Not Just Depressed: Bipolar Disorder Prediction on Reddit

Ivan Seuklić Matej Gjurković Jan Šnajder
Text Analysis and Knowledge Engineering Lab
Faculty of Electrical Engineering and Computing, University of Zagreb
Unska 3, 10000 Zagreb, Croatia
{ivan.sekulic,matej.gjurkovic,jan.snajder}@fer.hr

Abstract

Bipolar disorder, an illness characterized by manic and depressive episodes, affects more than 60 million people worldwide. We present a preliminary study on bipolar disorder prediction from user-generated text on Reddit, which relies on users' self-reported labels. Our benchmark classifiers for bipolar disorder prediction outperform the baselines and reach accuracy and F1-scores of above 86%. Feature analysis shows interesting differences in language use between users with bipolar disorders and the control group, including differences in the use of emotion-expressive words.

1 Introduction

World Health Organization's 2017 and Wykes et al. (2015) report that up to 27% of adult population in Europe suffer or have suffered from some kind of mental disorder. Unfortunately, as much as 35–50% of those affected go undiagnosed and receive no treatment for their illness. To counter that, the WHO's Mental Health Action Plan's (Saxena et al., 2013) lists as one of its main objectives the gathering of information and evidence on mental conditions. At the same time, analysis of texts produced by authors affected by mental disorders is attracting increased attention in the natural language processing community. The research is geared toward a deeper understanding of mental health and the development of models for early detection of various mental disorders, especially on social networks.

In this paper we focus on bipolar disorder, a complex psychiatric disorder manifested by uncontrolled changes in mood and energy levels. Bipolar disorder is characterized by manic episodes, during which people feel abnormally elevated and energized, and depression episodes, manifested in decreased activity levels and a feeling of hopelessness. The two phases are recurrent and differ in intensity and duration, greatly affecting the person's capacity to carry out daily tasks. Bipolar disorder affects more than 60 million people, or almost 1% of the world population (Anderson et al., 2012). The suicide rate in patients diagnosed with bipolar disorder is more than 6% (Nordentoft et al., 2011). There is thus a clear need for the development of systems capable of early detection of this illness.

As a first step toward that goal, in this paper we present a preliminary study on bipolar disorder prediction based on user-generated texts on social media. The main problem in detecting mental disorders from user-generated text is the lack of labeled datasets. We follow the recent strand of research (Gkotsis et al., 2016; De Choudhury et al., 2016; Shen and Rudzicz, 2017; Gjurković and Šnajder, 2018) and use Reddit as a rich and diverse source of high-volume data with self-reported labels. Our study consists of three parts. First, we test benchmark models for predicting Reddit users with bipolar disorder. Second, we carry out a feature analysis to determine which psycholinguistic features are good predictors of the disorder. Lastly, acknowledging that emotional swings are the main symptom of the disorder, we analyze the emotion-expressive textual features in bipolar disorder users and the non-bipolar control group of users.

2 Related Work

Psychologist have long studied the language use in patients with mental disorders, including schizophrenia (Taylor et al., 1994), suicidal tendencies (Thomas and Duszynski, 1985), and depression (Schnurr et al., 1992). Lately, computer-based analysis with LIWC (Linguistic Inquiry and Word Count) (Pennebaker et al., 2001) resource was used to extract features for various studies regarding mental health (Pennebaker and King, 1999). For example, Stirman and Pennebaker (2001) found

Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, pages 72–78
Brussels, Belgium, October 31, 2018.
https://doi.org/10.18653/v1/P17

the increased use of the first-person singular pronouns (*I, me, my*) in poems to be a good predictor of suicidal behavior, while Rude et al. (2004) detected an excessive use of the pronoun *I* in essays of depressed psychology students. In a recent study, however, Tackman et al. (2018) suggest that first-person singular pronouns may be better viewed as a marker of general distress or negative emotionality rather than as a specific marker of depression.

A number of studies looked into the use of emotion-expressive words. Rude et al. (2004) found that currently depressed students used more negative emotion words than never-depressed students. Halder et al. (2017) tracked linguistic changes of social network users over time to understand the progression of their emotional status. Kramer et al. (2004) found that conversations in bipolar support chat rooms contained more positively valence words and slightly more negatively valenced emotions than casual conversations.

