Online Behaviour Indicator Extraction for Enhanced Cancer Patient Management using Real-World Data

Georgia Pantalona
CERTH-ITI
Thessaloniki, Greece
georgia.pant@gmail.com

Anna Barachanou
CERTH-ITI
Thessaloniki, Greece
barachanou@iti.gr

Nikolaos Loukas
CERTH-ITI
Thessaloniki, Greece
nickloukas@iti.gr

Lazaros Apostolidis
CERTH-ITI
Thessaloniki, Greece
laaposto@iti.gr

Filareti Tsalakanidou
CERTH-ITI
Thessaloniki, Greece
filareti@iti.gr

Symeon Papadopoulos
CERTH-ITI
Thessaloniki, Greece
papadop@iti.gr

ABSTRACT

Mental health issues of breast cancer patients and survivors are detrimental to their recovery, reintegration to normal living, and quality of life. With the increasing popularity of social media platforms and online browsing, combined with the rise of Natural Language Processing (NLP), there exists great potential in analyzing textual information of various online user activities to infer important indicators about a person's online behaviour and emotional well-being. This work proposes a NLP-based framework for emotion extraction from social media posts and for user interest identification from browser history data. We experimented with contextual representations by fine-tuning pre-trained models like BERT combined with a fully connected layer, stacked LSTM and Bi-LSTM layers, and external resources in the form of emotion lexicons and self-attention. We evaluated our models on curated datasets derived from the combination of open datasets and achieved state-of-the-art performance. Overall, our work provides evidence of the potential for non-obtrusive extraction of online behaviour indicators that reflect the emotional well-being of patients.

CCS CONCEPTS

• **Computing methodologies** → **Natural language processing**.

KEYWORDS

real-world data, emotional well-being, online behaviour, emotion recognition, interest identification, natural language processing

ACM Reference Format:
Georgia Pantalona, Anna Barachanou, Nikolaos Loukas, Lazaros Apostolidis, Filareti Tsalakanidou, and Symeon Papadopoulos. 2023. Online Behaviour Indicator Extraction for Enhanced Cancer Patient Management using Real-World Data. In *20th International Conference on Content-based Multimedia Indexing (CBMI 2023), September 20–22, 2023, Orléans, France.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3617233.3617264

CBMI 2023, September 20–22, 2023, Orléans, France

ACM ISBN 979-8-4007-0912-8/23/09…$15.00
https://doi.org/10.1145/3617233.3617264

1 INTRODUCTION

Breast cancer is a major health concern affecting millions of women worldwide, impacting their mental health and quality of life during and after treatment [18]. Therefore, the ability to monitor and assess the emotional state of cancer patients and survivors is crucial for providing them effective care and support. Up until now, traditional self-reported questionnaires have been used, but the digital transformation of society offers new opportunities for observing and monitoring a patient's mental health. Real-world data, such as online activity and social media, can provide valuable insights into the mental health of patients and help medical professionals identify and address issues more effectively [22, 24]. In this paper, we aim to present the potential of leveraging real-world data for breast cancer patient management enhancement and improvement of the overall quality of life of patients and survivors.

Leveraging social media and web activity data presents us with the opportunity to extract a range of indicators that can provide mental health experts with insights into a patient's emotional state. Another benefit to monitoring online behaviour data, is that it allows for continuous evaluation of a patient's emotional state and thus uncovering their emotional ups and downs throughout their journey, unlike traditional methods like questionnaires that only sparsely capture the emotional state at given points in time. Introduction One of the challenges in the field of our research has been the lack of literature relative to techniques to be implemented and 24/7 monitoring of users' and, specifically, patients' online activity. There is also a lack of high quality open-source datasets; existing datasets are often unbalanced and poorly annotated, while in addition there is a lack of datasets focused on health.

