

Using Noisy Self-Reports to Predict Twitter User Demographics

Zach Wood-Doughty*, Paiheng Xu*, Xiao Liu, Mark Dredze

Department of Computer Science

John Hopkins University, Baltimore, MD 21218

zach@cs.jhu.edu, paiheng@jhu.edu, xliu119@jhu.edu, mdredze@cs.jhu.edu

Abstract

Computational social science studies often contextualize content analysis within standard demographics. Since demographics are unavailable on many social media platforms (e.g. Twitter), numerous studies have inferred demographics automatically. Despite many studies presenting proof-of-concept inference of race and ethnicity, training of practical systems remains elusive since there are few annotated datasets. Existing datasets are small, inaccurate, or fail to cover the four most common racial and ethnic groups in the United States. We present a method to identify self-reports of race and ethnicity from Twitter profile descriptions. Despite the noise of automated supervision, our self-report datasets enable improvements in classification performance on gold standard self-report survey data. The result is a reproducible method for creating large-scale training resources for race and ethnicity.

1 Introduction

Contextualization of population studies with demographics forms a central analysis method within the social sciences. In domains such as political science or public health, standard demographic panels in telephone surveys enable better analyses of opinions and trends. Demographics such as age, gender, race, and location are often proxies for important socio-cultural groups. As the social sciences increasingly rely on computational analyses of online text data, the unavailability of demographic attributes hinders comparison of these studies to traditional methods (Al Baghal et al., 2020; Amir et al., 2019; Jiang and Vosoughi, 2020).

Computational social science increasingly utilizes methods for the automatic inference of demographic attributes from social media, such as Twitter (Burger et al., 2011; Chen et al., 2015; Ardehaly and Culotta, 2017; Jung et al., 2018; Huang and Paul, 2019). Demographics factor into social media studies across domains such as health, politics, and linguistics (O’Connor et al., 2010; Eisenstein et al., 2014). Off-the-shelf software packages support the inference of gender and location (Knowles et al., 2016; Dredze et al., 2013; Wang et al., 2019).

Unlike age or geolocation, race and ethnicity are sociocultural categories with competing definitions and measurement approaches (Comstock et al., 2004; Vargas and Stainback, 2016; Culley, 2006; Andrus et al., 2021). Despite this complexity, understanding race and ethnicity is crucial for public health research (Coldman et al., 1988; Dressler et al., 2005; Fiscella and Fremont, 2006; Elliott et al., 2008, 2009). Analyses that explore mental health on Twitter (Loveys et al., 2018) should consider racial disparities in healthcare (Satcher, 2001; Amir et al., 2019) or online interactions (Delisle et al., 2019; Burnap and Williams, 2016). Despite the importance of race and ethnicity in these studies, and multiple proof-of-concept classification studies, there are no readily-available systems that can infer demographics for the most common United States racial/ethnic groups. This gap arises from major limitations for all publicly-available data resources.

A high-quality dataset for this task has several desiderata. First, it should cover enough categories to match standard demographics panels. Second, the dataset must be sufficiently large to support training

* Equal contribution

Citation	Annotation	% Missing	# Users	% W	% B	% H/L	% A
Preoțiu-Pietro et al. (2015)	Survey	4.7	3572	80.8	9.5	6.1	3.6
Culotta et al. (2015)	Crowdsourced	60.0	308	50.0	19.5	30.5	0
Volkova and Bachrach (2015)	Crowdsourced	36.5	3174	48.0	35.8	8.9	3.0
Total Matching Users	Self-report	-	2.50M	26.8	53.8	11.3	8.1
Query-Bigram	Self-report	8.1	112k	51.2	40.8	1.4	6.6
Heuristic-Filter	Self-report	40.6	135k	42.2	45.9	5.6	6.4
Balanced-Group-Person	Self-report	0.0	31k	25.0	25.0	25.0	25.0

Table 1: Previously-published Twitter datasets annotated for race/ethnicity and datasets collected in this work. “% Missing” shows the percent of users that could not be scraped in 2019. “# Users” shows the number users that are currently available. The abbreviations W, B, H/L, and A corresponds to White, Black, Hispanic/Latinx, Asian respectively, which we use for the rest of the paper. Per-group percentages are from non-missing data.

accurate systems. Third, the dataset should be reproducible; Twitter datasets shrink as users delete or restrict accounts, and models become less useful due to domain drift (Huang and Paul, 2018).

We present a method for automatically constructing a large Twitter dataset for race and ethnicity. Keyword-matching produces a large, high-recall corpus of Twitter users who potentially self-identify as a racial or ethnic group, building on past work that considered self-reports (Mohammady and Culotta, 2014; Beller et al., 2014; Coppersmith et al., 2014). We then learn a set of filters to improve precision by removing users who match keywords but do not self-report their demographics. Our approach can be automatically repeated in the future to update the dataset. While our automatic supervision contains noise – self-descriptions are hard to identify and potentially unreliable – our large dataset demonstrates benefits when compared to or combined with previous crowdsourced datasets. We validate this comparison on a gold-standard survey dataset of self-reported labels (Preoțiu-Pietro and Ungar, 2018). We release our code publicly¹. We also release our collected datasets and trained models to researchers with approval from an IRB or similar ethics board, contingent on compliance with our data usage agreement².

2 Ethical Considerations

Complexities of racial identity raise ethical considerations, requiring discussion of the

benefits and harms of this work (Benton et al., 2017). The benefits are clear in settings such as public health; many studies use social media data to research health behaviors or support health-based interventions (Paul and Dredze, 2011; Sinnenberg et al., 2017). These methods have transformed areas of public health which otherwise lack accessible data (Ayers et al., 2014). Aligning social media analyses with traditional data sources requires demographic information.

The concerns and potential harms of this work are more complex. Ongoing discussions in the literature concern the need for informed consent from social media users (Fiesler and Proferes, 2018; Marwick and boyd, 2011; Olteanu et al., 2019). Twitter’s privacy policy states that the company “make[s] public data on Twitter available to the world,” but many users may not be aware of the scope or nature of research conducted using their data (Mikal et al., 2016). Participant consent must be *informed*, and we should study users’ comprehension of terms of service when conducting sensitive research. IRBs have applied established human subjects research regulations in ruling that passive monitoring of social media data falls under public data exemptions.

