

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

---

12-2017

### Inferring social media users' demographics from profile pictures: A Face++ analysis on Twitter users

Soon-Gyo JUNG

Jisun AN

*Singapore Management University, jisunan@smu.edu.sg*

Haewoon KWAK

Joni SALMINEN

Bernard J. JANSEN

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Databases and Information Systems Commons](#), [Numerical Analysis and Scientific Computing Commons](#), and the [Social Media Commons](#)

---

#### Citation

1

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylids@smu.edu.sg](mailto:cherylids@smu.edu.sg).

Jung, S.G., An, J., Kwak, H., Salminen, J. & Jansen, B.J. (2017). Inferring social media users' demographics from profile pictures: A Face++ analysis on Twitter users. In *Proceedings of The 17th International Conference on Electronic Business* (pp. 140-145). ICEB, Dubai, UAE, December 4-8.  
<https://aisel.aisnet.org/iceb2017/22/>

## Inferring Social Media Users' Demographics from Profile Pictures: A Face++ Analysis on Twitter Users

Soon-Gyo Jung\*, [sjung@hbku.edu.qa](mailto:sjung@hbku.edu.qa)<sup>1</sup>

Jisun An, [jan@hbku.edu.qa](mailto:jan@hbku.edu.qa)<sup>1</sup>

Haewoon Kwak, [hkwak@hbku.edu.qa](mailto:hkwak@hbku.edu.qa)<sup>1</sup>

Joni Salminen, [jsalminen@hbku.edu.qa](mailto:jsalminen@hbku.edu.qa)<sup>1</sup>

Bernard J. Jansen, [jjansen@acm.org](mailto:jjansen@acm.org)<sup>1</sup>

<sup>1</sup>Qatar Computing Research Institute, Hamad Bin Khalifa University, Qatar,

### ABSTRACT

In this research, we evaluate the applicability of using facial recognition of social media account profile pictures to infer the demographic attributes of gender, race, and age of the account owners leveraging a commercial and well-known image service, specifically Face++. Our goal is to determine the feasibility of this approach for actual system implementation. Using a dataset of approximately 10,000 Twitter profile pictures, we use Face++ to classify this set of images for gender, race, and age. We determine that about 30% of these profile pictures contain identifiable images of people using the current state-of-the-art automated means. We then employ human evaluations to manually tag both the set of images that were determined to contain faces and the set that was determined not to contain faces, comparing the results to Face++. Of the thirty percent that Face++ identified as containing a face, about 80% are more likely than not the account holder based on our manual classification, with a variety of issues in the remaining 20%. Of the images that Face++ was unable to detect a face, we isolate a variety of likely issues preventing this detection, when a face actually appeared in the image. Overall, we find the applicability of automatic facial recognition to infer demographics for system development to be problematic, despite the reported high accuracy achieved for image test collections.

*Keywords:* Face++, Twitter, demographic inference, social media, demographics, user attributes, personas.

---

\*Corresponding author

### INTRODUCTION

Social media platforms have become a major resource for population or customer segmentation analysis in many fields, including sociology, marketing, and public policy (Zagheni, Garimella & Weber, 2014). Often, these analyses require the use of users' demographic information, many times gathered from data available from social media platforms. Major social media platforms, such as YouTube and Facebook, share aggregated demographic data, in a privacy-preserving manner. Access to this data is provided via application program interfaces (APIs) to account holders. The APIs enable one to retrieve and utilize demographic information related to user interactions with posts in these accounts. However, not all social media platforms provide such post-level information. For example, Twitter does not provide such demographic data at a tweet level. While Twitter does provide information about followers as an overall audience statistic, more granular, tweet-level information is necessary for many purposes, such as analyzing the performance of content among different target groups, or for automatic persona generation that associates demographic groups with content-level interaction (An *et al*, 2016).

Because accurately gathering demographic information from these platforms, particularly Twitter, has numerous obstacles (Salminen *et al*, 2017), researchers have proposed methods for automatically inferring demographic attributes (e.g., by leveraging users' behavioral or profile information) (Alowibdi, Buy & Yu, 2013a; 2013b; Asoh & Ono, 2012). One of the most promising methods is the analysis of profile pictures. In particular, Face++ provides technologies for facial analysis within images and it is currently used by many researchers and industry practitioners for automating the analysis of facial images (Asoh & Ono, 2012).

