

Selfie-Presentation in Everyday Life: A Large-scale Characterization of Selfie Contexts on Instagram

Julia Deeb-Swihart, Christopher Polack, Eric Gilbert and Irfan Essa

Georgia Institute of Technology

North Ave NW

Atlanta, GA 30332

{jdeeb3, cfpolack, gilbert, irfan}@gatech.edu

Abstract

Carefully managing the presentation of self via technology is a core practice on all modern social media platforms. Recently, *selfies* have emerged as a new, pervasive genre of identity performance. In many ways unique, selfies bring us full-circle to Goffman—blending the online and offline selves together. In this paper, we take an empirical, Goffman-inspired look at the phenomenon of selfies. We report a large-scale, mixed-method analysis of the categories in which selfies appear on Instagram—an online community comprising over 400M people. Applying computer vision and network analysis techniques to 2.5M selfies, we present a typology of emergent selfie categories which represent emphasized identity statements. To the best of our knowledge, this is the first large-scale, empirical research on selfies. We conclude, contrary to common portrayals in the press, that selfies are really quite ordinary: they project identity signals such as wealth, health and physical attractiveness common to many online media, and to offline life.

Introduction

Recently, a Thai photographer named Chompoo Baritone unveiled a series of photos revealing the “messy reality we usually crop from our social media shots” (Merelli 2015). In one piece, we see a woman holding up a friend by her feet; yet in the version shared to Instagram, the supportive friend is cropped out, making it appear as if the subject does a handstand by herself. In another, the Instagram photo shows a Macbook surrounded by a few tasteful knick knacks; outside the view of the camera, however, we see that the room is a mess, with clothes strewn everywhere and the bed unmade.

While clearly not meant as a comprehensive portrayal of Instagram, the art project cleverly showcases a core practice underlying all social media platforms: carefully managing the *presentation of self* via technology (e.g., (Donath and others 1999; Farnham and Churchill 2011; Newman et al. 2011)). It is important, in part, because we make judgments of other people—sometimes quick ones (Sunnafrank and Ramirez 2004)—based on these portrayals (Donath and others 1999; Shami et al. 2009; Smith and Collins 2009). The practice takes different forms on different platforms. On

Facebook, we curate profile elements to draw a unique portrait of ourselves (i.e., our hometowns, the bands we like, etc.) (Lampe, Ellison, and Steinfield 2007), in addition to the identity work that goes on through status updates (Silfverberg, Liikkanen, and Lampinen 2011). On Twitter, we may signal identity through the links we choose to share. On Instagram, we portray the self we want to share (and perhaps want to *be*) through the images we take. While at times this practice may appear thoroughly contemporary, upon reflection we can see it everywhere in offline life as well—most famously described in Erving Goffman’s *The Presentation of Self in Everyday Life* (Goffman and others 1959). From the clothes we choose to put on in the morning, to social roles we inhabit when asked, we continually construct and control the version of ourselves that others see. In Goffman’s words, “the world, in truth, is a wedding,” with its concomitant costumes, rituals and roles—all acting to shape how others perceive us (Goffman and others 1959).

Amid this mix of online self-presentation practices, *selfies* have emerged as a new, pervasive genre. The Oxford Dictionary’s “Word of the Year” in 2013 (Brumfield 2013), a selfie is “a photograph that one has taken of oneself, typically taken with a smartphone or webcam and uploaded to a social media website.” Unlike many earlier forms of identity work, however, selfies are particularly “sociomaterial” (Svelander and Wiberg 2015) in that they emerge at the intersection of our social worlds (i.e., Instagram and similar sites) and the unique affordances of a smartphone’s camera. Unlike cameras that came before it, with a smartphone’s front-facing camera we can easily place ourselves in the frame, thereby carefully crafting the contexts in which people see and come to understand us. By placing ourselves, our bodies, in the photos we take and share via social media, we come full-circle to Goffman, in some sense blending the online and offline selves together in this new practice.

In this paper, we adopt a Goffman-inspired perspective on the phenomenon of selfies: we conduct a large-scale, mixed-method analysis of the emphasized identity statement categories in which people take and post pictures of themselves on Instagram—an online community with over 400 million users (Instagram 2015). (Throughout this paper, we will refer to a selfie’s emphasized identity statement categories as both the setting in which the person appears, as well as personal attributes present in the photo.) The primary contribu-

tion of this work is a typology of the emergent categories in which people post selfies on Instagram—and by its exhaustiveness, the categories where selfies *do not appear*. Specifically, we collected 2.5 million Instagram photos over a three-month period with the associated tag #selfie. Applying computer vision techniques, we eliminated those photos that did not contain faces (surprisingly, nearly half did not). Next, we applied a network analysis technique known as *community detection* (Peixoto 2014) to uncover emergent contexts as revealed by associations between tags. With this approach, for example, we discovered a popular context centered around showing off luxury goods in a fashion setting, comprising tags like #diamond, #armani and #king. Then, we are able to infer (again, with computer vision) the distribution over age and gender by selfie context; we find, for example, that women are more likely to take selfies that project a healthy lifestyle, but that men and women post travel selfies in equal numbers.

We believe this is the first empirical, scholarly work on selfies, and among the first in a new thread of social multimedia scholarship (e.g., (Abdullah et al. 2015; Bakhshi, Shamma, and Gilbert 2014)). We see the implications of our work as speaking to social media research, but also a long lineage of identity-centric social science. This work primarily addresses the “What?” and “Who?” questions surrounding selfies; however, it may enable others to answer the “How?”, “Where?”, and “Why?” questions raised by this study.

Literature Review

Next, we explore the related literature informing the present work, centered around *Identity* and *Instagram*. The former provides greater theoretical depth for interpreting our findings, while the latter concentrates on contemporary work profiling and examining Instagram, as well as similar photo-sharing social network sites.

