

COVID-19: Detecting Depression Signals during Stay-At-Home Period

Jean Marie Tshimula,¹ Belkacem Chikhaoui,^{1,2} Shengrui Wang¹

¹Department of Computer Science, Université de Sherbrooke, QC J1K 2R1, Canada

²LICEF Research Center, Université TÉLUQ, QC H2S 3L5, Canada

{kabj2801, shengrui.wang}@usherbrooke.ca

belkacem.chikhaoui@teluq.ca

Abstract

The new coronavirus outbreak has been officially declared a global pandemic by the World Health Organization. To grapple with the rapid spread of this ongoing pandemic, most countries have banned indoor and outdoor gatherings and ordered their residents to stay home. Given the developing situation with coronavirus, mental health is an important challenge in our society today. In this paper, we discuss the investigation of social media postings to detect signals relevant to depression. To this end, we utilize topic modeling features and a collection of psycholinguistic and mental-well-being attributes to develop statistical models to characterize and facilitate representation of the more subtle aspects of depression. Furthermore, we predict whether signals relevant to depression are likely to grow significantly as time moves forward. Our best classifier yields F-1 scores as high as 0.8 and surpasses the utilized baseline by a considerable margin, 0.173. In closing, we propose several future research avenues.

1 Introduction

The ongoing coronavirus outbreak has been officially defined a global pandemic by the World Health Organization (WHO) on March 11, 2020. Coronavirus disease 2019 (COVID-19) is an infectious disease caused by a newly discovered coronavirus (WHO, 2020). COVID-19 causes a respiratory illness characterized by symptoms such as cough, fever, difficulty breathing, and pneumonia in both lungs. These symptoms may take up to 14 days to appear after exposure to COVID-19. COVID-19 spares no one and infects people of all ages. Older people and those with

pre-existing medical conditions like cardiovascular disease, diabetes, chronic respiratory disease, and cancer appear to be more vulnerable to becoming severely ill with COVID-19 (WHO, 2020; Canada, 2020).

WHO has reported a drastic increase in confirmed cases and deaths all over the world. To mitigate the rapid spread of COVID-19, many countries have forbidden indoor and outdoor gatherings in excess of particular numbers of people; asked non-essential services, nonprofit entities, and retail businesses to close; issued stay-at-home orders for their residents; and advised them to practice social distancing and avoid all non-essential travel abroad. We are living through a pivotal moment in history. The onslaught of the pandemic has severely challenged our economic systems (McKibbin and Fernando, 2020) and caused substantial changes to people's daily routine. The current pandemic can affect people both physically and psychologically (Wang et al., 2020). For example, in China, 96.2% of clinically stable COVID-19 patients in the early recovery phase reported significant post-traumatic stress disorder (PTSD) symptoms (Bo et al., 2020). Psychological distress is increasing worldwide and may have long-lasting consequences and repercussions on mental health (Brooks et al., 2020; Gunnell et al., 2020; Li et al., 2020; Meng et al., 2020).

Given the developing situation with the pandemic, social media allows people to inform themselves and get updates from official sources. People may naturally panic when seeing headlines announcing bad news and numbers of cases. This may affect ways in which individuals express themselves and share opinions, thoughts, and personal experiences with others. The emotion and language in social media postings may potentially indicate feelings such as loneliness (Guntuku et al., 2019), anxiety, anger and stress, among others (De Choudhury, 2013). For instance, a per-

son may express emotional reactions that can be unpleasant, disturbing, and overwhelming. Emotional problems like anxiety and depression manifest themselves as feelings of inner emotional distress. Mental health issues can comprise a wide range of disorders that affect mood, thinking, and behavior. Some examples of mental illness include PTSD, depression, anxiety disorders, addictive behaviors, etc. In this paper, our primary interest is in depression. Depression is a serious condition that can cause a persistent feeling of sadness and loss of interest and can affect a person's daily life (Kanter et al., 2013). Survey research conducted by Mental Health Research Canada found that feelings of depression are rising constantly (MHRC, 2020). Before the pandemic, 7% of Canadians reported high levels of depression. This rate has risen to 16% during the stay-at-home period and 22% predict high levels of depression if social isolation continues for two more months.