In this work, we present a framework that has been developed in the scope of the REBECCA[1] H2020 project, which aims to tap into the potential of real-world Data to support clinical research and to improve existing clinical workflows. This framework utilizes social media data and web activity data to extract online behavior indicators relevant to the emotional well-being and quality of life of breast cancer patients during the post-treatment period. The framework comprises two modules: emotion extraction from social media posts and user interest identification from web browsing activity. Due to the aforementioned challenges, this kind of framework is really hard to be evaluated and the results have to be taken

[1] REsearch on BrEast Cancer induced chronic conditions supported by Causal Analysis of multi-source data, https://rebeccaproject.eu/

into consideration with caution. However, our emotion and topic classification models achieved performance close to the best results in the literature and showed promising results when tested on actual samples of social media and web activity user data. Subsequently, the results obtained from these models can be utilized for further statistical and causal analysis, as well as for visualizations on clinical dashboards.

In summary, our work makes the following contributions:

- We present a pipeline to extract online behavior indicators from social media and web activity data for extracting insights into the emotional well-being and quality of life of breast cancer patients.
- We experiment with various established machine learning (ML) models and deep learning models for emotion classification and user interest identification, achieving results close to the state-of-the-art.
- Our proposed models offer a complementary tool for healthcare providers to better understand and monitor the emotional well-being and quality of life of breast cancer patients during the post-treatment period.

2 RELATED WORK

2.1 Emotion analysis from text

Natural Language Processing (NLP) has gained importance in emotion recognition from text. Researchers have used lexicons and various NLP techniques for the identification of emotions from text. For example, Kusen et al. [8] combined NLP techniques to produce linguistic features and found that the NRC lexicon performs better in identifying emotions like anger, fear and joy, whereas DepecheMood identifies sadness better. In another study, Bandhakavi et al. [4] created a domain-specific emotion lexicon (DSEL) using a generative unigram mixture model and found that DSEL outperformed other methods such as PMI and sLDA. Researchers have also used deep learning models for emotion recognition, such as Bi-LSTM and DNN and GRU and DNN.

Winata et al. [9] compared the performance of several pre-trained word and emotion embeddings including Emo2vec, BERT, Emoji2Vec, DeepMoji, ELMo and GLoVe, and found that the best performance was observed from the DeepMoji since it was trained on a large emotion corpus. Gievska et al. [20] combined three different lexicons and assessed several classification algorithms, out of which SVM gave the best results. Haggag [12] proposed a knowledge-based artificial neural network (KBANN), which stores semantic and syntactic information and the emotions are recognized via a matching process; results showed that the proposed model outperformed other approaches such as supervised learning models or keyword spotting.

Shaheen et al. [21] combined a rule-based approach with a learning-based approach, created an annotated reference that is used to describe the emotion of a sentence called emotion extraction rules (EERs), and then compared the EER of the new sentence with the annotated EERs using a KNN classifier. Barron Estrada et al. [16] focused their research on learning environments, more specifically, emotion recognition from student feedback. They tried to identify learning-centered emotions using various algorithms. The best performing algorithms were BERT, EvoMSA, B4MSA, and SVM.

2.2 Mental health analysis from text and online behaviour

Two approaches are followed in mental health research: a generic approach that focuses on population level surveillance and another that focuses on individuals and early identification of mental issues. Skaik and Ipken [23] analyzed 110 publications that use social media posts to predict mental issues, focusing on suicide and depression using ML and NLP techniques. Fine et al. [5] focused their research on a population segment, healthcare providers, and the impact of specific events, such as Covid-19 and George Floyd's death, on their mental health. Coppersmith et al. [10] created a model that predicts suicide risk using public data from social media interactions, including word embeddings, LSTM, and attention layers. The performance of their model is quite high, indicating that signals that can indicate suicide risk exist in social media.