In the research reported here, we evaluate the applicability of using profile pictures combined with facial recognition of these images to infer the demographic attributes of gender, race, and age of the account owners. We use Twitter (<https://www.twitter.com>), a popular social media platform, to retrieve users' profile pictures, due to three reasons: 1) our system leverages tweets as a part of the automatic persona generation methodology, 2) the profile pictures on Twitter can be readily accessed via an API, and 3) Twitter is a highly popular social media platform and therefore impactful research context. We aim to first evaluate whether or not profile pictures are a reliable source for determining demographics automatically. For this, we leverage a commercial and well-known facial recognition service in images, specifically Face++ (<https://www.faceplusplus.com/>). Face++ is chosen because it represents a state-of-the-art facial recognition technology and is

widely cited in academic works (An & Weber, 2016; Azzopardi, Greco & Vento, 2016; Vikatos *et al.*, 2017). However, as far as we are aware, its performance has not been validated in the context of social media profile pictures. While Face++ has been found to perform well with clear portrait images (Fan *et al.*, 2014; Messias, Vikatos & Benevenuto, 2017), social media profile pictures may be noisier than test collections, involving attributes that may hinder automatic retrieval of demographic information.

Yet, accurately determining user demographics automatically has applicability in several areas, such as audience segmentation. We are specifically interested in determining user demographics for automatic persona generation in which we are utilizing major social media platforms' demographic data to generate personas representing the core customer segments of various organizations (An *et al.*, 2016; Jung *et al.*, 2017; An, kwak & Jansen, 2017). As part of an on-going research project, we have developed and are continuously improving a methodology for persona generation that uses behavioral and demographic social media data (Alowibdi, Buy & Yu, 2013b; Asoh & Ono, 2012; An & Weber, 2016). The methodology so far has shown that personas can be automatically generated using real user data aggregated from major online social media platforms, such as YouTube and Facebook. Even though the approach is flexible and resilient for application in a variety of contexts, it requires aggregated demographic information of the users' interaction with individual content, which is only provided by some social media channels (e.g., YouTube). Some social media channels, particularly Twitter, do not provide the aggregated users demographics at the individual content level, which are required to automatically generate personas. Therefore, although we can gather behavioral data from Twitter, there is a limitation in deriving demographical data. To overcome this restriction, we are exploring alternative ways to get users demographics from these channels. While some of the prior studies dealing with inferring demographic information from Twitter profiles have used profile images and facial recognition technologies, such as Face++, determining the reliability and accuracy of these techniques remains challenging because of the unavailability of ground-truth for demographic classification (Dhifli & Diallo, 2016) Yet, prior to using methods for determining user demographics, we must assess the reliability and accuracy of such methods.

In this study, we focus on evaluating the effectiveness of facial recognition for reliably inferring three demographic attributes; gender, age, and race, from profile pictures of users' social media profiles. We use images from Twitter accounts in order to determine if profiles images are a reliable source of actual people pictures. Included in this inquiry, we evaluate the ability of the well-known facial recognition technology, Face++, to analyze profile pictures to identify faces within the images and to generate accurate predictions of demographic attributes from those photos that we know do contain images of people. Using the Twitter profile pictures, we determine what percentage of these actually contains pictures of people. We use both Face++ and human annotators in our analysis, comparing the results of each.

### **PRIOR WORKS**

Given the need for gathering demographic information, researchers have investigated methods for automatic demographic determination, with efforts focused on different social media platforms, especially Twitter (Alowibdi, Buy & Yu, 2013a; 2013b; An & Weber, 2016). Prior work in this area of automatic demographic recognition has utilized user posts (Alowibdi, Buy & Yu, 2013a; 2013b; Asoh & Ono, 2012), profile names (Alowibdi, Buy & Yu, 2013a; 2013b), user locations (Beretta *et al.*, 2015), social networks (Bi *et al.*, 2013), and preferences (Huang, Weber & Vieweg, 2014). In this research, we are specifically interested in those focusing on gender, age, and racial identification.