Recognizing early signs of depression is of critical importance and can aid mental health services in assessing the impact of the pandemic on the population and implementing healthier coping strategies to build personal resilience. In addition, appropriate services can be provided for those in need. In this paper, we leverage social media postings to detect signals relevant to depression due to COVID-19. To this end, we build a corpus of postings shared on Twitter during the stay-at-home period. We make use of a topic modeling approach to generate topics addressed by individuals and evaluate language features from topic words to determine whether they indicate signals for depression. It should be noted that we retain solely depression-indicative topics and collect individuals who engage with these topics to investigate their posting histories since the onset of the stay-at-home order. Specifically, this work makes the following contributions:

- We demonstrate the effectiveness of our data collection and data pre-processing strategy to gather social media postings containing signals relevant to depression.
- We capture evidence from a corpus of postings and potential individuals who manifest signals for depression and consider them as an experimental group. We measure the similarity between different topics addressed by individuals in the experimental group to dis-

cover their overlapping behavioral characteristics and understand their linguistic idiosyncrasies.

- We develop models to predict whether signals relevant to depression are likely to grow significantly as time moves forward.

2 Related Work

The role of social media in mental health has been explored by De Choudhury (De Choudhury, 2013). The study suggested a guideline that emphasizes the use of social media postings to gauge what the pertinent mental literature would predict at the individual- and population-levels. This could allow the identification of depressed or otherwise at-risk individuals through the large-scale passive monitoring of social media (Guntuku et al., 2017). Recently, research has associated social media with several mental health conditions, including stress (Guntuku et al., 2019; Saha and De Choudhury, 2017; Thelwall, 2017), post-traumatic stress disorder (Coppersmith et al., 2014a; Coppersmith et al., 2014b; He et al., 2012) and depression (Guntuku et al., 2017; Cacheda et al., 2019; Coppersmith et al., 2015; De Choudhury et al., 2013; Jamil et al., 2017; Resnik et al., 2015a; Resnik et al., 2013; Resnik et al., 2015b; Sadeque et al., 2018; Schwartz et al., 2014; Shen et al., 2017; Tsugawa et al., 2015).

To quantify depression from texts, De Choudhury et al. proposed a social media depression index to identify levels of depression among individuals and predict social network behavior changes related to post-partum depression using several features, including structural properties of social networks (De Choudhury et al., 2013). While some studies rely exclusively on open-vocabulary analysis and lexicon-based techniques such as Linguistic Inquiry and Word Count (LIWC) (Pennebaker et al., 2015) to build a classifier, other studies couple LIWC with topic modeling features (Resnik et al., 2013; Stark et al., 2012; Tadesse et al., 2019; Zhai et al., 2012). For instance, Coppersmith et al. used LIWC to demonstrate characteristic differences in language use for mental disorders (Coppersmith et al., 2014a). Their approach utilizes uni-grams and 5-grams to indicate the presence of mental health conditions. Stark et al. (2012) combined LIWC and latent Dirichlet allocation (LDA)-based features in the classification of social relationships. Resnik et

al. (2013) explored the value-add of topic modeling in text analysis for depression and showed that topic models can take us beyond the LIWC categories to relevant themes related to depression and neuroticism as a strongly associated personality measure. Another work of Resnik et al. (2015b) investigated the use of supervised topic models in the analysis of linguistic signals for detecting depression. Tadesse et al. (2019) demonstrated that multiple feature combinations (LIWC+LDA+bi-gram) can yield competitive results. In this paper, we take a step forward by combining LDA with bi-gram, LIWC and other psycholinguistic dictionary-based features to identify depression-indicative topics, in order to facilitate the investigation of signals relevant to depression. The rationale behind the incorporation of additional features is to enrich the model to be able to capture depression-related terms and patterns that may escape the LIWC dictionary. We utilize correlation metrics to compare the performance of the proposed features with other alternative feature combinations.

3 Detection of depression signals

Dataset during the stay-at-home period. All data we obtained is public, posted between 12 March 2020 and 25 May 2020,¹ and made available from Twitter. Specifically, we extracted tweets bearing the words or hashtags: COVID, coronavirus, #StayAtHome, or #StayHome. For privacy and ethical reasons, we avoid displaying personally identifiable information, especially names and pseudonyms. Therefore, we randomly replaced such information to ensure the anonymity and privacy of the data.

To preprocess the data, we limited our set to Canadian users and removed tweets written in a language other than English or French. Additionally, we discarded redundant tweets, retweets without comments, tweets containing only the keyword (i.e., words or hashtags utilized for extraction), and multimedia such as image and video. We removed links in tweets, but kept emojis, since research has proven that emotions within a text can be expressed through the use of emojis (Hauthal et al., 2019). We used the Python *Googletrans*² implementation package to translate tweets from

¹This corresponds to the onset of lockdown and the date on which COVID-19 lockdown restrictions began slowly being relaxed across the country.

