You are what emojis say about your pictures

Bizhan Alipour Pijani, Abdessamad Imine, Michaël Rusinowitch

You Are What Emojis Say About Your Pictures*

Language-independent Gender Inference Attack on Facebook

Bizhan Alipour Pijani
Lorraine University, Cnrs, Inria, Loria
Nancy, France
bizhan.alipourpijani@loria.fr

Abdessamad Imine
Lorraine University, Cnrs, Inria, Loria
Nancy, France
abdessamad.imine@loria.fr

Michaël Rusinowitch
Lorraine University, Cnrs, Inria, Loria
Nancy, France
michael.rusinowitch@loria.fr

ABSTRACT

The picture owner's gender has a strong influence on individuals' emotional reactions to the picture. In this study, we investigate gender inference attacks on their owners from pictures meta-data composed of: (i) alt-texts generated by Facebook to describe the content of pictures, and (ii) Emojis/Emoticons posted by friends, friends of friends or regular users as a reaction to the picture. Specifically, we study the correlation of picture owner gender with alt-text, and Emojis/Emoticons used by commenters when reacting to these pictures. We leverage this image sharing and reaction mode of Facebook users to derive an efficient and accurate technique for user gender inference. We show that such a privacy attack often succeeds even when other information than pictures published by their owners is either hidden or unavailable.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability;

KEYWORDS

Social network, privacy, inference attack, gender inference, picture, Emojis

1 INTRODUCTION

Attribute inference from social network profiles and behaviors is a powerful mean to breach user privacy for malicious purpose or targeted advertisements. Attribute inference amounts to derive private attributes of a target user (such as gender, age, political view, or sexual orientation) from publicly available data.

Recent works have investigated two types of attribute inference attacks on Facebook: friend-based [15] and behavior-based [1] inference attacks. Friend-based attacks follow the intuition that *you are who you know*. They work in two steps: the attacker first collects the friend list of the target user, and then from the target user and his/her friend's available data infers target hidden attributes. Behavior-based attacks follow the intuition that *you are how you behave*. In this attribute inference attack, the attacker monitors user behavior such as liked pages and joined groups to infer his/her private attributes. Most existing inference techniques proceed by analyzing data directly generated by the target user, or data obtained by crawling the user vicinity network. However, in a real scenario, the amount of available information to an attacker is rather small.

Unlike previous studies, we show how to detect Facebook user's gender through his/her shared images. With the huge amount of available information on Facebook, identifying user's gender from their online activities and shared data is an essential mechanism for targeted advertising or privacy breaking [6]. Gender is a valuable information source in developing more accurate classifiers for inferring other private attributes such as age [23]. In [13], the authors investigated *479k* Facebook users to determine the level of privacy awareness. They showed that about one-half of their collected Facebook users hide their gender. Facebook users prefer to hide their gender for two reasons. First, they want camouflage against sexual harassment and stalking. The Facebook search bar lets users track down pictures of their female friends, but not the male ones [17]. Second, they want to reduce discrimination. Gender is the direct beneficial information that helps the private sector to present personalized services. Facebook faced criticism for enabling biased discrimination and misinformation. The American Civil Liberties Union (ACLU)[1] accused Facebook of enabling employers to use targeting technology that excludes a woman from receiving job ads for some positions. Additionally, [7] studied how different kinds of self-presentation information on Facebook interpreted by employers and the subsequent attraction in hiring decisions.

While many Facebook users hide their sensitive attributes (e.g., gender, age, political view), pictures are still available to the public. A social media sharing analysis conducted by *The New York Times* revealed that *68%* of their respondents share images to give people a better sense of *who they are* and *what they care about* [27]. Users in social media share pictures to receive feedback for their activities, especially from friends, and acquaintances, provide a great sense of connectedness. However, they lose privacy control on their posted

pictures due to extra information (i.e., *meta-data*) added by third-party during the publication process. Let us review this added *meta-data* that we consider in our attack.

(i) Generated alt-text. Facebook has designed and deployed automatic alt-text, a system to identify faces, objects, and themes from photos by applying computer vision technology. This system is proposed to help blind people to feel more connected and involved in Facebook. The alt text generates a summary of the existing content for the image automatically. The technology can reliably recognize a list of *97* concepts (tags), including people (e.g., people count, smiling, child, baby), objects (e.g., car, building, tree, cloud, food), settings (e.g. inside restaurant, outdoor, nature) and themes (e.g., close-up, selfie, drawing) [29].

(ii) Emoji. Users in social media use Emoji to express their feelings directly. Since the *2010*s, Emojis emerged into communication where Oxford Dictionaries[2] announced 😂, commonly known as *FACE WITH TEARS OF JOY*, as the word of the year.

(iii) Emoticons. An Emoticon[3] is a representation of human facial expression using only keyboard characters such as letters, numbers, and punctuation marks. They express emotions differently through facial gestures inside text-based communication.

1.1 Motivation

To increase awareness of Facebook users about threats on their privacy, we show that from very limited information, even when the user hides his/her owned comments, we can infer the user gender. Previous gender inference attacks on Facebook have two main limitations. First, users friend-based and behavior-based data are extensively considered in the attack process, degrading prediction accuracy in the case of unavailability. Second, it is limited to text-based knowledge. For example, a person's gender identity is constructed through language by using linguistic features associated with male or female writing style on social media, decreasing the prediction accuracy, when texts are multilingual or unavailable. In this work, we relax these two concrete limitations as follows:

1) We exploit non-user generated data (i) alt-text which is computed by Facebook, and (ii) Emojis/Emoticons added by other Facebook users while commenting on the picture.

2) We rely on Emojis/Emoticons as they are not limited to a specific language.

The advantage of Emojis and Emoticons are twofold: it is a universal language, and it is a non-verbal communication way. On the other hand, alt-text is a widely available description of the picture that saves one from image processing tasks.

