

Detecting Anorexia in Spanish Tweets

Pilar López-Úbeda, Flor Miriam Plaza-del-Arco, Manuel Carlos Díaz-Galiano,
L. Alfonso Ureña-López, Maria-Teresa Martín-Valdivia

Department of Computer Science, Advanced Studies Center in ICT (CEATIC)
Universidad de Jaén, Campus Las Lagunillas, 23071, Jaén, Spain
{plubeda, fmplaza, mcdiaz, laurena, maite}@ujaen.es

Abstract

Mental health is one of the main concerns of today's society. Early detection of symptoms can greatly help people with mental disorders. People are using social networks more and more to express emotions, sentiments and mental states. Thus, the treatment of this information using NLP technologies can be applied to the automatic detection of mental problems such as eating disorders. However, the first step for solving the problem should be to provide a corpus in order to evaluate our systems. In this paper, we specifically focus on detecting anorexia messages on Twitter. Firstly, we have generated a new corpus of tweets extracted from different accounts including anorexia and non-anorexia messages in Spanish. The corpus is called SAD: Spanish Anorexia Detection corpus. In order to validate the effectiveness of the SAD corpus, we also propose several machine learning approaches for automatically detecting anorexia symptoms in the corpus. The good results obtained show that the application of textual classification methods is a promising option for developing this kind of system demonstrating that these tools could be used by professionals to help in the early detection of mental problems.

1 Introduction

Mental health is one of the main concerns of today's society. The World Health Organisation estimates that 1 in 4 individuals experience mental disorders at some stage of their lives. Globally, it is estimated that about 450 million people worldwide are mentally ill, with this kind of ill-

ness making up 13% of diseases around the world (Vos et al., 2015).

Traditionally, mental health evaluation is based on face-to-face interviews, self-reported issues or the distribution of questionnaires, which is usually labor-intensive and time consuming. However, in recent years several studies have used different technologies to improve the detection of mental health issues. Specifically, some interesting studies explore the relationship between data from online social networks and users' mental conditions (Rahman et al., 2018). Some of them focus on stress (Thelwall, 2017; Lin et al., 2017), depression (Tsugawa et al., 2015), suicide (O'Dea et al., 2015; Astoveza et al., 2018) or anxiety (Shen and Rudzicz, 2017), and most of them use and extract data from Twitter, probably because the information is open and more accessible than on other platforms, and also because it is one of the most popular social networks among young people. In this paper we focus on mental health problems related to eating disorders because they exhibit the highest mortality rate of any mental illness and 20% of all deaths from anorexia are the result of suicide (Arcelus et al., 2011).

Eating disorders are complex mental disorders considered serious and often fatal illnesses associated with severe disturbances in people's eating behaviors and related thoughts and emotions (Prieto et al., 2014). Common eating disorders include anorexia nervosa, bulimia nervosa, and binge-eating disorder and affect both females and males although they are most usual among young women.

The early detection of eating disorders can increase the chances of recovering, and technology can be applied to developing systems to help professionals. Different approaches to text and data mining methods can be applicable to social media data and may prove invaluable for health moni-

toring and surveillance. Specifically, Natural Language Processing (NLP), also known as Language Technologies (LT) can be used to generate systems for early anorexia detection. One of the main problems is the lack of resources to train systems and more if we focus on a language other than English.

The main goal of this paper is to develop a system for the automatic detection of anorexia in textual information. For this, we first generated a corpus with tweets written in Spanish including both people talking about anorexia and people talking about healthy food habits. The corpus is called SAD (Spanish Anorexia Detection). Using the SAD corpus, we have developed different models based on Machine Learning approaches that integrate several linguistic features. We have analyzed the results and compared the different approaches.

The rest of the paper is structured as follows: In Section 2 we comment on some related studies. The SAD corpus is described in Section 3, and present some interesting statistics. The different machine learning approaches and the results obtained are shown in Section 4. Finally, the analysis of errors is conducted in Section 5 and conclusions are presented in Section 6.

2 Related Work

The detection of mental health issues using textual information is a recent task mainly inspired by the massive use and access to social networks. People have become accustomed to using social networks to express all kinds of opinions, feelings and emotions. This valuable information can be captured and treated by an automatic system to learn how people with some health problems use language to express the frustration, depression or bad feelings. In this way, NLP can help to build systems to detect early on health problems such as eating disorders, depression or suicidal tendencies.