²<https://py-googletrans.readthedocs.io/en/latest/>

French to English. We removed tweets in which the word COVID or coronavirus occurs simultaneously with the term mental health or depression. We believe that people reacting emotionally may avoid combining the two words in a single tweet when it conveys a personal account. Consequently, we assume that these kinds of tweets are more likely to convey information or warnings about mental health. We eliminated stopwords but kept pronouns.³ Pronouns reveal information on people’s emotional state, thinking, and personality (Pennebaker et al., 2015). Chung and Pennebaker (2007) discovered that individuals susceptible to mental illness such as depression more frequently use first-person pronouns, suggesting higher self-attention focus.

To concentrate exclusively on data containing signals relevant to depression, we quantified different aspects of the language usage and patterns of individuals, using automated methods in order to extract features indicative of depression in tweets.

Dataset before the stay-at-home order. We replicated and applied the same logic as above to collect tweets posted before the stay-at-home order, that is, from 1 January 2020 to 11 March 2020. In total, we extracted 1,006,941 tweets and 161,327 distinct users, that is, users who had at least five tweets.

3.1 Feature Design

Bi-gram features. We extracted bi-grams from tweets by leveraging the vectors based on the term frequency-inverse document frequency (TF-IDF) approach (Ramos et al., 2003; Tadesse et al., 2019). We used TF-IDF as a statistical measure to evaluate how important a word is to each tweet in the corpus. We convert each tweet into its bag-of-word representation and calculate the TF-IDF value of each word utilizing the standard formula (Equation 1).

$$\text{TF-IDF} = (1 + \log n_{w,t}) \times \log \frac{T}{T_w} \quad (1)$$

where the TF-IDF value of word w in tweet t is the log normalization of the number of times the word occurs in the tweet ($n_{w,t}$) times the inverse log of the number of tweets T and T_w the number of tweets containing word w .

³I, you, she, he, we and they (see Table 2)

Table 1: Prediction quality for depression, for different feature sets and all combinations, as measured using the Pearson r . For LIWC features, we consider one feature per category and for LDA features, we take one feature per topic.

Feature set	r
LIWC	0.286
LIWC+LDA	0.342
LIWC+bi-gram	0.313
LIWC+bi-gram+LDA	0.371
LIWC+PLUS+bi-gram+LDA	0.506

LIWC features. The Linguistic Inquiry and Word Count (LIWC) dictionary is a widely used psychometrically validated system for psychology-related analysis of language and word classification (Pennebaker et al., 2015). LIWC includes word categories that have pre-labeled meanings. For each tweet, we calculated the number of observed words, using the LIWC dictionary and focusing on three LIWC categories: linguistic dimensions, psychological processes, and personal concerns. For the psychological processes and personal concerns categories, we utilized all of their subcategories, while for the linguistic dimensions category, we exclusively measured the proportion of first-person pronouns in the tweet.

PLUS features. We extracted depression-related features from the MRC psycholinguistic database (Wilson, 1988), the WHO glossary of psychiatric and mental health terms (WHO, 1994), and the NRC emotion lexicons (Mohammad and Turney, 2013). The NRC emotion lexicon is a list of English words and their associations with eight basic emotions (*anger*, *fear*, *anticipation*, *trust*, *surprise*, *sadness*, *joy*, and *disgust*) and two sentiments (*negative* and *positive*). MRC provides information about 26 different linguistic properties and includes more than 150,000 words with linguistic and psycholinguistic features of each word. For each tweet, we identified depression-related words using the WHO glossary and verified whether these words fall into the NRC emotion lexicons. Specifically, we discarded all the words that imply “joy” as the emotional state. Each MRC feature was computed by averaging the scores of all the depression-related words found in the database.

Table 2: Top fifteen words for the first five of the 38 validated depression-indicative topics.

Topic	Words
1	limit, alone, bad, I, bored, hard, when, time, wash, hand, tired, isolation, abuse, social, paper
2	feeling, myself, mask, mood, extremely, time, affect, out, crisis, mind, bad, finish, way, I, worse
3	friends, sleep, I, life, suffer, miss, shit, always, dull, long, end, back, family, hopeless, change
4	disgust, hell, freaking, I, enemy, worry, care, moment, invisible, difficult, feel, bad, health, home, sick
5	time, sad, home, close, depressed, hard, move, limited, boring, unhappy, stay, services, weird, feel, park

LDA features. We utilized LDA (Blei et al., 2003) to learn the topics addressed from the tweets. LDA is a probabilistic model that discovers latent topics in a text corpus and can be trained using collapsed Gibbs sampling. A topic is a distribution over a fixed vocabulary. As the parameters of LDA, we set α and β to 0.01. All extracted topics were used as features.

