

Question Answering for Suicide Risk Assessment using Reddit

Amanuel Alambo*, Manas Gaur*, Usha Lokala*, Ugur Kursuncu*, Krishnaprasad Thirunarayan*,
Amelie Gyrard*, Randon S. Welton[†], Jyotishman Pathak[‡] Amit Sheth*

* Kno.e.sis Center, Wright State University, Dayton, Ohio

{ammanuel, manas, usha, ugur, tkprasad, amelie, amit}@knoesis.org

[†]Department of Psychiatry, Wright State University, Dayton, Ohio
randon.welton@wright.edu

[‡]Cornell University, New York, NY
jyp2001@med.cornell.edu

Abstract—Mental Health America designed ten questionnaires that are used to determine the risk of mental disorders. They are also commonly used by Mental Health Professionals (MHPs) to assess suicidality. Specifically, the Columbia Suicide Severity Rating Scale (C-SSRS), a widely used suicide assessment questionnaire, helps MHPs determine the severity of suicide risk and offer an appropriate treatment. A major challenge in suicide treatment is the social stigma wherein the patient feels reluctance in discussing his/her conditions with an MHP, which leads to inaccurate assessment and treatment of patients. On the other hand, the same patient is comfortable freely discussing his/her mental health condition on social media due to the anonymity of platforms such as Reddit, and the ability to control what, when and how to share.

The popular “SuicideWatch” subreddit has been widely used among individuals who experience suicidal thoughts, and provides significant cues for suicidality. The timeliness in sharing thoughts, the flexibility in describing feelings, and the interoperability in using medical terminologies make Reddit an important platform to be utilized as a complementary tool to the conventional healthcare system. As MHPs develop an implicit weighting scheme over the questionnaire (i.e., C-SSRS) to assess suicide risk severity, creating a relative weighting scheme for answers to be automatically generated to the questions in the questionnaire poses as a key challenge.

In this interdisciplinary study, we position our approach towards a solution for an automated suicide risk-elicitation framework through a novel question answering mechanism. Our two-fold approach benefits from using: 1) semantic clustering, and 2) sequence-to-sequence (Seq2Seq) models. We also generate a gold standard dataset of suicide posts with their risk levels. This work forms a basis for the next step of building conversational agents that elicit suicide-related natural conversation based on questions.

Index Terms—Semantic Social Computing, Suicide Risk Assessment, Reddit, C-SSRS, The Diagnostic and Statistical Manual of Mental Disorders – Fifth Edition (DSM-5), Web-based Intervention, Semantic Machine Learning

I. INTRODUCTION

An article published by *PsychiatryAdvisor* in 2016 mentions 6 suicide assessment scales that can be used for suicide risk evaluation of patients¹. A large number of suicide risk evaluations are conducted through direct contact between an

MHP and a patient where the former performs assessments using rating scales or questionnaire [1]. While rating scales such as PHQ-9², or C-SSRS³ are widely used by the clinical community to assess suicide risk, they can be prone to subjective bias among MHPs. A commonly used suicide risk assessment scale, C-SSRS, does not adequately address completeness and relative weighting of questions which can lead to inaccurate risk evaluation [2]. MHPs use their personal experience to understand and give different levels of severity to patient responses to questions in their suicide risk evaluation task. Further, the public social stigma on the subject of suicide deters patients from disclosing their symptoms where a questionnaire is used in a direct MHP-to-patient interaction.

Reddit is an invaluable resource of symptom markers (e.g. signs of hopelessness, anger, anxiety) that can be used to augment the passive assessment of patients by MHPs [3]. Reddit has been used by individuals to freely share their symptoms, disorders, and experiences with no thoughts of social stereotypes. The fact users anonymously and candidly disclose their feelings provide the convenience of having timely healthcare signals [4]. One of these signals shared on Reddit is related to suicidal thoughts. This makes the data more reflective of the suicidal marks the user has experienced at the specified point in time making intervention effective. Thus, Reddit is an untapped resource for monitoring and studying suicide at a community level, offering benefits related to enforcing policies at a large scale. A framework that uses and analyzes data from Reddit enables MHPs to quantitatively assess individuals’ level of suicide risk. We address the challenge of rating the suicide risk of Reddit online posts as a Question Answering problem using sets of questions developed by practicing clinical psychiatrists. Since C-SSRS is developed for an MHP interacting with a patient and is limited in scope [2], we had our clinical psychiatrist collaborators produce a novel questionnaire that is complete and linguistically adaptable with Reddit content [5]. We approach the Reddit posts elicitation study using our suicide risk

