. Looking across image data collected from several online communities, McDonald [135] identified four types of visual conversation styles evident through posted images (positional play, image quote, text- in-picture, and animation). Observations of the social role of images have also been made by Clarke et al. [50] who looked at the role of photo-sharing in rebuilding lives after domestic violence; O'Hara et al. [154] who studies photo mementos as social mediation during family dinners. Drew et al. [62] used visual storytelling as means of personal elicitation for adolescents faced with chronic disease management and found that this approach increased positive feelings, was considered fun by their participants, and promoted self-understanding and focus during interviews.

Carter and Mankoff [44] point out that allowing participants to control the timing and

method of data capture provides a distinct opportunity for researchers to reflect on the ways in which modality influences the construction of participant narratives. In a diary study using photos, audio clips, and physical objects, they found that images led to more specific recall than any other medium and observed that participants would often post photographs that were not meaningful on their own, but served as recall cues or pointers that signaled specific events.

Photo-sharing studies are often supported by online and mobile applications designed to assist users in documenting and sharing events they identify as important. These apps also enable image-based ecological momentary assessments of transitory states such as mood or pain level [168]. For example, passive capture or contextual photography tools provide users with cameras that can be worn around the neck, automatically taking a certain number of photos during the day at regular intervals [75, 76, 87, 94, 130, 167]. These studies can result in thousands of images, often annotated with caption and comments. Faced with these very large collections of images, researchers need techniques to identify images that are not only interesting from a research perspective but are also meaningful to participants, especially in cases where the images will be used to prompt follow-up interviews [75, 76].

Given the ubiquitous presence of digital cameras in our daily lives, researchers interested in the social aspects of photo sharing have also devoted attention to attitudes regarding ethics, privacy and ownership of images posted online. Gulotta et al. [86] conducted photo-elicitation interviews to explore user motivations for posting provocative, controversial, or deeply personal images, finding that the reasons participants had for sharing these types of images were idiosyncratic and personal. Ahern et al. [9] used context-aware camera phones to examine privacy decisions in mobile and online photo sharing. They identified common themes in users privacy considerations including security, social disclosure, identity, and convenience. In many of the cases described here, researchers took what multi-modal analysts would call a social semiotic [68, 115, 208] approach to understanding the role of images in participant interactions. Images are viewed as having the ability to perform multiple semiotic, or meaning making, roles including (1) representing ideas; (2) mediating interactions between makers and viewers; (3) providing genre specific cues for meaning

making activities [116].

2.3 Chapter Summary

In this chapter, I presented two groups of previous work. The theoretical work that discusses perceptions of various content types in images, the relevant literature on multimedia-mediated communication and its effectiveness. The empirical literature in this chapter provided an overview of research on images in the social setting, previous work on engagement and diffusion of content and value of content in engaging users. The work on multimedia and engagement provide distinct benefits and insights in working with images, however they are not integrated. The research that investigates the images and perceptions of multimedia is often detached from the social meanings of the photograph and the research that focuses on the social behavior often ignores the content of images. My research bridges this gap by connecting social and visual research through studies of image engagement.

CHAPTER III

COLORS AND DIFFUSION OF IMAGES

In 1979, a director at the American Institute for Biosocial Research began observing curious psychological variation in his patients—variation seemingly rooted in what colors he showed them. To test his theory, he convinced the directors of a naval prison to paint their cells pink, believing pink would calm the inmates. What he found was fascinating. Rates of violent behavior fell dramatically after exposure to the plain pink walls. According to the Navy's follow-up report, "Since the initiation of this procedure ...there have been no incidents of erratic or hostile behavior" [42].

In addition to inducing calm, colors can evoke powerful reactions like warmth, relaxation, danger and energy [89, 111, 121, 132]. In short, they have remarkable power to move us emotionally. For example, prior work has shown that red is associated with excitement, yellow with cheerfulness and blue with comfort [216]. In this chapter, I ask: Might these phenomena documented in lab experiments also affect large-scale online behavior? Could color drive how we act on social media?

Recently, we have seen image-sharing communities truly take off—sites like Pinterest, Imgur and Tumblr, just to name a few. A key research-level challenge for communities like these is uncovering the mechanisms by which content spreads from person to person (or "diffuses", adopting the term from the academic literature). For example, a study of the most widely shared New York Times stories found that they tend to "inspire awe" in their readers [25]. While we have results like this for text and network structure (e.g., [20, 195]), as far as we know, we have no such similar results on what makes images diffuse widely. It is within this context that I make the leap from color to diffusion: Is there a link between them?

In this chapter, I aim to answer this question. I adopt Pinterest as the research site; at the time of this writing, it is the fastest growing social network site on the web, recently

growing to over ten million users. Drawing on a corpus of one million images crawled from the site, I find that color significantly drives how far an image diffuses (to what extent it is adopted by other users), even after partially controlling for user activity and network structure. Specifically, red, purple and pink seem to promote diffusion, while green, blue and yellow suppress it. This is the first result describing how image features affect diffusion. This work bridges the gap between online user behavior and psychological studies of color. In addition to speaking to the ongoing research conversation surrounding diffusion, these findings suggest deep future research to both uncover the impact of color on other aspects of online user behavior, and also to use more sophisticated computer vision techniques.

For designers, the present findings shed light on engagement on large social sites. Recently, for example, many mobile applications let users transform their photos with image filters. Instagram popularized this technique, but similar features can be found in Flickr's latest mobile app, for example. The filters often change saturation, brightness, and color distributions. The present results can be used to guide the design of these filters. For example, filters that increase an image's saturation or enhance its warmness will likely increase diffusion, a highly sought-after form of engagement.

3.1 Methods

I take a quantitative approach in this chapter to investigate how color might shape diffusion. In this section, I first provide a short background on Pinterest. I then describe the data I collected and the statistical methods I used. I adopted the Munsell color system to identify and categorize colors; it is a widely used system in psychological and physiological studies of color.

3.1.1 Pinterest

For the purpose of this research, I take the data from Pinterest, a pinboard-style, photo-sharing social network site that allows users to upload images or bookmark them from external websites. It allows users to save and share images and categorize them to "boards". Pinterest is the fastest growing website in the web's recent history, with 429% growth from September to December 2011 [4].

Data Collection:　　**Pixel by pixel analysis:**　　**Color analysis:**

Dominant Hue

Mean Saturation

Mean Brightness

Pin

Pinner

Negative binomial regression analysis

Figure 4: A flowchart of steps taken in this chapter to prepare data for analysis. I collected a random set of pins from Pinterest. Each pin's image was analyzed pixel by pixel, resulting in a dominant hue, mean saturation and mean brightness. Negative binomial regression was used to characterize the effects of colors on repins.

In Pinterest's own words, the site's "goal is to connect everyone in the world through the *things* they find interesting. We think that a favorite book, toy, or recipe can reveal a common link between two people"[1]. This focus on *things* has made Pinterest a favorite among retailers and marketers: for example, a recent market survey showed that Pinterest users are more likely to click through to e-commerce sites, and when they go there, they spend significantly more money than people who come from sites like Facebook or Twitter [5]. At the time of this writing, Pinterest has 70 million users and 2.5 million monthly page views.

3.1.2　Data

While my goal was to obtain a random sample of pins and pinners, there is no publicly available means to do so. Instead, I developed a web crawler to approximate a random sample. Pinterest does not have a page that lists all pins or pinners; however, it does provide a list of "popular" pins[2]. In June 2012, I ran the web crawler to collect popular pins using 884

---

[1]http://pinterest.com/about

[2]http://pinterest.com/popular

machines allocated via PlanetLab[3]. This gave me a "seed set" of approximately 1,000 pins. Next, I wanted to expand this sample beyond "popular pins." I experimented with a variety of strategies, such as snowball sampling (which brings with it certain complications [29]). Eventually I discovered, however, that I could reverse-engineer a pins URL identifier structure to generate an apparently random view of the universe of Pinterest pins. I observed that adding small integers to the identifier in a pins URL would find new pins created slightly later. Via trial and error, I discovered that the trailing five digits of the identifier encoded a pins creation time. Therefore, I generated new pins by adding random integers in the range [1, 10,000] to the seed set of popular pins. I found no correlations among these pins in terms of boards (i.e., categories), pinners, popularity, location or time. Obviously, I cannot make any absolute claims about the randomness of this sample. However, by inspection it looked like a "public timeline" view of Pinterest, and seemed preferable to a snowball sample crawl. Using this method, I collected a total of one million pins and their associated data as noted above. I also collected the pinners for these pins, ending up with a set of 989,355 pinners. The final dataset is uniformly distributed across months of the year and does not show any seasonal bias. Data spans all 12 months of 2009-2011 collected in June 2012. Figure 4 presents an overview of the the data collection and analysis process. Table 1 describes the variables used in this chapter. The distributions of each numerical variable is marked by its mean (red line) and its median (blue line).

Response Variable (Dependent measure). I used the number of repins (i.e., the number of times the pin was shared) as the dependent variable in this chapter. Repins is a measure of the pin's diffusion.

Predictor Variables. I have grouped the predictor variables into two categories: the control features, and the features related to the image's color.

Control features. A pin is composed of an image (or sometimes a video) that links outside Pinterest. Users can upload their images or use Pinterest's bookmarklet on other websites to create a pin. All pins on Pinterest link back to their source and they can be repinned by other users. Pinterest users can organize their pins by topic into self-defined boards.

---

[3]https://www.planet-lab.org

Table 1: Distributions of quantitative variables used in this chapter. I used a non parametric smoothing process, Kernel Density Estimators, to present the probability density functions. Variables marked with '*' are log transformed and x axis represents the log scale. The red and blue lines identify mean and median of the distribution, respectively.

Variable	Distribution
pinner pins*	
followers*	
mean saturation	
mean brightness	
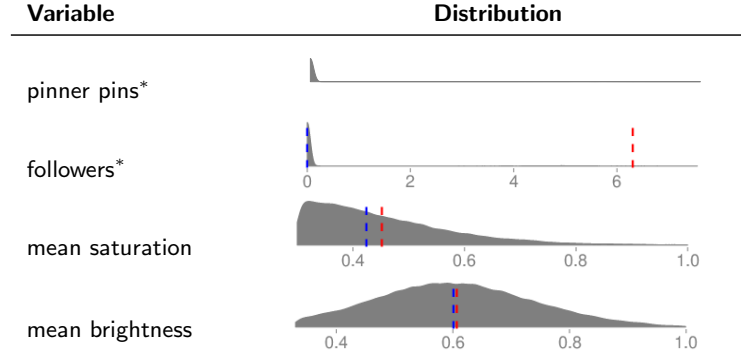
Boards can be set to private or public, and they may attract followers independent of user's other boards.

On Pinterest, the relationship between users mimics Twitter: users form into social networks based on "follow" relationships. When A follows B, B's pins will show up in A's Pinterest feed. In an extension of the Twitter model, the following relationship can be limited to certain boards so users don't get updates from all pins. For the purpose of this research, I was interested in some pinner features that signal activity and social network reach. For activity, I use the *number of pins* on a user's profile. The *number of followers* is the measures of the user's audience size. This is a powerful and intuitive control, as we would expect pinners with larger audiences to have a higher baseline probability of being repinned by their followers. Table 1 lists distributions of these variables.

Color features. Images, however, are the focal point of this study. For every pinned image, my code traversed each pixel. Each pixel was represented by a triple of H (hue), S (saturation or chroma) and V (value or brightness). Using the Munsell color system, I mapped each pixel to one of the ten major hues. (The Munsell system is described in greater detail next). I then find the most dominant color in the image (mode of the distribution) and use it in the model. I also calculate the mean saturation and brightness of the image and use them as predictors for repin model. The saturation and brightness are on scale of 0 to 10 in Munsell Color System (See next section). For simplicity, I convert saturation and
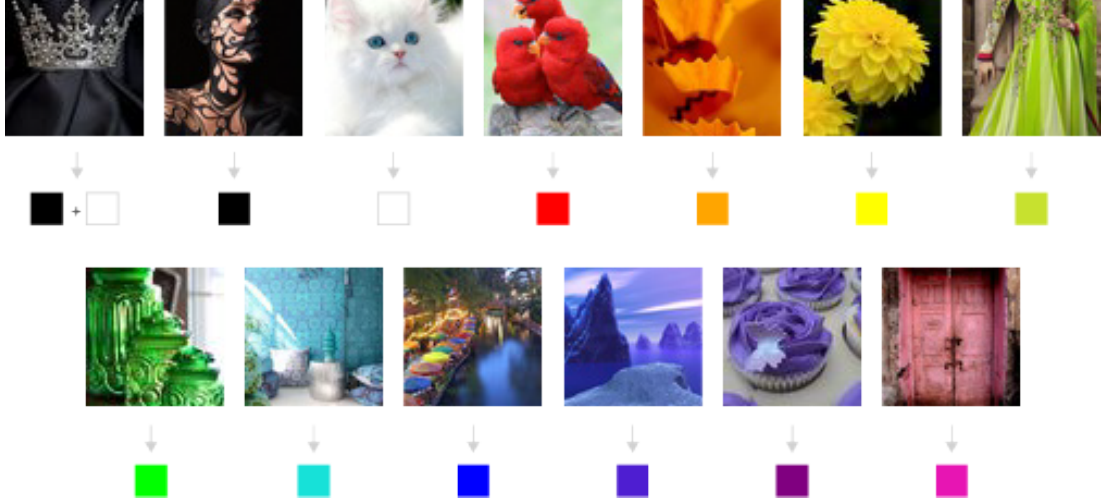
Figure 5: Example Pinterest images belonging to each of the 10 classes of hues and 3 achromatic categorizations derived from the Munsell system.

brightness to the scale of 0 to 1. Table 1 shows the distribution of saturation and brightness for this dataset.

Other than the 10 major hues, I also identified the images which were mostly consisting of white or black. I also used a binary feature which identifies whether the image is only consisted of black and white colors. If the pin's image is black and white this variable is set to 1. Figure 5 illustrates randomly selected Pinterest images from this dataset, along with their algorithmically assigned hue class.

3.1.3   The Munsell System

Color representations allow us to specify or describe colors as a low-dimensional projection—useful for meaningful modeling. We usually identify colors with simple names, but names can also be subjective or culture dependent. In this chapter, I use the Munsell system [149] to characterize colors, a widely used system in applications requiring precise specification of colors, including psychological studies [207]. According to this system, each color has three basic attributes: hue, chroma (saturation) and value (brightness).

Hue is a color's pigment, what we normally understand as blue, red, yellow, etc. There are ten hues in the Munsell color system, five of which are identified as principal hues (i.e.
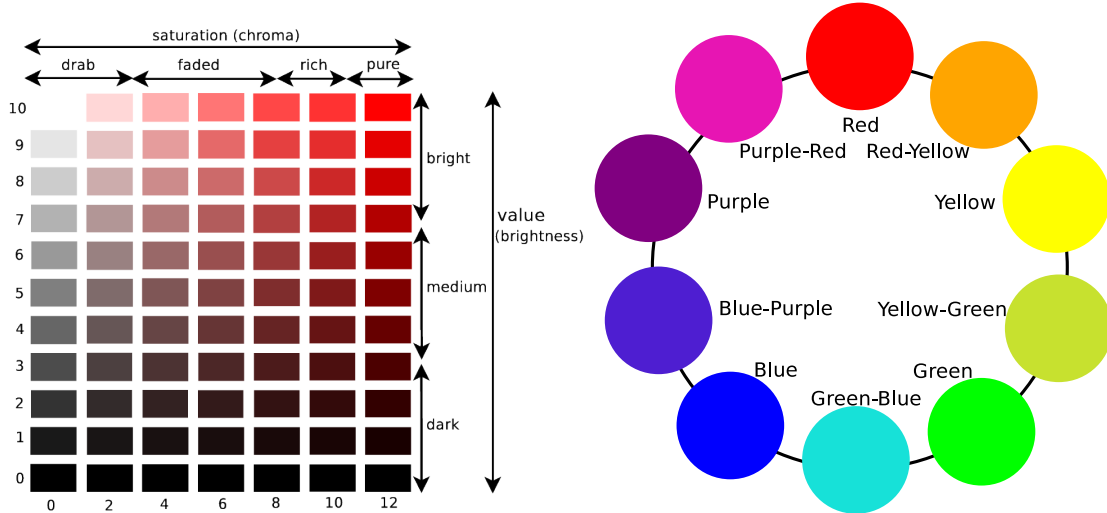
Figure 6: The Munsell color system specifies colors based on hue, chroma (saturation) and brightness (value). I used this system to classify color of images in this chapter. The left image shows the combinations of brightness and chroma for the color red. The right image illustrates how the Munsell system divides the hue space into 10 different colors.

red, yellow, green, blue, and purple). The other five colors are the intermediate hues (red-yellow, yellow-green, green-blue, blue-purple and purple-red). Figure 6 demonstrates the classification of ten different hues.

Chroma refers to saturation, the degree of purity or vividness of the hue. Highly saturated colors have higher proportions of pigment in them and contain less gray. Low chroma colors are drab and dull, while the high chroma colors are rich and pure. Brightness, on the other hand, refers to the degree of darkness or lightness of the color. Low brightness colors look *blackish* while high brightness colors look *whitish*. In the Munsell system, a brightness of 0 is used for pure black, while a brightness of 10 is used for pure white. Black, white and the shades of gray are called neutral (achromatic) colors. Figure 6 illustrates all three dimensions of color. The leftmost image shows saturation versus brightness, while the right demonstrates hue classification.

3.1.4 Statistical Methods

The dependent variable (the number of repins) is a *count* variable. I model the number of repins using *Negative Binomial* regression, on two classes of independent variables: control

atrributes and color attributes. Negative binomial regression is well-suited for *over-dispersed* distributions of *count* dependent variable [43]. I use negative binomial regression instead of Poisson regression since the variance of the dependent variable is larger than the mean. I use over-dispersion to test whether Poisson or negative binomial regression should be used. This test was suggested by Cameron and Trivedi [43], and involves a simple least-squares regression to test the statistical significance of the over-dispersion coefficient. The negative binomial regression models the *expected* number of repins for a pin as a function of control and color independent variables.

The regression coefficients β allow us to understand the effect of an independent variable on the number of repins (note that to be able to compare coefficients, I z-score all numerical variables before performing regression). I use the Chi-squared statistics to find the statistical significance of the regression models, computing the reduction in deviance as compared to a null model.

3.2 Method Validation

As I described in the previous section, I used a pixel-based method to find the most dominant color in the image: the algorithm picks the modal hue class. To make sure that the dominant color found in the image via this method matches what people actually perceive, I performed an evaluation experiment on Mechanical Turk.

I randomly selected a subset of 2,000 images from the dataset and asked Mechanical Turk workers to identify the dominant hue they perceived. I also asked them to mark the image as black-and-white if they recognize it as a black-and-white image. I provided them with the Munsell color chart as shown in figure 6 and asked workers to choose the closest hue in the Munsell system that matches the dominant color in the image. I also gave them the option to choose black and white, for images with those colors as the dominant hue. Figure 7 shows a sample task delivered to Turkers.

All 2,000 images were evaluated by at least 5 Turkers. I recorded the perceived dominant color as the one at least 3 of the participants agreed on. In every single case, at least 3 Turkers agreed on a dominant hue. I then compared the dominant color identified by

33

Task: Identify the dominant color
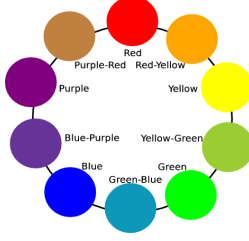
The reference colors are presented on the left. For the image shown below, choose the most dominant color that is closest to one the 10 left hues. If the most dominant color is black or white choose accordingly. If the image is only black and white, select Yes for the second question.

Q: What is the most dominant color in this image?
A: choose between 10 colors above, black or white.

Q: Is this image black and white?
A: choose between Yes or No

Figure 7: A sample mechanical turk task for evaluating the color extraction method.

Turkers with the one the algorithm found. Of the 2,000 in the test set, 1,880 of the images (94%, margin of error 0.55) were in agreement with the algorithm's decision. Table 2 shows the percentage of dominant hues validated in each hue class, with the associated margin of error.

3.3 Results

Table 3 summarizes the β coefficients of the negative binomial regression model for repin counts. I use the Chi-squared Test to find the significance of the regression model, by computing the reduction in deviance from a *null model*. The reduction in deviance is χ^2 = 3.1M - 1.42M, or a 55% drop, on 17 degrees of freedom. The test rejected the null hypothesis, $p < 10^{-15}$; hence, the regression model is well-suited to characterize the effects of the independent variables.

3.3.1 Control Effects

As expected, a pinner follower count has a large positive effect on the virality of a pin, $\beta = 0.82$, $p < 10^{-15}$. This means the higher the number of followers, the more likely it is for the pin to get repinned. On the other hand, the number of pins someone has on his/her boards shows a negative effect on repins, $\beta = -0.76$, $p < 10^{-15}$. The number of

Table 2: Results of mechanical turk evaluation for dominant hue detection. Margin of Error (MOE) is computed for 95% confidence.