4 Experimental Setup and Results

Prediction of depression during the stay-at-home period. We generated 50 topics overall, of which we especially examined topics containing words related to mental health. To this end, we combined PLUS, bi-gram, and LIWC features to identify topics containing depression-related words. The depression-indicative topics were validated by clinical psychologists. Next, we took users who engaged with the 38 depression-indicative topics (see Table 2) and collected all tweets of these users from 12 March 2020 to 25 May 2020. We kept users who had at least five tweets and considered these users as an experimental group. In total, we were left with 87,236 distinct users and 857,294 tweets. We performed linear regression with elastic-net regularization to predict depression signals derived from previous features and evaluated the quality of prediction using the Pearson correlation (r). We stratified the dataset for 10-fold cross-validation to separate our training and testing sets. Table 1 shows that all of the feature sets combined (LIWC+PLUS+bi-gram+LDA) produce much stronger correlations ($r = 0.506$, $p < 0.001$) with depression than

Table 3: Prediction performances over time. Bold font indicates the best result for each feature set.

Feature set	SVM	LR	SVM
LIWC	0.629	0.611	0.623
LIWC+LDA	0.652	0.647	0.654
LIWC+bi-gram+LDA	0.706	0.718	0.715
LIWC+PLUS+bi-gram+LDA	0.802	0.800	0.780

other alternative combinations or LIWC alone, and perform reliably well at predicting depression. We report that all correlation coefficients meet ($p < 0.05$). We observe that adding PLUS features improves significantly on the results yielded by LIWC+bi-gram+LDA by a considerable margin. It should be noted that Pearson correlations between behavior (such as language use) and psychologically-based features rarely surpass an r of 0.4 (Meyer et al., 2001).

To make predictions over time for signals relevant to depression, we divided our data (857,294 tweets) into one-week periods. Specifically, we separately derived 50 topics from each subset. We prepared the training set using topics from the first to the penultimate week and took topics from the last week as the test set. We utilized three different classifiers: support vector machine (SVM), logistic regression (LR), and random forest (RF). We trained our classifiers with the three feature sets which achieved the highest Pearson’s (r) results in Table 1: LIWC+LDA, LIWC+bi-gram+LDA, and LIWC+PLUS+bi-gram+LDA. We considered the feature set LIWC itself as a baseline. For SVM, we set the regularization parameter $\lambda = 0.0001$ and the value γ of the radial basis function kernel to 0.5 and for RF, we set the number of trees to 500 and the maximum depth and number of features to 3 and 30, respectively. The prediction performances are reported as F-1 scores, i.e., the harmonic mean of precision and recall.

Table 3 shows the results for depression prediction over time. We see that the F-1 scores achieved with SVM, LR, and RF over the used feature sets are significantly higher than 0.5. We observe that SVM yielded the best performance over LIWC+PLUS+bi-gram+LDA features (0.802), surpassing the baseline (0.629) with a substantial improvement of 0.173. We note that the smallest result achieved with LIWC+PLUS+bi-gram+LDA (0.780) is superior to the performance of our second-best features, LIWC+bi-gram+LDA (0.718). These results in-

Table 4: Similarity between different depression-related topics addressed by individuals between before and during the stay-at-home period.

	Similarity Before During		
LIWC+LDA	JS	0.005	0.327
	KL	0.017	0.403
LIWC+bi-gram+LDA	JS	0.022	0.341
	KL	0.02	0.335
LIWC+PLUS+bi-gram+LDA	JS	0.025	0.478
	KL	0.027	0.290

dicates that LIWC+PLUS+bi-gram+LDA can detect signals relevant to depression more effectively than other features. LIWC+bi-gram+LDA features resulted in better results than LIWC features alone (0.629) or the combination of LIWC and LDA (0.654). We note that prediction quality depends heavily on complementary features, that is, the more a combination includes several features, the more it yields significantly better results.

$$\text{KL}(P\|Q) = \sum_{i \in [n]} p_i \times \log\left(\frac{p_i}{q_i}\right) \quad (2)$$

$$\text{JS}(P\|Q) = \frac{1}{2}\text{KL}(P\|M) + \frac{1}{2}\text{KL}(Q\|M) \quad (3)$$

Similarity between topics before and during stay-at-home restrictions. To discover overlapping behavioral characteristics of depression-related terms, we experimented with 50 topics on each one-week subset of the data as divided above. Each topic was represented by the top fifteen highest-probability words, out of which we retained solely the top ten depression-related words. We computed topic similarity using measures based on topic word probability distributions (Aletras and Stevenson, 2014) (such as Kullback-Leibler divergence (KL) (Kullback and Leibler, 1951)) and topic word sets (Mäntylä et al., 2018) (such as Jaccard similarity (JS) (Jaccard, 1912)).