The identification of depression in social media has been researched extensively, with various approaches being used, such as lexical and emotion-based analysis. Jamil et al. (2019) [28] and Peng et al. (2015) [14] utilized ML models such as SVM and multi-kernel SVM to identify depression in users based on their tweets, user profiles, and behaviors. Wolohan et al. (2018) [26] analyzed a corpus containing depression-related posts and found that lexical models still perform adequately without them. Chen et al. [27] incorporated emotion-based features produced by the EMOTIVE algorithm and temporal analysis of these features to improve the classification task. Gallegos Salazar et al. (2020) [11] went a step further by including sentiment, sarcasm, intent, and abuse in their analysis, which outperformed several state-of-the-art classifiers reported in the literature, and produced more interpretable results.

3 PROPOSED FRAMEWORK

The goal of the proposed framework is to extract online behaviour indicators from the collected raw online activity data (Figure 1). Textual data from social media and web browsing sessions are fed as input and are first preprocessed. Subsequently, features are extracted using different NLP techniques such as word embeddings, use of lexicons, pre-trained models and others. Those features are then used as input to different classification models such as SVM, LSTM, and CNN, in order to extract indicators that can provide insights about users' emotional well-being. These indicators include among others sentiment or emotion scores as well as indicators of anxiety, depression and stress levels. Additionally, the website visits and online search queries of the patient can be classified to topics of interest, such as health, news, shopping, sports, etc.

All of the employed NLP methods leverage embeddings from pre-trained text-based models as their input, since such models are important components of many approaches that attain state-of-the-art performance in downstream NLP tasks, particularly in emotion classification tasks. These pre-trained models offer text representations that incorporate details about every word's context, as well as the overall context of full sentences.

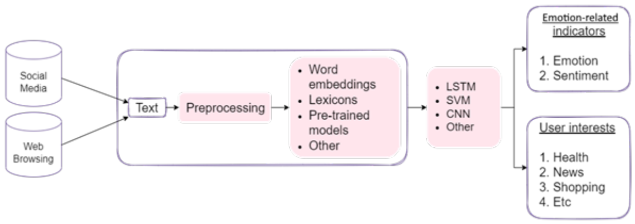

Figure 1: Online activity data processing pipeline.

3.1 Browser plugin

The raw online activity data is collected via a browser plugin specifically designed for that purpose. The plugin tracks the patient's online activities, such as website visits, search engine queries, and social media interactions (e.g., likes, comments, etc.) on platforms like YouTube and Facebook. To protect the privacy of the patients, all collected data points are pseudonymized, while patients also have the ability to customize their online activity tracking. The plugin has been designed to be user-friendly and easy to install, with a focus on the security and privacy of collected data. It operates in a non-invasive way, with minimal impact on the user's browsing experience, and works seamlessly in the background with little user intervention.

3.2 Online behaviour indicator extraction

Our main goal is to quantify the emotional well-being of patients using information extracted from their online activities and calculate indicators for their emotional state and interests in topics they search or read about on the Web. The data is collected using the plugin and is processed through pre-processing, feature extraction, and classification to extract relevant indicators.

3.2.1 Emotion extraction. Several techniques have been developed for the extraction of emotion from social media posts or comments. Posts can convey the emotions of the patient, offering insights on how they are feeling about what they are experiencing.

Our emotion extraction module is employing state-of-the-art technologies to classify text to a specific emotion, using as input the text collected from social media sources using the browser plugin; more specifically, the posts and comments made by the user in Facebook as well as YouTube. Each post/comment is processed separately and an emotion is assigned to it. The final output of this process is an aggregation of the emotions identified in a period of time. For the emotion extraction task, two different approaches were studied: a multi-class approach, where the classifier assigns one emotion label to each text (e.g., happiness, sadness, worry, etc.) and the multi-label approach where one or more labels can be assigned to each text.

In order to make the predictions, we used the BERT representation with LSTM and Bi-LSTM layers, emotion-related lexicons, and self-attention. The output probabilities were transformed into predictions using a threshold for the multi-label problem and selecting the class with the highest probability for the multi-class problem.

