Machine learning of language use on Twitter reveals weak and non-specific predictions

Sean W. Kelley^{*1,2}, Caoimhe Ní Mhaonaigh¹, Louise Burke¹, Robert Whelan^{1,2,3}, and Claire M. Gillan^{1,2,3}

¹School of Psychology, Trinity College Dublin, Dublin, Ireland

²Trinity College Institute of Neuroscience, Trinity College Dublin, Dublin, Ireland

³Global Brain Health Institute, Trinity College Dublin, Dublin, Ireland

*Correspondence: sekelley@tcd.ie

Abstract

Background: Depressed individuals use language differently than healthy controls and it has been proposed that social media posts could therefore be used to identify depression. But much of the evidence behind this claim relies on indirect measures of mental health that are sometimes circular, such as statements of self-diagnosis (“Got an OCD diagnosis today”) on social media or membership in disorder-specific online forums. Relatedly, few studies have tested if these language features are specific to depression versus other aspects of mental health.

Methods: We analyzed the Tweets of 1,006 participants who completed questionnaires assessing symptoms of depression and 8 other mental health conditions. Daily Tweets were subjected to textual analysis and the resulting linguistic features were used to train an Elastic Net model on depression severity, using nested cross validation. We then tested performance in a held-out test set (30%), comparing predictions of depression versus 8 other aspects of mental health.

Results: The depression trained model had only modest predictive performance when tested out of sample, explaining just 2.5% of variance in depression symptoms ($R^2 = 0.025$, $r = 0.16$). The performance of this model was as-good or superior when used to identify other aspects of mental health: schizotypy ($R^2 = 0.035$, $r = 0.19$), social anxiety ($R^2 = 0.025$, $r = 0.16$), eating disorders ($R^2 = 0.011$, $r = 0.12$), generalized anxiety ($R^2 = 0.041$, $r = 0.21$), above chance for obsessive-compulsive disorder ($R^2 = 0.011$, $r = 0.12$),

45   apathy ($R^2 = 0.008$, $r = 0.11$), but not significant for alcohol abuse ($R^2 = -0.012$, $r = 0.04$)
46   or impulsivity ($R^2 = -0.001$, $r = 0.08$).

47   **Conclusions:** Machine learning analysis of social media data, when trained on well-
48   validated clinical instruments, could not make meaningful individualized out-of-sample
49   predictions regarding mental health status of users. For the small effects observed,
50   language use associated with depression was non-specific, having similar performance
51   in predicting other mental health problems.

52

53

Introduction

Background

Approximately 20% of adults will experience a mental illness in any given year[1]. But our ability to treat those affected is hampered by the fact that patients present to clinics relatively infrequently[2], and when they do so, it is often belated, making their symptoms more difficult to treat[3]. For this reason, efforts to detect mental illness early and predict individual vulnerability is a key focus of research. Of course, this is challenging in real-world settings, because it is unclear what sources of data can and should be utilized to make these predictions. It has been proposed that one way to overcome this difficulty is to use other sources of data that the general public produce regularly, such as social media data, to detect, predict and better understand mental health in the population. Social media adoption is widespread with approximately 72% of U.S. adults using at least 1 social media platform offering a unique opportunity for gathering information about mental health[4].

Recent studies have suggested that social media data can be used to recognise a broad range of mental health problems in the general public including depression[5-10], eating disorders[11-14], schizophrenia[15-17], and suicide[18-21]. A key premise of such work is that these data could be used to facilitate early intervention, for example by providing users with personalised risk scores for having a mental illness and/or developing one in the near future. Inherent in that is the assumption that such models are (i) accurate enough to be clinically actionable and (ii) precise enough to detect one illness from another. In the present paper, we investigated the extent to which models based on social media data meet these criteria.

Specificity of Language Use in Mental Health Conditions

There is now a wealth of data supporting the notion that people with depression use language differently than those without depression. For example, depressed individuals use more first person singular pronouns [22], obscenities [6] and express more negative emotions [5] in their language. This language occurs in a variety of settings including semi-structured interviews [23, 24], journal entries [22, 25], and critically, social media posts [7, 9, 26]. However, it is not clear if these language patterns are specific to depression. Shared variance between disorders presents a challenge to identifying what aspects of language use are specific to a particular disorder. Indeed, because mental health disorders tend to co-occur in the same individuals and our existing diagnostic system lacks clear separation between disorders[27, 28], language-based models are unlikely to have high specificity when trained on summed scores or diagnostic categories. Studies typically compare data from depressed individuals to that of healthy controls, but do not test if language patterns discriminate among psychiatric disorders. There are few clear distinguishing features [29-32] among multiple mental health conditions the few papers that have studied multiple groups. For example, although elevated first person singular pronoun usage is considered to be a defining feature of language in depression, – reflecting an increase in self-focused attention – it is also elevated in people with obsessive-compulsive disorder [30, 31], anxiety [30, 31], eating disorders [11-13, 30, 33], and schizophrenia [15, 16, 30, 34, 35]. Without accounting for comorbidity among disorders, it is not possible to discern whether first person singular pronouns are unique to depression, a transdiagnostic marker of mental illness, or specific to another aspect of mental health entirely. Tackling the issue of specificity

more directly, one study found greater evidence of third person plural pronouns (they, them) in those who participated in Schizophrenia discussion forums versus other sorts of mental health forums, a putative marker of persecutory delusions [31]. However, the topic of the discussion forum from which language-use was gathered is a major confounding factor. That is, the content in these forums may not reflect speech patterns of persons with schizophrenia in their everyday life, when not discussing their illness. Moreover, no clinically validated screening tools were used to define cases, rather, participation in these forums and explicit statements of self-diagnosis were used to identify patients.

Defining Cases of Mental Illness in Social Media

As discussed in a recent review[36], this is a common approach and has been applied to study language-use on more generic social media outlets like Twitter, using 'statements of diagnosis' e.g., "I have PTSD" [29, 32], to define cases of mental illness, rather than validated clinical instruments. In addition to the issue of diagnostic validity, this approach is limited by the fact that a person who openly reveals a diagnosis on Twitter is not someone trying to conceal it and is probably not part of the cohort of undiagnosed/untreated individuals that such methods may wish to identify. Moreover, they may be more likely to tweet about disorder-relevant topics, which could create circularity, inflating effect sizes, leading us to conclude that mental health status is more readily detected from social media data than it actually is. For example, Coppersmith et al., reported 85% precision at detecting generalized anxiety disorder on Twitter when allowing for false positive rate of 10% [29] using this method. To remove these potential circularities, studies are moving toward less biased methods, where mental health

status is not defined by the same or similar content that is ultimately used to study language-use. For example, separating the content used to define disorder status (e.g. membership of a mental health forum) from content used to characterize language use (posts by those users on other forums) [37]. When this approach was taken, accuracy was substantially worse. Using a range of machine learning algorithms, the F1 score (average of precision and recall) rarely exceeded 0.5 [30]. This poor performance could be because the signal is weak, or the diagnoses are not accurate. However, even in studies with a more rigorous definition of disorder status because the cases are still binary, evaluating how specific language use is to that disorder and not another condition remains an issue. This is due to the lack of multiple continuous measures of mental health in the sample participants. In the few studies that have administered self-report questionnaires to consenting participants, performance was again modest, albeit somewhat improved [9], but the specificity of the findings to the disorder of study (here, depression) was not examined.

The Present Study

Thus, while social media makes substantial amounts of language data available to researchers, a caveat to much of this research is the acquisition of high-quality mental health data. The notion that we can detect mental illness from social media posts presents opportunities for novel public health interventions, but with this comes significant privacy concerns and potential for discriminatory practices to emerge[38]. But are these opportunities and concerns overstated? To date there is little evidence that mental health status can be detected accurately, and even less evidence for individuals who do not choose to openly disclose/discuss mental health online disorder status

online. Moreover, if such predictions could be made with any fidelity, it is unclear if they can be in any way specific, which is crucial if these indicators are to be used to guide the choice of intervention. The present study sought to address these issues, determining (i) the specificity of language patterns to different aspects of mental health and (ii) providing an estimate of the performance of these models on unseen data. To do this, we acquired Twitter data over the past year from over 1000 individuals who completed 9 different self-report mental health related questionnaires and consented for us to link that to their Tweets. We tested the performance of a machine learning algorithm trained on a gold-standard ground truth measure of depression symptomatology when applied to unseen data. To test its specificity, we then applied this depression model to predicting scores of a range of mental health phenotypes. Finally, we trained a machine learning model on to predict the residuals of 3 transdiagnostic dimensions of mental health (after controlling for one-another), allowing us to identify text features are specific to a particular dimension after removing the shared variance between disorders.

Methods

Participants

We recruited 1,450 participants for this study. The majority of participants were recruited on Clickworker (N = 1,395), an online worker platform, and were paid €2.5 for their participation. A smaller number participated voluntarily (i.e., without payment) and were recruited through general advertising on Twitter and in print media (N = 55). Participants were included for analysis if they were at least 18 years old and had a Twitter account with at least 5 days of tweets and if at least 50% of their tweets were in English. They were also required to pass an attention check, a combination of a captcha and an item with an obvious correct response (“Please select ‘A little’ if you are paying attention”). Of the 1,450 participants recruited, 99 were excluded due to failing the attention check and a further 345 participants were excluded for either not having at least 5 days of tweets or fewer than 50% of their tweets were in English. After excluding these participants, 1,006 participants were brought forward for analysis. Participants had a mean age of 30.5 years (SD: 10.1, range: 18-68), a majority were female (66.4%), currently employed (63.8%), and resided in either the U.K. (41%) or U.S. (46.9%). Participants tweeted an average of 21,126 words (SD: 30,204), median of 6,432 words, and a range of 43 – 163,700. In total, there were 21,252,845 words posted across the 1,006 participants.

