

# Minority Stress Experienced by LGBTQ Online Communities during the COVID-19 Pandemic

Yunhao Yuan<sup>1</sup>, Gaurav Verma<sup>2</sup>, Barbara Keller<sup>1</sup>, Talayeh Aledavood<sup>1</sup>

<sup>1</sup> Aalto University

<sup>2</sup> Georgia Institute of Technology

yunhao.yuan@aalto.fi, gverma@gatech.edu, barbara.keller@aalto.fi, talayeh.aledavood@aalto.fi

## Abstract

The COVID-19 pandemic has disproportionately impacted the lives of minorities, such as members of the LGBTQ community (lesbian, gay, bisexual, transgender, and queer) due to pre-existing social disadvantages and health disparities. Although extensive research has been carried out on the impact of the COVID-19 pandemic on different aspects of the general population's lives, few studies are focused on the LGBTQ population. In this paper, we develop and evaluate two sets of machine learning classifiers using a pre-pandemic and a during-pandemic dataset to identify Twitter posts exhibiting minority stress, which is a unique pressure faced by the members of the LGBTQ population due to their sexual and gender identities. We demonstrate that our best pre- and during-pandemic models show strong and stable performance for detecting posts that contain minority stress. We investigate the linguistic differences in minority stress posts across pre- and during-pandemic periods. We find that anger words are strongly associated with minority stress during the COVID-19 pandemic. We explore the impact of the pandemic on the emotional states of the LGBTQ population by adopting propensity score-based matching to perform a causal analysis. The results show that the LGBTQ population have a greater increase in the usage of cognitive words and worsened observable attribute in the usage of positive emotion words than the group of the general population with similar pre-pandemic behavioral attributes. Our findings have implications for the public health domain and policy-makers to provide adequate support, especially with respect to mental health, to the LGBTQ population during future crises.

## Introduction

The onset of the COVID-19 pandemic has profoundly impacted health, economy, and societies as a whole. It leads to increased anxiety, depression, and stress as people cope with fear of contagion, and uncertainty about the future. Across the world, a wide range of substantive strategies, such as self-isolation, social distancing, and school closures, have been implemented to mitigate the spread of the coronavirus. These policies, which effectively slow the spread of the virus, have exacerbated mental health issues for individuals and communities.

Isolation from social relationships and disengagement from social activities can lead to adverse mental health outcomes, such as depression, anxiety, and loneliness (Pietrabissa and Simpson 2020). Recent studies on the effects of the COVID-19 pandemic show a considerable increase in psychological expressions, including expressions related to mental health issues and support, among the general population (Saha et al. 2020; Gallagher et al. 2020). The authors highlight major concerns about personal challenges, health issues, and uncertainties surrounding the pandemic. The adverse effects of these policy interventions are shown to have a disproportional impact on different groups and communities within society. For example, recent work by Pierce et al. (2020) assesses the level of mental distress in the UK population before and during the lockdowns. Their research indicates deteriorated mental health during the pandemic, especially among women, youth, and children.

Due to pre-existing social and structural inequalities, the COVID-19 pandemic is expected to have a greater impact on minority groups, including lesbian, gay, bisexual, transgender, and queer (LGBTQ) individuals. Even in the best of times, LGBTQ community members have experienced worse mental health outcomes resulting from persistent discrimination and stigmatization (Wang et al. 2020). This community struggles with health disparities, such as higher occurrences of long-term chronic illnesses and unique healthcare needs, compared to the straight and cisgender population (Zeeman et al. 2019). During the COVID-19 pandemic, Kneale and Becares (2020) report an increased level of stress and depressive symptoms among the LGBTQ community, particularly among transgender people. It is suggested that the closing of educational institutions and mandatory isolation during the pandemic may expose young members of the LGBTQ community to verbal abuse and domestic violence (Salerno, Williams, and Gattamorta 2020). However, research to date mainly relies on diagnostic reports and psychological questionnaires to analyze the mental health of the LGBTQ population. Due to the fear of cultural discrimination and stigmatization, a large number of LGBTQ individuals are hard to reach, limiting generalizability. To fill this gap, our study aims to investigate the mental health of the LGBTQ population during the pandemic by leveraging a large sample of social media data.

Social media are known to provide spaces where indi-

viduals and communities with stigmatized identities can express themselves, talk about their experiences, and seek support, with less fear of offline harm (Andalibi et al. 2016). Data collected from different social media platforms have already been used to investigate stigmatized experiences in various domains such as drug misuse (Garg et al. 2021), sexual abuse (Andalibi et al. 2016), and mental health (Coppersmith et al. 2018). Similarly, social media offers an easily accessible and semi-anonymous avenue for LGBTQ individuals to connect with others, seek support and share their concerns (Cannon et al. 2017). Consequently, social media data provide an unparalleled opportunity to study the concerns of the LGBTQ community and their evolution over time. Harnessing this opportunity is especially important in the context of the COVID-19 pandemic, when mental health issues are known to have grown significantly among different populations.

In this work, we investigate the impact of COVID-19 on the LGBTQ community in the United States via Twitter data. We study *minority stress*, which is defined as the chronic stress derived from stigmatized social status in the context of a heterosexist society faced by the LGBTQ community members due to their sexual and gender identities (Meyer 1995). We address the following three research questions to better understand shifts in the minority stress patterns before and after the breakout of the COVID-19 pandemic and its disproportionate impacts on the LGBTQ population as compared to the general population:

**RQ1:** *How can we measure minority stress on social media using computational methods?*

**RQ2:** *Are there linguistic differences in Twitter posts (tweets) containing minority stress in the time prior to the pandemic compared to minority stress posts during the pandemic?*

**RQ3:** *Are there linguistic differences in tweets between the LGBTQ and the general population during the pandemic?*

To answer these research questions, we collect 20,593,823 tweets from 7,241 US-based LGBTQ Twitter users who disclose their identities on Twitter from June 2018 to June 2021. We manually annotate 2,800 tweets and divide them into two datasets. One set contains tweets posted before the COVID-19 pandemic, and the other set includes tweets posted during the pandemic. We train multiple machine learning classifiers separately on pre- and during-pandemic datasets, and evaluate the best-performing classifiers on different feature sets. Our best pre-pandemic and during-pandemic models show stable and robust performance with an area under the receiver operating characteristic (AUC) curve of 0.838 and 0.845, respectively. We investigate the top statistically significant features of the pre-pandemic and during-pandemic models to compare the emotions behind the minority stress posts on Twitter. We find that anger has become strongly associated with minority stress during the COVID-19 pandemic. Lastly, we apply a stratified propensity score analysis to compare the COVID-19 pandemic effects on the LGBTQ population and the general population on social media. We observe that

the COVID-19 pandemic has resulted in a higher usage of cognitive words and a reduced usage of positive emotion words for the LGBTQ population.

**Privacy and Ethics** Despite the public accessibility of Twitter data, we adopt several approaches to protect the privacy of the data owners. We remove any information related to personal identity and paraphrase all quotations in this paper to avoid traceability. One of the authors is a member of the LGBTQ community. This helps us better understand the privacy needs and concerns of this community.

