

Analysing the Causes of Depressed Mood from Depression Vulnerable Individuals

Noor Fazilla Abd Yusof, Chenghua Lin and Frank Guerin

Computing Science, University of Aberdeen, Aberdeen UK

{noorfazilla.yusof, chenghua.lin, f.guerin}@abdn.ac.uk

Abstract

We develop a computational model to discover the potential causes of depression by analysing the topics from user-generated contents. We show the most prominent causes, and how these causes evolve over time. Also, we highlight the differences in causes between students with low and high neuroticism. Our studies demonstrate that the topics reveal valuable clues about the causes contributing to depressed mood. Identifying causes can have a significant impact on improving the quality of depression care; thereby providing greater insights into a patient's state for pertinent treatment recommendations. Hence, this study significantly expands the ability to discover the potential factors that trigger depression, making it possible to increase the efficiency of depression treatment.

1 Introduction

Depression is one of the most common mental disorders that can affect people of all ages. It is the leading cause of disability and requires significant health care cost to treat effectively (Smith et al., 2013). Early detection and treatment has a profound impact to encourage remission, and prevent relapse (Halfin, 2007). However, it is rather common that the stigma associated with mental illness makes patients reluctant to seek help, or makes them tend to answer questions in a manner that will be viewed favourable by the clinician (Chandra and Minkovitz, 2006). In addition, clinical diagnosis depends on the hypothetical or retrospective self-reports of behaviour, requiring patients to reflect on what they were doing and thinking sometime in the past, which may have become obscured over time. Hence, it is difficult for physi-

cians to capture first-hand the patients' own experiences, an important factor for diagnosis and providing the most appropriate treatment at the point of care (Kosinski et al., 2015).

Social media sites provide great venues for people to share their experiences, vent emotion and stress, and seek social support. Therefore, mental health studies based on social media present several advantages (Inkster et al., 2016). For instance, these digital footprints contain vast amounts of implicit knowledge, which are useful for medical practitioners to understand patients' experiences outside the controlled clinical environment. In addition, information captured during clinical consultation generally reflects only the situation of the patient at the time of care. In contrast, data collected from social media is dynamic, thereby providing opportunities for observing and recognising critical changes in patients' behaviour and permitting certain interventions in real time (Inkster et al., 2016).

Due to the advantages identified above, social media and natural language processing techniques are increasingly used in a wide range of mental health related studies. This includes works that can detect (Coppersmith et al., 2014; Resnik et al., 2015) and measure the degree of depression (Schwartz et al., 2014), identify depressive symptoms (Mowery et al., 2016), and detect the behavioural changes associated with onset of depression (De Choudhury et al., 2013b). There are also works focus on predicting personality traits such as neuroticism (Resnik et al., 2013), which is known to be highly associated with depression.

However, the above mentioned works mainly focus on recognising depression, with the potential causes of depression being ignored. Discovering the potential causes of depression is a worthwhile aspect of psychiatric diagnosis in order to offer the individual patient with solution-

specific, interpersonal or psychodynamic therapy for the best treatment outcome (Nathan and Gorman, 2015). For example, interpersonal psychotherapy can benefit patients who have recognised interpersonal problems as the cause of their depression. Furthermore, identifying causes can improve the efficiency of the treatment plan, as it is normally involved in a patient’s diagnostic evaluation and can help recognise possible barriers to treatment support (Gilman et al., 2013).

In this paper, we tackle the research challenge of discovering potential causes of depression by analysing the topics from user-generated contents. We approach the problem by developing a computational model which extends the dynamic topic model (He et al., 2012). In order to extract coherent sentiment-bearing topics that are indicative for identifying the causes of depression, we develop the major categories of depression cause based on two well-known resources i.e., DSM-IV (Gilman et al., 2013) and Crisis Text Line (www.crisistextline.org). We also propose a mechanism to incorporate domain knowledge of depression causes into our model for guiding the model inference procedure, which helps us to extract depression related and meaningful topics. Experimental results show that our approach can extract topics revealing valuable clues and risk factors about the causes contributing to depression based on informal user-generated data; thereby providing deep insights into a patient’s state for pertinent treatment recommendations.