Dominant Hue	Sample size	Correct %	MOE
red	246	91%	2.01%
red yellow	145	98%	1.19%
yellow	125	97%	1.6%
yellow green	136	96%	1.81%
green	212	94%	1.65%
green blue	139	94%	2.06%
blue	257	94%	1.52%
blue purple	193	96%	1.37%
purple	118	92%	2.6%
purple red	126	92%	2.57%
black	87	94%	2.61%
white	98	92%	2.97%
black and white	118	91%	2.76%

pins is an indicator of activity on Pinterest. As we can see in the results (table 3) the higher the activity (number of pins) the lower chances of receiving repins. This might also be interpreted another way: the more pins a pinner has, the lower probability any single one has of being repinned. This predictor acts as a drag coefficient in the model.

As expected, we can see that the number of followers strongly influences repins. This is intuitive, since a user with more followers will have the pin routed to more feeds. Having a larger audience increases the likelihood of a repin, a common sense fact realized in this model. Furthermore, the activity level is negatively correlated with repins. The more pins a user posts, the less likely it is that her pins receive repins. As I mentioned earlier, this most likely represents the intuition that the more pins any user makes, the less likely any one of them is to be highly repinned.

3.3.2 Color Effects

All other predictors in the model are related to colors. The goal of this chapter is to quantify the effect of visual features and their comparative effects on diffusion. I used a color variable that represents the dominant color of the image, as described previously.

Table 3: The results of negative binomial regression with number of repins as the dependent variable. The model is significant and reduces 55% of deviance.

type	Variable	β	$Pr(> |z|)$
Controls	pinner followers	0.82	$< 10^{-15}$
	pinner pins	-0.76	$< 10^{-15}$
Color	saturation	0.25	$< 10^{-8}$
	brightness	0.02	$< 10^{-12}$
	red	0.49	$< 10^{-15}$
	red yellow	0.02	$< 10^{-4}$
	yellow	-0.08	$< 10^{-4}$
	yellow green	-0.10	$< 10^{-12}$
	green	-0.13	$< 10^{-8}$
	green blue	-0.01	$< 10^{-4}$
	blue	-0.17	$< 10^{-12}$
	blue purple	0.03	$< 10^{-8}$
	purple	0.31	$< 10^{-12}$
	purple red	0.16	$< 10^{-15}$
	black	-0.12	$< 10^{-11}$
	white	0.04	$< 10^{-15}$
	black and white	-0.12	$< 10^{-8}$
	(Intercept)	0.47	$< 10^{-15}$
	Null deviance	3.1M	
	Residual deviance	1.42M	

We can see from table 3 that the highest effect on diffusion among image features belongs to color red, $\beta = 0.49$, $p < 10^{-15}$. Purple ($\beta = 0.31$, $p < 10^{-12}$) and Purple-Red ($\beta = 0.16$, $p < 10^{-15}$) are also significant contributors to the number of repins. Note that purple-red is usually referred as pink. White ($\beta = 0.04$, $p < 10^{-15}$), Blue-Purple ($\beta = 0.03$, $p < 10^{-8}$), Red-Yellow ($\beta = 0.02$, $p < 10^{-4}$) and Green-Blue ($\beta = -0.01$, $p < 10^{-4}$) have much smaller β coefficients, representing very small effect sizes.

Yellow ($\beta = -0.08$, $p < 10^{-4}$), Yellow Green ($\beta = -0.10$, $p < 10^{-12}$), Green ($\beta = -0.13$, $p < 10^{-8}$), Black ($\beta = -0.12$, $p < 10^{-11}$) and Blue ($\beta = -0.17$, $p < 10^{-12}$), on the other hand, all negatively affect diffusion.

The overall results suggest that while the images with dominant color red, pink and purple increase the chances of getting repins, while the images which are mostly Black,
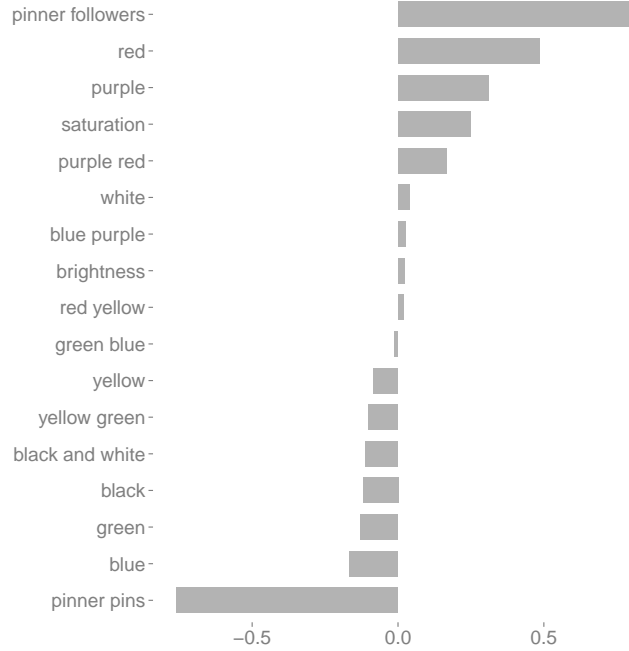
Figure 8: Visualization of β coefficients in the negative binomial model. Results show that after pinner followers, red, purple, saturation and purple-red (pink) have the highest effect on repins.

Yellow, Yellow-Green, Green or Blue decrease the chances of diffusion. Figure 8 visualizes the color coefficients.

I also identified the images that are black-and-white using a binary variable in the model. The model (table 3) shows that when an image is black-and-white, the chance of getting repins decreases ($\beta = -0.12$, $p < 10^{-8}$). Saturation ($\beta = 0.25$, $p < 10^{-8}$) has a large positive effect on repins, suggesting that highly saturated photos are more likely to be repinned on Pinterest. On the other hand, brightness ($\beta = 0.02$, $p < 10^{-12}$) does not seem to impact diffusion.

3.4   Discussion

In this work, I take a first step towards understanding how image content affects diffusion. I have shown the effect of colors, brightness and saturation on the diffusion of pins on Pinterest. There has been extensive research on factors affecting virality in social media, but this is the first result on the diffusion of images. Now, I revisit and contextualize the

results, offering implications for theory and practice, as well as points of departure for future work, both in design and theory.

3.4.1 Images with color diffuse farther than those without

The results of this work show that there is an evident advantage in using color in images when it comes to spreading them. Black-and-white images are not shared as widely as colored images. This finding is in agreement with prior work where researchers found that magazine advertisements shown in color had a larger effect on readers than the black-and-white ads [140, 190]. Because images that contain color are more likely to render the objects in the image more pleasing [33, 51], it follows that these images are more likely to be viewed more favorably (and as a result shared more often) when they appear in color than black and white [160]. Thus, under low processing motivation, users are more likely to be persuaded by images that make use of color, as it is the case with colored ads.

When the processing motivation is high, however, users are expected to get engaged in more effortful processing, a hypothesis confirmed by studies on advertisements [165], allotting a sizable portion of their resource capacity to processing the image. Like less motivated users, the motivated viewers may initially attend to the photo; yet, they should go beyond this by processing the objects and content of the photo. The current study is focused more on the presentation of the photo as collection of pixels and what immediately appears to the viewer's eyes.

3.4.2 Red, purple and pink drive diffusion

The major finding of this chapter is that colors affect diffusion. This effect is large for several colors, including red, purple and pink. Previous psychological studies confirm that red confers [24, 142]. In a recent empirical work, Elliot et al. focused on the influence of the color red on performance in achievement situations [65]. They posited that red carries the meaning of failure in achievement situations and therefore evokes avoidance motivation in such situations. Although red can have inimical implications for psychological functioning, it was shown that it leads men to view women as more attractive as an aphrodisiac for men because it carries the meaning of sex and romance in the context of heterosexual

interaction [66].

Previous empirical work has supported the idea that red has amorous meaning, as studies of color associations have indicated that people tend to connect red to carnal passion, lust, and romantic love [17, 104, 111, 150]. In some of the earliest rituals known to anthropologists, red ochre was used as face and body paint on females to symbolize the emergence of fertility [114, 124]. Red often appears as a symbol of passion, lust, and fertility in ancient mythology and folklore [23, 69, 100, 101, 108].

Mehta and Zhu show that color can influence the task performance with red working better with detailed oriented task [139]. Another study published in the journal nature shows that a similar effect can influence the outcome of physical contests in humans across a range of sports [91]. They found that wearing red is consistently associated with a higher probability of winning. These results indicate not only that sexual selection may have influenced the evolution of human response to colors, but also that the color of sportswear needs to be taken into account to ensure a level playing field in sport.

In animal experiments, some behavioral and morphological changes were reported by Salterelli and Coppola when mice were exposed to pink light (45.7 cm fluorescent lamps (F15T8, 15W, General Electric); 550-700 nm, 620 peak) [176]. Pink light, they reported, increased the weight of the adrenal compared to all other light conditions ($p < 0.05$) when the mice were exposed for 12 hours each day for a total of 30 days. In 1978, Schauss demonstrated the Kinesoid experiment, involving pink color, to a series of classes on innovative treatment techniques in corrections [179].

The findings of this chapter on Pinterest confirm previous findings on colors, showing that images with a dominant color of red, purple and pink have higher chances of diffusion on Pinterest. In other words, images that contain red and purple are more likely to be shared by other users.

3.4.3 Blue, green, yellow and black suppress diffusion

Blue, green, black, yellow green and yellow all negatively contribute to a pin's diffusion. Interestingly, research suggests that both blue and green primarily have associations with

calm and relaxing emotions. Yellow is a cheerful color, and blue is associated with comfort and security [216]. One interpretation is that Pinterest users are more interested in sharing exciting (red-colored images) and elegant (purple colored images) images than cool, cheerful and relaxing ones (blue, green, yellow). We need more work here to discover *why* these colors matter and affect online communities this way. Drawing conclusions about underlying motivations and mechanisms in all observational, statistical work is limited.

3.4.4 High saturation drives diffusion

The findings suggests that a strong positive correlation exists between an image's saturation and it's degree of diffusion. Previous research has shown that highly saturated colors can be more exciting [8, 82, 207], as well as more widely liked [84, 85, 136, 186]. This is consistent with the current findings.

3.5 Implications

This work opens a new line of study: the present findings introduce new factors never before considered influential for diffusion. As I explain more fully next, I believe a deep thread of research can emanate from this work.

The current results connect to psychological studies of color, and they highlight the importance of evocative content. This work echoes the findings of earlier text-based studies: emotional activation is an important underlying diffusion driver within online social networks. In other words, online content that evokes specific emotions are often more viral. In color theory, purple and red are known as colors that elicit feelings of arousal (either in positive or negative direction), and results from this work seem to reinforce this.

For practitioners rather than theorists, the current findings shed light on how to construct viral content. Although I can't claim that every image draped in red or purple will be shared significantly more, on average, warm and exciting colors seem to affect the recipient's likelihood of sharing the image. The current results suggest that using warm, saturated colors can increase chances of diffusion compared to images with cool and relaxing theme.

Moreover, Recently many mobile applications provide users with photo-editing tools.

One of the popular ways of editing photos is to apply filters to them. Filters can change saturation, brightness, and color distribution of the image. The current results can be used to design new filters for images. Filters that increase saturation or enhance the warmness of the image will likely increase engagement with the photo. In other words, a filter that saturates the image and adds red tint will probably be better than a blueish one, in terms of virality.

3.6   Chapter Summary

Many research studies have documented that colors are emotional stimuli. In this chapter, I investigated whether color influences diffusion in large image-sharing communities. I used a dataset of over 1M images from Pinterest. I found that color significantly drives how far an image diffuses, after controlling for network structure and activity. Specifically, red and purple promote diffusion, while blue, green and yellow hinder it. This is the first study incorporating image features into understanding behavior on social media. In addition to speaking to ongoing diffusion research, these findings suggest deep future work using more sophisticated techniques.

CHAPTER IV

WHY WE USE PHOTO FILTERS AND HOW THEY IMPACT
SOCIAL ENGAGEMENT

Mobile phone photography has dramatically risen in popularity recently. For example, the various iPhone models have been Flickr's most popular cameras[1] for years. In its first year, Instagram saw over 150 million photos uploaded. Now, at the time of this writing, Instagram users upload 60 million photos a day[2] on average. Part of the success of mobile camera phone sharing is attributed to the use of on-camera visual effects. These effects, or filters, provide a quick preset path to an artistic rendering of the photo. Mobile photo-sharing sites, such as Instagram and Flickr, provide several filter options; yet, despite their widespread use and the HCI community's interest in mobile photography [10, 21, 113], there is little work—scholarly or otherwise—around filters, their use, and their effect on photo-sharing communities.

Filters are tools that give users the opportunity to enhance their photos. Most photo filters manipulate colors, saturation, light exposure or simulate a change in focus. Filters can age a photo, make colors more vibrant, or give photos a cooler color temperature. Some filters overlay masks in the image or add borders or frames [137]. Overall the goal of filters is to give photos a better exposure or stylized look without knowledge of photo processing. The literature hints to us that digital artifacts like these that are self-made should be more deeply cherished [80, 153, 164, 175].

Vitally important to the sites that play host to filtered images, this chapter presents the results of a mixed-method study exploring the motivations behind filter use, and their impact on engagement. Why do people filter their photos before sharing them? What happens after sharing: Are filtered photos more or less engaging than original ones? Why? While

---

[1]`http://www.flickr.com/cameras/` (accessed 6/2014)
[2]`http://instagram.com/press/` (accessed 6/2014)

existing scholarly work has extensively examined textual content in social networks [102], I know of a very few similar studies on what makes images engaging and interesting [19]. This work is based on interviews with 15 Flickr mobile users and quantitative analysis of 7.6M Flickr photos (some of which were cross posted from Instagram) illuminate the practices of post-processing and how viewers engage with filtered photos.

In this study, I find that both serious and casual photographers like to apply filters on their photos, but have very different motivations in mind. Serious photography hobbyists use filters to improve aesthetics, manipulate colors and highlight certain objects. Although some of these motivations are common among casual photographers as well, the main motive among them is to make their photos look more fun, unique and special by adding artificial effects. Note that this suggests that filters play a social role (i.e., uniqueness) as well as their more obvious visual one. Taking a look at filters from an audience point of view, I find that filtered photos are considerably more engaging than original ones—with filters that increase warmth, exposure and contrast boosting engagement the most. This work explores a new path toward understanding social aspects of photowork. The findings of this work also provide several design implications, such as designing effective filters for both serious and casual use cases.

4.1   Filters and Visual Effects

Recent work on creativity found that digital artifacts that are special are often self-made, such as presentations, animations and photo montages [164]. These results have shown that crafting and making with digital media can make these media more special or cherished, and in fact, being self-made or augmented appears to be one of the main reasons people cherish their digital possessions [80, 153, 175].

Image editing, also known as photo retouching, is the process of altering images. Prior to widespread use of mobile camera phones, filters were mostly used by photographers to correct problems, apply effects or simply to make a photo look different. Even before the age of digital photography, people used to make changes in the dark rooms. They could adjust the contrast, use different papers and different tones on those papers to make their

Figure 9: Example of an original photo (top left photo) and many filtered variations. Some filters change contrast, brightness, saturation; some add warmth or cool colors or change the borders.

photos look better [182].

Digital images are stored in the form of picture elements or pixels. Process of editing can change the pixels to enhance the image in many ways. The pixels can be changed in a group or individually. Camera or computer image editing programs often offer basic automatic image enhancement features that correct color hue and brightness imbalances as well as other image editing features, such as red eye removal, sharpness adjustments, zoom features and automatic cropping. These tools are complex and more catered to experts rather than end users. Filters, on the other hand, are tools that give users the opportunity to enhance their photos, without the need to go through professional software. Most photo filters manipulate colors, saturation, light exposure or simulate a change in focus. There are different use cases introduced by each filter. Filters can age a photo, make colors more vibrant, or give photos a cooler color temperature. Some filters overlay masks in the image or add borders or frames [137]. Figure 9 shows some of the filter effects on a sample photo. Instagram and Flickr are among the most popular mobile photo-sharing platforms that provide users with a variety of filter choices.

The current work is the first to look at filtering and visual post processing practices on photo sharing communities, it is also the only study looking at the impacts of these visual effects in engaging users.

4.2 Method

I use a mixed-method approach to understand users motivations behind filter use and look for the impact of filters on those who view the photos. My approach is two-fold: first, I take a qualitative approach and interview 15 users of Flickr mobile to gain deeper insight into motivations behind using filters. Then, I focus on viewers impression of filtered photos by looking at quantitative data obtained from Flickr.

4.2.1 Qualitative Approach: Participants

To recruit users for the study, I emailed a pre-screening survey to a database of potential study participants. I screened for Flickr users who used the mobile app at least a few times a month. This assured that they were familiar with the app and the filter options through

the mobile interface. Nine of all participants mentioned that they use the web interface, though interviews didn't focus on this usage. Those who matched the recruitment criteria were scheduled for a 60-minute interview session.

I recruited 15 participants. Of the 15 participants, 10 were interviewed in person and 5 were interviewed remotely through telephone and Skype. In all the sessions, participants were video-recorded. All participants were located in the United States. Those interviewed in person were located in the San Francisco Bay Area. The sample consisted of 5 women and 10 men ranging between 29 and 53 years old, with a median of 35. During the interview sessions, I mostly asked about participants use of camera phone, Flickr app and other photo sharing apps such as Instagram. I then asked participants to choose random photos from their photostream and explain why they filtered it or left it as original. Additionally I asked them to apply different filters on their choice of unfiltered photos and explain which filters work better. At times, I asked them to go back to their photos and describe how filtering changed the ways the photo looked. Other than the mobile phone camera, 8 of the participants stated that they own digital SLR cameras. The other 7 used their phone as their primary photo taking device.

4.2.2 Quantitative approach: Flickr Mobile Data

To understand the role of filters on engaging users, I perform a large scale quantitative study on Flickr. Along with the qualitative study, I also collected public photo metadata from Flickr; these photos were identified by the API as having been uploaded from Flickr's or Instagram's mobile app. In total, the dataset consists of over 7.6 million photos, approximately 4.1 million posted from Instagram and 3.5 million from Flickr's mobile app, uploaded between early 2012 and mid 2013. I identified whether they were posted as original or filtered by checking their machine tags, auto-generated tags from the uploading application. For this study, I wish to predict engagement using two dependent variables: implicit usage (viewing a photo) and explicit actions (commenting on a photo) [221].

Implicit dependent variable. I use the number of *views* of each photo as an implicit measure of behavior. It quantifies the number of distinct users who viewed the photo.

Explicit dependent variable. *Comments* are explicit forms of actions taken on each photo. The number of comments quantifies the number of distinct comments users posted on the photo.

Control features. Every user on Flickr has a *photostream* which can be viewed by other users, depending on the privacy settings on a per photo basis. Tags on Flickr are used by the search index to help people find photos. I use the photostream views and tags as controls for finding photos, either through user's stream or direct search.

Like Twitter, the Flickr relationship model between people is asymmetric: users form into social networks based on "follow" relationships. The number of followers is the measure of the user's audience size. This is a powerful and intuitive control, as we would expect users with larger audiences to have higher baseline probability of being viewed or commented on by their followers. I also use the *account's age* measured in months as a measure of seniority of the user on the site. The longer a user is on Flickr, the larger this value.

Filter feature. Filters are the focal point of this study. For every image, I identified whether it was shared as original or was filtered before shared by checking their machine tags. I coded *is filtered* as a factor variable, with a value of 1 for filtered photo and 0 otherwise. Figure 9 shows an original photo and examples of variety of filters applied to it.

Further, I reverse engineered each filter by comparing the red, green, blue, and luminosity histogram channels of the uploaded no-filter image with a filtered one. This allows us to describe a filter as a change in saturation, contrast, color temperature, and/or exposure. The filters which added an aging effect, through the introduction of dust, scratches, and noise, were visually identified as well. I use these features as descriptors based on the theoretical frameworks around colors, saturation and lighting effects [82, 84, 85, 89].

4.3   Results

I find that mobile phone users utilize their phones in various ways, one of them is to take photos and capture the moments during their daily activities. All participants explicitly mentioned that their camera phones are specially useful in daily events and when they do not have access to any other cameras. Among the users of the mobile app, I find two main

Table 4: Summary of users motives in applying filters on their photos. I provided an example for each motive and the number of participants who mentioned that motive.

Motive	Example	# Participants
Improving Aesthetics	make the clouds look distinct from the sky	13
Adding artificial Effects	give an old look to an old theater photo	12
Highlighting certain objects	focus the attention on the face	9
Manipulating colors	improve the colors in the image	8
Making photos appear more fun and unique	emulate film by removing colors from a portrait of an old man	5

Figure 10: Example photo where the user applied a filter (Toucan filter from Flickr mobile app) to make the clouds stand out. (Photo: ©①◎ *skelington* `https://flic.kr/p/eZ2BA9`.)

groups: (1) Those who consider photography as a serious hobby. These users usually own a professional or semi-professional Digital SLR camera that they use aside from their mobile devices. They use their mobile phone cameras for unplanned events, whenever carrying a camera is intrusive or when they do not need high resolution. (2) Those who take casual photos mainly for the purpose of documenting objects, events and people and sharing them with their family and friends. I will discuss the differences in filter use between these two groups later in this section.