Female and male posted pictures can receive non-English comments as in Figure 1(a), or Emoji only comments as in Figure 1(b) and analyzing this non-user generated data still is sufficient to launch gender inference attack. Note, we selected these female-owned pictures randomly from Facebook. For the sake of privacy, we blur the pictures and commenter's name. Consider the example in Figure 1(a), where the target user can be inferred as male since the picture received the Spanish word *lindo*. However, a word like *linda* may lead the inference process to an incorrect result. A possible remedy is to consider other available meta-data such as alt-text and Emojis

(a)

(b)
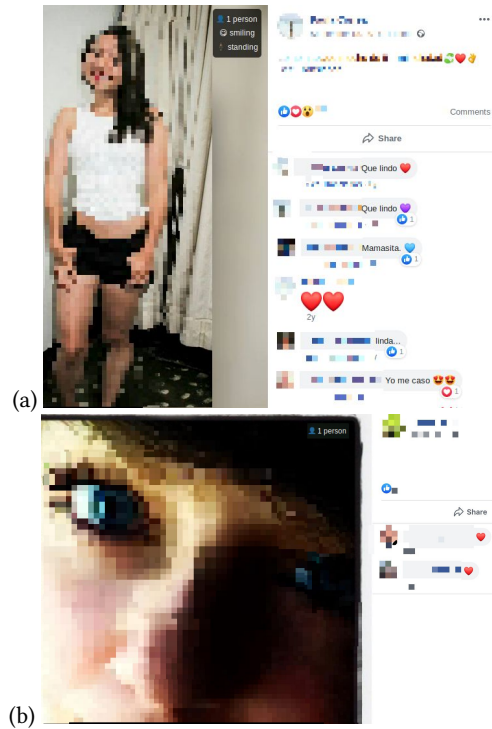
Figure 1: Target user received comments: (a) Emojis and non-English words (b) only Emojis.

into an end-to-end learning system. Figure 1(b) shows an example when the Emojis and generated alt-text are the only available data to the attacker. Emojis/Emoticons are language-independent and make the inference attack possible even when the received comments only consist of Emojis/Emoticons.

Our work shows that gender inference attack is possible even when, as in previous examples, essential information from the target user and his/her vicinity network is not available.

1.2 Problem Description

Using limited amounts of available data, we propose to investigate gender inference attack by leveraging non-user generated data. Therefore our attack is an indirect attack that targets Facebook users even when they are cautious about their privacy, and hides direct generated data such as friend list, liked pages, groups, writing style (e.g., comments), and profile attributes. Inferring gender from non-user generated data (e.g., other Facebook user Emoji preferences) may be deceiving. As shown in [28], some categories of pictures (e.g., baby and animal images) are more likely to receive strong emotional responses than other pictures. Below, we present an example that the Emojis alone have biased gender inference towards one gender, but reviewing alt-text has fixed the initial wrong guess.

Image 1:
Generated alt-text: 1 person, child, sleeping and bedroom
Comment: Precious!!! 💙 ❤️ 💙
Comment: Priceless moments! Love y'all 😘
Comment: I love this!!! 💕

The presence of *child* in the generated alt-text hints that the posted comments, other Facebook users Emoji preference, are not related to the picture owner gender. To circumvent this misleading information in our inference attack process, we filter pictures by using alt-text. We define the picture filtering rules in Sub-section 3.2.2.

1.3 Contributions

We are interested in answering the following questions: Is there a significant difference in the usage of Emoji/Emoticon for commenting female and male-owned pictures? Do female and male receive different Emojis/Emoticons for pictures with the same theme and settings (similar alt-text tags)? Do female and male share picture with similar themes, setting, and objects?

To answer these questions, we consider three different scenarios. In the first scenario, we purely rely on the Emojis/Emoticons regardless of the language of the comments as input data to infer the target user gender. Each Emoji/Emoticon is used without any meaning. In the second scenario, we study only the generated alt-text as input data, in order to discover the female and male preferable pictures sharing style on Facebook. We analyze which tags of alt-text generated more for female and male-owned pictures by Facebook. This analysis enables us to show the preferences of picture sharing style between female and male users. As for the third scenario, we consider the correlation of Emojis/Emoticons and tags of alt-text. In this scenario, we aim to discover the preference Emoji/Emoticon usage of other Facebook users by considering picture owner gender and generated alt-text (tags) of that picture.

Our work is based on the following assumptions:
1) We consider the commenters gender as hidden data which is not accessible.
2) We ignore the comments written by the target user. We assume the target user is careful to hide leaking information from his/her posted comments.
3) We do not know the relation between the target user and his/her commenters.
4) We do not consider user profile name for two reasons. First, although some names used only for a particular gender, it is known that the cultural and geographic origin of names have a large impact on the reliability of gender inference methods[25]. Second, Facebook users may use the shortened name due to privacy concerns. It is a popular tactic to be identifiable only to friends, but not so easily to a stranger.

To the best of our knowledge, we present the first study of gender inference attack on Facebook using target user friends, friends of friends, or ordinary users Emoji preferences. Other users Emojis/Emoticons preferences, in any social media, as a reaction to the observed pictures, posts, or tweets have not studied before this paper.

In the following, the essence of our contributions and improvements over the previous works:
1) Rather than considering the friend-based and behavioral-based data, which might be costly and unavailable in the real scenario, we provide a new approach for gender inference attack by considering picture meta-data.
2) We use Emojis/Emoticons as universal and powerful language

to infer the owner picture gender. This inference has an advantage over any text-based inference attack as it is independent of any language restrictions.
3) We conduct experiments to analyze the Emojis/Emoticons usage for commenting female and male posted pictures on Facebook. These experiments enable us to know the commenter preferences in terms of Emojis/Emoticons that are sufficient to infer gender from pictures.
4) We further analyze the relation between Emoji/Emoticon usage and alt-text for female and male-owned pictures. The recognized relationships then help us to devise attacks with higher inference accuracy. This type of attacks takes advantages of generated picture description instead of relying purely on Emojis/Emoticons.

Outline. The paper is organized as follows: we review related work in Section *2*. In section *3*, we describe our data set preparation. Section *4* presents the gender differences in receiving Emojis. In Section *5*, we explain our selected features. Section *6* presents our experimental results. Section *7* discusses the attack process, and we conclude the paper in Section *8*.

2 RELATED WORK

In this section, we review recent works that are related to our research. For that, we consider two aspects: attribute inference attack and usage preference of Emojis.

2.1 Gender Inference Attack on Social Media

Profiling users based on their activities has obtained great attention in the past decade. Especially, user profiling based on gender is important for recommendation systems. Recently, researchers have investigated on popular social media platforms to distinguish male and female based on content sharing [12] and behavior [20]. Prior works claimed that gender prediction is possible from the writing style of the target user [14], word usage [26] and phrase choice [24]. Gender inference attack by evaluating the target user name performed by [16] across major social networks. However, [25] proved that the performance of this type of attack is biased towards countries of origin. The authors of [11] propose user gender identification through user shared images in Fotolog and Flickr, two image-oriented social networks. They perform image processing task on each crawled image (in the offline mode), which is not feasible in an online attack.

