JA-NLP@LT-EDI: Empowering Mental Health Assessment: A RoBERTa-Based Approach for Depression Detection

Jyoti Kumari
Department of CSE
Siksha 'O' Anusandhan,
Deemed to be University
Bhubaneswar, India
j2kumari@gmail.com

Abhinav Kumar
Department of CSE
Motilal Nehru National Institute
of Technology Allahabad,
Prayagraj, India
abhinavanand05@gmail.com

Abstract

Depression, a widespread mental health disorder, affects a significant portion of the global population. Timely identification and intervention play a crucial role in ensuring effective treatment and support. Therefore, this research paper proposes a fine-tuned RoBERTa-based model for identifying depression in social media posts. In addition to the proposed model, Sentence-BERT is employed to encode social media posts into vector representations. These encoded vectors are then utilized in eight different popular classical machine learning models. The proposed fine-tuned RoBERTa model achieved a best macro F_1-score of 0.55 for the development dataset and a comparable score of 0.41 for the testing dataset. Additionally, combining Sentence-BERT with Naive Bayes (S-BERT + NB) outperformed the fine-tuned RoBERTa model, achieving a slightly higher macro F_1-score of 0.42. This demonstrates the effectiveness of the approach in detecting depression from social media posts.

1 Introduction

Depression is a prevalent mental health disorder affecting people of all ages from diverse backgrounds, irrespective of their socioeconomic status or cultural circumstances. The World Health Organization (WHO) approximates that there are over 264 million individuals globally who suffer from depression, underscoring its significant impact on public health. Depression exhibits a higher occurrence rate in women, surpassing that in men by approximately 50%. Alarmingly, suicide claims the lives of over 700,000 people each year, positioning it as the fourth leading cause of death among individuals aged 15 to 29[1] (Vioules et al., 2018).

The emergence of social media platforms has revolutionized online communication, information sharing, and self-expression. By 2021, Facebook alone had reported a staggering 2.8 billion monthly active users, while other platforms like Twitter, Instagram, and Reddit also maintained substantial user bases. This widespread adoption of social media has presented researchers with an extensive pool of user-generated content to investigate, including its application in mental health analysis and other diverse endeavors (Guntuku et al., 2017; Kumar and Kumari, 2021; Kumari and Kumar, 2021b; Gkotsis et al., 2017). There are several advantages of detecting signs of depression from social media posts. Firstly, it offers a large-scale and easily accessible data source, allowing for real-time analysis of a significant number of individuals. Secondly, social media text often reflects individuals' genuine emotions and experiences, as they express themselves freely and spontaneously on these platforms. Additionally, social media data can be collected over time, enabling the monitoring of changes in individuals' mental well-being as it evolves.

In recent years, there has been increasing interest in analyzing social media texts to detect signs of depression. Research has revealed that individuals suffering from depression often demonstrate specific language patterns in their online posts (Guntuku et al., 2019; Schwartz et al., 2013). Studies have found that depressed individuals tend to utilize more self-referential pronouns (such as "I" and "me"), expressing a higher frequency of negative emotions and words, and exhibit reduced engagement in positive social interactions (Coppersmith et al., 2014; Park et al., 2012).

However, the detection of depression signs from social media texts is not without its challenges. While specific linguistic patterns associated with depression have been identified (Coppersmith et al., 2014; Park et al., 2012; Guntuku et al., 2017), it is important to consider individual differences, personality traits, cultural variations, and context-

[1] https://www.who.int/news-room/fact-sheets/detail/depression

LT-EDI 2023 – Third Workshop on Language Technology for Equality, Diversity and Inclusion, pages 89–96, Varna, Bulgaria, Sep 7, 2023.
https://doi.org/10.26615/978-954-452-084-7_013

specific factors that can influence the use of language (Schwartz et al., 2013; De Choudhury et al., 2014). By exploring diverse approaches, techniques, challenges, and ethical considerations (Al-Sagri and Ykhlef, 2020; Kumari and Kumar, 2021a; He et al., 2022), our aim is to contribute to the continuous efforts of enhancing mental health outcomes through the innovative and responsible use of social media data. In accordance with existing research, this study introduces a refined RoBERTa model to detect depression in social media posts. Additionally, Sentence-BERT is employed to convert social media posts into a consistent encoded vector. These vectors are subsequently utilized as input for various well-known classical machine learning classifiers.

The remaining paper is organized as follows: Section 2 briefs the related work, Section 3 discusses the methodology adopted to perform the work, Section 4 highlights the results obtained, and finally Section 5 concludes the overall work.

2 Related Works

This section provides insights into the field of detecting signs of depression from social media texts. These highlights the use of linguistic analysis, sentiment analysis, machine learning, psychological frameworks, and natural language processing, providing a comprehensive understanding of the field (Park et al., 2012; Guntuku et al., 2017; Schwartz et al., 2013; Saumya et al., 2021). Machine learning algorithms play a vital role in identifying depression through text extracted from social media. Various supervised learning techniques, such as Support Vector Machines (SVM), Naïve Bayes, and Random Forests, have been utilized to classify text as depressed users or non-depressed (AlSagri and Ykhlef, 2020; Angskun et al., 2022). Moreover, deep learning approaches like Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs) have demonstrated encouraging outcomes in capturing intricate linguistic patterns to detect depression (Tadesse et al., 2019).