## Background and Related Work

### Social Media Use During the COVID-19 Pandemic

Due to the need for social distancing during the COVID-19 pandemic, people increasingly use social media not only to consume and share information, but also to engage in seeking help, advice, and providing supportive information (Saud, Mashud, and Ida 2020). Recent work by Tsao et al. (2021) highlights various facets of public health research on social media during the pandemic, such as psychological reactions, government responses, and public attitudes.

Xiang et al. (2021) collect data from Twitter and DeChose to analyze the public discourse and sentiments relative to older people and the extent of ageism among social media users. Zhou et al. (2021) build a depression classifier by extracting multi-modal features from Twitter data, such as emotion, topic, and domain-specific features, and analyze the local dynamics during the COVID-19 in Australia. Verma et al. (2022) investigate the 80 million Twitter posts and demonstrate that users sharing COVID-19 related misinformation will experience exacerbated anxiety.

### LGBTQ Community on Social Media

Social media plays a prominent role in the exchange of information. It also provides opportunities and platforms for minorities and communities, such as LGBTQ individuals, to engage in online activities, such as seeking out social support, exchanging information, and maintaining social relationships (Steinke et al. 2017). Prior research suggests that LGBTQ youths use social media as a primary method of socializing and spend more time online than their counterpart cisgender individuals (Steinke et al. 2017).

Given the significant growth in social media usage, data from LGBTQ online communities has been used to analyze social media's effects on the LGBTQ population. One study (Craig and McInroy 2014) illustrates that LGBTQ youth use social media platforms as a tool for actively coming out as they are perceived as safe places. In another study, researchers demonstrate the feasibility of evaluating equality of hospital care for LGBT individuals through analyzing online discussions related to LGBT hospital care on Twitter (Hswen et al. 2018). Haimson (2019) show that the positive sentiment increases over time on average after disclosing their identities to people with close relationships by analyzing posts from Tumblr. Overall, existing research indicates a need to understand better the relationship between the LGBTQ community and their use of social media.

## Minority Stress Theory

The minority stress theory (Meyer 1995) provides a full conceptual framework for understanding how minority stressors impact the mental health of LGBTQ individuals. It defines minority stressors as chronic stress derived from stigmatized social status in the context of a heterosexist society. Meyer (1995) suggests three types of minority stressors: internalized homophobia, perceived stigma, and prejudice events. Stressors specific to LGBTQ individuals typically contain social stigma, prejudice, and discrimination, such as verbal abuse, harassment, and physical violence.

The minority stress theory has been utilized as a theoretical framework for numerous studies to examine the mental health of gender and sexual minorities. Previous studies have found widespread mental health disparities between the LGBTQ and the broader population in generalized anxiety symptoms (Borgogna et al. 2019), depression (Kulick et al. 2017), and substance abuse (Goldbach, Mereish, and Burgess 2017). Lehavot and Simoni (2011) analyze the links of minority stress with mental health and substance use among women belonging to a sexual minority. Frost, Lehavot, and Meyer (2015) find that minority stress impacts the physical health of the LGBTQ population.

## Data

We leverage Twitter timeline data of users who disclose their identities on Twitter as LGBTQ members. The data collection includes four steps: 1) curating a list of LGBTQ related hashtags; 2) identifying LGBTQ users; 3) collecting the timeline data of these self-disclosing LGBTQ users; 4) building a control dataset from non-LGBTQ users. We apply a codebook from (Saha et al. 2019) to annotate tweets in the LGBTQ dataset.

**LGBTQ Related Hashtag List** The first step in this process is seeding our sample with two popular LGBTQ relevant hashtags, #gay and #lgbt. We use the Twitter Application Programming Interface (API) to collect the latest 10,000 tweets posted in the US for each seeding hashtag and rank the top 100 co-occurring hashtags by their frequencies. By manually examining the hashtags and removing those with wide-ranging or ambiguous meanings, we find 60 LGBTQ relevant hashtags (Full list in Supplemental Material). After collecting 10,000 tweets containing at least one of the LGBTQ relevant hashtags, we get a corpus of 600,000 LGBTQ relevant tweets.

### Twitter Data of Members of the LGBTQ Community

We fetch the bios of the accounts which posted these 600,000 LGBTQ relevant tweets. Bios are public descriptions of Twitter accounts with a maximum length of 160 characters, where users typically disclose information about who they are. We search these bios using regular expression matching (e.g., gay, lesbian, and bisexual) and pride-related emojis, such as the rainbow flag, pride flag, and transgender symbol. We retrieve 7,241 distinct accounts belonging to self-disclosing LGBTQ Twitter users. We randomly select 40 users from this set and manually check if their latest 50 tweets explicitly express their identities. Based on manual

verification done by the authors, we confirm 32 users to be LGBTQ individuals. The bios of the eight remaining users contain terms with LGBTQ keywords or emojis, which are either supportive expressions (e.g., *LGBTQ life matters*) or declare their relationship with the LGBTQ community (e.g., *LGBTQ\* Advisory Committee in College*).

**Compiling the LGBTQ Dataset** From the 7,241 unique LGBTQ users, we collect Twitter metadata, including the number of tweets, likes, followers, followees, and the account creation time. Inspired by Pavalanathan and Eisenstein (2016), we filter out untypical users, such as inactive users and advertising accounts by removing users who have more than 5000 followers/followees or who posted fewer than 200 or more than 30,000 tweets. For the remaining 5,708 users, we collect their timeline data from June 1st, 2018 to June 1st, 2021. In the end, the *LGBTQ* dataset contains 5,708 users. Among these users, 972 users identify as gay, 173 users as lesbian, 1,424 users as transgender, and 95 users as bisexual. The remaining 3,044 users, do identify as part of the LGBTQ community but are not categorized into one of the former categories.

**Compiling the Control Dataset** To compare the linguistic differences between the LGBTQ and the general population during the pandemic, we build a *Control* dataset of users who do not use LGBTQ keywords in their bios. Recall that our RQ3 is to compare the linguistic difference between the general and the LGBTQ population. We seek to isolate the effect of the pandemic on two groups of users with similar pre-pandemic attributes. To find non-LGBTQ users who have similar linguistic patterns to the LGBTQ ones, we rank the top 100 hashtags from all tweets in the LGBTQ dataset and manually remove hashtags that occur in our LGBTQ related hashtag list, which results in 79 hashtags (Full list in Supplemental Material). For each of these 79 hashtags, we collect 10,000 tweets and bios of their authors. After removing users whose bios contain LGBTQ relevant keywords and/or LGBTQ emojis, we collect the Twitter metadata of the users and limit the dataset using the same methods as for the LGBTQ dataset. We collect the timeline data of the remaining 36,030 users to build the *control* dataset.

### Annotating Minority Stress on Social Media

We use a codebook from previous research (Saha et al. 2019) to annotate minority stress on our LGBTQ dataset. The codebook is developed from Meyer's minority stress framework (Meyer 1995) to identify minority stress in social media. It conceptualizes minority stress as three categories: prejudice events, containing discrimination and violent experiences; stigma, involving expectations of being rejected because of one's own minority identity; and internalized homophobia, related to the negative attitudes toward oneself. Table 1 presents some paraphrased tweet examples classified into these categories of minority stress.