Although this task is relatively new, some challenging workshops and shared tasks related to the detection of health conditions have been proposed in recent years. For example, Social Media Mining for Health Applications (SMM4H) is a workshop and shared task that has been held since 2016 (Sarker et al., 2016) and continues every year. The main goal is to attract researchers interested in automatic methods for the collection, extraction, representation, analysis, and validation of so-

cial media data for health informatics. Furthermore, eRisk (Losada et al., 2017) is a challenging workshop focused on mental health disorders and it has been held from 2017 in the framework of the well-known international conferences CLEF¹. eRisk explores the evaluation methodology, effectiveness metrics and practical applications (particularly those related to health and safety) of early risk detection on the Internet. The different tasks proposed include depression and anorexia detection.

Concerning to mental health, we can find some interesting papers studying NLP techniques for treating textual information. (Rahman et al., 2018) review several studies focused on detecting mental health using and analyzing the information extracted from social networks. After analyzing several methods, machine learning algorithms, languages and sources of information, the authors conclude that machine learning is the most frequently used method used for mental health detection, with Support Vector Machine (SVM) presenting the best results. In addition, the study shows that Twitter is the major data source from social networks and English is the main language studied in the different papers. In (Prieto et al., 2014) four different health conditions are classified using machine learning methods over a corpus of tweets extracted by applying a set of crafted regular expressions. They integrate some relevant features in order to improve the final system. In addition, this is one of the few papers which center on languages other than English. Specifically, the authors work on Spanish and Portuguese tweets and the results indicate that the approach is a feasible option for tracking health information on social networks.

Regarding eating disorders, we can also find some recent studies. For example, (De Choudhury, 2015) focuses on detecting anorexia on the social network Tumblr using different affective, social, cognitive, and linguistic features. They also analyze the clinical implications of detecting anorexia related content on social media. (Chancellor et al., 2016a) use Instagram in order to study the eating disorders community and propose a statistical model combining topic modeling and clinical annotations. Finally, (Wang et al., 2017) center on Twitter generating a corpus by collecting eating disorders and non-eating disorders data. Then

¹<http://www.clef-initiative.eu/>

they train a SVM classifier, obtaining promising results. The high performance achieved suggests that it is feasible to design automatic text analysis tools that give early warnings of signs of eating disorders. However, this study only focuses on English and it is important to prove that the systems can also be applied to other languages. Thus, in this paper we create a Spanish corpus from Twitter with information concerning of anorexia and non-anorexia data. Then we apply several machine learning algorithms in order to demonstrate the feasibility of implementing systems to automatically detect sings of anorexia in Spanish messages written on social networks.

3 SAD Corpus

Anorexia and bulimia are two of the most worrisome eating disorders, affecting adolescents and young people the most. "Ana y mia" are the names used on the web pages that promote anorexia and bulimia to identify themselves. "Ana" is anorexia and "mia" is bulimia. But it is not a recent phenomenon, it began to become popular on the Internet in 2004 (Campos Rodríguez, 2007). Today, they have millions more pages and loyal followers, and the Internet has connected thousands of people with eating disorders. For this reason there are currently several studies of this disease (Moessner et al., 2018; Bermejo et al., 2011; Chancellor et al., 2016b). Specifically, for Spanish there is no set of Twitter messages concerning this problem, and for this reason we have compiled our own corpus, SAD (Spanish Anorexia Dataset) in order to accomplish the experiments.

3.1 Data Collection

We decided to use the social network Twitter because it is currently one of the most common sites for sharing information. This social network allows people to freely post short messages (called tweets) of up to 280 characters. Twitter has rapidly gained popularity worldwide, with more than 326 million active users generating more than 500 million tweets daily.

The task of downloading tweets has been performed through the Application Programming Interface (API) offered by Twitter. The API allows us to download messages using a query in a specific language. Our retrieving system always sets the option to Spanish, thus our classification system only works on tweets in Spanish. However,

our method can easily be adapted to other languages since the Twitter API allows specification of the language of the posts retrieved.

