

# Small but Mighty: Affective Micropatterns for Quantifying Mental Health from Social Media Language

**Kate Loveys**  
Qntfy

kate@qntfy.com

**Patrick Crutchley**  
Qntfy

patrick@qntfy.com

**Emily Wyatt**  
Qntfy

emily.wyatt@qntfy.com

**Glen Coppersmith**  
Qntfy

glen@qntfy.com

## Abstract

Many psychological phenomena occur in small time windows, measured in minutes or hours. However, most computational linguistic techniques look at data on the order of weeks, months, or years. We explore *micropatterns* in sequences of messages occurring over a short time window for their prevalence and power for quantifying psychological phenomena, specifically, patterns in affect. We examine affective micropatterns in social media posts from users with anxiety, eating disorders, panic attacks, schizophrenia, suicidality, and matched controls.

## 1 Introduction

Mental illness and suicide pose a significant public health problem. Each year approximately 800,000 people will die by suicide, and an estimated 16 million suicide attempts will occur (World Health Organization, 2013). Mental illness is a similarly widespread problem, affecting almost one in four people worldwide during the course of their lifetime (World Health Organization, 2013). Mental illness (including suicide) detrimentally affects quality of life, ranking as the fourth-largest contributor to disability-adjusted life years (Vigo et al., 2016). Moreover, five of the top twenty causes of global disease burden were from mental illness (Vigo et al., 2016). Little progress has been made over the past fifty years in terms of improving these figures (Franklin et al., 2016).

A key step to reducing the global burden of mental illness and suicide deaths is to ensure that early risk detection and intervention occur (Insel, 2009). Current systems of care struggle with scalability and measures of long term efficacy. Given recent advances in many industries by ubiq-

uitous technology and data science, many hold out hope that a similar revolution is possible in mental health. *Digital phenotyping*, where data from everyday interactions with digital devices like smartphones and computers can be turned into quantifiable signals of mental health, holds promise for providing the real-time data needed for these advances. Real-time analysis of dispositional and discrete situational factors could help clinicians predict the onset or exacerbation of symptoms or suicidal behaviors (Nelson et al., 2017). This would transcend analysis and open the possibility for data-empowered interventions.

Generally, computational linguistics uses techniques that examine significant portions of a user’s data, spanning a long period of time. The few exceptions still only examine subsets of the data on the order of days or weeks (Resnik et al., 2015; Coppersmith et al., 2016; Mitchell et al., 2015). However, there are meaningful psychological phenomena occurring at much smaller time scales that slip past current methods (Nelson et al., 2017). Micropatterns, inspired by Bryan et al. (in press), are intended to focus on this neglected time window on the order of hours, by analyzing consecutive social media posts within such a window.

Here we examine affective micropatterns in language produced by individuals with a self-reported diagnosis of mental illness, a panic attack or suicide history, and neurotypical controls. We evaluate the affective valence of sequences of three consecutive tweets produced by individuals in each user group to identify micropatterns characteristic of each group. We compared suicide, panic attack, and mental illness group micropatterns to those of neurotypical controls. We address two questions:

- [1] Are there meaningful signals in affective micropatterns relevant to mental health?
- [2] Do micropatterns hold more information than the labels that make up their components?

This paper is the first time that affective micropatterns are examined directly, rather than as a component of a more complex learning system. This is also the first time that the relative power of micropatterns is explored beyond suicide risk.

### 1.1 Why Social Media?

One particularly compelling and rich source of data for digital phenotyping is language. Language provides a window into the perception, cognition, and other psychological processes at work in a person, and thus provides a useful lens through which we can understand, quantify, and eventually improve mental health. Social media, in particular, provides a trove of language data in a form conducive to computational analysis. Critically for this work, it also includes the time that a particular piece of language was authored by the user. Social media is, thus, one data source through which the early signs of mental illness and suicide can be detected (Reece et al., 2016; Coppersmith et al., 2016; Bryan et al., in press). Quantifiable signals for a wide range of behavioral health conditions have been uncovered recently, and this provides a foothold into analysis and intervention empowered by data science. A wide array of conditions have been studied including major depressive disorder (Chung and Pennebaker, 2007; De Choudhury et al., 2013), post-traumatic stress disorder (Coppersmith et al., 2014b, 2015b; Resnik et al., 2015; Preotiuc-Pietro et al., 2015; Pedersen, 2015), schizophrenia (Mitchell et al., 2015), eating disorders (Walker et al., 2015; Chancellor et al., 2016), generalized anxiety disorder, bipolar disorder (Coppersmith et al., 2014a), suicide (Coppersmith et al., 2015c; Kumar et al., 2015; Wood et al., 2016; Kiciman et al., 2016), borderline personality disorder, and others (Coppersmith et al., 2015a).