While our data usage agreement prohibits such behavior, a malicious actor could attempt to use predicted user demographics to track or harass minority groups. Despite the severity of such a worst-case scenario, there are two arguments why the benefits may outweigh the harms. First, if open-source methods and models were used for such malicious behavior, platform moderators could simply

¹<https://bitbucket.org/mdredze/demographer>

²<http://www.cs.jhu.edu/~mdredze/demographics-training-data/>

incorporate those tools into combatting any automated harassment. Second, harassment against historically disenfranchised groups is already extremely widespread. Open-source tools would provide more good than harm in the hands of researchers or platform moderators (Jiang and Vosoughi, 2020). Recent work has shown that women on Twitter, especially journalists and politicians, receive disproportionate amounts of abuse (Delisle et al., 2019). On Facebook, advertisers have used the platform’s knowledge of users’ racial identities to illegally discriminate when posting job or housing ads (Benner et al., 2019; Angwin and Parris Jr, 2016). To protect against misuse of our work, we follow Twitter’s developer terms which prohibit efforts to “target, segment, or profile individuals” based on several sensitive categories, including racial or ethnic origin, detailed in our data use agreement. Predictions should not be analyzed to profile individual users but rather must only be used for aggregated analyses.

Another concern of any predictive model for sensitive traits is that a descriptive model could be interpreted as a prescriptive assessment (Ho et al., 2015; Crawford, 2017). Individual language usage may also differ from population-level demographics patterns (Bamman et al., 2014). Additionally, our datasets and models do not cover smaller racial minorities (e.g. Pacific Islander) or the fine-grained complexities of mixed-race identities. More fine-grained methods are needed for many analyses, but current methods cannot support them.

Finally, we distinguish between biased models and biased applications. Our models are imperfect; if we only analyze a small sample of users and our models have high error rates, a difference that appears significant may be an artifact of misclassifications. Any downstream application must account for this uncertainty.

On the whole, we believe demographic tools provide significant benefits that justify the potential risks in their development. We make our data available to other researchers, but with limitations. We require that researchers comply with a data use agreement and obtain approval by an IRB or similar

ethics committee. Our agreement restricts these tools to population-level analyses³ and **not** the analysis of individual users. We exclude certain applications, such as targeting of individuals based on race or ethnicity. Any future research that makes demographically-contextualized conclusions from classifier predictions must explicitly consider ethical trade-offs specific to its application. Finally, our analysis of social media for public health research has been IRB reviewed and deemed exempt (45 CFR 46.101(b)(4)).

3 Datasets for Race and Ethnicity

Our tools and analysis focus on the United States, where recognized racial categories have varied over time (Hirschman et al., 2000; Lee and Tafoya, 2006). Current US census – and many surveys – record self-reported racial categories as White, Black, American Indian, Asian, and Pacific Islander. Surveys often frame ethnicity as Hispanic/Latinx origin or not; however, there is not necessarily a clear distinction between race and ethnicity (Gonzalez-Barrera and Lopez, 2015; Campbell and Rogalin, 2006; Cornell and Hartmann, 2006). Individuals may identify as both a race and an ethnicity, and 2% of Americans identify as multi-racial (Jones and Smith, 2001). Because of the limited data availability, we only consider the four largest race/ethnicity groups, which we model as mutually exclusive: White, Black, Asian, and Hispanic/Latinx. Our methodology could be extended to be more comprehensive, but we do not yet have the means to validate more fine-grained or intersectional approaches.

Table 1 lists three published datasets for race/ethnicity. Since only user ids can be shared, user account deletions over time cause substantial missing data. Past work has taken varied approaches to annotate racial demographics. Culotta et al. (2015) and Volkova and Bachrach (2015) relied on manual annotation, noting inter-annotator agreement estimated at 80% and Cohen’s κ of 0.71, respectively. Crowdsourced annotation

³Twitter’s API “restricted use cases” explicitly permit aggregated analyses.

	Raw	Color	Plural	Bigram	Quote	All
Precision	76.7	78.6	76.7	82.5	78.6	86.8
Removed by filter	-	314k	212k	281k	4k	784k

Table 2: Applying our HF filters (§ 4) individually and together. Precision is on dev set from Appendix B, after thresholding on self-report score.

assumes that racial identity can be accurately perceived by others, an assumption that has serious flaws for gender and age (Flekova et al., 2016; Preotiuc-Pietro et al., 2017). Rule-based or statistical systems for data collection can be effective (Burger et al., 2011; Chang et al., 2010), but raise concerns about selection bias: if we only label users who take a certain action, a model trained on those users may not generalize to users who do not take that action (Wood-Doughty et al., 2017).

Gold-standard labels for sensitive traits requires individual survey responses, but this yields small or skewed datasets due to the expense (Preotiuc-Pietro and Ungar, 2018). Our approach instead relies on automated supervision from racial self-identification and minimal manual annotation to refine our dataset labels. We are not the first to use users’ self-identification to label Twitter users’ demographics, but past work has relied heavily either on restrictive regular expressions or manual annotation (Pennacchiotti and Popescu, 2011; Mohammady and Culotta, 2014). Such work has also been limited to datasets of under 10,000 users. We expand on previous work to construct a much larger dataset and evaluate it via trained model performance on ground-truth survey data.

4 Data Collection of Self-Reports

We construct a regular expression for terms associated with racial identity. We select tweets from Twitter’s 1% sample from July 2011 to July 2019 in which the user’s profile description contains one of the following racial keywords in English: `black`, `african-american`, `white`, `caucasian`, `asian`, `hispanic`, `latin`, `latina`, `latino`, `latinx`. While there are other terms that signify racial identity, these match common

survey panels (Hirschman et al., 2000) and our empirical evaluation is limited because our survey dataset only covers four classes. We omit self-reports that indicate a country of origin (e.g. “Colombian” or “Chinese-American”), smaller racial minorities (e.g. “Native American” or “two or more races”), or more ambiguous terms, leaving such groups for future work. If a user appears multiple times, we use their latest description.

We select users whose profile descriptions contain a query keyword, which heavily skews towards color terms (“white”, “black”). This produces 2.67M users, 2.50M of which match exactly one racial/ethnic class (Table 1, “Total Matching Users”). While this is several orders of magnitude larger than existing datasets, many user descriptions that match racial keywords are not racial self-reports. We next consider approaches to filter these users’ profile descriptions to obtain three self-report datasets of different sizes and precisions.

For all three datasets, we use a model that assigns a “self-report” score based on the likelihood that a profile contains a self-report. We then use a binary cutoff to only include users with a high enough self-report score. We obtain this score by leveraging lexical co-occurrence, an important cue for word associations (Spence and Owens, 1990; Church and Hanks, 1989). We combine relative frequencies of co-occurring words within a fixed window, weighed by distance between query and co-occurring self-report words. For example, if “farmer” is a self-report word, then “Black farmer” should score higher than “Black beans farmer” since the query and self-report word are closer. We choose the window size and threshold for this score function on a manually-labeled tuning set, after which our scoring function achieves 72.4% accuracy on a manually-labeled test set. Details on preprocessing and our self-report score are in Appendices A and B.