It is notable that a large amount of methods that achieve state-of-the-art performance in downstream NLP tasks and specifically in the emotion classification task, heavily rely on transfer learning and on the usage of pre-trained models. These pre-trained models provide text representations, which include information regarding all surrounding words and the general context of the text. Therefore, all of the employed architectures in this work rely on the output of the pre-trained transformer models, which is in short, the contextual representation of the input text. More specifically, we use the pre-trained BERT model (bert-base-uncased), which was first published by Devlin et al. [13] in 2019 and is publicly available.

3.2.2 User interest extraction. Regarding user interest identification, we have developed novel NLP models for the classification of visited URLs to topics (e.g., health, news, sports, shopping, etc.). These include traditional ML approaches, namely KNN, Multinomial Naïve Bayes, Logistic Regression, Random Forest, and Linear SVM, but also deep learning approaches like CNN and fine-tuned BERT embeddings with a linear classifier. The browsing data can also provide hints into the patient's emotional well-being. For example, if a patient visits health-related websites with high frequency or higher frequency compared to previous periods, this could possibly be an indication of increased stress or anxiety with regard to their health condition. For this specific problem, we used a multi-class approach.

3.3 Visualization of indicators

Figure 2 illustrates an example of how the results of our models could be visualized in the form of a dashboard, allowing experts to interpret these results and extract insights.

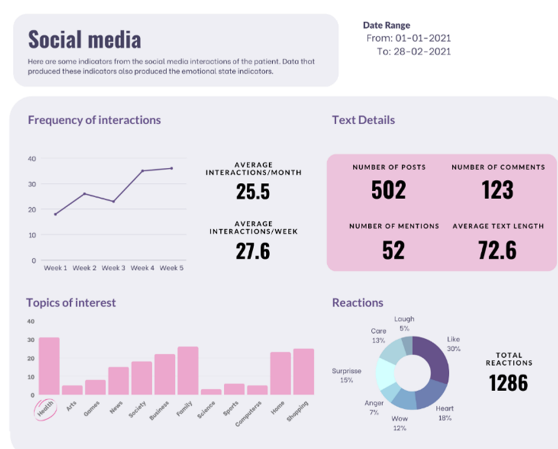

Figure 2: Visualisation of online activity indicators extracted from social media.

4 EXPERIMENTS

4.1 Emotion extraction

4.1.1 Datasets. For training and testing of our emotion recognition models, four datasets were used [1, 6, 7, 17], which contained tweets labelled with various emotions. However, when experimenting with these datasets for the multi-label problem, satisfactory results were not obtained. As a result, a new dataset was curated by merging datasets [1], [6] and [17], with the aim of retaining six specific emotions: happiness, sadness, worry, love, anger, and surprise.

4.1.2 *Architectures.*

Three different model architectures for emotion classification have been employed. The first is a baseline method using contextual sentence representation. The second involves stacked LSTM and Bi-LSTM to incorporate contextual information and the order of text. The third model includes external emotion-related resources and a self-attention mechanism to boost performance. All models can be used for both multi-class and multi-label problems by adjusting the transformation of probabilities to predictions. The multi-label problems are treated as multiple binary classification problems and the multi-class problems assign the highest probability class to the input text.