2 Related Work

A wide range of risk factors are associated with the development and persistence of depression (e.g., biological, psychological or cognitive), however, psychosocial are among the strongest (Slavich and Irwin, 2014). To examine depressive disorder, one effective means is via language analysis, as the use of language can be linked to important information about people’s behaviours and psychological insights (Pennebaker et al., 2003).

De Choudhury et al. (2013a) showed that Support Vector Machines (SVM) with Radial Basis Function (RBF) kernel could predict depression signs from Twitter posts. A similar approach was applied to Japanese Twitter posts for investigating the correlations between users’ activities and depression (Tsugawa et al., 2015). Copper-smith et al. (2014) used language models and Lin-

guistic Inquiry Word Count (LIWC) (Pennebaker et al., 2007), a psychometrically validated analysis tool, to explore the language differences of Post-Traumatic Stress Disorder users. Schwartz et al. (2014) built a regression model to predict the degree of depression across seasons based on language features on Facebook.

In contrast to the works above which analyse static data, there has also been research in examining changes in behavioural patterns relating to onset of depression. For instance, De Choudhury and Counts (2013) analysed the behavioural changes of new mothers who are at risk of postpartum depression following childbirth. In the subsequent work, De Choudhury et al. (2013b) further predicted whether one is likely to have depression in the future by examining the patterns of one’s Twitter postings in a one-year time frame. Both studies showed that significant changes in social media activities could be the potential measures for predicting depression.

Another stream of works employ the Big Five personality traits (John and Srivastava, 1999) in depressive illness related studies. The Big Five personality traits define five different personality characteristics i.e., *extroversion*, *agreeableness*, *conscientiousness*, *neuroticism*, and *openness*. Among these personal traits, neuroticism is known to have a substantial correlation to the prior development of common depressive illness and psychological distress (Fanous et al., 2007). Schwartz et al. (2013) explored an open-vocabulary approach to gain psychological insights based on the demographics and personality traits framework. The works of Resnik et al. (2013, 2015) are most closely related to ours, as they explored topic modelling to automatically identify depressive-related language. They showed that using topic models provides better predictive performance than solely relying on pre-defined lexical features. They also highlighted that the topics extracted by Latent Dirichlet Allocation (LDA) (Blei et al., 2003) are meaningful and psychologically relevant. Specifically, Resnik et al. (2013) combined lexical features with features extracted by topic models, which improves the prediction of neuroticism and depression on student essay data. In the more recent work, Resnik et al. (2015) further explored using Supervised LDA (Mcauliffe and Blei, 2008) and Supervised Anchor model (Arora et al., 2013) to analyse the linguistic

signal for detecting depression.

To the best of our knowledge, no studies have explored the research problem of automatically identifying the causes of depression using natural language processing techniques. We envisage that by addressing this problem, our work would be useful for both individual and population-level mental health monitoring and prevention.

3 Methodology

In this section, we first describe how we acquire the categories of causes of depression, and then describe the computational framework for automatically extracting coherent topics that are indicative for identifying the causes of depression.

3.1 Development of Cause Categories for Depression

For extracting potential causes of depression from text, we first develop the major categories of depression cause based on two well-known resources.

First, we construct the primary list based on the risk factors outlined in the description of Axis IV in DSM-IV (Gilman et al., 2013). DSM-IV is a standard diagnostic manual of mental disorders, which defines nine broad categories that increase the risk of developing depression. The broad categories include problems related to *primary support group, social environment, occupational, economic, educational, housing, accessing to health-care services, and legal/crime*. However, some of the categories are too broad or do not state precisely enough for the causes, such as “primary support group” and “social environment”. Therefore, in order to obtain a more comprehensive list of depression causes, we further make use of the resources available from Crisis Text Line¹.

Basically, Crisis Text Line is one of the largest crisis counselling services which supports a wide range of issues from relationship concerns to depression to suicidal thoughts. We utilise the trends list of 17 issues prevalent to depression: *anxiety, bereavement, bullying, eating disorders, family issues, friend issues, health concerns, isolations, LGBT issues, physical abuse, relationships, school problems, self-harm, sexual abuse, stress, substance abuse, and suicidal thoughts*.