Much recent work has leveraged social media as a source of user-generated text for mental health profiling (Park et al., 2012). Most studies used Twitter data; e.g., De Choudhury et al. (2013) predicted depression in Twitter users, while CLPsych 2015 shared task (Coppersmith et al., 2015b) addressed depression and post-traumatic stress disorder (PTSD). Bipolar disorder on Twitter is usually classified alongside other disorders. E.g., Coppersmith et al. (2014, 2015a) achieved a precision of 0.64 at 10% false alarms, while Benton et al. (2017) used multi-task learning and achieved an AUC-score of 0.752.

Reddit has only recently been used as a source for the analysis of mental disorders. Gkotsis et al. (2016) analyzed the language in different subreddits related to mental health, and showed that linguistic features such as vocabulary use and sentence complexity vary across different subreddits. De Choudhury et al. (2016) explored the methods for automatic detection of individuals which could transit from mental health discourse to suicidal ideas. Shen and Rudzicz (2017) used topic modeling, LIWC, and language models to predict whether a Reddit post is related to anxiety. To our knowledge, there is no previous study on the analysis of bipolar disorder of Reddit users.

3 Dataset

Reddit is one of the largest social media sites in the world, with more than 85 million unique visitors per month.[1] Reddit is suitable for our study not only because of its vast volume, but also because it offers user anonimity and covers a wide range of topics. Registered users can anonymously discuss various topics on more than 1 million subpages, called "subreddits". A considerable number of subreddits is dedicated to mental health in general, and to bipolar disorder in particular. All comments between 2005 and 2018 (more than 3 billion) are available as a Reddit dump database via Google Big Query, which we used to obtain the data.

Bipolar disorder users. To obtain a sample of users with bipolar disorder, we first retrieved all subreddits related to the disorder, i.e., *bipolar*, *bipolar2*, *BipolarReddit*, *BipolarSOs*, *bipolarart*, as well as the more generic *mentalhealth* subreddit. Then, following Beller et al. (2014) and Coppersmith et al. (2014), we looked for self-reported bipolar users by searching in the user's comments for the string *"I am diagnosed with bipolar"* and its paraphrased versions. In addition, following Gjurković and Šnajder (2018), we inspect users' *flairs* – short descriptive texts that the users can set for certain subreddits to appear next to their names. While a flair is not mandatory, we found that many users with bipolar disorder do use flairs on mental health subreddits to indicate their disorder.

The acquisition procedure yielded a set of 4,619 unique users with self-reported bipolar disorder. The users generated around 5 million comments, totaling more than 163 million tokens. To get an estimate of labeling quality, we randomly sampled 250 users and inspected their labels and text. As we found no false positives (i.e., all 250 users report on being diagnosed a bipolar disorder), we gauge that the dataset is of high precision. The true precision of the dataset depends, of course, on the veracity of the self-reported diagnosis.

To make the subsequent analysis reliable and unbiased, we decided to additionally prune the dataset as follows. To mitigate the topic bias, we removed all comments by bipolar disorder users on bipolar subreddits, as well as on the general mental health subreddit. Additionally, any comment on any subreddit that mentions the words *bipolar* or *BP* (case insensitive) was also excluded. Finally, to increase the reliability, we retained in our dataset only the users who, after pruning, have at least 1000 word remaining. The final number of users in our dataset is 3,488.

[1] https://www.alexa.com/siteinfo/reddit.com

Category	# bipolar	# control
Animals	397	898
AskReddit	1797	2767
Gaming	489	1501
Jobs and finance	293	586
Movies/music/books	502	1606
Politics	332	2445
Religion	264	700
Sex and relationships	948	1000
Sports	156	785
All	3488	3931

Table 1: The number of unique bipolar disorder and control group users broken down by topic categories

Control group. The control group was sampled from the general Reddit community, serving as a representative of the mentally healthy population. To ensure that the topics discussed by the control group match those of bipolar disorder users, we sampled users that post in subreddits often visited by bipolar disorder users (i.e., subreddits where posting frequency of bipolar disorder users was above the average). To also ensure that the control group is representative of the mentally healthy Reddit population, we removed all users with more than 1000 words on mental health related subreddits. As before, we only retained users that had more than 1000 words in all of their comments. The final number of users in the control group is 3,931, which is close to the number of bipolar users, with the purpose of having a balanced dataset. The total number of comments is about 20 million, which is four times more than for the bipolar disorder users.