¹<https://bit.ly/2rLLTTg>

²<https://bit.ly/2A7styu>

³<http://cssrs.columbia.edu/>

assessment questionnaire as text skimming problem using our novel semantic clustering technique. We build four semantic clusters where semantically similar sentences in a Reddit post and questions in a questionnaire are grouped together. The semantic clusters we build are inspired by the scope of the suicide lexicons used in this study (suicide indicator, suicide ideation, suicide behavior, suicide attempt). We weight our semantic clusters based on their suicide grades which are used in assessing the aggregate suicide risk severity of a Reddit post.

For our semantic clustering task, we use background knowledge sources built from SNOMED-CT [6], ICD-10 [7], UMLS [8], and DataMed [9] to determine the classification of a sentence in a Reddit post and the questions in our questionnaire. The background knowledge sources we used as a part of this study are Suicide Severity lexicons⁴(see footnote) and DSM-5 lexicons[10]. The key contributions of this study are: 1) an automated social media content elicitation framework based on semantic clustering and background knowledge; 2) a cohort of 16 questions developed by practicing clinical psychiatrists to assess suicide risk on social media; and 3) a gold standard dataset of 4992 Reddit suicide posts and their suicide risk level.

In this position paper, we begin with motivating scenarios where challenges to process social media text are discussed through examples on suicide and related mental disorders (Section II). We then briefly go through previous studies conducted that align with our work as stated in the Related Work (Section III). We proceed onto a detailed description of the data used in our two-fold approach (Section IV). Finally, we lay out the methodology to address the challenges stated in the motivating scenario (Section V). We conclude with the broader implications of our study (Section VI).

II. MOTIVATING SCENARIOS

Our study is motivated by examples that portray main challenges of working with social media text. The following examples describe social media data showing signals of mental health disorders related to suicide⁵. We picked the following four DSM-5 categories to show the examples of strong signals of mental health-related disorders in social media.

a) **Substance Use and Addictive disorder:** *“I’m ready to overdose. I recently got prescribed oxycodone for a car accident I was in a couple weeks ago. I have about 10 pills left...But my family deserve someone honest and not a piece of trash like myself.”*

b) **Depressive disorder:** *“My depression royally screwed up my relationship and I have never felt so alone. I want to die so badly.”*

c) **Disruptive, Impulse Control and Conduct disorder:** *“I live in anxiety and fear. I just want to let you know that I have suffered from anxiety, bipolar and depression myself and while I’m not saying in any way I know exactly how you feel,*

I know what it’s like to feel lost, helpless, I know what it’s like to have your house feel like it’s not a home, and I know what it’s like to be so very scared by the world around me, and I know what it’s like to feel like killing yourself is the option.”

d) **Schizophrenia Spectrum:** *“I’ve been taking Wellbutrin after getting off Prozac and Zoloft. And at first it felt great. Now I feel absolutely TERRIBLE I feel like I’m dreaming all the time Nothing feels real I don’t want to try I feel like a failure I want to give up I know my parents hate me. I hate living.”*

The below set of examples are challenges in understanding the context and real sense of suicide posts where just syntactic lexical match is not enough and demands more understanding in terms of temporality, and context. We also address NLP-specific issues such as negation detection, negation combined with antonyms, and sarcasm detection.

e) **Express Sarcasm/Sentiment:** *“I love my company but I suck at my job and it makes me miserable. I can’t even keep fucking plants alive, I fuck up everything I touch. Everywhere I go, I’m the problem. I just want to quit, break up, cancel and unsubscribe everything.”*

f) **Suicide tendency from Inspiration of Celeb Suicide and not directly about him/her (Werther Effect):** In this, a person is actually feeling suicidal not because of his actions or situation but inspired from a recent celebrity suicide and feels that makes him think of suicide. This kind of analysis is difficult to be done with simple lexical matches.