To sum up, the above works depend on the availability of user-generated data, which is costly in a real scenario. In contrast, we perform gender inference attacks by relying only on small information that is not under the direct possession of the user. We do not explore the user network, which has two advantages: (i) makes the attack robust even when the entire personal data and his/her vicinity network is unavailable, and (ii) makes the attack suitable for online mode. Additionally, our attack is not limited to textual language as we use Emoji/Emoticon, a universal language. We showed the benefit of non-user generated data analysis to infer the picture owner gender by relying on the textual part of the comments, regardless of the Emoji/Emoticon usage [3]. However, this work is complementary work to our previous work.

2.2 Emoji Usage Analysis

Several works have analyzed Emoji usage in recent years. Researchers have studied the individual intercept on messages containing Emojis [8]. They have performed experiments on how people use Emojis, an emerging universal language for stating emotions in different countries [18] and culture [5]. Emoji is a rich resource for sentiment analysis and emotion measurement. For example, [2] performs the first quantitative study to correlate Emoji usage to its semantic. Additionally, [4] analyzed messages of Wechat [4], and IM APP users in China, to learn the diversity of usage preferences of Emoji in frequency, type, and sentiment. The diversity and global usage of Emojis lead researchers to perform analysis of Emoji usage according to gender [9]. This study collected the data through the *Kika Keyboard*, and they rely on the usage preference of the user himself. This method may be affected in two way: (i) if the user interacts more with opposite-gender friends, his/her Emoji usage may have affected by them [22], and (ii) if the user is careful in choosing the Emojis. Our work is different in two senses. First, we skip the user Emoji usage and rely on other Facebook users Emotional response to solve the above limitations. Second, we engage the content of the picture as a powerful impact on individuals' emotional responses. Emoji can be interpreted differently according to the platform, which might influence communication [21]. Besides, some researchers investigated the power of Emoji in the cross-lingual sentiment classification task [10] and have performed large scale empirical study on how developers used Emoji on GitHub [19].

To conclude, all these approaches depend on the target user Emoji's usage. It might be straightforward to guess the Emoji publisher's gender. In contrast, we study gender inference attacks on Facebook by considering Emojis/Emoticon's preferences of other Facebook users (e.g., friends) while commenting on target-owned pictures. Although this approach is more complicated, it has two advantages over previous works: (i) target user personality does not affect the performance, (ii) the attack is still possible even when the target user is careful enough to manipulate Emoji/Emoticon neutrally.

3 COLLECTED DATA

In this section, we illustrate in detail the data set and the pre-processing steps.

3.1 Data Set

We launch our gender inference attack by collecting picture meta-data. To that end, we extract each picture meta-data, alt-text, and Emoji(s)/Emoticon(s) of the comments, from the related HTML file. We extract the user gender, when available, to create labeled data sets to be exploited by our supervised machine learning algorithms. We use two labels *female* and *male*, corresponding to biological sex. Let $U = \{u_1, u_2, u_3, .., u_m\}$ be the set of target user pictures. For every picture we collect $u_i = < a_i, e_i >$ where a_i is an alt-text generated by Facebook and e_i is the set of Emojis/Emoticons posted by other Facebook users for that picture. We denote (i) $A = \bigcup_{0 < i \leq m} a_i$ the set of all extracted alt-texts, and (ii) $E = \bigcup_{0 < i \leq m} e_i$

the set of all extracted Emojis/Emoticons. Later, we create $O = \bigcup_{0 < i \leq m} a_i e_i$ which is the co-occurrence of a_i and e_i of picture i.

3.2 Data Pre-processing

We perform three pre-processing steps to clean the collected data as follows:

3.2.1 Cleaning alt-text. Facebook uses *97* different tags to describe the content of each uploaded picture. We reformulate the generated alt-text in three steps, which contain:

Conjunction: Facebook uses two conjunctions inside their generated alt-text, which are *and*, and *or*. For example, Facebook produces tag like *one or more people*, if their algorithm is uncertain about the number of people inside the picture. We reformulate this tag to *more people*.

Redundant Tag: In some pictures, the alt-text comprises repeated tags. Consider the following example where the generated alt-text contains: *3 people, people smiling, people standing, and hat.* The *people* in the first tag refers to the number of people inside the picture, which is *3*. However, *people* in the second and third tags *{people smiling, people standing}* respectively, is the redundant word. As a result, we reformulate the alt-text to *3 people, smiling, standing, hat.*

Text: Facebook in some condition appends the text messages to the generated alt-text if the image has text in it. For example, *4 people, meme, text that says 'WHO ARE YOU IN THE DIFFERENT WORLDS?.* We reconstruct the alt-text to *4 people, meme.*

These steps aim to clean the alt-text in order to create proper *n-grams* and co-occurrence, respectively, in the second and third scenarios. These scenarios are detailed in Section 1.3. Note, we do not reorder the alt-text proposed and generated by Facebook.

3.2.2 Picture Filtering. Facebook generates a description for each picture by using four categories {*people*, *objects*, *setting*, and *theme*} (detailed in Section 1). To address the problem of selecting useful picture meta-data, we define two rules on the generated alt-text to filter pictures. These rules mean to preserve pictures meta-data, alt-text, and Emoji/Emoticon, in the inference process.

General Rule: We keep the picture meta-data if there is no *animals* in *objects tag*, and no *child* inside the alt-text. In this rule we do not consider the number of recognized *people*, in *people tag*, because we are interested in the Facebook users emotional response while observing female and male-owned pictures with similar tags (e.g., *2 people*). As a consequence of this rule, the number of people inside the picture can mislead the attack process, detailed [3]. To solve this problem, we set the second rule on the *people tag*.

Restricted Rule: We keep the picture meta-data if the generated alt-text contains *1 person*. This rule is satisfied when there is only *1 person* in *people tag*, no *animals* in *objects tag*, and no *child* in the generated alt-text.

We apply our picture filtering rules only on the third data set (*O*). Our analysis is regardless of language, and we do not analyze the words which give a clue about the gender disparity inside the posted comments, like *she* and *he*. As a result, we filter the pictures by the alt-text without analyzing the received Emojis.