4.3.1 Why Use Filters?

I find a variety of reactions toward filters among participants; six of them almost always filter their photos, six of them filter their photos occasionally (depending on the photo) and three participants are familiar with filters but do not use them. Table 4 summarizes the

reasons participants mentioned as their motivations for filter use.

Improving aesthetics.

One of the main motivations for filter use is to enhance a photo and correct for brightness, saturation, contrast and focus. This reason is specially popular among the serious photography hobbyist because mobile phones do not give many options to control for such factors inside the camera. The user can upload the photo through the app and use filters to apply certain effects; for example increasing the contrast so the clouds can be more visible in the sky (See figure 10):

> "Sometimes you want to, where you'd take a picture of clouds, to show the clouds, then you have to somehow enhance those little differences between the sky and cloud, so you would enhance the contrast, then I look for those filters."
> (P6)

The filters that are used for enhancement are usually milder in the effect intensity and are applied to enhance the photo while keeping the main imagery or the subject with minimal alterations. Some of the participants who were interested in applying filters to their photos for the purpose of enhancing their photos shared the concern that having strong filters might devalue the main image.

Adding artificial Effects.

Sometimes users use filters to give their photos a look and feel that is relevant but non-existent. For example a few of participants mentioned that they like to make their photos black and white whenever the existence of the color is not necessary to the aesthetics of the photo. One of the popular subjects of black and white filters are photos that want to bring out attention to a certain texture rather distract viewers with colors (see figure 12). Another reason mentioned by participants for applying black and white filters is to give an older look to the image or emulate film:

> "It's my favorite bar in San Jose, it's called Singlebarrel and it's kind of an old speakeasy theme, like in the 20s. I just thought that that [black and white] fit the theme of bar better." (P3)

49

Aged look is another popular feature introduced by the filter apps. The sentiments toward using this feature is highly variant across participants. Some participants stated that they love using the sepia tones and age look effects on their photos, while some other (more of serious hobbyist group) found it too artificial:

> "I like some filters for the effects, like either making it starkly black and white, or giving it an old look, like the Antique look, or the Throwback look. I also like the vignetting or tilting, I like those features, in the set of filters." (P2)

Highlighting certain objects.

Sometimes filters can help bringing out the focus to a certain object in the photo. This feature in filters is very popular, specially for photos of people where the focus is mostly the person and less the surroundings. Sometimes the photographer wants to bring out certain aspects of the landscape that in the original photo might not be recognizable, or remove certain distractions so that the feature stands out. Figure 12 is an example photo in which one of the participants applied black and white filter to emphasize the roughness of the landscape:

> "Its a lake that had no water in the winter. This is now covered with water and this [the hill area] was all brown and this [trees] was green. I think it had more to say with the roughness of the landscape. I didn't want to actually show the soil and the trees. I wanted to show the roughness of the landscape, the reflection of the water, the fog and this." (P6)

Another example where the participant mentioned how filters helped him to focus the attention of the viewer on the main subject of the photo:

> "I have a shot of my children from the back and they are looking at something, so I use the filter to give it kind of a very subtle kind of border that focuses their attention a little bit more on the foreground and less on the background." (P12)

Figure 11: Participant describes how the black and white filter helped him show the roughness of the landscape (Photo used by permission of the participant).

Manipulating colors.

Many of the current filters on Flickr and Instagram apps manipulate the colors by adjusting saturation, brightness, contrast or simply by changing the color gradients into warmer or cooler colors. Some participants mentioned that one of the main uses of filters is to either emphasize on certain colors or reduce the diversity of colors in a photo:

"Yeah. I tend to filter just about almost everything, depending on what Im trying to draw attention to in the picture. Say, I'm trying to bring out the color of something, I'll use a filter that does that, or if I wanted to make it look like aged or something, I'll do something that way." (P3)

Some of the participants also mentioned how they use filters to adjust for the right levels of saturation and brightness in colors:

"I'd just try to play around with the saturation to get all the bright colors. I like to do this once in a while. I just go through all the different filters and see which one looks better." (P11)

Making the photos appear more fun and unique.

Figure 12: Participant describes how the filter surprised her with a new look for the photo that she did not imagine while taking the photo (Photo used by permission of the participant).

Other than primarily photographic motives, sometimes filters help users give their photos a fun and unique look that they could not capture through the camera. This motive is a more social motive which is also more popular among the casual photographers. The casual photographers who do not have much knowledge of photography as an art, described the filters primarily as tools that make the photos more special:

> "They make some pictures more fun and more interesting and more unique. That being said, because most of the pictures that I'm sharing are just get across pictures of my baby to show my parents or whatever. I don't feel like they're the target audience. They wouldnt get it, so...that's why I'm not using them as much." (P5)

Sometimes these fun and unique looks are not imagined by the photographer or not intended until when the filter is applied:

> "Sometimes, it's hard to imagine what I can do with the filters until I get to it. Even when I'm taking the picture. For instance, this picture was taken at

Central Park. When I clicked it, I just clicked it because it was looking good that day. I came back and I tried that, displaying it within filters, and I realized that this filter looks good. It gives it a particular look that I could not have even thought of before I applied it." (P2)

4.3.2 Serious Hobbyist vs. Casual Photographer

As I briefly mentioned earlier, there are differences in filter use among serious hobbyist photographers and the casual photographers. While serious hobbyists like to use filters as enhancement methods, casual photographers take advantage of filters to make their photos more special and fun. On the other hand serious hobbyists are interested in subtle changes where the value is mostly in the imagery itself while mild filters made some of the details more visible. They express their dislike in Instagram like filters where the image is completely changes by filters:

"I don't want the treatment of the image to detract from what's happening in the photograph. A lot of these apps, they just pile stuff on top of stuff on top of stuff, so they have scratchy lens, scratchy film, vignetted, soft on the edges, hyper saturated, super desaturated, super high contrast. Basically, pardon my French, they're taking a really shitty photograph, and they're putting so much stuff on top of it that it doesn't really matter anymore. You don't even see the image." (P10)

Some of them also expressed concern in how filters might devalue the art in the image by making it easy for casual photographers to create beautiful photos:

"My 10-year-old cousin, he takes the app. He takes the photo. He passes it through filters and it's beautiful. You feel great and you feel a bit sad. Sad because the actual art in it is lost into the filter" (P6)

For casual photographers, the act of sharing is the primary intention when taking mobile photos. Participants in this study confirmed that when they take photos with their mobile phones they usually think of their audience interest on social media or Flickr.

4.3.3 When is Original Better than Filtered?

There are a lot of cases when users prefer to share the photo without filtering it. A few of the users mentioned that the practice of applying takes times and effort and so they do not always want to put the effort into filtering photos. These set of users only filter those photos that they think are worth the effort. Photos that have special people or subjects in them are among those. Other times the original photo is of good quality and the photographer does not want to change the content of the photo with filters because the details captured in the photo might be lost.

> "This is a photo taken from above a little cove and it's a pretty fast shot at waves breaking. There's enough going on in the photo where it's not a static landscape or anythg. The movement was interesting and it isn't too washed out color-wise. I felt it was enough was good about it that I didn't need to start messing with it." (P1)

On the other hand some participants, specially those with lower tendency of using filters, mentioned that some of the filter effects are too bold and too much to the extent that it detracts from the image itself.

> "In some ways when I think of Instagram, I think about it as very filtered pictures. The sort of pictures I'm taking, because I would like to be a bit artsy, I'm tending to find a clean shot of an object or an angle is a little bit more my language for expression. I'm intrigued by the filters. If I can add that in a way that makes sense in due course, then I'm definitely interested in." (P15)

Another reason that users might not use filters on their photos is the subject of the photo. Sometimes the subject is important to be captured in its reality and without alterations. These subjects usually have memory values so that the user wants to remember the subject in its original way:

> "Like if I'm taking a picture of my frog, I want like that actual frog. I don't want a filter. Like usually if it's like something natural or organic, I don't want

a filter applied...For example, this was a private art sale. For this, I didn't use any filter because I always wanted to remember that painting the way it is in its natural state with good light in Carmel, like a painting. Something like a painting because you really want to capture the painting as it as intended to be viewed." (P4)

Sometimes the photo is taken to document a certain event or people and so alteration would spoil the original purpose:

"Generally, if it's in a group shot, it's about capturing the memory of that moment, I wouldn't apply filters on it. I would try to keep it as original as possible. It's not for the art. It's for capturing the moment."(P2)

To summarize, I found that two main groups of mobile photography users, the serious hobbyist and the casual photographer, use filters to improve aesthetics of their photos, manipulate the colors or highlight certain objects. I also find that although more common among casual photographers, mobile app users like to apply filters to transform their photos into fun and unique looks or add artificial effects. Next, I describe how filters are perceived by the large scale audience on photo sharing communities like Flickr. Are they engaging the viewer more or less than the original camera photos?

4.4 Filtered Photos and Engagement

In the previous sections I asked people about their own intentions and motivations in filter use. To broaden this view, I next consider how viewers engage with filtered photos. Particularly, I ask what effect their choices to use filters have on people who view those photos. While the photo creator's motivations are essential to understanding filter work, it is also important to understand how the general audience engages with such content.

For this purpose, I conduct a large scale study on photos shared by Instagram and Flickr apps on Flickr and evaluate the role of filters in engaging users. A large quantitative study here helps us evaluate the role of filters on the general Flickr users. I use a

Table 5: Results of negative binomial regression with number of views and number of comments as dependent variables. For all coefficients Std.Err $< 10^{-2}$ and $p < 10^{-15}$. The results show that a filtered photo is 21% more likely to receive views and 45% more likely to receive comments. Other variables are used as controls.

Variable	β_{views}	$\beta_{comments}$
is filtered	0.19	0.37
followers	0.58	1.21
photostream views	0.59	0.46
tags	0.28	0.003
account's age	0.04	-0.18
(Intercept)	2.26	-2.77

dataset consisting of 7.6 million photos meta-data collected from Flickr. I use two regression models to study the effect of filters on engagement. Each model takes the number of views and comments as dependent variables, considering control and filter features as predictors. Specifically, I use *negative binomial* estimators, because views and comments are highly skewed count measurements that cannot include negative values. The regression coefficients β, table 5, allow us to understand the effect of an independent variable on the number of views/comments. To be able to compare coefficients, I z-score all numerical variables before performing regression.

Both engagement models show significance, $p < 10^{-15}$, for all predictors. I use the Chi-squared Test to find the significance of the regression model, by computing the reduction in deviance from a *null model*. For this model for the views, I found the reduction in deviance $\chi^2 = 7.4M - 5.9M$, or a 19% drop, on 5 degrees of freedom. For the comments model, I find reduction of $\chi^2 = 1.3M - 0.94M$, or 28% drop, on 5 degrees of freedom. The test rejected the null hypothesis for both models, $p < 10^{-15}$; hence, the regression models are well-suited to characterize the effects of the independent variables.

4.4.1 Photo Views

As expected, the views model shows that the largest coefficient belongs to photostream views: as more users view a photostream, it is more likely for a photo in that photostream to be viewed as well. It is also expected that the number of followers is another large

contributor to the number of views. The follower/following relationship on Flickr allows followers of a user to see all public photo updates of that user in their feed. The larger the number of followers, the greater the audience of the shared photo and hence it is more likely for the photo to be viewed.

Tags on Flickr have an important role in finding photos as they are highly used for search. From the views model, I conclude that the more tags a photo is associated with, the higher the number of views. Overall, all three ways of searching that leads to viewing a photo, either through photostream views, followers list or tags list are common ways to increase the number of views. The current results also show that the account's age has a positive but small role in the change in views.

The main objective is to explore the impact of filters on photo engagement. With regards to view count, the current results show that filters are indeed strongly positively correlated with the number of views. Existence of a filter can increase the chances of photo being viewed by other users by 21%. I calculate the percentage by replacing the coefficients in the negative binomial equation.

4.4.2 Photo Comments

In the comments model, the photostream views has a large positive coefficient but ranks second after the number of followers among predictors. This shows that the number of followers is far more effective in explaining the comments than the photostream views. We also observe from the model that the effect of tags ($\beta = 0.003$) is negligible, giving most of the credit to the number of followers. Given that commenting is a more social type of engagement compared to views, it seems intuitive that the followers contributes the most to it's variance. On the other hand, the current results from the comments model show that the account's age effect is negative ($\beta = -0.18$). In effect, the long term site users are less likely to receive comments on their mobile posts than the new users. This is surprising and suggests future work to further explain these observations. There is a strong positive relationship ($\beta = 0.37$) between comments and filtered photos. This means filtered photos receive 45% more comments than the original ones.

Table 6: Results of negative binomial regression with number of views and number of comments as dependent variable. The filter features are used as independent variables. For all coefficients Std.Err < 0.02 and $p < 10^{-10}$. The results show that warm temperature, higher contrast, and higher exposure increase chances of receiving views and comments.

Variable	β_{views}	$\beta_{comments}$
warm temperature	0.06	0.23
increase saturation	-0.02	0.07
increase contrast	0.08	0.13
increase exposure	0.10	0.16
age effect	0.06	-0.08
(Intercept)	2.04	-2.69

4.5 Which Filters Impact Engagement?

So far I find that why people use filters and how it impacts the photos engagement. Now, I ask: What makes filters engaging? Are all filters equally engaging? What photo transformations increase the likelihood of being viewed and commented on? In this section, we investigate the properties of filters that make them engaging.

I extract certain transformations, commonly used in photo filters, and use them as predictors in the models. The features identify whether the filter has a warming effect, aging effect or adds saturation, contrast or exposure. I use views and comments as proxies for engagement and I construct *negative binomial regression* models. Both models rejected the null hypothesis of Chi-Squared test, with $p < 10^{-15}$. Table 6 summarizes the β coefficients for both models.

The results show that an increase in contrast and exposure positively affects the number of views and comments. This confirms some of the findings in the motivation section where participants mentioned that they use filters to bring out colors or concentrate on certain objects in the photo. Many of the participants in the qualitative study mentioned that they prefer filters that bring out contrast more often:

> "Graphite and Noir are black and white and they bring out the contrast more than the other pictures. I like high contrast pictures." (P2)

Filters with warm temperature significantly increase number of comments and their

effect on number of views is also positive. The aging effect seem to increase the views but decrease the number of comments. The participants also showed interest in some of the filters that add aging effects and introduce warm colors:

> "I like Sepia looks. Sepia looks with high contrast. Sepia looks usually have low contrast, whereas I like Sepia looks with high contrast on similar image like this. For that, I have to play around with, to get that look." (P2)

I also find that effect of saturation on views is small and negative while on comments, it has a positive impact. Figure 13 shows examples of some filters that are engaging and some that are less-engaging.

Photographically speaking, filters which auto-enhance a photo (e.g. correct for contrast and exposure) drive more engagement. I find the less-engaging filters exhibit transformation effects which are exaggerated and often cause photographic artifacts and/or loss of highlight details. The exception being filters which make a photo look antique.

4.6 Discussion

Through conducting interviews with 15 Flickr mobile users I find that most users like to apply filters on their photos. I find two groups of Flickr mobile users; the first group are serious photography hobbyists who own professional cameras aside from their mobile devices. This group of users use their mobile cameras for convenience and availability. Those of them who apply filters on their photos use the mild effects mostly to correct errors, enhance the photo or bring out colors or objects. Those who do not use filters as much are mostly concerned about the artifacts and artificial look of filters on their photos.

The second group of Flickr mobile mainly rely on their phones for taking photos. They are casual photographers who enjoy recording moments or events in their daily lives. The casual photographers like to give cool and unique look to their photos by applying filters. Filters help them make their photos appear cooler and sometimes more fun. Some of them do not use filters due to them taking more effort and time. I find that sharing is main purpose of filtering for casual photographers and as we saw in the quantitative analysis part of the chapter the filtered photos are more engaging.

(a)

(b)

(c)

Figure 13: Examples of (a) a raw image and (b) an engaging filter that adds contrast and warmth, and (c) a less-engaging filter, which introduces artifacts and adds a cooler temperature.

I find that filtered photos attract more implicit usage as well as explicit action. Filtered photos are 21% more likely to be viewed and 45% more likely to receive comments, compared to non-filtered ones. Specifically, I find that filters that impose warm color temperature, boost contrast and increase exposure, are more likely to be noticed. Other filters showed no significant effects. See figure 13. I control for several features that might affect photos engagement. I find that the number of followers strongly influences both views and comments. This is intuitive, as a higher follower count increases photo distribution. Similarly, we also see photostream views are positively correlated with photo views and comments. The number of tags is also a driving factor for views but not for comments. The relative importance of followers compared to photostream views is stronger in comments than views, suggesting stronger social dependency in comments.

These results highlight the motivations behind using filters and how it might impact the perceptions of the viewers. I have shown that prevalence of mobile photography has made the visual processing more popular to the level that many casual photographers enjoy post processing their photos. There is some concern that pre-made filters reduce the value of the art presented in the imagery; however, I also find that the serious photographers are able to distinguish between photos filtered in an artsy way compared to the ones filtered for artificial cool looks. This work is the first to deeply dive into understanding filters.

On the other hand the current findings show that filters might directly impact the level of engagement on photos. The results also connect to psychological studies of color [84, 85, 89], and they emphasize the importance of emotionally evocative visual content. This work echoes the findings of earlier text-based studies: emotional activation is an important underlying driver of engagement within online social networks [102]. In color theory, warm colors such as red and yellow are known to elicit feelings of arousal and cheerfulness, and my results seem to echo this [214]. The current findings shed light on how to construct engaging content; filters influence the engagement on Flickr. Although I do not claim that every image filtered will be viewed significantly more, on average filtered image seem to affect an observer's likelihood of engagement.

4.6.1 Design Implications

The current findings provide several design implications for mobile photo-sharing communities such as Flickr and Instagram. I showed that the mobile photographers belong to the two groups of serious hobbyist and the casual photographers. While there are many similarities in their use of mobile phones for photography, many of their editing needs and aesthetics values differ. Considering these two groups of users and their expectations from the filters can help designing for better user experience. For example, the app can offer tuning options on each filter. As I mentioned earlier, many of the serious photography hobbyists prefer their effects mild and filters to be less noticeable. On the other hand the casual photographers like the artifacts that make their photos visibly different and unique. Allowing users to tune the changes of a filter on their photos will help both groups:

> "One of the things I like to do is to dial down the amount of the the filters that's been applied which let's me cheat so I know that the filter's been there, but it doesn't look so much like there's been a filter there." (P1)

The current findings can also be used to improve filter construction. While filters seem to be related to content usage and social engagement, not every filter works equally in driving views and comments. The filters which increase the saturation, for example, do not drive engagement as much as warm temperature, high exposure, high contrast filters. Designers can use the findings of this work to build photo feeds that takes advantage of photos with such filters, or include the findings of this chapter in algorithms that decide what is trending or popular, or design filters with engagement in mind.

Additionally, 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 views and comments.

4.6.2 Theoretical Implications

Above, I explore design-oriented research questions that may help shed light on the findings. I believe the present research also suggests new directions for Computer Mediated

Communication (CMC) theory. Researchers have found that a creative touch on digital artifacts such as photo montages makes them more special and cherished [164]. In this work I show in a large scale that filtered photos are more viewed and commented on. Could this be due to aesthetics enhancements of filters on the photos or due to value added by personal creative touch? We like the self-made or augmented artifacts more, but could this also affect the viewers on social media? Researchers might leverage this work to investigate the role of creativity and personal touch on large scale viewer's engagement.

This work is a first step opening a larger set of research directions and areas of investigation. While I find how people use filters and that filtered photos are more likely to engage users, I can't say how much of this effect comes from the content of the photo. For example, is it that users filter their engaging content before sharing, or that filters increase the engagement of the photo? Regardless of the cause, the current findings invite deeper analysis of content of the photo and its impact.

We see in this work that users like to apply filters on their photos even though it is a time-consuming process and requires spending more effort. This suggests that post-processing as a tool to enhance photos can motivate creations. Many of participants in this study mentioned that the changing their photos through filters makes those photos more special and fun. We also see that the viewers tend to engage with those photos more than the original snapshots. This opens a new question for researchers of CMC: Are filtered photos more engaging and fun because of the filters or simply because they are result of personal creations?

Future work can also look at other visual characteristics of multimedia and study their impact on online behavior. For example, computer vision techniques, from visual features to scene detection, can be used to further design of such applications and improve understanding of photo engagement. How does photo content change other explicit social behaviors, such as likes/favorites or sharing to other communities? Can we suggest more appropriate filters based on the color composition of photos? For example recommending filters that enhance brightness when the photo is too dark.