Our first dataset selects users with a bigram containing a racial keyword followed by a “person keyword.” Our person keywords are: `man`, `woman`, `person`, `individual`, `guy`, `gal`, `boy`, and `girl` so this method matches users with descriptions containing bigrams such as

“Black woman” or “Asian guy.” We expect this method to have high precision, but it has extreme label imbalance; 91% of the users are labeled as either white or black. From the Twitter 1% sample, this dataset contains 122k users, but only 112k users could be re-scraped in 2019. We refer to this dataset as **Query-Bigram (QB)**.

As **QB** contains only 112k users, we consider a less restrictive approach. Our second dataset uses four heuristic filters to remove false positives from the original 2.67M users. Many descriptions spuriously match “black” and “white” in addition to other colors, so we filtered out all words from a color-list (Berlin and Kay, 1991). Second, we filter out racial keywords followed by plural nouns (e.g. “white people”), using NLTK `TweetTokenizer` (Bird et al., 2009) to obtain part-of-speech tags. We curate a list of 286 Google bigrams that frequently contain a query but are unlikely to be self-reports (e.g. “black sheep,”) (Michel et al., 2011). Finally, we ignore query words that appear inside quotation marks. Table 2 shows how precision and dataset size change as we apply these filters. Applying all four gives a total of 1.72M users; after thresholding on self-report score we are left with 228k users. 135k such users could be scraped in 2019, producing our **Heuristic-Filtered (HF)** dataset.

As **QB** and **HF** are quite imbalanced, we design a third dataset to equally represent all four classes. Across both our **QB** and **HF** datasets we have only 7,756 Hispanic/Latinx users that we could scrape in 2019, making it our smallest demographic class. We thus use our self-report scores to select the highest-scoring 7,756 users from each of other classes, producing our **Class-Balanced (CB)** dataset of 31k users.

5 Experimental Evaluation

We now conduct an empirical evaluation of our noisy self-report datasets. Showing that our datasets produce accurate classifiers demonstrates the value of our noisy self-report method for dataset construction. We train supervised classifiers on both our and existing datasets, comparing classifier performance in two evaluation settings.

We divide the six datasets described in Table 1 into training, dev, and test sets. We use the gold-standard self-report survey data from Preoțiuc-Pietro et al. (2015) as our held-out test set for evaluating all models. We combine the crowdsourced data from Volkova and Bachrach (2015) and Culotta et al. (2015) into a single dataset containing 3.5k users, which we then split 60%/40% to create a training and development set. The training set is our baseline comparison, referred to as **Crowd** in our results tables. We also create class-balanced versions of the dev and test sets with 156 and 452 users, respectively. Finally, we use each of our three collected datasets (**QB**, **HF**, **CB**) as training sets, and use a combination of each with the **Crowd** training set. Thus in total, we have seven training datasets, which make up the bottom seven rows of our results in Table 3, below. These results show our three models evaluated on the imbalanced and balanced test sets.

The balanced and imbalanced dev sets are used for all model and training set combinations in Table 3, which controls for the effect of model hyper-parameter selection. Cross-validation could be used in practical low-resource settings, but we use a single held-out dev set, which we subsample in the balanced case.

5.1 Demographic Prediction Models

We consider three demographic inference models which we train on each training set. The first follows Wood-Doughty et al. (2018) and uses a single tweet per user. A character-level CNN maps the user’s name to an embedding which is combined with features from the profile metadata, such as user verification and follower count. These are passed through a two fully-connected layers to produce classifications. This model is referred to as “Names” in Table 3. The second model from Volkova and Bachrach (2015) uses a bag-of-words representation of the words in the user’s recent tweets as the input to a sparse logistic regression classifier. The vocabulary is the 77k non-stopwords that occur at least twice in the dev set. We download up to the 200 most recent tweets for each user from the Twitter API. This model is referred to as “Unigrams” in Table 3. The third model

Dataset/Baseline	Imbalanced prediction						Balanced prediction					
	Names		Unigrams		BERT		Names		Unigrams		BERT	
	F1	Acc%	F1	Acc%	F1	Acc%	F1	Acc%	F1	Acc%	F1	Acc%
Random	.250	25.0	.250	25.0	.250	25.0	.250	25.0	.250	25.0	.250	25.0
Majority	.224	80.8	.224	80.8	.224	80.8	.100	25.0	.100	25.0	.100	25.0
Crowd	.268	74.9	.432	83.2	.402	74.8	.213	.322	.343	40.9	.402	43.7
QB	.335	71.7	.394	71.4	.371	61.0	.316	.377	.406	46.5	.461	48.3
Crowd+QB	.331	74.3	.460	78.4	.383	62.4	.276	.344	.453	47.6	.484	50.1
HF	.324	64.4	.401	72.4	.346	62.3	.308	.377	.418	47.3	.408	44.1
Crowd+HF	.198	54.0	.449	76.9	.360	62.1	.149	.233	.466	50.9	.441	47.4
CB	.299	49.4	.300	43.3	.285	39.0	.379	.381	.463	48.9	.474	49.0
Crowd+CB	.249	35.9	.449	74.6	.349	52.0	.386	.390	.465	48.9	.514	52.6

Table 3: Experimental results for baseline methods, models trained on the crowdsourced datasets, and models trained on our self-report datasets. The best result in each column is in bold. Dataset abbreviations are defined in § 4. “+” indicates a combined dataset of crowdsourced data plus our self-report data. Section 5 and Appendix C contain the training and evaluation details.

Method	Imbalanced			
	W	B	H/L	A
Random	25.0	25.0	25.0	25.0
Majority	100.	-	-	-
Crowd	95.1	49.8	0.9	19.1
QB	77.7	74.0	5.4	30.1
Crowd+QB	86.5	66.5	13.7	29.2
HF	78.9	74.3	7.4	25.6
Crowd+HF	84.2	72.1	14.7	24.8
CB	41.1	77.1	16.7	51.3
Crowd+CB	81.1	68.7	20.1	30.1
Method	Balanced			
	W	B	H/L	A
Random	25.0	25.0	25.0	25.0
Majority	100.	-	-	-
Crowd	95.6	51.3	15.0	1.8
QB	75.2	75.2	5.3	30.1
Crowd+QB	76.1	67.3	25.6	21.2
HF	77.9	77.0	8.9	25.6
Crowd+HF	87.6	73.5	15.9	26.5
CB	41.6	82.3	20.4	51.3
Crowd+CB	72.6	72.6	19.5	31.0