Topic categories. Topic of discussion may affect the language use, including the stylometric variables (Mikros and Argiri, 2007), which means that topic distribution may act as a confounder in our analysis. To minimize this effect, we split the dataset into nine topic categories, each consisting of a handful of subreddits on a similar topic. Table 1 shows the breakdown of the number of unique users from both groups across topic categories. *AskReddit* is the biggest subreddit and not bound to any particular topic; in this category, we also add other subreddits covering a wide range of topics, such as *CasualConversation* and *Showerthoughts*. To be included in a category, the user must have had at least 1000 words on subreddits from that category.

4 Bipolar Disorder Prediction

Feature extraction. For each user, we extracted three kinds of features: (1) psycholinguistic features, (2) lexical features, and (3) Reddit user features. For the psycholinguistic features, in line with much previous work, we used LIWC (Pennebaker et al., 2015), a widely used tool in predicting mental health, which classifies the words into dictionary-defined categories. We extracted 93 features, including syntactic features (e.g., pronouns, articles), topical features (e.g., work, friends), and psychological features (e.g., emotions, social context). In addition to LIWC, we used Empath (Fast et al., 2016), which is similar to LIWC but categorizes the words using similarities based on neural embeddings. We used the 200 predefined and manually curated categories, which Fast et al. have found to be highly correlated with LIWC categories (r=0.906).

The lexical features are the tf-idf weighted bag-of-words, stemmed using Porter stemmer from NLTK (Bird et al., 2009). Finally, Reddit user features are meant to model user's interaction patterns. These include post-comment ratio, the number of *gilded* posts (posts awarded with money by other users), average controversiality, the average difference between *ups* and *downs* on user's comments and the time intervals between comments (the mean, median, selected percentiles, and the mode).[2]

Experimental setup. We frame bipolar disorder prediction as a binary classification task, using the above-defined features and three classifiers: a support vector machine (SVM), logistic regression, and random forest ensemble (RF). We evaluated our models and tune the hyperparameters using 10×5 nested cross validation. To mitigate for class imbalance, we use class weighting when training classifiers on the dataset split into categories. As baselines, we used a majority class classifier (MCC) for evaluating the accuracy score and a random classifier with class priors estimated from the training set for evaluating the F1-score (F1-score is undefined for MCC). For implementation, we used Scikit-learn (Pedregosa et al., 2011). We use a two-sided t-test for all statistical significance tests and test at p<0.001 level.

Results. Table 2 shows the accuracy and F1-scores for the different classifiers. Random forest

[2]Users with bipolar disorder often experience sleep disturbance, which can make their usage patterns deviate from that of other users. Unfortunately, timestamps in Big Query are in UTC, not in users' local times, thus determining the time zone would require geolocalization. We leave this for future work.

	Acc	F1
MCC	0.529	–
Random	0.546	0.453
SVM	0.865	0.853
LogReg	0.866	0.849
RF	**0.869**	**0.863**

Table 2: Prediction accuracy and F1-scores

	LIWC	Empath	Tf-idf	All
SVM	0.837	0.782	**0.865**	0.838
LogReg	0.841	0.819	**0.866**	0.862
RF	0.829	0.825	0.869	**0.869**

Table 3: Prediction accuracy for the different models and feature sets

classifier achieved the best results, with accuracy of 0.869 and an F1-score of 0.863. All models outperform the baseline accuracies of 0.529 and 0.546, and the baseline F1-score of 0.453.

Table 3 shows the accuracy of the models using different feature sets. We observe two trends: Empath generally performs worse than LIWC, and tf-idf features perform better than LIWC. However, looking at the scores of the random forest classifier as the best model, we find that there is no significant difference between LIWC and Empath. Tf-idf does perform significantly different than both LIWC and Empath, while all features combined (including Reddit user features) do not differ from tf-idf alone. We speculate that tf-idf might yield better results in this case because essentially all the words that LIWC and Empath detect also exist as individual features in tf-idf. Similarly, Coppersmith et al. (2014) achieve better results using language models than LIWC, arguing that many relevant text signals go undetected by LIWC.