### 1.2 Social Media Micropattern Analysis

Micropatterns in short sequences of emotion, cognition, behavior and symptoms relevant to specific psychological states may be evident in social media data, reflecting dynamic shifts in internal situational factors. Many social media users report enough personal information on public feeds to be able to capture brief shifts in behaviors, cognitions, emotions, and symptoms relevant to particular psychological states. This information has been used to assess whether a user is declining into a suicidal state (Bryan et al., in press). Bryan

et al. (in press) found that distinct micropatterns in content of social media posts were predictive of proximity to a suicide death. One month prior to a suicide death, a seesaw-like effect was observed between social media posts about a maladaptive coping behavior and a negative belief, and at one week prior to a suicide death, this negative relationship grows stronger. Bryan et al. (in press) detected micropatterns from human-labeled posts and a complex model informed by dynamic systems theory. Here, we complement this work by adding automation to the labeling and exploring the micropatterns directly, rather than embedded in a larger system. No prior research has evaluated micropatterns in social media post content for psychological disorders other than suicidality.

This technique of looking at short subsequent posts and the psychological phenomena present therein is relatively new, so we aim for simplicity and straightforwardness in our experimental design and features. While there are a number of potentially more interesting avenues of exploration involving fine-grained emotions, psychologically meaningful events, coping mechanisms, and decompensation, we eschew the added complexity in favor of exploring a fundamental unanswered question: Is there meaningful signal in the micropatterns relevant to mental health?

### 1.3 Symptom Dynamics

Broadly, the motivation for exploring micropatterns and data on the timescale of minutes and hours stems from the importance of temporal information in the assessment of psychological symptoms. Knowledge of symptom co-occurrence over specified time periods can determine whether a mental illness diagnosis is received, as well as inform assessments of treatment responsiveness and relapse (American Psychiatric Association, 2013; Nelson et al., 2017). Temporal information is essential to detecting ongoing fluctuations in psychological symptoms, which may be key to predicting the onset of psychological disorders or increased suicide risk (McGorry and van Os, 2013).

Emotions, behavior, and cognitions fluctuate rapidly as an individual interacts with the environment (van Ockenburg et al., 2015; van Os, 2013). People have tendencies to behave, think, or feel certain ways, however, conditions and interactions fluctuate and one might have a markedly differ-

ent reaction to the same environment on a different day. These brief shifts in behaviors, emotions, cognitions, and physical symptoms relative to one another in an environment, over the course of seconds to hours, can determine a person's present-moment psychological state (van Os, 2013). The Fluid Vulnerability theory encapsulates this idea, suggesting that daily perturbations in situational factors interact with dispositional factors to trigger present-moment psychological states (Rudd, 2006). Dispositional (or distal) factors establish baseline risk, and are relatively fixed variables such as demographics, trait characteristics, beliefs or life histories, which tend to indicate stable predispositions toward experiencing particular psychological states or disorders. Conversely, situational (or proximal) factors indicate the likelihood that a person experiences a mental illness episode or engages in self-harming behavior at a specific point in time. Examples could include events such as the onset of a troubling thought or an unpleasant social interaction in the workplace. The Fluid Vulnerability theory suggests that for individuals with low baseline risk, even a severe stressor will not elicit suicidality or exacerbations in mental illness symptoms; alternatively, for people with high baseline risk, situational factors conducive to suicidality or mental illness episodes need not be as high for an episode to be triggered (Rudd, 2006). Most work at the intersection of natural language processing and social media has focused on assessing dispositional factors through examination of a large corpus of posts. However, assessing more situational risk factors will require a different set of methods. While existing bag of words approaches evaluate dispositional risk factors, temporal analyses are necessary to detect brief fluctuations in situational risk factors.

## 2 Data

We briefly explain the data collection method here, but we refer the interested reader with further questions on the methodology to Coppersmith et al. (2016) for the suicide attempt data and Coppersmith et al. (2014a) for all other conditions.

The data for these analyses are Twitter posts collected via two methods. Most of the data come from users who have publicly discussed their mental health conditions. These users are frequently referred to as “self-stated diagnosis” users, as they state publicly something like “I was diagnosed

with schizophrenia”, or “I’m so thankful to have survived my suicide attempt last year”. The data for users with a suicide attempt was supplemented by data from `OurDataHelps.org`, a data donation site where people provide access to their public posts and fill out a short questionnaire about their mental health history. Data are then de-identified and made available to researchers addressing questions of interest to the mental health community. Donors provide consent for their data to be used in mental health research upon sign-up. Of the users who attempted suicide, 146 came from `OurDataHelps.org`.

Specifically, we examine generalized anxiety disorder, eating disorders, panic attacks, schizophrenia, and attempted suicides. These conditions were selected based on the theory that there are important timing aspects to their symptoms – ebbing and flowing of symptoms as treatment is effective (especially schizophrenia), onset and exacerbation of symptoms by external events and stress, and punctuated events in time of psychological symptoms (suicide attempts, panic attacks, and bingeing/purging behavior with eating disorders).

