

2020

Assessing the Severity of Health States Based on Social Media Posts

Shweta Yadav

Joy Prakash Sain

Amit P. Sheth

University of South Carolina - Columbia, amit@sc.edu

Asif Ekbal

Sriparna Saha

See next page for additional authors

Follow this and additional works at: https://scholarcommons.sc.edu/aii_fac_pub



Part of the [Computer Engineering Commons](#), and the [Electrical and Computer Engineering Commons](#)

Publication Info

Preprint version *25th International Conference on Pattern Recognition*, 2020.

© The Authors, 2020

This Conference Proceeding is brought to you by the Artificial Intelligence Institute at Scholar Commons. It has been accepted for inclusion in Publications by an authorized administrator of Scholar Commons. For more information, please contact dillarda@mailbox.sc.edu.

Author(s)

Shweta Yadav, Joy Prakash Sain, Amit P. Sheth, Asif Ekbal, Sriparna Saha, and Pushpak Bhattacharyya

Assessing the Severity of Health States based on Social Media Posts

Shweta Yadav
Department of Computer Science
Wright State University
Dayton, Ohio, USA
shweta@knoesis.org

Joy Prakash Sain
Department of Computer Science
Wright State University
Dayton, Ohio, USA
sain.9@wright.edu

Amit Sheth
Artificial Intelligence Institute
University of South Carolina
Columbia, South Carolina, USA
amit@sc.edu

Asif Ekbal*, Sriparna Saha[†] and Pushpak Bhattacharyya[‡]

Department of Computer Science & Engineering
Indian Institute of Technology Patna

Bihar, India

Email: *asif@iitp.ac.in, [†]sriparna@iitp.ac.in, [‡]pb@iitp.ac.in,

Abstract—The unprecedented growth of Internet users has resulted in an abundance of unstructured information on social media including health forums, where patients request health-related information or opinions from other users. Previous studies have shown that online peer support has limited effectiveness without expert intervention. Therefore, a system capable of assessing the severity of health state from the patients’ social media posts can help health professionals (HP) in prioritizing the user’s post. In this study, we inspect the efficacy of different aspects of Natural Language Understanding (NLU) to identify the severity of the user’s health state in relation to two perspectives(tasks) (a) Medical Condition (i.e., Recover, Exist, Deteriorate, Other) and (b) Medication (i.e., Effective, Ineffective, Serious Adverse Effect, Other) in online health communities. We propose a multiview learning framework that models both the textual content as well as contextual-information to assess the severity of the user’s health state. Specifically, our model utilizes the NLU views such as sentiment, emotions, personality, and use of figurative language to extract the contextual information. The diverse NLU views demonstrate its effectiveness on both the tasks and as well as on the individual disease to assess a user’s health.

Index Terms—Natural Language Understanding, Social Media, Biomedical Natural Language Processing

I. INTRODUCTION

The volume of patient-generated healthcare data is experiencing an immense growth. The primary contributors to this enormous amount of data are social networks, forums, and blogs where patients share their medical problems and treatment experiences, including adverse reactions to medical products. According to Pew Internet & American Life Project [1]–[4], almost 80% of Internet users in the US explore health-related topics in online health forums. Among them, 63% look for information about specific medical problems, and nearly 47% look for medical treatments or procedures.

In the online health communities (OHC) and support groups, healthcare professional (HPs) provide clinical intervention when required, for example, when patients need any interpretation or details of clinical concepts or medical consultation. In these settings, patients benefit from the knowledge of both

peer-patients’ and HPs simultaneously [5]–[7]. [8] reports that the users on health forums expect the involvement of HPs for a quality suggestion and virtual observation. Another survey carried out by Pew Internet Research¹, shows that more than 80% of the patients preferred to consult the HPs rather than peer-patients for the information on prescription drugs, medical diagnoses, and alternative treatment options. However, the participation of HPs in the large-scale discussion forums is time-consuming. A study [9] on the 6 most active communities on WebMD.com data showed the level of HPs participation was observed to be extremely low (only 4.7% of the posts were answered). Hence, novel strategies are necessary for prioritization of the blog-posts based on the severity of ones’ health states that could assist the HPs to efficiently select the posts that need their expertise for making an effective and timely response. Specifically, we explore two important facets of the health state as proposed by [10] that can help in assessing the severity of ones’ health states based on their social media posts discussed as follows:

- **Status of the medical/health condition** (e.g., a patient whose medical condition is deteriorating even after clinical trials tends to be more severe than the patient who has started experiencing the medical symptoms.)
- **Consequences of the medication/treatment** (e.g., a patient reporting an adverse drug effect will be more severe than a patient whose treatment was ineffective.)

In Table-I, we provide the examples and description of the classes associated with the above described facets.

To evaluate our study, we have used the benchmark dataset made available through the LRE map [10]. The dataset utilized the blog-posts from two popular medical discussion forums, namely ‘patient.info’ and ‘dailystrength.org’ over four groups, namely asthma, allergy, depression and anxiety. The prior studies [10], [11] have treated this task as a traditional document classification tasks and propose the neural network framework