Concerning gender, prior studies concerning users' posts/tweets proposed gender recognition methods reflecting linguistic patterns (Alowibdi, Buy & Yu, 2013b). Some studies utilized users' names and proposed gender inference methods in linking names to ground truth data (Longley, Adnan & Lansley, 2015). Some studies used several attributes of users in combination to infer gender, such as profile image, user description and so on (Alowibdi, Buy & Yu, 2013b). Age inference has shown to be more challenging compared to gender inference (Azzopardi, Greco & Vento, 2016). Some authors recommend human annotation as a reliable method (Zhong *et al.*, 2015), while others have utilized Face++ in their research (An & Weber, 2016). More simple approaches used mentions of age or birthday from users' tweets or profile information (Asoh & Ono, 2012). Moreover, there are studies to infer race or ethnicity using several approaches. Some studies used supervised learning approaches (McCormick *et al.*, 2015). Similarly, some of the studies used facial recognition using profile images (An & Weber, 2016). Some of the studies utilized username, posts/tweets, and description separately or in combination (Zhong *et al.*, 2015).

The traditional definition of race and ethnicity is related to physical and cultural factors respectively (Alowibdi, Buy & Yu, 2013b). There are some face inference studies that use a vague definition of race or ethnicity (Dhifli & Diallo, 2016). One study, for instance, included both racial and ethnic classification to refer to ethnicity (Messias, Vikatos & Benevenuto, 2017). There is obviously some confusion in the literature with the use of terms. In our study, we use race, defined as the three major groupings of human (White, Black, Asian).

### **RESEARCH OBJECTIVES**

There are times that one wants to know demographics and other attributes of users or customers. Social media platforms are one avenue to get such data. However, many social media do not provide such demographic and related information on an individual post or user level. Additionally, there is the need to infer demographics of social media users when one cannot access to the in-house data from the social media platforms. Twitter is one of the social media platforms that does not provide aggregate

demographic information at the individual content level via the APIs. Therefore, we need to leverage other methods to automatically infer the demographic data, specifically from profile images of Twitter users. Among other methods, we use profile images because this single method can infer all three demographic traits (age, gender, and race) at once.

In support of this a broader effort, this research examines:

- Research Question 1: What is the effectiveness of user's social media account profile pictures as a source to infer user demographic attributes of gender, age, and race?
- Research Question 2: What is the effectiveness of using facial recognition of user's social media account profile pictures to infer user demographic attributes of gender, age, and race?

For research question 1, we investigate if users' social media profile images are a reliable source to determine demographic information. To address this question, we take a large sample of Twitter account profile pictures and determine which images contain the picture of a person. Obviously, if a sizable portion of profile images does not at least contain a picture of a person, use of social media profile pictures is not an effective method to achieve our goal. For research question 2, we investigate if the profile pictures containing an image of a person actually represent the user. To address this question, we manually review thousands of user profile images, annotating those that are obvious fakes, drawings, images of celebrities, etc. For both investigations of research questions 1 and 2, we employ the facial recognition service, Face++.

### BACKGROUND – FACIAL RECOGNITION WITH FACE++

Face++ is a platform offering computer vision technologies. It made great achievements in facial recognition research field (Fan *et al*, 2014; Messias, Vikatos & Benevenuto, 2017), with accuracy reportedly in the high ninety percent. Moreover, there are some studies to infer demographic of social media users by utilizing Face++ (Asoh & Ono, 2012; An & Weber, 2016), with this prior work showing that Face++ is useful in inferring demographic of users with user related images.




Figure 1: An example of the data Face++ returns.

Face++ provides several facial recognition APIs, such as face detection, face search, face landmark, face mask. In this study, we utilize the face detection API provided by Face++. The face detection API detects and analyzes human faces within a provided image file or URL to an image file and returns a result set in JSON format. The results include a list of faces detected within the image. Each face object consists of several objects, such as gender, age, and race. Each object within the result set has its relevant value/name pairs. The example in Figure 1 shows the data returns from the Face++ face detection API. The JSON results show age, gender, and ethnicity, although a correct term would be race. Therefore, we use the term race from this point onward. Face++ has three classification values for race, which are White, Black, and Asian. The labels and the values are those used by Face++.