It is also worth considering how filters are used in different subjects. For example, are

people using highly saturated filters on photos of food, while using aging filters on street photography subjects? Another direction for future work is to examine the intensity of filters on how they engage users. Are milder effects more likely to be engaging than the bold filters?

4.7 Chapter Summary

Filters are becoming increasingly popular among users of mobile photo sharing tools and sites. In this work, I take a first step towards understanding motivations behind filter use and their impact on user behavior. My contributions are two-fold. First I perform a qualitative study to understand motivations of users in applying filters on their photos. I find that both serious and casual photographers use mobile app filters on their photos. The serious hobbyists apply filters to enhance and correct their photos, expose certain objects or manipulate the colors. More so casual photographers like to add artificial effects to their photos and make them more playful and unique. My second contribution is an empirical study on 7.6 million mobile uploaded photos to analyze the effect of filters on users' engagement. I find that filtered photos are more likely to be viewed and commented on. This work has several implications both for theory and design of technology.

CHAPTER V

HUMAN FACES IN THE PHOTOS AND HOW THEY IMPACT
SOCIAL ENGAGEMENT

Even as babies, humans love to look at faces; infants, barely minutes old, turn toward faces, sensing that they are important [148]. It is widely accepted in neuroscience that face perception is perhaps the most highly developed human visual skill [88]. Faces are also powerful channels of nonverbal communication [197]. 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 [79].

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" [112].

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 site[1]. 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 [16, 25, 40, 41, 105, 144], there are few studies on what makes visual content socially engaging online. To investigate this, I ask the following research questions in this chapter, driven by social psychology work on face perception:

---

[1]`http://instagram.com/press` (Accessed 9/2013)

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?

I adopt Instagram as the site of study in order to answer these 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 image [2]. Using a corpus of 1 million images from the community, I 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, I find that the number of faces in the photo, their age and their gender do not impact engagement.

This 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, 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.

5.1 Instagram

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 [59]. 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;

---

[2] I use the terms image and photo interchangeably throughout this chapter to refer to the images on Instagram.

Original photo **Face++ API results** **Our constructed variables**

male, between 18 and 35 years old
male, <18 years old   female, between 18 and 35 years old

female, <18 years old

has face: YES
has female: YES
has male: YES
has face <18 years old: YES
has face between 18 and 35 years old: YES
has face >35 years old: NO
number of faces: 4
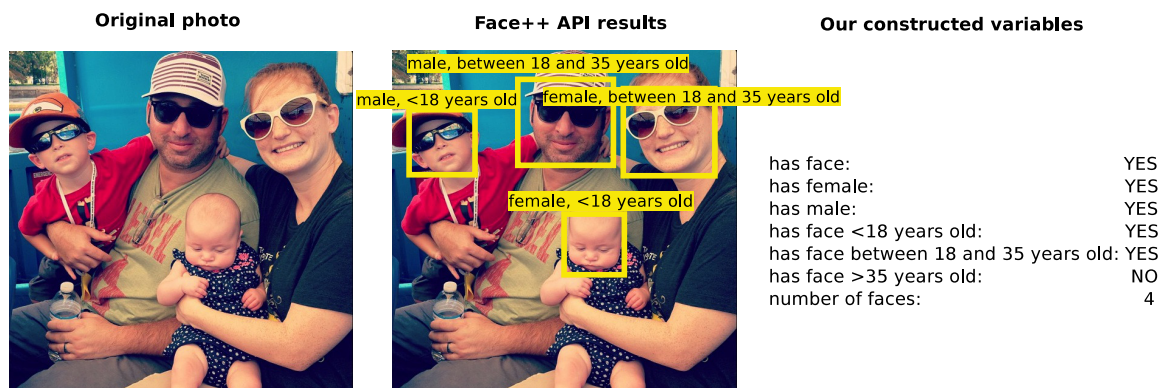
Figure 14: Example Face++ face detection and how I construct the variables. The photo used in this example is a photo under Creative Commons license from Flickr.

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. [93] 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.

5.2   Method

I take a quantitative approach in this chapter to investigate the relationship between faces and engagement. While engagement can be quantified in various ways, I 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, I describe the data I collected from Instagram and how I detected faces and their age and gender; followed by clarifying the statistical methods and analysis process.

data collection	evaluation	analysis
Using Instagram API, collected 2,000 popular images.	Randomly sampled 2,000 images for method evaluation.	For all 1M images, created binary features for face, age, gender.
Via snowball sampling, collected 23M more images.	For each image, 5 Tukers judged algorithm's face & age results.	Negative binomial model assessed effect on engagement.
Randomly sampled to 1M.	Face, age, gender results validated.	
Each image augmented with profile data & analyzed for faces.		

Figure 15: An overview of the steps taken in this chapter for collection, evaluation and analysis of the data.

5.2.1 Face Detection

Face detection and recognition from images or video is a popular topic in vision research and it has received lots of attention [211, 219]. 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 [107], I did not have access to an implementation that can work for large scale image analysis. I therefore used a publicly available face detection API developed by Face++ [7]. I only use the detection modules, as the goal of this chapter is to find relationship between existence of faces and the social engagement. Face++ provides a set of compact, powerful, and cross-platform vision services, which enabled us to use cutting-edge vision techniques. The API does not provide an estimation of accuracy, so I turn to a crowd-sourced validation method to confirm the accuracy of the face detector described later in the validation section.

Table 7: Distributions of quantitative and binary variables used in this chapter. 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 the dependent measures. Audience and activity variables are used as controls, and faces variables are the focal point of this study.

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.	

Face++ provides 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. I then reduce the dimensionality of data by converting the results into a binary space, where I mark only when there is a face in an image. I also identify whether any of the faces in the image belong to certain age ranges. The three age ranges that I consider in this chapter 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, I construct another binary feature which determines whether at least one female or one male face is in the image. Figure 14 shows an example Face++ detection and how I construct the variables.

5.2.2   Data

My goal is to obtain a random sample of photos from Instagram. Even though Instagram provides a publicly available API, gathering a random subset of photos is a challenging task. One can either search for photos by location or query on the list of most recent popular photos. I opted to start with a set of 2,000 popular Instagram photos, collected on November 2012. I then used snowball sampling [81] to collect the users and their followers as well as a random set of their photos. This dataset consists of 23 million Instagram photos and over 3 million Instagram users. To soften biases due to snowball sampling, I 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 15 shows a detailed flowchart of data collection, evaluation and analysis processes.

Response variables (dependent measures)

The goal of this chapter is to investigate the role of photos in predicting user engagement on Instagram. I choose 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 7.

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 [221] 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 chapter, I 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 I 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.

Since faces are found to be effective stimuli [38, 103] in attracting people's attention. I use a binary variable as our predictor to account for presence of a face in the photo.

Has face. For each Instagram photo, I 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, I consider variables identifying age and gender of them. The age and gender variables are derived using face detection method.

Has children and teens- has face < 18 years old. I 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. The 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.

Look at the following image and find the clearly visible faces
Provide the number of faces present in the image
Identify gender of each face
Provide the number of females and males
Identify age group of each face
Provide the numer of faces in each age group

Number of faces in the image: 2
Number of females in the image: 1
Number of males in the image: 1
Number of faces with estimated age under 18:   1
Number of faces with estimated age between 18 and 35: 0
Number of faces with estimated age over 35:   1

Figure 16: A sample mechanical turk task for validating results of face detection API. Photo sample is taken from Flickr with creative common license. Turkers identify number of faces and their gender and age group.

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 7.

5.2.3 Statistical Methods

Next, I present statistical methods I used to model the two dependent variables, number of likes and number of comments. Both dependent variables are *count* variables. I 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 the 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 [43]. I 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. I use over-dispersion to test whether Poisson or negative binomial regression should be used. This test was suggested by Cameron and Trivedi [43], 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, I z-score all numerical variables before performing regression). For the variables with heavy tail distribution, such as followers count and photos count, I log transformed the variables before performing regression. I use Chi-squared statistics to find the statistical significance of our regression models, computing the reduction in deviance as compared to a null model.

5.3 Results

I use negative binomial regression to model the number of likes and comments on photos. The results of the regression are presented in table 10. I use the Chi-squared Test to find the significance of the regression model, by computing the reduction in deviance from a null model. For the likes model, the reduction in deviance is $\chi^2 = (5.2M - 1.2M)$, 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 the 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

Table 8: 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.

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%

variables.

I 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.

5.3.1 Effect of Control Variables

I use number of followers and number of photos as control variables. As expected, the followers count has a large positive effect on the number of likes and comments. This 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 the results (table 10) 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.

Table 9: Negative binomial regression results for number of likes.

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

Table 10: Negative binomial regression results for number of comments

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

5.3.2 Effect of Faces

All other predictors in the model come from the face detection results. The goal of this chapter is to quantify the effect of faces and their comparative importance on social engagement. I use a binary variable that reflects the existence of a face in the image. We can see in table 10 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

Figure 17: Visualization of negative binomial regression coefficients with dependent variables of likes and comments. Results show that photos with faces are significantly more likely to receive likes and comments.

$(IRR = 0.32)$ to receive comments[3]).

I 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. I did not include the number of faces in the final model to avoid co-linearity of the predictor variables.

5.3.3 Effects of Age and Gender

To test whether the demographic of users [6] biases toward photos with younger face groups, I considered using three binary variables each identifying the age of a face. Table 10 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.

Gender of the faces in the image does not show any strong effect on the image's engagement. Table 10 shows that the β coefficients for gender variables are of negligible size

---

[3]I use IRR to refer to *Incidence Rate Ratio*. I 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.

compared to some other features such as existence of a face.

5.4 Discussion

Using Instagram as research context, I set out to investigate how photos with faces relate to engagement, as measured as the number of likes and comments, compared to those without. I considered presence of a face in a photo, it's gender and age as predictors, controlling for social network reach and activity. From this I 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?

I find that among the factors I 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 the models. Furthermore, I 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 I 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.

5.4.1 Faces Engage Us.

The major finding of this chapter 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%. I 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. The current 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.

5.4.2 Age and Gender Do Not Impact Engagement.

The results of this work show that the age and gender of faces in the photo does not seem to drive or hinder it's 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.

5.5 Implications

This work raises many fundamental questions about the nature of social interaction around multimedia content. I believe that this is an initial step and there is a rich landscape 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. In this work I show 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 [103]. Emotional expressions in faces are known to activate several areas of the brain [79]. 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?

The current 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 this work is based on quantitative studies and observational data, I cannot make any strong causal claims. More experimental work needs to corroborate these findings. Further, the statistical methods I 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 [185]. 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. The current 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 [127], 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.

5.6 Chapter Summary

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 chapter, I took a first step toward uncovering an important feature of some of images, the human faces. I find that photos with faces are 38% more likely to be liked and 32% more likely to be commented on. The current 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, this findings open a new thread of future work, suggesting research in visual analysis.

CHAPTER VI

IMPACT OF FILTERED IMAGE CONTENT ON ENGAGEMENT

Henri Cartier-Bresson, the father of modern photo-journalism, and master of candid photography, showed a considerable lack of interest in the process of photography in general, likening photography with the small camera to an "instant drawing" [35]. Fred Lyon, a well known photographer who has been shooting for 70 years, argues for opposite:

> "The feeling is that if you are a "real" photographer it must be "pure."... I can absolutely say that I can make indistinguishable prints from gelatin silver and in some cases, better! I have a lot more control. Some people think that its fudging (to alter them digitally) but how is that different from a painter who picks up a brush and fixes or changes something? They would have a good argument if I insisted that my photography is all facsimile, is all "truth" but I have never presented any of it that way. My feeling is-whatever please people." [1]

Post-processing, once only done by photography artists, is now available to the public by the photo sharing mobile services such as Instagram and Flickr. The rise of mobile phone photography in popularity, has made post processing more available to the unprofessional photographer. We have seen the various iPhone models being Flickr's most popular cameras [2]. This presents a key research challenge for photo sharing communities like Flickr (cf. Instagram, 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 [16, 25, 40, 41, 105, 144], there are few studies on what makes visual content socially engaging online.

---

[1]`http://therumpus.net/2012/03/the-rumpus-interview-with-fred-lyon/` (accessed 11/2013)
[2]`http://www.flickr.com/cameras/` (accessed 9/2013)
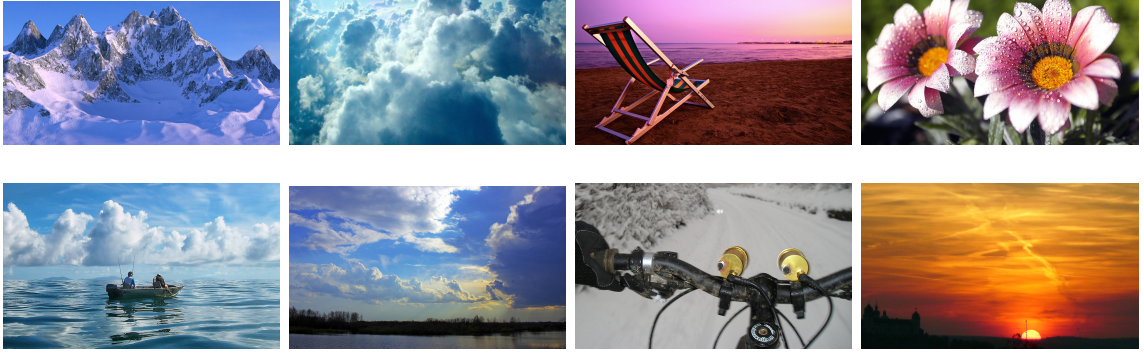
(a) Visually tagged as Nature: mountain (shared by *glenn-@flickr*), clouds (shared by *nirak@flickr*), beach (shared by *aigle_dore@flickr*), flower (shared by *Alice1664@flickr*), ocean (shared by *em-photos103@flickr*), sky (shared by *ponomarevsergey@flickr*), snow (shared by *restlessglobetrotter*), sunset (shared by *axel-d*)

(b) Visually tagged as People: portrait (shared by *21644167@N04@flickr*), face (shared by *through-paintedeyes@flickr*) , groupshot (shared by *abuaiman@flickr*)

(c) Visually tagged as outdoor (shared by *thartz00@flickr*), text (shared by *vblibrary@flickr*), low-light (shared by luchilu@flickr) and food (shared by gpeters@flickr)

Figure 18: Examples of (a) photos of natural scene (b) photos of people, and (c) photos taken of outdoor, text, foor or in low-light.

The content of a photo is of vital importance to how it is perceived and consumed by online users. For example, photos of faces are popular in social image sharing sites such as Flickr and Instagram [19]. Faces are specifically interesting as photo content, because humans love to look at faces; infants, barely minutes old, turn toward faces, sensing that they are important [148]. Research has also found that viewing nature can help reduce stress or help wellness [204]. Overall, communication through images has become so popular that sometimes users share quotes or other textual information in the format of an image. As Kelsey puts it [112], "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."

Vitally important to the sites that play host to images, this paper considers how filters impact engagement on various types of photo content. Whether existence of faces, natural scenes, textual content, etc. improve/hinder online experiences in the forms of engagement? While existing scholarly work has examined textual content in social networks [102], the research question around images has not received much attention. Working from a dataset of 4.9M Instagram and Flickr photos and 15 interview studies with Flickr mobile users, this work takes a new perspective in understanding practices of photo sharing, and how users engage with visual content. That is, I use filters as proxies for overall visual features to take a step towards uncovering the impact of visual stimuli on online user behavior. In particular, confirming results of chapter 4, I find that filtered photos are considerably more engaging than original ones. However, this engagement differs across different categories of content. For example photos of natural scenes are more engaging when posted as original rather than filtered and photos of food are more engaging when they are filtered.

6.1 Method

I use a mixed-method approach to understand the role of filters on engaging values of various types of content. My approach is two-fold; first, I take a quantitative approach and analyze data collected from Flickr mobile app in order to understand the role of filters in engaging users across various type of content; Next I perform a qualitative study and interview 15 of Flickr mobile users in order to gain deeper insight into why certain types of content are more popular with/without filters.

6.1.1 Qualitative Approach: Participants

To recruit users for the study, I emailed a pre-screening survey to a database of potential study participants. I screened for Flickr users who used the mobile app at least a few times a month. This assured that they were familiar with the app and the filter options through the mobile interface. Nine of all participants mentioned that they use the web interface, though interviews didn't focus on this usage. Those who matched the recruitment criteria were scheduled for a 60-minute interview session.

I recruited 15 participants. Of the 15 participants, 10 were interviewed in person and 5 were interviewed remotely through telephone and Skype. In all the sessions, participants were video-recorded. All participants were located in the United States. Those interviewed in person were located in the San Francisco Bay Area. The sample consisted of 5 women and 10 men ranging between 29 and 53 years old, with a median of 35. During the interview sessions, I mostly asked about participants use of camera phone, Flickr app and other photo sharing apps such as Instagram. I then asked participants to choose random photos from their photostream and explain why they filtered it or left it as original. Additionally I asked them to apply different filters on their choice of unfiltered photos and explain which filters work better. At times, I asked them to go back to their photos and describe how filtering changed the ways the photo looked. Other than the mobile phone camera, 8 of the participants stated that they own digital SLR cameras. The other 7 used their phone as their primary photo taking device.

6.1.2 Quantitative approach: Flickr Mobile Data

In this chapter, I aim to understand the impact of filters and content of photos on social engagement as measured in photo views, comments and favorites. I am specifically interested in the differences in filtered engagement based on the content of the photo. In this section, I describe the data collection process and summarize the descriptive statistics of the meta data.

I collected public photo meta-data from Flickr in November 2013. These photos were identified by Flickr as having been uploaded from Flickr's mobile app. In total, the dataset consists of over 4.9 million photos. I identified whether the photos were posted as original or filtered by checking their machine tags, auto-generated tags from the uploading application. 3.5 million of the photos are uploaded using the ios apps and 1.4 million are uploaded using android apps. I also identified the type of the filter that was applied.

Flickr auto-tags images using advanced vision techniques. The visual analysis detects certain objects (e.g. car, food, flower and plant), certain scenes (e.g. beach, clouds, mountain, ocean, sky, snow, street and sunset) and faces of people (including people, portrait and

groupshot). The visual analysis also determines whether the photo is taken at night, has text or is taken outdoors. With every visual tag detected in the image, there is a confidence level reported. For example if the visual analysis has detected a face with 99% confidence, the visual tag shows the level of the confidence along with the visual tag of the face. In this study I only use visual tags with level of confidence higher than 95%.

6.1.3 Dependent features.

Measures of engagement with photo are photo *views*, photo *comments* and photo *favorites*.

- *Implicit dependent variable.* I use number of views of each photo as an implicit measure of behavior. It quantifies the number of distinct users who viewed the photo. The higher number of views suggests that the photo was interesting to more number of people.

- *Explicit dependent variable.* Comments are explicit forms of actions taken on each photo. They quantify the number of distinct comments users posted on the photo. Comments is an explicit action [221] taken on the photo and higher number of comments shows that the photo more explicitly got attention. One of the participants explains her perception of comments in the following:

  "Sometimes, some people comment it just to get attention, so that you click on their photostream. Sometimes, they comment it if they already know you and they want to tell you something. Like, Oh yeah, I saw that too, or if it's a part of a conversation. Sometimes, there are assignments in groups and it's a conversation."(P2)

- *Interest dependent variable.* Favorites on a photo quantify the amount of interest in the photo. The number of favorites is the number of distinct users who favorited the photo. On Flickr, favorites are used both for curation of content and also as a social signal. Users may favorite photos to create a collection of their own or to show their interest in a certain photo; For example one of the participants mentioned that he uses favorites mostly to show that he likes a photo:

"It's more of like I think this is a good photo which I guess is sort of a like, but you can like for all sorts of other reasons on Facebook. For me, I'm using the favorite ads especially since that's the way that they use to operate more like it was just more obvious that I favored on Facebook, I mean I favor on Flickr was: Hello, I don't know you, but I like your photo. That's how I still use them and to get back to them if I want to show it to someone like that." (P1)

On the other hand many of the participants mentioned that he only uses favorites to bookmark the photos he wants to go back to later. These photos are collected under user's favorites and the collection is a form of content curation:

"I favorite to keep a bookmark of the picture. When I find a picture with me, I want to keep it, if I find the picture nice for some reason, I want to keep it and to watch it."(P7)

Although some use favorites as a social signal, and some only as a collection method, most of the participants mentioned both uses cases:

"I favorite a photo to, one, show my appreciation that I liked it and tell that person I liked it. Its a good photo. Secondly, to bookmark it. Sometimes, when a photos good and I want to come back and look at it and if I dont fave it, its gone. I dont know where it is. Those two reasons."(P2)

There are differences in how users perceive and adopt each of these social engagement metrics. These three engagement metrics are not necessarily related to each other, a photo might have high number of views but few favorites (See figure 19 as an example), or many favorites but not comments:

"I've noticed that to get the maximum views, the photo has to have something interesting. In a group for example, it's among the pool and among many other photos. My photo comes up as so small. It has to stand out in that small size

Figure 19: Participant describes how an interesting subject in the photo can get a lot of views but not necessarily favorites. Photo is shared with participant's permission.

that people see. If thats something unique, it really comes up. For instance, this picture [figure 19] is a very simple picture, but this has 600 views and this is one of my most popular photos in terms of views, but it is not popular in terms of favorites. Once people click on it and see, Oh, its just a comb. Eh. Nothing big deal about it, and they dont favorite it, I think. Because its not something that you want to set as wallpaper, I think, but there are other pictures, which didnt look as interesting, doesnt catch your eye, but once you click it, youre like, Oh, this is really good and I want to probably save it or wallpaper it or favorite it. Thats when they favorite it." (P2)

6.1.4 Control features.

In this subsection I describe the features that I use as controls. I consider four features to be able to control some of the scenarios that might affect the engagement values on the photo and also to compare the effects of filters and content with them.