¹<https://pewrsr.ch/2wEL6JU>

Task 1: Medical Condition		Task 2: Medication	
Health Blog-post	Class-labels	Health Blog-post	Class-labels
<i>“my high resolution CT scan came back normal...I’m doing better after a long, breathless blue journey”</i>	Recover	<i>“I think the plaquenil is helping- been on it for almost 3 months”</i>	Effective
<i>“Been having eye problems ...lots of swelling redness and eye discharge.”</i>	Exist	<i>“I have been on Hizentra for almost a year...I don’t seem to be getting sick as much or as bad (yay) but things aren’t normal, that’s for sure”</i>	Ineffective
<i>“It’s been just over three months and I’m actually feeling worse! my IgG levels are rising from 766 to 1423, but I don’t feel good...”</i>	Deteriorate	<i>“I was given propranolol for migraine associated vertigo. I’ve only taken 5 mgs for last 2 nights and I’ve had bad nausea since starting it, I want to know if I can just stop taking it now without any problems ”</i>	Serious Adverse Effect
<i>“Not everyone uses or likes Facebook. Let’s remember there are so many who are looking for information and support on this covid forum.– Our voice can help Other”</i>	Other	<i>“Hello ladies, I am curious if anyone has any knowledge or experience with kidney symptoms resulting from using either Gammagard or other IG therapy brands”</i>	Other

TABLE I

EXAMPLES AND DESCRIPTION OF CLASSES FOR THE TASK 1: MEDICAL CONDITION AND TASK 2: MEDICATION. THE ‘EXIST’ CLASS BELONGS TO THOSE USERS’ WHO ARE REPORTING THEIR MEDICAL PROBLEM (SYMPTOMS) BUT ARE NOT UNDER ANY FORM OF TREATMENT/MEDICATION. THE ‘DETERIORATE’ CLASS CONSIST OF THOSE USERS’ WHOSE MEDICAL PROBLEM HAS WORSEN EVEN AFTER TAKING THE TREATMENT/MEDICATION.

which is able to capture the content level information from the user’s blog-post. However, the social-media texts often carry slang terminology, grammatical errors, figurative languages and hold in the information which is highly contextual. In such situations, mining only linguistic information turns out to be generally inefficient and urges for extra information/clues. The study conducted by [12] also demonstrates this requirement by showing how traditional classifier fails in instances where humans need additional context. They further illustrate the importance of speaker and topical information associated with the text to incorporate such context.

Advancement in NLU technology, is one of the most promising avenues for discovering vital contextual information from such data. Motivated by that, in this study we hypothesize that the NLU views such as emotions, sentiments, personality and usage of figurative language can help in understanding the vital contextual clues and discourse of the blog post required for detecting the health states. Ahead in this section, we have provide the detailed motivation behind utilizing the various NLU views in the multi-view learning framework.

The emotions that patients express towards their personal situation could be an important indicator to understand their health. Understanding the emotion can capture a user’s mental and physical health and can be applied to microblog posts [13] for behavioral decision making. Similarly, in mental health diagnosis, certain personality traits correlate with the diagnosis. Personality can be defined as characteristics patterns of an individual’s thinking, behaving, and emotional feeling. Hence, automatic identification of an individual’s personality type can have a wide range of applications in personalized health diagnosis [14] and to discover the user’s behavior.

We also explore the sentiments, which focus on extracting opinions and affects from the textual content [15], it seems natural that incorporating these knowledge can be helpful to discover emotional statement in the area of online health text classification. Previous research has provided evidence

to suggest that people’s mental and physical health can be predicted by analyzing their sentiment [10], [16] and the word they use [17], [18].

Similarly, in sentiment analysis, the presence of figurative language (FL) such as sarcasm in a text can work as an unexpected polarity reverser, which may undermine the accuracy of the system, if not addressed adequately [19].

In this work, we propose a multi-view learning framework [20] that jointly models both the content and the contextual-information specific for the tasks. It begins by processing contextual information using several NLU views such as emotions, personality, sentiments, and use of figurative language (FL). Following the contextual modeling phase, we perform the content modeling using the Bidirectional Encoder Representations from Transformers (BERT) to extract the task-agnostic view. The task-agnostic view is then fused with the various NLU derived views to obtain the final representation used for predicting the health of patients. Our contributions include:

- 1) A multi-view learning framework that allows the integration of different semantics captured from social media texts to support several aspects of natural language understanding (NLU) required by the health severity assessment task.
- 2) A comprehensive evaluation of the framework by testing it on a publicly available dataset and comparing the performance against the state-of-the-art baselines.
- 3) A demonstration of the effectiveness of various views such as emotion, sarcasm, personality, and sentiment on the diseases/disorders which provides the complementary information for assessing the severity of health states.

II. RELATED WORK

Literature shows increased attention on the OHCs for computationally discovering patient health [21]. A majority of

these studies follows a qualitative approach based on the manual categorization of posts by the domain experts. The categorization includes: (i) the type of support [22], [23], (ii) the type of emotion and sentiment expressed [24], and (iii) discussion on other illness specific topics or adverse drug effects [2], [25]. Below, we describe some of the prior research that utilized social media text of patients.

Several techniques have been devised to automatically process the OHCs content, uncover the user behaviours [9] and characteristics [26]. [27] in their study formulated the OHCs from a social support viewpoint and defined three variables : type of support, source of support, and setting in which the support is exchanged on online cancer communities. [28] developed a new content analysis method to (a) recognize various healthcare participants, (b) discover currents trends, and (c) analyze the sentiments expressed by different healthcare participants from lung cancer, diabetes, and breast cancer forums. [29] proposed a text mining technique to classify the user’s participation based on the different types of social support, such as informational support, emotional support, and companionship. Further, they developed a supervised machine learning approach to predict whether user will be churn from OHC. The challenge defined in [30] aimed to automatically classify the user posts from an online mental health forum into four different categories (crisis/red/amber/green) according to need of urgent attention. [5] developed a text classification technique for assisting the moderators in OHCs. Specifically, they devised a classification scheme to automatically categorize the blog-post over the codes such as: ‘Asking for medical information’, ‘Asking for peer-patients’ experience’, ‘General chatting’, and ‘Miscellaneous’, using WebMD’s online diabetes community data. Some of the other prominent research in this area includes work of [31], [32]. Recently, [10] has introduced the novel annotation scheme for analyzing medical sentiment on social media text that can capture the severity of the user’s health states. They utilized a CNN to understand the possible sentiment. [11] further extended the study utilizing a multitask learning framework to capture the multiple facets of medical sentiments simultaneously.