- **BERT Fine-tuning with a fully connected layer:** To set the baseline performance for the methods using pre-trained models, we first experiment with a simple fine-tuning architecture. Our approach involves passing the tokenized input through a pre-trained BERT model and extracting the contextual representation of the input text by using the last hidden state of the classification token. This representation is then fed through a fully connected layer followed by a ReLU activation layer. We also incorporate dropout to reduce overfitting and improve the generalization error of the neural network during training. The output of this layer is passed through another fully connected layer resulting in logits, which are then transformed into probabilities using a sigmoid layer for multi-label problems or a softmax layer for multi-class problems.
- **Stacked LSTM and Bi-LSTM:** We used neural architectures such as LSTM and Bi-LSTM for emotion classification tasks, as they perform well in understanding temporal features and incorporating context. We used word vectors from the pre-trained models and passed them through different RNN layers with varying parameter settings. We found that the best results were achieved with a Bi-LSTM layer with 200 hidden states and a dropout of 0.3, in addition to two fully connected layers.
- **External resource boosting with self-attention mechanism:** We propose a model that combines the contextual representation of the text with the emotional knowledge of two lexicons. The emotion knowledge from the two lexicons transforms the input text to two feature vectors that are fed to a fully connected layer. The output of this layer is concatenated with the final hidden state of the classification token to create the Key of an attention mechanism. The attention mechanism uses the final hidden state of the pre-trained model as the Query, and multiplies the Key and Query to provide the attention score. The softmax of the attention score is then multiplied with the Key to output the final representation of the input text, which is passed through fully connected layers and a dropout layer to finally output the logits for the classification task. This model is based on the Sentence-level KEA proposed by Suresh and Ong [25].

Unique multi-class emotion classification architectures. We experimented with different models to establish a baseline for the multi-class emotion classification problem. We initially used established machine learning algorithms with features created from the text using the TF-IDF representation. Then we experimented with CNNs using FastText word embeddings and LSTMs using the same FastText word embeddings. The CNNs and LSTMs showed promise in improving the performance of the traditional models by detecting relevant patterns in the textual data regardless of their position in the sentence and understanding and interpreting temporal features, respectively. Additionally, we fine-tuned the T5 [19] model on the merged dataset; T5 is an encoder-decoder model that can convert all NLP problems into a text-to-text task, providing the same model, loss function, and hyperparameters for all NLP tasks. Finally, we tested ALBERT [15], a light version of the BERT model, and we noticed that it could match BERT's performance with approximately a quarter of the parameters.

4.1.3 *Results.*

Multi-class emotion classification. As can be seen in Table 1, the traditional machine learning models that use TF-IDF features are the worst performing, but the linear SVM classifier actually has comparable results to the other methods. Additionally, amongst all experiments with CNN and LSTM architectures using different word embeddings, FastText performed the best. Combining BERT embeddings and LSTM increased F1 by 1% compared to the FastText-LSTM combination. Finally, the best results of this set of experiments with the BERT model were produced when using the embeddings from the pre-trained BERT model and including a boosting of the emotional context using external lexicon resources. This architecture improved results by 2% comparing to the FastText-LSTM results and by 8% comparing to the linear SVM. The overall best F1 score was obtained by the T5 model, but at a considerable computational cost, since T5 is twice as big as the BERT model. In Table 2, we provide some sample output predictions (both true and false) for the test dataset using the trained T5 model. Observing the results, we notice that in some cases the model misclassifies the text; however the predicted label seems to be relevant in a wider context or as a second attribute.

Model type	M-F1
Logistic Regression	0.5776
Random Forest	0.5388
Linear SVM	0.6127
fasttext - CNN	0.6626
fasttext - Bi-LSTM	0.6725
fasttext - LSTM	0.6755
BERT	0.6773
BERT – LSTM	0.6845
BERT – Lexicon features	0.6944
T5	**0.7144**
ALBERT	0.6793

Table 1: Performance of all models on the merged dataset for the multiclass problem, evaluated using M-F1 score.

Multi-label emotion classification. For this problem, we experimented with various methods including different activation functions and schedulers. We concluded that the combination of

Tweet	Real	Predicted
Your future is bright. #Remember	Happiness	Happiness
Feel the wrath of my air-dropped pamphlets!	Happiness	Anger
mmmm i'm kinda sad i hope i can shake this before school	Sadness	Worry

Table 2: Sample outputs of the emotion detection module

the chained scheduler with a softmax function in the attention mechanism gives better results in 16 out of 27 emotion categories, whereas the next best combination gives the best results for only nine categories. Additionally, the modified model performs better in the under-represented emotion categories such as embarrassment, grief, nervousness, pride, and relief compared to the baseline model.