“I liked Linkin Park when I was younger, but never felt a connection with Chester and I still don’t. But the more I see it, the more I read about it, the more it makes me think of suicide myself.”

g) **Past Temporal ideation leading to present suicide behavior:** The time series of individual posts can be studied and analyzed in a way that past indicators lead to current behavior and may lead to a future attempt which is shown in the post below. In some cases, the individual is only talking about his/her childhood experiences and past actions which may not reflect current state and hence clinical visit may not be warranted. Often, a straightforward lexical analysis can misinterpret the context of the post.

“Holy shit dude I remember thinking how bizarre the idea of suicide was back when i was a kid and now I’m here longing it more than ever.”

h) **A strong sign of Suicide Attempt but still not about the person who posted:** The post below is a clear example of a case that if the question-answering automation is done purely based on knowledge sources to identify the suicide severity as attempt, critical things like a person addressing someone else is missed resulting in false positives.

“so my friend has this boyfriend...before they started dating he was planning to commit suicide. (he told me this after they started dating) and when she asked him out he felt like he owed her his life because he hasn’t wanted to do any self-harm or anything since and when my best friend set me up with my boyfriend the thoughts of suicide didn’t stop. so is my friend’s bf hiding something or am I just messed up?”

⁴<http://knoesis.org/node/2908>

⁵All subjects chosen are anonymous and not personally identifiable, and hence supports human subject protection.

III. RELATED WORK

There are various works on assessing suicide risk using different rating scales. Louzon et al. [11] used item-9 of PHQ-9 to measure suicide ideation for assessing suicide risk in VA patients. Runeson et al. [12] conducted assessments of diagnostic accuracy of 8 instruments for suicide risk grading and bias. Coppersmith et al. [13] built a machine learning model that classifies users who would go on to attempting suicide from those who would not. They built a text classification model that would predict the suicide risk of an individual based on the aggregate scores from many posts of an individual user. They leveraged Glove embeddings for representing individual words in a post. Bi-LSTM is used over sequences of word vectors to capture contextual information between words. Shing et al. [1] performed suicide risk assessment of Reddit online posts for specifically subreddit ‘SuicideWatch’. They created a dataset of users who posted on subreddit SuicideWatch and their suicide risk level. They defined a four-level classification of suicide risk level based on Corbitt-Hall definitions of risk categories [14]. De Choudhury et al. [15] employed statistical techniques to study the transitions from mental disorder to suicidal ideation on Reddit. They particularly built a model that predicts users on Reddit who have mental disorders and would later proceed onto developing suicidal ideation. Johnson et al. [16] used Twitter to identify suicide-warning signs for individuals. They employed the martingale framework to capture the behavioral changes present in Twitter posts through a combination of textual and behavioral features. Du et al. [17] used deep learning methods to detect psychiatric stressors leading to suicide. They built binary classifier for identifying suicidal tweets from non-suicidal tweets using Convolutional Neural Networks (CNNs). Once suicidal tweets are detected, they performed Named Entity Recognition (NER) using Recurrent Neural Networks for tagging psychiatric stressors in a tweet classified as suicidal.

IV. DESCRIPTION OF DATA AND RESOURCE

We discuss the dataset and background knowledge used to define the position of our study. Our dataset consists of 4992 posts of 500 Reddit users spanning over a period of 2005-2016. There are 28997 sentences in total in our dataset⁶. There are on average 6 sentences in a post and 58 posts per Reddit user. The tables below show the statistics of the lexicons (Suicide Lexicons (see Table I), and DSM-5 Lexicons (see Table II)) used for this study.

Further DSM-5 classes mentioned in Table II encompass various mental disorders [10]. Mental disorders such as Borderline Personality Disorder (BPD) are categorized under class DICD. Similarly, Bipolar Disorders are categorized under class Depressive disorder.

V. METHODOLOGY

We conducted our study by dividing the problem into three phases as shown in Figure 1. This architecture is inspired by

⁶<http://knoesis.org/node/2908>

Suicide Risk Class	Number of Entities	Sample Medical Phrases
Suicide Indicator	1472	Severe mood disorder with psychotic feature; severe major depression; Family history of suicide; Sedative
Suicide Ideation	409	Bipolar affective disorder; Borderline Personality; Depressive conduct disorder; Sexual maturation disorder
Suicide Behavior	145	Suicidal behavior; Intentional self-harm; Incomplete attempt; Threatening suicide
Suicide Attempt	123	Attempt actual suicide; Attempt physical damage; Intensive care; Second-degree burns