To ensure the quality of the cause categories developed based on the above resources, we also

¹www.crisistrends.org

consulted a physician who had an extensive experiences dealing with depression cases. We worked together to refine the list, taking into consideration the leading factors contributing to depression. We discarded the ones which are not quite related to the causes of depression, i.e., suicidal thoughts and self-harm. We also added to the list two other common issues that reinforce negative thoughts or emotions. For instance, body image (e.g. body-hatred, overweight, underweight) and homesickness have been considered associated with psychological disturbance, especially among young adults.

We present the categories of depression causes in Table 1. While this is by no means an explicit list of causes of depression. Indeed, there can be as many different causes of depression as possible. Our argument that this is an initial development and we account for their relative significance. The list can be extended in our future research.

Bullying	Family issues
Housing	Health concerns
Body image	Substance abuse
Bereavement	Occupation
Homesickness	Academic
Relationships	Economic
Discrimination	Sexual abuse
Physical abuse	

Table 1: Depression cause categories.

3.2 The dynamic Joint Sentiment-Topic model

We employ the dynamic Joint Sentiment-Topic (dJST) model (He et al., 2012) to extract coherent sentiment-bearing topics that are indicative for identifying the causes of depression. In addition, the model is also capable to track how the topics evolve over time, permitting investigation of the prominence of depression causes.

The dJST model, as shown in Figure 1, assumes that the current sentiment-topic words distributions are generated by the word distributions at the previous epochs. Each document d at epoch t is represented as a vector of word tokens, $w_d^t = (w_{d_1}^t, w_{d_2}^t, \dots, w_{d_{N_d}}^t)$. By assuming that the documents at current epoch are influenced by documents at past, the current sentiment-topic specific word distributions $\varphi_{l,z}^t$ at epoch t are generated according to the word distribu-

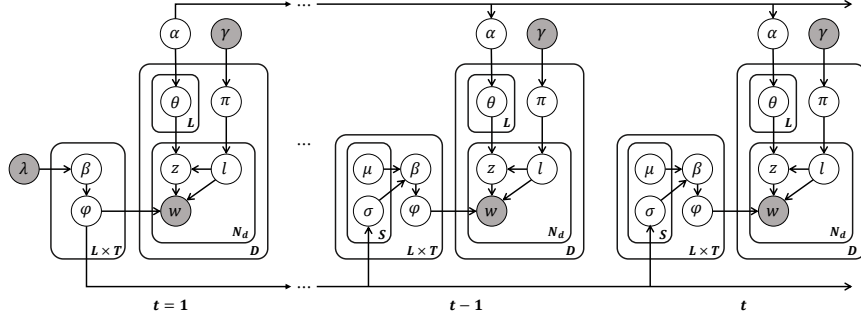


Figure 1: dJST model.

tions at previous epochs. In particular, an evolutionary matrix of topic z and sentiment label l , $E_{l,z}^t$ where each column in the matrix is the word distribution of topic z and sentiment label l , $\sigma_{l,z,s}^t$, generated for document streams received within the time slice specified by s , where $s \in \{t-S, t-S+1, \dots, t-1\}$, the current sentiment-topic-word distributions are dependent on the previous sentiment-topic specific word distributions in the last S epochs.

We then attach a vector of S weights $\mu_{l,z}^t = [\mu_{l,z,0}^t, \mu_{l,z,1}^t, \dots, \mu_{l,z,S}^t]^T$, each of which determines the contribution of time slice s in computing the priors of $\varphi_{l,z}^t$. Hence, the Dirichlet prior for sentiment-topic-word distributions at epoch t is $\beta_{l,z}^t = E_{l,z}^{t-1} \mu_{l,z}^t$. Assuming we have already calculated the evolutionary parameters $\{E_{l,z}^{t-1}, \mu_{l,z}^t\}$ for the current epoch t , the generative story of dJST as shown in Figure 1 at epoch t is given as follows:

- For each sentiment label $l = 1, \dots, L$
 - For each topic $z = 1, \dots, T$
 - * Draw $\alpha_{l,z}^t | \alpha_{l,z}^{t-1} \sim \Gamma(\nu \alpha_{l,z}^{t-1}, \nu)$
 - * Compute $\beta_{l,z}^t = \mu_{l,z}^t E_{l,z}^t$
 - * Draw $\varphi_{l,z}^t \sim \text{Dir}(\beta_{l,z}^t)$.
- For each document $d = 1, \dots, D^t$
 - Choose a distribution $\pi_d^t \sim \text{Dir}(\gamma)$.
 - For each sentiment label l under document d , choose a distribution $\theta_{d,l}^t \sim \text{Dir}(\sigma^t)$.
 - For each word $n = 1, \dots, N_d$ in document d
 - * Choose a sentiment label $l_n \sim \text{Mult}(\pi_d^t)$,
 - * Choose a topic $z_n \sim \text{Mult}(\theta_{d,l_n}^t)$,
 - * Choose a word $w_n \sim \text{Mult}(\varphi_{l_n,z_n}^t)$.