3.3 Emoticons

An Emoticon is a digital icon that conveys a human expression. It usually contains punctuation marks, numbers, and letters. According to Wikipedia[5] *330* Emoticons, most used Emoticons in different categories, have their corresponding Emoji(s). For example, ; −), and ;) Emoticons have the corresponding Emoji 😉. To reduce the complexity of the model and make the inference attack effective, we find the best match for each Emoticon and replace it by Emoji. We observe *120* different Emoticons in our data set. To perform this task, we used online services such as *The Smiley Dictionary*[6], *urbandictionary*[7], *Emojicodes*[8], *Emoticonr*[9], and *IM Emoticons*[10]. For the rest of the paper, we use Emoji to show Emoji/Emoticon.

4 GENDER BIAS IN RECEIVED EMOJIS

In this section, we discuss the picture owner gender impact on Facebook users emotional responses. We hypothesized that the content of the image along with picture owner gender shape Facebook users response, which evoke a specific emotion in commenting for pictures.

4.1 Emoji Popularity

We investigate other Facebook users Emoji(s) preferences while observing female and male-owned pictures with similar alt-text (tags). Developing a comprehensive analysis on Facebook users emotional response while commenting for female and male-owned images drive to gender inference attack. To understand the differences, we collect the top *4* tags of alt-text generated by Facebook for female and male published pictures. Later, we study top *10* Facebook users Emoji preferences in commenting on photos containing these *4* tags (by considering the owner gender). Table 1 indicates that Facebook users used more smiley Emojis (e.g., 😂) for male images, while emotional Emojis (e.g., ❤️) are commented more for female pictures. The order of these *10* Emojis preferences are from left to right. To avoid the difference in male and female posted pictures number affecting the result, we apply sub-sampling over the female pictures, the dominant one, and calculate the result over the random sub-samples. We discover Facebook users used *1291* different types of Emoji in commenting to female pictures, while they used *877* for male pictures. To conclude, we identify that gender and content of the image affect Facebook users Emoji usage. They have a higher tendency of posting emotional Emoji for female than male pictures.

4.2 Emoji Categories and Reaction Modes

4.2.1 Emoji Categories. We use Emoji categories [11] to demonstrate the influence of picture owner gender on Facebook users responses. To that end, we show the Emoji preferences of other Facebook users according to *10* given categories. We discover Facebook users frequently use Emojis from *Smileys* and *Symbols* categories for commenting in male and female photos, respectively, more than other

Table 1: Emoji preferences of Facebook users in commenting female and male-owned pictures with specific alt-text tags.

alt_text	Gender	Top 10 Facebook users Emojis usage
1 person	Male	😂❤️😎😘😁😂👍😍😂
	Female	😍❤️😘🔥💕💙😊😉😘💜💛
smiling	Male	😂❤️😎😘😂👍😊😂🙏😂
	Female	❤️😍😘😂💕😊🔥❤️💜💛
closeup	Male	😂❤️😎😘😁😂😊😉👍😂
	Female	😍❤️😘😂💙💕😊💜💜😉
outdoor	Male	😂❤️👍😂😁😘😊😊😉😂
	Female	❤️😍😘😁💜💕😊🔥💜💜🎉

categories. As illustrated in Figure 2(a), Facebook users use more emotional Emojis from *Symbol* category, which contains heart-based Emojis, to express their feeling to female-owned images. While, they have a higher tendency in using Emojis from *Smileys* category, which holds face-based Emojis, to comment to male-owned pictures.

4.2.2 Reactions Modes. Following the previous findings, we go further and plot Facebook users reaction to the observed pictures. In general, there are six different reactions on Facebook (as it is shown in Figure 2(b)). The result shows that Facebook users *Like* male and female posted images frequently, in comparison to other reaction types. Although this reaction is very close, female pictures received slightly more *Like* than male pictures. According to Figure 2(b), Facebook users react more by *Haha* to male-owned photos, while they use more *Love* to react to female-owned photos. This outcome is additional evidence to the previous result which commits the Facebook users different emotional response to male and female posted pictures.

To sum up, there is a significant disparity in using Emoji, based on the categories, and reaction when it comes to commenting, and reacting to female and male-owned pictures. These differences drive to gender inference attack. As such, the evaluations of the picture owner gender in forming others emotional responses, and the association of these responses with the content of the picture are two empirical questions that need to be investigated. To answer these questions, we apply Mutual Information (MI) in all three scenarios (See Section 1.3) to effectively conduct the differences, which measures (i) the mutual dependence between gender and received Emojis in the first scenario, (ii) the mutual dependence between picture owner gender and generated alt-text for that picture in the second scenario, and (iii) the correlation of picture owner gender with generated alt-text and received Emojis in the third scenario.

Let e and a denote the Emoji and alt-text tag respectively. Let X be a random variable that takes values $x = 1$ (the posted photo contain e in the first scenario, a in the second scenario, and a and e in the third scenario) and $x = 0$ (the posted photo does not contain e in the first scenario, a in the second scenario, and a and e in the third scenario), and Y is a random variable that shows the picture owner gender, where it takes values $y = 1$ for female and $y = 0$ for male. Note, there is no value for a in the first scenario, and e in the second scenario as we only consider the Emoji, and alt-text as the input data to our attack process respectively. Then, we compute the MI as follows:

$$I(X;Y)_{ae} = \sum_{x \in \{0,1\}} \sum_{y \in \{0,1\}} P(X=x, Y=y)_{ae} \, log_2 \frac{P(X=x, Y=y)_{ae}}{P(X=x)_{ae}, P(Y=y)_{ae}} \quad (1)$$

Figure 2: Received emotional responses: (a) Emojis based on categories (b) Reactions.