Table I: Suicide Lexicon

DSM-5 Category	Number of Entities	Sample Medical Phrases
Substance Use and Addictive Disorder	76	Psychoactive substance use disorder; Addictive behavior with potential to damage skin; Dependent drug abuse; Glue sniffing dependence
Depressive Disorder	148	Mild bipolar II disorder; most recent episode major depressive; Persistent somatoform pain disorder; Bipolar affective disorder
Disruptive Impulse Control and Conduct disorder	58	antisocial personality; attention deficit hyperactivity disorder; Stress reaction causing mixed disturbance of emotion and conduct; Conduct disorder in remission
Schizophrenia Spectrum	125	Paranoid schizophrenia; Disorganized schizophrenia in remission; Residual schizophrenia; Organic hallucinosis

Table II: DSM-5 Lexicon

the notion of web-based intervention where digital markers from social media are used for suicide prevention.

A. Semantic Layer

In the semantic layer, we perform tasks such as filtering of content from social media, followed by a grouping of semantically similar sentences in a Reddit post and a question from a questionnaire. We use two background knowledge sources to assess the entity-based semantic similarity of sentences and questions. For this task, we create four suicide-related semantic clusters namely *indicator*, *ideation*, *behavior*, and *attempt*; each semantic cluster (or contextual cluster) represents different states(characteristics) of suicide or suicidal

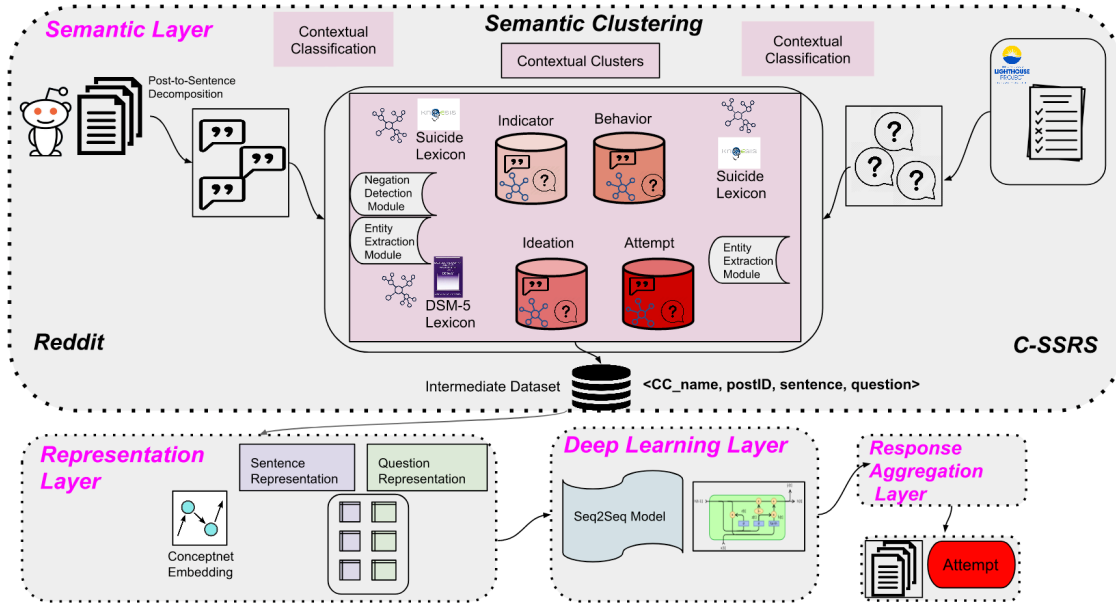


Figure 1: Contextual Elicitation Framework: Semantic Layer deals with semantic clustering of sentences in a post and questions based on background knowledge; Representation Layer is used to generate representations (embeddings) for sentences and questions and combine the representations; Deep Learning Layer is used to generate responses to the representations fed as an input feature.