Twitter Behaviour		Mean (SD)
<i>No. of Tweets</i>		231.3 (489.4)
<i>No. of Retweets</i>		173.5 (416)
<i>No. of Likes</i>		768.4 (1041.3)
<i>Word count per day</i>		96.5 (119.9)
Age		Mean (SD)

<i>Years</i>	30.5 (10.1)
Gender	N (%)
<i>Male</i>	312 (31%)
<i>Female</i>	668 (66.4%)
<i>Transgender Male</i>	6 (0.6%)
<i>Transgender Female</i>	1 (0.1%)
<i>Non-Binary</i>	16 (1.6%)
<i>Other</i>	3 (0.3%)
Country	N (%)
<i>Ireland</i>	32 (3.2%)
<i>United Kingdom</i>	412 (41%)
<i>United States</i>	472 (46.9%)
<i>Canada</i>	52 (5.2%)
<i>Australia</i>	24 (2.4%)
<i>Other</i>	14 (1.4%)
Education	N (%)
<i>Less than high school</i>	22 (2.2%)
<i>High School</i>	220 (21.9%)
<i>Some University</i>	301 (29.9%)
<i>Bachelor's degree</i>	325 (32.3%)
<i>Master's degree</i>	108 (10.7%)
<i>Professional degree</i>	17 (1.7%)
<i>Doctorate</i>	13 (1.3%)
Employment Status	N (%)
<i>Currently Employed</i>	642 (63.8%)

Table 1: Twitter Use and Demographics of Sample

Procedure

After providing informed consent, participants were asked to complete a self-report questionnaire and provide their Twitter handle which was used to collect the most recent (max. 3,200) tweets and (max. 3,200) likes from their account. Tweets were collected using a data collection app written in Python using the Twitter developer's

Application Programming Interface. Participants were asked to provide their age, gender, country of residence, current employment status, and highest educational attainment. Participants then completed 9 different psychiatric questionnaires including the Zung depression scale (SDS)[39], Short Scales for Measuring Schizotypy (SSMS)[40], Obsessive Compulsive Inventory Revised (OCI-R)[41], Eating Attitudes Test (EAT-26)[42], Barratt Impulsiveness Scale (BIS-11)[43], Alcohol Use Disorders Inventory Test (AUDIT)[44], Apathy Evaluation Scale (AES)[45], Liebowitz Social Anxiety Scale (LSAS)[46], State-Trait Anxiety Inventory (STAI)[47] (Figure S1).

Pre-processing and Text Analysis

We restricted our analysis to tweets published in the 12 months prior to survey completion. Before text analysis, extraneous information was removed from tweets including reply symbol (@), hashtag symbol (#), emojis, punctuation, links (URLs), and all other non-alphanumeric characters. Periods, exclamation points, and question marks were the only punctuation retained because they are necessary to calculate the number of words per sentence. Tweets were aggregated into daily bins and text analysis was then performed on all tweets published per day per user. Daily observations were chosen to increase the amount of text for reliable estimation of text features. Text analysis of daily Tweets was carried out using the Linguistic Inquiry and Word Count (LIWC 2015) dictionary [48]. The LIWC is a dictionary comprised of approximately 6,400 words and word-stems with 90 different output variables including: linguistic characteristics (e.g., articles and pronouns), psychological constructs (e.g., sadness and positive emotions), and general text information (e.g., punctuation and word count).

211 In addition to text features, we carried out some additional analyses using Twitter
212 metadata variables including number of followees and followers, replies per day,
213 number of tweets per day, and the insomnia index. The insomnia index is the relative
214 difference in percentage of tweets tweeted during the day (6:01AM to 8:59PM) versus
215 the night (9PM to 6AM). Previous research has found that people with depression tend
216 to tweet more at night [26, 49].

217 **Univariate Associations with Mental Health Symptoms**

218 We focus on depression at the outset because it is the most commonly studied disorder
219 in this field of research and therefore several benchmark studies exist. To examine the
220 specificity of text features to depression, in the first instance, we report univariate
221 associations between the total score for each psychiatric disorder and the top 10 text
222 features associated with depression severity including: word count, negative emotions,
223 focus on present, verbs, auxiliary verbs, adverbs, tone, analytic, six letter words, and
224 leisure words. Each linear model contained just one text feature and controlled for the
225 effects of both age and gender (e.g., depression ~ adverbs + age + gender), both of
226 which showed associations with mental health symptoms consistent with prior work
227 (Figure S2).

228 **Machine Learning**

229 Next, we trained a model to predict depression scores from LIWC text features using
230 Elastic Net regularization[50]. Elastic Net is a combination of L-1 and L-2 norm
231 regularization, performing both feature selection and regularization which results in a
232 sparse solution when features are correlated with each other. We chose Elastic Net

233 because the input text features are highly correlated and has been shown to make
234 accurate prediction with small effect sizes in samples larger than 400[51]. Another
235 advantage of Elastic Net is that the output (regression coefficients) is easily
236 interpretable. It is thus possible to directly compare the relative importance of input
237 features and see how they contribute to predictive performance.

238 We tested the model's performance in predicting out-of-sample depression scores, and
239 to assess specificity, we also tested it on out-of-sample scores on 8 other psychiatric
240 scales (which we not using in training). The data was split into training (70%) and test
241 (30%) sets, stratified by gender to ensure equal proportions of gender categories
242 between the two sets. Nested cross-validation was performed within the training data,
243 using 10 outer loops stratified by gender and 5 inner loops, with optimization of elastic
244 net hyperparameters (alpha and the l1 ratio) within the inner loops. We repeated this
245 process 100 times to select the best elastic net model to take forward to test on the 30%
246 of data we held-out. We first tested it on depression scores to determine predictive
247 power and then on the 8 other psychiatric scales. Random label permutation was used
248 to determine the predictive value of these models. To ensure that the apparent
249 predictive power of LIWC was not simply the result of confounded associations between
250 age, gender, mental health, and language use. We preformed control analyses that
251 included: LIWC text features plus age and gender as features and compared it to a
252 model with just age and gender (and randomly permuted LIWC features). This allowed
253 us to determine if there was a marginal benefit of text features above and beyond the
254 predictive ability of basic demographics[52]. Finally, we sought to identify reasons why
255 our model made poor predictions for many individuals. To do this, we examined

256 depression residuals and tested if these related to how participants use and interact
257 with Twitter. We associated depression residuals from the held-out test set of the LIWC
258 text feature only model with the z-score of i) mean word count, ii) total number of
259 Tweets, iii) tweet volume, iv) total number of replies, v) number of followers, and vi)
260 number of followees.

261 Selecting a machine learning method based in its performance is a form of overfitting.
262 For this reason, we chose to work with the Elastic Net *a priori*, a method well suited to
263 continuous prediction problems. However, for the purposes of comparison to other
264 studies, and to ensure our results are not specific to the Elastic Net, we repeated our
265 analysis pipeline using two alternative classification models. AUC, is the primary
266 measure of predictive performance in most studies in the literature, so to enable
267 comparisons between our results and those of previous studies, we binarised
268 depression scores and classified participants as either depressed or non-depressed.
269 Participants with depression scores above 50 were classed as depressed, while those
270 with scores under 50 were non-depressed[39]. We restricted our analysis to depression
271 because that is the primary disorder of interest in the current study and several
272 questionnaires used do not have established clinical cutoffs. Consequently, it would not
273 be possible to directly apply the depression trained model to the other conditions as we
274 did for the Elastic Net model. We next ran four separate models based on varying
275 combinations of sample size and model type to assess performance. We used two
276 classification models: Random Forest (RF) and Support Vector Machine (SVM) with a
277 radial kernel. Both models were validated with 10-fold cross validation and 100
278 experimental runs using either i) the full sample of participants (n = 1,006) or ii) the top

476 participants by word count. We chose the second sample size to be the same as that from de Choudhury et al., (2013)[5], so we could make a direct comparison between our classification performances. Our sample had fewer mean posts over the previous year compared to that of de Choudhury et al., (2013), a mean of 4,533.4 posts per user compared to 1,173.2 posts per user. By selecting the top 476 users by word count the number of posts per user increased to 2,277.1, thereby making the samples more comparable.

Comparison of Validated Self-Report vs Twitter-Derived Ground Truth

We compared the performance of a model trained self-reported depression versus depression keywords extracted from Twitter. Using a regular expression, i.e., depress*, we identified any posts that contained depression relevant keywords and phrases. We found that approximately 2.0% of all days with Tweets had at least 1 keyword matching the depression regular expression. We then classified participants as either depressed or not depressed based on whether they had at least 1 post with a depression keyword present; with this approach, 44% of participants were classified as depressed. Days with Tweets that contained the depression keyword were omitted to ensure independence between the testing and training data. Subsequently, we trained a random forest classification model on the depression keyword outcome and compared it to the model train on binarized self-reported depression. Finally, we evaluated all our models' performance using AUC, F1 score, accuracy, sensitivity, and specificity.