Table 4: Class-specific accuracy for Unigram models. Dashes indicate 0% accuracy. In general, the more class-imbalanced a dataset is, the worse it does on the smaller classes. In the imbalanced setting, the Unigram model trained on the Crowd dataset achieves the best accuracy solely due to its 95.1% accuracy on the users labeled as White.

uses DistilBERT (Sanh et al., 2019) to embed those same 200 tweets into a fixed-length representation, which is then passed through logistic regression with L2 regularization to

	W	B	H/L	A
White	12.7	4.0	3.6	4.9
Black	3.3	16.9	1.8	3.1
Hispanic/Latinx	7.6	4.0	6.5	6.7
Asian	6.2	2.2	1.8	14.7

Table 5: Balanced confusion matrix for BERT on Crowd+CB. Rows show true labels and columns predictions. Each cell shows test set percentage.

produce a classification. This model is referred to as “BERT” in Table 3. For all models we tune hyperparameters using the crowdsourced dev set. Training details for all models are in Appendix C and released in our code.

5.2 Evaluation and Baselines

We consider multiple evaluation setups to explore the extreme class imbalance of the survey and crowdsourced datasets (Table 1). First, we evaluate both total accuracy and macro-averaged F1 score, which penalizes poor performance on less-frequent classes. Second, we separately evaluate tuning and testing our models on either imbalanced or balanced dev and test sets, to see how it affects per-class classifier accuracy. Finally, we train our unigram and BERT models to reweigh examples with the inverse probability of the class label in the training data.

We also show the performance of two naïve strategies: randomly guessing across the four demographic categories, and deterministically guessing the majority category. These

	# Users	LD	CPT	TTR	HPT	Formality	Politeness	Top SAGE Keywords
A	9442	.751	.075	.533	.155*	-1.770	.4595	liked, visit, hahaha, art, youtube
B	70838	.747	.067	.532	.096 [†]	-1.750	.4584	avrillavigne, ni**as, black, ni**a, wit
H/L	8349	.731	.051	.563	.145*	-1.802	.4609	justinbieber, justin, online, follow
W	57724	.759	.085	.510	.081 [†]	-1.697	.4614	bc, realdonaldtrump, snapchat, dog, holy

Table 6: Comparison of the mean values for each numerical feature between groups. The last column has the top keywords per group as differentiated according to the SAGE model. Methods are described in § 7. Abbreviations: LD, Lexical Diversity; CPT, Contractions/tweet; TTR, Type-Token Ratio; HPT, Hashtags/tweet. Almost all differences are significant; only those numbers that share superscript symbols are **not** significantly different at a 0.05 confidence level when using a Mann-Whitney U test.

baselines highlight the trade-offs between accuracy and F1. Because the imbalanced test set is so imbalanced, the “Majority” baseline strategy can achieve high overall accuracy, but very low F1. The Random baseline has low overall accuracy but slightly better F1 than the Majority strategy. These two baselines provide the first two rows of Table 3.

We stress these evaluation details because the class-imbalance may have serious implications for downstream applications. Models trained to do well on the majority class at the expense of minority classes could bias downstream analyses by under-representing minority groups. In public health applications with disparities between groups (LaVeist, 2005), not accounting for imbalances between the training and test datasets could exacerbate rather than ameliorate inequalities.

6 Experimental Results and Discussion

Table 3 shows several trends. The BERT and Unigram models, using 200 tweets per user, generally outperform the single-tweet Names models. In the imbalanced evaluations, we see a large trade-off between accuracy and F1, with models achieving higher overall accuracy when they ignore the smaller Asian and Hispanic/Latinx classes. Even the trivial “Majority” baseline is competitive due to the extreme class-imbalance. While models trained only on **Crowd** achieve significantly higher accuracy on the imbalanced test set than models trained on our datasets, this is only because of their excellent performance on White users. Table 4 shows the class-specific accuracy of Unigram models; the model trained only on the imbalanced **Crowd** dataset achieves 95.1% accuracy on

White users, but lower than 50%, 1%, and 20% accuracy on Black, Hispanic/Latinx, and Asian users. While more sophisticated approaches to addressing the extreme class imbalance could close the gap between training on **Crowd** alone and using our noisy datasets, we can see the benefits of our data in the balanced evaluation.

Across all balanced evaluations, all but one of the models trained with our collected datasets outperform models trained only on **Crowd** in both accuracy and F1. Several models improve by more than .10 F1 over models trained only on **Crowd**. The BERT models achieve the best performance in the balanced evaluation, while performing relatively poorly on imbalanced data. This occurs because the BERT models achieve high accuracy on the Black and Asian classes, which are underrepresented in our imbalanced test set. We show a confusion matrix for our best balanced model in Table 5.

These models are quite simple, and more complex models could improve performance independent of the dataset. However, by limiting ourselves to simpler models, we can demonstrate that for learning a classifier that performs well on four-class classification of race and ethnicity, our noisy datasets are clearly beneficial. While the self-reports are noisy, we collect enough data to support better classifiers on held-out, gold-standard labels. Despite this experimental improvement, real-world applications may require more accurate classifiers or may need to prioritize classifiers with high precision or recall for a particular group. Such research requires a careful contextualization of what conclusions can be drawn from the available data and models; classifier error may exaggerate

	Asian	Black	Hispanic/Latinx	White	Random
% users in dataset	6.71	49.44	5.83	38.02	–
% users with 1+ tweets from Android	38.95* [†]	38.33*	39.41 [†]	25.46	–
% users with 1+ tweets from iPhone	60.28	58.21	54.89	75.37	–
% users with 1+ tweets from Desktop	43.34	30.59	44.87	31.04	–
% users with profile URL	34.09*	29.71	34.75*	24.78	20.8
% users with custom profile image	98.83	99.29* [†]	99.24* [†]	99.33 [†]	95.4
% users with geotagging enabled	48.65*	53.27	49.54*	56.04	33.1
% users with 1+ geotagged tweet	8.35*	6.46	7.81*	5.43	7.9
Average statuses count	11974	18709	12449	14177	–
Average tweets per month	177.83	255.41	182.13	200.85	739
(m) % tweets that mention a user	59.73	58.71	60.44*	61.77*	22.3
(m) % tweets that include an image	20.44*	17.20	18.39	19.17*	33.9
(m) % tweets that include a URL	20.99	21.64	24.01	17.22	–

Table 7: Profile Behavioral Features. The first four columns show our **HF** users, the fifth shows a random sample of 1M users reported in (Wood-Doughty et al., 2017), when available. (m) indicates micro-averaging; all others are macro-averaged across users. Almost all differences between **HF** groups are statistically significant according to a Mann-Whitney U Test. However, if two entries in the same row share a superscript, they are not significantly different at a 0.05 confidence level. We cannot test significance against the random sample.

differences between groups.