Let us look at two discrete probability distributions $P = \{p_i\}_{i \in [n]}$ and $Q = \{q_i\}_{i \in [n]}$ supported on $[n]$. KL measures the difference between two probability distributions (Equation 2). Equation 2 determines how the Q distribution is different from the P distribution. KL is a non-negative, asymmetric distance (i.e., $\text{KL}(P\|Q) \neq \text{KL}(Q\|P)$) which yields zero if the two distributions are identical and can potentially equal infinity (Shlens, 1912). For JS, we measured the similarity between all possible topic pairs. JS is a symmetrized, smoothed version of KL which




Figure 1: Monthly trends of similarity between depression-related topics addressed by individuals. Note that we utilize LIWC+PLUS+bi-gram+LDA.




Figure 2: The number of individuals who have participated in depression-related topics. We make a weekly count of these individuals in the months before and during the stay-at-home order. For instance, the blue bar in Jan (January) is associated with the first week (W1), the red bar with the second week (W2), and so on.

measures the total KL divergence from the average mixture distribution, $M = \frac{(P+Q)}{2}$ (Equation 3). Some salient features of JS are that it is always defined, bounded and symmetric, and only vanishes when $P = Q$. When all the top words of a pair of topics are different, JS may result in 0. We found that some topic pairs bear words that include different spellings but are synonyms. To harmonize topic pairs that fall into that situation, we manually replaced synonyms with a single word on either side. We calculated the average JS and KL yielded from different time periods and found that depression-related words were overlapping from one topic to another during the stay-at-home period, and were slightly overlapping before

the stay-at-home order (see Table 4).

The Spearman correlation (ρ) between the two-similarity metrics is presented. We obtain $\rho = 0.839$ for LIWC+LDA, $\rho = 0.873$ for LIWC+bi-gram+LDA, and $\rho = 0.930$ for LIWC+PLUS+bi-gram+LDA during the stay-at-home period; and $\rho = 0.011$ for LIWC+LDA, $\rho = 0.016$ for LIWC+bi-gram+LDA, and $\rho = 0.02$ for LIWC+PLUS+bi-gram+LDA before the stay-at-home order. We report that all correlations are statistically significant ($p < 0.001$) and superior to 0.820 during the stay-at-home; and all correlations are not significant before the stay-at-home order ($p > 0.05$). In Figure 1, we utilize LIWC+PLUS+bi-gram+LDA. It should be recalled that the stay-at-home was is-

sued on March 12. Consequently, we combine all the data of March to measure the similarity. Specifically, January and February are fully comprised in the data before the stay-at-home. We obtain a KL of 0.024 and 0.035 in January and February ($p > 0.05$), respectively; 0.29 and 0.3 in April and May ($p < 0.001$), respectively; and 0.27 in March ($p < 0.05$). We get a JS of 0.026 and 0.0249 in January and February ($p > 0.05$), respectively; 0.48 and 0.5 in April and May ($p < 0.001$), respectively; and 0.39 in March ($p < 0.05$).

These results indicate strong and meaningful correlations between depression-indicative topics addressed during the stay-at-home. The language in these topics appears to be somewhat similar and recurs from one period to another during the stay-at-home period. This suggests that we should give more attention to this vocabulary when predicting depression from the individual-level.

Figure 2 shows the trend of individuals who have participated in depression-related topics. We observe a rise of participants within the second week of March, which symbolizes the onset of lockdown; and we note that the number substantially decreased within the fifth week of May, which represents the date on which COVID-19 lockdown restrictions began slowly being relaxed across the country. We calculated the percentage that individuals who have participated in depression-related topics represents to the overall number of individuals collected for each month. We found that 6.9%, 7.7%, 28.4%, 36.4% and 30.1%, respectively, for January, February, March, April and May.

