

Gender and Racial Fairness in Depression Research using Social Media

Carlos Aguirre, Keith Harrigian, Mark Dredze

Johns Hopkins University

caguirr4@jhu.edu, kharrigian@jhu.edu, mdredze@cs.jhu.edu

Abstract

Multiple studies have demonstrated that behavior on internet-based social media platforms can be indicative of an individual's mental health status. The widespread availability of such data has spurred interest in mental health research from a computational lens. While previous research has raised concerns about possible biases in models produced from this data, no study has quantified how these biases actually manifest themselves with respect to different demographic groups, such as gender and racial/ethnic groups. Here, we analyze the fairness of depression classifiers trained on Twitter data with respect to gender and racial demographic groups. We find that model performance systematically differs for underrepresented groups and that these discrepancies cannot be fully explained by trivial data representation issues. Our study concludes with recommendations on how to avoid these biases in future research.

1 Introduction

Work from De Choudhury et al. (2013) and Coppersmith et al. (2014), showing that an individual's mental health can be evaluated based on the language they generate on social media platforms, has served as the basis for a substantial amount of computational research over the last decade. Subsequent studies have examined an even wider range of mental health conditions, social media platforms, and types of online behavior at both the individual and population level (Coppersmith et al., 2015b; Lynn et al., 2018; De Choudhury et al., 2016). Vast potential for societal benefits underlies this work, as conservative estimates suggest that 8.1% of American adults suffer from major depressive disorder at any given time and up to 16.2% of individuals will experience at least one major depressive episode during their lifetime (Kessler et al., 2003; Brody et al., 2018; Hasin et al., 2018).

Mental health services are transitioning to online mediums at a rapid pace, with the recent COVID-19 pandemic dramatically further accelerating this trend (Zhou et al., 2020; Ohannessian et al., 2020). Thus, analysis of online language may play a key role in mental health treatment in the future.

Nonetheless, care must be taken to understand potential biases inherent in this research before any technologies are deployed in a clinical setting. For instance, previous work has found that Black and Hispanic/Latinx individuals are less likely to be treated for depression than White individuals (Simpson et al., 2007). Possibly a result of this underlying bias, recent studies of the US population have concluded that baseline rates of depression vary depending on demographics (Brody et al., 2018; Hasin et al., 2018) — major depressive disorder was found to be more prevalent in females and White adults. Yet, it remains unclear whether these supposed differences in depression prevalence between gender and racial/ethnic demographic groups are the result of measurement error or other confounders. Various psychological studies have found mental health disorders, including depression, may manifest differently depending on cultural background and thus make uniform diagnosis a difficult proposition (Blanchard et al., 2020; Henrich et al., 2010). These ambiguities were highlighted by recent computational research from Amir et al. (2019), which found predictive rates of depression inferred using classifiers for social media data to not match previous US depression estimates. Indeed, the authors actually find that Black and Hispanic/Latinx individuals are *more* likely to be affected by depression than White individuals.

Additionally, NLP and other data-driven algorithms have been shown to suffer from *content* biases; that is, undesirable group-wise differences with respect to protected groups, such as race/eth-

nicity or gender (Johannsen et al., 2015; Hovy and Søgaard, 2015; Bolukbasi et al., 2016; Gonen and Goldberg, 2019; Rudinger et al., 2018). Therefore, in consideration of the social impact of NLP research (Hovy and Spruit, 2016), in the area of mental health content analysis it stands to reason that we should also look for *population* biases, as they pertain to protected groups, and the way these might affect NLP algorithms' fairness.

Previous research has utilized user demographics within social media mental health studies to construct control groups (Coppersmith et al., 2014), to enhance classifier performance through additional features (Preoțiuc-Pietro et al., 2015), and to analyze trends amongst specific populations (De Choudhury et al., 2014). In an attempt to preemptively address *population* biases, Amir et al. (2019) proposed a cohort-based sampling approach to collect representative measures of wellness amongst the general population. However, as noted in a recent literature reviews (Chancellor and De Choudhury, 2020; Harrigian et al., 2020), no previous computational mental-health study has accounted for differences in population-level depression rates nor explored performance variations across demographic subgroups at training time. Therefore, little is known about the fairness of these automated systems. Are models trained for mental health fair across demographic groups? Are current datasets demographically representative? If bias exists, what is its source?

In this study, we analyze two common depression-inference datasets and explore the susceptibility of different computational methods to demographic biases. We find that existing datasets are **not demographically representative** and that, without accounting for this, we find degradation in model performance for underrepresented groups. We explore the possible sources of this bias and conclude with recommendations for future research that may address these issues.

2 Mental Health and Social Media

Challenges obtaining mental health annotations for social media data have thus far constrained the size and quality of existing datasets. For instance, manual annotation of mental health status generally requires expert domain knowledge, while the sensitive nature of such annotations limit multi-institutional data sharing (Arseniev-Koehler et al., 2018). Consequently, most datasets rely on labels

based on behavioral proxies or self-reported diagnoses, which more easily scale, but introduce problematic *self-disclosure* bias and label noise. Furthermore, as our understanding of mental health is continually evolving, studies have used different and sometimes conflicting guidelines for annotation (Brody et al., 2018; Hasin et al., 2018). With these challenges at the forefront of dataset curation, issues surrounding demographic balance and representation have been largely kicked down the road of the research domain.

Challenges accounting for demographics go beyond the computational research space and are well-illustrated by disparities between two recent surveys of depression prevalence. The Centers for Disease Control and Prevention (CDC) found depression prevalence between race/ethnicity groups did not differ (Brody et al., 2018), while a study using the results of the National Epidemiologic Survey on Alcohol and Related Conditions III (NESARC-III) found depression to be more prevalent in White Americans versus minorities (Hasin et al., 2018).

3 Ethical Considerations

The sensitive nature of mental health research and individual demographics requires us to consider possible benefits of this study alongside its potential harms. Specifically, we must evaluate the cost-benefit trade-off of inferring and/or securing three highly-personal individual attributes: (*depression diagnoses*: Benton et al., 2017a; *gender identity*: Larson, 2017; *race/ethnicity identity*: Wood-Doughty et al., 2020).

The potential immediate benefit of this study is a better understanding of demographic bias in computational mental health research. A potential secondary benefit is the mitigation of extant clinical treatment disparities (Simpson et al., 2007). As mental health treatment increasingly adopts an online delivery mechanism, this research is uniquely situated to inform the development of new AI systems and public policy in the area.

However, we are cognizant of the potential harms from our work. Mental health status and demographic identities are both sensitive personal attributes that could be used to maliciously target individuals on publicly-facing online platforms. Therefore, we follow the guidelines of Benton et al. (2017a) and Ayers et al. (2018) on data use, storage, and distribution. All analysis was conducted on de-identified versions of data, with any identifiable

information being used only during intermediate data-processing subroutines that were hidden from researcher interaction and approved by the original dataset distributors. Our study was exempted from review by our Institutional Review Board under 45 CFR § 46.104.