7 Twitter Behaviors across Groups

Our experiments show that our datasets enable better predictive models, but say nothing about *how* self-reporting users use Twitter. Do different groups in our dataset differ in other behaviors? We explore this using a variety of quantitative analyses of Twitter user behavior, following similarly-motivated public health research (Coppersmith et al., 2014; Homan et al., 2014; Gkotsis et al., 2016). Two interpretations are possible for these group-level differences: either user behavior correlates with demographic categories (Wood-Doughty et al., 2017), or the *choice to self-report* correlates with these behaviors. These can both be true, and our current methods cannot distinguish between them. While our empirical evaluation shows that our data is still useful for training classifiers to predict gold-standard labels, possible selection bias may influence real-world applications.

Lexical features are widely used to study Twitter (Pennacchiotti and Popescu, 2011; Blodgett et al., 2016). For each user in our dataset, we follow §3.1 of Inuwa-Dutse et al. (2018) and calculate Type-Token Ratio⁴,

⁴The number of unique tokens in a tweet divided by the total number of tokens in the tweet.

Lexical Diversity⁵ (Tweedie and Baayen, 1998), and the number of hashtags and English contractions they use per tweet. We then use existing trained models for analyzing formality and politeness (Pavlick and Tetreault, 2016; Danescu-Niculescu-Mizil et al., 2013) of online text. The formality score is estimated with a regression model over lexical and syntactic features including n-grams, dependency parse, and word embeddings. The politeness classifier uses unigram features and lexicons for gratitude and sentiment. We use the published implementations.^{6,7} For both trained models, we macro-average over users’ scores to obtain a value for each demographic group. We also use a SAGE (Eisenstein et al., 2011) lexical variation implementation to find the words that most distinguish each group. The means of the six quantitative features and the top five SAGE keywords for each group is shown in Table 6.

We then consider a few basic measures of Twitter usage, computed from the profile information of each user. Table 7 contains the mean value of these features, describing the broad range of basic user behaviors on the Twitter platform. Almost all differences

⁵The total number of tokens in a tweet without URLs, user mentions and stopwords divided by the total number of tokens in the tweet.

⁶<https://github.com/YahooArchive/formality-classifier>

⁷<https://github.com/sudhof/politeness>

in these behavioral features are significant across groups. Device usage shows the biggest difference; White users are much more likely to have used an iPhone than an Android to tweet. In past work, [Pavalanathan and Eisenstein \(2015\)](#) demonstrated that the use of Twitter geotagging was more prevalent in metropolitan areas and among younger users. Table 7 follows [Wood-Doughty et al. \(2017\)](#) which calculated these features for a sample of 1M Twitter users. Users in our datasets comparatively more often customize their profile image or URL or enable geotagging. More bots or spam in the random sample may partially account for these differences ([Morstatter et al., 2013](#)). Table 8 in Appendix D also compares lists of the most common common emojis, emoticons, and part-of-speech tags within each group.

These analyses show substantial differences between the groups labeled by our self-report methods, suggesting our noisy self-reports correlate with actual Twitter usage behavior. However, it cannot reveal whether these differences primarily correlate with racial/ethnic groups or whether these differences appear from how users decide whether to self-report a race/ethnicity keyword. Researchers working on downstream public health applications (e.g. [Gkotsis et al. \(2016\)](#)) may want to account for these empirical differences between groups in our training datasets when drawing conclusions about users in other datasets.

8 Limitations and Future Work

We have presented a reproducible method for automatically identifying self-reports of race and ethnicity to construct an annotated dataset for training demographic inference models. While our automated annotations are imperfect, we show that our data can replace or supplement manually-annotated data. Our data collection methodology does not rely on large-scale crowd-sourcing, making it more reproducible and easier to keep datasets up-to-date. These contributions enable the development and distribution of tools to facilitate demographic contextualization in computational social science research.

There are several important extensions to

consider. First, our analysis focuses on the United States and English-language racial keywords; most countries have a unique cultural conceptualizations of race/ethnicity and unique demographic composition, and may require a country-specific focus. We only cover four categories of race/ethnicity, ignoring smaller populations and multi-racial categories ([Jones and Smith, 2001](#)). We use a limited set of query terms, which ignores the diversity of how people may choose to self-report their identities. While our methods scale easily to additional categories and/or racial keywords, our evaluation method requires a gold-standard test set that covers those groups. For specific applications, a domain expert might prioritize precision or recall for a specific demographic class. This may involve fine-tuning a classifier on a dataset constructed with a particular class-imbalance; the details of that imbalance should be contextualized with the general class distribution of the population on Twitter. Our analyses could be compared against human perceptions of users' racial identity, though past work has suggested such perceptions have underlying biases ([Preoțiuc-Pietro et al., 2017](#)). Finally, past work has highlighted various biases in demographic inference ([Pavalanathan and Eisenstein, 2015](#); [Wood-Doughty et al., 2017](#)), and our analyses cannot fully rule out the presence of such biases in our data or models. In future work, we strongly encourage the study of racial self-identities and social cultural issues as supported by computational analyses. These issues should be viewed from a global perspective, especially with regards to biases in our collection methods ([Landeiro and Culotta, 2016](#)).

We release our code in the `Demographer` package to enable training new models and constructing future updated datasets. We also release our trained models and annotated Twitter user ids for academic researchers that agree to the data use agreement and obtain approval from an ethics board.