- *Photostream views.* Every user on Flickr has a *photostream* which can be viewed by other users, depending on the privacy settings on a per photo basis[3]. Photostream

---

[3]Note that for the purpose of this study I only consider photos that are public and visible to everyone.

quantifies the level of implicit action mainly as a function of the user's popularity. Views of photosteam are usually obtained by direct views of user's profile and photostream, so when someone views a user's photostream, they have to click on the user's profile and as a result I assume that the profile and the user's collection were of interest.
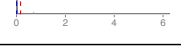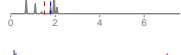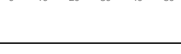
- *Followers.* Like Twitter, the Flickr relationship model between people is asymmetric: users form into social networks based on "follow" relationships. The number of followers is a measure of the user's audience size. This is a powerful and intuitive control, as we would expect users with larger audiences to have higher baseline probability of being viewed, commented on or favorited by their followers.

- *Tags.* Tags on Flickr are used by the search index to help people find photos. The higher number of tags usually imply that the photo will appear in more relevant searches and as a result may have higher chances of being viewed. I use the number of tags as control for searchability of a photo.

- *Number of photos.* I use the number of photos as a control variable for level of activity. Users who post more photos on Flickr are considered more active than those with fewer photos. The higher number of photos posted on one's profile is usually related to lower likelihood of a single photo being viewed.

6.1.5 Filter and Content features

My focus in this study is to determine the impact of filters on each type of photo content, and understand how they drive/ hinder engagement.

- *Filter feature.* For every image, I identify whether it was shared as original or was filtered before shared. I do so by checking the photos machine tags which are created by the uploading app. I code a new variable *is filtered* as a factor variable, with a value of 1 for filtered photo and 0 otherwise.

- *Content type.* Flickr determines whether a photo contains several scenes or objects. I group these tags into natural scenes (beach, clouds, flower, mountain, ocean, sky,

Table 11: Distributions of quantitative and categorical variables used in this paper. Variables marked with '*' are log transformed. The red dashed line identifies the mean and the blue line identifies the median of the distribution.

Variable	min	mean	median	max	std.dev	Distribution
views*	0	49.87	4	100.5K	351.74	
comments*	0	0.11	0	365	1.40	
favorites*	0	0.36	0	738	3.76	
tags	0	5.10	5	80	2.70	
photostream views*	0	11.6K	92	9.52M	73.5K	
photos*	0	3.3K	158	317.6K	12.03K	
followers*	0	38	5	15.8K	150.30	

Categorical			
filter	1M is filtered	1.9M is not filtered
content	900K has face	2M does not have face

snow or sunset), people (face, portrait, groupshot), outdoor, text, night and food. I use a categorical variable that identifies the photos with any of these categories of photos. Figure 18 shows examples of content types determined with high confidence ($> 95\%$) on Flickr.

6.2 Modeling engagement

The number of views, comments and favorites are all count variables. I model them using negative binomial regression, on two classes of independent variables: control features and features of interest (filters and content). Negative binomial regression is well-suited for *over-dispersed* distributions of count dependent variable [43]. I use negative binomial regression instead of Poisson regression since the variable of the dependent variable is larger than the mean for both views ($\mu = 49.87$, $\sigma = 351.74$), comments ($\mu = 0.11$, $\sigma = 1.40$) and favorites ($\mu = 0.36$, $\sigma = 3.76$). I use over-dispersion to test whether Poisson or negative binomial regression should be used. This test was suggested by Cameron and Trivedi [43],

and involves a simple least-squares regression to test the statistical significance of the overdispersion coefficient.

The negative binomial regression models the *expected* number of views, y_{views}, comments, $y_{comments}$, or favorites, $y_{favorites}$, for a photo as a function of control and interest features. For each dependent measure, I construct two regression models to evaluate the *impact* of control and interest variables: first to model control variables alone (control model), and the second to model both control and interest variables (full model). The reduction in deviance from the full model to the control-only model shows the significance of filters and content on describing the number of views, comments and favorites.

The first model uses control attributes as predictors of the number of views, comments or favorites on a photo.

$$ln(y_{views}) = I + \Sigma_i^{x_i \in controls} \beta_i x_i$$
$$ln(y_{comments}) = I + \Sigma_i^{x_i \in controls} \beta_i x_i \tag{1}$$
$$ln(y_{favorites}) = I + \Sigma_i^{x_i \in controls} \beta_i x_i$$

where I is the intercept for the model and the control sum is computed using the following control attributes:

$$\Sigma_i^{x_i \in controls} \beta_i x_i = \beta_{photostream} x_{photostream} + \beta_{tags} * x_{tags}$$
$$+ \beta_{followers} * x_{followers} + \beta_{no\_photos} * x_{no\_photos} \tag{2}$$

This model allows us to understand the effect on the number of views, comments and favorites of control variables alone. I then model the impact of filters and content types on the number of views, comments and favorites. I construct a second model that includes both control and interest attributes.

$$ln(y) = I + \Sigma_i^{x_i \in controls} \beta_i x_i + \Sigma_j^{x_j \in interest} \beta_j x_j \tag{3}$$

Where, the controls sum is taken from equation 2 and interest sum is computed using filter and face variables and their interaction.

Table 12: Summary of the models from equation 2 and 4 for the number of views. θ is the shape parameter of negative binomial distribution, Resid. df is the residuals degree of freedom for the fitted model. The chi-square test rejects the hypothesis and so the full model is significant.

Model	θ	Resid. df	2 x log-lik.
control model	0.43	2946943	-29195091.96
full model	0.44	3030060	-29158614.06
Summary			
LR.stat		22993	
degrees of freedom		18	
Pr(>Chisq)		$< 10^{-15}$	

Table 13: Summary of the models from equation 2 and 4 for the number of comments. θ is the shape parameter of negative binomial distribution, Resid. df is the residuals degree of freedom for the fitted model. The chi-square test rejects the hypothesis and so the full model is significant.

Model	θ	Resid. df	2 x log-lik.
control model	0.07	414999	-1575770.07
full model	0.07	477135	-1569749.08
Summary			
LR.stat		62136	
degrees of freedom		18	
Pr(>Chisq)		$< 10^{-15}$	

$$\Sigma_j^{x_j \in interest} \beta_j x_j = \beta_{is\_filtered} x_{is\_filtered} + \beta_{content} x_{content}$$

$$+ \beta_{is\_filtered\_and\_content} x_{is\_filtered} x_{content} \tag{4}$$

The regression coefficients β allow us to understand the effect of an independent variable on the number of views, comments and favorites (note that to be able to compare coefficients, I z-score all numerical variables before performing regression).

In order to choose which subset of independent variables should be included in the model, I use the *Akaike Information Criterion (AIC)*. AIC is a measure of the relative

Table 14: Summary of the models from equation 2 and 4 for the number of favorites. θ is the shape parameter of negative binomial distribution, Resid. df is the residuals degree of freedom for the fitted model. The chi-square test rejects the hypothesis and so the full model is significant.

Model	θ	Resid. df	2 x log-lik.
control model	0.18	1269819	-3259137.89
full model	0.18	1376276	-3255783.73
Summary			
LR.stat		106457	
degrees of freedom		18	
Pr(>Chisq)		$< 10^{-15}$	

quality of one model against another, and is defined as following:

$$AIC = -2L + 2k$$

where, k is the number of parameters and L is the maximum log-likelihood of the model. The smaller the value of AIC, the better the fit of the model. Starting with a full set of independent variables listed in the data section, and all possible interactions of those, I use a step-wise procedure to select the model that minimizes AIC. Using the model with minimum AIC also reduces the chances of choosing a model that overfits the data.

I test coefficients of all independent variables for the null hypothesis of a zero-valued coefficient (two-sided). This method is based on standard errors of coefficients, which is analogous to the *t-test* used in conventional regression analyses. I use a Chi-squared test with one degree of freedom to test the hypothesis that each coefficient β_j is zero. To do this, I compute the following term:

$$\chi^2 = \frac{b_j^2}{(SE_j)^2}$$

where, b_j is the estimate of β_j and SE_j is the standard error of the coefficient β_j. Table 15 shows the β coefficients and the *p*-values from the Chi-squared test. We see that all independent variables have coefficients that are statistically significant.

I use the deviance goodness of fit test to assess the regression fit [90]. The deviance is expressed as:

$$D = 2\sum_{i=1}^{n}(\zeta(y_i; y_i) - \zeta(\mu_i; y_i))$$

92

with $\zeta(y_i; y_i)$ indicating a log-likelihood function with every value of μ given the value y in its place. The $\zeta(\mu_i; y_i)$ is the log-likelihood function for the model being estimated.

The deviance is a comparative statistic. I use the *Chi-square test* to find the significance of the regression model, with the value of deviance and the degrees of freedom as two Chi-square parameters. The degrees of freedom is the number of predictors in each model. Tables 12, 13 and 14 summarize the model parameters and the goodness of fit test results, showing that the regression models are a good fit for the data.

6.3   Effect of Control Variables

The first class of variables I study are the control variables. In particular I look at the effect of *photostream views*, *number of photos* posted by user, *number of tags* on the photo and *number of followers* of the user. Tables 15, 16 and 17 summarize the β coefficients for both control and full models with views, comments and favorites as dependent variables. The models significance are summarized in tables 12, 13 and 14. All control models for views, comments and favorites show significance, $p < 10^{-15}$, for all predictors.

I use the *Chi-square* Test to find the significance of the regression model, by computing the reduction in deviance from a *null model*. Tables 12, 13 and 14 summarize the significance of each model. Now, I discuss the effect of each of the control variables.

- *Effect of Photostream views.*

  I use Photostream views as a control measure for the amount of engagement a photo receives. Photostream views control for how many views, comments or favorites were created as a result of viewing one's phototstream. As expected, photostream views is a strong predictor for the number of views, comments and favorites on Flickr. The relationship between photostream views and views is strong ($\beta = 0.89$, $p < 2e - 16$) and expected. As the higher volume of visits to one's profile, or higher aggregated views implies the higher likelihood of each photo having been viewed. I find that the size of β coefficient compared to other control variables such as followers is much larger in views model compared to comments and favorites model. In the comments model photostream seems to be significant ($\beta = 0.46$, $p < 2e - 16$) but smaller than

1/3 of the followers effect. This is while the size of photostream views coefficient was in the same order as the followers for the views. We see similar pattern in photostream views effect on favorites ($\beta = 0.65, < 2e - 16$). This means that while photostream views is strong predictor for the number of views, it's role in predicting the number of comments or favorites is significantly lower than the effect of number of followers.

- *Effect of Followers.*

  I use followers as a feature to control for how influential the user is or how large the audience set is. The network structure, or the size of audience following the user, may have strong impact on how many views, comments and favorites a photo receives. The follower/following relationship on Flickr allows followers of a user to see all public photo updates of that user in their feed. The larger the number of followers, the greater the audience of the shared photo and hence it is more likely for the photo to be viewed.

  From the views model, table 15 we see that the followers feature is a largest contributor to the number of views($\beta = 0.93, p < 2e - 16$). In the comments model, the followers seems to show larger effect size in describing the variance in the dependent variable ($\beta = 1.59, p < 2e - 16$). The role of followers in relationship to favorites is strong as well ($\beta = 1.68, p < 2e - 16$). One reason could be the significant role of number of followers in social network reach for certain user. While views is an implicit engagement metric and is more a metric of consumption, comments and favorites have more social identity and so the photos discovered through the social network (followers) are more likely to be commented on, or favorited while the photos discovered through the user's photostream are more likely to be consumed passively.

- *Effect of Tags.*

  Tags on Flickr have an important role in finding photos as they are highly used for search. I use tags as a measure to control for the fact that photos that are associated with more number of tags, might be easier to find.

  Engagement models of views, comments and favorites show that the role of tags in

predicting views, comments and favorites is not as strong as other control features such as photostream views and followers. While tags are positively related with more number of views ($\beta = 0.14$, $p < 2e - 16$) this relationship is not strong. The same trend is visible among models of comments ($\beta = 0.17$, $p < 2e - 16$) and favorites ($\beta = 0.22$, $p < 2e - 16$).

Many of participants mentioned that they seldom use search on Flickr for finding photos. Those who use search for viewing photos are usually looking for specific type of content on Flickr.

- *Effect of Photos.*

  I also control for the user's level of activity by considering the number of photos as an independent variable. The effect of activity on all dependent variables is negative, in all models of views ($\beta = -0.84$, $p < 2e - 16$), comments ($\beta = -1.21$, $p < 2e - 16$) and favorites ($\beta = -1.47$, $p < 2e - 16$) it is a significant and large effect. One explanation for negative relationship between engagement and number of photos is that the likelihood of per photo engagement decreases with the increase in the number of photos.

6.4   Effect of Filters and Content

My main objective in this chapter is to evaluate the impact of filters, photo content and the interplay between the two on photo engagement. Here, I discuss the results on effect of filters, photo content and filtered content on engagement.

6.4.1   Effect of Filters.

I first summarize the effects of filters on different engagement metrics. With regards to views, the results show that filters are indeed strongly positively correlated with the number of views ($\beta = 0.19$, $p < 2e - 16$). For the categorical variables such as filters I calculate the *Incidence Risk Ratio* to quantify the effect with respect to reference category. For the views we have: $IRR = 21\%$ which means that filtered photos are 21% more likely to be viewed compared to the original photos.

Table 15: Results of negative binomial regression with number of views as dependent variables.

	control model			full model		
Variable	β	Std.Err	p	β	Std.Err	p
(Intercept)	3.17	0.00	$< 2e-16$	3.14	0.00	$< 2e-16$
tags	0.14	0.00	$< 2e-16$	0.16	0.00	$< 2e-16$
photostream views	0.89	0.00	$< 2e-16$	0.88	0.00	$< 2e-16$
photos	-0.84	0.00	$< 2e-16$	-0.83	0.00	$< 2e-16$
followers	0.93	0.00	$< 2e-16$	0.93	0.00	$< 2e-16$
is filtered				0.19	0.03	$< 2e-16$
content:food				0.29	0.01	$< 2e-16$
content:nature				0.37	0.00	$< 6e-15$
content:low-light				-0.28	0.01	$< 2e-16$
content:outdoor				0.11	0.00	$< 2e-16$
content:people				0.23	0.00	$< 2e-16$
content:text				-0.27	0.01	$< 2e-16$
is filtered & content:food				0.15	0.04	$< 7e-4$
is filtered & content:nature				-0.17	0.03	$< 2e-4$
is filtered & content:low-light				-0.23	0.04	$< 2e-10$
is filtered & content:outdoor				-0.21	0.03	$< 5e-11$
is filtered & content:people				0.24	0.03	$< 4e-14$
is filtered & content:text				-0.17	0.03	$< 6e-8$

Table 16: Results of negative binomial regression with number of comments as dependent variables.

	control model			full model		
Variable	β	Std.Err	p	β	Std.Err	p
(Intercept)	-3.27	0.00	$< 2e-16$	-3.16	0.02	$< 2e-16$
tags	0.17	0.00	$< 2e-16$	0.19	0.00	$< 2e-16$
photostream views	0.46	0.01	$< 2e-16$	0.46	0.01	$< 2e-16$
photos	-1.21	0.01	$< 2e-16$	-1.21	0.01	$< 2e-16$
followers	1.59	0.00	$< 2e-16$	1.59	0.00	$< 2e-16$
is filtered				0.37	0.09	$< 2e-16$
content:food				0.28	0.03	$< 2e-16$
content:nature				0.16	0.02	$< 2e-16$
content:low-light				-0.06	0.03	0.02
content:outdoor				0.15	0.02	$< 1e-13$
content:people				0.20	0.02	$< 2e-16$
content:text				-0.18	0.02	$< 2e-16$
is filtered & content:food				0.11	0.03	0.04
is filtered & content:nature				-0.31	0.09	$< 1e-3$
is filtered & content:low-light				-0.40	0.12	$< 9e-4$
is filtered & content:outdoor				-0.32	0.10	$< 8e-4$
is filtered & content:people				0.37	0.10	$< 2e-5$
is filtered & content:text				-0.40	0.10	$< 6e-5$

The relationship among filtered photos and comments is strong and positive as well ($\beta = 0.37$, $p < 2e-16$) $IRR = 45\%$. It is 1.41 times more likely for a filtered photo to

Table 17: Results of negative binomial regression with number of favorites as dependent variables.

Variable	control model			full model		
	β	Std.Err	p	β	Std.Err	p
(Intercept)	-2.33	0.00	$2e-16$	-2.30	0.01	$< 2e-16$
tags	0.22	0.00	$2e-16$	0.22	0.00	$< 2e-16$
photostream views	0.65	0.00	$2e-16$	0.65	0.00	$< 2e-16$
photos	-1.47	0.00	$2e-16$	-1.48	0.00	$< 2e-16$
followers	1.68	0.00	$2e-16$	1.67	0.00	$< 2e-16$
is filtered				0.64	0.06	$< 2e-16$
content:food				0.15	0.02	$< 2e-16$
content:nature				0.08	0.01	$< 2e-16$
content:low-light				-0.01	0.02	0.03
content:outdoor				0.10	0.12	$< 4e-16$
content:people				0.04	0.01	$< 4e-5$
content:text				0-0.01	0.01	0.41
is filtered & content:food				0.62	0.09	$< 8e-12$
is filtered & content:nature				-0.20	0.06	$< 1e-3$
is filtered & content:low-light				-0.30	0.08	$< 2e-4$
is filtered & content:outdoor				-0.30	0.06	$< 5e-6$
is filtered & content:people				0.52	0.06	$< 2e-15$
is filtered & content:text				-0.27	0.06	$< 3e-5$

receive comments compared to an original photo. In the case of favorites the relationship is strong as well ($beta = 0.64$, $p < 2e-16$) and $IRR = 90\%$), suggesting that it is 90% more likely for filtered photos to receive favorites compared to original photos.

While the above findings are from the viewers point, I also observed that many of participants who are Flickr mobile app users expressed interest in applying filters:

> "I like some filters for the effects, like either making it starkly black and white, or giving it an old look, like the Antique look, or the Throwback look. I dont necessarily have to use the filter, but I like those features, basically, the set of filters." (P2)

6.4.2 Effect of Content.

I also use a categorical variable to quantify effect of different content types on photo engagement. My results in tables 15, 16 and 17 summarize the effect of photo content on engagement. I find that some groups of the photo content are more likely to be engaging than others, for example food, outdoor and people. I also analyze the impact of each group

of photos when they are filtered and find that some photo contents are more popular when filtered. For example photos of food and people are significantly more likely to be engaging when they are filtered. I describe the results of each photo content in the following.

- *Content:Food*

  Photos of food are highly likely to be viewed by Flickr users ($\beta = 0.29$, $p < 2e - 16$, $IRR = 34\%$). They are highly likely to be commented on ($\beta = 0.28$, $p < 2e - 16$, $IRR = 32\%$) and favorited ($\beta = 0.15$, $p < 2e - 16$, $IRR = 16\%$) as well. This suggests that photos of food are more popular than other types of content on Flickr and they are 34% more likely to be viewed, 32% more likely to be commented on and 16% more likely to be favorited.

  On the other hand, photos of food that are filtered are more likely to be viewed ($\beta = 0.15$, $p < 7e - 4$, $IRR = 16\%$) compared to the ones that are posted as original. I don't see significant effects of filtered photos of food on comments and favorites and so I cannot make claims on those types of engagements. It is however interesting that filtered photos of food are 16% more likely to be viewed by Flickr users.

  > "Photos of food, I'm more likely to filter specially if the whole plate looks like its one color of food, so I was like Oh, there's chicken fingers and there's fries on a yellow plate, so yes, definitely more than the people. The one thing that I ... Again, it's not really filtering, but I will take a whole bunch like a whole burst of photos that wants for ad or group shot or or for selfies just because I hate when they're blurry. I do not really filter so much like if it's a clear shot that's not shaky to me, that's good enough. But food is definitely, yeah. That also, in some situation where the camera's setting right in, I pay a lot more attention to the composition there. Again, it might just me that I'm more sensitive to, it's a thing and not a person like I'm more experienced in the photo. I don't know. Probably all at once. "

- *Content:Nature*

  Photos of nature are highly likely to be viewed ($\beta = 0.37$, $p < 6e - 15$, $IRR = 45\%$)

on Flickr. Nature photos are also likely to be commented on ($\beta = 0.16$, $p < 2e - 16$, $IRR = 17\%$) but not very likely to receive more favorites ($\beta = 0.08$, $p < 2e - 16$, $IRR = 8\%$) compared to other photo contents.