We use the Twitter streaming API to collect a sample of users who used a series of mental health words or phrases in their tweet text (e.g., ‘schizophrenia’ or ‘suicide attempt’). Each tweet that uses one of these phrases is examined via regular expression to indicate that the user is talking about themselves. Finally, those tweets that pass the regular expression are examined by a human to confirm (to the best of our ability) that their self-statement of diagnosis appears to be genuine.

This results in a dataset with users that have a self-stated diagnosis of generalized anxiety disorder ( $n = 2408$ ), an eating disorder (749), panic attacks (263), schizophrenia (350), or someone who would go on to attempt suicide (423). Some of these users do not exhibit the sort of posting behavior required to create micropatterns (i.e., they rarely post multiple times within a 3 hour time window). We exclude these users from our analysis, which is 5-9% of users for most conditions, with the exception of those with a suicide attempt, where a little over half the users do not exhibit this posting behavior. The resultant dataset used for analyses is: generalized anxiety disorder ( $n = 2271$ ), eating disorders (687), panic attacks (247), schizophrenia (318), suicide attempts (157).

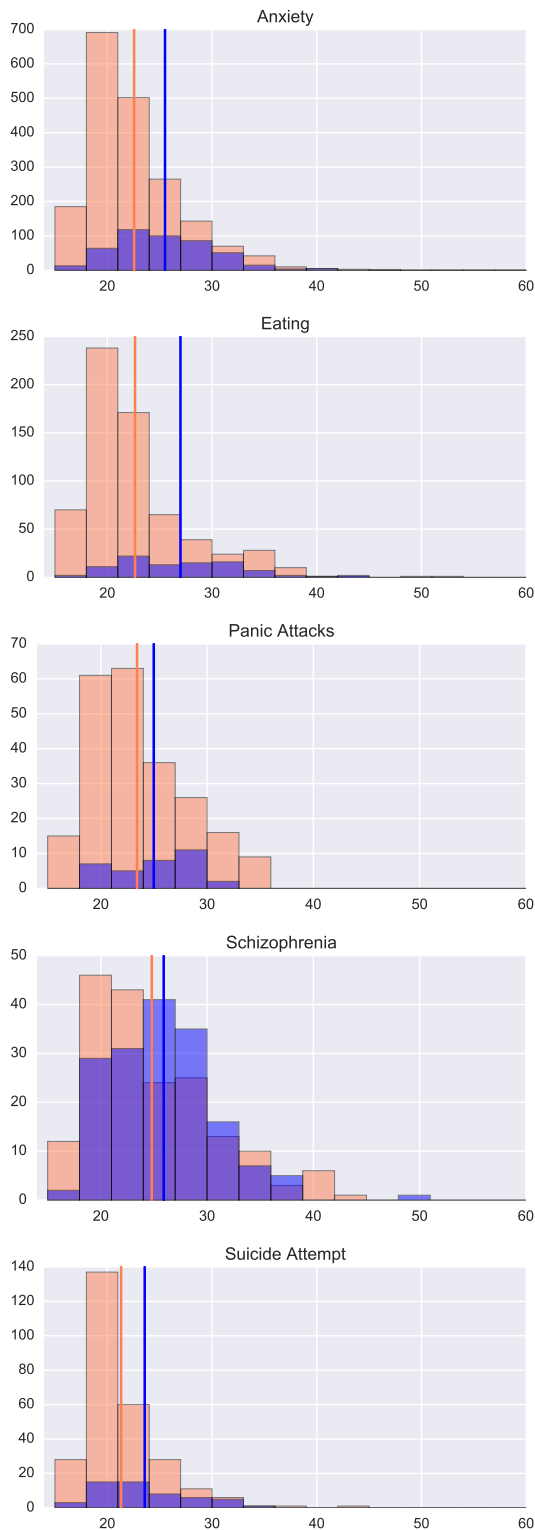


Figure 1: Histograms of age distributions for each condition. Females are in coral and males are in blue. The mean of each gender is denoted by the corresponding vertical line.

In order to allow comparisons of each condition to control users, we gather a random sample of 10,000 Twitter users for whom at least 75% of

their posts are identified by Twitter as English. All the users with a self-stated diagnoses and all members of this control population have their age and gender estimated according to Sap et al. (2014). For each user with a self-stated diagnosis, we find a matched control through the following procedure: create a pool of users where the estimated gender matches and the estimated age is within the same 10-year bracket (the suggested accuracy of the age estimator). From that pool of age- and gender- matched users, we select the user whose tweets start and end over the most similar time-frame. We will refer to these age-, gender-, and time-matched controls simply as “matched controls” throughout the rest of this paper.