3.2.1 Incorporating Domain Knowledge

In order to extract depression cause relevant and semantically meaningful topics, we incorporate two types of domain knowledge into our model for guiding the model inference procedure.

The first type of domain knowledge is a general sentiment lexicon consists of the Multiperspective Question Answering (MPQA) subjectivity lexicon² and the appraisal lexicon (Bloom et al., 2007). In total, there are 1,511 positive and 2,542 negative words, respectively. The second type of domain knowledge is seed words relating to depression cause categories described in Section 3.1. To acquire the seed words, we make use of OneLook Dictionary³, a search engine for words and phrases. Precisely, for each of the depression category, we query OneLook by searching for words related to the depression cause category label. For instance, for category “bullying”, we obtain 539 words related to bullying e.g., *intimidation, harrasment, brutal*, etc. We filter and retain the top 50 most relevant words to the search query. We use those words to be the seed words as prior knowledge for the model training. The total number of seed words contains 750 words.

At epoch 1, the Dirichlet priors β of size $L \times T \times V$ are first initialized as symmetric priors of 0.01 (Steyvers and Griffiths, 2007), and then modified by a transformation matrix λ of size $L \times V$. We encode the word prior knowledge in the way that elements of β corresponding to positive sentiment words, e.g., *good*, will have small values for topics associated with negative sentiment labels, and vice versa for the negative sentiment words. For subsequent epochs, if any new words encountered, the prior knowledge will be incorporated in a similar way. But for existing words, their

²<http://www.cs.pitt.edu/mpqa/>

³<https://www.onelook.com>

Dirichlet priors for sentiment-topic-word distributions are obtained using $\beta_{l,z}^t = \mathbf{E}_{l,z}^{t-1} \boldsymbol{\mu}_{l,z}^t$.

3.2.2 Online Inference

We employ a stochastic EM algorithm (He et al., 2012) to sequentially update model parameters at each epoch. At each EM iteration, we infer latent sentiment labels and topics using the collapsed Gibbs sampling and estimate the hyperparameters using maximum likelihood.

We set the symmetric prior $\gamma = (0.05 \times \text{average document length})/L$, where L is the total number of sentiment labels and the value of 0.05 on average allocates 5% of probability mass for mixing. Moreover, there are two sets of evolutionary parameters to be estimated, the weight parameters $\boldsymbol{\mu}$ and the evolutionary matrix \mathbf{E} . We set $\boldsymbol{\mu}$ using an exponential decay function $\boldsymbol{\mu}^t = \exp(-\kappa t)$, so that more recent documents would have a relatively stronger influence on the model parameters in the current epoch compared to earlier documents. We set $\kappa = 0.5$ in the experiments. The derivation of the evolutionary matrix \mathbf{E} requires the estimation of each of its elements, $\sigma_{l,z,w,s}$, i.e., the word distribution of word w in topic z and sentiment label l at time slice s , which is defined as $\sigma_{l,z,w,s}^t = \frac{C_{l,z,w,s}^t}{\sum_w C_{l,z,w,s}^t}$. Here $C_{l,z,w,s}$ is the expected number of times word w was assigned to sentiment label l and topic z at time slice s . Each time slice s is equivalent to an epoch t , thus $C_{l,z,w,s}$ can be obtained directly from $N_{l,z,w,t}$ by setting $s = t$.

4 Experimental Setup

Dataset. We conducted experiments on a real-world dataset, namely, the student essay dataset. This dataset is publicly available and has been used in a number of mental health studies (Resnik et al., 2013). It contains 6,459 stream-of-consciousness essays collected between 1997 and 2008, and each essay is labelled with Big-5 personality traits scores. As discussed earlier, neuroticism (negative affectivity) is a factor that strongly associates with high risk of depressive disorders. We divided the dataset into two categories based on the neuroticism scores, i.e, essays with positive scores are classified as high neuroticism, and negative score as low neuroticism. Any essays missing personality traits scores were eliminated from the dataset. The final dataset contains a total number of 4,954 essays with 2,566 asso-




Figure 2: Number of stream-of-consciousness essays over 11 years (1997-2008).

ciated with high neuroticism and 2,388 associated with low neuroticism, as shown in Figure 2.