Transdiagnostic Psychiatric Dimensions

To test if performance might be more specific when using transdiagnostic psychiatric dimensions, rather than these questionnaire total scores, the individual answers to the 209 questions in our survey were transformed into 3 transdiagnostic dimensions. These were dimensions previously identified using factor analysis of these 9 questionnaires[53], corresponding to ‘anxious depression’, ‘compulsivity and intrusive thoughts’, and ‘social withdrawal’. We then used the weights derived from that independent study to construct the 3 transdiagnostic dimensions in this study. The transdiagnostic dimensions are designed to reduce collinearity across these questionnaires and have been shown to relate to cognitive test performance and brain signatures in a stronger and more specific manner than the original questionnaire total scores[53-55].

Comparison of Tweets, Retweets, Likes

We also ran additional analyses to understand the influence of Tweet type i.e., Tweet, Retweet, Like, and the amount of Twitter data on model performance. We trained depression models using text features that included only Tweets, Retweets, or Likes and then tested the models on all 9 questionnaire total scores. By training models separately on each type of Twitter data, we could determine whether each Tweet type is independently predictive of depression. We subsequently split the text feature data containing, with Tweets, Retweets, and Likes merged together, into quartiles based on the total number of Tweets. Then, we trained a depression model on data from each quartile and determined its predictive performance. Splitting the data into quartiles also allowed us to control for sample size, such that any differences in performance are solely caused by amount of text data. We expected that models trained on data from the

upper quartiles (with the most twitter data) would have better performance than the lower quartiles.

Similarity and specificity of text features across mental health phenotypes

To probe specificity in more detail and with less of a central focus on depression, we generated predictive models for each psychiatric questionnaire separately, using the procedure outlined above. The Elastic Net can assign a weight of 0 to variables that are not predictive. To better understand our results, we used feature selection derived from the Elastic Net regularization by summing each of the text feature's non-zero weights within each main fold and then averaged this value over the 100 iterations. The 'selection frequency' provides a useful heuristic for the importance of each text feature as a predictor of that clinical phenotype. Selection frequency is a good measure of a variable's importance since text features that are frequently included are likely to be a truly associated with the target outcome. After generating text feature selection frequencies for each questionnaire, we applied a hierarchical clustering algorithm, using Ward's method, to determine how similar language use is between mental health conditions[56]. To examine potential for specificity in a highly correlated space, we trained an Elastic Net model on i) each of the 3 transdiagnostic dimensions and ii) the residuals of each transdiagnostic dimension after controlling for the other 2 dimensions. That is, the anxious-depression residual is derived from a linear model as follows:
$$\text{anxious-depression} \sim \text{compulsivity} + \text{social withdrawal}.$$
We then used the selection frequencies in the same manner as above; to identify text features that were predictive of the residual and if so, whether or not that text feature was specific to that dimension, after removing the shared variance with the other dimensions.

The effect of number of words per user on predictive performance

As a control analysis, we tested the effect of the minimum number of words per user on predictive performance. For our main analyses we chose a relatively low word count threshold per user to maximize our sample size, including participants with at least 5 days of Tweets [9]. However, there is evidence recommending that a minimum of 200 words per user be used in order to achieve stable predictive performance[57]. While several other studies have used a minimum of 500 words per user[8, 58]. We thus tried three additional minimum word count per user threshold for inclusion: 200, 400, and 500 words per user. Sample size was not substantially affected by the additional inclusion criterion such that excluding participants with 200 words reduced the sample size to 945 participants, 400 words reduced the sample to 866 participants, and 500 words to 836 participants. However, to ensure that sample size differences across these minimum-word thresholds did not affect our results, we down-sampled our data to smallest sample size of 836 participants from the 500-word threshold and carried out all analyses on these subjects.

Statistical Power

Finally, we tested if we were sufficiently powered to find an effect size greater than the reported value. To interrogate this possibility, we simulated three types of datasets with 99 input features and 1 continuous target outcome with a sample size of either 1,000 or 3,000. We set the correlation between either 1, 10, or 20 features with the target variable at $r = 0.32$, while the other input variables had no association with the target. We chose to set $r = 0.32$, because that is approximately twice the observed effect size obtained from our depression model. Furthermore, for the datasets with greater than 1

feature associated with the target, we simulated multicollinearity among the relevant features by setting the correlation between those features at $r = 0.50$. We then ran each dataset through our Elastic Net analysis pipeline, 10-fold nested cross validation with 100 experimental runs and reported both R^2 and the mean absolute error as measure of model fit. To test for the likelihood that the true effect size is larger than we report (i.e., that it is in fact $r=.32$) and we are missing it, we plot the proportion of cases in the 1,000-person sample that performed worse than our reported predictive performance (Figure S3).

Results

Age was significantly negatively associated with all psychiatric questionnaires, except alcohol abuse (all $\beta < 0.07$, $p < 0.05$). Female participants had significantly elevated eating disorder symptoms ($\beta = 0.34$, $SE = 0.07$, $p < 0.001$), social anxiety ($\beta = 0.38$, $SE = 0.07$, $p < 0.001$), generalised anxiety ($\beta = 0.28$, $SE = 0.07$, $p < 0.001$), and depression ($\beta = 0.35$, $SE = 0.07$, $p < 0.001$) than men. Male participants had a significantly higher rate of alcohol abuse symptoms ($\beta = 0.31$, $SE = 0.01$, $p < 0.001$) (Figure S2). As expected, all psychiatric questionnaires were positively correlated with each other (Figure S4).

Univariate Associations with Mental Health Symptoms

The top ten text features associated with depression were word count, negative emotions, focus on present, verbs, adverbs, auxiliary verbs (all positively associated with depression severity, $\beta > 0.08$, $p < 0.05$) and tone, the analytic summary variable, number of six letter words, and leisure words (all negatively associated with depression

severity, $\beta < -0.07$, $p < 0.05$). These effects were non-specific. Negative emotions (all $\beta > 0.08$, $p < 0.001$) were significantly positively associated with all aspects of mental health studied, except alcohol abuse ($\beta = 0.05$, $SE = 0.03$, $p = 0.05$) and obsessive-compulsive disorder ($\beta = 0.04$, $SE = 0.03$, $p = 0.11$). Schizotypy, social anxiety, and generalised anxiety were significantly associated with all 10 text features (all $\beta > |0.06|$, $p < 0.05$), except for the associations between social anxiety with tone ($\beta = -0.03$, $SE = 0.03$, $p = 0.33$) and leisure ($\beta = -0.04$, $SE = 0.03$, $p = 0.12$), which were non-significant. None of the alternative questionnaires were significantly associated with a text feature in the opposite direction of depression. Individual text features were thus not specific to depression but broadly associated with other psychiatric dimensions (Figure 1).

Figure 1: Associations between 9 self-reported psychiatric questionnaires and mean values over the past year of the top 10 LIWC text features associated with depression severity, controlling for age and gender. Dashed lines indicate p-values below 0.05.

In terms of Twitter metadata, participants with elevated obsessive-compulsive symptoms followed more accounts ($\beta = 0.03$, $SE = 0.01$, $p = 0.01$), while participants who scored higher on eating disorder severity had a larger number of followers ($\beta = 0.02$, $SE = 0.01$, $p = 0.02$). Participants scoring high on depression, apathy, impulsivity, obsessive-compulsive disorder, and schizotypy tended to tweet more at night i.e. higher insomnia index (all $\beta < -0.06$, $p < 0.05$). Replies to Tweets (all $\beta > 0.08$, $p < 0.05$) and volume of Tweets (all $\beta > 0.03$, $p < 0.05$) were positively associated with all aspects of mental health recorded, except apathy and eating disorders (Figure 2).

Figure 2: Associations between 9 self-reported psychiatric questionnaires and Twitter meta-data features.

Replies (all $\beta > 0.08$, $p < 0.05$) and volume (all $\beta > 0.03$, $p < 0.05$) of tweets were significantly elevated across all aspects of mental health studied, except for alcohol abuse and eating disorders. Participants with elevated obsessive-compulsive symptomology tended to follow more accounts ($\beta = 0.03$, $SE = 0.01$, $p = 0.01$). While participants with more eating disorder symptoms had significantly

more followers ($\beta = 0.02$, $SE = 0.01$, $p = 0.02$). Participants with greater depression, apathy, impulsivity, obsessive-compulsive, and schizotypy symptoms tweeted more at night than during the day (all $\beta < -0.06$, $p < 0.05$).

Dashed lines indicate p-values below 0.05.

Machine Learning

We trained an Elastic Net model on depression symptoms and tested it on unseen data. The model of depression symptomatology had an R^2 of 0.025 ($r = 0.16$) vs. $R^2 -0.040$ ($r = -0.16$) for the null model. A model trained on text features plus age and gender ($R^2 = 0.045$, $r = 0.22$) performed better than a model with randomised text features plus age and gender ($R^2 = 0.039$, $r = 0.20$). Our simulation results also demonstrated that we were sufficiently powered to detect a larger signal, if it was truly present (see Supplementary Material).