All tweets were publicly posted by their author (i.e., no users marked at “protected” or “private” were included). On average, users had 2949 tweets. The distribution of estimated age and genders for users with each self-stated condition can be seen in Figure 1. For most conditions, the population skews female, though for schizophrenia the genders are roughly balanced. The average age tends to be in the early-to-mid 20s.

## 2.1 Caveats

All of the following analysis is subject to a few caveats emergent from the data and how the data were collected. The users with mental health conditions are all found data of one sort or another, so there are some inherent biases. We prefer to express these biases rather than add complexity by attempting to cleverly correct for them. Many of these users talk publicly about their mental health, which given the stigma and discrimination they face, is likely a distinct subpopulation of those with mental health conditions. It is possible that users with a psychological disorder or suicide history who did not publicly disclose this information could have been included in the control group for analyses, which may have the effect of artificially lowering the estimated power of any emergent differences. Users who donated data through OurDataHelps.org are likely biased differently, with over representation of altruism, since they are willing to do things for the public good without any obvious self gain. Another consideration is that all users who reported a suicide attempt within our dataset survived. There is a possibility that characteristic differences also exist between individuals who do and do not die by




Figure 2: Histograms of micropattern per day average for each condition. The median for each condition is denoted by the vertical blue line. Note that neurotypicals generally generate micropatterns at a rate lower than the mental health conditions, with the exception of users who would go on to attempt suicide.

a suicide attempt. Note that this research was conducted on English-speaking social media users. The content of social media post micropatterns for psychological disorders and suicidality could differ between cultural contexts, due to differences in cross-cultural expressions of mental illness (Chentsova-Dutton et al., 2007). These are active Twitter users, which imparts a demographic skew compared to the rest of the world (in particular, these users skew young). We see more females in our user populations than the rough gender balance observed for general Twitter users (Greenwood et al., 2016). The language data itself is meant for public consumption, and may reflect how the authors wish to be perceived, and not what one would get from a more traditional journal study of internal and private thoughts and feelings. Finally, we **include** users who had a concomitant or comorbid mental health condition. Thus a small number of users appear in more than one category.

### 3 Methods

This study aimed to examine the prevalence of affective *micropatterns* in social media posts and highlight differences in micropattern occurrence that might be relevant to quantifying mental health. Primarily, we do this through comparison of users with anxiety disorders, eating disorders, schizophrenia, suicide attempt history, and their matched controls.

We use a straightforward and well-understood method for sentiment analysis, VADER (Hutto and Gilbert, 2014), to produce a trinary label for each message: *positive*, *neutral*, or *negative*. VADER outputs a  $[0, 1]$  score for each sentiment label; we use the label with the maximum score.

Specifically, we examined trajectories of posted emotional content in three subsequent tweets, no more than three hours from earliest to latest. The same tweet will be counted in more than one over-

lapping micropattern if more than three tweets occur in the three-hour time window – so if 5 tweets occur in 3 hours, 3 micropatterns will be recorded from those 5 tweets, likewise for 4 tweets, 2 micropatterns will be recorded. The potential overlap exists for both patients and neurotypical users, and subsequent analyses (e.g., classifying users based on proportion of micropatterns) were designed to be robust to this property of overlapping micropattern generation. The number of sequential tweets to examine was chosen to minimize the complexity of the analysis while allowing significant variability to be observed. Critically, we aimed for the resulting dimensions (i.e., number of distinct micropatterns) to be small enough for meaningful interpretation by clinical psychologists.

## 4 Results

Our results collectively suggest that (1) micropatterns are not random (2) there are some significant differences in the occurrence of micropatterns between users who have a given mental health condition and their matched controls and (3) there is some quantifiable predictive power for separating users with mental health conditions from their matched controls captured by the micropatterns, in excess of what power the labels that underlie the micropattern have alone.

### 4.1 Micropatterns are not Randomly Distributed

Before any analysis of the differences in micropattern occurrence between users with mental health conditions and their matched controls, we demonstrate that these micropatterns are not randomly distributed, nor are they an artifact of the different base rate of users with mental health conditions expressing negative sentiment more often.

Previous work indicates that there are some expected variability in the proportion of messages from users in each condition, and significantly different from their matched control users (Coppersmith et al., 2015a). Specifically, it has been widely reported that users with certain behavioral health conditions use more words from the LIWC category `Negative Emotion` (Chung and Pennebaker, 2007; Park et al., 2012; De Choudhury et al., 2012; Coppersmith et al., 2015a), which in this case would have the effect of inflating the number and proportion of micropatterns involving `negative` labels, simply because

the prevalence of these labels were higher.