Table 2: Discriminative Emojis for female and male-owned picture

Emoji	MI	p(female|Emoji)	p(male|Emoji)
❤️	0.026	**0.84**	0.16
😍	0.025	**0.89**	0.11
🥰	0.015	**0.86**	0.14
💕	0.011	**0.91**	0.09
💙	0.008	**0.87**	0.13
🔥	0.008	**0.92**	0.08
😊	0.006	**0.82**	0.18
💗	0.006	**0.92**	0.08
💜	0.005	**0.91**	0.09
💜	0.005	**0.90**	0.10

where $P(X = x)_{ae}$, $P(Y = y)_{ae}$ are the marginal probabilities of x and y, and $P(X = x, Y = y)_{ae}$ is the joint probability of x and y. For example, the join probability of $P(1, 1)_{\{1person, 😍\}}$ is the probability that female-owned picture contains *1 person* in the generated alt-text, and other Facebook users used 😍 for commenting for that picture. Additionally, we compute the probability of a person being male or female, given the picture generated alt-text, and commenting Emojis. For example, $p(female|\{1person, 😍\})$ is the probability that the user is female if her picture contains *1 person* in Facebook generated alt-text and received comments with 😍 Emoji from other Facebook users. Note, we compute the same probability for the first, and second scenario. For example, $p(female|😍)$ is the probability that the user is female if she received comments with 😍 Emoji from other Facebook users. Also, $p(male|beard)$ shows the probability that the target user is male if the generated alt-text for his picture contains *beard*.

Table 2 illustrates the top *10* discriminative Emojis used by other Facebook users in commenting to female and male-owned pictures. Table 3 shows the discriminative alt-texts generated for female and male-owned pictures, while Tables 4, and 5 take into count the correlation of received Emojis and generated alt-text. Note that in Table 2, we draw inspiration from [9], and all entries in Table 2 are discriminative Emojis for female. From the same spirit, we show the top 10 discriminative alt-text, and correlation of alt-text and Emoji for female and male in table 3, 4, and 5.

Table 3: Discriminative alt-text for female and male-owned picture

alt_text	MI	p(female|alt_text)	p(male|alt_text)
closeup	0.015	**0.80**	0.20
smiling	0.011	**0.75**	0.25
1 person	0.006	**0.70**	0.30
1 person smiling	0.006	**0.80**	0.20
smiling closeup	0.006	**0.89**	0.11
1 person closeup	0.004	**0.85**	0.15
beard	0.004	0.28	**0.72**
2 people smiling	0.004	**0.74**	0.26
car	0.003	0.28	**0.72**
selfie closeup	0.003	**0.82**	0.18

Table 4: Discriminative Emojis and alt-text for female-owned picture (Restricted rule)

Emoji + alt_text	MI	p(female|Emoji)	p(male|Emoji)
(1 person, 😎)	0.044	**0.89**	0.11
(1 person, ❤️)	0.036	**0.87**	0.13
(closeup, 😍)	0.027	**0.93**	0.07
(closeup, ❤️)	0.021	**0.92**	0.08
(1 person, 🥰)	0.020	**0.85**	0.15
(1 person, 💕)	0.019	**0.86**	0.14
(1 person, 🔥)	0.017	**0.93**	0.07
(smiling, 😍)	0.016	**0.93**	0.07
(closeup, 🥰)	0.015	**0.92**	0.08
(smiling, ❤️)	0.013	**0.90**	0.10

Table 5: Discriminative Emojis and alt-text for male-owned picture (Restricted rule)

Emoji + alt_text	MI	p(female|Emoji)	p(male|Emoji)
(beard, 😂)	0.009	0.09	**0.91**
(beard, ❤️)	0.003	0.17	**0.83**
(beard, 🥰)	0.002	0.11	**0.89**
(hat, 😂)	0.002	0.22	**0.78**
(beard, 😊)	0.002	0.08	**0.92**
(outdoor, 😍)	0.002	0.22	**0.78**
(outdoor, 👍)	0.002	0.37	**0.63**
(smiling, 🤣)	0.001	0.23	**0.77**
(sky, 😂)	0.001	0.40	**0.60**
(standing, 😎)	0.001	0.33	**0.67**

5 FEATURES

In this section, we discuss our selected features and illustrate them with an example. Then we sketch the feature extraction algorithms.

5.1 Feature Selection

Feature selection is the process of identifying and selecting relevant features correlated to variables of interest (in our case gender). The purpose of feature selection is three-fold: promoting the model prediction performance, providing faster and efficient classifiers, and reducing the data dimensionality that decreases the complexity of the model. We extract features in four different categories, which consist of:

5.1.1 N-grams. To capture syntactic similarities, we employed n-grams on Facebook generated alt-text. N-grams are a set of co-occurring words within a given window size (*n*). Our previous experiments showed that introducing *6-grams* on *alt-texts* degraded the performance, and we kept the lengths up to *5-grams* for *alt-texts* dataset [3]. In this work, we use our previous experimental result and apply *5-grams* on the alt-text. Additionally, we run *1-grams* on Emojis, as other *n-grams* do not perform best on Emoji. We keep

Table 6: Pattern-based features

Non-textual
Single Emoji without text. For example, Figure 1(b)
Repeated Emoji without text. For example, Figure 1(a) forth comments.
Textual
Single Emoji. For example, the third and fourth comments of Image 1.2.
Single repeated Emoji. For example, the last comment of Figure 1(a).
Several repeated/non-repeated Emojis. For example, the first comment of Image 1.2.

the result of *n-grams* that appear more than *50* times in total. With that, we collected *894* features.

5.1.2 Patterns. The different Emojis preferences of Facebook users in commenting on female or male pictures follow some patterns, from which we derive features. We divided these patterns into non-textual and textual categories. The non-textual category contains comments which have only Emoji (no words), where the textual category contains words and Emoji. Table 6, defines our five pattern-based features.

5.1.3 Emoji Usage. As presented in Table 1, Emoji preferences in commenting pictures are significantly different according to picture owner gender. We compute the frequency of each Emoji concerning all Emojis. To reduce the number of selected features, and the complexity, we consider only Emojis that Facebook users used more than *50* times in total. By considering that, we extracted *735* features.

5.1.4 Correlation of alt-text and Emoji. Tables 4 and 5 confirm the difference in alt-text and Emoji correlation for male and female posted pictures. These differences lead us to consider pairs of alt-text and Emoji as features. To that end, we constructed the co-occurrence network for Emojis and alt-text. Note that we drop rare co-occurrence pairs that appear less than *50* times in all the pictures. Finally, we collected *1363* features from all the possible combination of Emojis, and alt-text in our data set.

In total, we selected *2992* features from the above different categories. After selecting these features, we needed to apply feature extraction algorithms to prune and reduce the selected features.