After establishing there was modest, but non-zero signal, we applied the depression trained model using LIWC text features only to the other 8 psychiatric scales to test for specificity (Figure 3). The depression model had above zero predictive power for all other aspects of mental health except impulsivity and alcohol abuse. Nominally, the depression model performed somewhat worse when tested on apathy ($R^2 = 0.008$, $r = 0.11$), alcohol abuse ($R^2 = -0.012$, $r = 0.04$), eating disorder symptoms ($R^2 = 0.011$, $r = 0.12$), and obsessive-compulsive disorder symptoms ($R^2 = 0.011$, $r = 0.12$), predictive ability was identical to depression for social anxiety ($R^2 = 0.025$, $r = 0.16$) and the model performed nominally better in predicting schizotypy ($R^2 = 0.036$, $r = 0.19$) and generalized anxiety ($R^2 = 0.041$, $r = 0.21$) scores. Alcohol abuse and impulsivity were the only aspects of mental health that had negative R^2 values for the non-random

models. Increasing the number of words per user, with a constant sample size, did increase the depression trained model's predictive performance. Increasing the threshold from 5 days of Tweets (minimum of 43 words) to 500 words per user caused the R^2 to increase from 0.010 to 0.034 (Table S1). However, there was substantial variation at the lower word count thresholds with $R^2 = -0.001$ at 200 words per user up to a maximum of $R^2 = 0.044$ at 400 words per user.

Prior research has suggested that partially dissociable transdiagnostic dimensions of mental health may be a better fit to the underlying neurobiology of mental illness. This was not true of its fit to language patterns in Twitter data assessed here. A model trained to predict 'anxious-depression' scores performed nominally worse ($R^2 = 0.016$) than the models trained and tested on the depression ($R^2 = 0.025$) or generalised anxiety ($R^2 = 0.045$) questionnaires, as reported above (Figure S5). The anxious depression model was also non-specific, having modest but non-zero predictive power for both 'compulsivity and intrusive thought' ($R^2 = 0.025$) and social withdrawal ($R^2 = 0.014$). There were no significant associations between any aspect of Twitter use, e.g., number of Tweets, and depression residuals from the held-out test set (all $|\beta| > 0.02$, $p > 0.05$) (Figure S6A). Additionally, depressed residuals were normally distributed and centered on zero (Mean = -0.05, $t = -0.84$ (df = 301), $p = 0.40$) (Figure S6B). Deviations from true depression scores, i.e., depression residuals, within the LIWC text features model are not due to any systematic differences in participant engagement on Twitter. Finally, we evaluated the predictive performance of a depression trained classification model (Table 2). The best performing model was a SVM trained on data from the top 476 users by word count which had an AUC and accuracy of 0.59. There was no

difference in performance between the SVM and RF models, and both models performed worse than the 68% accuracy found by [5]. Decreasing the sample size to included only the top participants also had no effect on predictive performance, although there was a substantial increase in sensitivity with a slight decline in specificity. The depression-keyword model had an 83.6% accuracy, 0.83 AUC, 76.9% sensitivity, and 88.9% specificity, which substantially outperformed the classification model trained and tested on self-reported depression (57% accuracy, 0.57 AUC, 52% sensitivity and 63% specificity). While participants with depression relevant keywords do have greater depression severity ($\beta = 0.26$, $SE = 0.016$, $z = 4.00$ ($df = 1,004$), $p < 0.001$), the use of keywords within posts to define cases of mental illness demonstrates a substantial overestimation of model performance compared to validated self-report questionnaires.

Figure 3: Elastic Net predictive performance of a depression model on itself, and 9 other aspects of mental health

Predictive performance (R^2) from an Elastic Net model trained on depression and tested on each of the other aspects of mental health recorded for randomised text

features (red), text features only (blue), age and gender only (pink), and text features plus age and gender (green).

Model	AUC	Accuracy	F1	Sensitivity	Specificity
SVM (n = 1,006)	0.56	0.59	0.42	0.33	0.79
SVM (n = 476)	0.59	0.59	0.54	0.52	0.66
RF (n = 1,006)	0.56	0.58	0.44	0.38	0.75
RF (n = 476)	0.57	0.57	0.53	0.52	0.63
De Choudhury SVM (n = 476)	NA	0.68	NA	0.58	NA

Table 2: Depression Classification Performance of a Support Vector Machine (SVM) and Random Forest (RF) model with varying sample size

Depression classification performance was similar between SVM and RF regardless of sample size. In the reduced sample size (n = 476) of participants with high word counts there was an increase in the F1 score and sensitivity for both the SVM and RF models. However, in both models the increase in sensitivity was accompanied by a reduction in specificity. Neither model improved over the classification performance of de Choudhury et al. (2013).

Figure 5: Comparison of text feature only random forest models trained on either i) a depression self-report questionnaire or ii) depression relevant keywords.

A) Receiver operator curves for depression self-report (AUC = 0.56) and depression keyword (AUC = 0.83) trained models, dashed line indicates chance level performance. B) Top 10 text features for the depression self-report model. C) Top 10 text features for the depression keyword model.

Comparison of models trained on Tweets, Retweets, Likes

By training the depression model on Tweets, Retweets, and Likes separately, we found that the Likes model had the greatest predictive value ($R^2 = 0.026$) compared with an R^2 of 0.010 for the Tweets only model (Figure S7A). There was, however, the improved predictive power may be because there was simply more data for Likes ($M=798.4$, $SD=1041.3$) compared to Tweets ($M=231.3$, $SD=489.4$) and Retweets ($M=173.5$, $SD=416$) (Table 1). Indeed, when we split data based on quantity, predictive power was greatest in the 4th quartile of Tweet data ($R^2 = 0.043$) with negative R^2 values from models trained on the 1st and 2nd quartiles of Tweets (Figure S7B). Therefore, improved performance can be achieved when the corpus is particularly large, though it clearly remains modest, explaining just 4% of variance in depression.

Similarity and specificity of text features across mental health phenotypes

To compare the content of models developed for each of our 9 aspects of mental health, we generated predictive models for each separately and examined how similar

the most predictive text features were. In the depression trained model, we found that 'focus on present' and 1st person plural pronouns were selected in 100% of models. While negative emotions and 1st person singular pronouns, text features previously found to be associated with depression, had selection frequencies of 0.951 and 0.659 respectively (Figure 4A). First person singular pronouns were selected more often in models related to schizotypy (0.963), social anxiety (0.658), eating disorders (0.999) and generalised anxiety (1.0) than depression. Negative emotions were slightly more specific to depression than 1st person pronouns with higher selection frequencies in only schizotypy (0.963) and generalised anxiety (1.0). Among the top 20 text features by selection frequency, only affiliation words were unique to depression. Eating disorders and alcohol abuse had the largest number of unique text features with five and six respectively. The generalised anxiety trained model had the highest predictive performance of all models ($R^2 = 0.045$), followed by schizotypy ($R^2 = 0.037$). For impulsivity, eating disorders, alcohol abuse, and apathy the percent of variance explained was below 1% (Figure 4B).

Hierarchical clustering revealed, unsurprisingly given the high correlation across these questionnaires, depression language use was most similar to generalized anxiety followed closely by schizotypy (Figure 4C). Obsessive compulsive disorder was most closely related to apathy in terms of language use patterns and are slightly more dissimilar to each other compared to the depression, generalised anxiety, and schizotypy cluster. Perhaps most interestingly, alcohol abuse and eating disorders formed their own cluster largely separate from other disorders, indicating substantial

540 differences in language use both between those disorders and relative to the other
541 disorders considered.

A

B

C

Figure 4: Depression model text feature selection frequencies and dendrogram of language similarities between aspects of mental health

A) Model selection frequencies for the top 20 text features in models of 9 psychiatric questionnaires over 100 iterations. Darker colours indicate text features that appear among the top 20 text features across fewer questionnaires, i.e., are more specific. The mean direction of association between each text feature and the target outcome is denoted by a + (positive) or – (negative) next to the label. B) Predictive performance of models trained and tested on each self-report questionnaire C) Hierarchical clustering dendrogram from a model trained on mean text feature selection frequencies.

To examine specificity, we focused on a smaller set of three transdiagnostic dimensions of mental health that can be derived from the larger set of questionnaires: ‘anxious-depression’, ‘compulsivity and intrusive thought’ and ‘social withdrawal’. Much like the analysis of the original questionnaire total scores, the top text features by selection frequency in each transdiagnostic dimension were not specific to any one dimension (Figure 5A). Similar to models trained on each questionnaire individually, 1st person plural pronouns and focus on present words were among the top text features associated with not just anxious depression, but also and compulsivity and intrusive thought dimensions. In each dimension, less than 45% of top 20 text features were specific to the top 20 for that particular dimension. However, after removing the shared variance between dimensions, at least 80% of text features were specific to each dimension (Figure 5B). No text feature was present in the top 20 of all 3 dimensions.

Anger words, 1st person singular pronouns, and family words were the top text feature associated with anxious depression, compulsivity and intrusive thought, and social withdrawal respectively. Importantly each of these text features was specific to that particular dimension, i.e., not found in the top 20 for any other dimension.

Figure 5: Transdiagnostic dimension text feature selection frequencies

A) Model selection frequencies for the top 20 text features in models of 3 transdiagnostic dimensions over 100 iterations. B) Model selection frequencies for the top 20 text features in models trained on the residuals of each transdiagnostic dimension after controlling for the shared variance due to the other dimensions

Discussion

Principal Findings

There is growing interest in the power of artificial intelligence for improving healthcare provision, early intervention, and diagnosis. But large amounts of data are needed to develop and train models, which can be arduous to gather. Social media data has been suggested to be a convenient and readily available source of such data. This is because social media platforms are in widespread use, users produce sometimes exceedingly high volumes of data, regularly and spanning many years, and these data often contains rich personal and emotional information of putative relevance to their mental state. Although several studies have examined this in recent years, there are substantial limitations to the methods in widespread use[36, 59], including but not limited to the validity of diagnostic classifications employed and the rigor of the machine learning methods employed. Here, we collected Twitter data and 9 validated self-report questionnaires from over 1,000 participants assessing their mental health. We used gold-standard machine learning methods with out-of-sample testing to establish the predictive power of models trained to predict depression and other aspects of mental health, using linguistic features derived from Tweets.