To analyze the linguistic differences in minority stress tweets between the time span before and during the COVID-19 pandemic, it is essential to keep the number of tweets in the training dataset balanced between the two time periods. We randomly selected 2800 tweets for annotation, with 1400

Category	Tweet Example
<b>Prejudice Events</b>	<i>I was feeling really good until I was called, and then the verbal threats started once they knew I was trans. I'm at a loss for words.</i>
<b>Perceived Stigma</b>	<i>I believe that my desire to be visibly gay in games is interfering with my enjoyment of the ones I have. I'm sick of modifications stitching things together and half-measures. I just want to play games and not wear a mask of myself.</i>
<b>Internalized LGBTQphobia</b>	<i>Being transgender is something I despise. Why couldn't I just be a cisgirl from the start? Every ray of optimism that you have ends up destroying you. Either you let it kill you for not being yourself, or you let the world kill you for being yourself.</i>

Table 1: Slightly paraphrased sample tweets that belong to different categories of minority stress.

tweets created between June 1st, 2018 and November 30th, 2019, referred to as the **pre-pandemic dataset**, and 1400 tweets created between 1st December, 2019 and June 1st, 2021, referred to as the **during-pandemic dataset**.

**Background of the Annotators** Four authors of this paper and two recruited annotators participate in the annotation process. The group consists of four females and two males, each with diverse cultural background. Besides, one of the annotators belongs to the sexual minority community, which might help in explaining the highly subjective, intrinsic, and complicated context of the tweets to be annotated and the final goals of the annotation task.

**Annotation Task** Four authors independently annotate 200 randomly-sampled tweets and revise the codebook by discussing the discrepancies. Together, the annotators decide only to annotate tweets discussing recent and present experiences pertaining to minority stress. Lacking the context of URLs and the complexity of checking the contents in URLs, the URLs in posts are removed, but the remaining content of the tweet is included in the dataset.

After establishing the annotation rules, two university students are recruited to annotate the remaining 2,600 tweets. They are asked to learn the codebook with the code explanations, discrepancies, and tweet examples. The two annotators annotate the same sample sets to avoid bias and increase accuracy. In case annotators have any uncertainties about a tweet, they independently discuss it with the authors. The discussion leads to substantial inter-rater reliability, with an overall Cohen’s kappa ( $\kappa$ ) coefficient (Landis and Koch 1977) of 0.727, which is an improvement over the pre-discussion ( $\kappa$ ) of 0.542. If there are any remaining tweets with no agreement after discussion, the annotators and authors examine them together, and the final labels are decided by majority voting. Out of the 2,800 annotated tweets, 684 are categorized as prejudice events, 306 as perceived stigma, and 36 tweets as internalized LGBTQ phobia. Minority stress subcategories are not mutually exclusive and tweets do not provide the full information on the context. So, we subsequently focus on the binary categorization of “minority stress” and “non-minority stress”.

## Methods

### Study Design and Rationale

Our study aims to answer three research questions — measuring minority stress using social media posts (RQ1), capturing the linguistic differences in minority stress posts by LGBTQ members before and during the COVID-19 pandemic (RQ2), and contrasting the linguistic difference of LGBTQ members with the general population to explain their emotional reactions to the pandemic (RQ3). For RQ1, we conduct annotations based on the existing taxonomy of minority stress and develop machine learning classifiers using the annotated data. For RQ2, we calculate the coefficient scores for every input feature and analyze the most discriminative features for the best-performing models. RQ3, however, involves additional considerations — given that the linguistic difference of LGBTQ members can precipitate due to a variety of reasons beyond their sexual and gender identities, it is essential to contrast these differences with carefully chosen members of the general population. By closely matching a series of pre-pandemic attributes between members of the general population and members of the LGBTQ population, we try to isolate the effect of belonging to a minority group on the respective emotional response to the pandemic. To this end, we adopt a causal inference framework (Imbens and Rubin 2015) to address RQ3. The schematic diagram of our approach is shown in Figure 1. Our approach first matches members of the general population (i.e., the *control* group) with members of the LGBTQ group (i.e., the *treatment group*<sup>1</sup>) based on pre-pandemic behavioral attributes. For this, we train a machine learning classifier to estimate the likelihood of an individual being assigned to either the LGBTQ or control group (i.e., *propensity*) based on covariates and perform matching across groups using estimated propensity scores. Within matched groups of control and LGBTQ groups, we analyze the linguistic difference during the pandemic to compare the behavior of similar individuals who differ in terms of their

<sup>1</sup>LGBTQ group: Conventionally, the causal inference terminology uses “treatment” to refer to intervention or exposure of interest that is being studied, but we do not insinuate that someone of this group needs treatment or has been treated. In this paper, we use “LGBTQ group” to replace the term “treatment group”.

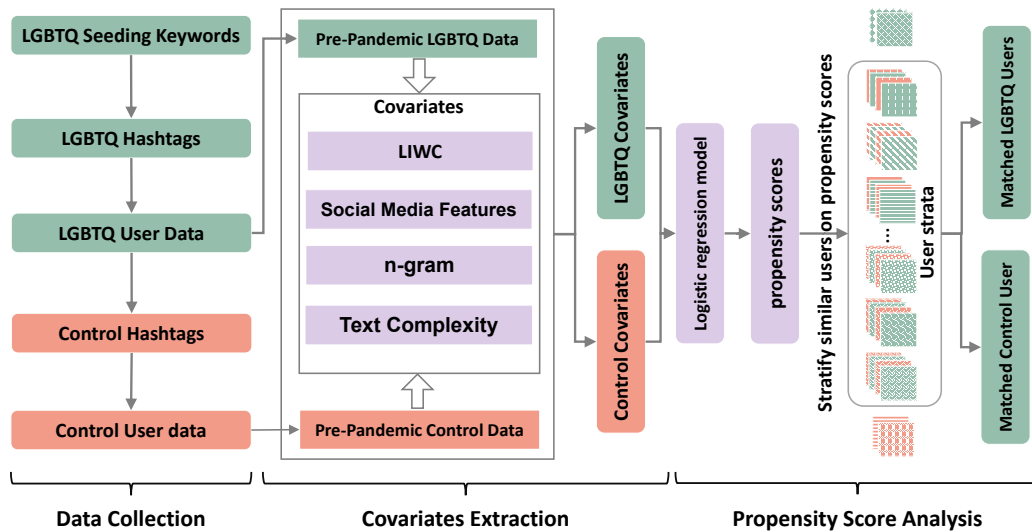


Figure 1: Schematic diagram of propensity score matching between LGBTQ users and general users.

sexual and/or gender identities. In sum, our approach ensures that members of the LGBTQ population and the general population who are being compared have similar behaviors prior to the pandemic. This gives us the means to analyze the linguistic differences between LGBTQ members, compared to the members of the general population.

### Machine Learning Classifier for Evaluating Minority Stress

Our methodological approach describes the way that we conceptualize the intrinsic language features within tweets and apply them to a series of machine learning models.