5 Conclusion

This study focuses on detecting depression from social media postings, computes the language similarity between all possible topic pairs addressed by individuals, and predicts the evolution of depression over time. Our best classifier achieves F-1 scores as high as 0.8, which is a 0.173 relative the improvement over the baseline features. The proposed features yield a higher Pearson correlation ($r = 0.506$) than other alternative feature combinations and the improvement is statistically significant ($p < 0.001$). Prior work found that Pearson correlations between language use and psychologically-based features rarely exceed a value of $r = 0.4$, while our result has surpassed this value by 0.106. We measure the similarity

between different topics addressed by individuals to discover overlapping behavioral characteristics of depression-related words. We report that the Spearman correlations for this task are statistically significant for all the features utilized, and the proposed features specifically achieve the strongest Spearman correlation. In future work, we aim to include socioeconomic and demographic attributes with network and language information to predict depression at the regional level. Additionally, we would like to investigate affinity relationships between individuals who manifest signs of depression (Tshimula et al., 2020; Tshimula et al., 2019).

References

- [WHO2020] World Health Organization. 2020. *Coronavirus disease (COVID-19) pandemic*. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019> (2019, accessed 13 September 2020).
- [Canada2020] The Government of Canada. 2020. *Health: Diseases and conditions*. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019> (2019, accessed 13 September 2020).
- [McKibbin and Fernando2020] Warwick J. McKibbin and Roshen Fernando. 2020. *The global macroeconomic impacts of COVID-19: Seven scenarios*. CAMA Working Paper No. 19/2020. Available at SSRN, p. 45.
- [Wang et al.2020] Cuiyan Wang, Riyu Pan, Xiaoyang Wan, Yilin Tan, Linkang Xu, Cyrus S. Ho and Roger C. Ho. 2020. *Immediate psychological responses and associated factors during the initial stage of the 2019 coronavirus disease (COVID-19) epidemic among the general population in China*. *Int. J. Environ. Res. Public Health* 2020, 17(5):1729.
- [Bo et al.2020] Hai-Xin Bo, Wen Li, Yuan Yang, Yu Wang, Qinge Zhang, Teris Cheung, Xinjuan Wu and Yu-Tao Xiang. 2020. *Posttraumatic stress symptoms and attitude toward crisis mental health services among clinically stable patients with COVID-19 in China*. *Psychological Medicine*, 1–2.
- [Brooks et al.2020] Samantha K Brooks, Rebecca K Webster, Louise E Smith, Lisa Woodland, Simon Wessely, Neil Greenberg and Gideon James Rubin. 2020. *The psychological impact of quarantine and how to reduce it: Rapid review of the evidence*. *The Lancet*, 395:912–920.
- [Gunnell et al.2020] David Gunnell, Louis Appleby, Ella Arensman, Keith Hawton, Ann John, Nav Kapur, Murad Khan, Rory C O’Connor, Jane Pirkis and