General Settings. Each dataset underwent pre-processing including conversion to lowercase, removal of non-alphanumeric characters, and removal of stop words. We empirically set the number of topics to 20 for the 2 sentiment labels (i.e., positive and negative), which is equivalent to a total of 40 sentiment-topic clusters. The number of time-slices set to 4.

5 Experimental Results

In this section, we present our results on the experimental datasets. In particular, we aim to investigate the following two research questions: (i) what are the most prominent causes for neuroticism among college students and how do these causes evolve over time; and (ii) what are the differences of the topics extracted from essays written by of students from low and high neuroticism groups.

5.1 Analysing the Depressing-related Issues from the High Neuroticism Group

In this section, we present the results on extracting topics for depression-related causes analysis. Figure 3 illustrates the example topics from high neuroticism (negative sentiment) and low neuroticism (positive sentiment). The topics extracted are in agreement with those reported in the risk factors that contribute to stress experienced by students (Robotham and Julian, 2006). We discuss some of the topics in details below.

Academic. Unsurprisingly, one of the prominent causes among the students with high-neuroticism score is related to academic studies, consistent with the report of primary causes of stress among




Figure 3: Example topics from high and low neuroticism. The upper panels show the topics under negative sentiment (high neuroticism); the lower panel under positive sentiment (low neuroticism). Note: The connecting arrows describe the sentiment differences towards the same topic theme.




Figure 4: Frequency distribution of top 5 sentiment topics across years (a) High neuroticism; (b) Low neuroticism.

students⁴. Academic topic words in Figure 3 (e.g., “study, test, bad, grades, worried”) show that students express a very stressful experiences in study. For example, “I am scared that my grades

won’t be able to cut it though”, “I am pretty much worried about my classes and what grades I will get in them”.

Relationship and homesickness. Our analysis found that relationship problems and homesickness are also widely common among stu-

⁴<https://yougov.co.uk/news/2016/08/09/quarter-britains-students-are-afflicted-mental-hea/>

dents, as shown in the Relationship and Homesickness topics. Indeed, intimate relationships with a partner can be a great source of love, support and excitement. However, relationships can also be a source of grief and anguish if they go wrong. University students are in a period of personal change, which can then make them feel less sure or what they want or how to cope with relationship problems. For examples, *“I feel insecure about our relationship myself and in some way feel like I am not worthy of someone liking me”*. Research by the National Union of Students⁵ shows that 50-70% of new students suffer from homesickness to some extent within in their first two or three weeks. Although most students find their symptoms begin to fade after a few weeks, the symptoms tend to stay longer for students with high neuroticism.

Housing. Another interesting finding is that bad accommodation seems to have strong negative impact on the socio-emotional development and psychological distress of students. For instance, there are lot of complaint about the condition of the university accommodation, e.g., *“Someone in this room needs to buy a mop because our floor is getting really gross”*, *“Our shower is very small in the first place and combined with being dirty, well that’s just plain bad”*. There is a strong link of mental health problems with insecure, and the chaotic way of living (Tight, 2011).

Body image and health concerns. Individuals with high level of neuroticism seem less satisfied with their body image and health as implied in topics Body Image and Health concerns. For instance, messages similar to *“I am so nervous about gaining weight. I always watch what i eat”* appear quite often in the dataset. This shows that physiological factors do promote to dissatisfaction in students’ life, which lead to low self-esteem (correlated with high neuroticism).

5.2 Comparing the High and the Low Neuroticism Groups.

We estimate the probability of top 5 topics across years using the relative frequency technique. We calculate the ‘relative observed frequencies’ of a topic, and divide the number of occurrences of the topic by the number of documents for that particular year. Figure 4 shows the distribution of top 5 topics of high and low neuroticism. We found

⁵<https://www.nus.org.uk/>

that the topic Relationship is very common among students from both groups. The trend reflects the fact that students are prone to stress in student life, often caused by the poor relationship issues, which also lead to struggles in academic, social adjustment, and individual self-esteem.