### DATA AND METHODOLOGY

To address our research questions, we use a dataset of profile images of Twitter users who follow a major media corporation. We selected this group for two reasons. One, the media corporation has a worldwide audience, so the profile images are likely to be varied. Two, we wanted to demonstrate the applicability of this method for actual deployment of a system used by an organization. We then get demographic results from a Face++ analysis of these photos. For the dataset, we compare the results of Face++ to those of human annotators.

Using the Twitter REST APIs, we first retrieve tweets published in a given period by the major media channel, also collecting retweets and replies associated with these tweets. From this collection of tweets, we identify distinct users. Resulting from these steps, 12,917 interactions were collected, which consist of 7,901 retweets and 5,016 replies. We obtained 10,309 distinct users

who are involved in the interaction with the media corporation account and collected their profile images. Figure 2 shows some sample profile images to show the variety of profile images used and illustrating that profile images are not always the expected headshot. As shown in Figure 1, some users use cartoon images like (a), or celebrities' images like (f). Even if some users use their own photo, the photo might not include their face. Figure 2e, for instance, shows user's legs, and Figure 2g includes two faces repeated multiple times.

In evaluating demographic recognition, we use 6 categorized age bins, like most of the social media analysis tools provide. Those with an age inferred or from ground-truth between 13 to 17 (denoted as 13-17), 18 to 24 (18-24), 25 to 34 (25-34), 35 to 44 (35-44), 45 to 54 (45-54), 55 to 64 (54-64), and 65 or older (65+). Gender is labeled as male or female. For race, Asian, Black, and White comprise the label set, which are three racial labels that Face++ provides.



Figure 2: Sample profile pictures of Twitter users. Note: Black bars added for privacy preservation.

## RESULTS

Using the 10,309 profile pictures of Twitter users, we submitted the URLs to the Face++ face detection API. Of these 10,309 images, the Face++ API detected one or more faces in 3,566 (34.6%) pictures, with the remaining 6,743 (65.4%) containing no face detectable by the Face++ service. Already, focus back on our research question one, it is questionable whether profile images are a reliable source for determining user demographics, with only 34.6% of the profile pictures have a detectable facial image. As we shown in Figure 2 above, there are many profile pictures that do not contain an image of the user or even a person. So, it makes us question whether social media profile pictures are useful as a source of demographic data or not for the actual application development. However, we explore the images more to get more detailed results. Specifically, we manually check the two sub-datasets, the profile pictures in which Face++ detected a face and the profiles pictures where Face++ reported no face.

### Profiles Pictures With Faces Detected by Face++

In the manual coding of the 3,566 pictures where Face++ had detected, 3,177 (30.8% of the original dataset) had only one face. From our review, determining the demographics of the user from 389 images with multiple faces was problematic, as nearly all of these images contained multiple people making a determination of the actual account holder impossible from just the profile picture. During the process of the manually coding, we determined that 775 (24.4%) of the account holders had changed their profile image, so profiles images are a fairly unstable source of demographic information.

However, we then manually checked 2,402 (75.6%) of the 3,177 pictures having only one recognized face by Face++ to gauge the reliability of facial recognition technology for social media images. From our manual analysis, 80% (2,039) of the pictures had an image of a single person that could reasonably be the account holder (i.e., a real person, not a celebrity, etc.). So, this set would be utilizable for inferring user's demographic. Interestingly, even though Face++ identifies only one face in this set of profile pictures, there were still pictures with multiple people (180 images, 7.4%). Plus, there is a lot of non-user face pictures, especially memes and celebrities (62 images, 2.5%). Also, there were pictures which are not real human being (i.e. Superman, Wonder Woman, etc.) (121 images, 5.0%). These results indicate a problem utilizing automatic face recognition services because the image may not be the actual Twitter user. However, the overall accuracy of at least detecting a face was about 80%.