In order to obtain enough tweets, we had to download messages from past years, more concretely, in a date range of February 2014 to March 2019.

To make the corpus more interesting, we used as a query different hashtags related to food, nutrition, diet and healthy living in a converse way to anorexia. We collected data referring to anorexia using as query the hashtag *#anaymia* on Twitter. In addition, we collected three sets of reference data as negative samples using the hashtag *#realfood* *#comidareal* and *#fitness*.

Label 1 (anorexia) has been assigned to tweets that satisfy the query *#anaymia*, label 0 (control) for the other cases. Different messages are shown in Table 1 and in Table 2 we can see the English translation.

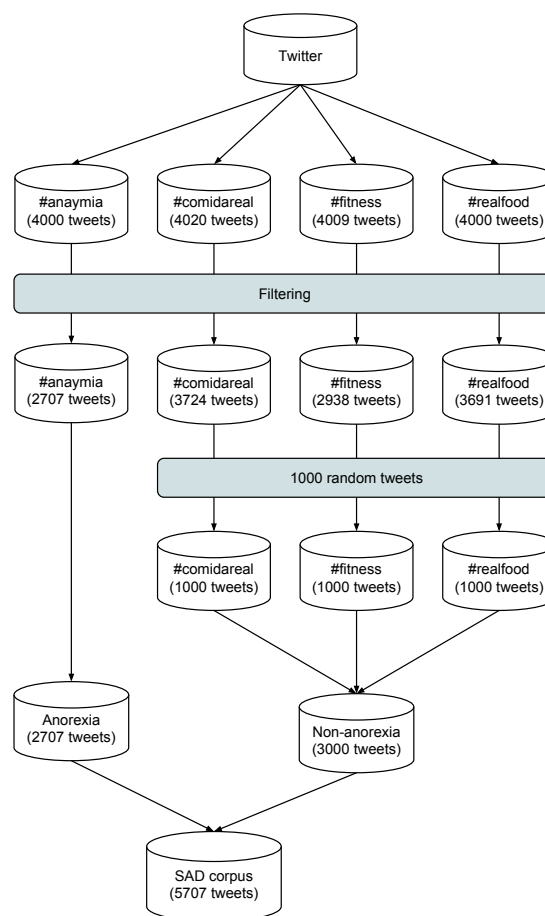


Figure 1: Process of generating the corpus from Twitter.

| <i>Tweet</i> | <i>Label</i> |
|---|--------------|
| Solo quiero llegar a mi casa a comer csm, no puedo más con esta hambre. Pero el hambre es belleza entrando a tu cuerpo. | 1 |
| La comida de hoy es ligera, muy ligera. Alcachofas al horno, simplemente llevan ajo, aceite, perejil y sal. Mmmm. #masendocrino #alcachofas #dietasana #dietamediterranea #aove #aceitedeoliva #hungry #cocinaespañola #comidacasera #foodpic #banquetesv | 0 |

Table 1: Examples of Spanish tweets tagged in SAD corpus.

| <i>Tweet</i> | <i>Tag</i> |
|---|------------|
| I just want to get home to eat, I can no longer cope with this hunger. But hunger is beauty coming into your body. | 1 |
| Today’s food is light, very light. Baked artichokes, simply with garlic, oil, parsley and salt. Mmmm. #masendocrino #alcachofas #dietasana #dietamediterranea #aove #oaceit-eoliva #hungry #cocinaespañola #comidacasera #foodpic #banquetesv | 0 |

Table 2: Example of translated tweets.

3.2 Data Filtering

Secondly, the extracted data is very noisy, so the set requires thorough cleaning before any analysis can be carried out. The language used by Twitter users contains some attributes that we had to remove to provide useful information for the classification process. This filtering that was performed:

- Repeat - the first filter to be performed was the removal of repeated tweets. Repeated tweets do not bring new information to the collection.
- Hashtag queries - we removed from the tweets the hashtag that we used as a query for downloading messages.
- All hashtag - we also removed tweets that only contained hashtags in the message. This step was followed since the experiments described in Section 4 were carried out without using hashtags.
- Short tweet - finally, tweets containing fewer than four words were removed since we consider that they do not provide enough representative information.