Identity, Online and Offline

As discussed above, Goffman’s self-presentation work forms the central axis around which much of this work revolves (Goffman and others 1959). While it is impossible to summarize all the claims made in *The Presentation of Self in Everyday Life*, one of the most central and relevant is the idea that people inhabit roles (much as if they were performers in a play), and ask that others believe those portrayals. In a well-known passage, Goffman defines the materials and methods used to construct these portrayals, known as the “front:”

That part of individual’s performance which regularly functions in a general and fixed fashion to define the situation for those who observe the performance. Front, then, is the expressive equipment of a standard kind intentionally or unwittingly employed by the individual during his performance. (p. 22)

Here, Goffman works to set up the dramatic metaphor structuring his work: there are settings (props, locations, scenery), clothes, and appearances that function within the scene being acted out. For the purposes of this work, we

might ask ourselves: What “expressive equipment” do Instagram users have at their disposal when they post selfies? Elements such as the bounding box of the camera’s viewpoint, the location from which they post the photo (perhaps GPS tagged), the clothes in which they appear (or, some cases, those clothes they choose to leave off), attributes of the body given particular prominence (i.e., tattoos, duck face, etc.), all spring to mind. An essential idea inherited here from Goffman is that people willfully and purposefully control and craft these elements; we make heavy use of these ideas in both the construction of our dataset (i.e., what tags we choose to include) and in the interpretation of our findings.

It is important to point out that Goffman has been hugely influential in social computing, and HCI more broadly. At the time of this writing, for instance, more than 400 articles from SIGCHI-related conferences invoke Goffman (by name) to somehow inform their work¹. A necessarily brief tour, Goffman’s concepts of identity work have informed scholarship ranging from expertise location (Shami et al. 2009), to identity fragmented across multiple platforms (Farnham and Churchill 2011), to health information-seeking behavior (Newman et al. 2011).

The present work builds on this thread. It is an empiricist’s take on selfies through the lens of Goffman: we seek to understand, quantitatively and at large-scale, the contexts in which people portray the characters they want to be.

Signalling Theory. Given that people so carefully architect their identities in both offline and online settings, we must assume that they sometimes lie. But how often and about what? The core idea of *signalling theory* is that identity signals are differentially costly to fake, and therefore differentially reliable (Donath and others 1999). This is most easily grasped through an example. In Donath’s work, for instance, she discusses that goal of signalling oneself as a strong person. Which signal most convincingly demonstrates that: wearing a Gold’s Gym t-shirt, or lifting something heavy? Clearly, the latter *costs more* to perform (i.e., all the time required in the gym beforehand), and is therefore more reliable. The former is, by comparison, much easier to fake.

In the context of the present work, signalling theory might inform how we reason about the emergent categories in which we find selfies. For example, assuming that a person wants to project *wealth*, which signal is more reliable: the photo taken inside a Lexus, or the one taken standing next to it? Clearly the former, as anyone could walk up to a car in a parking lot and pose with it. We will invoke some of the concepts as we interpret the categories we discover.

Instagram

Next, we turn to reviewing the emerging body of work looking at Instagram from multiple angles. The photo-sharing social network site Instagram, which has grown to over 400M users (Instagram 2015), is the subject of a number of studies looking at social network structure. For example, there is significant spatial and temporal data available on Instagram that can be analyzed in order to create profiles

¹<http://dl.acm.org/results.cfm?query=goffmandl=GUIDE&dimgroup=5dim=3206>

of the habits, culture, points of interest, and photographic trends of a given area (Hochman and Manovich 2013; Hochman and Schwartz 2012; Silva et al. 2013) or to determine the strength of offline social group ties (Scellato et al. 2011). On a smaller scale, observing the patterns of individual users provides insight into why people use Instagram. As might be expected from such a large online community, Instagram usage can vary greatly across demographic groups: for example, teenagers are more likely to post a selfie than adults (Jang et al. 2015). Despite these differences, it is possible to create generalized types of users by clustering posts into categories including *fashion* and *food* (Hu et al. 2014).

Particularly relevant to work presented here, the presence of a face in an Instagram photo dramatically increases its likelihood of receiving likes and comments: posts with faces see a boost of more than 38% and 32% percent for likes and comments respectively (Bakhshi, Shamma, and Gilbert 2014). Selfiecity, one of the few existing projects to look at selfies, is an Instagram selfie visualization project (Manovich et al. 2014). It collected Instagram photos from six major cities across the world, analyzed various attributes of selfie posts—including demographics, pose, features, and mood. The selfie has also been used as a tool for personality prediction (Qiu et al. 2015; Guntuku et al. 2015).

To the best of our knowledge, the present work is the first empirical research on selfies, and serves to broaden what we know about practices on Instagram.

Methods and Data

For this study, we are concerned with finding what identity statements users are emphasizing through the medium of selfies. Once these identity statement categories are discovered, we also work to evaluate the believability of the performances in each category. Thus we structure our methodology to first empirically discover the high-level selfie categories and then design a study to evaluate the perceptions surrounding those categories.

To begin this section, we will first describe the data we collected from Instagram, how we detected age and gender information, and how we constructed a typology of selfies (see Figure 2 for an overview). For the purposes of this paper, we define a selfie as an image tagged #selfie that also contains at least one human face. Note that it is not a per-

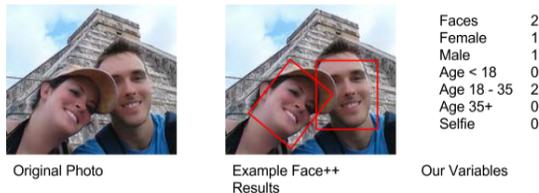


Figure 1: Example of how we construct our variables from the facial recognition from Face++. Image: 📷👤📷 Kyla Heineman on Flickr.

fect sample—as certainly many actual selfies do not carry the #selfie tag—but this seemed to us a reasonable tradeoff given the constraints of the Instagram and computer vision APIs (e.g., search, rate limits). Finally, we distinguish between the traditional selfie, which contains only one face, and the group selfie, which contains two or more faces.

We chose Instagram as the site of study for three reasons: first because of its widespread, cross-cultural usage; second, its large collection of selfies (as of writing this paper, Instagram reports over 200 million posts tagged #selfie); and third, unlike many other photo-sharing sites, Instagram has a publicly available, documented API.