5.2 Feature Extraction

The goal of feature extraction is to downsample the selected features while keeping those that contribute more to predicting the final result (in our case female and male). We examined four feature extraction methods (*Chi-Square*[12], *Information Gain*[13], *Feature importance*[14], and *Univariate feature selection*[15]) in our previous work. We evaluated all the possible individual and combined feature extraction methods to achieve the best features set [3]. In this study, we used our previous algorithm with a slight change. We separately run the feature extraction methods on *A*, *E*, and *O* data sets, and we did not merge the output features set of each data set. As a result, we only run the feature extraction methods one time.

---

6 EXPERIMENTS

In this section, we demonstrate, and evaluate the experimental result of our approach for all three scenarios.

6.1 Data Set

Using a Python crawler, we collected a set of *141,812* pictures and their *446,655* messages. Our statistic showed that *1291* different types of Emoji appear in our data set. We randomly selected Facebook users to not biased the data collection by region or country usage preferences. Facebook was unable to generate alt-text for *13000* pictures. We kept those pictures for our first attack scenario, where we purely rely on Emojis. Note that although most collected data set comprises textual content, we only considered the Emojis in our study.

6.2 Results

The experimental results are achieved by applying the classifiers from the Python library *scikit-learn*. We model gender inference attack as a binary classification problem. To achieve robust results, we apply several supervised machine learning algorithms such as *Logistic Regression*, *Random Forests*, *K-Nearest Neighbors*, *Support Vector Machine*, *Naive Bayes* and *Decision Tree*. To evaluate the classifier, we select the same number of male and female to prevent biased classification. Train-test splitting was preferable in this study as it runs k-times faster than k-fold. We vary the size of the training set to measure the difference in the attack accuracy. Finally, we choose the train-test size of *75-25*, which gives the best accuracy. To address the problem of fairly estimating the performance of each classifier, and make sure the classifiers can generalize to unseen data, we split the training data set of *train-test* into training, and validation data set *75-25*. We randomly select the training data set of *train-validation* and train the classifiers by using these data set. Later, we record the performance on the validation data set and adjust the hyper-parameters to optimize the performance of the classifiers. Eventually, we evaluate the classifiers on the test data set of *train-test*. Considering the extracted gender as the ground-truth, to evaluate our attack, we compute the standard *Accuracy*, *Precision*, *Recall*, and *F1_score* metrics. In Tables 7 and 8, we report the results of the test data set for gender inference attack. In these two tables, we compare our six classifiers on three different scenarios.

In the first scenario, we purely conduct an inference attack by using other Facebook users Emoji preferences in commenting on the target user pictures. As we mention in Section 5.2, we use our previous feature selection algorithm to select the best features. The *Feature importance* method performs the best and generates *258* features. We evaluate the effectiveness of the selected features by training the classifiers on these features. Selecting the best machine learning algorithms based on evaluation metrics depends entirely on the problem. We consider the gender inference attack as a binary classification problem and since we are dealing with a balanced data set, we used *Accuracy* as important metrics to evaluate each classifier performance. As illustrated in Table 7 (left section), *Logistic Regression* classifier outperforms other classifiers in this scenario. According to the result, *Logistic Regression* can infer the target user gender with accuracy *76%*. *Logistic Regression* is a discriminative

model which is appropriate to conduct when the dependent variable is binary. So, it learns better than other classifiers between the dependent and independent variable in the first scenario. We can observe that *Logistic Regression* performs better than other supervised classifiers in the first scenario based on *Accuracy*, *Recall* and *F1-Score*. *Logistic Regression* also performs slightly worse than *Support Vector Machine* and *Naive Bayes* classifiers in *Precision*. The result confirms the preference Emojis usage of other Facebook users is enough to implement a machine learning classifier to infer the target user gender.

In the second scenario, we rely on Facebook generated alt-text to infer the picture owner gender. Table 7 (right section) displays our inference results on the extracted alt-texts features. The intersection of *Feature importance* and *Chi-Square* perform the best and generate 316 features. To measure the efficiency of the selected features, we train the classifiers by using only these features. According to Table 7 (right section), *Support Vector Machine* performs better than *Logistic Regression* in *Accuracy*. Based on the result, an attacker can infer the target user gender with an accuracy of *80%*. we can observe that *Logistic Regression* performs slightly better than *Support Vector Machine* in *Precision* and *Recal*. Moreover, *Support Vector Machine* also perform slightly worse than *Logistic Regression* and *Naive Based* in *F1-Score*. *Support Vector Machine* is a discriminative classifier defined by a separating hyperplane. The goal of the algorithm is to determine a hyperplane in N-dimensional space where N is the number of features, in our case *316*, that precisely classify the new example. In the second scenario, *SVM* outputs an optimal hyperplane that classifies *female* from *male* better than other classifiers.

In the third scenario, we train the machine learning classifiers by using the co-occurrence of alt-texts and Emoji preferences of other Facebook users. We show the result of each classifier for the *General rule*, and *Restricted rule* separately. The intersection of *Feature importance* and *Univariate feature selection* performs the best and creates *224*, and *211* features for *General rule*, and *Restricted rule*, respectively. Table 8 (left section) displays our inference attack results by considering the *General rule*. It shows that the *Logistic Regression* performs better than other classifiers in *Accuracy*, *Recall*, and *F1-Score* when using these *224* features. *Naive Bayes* classifier received a slightly higher *Precision* score than *Logistic Regression*. The *Logistic Regression* model, which had *76%* accuracy in the first scenario, and *80%* in the second scenario gets a *7%*, and *3%* accuracy boost in this scenario, which is a fairly substantial gain in accuracy. On the other hand, Table 8 (right section) shows the result of the *Restricted rule*. As illustrated, the accuracy increased *4%* in compare to *General rule*. Similar to the *General rule* result, *Logistic Regression* performs better than other classifiers in *Accuracy*, *Recall*, and *F1-Score*. *Naive Bayes*, and *Support Vector Machine* perform best in *Precision*. Note, by setting the *Restricted rule*, we filter *30,595* pictures more than the *General rule*. The results in Table 8 shows the effect of third scenario in increasing attack accuracy.