**Feature Engineering** Each tweet in the dataset is described as a vector of the following feature sets:

- **Latent Semantics (Word Embeddings):** To capture language semantics beyond raw keywords, we use word embeddings, which are essentially vector representations of words in latent semantic dimensions. In particular, we use pre-trained word embeddings (GloVe from StanfordNLP (Pennington, Socher, and Manning 2014)) in 50 dimensions that are trained on word-word co-occurrences in a Wikipedia corpus of 6B tokens.
- **Psycholinguistic Attributes (LIWC):** The use of Linguistic Inquiry and Word Count lexicon (Tausczik and Pennebaker 2010) is a well-established approach in extracting psycholinguistic features of texts in multiple dimensions, consisting of affection, cognition and perception, interpersonal focus, temporal references, lexical density and awareness, biological concerns, and social and personal concerns (De Choudhury and De 2014).
- **Sentiment:** Stanford CoreNLP’s deep learning based sentiment analysis tool (Manning et al. 2014) is used to measure the sentiment of the tweets. We use this tool to calculate the positive, negative, and neutral scores for each tweet.

- **Open Vocabulary (n-grams):** Open Vocabulary has been used in a previous study (Kim et al. 2021) to examine the psychological features of individuals. Similar to their approach, we apply top 500 uni- and bi-grams as open vocabulary features in our dataset.

**Machine Learning Models** Using the pre-pandemic and during-pandemic datasets, we build two sets of binary classifiers for detecting tweets with minority stress. We refer to the models trained on the pre-pandemic dataset as **pre-pandemic models** and the ones trained on the during-pandemic dataset as **during-pandemic models**.

Similar to (Saha et al. 2019), we implement and evaluate six models for each training dataset, including a dummy classifier (baseline), naïve Bayes, Logistic Regression, Decision Tree, Support Vector Machine (SVM), and Multilayer Perceptron (MLP) algorithm. For hyperparameter tuning, all classifiers go through  $k$ -fold cross-validation ( $k = 10$ ).

**Feature Importance** To understand which features contribute the most to classification, we fit the best-performing model to the two training datasets (pre-pandemic and during-pandemic) separately. It turns out that logistic regression models show the best performance. We, therefore, retrieve coefficient scores for each input variable and compare their importance. In the Results section, we focus on the interpretable feature sets, namely LIWC, and n-grams. These two feature sets turn out to be the most important feature sets across all feature sets to predict minority stress. We rank all features with their weights. Features with positive weights contribute to identifying the minority stress tweets, while features with negative weights contribute to the opposite decision of “no minority stress”.

### Matching For Causal Inference

**Matching Covariates** Our goal is to compare the linguistic difference in tweets between members of the LGBTQ

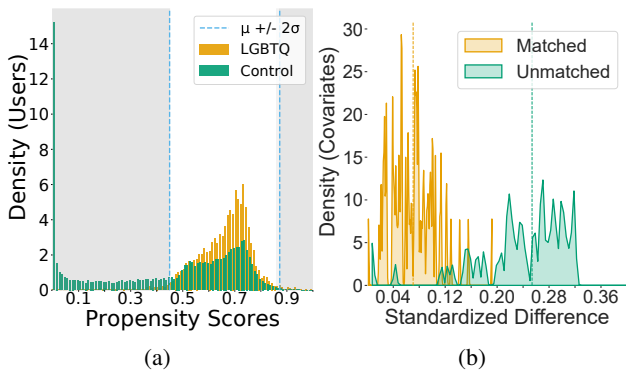


Figure 2: (a) Propensity score distribution (shaded region are the dropped strata), (b) Quality of matching.

population and the general population during the COVID-19 pandemic. To minimize the possible occurrence of selection bias, matching is an efficient strategy in case-control studies to estimate causal effects. In our case, we build several covariates from the LGBTQ and control datasets using data before December 1st, 2019, to control for similar pre-pandemic behavior on social media. The first set of covariates comprises Twitter users’ social media features (the number of tweets, likes, followers, followees, account creation time, and posting frequency). The second set contains the distribution of word usage in Twitter timelines. We extract the top 500 unigrams as the second covariates set. The third set includes the psycholinguistic features in the timeline data by measuring the word distribution in the LIWC lexicon (Wei et al. 2021). As previous research has indicated that mental health is correlated with text readability (Tausczik and Pennebaker 2010), we add several metrics of readability into the last set.

**Propensity Score Analysis** To ensure the similarities between our LGBTQ dataset and the Control dataset, matching is used to pair LGBTQ users and control users, whose covariates are similar to each other. We implement a logistic regression classifier to predict the likelihood of a user belonging to the LGBTQ group or control group based on their covariates. We refer to the likelihood as the propensity score.

We divide the propensity score distribution into 100 strata with equal width. The users with similar propensity scores are grouped into the same stratum (Kiciman, Counts, and Gasser 2018). This helps us evaluate the effects of the COVID-19 pandemic within each stratum, where the control group users are matched to the LGBTQ users based on the pre-pandemic behavioral traits. We remove the users with propensity scores falling within two standard deviations from the mean. We drop the strata failing to satisfy the minimum sample size within each stratum based on previous causal inference research (De Choudhury et al. 2016). By ensuring that there are at least 50 users per group in each stratum, this approach results in 33 strata, containing 4,107 LGBTQ and 7,129 Control users (Figure 2a).

**Quality Assessment of Covariate Matching** To assess whether users in the LGBTQ and control groups are statistically comparable, we measure the balance of the covariates. Comparisons between the two groups in each of the 33 valid strata are made by calculating the standardized mean difference (SMD). SMD is the difference in the mean covariate values between the LGBTQ and control groups divided by the pooled standard deviation of the two groups. A previous study (Kiciman, Counts, and Gasser 2018) has indicated that two groups are balanced if the absolute SMD of all covariates is less than 0.2. For the unmatched dataset, the maximum absolute SMD is 0.8, and the mean is 0.25, while the maximum absolute SMD is 0.19 and the mean is 0.07 in the matched dataset (Figure 2b). This satisfies the suggested threshold of less than 0.2.

**Estimating the Average Treatment Effects** To estimate the effect of the COVID-19 pandemic, we compute the change in word usage in the LIWC variables. The treatment effect<sup>2</sup>(TE) of an outcome variable  $S$  is calculated in stratum  $k$  by the following equation:

$$TE_k = \frac{\sum_i^{N_k} (S_{L_i}^{\text{during}} - S_{L_i}^{\text{before}})}{N_k} - \frac{\sum_j^{M_k} (S_{C_j}^{\text{during}} - S_{C_j}^{\text{before}})}{M_k}, \quad (1)$$

where  $N_k$  is the number of LGBTQ users in stratum  $k$ , and  $M_k$  is the number of matched control users in the same stratum.  $S_{L_i}^{\text{before}}$  and  $S_{C_j}^{\text{before}}$  are the word usage in the LIWC variables measured for LGBTQ user  $L_i$  and control user  $C_j$  before the pandemic.  $S_{L_i}^{\text{during}}$  and  $S_{C_j}^{\text{during}}$  are the word usage in the LIWC variables measured for LGBTQ user  $L_i$  and control user  $C_j$  during the pandemic. After calculating the average TE across all strata, we obtain the mean TE for the pandemic per LIWC variable. The outcome is interpreted as a higher increase (greater than 0) or higher decrease (less than 0) of observable attributes during the pandemic of LGBTQ users compared to the control group with similar pre-pandemic attributes.