Data Collection

Over a period of three months, we scraped roughly 2.5 million public posts from Instagram that had been tagged by the user as #selfie. For this study, we only consider images that are posted publicly and do not consider videos. In our dataset for this study, a post contains an image URL and an associate tag-list. However, we also collected additional information concerning each post such as user profile information, date of post, and if available, the GPS information of the post.

The data represents posts from the following 12 days: June 29–30, July 5–7, and July 11–17. Collection was restarted 4 different times throughout the process to collect a larger variety of data that includes all weekdays and weekends. It is important to note that the following holidays occurred during data collection and were specifically mentioned in the dataset: Ramadan (June 18–July 17) and American Independence Day (July 4).

Facial Recognition

After we collected the dataset, we examined the posts collected and noticed that a high number of posts contained spam-related images (such as blank images, or images with text asking for followers). Thus we filtered the dataset automatically using facial recognition and removed posts from the dataset that did not contain at-least one face.

Facial recognition from images is a widely studied topic in computer vision (Zhu and Ramanan 2012). The current state of the art in facial recognition demonstrates high accuracy and is known to even exceed human performance on similar tasks (Lu and Tang 2014). Facial recognition problems are constructed as follows: given an image of a scene, identify if the image contains a face - and if it does, where the face is positioned within the image.

We used the publicly available API developed by Face++, a cloud-based facial recognition system, to first filter out images collected that do not contain a face. Given the popularity of the tag #selfie, part of the dataset represents spam posts (e.g., images of text asking for followers or likes, images of products for sale, etc.) or unrelated images (e.g., images of just pets, pictures of body parts, etc.). Our dataset is too large to manually filter out non-selfies; thus, Face++ provides us a highly accurate and feasible alternative for filtering.

Face++ provides a service that accepts the URL of an Instagram photo, and returns whether or not the photo contains a face, and if so, information concerning the number of

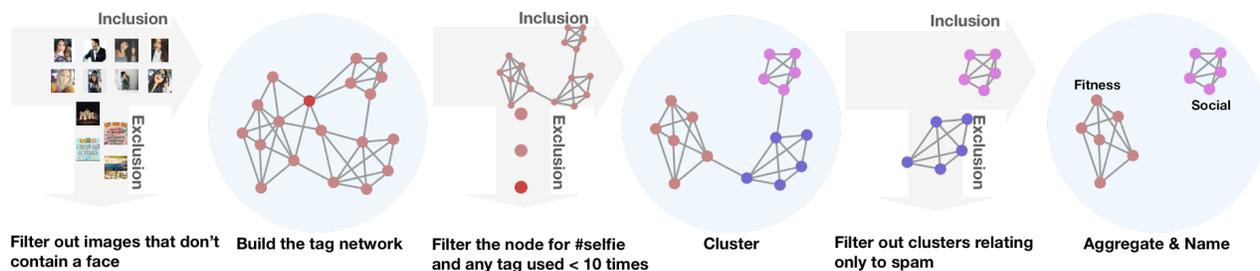


Figure 2: An overview of the steps taken to collect, process, and analyze the data used in this work.

faces in the photo, demographic information of the individuals in the photos, and the position and size of each face in the image. Face++ reports a 99.5% accuracy on an established facial recognition benchmark (Zhou, Cao, and Yin 2015); this accuracy is further supported by the results in (Bakhshi, Shamma, and Gilbert 2014), which reports a 97% +/- 5% accuracy on similar Instagram photos.

For the images in the dataset which contain faces, we quantify the demographic and facial results from Face++ into a binary feature vector concerning the gender, age group, and type of post. We discretize the age ranges as younger than 18 (Minor), 18-35 (Young Adult), and 35 or higher (Adult), the gender as male or female, and the post type as selfie or group selfie. We also extract information regarding the percentage of the image each face occupies (this is represented as a percentage of pixels within the image frame that represents the face). An example of how we constructed these variables can be seen in Figure 1.

Network Construction and Clustering

We approached the problem of quantitatively detecting selfie contexts as a community detection or network clustering problem. Community detection problems are typically formulated as follows: given a network (represented with a set of nodes and a list of edges that form between pairs of nodes), divide the nodes into groupings where the edges within a grouping are much denser than edges between groupings. We hypothesized that the groupings would correspond to identity contexts related to selfies, something we verify and present later in this section. By modeling the relationships between the tags, we can cluster the tags based on co-occurrence and thus discover the identity contexts within the data itself related to the term *selfie*.

For our purposes, we represent the filtered dataset as a tag-relationship network where each node is a unique tag (not including the tag #selfie, which is part of every image) and each edge represents a co-occurrence relationship between two tags. We define two tags as co-occurring if both appear in the tag list for the same image. The edges in the network are weighted by the Jaccard Coefficient, a metric used to measure the association between two tags (Frakes 1992).

Once constructed, the network models 999,901 unique tags and 29 million relationships between them. To reduce noise, we filter the network and remove nodes that represent tags used in 10 or fewer posts. This reduces the network

to 64,782 unique tags and 13 million tag relationships. As noted before parenthetically, we also remove the node associated to the tag #selfie so that the underlying structure can be detected.

We then cluster the network using a variant of the community detection called the stochastic blockmodel method (Peixoto 2014) to detect the underlying community structure. This results in 489 clusters spanning a variety of topics such as weightlifting, drug usage, and fashion. We chose the stochastic blockmodel because this particular model allowed us to find groupings of arbitrary size and number. This method, using likelihood estimation, discovers the number of groupings from the data, and thus we did not have to specify the total number of groupings beforehand, an important methodological advantage.

Cluster Curation and Evaluation

Our overarching goal in this paper is to investigate how people project an identity online through the medium of selfies. However due to the widespread usage of #selfie on Instagram, posts returned to us in this dataset were highly noisy. Even if the photo contained a face, the image might be a meme or an advertisement for a product.