References

- Tarek Al Baghal, Luke Sloan, Curtis Jessop, Matthew L Williams, and Pete Burnap. 2020. Linking twitter and survey data: The impact of survey mode and demographics on consent rates across three uk studies. *Social Science Computer Review*, 38(5):517–532.
- Silvio Amir, Mark Dredze, and John W. Ayers. 2019. Population level mental health surveillance over social media with digital cohorts. In *CLPpsych*.
- McKane Andrus, Elena Spitzer, Jeffrey Brown, and Alice Xiang. 2021. [What we can't measure, we can't understand: Challenges to demographic data procurement in the pursuit of fairness](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 249–260, New York, NY, USA. Association for Computing Machinery.
- Julia Angwin and Terry Parris Jr. 2016. Facebook lets advertisers exclude users by race. *ProPublica blog*, 28.
- Ehsan Mohammady Ardehaly and Aron Culotta. 2017. Co-training for demographic classification using deep learning from label proportions. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 1017–1024. IEEE.
- John W Ayers, Benjamin M Althouse, and Mark Dredze. 2014. Could behavioral medicine lead the web data revolution? *Jama*, 311(14):1399–1400.
- David Bamman, Jacob Eisenstein, and Tyler Schnoebelen. 2014. Gender identity and lexical variation in social media. *Journal of Sociolinguistics*, 18(2):135–160.
- Charley Beller, Rebecca Knowles, Craig Harman, Shane Bergsma, Margaret Mitchell, and Benjamin Van Durme. 2014. I'm a believer: Social roles via self-identification and conceptual attributes. In *ACL*.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Katie Benner, Glenn Thrush, and Mike Isaac. 2019. Facebook engages in housing discrimination with its ad practices, us says. *The New York Times*, 28:2019.
- Adrian Benton, Glen Coppersmith, and Mark Dredze. 2017. [Ethical research protocols for social media health research](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 94–102, Valencia, Spain. Association for Computational Linguistics.
- Brent Berlin and Paul Kay. 1991. *Basic color terms: Their universality and evolution*. Univ of California Press.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. "O'Reilly Media, Inc."
- Su Lin Blodgett, Lisa Green, and Brendan O'Connor. 2016. [Demographic dialectal variation in social media: A case study of African-American English](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1119–1130, Austin, Texas. Association for Computational Linguistics.
- John D Burger, John Henderson, George Kim, and Guido Zarrella. 2011. Discriminating gender on twitter. In *EMNLP*.
- Pete Burnap and Matthew L Williams. 2016. Us and them: identifying cyber hate on twitter across multiple protected characteristics. *EPJ Data Science*, 5(1):11.
- Mary E Campbell and Christabel L Rogalin. 2006. Categorical imperatives: The interaction of latino and racial identification. *Social Science Quarterly*, 87(5):1030–1052.
- Jonathan Chang, Itamar Rosenn, Lars Backstrom, and Cameron Marlow. 2010. epluribus: Ethnicity on social networks. In *ICWSM*.
- Xin Chen, Yu Wang, Eugene Agichtein, and Fusheng Wang. 2015. A comparative study of demographic attribute inference in twitter. *ICWSM*, 15:590–593.
- Kenneth Ward Church and Patrick Hanks. 1989. [Word association norms, mutual information, and lexicography](#). In *27th Annual Meeting of the Association for Computational Linguistics*, pages 76–83, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Andrew J Coldman, Terry Braun, and Richard P Gallagher. 1988. The classification of ethnic status using name information. *Journal of Epidemiology & Community Health*, 42(4):390–395.
- R Dawn Comstock, Edward M Castillo, and Suzanne P Lindsay. 2004. Four-year review of the use of race and ethnicity in epidemiologic and public health research. *American journal of epidemiology*, 159(6):611–619.

- Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. Quantifying mental health signals in twitter. In *CLPsych*.
- Stephen Cornell and Douglas Hartmann. 2006. *Ethnicity and race: Making identities in a changing world*. Sage Publications.
- Kate Crawford. 2017. The trouble with bias. In *Conference on Neural Information Processing Systems, invited speaker*.
- Lorraine Culley. 2006. Transcending transculturalism? race, ethnicity and health-care. *Nursing Inquiry*, 13(2):144–153.
- Aron Culotta, Nirmal Kumar, and Jennifer Cutler. 2015. Predicting the demographics of twitter users from website traffic data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.
- Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. [A computational approach to politeness with application to social factors](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 250–259, Sofia, Bulgaria. Association for Computational Linguistics.
- Laure Delisle, Alfredo Kalaitzis, Krzysztof Majewski, Archy de Berker, Milena Marin, and Julien Cornebise. 2019. [A large-scale crowdsourced analysis of abuse against women journalists and politicians on twitter](#). *CoRR*, abs/1902.03093.
- Mark Dredze, Michael J Paul, Shane Bergsma, and Hieu Tran. 2013. Carmen: A twitter geolocation system with applications to public health. In *AAAI workshop on expanding the boundaries of health informatics using AI (HIAI)*, volume 23, page 45. Citeseer.
- William W Dressler, Kathryn S Oths, and Clarence C Gravlee. 2005. Race and ethnicity in public health research: models to explain health disparities. *Annu. Rev. Anthropol.*, 34:231–252.
- Jacob Eisenstein, Amr Ahmed, and Eric P. Xing. 2011. [Sparse additive generative models of text](#). In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 1041–1048. Omnipress.
- Jacob Eisenstein, Brendan O’Connor, Noah A Smith, and Eric P Xing. 2014. Diffusion of lexical change in social media. *PloS one*, 9(11):e113114.
- Marc N Elliott, Allen Fremont, Peter A Morrison, Philip Pantoja, and Nicole Lurie. 2008. A new method for estimating race/ethnicity and associated disparities where administrative records lack self-reported race/ethnicity. *Health services research*, 43(5p1):1722–1736.
- Marc N Elliott, Peter A Morrison, Allen Fremont, Daniel F McCaffrey, Philip Pantoja, and Nicole Lurie. 2009. Using the census bureau’s surname list to improve estimates of race/ethnicity and associated disparities. *Health Services and Outcomes Research Methodology*, 9(2):69–83.
- Casey Fiesler and Nicholas Proferes. 2018. ‘participant’ perceptions of twitter research ethics. *Social Media+ Society*, 4(1):2056305118763366.
- Kevin Fiscella and Allen M Fremont. 2006. Use of geocoding and surname analysis to estimate race and ethnicity. *Health services research*, 41(4p1):1482–1500.
- Lucie Flekova, Jordan Carpenter, Salvatore Giorgi, Lyle Ungar, and Daniel Preoțiu-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 Linguistics (Volume 1: Long Papers)*, pages 843–854, Berlin, Germany. Association for Computational Linguistics.
- George Gkotsis, Anika Oellrich, Tim Hubbard, Richard Dobson, Maria Liakata, Sumithra Velupillai, and Rina Dutta. 2016. [The language of mental health problems in social media](#). In *Proceedings of the Third Workshop on Computational Linguistics and Clinical Psychology*, pages 63–73, San Diego, CA, USA. Association for Computational Linguistics.
- A Gonzalez-Barrera and MH Lopez. 2015. Is being hispanic a matter of race, ethnicity or both? *Pew Research Center*.
- Charles Hirschman, Richard Alba, and Reynolds Farley. 2000. The meaning and measurement of race in the US census: Glimpses into the future. *Demography*, 37(3).
- Arnold K Ho, Steven O Roberts, and Susan A Gelman. 2015. Essentialism and racial bias jointly contribute to the categorization of multiracial individuals. *Psychological Science*, 26(10):1639–1645.
- Christopher Homan, Ravdeep Johar, Tong Liu, Megan Lytle, Vincent Silenzio, and Cecilia Ovesdotter Alm. 2014. [Toward macro-insights for suicide prevention: Analyzing fine-grained distress at scale](#). In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 107–117, Baltimore, Maryland, USA. Association for Computational Linguistics.