To facilitate any form of statistical analysis, we also need to formalize gender and race/ethnicity. We seek a balance between the limitations of demographic inference systems¹ and alignment to demographic categories conventions used by the mental health literature (Brody et al., 2018; Hasin et al., 2018), versus propagating demographic definitions that exacerbate existing biases towards gender and racial/ethnic minorities. We consider the ‘folk conception’ of gender as described in Larson (2017) and prominently leveraged in traditional depression research in the United States — we use the sex categories *male* and *female* to denote the corresponding gender categories *masculine* and *feminine*. However, many individuals do not fit in these gender categories, some present a gender online inconsistent to their true identity (Nilizadeh et al., 2016), and they often experience depression and other mental health conditions at a higher rate (McDonald, 2018). For race/ethnicity labels, we consider the mutually-exclusive labeling conventions invoked by Brody et al. (2018) and Wood-Doughty et al. (2020): non-Hispanic White, non-Hispanic Black, non-Hispanic Asian and Hispanic/Latinx, as they are representative of the majority of racial and ethnic identities in the US. Our racial/ethnic categories do not capture multiracial individuals or those with a race/ethnicity outside this group.

We acknowledge these important limitations, but at the same time, there is an urgency to the questions we pose. Computational methods for monitoring mental health have already been deployed by digital surveillance companies (Bark), while analytics dashboards based on these methods are gradually making their way into patients’ (Yoo and De Choudhury, 2019) and providers’ hands (Yoo et al., 2020). The question is: should we avoid asking these questions about current datasets because we cannot produce clear answers, or should we conduct analyses with acknowledged limitations to learn what we can about research that is already being moved into products? We firmly believe the latter. Our hope is that this paper causes re-

searchers to carefully consider these issues, elevate the need for further work, and produce studies that go beyond our study’s limitations with new data and methods. This should not be the last study on this topic; rather, we hope it is the first step which can inform further critical analyses of work in this area.

4 Datasets

We select the task of depression inference for this study, as it is the most widely studied mental health condition in social media research (Harrigian et al., 2020). We consider two Twitter datasets: CLPSYCH (Coppersmith et al., 2015b) and MULTITASK (Benton et al., 2017b).

4.1 CLPSYCH

CLPSYCH was introduced by Coppersmith et al. (2014) and subsequently used in the CLPsych 2015 shared task (Coppersmith et al., 2015b). Tweets were publicly posted between 2008 and 2013. Users who self-disclosed a depression diagnosis were identified using regular expressions (e.g. “I have been diagnosed with disorder”) and then manually reviewed by a team of clinical and computational researchers to verify authenticity of matched disclosures. The control group was sampled from a random pool Twitter users so that the joint distribution of inferred age and gender attributes closely resembled that of users with self-disclosed diagnoses. The 3000 most recent tweets from each user (as of the original dataset collection date) were retrieved. To reduce ambiguity in model performance that arises due to data insufficiencies, we isolate individuals with at least 100 tweets, leading to a final dataset size of 475 depressed individuals and their matched controls (i.e. 950 total users).

4.2 MULTITASK

Benton et al. (2017b) constructed a Twitter dataset (MULTITASK) combining a subset of CLPSYCH with datasets annotated using the same procedure from Coppersmith et al. (2015a,c). In addition to an expanded number of unique individuals (1400 depression, 1400 control), MULTITASK also boasts a more robust historical timeline of tweets for each user.

¹We infer gender and race/ethnicity labels using a content classifier and explore additional limitations in Section 5

5 Demographic Labels

Only age and gender (Schwartz et al., 2013) attributes are available in the originally distributed form of CLPSYCH and MULTITASK, both of which were inferred using now-outdated models. All identifying metadata was either redacted or obfuscated to preserve the privacy of individuals in these datasets. Accordingly, we are confronted immediately by the challenge of securing accurate demographic information to facilitate a robust analysis of any potential gender and racial/ethnic biases. Fortunately, this problem has been tackled using a multitude of different techniques across multiple studies specific to mental health (Yazdavar et al., 2020; Amir et al., 2017; Preotjuc-Pietro et al., 2015; Coppersmith et al., 2015a) and social media applications in general (Volkova et al., 2014; Burger et al., 2011; Fink et al., 2012; Rao et al., 2011).

We obtain race labels using a unigram model from Wood-Doughty et al. (2020), who combine multiple crowd-sourced and self-reported datasets to train classifiers for 4 demographic groups in line with the CDC’s conventions (Brody et al., 2018): non-Hispanic Asian American (A), non-Hispanic African American (B), non-Hispanic White (W) and Hispanic/Latinx (H/L). Their classifier achieves an accuracy of 82.3% within intrinsic evaluations and shows even more promise as high-confidence thresholds are applied. To validate and further reduce noise in previously-inferred gender attributes, we train a new gender inference model on data from Burger et al. (2011) using the same architecture of Wood-Doughty et al. (2020). Our classifier obtains an accuracy of 83.3% amongst within-distribution data and outputs a distribution of inferred gender attributes that strongly aligns with that of the original datasets.

Although each of these procedures has strong internal validity, we recognize that inference errors incurred during this stage may confound and complicate downstream analysis of demographic bias. To mitigate this potential noise, we also de-anonymize a subset of CLPSYCH with the permission of Coppersmith et al. (2014) and apply name-based demographic classifiers (Wood-Doughty et al., 2018, 2020) to each user’s profile to obtain “alternative” age and race attributes.

Between our content and name based classifiers, we are afforded the opportunity to perform downstream analysis of demographic bias based on the attributes derived using the following mechanisms:

- **High Confidence Filter:** Only considers users whose most probable demographic class based on unigram classifier has a confidence $> .95$.
- **Random Sampling:** Considers all available users; randomly split each individual’s tweets into two independent pools so that demographic and mental health inferences are based on separate sets of data.
- **Name Labels:** Only considers users from CLPSYCH who could be de-anonymized; demographics annotated using name-based gender (Wood-Doughty et al., 2018) and ethnicity classifiers (Wood-Doughty et al., 2020).

While we find some variation in the individual-level demographic labels when using the three techniques, the downstream mental health models perform similarly: see details in Appendix A. For the experiments discussed below, we report results from the most computationally-efficient approach, high confidence filtering.

6 Analysis

We conduct an analysis of these datasets and depression models trained on these datasets to answer the following questions:

1. Are depression datasets demographically representative?
2. Do depression classifiers perform similarly across demographic groups?
3. Can we mitigate demographic biases by changing characteristics of the dataset?
4. Do differences in features between demographic groups account for classifier biases?

6.1 Are depression datasets demographically representative?

Before we can empirically measure if these datasets are demographically representative, we must first establish the expected distribution of a representative dataset. While the demographic groups distribution should match the true population (Twitter users with depression), there are no estimates of depression prevalence on Twitter. Thus, we use the Twitter US population as our baseline, and combine it with estimated prevalence of depression among

US demographic groups.²

Methods. Brody et al. (2018) found in a study of adults (>20yrs.) that women are almost twice as likely to be diagnosed with depression compared to men (1.89×) using Patient Health Questionnaires (PHQ-9). We refer to this study as ‘CDC.’ Similarly, Hasin et al. (2018) used a national survey of adults (>18yrs.) and the DSM-5 standard for major depressive disorder (MDD) to estimate that women are almost twice as likely to be diagnosed with depression compared to men (1.86×). We refer to this study as ‘NESARC.’

