

Characterizing the Demographics Behind the #BlackLivesMatter Movement

Alexandra Olteanu
EPFL
alexandra.olteanu@epfl.ch

Ingmar Weber
QCRI
iweber@qf.org.qa

Daniel Gatica-Perez
Idiap and EPFL
gatica@idiap.ch

Abstract

The debates on minority issues are often dominated by or held among the concerned minorities: gender equality debates have often failed to engage men, while those about race fail to engage the dominant group. To test this observation, we study the **#BlackLivesMatter** movement and hashtag on Twitter—that has emerged and gained traction after a series of events typically involving the death of African-Americans as a result of police brutality—aiming to quantify the *population biases* across user types (individuals vs. organizations), and (for individuals) across 3 demographics factors (*race, gender and age*). Our results suggest that more African-Americans engage with the hashtag, and that they are also more active than other demographic groups. We also discuss ethical caveats with broader implications for studies on sensitive topics (e.g. mental health or religion) that focus on users.

Introduction

While the growing number of discussions about minority¹ issues—including gender (O’Brien and Kelly 2013), income (Moodie-Mills 2015), or race (Lashinsky 2015)—is good news, empirical evidence suggests that they are held mainly among the discriminated group: women dominate the debate on gender (Royles 2014), while African-Americans dominate the one on race (Pettit 2006). Although social media has led to a paradigm shift for advocacy by increasing the effectiveness, the speed and the outreach of social campaigns, *many still fail to reach far beyond the communities for which they advocate*.

In this paper, we explore this observation in the context of the **#BlackLivesMatter** movement² on Twitter. We want to gain insights into the level of involvement across user demographics. What can be said about the demographic composition of the communities engaged in the discussions? Does the discriminated group dominate the debate? Ultimately, engaging diverse stakeholder groups is beneficial for the social campaign’s success (Ward 2013), and knowing the extent to which they contribute to the debate is helpful in learning how to alter the message to appeal to them.

Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Throughout the paper, by *minority* we refer to a group that is subordinate to a more dominant group in society.

²<http://blacklivesmatter.com/contact/>

#BlackLivesMatter is a movement (and a hashtag) created after the killing of Trayvon Martin in 2012, with over 1,000 demonstrations being held since then.³ The hashtag has been used during a number of events involving disproportionate police violence against African-Americans, as well as disproportionate reaction of mainstream media when terror attacks occur in Western countries compared to when they occur in African countries (Zuckerman 2015).

Contributions. Our main contribution is a *demographic characterization of users involved in the #BlackLivesMatter movement on Twitter*. Our findings suggest that African-Americans are both more numerous and active than other demographic groups. Young females are more likely to actively engage in the debate than men, yet, the proportions of white and African American females are similar. Looking at male users, we see a slightly different pattern: young adults still dominate the discussions, but they are largely African Americans. Contrasting individuals to organizations, amounting for ~5% of profiles, we see a 3 times higher tweeting rate.

To run this study, we also created a *collection of about 6,000 Twitter users annotated with demographic information such as race, age, and gender*. In contrast with previous work that reports demographic information by automatically predicting demographic factors for each user based e.g. on their profile picture or name (Minkus, Liu, and Ross 2015; Zagheni et al. 2014; Bakhshi, Shamma, and Gilbert 2014; Mislove et al. 2011), we crowdsourced these annotations. Although more expensive, we do so to work around known pitfalls of automated user classification such as low recall (Minkus, Liu, and Ross 2015) and classification errors (Yadav et al. 2014).

Limitations and Ethical Challenges. We note that such an endeavor is not without caveats. First, there are intrinsic issues with hashtag-based analyses, and the reliance on a single media platform and public APIs (Tufekci 2014; Boyd and Crawford 2012): The hashtag we focus on does not cover all the discussions and contributions around the issue at core. The movement and hashtag use are recent and we cannot capture the long-term evolution of the demographics behind the core debate.

Second, there are important ethical challenges (Boyd and Crawford 2012): Although publicly available, user profile

³<https://www.elephrame.com/textbook/protests>

Table 1: The basic stats of our dataset.