## Results

In this section, we examine if different feature sets affect the preference of the models (RQ1). We investigate the linguistic difference before and during the pandemic in the context of minority stress (RQ2). We compare the linguistic differences in tweets during the COVID-19 pandemic between the general and the LGBTQ population (RQ3).

### RQ1: Machine Learning Classifier Performance

We implement and evaluate multiple classifiers in the pre-pandemic training dataset and during-pandemic training dataset, including a dummy classifier (baseline), logistic regression, random forest, support vector machine, naïve

<sup>2</sup>treatment effect: The phrase “treatment effect” is used in causal inference to describe the causal impact of a specific intervention or exposure that is being investigated. It’s important to note that we do not intend to imply that individuals in the studied group require or have received any form of treatment.




Figure 3: (a) Confusion matrices of the best-performing pre-pandemic model, (b) Confusion matrices of the best-performing during-pandemic model.




Figure 4: (a) ROC curves of the best-performing pre-pandemic-model (b) ROC curves of the best-performing during-pandemic model.

bayes, and multilayer perceptron algorithm. Table 2 summarizes the performance of our models and includes the average precision, recall, F1, accuracy, and AUC. We find that logistic regression models show strong and stable performance to indicate minority stress. The best pre-pandemic model shows 75.3% accuracy with an AUC of 0.838, while the best during-pandemic model shows 76.5% accuracy with AUC of 0.845, which are used for further analysis.

We investigate which feature sets in practice contribute the most to the preference of our classifier. We evaluate our model under five different feature sets: a full feature set as well as each of the four sets excluding one of the feature sets. Table 3 summarizes the performance of each feature set in logistic regression models trained on pre-pandemic and during-pandemic datasets. We find that all pre-pandemic and during-pandemic models perform considerably better than the baseline. This reveals that our feature sets do adequately capture the minority stress in tweets, and these different features in the training set do boost the performance of machine learning classification. In addition, we generate the confusion matrix (Figure 3) and the receiver operating characteristic (ROC) curves (Figure 4) of the best-performing models. We see that the during-pandemic model shows a higher false positive score than a false negative score, but the pre-pandemic model has the opposite.

Model	Prec	Rec	F1	Accr	AUC
<b>Pre-pandemic</b>					
Baseline	0.269	0.500	0.350	0.539	0.500
Decision Tree	0.654	0.653	0.652	0.655	0.653
Naïve Bayes	0.697	0.675	0.673	0.687	0.749
MLP	0.706	0.706	0.705	0.707	0.802
SVM	0.737	0.736	0.736	0.738	0.812
Logistic Regression	<b>0.753</b>	<b>0.749</b>	<b>0.749</b>	<b>0.753</b>	<b>0.838</b>
<b>During-pandemic</b>					
Baseline	0.265	0.500	0.346	0.530	0.500
Decision Tree	0.651	0.651	0.651	0.654	0.651
Naïve Bayes	0.706	0.701	0.691	0.693	0.764
MLP	0.713	0.711	0.710	0.713	0.788
SVM	0.752	0.751	0.750	0.752	0.835
Logistic Regression	<b>0.764</b>	<b>0.766</b>	<b>0.764</b>	<b>0.765</b>	<b>0.845</b>

Table 2: Average metrics of the pre-pandemic-model and during-pandemic-model in k-fold (k=10) cross-validation.

Features	Prec	Rec	F1	Accr	AUC
<b>Pre-pandemic</b>					
Full	<b>0.753</b>	<b>0.749</b>	<b>0.749</b>	<b>0.753</b>	<b>0.838</b>
-Embeddings	0.746	0.742	0.742	0.746	0.832
-Sentiment	0.735	0.732	0.732	0.735	0.816
-LIWC	0.737	0.735	0.735	0.738	0.824
-n-grams	0.738	0.736	0.736	0.739	0.824
<b>During-pandemic</b>					
Full	<b>0.764</b>	<b>0.766</b>	<b>0.764</b>	<b>0.765</b>	<b>0.845</b>
-Embeddings	0.757	0.759	0.757	0.759	0.842
-Sentiment	0.756	0.758	0.756	0.758	0.834
-LIWC	0.750	0.751	0.749	0.750	0.833
-n-grams	0.726	0.727	0.726	0.728	0.825

Table 3: Logistic regression model performance of the pre-pandemic-model and during-pandemic-model across different feature sets.

We analyze the misclassified posts to isolate factors contributing to misclassifications. We sample several posts from test sets with the predicted labels by pre-pandemic and during-pandemic models. We find several false positive posts containing the condemnation of discriminatory behavior. For example, “*The Supreme Court ruled that companies do not have a right to discriminate against LGBTQ people in the workplace. This historic decision holds firm that LGBTQ people are, and should be, protected from discrimination under federal law.*” This post is annotated as a non-minority stress post, suggesting the advancement of LGBTQ rights. Since there are implicit references to past discrimination behaviors in employment for the LGBTQ population, the models fail to distinguish it from minority stress posts. We notice some false negatives when posts mention minority stress without further details, such as: “*It’s sad that we as LGBTQ Americans still cannot help save our own neighbors’ lives! I feel as though it is plain and simple discrimination!*”.

	Feature	Pre	$\Delta$
Minority stress	cause	3	593 ↓
	certain	6	220 ↓
	family	7	398 ↓
Non-minority stress	number	MR	498 ↑
	discrepancy	MR-13	437 ↑
	perception	MR-15	521 ↑
	Feature	During	$\Delta$
Minority stress	anger	2	40 ↓
	inhibition	24	22 ↓
	see	45	509 ↓
Non-minority stress	leisure	MR-1	96 ↑
	cause	MR-8	593 ↑
	posEmo	MR-11	497 ↑

Table 4: Top three LIWC features for the best-performing pre-pandemic and during-pandemic model with its ranking among 607 features. MR indicates the Maximum Rank (= 607). The  $\Delta$  indicates the change of rank from one model to another with its change direction.

## RQ2: Linguistic Differences of Minority Stress

To answer RQ2, we examine the most important features of the best minority stress classifiers. We explore which features in tweets contribute to identifying the difference between minority stress in posts before and during COVID-19. These features indicate that the models consider them to be associated with the labels of “minority stress” assigned by human annotators and provide further insights into inferring minority stress from tweets.

Tables 4 and 5 report the top three LIWC features and top five n-gram features, indicating minority stress tweets and non-minority stress tweets with their ranking among all features. MR indicates the maximum rank (607) and  $\Delta$  indicates the change of rank with its change direction. For example, feature *number* with the rank MR in the pre-pandemic model indicates that its rank is 607, which is the most important feature to classify non-minority stress tweets.