III. MATERIALS AND METHODS

First, we define the task and detail the approach used in the study. Later, we describe the dataset and provide the training details, followed by the baseline models.

A. Task Definition:

For a given medical forum post M , consisting of n sentences, i.e., $M = \{s_1, s_2, s_3, \dots, s_n\}$, the task is to predict the two aspects of health status ‘ Y ’ & ‘ Z ’ from a discrete set of medical conditions ‘ $Y = \{Recover, Exist, Deteriorate, Other\}$ ’ and medications ‘ $Z = \{Effective, Ineffective, Serious Adverse Effect, Other\}$ ’.

B. Summary of the Approach

Given the user’s forum post M to be classified, the proposed system leverages both task-agnostics and context specific

views from the text. In order to capture task-agnostic view, we model M using BERT model to obtain a vector representation of the medical blog post. BERT generate the abstract representation of words by capturing bi-directional context in the input sentence [33]. Basically, BERT aims to learn the deep bidirectional representations from the Transformer stack [34]. For modeling the contextual views, the proposed system utilizes the information of multiple NLU based features/views (emotion, sarcasm, personality, and sentiment) extracted from the user forum post. For extracting emotion, sarcasm and personality views we adopt the domain adaptation approach [35] where we train a model in one domain (i.e., publicly available gold data), and extract the features on another domain (i.e., our dataset). Finally, we fused the NLU views with the BERT generated task-agnostic views which are used to categorize the medical forum post. Figure-1 shows the architecture of our approach.

C. Task-agnostic View

We generated the task-agnostic views as follows: The task-agnostic view are generated using the BERT network. We employ the pre-trained BERT model² having 12 Transformer layers (L), each having 12 heads for self-attention and hidden dimension of 768, to extract the feature representation of the medical forum post. The pre-trained model has shown the state-of-the-art performance in various natural language processing tasks [36], [37]. The pre-trained BERT model is highly efficient in generating the task-agnostic input representation from the transformer architecture [38]. This enables even the low-resource tasks to benefit from deep bi-directional architectures [39] and the unsupervised training framework to obtain the pre-trained network. We perform extensive experiment and to obtain the effective representation of the medical forum post representation. Our experimental studies conclude that the aggregating the last three layers of the BERT model achieves the best result in our experiment. Given a forum post M consisting of n tokens $\{t_1, t_2, t_3, \dots, t_n\}$. We use the Word Piece tokenizer [40] to tokenize the sentence. We use the special $[CLS]$ token representation as the task-specific feature for the medical forum post M . We represent this task-agnostic view as the content feature \hat{d}_{cont} .

D. Contextual (NLU) View

In this section, we will discuss in details all the generated NLU (emotion, sarcasm, personality, and sentiment) views:

1) **Emotion View:** Humans are emotional beings; emotion carries an intrinsic role in human life. It influences our decision making [41], shapes our behavior [42], and affects mental and physical health [43].

Towards this, we study the user’s forum post on five primary emotions such as ‘*anger*’, ‘*disgust*’, ‘*joy*’, ‘*fear*’, and ‘*sadness*’ that can assist the model in capturing the overall health condition of an individual. We also examine the post on more fine-grained emotions (i.e., ‘*valence*’, ‘*arousal*’, and

²shorturl.at/nDJPY

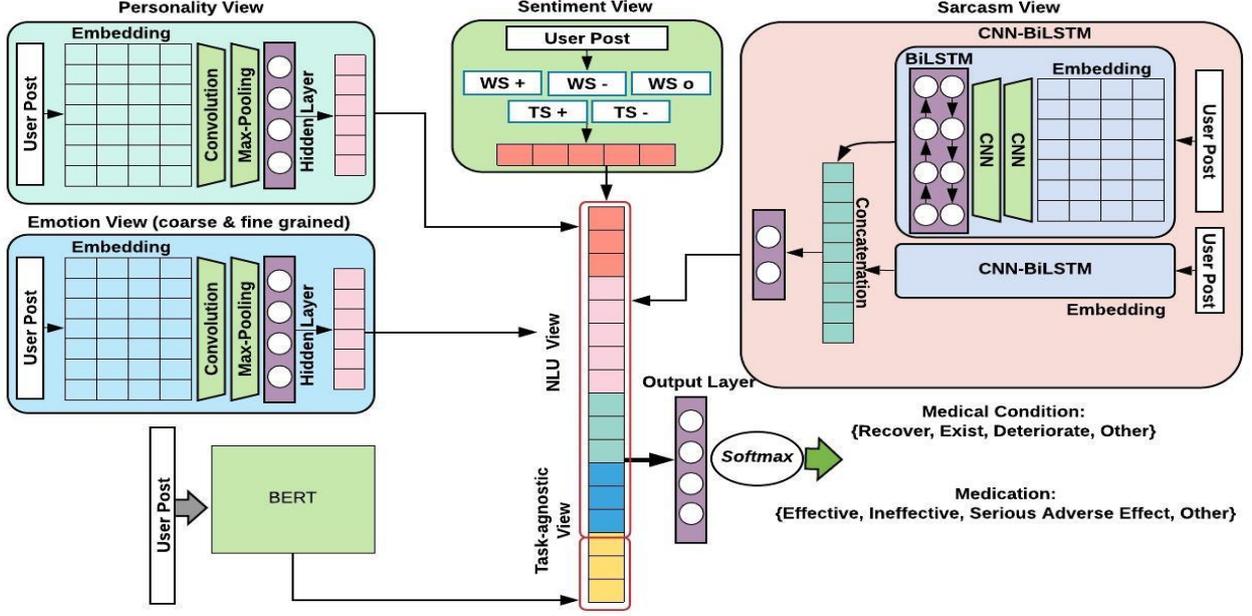


Fig. 1. Proposed architecture for predicting the health severity assessment task.