On the contrary, we see that filtered photos of nature are less likely to be viewed ($\beta = -0.17$, $p < 2e - 4$, $IRR = 18\%$), less likely to be commented on ($\beta = -0.31$, $p < e - 3$, $IRR = 36\%$) and less likely to be favorited ($\beta = -0.20$, $e - 3$, $IRR = 22\%$). This finding suggests that filtered photos of nature (sky, clouds, mountains, beaches, etc) are 18% less likely to be viewed, 36% less likely to be commented on and 22% less likely to be favorited, implying that photos of nature are more popular when posted as original.

Some of the participants describe that the photos of natural scenes are usually taken outdoor and so the lighting is sufficient.

> "If the thing about the nature or subject about the nature is the colors, which happens to be a lot of the time in my pictures, like fall or spring, so it's the colors that I want to capture. If it changes the color too much from what it actually was, I wouldn't use it."(P4)

Sometimes the photographers want to capture beauty in nature that is better presented intact rather than filters:

> " I think this has no filter. The reason was that this is a color photo. I don't want to make it black-and-white. Black-and-white filter would definitely not work. I also like the shade. Actually the pictures specially for the colors of the flower. Applying any filter would change that.I did not want that to happen."(P8)

- *Content:low-light*

  Photos taken in low-light are less likely to be viewed by Flickr users ($\beta = -0.28$, $p < 2e - 16$, $IRR = 32\%$) but their relationships with comments ($\beta = -0.06$, $p = 0.02$)

and favorites ($\beta = -0.01$, $p = 0.03$) are not significant. This could be due to low quality of photos taken by mobile cameras in low-light settings:

> "We went to somebody's birthday party a couple weeks ago, and there was somebody eating fire. I had a little film camera with me, and I had my iPhone, and I shot some photos with my iPhone. It was completely dark in there except for somebody eating fire, and it looked great on black and white, because it didn't look great with color, because the light was just gross and muddy and orange and weird. "(P10)

We also see that filtered photos taken at night are less likely to be viewed ($\beta = -0.23$, $p < 2e - 10$, $IRR = 26\%$), less likely to be commented on ($\beta = -0.40$, $p < 9e - 4$, $IRR = 49\%$) and less likely to be favorited ($\beta = -0.30$, $p < 2e - 4$, $IRR = 35\%$). This suggests that filters are contradictory to engagement for photos taken at night. In terms of filtering the low-light photos, we have to take into account the original photos are not great at the first places. There are not many options to modify a photo that is dark and does not have many colors.

- *Content:Outdoor*

  The outdoor photos on Flickr are more likely to be viewed ($\beta = 0.11$, $p < 2e - 16$, $IRR = 12\%$), more likely to be commented on ($\beta = 0.15$, $p < 1e - 13$, $IRR = 16\%$) and more likely to be favorited ($\beta = 0.10$, $p < 4e - 16$, $IRR = 10\%$). On the other hand filtered photos of outdoor are less likely to be viewed ($\beta = -0.21$, $p < 5e - 11$, $IRR = 23\%$), less likely to be commented on ($\beta = -0.32$, $p < 8e - 4$, $IRR = 38\%$) and less likely to be favorited ($\beta = -0.30$, $p < 5e - 6$, $IRR = 35\%$).

  One of the main characteristics of outdoor photos is the natural light. This factor on it's own can enhance a photo:

  > "I don't filter some of my photos. If the light is very good, and usually it's going to be something during the day. Likely outdoor. If the blue is perfect, you don't want to degrade it. "(P7)

That could be why photos of outdoors are engaging on Flickr. Although I did not find a correlation between photos of nature and photos of outdoor, most photos of natural scenes are taken outdoors.

- *Content:People*

  Photos of people are more likely to be viewed ($\beta = 0.24$, $p < 4e - 14$, $IRR = 27\%$), more likely to be commented on ($\beta = 0.20$, $p < 2e - 16$, $IRR = 22\%$). The effect of photos of people on favorites is positive but small ($\beta = 0.04$, $p < 4e - 5$, $IRR = 4\%$). When photos of people are filtered they are significantly more likely to be viewed ($\beta = 0.24$, $p < 4e - 14$, $IRR = 27\%$), more likely to be commented on ($\beta = 0.37$, $p < 2e - 5$, $IRR = 45\%$) and more likely to be favorited ($\beta = 0.52$, $p < 2e - 15$, $IRR = 68\%$). This suggests that photos of people are generally more likely to be engaging for Flickr users but also if they are filtered they are significantly more engaging than when they are posted as original.

- *Content:Outdoor*

  The outdoor photos on Flickr are more likely to be viewed ($\beta = 0.11$,$p < 2e - 16$,$IRR = 12\%$), more likely to be commented on ($\beta = 0.15$, $p < 1e - 13$,$IRR = 16\%$) and more likely to be favorited ($\beta = 0.10$,$p < 4e - 16$,$IRR = 10\%$). On the other hand filtered photos of outdoor are less likely to be viewed ($\beta = -0.21$,$p < 5e - 11$,$IRR = 23\%$), less likely to be commented on ($\beta = -0.32$,$p < 8e - 4$,$IRR = 38\%$) and less likely to be favorited ($\beta = -0.30$,$p < 5e - 6$,$IRR = 35\%$).

- *Content:People*

  Photos of people are more likely to be viewed ($\beta = 0.24$,$p < 4e - 14$,$IRR = 27\%$), more likely to be commented on ($\beta = 0.20$,$p < 2e - 16$,$IRR = 22\%$). The effect of photos of people on favorites is positive but small ($\beta = 0.04$,$p < 4e - 5$,$IRR = 4\%$). When photos of people are filtered they are significantly more likely to be viewed ($\beta = 0.24$,$p < 4e - 14$,$IRR = 27\%$), more likely to be commented on ($\beta = 0.37$,$p < 2e - 5$,$IRR = 45\%$) and more likely to be favorited ($\beta = 0.52$,$p < 2e - 15$,$IRR = 68\%$). This suggests that photos of people are generally more likely to be engaging for Flickr

users but also if they are filtered they are significantly more engaging than when they are posted as original.

> "here's the one that I didn't have to do a lot too because it was taken from Vista tower with a bright ... I was very dark and the background was very bright so I did have to filter that one just so you could see my face and the whole point of the selfie was that I was ... about to say."(P1)

- *Content:Text*

  Photos that contain text are less likely to be viewed ($\beta = -0.27$,$p < 2e - 16$,$IRR = 31\%$), less likely to be commented on ($\beta = -0.18$,$p < 2e - 16$,$IRR = 20\%$). The p-value for effect of text content on favorites is large ($p = 0.41$) and so I cannot claim anything on this effect. When photos of text content are filtered they are less likely to be viewed ($\beta = -0.17$,$p < 6e - 8$,$IRR = 19\%$), less likely to be commented on ($\beta = -0.40$,$p < 6e - 5$,$IRR = 49\%$) and less likely to be favorited ($\beta = -0.27$,$p < 3e - 5$,$IRR = 31\%$). The findings imply that photos that contain text are generally less engaging specially if they are filtered.

6.5 Discussion

Filters are becoming increasingly popular among users of mobile photo sharing tools and sites. Currently they are provided on mobile apps as generic tools which can be applied on the photos. In this work, not only I investigate the role of filters in engaging users with photos uploaded on Flickr, I also take a first step evaluating this role on each different type of content. Using advanced vision techniques provided by Flickr algorithms, I detect several groups of photo content, including people, nature, outdoor, food, text and night and evaluate their role as original posts and filtered photos in engaging users. I summarize the findings here and discuss them in detail.

6.5.1 Social Network Size is the Main Contributor to the Photo Engagement

With analysis of the control variables (photostream views, followers, photos and tags), I find that the size of social network audience is the main contributor to engagement values of

a photo. I found similar patterns across photos uploaded to Instagram (See chapter 5). On the other hand some of the work conducted on other social networks such as Twitter [47] found that the number of followers is the not the only factor shaping the influence of a user and the type of topics and content of posts are also important factors in determining popularity levels of the posts. Here, I find strong evidence that the social network reach, the number of followers, is a significant predictor to the engagement factors around the image but not the only one. For example we see that photostream views is also strongly related to the number of views, however the level of engagement predicted by the photostream views on comments and favorites is not as strong. This suggests that while social network reach is the main factor contributing to all types of engagement, it is relatively more impactful on comments and favorites compared to the photostream views. One scenario that can explain such observation is photo discovery through social network followers, or photo discovery through Flickr profile. While both seem to be highly impacting the views, the number of comments and favorites are more influenced by the social network.

I also find that tags are not as significant contributors to engagement as photostream views or followers. Tags are usually used as another way to discover a photo on Flickr. Users search for a specific content by a tag name and the photos associated with that tag appear in the search results. From the comparison between effect sizes of tags, photostream views and followers I hypothesize that tags are not as effective in discovery of a photo and engaging the user with the photo as the other two methods are. Finally, the number of photos are negatively related to the engagement level they receive. This suggests that the likelihood of a photo being noticed among pool of photos decreases with higher number of photos on the profile. The more number of photos shared by the user, the less likely each photo is to receive views, comments and favorites.

6.5.2 Photos of Natural Scenes and Outdoors

I show in this chapter, that user engagement differs across different types of shared content. Specifically, I find significant differences between photos of natural scenes and other content types. On Flickr, photos of nature are 45% more likely to receive views, 17% more likely

to receive comments and 8% more likely to receive favorites compared to all other type of photos. The significant likelihood increase in engagement suggests that Flickr users are drawn and engaged to nature photos. This might be a community specific behavior, specially since many of Flickr users are passionate about photography practices.

On the contrary, photos of nature do not seem to encourage engagement when they are filtered. Filtered photos of nature are 18% less likely to be viewed, 36% less likely to be commented on and 22% less likely to be favorited. This shows that Flickr users prefer the photos of nature to be posted as original and not filtered. One explanation is that this type of photos are more appreciated when presented as pure rather than post-processed. Additionally, one could argue that the inherent value of these kind of photos is more visible when they are presented as original. My findings here, specially sheds light on the user base of Flickr and their taste in photography.

I show that in general photos taken outdoors are more engaging, with 12% higher likelihood of being viewed, 16% higher likelihood of being commented on and 10% higher likelihood of being favorited. Similar to photos of nature, outdoor photos are less engaging when they are filtered. When outdoor photos are filtered, they are 23% less likely to be viewed, 38% less likely to be commented on and 35% less likely to be favorited.

6.5.3 Photos of Food and People

Photos of food content are highly likely to be engaging on Flickr. I find that photos of food are 34% more likely to ve viewed, 32% more likely to be commented on and 16% are more likely to be favorited. This suggests that photos of food are highly engaging on Flickr, they are also 16% more likely to be viewed when they are filtered. Based on findings of this research, one can design new features where the filters are suggested to the right type of content where it can improve engagement.

I also investigated how photos with people relate to engagement, as measured as the number of views, comments and favorites, compared to those without. I considered presence of a face in a photo, controlling for social network reach and activity. I also considered the interaction factor between faces and filters as a predictor.

The results of photos of people on Flickr is consistent with the study conducted on Instagram [19] which is summarized in chapter 5 where I found that photos of human faces are more engaging on Instagram than those photos that don't have faces in them. The current results show that photos of people are more likely to be engaging on Flickr as well. They are 27% more likely to be viewed and 22% more likely to be commented. This could be explained by the fact that humans are naturally drawn to faces and they like to view photos of themselves, their friends and even faces of strangers.

6.5.4   Photos of Text and Photos Taken in low light

The present findings show that photos of text are not engaging on Flickr. These photos are 31% less likely to be viewed and 20% less likely to be commented on. Filtering such photos does not impose a positive effect. Filtered photos of textual content are 19% less likely to be viewed, 49% less likely to be commented on and 31% less likely to be favorited.

I also find that photos taken at night are more likely to be viewed (32%). Despite their higher likelihood of view, when they are filtered they are 26% less likely to be viewed, 49% less likely to be commented on and 35% less likely to be favorited. This is perhaps related to the photo transformation used in filters. Most filters manipulate colors to add effects. Photos taken at night are usually dark and so change of color might not make it more professional looking.

6.5.5   Implications and Future Work

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. The practical implication of social engagement in online photo sharing lies strongly in search and recommendation. Knowing photos with filters increase engagement suggests one could increase their search ranking to keep people on site and active. The current results highlight the importance of effective methods that take advantage of filtered photos for personalization of site content. Additionally, designers can take advantage of the findings of this research on various types of content, and customize filters based on the content. For example I showed that filtered

photos of food are more likely to be engaging while filtered photos of nature are less likely to attract user engagement. Based on these findings one can suggest filters on the right type of content.

Future work can look at other visual characteristics of multimedia and study their impact on online behavior. It is also worth considering how filters are used in different contexts. For example, are people using highly saturated filters on photos of food, while using aging filters on human faces?

6.6 Chapter Summary

Filters are used to add effects to images before sharing them online. 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 chapter, I took a step toward understanding the role of photo content and visual filters on engaging users. I find that photos with filters have higher chances of being viewed, commented on, and favorited on Flickr but I don't claim if filters are the exact cause of this. I took a step further and evaluated the role of photo content. The results show that filters impact different types of content differently.

CHAPTER VII

REFLECTIONS AND FUTURE WORK

In this chapter, I revisit the dissertation's research questions and describe the contributions to the research community. I then summarize the contributions of this dissertation and argue for the importance of image presentation and content in understanding user behavior and social media practices. Finally, I reflect on design implications and future directions from my research.

7.1 Research Questions Revisited

The thesis of this dissertation has been around the composition and content of images, which impact the user engagement and content diffusion. To investigate this thesis, the guiding question has been what in images attracts users' attention and get them engaged with the content. I studied widely used image sharing sites in this dissertation. My research questions are developed around two main themes: pixel and macroscopic composition, and the content of the image. The first two research questions in this dissertation investigate the role of color composition: at a microscopic level of pixels and macroscopic level of filters. The next two research questions are more focused on the subjects in the image. I also investigate the interaction between the macroscopic composition and the content of image.

RQ1: What is the role of colors in how much the image diffuses in the network?
To answer RQ1, I conducted a quantitative study on social networking site Pinterest. I collected the images and their associated metadata from Pinterest and built a statistical model for the relationship between the number of repins and the dominant color in the image. The results from this work are reported in Chapter 3. I found four high level patterns of diffusion with respect to colors:

- Images with color diffuse more than those that are black and white.

- Red, purple and pink drive diffusion. Images with these dominant colors are more

likely to be shared (repinned) on Pinterest.

- Blue, green, yellow and black suppress diffusion. Images with these dominant colors are less likely to be shared (repinned) on Pinterest.

- Saturated images are more likely to be shared (repinned) on Pinterest.

RQ2: Why do we use filters and how do they impact engagement? To explore RQ2, I conducted two studies. The first one is a qualitative study where I interviewed 15 Flickr mobile app users to understand how and why they filter their photos. Then, I performed a large scale analysis of Flickr mobile photos to understand how filters impact photo engagement. My results show that:

- There are two groups of Flickr mobile users. The first group are those who are serious photography hobbyists and own professional cameras in addition to cameras in their mobile devices. The second group are users who use their mobile phone cameras as their primary photography device.

- Serious photography hobbyists use filters to correct errors and enhance their photos. They are mainly interested in mild filters and do not want to create an artificial look to their photos.

- Casual photographers like to add cool and unique look to their photos by applying filters. Filters help them make their photos more cool and fun.

- Filtered photos are more engaging than original photos. Filtered photos are 21% more likely to be viewed and 45% more likely to receive comments, compared to non-filtered ones.

- Filters that impose warm color temperature, boost contrast and increase exposure are more likely to be noticed. Other filters showed no significant effects on engagement.

RQ3: How do faces in photos impact engagement? To explore RQ3, I used a large dataset from Instagram photos. I used computer vision algorithms to detect faces and their

age and gender in the image. I then modeled engagement (likes and comments on the photo) using several control features and face attributes. I found the following:

- Existence of a face in a photo significantly affects its social engagement. The effect is substantial, with photos of faces being 38% more likely to be liked and 32% more likely to receive comments.

- The age and gender of faces in the photo does not seem to drive or hinder its engagement value.

RQ4: How does the content of a photo impact its engagement and how do different categories of content interact with filters? Inspired by findings of RQ2 and RQ3, I designed a more general study where I investigated the role of different content categories on engagement. I also considered each of these categories of content in the presence and absence of filters. I used a mixed-method approach to perform this study: a large-scale analysis of Flickr mobile photos complemented by interviews with 15 Flickr mobile users. My results showed that:

- Photo engagement differs with the category of shared content. Specifically, I find significant differences between photos of natural scenes and other content categories. On Flickr, photos of nature are 45% more likely to be viewed, 17% more likely to receive comments and 8% more likely to be favorited. On the contrary, photos of natural scenes are less engaging when they are filtered. Filtered photos of nature are 18% less likely to be viewed, 36% less likely to be commented on and 22% less likely to be favorited.

- Generally, photos that are taken outdoors are more engaging, with 12% higher likelihood of being viewed, 16% higher likelihood of being commented on and 10% higher likelihood of being favorited. When outdoor photos are filtered, they are 23% less likely to be viewed, 38% less likely to be commented on and 35% less likely to be favorited.

- Photos of food content are highly likely to be engaging on Flickr. I found that photos of food are 34% more likely to be viewed, 32% more likely to be commented on and 16% are more likely to be favorited. They are also 16% more likely to be viewed when they are filtered.

- Results on engagement of photos containing people were in agreement with the results found on Instagram. On Flickr, photos of people are 27% more likely to be viewed and 22% more likely to receive comments.

- Photos that contain text are 31% less likely to be viewed and 20% less likely to be commented on. Filtered photos of text are 19% less likely to be viewed, 49% less likely to be commented on and 31% less likely to be favorited.

- Photos taken in low light are 31% more likely to be viewed. Despite their higher likelihood of receiving views, when they are filtered they are 26% less likely to be viewed, 49% less likely to be commented on and 35% less likely to be favorited.

7.2   Contributions

This dissertation opens a new line of research towards understanding image properties that impact the image's engagement and diffusion. The thesis explores the critical relationship between image content and user engagement and adoption of content. I first ask "whether" there is a relationship, and second, "under what circumstances". By considering different forms of image content and different presentations of the image, this dissertation sheds light on some of the important aspects of online user behavior, which have not been considered in prior research and design works. The dissertation contrasts offline findings from cognitive and psychological studies of human behavior with basic interaction levels with online image content – allowing a rich, multilevel picture of the relationship between users and their engagement with images.

This dissertation contributes to human-computer interaction (HCI) theory and practice. At the theoretical level, the work confirms previous findings on interactivity, user engagement and understanding of people in a new setting. It clarifies how colors and visual

presentation of images could be related to evocative content, connecting to psychological studies of color, ownership of digital artifacts, and perceptions of faces and natural scenes. It adds to our understanding of online content consumption and how attention depends on the category of image content. I show that different types of content receive different responses from the social media users, a finding that illustrates the need for a deeper understanding of the image itself. This dissertation contributes a detailed examination of the effect of pictorial content on social actions and how image-mediated communication takes place on social platforms.

At a practical level, this work builds models of engagement and social adoption of content that can be deployed by image sharing communities. The models are simple and computationally inexpensive compared to blackbox models used in the industry. The models developed in this thesis can be used to tailor social image feed and advertisements, by showing people the content they most care about and are likely to engage with. Current recommendation models are mainly based on tags, text description or category labels. The models developed in this thesis can enhance such recommendation platforms by pre-populating tag lists according to content of the image, colors and visual information. They can associate weights to colors and content types that might be more engaging or interesting to the user, easing the burden on users to classify and tag their content.

Finally, this dissertation presents classes of image features that are associated with improvements in social capital, enabling users to interact with their content of interest and eventually have a better experience. Understanding how these classes generally operate is critical to designing new features on image sharing communities. Results presented in this thesis suggest what type of colors help in designing viral content. Moreover, the results can be used to design new filters for images. For example, filters that increase saturation or enhance the warmness of the image will likely increase engagement with the photo. In other words, a filter that saturates the image and adds red tint will probably be better than a blueish one, in terms of virality.

7.3 Design Implications

In this section I summarize some of the design implications of the findings in this dissertation.