Another interesting key difference to highlight is on the social engagement level. Students with low neuroticism are much more active and show more interest in various social activities (e.g., music and sports). Similar findings are consistent with research (Afshar et al., 2015), in that the individuals with low-neuroticism are more likely to utilise relaxation (music, meditation, yoga, etc.) and physical recreation (regular exercise, sports, running, etc.) as coping mechanisms.

Sentiment analysis. We discuss the sentiment differences on the same topic between two groups, as shown in Figure 3. The topics about academic and relationship are both prominent among these groups, however, there are differences in perceptions and emotional reactivity towards these topics. Specifically, students with high-neuroticism respond poorly to environmental stress and interpret ordinary situations as threatening and experience minor frustrations as hopelessly overwhelming. For example, they seem to have difficulty in coping the issues and challenges from academic studies. Below are some examples,

- *“I hate myself for not doing well in some other class. Its a vicious cycle that I can’t seem to get out of. I do bad in one class because I focus on all the wrong things and then it carries over to ever other class, which in turn makes my academics suffer”*.
- *“I’m worried about studying for psychology. It’s my first collegiate test. I’ll probably do terrible. or at least far less than my expectations”*.

Whereas, students with low neuroticism seem to have higher stress tolerance when dealing with academic pressures. They are more resilient to challenges, embrace and overcome obstacles in a positive way. Take for examples the following,

- *“I am eager for the future and ecstatic for what is yet to come. I hope I am joining the right organizations that appeal to me. I also hope I will stay academically strong as I was in high school”*.
- *“Its going to be my first official test in college. I just can’t imagine me taking a test. I just need to relax, study the best I can, and be optimistic about my academics”*.

6 Conclusion

The causes of depression can vary greatly from person to person. It is a great challenge for clinical practice in the recognition and treatment of depression, particularly when there are barriers in getting the appropriate support, e.g., time constraints in primary care, a strong social stigma attached to mental illness and discrimination. Our approach on topic modelling for classification certainly bridges the gap and significantly expands the access in identifying the possible factors that trigger depression in individuals. Identifying the causes of depression increases the accuracy of selecting the most appropriate treatment and improves the quality of depression care. Therefore, further research should be undertaken to optimise topic models for drawing out potential causes of depression from social media data. Furthermore, a dataset with ground truth that covers wider causes of depression such as financial and occupation should be explored in the future, to cater for different groups of people e.g., employee, housewives, etc.

Acknowledgments

This work is supported by the awards made by the UK Engineering and Physical Sciences Research Council (Grant number: EP/P005810/1). Lead author would like to thank Ministry of Higher Education Malaysia and Technical University of Malaysia Malacca for PhD scholarship award.