‘dominance’) that reveals the user’s state of feeling [44].

In order to extract the emotion views from forum post, we leverage a system similar to the CNN network [45]. We utilize the benchmark emotion EmoBank-2017 dataset [46] and EmoInt-2017 dataset [47] to train the model for fine and coarse-grained emotion analysis, respectively. After training the model, we extract the feature representation for each forum post.

Given a forum post M consisting of n words, after extracting the pooled representation, we pass it to a hidden layer to generate the view representation as follows:

$$F_e = \varphi(W_e \hat{d}_e + b_e) \quad (1)$$

where φ is a non-linear activation function, W_e , and b_e are the weight, and bias of the last hidden layer. The pooled feature with the given window size k is generated as follows:

$$d_e = \text{pool}(\text{conv}(\text{emb}(M), k; \theta)) \quad (2)$$

where $\text{emb}(\cdot)$, $\text{conv}(\cdot)$ and $\text{pool}(\cdot)$ are the embedding, convolution and pooling operations similar to standard CNN model. θ is the model parameter. The final extracted feature is \hat{d}_e is obtained by concatenating the multiple features of different window size. The fine and coarse grained emotion views F_e^f and F_e^c are extracted using the aforementioned equation.

2) **Sarcasm View:** Humans seamlessly perform high-level semantic tasks by subconsciously utilizing a vast collection of composite linguistic units along with their background knowledge to visualize the reality. Social media texts often contain FL. The presence of FL makes it challenging to process for any NLP applications, sentiment analysis in particular. In sentiment analysis, the presence of FL such as sarcasm

in a text can work as an unexpected polarity reverser, which may undermine the accuracy of the system, if not addressed adequately [19]. In medical forums, patients seeking support for their medical problems, often use sarcasm to express their emotion.

In the following example, one patient with anxiety problem describes her health condition sarcastically, writing as: “... It goes from pain to slight discomfort.. I cant move. Great way to start the day !”

Here, the phrase “Great way to start the day !” is presented in the sarcastic sense to express the medical condition of not feeling good. To model this feature for each forum post in our framework, we utilize the approach proposed by [48]. The model utilizes 2-layers of CNN followed by a Long Short Term Memory (LSTM) [49] network. We compute the sarcastic scores as follows:

$$F_s = \sigma(W_s \cdot V_c + b_s) \quad (3)$$

where σ represents the *sigmoid* activation function. V_c , W_s and b_s are the input, weight, and bias term of the output layer. The input vector of the output layer V_c is the concatenated feature vector that combines the features extracted by the LSTM layers, which capture the state of mind of a user.

3) **Personality View:** One’s behavior characterizes personality, sympathy, emotion, thought process, and motivation. Our personality impacts many preferences in our lives such as decision making, life choices, physical and mental health. For extracting personality view, we employ an approach similar to the one used for emotion view. We utilize the corpus curated by [50], which contains 2,400 essays labeled with big-five personality traits, to train the model. We utilize the CNN

as a learning model for the multi-label personality detection task. After the training, the CNN model is used to infer the personality traits present in each medical forum post by extracting the activation of the CNN's last hidden layer vector which we call as the *personality vector*. After the training, the model is used for extracting the feature representation. The extracted personality vector can be represented as:

$$F_p = \varphi(W_p \hat{d}_p + b_p) \quad (4)$$

where φ is a non-linear activation function, and \hat{d}_p (i.e., extracted pooled feature), W_p , and b_p are the input, weight, and bias of the last hidden layer.

4) **Sentiment View:** We generate the sentiment views from the forum posts as described below:

1) Word-level Sentiment (WS): Sentiment clue words provide important features in deciding the sentiment of the users. Besides, the inclusion of negation to the sentiment word can change the polarity. For example, there is positive sentiment in "I'm **stable**" but after including negation like "I'm **not stable**", the sentiment polarity changes. Briefly, there are two types of sentiment events by which we can capture the sentiments of users: occurrences of sentiment words (SW), occurrences of sentiment words with negation (NSW). This feature calculates the positive ($SW_{(+)}$), negative ($SW_{(-)}$) and objective ($SW_{(O)}$) score for each word by capturing the sentiment event [51]. Publicly available SentiWordNet (SWN)³ is used to calculate the score for each word as follows:

$$SW_+ = (tf(SW) * f_+(SW) + tf(NSW) * f_-(SW)) * idf(SW) \quad (5)$$

$$SW_- = (tf(SW) * f_-(SW) + tf(NSW) * f_+(SW)) * idf(SW) \quad (6)$$

$$SW_O = tfidf(SW) * f_O(SW) \quad (7)$$

Here, tf and idf represent the term and inverse document frequencies, respectively. $f_+(SW)$, $f_-(SW)$ and $f_O(SW)$ are positive, negative and objective scores, respectively obtained from the SentiWordNet. The word-level sentiment feature of a forum post having n words is obtained as follows:

$$WS_+ = \frac{\sum_{SW \in n} SW_+}{n} \quad (8)$$

$$WS_- = \frac{\sum_{SW \in n} SW_-}{n} \quad (9)$$

$$WS_O = \frac{\sum_{SW \in n} SW_O}{n} \quad (10)$$

2) Target-specific Sentiment (TS): After analyzing the validation data, we observe that approximately 92% of the posts depict sentiments in the context of a certain stative verbs such as 'feel', 'suffer', 'experience'. We design this feature

by considering a context window of [-5,5] words and selecting the most effective stative verb. After that, negative and positive densities of a post are calculated by the frequency of the clue words to the number of words in the context (i.e., 10 in this case). For example, if a post contains more than one instance of 'feel' term, we calculate the score individually and consider the maximum one. If the word 'feel' appears at the i^{th} position in a forum post then the score is calculated as follows:

$$Score(+)=\sum_{m=i-k}^{m=i-1} w_m \times f_+(SW_m) + \sum_{n=i+1}^{n=i+k} w_n \times f_+(SW_n) \quad (11)$$

$$Score(-)=\sum_{m=i-k}^{m=i-1} w_m \times f_-(SW_m) + \sum_{n=i+1}^{n=i+k} w_n \times f_-(SW_n) \quad (12)$$

where, k is context window size, weight $w_m = m + k - i + 1$ and $w_n = k - n + i + 1$. The aggregate scores TS_+ and TS_- of a forum post are calculated as follows:

$$TS_+ = \max(Score_{t=0}(+), Score_{t=1}(+), \dots, Score_{t=T}(+)) \quad (13)$$

$$TS_- = \max(Score_{t=0}(-), Score_{t=1}(-), \dots, Score_{t=T}(-)) \quad (14)$$

where, T is the number of sentiment bearing words in the post.

E. Multi-view Fusion Layer

We take a multi-view learning approach to combine the various views discussed above into a comprehensive embedding for each medical forum post. We use an extension of Canonical Correlation Analysis (CCA) [52] to perform fusion from multiple views. The extended CCA captures maximal information between multiple views and creates a combined representation. The extension of CCA is called the Generalized CCA (GCCA) [53], which has been used in the literature to fuse the multiple sources of information into a single source. GCCA finds G, U_i by solving the optimization problem

$$\arg \min_{G, U_i} \sum_i \|G - X_i U_i\|_F^2 \quad (15)$$

such that $G^T G = I$.

where, $G \in \mathbb{R}^{m \times k}$ contains the fused feature representation matrix, $X_i \in \mathbb{R}^{m \times d_i}$ corresponds to the data matrix for the i^{th} view and $U_i \in \mathbb{R}^{d_i \times k}$ maps from latent space to observable view i . However, since all the views are not equally important, we employ the weighted GCCA (wGCCA) [54]. In this representation, we add a weight term to the above equation as follows:

$$\arg \min_{G, U_i} \sum_i w_i \|G - X_i U_i\|_F^2 \quad (16)$$

such that $G^T G = I$ and $w_i \geq 0$ and represents the importance of the i^{th} view in the fusion process. The columns of G are the eigenvectors of $\sum_i w_i X_i (X_i' X_i)^{-1} X_i'$, and the solution

³<http://sentiwordnet.isti.cnr.it/>

Dataset Categories	Classes (# of instances)				# blog-post	# avg sentences/blog	# avg words/blog
	Exist (2,396)	Recover (703)	Deteriorate (2,089)	Other (432)			
Medical Condition	Effective (462)	Ineffective (613)	Serious Adverse Effect (1,226)	Other (432)	5,621	10	194
Medication					2,734	9	177

TABLE II
DATASET STATISTICS FOR BOTH THE CATEGORIES OF MEDICAL SENTIMENT

Models	Techniques Used	Medical Condition			Medications		
		Precision	Recall	F-Score	Precision	Recall	F-Score
Baseline 1	BERT	72.70	73.15	72.89	86.64	87.55	86.81
Baseline 2	BioBERT	72.42	72.30	72.28	86.68	86.97	86.76
Baseline 3	MTL [11]	66.71	64.33	65.5	85.33	81.90	83.58
Proposed Approach	NLU based Multi-view Learning	75.52	80.25	77.45	89.52	89.91	89.57

TABLE III
PERFORMANCE COMPARISON OUR PROPOSED MODEL WITH THE BASELINES ON BOTH THE DATASETS

Index	View	Medical Condition	Medications
(1)	All	77.45	89.57
(2)	- Emotion (coarse)	75.08	87.32
(3)	- Emotion (fine)	75.44	88.61
(4)	- Sarcasm	77.10	86.82
(5)	- Personality	74.91	87.66
(6)	- Word-level Sentiment	74.54	85.94
(7)	- Target-specific Sentiment	75.85	85.27

TABLE IV
FEATURE ABLATION STUDY: THE SYSTEM PERFORMANCE (IN F1 SCORE) BY REMOVING ONE VIEW AT A TIME

for $U_i = (X_i' X_i)^{-1} X_i' G$.

We use wGCCA, to obtain the final feature G representation. Finally, the classification of forum post is carried out by the following equation:

$$p(Y = y|M, G) = \text{softmax}_y(G^T W + b) = \frac{e^{G^T W_y + b_y}}{\sum_{k=1}^K e^{G^T W_k + b_k}} \quad (17)$$

F. Dataset

We evaluate the performance of the system on the publicly available dataset [10] obtained from popular online health forum⁴. The forum posts were collected from four online discussion groups: *Depression*, *Allergy*, *Asthma*, and *Anxiety*. The dataset consists of 5,621 medical forum posts related to *Task 1: medical conditions* and 2,734 forum posts related to the category of *Task 2: medication*. We have extended the previous dataset by including 1 more class ('Other'), which consider the miscellaneous blog-post. These type of blog post does not explicitly provide any information regarding their medical condition or treatment but are more sort of general enquiry. The more detailed description for each class can be obtained from [10]. The detailed description of dataset statistics is provided in Table-II. We perform a 10-fold cross-validation experiment on both the datasets.