- the COVID-19 Suicide Prevention Research Collaboration. 2020. *Suicide risk and prevention during the COVID-19 pandemic*. The Lancet Psychiatry, 7(6):468–471.
- [Li et al.2020] Sijia Li, Yilin Wang, Jia Xue, Nan Zhao and Tingshao Zhu. 2020. *The impact of COVID-19 epidemic declaration on psychological consequences: A study on active Weibo users*. International Journal of Environmental Research and Public Health, 17(6):2032.
- [Meng et al.2020] Hui Meng, Yang Xu, Jiali Dai, Yang Zhang, Baogeng Liu and Haibo Yang. 2020. *Analyze the psychological impact of COVID-19 among the elderly population in China and make corresponding suggestions*. Psychiatry Research, 289:112983.
- [Guntuku et al.2019] Sharath Chandra Guntuku, Rachele Schneider, Arthur Pelullo, Jami Young, Vivien Wong, Lyle Ungar, Daniel Polsky, Kevin G Volpp and Raina Merchant. 2019. *Studying expressions of loneliness in individuals using twitter: an observational study*. BMJ Open, 9:e030355.
- [De Choudhury2013] Munmun De Choudhury. 2013. *Role of social media in tackling challenges in mental health*. In Proc. of the 2nd International Workshop on Socially-Aware Multimedia, p. 49–52.
- [Kanter et al.2013] Jonathan W Kanter, Andrew M Busch, Cristal E Weeks, and Sara J Landes. 2008. *The nature of clinical depression: Symptoms, syndromes, and behavior analysis*. The Behavior Analyst, 31(1):1–21.
- [MHRC2020] Mental Health Research Canada (MHRC). 2020. *Mental health in crisis: How covid-19 is impacting Canadians*. <https://bit.ly/2UUvYKU> (2020, accessed 15 September 2020).
- [Guntuku et al.2017] Sharath Chandra Guntuku, David B Yaden, Margaret L Kern, Lyle H Ungar and Johannes C Eichstaedt. 2017. *Detecting depression and mental illness on social media: an integrative review*. Current Opinion in Behavioral Sciences, 18:43–49.
- [Guntuku et al.2019] Sharath Chandra Guntuku, Anneke Buffone, Kokil Jaidka, Johannes Eichstaedt and Lyle Ungar. 2019. *Understanding and measuring psychological stress using social media*. In ICWSM, p. 214–225.
- [Saha and De Choudhury2017] Koustuv Saha and Munmun De Choudhury. 2017. *Modeling stress with social media around incidents of gun violence on college campuses*. In Proc. of the ACM on Human-Computer Interaction, 92.
- [Thelwall2017] Mike Thelwall. 2017. *TensiStrength: Stress and relaxation magnitude detection for social media texts*. Information Processing & Management, 53(1):106–121.
- [Coppersmith et al.2014a] Glen Coppersmith, Mark Dredze and Craig Harman. 2014. *Quantifying mental health signals in Twitter*. In Proc. of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality, p. 51–60.
- [Coppersmith et al.2014b] Glen Coppersmith, Craig Harman and Mark Dredze. 2014. *Measuring post traumatic stress disorder in Twitter*. In ICWSM, p. 23–45.
- [He et al.2012] Qiwei He, Bernard Veldkamp and Theo de Vries. 2012. *Screening for posttraumatic stress disorder using verbal features in self narratives: A text mining approach*. Psychiatry Research, 198(3):441–447.
- [Cacheda et al.2019] Fidel Cacheda, Diego Fernandez, Francisco J Novoa and Victor Carneiro. 2019. *Early detection of depression: Social network analysis and random forest techniques*. Journal of Medical Internet Research, 21(6):e12554.
- [Coppersmith et al.2015] Glen Coppersmith, Mark Dredze, Craig Harman, Kristy Hollingshead and Margaret Mitchell. 2015. *CLPsych 2015 shared task: Depression and PTSD on Twitter*. In Proc. of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality, p. 31–39.
- [De Choudhury et al.2013] Munmun De Choudhury, Scott Counts and Eric Horvitz. 2013. *Social media as a measurement tool of depression in populations*. In Proc. of the 5th Annual ACM Web Science Conference, p. 47–56.
- [Jamil et al.2017] Zunaira Jamil, Diana Inkpen and Prasadith Buddhitha. 2017. *Monitoring tweets for depression to detect at-risk users*. In Proc. of the Fourth Workshop on Computational Linguistics and Clinical Psychology, p. 32–40.
- [Resnik et al.2013] Philip Resnik, Anderson Garron and Rebecca Resnik. 2013. *Using topic modeling to improve prediction of neuroticism and depression in college students*. In EMNLP, p. 1348–1353.
- [Resnik et al.2015a] Philip Resnik, William Armstrong, Leonardo Claudino and Thang Nguyen. 2015. *The University of Maryland CLPsych 2015 shared task system*. In Proc. of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality, p. 54–60.
- [Resnik et al.2015b] Philip Resnik, William Armstrong, Leonardo Claudino, Thang Nguyen, Viet-An Nguyen and Jordan Boyd-Graber. 2015. *Beyond LDA: Exploring supervised topic modeling for depression-related language in Twitter*. In Proc. of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality, p. 99–107.