For each condition, we observe the distribution of labels for all messages from each condition. This establishes the base rate of each label occurring for that condition. Using these base rates, we randomly generate a label for each message from each user according to the base rate (i.e., respecting the timestamps of each post, but randomly assigning a label rather than what VADER predicted from the text). We then, for each user, examine the observed micropatterns with these randomly-assigned labels. We repeat this procedure 10,000 times, thus providing a null distribution of what we would expect the number and proportion of micropatterns to be if the underlying sentiment labels were randomly distributed. When we compare the observed value from real data to this randomly-generated population, the differences are stark and large. The observed  $z$ -scores for each micropattern’s deviation from normal range from 13.3 to 423859.1, with a median of 895.5. Since the significance for a  $z$ -score (at the  $p < 0.05$  level) is 1.96, we can safely assume that the observed population of labels was not likely the result of a random process. This strongly suggests that the differences observed are not attributable merely to random fluctuations and a different base-rate of the underlying labels.

### 4.2 Differences in Micropatterns

Figure 3 shows the deviation in each micropattern for users with mental health conditions relative to their matched neurotypical controls. This, taken with significant differences observed in matched-sample  $t$ -tests (omitted for brevity), clearly indicates that there are significant differences in micropatterns for a range of mental health conditions. While there are some observed similarities between the changes in micropatterns across conditions, significant differences exist between the various mental health conditions and their deviations from controls.

Note that the vast majority of the micropatterns observed in all conditions ( $> 80\%$ ) are (`neutral, neutral, neutral`). This is likely an overestimate of the number of neutral messages present, due to the closed-vocabulary nature of our lexicon-based labeling approach. Specifically, VADER depends on a lexicon of words and associated scores, and lexicon-based approaches generally provide higher precision




Figure 3: Change in micropattern frequency relative to age-, gender-, and time-matched controls for each condition. Red cells indicate lower frequency in users with a given mental health condition versus neurotypical, blue cells indicate higher frequency in users with a mental health condition versus neurotypical. Emoticons below the columns indicate the patterns in sentiment: far left is (negative, negative, negative), second to left is (negative, negative, neutral), and far right is (positive, positive, positive).

(i.e., fewer false alarms, which means fewer neutral messages tagged as valenced) at the cost of significantly decreased recall (i.e., many valenced messages are tagged as neutral). This is exacerbated by the fact we are scoring individual tweets, which contain relatively few words. Thus, while there are often some parameters to adjust around the sensitivity of classifiers, the combination of the lexicon approach and the short document makes for a very sparse set of features to score from. In turn, this tends to create more neutral labeled messages.

Some observed deviations line up with current psychological literature, providing some face validity to this approach. First, all mental health conditions show an increase in the number of (negative, negative, negative) affect micropatterns. This is consistent with the widely-found phenomenon that those with mental health conditions tend to experience greater negative affect (Chung and Pennebaker, 2007; Park et al., 2012; De Choudhury et al., 2012; Coppersmith et al., 2015a). This does suggest, though, that these are not necessarily randomly distributed negative posts, but in fact they are more likely to have concentrated and subsequent strings of negative posts. Second, users with schizophrenia were less likely than neurotypicals to show affect or affective variability between posts. This reflects research suggesting that individuals with schizophrenia display deficits in affective expression; a common negative symptom triggered by both disease pathophysiology and use of

antipsychotic medication (Messinger et al., 2011). Third, we see increases in affective volatility by users prior to a suicide attempt (as evidenced by (positive, negative, positive) and (negative, positive, negative) micropatterns, consistent with many as-of-yet unpublished findings from the Jelenik Summer Workshop at Johns Hopkins University (Hollingshead et al., in prep.). Fourth, users with an anxiety disorder were less likely than neurotypical controls to post consecutive positively-valenced tweets. This may be reflective of a negative attentional bias often associated with anxious emotion (Bar-Haim et al., 2007).

### 4.3 Separating Users

We also aim to understand if micropatterns convey some additional information about mental health and mental health status, above and beyond the labels that go into the micropattern (in this case, positive, negative, and neutral sentiment labels). Ideally, we would examine how well micropatterns could predict meaningful psychological events, but we lack significant data to do this more than anecdotally. Instead, we continue in line with previous work and compare performance on a binary prediction task. The task is to separate users with mental health conditions from their matched controls. Rather than examining absolute performance of this task as if it were a real world scenario, we aim to examine the *relative* performance of the micropatterns, the underlying sentiment labels, and a combination of the two, as a




Figure 4: Prediction accuracy for separating users with mental health conditions from their matched controls by base rate occurrence of sentiment labels alone (blue) and occurrence of micropatterns alone (green) and both features together (coral). Chance is 0.5 and is denoted by a black dotted vertical line.

way of assessing how much unique information the micropatterns themselves impart<sup>1</sup>.