A model developed to predict individual differences in self-report depression explained 2.5% of variance when tested out of sample. The age and gender model, however, slightly outperformed the text feature only model, illustrating that a similar level of depression prediction can be achieved using just these two data points. It is worth noting that age and gender is not routinely available on Twitter but was gathered as part of our survey. When age, gender, and text features were included in the same model, it

602 still only explained approximately 4% of variance in depression severity. We examined
603 the specificity of this depression model, on 8 other questionnaire total scores gathered
604 from the same participants. We found that although the model had some small
605 predictive value for 6 other aspects of mental health studied here, generalized anxiety,
606 schizotypy, obsessive-compulsive disorder, eating disorder, apathy and social anxiety, it
607 was not able to explain variance in alcohol abuse and impulsivity. Furthermore, we
608 found that there were no associations between any aspect of Twitter use, e.g., word
609 count, and the residuals of the depression model's predictions. Failures in model
610 performance, therefore, seem to be random and not explained by lower engagement
611 nor number of social connections.

612 We tested if previously identified transdiagnostic symptom dimensions, which tend to
613 perform better than these questionnaires in fitting cognitive test performance[54], might
614 improve signal and/or specificity. This was not the case. However, after controlling for
615 shared variance among the transdiagnostic dimensions, we found that most text
616 features were specific to the residuals of each dimension. Perhaps most strikingly, 1st
617 person singular pronouns have been consistently found to be a key characteristic of
618 depression relevant language[22, 60], but when controlling for shared variance, we
619 found that an increased use of 1st person pronouns was actually most associated with
620 the compulsivity and intrusive thought dimension. Overall, generalised anxiety and
621 schizotypy were the best performing models while the alcohol abuse model had close to
622 zero out of sample performance. Hierarchical clustering revealed that language use
623 associated with alcohol abuse and eating disorders were most dissimilar to the other
624 disorders. Most prior social media research has focused on associations between

language use and alcohol usage at the sub-national, rather than individual, level[61]. A possible explanation for the low predictive value of the alcohol model is that few people in our study scored high enough to qualify as alcohol dependent.

A depression classification model trained on the presence of depression relevant keywords had substantially better predictive performance compared to a model trained on dichotomized self-reported depression. Prior studies have shown that regular expression, i.e., keywords, can be used to identify depression with a high degree of accuracy[62, 63]. To our knowledge, however, no studies have compared relative predictive performance of a depression-keyword trained model to one trained on depression self-report scores within the same sample. Although self-report measures are more difficult to pragmatically acquire from a large sample, they represent an important and clinically validated ground truth. Our results indicate the potential pitfalls of defining cases of mental illness through keyword-based methods, that is, a sort of content-based circularity can arise when social media posts are used to define caseness, train and evaluate machine learning models. Our data suggest that persons more likely to discuss depression in Tweets, have a distinct pattern of associated language use, but they do not necessarily suffer from clinical depression, with only 50% of these participants meeting the clinical cut-off for depression. These findings underscore the need to use valid ground truth estimates of mental health in developing models of clinical relevance.

Twitter Metadata

Exploratory analyses found that elevated rates of replying and Tweeting were broadly associated with mental health, correlating all questionnaire total scores, except alcohol

abuse and eating disorders. Inconsistent evidence exists around whether people with greater depression severity are more [8] or less [5] active on social media. We found that participants with greater obsessive-compulsive severity had more followees while people with more severe eating disorder symptoms had more account followers. Depressed individuals have consistently been shown to Tweet more at night during the day [5, 49, 64]. Later Tweet times were associated with depression severity, but also apathy, impulsivity, obsessive compulsive, and schizotypy. Impulsivity had the strongest association with the insomnia index, in line with prior research showing a positive association between impulsivity and sleep disturbances[65, 66]. Besides findings related to the number of followees and followers, Twitter metadata, like language use, was generally not specific to any one mental health condition.

Comorbidity

People with depression have been found to use language differently from healthy controls. Most studies, however, compare people with one mental health disorder with healthy controls[22, 25, 67-69]; few have examined the specificity of different aspects of language use across disorders. The non-specific patterns of language use observed here, both in prior work[29-31] and the current study, is likely related to the high comorbidity rates among disorders. We found that only by removing the shared variance among disorders could we identify which aspects of language use were specific to each mental health dimension. Major depressive disorder is positively associated with a variety of other mental health conditions including: panic disorder, agoraphobia, generalised anxiety disorder, post-traumatic stress disorder, obsessive compulsive disorder, and separation anxiety disorder[70]. For example, a patient

diagnosed with major depression is 8.2 times more likely to have a concurrent diagnosis of generalised anxiety than someone without depression[71]. In our study, we found that depression and anxiety had the most similar language use of any pair of disorders. Depression symptoms overlap strongly with other disorders and are associated numerous symptoms in other diagnostic categories[72]. In a network of Diagnostic and Statistical Manual of Mental Disorders-IV symptoms, depression symptoms (insomnia, psychomotor agitation/retardation, and depressed mood) were the most connected symptoms with connections to over 28% of other symptoms in the network[73]. The spread of symptoms across disorders makes it unlikely that individual text features or even combinations of text features could ever be specific to categorical disorders, a finding in line with growing consensus that these diagnostic categories are overlapping and warrant revision[28].

Social Media Interaction Effects

Social media is not a one-way street. While the content of social media posts reflects the underlying mental health of the user, interactions, both passive and active, on the platforms can act to either improve or worsen mental health. When users experience stressful event, they are more likely to disclose this information on social media. Self-disclosure was shown to subsequently moderate the adverse effects of a stressful event and led to enhanced life satisfaction and lower depression via enhanced social support[74]. However, in a separate study, Reddit users who transitioned to talking about suicide had elevated levels of self-disclosure but received less social support and engagement than users who did not[75]. Furthermore, specific types of social support are more likely to lead to improvements in mental health, e.g., use of the phrase 'be

tough'[76]. Increasing awareness about these types of comments would help friends, family, and content moderators to know what to say to and what not to say to someone experiencing mental health difficulties. While there are benefits to self-disclosure these can only be realized if the user is able to communicate free of stigma and receive adequate support. The effects of self-disclosure on social media highlight the need to follow users longitudinally and consider factors beyond just language use, i.e., social network structure, when predicting mental health. Considering the availability of online social support could help triage users with the same predicted risk of mental illness; users with less social support should be prioritized for receiving help.

Practical Utility

Most Twitter data is generated by a small subset of users, 80% of Tweets are written by only 10% of users[77]. We found some evidence that machine learning language models perform slightly better when trained on subsets of users with more Tweets. This might suggest that in an even more select sample, those in the top 10% of users overall, one could produce more reliable predictions. However, two things are important to remember here. First, even in our top quartile, the variance explained only rose to a high of 4.3%, additional gains are unlikely to take this to the realm of practice. Similarly increasing the minimum word count per user only slightly increased the percent variance explained. At a minimum threshold of 400 words, 6.4% of variance was explained, while a threshold of 500 words was slightly worse at 3.4%. Second, those users are not representative of social media users in general, so even if such performance could be achieved, these models are unlikely to be generalizable. An interesting possibility is that the signal may be more meaningfully improved if private

sources of text could be harnessed such as text messages. This would have the additional benefit of increasing the amount of data available for each user while simultaneously being more relevant to a user's true mental health status.

Although we demonstrated that social media data has low predictive power at an individual level, this should be contextualized as part of the broader landscape of effect sizes in mental health science. For example, well-established correlates of mental health problems such as adverse childhood experiences only yield an area under the curve of 0.58 in predicting mental health problems at age 18[78]. A recent preprint showed that resting-state and structural brain-wide associations to psychopathology are exceedingly small, with no reliable correlation exceeding 0.16 [79]. Because these observations do not have value as individual predictors, does not make the observation devoid of meaning. Mental health is exceedingly complex and likely combinations of a range of sources of multimodal data will be required to take these small effects and transform them into meaningful N-of-1 predictions. Twitter data, by itself, has already proven an interesting testbed for nascent theories of mental health such as network theory, which for example, has struggled to acquire large enough longitudinal datasets to test some of its core predictions[80]. We recently found for example that using social media posts as a proxy for experience sampling allowed us to study a large cohort of individuals through a transition to a depressed state, detecting subtle network signatures of depression vulnerability[81].