Reddit Post	Annotator Label
I feel like such a crap for saying it, but the people who are in my life and who love me just dont seem like enough because I cant fully connect with any of them. I feel like I bring everyone down when Im with them.	Suicide Indicator
My little brother possibly killed himself and let me tell you, its been months and I havent gone a day without sobbing and considering suicide and feeling like my ribs were splitting apart.	Suicide Ideation
If I had a gun Id be dead already.A good shotgun or high-caliber handgun would do the job, and Id do it without hesitation. Killing myself would be far, far easier than facing, let alone actually dealing with, any of the bullshit.	Suicide Behavior
So heres why I want to do it:Ive been Suicidal since I was in second grade because it (being Suicidal) runs in my family. I started really thinking deeply about it in fourth grade but nothing ever happened. I tried to kill myself three times that same year	Suicide Attempt

Table III: Annotated sample Reddit posts by clinical psychiatrists.

thoughts. For instance, states related to actions are identified as attempt while states related to planning or intent are linked to behavior. This helps approach the problem of answering a question in a questionnaire from a social media text (Reddit post) using the divide-and-conquer approach. Thus, a Reddit post is broken down into individual sentences, followed by extraction of entities using background knowledge (suicide lexicons and DSM-5 lexicons) for classification of a sentence into a semantic cluster (CC_name). Similarly, entities in a *question* in a questionnaire are extracted and using background knowledge, the corresponding question is placed into a se-

semantic cluster. For this study, we leverage the frequency of occurrences of a sentence’s entity in the suicide lexicons to determine which semantic cluster a sentence is classified into. *Sentences* in a Reddit post that do not have entities present in any of the suicide lexicons are dropped. *Post ID* is used to keep track of a sentence in the complete architecture as multiple sentences belonging to the same post will eventually be combined to generate an aggregate suicide risk severity response. The intermediate dataset generated after this stage is a list of $\langle CC_name, postID, sentence, question \rangle$.

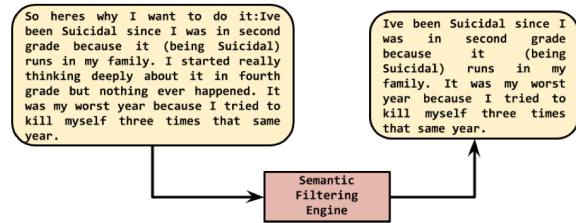


Figure 2: Semantic filtering based on the presence of entities in background knowledge: Semantic Filtering Engine is used to filter Reddit posts and generate preprocessed posts (sentences with entities present in suicide lexicons are passed).

Figures 2 and 3 shows how the filtering engine is used to extract relevant sentences in a Reddit post that are used to assess the suicide risk severity. Sentences that don’t have entities present in background knowledge sources (lexicons) or have entities negated are filtered and not passed through to the downstream tasks (semantic clustering and subsequent intermediate dataset generation). Figure 2 shows a single Reddit post passed through the semantic filtering engine to

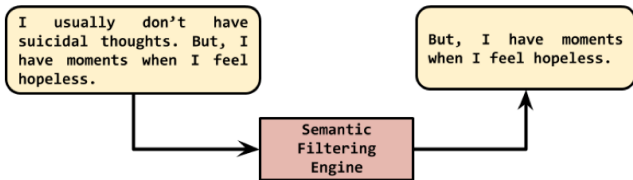


Figure 3: Semantic filtering augmented with negation detection: With negation detection, sentences in a post with entities present in suicide lexicons but negated are filtered and are not part of the preprocessed post

yield the preprocessed post on the right. Figure 3 shows semantic filtering process augmented with negation detection to filter entities that are negated; and hence minimizing false positives.

Cluster Name	Post ID	Sentence in a Post	Question
Suicide Indicator	118	I lost three months last year to drugs and cant remember what happened my friends just said I made a lot of bad decisions and made a lot of people hate me	Has subject had any history of mental disorder?
Suicide Ideation	4513	I just couldnt deal with my problems and there was nobody to help me, I sometimes wish I was not alive	Has subject ever wished he was dead?
Suicide Behavior	4436	I locked myself in the bathroom for the day	Have subject done anything to harm himself?
Suicide Attempt	3417	A botched attempt is likely to make your life a living hell and not only mentally physically as well	Has subject ever had life-threatening attempt?

Table IV: Examples of Sentence-Question pairs for each semantic cluster.

Having performed semantic clustering taking into the account the examples shown above, we have generated an intermediate dataset composed of fields: 1) cluster name (*CC_name*), 2) post ID, 3) sentence, and 4) question. Sample combinations of *cluster name*, *postID*, *sentence*, *question* are shown in Table IV.