To conclude, *Logistic Regression* performs best in first and third scenarios, and *Support Vector Machine* was the suitable classifier for the second scenario. The results confirm our hypothesis that the gender and contents of the picture have an impact on Facebook users emotional responses. Moreover, there is a substantial difference in receiving Emojis between female and male posted pictures

| | Emojis | | | | alt_text | | | |
	Accuracy	Precision	Recall	Fscore	Accuracy	Precision	Recall	Fscore
LR	**76.21**	81.57	71.54	76.22	79.16	82.19	77.74	79.90
KNN	68.38	73.08	64.49	68.49	73.54	76.34	72.93	74.59
SVM	76.09	83.35	68.93	75.45	**80.04**	82.07	75.71	78.76
NB	70.85	83.09	56.90	67.53	68.50	65.08	56.62	79.13
DT	65.16	69.50	61.78	65.41	69.61	71.64	71.12	71.38
RF	69.40	73.02	67.59	70.17	74.12	79.16	69.89	74.19

Table 7: Machine Learning classifiers performance with optimal hyper-parameters.

| | General Rule | | | | Restericted rule | | | |
| | Emojis + alt_text | | | | Emojis + alt_text | | | |
	Accuracy	Precision	Recall	Fscore	Accuracy	Precision	Recall	Fscore
LR	**83.89**	86.80	80.12	83.31	**87.05**	85.68	79.64	82.55
KNN	77.69	83.48	72.49	77.59	76.79	83.67	70.17	76.32
SVM	81.36	85.92	77.79	81.65	81.09	86.34	76.68	81.21
NB	79.29	86.90	72.00	78.75	78.55	86.82	70.45	77.78
DT	69.39	72.84	68.07	70.32	72.13	75.29	71.14	73.11
RF	76.72	81.08	73.48	77.09	76.80	82.17	72.16	76.83

Table 8: Machine Learning classifiers performance with optimal hyper-parameters.

on Facebook. As a result, an attacker can train standard classifiers by using non-user generated data (Emoji preferences of other Facebook users and generated alt-text) to infer the picture owner gender. Note, as we relied solely on non-user generated data, the results cannot be compared to previous works that used user-generated data.

7 DISCUSSION

Based on our analysis, the best scenario for the attacker is the third scenario, *Restricted rule*, when he has access to Facebook generated alt-text and other Facebook users posted Emojis. Selecting *General rule* and *Restricted rule* depends on the quantity of the pictures, after setting filtering rules. Initially, the attacker applies the *Restricted rule*, and verify if the number of filtered images are suitable to continue the attack. Otherwise, the attacker picks the *General rule* to increase the number of pictures in his analysis. The attacker uses the *General rule* if the number of crawled images after filtering pictures are not sufficient to be analyzed. The second scenario is suitable when the crawled pictures contain only generated alt-text and no Emoji(s) commented by target friends, friends of friends, or ordinary users. In this case, the generated alt-text can help the attacker to infer the target user gender. The first scenario is useful when the Facebook algorithm, in rare cases, is enabled to generate alt-text for the target user pictures. In this case, the attacker has the advantage of using Emojis, the universal language, and launch the gender inference attack.

A limitation of our work concerns the type of data that we collect. As we mentioned earlier, we only collect non-user generated data (picture meta-data). However, this data set might not be suitable for another social network such as Twitter. For example, Facebook generates a description of posted pictures (alt-text), which is not yet accessible on Twitter. As a result, it disables a more comprehensive analysis of other social media. Additionally, as the meaning of Emojis may differ on each platform, we are not sure that our analysis remains valid for Emoji usage by Twitter users.

8 CONCLUSION

Identifying users gender from their online activities and data sharing behavior is an important topic in the growing research field of social networks. It provides an opportunity for targeted advertising, profile customization, or privacy risks. This study has investigated *141,812* images and their *446,655* comments. Based on the intensive analyses of the shared images, this work has demonstrated a new perspective of gender inference attack on Facebook users by relying on non-user generated data. We have shown the possibility of gender inference attack even when all user attributes/activities are hidden, such as profile attributes, friend list, liked pages and joined groups. Our experimental results showed that on average female posted pictures receive more Emojis-based comments than male posted pictures. Additionally, we have shown alt-text gives extra free information that boosts inference accuracy. We have noticed that other Facebook users use more emotional Emojis to comment on female posted images with a particular theme and setting.

As future work, we plan to (i) apply sentiment analysis and consider the meaning of Emojis/Emoticons in the classification task, (ii) compare Deep Learning approaches with current Machine Learning algorithms, (iii) create an application to deal with online gender inference attacks, and (iv) propose counter-measures to picture owners, considering a trade-off between privacy risks and comments-based social benefits.

REFERENCES

[1] Chaabane Abdelberi, Gergely Ács, and Mohamed Ali Kâafar. 2012. You are what you like! Information leakage through users' Interests. In *19th Annual Network and Distributed System Security Symposium, NDSS*. The Internet Society, San Diego, California, USA.

[2] Wei Ai, Xuan Lu, Xuanzhe Liu, Ning Wang, Gang Huang, and Qiaozhu Mei. 2017. Untangling Emoji Popularity Through Semantic Embeddings. In *Proceedings of the Eleventh International Conference on Web and Social Media, ICWSM*. AAAI Press, Montréal, Québec, Canada, 2–11.

[3] Bizhan Alipour, Abdessamad Imine, and Michaël Rusinowitch. 2019. Gender Inference for Facebook Picture Owners. In *International Conference on Trust, Privacy and Security in Digital Business, TrustBus*. Springer, Linz, Austria, 145–160.

[4] Jiaxin An, Tian Li, Yifei Teng, and Pengyi Zhang. 2018. Factors Influencing Emoji Usage in Smartphone Mediated Communications. In *Transforming Digital Worlds - 13th International Conference, iConference*. Springer, Sheffield, UK, 423–428.

[5] Francesco Barbieri, Germán Kruszewski, Francesco Ronzano, and Horacio Saggion. 2016. How Cosmopolitan Are Emojis?: Exploring Emojis Usage and Meaning over Different Languages with Distributional Semantics. In *Proceedings of the Conference on Multimedia Conference, MM*. ACM, Amsterdam, The Netherlands, 531–535.

[6] Tena Belinic. 2009. Personality profile of social media users how to get maximum from it. https://medium.com/krakensystems-blog/personality-profile-of-social-media-users-how-to-get-maximum-from-it-5e8b803efb30.

[7] Robyn L. Brouer, Michael Stefanone, Rebecca L. Badawy, and Michael J. Egnoto. 2017. Gender (In)Consistent Communication via Social Media and Hireability: An Exploratory Study. In *50th Hawaii International Conference on System Sciences, HICSS*. ScholarSpace / AISeL, Hawaii, USA, 1–10.