Since the multi-label classification did not show satisfying performance compared to its multi-class counterpart, we decided to focus our analysis on the latter.

4.2 User interest extraction

4.2.1 Datasets. For the URL classification problem, we created a dataset of 24,000 URLs by randomly selecting URLs from two publicly available datasets and assigning them to 16 categories. The first dataset contains around 1.56 million labelled URLs with five categories [2], while the second dataset contains around 31,000 labelled URLs with 24 categories [3]. However, assigning categories to the URLs based solely on their path might not always be sufficient. For example, a URL like https://shoppinghouse.gr/ provides hints that it might be an online shopping website, while a URL like https://www.bbc.com/ does not provide any information about the website's topic. To overcome this, additional information about the website, such as its title, description, and keywords, were gathered from the source page of the URL.

4.2.2 Architectures. In our experiments for URL topic classification, we tested various machine learning and deep learning models. The traditional machine learning models we tested include K Nearest Neighbors, Multinomial Naïve Bayes, Logistic Regression, Random Forest, and Linear Support Vector Classifier. We also used two ways of representing document vectors, namely using words as tokens and using character n-grams as tokens. To improve classification performance, we implemented a simple CNN with trainable embeddings on the corpus and another CNN with pre-trained GloVe embeddings. Finally, we used BERT for fine-tuning in the same architecture as the simple CNN. All models utilized the website title, description, and keywords as input, excluding the URLs. Our results show that deep learning models, particularly those with pre-trained embeddings, outperform traditional machine learning models.

4.2.3 Results. Initially, we only used the URLs as input to the models, but we found that enriching the dataset with additional information such as the title, description, and keywords of the web pages led to better performance. We performed several experiments using different machine learning algorithms and enriched datasets; starting with traditional models with the best performing model being a Linear SVC with (3,6) TF-IDF n-gram features, achieving a Macro F1 score of 0.68 (see Table 3). More advanced models such as CNNs and BERT were also tested, with fine-tuned BERT achieving the highest score of 0.69.

Model type	M-F1
Multinomial Naïve Bayes	0.64
CNN	0.64
Multinomial Logistic Regression	0.65
Linear SVC	0.68
BERT fine-tuned	**0.69**

Table 3: Model performance on topic classification, evaluated using M-F1 score.

We found out that by analyzing the type of websites visited by a user we can gain a better understanding of their online activities. Moreover, as a next step, classifying health topics into specific health subcategories (e.g., treatment plan, medical advice, test result analysis, etc.) to gain insights into potential health concerns of the patient can be valuable in developing personalized healthcare interventions and improving patient outcomes.

5 CONCLUSIONS AND FUTURE WORK

This analysis provided the necessary foundations extracting online user activity indicators that will allow to reveal useful insights into breast cancer patients' emotional well-being during the post-treatment period. Our research focused on two distinct research topics: emotion extraction from social media posts and user interest identification from Web browsing activity. Social media posts can convey the emotions of patients, revealing how they are feeling about what they are experiencing, while the sites the user visited can also provide hints into the patient's emotional well-being by revealing what their interests are and how they change over time.

With regard to the emotion extraction module, we focused on two types of emotion classification: multi-class and multi-label classification. Regarding the first problem, experimentation began from the traditional ML models after the documents were transformed into vectors using TF-IDF, and then while aiming to improve classification accuracy, deep learning models like CNN, LSTM and bi-LSTM combined with utilizing FastText word embeddings were implemented. Additionally, we used embeddings from a pre-trained BERT model, including external lexicon features. The best performing model was the T5 model.

Concerning the multi-label emotion classification, inspired by models in literature such as BERT, we made an effort to improve classification results by using weighted loss for imbalanced datasets, a scheduler, and sparsemax. The optimal model turned out to be BERT using external emotion lexicons, a chained scheduler and a softmax function in the attention mechanism.