References

- H Afshar, H Roohafza, A Keshteli, M Mazaheri, A Feizi, and P Adibi. 2015. The association of personality traits and coping styles according to stress level. *IJRMS* 20(4):353.
- S Arora, R Ge, Y Halpern, D Mimno, A Moitra, D Sonntag, Y Wu, and M Zhu. 2013. A Practical Algorithm for Topic Modeling with Provable Guarantees. In *Int. Conf. Mach. Learn.*, pages 280–288.
- D Blei, Andrew Ng, and M Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3(Jan):993–1022.
- K Bloom, S Stein, and S Argamon. 2007. Appraisal extraction for news opinion analysis at NTCIR-6. In *Proc. NTCIR-6, National Institute of Informatics*.
- A Chandra and C Minkovitz. 2006. Stigma starts early: Gender differences in teen willingness to use mental health services. *J. Adolesc. Heal.* 38(6):754–e1.
- G Coppersmith, C Harman, and M Dredze. 2014. Measuring Post Traumatic Stress Disorder in Twitter. In *Proc. ICWSM*, pages 579–582.
- M De Choudhury and Sand Horvitz E Counts. 2013. Predicting postpartum changes in emotion and behavior via social media. In *Proc. SIGCHI-13*, pages 3267–3276.
- M De Choudhury, S Counts, and E Horvitz. 2013a. Social media as a measurement tool of depression in populations. In *Proc. ACM WebSci*, pages 47–56.
- M De Choudhury, M Gamon, S Counts, and E Horvitz. 2013b. Predicting Depression via Social Media. In *Proc. ICWSM*, volume 2, page 2.
- A Fanous, M Neale, S Aggen, and K Kendler. 2007. A longitudinal study of personality and major depression in a population-based sample of male twins. *Psychol. Med.* 37(08):1163.
- S Gilman, N Trinh, J Smoller, M Fava, J Murphy, and J Breslau. 2013. Psychosocial stressors and the prognosis of major depression: a test of Axis IV. *Psychol. Med.* 43(2):303–316.
- A Halfin. 2007. Depression: the benefits of early and appropriate treatment. *Am. J. Manag. Care* 13(4 Suppl):S92–7.
- Y He, C Lin, and A Cano. 2012. Online sentiment and topic dynamics tracking over the streaming data. In *Int. Conf. Soc. Comput.*, pages 258–266.
- B Inkster, D Stillwell, M Kosinski, and P Jones. 2016. A decade into Facebook: where is psychiatry in the digital age? *Lancet Psychiatry* 3(11):1087–1090.
- O John and S Srivastava. 1999. The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handb. Personal. Theory Res.* 2(1999):102–138.
- M Kosinski, S Matz, S Gosling, V Popov, and D Stillwell. 2015. Facebook as a research tool for the social sciences: Opportunities, challenges, ethical considerations, and practical guidelines. *Am. Psychol.* 70(6):543–556.
- J Mcauliffe and D Blei. 2008. Supervised topic models. In *Adv. Neural Inf. Process. Syst.*, Curran Associates, Inc., pages 121–128.
- D Mowery, A Park, M Conway, and B Craig. 2016. Towards Automatically Classifying Depressive Symptoms from Twitter Data for Population Health. In *Proc. PEOPLES 2016*, pages 182–191.
- P Nathan and J Gorman. 2015. *A guide to treatments that work*. Oxford University Press.
- J Pennebaker, C Chung, M Ireland, A Gonzales, and R Booth. 2007. The development and psychometric properties of liwc2007: Liwc. net.

- J Pennebaker, M Mehl, and K Niederhoffer. 2003. Psychological aspects of natural language use: Our words, our selves. *Annu. Rev. Psychol.* 54(1):547–577.
- P Resnik, W Armstrong, L Claudino, T Nguyen, V Nguyen, and J Boyd-graber. 2015. Beyond LDA : Exploring Supervised Topic Modeling for Depression-Related Language in Twitter. In *Proc. CLPsych.* volume 1, pages 99–107.
- P Resnik, A Garron, and R Resnik. 2013. Using Topic Modeling to Improve Prediction of Neuroticism and Depression in College Students. In *Proc. EMNLP.* pages 1348–1353.
- D Robotham and C Julian. 2006. Stress and the higher education student: a critical review of the literature. *J. Furth. High. Educ.* 30(2):107–117.
- H Schwartz, J Eichstaedt, M Kern, L Dziurzynski, S Ramones, M Agrawal, A Shah, M Kosinski, D Stillwell, M Seligman, and L Ungar. 2013. Personality, Gender, and Age in the Language of Social Media: The Open-Vocabulary Approach. *PLoS One* 8(9):e73791.
- H Schwartz, J Eichstaedt, M Kern, G Park, M Sap, D Stillwell, M Kosinski, and L Ungar. 2014. Towards Assessing Changes in Degree of Depression through Facebook. In *Proc. CLPsych.* pages 118–125.
- G Slavich and M Irwin. 2014. From stress to inflammation and major depressive disorder: A social signal transduction theory of depression. *Psychol. Bull.* 140(3):774.
- K Smith, P Renshaw, and J Bilello. 2013. The diagnosis of depression: Current and emerging methods. *Compr. Psychiatry* 54(1):1–6.
- M Steyvers and T Griffiths. 2007. Probabilistic topic models. *Handb. Latent Semant. Anal.* 427(7):424–440.
- M Tight. 2011. Student accommodation in higher education in the United Kingdom: changing postwar attitudes. *Oxford Rev. Educ.* 37(1):109–122.
- S Tsugawa, Y Kikuchi, F Kishino, K Nakajima, Y Itoh, and H Ohsaki. 2015. Recognizing Depression from Twitter Activity. In *Proc. ACM CHI.* pages 3187–3196.