Finally, Table 4 shows the accuracy across topic categories for the MCC baseline and the best classifier in each category. Our models outperform MCC in all categories, and the differences are significant for all categories except *Sports*.

5 Feature Analysis

We analyze the merit of the psycholinguistic features using a two-sided t-test, with the null hypothesis of no difference in feature values between users with bipolar disorder and control users. The lower the p-value, the higher the merit. We analyzed the features separately on the entire dataset and on the dataset split into categories.

	MCC	Our models
Animals	0.693	0.807*
AskReddit	0.606	0.856*
Gaming	0.754	0.777*
Jobs and finance	0.665	0.752*
Movies/music/books	0.761	0.817*
Politics	0.880	0.882*
Religion	0.724	0.784*
Sex and relationships	0.513	0.801*
Sports	0.832	0.837

Table 4: Accuracy of the MCC baseline and our models across topic categories. Accuracies marked with "*" are significantly different from the baseline.

Between-group analysis. Ten LIWC features with the lowest p-value on the entire dataset are presented in Table 5, together with feature value means for the two groups. The values in the table are percentages of words in text from each category, except *Authentic* and *Clout*, which are "summary variables" devised by Pennebaker et al. (2015). Personal pronouns, especially the pronoun *I*, are used more often by bipolar disorder users. This observation is in accord with past studies on language of depressed people, which we can compare to because a bipolar depressive episode is almost identical to major depression (Anderson et al., 2012). Coppersmith et al. (2014) also report a significant difference in the use of *I* between Twitter users with bipolar disorder and the control group. The *Authentic* feature of Newman et al. (2003) reflects the authenticity of the author's text: a higher value of this feature in bipolar disorder users may perhaps be explained by them speaking about personal issues more sincerely, though further research would be required to confirm this. We also observe a higher use of words associated with feelings (*feel*), *health*, and biological processes (*bio*). Kacewicz et al. (2014) argue that pronoun use reflects standings in social hierarchies, expressed through *Clout* and *power* features: we observe a lower use of these words in users with bipolar disorder, which might indicate they think of themselves as less valuable members of society. The analysis of Empath features yielded similar findings: *health*, *contentment*, *affection*, *pain*, and *nervousness* have higher values in users with bipolar disorder.

Per-category analysis. Significant features in specific categories follow a pattern similar to the features on the complete dataset. Pronoun *I* is statistically significant in all of the categories, as well as features *Clout* and *Authentic*.

Feature	bipolar μ	control μ
Authentic	52.65	32.92
ppron	10.69	8.66
i	5.84	3.38
health	0.96	0.50
feel	0.69	0.48
power	2.11	2.58
pronoun	16.87	14.86
bio	2.65	1.90
Clout	48.51	58.03
article	5.88	6.55

Table 5: Mean values of most significant LIWC features for both groups

	Bipolar	Control
posemo	3.899 ± 1.02	3.442 ± 0.78
negemo	2.432 ± 0.67	2.569 ± 0.70
anxiety	0.367 ± 0.19	0.266 ± 0.10
anger	0.818 ± 0.39	1.128 ± 0.52
sad	0.455 ± 0.21	0.363 ± 0.11
affect	6.415 ± 1.22	6.074 ± 1.12

Table 6: Means and standard deviations of LIWC emotion categories for bipolar and control group

6 Emotion Analysis

As emotional swings are of the main symptoms of bipolar disorder, we expect that there will be a difference in the use of emotion words between users with bipolar disorder and the control group. We report the results for LIWC, as Empath gave very similar results.

Between-group differences. Table 6 shows means and standard deviations of the values of six LIWC emotion categories (*posemo*, *negemo*, *anxiety*, *anger*, *sad*, and *affect*) for the users with bipolar disorder and the control group. Users with bipolar disorder use significantly more words linked with general affect. Furthermore, we observe increased use of words related to sadness, while the control group uses more anger-related words. The results for *sadness* are in line with previous work on depressed authors. In addition, we find significant use of *anxiety* words in users with bipolar disorder, similar to the findings of Coppersmith et al. (2014). Surprisingly, we find that users with bipolar disorder use more positive emotion words than the control group. This is in contrast to findings of Rude et al. (2004), who report no statistical significance in the use of positive emotion words in depressed authors. We speculate that this difference may be due to the characteristics of manic episodes, which do not occur in clinically depressed people.