In table 4, we observe that cognitive attributes, including *causation*, *certainly* and *inhibition*, contribute significantly to the classification. *Cognitive attribute* in LIWC includes words such as “cause”, “know”, and “ought”. Cognition ability is linked to re-structure and understanding the emotional experience. Previous studies suggest that increasing the use of cognitive words is linked to health improvement with repeated writing and cognition ability may be essential to adaptation to traumatic or stressful experiences (Klein and Boals 2001). This is also consistent with our annotated dataset. For example, “*It took me more than a decade to embrace my homosexuality and that it can never be changed. Even with that knowledge, I am still depressed because sexuality is completely useless to me as an introvert depressed social anxiety sufferer.*” Cognitive words, including *discrepancy* and *cause*, also contribute to identify the non-minority stress tweets. This might be because some traumatic expe-

	Feature	Pre	$\Delta$
Minority stress	take	1	548 ↓
	get	2	87 ↓
	lgbtq rights	4	504 ↓
	anti	5	177 ↓
	discrimination	8	3 ↓
Non-minority stress	support	MR-1	456 ↑
	right	MR-3	38 ↑
	rt <sup>3</sup>	MR-4	0 –
	rights in	MR-5	497 ↑
	important	MR-6	192 ↑
	Feature	During	$\Delta$
Minority stress	being	1	62 ↓
	anti trans	3	494 ↓
	discrimination	4	468 ↓
	homophobia	5	216 ↓
	2020	6	236 ↓
	allow	6	236 ↓
Non-minority stress	bigotry	MR	38 ↑
	discrimination	MR-2	356 ↑
	music	MR-3	235 ↑
	for	MR-4	0 –
	rt	MR-5	199 ↑
	all the	MR-5	199 ↑

Table 5: Top five n-gram features for the best-performing pre-pandemic and during-pandemic model with its ranking among 607 features. MR indicates the Maximum Rank (= 607). The  $\Delta$  indicates the change of rank from one model to another with its change direction.

riences are caused by their stigmatized identities but other heartbreaking events. Additionally, we find that the during-pandemic model heavily relies on *anger*, which contains words describing violent actions or angry emotions (e.g., “hate”, “kill”, and “annoyed”). For instance, “*I didn’t want others to assume i was gay then try to harm me or figure out where I live for future attack*”.

We find many discriminating n-grams that are significant features, such as *anti/anti trans*, *discrimination*, *discrimination homophobia*. Additionally, we also observe that the n-gram *lgbtq rights* frequently occurs when complaining about Trump’s statements about LGBTQ. For example, “*Since Donald Trump took office, his government has attacked LGBTQ advancement. Turnout today will determine the future of LGBTQ rights. Please VOTE*”. One interesting observation is that the pre-pandemic classifier uses *anti trans* and *discrimination homophobia* to predict non-minority stress tweets. This can be explained by non-minority tweets in our dataset, which use these words as hashtags to advocate advancing LGBTQ rights.

## RQ3: Linguistic Analysis of Matched Result

To investigate if the COVID-19 pandemic has a statistically significant effect on LGBTQ users, we estimate the effect

<sup>3</sup>rt refers to a re-posted tweet.



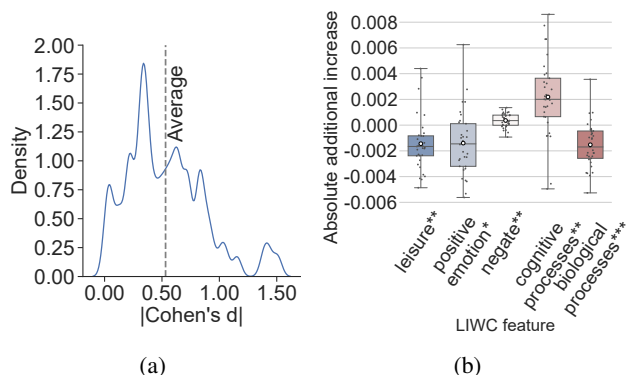


Figure 5: (a) Distribution of effect size magnitude in the outcome change between LGBTQ and Control users; (b) Distribution of treatment effect across all the LGBTQ users. Statistical significance is conducted using Welch’s t-test and  $p$ -values are reported after Bonferroni correction (\* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ ).

size (Cohen’s  $d$ ) in LIWC word usage changes between the LGBTQ and the control group per stratum. As shown in Figure 5a, we find that the average Cohen’s  $d$  in the outcome changes across all valid strata is 0.53, which illustrates medium to large differences in LIWC word usage. A Welch’s t-test reveals the statistical significance of these differences ( $t \in [-5.64, 3.38]$ ,  $p < 0.05$ ), which confirms the significant changes in these outcomes.

We analyze the effect of the pandemic. Figure 5b shows the distribution of the treatment effect across all strata. The average treatment effects greater than 0 indicate a greater increase in word usage following the impact of the pandemic. We observe that the pandemic causes greater increases for the LGBTQ group on the outcomes of *negation* ( $t = 3.38$ ,  $p < 0.01$ ) and *cognitive processes* ( $t = 3.38$ ,  $p < 0.01$ ). *Negation* is a category in LIWC containing “no”, “not”, “never”, and so on. One explanation is that because of the disproportionate impact of the pandemic on LGBTQ groups, LGBTQ users use more cognitive words to express their traumatic experiences. We observe the average treatment effect of *positive emotion* ( $t = 1.21$ ,  $p < 0.05$ ) is less than 0, which is translated to worsened observable attribute in the usage of positive words, such as “live”, “nice”, and “sweet”. We find the worse observable attributes of ( $t = 1.21$ ,  $p < 0.05$ ), *leisure* ( $t = -4.11$ ,  $p < 0.01$ ) and *biological processes* ( $t = -5.64$ ,  $p < 0.001$ ). The *leisure* in LIWC contains words related to lifestyles, such as “cook”, “chat”, and “movie”, and the *biological processes* contains “eat”, “clinic”, “horny”, and so on. Interestingly, there is no significant difference in the usage of *negative emotion*, which suggests that the pandemic has no disproportionate impact for the LGBTQ group on that.

## Conclusion and Discussion

**Summary:** In this work, we develop theory-driven computational approaches to quantify minority stress from social media data. Our classifiers demonstrate strong capabilities

to detect minority stress from tweets. Following this, we investigate the differences in language-related attributes of minority stress before and during the COVID-19 pandemic. We observe the increased prominence of anger words as predictive attributes of minority stress during the pandemic. To explain the differential impact on their emotions, we also contrast the linguistic difference in tweets between the LGBTQ population and the general population during the pandemic by adopting a propensity score-based matching approach. Our findings reveal that during the pandemic, the LGBTQ population is faced with a disproportionate impact on their cognitive and emotional states. This is demonstrated in their worsened observable outcome in the usage of words that convey positive emotions and a greater increase in the usage of cognitive words.