The objective was to obtain as balanced a corpus as possible. For cases of anorexia all tweets were incorporated. For the negative cases, we followed a different strategy, with 1000 random

tweets being taken from each hashtag (*#comidareal*, *#fitness* and *#realfood*), in this way, the corpus contains 2707 tweets annotated as positive (anorexia) and 3000 tweets annotated as negative (control). Figure 1 shows the number of tweets downloaded and how the collection decrements at each step.

3.3 Corpus Statistics

In this Section we will focus on obtaining statistics referring to the corpus containing relevant information. These statistics refer to the number of words, stopwords, hashtags, and part-of-speech tagging, among others.

The first study carried out consisted of obtaining the number of tweets, the number of words, the number of users and the number of stopwords in Spanish that exist in the corpus. This is shown in Table 3, where we can find the difference between the messages annotated with anorexia and those annotated as control.

It is interesting to see how the percentage increase in controlled tweets is 44% greater than the anorexia vocabulary, taking into account the number of total words as it can be seen in Table 3. But this information is reasonable because the average of tweet words is higher in controlled cases.

The grammatical labelling can be found in the Table 4. For this study we have used the spaCy²

²<https://spacy.io/>

| | <i>Total</i> | <i>Anorexia</i> | <i>Control</i> |
|--------------------------------|--------------|-----------------|----------------|
| Number of tweet | 5707 | 2707 | 3000 |
| Number of different users | 2585 | 1120 | 1466 |
| Number of total words | 122798 | 43179 | 79619 |
| Number of different words | 24635 | 8761 | 18515 |
| Average of tweet words | 21.52% | 15.95% | 26.54% |
| Number of total stop words | 30619 | 13118 | 17501 |
| Number of different stop words | 207 | 183 | 185 |
| Average of tweet stop words | 5.37% | 4.85% | 5.83% |

Table 3: Linguistic statistics on SAD corpus.

library with the module *es_core_news_sm*³. spaCy is a free open-source library for NLP in Python.

Table 4 shows the statistics obtained, and in it we can see relevant information on verbs, nouns, adjectives and adverbs. We found special interest in the high number of verbs and nouns used in annotated tweets without anorexia.

We wanted to obtain some statistics about the mood of users and how they express themselves through social networks. To obtain this information we used the resource iSOL (Molina-González et al., 2013). This resource has a list of opinion indicator words in Spanish independent of the domain. The list consists of 2,509 positive words and 5,626 negative. The results are described in Table 5. This table shows that users with anorexia problems use more negative language than users without anorexia. The same happens in the opposite case, whereby the tweets annotated as controlled are written with more positive words.

Finally, Table 6 shows some data about the use of hashtags in the messages collected. We can observe that the number of hashtags used in controlled tweets is much higher than on the contrary, and consequently there is also more variety of hashtags in messages annotated without anorexia.

4 Experiments and Results

In this section, we describe different experiments we carried out to test the validity of the SAD corpus. In particular, we trained several classifiers based on machine learning.

4.1 Pre-Processing

Pre-processing the data is the process of cleaning and preparing the text for classification. It is

³https://github.com/explosion/spacy-models/releases/tag/es_core_news_sm-2.1.0

one of the most important steps because it should help improve the performance of the classifier and speed up the classification process. Online texts usually contain lots a great deal of noise and uninformative parts which increases the dimensionality of the problem and hence makes the classification more difficult. For this reason, we applied pre-processing techniques in order to prepare the data for the text classification. In particular, we preprocessed the tweets of the SAD Dataset following these steps: The tweets were tokenized using NLTK TweetTokenizer⁴ and all hashtags were removed.

Features in the context of text classification are the words, terms or phrases that express the opinion of the author. They have a higher impact on the orientation of the text. There are several ways to assess the importance of each feature by attaching a certain weight to it in the text. We use the most popular: The Term Frequency Inverse Document Frequency scheme (TF-IDF). Specifically, using this scheme each tweet is represented as a vector of unigrams.

4.2 Machine Learning Algorithms

Machine learning techniques are popular in the binary classification. For this reason, we decide to employ different machine learning algorithms in order to classify the corpus in anorexic and non anorexic tweets. The algorithms are Support Vector Machine (SVM), Naive Bayes (NB), Random Forest (RF), Multilayer Perceptron (MLP), Logistic Regression (LR) and Decision Tree (DT).