Ethical Implications

Mental health detection from social media offers the potential for generating continuous insights into mental health at the population and individual level, but also poses a

740 unique set of ethical challenges. Large scale analyses of social media data are typically
741 exempt from requiring participant consent due to the public nature of data and lack of
742 experimental intervention. Because of this exemption, social media users are often
743 unaware whether or not their data is included in research and when asked tend to be
744 uncomfortable with the idea that their Twitter data could be used for research purposes
745 without their knowledge[82]. While it is impractical to ask for consent in all
746 circumstances, requiring consent whenever possible ensures that participants have
747 safeguards for how their data is used. Predicting an individual's mental health outside of
748 a clinical context inherently poses the question of how to act on that information and
749 whether there is in fact an obligation to act[83]. Unlike clinicians, software developers
750 are not obligated to intervene if their algorithm detects that a person is struggling with
751 their mental health. If they developers are not obligated to intervene, would the burden
752 fall on family members, friends or the individuals themselves? Even if a patient
753 consented to having their social media feed monitored by their physician, a high rate of
754 false positives would overwhelm a clinician and impede their ability to effectively
755 allocate care. Furthermore, there is a potential for misuse of mental health predictions
756 by bad actors who do not consider the best interests of the user. Passive and automatic
757 detection of mental illness could lead to targeted advertisements of prescription
758 medication[84] or result in an increase in health insurance premiums. A final concern
759 relates to algorithmic bias based on the data used to train these models. Social media
760 users tend to be younger, more affluent, and hold more left leaning political views than
761 the general population[77, 85]. Furthermore, social media research is strongly focused
762 on predominantly English-speaking countries yet there is evidence that people from

different cultures behave differently online, for example users from China and India post questions online more frequently than users from the US and UK[86]. Extrapolating models to very different users than the models were trained on could lead to systematic biases that impact the predictive performance for groups not included in the training data.

Limitations

Prior studies have had larger Twitter datasets in terms of the number of posts per user. For example, de Choudhury et al [5] had a mean of 4,500 posts per user in a 1-year period, while in our study had a mean of about 1,100 posts, including likes, per user. As mentioned above, models perform better when provided more training data per user. Indeed, this study achieved greater predictive power than reported here. However, there were other differences across our studies too; our sample was twice as big, and we used an independent training set to build our model and then evaluate it on an independent test set. Compared to simple K-fold cross validation using the entire dataset [5], this procedure is less likely to overfit the data and overestimate predictive performance. Another potential limitation to our study is that our text features analysis was restricted to only using categories from the LIWC library. Some evidence exists that more data-driven approaches, e.g., topic analysis, could slightly improve predictive ability over closed libraries[18, 87]. More sophisticated machine learning models, such as convolutional neural networks have the potential to make superior predictions than more commonly used algorithms, although with the limitation of needing substantially more data[88]. While these methods might indeed yield improvements in performance, the use of LIWC has key advantages. LIWC is a closed library that has been well-

786 validated and studied across a range of communication media from diary entries[22] to
787 spoken word[89]. This means that the numerical values and classifications assigned to
788 individual words in LIWC does not change from dataset to dataset, as is often the case
789 with topics and neural networks[90, 91]. This makes the insights derived here more
790 reproducible and generalizable to new datasets that may be of keen interest in future,
791 such as text messages and email communications.

792 Regarding the choice of social media platform, it is nonetheless a limitation that our
793 study was confined to Twitter. Recent evidence has also shown that Facebook may be
794 more predictive of mental health conditions than Twitter[57]. We selected Twitter
795 because it is the most used social media platform for studying mental health, comprising
796 approximately 40% of studies on the subject, while Facebook makes up only about
797 8%[36]. It remains a limitation that these results could reflect a relative lack of predictive
798 performance that is particular to Twitter. Because we did not have binary diagnostic
799 information, we did not attempt to classify participants with either depression vs.
800 anxiety, obsessive-compulsive disorder etc., i.e., multi-class classification of mental
801 health diagnoses. Instead, we tried to continuously predict a participant's score on a
802 range of self-report questionnaires probing different aspects of mental health.
803 Therefore, rather than differentiating users with one diagnosis or another, we instead
804 attempted to quantify the similarity of language use between self-report symptoms of
805 highly comorbid conditions. We think this dimensional approach has many advantages,
806 but this creates a limitation in how directly these data can be applied to diagnoses
807 assigned by a clinician. Finally, subjects in this study reported mental health symptoms
808 at the point of study entry, and we analyzed data corresponding to the 12-month directly

prior to this. This necessitates taking a ‘trait’ perspective on the mental health symptoms we assessed and it is likely that our model is diluted by variations in state/episodic features of depression. However, in a recent study we found that individuals’ use of depression-relevant text features in fact didn’t change significantly across within-subject periods of mental health and wellness, suggesting this may not be a major issue[81].

Conclusions

We found that language use patterns on Twitter that relate to depression symptom severity cannot be used to develop predictive models with high accuracy on an individual subject-basis. A model trained to predict depression is also non-specific, being additionally predictive of several mental health symptom profiles. Although performance was poor at the individual subject level, the effect sizes observed are not out of proportion with other routinely studied cross-sectional observations in psychiatry. The addition of age and gender improved performance of our depression model, suggesting that the combination of various sources of multimodal data (with individually small effect sizes) is a viable path forward to improve predictive power of these class of models. Furthermore, controlling for other mental health conditions and training models on the resultant residuals is a promising method for finding language use specific to that condition. To our knowledge, we are the first study to train machine learning algorithms on the residuals of mental health dimensions in order to identify unique patterns of language. This approach highlights the benefits of using self-report questionnaires to measure mental health since it is not possible for studies with a binary classification of cases, i.e., healthy control vs. case, to account for shared variance between disorders. Although classification studies are able to identify cases of mental illness, they are

832 unlikely to be able to determine specifically what aspects of language are different and
833 unique to a particular condition. Determining determine specific changes in language
834 patterns and use is crucial for the utility of using text data for diagnostic purposes,
835 regardless of data source.

836 Nevertheless, we do not believe that social media should be used in a diagnostic setting
837 both for privacy concerns on behalf of the user and the relatively low quality of
838 predictions that would limit clinical utility. Despite the low signal, by virtue of the
839 availability of large amounts of data, the analysis of social media data remains a useful
840 tool to test theories of mental health that are difficult to test using conventional means.
841 Should people be concerned that their mental health status can be unintentionally
842 revealed by the content of their Tweets? We think the data do not support this as a
843 meaningful risk at present.

844 **Declarations**

845 **Ethics approval and consent to participate**

846 Approved was granted by the Trinity College Dublin Department of Psychology
847 Research Ethics Committee (Approval ID: SPREC112018-32).

848 **Consent for publication**

849 Not applicable.

850 **Availability of data and materials**

851 Processed and anonymized datasets will be made available from the corresponding
852 author upon reasonable request from other researchers, but raw tweets cannot be
853 made available due to the potential for re-identifying research participants. The code
854 used to analyse the data in the current study is available at:
855 <https://github.com/seanwkelley/TwitterLanguageSpecificity>

856 **Competing interests**

857 The authors declare that they have no competing interests.

858 **Funding**

859 This study was funded by the ‘Institutional Strategic Support Fund’ grant
860 (204814/Z/16/A) to Trinity College Dublin, funded by the SFI-HRB-Wellcome Trust
861 partnership. SWK is funded by a Provost PhD Project Award awarded to CMG. CMG
862 holds funding from Science Foundation Ireland’s Frontiers for the Future Scheme
863 (19/FFP/6418), MQ: transforming mental health (MQ16IP13) and a European Research
864 Council Starting Grant (ERC-H2020-HABIT).

865 **Authors' contributions**

866 SWK: Conceptualization, methodology, study design, statistical analysis, writing
867 (drafting and editing). CNM: Methodology. LB: Methodology. RW: methodology,
868 statistical analysis, writing (editing). CMG: Conceptualization, methodology, study
869 design, statistical analysis, writing (drafting, editing, and supervision).

870

References

1. Abuse S: **Mental Health Services Administration.(2019). Key substance use and mental health indicators in the United States: Results from the 2018 National Survey on Drug Use and Health (HHS Publication No. PEP19-5068, NSDUH Series H-54).** Rockville, MD: Center for Behavioral Health Statistics and Quality. *Substance Abuse and Mental Health Services Administration* Retrieved from <https://www.samhsa.gov/data> 2019.
2. Lépine J-P, Gastpar M, Mendlewicz J, Tylee A: **Depression in the community: the first pan-European study DEPRES (Depression Research in European Society).** *International clinical psychopharmacology* 1997.
3. Ghio L, Gotelli S, Marcenaro M, Amore M, Natta W: **Duration of untreated illness and outcomes in unipolar depression: a systematic review and meta-analysis.** *Journal of affective disorders* 2014, **152**:45-51.
4. Perrin A: **Social media usage.** *Pew research center* 2015, **125**:52-68.
5. De Choudhury M, Gamon M, Counts S, Horvitz E: **Predicting depression via social media.** In: *Seventh international AAAI conference on weblogs and social media: 2013; 2013.*
6. De Choudhury M, Counts S, Horvitz EJ, Hoff A: **Characterizing and predicting postpartum depression from shared facebook data.** In: *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing: 2014; 2014:* 626-638.
7. Tsugawa S, Kikuchi Y, Kishino F, Nakajima K, Itoh Y, Ohsaki H: **Recognizing Depression from Twitter Activity.** In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15.* 2015: 3187-3196.
8. Eichstaedt JC, Smith RJ, Merchant RM, Ungar LH, Crutchley P, Preotiuc-Pietro D, Asch DA, Schwartz HA: **Facebook language predicts depression in medical records.** *Proc Natl Acad Sci U S A* 2018, **115**(44):11203-11208.
9. Reece AG, Reagan AJ, Lix KLM, Dodds PS, Danforth CM, Langer EJ: **Forecasting the onset and course of mental illness with Twitter data.** *Sci Rep* 2017, **7**(1):13006.
10. Reece AG, Danforth CM: **Instagram photos reveal predictive markers of depression.** *EPJ Data Science* 2017, **6**(1).
11. De Choudhury M: **Anorexia on Tumblr.** In: *Proceedings of the 5th International Conference on Digital Health 2015.* 2015: 43-50.
12. Wolf M, Theis F, Kordy H: **Language Use in Eating Disorder Blogs.** *Journal of Language and Social Psychology* 2013, **32**(2):212-226.
13. Wang T, Brede M, Ianni A, Mentzakis E: **Social interactions in online eating disorder communities: A network perspective.** *PLoS One* 2018, **13**(7):e0200800.
14. Chancellor S, Lin Z, Goodman EL, Zerwas S, De Choudhury M: **Quantifying and predicting mental illness severity in online pro-eating disorder communities.** In: *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing: 2016; 2016:* 1171-1184.
15. Mitchell M, Hollingshead K, Coppersmith G: **Quantifying the language of schizophrenia in social media.** In: *Proceedings of the 2nd workshop on Computational linguistics and clinical psychology: From linguistic signal to clinical reality: 2015; 2015:* 11-20.
16. Birnbaum ML, Ernala SK, Rizvi AF, De Choudhury M, Kane JM: **A Collaborative Approach to Identifying Social Media Markers of Schizophrenia by Employing Machine Learning and Clinical Appraisals.** *J Med Internet Res* 2017, **19**(8):e289.
17. McManus K, Mallory EK, Goldfeder RL, Haynes WA, Tatum JD: **Mining Twitter data to improve detection of schizophrenia.** *AMIA Summits on Translational Science Proceedings* 2015, **2015**:122.