While there were only small incongruencies between these studies in estimated prevalence of depression as a function of gender, there were significant discrepancies between studies with respect to estimated prevalence as a function of race/ethnicity. Specifically, CDC found that rates of depression were not statistically different between groups, whereas NESARC found a greater prevalence of depression among Whites compared to African Americans, Hispanics/Latinx and Asian Americans (1.25x).

We project these estimates of depression prevalence to the general Twitter population, where both males and females are estimated to participate equally (approximately matching the US population). While there is a slight under-representation of White individuals in Twitter compared to US population (60% vs 64%), Black and Hispanic/Latinx individuals are well represented (Wojcik and Hughes, 2019). Thus, barring slight variations, Twitter roughly mimics the demographic composition of the United States demographics fairly with respect to gender and race/ethnicity.

We combine the Twitter population estimates with depression rates in the US to get the target distributions of demographic users that we expect to observe in our datasets. Figure 1 shows differences between the expected, representative distribution and our complete Twitter datasets.

Results. Based on these estimates, are the depression datasets demographically representative? Figure 1 shows that CLPSYCH and MULTITASK are **not demographically representative** with respect to either gender or race/ethnicity. White individuals are over-represented, while Hispan-

²Our study uses mental health statistics from the United States since they are extensive and widely available. However, due to data anonymization we could not filter our data based on residence in the US. Since these datasets are filtered to focus on English accounts, US accounts likely dominate.

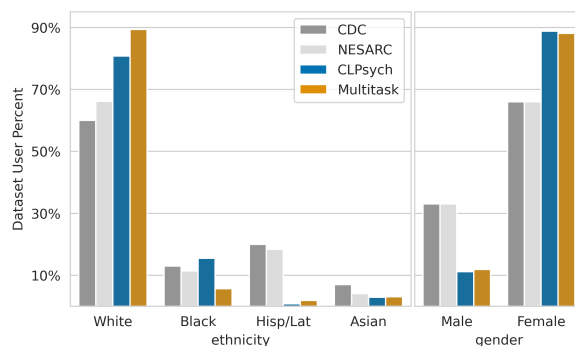


Figure 1: Hispanic/Latinx and Male individuals are underrepresented in both of our Twitter datasets (CLPSYCH & MULTITASK) compared to the Twitter US population estimates (CDC & NESARC).

ic/Latinx individuals are the most underrepresented. In fact, there are *no* male H/L individuals represented in the train split of CLPSYCH. MULTITASK exhibits a larger population bias against minorities compared to CLPSYCH; White individuals are over-represented and Black individuals are under-represented. With respect to gender, both CLPSYCH and MULTITASK have similar distributional skews – females are over-represented compared to the depression adjusted general US population. At the user level, we found no major differences on number of tweets and vocabulary size between demographics: see details in Appendix B.

Overall, CLPSYCH and MULTITASK are **not demographically representative** with respect to US depression rates projected on Twitter demographic estimates.

6.2 Do depression classifiers perform similarly across demographic groups?

We consider this question through experimentation on our datasets, CLPSYCH and MULTITASK.

Methods. We train a depression classifier on CLPSYCH and MULTITASK datasets.

We follow standard pre-processing procedures and filter numeric values, username mentions, retweets and urls from the raw tweet text. We use ℓ_2 -regularized logistic regression models for all of our experiments. TF-IDF vectors are used to represent text in and across tweets, along with mean-pooled 200 dimensional GloVe embeddings pretrained with 2B tweets (Pennington et al., 2014). The vocabulary is pre-filtered per training, as each unigram must appear at least 10 times across all the individuals in training data.

We also experimented with Linguistic Inquiry

Word Count (LIWC) features, a closed-vocabulary English lexicon containing 64 categories (excluding punctuation categories), ranging from linguistic dimensions to psychological processes covering emotions and personal concerns, traditionally used in psychological studies (Pennebaker et al., 2007). In social media analysis, LIWC has been shown to contain signals for mental health disorders (Ireland and Iserman, 2018; Wolohan et al., 2018; Mitchell et al., 2015), including CLPSYCH and MULTITASK.

We also use features based on topic distributions learned via Latent Dirichlet Allocation (LDA) (Blei et al., 2003), following its implementation for Twitter data as specified in Mitchell et al. (2015) where all the tweets for an individual are combined into a “document” and we infer “topics” ($K = 50$ topics). All of the models in our experiments use all four feature groups: TF-IDF, GloVe embeddings, LIWC and LDA. We considered using demographic labels as features, which have been shown to capture signals for depression in Twitter (Preoțiuc-Pietro et al., 2015), but found no significant impact on our analysis or model performance; since demographic labels are not normally available, we do not include them in our analysis. See Appendix C for further implementation details.

To measure the performance bias across demographic groups we report performance on each demographic group. However, the racial/ethnic minority groups in the data are vastly underrepresented. While we address this by combining them into a ‘persons of color’ (PoC) category, the PoC group is still small and limits the reliability and extension of our analysis in this data.

For each dataset, we randomly sample individuals with repetition to construct a training set (bootstrap method) and subsequently obtain a distribution of F1 scores (100 repetitions) followed by one way ANOVA and pairwise T-Tests for each demographic group pair. Motivated by Simpson’s Paradox (Blyth, 1972) and the Matrix of Domination (Costanza-Chock, 2018), we combine the gender and race/ethnicity labels to create a matrix of demographics and report mean F1 scores and 95% confidence interval of each demographic subcategory. Additionally, we seek a metric to measure fairness in performance across demographic groups — our criterion is that model performance should be independent of the demographic labels. Hardt et al. (2016) introduce *equal odds* and *equal*

		CLPSYCH		MULTITASK	
Female	White	0.77	± 0.005	0.84	± 0.002
	PoC	0.41	± 0.013	0.91	± 0.003
Male	White	0.74	± 0.008	0.83	± 0.005
	PoC	0.76	± 0.035	0.45	± 0.016
Equal Odds		0.21	± 0.023	0.13	± 0.013
Opportunity		0.25	± 0.039	0.18	± 0.010

Table 1: Avg. F1 with 95% conf. interval from bootstrap across gender and ethnicity groups (*italics: not significant*), and avg. equal odds and equal opportunity differences. Models underperform for PoC in general, sometimes male PoC (MULTITASK) or female PoC (CLPSYCH).

opportunity, two criteria that seek to equalize the FPR and TPR, or just FPR for the latter, across the protected attributes — these are also known as ‘error rate balance’ (Chouldechova, 2017), ‘conditional procedure accuracy equality’ (Berk et al., 2018) and ‘classification parity’ (Corbett-Davies and Goel, 2018). We compute the average pairwise equal odds and equal opportunity difference, a score of 0 means overall fairness, across the demographic groups in our bootstrap sampling splits and report 95% confidence interval.

Results. Table 1 shows performance of classifiers trained on CLPSYCH and MULTITASK by demographic group. Models trained on CLPSYCH tend to perform worse on female PoC users compared to all other demographic groups. While we observe higher model performance for MULTITASK in general, models trained on MULTITASK tend to perform worse on male PoC users, compared to all other demographic groups. CLPSYCH is scored worse with the fairness metrics compared to MULTITASK.