**Implications:** Past research work on LGBTQ communities has acknowledged the difficulty of accessing health-related data (McDermott, Roen, and Piela 2013). Most existing LGBTQ health-related research is limited to diagnostic reports and psychological questionnaires. We believe that the methodologies and results from our study can provide policy-makers with important health and well-being information to support evidence-based decisions for a better understanding of the development of mental health and minority stress among the LGBTQ population. The need for such targeted approaches is more pertinent than ever in times of a global pandemic, as our findings also illustrate the disproportionate impact on the emotional and cognitive states of the LGBTQ population compared to the general population. We find that, on average, the LGBTQ population has a greater increase in the usage of cognitive words and uses positive words less often during the COVID-19 pandemic as compared to a group of the general population with comparable pre-pandemic behavior traits. Prior studies have shown that the mental health of individuals within the LGBTQ community is adversely affected due to the discrimination and bias they face. These pre-existing mental health conditions could make them even more vulnerable to the adverse outcomes of the pandemic, leading to threatening situations.

We demonstrate the opportunities that machine learning techniques and social media data can offer to public health professionals. Our methods can be used to build tools allowing social media platforms to measure the relative increase of minority stress among LGBTQ individuals during public crises. This aligns with the American Counseling Association’s Code of Ethics (Association et al. 2014), which suggests that counselors need to gain knowledge, sensitivity and skill to work with diverse backgrounds. Our findings can provide counselors with sensitivity and awareness about the discrimination and stigmatization that LGBTQ individuals face to build client-centered practices.

We believe that our research can complement the tools to develop the minority stress theory further. Broadly, several studies have discussed increased fear, anger, and trauma during the COVID-19 pandemic. Aligned with these broader findings, our work demonstrates an increased association of anger with minority stress in LGBTQ groups. Our results can be used to build tools that reveal the feelings and emo-

tions behind the language of LGBTQ people. Understanding the social media communication from the LGBTQ population can offer a new lens onto minority stress at the population level, which can be utilized to verify the existing research on minority stress or identify the under-observed variables of minority stressors.

**Limitations:** Our work is grounded in Meyer's theory of minority stress. Although it is widely used for various empirical investigations on the health of the minority population, this framework has its own limitations. As acknowledged by Meyer (2003), there is a necessity to expand this framework and distinguish between different stigmatized statuses. Future work could dig deeper into this issue by adding different psychological scales for different groups.

Another limitation of our study is selection biases. We search Twitter users using hashtags, and collect self-disclosed identity information in their bios. This is likely to be influenced by self-selection bias as we can only collect data from those who expressed themselves on social media. This is especially true when considering that LGBTQ people have long been exposed to stigmatization and prejudice. Similarly, this research exclusively concentrates on Twitter, and this could lead to biased or insufficient viewpoints. Consequently, our findings may not be generalized to other online communities.

A further note is that we only use features inside the tweets to measure minority stress. Other features, such as frequency of posting and author metadata, could be incorporated to analyze the minority stress comprehensively. Another research direction is to explore posts on minority stress across different social media platforms to build accurate, comprehensive, and up-to-date feature sets and gain a better and more thorough understanding of this topic.

Lastly, our analysis does not differentiate results across subcategories of minority stress due to the difficulty of extracting the context from a single tweet. Future works can analyze minority stress from the conversation level on Twitter to better understand minority stress.

## Acknowledgements

The calculations presented above were performed using computer resources within the Aalto University School of Science "Science-IT" project.

## Supplementary Material

**LGBTQ Related Hashtag List:** *lgbtq, gay, pride, lgbtqia, pridemonth, love, queer, trans, pridemonth2021, pride2021, mmromance, loveislove, gaymer, lesbian, gayromance, nonbinary, transgender, hiv, gaydar, equality, bisexual, gayrights, paranormal, comingout, pansexual, transrightsarehumanrights, transisbeautiful, homophobia, lgbtpride, rainbow, lgbtpride, kpopstans, lgbt, clubsilver, gaytravel, lgbttravel, instagay, gayboy, gayart, dragqueen, gaytwitter, gayguy, bara, yaoi, bear, sexy, gayman, queen, gayuk, furry, gaylove, bi, queerart, grindr, queenofpop, gaybear, gayfiction, gaymen, scruff, queerfilmrec.*

**Non-LGBTQ Related Hashtag List:** *sunday, birthday, yelp, covid19, florida, model, smile, losangeles, fitness,*

*coronavirus, dogsofinstagram, tiktok, family, makeup, faith, catsofinstagram, blacklivesmatter, newprofilepic, style, foodporn, nowplaying, chicago, summer, grateful, truth, travel, fun, life, photography, nyc, foodie, picoftheday, latergram, nofilter, fashion, artist, halloween, friends, wynonnaearp, music, art, shopmycloset, beach, neasummit, disney, vote, podcast, happy, texas, miami, food, sundayfunday, beautiful, nature, blessed, home, inspiration, usa, flowers, happybirthday, instagram, california, selfie, soundcloud, vacation, vegan, poshmark, twitch, sunset, beauty, photooftheday, christmas, goodmorning, repost, realestate, motivation, instagood, happiness, newyork.*

## References

- Andalibi, N.; Haimson, O. L.; De Choudhury, M.; and Forte, A. 2016. Understanding social media disclosures of sexual abuse through the lenses of support seeking and anonymity. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, 3906–3918.
- Association, A. C.; et al. 2014. American Counseling Association code of ethics. *Alexandria, VA: Author.*
- Borgogna, N. C.; McDermott, R. C.; Aita, S. L.; and Kridel, M. M. 2019. Anxiety and depression across gender and sexual minorities: Implications for transgender, gender nonconforming, pansexual, demisexual, asexual, queer, and questioning individuals. *Psychology of sexual orientation and gender diversity*, 6(1): 54.
- Cannon, Y.; Speedlin, S.; Avera, J.; Robertson, D.; Ingram, M.; and Prado, A. 2017. Transition, connection, disconnection, and social media: Examining the digital lived experiences of transgender individuals. *Journal of LGBT Issues in Counseling*, 11(2): 68–87.
- Coppersmith, G.; Leary, R.; Crutchley, P.; and Fine, A. 2018. Natural language processing of social media as screening for suicide risk. *Biomedical informatics insights*, 10: 1178222618792860.
- Craig, S. L.; and McInroy, L. 2014. You can form a part of yourself online: The influence of new media on identity development and coming out for LGBTQ youth. *Journal of Gay & Lesbian Mental Health*, 18(1): 95–109.
- De Choudhury, M.; and De, S. 2014. Mental health discourse on reddit: Self-disclosure, social support, and anonymity. In *Eighth international AAAI conference on weblogs and social media.*
- De Choudhury, M.; Kiciman, E.; Dredze, M.; Coppersmith, G.; and Kumar, M. 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*, 2098–2110.
- Frost, D. M.; Lehavot, K.; and Meyer, I. H. 2015. Minority stress and physical health among sexual minority individuals. *Journal of behavioral medicine*, 38(1): 1–8.
- Gallagher, M. W.; Zvolensky, M. J.; Long, L. J.; Rogers, A. H.; and Garey, L. 2020. The impact of Covid-19 experiences and associated stress on anxiety, depression, and functional impairment in American adults. *Cognitive Therapy and Research*, 44(6): 1043–1051.