- 918 18. Nobles AL, Glenn JJ, Kowsari K, Teachman BA, Barnes LE: **Identification of Imminent Suicide Risk**
919 **Among Young Adults using Text Messages.** *Proc SIGCHI Conf Hum Factor Comput Syst* 2018,
920 **2018.**
- 921 19. Coppersmith G, Ngo K, Leary R, Wood A: **Exploratory analysis of social media prior to a suicide**
922 **attempt.** In: *Proceedings of the third workshop on computational linguistics and clinical*
923 *psychology: 2016*; 2016: 106-117.
- 924 20. Cheng Q, Li TM, Kwok CL, Zhu T, Yip PS: **Assessing Suicide Risk and Emotional Distress in**
925 **Chinese Social Media: A Text Mining and Machine Learning Study.** *J Med Internet Res* 2017,
926 **19(7):e243.**
- 927 21. Coppersmith G, Leary R, Crutchley P, Fine A: **Natural Language Processing of Social Media as**
928 **Screening for Suicide Risk.** *Biomed Inform Insights* 2018, **10:1178222618792860.**
- 929 22. Rude S, Gortner E-M, Pennebaker J: **Language use of depressed and depression-vulnerable**
930 **college students.** *Cognition & Emotion* 2004, **18(8):1121-1133.**
- 931 23. Zimmermann J, Brockmeyer T, Hunn M, Schauenburg H, Wolf M: **First-person Pronoun Use in**
932 **Spoken Language as a Predictor of Future Depressive Symptoms: Preliminary Evidence from a**
933 **Clinical Sample of Depressed Patients.** *Clin Psychol Psychother* 2017, **24(2):384-391.**
- 934 24. Zimmermann J, Wolf M, Bock A, Peham D, Benecke C: **The way we refer to ourselves reflects**
935 **how we relate to others: Associations between first-person pronoun use and interpersonal**
936 **problems.** *Journal of Research in Personality* 2013, **47(3):218-225.**
- 937 25. Molendijk ML, Bamelis L, van Emmerik AA, Arntz A, Haringsma R, Spinhoven P: **Word use of**
938 **outpatients with a personality disorder and concurrent or previous major depressive disorder.**
939 *Behav Res Ther* 2010, **48(1):44-51.**
- 940 26. De Choudhury M, Counts S, Horvitz E: **Social media as a measurement tool of depression in**
941 **populations.** In: *Proceedings of the 5th Annual ACM Web Science Conference: 2013*; 2013: 47-
942 56.
- 943 27. Kessler RC, Magee WJ: **Childhood adversities and adult depression: basic patterns of**
944 **association in a US national survey.** *Psychological medicine* 1993, **23(3):679-690.**
- 945 28. Insel T, Cuthbert B, Garvey M, Heinssen R, Pine DS, Quinn K, Sanislow C, Wang P: **Research**
946 **domain criteria (RDoC): toward a new classification framework for research on mental**
947 **disorders.** In.: Am Psychiatric Assoc; 2010.
- 948 29. Coppersmith G, Dredze M, Harman C, Hollingshead K: **From ADHD to SAD: Analyzing the**
949 **language of mental health on Twitter through self-reported diagnoses.** In: *Proceedings of the*
950 *2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to*
951 *Clinical Reality: 2015*; 2015: 1-10.
- 952 30. Cohan A, Desmet B, Yates A, Soldaini L, MacAvaney S, Goharian N: **SMHD: a large-scale resource**
953 **for exploring online language usage for multiple mental health conditions.** *arXiv preprint*
954 *arXiv:180605258* 2018.
- 955 31. Lyons M, Aksayli ND, Brewer G: **Mental distress and language use: Linguistic analysis of**
956 **discussion forum posts.** *Computers in Human Behavior* 2018, **87:207-211.**
- 957 32. Coppersmith G, Dredze M, Harman C: **Quantifying mental health signals in Twitter.** In:
958 *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic*
959 *signal to clinical reality: 2014*; 2014: 51-60.
- 960 33. Wolf M, Sedway J, Bulik CM, Kordy H: **Linguistic analyses of natural written language:**
961 **unobtrusive assessment of cognitive style in eating disorders.** *Int J Eat Disord* 2007, **40(8):711-**
962 **717.**
- 963 34. Ernala SK, Rizvi AF, Birnbaum ML, Kane JM, De Choudhury M: **Linguistic Markers Indicating**
964 **Therapeutic Outcomes of Social Media Disclosures of Schizophrenia.** *Proceedings of the ACM*
965 *on Human-Computer Interaction* 2017, **1(CSCW):1-27.**

- 966 35. Zomick J, Levitan SI, Serper M: **Linguistic analysis of schizophrenia in reddit posts**. In: *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology: 2019*; 2019: 74-83.
- 967
- 968
- 969 36. Chancellor S, De Choudhury M: **Methods in predictive techniques for mental health status on social media: a critical review**. *NPJ digital medicine* 2020, **3**(1):1-11.
- 970
- 971 37. Ireland M, Iserman M: **Within and between-person differences in language used across anxiety support and neutral reddit communities**. In: *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic: 2018*; 2018: 182-193.
- 972
- 973
- 974
- 975 38. Brundage M, Avin S, Clark J, Toner H, Eckersley P, Garfinkel B, Dafoe A, Scharre P, Zeitzoff T, Filar B: **The malicious use of artificial intelligence: Forecasting, prevention, and mitigation**. *arXiv preprint arXiv:180207228* 2018.
- 976
- 977
- 978 39. Zung WW: **A self-rating depression scale**. *Archives of general psychiatry* 1965, **12**(1):63-70.
- 979 40. Mason O, Linney Y, Claridge G: **Short scales for measuring schizotypy**. *Schizophrenia research* 2005, **78**(2-3):293-296.
- 980
- 981 41. Foa EB, Huppert JD, Leiberg S, Langner R, Kichic R, Hajcak G, Salkovskis PM: **The Obsessive-Compulsive Inventory: development and validation of a short version**. *Psychological assessment* 2002, **14**(4):485.
- 982
- 983
- 984 42. Garner D, Olmsted M, Bohr Y, Garfinkel P: **The eating attitudes test: Psychometric features**. *Psychological medicine* 1982, **12**:871-878.
- 985
- 986 43. Patton JH, Stanford MS, Barratt ES: **Factor structure of the Barratt impulsiveness scale**. *Journal of clinical psychology* 1995, **51**(6):768-774.
- 987
- 988 44. Saunders JB, Aasland OG, Babor TF, De La Fuente JR, Grant M: **Development of the alcohol use disorders identification test (AUDIT): WHO collaborative project on early detection of persons with harmful alcohol consumption-II**. *Addiction* 1993, **88**(6):791-804.
- 989
- 990
- 991 45. Marin RS, Biedrzycki RC, Firinciogullari S: **Reliability and validity of the Apathy Evaluation Scale**. *Psychiatry research* 1991, **38**(2):143-162.
- 992
- 993 46. Liebowitz MR: **Social phobia**. *Modern problems of pharmacopsychiatry* 1987.
- 994 47. Spielberger CD: **State-trait anxiety inventory for adults**. 1983.
- 995 48. Pennebaker JW, Boyd RL, Jordan K, Blackburn K: **The development and psychometric properties of LIWC2015**. In.; 2015.
- 996
- 997 49. Chen X, Sykora MD, Jackson TW, Elayan S: **What about mood swings: Identifying depression on twitter with temporal measures of emotions**. In: *Companion Proceedings of the The Web Conference 2018: 2018*; 2018: 1653-1660.
- 998
- 999
- 1000 50. Zou H, Hastie T: **Regularization and variable selection via the elastic net**. *Journal of the royal statistical society: series B (statistical methodology)* 2005, **67**(2):301-320.
- 1001
- 1002 51. Jollans L, Boyle R, Artiges E, Banaschewski T, Desrivieres S, Grigis A, Martinot J-L, Paus T, Smolka MN, Walter H: **Quantifying performance of machine learning methods for neuroimaging data**. *NeuroImage* 2019, **199**:351-365.
- 1003
- 1004
- 1005 52. Dinga R, Schmaal L, Penninx BW, Veltman DJ, Marquand AF: **Controlling for effects of confounding variables on machine learning predictions**. *BioRxiv* 2020.
- 1006
- 1007 53. Rouault M, Seow T, Gillan CM, Fleming SM: **Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance**. *Biological psychiatry* 2018, **84**(6):443-451.
- 1008
- 1009
- 1010 54. Gillan CM, Kosinski M, Whelan R, Phelps EA, Daw ND: **Characterizing a psychiatric symptom dimension related to deficits in goal-directed control**. *Elife* 2016, **5**:e11305.
- 1011
- 1012 55. Seow TX, Gillan CM: **Transdiagnostic phenotyping reveals a host of metacognitive deficits implicated in compulsivity**. *Scientific reports* 2020, **10**(1):1-11.
- 1013