In short, we observe that depression classifiers perform **worse on people of color**, specifically female PoC in CLPSYCH and male PoC MULTITASK.

6.3 Can we mitigate demographic biases by changing characteristics of the dataset?

Why do depression classifiers perform worse/inconsistently for PoC individuals? We conduct two analyses that investigate how the datasets may cause disparities in fairness.

6.3.1 Data Size

Perhaps the classifier performs worse on demographic groups because we have insufficient train-




Figure 2: Log learning curve of varying dataset sample size on CLPSYCH. Dashed vertical lines represent the total size of the demographic group. Curves suggest more data is helpful to close model performance gap between demographic groups.

ing data. In Section 6.2, we observed fairer results with more data on MULTITASK compared to CLPSYCH. We perform a dataset size experiment to verify the effect on model performance across demographics.

Methods. How do results change with increased amounts of training data? To evaluate this gradient, we consider sampling an equal number of individuals from each demographic group and gradually increase overall dataset size until all available individuals available have been considered. At each dataset size step, we employ a similar bootstrap procedure to the one discussed in Section 6.2, sampling from the available user pool and training a classifier 25 times before moving on to the next dataset size. We continue adding data after a demographic group has been fully saturated to understand how information from overly-represented groups can generalize to under-represented groups.

Results. Figure 2 shows models performance as we increase training dataset size within the CLPSYCH dataset; results for MULTITASK are included in Appendix D and lead us to similar conclusions. As expected, overall performance across all classes improves with additional training data. Interestingly, even when the same amount of data is present for each demographic group, error rates remain higher for PoCs than for White users. This suggests that other factors beyond superficial representation are to blame for model degradation. It is also worthwhile noting that performance continues to improve for underrepresented groups after they have been fully saturated, thus implying that at least some signal generalizes between demographics.

		MULTITASK			
		full	NESARC	CDC	even
Female	White	0.84 ± 0.002	0.75 ± 0.006	0.75 ± 0.008	0.69 ± 0.013
	PoC	0.91 ± 0.003	0.87 ± 0.005	0.86 ± 0.005	0.82 ± 0.009
Male	White	0.83 ± 0.005	0.74 ± 0.007	0.74 ± 0.008	0.68 ± 0.011
	PoC	0.45 ± 0.016	0.56 ± 0.029	0.54 ± 0.033	0.48 ± 0.030
Equal Odds		0.13 ± 0.013	0.14 ± 0.021	0.14 ± 0.019	0.12 ± 0.014
Opportunity		0.18 ± 0.010	0.16 ± 0.032	0.18 ± 0.036	0.12 ± 0.027

Table 2: Avg. F1 with 95% conf. interval from bootstrap across gender and ethnicity groups, and absolute avg. equal odds and equal opportunity differences. Balanced models close the performance difference gap at the cost of overall model performance.

6.3.2 Data Balance

What other factors could account for the difference in model performance on PoC? Below, we examine the effect of balancing the training data for the demographic groups.

Methods. We consider MULTITASK when constructing demographically-balanced datasets. As explored in Section 6.1, there are two different estimates of depression rates in demographic groups: CDC and NESARC. We balance MULTITASK to match the depression rates of both estimates, and name the models trained on those datasets MULTITASK CDC and MULTITASK NESARC respectively. Additionally, we compare these with an *even* balanced distribution. Models are trained following the methodology in Section 6.2.

Results. Table 2 shows the average F1 score of classifiers across gender and race/ethnicity groups. We copy MULTITASK column from Table 1 (labeled as *full*) for ease of comparison. There is a performance difference between male PoC users and the rest of the groups in models trained on MULTITASK balanced datasets, similar to MULTITASK full. However, the performance difference is smaller on models trained on balanced datasets. We observe no difference between balancing datasets according to NESARC or CDC, despite the 1.25x White user population increase in CDC. While fairness performance of both NESARC or CDC are similar to the full dataset, the *even* dataset shows considerable improvement for both fairness metrics at the cost of model performance.

Our experiments with both dataset size and balance show that it matters when datasets are not demographically representative, and as shown in section 6.1, they are not.




Figure 3: Percent of tweets containing first-person pronoun (*Pro1*) previously shown to correlate with the depression group; shaded notches show median CI (95%) with bootstrap $n = 10000$. Correlations shown by other work are not universal among demographic groups.

6.4 Can differences in features between demographic groups account for classifier biases?

We have demonstrated a demographic bias in classifiers trained on CLPSYCH and MULTITASK. Perhaps differences in feature representations between the groups can explain some of this bias. We examine LIWC features, which previous research has identified as useful in depression classification (Coppersmith et al., 2015c, 2014), in addition to performance analysis in Appendix E.

Methods. Previous research using the CLPSYCH and MULTITASK datasets has identified LIWC dimensions that over-index amongst depressed individuals: *Negative Emotion* (negemo), *Swearing* (swear), *Anger* (anger), *Anxiety* (anx) and *First-person Pronoun Usage* (Pro1) (Coppersmith et al., 2014).³ We evaluate whether this finding holds within each demographic group independently and whether there exist shifts between demographic groups.

Results. Figure 3 shows the distribution over users of percentage of tweets with at least one word matching the *Pro1* category across demographic groups, with shaded notches showing median confidence interval. Results for other LIWC categories associated with depression are similar to those for

³*Pro1* is constructed by combining the *i* and *we* LIWC categories.

Pro1 (see Appendix F).

From previous research, we expect to observe a greater *Pro1* prevalence in depression groups compared to controls across all demographics i.e. shaded notches of depression box should not overlap with control in each demographic category. However, in CLPSYCH, we do not observe any difference in prevalence of *Pro1* in the PoC groups, and in MULTITASK we do not observe any difference in prevalence in the male PoC group, both contradicting previous results. We also observe a correlation between prevalence of these LIWC categories in the depression group and downstream model performance for each demographic group, corroborating previous findings of the correlation of LIWC categories and depression signals (Coppersmith et al., 2015c, 2014). In general, female groups for both the control and depression sets tend to have a higher prevalence of *Pro1* compared to their male counterparts, suggesting a difference in language between the groups.

In short, LIWC correlations with depression are not universal across demographic groups. Furthermore, a closed vocabulary feature, such as LIWC, may contribute towards bias against some demographic groups.

7 Limitations

Depression and Control Groups. The method used to curate the depression group in these datasets is susceptible to *self-selection* bias, as noted by Coppersmith et al. (2014) and Amir et al. (2019), as it likely over-represents individuals who are more vocal about their condition. Therefore, differences in use of social media and cultural perceptions around mental health may introduce biases in these datasets. Further, while expert annotators identified non-genuine disclosures depression and removed these individuals from the CLPSYCH and MULTITASK datasets, they did not verify the authenticity of the diagnosis. Similarly, individuals in the control group may have been actually diagnosed with depression, but did not disclose their condition anywhere in their public timeline. Thus, labels for both the the depression and control groups are bound to be noisy.