4.3 Results

In this subsection, we report and discuss the performances of our systems on the Spanish anorexic

⁴<https://www.nltk.org/api/nltk.tokenize.html>

| | <i>Total</i> | <i>Anorexia</i> | <i>Control</i> |
|--|--------------|-----------------|----------------|
| Adjectives in corpus | 15332 | 3996 | 11336 |
| Nouns in corpus | 28594 | 8536 | 20058 |
| Verbs in corpus | 13647 | 5592 | 8055 |
| Adverbs in corpus | 5326 | 2518 | 2808 |
| Number of different adjectives in corpus | 4786 | 1493 | 3638 |
| Number of different nouns in corpus | 7326 | 2769 | 5449 |
| Number of different verbs in corpus | 4990 | 2342 | 3256 |
| Number of different adverbs in corpus | 622 | 296 | 455 |
| Average adjectives in tweet | 2.69% | 1.48% | 3.78% |
| Average nouns in tweet | 5.01% | 3.15% | 6.69% |
| Average verb in tweet | 2.39% | 2.07% | 2.69% |
| Average adverbs in tweet | 0.93% | 0.93% | 0.94% |

Table 4: Part-of-speech tagging statistics on SAD corpus.

| | <i>Total</i> | <i>Anorexia</i> | <i>Control</i> |
|--|--------------|-----------------|----------------|
| Negative words in corpus | 1549 | 1070 | 479 |
| Positive words in corpus | 2530 | 807 | 1723 |
| Different negative words in the corpus | 456 | 319 | 236 |
| Different positive words in the corpus | 460 | 227 | 358 |
| Average of negative words in tweet | 0.44% | 0.30% | 0.57% |
| Average of positive words in tweet | 0.27% | 0.40% | 0.16% |

Table 5: Statistics about positive and negative words in the corpus.

classification task on the SAD corpus. In order to evaluate and compare the results obtained by our experiments, we use the usual metrics in text classification, called precision (P), recall (R), F-score (F_1) and Accuracy (Acc).

To determine the optimal classification algorithm, we conducted experiments with the six classification models. We used 10-fold cross validation to evaluate the machine learning classification approaches including: The *SVM* classifier, the *Naive Bayes* classifier, the *Random Forest* classifier, the *Decision Tree* classifier, *Logistic Regression* and the *Multilayer Perceptron* classifier. The test results achieved by these algorithms on the SAD corpus are shown in Table 7.

The classifiers with the best performance were SVM and MLP with the default settings in the Scikit-learn 0.19.1 package (Pedregosa et al., 2011). The other classifiers also showed good results, all achieving an accuracy score superior to 80%. It should be noted that they performed well in both classes (anorexia and control) because the corpus is well balanced.

5 Error Analysis

The main purpose of this section is to carry out an error analysis to identify the weaknesses of our system. For this, we analyze some of the tweets not correctly classified by our system.

Of the total number of tweets (5707), 478 were not correctly classified, only 8.38% of the total tweets. In Figure 2, the confusion matrix of our system can be seen. It shows that there were more false positives (300) than false negatives (178). Therefore, we analyzed some of these tweets manually in order to find the main reasons why our system can be confused.

Table 2 presents some examples of tweets incorrectly classified by our system and Table 1 shows the corresponding translation. Specifically, two of the tweets are false positives and the other two false negatives. On the one hand, if we look at the false positives, we can see that two of the reasons why our system can be wrong is because it detects that there are words related to food and also that the vocabulary of the other tweets labeled as control is very similar to the vocabulary used in anorexia. Therefore, the system is sometimes con-

| | <i>Total</i> | <i>Anorexia</i> | <i>Control</i> |
|------------------------------|--------------|-----------------|----------------|
| Hashtags in corpus | 25133 | 5037 | 20096 |
| Different hashtags in corpus | 7479 | 1282 | 6341 |
| Average hashtags in tweet | 4.40% | 1.86% | 6.70% |

Table 6: Statistics about hashtag in the corpus.