Ethics: Our project involves analysis of anonymized data that

is publicly available and used by the other publication. It does not involve any direct interaction with any individuals or their personally identifiable data. Thus, this study was reviewed by the Wright State University IRB and received an exemption determination.

IV. EXPERIMENTAL RESULTS

Here, we present results on the severity assessment task. Thereafter, we will provide technical interpretation of the results followed by ablation study. We used Recall, Precision and F₁-Score to evaluate our proposed task against state-of-the-art relation extractor. As a baseline model, we used **BERT**, **BioBERT** [55], and multi-task adversarial learning framework [11] to compare our proposed model.

We report the performance of our proposed approach along with other baselines in Table III on task-1 (Medical Condition), and task-2 (Medication). The obtained results shows that BERT model is the best among all the baselines models. The proposed approach achieves the improvement of 4.56, 5.17, and 11.95 F-Score points for task-1, and 2.76, 2.81, and 5.99 F-Score for task-2 over the baseline 1, 2, and 3 respectively. Statistical significance test (t-test) shows that improvements over the baselines are significant (p -value < 0.05).

To prove the effectiveness of each view, we conduct the ablation experiments on our proposed model. As shown in Table-IV, the performance of the model shows varying degrees of decline when we remove different view from the model. All the declines are significant with $p \leq 0.05$ under the t-test.

On the task 1(medical condition), sentiment view seems to be most crucial view, as removal of the word level sentiment drops the F-Score points by 2.91. Similarly, for the target-specific sentiment, we observe a decline of 1.60. On the medication task, again sentiment view is found to be important. Removal of word-level and target-specific sentiment view declines the model performance by 3.63 and 4.30 F-Score point respectively. This shows that primarily the sentiment views contribute to determining the health states. The two other significant views affecting the final predictions are emotion

⁴<https://patient.info>

(fine) and personality. As they directly reflect the behavior of a user, the exclusion leads to a decline in the performance. The impact of sarcasm view is found to be smaller as compared to other views, as its removal drops the performance by 0.35 and 2.75 F-score points on Medical condition and Medications dataset respectively. Although our analysis shows that the identification of sarcasm is crucial, the little impact could be because of our learning strategy. Since the use of FL in the medical domain is quite different than the general domain.

V. DISCUSSION

In this section, we study a couple of cases from both the datasets, where our model correctly identifies various aspects of the health states with the help of NLU views.

- **Case 1: Effect of the emotion view**

Consider the following example from task-1:

“why does anxiety feel like you have to make yourself breathe instead of letting your body breathe on its own. Am’i the only one.”

In the absence of the emotion views, the system misclassifies the post as ‘Other’. However, the inclusion of the emotion views assists in understanding the users’ implicit states of mind, and classifies correctly as ‘Exist’. The system captures the anger and disgust emotions present in the text, which are highly correlated with the emotion distribution associated with the medical condition category ‘Exist’.

- **Case 2: Effect of the sarcasm view**

In the following post, the user sarcastically expresses his/her condition:

“Lol I’m just a big ball of anxiety fun.”

The system misclassifies the post as ‘Recovered’ in the absence of sarcasm view, which may be due to the presence of positive sentiment-bearing words. However, sarcasm view helps the model to predict the class ‘Exist’ correctly. The contextual cues extracted from a post are not always enough to understand someone’s feelings and require common sense and background knowledge about the topic of discussion. Such situations are prevalent in the forum posts with very long sentences.

- **Case 3: Effect of the personality view**

In our study we find the personality view to be very useful in understanding the users having an anxiety disorder. Consider the following example from the anxiety group of task-1:

“cutting open my arm. The sensible bit of me says NO just a reaction to my new meds, the other half want to self destruct... I don’t know which one is going to win...”

The personality view identified signs of neurotic personality in the above post, which overlaps with the symptoms of anxiety. This helped the system in correctly classifying the user’s mental state as having the anxiety disorder.

- **Case 4: Effect of the sentiment view**

As shown in Table IV, the sentiment views help the model in boosting the performance by nearly 2% F-Score. In the following example:

“I think it’s because I’m afraid of feeling ill when I’m out. This past week I increased citalopram to 20 mg and zI don’t know if it’s making me feel worse.”

The sentiment views capture the negative sentiment-bearing words (i.e., afraid, ill, and worse) and encode this information to assist the model for correctly predicting the class as ‘Serious Adverse Effect’.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we identify the severity of a user’s health state by analyzing different medical aspects (such as medical condition and outcome of treatment) from their social media texts. We validate our study by utilizing a benchmark dataset curated from medical web forums. We propose a deep learning model leveraging various NLU views such as emotion, sarcasm, personality, and sentiment along with the textual content for classifying the medical forum posts. The evaluation shows that combining the content view to context views is an effective way to boost the classification performance. In the future, we would like to explore the other facets of a user’s health state like ‘Consequence of a treatment’ and ‘Certainty of a diagnosis’. In addition to sarcasm, we would also like to model other forms of figurative languages like ‘metaphor’ and ‘irony’ which are widely used in social media texts.

VII. ACKNOWLEDGEMENT

Amit Sheth acknowledged partial support from NMH award R01MH105384 “Modeling Social Behavior for Healthcare Utilization in Depression. All findings and opinions are of authors and not sponsors. Sriparna Saha would like to acknowledge the support of SERB WOMEN IN EXCELLENCE AWARD 2018.