	Bipolar	Control	p-value
posemo	0.00272*	0.00166	0.00272
negemo	0.00583*	0.00379	0.00583
anxiety	0.00765*	0.00627	0.00765
anger	0.01745	0.01422	0.01745
sadness	0.00695*	0.00572	0.00695

Table 7: Averages of standard deviations in the use of emotion-expressive words for the two groups. All differences are significant except for “anger”.

Per-category differences. The difference between users with bipolar disorders and the control group in *AskReddit*, *Animals*, *Movies/music/books*, and *Sex and relationships* categories is significant in words related to sadness, anxiety, anger, and positive emotions. However, there is no significant difference in positive and negative emotions in categories *Jobs* and *Politics*, while *Sports*, *Gaming*, and *Religion* differ only in positive emotions.

User-level variance. We hypothesize that, due to the alternation of manic and depressive episodes, users with bipolar disorder will have a higher variance across time in the use of emotion words than users from control group. To verify this, we randomly sampled 100 users with bipolar disorder and 100 control users from all the users in our dataset with more than 100K words and split their comments into monthly chunks. For each of the 200 users, we calculated the LIWC features for each month and computed their standard deviations. We then measured the difference between standard deviations for the two groups. Table 7 shows the results. We find that bipolar users have significantly more variance in most emotion-expressive words, which confirms our hypothesis.

7 Conclusion

We presented a preliminary study on bipolar disorder prediction from user comments on Reddit. Our classifiers outperform the baselines and reach accuracy and F1-scores of above 86%. Feature analysis suggests that users with bipolar disorder use more first-person pronouns and words associated with feelings. They also use more affective words, words related to sadness and anxiety, but also more positive words, which may be explained by the alternating episodes. There is also a higher variance in emotion words across time in users with bipolar disorder. Future work might look into the linguistic differences in manic and depressive episodes, and propose models for predicting them.

References

Ian M. Anderson, Peter M. Haddad, and Jan Scott. 2012. Bipolar disorder. *BMJ: British Medical Journal (Online)*, 345.

Charley Beller, Rebecca Knowles, Craig Harman, Shane Bergsma, Margaret Mitchell, and Benjamin Van Durme. 2014. Ima belieber: Social roles via self-identification and conceptual attributes. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 181–186.

Adrian Benton, Margaret Mitchell, and Dirk Hovy. 2017. Multi-task learning for mental health using social media text. *arXiv preprint arXiv:1712.03538*.

Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. " O'Reilly Media, Inc.".

Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. Quantifying mental health signals in Twitter. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 51–60.

Glen Coppersmith, Mark Dredze, Craig Harman, and Kristy Hollingshead. 2015a. From ADHD to SAD: Analyzing the language of mental health on Twitter through self-reported diagnoses. In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 1–10.

Glen Coppersmith, Mark Dredze, Craig Harman, Kristy Hollingshead, and Margaret Mitchell. 2015b. CLPsych 2015 shared task: Depression and PTSD on Twitter. In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 31–39.

Munmun De Choudhury, Scott Counts, and Eric Horvitz. 2013. Social media as a measurement tool of depression in populations. In *Proceedings of the 5th Annual ACM Web Science Conference*, pages 47–56. ACM.

Munmun De Choudhury, Emre Kiciman, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, pages 2098–2110. ACM.

Ethan Fast, Binbin Chen, and Michael S Bernstein. 2016. Empath: Understanding topic signals in large-scale text. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4647–4657. ACM.

Matej Gjurković and Jan Šnajder. 2018. Reddit: A gold mine for personality prediction. In *Proceedings of the Second Workshop on Computational Modeling of Peoples Opinions, Personality, and Emotions in Social Media*, pages 87–97.

George Gkotsis, Anika Oellrich, Tim Hubbard, Richard Dobson, Maria Liakata, Sumithra Velupillai, and Rina Dutta. 2016. The language of mental health problems in social media. In *Proceedings of the Third Workshop on Computational Lingusitics and Clinical Psychology*, pages 63–73.