For each user, we created a feature vector where each entry was the proportion of micropatterns that a particular micropattern made up. Similarly, we made a feature vector for the proportion of sentiment labels that each sentiment label made up (the base rate). Figure 4 shows the accuracy results of a 10-fold cross validation binary classification experiment (balanced samples) using a random forest classifier. In all cases, the micropatterns outperform the base rate, which is often little better than chance. In most cases, using both signals together (by concatenating the feature vectors) provides no significant gain in performance over either one alone. This suggests that for most conditions, most of the information from the sentiment labels are captured as part of the micropat-

<sup>1</sup>From an information theoretic perspective, it may be more appropriate to say how much information is lost by ignoring the ordering of the labels (in going from the micropatterns to simply the sentiment labels).

terns, but not all of it. Thus, we are led to conclude that micropatterns do provide additional information over the base rate of the sentiment labels alone.

## 5 Discussion

This paper presents foundational analysis of a relatively novel computational linguistic method that incorporates temporal information over short durations. Micropattern analysis provides information about common shifts in language content which may be useful for helping to distinguish between people with and without a psychological disorder or suicide risk. This study demonstrated that micropatterns in social media posts hold some power to distinguish between users who have a mental health condition or a history of suicide attempts or panic attacks from their matched controls.

Despite potential limitations, this study provides promising evidence in support of using micropattern analysis to detect progressions in suicide risk and symptoms of psychological disorders in future research. While the present study demonstrated that differences in micropatterns exist between users with and without a particular psychological disorder, information was not gathered on whether specific micropatterns can indicate the severity of a psychological disorder. We also did not assess whether micropatterns can distinguish between clinical conditions, and this is a likely next step for future research.

While there are a number of potentially more interesting avenues of exploration involving more fine-grained emotions, psychologically meaningful events, sleep disturbance, physical symptoms, coping mechanisms, decompensation, and their interplay, these bring with them an exponential complexity. We have done some preliminary examination of more fine-grained emotional labels, and found that interpretation and assessment was unwieldy and too complex for a reasonable human to undertake – 27 possible micropatterns are observed here (three labels, observed over three subsequent messages:  $3^3 = 27$ ). Extending this to the emotion classifier from [Coppersmith et al. \(2016\)](#), for example, would bring this to  $8^3 = 512$  micropatterns. Careful thought is required for analysis as the depth of possible labels grows.

Many avenues for future work seem apparent, as the veritable panoply of labels to augment the



straightforward VADER sentiment labels opens up. However, first and foremost of those possibilities is to directly replicate the work of Bryan et al. (in press) and extend it to non-military populations, and populations of different demographics to assess generalizability. This paper strongly suggests that micropatterns hold power for a wide range of mental health conditions, not just suicide risk. Specifically, including some of the known-relevant psychological phenomena that can be inferred from explicit self-reports seem a worthwhile next step, including: cognitive symptoms, physical symptoms, sleep disturbance, coping behavior, and suicidal thoughts and behavior.

Ultimately, technology is only a small part of the solution, since humans, workflows, and incentives that make up the existing system of care will need to integrate these technological solutions into their processes.

### 5.1 Ethics and Privacy

We gave careful consideration to the ethics and privacy surrounding this work, and employed the ethical guidelines from Benton et al. (2017), and used social media data donated with consent for use in mental health research from OurDataHelps.Org. We strongly encourage researchers interested in working in this space to consider the ethical implications from the outset, both of the research itself and also for the possible resultant technology. Recently, Mikal et al. (2016) conducted focus groups around their perception of this vein of work, which has greatly informed our work, and we heartily recommend it for informing ethical discussions.

## 6 Conclusion

We present evidence that quantifiable information relevant to mental health can be found in examining subsequent posts in relatively short order (so-called *micropatterns*). Furthermore, we demonstrate that even with a simple and straightforward lexicon approach, significant deviations in micropatterns can be found between users who have mental health conditions and their matched controls. While some of the observable differences have face validity and align with existing psychological literature, some remain unexplained. Moreover, micropatterns hold more predictive power than the sentiment labels that they rely upon, which suggests that they are capturing

important information not captured by the sentiment of the message alone. The results here were presented on simple and straightforward lexicon-based linguistic analysis, but the evidence strongly suggests that increasing the variety of psychologically meaningful (e.g., life changing events, coping mechanisms, decompensation) will lead to additional fruitful insights. Challenges remain about the sheer dimensionality of these more complex micropatterns, and how they should be best interpreted for synthesis with the psychological literature. While there is significant future work to understand why these micropatterns emerge and what value they hold for psychological understanding and intervention, we see this as a promising step, and a worthy avenue of future study.

## Acknowledgments

The authors would like to acknowledge the support of the 2016 Jelinek Memorial Workshop on Speech and Language Technology, at Johns Hopkins University, for providing the concerted time to perform this research. The authors would like to especially thank Craig and Annabelle Bryan for the inspiration for this work and the generosity with which they shared their time to mutually explore results. Finally and more importantly, the authors would like to thank the people who donated their data at [OurDataHelps.org](http://OurDataHelps.org) to support this and other research endeavors at the intersection of data science and mental health.