The user interest identification module exploited the patient's browsing activity, including visited URLs and submitted search queries, which were classified into 16 categories/topics. The multi-class topic classification problem was initially approached by using traditional machine learning models utilizing character N-grams and TF-IDF. In order to improve topic classification results, deep learning techniques were also implemented, such as CNN, CNN using GloVe embeddings and BERT. The input to this model consists of the text coming from the meta tags of title, description and keywords of the sites visited. Here, BERT obtained the best performance.

Expanding from the concept of emotion extraction, a natural future step will also be to examine NLP techniques for stress, anxiety and depression indicator extraction from the patient's social media behavior. Additionally, we will focus on leveraging the temporal dimension of online activity to enhance indicators and analyze the changing emotional well-being of cancer patients during the post-treatment period.

It is vital to note that the accuracy of the aforementioned insights may be hindered by factors such as limitations in data collection, processing, and interpretation, as well as different patient coping strategies, including approach-based and avoidance-based strategies. Therefore, while online behaviour monitoring data offers potential benefits, it is important to carefully consider its limitations when using it to assess patients' emotional well-being.

ACKNOWLEDGMENTS

This work has been partially funded by the H2020 project "REBECCA: REsearch on BrEast Cancer induced chronic conditions supported by Causal Analysis of multi-source data" under Grant Agreement no. 965231 (https://rebeccaproject.eu/).

REFERENCES

[1] [n. d.]. Emotions dataset for NLP. ([n. d.]). https://www.kaggle.com/datasets/praveengovi/emotions-dataset-for-nlp

[2] [n. d.]. URL classification dataset. ([n. d.]). https://www.kaggle.com/code/shawon10/url-classification-by-naive-bayes/input

[3] [n. d.]. URL clategorization dataset. ([n. d.]). https://data.world/crowdflower/url-categorization

[4] J. Blase J. Carroll A. Fine, P. Crutchley and G. Coppersmith. Jul. 2017. Lexicon based feature extraction for emotion text classification. *Pattern Recognit. Lett.* vol. 93 (Jul. 2017), 133–142. https://doi.org/10.1016/j.patrec.2016.12.009

[5] J. Blase J. Carroll A. Fine, P. Crutchley and G. Coppersmith. Nov. 2020. Assessing population-level symptoms of anxiety, depression, and suicide risk in real time using NLP applied to social media data. (Nov. 2020), 50–54. https://doi.org/10.18653/v1/2020.nlpcss-1.6

[6] crowdflower. [n. d.]. Sentiment Analysis in Text. ([n. d.]). https://data.world/crowdflower/sentiment-analysis-in-text

[7] Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. GoEmotions: A Dataset of Fine-Grained Emotions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 4040–4054. https://doi.org/10.18653/v1/2020.acl-main.372

[8] K. Figl M. Conti E. Kušen, G. Cascavilla and M. Strembeck. Aug. 2017. Identifying Emotions in Social Media: Comparison of Word-Emotion Lexicons. (Aug. 2017), 132–137. https://doi.org/10.1109/FiCloudW.2017.75

[9] G. I. Winata et al. [n. d.]. CAiRE-HKUST at SemEval-2019 Task 3: Hierarchical Attention for Dialogue Emotion Classification. ([n. d.]). http://arxiv.org/abs/1906.04041

[10] P. Crutchley G. Coppersmith, R. Leary and A. Fine. Jan. 2018. Natural Language Processing of Social Media as Screening for Suicide Risk. *Biomed. Inform. Insights* vol. 10 (Jan. 2018), 1178222618792860. https://doi.org/10.1177/1178222618792860

[11] Medina-Pérez Gallegos Salazar, Loyola-González. [n. d.]. An Explainable Approach Based on Emotion and Sentiment Features for Detecting People with Mental Disorders on Social Networks. *Appl. Sci.* vol. 11 ([n. d.]). https://doi.org/10.3390/app112210932