We also discovered that the communities detected by the algorithm sometimes represented smaller niche communities. For example, one cluster had the following tags: #muscles, #beast, #dedication, #bodybuilding. Whereas another separate cluster had the following tags: #interval, #running, #instarun. The first cluster uses tags discussing bodybuilding and the other running. Despite their similarity, these two clusters are returned to us as separate because those two clusters represent different communities on Instagram that use different language. However, at a global level both do concern the same overarching theme of fitness.

Thus we turn to using qualitative coding processes to overcome the algorithmic shortcomings of the automatic clustering and filtering processes.

To do this, we first began by using a process referred to as inductive, open coding. Two researchers went through the clusters and through an iterative process generated a list of category labels that represent high level behavior patterns. Initially, we discovered 33 different categories - however due to the high amount of overlap between the categories we iteratively paired down that list to the 16 categories presented in the results section.

One of these 16 categories represented clusters to be filtered from the dataset. Given the framework through which we are analyzing these contexts, we filtered out clusters that do not concern identity or behavior characteristics. These filtered clusters comprised spam related tags - such as those relating to follower and like requests. Unfortunately due to the language restrictions of the researchers in this study, we also filtered out clusters with primarily non-English tags. The filtered clusters comprised 17.15% of the dataset.

Once the codebook was established, two volunteers separately and independently coded the dataset into these 16 categories where the first category represented clusters to be filtered from the dataset. These two researchers established a Cohen's κ of 79.16%, regarded as high agreement in the social sciences literature (Viera, Garrett, and others 2005). We used a third researcher to break ties. Results from this process can be seen in Figure 3, which visualizes the emphasized identity statements along with representative clusters. With the spam related cluster removed, the dataset consisted of close to 1.4 million images.

Assigning Cluster Membership

Once the clusters are established, we classify each post in the dataset as belonging to a particular cluster (as tags belong to clusters, not posts). To do this, we use the tag list accompanying a post and allow each tag in the list to vote for membership to a particular cluster. The post is then assigned to the cluster with the highest number of votes. In the cases of ties between clusters, a final label is assigned from the list of ties randomly. The demographic information about each post is then summarized at the cluster level - results of which can be seen in Table 1.

Results

In this section we highlight all the emphasized identity statements made through selfies as well as analyze the reliability of each of these signals. As part of the analysis we present a table of corresponding demographic information (Table 1). This demographic information helps identify various aspects of the personal front shown in a selfie in a typical context.

Our results in this section answers the *what* identity statements individuals make online through the medium of selfies. We use the framework established by Goffman for our analysis in this section.

We discover 15 emphasized identity statement categories of selfies, which we will now give a high level overview of. Examples of the clusters that form these categories can be best seen in Figure 3. The *Appearances* category contains posts concerning physical appearance of an individual, the clothing and decoration they chose to wear, and outward descriptions of an individuals status and wealth. The *Social* category describes posts concerning social events and interacting with friends, family, significant others, and pets. The *Health & Fitness* category describes aspects of maintaining overall personal health such as weight loss, positive mental attitude, fitness, and healthy diet. *Ethnicity* concerns posts describing the person's ethnic or national identity. The *Travel* category concerns posts about travel related activities and locations a person is traveling in. *Hobbies* refers

to hobby and pastime related posts including gaming, cosplay, art, music, and motorcycles. The *Teen/Young Adult* category refers to posts concerning issues specific teenagers and young adults as well as specific mentions of the age group. The *Food* category contains posts referring to specific food items. *Gender & Sexuality* contains posts describing one's gender identity or posts concerning expressions of sexual orientation. The *Celebrity & Entertainment Industry* describes posts relating to those who work in the entertainment industry as well as those who are fans of members of the industry. *Drugs* refers to posts about drug usage especially concerning marijuana usage. The *Japanophile* category refers to posts talking about the sub-culture (this sub-culture is also sometime referred to within the community as the otaku sub-culture). Similarly the *Alternative Culture* posts refer to the alternative sub-cultures including goth and punk. Finally, the *Work* category contains posts taken in the work environment.

From the demographic information, we can see that in our dataset, selfies are posted primarily by young adult women. Though this is influenced heavily by the demographics of Instagram itself which from a recent Pew Internet study found that the majority of users were between the ages of 18-19 and that the percentage of women outnumbered men (Dugan 2015).

Though it is important to note that our reported demographic information is limited to reporting binary genders since it uses the facial characteristics to report gender. In our findings, the gender and sexuality category deals with clusters associated with gender and sexual identity as well as activism thereof (*#prideparade*, *#pansexualpride*, *#translivesmatter*). This includes clusters concerning gender minorities (*#trans*, *#nonbinary*, *#genderfluid*, *#agender*) which means that our reported demographics for this category are inaccurate.

Four categories had especially high number of selfies (between 81% and 85%) as compared to other context families - indicating that posts primarily concern the poster them-self versus a group of individuals. These categories were *Alternative Culture*, *Gender & Sexuality*, *Japanophile*, and *Hobbies*.

By and large, the appearance category represented the largest emphasis for an identity statements, comprising 51.75% of posts in the dataset. This perhaps is not surprising given that a persons face represents the focus point of most of the images in the dataset, and thus the appearance of the face would therefore be a strong focal point for making identity statements. This is further supported by the idea that appearance is major part of a performance according to Goffman who refers to these aspects as the personal front (Goffman and others 1959). We also see as part of this personal front, users posing and drawing attention to luxury brands (*#armani*, *#alexandermcqueen*, *#maccosmetics*, *#revlon*, *#sephora*, *#loreal*) as a way of establishing rank and prestige (see Figure 4 for examples).

The appearance category also introduces a specific selfie posing behavior pattern referred to as "carfies" or selfies taken inside of a car, typically in the driver's seat. Examples of posts from these contexts can be seen in Figure 5.