REFERENCES

- [1] S. Fox, *The social life of health information 2011*. Pew Internet & American Life Project Washington, DC, 2011.
- [2] S. Yadav, A. Ekbal, S. Saha, and P. Bhattacharyya, “A unified multi-task adversarial learning framework for pharmacovigilance mining,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 5234–5245. [Online]. Available: <https://www.aclweb.org/anthology/P19-1516>
- [3] S. Yadav, A. Ekbal, S. Saha, P. Bhattacharyya, and A. Sheth, “Multi-task learning framework for mining crowd intelligence towards clinical treatment,” in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, Jun. 2018, pp. 271–277. [Online]. Available: <https://www.aclweb.org/anthology/N18-2044>
- [4] S. Yadav, A. Ekbal, S. Saha, and P. Bhattacharyya, “Medical sentiment analysis using social media: Towards building a patient assisted system,” in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Miyazaki, Japan: European Language Resources Association (ELRA), May 2018. [Online]. Available: <https://www.aclweb.org/anthology/L18-1442>
- [5] J. Huh, M. Yetisgen-Yildiz, and W. Pratt, “Text classification for assisting moderators in online health communities,” *Journal of biomedical informatics*, vol. 46, no. 6, pp. 998–1005, 2013.
- [6] S. A. Qureshi, S. Saha, M. Hasanuzzaman, and G. Dias, “Multitask representation learning for multimodal estimation of depression level,” *IEEE Intelligent Systems*, vol. 34, no. 5, pp. 45–52, 2019.