| Classifier | Anorexia | | | Control | | | Macro-avg | | | Acc |
|------------|----------|-------|-------|---------|-------|-------|-----------|-------|-------|-------|
| | P | R | F_1 | P | R | F_1 | P | R | F_1 | |
| SVM | 0.894 | 0.934 | 0.914 | 0.938 | 0.9 | 0.919 | 0.916 | 0.917 | 0.916 | 0.916 |
| MLP | 0.894 | 0.934 | 0.913 | 0.938 | 0.9 | 0.92 | 0.916 | 0.917 | 0.916 | 0.916 |
| RF | 0.837 | 0.895 | 0.865 | 0.899 | 0.842 | 0.87 | 0.868 | 0.868 | 0.867 | 0.867 |
| NB | 0.823 | 0.849 | 0.835 | 0.859 | 0.835 | 0.847 | 0.841 | 0.842 | 0.841 | 0.841 |
| LR | 0.846 | 0.898 | 0.871 | 0.902 | 0.853 | 0.878 | 0.874 | 0.875 | 0.874 | 0.874 |
| DT | 0.795 | 0.823 | 0.809 | 0.835 | 0.809 | 0.822 | 0.815 | 0.816 | 0.815 | 0.815 |

Table 7: Results obtained by different classifiers on the SAD corpus (10-fold cross validation).

fused when, for example, the user talks about sport in general. On the other hand, if we focus on false negatives, we see that one of the problems is the irony in the tweet and another of the problems is when the user is transmitting a negative emotion but does not say the cause.




Figure 2: Confusion matrix.

6 Conclusion

This article presents a new corpus in Spanish for detecting anorexia in social network messages. Several systems are also developed to test the performance of this task with different classifiers. The results obtained show that the performance is very similar in all systems, although SVM and MLP are the only ones that obtain accuracy above 0.9.

Error analysis reveals that there are cases where

classification systems do not work properly. Our next goal will be to apply other techniques (such as irony detection or sentiment analysis) in cases where textual information is poor or where rhetorical figures such as irony and sarcasm are used.

Acknowledgments

This work has been partially supported by Fondo Europeo de Desarrollo Regional (FEDER), LIVING-LANG project (RTI2018-094653-B-C21) and REDES project (TIN2015-65136-C2-1-R) from the Spanish Government.

References

- Jon Arcelus, Alex J Mitchell, Jackie Wales, and Søren Nielsen. 2011. Mortality rates in patients with anorexia nervosa and other eating disorders: a meta-analysis of 36 studies. *Archives of general psychiatry* 68(7):724–731.
- Ghelmar Astoveza, Randolph Jay P Obias, Roi Jed L Palcon, Ramon L Rodriguez, Bernie S Fabito, and Manolito V Octaviano. 2018. Suicidal behavior detection on twitter using neural network. In *TEN-CON 2018-2018 IEEE Region 10 Conference*. IEEE, pages 0657–0662.
- Belén G Bermejo, Luis Ángel Saúl, and Cristina Jenaro. 2011. La anorexia y la bulimia en la red. ana y mia dos malas compañías para las jóvenes de hoy [the anorexia and bulimia on the web: Ana and mia two “bad company” for youth today]. *Acción psicológica* 8(1):71–84.
- José Miguel Campos Rodríguez. 2007. Anorexia, bulimia e internet. aproximación al fenómeno pro-ana

| Tweet | True label | Predicted |
|--|------------|-----------|
| ”El físico no importa” | 1 | 0 |
| Momento de escuchar música para relajarme y olvidarme de la mierda de mundo en el que vivo | 1 | 0 |
| Hola @IKEASpain mi bebé de 9 meses come sólido, no ha comido un potito nunca y me parece injusto que a él le cobréis la comida | 0 | 1 |
| Rutina de ejercicios para glúteos | 0 | 1 |

Table 8: Examples of tweets badly classified by our system.

| Tweet | True label | Predicted |
|--|------------|-----------|
| ”The physical aspect doesn’t matter” | 1 | 0 |
| Time to listen to music to relax and forget about the shitty world I live in | 1 | 0 |
| Hello @IKEASpain my 9 month old baby eats solid food, and has never eaten a baby food and I think it’s unfair that you charge him for the food | 0 | 1 |
| Buttock Exercise Routine | 0 | 1 |

Table 9: Examples of translated tweets badly classified by our system.