### Profiles Pictures With No Faces Detected by Face++

We also manually examined 2,488 (36.8%) images among the 6,743 images that had no recognized face result from Face++. Among these 2,488 user images, 1,800 (72.4%) did not contain a face. However, there were 688 users (27.6%) who were most likely using their own images, even though the faces were not detected by Face++. We examined these images in more detail by conducting a failure analysis, in which we tagged the images with various attributes, such as glasses, sunglasses, half face, blurred, back, recognizable and so on, that may have caused Face++ to not recognize the images as that of a person. Figure 3 shows some example images. Thirty-six (36) users (1.4%) showed their backside on their pictures. Eighty-five (85) users (3.4%) were wearing sunglasses. One hundred and twenty-one (121) users (4.8%) were wearing glasses. Two hundred and fifty (250) (10.0%) were showing their half face. Fifty (50) users (2.0%) took their pictures with a very small face. Twenty-four (24) users (0.9%) were laying down. Only 59 users (2.3%) had a recognizable complete facial image of the person, but these faces were not recognized by Face++ for some non-obvious reason.

As shown in Figure 3, we see the difficulties in inferring accurate demographics. In many cases, even a human cannot accurately judge the user demographics. For example, the image tagged with 'back', in which we can possibly conclude the person's gender might be female judging by the person's hairstyle. The image tagged by 'half' seems difficult for Face++ to be recognized, although we can make a reasonable guess it is a male from human evaluation. Figure 3 provides other examples of difficult to classify profile images.







Image	Tags	Image	Tags
	back		blurred, glasses, blocked
	half		half, blocked, sun- glasses
	small, glasses		sunglasses

Figure 3: Some unrecognizable profile images of Twitter users with manual tags.

### DISCUSSION AND IMPLICATIONS

In this work, we examined the potential of Face++, a popular facial recognition technology, for inferring demographic information with four datasets. Returning to our first research question, (Are user social media profile images a reliable source to determine demographic information?), we have to conclude that profile images are by themselves not enough to be a reliable source of inferring demographic information, with only about thirty percent of users have an image with a face. Concerning our second research question (If the profile picture does contain an image of person, does the picture actual represent the user?), of the images that did contain a face, the automatic facial recognition technology was about eighty percent accurate it actually detects a face and these images were, generally, that of the account user. Overall, only using facial recognition technology would not be a good approach for inferring users' demographics. To effectively utilize overall demographics of the social media users, we need to employ various features like first name, color theme of the user profile, social media posts (e.g. tweets), user's preferences and so on what we introduced studies in the literature section.

### CONCLUSION

Our research investigates the effectiveness of facial recognition technology and the accuracy of Face++ in a real-world system development setting. Results show that only a minority of users use their own image as their profile pictures, dropping reliability of this method to about thirty percent. However, when there is an image with a face, Face++ has strong accuracy in determining a face and that image generally appears to be that of the account holder. In future work, we are exploring other methods of determining user demographics, which we will combine with image facial recognition to increase demographical identification.