- Xiaolei Huang and Michael J. Paul. 2018. [Examining temporality in document classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 694–699, Melbourne, Australia. Association for Computational Linguistics.
- Xiaolei Huang and Michael J. Paul. 2019. [Neural user factor adaptation for text classification: Learning to generalize across author demographics](#). In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)*, pages 136–146, Minneapolis, Minnesota. Association for Computational Linguistics.
- Isa Inuwa-Dutse, Bello Shehu Bello, and Ioannis Korkontzelos. 2018. Lexical analysis of automated accounts on twitter. *arXiv:1812.07947*.
- Jiachen Jiang and Soroush Vosoughi. 2020. Not judging a user by their cover: Understanding harm in multi-modal processing within social media research. In *Proceedings of the 2nd International Workshop on Fairness, Accountability, Transparency and Ethics in Multimedia*, pages 6–12.
- Nicholas A Jones and Amy Symens Smith. 2001. *The two or more races population, 2000*, volume 8. US Department of Commerce, Economics and Statistics Administration, US.
- Soon-Gyo Jung, Jisun An, Haewoon Kwak, Joni Salminen, and Bernard Jansen. 2018. Assessing the accuracy of four popular face recognition tools for inferring gender, age, and race. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12.
- Rebecca Knowles, Josh Carroll, and Mark Dredze. 2016. Demographer: Extremely simple name demographics. In *Proceedings of the First Workshop on NLP and Computational Social Science*, pages 108–113.
- Virgile Landeiro and Aron Culotta. 2016. Robust text classification in the presence of confounding bias. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Thomas A LaVeist. 2005. *Minority populations and health: An introduction to health disparities in the United States*, volume 4. John Wiley & Sons.
- Sharon M Lee and Sonya M Tafoya. 2006. Rethinking us census racial and ethnic categories for the 21st century. *Journal of Economic and Social Measurement*, 31(3-4):233–252.
- Kate Loveys, Jonathan Torrez, Alex Fine, Glen Moriarty, and Glen Coppersmith. 2018. Cross-cultural differences in language markers of depression online. In *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, pages 78–87.
- Alice E Marwick and danah boyd. 2011. I tweet honestly, i tweet passionately: Twitter users, context collapse, and the imagined audience. *New media & society*, 13(1):114–133.
- Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, et al. 2011. Quantitative analysis of culture using millions of digitized books. *science*, 331(6014):176–182.
- Jude Mikal, Samantha Hurst, and Mike Conway. 2016. Ethical issues in using twitter for population-level depression monitoring: a qualitative study. *BMC medical ethics*, 17(1):22.
- Ehsan Mohammady and Aron Culotta. 2014. [Using county demographics to infer attributes of Twitter users](#). In *Proceedings of the Joint Workshop on Social Dynamics and Personal Attributes in Social Media*, pages 7–16, Baltimore, Maryland. Association for Computational Linguistics.
- Fred Morstatter, Jürgen Pfeffer, Huan Liu, and Kathleen M Carley. 2013. Is the sample good enough? comparing data from twitter’s streaming api with twitter’s firehose. In *ICWSM*.
- Brendan O’Connor, Ramnath Balasubramanian, Bryan R Routledge, and Noah A Smith. 2010. From tweets to polls: Linking text sentiment to public opinion time series. *ICWSM*, 11(122-129):1–2.
- Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kiciman. 2019. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data*, 2:13.
- Michael Paul and Mark Dredze. 2011. You are what you tweet: Analyzing twitter for public health. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 5.
- Umashanthi Pavalanathan and Jacob Eisenstein. 2015. Confounds and consequences in geotagged twitter data. *arXiv:1506.02275*.
- Ellie Pavlick and Joel Tetreault. 2016. [An empirical analysis of formality in online communication](#). *Transactions of the Association for Computational Linguistics*, 4:61–74.

- Marco Pennacchiotti and Ana-Maria Popescu. 2011. A machine learning approach to twitter user classification. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 5.
- Daniel PreoŃiu-Pietro, Sharath Chandra Guntuku, and Lyle Ungar. 2017. Controlling human perception of basic user traits. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2335–2341, Copenhagen, Denmark. Association for Computational Linguistics.
- Daniel PreoŃiu-Pietro and Lyle Ungar. 2018. User-level race and ethnicity predictors from Twitter text. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1534–1545, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Daniel PreoŃiu-Pietro, Svitlana Volkova, Vasileios Lampos, Yoram Bachrach, and Nikolaos Aletras. 2015. Studying user income through language, behaviour and affect in social media. *PloS one*, 10(9):e0138717.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108v4*.
- David Satcher. 2001. Mental health: Culture, race, and ethnicity—supplement to mental health: A report of the surgeon general.
- Lauren Sinnenberg, Alison M Buttenheim, Kevin Padrez, Christina Mancheno, Lyle Ungar, and Raina M Merchant. 2017. Twitter as a tool for health research: a systematic review. *American journal of public health*, 107(1):e1–e8.
- Donald P Spence and Kimberly C Owens. 1990. Lexical co-occurrence and association strength. *Journal of Psycholinguistic Research*, 19(5):317–330.
- Fiona J Tweedie and R Harald Baayen. 1998. How variable may a constant be? measures of lexical richness in perspective. *Computers and the Humanities*, 32(5):323–352.
- Nicholas Vargas and Kevin Stainback. 2016. Documenting contested racial identities among self-identified latina/os, asians, blacks, and whites. *American Behavioral Scientist*, 60(4):442–464.
- Svitlana Volkova and Yoram Bachrach. 2015. On predicting sociodemographic traits and emotions from communications in social networks and their implications to online self-disclosure. *Cyberpsychology, Behavior, and Social Networking*, 18(12):726–736.
- Zijian Wang, Scott A. Hale, David Ifeoluwa Adelani, Przemyslaw A. Grabowicz, Timo Hartmann, Fabian Flöck, and David Jurgens. 2019. Demographic inference and representative population estimates from multilingual social media data. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 2056–2067. ACM.
- Zach Wood-Doughty, Praateek Mahajan, and Mark Dredze. 2018. Johns Hopkins or johnny-hopkins: Classifying individuals versus organizations on Twitter. In *Proceedings of the Second Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, pages 56–61, New Orleans, Louisiana, USA. Association for Computational Linguistics.
- Zach Wood-Doughty, Michael Smith, David Broniatowski, and Mark Dredze. 2017. How does Twitter user behavior vary across demographic groups? In *Proceedings of the Second Workshop on NLP and Computational Social Science*, pages 83–89, Vancouver, Canada. Association for Computational Linguistics.