- y mía desde la teoría subcultural. *Frenia. Revista de Historia de la Psiquiatría* 7(1):127–144.
- Stevie Chancellor, Zhiyuan Lin, Erica L Goodman, Stephanie Zerwas, and Munmun De Choudhury. 2016a. 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*. ACM, pages 1171–1184.
- Stevie Chancellor, Jessica Annette Pater, Trustin Clear, Eric Gilbert, and Munmun De Choudhury. 2016b. # thyhgapp: Instagram content moderation and lexical variation in pro-eating disorder communities. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. ACM, pages 1201–1213.
- Munmun De Choudhury. 2015. Anorexia on tumblr: A characterization study. In *Proceedings of the 5th international conference on digital health 2015*. ACM, pages 43–50.
- Huijie Lin, Jia Jia, Jiezhong Qiu, Yongfeng Zhang, Guangyao Shen, Lexing Xie, Jie Tang, Ling Feng, and Tat-Seng Chua. 2017. Detecting stress based on social interactions in social networks. *IEEE Transactions on Knowledge and Data Engineering* 29(9):1820–1833.
- David E Losada, Fabio Crestani, and Javier Parapar. 2017. erisk 2017: Clef lab on early risk prediction on the internet: experimental foundations. In *International Conference of the Cross-Language Eval-uation Forum for European Languages*. Springer, pages 346–360.
- Markus Moessner, Johannes Feldhege, Markus Wolf, and Stephanie Bauer. 2018. Analyzing big data in social media: Text and network analyses of an eating disorder forum. *International Journal of Eating Disorders* 51(7):656–667.
- M Dolores Molina-González, Eugenio Martínez-Cámara, María-Teresa Martín-Valdivia, and José M Perea-Ortega. 2013. Semantic orientation for polarity classification in spanish reviews. *Expert Systems with Applications* 40(18):7250–7257.
- Bridianne O’Dea, Stephen Wan, Philip J Batterham, Alison L Calear, Cecile Paris, and Helen Christensen. 2015. Detecting suicidality on twitter. *Internet Interventions* 2:183–188.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Víctor M Prieto, Sérgio Matos, Manuel Álvarez, Fidel Casheda, and José Luís Oliveira. 2014. Twitter: A good place to detect health conditions. *PLoS ONE* 9:e86191.
- Rohizah Abd Rahman, Khairuddin Omar, Shahrul Azman Mohd Noah, and Mohd Shahrul Nizam Mohd

- Danuri. 2018. A survey on mental health detection in online social network. *International Journal on Advanced Science, Engineering and Information Technology* 8(4-2):1431–1436.
- Abeed Sarker, Azadeh Nikfarjam, and Graciela Gonzalez. 2016. Social media mining shared task workshop. In *Biocomputing 2016: Proceedings of the Pacific Symposium*. World Scientific, pages 581–592.
- Judy Hanwen Shen and Frank Rudzicz. 2017. Detecting anxiety through reddit. In *Proceedings of the Fourth Workshop on Computational Linguistics and Clinical Psychology—From Linguistic Signal to Clinical Reality*. pages 58–65.
- Mike Thelwall. 2017. Tensistrength: Stress and relaxation magnitude detection for social media texts. *Information Processing & Management* 53(1):106–121.
- Sho Tsugawa, Yusuke Kikuchi, Fumio Kishino, Kosuke Nakajima, Yuichi Itoh, and Hiroyuki Ohsaki. 2015. Recognizing depression from twitter activity. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, pages 3187–3196.
- Theo Vos, Ryan M Barber, Brad Bell, Amelia Bertozzi-Villa, Stan Biryukov, Ian Bolliger, Fiona Charlson, Adrian Davis, Louisa Degenhardt, Daniel Dicker, et al. 2015. Global, regional, and national incidence, prevalence, and years lived with disability for 301 acute and chronic diseases and injuries in 188 countries, 1990–2013: a systematic analysis for the global burden of disease study 2013. *The Lancet* 386(9995):743–800.
- Tao Wang, Markus Brede, Antonella Ianni, and Emmanouil Mentzakis. 2017. Detecting and characterizing eating-disorder communities on social media. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, pages 91–100.