## REFERENCES

- [1] Alowibdi, J.S., Buy, U.A. & Yu, P. (2013). Empirical evaluation of profile characteristics for gender classification on twitter. In *Proceedings of the 12th International Conference on Machine Learning and Applications* (pp. 365-369). IEEE, Miami, FL, USA, December 4-7.
- [2] Alowibdi, J.S., Buy, U.A. & Yu, P. (2013). Language independent gender classification on Twitter. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining* (pp. 739-743). ACM, Niagar Falls, ON, Canada, August 25-28.
- [3] An, J. & Weber, I. (2016). #greysanatomy vs. #yankees: Demographics and hashtag use on Twitter. In *the 10th International AAAI Conference on Web and Social Media – ICWSM-16*. AAAI, Cologne, Germany, May 17-20.
- [4] An, J., Cho, H., Kwak, H., Hassen, M. Z. & Jansen, B. J. (2016). Towards automatic persona generation using social media. In *IEEE International Conference on Future Internet of Things and Cloud Workshops (FiCloudW)* (pp. 206-211). IEEE, Vienna, Austria, Aug 22-24.
- [5] An, J., Kwak, H. & Jansen, B.J. (2017). Personas for content creators via decomposed aggregate audience statistics. In *Proceedings of the International Conference on Advances in Social Network Analysis and Mining (ASONAM 2017)* (pp. 632-635). IEEE/ACM, Sydney, Australia, July 31-August 03.
- [6] Asoh, H., Ikeda, K. & Ono, C. (2012). A fast and simple method for profiling a population of Twitter users. In *Proceedings of the Third International Workshop on Mining Ubiquitous and Social Environments* (pp. 19-26). MUSE, Bristol, UK, September 24.
- [7] Azzopardi, G., Greco, A. & Vento, M. (2016). Gender recognition from face images using a fusion of svm classifiers. In *International Conference on Image Analysis and Recognition* (pp 533-538). ICIAR, Póvoa de Varzim, Portugal, July 13-15.
- [8] Beretta, V., Maccagnola, D., Cribbin, T. & Messina, E. (2015). An interactive method for inferring demographic attributes in Twitter. In *Proceedings of the 26th ACM Conference on Hypertext & Social Media* (pp. 113-122). ACM, Ankara, Turkey, September 01-04.
- [9] Bi, B., Shokouhi, M., Kosinski, M. & Graepel, T. (2013). Inferring the demographics of search users: Social data meets search queries. In *Proceedings of the 22nd International Conference on World Wide Web* (pp. 131-140). ACM, Brazil, May 13-17.
- [10] Dhifli, W. & Diallo, A.B. (2016). Face recognition in the wild. *Procedia Computer Science*, 96, 1571-1580.
- [11] Fan, H., Yang, M., Cao, Z., Jiang, Y. & Yin, Q. (2014). Learning compact face representation: Packing a face into an int32. In *Proceedings of the 22nd ACM international conference on Multimedia* (pp. 933-936). ACM, Orlando, FL, USA, November 03-07.
- [12] Huang, W., Weber, I. & Vieweg, S. (2014). Inferring nationalities of Twitter users and studying inter-national linking. In *Proceedings of the 25th ACM conference on Hypertext and Social Media* (pp. 237-242). ACM, Santiago, Chile, September 01-04.
- [13] Jung, S.G., An, J., Kwak, H., Ahmad, M., Nielsen, L. & Jansen, B.J. (2017). Persona generation from aggregated social media data. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (pp. 1748-1755). ACM, Denver, CO, USA, May 06-11.
- [14] Longley, P.A., Adnan, M. & Lansley, G. (2015). The geotemporal demographics of Twitter usage. *Environment and Planning A*, 47(2), 465-484.
- [15] McCormick, T.H., Lee, H., Cesare, N., Shojaie, A. & Spiro, E.S. (2015). Using Twitter for demographic and social science research: Tools for data collection and processing. *Sociological Methods & Research*, 46(3), 390-421
- [16] Messias, J., Vikatos, P. & Benevenuto, F. (2017). White, man, and highly followed: Gender and race inequalities in Twitter. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence (WI'17)* (8 pages). IEEE, Leipzig, Germany. August 23-26.
- [17] Salminen, J. O., Şengün, S., Kwak, H., Jansen, B. J., An, J., Jung, S. G., Vieweg, S. & Harrell, F. (2017). Generating cultural personas from social data: A perspective of middle eastern users. In *Proceedings of the Fourth International Symposium on Social Networks Analysis, Management and Security* (pp. 120-125). SNMAS, Prague, Czech Republic, Aug 21-23.
- [18] Vikatos, P., Messias, J., Miranda, M. & Benevenuto, F. (2017). Linguistic diversities of demographic groups in Twitter. In *Proceedings of the 28th ACM Conference on Hypertext and Social Media* (pp. 275-284). ACM, Prague, Czech Republic, July 04- 07.
- [19] Zagheni, E., Garimella, V.R.K. & Weber, I. (2014). Inferring international and internal migration patterns from twitter data. In *Proceedings of the 23rd International Conference on World Wide Web* (pp. 439-444). ACM, Seoul, Republic of Korea, April 07-11
- [20] Zhong, Y., Yuan, N.J., Zhong, W., Zhang, F. & Xie, X. (2015). You are where you go: Inferring demographic attributes from location check-ins. In *Proceedings of the 8th ACM International Conference on Web Search and Data Mining* (pp. 295-304). ACM, Shanghai, China, January 31-February 06.