## References

- American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders (5th Edition)*. Arlington, VA: American Psychiatric Publishing.
- Yair Bar-Haim, Dominique Lamy, Lee Pergamin, Marian J Bakermans-Kranenburg, and Marinus H Van Ijzendoorn. 2007. Threat-related attentional bias in anxious and nonanxious individuals: a meta-analytic study. *Psychological bulletin* 133(1):1.
- Adrian Benton, Glen Coppersmith, and Mark Dredze. 2017. Ethical research protocols for social media health research. *EACL 2017* page 94.
- C. J. Bryan, J. E. Butner, S. Sinclair, A. O. Bryan, C. M. Hesse, and A. E. Rose. in press. Predictors of emerging suicide death among military personnel on social media networks. *Suicide and Life-Threatening Behavior*.

- Stevie Chancellor, Tanushree Mitra, and Munmun De Choudhury. 2016. Recovery amid pro-anorexia: Analysis of recovery in social media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, pages 2111–2123.
- Yulia E. Chentsova-Dutton, Joyce P Chu, Jeanne L Tsai, Jonathan Rottenberg, James J Gross, and Ian H Gotlib. 2007. [Depression and emotional reactivity: variation among Asian Americans of East Asian descent and European Americans](#). *Journal of Abnormal Psychology* 116(4):776–785. <https://doi.org/10.1037/0021-843X.116.4.776>.
- Cindy Chung and James Pennebaker. 2007. The psychological functions of function words. *Social Communication*.
- Glen Coppersmith, Mark Dredze, and Craig Harman. 2014a. Quantifying mental health signals in Twitter. In *Proceedings of the ACL Workshop on Computational Linguistics and Clinical Psychology*.
- 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 Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. North American Chapter of the Association for Computational Linguistics, Denver, Colorado, USA.
- 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 Shared Task for the NAACL Workshop on Computational Linguistics and Clinical Psychology*.
- Glen Coppersmith, Craig Harman, and Mark Dredze. 2014b. Measuring post traumatic stress disorder in Twitter. In *Proceedings of the 8th International AAAI Conference on Weblogs and Social Media (ICWSM)*.
- Glen Coppersmith, Ryan Leary, Eric Whyne, and Tony Wood. 2015c. Quantifying suicidal ideation via language usage on social media. In *Joint Statistics Meetings Proceedings, Statistical Computing Section*. JSM.
- Glen Coppersmith, Kim Ngo, Ryan Leary, and Tony Wood. 2016. Exploratory data analysis of social media prior to a suicide attempt. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. North American Chapter of the Association for Computational Linguistics, San Diego, California, USA.
- Munmun De Choudhury, Michael Gamon, and Scott Counts. 2012. Happy, nervous or surprised? Classification of human affective states in social media. In *Proceedings of the 6th International AAAI Conference on Weblogs and Social Media (ICWSM)*.
- Munmun De Choudhury, Michael Gamon, Scott Counts, and Eric Horvitz. 2013. Predicting depression via social media. In *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media (ICWSM)*.
- Joseph C Franklin, Jessica D Ribeiro, Kathryn R Fox, Kate H Bentley, Evan M Kleiman, Xieying Huang, Katherine M Musacchio, Adam C Jaroszewski, Bernard P Chang, and Matthew K Nock. 2016. Risk factors for suicidal thoughts and behaviors: A meta-analysis of 50 years of research. .
- S Greenwood, A Perrin, and M Duggan. 2016. Social media update 2016: Facebook usage and engagement is on the rise, while adoption of other platforms holds steady.
- Kristy Hollingshead, H Andrew Schwartz, Glen Coppersmith, Fatemeh Almoradasi, Adrian Benton, Jeff Craley, Patrick Crutchley, Dirk Hovy, Molly Ireland, Bu Sun Kim, Leo Kim, Raina Merchant, Margaret Mitchell, Phillip Resnik, Masoud Rouhizadeh, and Lyle Ungar. in prep. Detecting risk and protective factors of mental health using social media. *Center for Language and Speech Processing Technical Reports*.
- Clayton J Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International AAAI Conference on Weblogs and Social Media*.
- Thomas R Insel. 2009. Translating scientific opportunity into public health impact: a strategic plan for research on mental illness. *Archives of General Psychiatry* 66(2):128–133.
- Emre Kiciman, Mrinal Kumar, Glen Coppersmith, Mark Dredze, and Munmun De Choudhury. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM.
- Mrinal Kumar, Mark Dredze, Glen Coppersmith, and Munmun De Choudhury. 2015. Detecting changes in suicide content manifested in social media following celebrity suicides. In *Proceedings of the 26th ACM conference on Hypertext and hypermedia*. ACM.
- Patrick McGorry and Jim van Os. 2013. Redeeming diagnosis in psychiatry: timing versus specificity. *The Lancet* 381(9863):343–345.
- Julie W. Messinger, Fabien Trémeau, Daniel Antonius, Erika Mendelsohn, Vasthie Prudent, Arielle D. Stanford, and Dolores Malaspina. 2011. [Avolition and expressive deficits capture negative symptom phenomenology: Implications for DSM-5 and schizophrenia research](#). *Clinical Psychology Review* 31(1):161–168. <https://doi.org/10.1016/j.cpr.2010.09.002>.