Without a semantic clustering scheme, where semantically similar sentences and questions are paired based on background knowledge, extracting an answer for a question from an entire Reddit post using mere lexical matches leads to inaccurate responses [18]. Our novel approach of dividing a post into sentences provides better results because 1) sentences in a post that don't convey relevant information are removed improving focus in response extraction and 2) better responses are generated for a question if answer is looked for in the most

relevant part of a Reddit post.




Figure 4: Semantic Clusters statistical composition: Each bar represents the size of a semantic cluster; the size of a semantic cluster is the number of sentence-to-question pairs it contains

The Figure 4 summarizes the statistics of the semantic clusters and intermediate dataset generated at this phase. It can be inferred from Figure 4 that semantic cluster *indicator* contains the largest number of sentence-to-question pairs (30K) followed by *ideation*. Semantic cluster *attempt* has the least number of sentence-to-question pairs. While C-SSRS does not include a section for *indicator*, our results can be used to expand the questions in C-SSRS and address suicidal indicator signs in individuals. This addresses issues of questionnaires such as C-SSRS where reliability and comprehensiveness concerns are raised [2].

B. Representation Layer

Having generated the intermediate dataset, we use conceptnet⁷ numberbatch embeddings which consists of 417193 concepts to represent a sentence in a post and a question. We generate a 300-dimensional embedding for each term in a sentence and take the average of the embeddings to generate the 300-dimensional embedding for the sentence. We did the same for a question where a 300-dimensional embedding for each term in a question is generated and took the average of the individual terms to compute the embedding representation of the question. We used the same pre-trained embedding model for both the sentences and the questions. After generating 300-dimensional embeddings for a sentence and a question, we performed dot-product of the embeddings representing a sentence and a question to generate a combined 300-dimensional embedding which is a representation of a sentence-question pair. This 300-dimensional embedding representation later is used as an input feature for the learning phase which we build.

C. Learning Layer

Once a sentence-question pair representation is generated, the next task is to build a model that can give responses to

⁷<http://conceptnet.io/>

the individual pairs. We run a series of machine learning and deep learning experiments and evaluate the performance of each model against a truth dataset of the same posts annotated by practicing clinical psychiatrists. Since our study deals with sequential stream of text data, where a sentence in a Reddit post and a question should both be accounted for the sequences of words, our Seq2Seq models perform better based on the evaluation metric we defined. We perform ranking of responses to individual sentence-question pairs to generate a response to the parent Reddit post.

We plan to conduct evaluation of our findings taking into account the following factors.

- The error margin a predicted Reddit post’s suicide risk severity level differs from the annotator’s label. For example, a model predicting a Reddit post *indicator* while an annotator labeled *behavior* is penalized more than a model predicting *ideation* where an annotator labeled *behavior*.
- The linguistic nuances our model was unable to detect in its prediction. This may include nuances related to temporality or pronouns usage. For instance, our model takes into account a subject sharing their suicidal thoughts and the suicidal tendencies of third party individuals.
- False negatives and false positives weightage. Since false negatives are more likely to cost lives than false positives, the weights used for false negatives are bigger than for false positives in our model error computation.
- Semantically-enhanced model performance. Since ordinary performance metrics including precision, recall, and f1 score do not capture the semantics of model error, we incorporate the role of semantics in our model’s performance computation.

VI. CONCLUSION AND FUTURE WORK

In this study, we present an automated solution that will complement the traditional diagnosis and treatment procedures followed by MHPs for assessing a patient’s suicide risk. We highlight the contribution of rich content of Reddit posts in understanding of suicide risk for patients. We further explain the technical challenges in processing and analyzing Reddit posts. As we present an automated suicide-risk elicitation framework for Reddit, we plan to apply our approach over other social media platforms such as Twitter as our future work. Furthermore, our framework can be extended to solutions for broader healthcare problems involving multi-modal data, and can also be integrated with intelligent conversational agents or chatbots to reduce patient’s suicide risk level [19].

VII. ACKNOWLEDGEMENT

We acknowledge partial support from the National Science Foundation (NSF) award CNS-1513721: “Context-Aware Harassment Detection on Social Media”, National Institutes of Health (NIH) award: MH105384-01A1: “Modeling Social Behavior for Healthcare Utilization in Depression”. Any opinions, conclusions or recommendations expressed in this

material are those of the authors and do not necessarily reflect the views of the NSF or NIH.