- [7] S. A. Qureshi, M. Hasanuzzaman, G. Dias, and S. Saha, "Improving depression level estimation by concurrently learning emotion intensity," *IEEE Computational Intelligence Magazine*, vol. accepted, 2020.
- [8] P. E. Kummervold, D. Gammon, S. Bergvik, J.-A. K. Johnsen, T. Hasvold, and J. H. Rosenvinge, "Social support in a wired world: use of online mental health forums in norway," *Nordic journal of psychiatry*, vol. 56, no. 1, pp. 59–65, 2002.
- [9] J. Huh, D. W. McDonald, A. Hartzler, and W. Pratt, "Patient moderator interaction in online health communities," in *AMIA Annual Symposium Proceedings*, vol. 2013. American Medical Informatics Association, 2013, p. 627.
- [10] S. Yadav, A. Ekbal, S. Saha, and P. Bhattacharyya, "Medical sentiment analysis using social media: Towards building a patient assisted system." in *LREC*, 2018.
- [11] S. Yadav, A. Ekbal, S. Saha, P. Bhattacharyya, and A. Sheth, "Multi-task learning framework for mining crowd intelligence towards clinical treatment," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, vol. 2, 2018, pp. 271–277.
- [12] B. C. Wallace, L. Kertz, E. Charniak *et al.*, "Humans require context to infer ironic intent (so computers probably do, too)," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, vol. 2, 2014, pp. 512–516.
- [13] W. Wang, L. Chen, K. Thirunarayan, and A. P. Sheth, "Harnessing twitter" big data" for automatic emotion identification," in *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*. IEEE, 2012, pp. 587–592.
- [14] F. Cortellese, M. Nalin, A. Morandi, A. Sanna, and F. Grasso, "Personality diagnosis for personalized ehealth services," in *International Conference on Electronic Healthcare*. Springer, 2009, pp. 157–164.
- [15] B. Pang, L. Lee *et al.*, "Opinion mining and sentiment analysis," *Foundations and Trends® in Information Retrieval*, vol. 2, no. 1–2, pp. 1–135, 2008.
- [16] J. Carrillo-de Albornoz, J. R. Vidal, and L. Plaza, "Feature engineering for sentiment analysis in e-health forums," *PloS one*, vol. 13, no. 11, p. e0207996, 2018.
- [17] L. A. Gottschalk and G. C. Gleser, *The measurement of psychological states through the content analysis of verbal behavior*. Univ of California Press, 1969.
- [18] W. B. Stiles, *Describing talk: A taxonomy of verbal response modes*. Sage Publications Newbury Park, 1992.
- [19] C. Bosco, V. Patti, and A. Bolioli, "Developing corpora for sentiment analysis: The case of irony and senti-tut," *IEEE Intelligent Systems*, vol. 28, no. 2, pp. 55–63, 2013.
- [20] R. Kumar, S. Yadav, R. Daniulaityte, F. Lamy, K. Thirunarayan, U. Lokala, and A. Sheth, "edarkfind: Unsupervised multi-view learning for sybil account detection," in *Proceedings of The Web Conference 2020*, 2020, pp. 1955–1965.
- [21] B. I. Davidson, S. L. Jones, A. N. Joinson, and J. Hinds, "The evolution of online ideological communities," *PloS one*, vol. 14, no. 5, p. e0216932, 2019.
- [22] T. O. Blank, S. D. Schmidt, S. A. Vangsness, A. K. Monteiro, and P. V. Santagata, "Differences among breast and prostate cancer online support groups," *Computers in Human Behavior*, vol. 26, no. 6, pp. 1400–1404, 2010.
- [23] N. S. Coulson, H. Buchanan, and A. Aubeeluck, "Social support in cyberspace: a content analysis of communication within a huntington's disease online support group," *Patient education and counseling*, vol. 68, no. 2, pp. 173–178, 2007.
- [24] I. Ruthven, S. Buchanan, and C. Jardine, "Isolated, overwhelmed, and worried: Young first-time mothers asking for information and support online," *Journal of the Association for Information Science and Technology*, 2018.
- [25] J. Huber, A. Ihrig, T. Peters, C. G. Huber, A. Kessler, B. Hadaschik, S. Pahernik, and M. Hohenfellner, "Decision-making in localized prostate cancer: lessons learned from an online support group," *BJU international*, vol. 107, no. 10, pp. 1570–1575, 2011.
- [26] J. Huh and W. Pratt, "Weaving clinical expertise in online health communities," in *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 2014, pp. 1355–1364.
- [27] S. Zhang, E. O'Carroll Bantum, J. Owen, S. Bakken, and N. Elhadad, "Online cancer communities as informatics intervention for social support: conceptualization, characterization, and impact," *Journal of the American Medical Informatics Association*, vol. 24, no. 2, pp. 451–459, 08 2016. [Online]. Available: <https://doi.org/10.1093/jamia/ocw093>
- [28] Y. Lu, Y. Wu, J. Liu, J. Li, and P. Zhang, "Understanding health care social media use from different stakeholder perspectives: a content analysis of an online health community," *Journal of medical Internet research*, vol. 19, no. 4, p. e109, 2017.
- [29] X. Wang, K. Zhao, and N. Street, "Analyzing and predicting user participations in online health communities: A social support perspective," *J Med Internet Res*, vol. 19, no. 4, p. e130, Apr 2017. [Online]. Available: <http://www.jmir.org/2017/4/e130/>
- [30] D. N. Milne, G. Pink, B. Hachey, and R. A. Calvo, "Clpsych 2016 shared task: Triaging content in online peer-support forums," in *Proceedings of the Third Workshop on Computational Linguistics and Clinical Psychology*, 2016, pp. 118–127.
- [31] K. O'Leary, S. M. Schueller, J. O. Wobbrock, and W. Pratt, "suddenly, we got to become therapists for each other": Designing peer support chats for mental health," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 2018, p. 331.
- [32] R. Balyan, S. A. Crossley, W. Brown III, A. J. Karter, D. S. McNamara, J. Y. Liu, C. R. Lyles, and D. Schillinger, "Using natural language processing and machine learning to classify health literacy from secure messages: The eclipse study," *PloS one*, vol. 14, no. 2, p. e0212488, 2019.
- [33] A. Jacovi, O. S. Shalom, and Y. Goldberg, "Understanding convolutional neural networks for text classification," *arXiv preprint arXiv:1809.08037*, 2018.
- [34] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [35] X. Glorot, A. Bordes, and Y. Bengio, "Domain adaptation for large-scale sentiment classification: A deep learning approach," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 513–520.
- [36] J. Maillard, S. Clark, and D. Yogatama, "Jointly learning sentence embeddings and syntax with unsupervised tree-lstm," *Natural Language Engineering*, vol. 25, no. 4, pp. 433–449, 2019.
- [37] J. Hewitt and C. D. Manning, "A structural probe for finding syntax in word representations," in *Proceedings of the 2019 Conference of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 4129–4138.
- [38] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017.
- [39] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [40] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey *et al.*, "Google's neural machine translation system: Bridging the gap between human and machine translation," *arXiv preprint arXiv:1609.08144*, 2016.
- [41] G. A. Van Kleef, C. K. De Dreu, and A. S. Manstead, "An interpersonal approach to emotion in social decision making: The emotions as social information model," in *Advances in experimental social psychology*. Elsevier, 2010, vol. 42, pp. 45–96.
- [42] R. F. Baumeister, K. D. Vohs, C. Nathan DeWall, and L. Zhang, "How emotion shapes behavior: Feedback, anticipation, and reflection, rather than direct causation," *Personality and social psychology review*, vol. 11, no. 2, pp. 167–203, 2007.
- [43] N. Extremera and P. Fernández-Berrocá, "Emotional intelligence as predictor of mental, social, and physical health in university students," *The Spanish Journal of Psychology*, vol. 9, no. 1, pp. 45–51, 2006.
- [44] A. Mehrabian and J. A. Russell, *An approach to environmental psychology*. the MIT Press, 1974.
- [45] Y. Kim, "Convolutional neural networks for sentence classification," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar: Association for Computational Linguistics, October 2014, pp. 1746–1751. [Online]. Available: <http://www.aclweb.org/anthology/D14-1181>
- [46] S. Buechel and U. Hahn, "Emobank: Studying the impact of annotation perspective and representation format on dimensional emotion analysis," in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, vol. 2, 2017, pp. 578–585.

- [47] S. M. Mohammad and F. Bravo-Marquez, "Wassa-2017 shared task on emotion intensity," *arXiv preprint arXiv:1708.03700*, 2017.
- [48] A. Ghosh and T. Veale, "Magnets for sarcasm: Making sarcasm detection timely, contextual and very personal," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 482–491.
- [49] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, pp. 1735–80, 12 1997.
- [50] G. Matthews and K. Gilliland, "The personality theories of hj eysenck and ja gray: A comparative review," *Personality and Individual differences*, vol. 26, no. 4, pp. 583–626, 1999.
- [51] T.-T. Dang and K. Shirai, "Machine learning approaches for mood classification of songs toward music search engine," in *Knowledge and Systems Engineering, 2009. KSE'09. International Conference on*. IEEE, 2009, pp. 144–149.
- [52] H. Hotelling, "Relations between two sets of variates," in *Breakthroughs in statistics*. Springer, 1992, pp. 162–190.
- [53] J. D. Carroll, "Generalization of canonical correlation analysis to three or more sets of variables," in *Proceedings of the 76th annual convention of the American Psychological Association*, vol. 3, 1968, pp. 227–228.
- [54] A. Benton, R. Arora, and M. Dredze, "Learning multiview embeddings of twitter users," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2016, pp. 14–19.
- [55] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang, "Biobert: a pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2020.