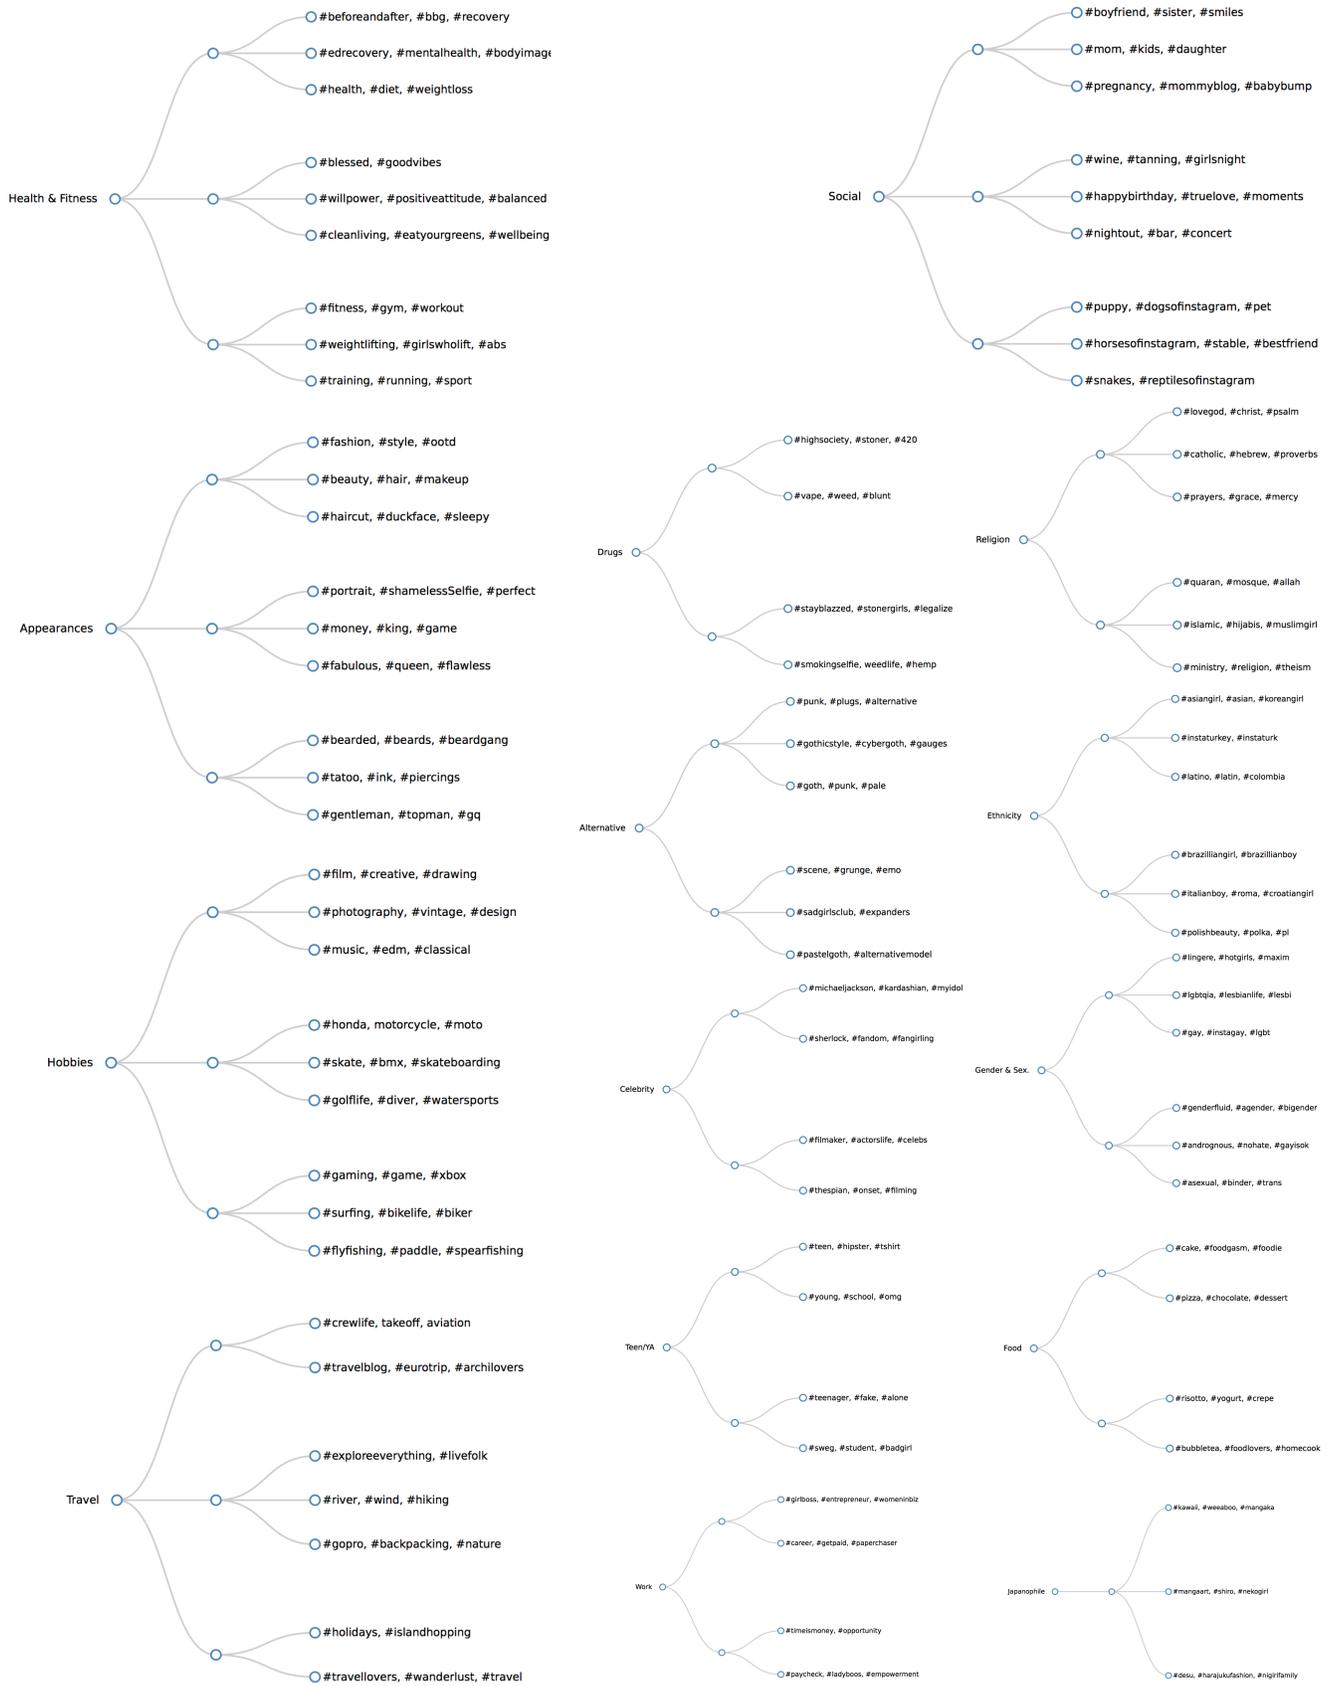


Figure 3: Hierarchy diagrams of each of the context families in the dataset, with their underlying contexts and three top tags.

Identity Emphasis	Age			Gender		Facial Area
	<18	18-35	35+	Male	Female	
Appearance (51.75%)	30.29%	58.88%	10.83%	33.43%	66.57%	6.40%
Social (14.38%)	31.95%	55.95%	12.52%	38.32%	61.68%	5.69%
Ethnicity (12.78%)	33.77%	55.40%	11.83%	33.38%	60.62%	5.31%
Travel (7.16%)	23.32%	57.23%	19.45%	43.75%	56.23%	5.67%
Health & Fitness (5.23%)	23.66%	60.49%	15.85%	45.78%	54.22%	5.70%
Hobbies (2.89%)	27.95%	56.86%	15.19%	46.56%	53.44%	5.16%
Gender & Sexuality (2.40%)	18.22%	66.10%	15.68%	68.84%	31.16%	7.59%
Teen + Young Adult (1.14%)	42.92%	49.92%	7.16%	34.18%	65.82%	6.33%
Celebrity & Entertainment Industry (0.72%)	28.47%	59.52%	12.00%	42.72%	57.28%	4.86%
Alternative Culture (0.68%)	34.80%	56.75%	8.44%	31.77%	68.23%	9.18%
Food (0.43%)	33.24%	54.80%	11.92%	37.80%	62.20%	4.97%
Religion (0.19%)	28.71%	56.53%	14.75%	44.42%	55.58%	4.55%
Drugs (0.14%)	25.48%	60.06%	13.56%	46.63%	53.37%	8.40%
Work (0.06%)	20.76%	59.98%	12.26%	45.70%	54.30%	5.96%
<i>Japanophile (0.03%)</i>	45.93%	48.64%	5.43%	17.75%	82.25%	8.20%

Table 1: Demographic information for all categories for emphasized identity characteristics - listed in order of size



Figure 4: Examples tagged with luxury brands, such as #armani.



Figure 6: Examples of fitness related mirror selfies

This perhaps represents a common performance in this category the purpose of which is to give off a candid impression so that these performances (despite their posed nature) represent an effortless slice-of-life shot.



Figure 5: Examples tagged #carfie.

Despite many of the selfies being used to document everyday life, we also see selfie being used for the purposes of activism and awareness. As mentioned earlier, the Gender & Sexuality category features posts discussing awareness and activism concerning issues facing gender minorities.

Different clusters also show other posing behavior patterns including the mirror selfie - categories such as Health & Fitness and Fashion with low average facial area typically represent categories where mirror selfies are particularly popular (see Figure 6). Perhaps because these context represent contexts in which a person might want to show off more of their body.

Discussion

In this paper, we apply Goffman-inspired empiricism on the phenomenon of selfies. With this approach, we discover the following 15 emphasized identity statement categories of selfies: Appearances, Social, Health & Fitness, Ethnicity, Travel, Hobbies, Teen/Young Adult, Food, Gender & Sexuality, Celebrity & Entertainment Industry, Drugs, Japanophile, Alternative Culture, and Work. Overall, people tended to post more 'solo' selfies in these contexts over group selfies. More women post selfies than men, and most selfies are posted by young adults between the ages of 18-35. Again, all of these demographic features are inferred with computer vision, an imperfect but now reliable tactic.