Representation. We balance the MULTITASK dataset to match depression rates in the US, which may not be representative of non-US populations. Additionally, we preserve the even depression/control splits for class balance in model training, in-

stead of using the true depression/control population rates (about 1/10). These splits are not reflective of the true depression prevalence and their use may need to be modified depending on downstream classifier use.

Demographic Labels. Due to dataset limitations, the ethnicity and gender labels for this study were inferred using the unigram model from (Wood-Doughty et al., 2020). This model considers only the four largest race/ethnicity groups in the US, which aligns with conventions from Brody et al. (2018) but ignores smaller populations and multi-racial categories. Further, in some of our analyses, we combine Asian, Black and Hispanic/Latinx individuals into people of color due to a lack of data. With respect to gender, the male and female labels used by this model do not consider individuals who fall outside of traditional binary gender. As our experiments rely on upstream demographic label inference, we cannot fully rule out confounding factors due to e.g. noisy labels in our experimentation, but we perform a high-confidence filter on demographic labels and statistical testing on results to strengthen our conclusions.

8 Conclusion and Recommendations

We examine whether datasets and the resulting trained classifiers for depression prediction are fair across demographic groups. Our analysis finds that (1) depression datasets are **not** demographically representative, in some cases excluding entire intersectional groups and (2) the resulting classifiers perform worse on people of color in general. In examining the reason for these differences, we find that performance difference could be improved after accounting for (3) the size of the dataset and balance across demographic groups. (4) Finally, we show that signals of depression found by previous work using e.g. LIWC features are *not* equally representative for all demographics.

These findings should give pause to researchers in this area. Since datasets and the resulting models are not demographically representative, advances in methods may be furthering biases towards some groups. Worse, since some intersectional demographic groups are not even represented in the data, compounded by the fact that most datasets do not have labels for demographic groups, we currently lack the means to even check how new methods perform on each group. Going forward, research in this area should include demographic analyses

so that improvements on the overall dataset can be contextualized by how they perform on each demographic group.

At the same time, there is reason for optimism. Our data balancing and dataset size experiments reduced demographic disparities of trained models. This suggests that research can continue with existing datasets, but with the modifications we proposed. We release the demographically balanced dataset from our experimentation upon appropriate terms of usage agreement.

Ultimately, the best approach will be to construct new datasets that better represent the population, especially underrepresented minorities who are most at risk from systematic bias. This may necessitate changes to the data collection methods themselves, which may bias collection against certain groups. For example, self-reports may be problematic as they rely on cultural attitudes towards the expression of mental health information. Further research is needed to understand if self-reports and other proxy-based methods for obtaining labels can be successfully adapted to include a more diverse population, e.g. do keywords used to collect tweets skew resulting user populations? Further, to produce more conclusive insights with respect to demographics, language-based classifiers for demographic labels need to be further improved. Alternatively, other data collection strategies, such as the cohort method of Amir et al. (2019), may be more successful at ensuring representative datasets.

Acknowledgments

The authors gratefully acknowledge Elizabeth Salesky, Zach Wood-Doughty, Rachel Wicks, Alexandra DeLucia and the CLSP NLP reading group for helpful feedback, and we thank the anonymous reviewers for their helpful comments.

References

- Silvio Amir, Glen Coppersmith, Paula Carvalho, Mário J. Silva, and Byron C. Wallace. 2017. [Quantifying mental health from social media with neural user embeddings](#).
- Silvio Amir, Mark Dredze, and John W. Ayers. 2019. [Mental health surveillance over social media with digital cohorts](#). In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pages 114–120, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alina Arseniev-Koehler, Sharon Mozgai, and Stefan

- Scherer. 2018. What type of happiness are you looking for?-a closer look at detecting mental health from language. In *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, pages 1–12.
- John W Ayers, Theodore L Caputi, Camille Nebeker, and Mark Dredze. 2018. Don’t quote me: reverse identification of research participants in social media studies. *NPJ digital medicine*, 1(1):1–2.
- Bark. [Best parental control monitoring app for iphone and android.](#)
- Adrian Benton, Glen Coppersmith, and Mark Dredze. 2017a. [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.
- Adrian Benton, Margaret Mitchell, and Dirk Hovy. 2017b. Multi-task learning for mental health using social media text. *ArXiv*, abs/1712.03538.
- Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2018. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, page 0049124118782533.
- Brittany E Blanchard, Angela K Stevens, Kenneth J Sher, and Andrew K Littlefield. 2020. Reexamining the psychometric properties of the substance use risk profile scale. *Assessment*, 27(3):454–471.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Colin R Blyth. 1972. On simpson’s paradox and the sure-thing principle. *Journal of the American Statistical Association*, 67(338):364–366.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, pages 4349–4357.
- Debra Brody, Laura Pratt, and Jeffery Hughes. 2018. Prevalence of depression among adults aged 20 and over: United states, 2013-2016. *NCHS data brief*, pages 1–8.
- John D. Burger, John Henderson, George Kim, and Guido Zarrella. 2011. [Discriminating gender on twitter.](#) In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1301–1309, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Stevie Chancellor and Munmun De Choudhury. 2020. Methods in predictive techniques for mental health status on social media: a critical review. *NPJ digital medicine*, 3(1):1–11.
- Hau-wen Chang, Dongwon Lee, Mohammed Eltaher, and Jeongkyu Lee. 2012. @ phillies tweeting from philly? predicting twitter user locations with spatial word usage. In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 111–118. IEEE.
- Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163.
- Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. [Quantifying mental health signals in twitter.](#) In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 51–60, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Glen Coppersmith, Mark Dredze, Craig Harman, and Kristy Hollingshead. 2015a. [From ADHD to SAD: Analyzing the language of mental health on twitter through self-reported diagnoses.](#) In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 1–10, Denver, Colorado. Association for Computational Linguistics.
- Glen Coppersmith, Mark Dredze, Craig Harman, Kristy Hollingshead, and Margaret Mitchell. 2015b. [CLPsych 2015 shared task: Depression and PTSD on twitter.](#) In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 31–39, Denver, Colorado. Association for Computational Linguistics.
- Glen Coppersmith, Ryan Leary, and Eric Whyne. 2015c. Quantifying suicidal ideation via language usage on social media.
- Sam Corbett-Davies and Sharad Goel. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*.
- Sasha Costanza-Chock. 2018. Design justice: towards an intersectional feminist framework for design theory and practice. *Proceedings of the Design Research Society*.
- Munmun De Choudhury, Scott Counts, Eric J. Horvitz, and Aaron Hoff. 2014. [Characterizing and predicting postpartum depression from shared facebook data.](#) In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW ’14*, page 626–638, New York, NY, USA. Association for Computing Machinery.
- Munmun De Choudhury, Michael Gamon, Scott Counts, and Eric Horvitz. 2013. [Predicting depression via social media.](#) AAAI.
- Munmun De Choudhury, Emre Kiciman, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. [Discovering shifts to suicidal ideation from mental](#)