[12] M. Haggag. [n. d.]. Frame Semantics Evolutionary Model for Emotion Detection. *Comput. Inf. Sci.* vol. 7, no. 1 ([n. d.]), 352–359. https://doi.org/10.5539/cis.v7n1p136

[13] K. Lee J. Devlin, M.-W. Chang and K. Toutanova. [n. d.]. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. ([n. d.]). http://arxiv.org/abs/1810.04805

[14] P. Yin Y. Zhou A. Yang L. Cao, S. Peng and X. Li. Dec. 2020. A Survey of Emotion Analysis in Text Based on Deep Learning. (Dec. 2020), 81–88. https://doi.org/10.1109/iSCI50694.2020.00020

[15] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. *CoRR* abs/1909.11942 (2019). arXiv:1909.11942 http://arxiv.org/abs/1909.11942

[16] R. Oramas Bustillos M. L. Barrón Estrada, R. Zatarain Cabada and M. Graff. [n. d.]. Opinion mining and emotion recognition applied to learning environments. *Expert Syst. Appl.* 150 ([n. d.]), 113265. https://doi.org/10.1016/j.eswa.2020.113265

[17] Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. SemEval-2018 Task 1: Affect in Tweets. (2018). https://aclanthology.org/S18-1001

[18] World Health Organisation. 2021. Breast cancer. (2021). https://www.who.int/news-room/fact-sheets/detail/breast-cancer

[19] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21, 140 (2020), 1–67. http://jmlr.org/papers/v21/20-074.html

[20] K. Koroveshovski S. Gievska and T. Chavdarova. [n. d.]. A Hybrid Approach for Emotion Detection in Support of Affective Interaction. ([n. d.]), 352–359. https://doi.org/10.1109/ICDMW.2014.130

[21] H. Hajj S. Shaheen, W. El-Hajj and S. Elbassuoni. [n. d.]. Emotion Recognition from Text Based on Automatically Generated Rules. ([n. d.]), 383–392. https://doi.org/10.1109/ICDMW.2014.80

[22] Jenesh Singh Shrestha, Alish Shrestha, Abja Sapkota, Rakshya Sharma, Samip Shrestha, Sudip Shestha, Kapendra Sekhar Amayta, and Madhav Gautam. 2017. Social support, quality of life and mental health status in breast cancer patients. *Cancer Rep Rev* 1 (2017).

[23] R. Skaik and D. Inkpen. Dec. 2020. Using Social Media for Mental Health Surveillance: A Review. *ACM Comput. Surv.* vol. 53 (Dec. 2020), 129:1–129:31. https://doi.org/10.1145/3422824

[24] K M Trevino, K Fasciano, S Block, and H G Prigerson. 2013. Correlates of social support in young adults with advanced cancer. *Support Care Cancer* 21 (2013). https://pubmed.ncbi.nlm.nih.gov/22790223/

[25] Suresh V. and On D. C. 2021. Using Knowledge-Embedded Attention to Augment Pre-trained Language Models for Fine-Grained Emotion Recognition. (2021). http://arxiv.org/abs/2108.00194

[26] Mukherjee Atreyee Sayyed Zeeshan Ali Wolohan JT, Hiraga Misato and Millard Matthew. [n. d.]. Detecting Linguistic Traces of Depression in Topic-Restricted Text: Attending to Self-Stigmatized Depression with NLP. vol. 10 ([n. d.]), 32–40. https://aclanthology.org/W18-4102

[27] Thomas W. Jackson Xuetong Chen, Martin D. Sykora and Suzanne Elayan. [n. d.]. What about mood swings? Identifying depression on Twitter with temporal measures of emotions. ([n. d.]). https://doi.org/10.1145/3184558.3191624

[28] P. Buddhitha Z. Jamil, D. Inkpen and K. White. 2017. Monitoring Tweets for Depression to Detect At-risk Users. vol. 10 (2017), 32–40. https://doi.org/10.18653/v1/W17-3104