We also find that the overwhelming majority of tags used alongside #selfie are positive, with the few negative tags (#suicide, #lonely, #sad). This seems to indicate that users are unlikely to post selfies concerning personal struggles and failures, but instead post the positive outcome after the fact. For example, you don't see a fitness related selfie of a person failing to lift a heavy weight—instead you see a selfie later on of that person either lifting that weight once they succeed, or a selfie of the muscles that person has gained once succeeding. As Goffman would predict, people want to appear *effortlessly* happy, healthy, and successful rather than showcasing the intermediate stages or the failures that had to occur to achieve that lifestyle, position, or appearance.

Distribution Over Contexts

The largest of these context families is appearances (52% of posts), and is more than double the size of the next largest, Social. As mentioned earlier, this category represents posts about personal appearance, fashion, style, and status. The clear popularity of this category seems to indicate that users posting selfies try to construct an identity that appears attractive, fashionable, wealthy, and/or important. This category is popular among bloggers, who would use this category to showcase their competence in the fashion and beauty industry.

The least popular context families in this dataset (representing 1% or less of the posts) are food, religion, drugs, work, and japanophile. Perhaps the most surprising is the relative unpopularity of food-related selfies despite tags such as #food, #foodporn, and #foodie representing nearly the same number of posts as #selfie (Instagram reports roughly 200 million posts tagged #selfie and 150M tagged #food as of time of writing). There are two explanations for this: first, the mechanics of taking a selfie while eating would be difficult given that the hands would be occupied by utensils, making holding a phone while eating difficult. However, as demonstrated earlier, posts in this category tend to have the individual posing with the food or leaving the food in the background. This seems to suggest that instead, the act of eating is a less popular context in which to take a selfie—perhaps because the act of eating is a less flattering context to appear in versus appearing with manicured food, or among friends in a social context.