- health content in social media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, page 2098–2110, New York, NY, USA. Association for Computing Machinery.
- Clayton Fink, Jonathon Kopecky, and Maksym Morawski. 2012. Inferring gender from the content of tweets: A region specific example. *ICWSM*, 12:459–462.
- Hila Gonen and Yoav Goldberg. 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. *arXiv preprint arXiv:1903.03862*.
- Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of opportunity in supervised learning. *arXiv preprint arXiv:1610.02413*.
- Keith Harrigan, Carlos Aguirre, and Mark Dredze. 2020. On the state of social media data for mental health research. *arXiv preprint arXiv:2011.05233*.
- Deborah S. Hasin, Aaron L. Sarvet, Jacquelyn L. Meyers, Tulshi D. Saha, W. June Ruan, Malka Stohl, and Bridget F. Grant. 2018. [Epidemiology of Adult DSM-5 Major Depressive Disorder and Its Specifiers in the United States](#). *JAMA Psychiatry*, 75(4):336–346.
- Joseph Henrich, Steven J. Heine, and Ara Norenzayan. 2010. [The weirdest people in the world?](#) *Behavioral and Brain Sciences*, 33(2-3):61–83.
- Dirk Hovy and Anders Søgaard. 2015. Tagging performance correlates with author age. In *Proceedings of the 53rd annual meeting of the Association for Computational Linguistics and the 7th international joint conference on natural language processing (volume 2: Short papers)*, pages 483–488.
- Dirk Hovy and Shannon L Spruit. 2016. The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598.
- Molly Ireland and Micah Iserman. 2018. [Within and between-person differences in language used across anxiety support and neutral Reddit communities](#). In *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, pages 182–193, New Orleans, LA. Association for Computational Linguistics.
- Anders Johannsen, Dirk Hovy, and Anders Søgaard. 2015. Cross-lingual syntactic variation over age and gender. In *Proceedings of the nineteenth conference on computational natural language learning*, pages 103–112.
- Ronald C Kessler, Patricia Berglund, Olga Demler, Robert Jin, Doreen Koretz, Kathleen R Merikangas, A John Rush, Ellen E Walters, and Philip S Wang. 2003. The epidemiology of major depressive disorder: results from the national comorbidity survey replication (ncs-r). *Jama*, 289(23):3095–3105.
- Brian Larson. 2017. [Gender as a variable in natural-language processing: Ethical considerations](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 1–11, Valencia, Spain. Association for Computational Linguistics.
- Veronica Lynn, Alissa Goodman, Kate Niederhoffer, Kate Loveys, Philip Resnik, and H. Andrew Schwartz. 2018. [CLPsych 2018 shared task: Predicting current and future psychological health from childhood essays](#). In *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, pages 37–46, New Orleans, LA. Association for Computational Linguistics.
- Kari McDonald. 2018. Social support and mental health in lgbtq adolescents: a review of the literature. *Issues in mental health nursing*, 39(1):16–29.
- Margaret Mitchell, Kristy Hollingshead, and Glen Coppersmith. 2015. [Quantifying the language of schizophrenia in social media](#). In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 11–20, Denver, Colorado. Association for Computational Linguistics.
- Shirin Nilizadeh, Anne Groggel, Peter Lista, Srijita Das, Yong-Yeol Ahn, Apu Kapadia, and Fabio Rojas. 2016. Twitter’s glass ceiling: The effect of perceived gender on online visibility. In *Tenth International AAAI Conference on Web and Social Media*.
- Brendan O’Connor, Michel Krieger, and David Ahn. 2010. Tweetmotif: Exploratory search and topic summarization for twitter. In *Fourth International AAAI Conference on Weblogs and Social Media*.
- Robin Ohannessian, Tu Anh Duong, and Anna Odone. 2020. Global telemedicine implementation and integration within health systems to fight the covid-19 pandemic: a call to action. *JMIR public health and surveillance*, 6(2):e18810.
- James W Pennebaker, Roger J Booth, and Martha E Francis. 2007. Operator’s manual: Linguistic inquiry and word count: Liwc2007. Austin, Texas: LIWC. net <http://homepage.psy.utexas.edu/HomePage/Faculty/Pennebaker/Reprints/LIWC2007.OperatorManual.pdf> (accessed 1 October 2013).
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Daniel Preotjiuc-Pietro, Johannes Eichstaedt, Gregory Park, Maarten Sap, Laura Smith, Victoria Tobolsky, H. Andrew Schwartz, and Lyle Ungar. 2015. [The role of personality, age, and gender in tweeting about mental illness](#). In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 21–30, Denver, Colorado. Association for Computational Linguistics.

- Juan Ramos et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, volume 242, pages 133–142. Piscataway, NJ.
- Delip Rao, Michael J Paul, Clayton Fink, David Yarowsky, Timothy Oates, and Glen Coppersmith. 2011. Hierarchical bayesian models for latent attribute detection in social media. *ICWSM*, 11:598–601.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. Gender bias in coreference resolution. *arXiv preprint arXiv:1804.09301*.
- H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*, 8(9):e73791.
- Sherri M Simpson, Laura L Krishnan, Mark E Kunik, and Pedro Ruiz. 2007. Racial disparities in diagnosis and treatment of depression: a literature review. *Psychiatric Quarterly*, 78(1):3–14.
- Nikhita Vedula and Srinivasan Parthasarathy. 2017. Emotional and linguistic cues of depression from social media. In *Proceedings of the 2017 International Conference on Digital Health*, pages 127–136.
- Svitlana Volkova, Glen Coppersmith, and Benjamin Van Durme. 2014. [Inferring user political preferences from streaming communications](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 186–196, Baltimore, Maryland. Association for Computational Linguistics.
- Stefan Wojcik and Adam Hughes. 2019. [How twitter users compare to the general public](#).
- JT Wolohan, Misato Hiraga, Atreyee Mukherjee, Zee-shan Ali Sayyed, and Matthew Millard. 2018. [Detecting linguistic traces of depression in topic-restricted text: Attending to self-stigmatized depression with NLP](#). In *Proceedings of the First International Workshop on Language Cognition and Computational Models*, pages 11–21, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Zach Wood-Doughty, Nicholas Andrews, Rebecca Marvin, and Mark Dredze. 2018. [Predicting twitter user demographics from names alone](#). In *Proceedings of the Second Workshop on Computational Modeling of People’s Opinions, Personality, and Emotions in Social Media*, pages 105–111, New Orleans, Louisiana, USA. Association for Computational Linguistics.
- Zach Wood-Doughty, Paiheng Xu, Xiao Liu, and Mark Dredze. 2020. [Using noisy self-reports to predict twitter user demographics](#).
- Amir Hossein Yazdavar, Mohammad Saeid Mahdavinejad, Goonmeet Bajaj, William Romine, Amit Sheth, Amir Hassan Monadjemi, Krishnaprasad Thirunarayan, John M Meddar, Annie Myers, Jyotishman Pathak, et al. 2020. Multimodal mental health analysis in social media. *Plos one*, 15(4):e0226248.
- Dong Whi Yoo, Michael L Birnbaum, Anna R Van Meter, Asra F Ali, Elizabeth Arenare, Gregory D Abowd, and Munmun De Choudhury. 2020. Designing a clinician-facing tool for using insights from patients’ social media activity: Iterative co-design approach. *JMIR Mental Health*, 7(8):e16969.
- Dong Whi Yoo and Munmun De Choudhury. 2019. Designing dashboard for campus stakeholders to support college student mental health. In *Proceedings of the 13th EAI International Conference on Pervasive Computing Technologies for Healthcare*, pages 61–70.
- Xiaoyun Zhou, Centaine L Snoswell, Louise E Harding, Matthew Bambling, Sisira Edirippulige, Xuejun Bai, and Anthony C Smith. 2020. The role of telehealth in reducing the mental health burden from covid-19. *Telemedicine and e-Health*, 26(4):377–379.