7.3.1 Constructing Engaging and Viral Content

At a high level, the dissertation sheds light on how to construct viral and engaging image content. Through findings of Chapter 3, I found that typically warm and exciting colors affect the recipient's likelihood of sharing the image. The results suggest that using warm, saturated colors can increase chances of diffusion compared to images with cool and relaxing themes. On the other hand, the findings of Chapter 4 suggest that filtered content is more engaging when the filters boost contrast and exposure, and introduce warm colors. I found that photos of people, food and nature are likely to be engaging on Flickr. Additionally, filters impact engagement of different categories of content differently. The findings can help designers in more efficient construction of engaging and viral content. For example, when engagement is a key objective, a user can take advantage of presence of faces in the photo, filter the photo with warmer color temperatures and post it on the blog or social network.

7.3.2 Filter Design

Designers can use the findings of this work to design new filters for images. Filters that increase saturation or enhance the warmness of the image will likely increase diffusion; and filters that boost contrast and exposure in the photo will likely have high engagement. In other words, a filter that saturates the image and adds red tint will probably be better than a blueish one, in terms of virality. Additionally, a filter that adds more contrast and exposure to the photo will be viewed more often.

Using findings of filter work, I showed that mobile photographers belong to two groups: serious hobbyists and casual photographers. While there are many similarities in their use of mobile phones for photography, many editing needs and aesthetic values differ between the two groups. Better user experience can be designed by considering these two groups of

users and their expectations from filters. For example, the app can offer tuning options on each filter. As I mentioned earlier, many of the serious photography hobbyists prefer their effects mild and filters to be less noticeable. On the other hand the casual photographers like filters that make their photos visibly different and unique. Allowing users to tune the changes of a filter on their photos will help both groups. Designers can also allow users to select parts of image for filtering, so that they can have more control over their subjects in the photo. Of course, these changes should not influence the complexity of the filter use, since one of the main reasons people like to filter their photos is the simplicity of using the filters.

7.3.3 Search and Recommendation

A practical implication of social engagement in online photo sharing is search and recommendation. With knowledge of the content of photos and what content positively impacts engagement, one could increase the search ranking of such photos to retain people on the site and make them active users. The results in this dissertation, especially Chapters 5 and 6, highlight the importance of effective methods that take advantage of the content of photos for user personalization. Additionally, while we have seen visual recognition applications for social media sites [127], these tools have been designed for the utility of retrieval and not for conversation and comments. The findings of this dissertation may shed light on how to prioritize and highlight photos from the global image stream, especially images are new and therefore haven't had the time to accumulate likes and comments.

7.3.4 Filter Customization

In Chapter 6, I showed how engagement varies across different types of content with and without filters. Photos of people and photos of food are more engaging when they are filtered. However, filters do not seem to improve engagement when applied to photos of natural scenes and photos taken at low light. Designers can take advantage of these findings in customizing filter experience. For example, a photo sharing app can detect the content of the photo and suggest sharing the original photo if it is a photo taken in low light, or suggest applying filters if the photo contains food. Such designs could help in improving

overall engagement on the site and encouraging social capital around the shared content.

7.4   Future Directions

Social computing is a nascent, yet rapidly emerging field. It has attracted interest across sub-disciplines of computer science as well as the social sciences. This dissertation uses a variety of tools and advances in computer science, spanning machine learning, data mining and computer vision, to better understand user behavior. It asks many fundamental questions about the nature of social interaction around image content. The dissertation connects to psychological studies of color and image perception, and work in mobile communication and multimedia-mediated communication. This section articulates potential follow-up research work in social computing, multimedia-mediated communication and understanding users.

User engagement is a key concept in designing user-centered social media applications. It refers to the quality of the user experience and emphasizes the positive aspects of the interaction. In this dissertation, my focus was mostly on a few aspects of engagement such as likes, comments, views and favorites. There is more about engagement than these signals. Future research can look at other aspects of engagement such as click-through rates, return ratios, sentiment in the interaction, amount of time spent on the content, etc. Furthermore, it is interesting to explore different dimensions along perception of content. For example, whether the interaction occurred because of interesting content or personal relationship with the owner of the content. Current models of user engagement do not consider properties of the image as influential factors in defining engagement. A future research direction could focus on improving such models. Additionally, we can benefit from understanding all aspects of engagement in a dynamic setting where the time of interaction is considered.

My dissertation mainly focuses on two aspects of user behavior: engagement and adoption of content. I showed that presentation and content of photos are strongly tied to their engagement and diffusion. What other aspects of online behavior could be dependent on the image content? Future work could look at how relationships form and change based on

the image content that is shared on social media. Another interesting direction could be to look at time-dependent patterns of behavior based on available content. Does certain image content drive users away from the site or encourages their return? Furthermore, research in the traditional setting may need to be revisited when looking at behavior on mobile devices. A possible direction could be to study image engagement, perception and in general, visual traces created by mobile users to better understand the limitations and challenges on mobile platforms.

With respect to content of images, my findings suggest using qualitative methods to answer "why" questions. For example, I show that images of people are more engaging. This opens up new research directions concerning 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?

The context in which people appear in the image also raises interesting questions about individual and group behavior. Are photos of friends groups more/less popular than the family groups? What about selfies and people's reaction to self-portraits? An open question is the cultural impact on photo sharing; for example, are group photos more engaging in collectivism cultures than individualistic ones?

In Chapter 5, I showed 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 expressions, gaze direction, as well as, body posture and movement. Although facial expressions reliably signal basic emotions such as fear or happiness, research has shown that human viewers are also surprisingly adept at making reliable judgments about social information from impoverished stimuli, such as faint changes in facial expressions [103]. Emotional expressions in faces are known to activate several areas of the brain [79]. A possible future direction is to look at emotional expressions of faces and explore these effects on user behavior. For example, are we more likely to comment on wry smiles or broad grins?

REFERENCES

[1] "150,000,000 photos." `http://blog.instagram.com/post/8758450298/150-000-000-photos`, August 2011.

[2] "Here's how to use instagram." `http://www.businessinsider.com/instagram-2010-11?op=1`, May 2011.

[3] "Testing, testing...." `http://blog.instagram.com/post/8758396471/testing-testing`, July 2011.

[4] "A review of social media's newest sweetheart." `http://www.engauge.com/assets/pdf/engauge-pinterest.pdf`, March 2012.

[5] "Who's driving shopping traffic for retail sites? facebook, pinterest and twitter compared." `http://richrelevance.com/downloads/richrelevance-social-infographic-us-2012-09.pdf`, September 2012.

[6] "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.

[7] "Face++." `http://en.faceplusplus.com/`, 2013.

[8] ADAMS, F. M. and OSGOOD, C. E., "A cross-cultural study of the affective meanings of color," *Journal of cross-cultural psychology*, vol. 4, no. 2, pp. 135–156, 1973.

[9] AHERN, S., ECKLES, D., GOOD, N. S., KING, S., NAAMAN, M., and NAIR, R., "Over-exposed?: privacy patterns and considerations in online and mobile photo sharing," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 357–366, ACM, 2007.

[10] AMES, M., ECKLES, D., NAAMAN, M., SPASOJEVIC, M., and VAN HOUSE, N., "Requirements for mobile photoware," *Personal and Ubiquitous Computing*, vol. 14, no. 2, pp. 95–109, 2010.

[11] ANDERSON, A. H., NEWLANDS, A., MULLIN, J., MARIE FLEMING, A., DOHERTY-SNEDDON, G., and VAN DER VELDEN, J., "Impact of video-mediated communication on simulated service encounters," *Interacting with Computers*, vol. 8, no. 2, pp. 193–206, 1996.

[12] ANDERSON, E. W., "Customer satisfaction and word of mouth," *Journal of service research*, vol. 1, no. 1, pp. 5–17, 1998.

[13] ANECKI, J., "Selling in cyberspace: Electronic commerce and the uniform commercial code," *Gonz. L. Rev.*, vol. 33, p. 395, 1997.

[14] APTER, M. J., "Some data inconsistent with the optimal arousal theory of motivation," *Perceptual and motor skills*, vol. 43, no. 3f, pp. 1209–1210, 1976.

[15] ARAL, S., MUCHNIK, L., and SUNDARARAJAN, A., "Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks," *Proceedings of the National Academy of Sciences*, vol. 106, no. 51, pp. 21544–21549, 2009.

[16] 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*, pp. 959–968, ACM, 2006.

[17] ASLAM, M. M., "Are you selling the right colour? a cross-cultural review of colour as a marketing cue," *Journal of marketing communications*, vol. 12, no. 1, pp. 15–30, 2006.

[18] BAKER, M. J. and CHURCHILL J., G. A., "The impact of physically attractive models on advertising evaluations," *Journal of marketing research*, pp. 538–555, 1977.

[19] BAKHSHI, S., SHAMMA, D. A., and GILBERT, E., "Faces engage us: Photos with faces attract more likes and comments on instagram," in *in Proceedings of the SIGCHI conference on Human Factors in computing systems,*, 2014.

[20] BAKSHY, E., KARRER, B., and ADAMIC, L. A., "Social influence and the diffusion of user-created content," in *Proceedings of the 10th ACM conference on Electronic commerce*, pp. 325–334, ACM, 2009.

[21] BALABANOVIĆ, M., CHU, L. L., and WOLFF, G. J., "Storytelling with digital photographs," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 564–571, ACM, 2000.

[22] BARRY, A. M., *Visual intelligence: Perception, image, and manipulation in visual communication.* SUNY Press, 1997.

[23] BARUA, H., *The red river and the blue hill.* Lawyer's Book Stall, 1956.

[24] BELLIZZI, J., CROWLEY, A. E., and HASTY, R. W., "The effects of color in store design," *Journal of retailing*, vol. 59, no. 1, pp. 21–45, 1983.

[25] BERGER, J. and MILKMAN, K. L., "What makes online content viral?," *Journal of marketing research*, vol. 49, no. 2, pp. 192–205, 2012.

[26] BERLYNE, D. E., *Conflict, arousal, and curiosity.* McGraw-Hill Book Company, 1960.

[27] BERRY, D. L., "Comprehension and recall of internet news: A quantitative study of web page design.," 1999.

[28] BHATT, C. A. and KANKANHALLI, M. S., "Multimedia data mining: state of the art and challenges," *Multimedia Tools and Applications*, vol. 51, no. 1, pp. 35–76, 2011.

[29] BIERNACKI, P. and WALDORF, D., "Snowball sampling: Problems and techniques of chain referral sampling," *Sociological methods & research*, vol. 10, no. 2, pp. 141–163, 1981.

[30] BIOCCA, F., "The cyborg's dilemma: Progressive embodiment in virtual environments," *Human Factors in Information Technology*, vol. 13, pp. 113–144, 1999.

[31] BIRDWHISTELL, R. L., "Field methods and techniques: Body motion research and interviewing," *Human organization*, vol. 11, no. 1, pp. 37–38, 1952.

[32] BIRREN, F., "The power of color," 1997.

[33] BOHLE, R. H. and GARCIA, M. R., "Reader reactions to color in newspapers.," 1986.

[34] BOWMAN, D. and NARAYANDAS, D., "Managing customer-initiated contacts with manufacturers: The impact on share of category requirements and word-of-mouth behavior," *Journal of Marketing Research*, pp. 281–297, 2001.

[35] BRESSON, H. C. and TÉRIADE, E., "The decisive moment," 1952.

[36] BROADBENT, D. E., "Successive responses to simultaneous stimuli," *Quarterly Journal of Experimental Psychology*, vol. 8, no. 4, pp. 145–152, 1956.

[37] BROWN, J. J. and REINGEN, P. H., "Social ties and word-of-mouth referral behavior," *Journal of Consumer Research*, pp. 350–362, 1987.

[38] BRUCE, V. and YOUNG, A., *In the eye of the beholder: the science of face perception.* oxford university press, 1998.

[39] 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*, pp. 945–954, acm, 2009.

[40] 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*, pp. 21–39, Springer, 2007.

[41] 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*, pp. 281–284, ACM, 2008.

[42] BYRNE, D., "Colors / pink." `http://cabinetmagazine.org/issues/11/pink.php`, 2003.

[43] CAMERON, C. A. and TRIVEDI, P. K., *Regression analysis of count data (econometric society monographs).* Cambridge university press, September 1998.

[44] CARTER, S. and MANKOFF, J., "When participants do the capturing: the role of media in diary studies," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 899–908, ACM, 2005.

[45] CASSELL, J. and OTHERS, "Nudge nudge wink wink: elements of face-to-face conversation for embodied conversational agents," *Embodied Conversational Agents*, pp. 1–27, 2000.

[46] CENTOLA, D., "The spread of behavior in an online social network experiment," *Science*, vol. 329, no. 5996, pp. 1194–1197, 2010.

[47] CHA, M., HADDADI, H., BENEVENUTO, F., and GUMMADI, P. K., "Measuring user influence in twitter: The million follower fallacy.," *ICWSM*, vol. 10, pp. 10–17, 2010.

[48] Chapanis, A., Ochsman, R. B., Parrish, R. N., and Weeks, G. D., "Studies in interactive communication: I. the effects of four communication modes on the behavior of teams during cooperative problem-solving," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 14, no. 6, pp. 487–509, 1972.

[49] Chen, W., Wang, Y., and Yang, S., "Efficient influence maximization in social networks," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 199–208, ACM, 2009.

[50] Clarke, R., Wright, P., Balaam, M., and McCarthy, J., "Digital portraits: photo-sharing after domestic violence," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2517–2526, ACM, 2013.

[51] Click, J. and Stempel, G. H., "Reader response to front pages with four-color halftones," *Journalism & Mass Communication Quarterly*, vol. 53, no. 4, pp. 736–738, 1976.

[52] Courtis, J. K., "Colour as visual rhetoric in financial reporting," in *Accounting Forum*, vol. 28, pp. 265–281, Elsevier, 2004.

[53] Crowley, A. E., "The two-dimensional impact of color on shopping," *Marketing letters*, vol. 4, no. 1, pp. 59–69, 1993.

[54] Daft, R. L. and Lengel, R. H., "Organizational information requirements, media richness and structural design," *Management science*, vol. 32, no. 5, pp. 554–571, 1986.

[55] Darwin, C., *The expression of the emotions in man and animals*. Oxford University Press, 1998.

[56] Davey, P., "True colours," *Architectural Review*, vol. 204, no. 1221, pp. 34–5, 1998.

[57] De Bruyn, A. and Lilien, G. L., "A multi-stage model of word of mouth through electronic referrals," *eBusiness Research Center, Working Paper*, 2004.

[58] Dehn, D. M. and Van Mulken, S., "The impact of animated interface agents: a review of empirical research," *International journal of human-computer studies*, vol. 52, no. 1, pp. 1–22, 2000.

[59] DesMarais, C., "Facebook's instagram says it has 90 million monthly active users." `http://www.techhive.com/article/2025801/facebooks-instagram-says-it-has-90-million-monthly-active-users.html`, february 2013.

[60] 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*, vol. 33, no. 2, p. 178, 1976.

[61] Doherty-Sneddon, G., Anderson, A., O'Malley, C., Langton, S., Garrod, S., and Bruce, V., "Face-to-face and video-mediated communication: A comparison of dialogue structure and task performance.," *Journal of Experimental Psychology: Applied*, vol. 3, no. 2, p. 105, 1997.

[62] DREW, S. E., DUNCAN, R. E., and SAWYER, S. M., "Visual storytelling: A beneficial but challenging method for health research with young people," *Qualitative Health Research*, vol. 20, no. 12, pp. 1677–1688, 2010.

[63] EHMKE, C. and WILSON, S., "Identifying web usability problems from eye-tracking data," in *Proceedings of the 21st British HCI Group Annual Conference on People and Computers: HCI... but not as we know it-Volume 1*, pp. 119–128, British Computer Society, 2007.

[64] EKMAN, P. and FRIESEN, W., *Pictures of facial affect.* consulting psychologists press, 1975.

[65] ELLIOT, A. J., MAIER, M. A., MOLLER, A. C., FRIEDMAN, R., and MEINHARDT, J., "Color and psychological functioning: The effect of red on performance attainment.," *Journal of Experimental Psychology: General*, vol. 136, no. 1, p. 154, 2007.

[66] ELLIOT, A. J. and NIESTA, D., "Romantic red: red enhances men's attraction to women.," *Journal of personality and social psychology*, vol. 95, no. 5, p. 1150, 2008.

[67] ELLISON, N., STEINFIELD, C., and LAMPE, C., "Connection strategies: social capital implications of facebook-enabled communication practices," *New media & society*, vol. 13, no. 6, pp. 873–892, 2011.

[68] EMMISON, M. and SMITH, P., "Researching the visual: Images, objects, contexts and interactions in social and cultural enquiry," 2000.

[69] ERDOES, R. and ORTIZ, A., *American Indian myths and legends.* Random House LLC, 1984.

[70] EVELAND JR, W. P. and DUNWOODY, S., "Examining information processing on the world wide web using think aloud protocols," *Media Psychology*, vol. 2, no. 3, pp. 219–244, 2000.

[71] FEHR, E., KIRCHSTEIGER, G., and RIEDL, A., "Gift exchange and reciprocity in competitive experimental markets," *European Economic Review*, vol. 42, no. 1, pp. 1–34, 1998.

[72] FESTINGER, L., RIECKEN, H. W., and SCHACHTER, S., "When prophecy fails.," 1956.

[73] FRENZEN, J. and NAKAMOTO, K., "Structure, cooperation, and the flow of market information," *Journal of Consumer Research*, pp. 360–375, 1993.

[74] GALE, S., "Human aspects of interactive multimedia communication," *Interacting with computers*, vol. 2, no. 2, pp. 175–189, 1990.

[75] GEMMELL, J., BELL, G., and LUEDER, R., "Mylifebits: a personal database for everything," *Communications of the ACM*, vol. 49, no. 1, pp. 88–95, 2006.

[76] GEMMELL, J., WILLIAMS, L., WOOD, K., LUEDER, R., and BELL, G., "Passive capture and ensuing issues for a personal lifetime store," in *Proceedings of the the 1st ACM workshop on Continuous archival and retrieval of personal experiences*, pp. 48–55, ACM, 2004.

[77] GODES, D. and MAYZLIN, D., "Firm-created word-of-mouth communication: Evidence from a field test," *Marketing Science*, vol. 28, no. 4, pp. 721–739, 2009.

[78] GODES, D., MAYZLIN, D., CHEN, Y., DAS, S., DELLAROCAS, C., PFEIFFER, B., LIBAI, B., SEN, S., SHI, M., and VERLEGH, P., "The firm's management of social interactions," *Marketing Letters*, vol. 16, no. 3-4, pp. 415–428, 2005.

[79] GOLDMAN, A. I. and SRIPADA, C. S., "Simulationist models of face-based emotion recognition," *Cognition*, vol. 94, no. 3, pp. 193–213, 2005.

[80] GOLSTEIJN, C., VAN DEN HOVEN, E., FROHLICH, D., and SELLEN, A., "Towards a more cherishable digital object," in *Proceedings of the Designing Interactive Systems Conference*, pp. 655–664, ACM, 2012.

[81] GOODMAN, L. A., "Snowball sampling," *The Annals of Mathematical Statistics*, vol. 32, no. 1, pp. 148–170, 1961.

[82] GORN, G. J., CHATTOPADHYAY, A., YI, T., and DAHL, D. W., "Effects of color as an executional cue in advertising: they're in the shade," *Management Science*, vol. 43, no. 10, pp. 1387–1400, 1997.

[83] GRANOVETTER, M. S., "The strength of weak ties," *American journal of sociology*, pp. 1360–1380, 1973.

[84] GUILFORD, J. P., "The affective value of color as a function of hue, tint, and chroma.," *Journal of experimental psychology*, vol. 17, no. 3, p. 342, 1934.

[85] GUILFORD, J. P. and SMITH, P. C., "A system of color-preferences," *The American Journal of Psychology*, vol. 72, no. 4, pp. 487–502, 1959.

[86] GULOTTA, R., FASTE, H., and MANKOFF, J., "Curation, provocation, and digital identity: risks and motivations for sharing provocative images online," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 387–390, ACM, 2012.

[87] HÅKANSSON, M., GAYE, L., LJUNGBLAD, S., and HOLMQUIST, L. E., "More than meets the eye: an exploratory study of context photography," in *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*, pp. 262–271, ACM, 2006.

[88] HAXBY, J., HOFFMAN, E., and GOBBINI, M., "The distributed human neural system for face perception," *Trends in cognitive sciences*, vol. 4, no. 6, pp. 223–233, 2000.

[89] HEMPHILL, M., "A note on adults' color–emotion associations," *The Journal of genetic psychology*, vol. 157, no. 3, pp. 275–280, 1996.

[90] HILBE, J. M., *Negative binomial regression*. Cambridge University Press, 2011.

[91] HILL, R. A. and BARTON, R. A., "Psychology: red enhances human performance in contests," *Nature*, vol. 435, no. 7040, pp. 293–293, 2005.

[92] HILL, S., PROVOST, F., and VOLINSKY, C., "Network-based marketing: Identifying likely adopters via consumer networks," *Statistical Science*, pp. 256–276, 2006.

[93] 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.

[94] Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., and Wood, K., "Sensecam: A retrospective memory aid," in *UbiComp 2006: Ubiquitous Computing*, pp. 177–193, Springer, 2006.