- Jude Mikal, Samantha Hurst, and Mike Conway. 2016. Ethical issues in using twitter for population-level depression monitoring: a qualitative study. *BMC medical ethics* 17(1):1.
- Margaret Mitchell, Kristy Hollingshead, and Glen Coppersmith. 2015. Quantifying the language of schizophrenia in social media. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. North American Chapter of the Association for Computational Linguistics, Denver, Colorado, USA.
- Barnaby Nelson, Patrick D McGorry, Marieke Wichers, Johanna TW Wigman, and Jessica A Hartmann. 2017. Moving from static to dynamic models of the onset of mental disorder: A review. *JAMA psychiatry*.
- Minsu Park, Chiyong Cha, and Meeyoung Cha. 2012. Depressive moods of users portrayed in Twitter. In *Proceedings of the ACM SIGKDD Workshop on Healthcare Informatics (HI-KDD)*.
- Ted Pedersen. 2015. Screening Twitter users for depression and PTSD with lexical decision lists. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. North American Chapter of the Association for Computational Linguistics, Denver, Colorado, USA.
- Daniel Preotiuc-Pietro, Maarten Sap, H. Andrew Schwartz, and Lyle Ungar. 2015. Mental illness detection at the World Well-Being Project for the CLPsych 2015 shared task. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. North American Chapter of the Association for Computational Linguistics, Denver, Colorado, USA.
- Andrew G. Reece, Andrew J. Reagan, Katharina L. M. Lix, Peter Sheridan Dodds, Christopher M. Danforth, and Ellen J. Langer. 2016. [Forecasting the onset and course of mental illness with Twitter data.](https://arxiv.org/abs/1608.07740) *arXiv:1608.07740 [physics]* ArXiv: 1608.07740. <http://arxiv.org/abs/1608.07740>.
- Philip Resnik, William Armstrong, Leonardo Claudino, Thang Nguyen, Viet-An Nguyen, and Jordan Boyd-Graber. 2015. The University of Maryland CLPsych 2015 shared task system. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. North American Chapter of the Association for Computational Linguistics, Denver, Colorado, USA.
- M. David Rudd. 2006. [Fluid Vulnerability Theory: A Cognitive Approach to Understanding the Process of Acute and Chronic Suicide Risk](#). In Thomas E. Ellis, editor, *Cognition and suicide: Theory, research, and therapy.*, American Psychological Association, Washington, pages 355–368. DOI: 10.1037/11377-016. <http://content.apa.org/books/11377-016>.
- Maarten Sap, Greg Park, Johannes C. Eichstaedt, Margaret L. Kern, David J. Stillwell, Michal Kosinski, Lyle H. Ungar, and H. Andrew Schwartz. 2014. Developing age and gender predictive lexica over social media. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pages 1146–1151.
- Sonja L. van Ockenburg, Sanne H. Booij, Harriette Riese, Judith G. M. Rosmalen, and Karin A. M. Janssens. 2015. [How to assess stress biomarkers for idiographic research?](https://doi.org/10.1016/j.psyneuen.2015.08.002) *Psychoneuroendocrinology* 62:189–199. <https://doi.org/10.1016/j.psyneuen.2015.08.002>.
- Jim van Os. 2013. The dynamics of subthreshold psychopathology: implications for diagnosis and treatment.
- Daniel Vigo, Graham Thornicroft, and Rifat Atun. 2016. [Estimating the true global burden of mental illness.](https://doi.org/10.1016/S2215-0366(15)00505-2) *The Lancet Psychiatry* 3(2):171–178. [https://doi.org/10.1016/S2215-0366\(15\)00505-2](https://doi.org/10.1016/S2215-0366(15)00505-2).
- Morgan Walker, Laura Thornton, Munmun De Choudhury, Jaime Teevan, Cynthia M Bulik, Cheri A Levinson, and Stephanie Zerwas. 2015. Facebook use and disordered eating in college-aged women. *Journal of Adolescent Health* 57(2):157–163.
- Anthony Wood, Jessica Shiffman, Ryan Leary, and Glen Coppersmith. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM.
- World Health Organization. 2013. *Mental health action plan 2013-2020*. Geneva: World Health Organization.