Movement	Tweets	Users	Start Day	End Day
#BlackLivesMatter	3.54M	0.88M	April 11, 2012	May 10, 2015

data is inherently sensitive as e.g. users might not anticipate a particular use of their data, especially when created in a context sensitive space and time. This becomes even more delicate when explicitly analyzing their demographic attributes. We discuss these challenges as we detail our methods and their implications.

Data Collection and Annotation

The Movement On Twitter. The #BlackLivesMatter hashtag (whose usage over time is shown in Figure 1) was first used on Twitter on April 2012 in relation to the killing of Trayvon Martin (Graeff, Stempeck, and Zuckerman 2014). Yet, it grew into a movement only after the acquittal of George Zimmerman (the man who fatally shot Martin) in July 2013,⁴ and got consistent traction after the killing of Michael Brown and with the Ferguson unrest.⁵ The movement gained momentum after the killing of Tamir Rice,⁶ a 12 year-old school boy, and the decision of a grand jury not to indict the officer that put Eric Garner in a chokehold.⁷ Since then, the movement periodically regained public attention after events involving police brutality, including the deaths of Walter Scott⁸ and Freddie Gray.⁹

Collecting Tweets. To collect tweets published from the day before the first use of the hashtag¹⁰ until May 10, 2015, we crawled Topsy¹¹ in April-May 2015—dataset figures in Table 1 and Figure 1. To maximize the coverage of our collection, we repeated the crawling with various time window sizes until its volume converged.

Data Collection and Annotation. User data (public profile data and crowdsourced annotations) were collected in June 2015. User profiles were annotated according to the entity behind the Twitter accounts via the crowdsourcing platform Crowdflower¹². We asked crowdworkers to categorize users as individuals, governmental agencies, NGOs, media, others; and, then, the individuals according to 3 perceived demographic attributes: race, age and gender. Crowdworkers were shown automatically generated screenshots of the upper part of users’ public profiles, including the picture banner, the profile picture, the name and profile description, and the last one or two tweets. The screenshots were provided via *short-lived* URLs in order to limit access to user profile information and minimize the risk of privacy violations.

⁴http://en.wikipedia.org/wiki/Black_Lives_Matter

⁵http://en.wikipedia.org/wiki/Shooting_of_Michael_Brown

⁶http://en.wikipedia.org/wiki/Shooting_of_Tamir_Rice

⁷http://en.wikipedia.org/wiki/Death_of_Eric_Garner

⁸http://en.wikipedia.org/wiki/Shooting_of_Walter_Scott

⁹http://en.wikipedia.org/wiki/Death_of_Freddie_Gray

¹⁰First tweet containing a term obtained via <http://ctrlq.org/first/>

¹¹<http://about.topsy.com/terms-and-conditions/>

¹²<https://crowdflower.com/>




Figure 2: The distribution of number of tweets per user.

We annotated about 6,000 users from 6 random samples with various characteristics (e.g. from all users, from highly active ones, from users tweeting about the topic even when the media attention fades away). We showed crowdworkers 5-6 user profiles at a time, out of which one profile was labeled by one of the authors (gold standard), and used to control the quality of the annotations. Given that we collect perceived attributes and some of them might be subjective, the profiles picked as gold standards were selected to be obvious cases for each of the categories. For all annotation jobs, we collected at least 3 independent annotations for each profile and categorization criteria, and kept the majority label. About 100 crowdworkers participated in each task. Full annotation instructions are included in our data release.

Exploratory Analysis

The users distribution according to the number of tweets¹³ is long tailed (Figure 2): most users post only a few tweets on the topic (e.g. ~62% of users have only one tweet in the collection), while only a few users post in the order of thousands of tweets (only 3 users with more than 10K tweets). This indicates that many users participate in the debate only incidentally. For our analysis, we split users according to their level of activity in 3 categories: a) non-active users—769,231 users with less than 5 tweets; b) moderately active users—96,905 users with 5 to 25 tweets; and c) highly active users—14,033 users with more than 25 tweets. We make this categorization as we conjecture that the activity w.r.t. a topic is a proxy for a user’s interest in the topic and her level of involvement, and we are interested in the interplay between the activity level and users demographics.