Gender

Our data paints a very traditional portrait of gender roles. Men tend to post in clusters concerning activities (Hobbies - especially concerning bikes and gaming- and Fitness) over personal appearances and fashion. Though the Appearance category also contains posts concerning male facial hair - the number of posts concerning facial hair is much smaller compared to the other tag clusters in the appearance category. This seems to indicate that men post selfies the most in categories that signal strength and skill, both traits typically associated with masculinity and being a viable mate (Donath and others 1999).

Women, on the other hand, dominate categories concerning personal appearances, fashion, and health. Many of these categories emphasize aspects of a person's appearance that indicates good health and attractiveness. For example, the appearance - especially concerning clusters about hair and makeup- category is the most popular category in which women post. Hairstyles have strong implications towards how healthy an individual appears (Mesko and Bereczkei 2004). Both sexes post equally to the status, travel, and work categories, all of which relate towards signaling status and wealth.

Activism in Selfies

Despite the reputation of selfies being associated with narcissism (Gregorie 2015), we see a few instances in our

dataset of selfies being used as way to make personalized activist statements. Especially in the Gender & Sexuality category, we found images associated with tags related to social movements and political ideology as well as tags relating to issues concerning minority groups. This usage case highlights the idea that you can be your own face of your ideals and that selfies help to facilitate this.

We also saw in the Health & Fitness category, users posting selfies about their struggles with both physical and mental illness. More specifically users were posting the selfies as a way to one acknowledge the existence of the illnesses and to combat the stigmas that surround them. By depicting their face alongside mentions of their struggles, they highlight the humanness in their struggles.

Instead of these ideals and issues being talked about in a nebulous way, through selfies the ideas are instead tied to a person and their identity. Selfies reveal the human aspects behind the activism as well as acknowledging underrepresentation of certain groups in a way that highlight's their humanness.

Pics or It Didn't Happen

Across all the selfie contexts, selfies typically contained more than just faces in the images (as evidenced by the fact that less 10% of the photo contained the face). Overall, the images were posed so as to contain props and backgrounds that visually related to the tag list. This indicated that it is not enough to simply tag a photo with an attribute for an audience to believe that you have that attribute - instead you need photographic evidence to support your claims. This strategy invokes the Internet saying "pics or it didnt happen" - which is to say that selfies serve as photographic evidence for a persons behavior or interests.



Figure 7: Examples of Social selfies



Figure 8: Examples of Travel selfies

For example, In the Travel category users often pose in such away that their face occupies are small region of the image, giving the background a larger prominence in the image. This allows the user to show off the location where the

person is at and gives viewers more of the background to be able to recognize the location (see Figure 8 for examples). This also means that the background serves as a prop to prove the authenticity of the selfie. And in the Social category (see Figure 7), we see a pattern of photos featuring more than one person in the shot. These seems to indicate that the most believable way to indicate that you are social is to take a photo with another person or a pet. Thus, a strategy users should employ for identity management is to place themselves in a photo with supportive props and settings.

Missing Contexts

By taking such a large-scale perspective in this work, we can also discuss what contexts we *do not see* in this dataset. This is particularly timely, as the selfie frequently finds itself maligned. As mentioned earlier, one category we do not see is unflattering posts or posts that depict personal failure. Despite mentions of the “ugly selfie” in popular media outlets (Bennett 2014), we found that users overwhelmingly attempt to appear attractive over appearing intentionally unattractive. It seems selfies follow conventional beauty standards, with individuals wishing to appear fashionable, clean, and put-together—even in the selfies tagged #ijust-wokeuplikethis.

Contrary to the way the press often portrays selfies, the overwhelming majority of selfies were typically taken in “appropriate” settings. We didn’t find, for example, funeral selfies (Post 2013) or divorce selfies (Dewey 2015). Which is to say, selfies are in general quite ordinary—depicting everyday life rather than the ridiculous and the improbable.

Limitations

It is important to note the limitations in our data and results. We will first highlight the limitations within our dataset and then move towards the limitations of the methodology. As mentioned earlier, we collected posts that were tagged #selfie from Instagram. Due to the API limitations set by Instagram, this seemed a reasonable trade-off to get a large sample of selfie images. Our data also comes from Instagram only, which means we may miss selfie categories that might be present on another platform. The images we collected were public only; thus there may be additional selfie behaviors that might differ if the user shares the photo privately to a small group of friends.

Since we used computer vision to infer the age and gender information from the facial characteristics, our demographic information report is not 100% accurate. Though the margin of error is quite small, there certainly are misclassifications. We know for certain that there exists misclassifications for non-binary gendered individuals since the algorithm we used can only infer binary genders. Despite filtering our spam tags and non-face images, the final dataset still contains some spam posts. This can occur for a number of reason including that memes and ads often have faces in them.

Conclusion

In this work, we present selfies as a new genre of identity performance which blends the offline and online selves. Using mixed-method analysis on 2.5 million selfies, we report a typology of emergent selfie categories which represent emphasized identity statements. By answering the “what?” and “who?” questions concerning selfie behaviors through our exhaustive typology, we envision our work enabling future research to answer the “how?”, “where?”, and “why?”. We will give a brief overview of possible future questions in these veins below.

Evaluation Audience Impressions

Using the Goffman framework, we evaluated the expressions given in the performance. In future work, perhaps we can evaluate the expressions given off - which is to say “how are these selfies in each category being perceived by the audience?”. To what degree are these performances believed by the audience? And are there selfie categories that are inherently more believable to their audience than others?

Examining Temporal Patterns

In this work, we examined each selfie individually, which provides a first impression of the person in the selfie. How might looking into all a user’s selfies effect these impressions? Do the type of selfies or number of selfies effect the audience’s impressions of the person? How might the audience’s impression change after subsequent viewings of the same selfie?