- Garg, S.; Taylor, J.; El Sherief, M.; Kasson, E.; Aledavood, T.; Riordan, R.; Kaiser, N.; Cavazos-Rehg, P.; and De Choudhury, M. 2021. Detecting risk level in individuals misusing fentanyl utilizing posts from an online community on Reddit. *Internet Interventions*, 26: 100467.
- Goldbach, J. T.; Mereish, E. H.; and Burgess, C. 2017. Sexual orientation disparities in the use of emerging drugs. *Substance Use & Misuse*, 52(2): 265–271.
- Haimson, O. L. 2019. Mapping gender transition sentiment patterns via social media data: toward decreasing transgender mental health disparities. *Journal of the American Medical Informatics Association*, 26(8-9): 749–758.
- Hswen, Y.; Sewalk, K. C.; Alsentzer, E.; Tuli, G.; Brownstein, J. S.; and Hawkins, J. B. 2018. Investigating inequities in hospital care among lesbian, gay, bisexual, and transgender (LGBT) individuals using social media. *Social science & medicine*, 215: 92–97.
- Imbens, G. W.; and Rubin, D. B. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- Kiciman, E.; Counts, S.; and Gasser, M. 2018. Using longitudinal social media analysis to understand the effects of early college alcohol use. In *Twelfth International AAAI Conference on Web and Social Media*.
- Kim, S.; Razi, A.; Stringhini, G.; Wisniewski, P.; and De Choudhury, M. 2021. You Don't Know How I Feel: Insider-Outsider Perspective Gaps in Cyberbullying Risk Detection. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, 290–302.
- Klein, K.; and Boals, A. 2001. Expressive writing can increase working memory capacity. *Journal of experimental psychology: General*, 130(3): 520.
- Kneale, D.; and Becares, L. 2020. The mental health and experiences of discrimination of LGBTQ+ people during the COVID-19 pandemic: initial findings from the Queerantime study. *MedRxiv*.
- Kulick, A.; Wernick, L. J.; Woodford, M. R.; and Renn, K. 2017. Heterosexism, depression, and campus engagement among LGBTQ college students: Intersectional differences and opportunities for healing. *Journal of homosexuality*, 64(8): 1125–1141.
- Landis, J. R.; and Koch, G. G. 1977. The measurement of observer agreement for categorical data. *biometrics*, 159–174.
- Lehavot, K.; and Simoni, J. M. 2011. The impact of minority stress on mental health and substance use among sexual minority women. *Journal of consulting and clinical psychology*, 79(2): 159.
- Manning, C. D.; Surdeanu, M.; Bauer, J.; Finkel, J. R.; Bethard, S.; and McClosky, D. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, 55–60.
- McDermott, E.; Roen, K.; and Piela, A. 2013. Hard-to-reach youth online: Methodological advances in self-harm research. *Sexuality Research and Social Policy*, 10(2): 125–134.
- Meyer, I. H. 1995. Minority stress and mental health in gay men. *Journal of health and social behavior*, 38–56.
- Meyer, I. H. 2003. Prejudice, social stress, and mental health in lesbian, gay, and bisexual populations: conceptual issues and research evidence. *Psychological bulletin*, 129(5): 674.
- Pavalanathan, U.; and Eisenstein, J. 2016. More emojis, less:) The competition for paralinguistic function in microblog writing. *First Monday*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Pierce, M.; Hope, H.; Ford, T.; Hatch, S.; Hotopf, M.; John, A.; Kontopantelis, E.; Webb, R.; Wessely, S.; McManus, S.; et al. 2020. Mental health before and during the COVID-19 pandemic: a longitudinal probability sample survey of the UK population. *The Lancet Psychiatry*, 7(10): 883–892.
- Pietrabissa, G.; and Simpson, S. G. 2020. Psychological consequences of social isolation during COVID-19 outbreak. *Frontiers in psychology*, 2201.
- Saha, K.; Kim, S. C.; Reddy, M. D.; Carter, A. J.; Sharma, E.; Haimson, O. L.; and De Choudhury, M. 2019. The language of LGBTQ+ minority stress experiences on social media. *Proceedings of the ACM on human-computer interaction*, 3(CSCW): 1–22.
- Saha, K.; Torous, J.; Caine, E. D.; and De Choudhury, M. 2020. Psychosocial Effects of the COVID-19 Pandemic: Large-scale Quasi-Experimental Study on Social Media. *Journal of medical internet research*, 22(11): e22600.
- Salerno, J. P.; Williams, N. D.; and Gattamorta, K. A. 2020. LGBTQ populations: Psychologically vulnerable communities in the COVID-19 pandemic. *Psychological Trauma: Theory, Research, Practice, and Policy*, 12(S1): S239.
- Saud, M.; Mashud, M.; and Ida, R. 2020. Usage of social media during the pandemic: Seeking support and awareness about COVID-19 through social media platforms. *Journal of Public Affairs*, 20(4): e2417.
- Steinke, J.; Root-Bowman, M.; Estabrook, S.; Levine, D. S.; and Kantor, L. M. 2017. Meeting the needs of sexual and gender minority youth: formative research on potential digital health interventions. *Journal of Adolescent Health*, 60(5): 541–548.
- Tausczik, Y. R.; and Pennebaker, J. W. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of language and social psychology*, 29(1): 24–54.
- Tsao, S.-F.; Chen, H.; Tisseverasinghe, T.; Yang, Y.; Li, L.; and Butt, Z. A. 2021. What social media told us in the time of COVID-19: a scoping review. *The Lancet Digital Health*.
- Verma, G.; Bhardwaj, A.; Aledavood, T.; De Choudhury, M.; and Kumar, S. 2022. Examining the impact of sharing COVID-19 misinformation online on mental health. *Scientific Reports*, 12(1): 8045.
- Wang, Y.; Pan, B.; Liu, Y.; Wilson, A.; Ou, J.; and Chen, R. 2020. Health care and mental health challenges for transgender individuals during the COVID-19 pandemic. *The Lancet Diabetes & Endocrinology*, 8(7): 564–565.

Wei, J.; Finn, K.; Templeton, E.; Wheatley, T.; and Vosoughi, S. 2021. Linguistic Complexity Loss in Text-Based Therapy. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4450–4459.

Xiang, X.; Lu, X.; Halavanau, A.; Xue, J.; Sun, Y.; Lai, P. H. L.; and Wu, Z. 2021. Modern senicide in the face of a pandemic: An examination of public discourse and sentiment about older adults and COVID-19 using machine learning. *The Journals of Gerontology: Series B*, 76(4): e190–e200.

Zeeman, L.; Sherriff, N.; Browne, K.; McGlynn, N.; Mirandola, M.; Gios, L.; Davis, R.; Sanchez-Lambert, J.; Aujean, S.; Pinto, N.; et al. 2019. A review of lesbian, gay, bisexual, trans and intersex (LGBTI) health and healthcare inequalities. *European journal of public health*, 29(5): 974–980.

Zhou, J.; Zogan, H.; Yang, S.; Jameel, S.; Xu, G.; and Chen, F. 2021. Detecting community depression dynamics due to covid-19 pandemic in australia. *IEEE Transactions on Computational Social Systems*.