- [Sadeque et al.2018] Farig Sadeque, Dongfang Xu and Steven Bethard. 2018. *Measuring the latency of depression detection in social media*. In Proc. of the Eleventh ACM International Conference on Web Search and Data Mining, p. 495–503.
- [Schwartz et al.2014] H. Andrew Schwartz, Johannes Eichstaedt, Margaret L. Kern, Gregory Park, Maarten Sap, David Stillwell, Michal Kosinski, Lyle Ungar. 2014. *Towards assessing changes in degree of depression through Facebook*. In Proc. of the ACL Workshop on Computational Linguistics and Clinical Psychology, 118–125.
- [Shen et al.2017] Guangyao Shen, Jia Jia, Liqiang Nie, Fuli Feng, Cunjun Zhang, Tianrui Hu, Tat-Seng Chua and Wenwu Zhu. 2017. *Depression detection via harvesting social media: A multimodal dictionary learning solution*. In Proc. of the Twenty-Sixth International Joint Conference on Artificial Intelligence, p. 3838–3844.
- [Tsugawa et al.2015] Sho Tsugawa, Yusuke Kikuchi, Fumio Kishino, Kosuke Nakajima, Yuichi Itoh and Hiroyuki Ohsaki. 2015. *Recognizing depression from Twitter activity*. In Proc. of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI'15, p. 3187–3196.
- [Pennebaker et al.2015] James W. Pennebaker, Ryan L. Boyd, Kayla N Jordan and Kate Blackburn. 2015. *The development and psychometric properties of LIWC2015*. Austin, TX: University of Texas at Austin.
- [Stark et al.2012] Anthony Stark, Izhak Shafran and Jeffrey Kaye. 2012. *Hello, who is calling?: can words reveal the social nature of conversations?*. In Proc. of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, p. 112–119.
- [Tadesse et al.2019] Michael M. Tadesse, Hongfei Lin, Bo Xu and Liang Yang. 2019. *Detection of depression-related posts in Reddit social media forum*. IEEE Access, p. 44883–44893.
- [Zhai et al.2012] Ke Zhai, Jordan Boyd-Graber, Nima Asadi and Mohamad L. Alkhouja. 2012. *Mr. LDA: a flexible large scale topic modeling package using variational inference in MapReduce*. In Proc. of the 21st international conference on World Wide Web, p. 879–888.
- [Hauthal et al.2019] Eva Hauthal, Dirk Burghardt and Alexander Dunkel. 2019. *Analyzing and visualizing emotional reactions expressed by emojis in location-based social media*. SPRS International Journal of Geo-Information, 8(3):113.
- [Chung and Pennebaker2007] Cindy Chung and James Pennebaker. 2007. *The Psychological functions of function words*. Social Communication, 343–359.
- [Ramos et al.2003] J. Ramos, J. Eden and R. Edu. 2003. *Using TF-IDF to determine word relevance in document queries*. In Proc. of the First Instructional Conference on Machine Learning, 242:133–142.
- [Wilson1988] Michael Wilson. 1988. *The mrc psycholinguistic database: Machine readable dictionary*. Behavioural Research Methods, Instruments and Computers, 20:6–11.
- [WHO1994] World Health Organization. 1994. *Lexicon of psychiatric and mental health terms*. 2nd ed. World Health Organization, <https://apps.who.int/iris/handle/10665/39342> (accessed 13 September 2020).
- [Mohammad and Turney2013] Saif M. Mohammad and Peter D. Turney. 2013. *Crowdsourcing a word-emotion association lexicon*. Computational Intelligence, 29(3):436–465.
- [Blei et al.2003] David M. Blei, Andrew Y. Ng and Michael I. Jordan. 2003. *Latent Dirichlet allocation*. Journal of Machine Learning Research, 3:993–1022.
- [Meyer et al.2001] G J Meyer, S E Finn, L D Eyde, G G Kay, K L Moreland, R R Dies, E J Eisman, T W Kubiszyn and G M Reed. 2001. *Psychological testing and psychological assessment*. A review of evidence and issues. American Psychologist, 56(2):128–165.
- [Aletras and Stevenson2014] Nikolaos Aletras and Mark Stevenson. 2014. *Measuring the similarity between automatically generated topics*. In Proc. of the 14th Conference of the European Chapter of the Association for Computational Linguistics, p. 22–27.
- [Kullback and Leibler1951] S. Kullback and RA Leibler. 1951. *On information and sufficiency*. Annals of Mathematical Statistics, 22(1):79–86.
- [Mäntylä et al.2018] Mika Mäntylä, Maëlick Claes and Umar Farooq. 2018. *Measuring LDA topic stability from clusters of replicated runs*. In Proc. of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement, p. 1–4.
- [Jaccard1912] P. Jaccard. 1912. *The Distribution of the flora in the alpine zone*. New Phytologist, 11(2):37–50.
- [Shlens1912] J. Shlens. 2014. *Notes on Kullback-Leibler divergence and likelihood theory*. In arXiv preprint arXiv:1404.2000.
- [Tshimula et al.2020] Jean Marie Tshimula, Belkacem Chikhaoui and Shengrui Wang. 2019. *HAR-search: A method to discover hidden affinity relationships in online communities*. In Proc. of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, p. 176–183.
- [Tshimula et al.2019] Jean Marie Tshimula, Belkacem Chikhaoui and Shengrui Wang. 2020. *A new approach for affinity relationship discovery in online forums*. Social Network Analysis and Mining, 10:40.