Examining Geographic Patterns

Our dataset includes the GPS coordinates where some of the photos were taken. Future work can look at the geographic patterns of selfie posting to look for cultural influences on the performances. For example, are there selfie categories established in this paper that are posted in certain locations over others?

Acknowledgements

Funding and sponsorship was provided by the U.S. Army Research Office (ARO) and Defense Advanced Research Projects Agency (DARPA) under Contract No. W911NF-12-1-0043 Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of our sponsors.

References

- Abdullah, S.; Murnane, E. L.; Costa, J. M.; and Choudhury, T. 2015. Collective smile: Measuring societal happiness from geolocated images. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW '15*, 361–374. New York, NY, USA: ACM.
- Bakhshi, S.; Shamma, D. A.; and Gilbert, E. 2014. Faces engage us: Photos with faces attract more likes and comments on instagram. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 965–974. ACM.

- Bennett, J. 2014. With some selfies, the uglier the better.
- Brumfield, B. 2013. Selfie named word of the year for 2013. *CNN Living*, November 20.
- Dewey, C. 2015. On the viral rise of divorce selfies (and the death of traditional marriage).
- Donath, J. S., et al. 1999. Identity and deception in the virtual community. *Communities in cyberspace* 1996:29–59.
- Duggan, M. 2015. Mobile messaging and social media 2015. *Pew Research Center*.
- Farnham, S. D., and Churchill, E. F. 2011. Faceted identity, faceted lives: Social and technical issues with being yourself online. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work, CSCW '11*, 359–368. New York, NY, USA: ACM.
- Frakes, W. 1992. Introduction to information storage and retrieval systems. *Space* 14:10.
- Goffman, E., et al. 1959. The presentation of self in everyday life.
- Gregorie, C. 2015. Study links selfies to narcissism and psychopathy.
- Guntuku, S. C.; Qiu, L.; Roy, S.; Lin, W.; and Jakhetiya, V. 2015. Do others perceive you as you want them to? modeling personality based on selfies.
- Hochman, N., and Manovich, L. 2013. Zooming into an instagram city: Reading the local through social media. *First Monday* 18(7).
- Hochman, N., and Schwartz, R. 2012. 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)*, 6–9.
- Hu, Y.; Manikonda, L.; Kambhampati, S.; et al. 2014. What we instagram: A first analysis of instagram photo content and user types. *Proceedings of ICWSM. AAAI*.
- Instagram. 2015. Celebrating a community of 400 million.
- Jang, J. Y.; Han, K.; Shih, P. C.; and Lee, D. 2015. Generation like: Comparative characteristics in instagram. In *ACM conference on Human factors in computing systems (CHI)*.
- Lampe, C. A.; Ellison, N.; and Steinfield, C. 2007. A familiar face(book): Profile elements as signals in an online social network. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '07*, 435–444. New York, NY, USA: ACM.
- Lu, C., and Tang, X. 2014. Surpassing human-level face verification performance on lfw with gaussianface. *arXiv preprint arXiv:1404.3840*.
- Manovich, L.; Baur, D.; Chow, J.; Goddemeyer, D.; Hochman, N.; Stefaner, M.; Tifentale, A.; and Yazdani, M. 2014. Selficity.
- Merelli, A. 2015. These pictures show what really happens in the background of stylized instagram shots.
- Mesko, N., and Bereczkei, T. 2004. Hairstyle as an adaptive means of displaying phenotypic quality. *Human Nature* 15(3):251–270.
- Newman, M. W.; Lauterbach, D.; Munson, S. A.; Resnick, P.; and Morris, M. E. 2011. It's not that i don't have problems, i'm just not putting them on facebook: Challenges and opportunities in using online social networks for health. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work, CSCW '11*, 341–350. New York, NY, USA: ACM.
- Peixoto, T. P. 2014. Efficient monte carlo and greedy heuristic for the inference of stochastic block models. *Physical Review E* 89(1):012804.
- Post, H. 2013. Funeral selfies are the latest evidence apocalypse can't come soon enough.
- Qiu, L.; Lu, J.; Yang, S.; Qu, W.; and Zhu, T. 2015. What does your selfie say about you? *Computers in Human Behavior* 52:443–449.
- Scellato, S.; Noulas, A.; Lambiotte, R.; and Mascolo, C. 2011. Socio-spatial properties of online location-based social networks. *ICWSM* 11:329–336.
- Shami, N. S.; Ehrlich, K.; Gay, G.; and Hancock, J. T. 2009. Making sense of strangers' expertise from signals in digital artifacts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09*, 69–78. New York, NY, USA: ACM.
- Silfverberg, S.; Liikkanen, L. A.; and Lampinen, A. 2011. "i'll press play, but i won't listen": Profile work in a music-focused social network service. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work, CSCW '11*, 207–216. New York, NY, USA: ACM.
- Silva, T. H.; de Melo, P. O.; Almeida, J. M.; Salles, J.; Loureiro, A.; et al. 2013. A picture of instagram is worth more than a thousand words: Workload characterization and application. In *Distributed Computing in Sensor Systems (DCOSS), 2013 IEEE International Conference on*, 123–132. IEEE.
- Smith, E. R., and Collins, E. C. 2009. Contextualizing person perception: distributed social cognition. *Psychological review* 116(2):343.
- Sunnafrank, M., and Ramirez, A. 2004. At first sight: Persistent relational effects of get-acquainted conversations. *Journal of Social and Personal Relationships* 21(3):361–379.
- Svelander, A., and Wiberg, M. 2015. The practice of selfies. *interactions* 22(4):34–38.
- Viera, A. J.; Garrett, J. M.; et al. 2005. Understanding interobserver agreement: the kappa statistic. *Fam Med* 37(5):360–363.
- Zhou, E.; Cao, Z.; and Yin, Q. 2015. Naive-deep face recognition: Touching the limit of lfw benchmark or not? *arXiv preprint arXiv:1501.04690*.
- Zhu, X., and Ramanan, D. 2012. Face detection, pose estimation, and landmark localization in the wild. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, 2879–2886. IEEE.