[95] Homans, G. C., "Social behavior as exchange," *American journal of sociology*, pp. 597–606, 1958.

[96] Honeyman, M. K., "Vegetation and stress: a comparison study of varying amounts of vegetation in countryside and urban scenes," *The role of horticulture in human well-being and social development*, pp. 143–145, 1992.

[97] Hoogeveen, M., "Toward a theory of the effectiveness of multimedia systems," *International Journal of Human-Computer Interaction*, vol. 9, no. 2, pp. 151–168, 1997.

[98] Hsia, H., "The information processing capacity of modality and channel performance," *AV Communication Review*, vol. 19, no. 1, pp. 51–75, 1971.

[99] Huang, J., White, R. W., and Dumais, S., "No clicks, no problem: using cursor movements to understand and improve search," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1225–1234, ACM, 2011.

[100] Hupka, R. B., Zaleski, Z., Otto, J., Reidl, L., and Tarabrina, N. V., "The colors of anger, envy, fear, and jealousy a cross-cultural study," *Journal of Cross-Cultural Psychology*, vol. 28, no. 2, pp. 156–171, 1997.

[101] Hutchings, J., "Colour in folklore and traditionthe principles," *Color Research & Application*, vol. 29, no. 1, pp. 57–66, 2004.

[102] Hutto, C., Yardi, S., and Gilbert, E., "A longitudal study of follow predictors on twitter," in *Proceedings of the sigchi conference on human factors in computing systems*, pp. 821–830, acm, 2013.

[103] 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*, vol. 98, no. 24, pp. 13995–13999, 2001.

[104] Jacobs, L., Keown, C., Worthley, R., and Ghymn, K.-I., "Cross-cultural colour comparisons: global marketers beware!," *International Marketing Review*, vol. 8, no. 3, 1991.

[105] 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*, pp. 32–38, ieee, 2009.

[106] JAMES, W. T. and DOMINGOS, W. R., "The effect of color shock on motor performance and tremor," *The journal of general psychology*, vol. 48, no. 2, pp. 187–193, 1953.

[107] JENKINS, R. and BURTON, A. M., "100% accuracy in automatic face recognition," *Science*, vol. 319, no. 5862, p. 435, 2008.

[108] JOBES, G., *Dictionary of Mythology, Folklore and Symbols. Part 3. Index*. Scarecrow Press, 1962.

[109] JOHNSON, M. H., DZIURAWIEC, S., ELLIS, H., and MORTON, J., "Newborns' preferential tracking of face-like stimuli and its subsequent decline," *cognition*, vol. 40, no. 1, pp. 1–19, 1991.

[110] 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*, vol. 17, no. 11, pp. 4302–4311, 1997.

[111] KAYA, N. and EPPS, H. H., "Relationship between color and emotion: a study of college students," *College student journal*, vol. 38, pp. 396–405, 2004.

[112] KELSEY, R. E. and STIMSON, B., *The meaning of photography*. Clark Art Institute, 2008.

[113] KIRK, D., SELLEN, A., ROTHER, C., and WOOD, K., "Understanding photowork," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pp. 761–770, ACM, 2006.

[114] KNIGHT, C., POWER, C., and WATTS, I., "The human symbolic revolution: a darwinian account," *Cambridge Archaeological Journal*, vol. 5, no. 1, pp. 75–114, 1995.

[115] KRESS, G. and VAN LEEUWEN, T. V., "Multimodal discourse: The modes and media of contemporary communication," 2001.

[116] KRESS, G. R. and VAN LEEUWEN, T., *Reading images: The grammar of visual design*. Psychology Press, 1996.

[117] KWAK, H., LEE, C., PARK, H., and MOON, S., "What is twitter, a social network or a news media?," in *Proceedings of the 19th international conference on world wide web*, pp. 591–600, acm, 2010.

[118] LANEY, R., "Does orange mean cheap?," *Forbes*, vol. 148, no. 14, pp. 144–146, 1991.

[119] LANG, A., "Defining audio/video redundancy from a limited-capacity information processing perspective," *Communication Research*, vol. 22, no. 1, pp. 86–115, 1995.

[120] LANG, A., "The limited capacity model of mediated message processing," *Journal of Communication*, vol. 50, no. 1, pp. 46–70, 2000.

[121] LANG, J., *Creating architectural theory: The role of the behavioral sciences in environmental design*. Van Nostrand Reinhold New York, 1987.

[122] LAUREL, B., "Interface agents: Metaphors with character," *Human values and the design of computer technology*, pp. 207–219, 1997.

[123] LEE, S. and BARNES, J. H., "Using color preferences in magazine advertising.," *Journal of advertising research*, 1989.

[124] LEE, Y., *Man as the prayer: The origin and nature of humankind.* Trafford Publishing, 2000.

[125] LEHMANN, J., LALMAS, M., YOM-TOV, E., and DUPRET, G., "Models of user engagement," in *User Modeling, Adaptation, and Personalization*, pp. 164–175, Springer, 2012.

[126] LEHMANN, J., LALMAS, M., YOM-TOV, E., and DUPRET, G., "Models of user engagement," in *User Modeling, Adaptation, and Personalization*, pp. 164–175, Springer, 2012.

[127] 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, (New York, NY, USA), pp. 651–654, ACM, 2011.

[128] LEVY, M. R., "Watching tv news as para-social interaction," *Journal of Broadcasting & Electronic Media*, vol. 23, no. 1, pp. 69–80, 1979.

[129] LIN, C. and FASTE, H., "Photographic social media: a framework for design," in *Proceedings of the Designing Interactive Systems Conference*, pp. 244–247, ACM, 2012.

[130] LINDLEY, S. E., RANDALL, D., SHARROCK, W., GLANCY, M., SMYTH, N., and HARPER, R., "Narrative, memory and practice: tensions and choices in the use of a digital artefact," in *Proceedings of the 23rd British HCI group annual conference on people and computers: celebrating people and technology*, pp. 1–9, British Computer Society, 2009.

[131] LIPPMANN, W., "The world outside and the pictures in our heads," *Public opinion*, vol. 4, pp. 1–22, 1922.

[132] MAHNKE, F. H., *Color, environment, and human response: an interdisciplinary understanding of color and its use as a beneficial element in the design of the architectural environment.* Wiley, 1996.

[133] MANNING, T., "Information environments for breast cancer education," *Health Promotion and Interactive Technology: Theoretical Applications and Future Directions*, p. 121, 2013.

[134] MARMOLIN, H., "Multimedia from the perspectives of psychology," in *Multimedia*, pp. 39–52, Springer, 1992.

[135] McDONALD, D. W., "Visual conversation styles in web communities," in *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*, pp. 76–76, IEEE, 2007.

[136] McMANUS, I. and JONES, A., "The aesthetics of color," *Perception*, vol. 10, no. 6, pp. 651–666, 1981.

[137] MEEHAN, J., *The Photographer's Guide to Using Filters*. Amphoto Books, October 1998.

[138] MEHRABIAN, A. and RUSSELL, J. A., "The basic emotional impact of environments," *Perceptual and Motor Skills*, vol. 38, no. 1, pp. 283–301, 1974.

[139] MEHTA, R. and ZHU, R. J., "Blue or red? exploring the effect of color on cognitive task performances," *Science*, vol. 323, no. 5918, pp. 1226–1229, 2009.

[140] MEYERS-LEVY, J. and PERACCHIO, L. A., "How the use of color in advertising affects attitudes: the influence of processing motivation and cognitive demands," *Journal of consumer research*, vol. 22, pp. 121–38, 1995.

[141] MEYERS-LEVY, J. and PERACCHIO, L. A., "Understanding the effects of color: How the correspondence between available and required resources affects attitudes," *Journal of Consumer Research*, pp. 121–138, 1995.

[142] MIDDLESTADT, S. E., "The effect of background and ambient color on product attitudes and beliefs," *Advances in consumer research*, vol. 17, no. 1, pp. 244–249, 1990.

[143] MIKELLIDES, B., "Color and physiological arousal.," *Journal of Architectural and Planning Research*, 1990.

[144] 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*, pp. 306–313, ACM, 2002.

[145] MILLER, M., "Cashing in on rich color schemes," *Business mexico*, vol. 4, no. 10, pp. 28–32, 1994.

[146] MITCHELL, C., *Doing visual research*. Sage, 2011.

[147] MOORE, S. G., *Some Things are Better Left Unsaid: How Word of Mouth Influences the Speaker*. PhD thesis, Duke University, 2009.

[148] MORTON, J. and JOHNSON, M. H., "Conspec and conlern: a two-process theory of infant face recognition.," *Psychological review*, vol. 98, no. 2, p. 164, 1991.

[149] MUNSELL, A. H., "A pigment color system and notation," *The American Journal of Psychology*, vol. 23, no. 2, pp. 236–244, 1912.

[150] NETO, F., "Colors associated with styles of love," *Perceptual and motor skills*, vol. 94, no. 3c, pp. 1303–1310, 2002.

[151] NICOLSON, R. I., PICKERING, S., and FAWCETT, A. J., "A hypercard spelling support environment for dyslexic children," *Computers & Education*, vol. 16, no. 2, pp. 203–209, 1991.

[152] O'BRIEN, H. L. and TOMS, E. G., "The development and evaluation of a survey to measure user engagement," *Journal of the American Society for Information Science and Technology*, vol. 61, no. 1, pp. 50–69, 2010.

[153] ODOM, W., PIERCE, J., STOLTERMAN, E., and BLEVIS, E., "Understanding why we preserve some things and discard others in the context of interaction design," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1053–1062, ACM, 2009.

[154] O'HARA, K., HELMES, J., SELLEN, A., HARPER, R., TEN BHÖMER, M., and VAN DEN HOVEN, E., "Food for talk: phototalk in the context of sharing a meal," *Human–Computer Interaction*, vol. 27, no. 1-2, pp. 124–150, 2012.

[155] OLMSTED, F., "The value and care of parks," in *Report to the Congress of the State of California (Reprinted in Landscape Architecture, 17: 20-23)*, 1865.

[156] OLNEY, T. J., HOLBROOK, M. B., and BATRA, R., "Consumer responses to advertising: The effects of ad content, emotions, and attitude toward the ad on viewing time," *Journal of Consumer Research*, pp. 440–453, 1991.

[157] OLSON, J. S., OLSON, G. M., and MEADER, D. K., "What mix of video and audio is useful for small groups doing remote real-time design work?," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 362–368, ACM Press/Addison-Wesley Publishing Co., 1995.

[158] OTT, J. N., "Influence of fluorescent lights on hyperactivity and learning disabilities," *Journal of learning disabilities*, vol. 9, no. 7, pp. 417–422, 1976.

[159] PAIVIO, A., *Mental representations*. Oxford University Press, 1990.

[160] PALLAK, S. R., "Salience of a communicator's physical attractiveness and persuasion: A heuristic versus systematic processing interpretation," *Social Cognition*, vol. 2, no. 2, pp. 158–170, 1983.

[161] PENNEY, C. G., "Modality effects in short-term verbal memory.," *Psychological Bulletin*, vol. 82, no. 1, p. 68, 1975.

[162] PETERS, K. and KASHIMA, Y., "From social talk to social action: shaping the social triad with emotion sharing.," *Journal of Personality and Social Psychology*, vol. 93, no. 5, p. 780, 2007.

[163] PETERSON, E. T. and CARRABIS, J., "Measuring the immeasurable: Visitor engagement," *Web Analytics Demystified*, 2008.

[164] PETRELLI, D. and WHITTAKER, S., "Family memories in the home: contrasting physical and digital mementos," *Personal and Ubiquitous Computing*, vol. 14, no. 2, pp. 153–169, 2010.

[165] PETTY, R. E., CACIOPPO, J. T., SCHUMANN, D., and OTHERS, "Central and peripheral routes to advertising effectiveness: The moderating role of involvement," *Journal of consumer research*, vol. 10, no. 2, p. 135, 1983.

[166] 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.

[167] PLODERER, B., LEONG, T., ASHKANASY, S., and HOWARD, S., "A process of engagement: engaging with the process," in *Proceedings of the Designing Interactive Systems Conference*, pp. 224–233, ACM, 2012.

[168] POLLAK, J. P., ADAMS, P., and GAY, G., "Pam: a photographic affect meter for frequent, in situ measurement of affect," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 725–734, ACM, 2011.

[169] RAMIREZ, A., "Lessons in the cracker market," *The new york times, july*, vol. 5, pp. 1–3, 1990.

[170] REESE, S. D., "Visual-verbal redundancy effects on television news learning," *Journal of Broadcasting & Electronic Media*, vol. 28, no. 1, pp. 79–87, 1984.

[171] REEVES, B. and NASS, C., *The Media equation: how people treat computers, television, and new media.* Cambridge University Press, 1997.

[172] RICHARDSON, M. and DOMINGOS, P., "Mining knowledge-sharing sites for viral marketing," in *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 61–70, ACM, 2002.

[173] RIME, B., MESQUITA, B., BOCA, S., and PHILIPPOT, P., "Beyond the emotional event: Six studies on the social sharing of emotion," *Cognition & Emotion*, vol. 5, no. 5-6, pp. 435–465, 1991.

[174] ROSE, G., *Visual methodologies: An introduction to researching with visual materials.* Sage, 2012.

[175] ROSNER, D. K. and RYOKAI, K., "Spyn: augmenting the creative and communicative potential of craft," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2407–2416, ACM, 2010.

[176] SALTARELLI, C. and COPPOLA, C., "Influence of visible light on organ weights of mice.," *Laboratory animal science*, vol. 29, no. 3, pp. 319–322, 1979.

[177] SANBONMATSU, D. M. and KARDES, F. R., "The effects of physiological arousal on information processing and persuasion," *Journal of Consumer Research*, pp. 379–385, 1988.

[178] SAURO, J. and DUMAS, J. S., "Comparison of three one-question, post-task usability questionnaires," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1599–1608, ACM, 2009.

[179] SCHAUSS, A. G., "Tranquilizing effect of color reduces aggressive behavior and potential violence," *Journal of Orthomolecular Psychiatry*, vol. 8, no. 4, pp. 218–221, 1979.

[180] SCHINDLER, P. S., "Color and contrast in magazine advertising," *Psychology and marketing*, vol. 3, no. 2, pp. 69–78, 1986.

[181] SCHMUTZ, P., HEINZ, S., MÉTRAILLER, Y., and OPWIS, K., "Cognitive load in ecommerce applications: measurement and effects on user satisfaction," *Advances in Human-Computer Interaction*, vol. 2009, p. 3, 2009.

[182] SCHOFIELD, J., *The Darkroom book: the comprehensive step-by-step guide to processing your colour or black-and-white photographs.* Watson-Guptill Pubns, 1985.

[183] SELLEN, A. J., "Remote conversations: The effects of mediating talk with technology," *Human-computer interaction*, vol. 10, no. 4, pp. 401–444, 1995.

[184] SERGENT, J., OHTA, S., and MDONALD, B., "Functional neuroanatomy of face and object processing: a position emission tomography study," *Brain*, vol. 115, no. 1, pp. 15–36, 1992.

[185] 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*, (New York, NY, USA), pp. 275–284, ACM, 2007.

[186] SHARPE, D. T., *The psychology of color and design.* Littlefield, Adams New Jersey, 1975.

[187] SHORT, J., WILLIAMS, E., and CHRISTIE, B., "The social psychology of telecommunications," 1976.

[188] SLATER, M., "Measuring presence: A response to the witmer and singer presence questionnaire," *Presence: Teleoperators and Virtual Environments*, vol. 8, no. 5, pp. 560–565, 1999.

[189] SMITH, K. C. P. and APTER, M. J., *A theory of psychological reversals.* Picton Pub., 1975.

[190] SPARKMAN JR, R. and AUSTIN, L. M., "The effect on sales of color in newspaper advertisements," *Journal of Advertising*, pp. 39–42, 1980.

[191] SPROULL, L., SUBRAMANI, M., KIESLER, S., WALKER, J. H., and WATERS, K., "When the interface is a face," *Human-Computer Interaction*, vol. 11, no. 2, pp. 97–124, 1996.

[192] STEUER, J. S., "Vividness and source of evaluation as determinants of social responses toward mediated representations of agency," 1995.

[193] STURKEN, M., CARTWRIGHT, L., and STURKEN, M., *Practices of looking: An introduction to visual culture.* Oxford University Press Oxford, 2001.

[194] 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)*, pp. 177–184, 2010.

[195] SUN, E., ROSENN, I., MARLOW, C., and LENTO, T., "Modeling contagion through facebook news feed," in *Proceedings of international aaai conference on weblogs and social media*, p. 22, 2009.

[196] SUNDAR, S. S., "Multimedia effects on processing and perception of online news: A study of picture, audio, and video downloads," *Journalism & Mass Communication Quarterly*, vol. 77, no. 3, pp. 480–499, 2000.

[197] 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*, pp. 187–193, ACM, 1993.

[198] Takeuchi, a. and naito, t., "Situated facial displays: towards social interaction," in *Proceedings of the sigchi conference on human factors in computing systems*, pp. 450–455, acm press/addison-wesley publishing co., 1995.

[199] Tang, J. C., Isaacs, E. A., and Rua, M., "Supporting distributed groups with a montage of lightweight interactions," in *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pp. 23–34, ACM, 1994.

[200] Tom, G., Barnett, T., Lew, W., and Selmants, J., "Cueing the consumer: the role of salient cues in consumer perception," *Journal of Consumer Marketing*, vol. 4, no. 2, pp. 23–27, 1987.

[201] Ulrich, R. S., "Visual landscapes and psychological well-being," *Landscape research*, vol. 4, no. 1, pp. 17–23, 1979.

[202] Ulrich, R. S., "Natural versus urban scenes some psychophysiological effects," *Environment and behavior*, vol. 13, no. 5, pp. 523–556, 1981.

[203] Ulrich, R. S., "Aesthetic and affective response to natural environment," in *Behavior and the natural environment*, pp. 85–125, Springer, 1983.

[204] Ulrich, R. S., "Effects of interior design on wellness: theory and recent scientific research," *Journal of health care interior design*, vol. 3, no. 1, pp. 97–109, 1991.

[205] Ulrich, R. S., Simons, R. F., Losito, B. D., Fiorito, E., Miles, M. A., and Zelson, M., "Stress recovery during exposure to natural and urban environments," *Journal of environmental psychology*, vol. 11, no. 3, pp. 201–230, 1991.

[206] Unnava, H. R., Burnkrant, R. E., and Erevelles, S., "Effects of presentation order and communication modality on recall and attitude," *Journal of Consumer Research*, pp. 481–490, 1994.

[207] Valdez, P. and Mehrabian, A., "Effects of color on emotions.," *Journal of experimental psychology: general*, vol. 123, no. 4, p. 394, 1994.

[208] Van Leeuwen, T. and Jewitt, C., *The handbook of visual analysis*. Sage, 2001.

[209] 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*, vol. 68, no. 10, pp. 719–728, 2010.

[210] Vedaldi, A. and Fulkerson, B., "Vlfeat: An open and portable library of computer vision algorithms," in *Proceedings of the international conference on Multimedia*, pp. 1469–1472, ACM, 2010.

[211] Viola, p. and jones, m. j., "Robust real-time face detection," *International journal of computer vision*, vol. 57, no. 2, pp. 137–154, 2004.

[212] Von Bergen, J., "What's your favorite color," *Calgary, H. C.*, vol. 1, 1995.

[213] 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*, pp. 85–91, ACM, 1994.

[214] WALTERS, J., APTER, M. J., and SVEBAK, S., "Color preference, arousal, and the theory of psychological reversals," *Motivation and emotion*, vol. 6, no. 3, pp. 193–215, 1982.

[215] WATTS, L. and MONK, A. F., "Remote assistance: a view of the work and a view of the face?," in *Conference Companion on Human Factors in Computing Systems*, pp. 101–102, ACM, 1996.

[216] WEXNER, L. B., "The degree to which colors (hues) are associated with mood-tones.," *Journal of applied psychology*, vol. 38, no. 6, p. 432, 1954.

[217] WICKENS, C. D., "Processing resources and attention," *Multiple-task performance*, pp. 3–34, 1991.

[218] WOJNICKI, A. and GODES, D., "Word-of-mouth as self-enhancement," *HBS marketing research paper*, no. 06-01, 2008.

[219] 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*, vol. 31, no. 2, pp. 210–227, 2009.

[220] YANG, S. and ALLENBY, G. M., "Modeling interdependent consumer preferences," *Journal of Marketing Research*, pp. 282–294, 2003.

[221] 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.

[222] YILDIZ, R. and ATKINS, M., "Evaluating multimedia applications," *Computers & Education*, vol. 21, no. 1, pp. 133–139, 1993.

[223] YOUNG, A., MCWEENY, K., HAY, D., and ELLIS, A., "Matching familiar and unfamiliar faces on identity and expression," *Psychological research*, vol. 48, no. 2, pp. 63–68, 1986.