Kishaloy Halder, Lahari Poddar, and Min-Yen Kan. 2017. Modeling temporal progression of emotional status in mental health forum: A recurrent neural net approach. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 127–135.

Ewa Kacewicz, James W. Pennebaker, Matthew Davis, Moongee Jeon, and Arthur C. Graesser. 2014. Pronoun use reflects standings in social hierarchies. *Journal of Language and Social Psychology*, 33(2):125–143.

Adam DI Kramer, Susan R. Fussell, and Leslie D. Setlock. 2004. Text Analysis as a tool for analyzing conversation in online support groups. In *CHI'04 Extended Abstracts on Human Factors in Computing Systems*, pages 1485–1488. ACM.

George K. Mikros and Eleni K. Argiri. 2007. Investigating topic influence in authorship attribution. In *PAN*.

Matthew L. Newman, James W. Pennebaker, Diane S. Berry, and Jane M. Richards. 2003. Lying words: Predicting deception from linguistic styles. *Personality and social psychology bulletin*, 29(5):665–675.

Merete Nordentoft, Preben Bo Mortensen, et al. 2011. Absolute risk of suicide after first hospital contact in mental disorder. *Archives of general psychiatry*, 68(10):1058–1064.

Minsu Park, Chiyoung Cha, and Meeyoung Cha. 2012. Depressive moods of users portrayed in Twitter. In *Proceedings of the ACM SIGKDD Workshop on healthcare informatics (HI-KDD)*, volume 2012, pages 1–8. ACM New York, NY.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, and Vincent Dubourg. 2011. Scikit-learn: Machine learning in Python. *Journal of machine learning research*, 12(Oct):2825–2830.

James W. Pennebaker, Ryan L. Boyd, Kayla Jordan, and Kate Blackburn. 2015. The development and psychometric properties of LIWC2015. Technical report.

James W. Pennebaker, Martha E. Francis, and Roger J. Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates*, 71(2001):2001.

James W. Pennebaker and Laura A. King. 1999. Linguistic styles: Language use as an individual difference. *Journal of personality and social psychology*, 77(6):1296.

Stephanie Rude, Eva-Maria Gortner, and James Pennebaker. 2004. Language use of depressed and depression-vulnerable college students. *Cognition & Emotion*, 18(8):1121–1133.

Shekhar Saxena, Michelle Funk, and Dan Chisholm. 2013. World health assembly adopts comprehensive mental health action plan 2013–2020. *The Lancet*, 381(9882):1970–1971.

Paula P. Schnurr, Stanley D. Rosenberg, and Thomas E. Oxman. 1992. Comparison of TAT and free speech techniques for eliciting source material in computerized content analysis. *Journal of personality assessment*, 58(2):311–325.

Judy Hanwen Shen and Frank Rudzicz. 2017. Detecting anxiety on Reddit. In *Proceedings of the Fourth Workshop on Computational Linguistics and Clinical Psychology—From Linguistic Signal to Clinical Reality*, pages 58–65.

Shannon Wiltsey Stirman and James W. Pennebaker. 2001. Word use in the poetry of suicidal and nonsuicidal poets. *Psychosomatic medicine*, 63(4):517–522.

Allison M. Tackman, David A. Sbarra, Angela L. Carey, M. Brent Donnellan, Andrea B. Horn, Nicholas S. Holtzman, To'Meisha S. Edwards, James W. Pennebaker, and Matthias R. Mehl. 2018. Depression, negative emotionality, and self-referential language: A multi-lab, multi-measure, and multi-language-task research synthesis. *Journal of personality and social psychology*.

Michael Alan Taylor, Robyn Reed, and Sheri A Berenbaum. 1994. Patterns of speech disorders in schizophrenia and mania. *Journal of Nervous and Mental Disease*.

Caroline B. Thomas and Karen R. Duszynski. 1985. Are words of the Rorschach predictors of disease and death? The case of "whirling.". *Psychosomatic medicine*.

Til Wykes, Josep Maria Haro, Stefano R. Belli, Carla Obradors-Tarragó, Celso Arango, José Luis Ayuso-Mateos, István Bitter, Matthias Brunn, Karine Chevreul, and Jacques Demotes-Mainard. 2015. Mental health research priorities for Europe. *The Lancet Psychiatry*, 2(11):1036–1042.