A Preprocessing, Tokenizing, and Tagging

We lowercase all descriptions and use NLTK Tweet Tokenizer (Bird et al., 2009) to get the PoS tags. Our candidate self-report words are scraped from 177M Twitter descriptions using the regex and PoS pattern, $\{I' / I a\}_m (+RB)(+DT)(+JJ)+NN$. We collect both adjectives and nouns from the pattern above, and refine the matches by keeping adjectives and nouns that match the majority tag in the Google N-gram corpus. We filter out plural words (e.g. “white **people**”) using a PoS tag pattern, $JJ + NNPS/NNS$, and refer to our set of self-report words as S .

B Calculating the “Self-Report” Score

To calculate the score described in § 4, we first obtain simple co-occurrence weighting by counting the occurrences $O_s(w_s)$ of word w_s as a self-report word, and its overall occurrences $O(w_s)$. Then:

$$R = \sum_{w_s \in S^{win}} \frac{1}{D(w_s, q)} \cdot \frac{O_s(w_s)}{O(w_s)},$$

where S^{win} is the self-report words in the fixed window size, $D(w_s, q)$ denotes the distance between w_s and query word q .

We also consider a TF-IDF weighting as:

$$R_{\text{tfidf}} = \sum_{w_s \in S^{win}} \frac{1}{D(w_s, q)} \cdot \frac{O_s(w_s)}{O(w_s)} \cdot \log \frac{\sum_{w \in S} O_s(w)}{O_s(w_s)}$$

To fine-tune our self-report score, three authors manually labeled a tuning set of 400 descriptions as to whether the user was self-reporting a matching query word, using a three-label nominal scale of “yes,” “no,” and “unsure.” We discarded 6 that we classified as organizations (Wood-Doughty et al., 2018), and had an Krippendorff α 0.8058 on the remaining 394. We use majority voting strategy to get binary labels and select the self-report score’s hyperparameters of window size and threshold, and whether to use the tf-idf weighting, based on the precision calculated on this tuning set.

To ensure that these chosen hyperparameters did not overfit to the tuning set, we sampled an additional 199 users from **HF**. Using a three-label nominal scale of “yes,” “no,” or “unsure,” the three annotators achieved a Krippendorff’s alpha of 0.625. After converting to binary “yes” and “no” by taking majority voting and discarding 7 users who were majority “unsure,” our best model achieves 72.4% accuracy on the test set with simple weighting, window size 5, and threshold of 0.35.

C Model Training Details

Our name model uses a CNN implementation released in Wood-Doughty et al. (2018). We use a CNN with 256 filters of width 3. The user’s name (not screen name) is truncated at 50 characters and embedded into a 256 dimensional character embedding. We fine-tuned the learning rate on our dev data, trained for 250 epochs, and used early-stopping on dev-set F1 to pick which model to evaluate on the test set.

Our unigram model follows Volkova and Bachrach (2015), using a simple sparse logistic regression. We use an implementation from Scikit-Learn, and tune the regularization parameter on the dev set. We introduce a hyperparameter to down-weight the contribution of our users compared to the baseline users; we also set that parameter on the dev set.

For BERT model, we first get embedding for every tweet by taking the vector with size 768 on special [CLS] token in the last hidden layer. The element-wise average of all tweet embeddings from one user is then passed through a logistic regression model with L2 regularization to make the classification. Similarly, the regularization parameter is tuned on the dev set. We fine-tuned DistilBERT model on tweets collected from training set split of the crowdsourced dataset. However, after observing limited performance improvement we just use pre-trained DistilBERT model.

Top k	Emojis		Hashtags		PoS bigrams	
	20	50	20	50	20	50
A v. B	-0.67	-0.26	-0.85	-0.87	0.29	0.19
A v. H/L	-0.10	-0.07	-0.84	-0.86	0.55	0.02
A v. W	-0.38	0.13	-0.83	-0.80	0.02	-0.02
B v. H/L	-0.65	-0.38	-0.83	-0.82	0.52	0.03
B v. W	-0.48	-0.16	-0.79	-0.72	0.04	0.24
H/L v. W	-0.40	-0.13	-0.91	-0.89	-0.17	-0.28

Table 8: Kendall’s τ correlation coefficients for top items of different list features. For hashtags in particular we see large negative coefficients.

D Additional Analyses of Twitter Behavior across Groups

This appendix contains an additional analysis following § 7.

In addition to the SAGE keyword comparison, we explore topical differences between groups by compiling ranked lists of common emojis, emoticons, and part-of-speech tags within each group. Table 8 shows a comparison of Kendall τ rank correlation between these. To compare across groups, we look at the top k items in each list and calculate Kendall τ rank correlation coefficients for each pair of demographic groups (Morstatter et al., 2013). Table 8 shows pairwise τ correlations. These coefficients vary between -1 and 1 for perfect negative and positive correlations. For emojis, all correlations are negative for $k = 20$, but increase at $k = 50$. For hashtags, however, correlations are strongly negative for all values of k , suggesting that groups labeled by our method substantially differ in the topics they discuss. While we use English keywords for data collection, topic difference may be confounded by users’ native language(s).

E Data Statement

Following Bender and Friedman (2018), we highlight characteristics of our collected noisy self-report data that may be important for mitigating ethical and scientific missteps.

Curation rationale Examples of Twitter users who self-report their racial identity using English-language keywords.

Language variety While our dataset contains predominantly English (en-US), there is substantial diversity in language due

to the international and due to the informal setting of Twitter. When we randomly sample 1000 users from our Heuristic Filter list and consider up to 100 tweets per user, we find that the Twitter-produced `lang` field indicates that 78.5% of the tweets are in English, with the next three most-common `lang` labels as Spanish (3.8%), Portuguese (3.7%), and Undetermined (3.3%).

Speaker demographics The speakers in our dataset are Twitter users. To be included in our initial dataset, users must use an English racial self-report keyword in their Twitter profile description, and must not be labeled as an organization by the classifier from Wood-Doughty et al. (2018). We then perform additional filtering of users, detailed in the paper, to improve the likelihood that a racial self-report keyword is actually self-reporting race.

Annotator demographics Our small manual annotation was conducted by three authors, Asian and White men, ages 20-30, with native languages of Chinese and English.

Speech situation Twitter user profiles and tweets.

Text characteristics Informal Twitter user descriptions and tweets. We make no restrictions on the content of the tweets.