- 1014 56. Murtagh F, Legendre P: **Ward's hierarchical agglomerative clustering method: which**
1015 **algorithms implement Ward's criterion?** *Journal of classification* 2014, **31**(3):274-295.
- 1016 57. Jaidka K, Guntuku S, Ungar L: **Facebook versus Twitter: Differences in Self-Disclosure and Trait**
1017 **Prediction.** In: *Proceedings of the International AAAI Conference on Web and Social Media:*
1018 *2018*; 2018.
- 1019 58. Schwartz HA, Eichstaedt J, Kern M, Park G, Sap M, Stillwell D, Kosinski M, Ungar L: **Towards**
1020 **assessing changes in degree of depression through facebook.** In: *Proceedings of the workshop*
1021 *on computational linguistics and clinical psychology: from linguistic signal to clinical reality:*
1022 *2014*; 2014: 118-125.
- 1023 59. Guntuku SC, Yaden DB, Kern ML, Ungar LH, Eichstaedt JC: **Detecting depression and mental**
1024 **illness on social media: an integrative review.** *Current Opinion in Behavioral Sciences* 2017,
1025 **18**:43-49.
- 1026 60. Edwards TM, Holtzman NS: **A meta-analysis of correlations between depression and first**
1027 **person singular pronoun use.** *Journal of Research in Personality* 2017, **68**:63-68.
- 1028 61. Curtis B, Giorgi S, Buffone AE, Ungar LH, Ashford RD, Hemmons J, Summers D, Hamilton C,
1029 Schwartz HA: **Can Twitter be used to predict county excessive alcohol consumption rates?** *PloS*
1030 *one* 2018, **13**(4):e0194290.
- 1031 62. Prieto VM, Matos S, Alvarez M, Cacheda F, Oliveira JL: **Twitter: a good place to detect health**
1032 **conditions.** *PloS one* 2014, **9**(1):e86191.
- 1033 63. Nakamura T, Kubo K, Usuda Y, Aramaki E: **Defining patients with depressive disorder by using**
1034 **textual information.** In: *2014 AAAI Spring Symposium Series: 2014*; 2014.
- 1035 64. Leis A, Ronzano F, Mayer MA, Furlong LI, Sanz F: **Detecting Signs of Depression in Tweets in**
1036 **Spanish: Behavioral and Linguistic Analysis.** *J Med Internet Res* 2019, **21**(6):e14199.
- 1037 65. Grant JE, Chamberlain SR: **Sleepiness and impulsivity: Findings in non-treatment seeking young**
1038 **adults.** *Journal of behavioral addictions* 2018, **7**(3):737-742.
- 1039 66. Van Veen M, Karsten J, Lancel M: **Poor sleep and its relation to impulsivity in patients with**
1040 **antisocial or borderline personality disorders.** *Behavioral Medicine* 2017, **43**(3):218-226.
- 1041 67. Fineberg S, Leavitt J, Deutsch-Link S, Dealy S, Landry C, Pirruccio K, Shea S, Trent S, Cecchi G,
1042 Corlett P: **Self-reference in psychosis and depression: a language marker of illness.**
1043 *Psychological medicine* 2016, **46**(12):2605.
- 1044 68. Bucci W, Freedman N: **The language of depression.** *Bulletin of the Menninger Clinic* 1981,
1045 **45**(4):334.
- 1046 69. Hswen Y, Gopaluni A, Brownstein JS, Hawkins JB: **Using Twitter to Detect Psychological**
1047 **Characteristics of Self-Identified Persons With Autism Spectrum Disorder: A Feasibility Study.**
1048 *JMIR Mhealth Uhealth* 2019, **7**(2):e12264.
- 1049 70. Kessler RC, Chiu WT, Demler O, Walters EE: **Prevalence, severity, and comorbidity of 12-month**
1050 **DSM-IV disorders in the National Comorbidity Survey Replication.** *Archives of general*
1051 *psychiatry* 2005, **62**(6):617-627.
- 1052 71. Kessler RC, Nelson CB, McGonagle KA, Liu J, Swartz M, Blazer DG: **Comorbidity of DSM-III-R**
1053 **major depressive disorder in the general population: results from the US National Comorbidity**
1054 **Survey.** *The British journal of psychiatry* 1996, **168**(S30):17-30.
- 1055 72. Boschloo L, van Borkulo CD, Rhemtulla M, Keyes KM, Borsboom D, Schoevers RA: **The Network**
1056 **Structure of Symptoms of the Diagnostic and Statistical Manual of Mental Disorders.** *PLoS One*
1057 *2015*, **10**(9):e0137621.
- 1058 73. Borsboom D, Cramer AO, Schmittmann VD, Epskamp S, Waldorp LJ: **The small world of**
1059 **psychopathology.** *PloS one* 2011, **6**(11):e27407.

74. Zhang R: **The stress-buffering effect of self-disclosure on Facebook: An examination of stressful life events, social support, and mental health among college students.** *Computers in Human Behavior* 2017, **75**:527-537.
75. De Choudhury M, Kiciman E, Dredze M, Coppersmith G, Kumar M: **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: 2016*; 2016: 2098-2110.
76. De Choudhury M, Kiciman E: **The language of social support in social media and its effect on suicidal ideation risk.** In: *Proceedings of the International AAAI Conference on Web and Social Media: 2017*; 2017.
77. Wojcik S, Hughes A: **Sizing up Twitter users.** *Washington, DC: Pew Research Center* 2019.
78. Baldwin JR, Caspi A, Meehan AJ, Ambler A, Arseneault L, Fisher HL, Harrington H, Matthews T, Odgers CL, Poulton R: **Population vs Individual Prediction of Poor Health From Results of Adverse Childhood Experiences Screening.** *JAMA pediatrics* 2021.
79. Marek S, Tervo-Clemmens B, Calabro FJ, Montez DF, Kay BP, Hatoum AS, Donohue MR, Foran W, Miller RL, Feczko E: **Towards reproducible brain-wide association studies.** *BioRxiv* 2020.
80. Mansueto AC, Wiers R, van Weert JC, Schouten BC, Epskamp S: **Investigating the Feasibility of Idiographic Network Models.** 2020.
81. Kelley S, Gillan C: **Within-subject changes in network connectivity occur during an episode of depression: evidence from a longitudinal analysis of social media posts.** 2020.
82. Fiesler C, Proferes N: **"Participant" perceptions of Twitter research ethics.** *Social Media+ Society* 2018, **4**(1):2056305118763366.
83. Chancellor S, Birnbaum ML, Caine ED, Silenzio VM, De Choudhury M: **A taxonomy of ethical tensions in inferring mental health states from social media.** In: *Proceedings of the conference on fairness, accountability, and transparency: 2019*; 2019: 79-88.
84. Ford E, Curlewis K, Wongkoblaph A, Curcin V: **Public opinions on using social media content to identify users with depression and target mental health care advertising: mixed methods survey.** *JMIR mental health* 2019, **6**(11):e12942.
85. Mellon J, Prosser C: **Twitter and Facebook are not representative of the general population: Political attitudes and demographics of British social media users.** *Research & Politics* 2017, **4**(3):2053168017720008.
86. Yang J, Morris MR, Teevan J, Adamic LA, Ackerman MS: **Culture matters: A survey study of social Q&A behavior.** In: *Fifth International AAAI Conference on Weblogs and Social Media: 2011*; 2011.
87. Resnik P, Armstrong W, Claudino L, Nguyen T, Nguyen V-A, Boyd-Graber J: **Beyond LDA: exploring supervised topic modeling for depression-related language in Twitter.** In: *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality: 2015*; 2015: 99-107.
88. Orabi AH, Buddhitha P, Orabi MH, Inkpen D: **Deep learning for depression detection of twitter users.** In: *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic: 2018*; 2018: 88-97.
89. Sun J, Schwartz HA, Son Y, Kern ML, Vazire S: **The language of well-being: Tracking fluctuations in emotion experience through everyday speech.** *Journal of Personality and Social Psychology* 2020, **118**(2):364.
90. Agrawal A, Fu W, Menzies T: **What is wrong with topic modeling?(and how to fix it using search-based se).** *arXiv preprint arXiv:160808176* 2016.
91. Greene D, O'Callaghan D, Cunningham P: **How many topics? stability analysis for topic models.** In: *Joint European conference on machine learning and knowledge discovery in databases: 2014*; Springer; 2014: 498-513.