Further, by briefly exploring the triggers behind the peaks of attention received by the movement,¹⁴ we find that most of them are generated by events involving killing of African-Americans by police in the US (when the debate focuses on the discrimination against African-Americans), see Figure 1. In addition, the attention peaks for a topic may be indicative of the topic entering and exiting the public debate: when the topic is in the spotlight, a larger community tends to get involved in the debate, yet, as the topic fades away, only the concerned community might care. To this end, we define a peak window (or the spotlight interval) as a 4 days interval including the day of the peak, the day before the peak, and two days after the peak. Using this definition, we found

¹³For simplicity, by tweet(ing) we refer to both the creation of an original tweet, as well as to passing on content, i.e. re-tweeting.

¹⁴To detect peaks we used a readily available implementation: <https://gist.github.com/endolith/250860#file-peakdet-m>




Figure 1: The distribution of the volume of tweets for #BlackLivesMatter per day over time.

Table 2: Accounts of organizations vs. of individuals across samples. Asterisks indicate stat. signif. differences w.r.t. the distribution of all users at $p < 0.01$ (**) and $p < 0.05$ (*)

	All Users	Peak	Non Peak	High Activ.	Mod. Activ.	Low Activ.
Org.	5.0%	4.6%	4.9%	11.1%	5.5%	4.2%
Indiv.	95.0%	95.4%	95.1%	88.9%	94.5%	95.8%

**
*
**


611,871 users tweeting in the peak times, as compared to less than half of that number being active before the topic “enters” or after it “exits” the public debate—268,298 users.

User Characterization


To study the demographic composition of users involved in the debate we extracted 6 random samples¹⁵: 2,000 users sampled from all users in our dataset, and 5 samples of 1,000 users from: users tweeting during peak times, users tweeting outside the peak times, highly active users, moderately active users, and non-active users. The samples were labeled in two rounds: the first annotation task aimed to distill the accounts of individuals from those of organizations, while the second task was designed to categorize accounts of individuals along 3 demographic criteria: race, gender and age.

Accounts of Organizations. Looking at the fraction of organization accounts w.r.t those of individuals, we notice that the sample drawn from highly active users contains twice as many organization accounts than other samples. The fraction of organization accounts seems typically higher within more active users: e.g. there are more organization accounts among moderately active users than among non-active users. This is largely explained by a higher fraction of accounts associated with NGOs (7.4%, 3.6%, 1% for highly active, moderately active and non-active users, and 2.2% across all users) and media organizations, which, however, attains the highest fraction among moderately active users (a possible artifact of the fact that media organizations tweet about many topics, while NGOs are typically focused on a handful of causes). Finally, accounts associated with governmental agencies account for less than half a percent in all samples.


¹⁵Due to technical limitations related to how the screenshots were displayed—resulting in profiles not being shown correctly for annotation—we were able to label only 5976 users.



(a) Distribution of users' age per sample.



(b) Distribution of users' race per sample.




(c) Distribution of users' gender per sample.

Figure 3: Race, age, and gender distribution across samples. Asterisks indicate stat. signif. differences w.r.t. the distribution of all users at $p < 0.01$ (**) and $p < 0.05$ (*). (best seen in color)

User Demographics. For individuals, we looked at the distribution of demographic factors. (*Age*) Figure 3(a) shows that the fraction of young adults is lower among highly active users, while the fraction of adults between 30 to 64 years old is lowest outside the peak times—these users engaging with the hashtag more actively during peak times when the topic is in the public spotlight. (*Race*) Figure 3(b) shows user distribution across racial groups and samples. We notice that the fraction of African-Americans is the highest within the sample of highly active users, and the smallest among the non-active users or during peak times. (*Gender*) Finally, in Figure 3(c) we see that the user distribution according to their gender is relatively stable across samples.

Next, we looked at the distribution of users across age



(a) Distribution of male users as a function of age and race. All cells sum to 100%.



(b) Male (M) to female (F) ratio. Red (resp. blue) indicates a higher fraction of female (resp. male) users w.r.t. the overall distribution (~ 0.78 marked by white in the colorbar). The percentages indicate the overall distribution of users.

Figure 4: Race and age distribution for female vs. male users. (best seen in color)

and race per gender—see Figure 4.¹⁶ We notice that the most active users are white and African-American adults between 18 to 64 years old. However, while for male African-American users the fraction of young adults (18 to 29 years old) is higher, for white users it is lower. Inspecting the differences between genders (Figure 4(b)), we see that women younger than 29 years old are more active than men in the same age category, while for users older than 30 years old, men tend to tweet more about the movement.