[8] Sarah E Butterworth, Traci A Giuliano, Justin White, Lizette Cantu, and Kyle C Fraser. 2019. Sender Gender Influences Emoji Interpretation in Text Messages. *Frontiers in psychology* 10 (2019), 784.

[9] Zhenpeng Chen, Xuan Lu, Wei Ai, Huoran Li, Qiaozhu Mei, and Xuanzhe Liu. 2018. Through a Gender Lens: Learning Usage Patterns of Emojis from Large-Scale Android Users. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW*. ACM, Lyon, France, 763–772.

[10] Zhenpeng Chen, Sheng Shen, Ziniu Hu, Xuan Lu, Qiaozhu Mei, and Xuanzhe Liu. 2019. Emoji-Powered Representation Learning for Cross-Lingual Sentiment Classification. In *The World Wide Web Conference, WWW*. ACM, San Francisco, CA, USA, 251–262.

[11] Ming Cheung and James She. 2017. An Analytic System for User Gender Identification through User Shared Images. *TOMCCAP* 13, 3 (2017), 30.

[12] Munmun De Choudhury, Sanket S. Sharma, Tomaz Logar, Wouter Eekhout, and René Clausen Nielsen. 2017. Gender and Cross-Cultural Differences in Social Media Disclosures of Mental Illness. In *Proceedings of the Conference on Computer Supported Cooperative Work and Social Computing, CSCW*. ACM, Portland, OR, USA, 353–369.

[13] Reza Farahbakhsh, Xiao Han, Ángel Cuevas, and Noël Crespi. 2017. Analysis of publicly disclosed information in Facebook profiles. *CoRR* abs/1705.00515 (2017).

[14] Lucie Flekova, Jordan Carpenter, Salvatore Giorgi, Lyle H. Ungar, and Daniel Preotiuc-Pietro. 2016. Analyzing Biases in Human Perception of User Age and Gender from Text. In *Proceedings of the 54th Annual Meeting of the Association for Computational*. ACL, Berlin, Germany, 843–854.

[15] Neil Zhenqiang Gong and Bin Liu. 2016. You Are Who You Know and How You Behave: Attribute Inference Attacks via Users' Social Friends and Behaviors. In *25th Security Symposium*. USENIX, Austin, TX, USA, 979–995.

[16] Fariba Karimi, Claudia Wagner, Florian Lemmerich, Mohsen Jadidi, and Markus Strohmaier. 2016. Inferring Gender from Names on the Web: A Comparative Evaluation of Gender Detection Methods. In *Proceedings of the 25th International Conference on World Wide Web, WWW*. ACM, Montreal, Canada, 53–54.

[17] Abbey Lenton. 2019. Facebook Wants You To Search Photos Of Your Female Friends At The Beach, But Not Your Male Mates. https://www.whimn.com.au/talk/people/facebook-wants-you-to-search-photos-of-your-female-friends-at-the-beach-but-not-your-male-mates/news-story/bbc21ee6883bd07bfbbbe76a0c8ca54c.

[18] Xuan Lu, Wei Ai, Xuanzhe Liu, Qian Li, Ning Wang, Gang Huang, and Qiaozhu Mei. 2016. Learning from the ubiquitous language: an empirical analysis of emoji usage of smartphone users. In *Proceedings of the International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp*. ACM, Heidelberg, Germany, 770–780.

[19] Xuan Lu, Yanbin Cao, Zhenpeng Chen, and Xuanzhe Liu. 2018. A First Look at Emoji Usage on GitHub: An Empirical Study. *CoRR* abs/1812.04863 (2018).

[20] Puneet Singh Ludu. 2014. Inferring gender of a Twitter user using celebrities it follows. *CoRR* abs/1405.6667 (2014).

[21] Hannah Jean Miller, Daniel Kluver, Jacob Thebault-Spieker, Loren G. Terveen, and Brent J. Hecht. 2017. Understanding Emoji Ambiguity in Context: The Role of Text in Emoji-Related Miscommunication. In *Proceedings of the Eleventh International Conference on Web and Social Media, ICWSM*. AAAI Press, Montréal, Québec, Canada, 152–161.

[22] Dong Nguyen, Dolf Trieschnigg, A. Seza Dogruöz, Rilana Gravel, Mariët Theune, Theo Meder, and Franciska de Jong. 2014. Why Gender and Age Prediction from Tweets is Hard: Lessons from a Crowdsourcing Experiment. In *25th International Conference on Computational Linguistics, Proceedings of the Conference, COLING*. ACL, Dublin, Ireland, 1950–1961.

[23] Claudia Peersman, Walter Daelemans, and Leona Van Vaerenbergh. 2011. Predicting age and gender in online social networks. In *Proceedings of the 3rd International Workshop on Search and Mining User-Generated Contents, SMUC*. ACM, Glasgow, United Kingdom, 37–44.

[24] Daniel Preotiuc-Pietro, Wei Xu, and Lyle H. Ungar. 2016. Discovering User Attribute Stylistic Differences via Paraphrasing. In *Proceedings of the Thirtieth Conference on Artificial Intelligence,*. AAAI Press, Phoenix, Arizona, USA, 3030–3037.

[25] Lucía Santamaría and Helena Mihaljevic. 2018. Comparison and benchmark of name-to-gender inference services. *PeerJ Computer Science* (2018).

[26] Maarten Sap, Gregory J. Park, Johannes C. Eichstaedt, Margaret L. Kern, David Stillwell, Michal Kosinski, Lyle H. Ungar, and H. Andrew Schwartz. 2014. Developing Age and Gender Predictive Lexica over Social Media. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP*. ACL, Doha, Qatar, 1146–1151.

[27] Shutterstock. 2015. The Psychology Behind Why We Share on Social Media. Retrieved January 13, 2015 from https://www.shutterstock.com/blog/the-psychology-behind-why-we-share-on-social-media

[28] Shutterstock. 2017. 6 Types of Images That Elicit an Emotional Response. Retrieved May 5, 2017 from https://www.shutterstock.com/blog/6-types-of-images-that-elicit-an-emotional-response

[29] Shaomei Wu, Jeffrey Wieland, Omid Farivar, and Julie Schiller. 2017. Automatic Alt-text: Computer-generated Image Descriptions for Blind Users on a Social Network Service. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing, CSCW*. ACM, Portland, OR, USA, 1180–1192.