REFERENCES

- [1] H.-C. Shing, S. Nair, A. Zirikly, M. Friedenber, H. Daumé III, and P. Resnik, “Expert, crowdsourced, and machine assessment of suicide risk via online postings,” in *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, 2018.
- [2] J. M. Giddens, K. H. Sheehan, and D. V. Sheehan, “The columbia-suicide severity rating scale (c-ssrs): Has the “gold standard” become a liability?” *Innovations in clinical neuroscience*, 2014.
- [3] M. De Choudhury, S. Counts, and E. Horvitz, “Social media as a measurement tool of depression in populations,” in *Proceedings of the 5th Annual ACM Web Science Conference*, 2013.
- [4] G. Gkotsis, A. Oellrich, S. Velupillai, M. Liakata, T. J. Hubbard, R. J. Dobson, and R. Dutta, “Characterisation of mental health conditions in social media using informed deep learning,” *Scientific reports*, 2017.
- [5] G. Gkotsis, A. Oellrich, T. Hubbard, R. Dobson, M. Liakata, S. Velupillai, and R. Dutta, “The language of mental health problems in social media,” in *Proceedings of the Third Workshop on Computational Linguistics and Clinical Psychology*, 2016.
- [6] K. Donnelly, “Snomed-ct: The advanced terminology and coding system for ehealth,” *Studies in health technology and informatics*, 2006.
- [7] O. mondiale de la santé, W. H. Organization, W. Staff, and WHO, *The ICD-10 classification of mental and behavioural disorders: clinical descriptions and diagnostic guidelines*, 1992.
- [8] O. Bodenreider, “The unified medical language system (umls): integrating biomedical terminology,” *Nucleic acids research*, 2004.
- [9] L. Ohno-Machado, S.-A. Sansone, G. Alter, I. Fore, J. Grethe, H. Xu, A. Gonzalez-Beltran, P. Rocca-Serra, A. E. Gururaj, E. Bell *et al.*, “Finding useful data across multiple biomedical data repositories using datamed,” *Nature genetics*, 2017.
- [10] M. Gaur, U. Kursuncu, A. Alambo, A. Sheth, R. Daniulaityte, K. Thirunarayan, and J. Pathak, ““ let me tell you about your mental health!” contextualized classification of reddit posts to dsm-5 for web-based intervention,” 2018.
- [11] S. A. Louzon, R. Bossarte, J. F. McCarthy, and I. R. Katz, “Does suicidal ideation as measured by the phq-9 predict suicide among va patients?” *Psychiatric Services*, 2016.
- [12] B. Runeson, J. Odeberg, A. Pettersson, T. Edbom, I. J. Adamsson, and M. Waern, “Instruments for the assessment of suicide risk: A systematic review evaluating the certainty of the evidence,” *PLoS one*, 2017.
- [13] G. Coppersmith, R. Leary, P. Crutchley, and A. Fine, “Natural language processing of social media as screening for suicide risk,” *Biomedical informatics insights*, 2018.
- [14] D. J. Corbitt-Hall, J. M. Gauthier, M. T. Davis, and T. K. Witte, “College students’ responses to suicidal content on social networking sites: an examination using a simulated facebook newsfeed,” *Suicide and Life-Threatening Behavior*, 2016.
- [15] M. De Choudhury, E. Kiciman, M. Dredze, G. Coppersmith, and M. Kumar, “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.
- [16] M. J. Vioulès, B. Moulahi, J. Azé, and S. Bringay, “Detection of suicide-related posts in twitter data streams,” *IBM Journal of Research and Development*, 2018.
- [17] J. Du, Y. Zhang, J. Luo, Y. Jia, Q. Wei, C. Tao, and H. Xu, “Extracting psychiatric stressors for suicide from social media using deep learning,” *BMC medical informatics and decision making*, 2018.
- [18] N. Limsopatham and N. H. Collier, “Normalising medical concepts in social media texts by learning semantic representation,” 2016.
- [19] A. P. Sheth, H. Y. Yip, U. Jaimini, D. Kadariya, V. Sridharan, R. Venkataramanan, T. Banerjee, K. Thirunarayan, and M. Kalra, “Augmented personalized health: Using semantically integrated multimodal data for patient empowered health management strategies,” 2018.