User Involvement. Finally, we checked if users belonging to specific demographic groups tend to be more vocal, or, in other words, if they generate more content on average. First, we find that organizations are more active than individuals (7:2). Then, depending on the demographic criteria, we see that: (a) African-Americans are most active, followed closely by white users; (b) women are more active than men (3.8:2.6); and (c) adults between 30 and 64 years old are the most active, followed by young adults (3.9:2.6:2).

Concluding Remarks

We started the study after one of the related events—the shooting of Walter Scott—and based on empirical evidence we hypothesized that the debate would be held largely among African-Americans. While our findings support this premise, African-Americans being the largest group (even up to 60% among highly active users), overall, whites make up about 40% of individuals and Asians 4%. Future work naturally includes an analysis of demographic factors across various movements related to minority groups issues in order to validate and broaden the observations we make here.

Parting Thoughts on Ethics. Although important, studies investigating social media to understand the public opinion and various narratives on minority issues across stakehold-

¹⁶This is based on users annotated along all demographic factors, as only some factors may be perceptible based on user profile info.

ers are scant, but growing. One reason are the limits in collecting and annotating users accurately and at scale (either manually or automatically). Yet, as we learn to work around these limits, *we also need to develop protocols to mindfully study such user collections while protecting the users.*

Data Release. The list of tweet ids are available for research purposes at <http://crisislex.org/>. The list of annotated users is available upon signing for **not** using it to study users in isolation or to single them out for their demographic attributes.

Acknowledgements. We thank Carlos Castillo for feedback on an early draft. A.O. was partially supported by the grant *Sinergia (SNF 147609)*.

References

- Bakhshi, S.; Shamma, D. A.; and Gilbert, E. 2014. Faces engage us: Photos with faces attract more likes and comments on instagram. In *CHI*.
- Boyd, D., and Crawford, K. 2012. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication and society*.
- Graeff, E.; Stempeck, M.; and Zuckerman, E. 2014. The battle for ‘Trayvon Martin’: Mapping a media controversy online and offline. *First Monday*.
- Lashinsky, A. 2015. Seven signs you are clueless about income inequality. <http://fortune.com/2015/03/20/anand-giridharadas-ted-inequality/>.
- Minkus, T.; Liu, K.; and Ross, K. W. 2015. Children seen but not heard: When parents compromise children’s online privacy. In *WWW*.
- Mislove, A.; Lehmann, S.; Ahn, Y.-Y.; Onnela, J.-P.; and Rosenquist, J. N. 2011. Understanding the demographics of twitter users. *ICWSM*.
- Moodie-Mills, D. 2015. Black lives matter: A tale of two covers. <http://www.nbcnews.com/news/nbcblk/black-lives-matter-tale-two-covers-n339796>.
- O’Brien, S., and Kelly, T. 2013. Gender equality won’t happen unless men speak up. <http://edition.cnn.com/2013/04/17/business/sandberg-gender-equality/>.
- Pettit, J. 2006. Can we talk about race? a few rules of engagement. http://articles.baltimoresun.com/2006-08-01/news/0608010135_1_racial-inequality-political-change-problem-of-racial.
- Royles, D. 2014. What’s missing from the debate about women leaders in the nhs? men. <http://www.theguardian.com/healthcare-network/2014/jan/08/female-managers-gender-equality-nhs>.
- Tufekci, Z. 2014. Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In *ICWSM*.
- Ward, J. A. 2013. The next dimension in public relations campaigns: A case study of the it gets better project. *Public Relations Journal*.
- Yadav, D.; Singh, R.; Vatsa, M.; and Noore, A. 2014. Recognizing age-separated face images: Humans and machines. *PLoS one*.
- Zagheni, E.; Garimella, V. R. K.; Weber, I.; et al. 2014. Inferring international and internal migration patterns from twitter data. In *WWW companion publication*.
- Zuckerman, E. 2015. Paying attention to Garissa. <http://www.ethanzuckerman.com/blog/2015/04/04/paying-attention-to-garissa/>.