A Demographic Labels Analysis

Due to data anonymization, we must utilize content-based demographic classifiers to infer labels. This introduces noise due to classification error to our analyses. In order to reduce these effects we consider 3 techniques: high confidence filter, name labels, and tweet sampling.

High Confidence Filter. We select users whose most probably demographic class for both gender and race is > 0.95 probability.

Random Sampling. Considers all available users; randomly split each individual’s tweets into two independent pools so that demographic and mental health inferences are based on separate sets of data.

Name Labels. With the permission of [Coppersmith et al. \(2014\)](#), we were able to collect name attributes for 622 of individuals in CLPSYCH. To obtain demographic labels from name attributes we used the demographer’s neural name classifier ([Wood-Doughty et al., 2018](#)), and ethnic/race name classifier ([Wood-Doughty et al., 2020](#)). Due to the small number of individuals that we could obtain name attributes, our analysis of this technique is limited.

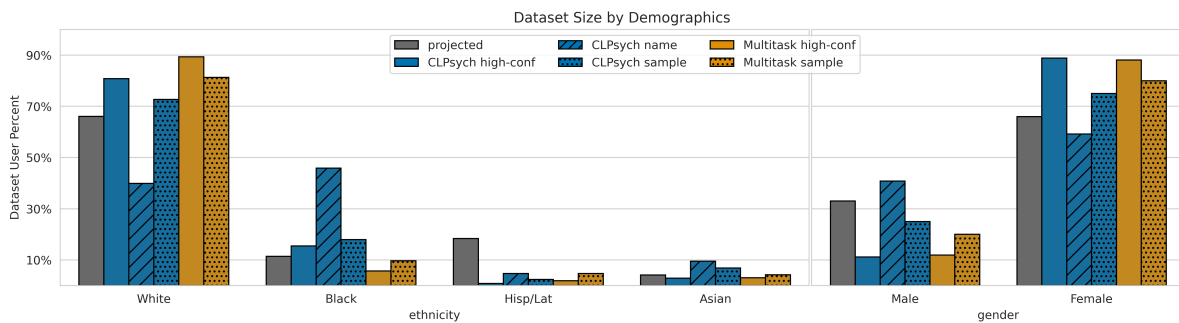


Figure 4: Hispanic/Latinx and Male individuals are underrepresented in both of our Twitter datasets (CLPSYCH & MULTITASK) compared to the Twitter US population estimates (NESARC) in all our labeling techniques.

Figure 4 shows the population percentage of the three techniques compared to the projected distribution. For all the techniques, some trends still hold, mainly Hispanic/Latinx group is underrepresented. Additionally, for the high confidence and split techniques, White and female groups are overrepresented and male underrepresented. In contrast, in the name based approach, the Black and male groups are overrepresented while the White and female groups are underrepresented.

While the name based technique shows more promising data distributions, the Hispanic/Latinx group is still vastly underrepresented. However, the rest of the demographic groups seem more fair and closer (sometimes better) than the target distribution, but do these distributions translate to greater downstream performance?

Table 3 shows the performance of mental health models trained on the different datasets obtained from our demographic labeling techniques. High confidence filtering has the largest performance difference and higher fairness metrics compared to random sampling and name based approaches. However, we still observe similar trends with PoC groups performing in general lower than White groups.

		CLPSYCH			MULTITASK		
		High Conf.	Rand. Sample	Name Labels	High Conf.	Rand. Sample	
Female	White	<i>0.77 ± 0.005</i>	0.77 ± 0.005	0.72 ± 0.006	0.84 ± 0.002	0.84 ± 0.003	
	PoC	<i>0.41 ± 0.013</i>	0.47 ± 0.020	0.44 ± 0.018	0.91 ± 0.003	0.87 ± 0.010	
Male	White	<i>0.74 ± 0.008</i>	0.74 ± 0.014	0.48 ± 0.032	0.83 ± 0.005	0.78 ± 0.011	
	PoC	<i>0.76 ± 0.035</i>	0.72 ± 0.040	0.56 ± 0.042	0.45 ± 0.016	0.78 ± 0.033	
Equal Odds		0.21 ± 0.023	0.13 ± 0.014	0.19 ± 0.023	0.13 ± 0.013	0.13 ± 0.021	
Opportunity		0.25 ± 0.039	0.16 ± 0.022	0.18 ± 0.031	0.18 ± 0.010	0.14 ± 0.034	

Table 3: Avg. F1 with 95% conf. interval from bootstrap across gender and ethnicity groups, and avg. equal odds and equal opportunity differences (*italics: not significant*). Models underperform for PoC in general, sometimes male PoC (MULTITASK) or female PoC (CLPSYCH).

B Tweets and Vocabulary in Demographics

Figure 5 shows the distribution of the average number of tweets and vocabulary size across all our demographic groups. There is no statistical significant difference in the means between demographic groups within the datasets. Very high variance is observed in the Asian group in multitask, although their vocabulary size is not as high variance. Additionally, MULTITASK dataset has on average higher number of tweets per user (as they were not limited to 3000 as in CLPSYCH) and consequently a higher average vocabulary size.




Figure 5: Average tweet number per user per demographic category, as well as average vocabulary size. Black bars show 95% CI.

C Model Specifications

Tokenization. Raw text within in Tweets was tokenized using a modified version of the Twokenizer (O’Connor et al., 2010). English contractions were expanded, while specific retweet tokens, username mentions, URLs, and numeric values were replaced by generic tokens. As pronoun usage tends to differ in individuals living with depression (Vedula and Parthasarathy, 2017), we removed any English pronouns from our stop word set (English Stop Words from nltk.org). Case was standardized across all tokens, with a single flag included if an entire post was made in uppercase letters.

Features. Text from all documents for an individual are concatenated together and tokenized as previously described. The vocabulary of each training procedure is fixed to a maximum of 100-thousand unigrams selected based on KL-divergence of the class-unigram distribution with the class-distribution of stop words (Chang et al., 2012). This reduced bag-of-words representation is then used to generate the following additional feature dimensions: a 50-dimensional LDA topic distribution (Blei et al., 2003), a 64-dimensional LIWC category distribution (Pennebaker et al., 2007), and a 200-dimensional mean-pooled vector of GloVe embeddings (Pennington et al., 2014). The reduced bag-of-words representation is transformed using TF-IDF weighting (Ramos et al., 2003).⁴

Hyperparameter Selection. Each model is trained using a hyperparameter grid search over the regularization strength {1e-3, 1e-2, 1-e1, 1, 10, 100, 1e3, 1e4, 1e5}, class weighting {None, Balanced}, and feature set standardization {On, Off}. Hyperparameters were selected to maximize held-out F1 score within a 20%-sied held-out split of the training data.

⁴All data-specific feature transformations (e.g. LDA, TF-IDF) are learned without access to development or test data.

D Multitask Size Experiment

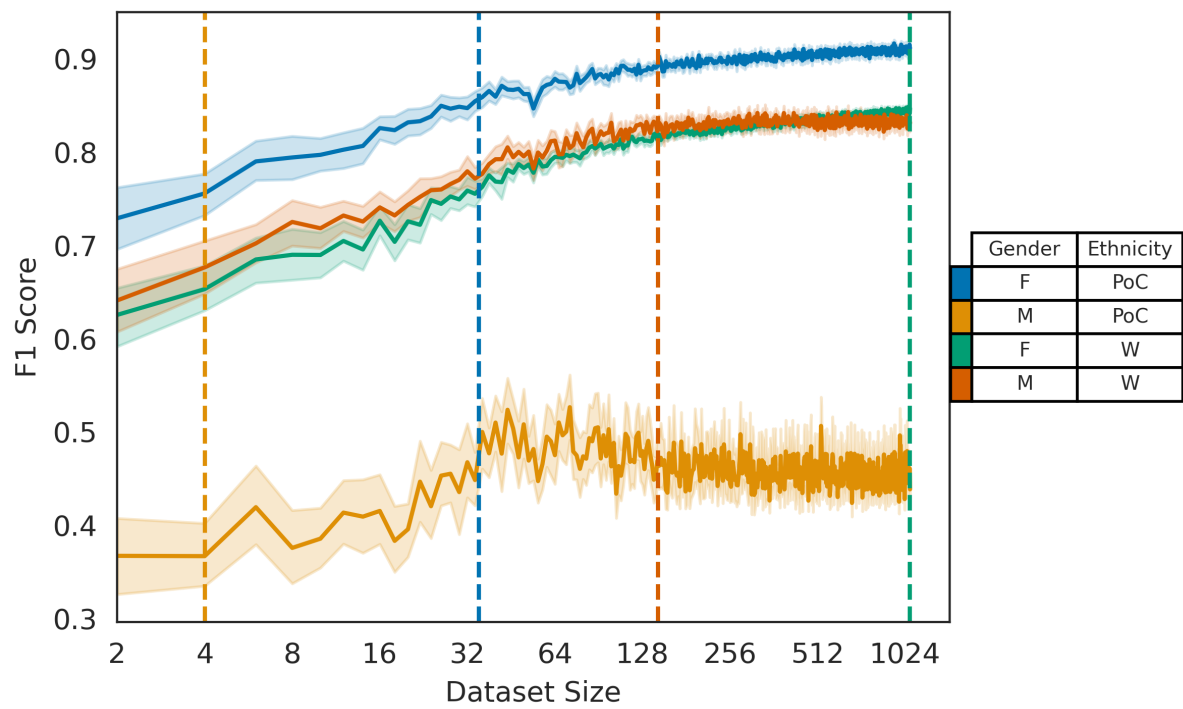


Figure 6: Log learning curve of varying dataset sample size on MULTITASK. Dashed vertical lines represent the total size of the demographic group.

E Feature Study

We observed performance difference between demographic groups for both our datasets with mental health models using the following feature groups: LIWC, LDA, GloVe and TF-IDF as specified in Appendix C. To explore the source of the performance difference observed in demographic groups, we train classifiers with each individual feature group and present average F1 score per demographic groups as well as our fairness metrics in Table 4.

TF-IDF and GloVe embeddings yield better model performance than the other feature groups at the expense of fairness as measure by our fairness metrics. The most fair feature set was LIWC, although it was also the least informative feature set resulting in worst performing models.

However, trends observed in Section 6.2 still apply in all feature groups, mainly models tend to underperform for PoC groups (female PoC groups in CLPSYCH).

		CLPSYCH									
		full		TF-IDF		LDA		LIWC		GloVe	
Female	White	0.77	± 0.005	0.77	± 0.005	0.70	± 0.008	0.68	± 0.010	0.71	± 0.007
	PoC	0.41	± 0.013	0.42	± 0.001	0.41	± 0.026	0.45	± 0.020	0.47	± 0.020
Male	White	0.74	± 0.008	0.74	± 0.008	0.65	± 0.019	0.68	± 0.021	0.68	± 0.019
	PoC	0.76	± 0.035	0.74	± 0.034	0.62	± 0.054	0.55	± 0.040	0.75	± 0.032
Equal Odds		0.21	± 0.023	0.20	± 0.023	0.21	± 0.031	0.17	± 0.021	0.21	± 0.022
Opportunity		0.25	± 0.039	0.25	± 0.040	0.24	± 0.039	0.17	± 0.028	0.18	± 0.029

Table 4: Avg. F1 with 95% conf. interval from bootstrap across gender and ethnicity groups, and absolute avg. equal odds and equal opportunity differences. PoC groups (female PoC) perform worse for models separately trained on each of our feature group.

F Additional LIWC Categories Figures

Previous research has identified specific LIWC dimensions that are of importance in depression groups. In addition to *First-person pronoun* (pro1), categories like *Negative Emotion* (negemo), *Swearing* (swear), *Anger* (anger) and *Anxiety* (anx) have been shown to be more prominent in the depression groups compared to control. We observe variation on prevalence on depression groups in all categories mentioned above among demographic groups, showing that LIWC features are *not* equally representative for all demographics




Figure 7: Negative Emotion Usage (*negemo*) LIWC category representation within each individual, previously shown to correlate with the depression group; statistical significance marked by not overlapping shaded notches. We observe median statistical difference only for **White male** individuals in CLPSYCH and **White groups** in MULTITASK.

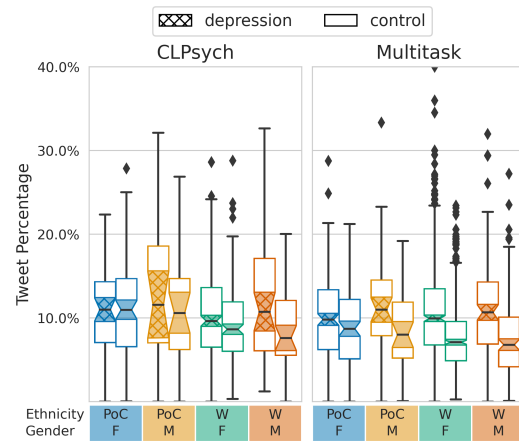


Figure 8: Anger (*anger*) LIWC category representation within each individual, previously shown to correlate with the depression group; statistical significance marked by not overlapping shaded notches. We only observe median statistical difference for **no groups** CLPSYCH and **White groups** in MULTITASK




Figure 9: Anxiety (*anx*) LIWC category representation within each individual, previously shown to correlate with the depression group; statistical significance marked by not overlapping shaded notches. We observe median statistical difference for the **male groups** in CLPSYCH and **all** demographic categories in MULTITASK




Figure 10: Swearing (*swear*) LIWC category representation within each individual, previously shown to correlate with the depression group; statistical significance marked by not overlapping shaded notches. We observe median statistical difference for **no groups** in CLPSYCH and **White groups** in MULTITASK.