Coppersmith et al. (2014) explores the use of Twitter data to quantify mental health signals, including depression. In this study, linguistic features including words, phrases, and linguistic structures are employed to detect mental health signals linked to depression, Post-traumatic stress disorder (PTSD), and other mental health conditions. To accomplish this, machine learning techniques are utilized to classify tweets and evaluate the models' accuracy in predicting mental health states. In addition, the paper discusses the challenges faced when quantifying mental health signals within Twitter data, such as the presence of inherent noise and bias in social media posts. Further, by incorporating various data sources, the research strives to provide a more comprehensive and robust analysis of mental health indicators.

The authors in (Park et al., 2012) with an objective of gaining insights into the expression and analysis of depressive emotions in the realm of social media, employed a large dataset of Twitter posts, and leverage natural language processing techniques to identify patterns and sentiments related to depressive language. Their study focuses on understanding how individuals express and communicate their feelings of depression through social media platforms like Twitter. The authors examine linguistic features in Twitter, such as words, emoticons, and grammatical structures, to understand the occurrence and expression of depressive moods. They discover distinct language patterns associated with depressive emotions. However, relying solely on Twitter data limits the study's representativeness and generalizability to diverse demographics.

In their study, Guntuku et al. (2017) conducts a systematic analysis of various studies utilizing social media data to detect depression and other mental health conditions. They provide a comprehensive overview of the research conducted in the field of detecting depression and mental illness through social media analysis by providing an in-depth assessment of their strengths and limitations. It examines the linguistic, behavioral, and contextual features utilized to detect signs of depression in social media posts from the application point of view of machine learning and natural language processing techniques. The challenges encompass concerns related to data privacy, the representativeness of social media users, biases arising from self-disclosure, and the requirement for more accurate ground truth labels to effectively train models.

Reece and Danforth (2017) takes a unique approach by leveraging the vast amount of user-generated data on Instagram. Their primary goal is to uncover predictive indicators of depression. To achieve this, the researchers meticulously analyzed a substantial number of Instagram photos posted by participants, along with the accompanying captions and metadata. The study revealed a significant pattern where participants diagnosed with depression exhibited a strong preference for darker, grayer, and bluer tones with less likely smiling faces in their Instagram photos, while those without depression tended to favor brighter, warmer colors. Further, the study also revealed that individuals with depression tended to engage more in frequent posting for social validation and support. Additionally, the researchers identified distinct textual markers like increased utilization of filters and more frequent references to feelings of sadness, anxiety, and loneliness. These findings provide valuable insights into the relationship between depression and online behavior, particularly in terms of text-based expression. The research primarily relies on self-reported diagnoses of depression and lacks a comprehensive clinical assessment conducted by mental health professionals. Furthermore, the study's findings are confined to data obtained solely from Instagram, which may not provide a complete representation of the broader population.

Gkotsis et al. (2017) focuses on characterizing mental health conditions by leveraging informed deep learning models to analyze user-generated content, including posts and interactions on social media platforms. The results demonstrate promising capabilities in detecting symptoms related to depression, PTSD, self-harm, and more. Additionally, the study uncovers distinct linguistic patterns and behavioral markers associated with different mental health conditions, encompassing variations in word usage, sentiment analysis, and linguistic styles employed by individuals facing diverse mental health challenges. However, the research relies on publicly available social media data, which may not fully represent the entire population. Furthermore, considerations regarding privacy, data accuracy, and generalizability need to be addressed.

Guntuku et al. (2019) adopts an innovative approach by examining the language patterns of adults with attention deficit hyperactivity disorder (ADHD) in their social media interactions. It uncovers distinct linguistic characteristics, including a higher frequency of self-referential pronouns, words related to time urgency, emotional language, and references to cognitive processes. These findings highlight the potential of language analysis to identify and understand communication patterns associated with ADHD in digital environments. However, the study relies on self-reported diagnoses and focuses on language patterns within specific social media platforms. Therefore, it may not fully encompass the experiences of all adults with ADHD or capture their interactions across a wide range of digital platforms.

Schwartz et al. (2013) explores the intriguing interplay between personality traits, gender, age, and language usage on social media platforms. Employing an open-vocabulary approach, the authors delve into an extensive analysis of large-scale social media data to uncover patterns and correlations between language use and individual characteristics. The linguistic analyses identified distinct linguistic markers associated with various personality traits, including extraversion, neuroticism, and agreeableness. Further, the study uncovers correlations between gender and linguistic choices, revealing variations in language use between male and female users. Notably, age-related disparities in language were also observed, indicating that language patterns on social media may undergo transformations as individuals grow older. The study relies on self-reported personality assessments that focus on language patterns within specific social media platforms. As a result, the study may not fully encompass the complete spectrum of personality traits, gender identities, or age groups within the broader population.

De Choudhury et al. (2014) aims to characterize and predict postpartum depression by utilizing data shared on Facebook. The authors explore different variables including language patterns, linguistic styles, social interactions, and behavioral cues exhibited by users, with the goal of identifying significant markers associated with postpartum depression. Further, they discover distinct language patterns like increased use of first-person pronouns, negatively affecting words, feelings of loneliness and isolation, etc. They also highlight the importance of social interactions and behavioral cues, such as changes in posting frequency and decreased engagement with friends, as potential indicators of postpartum depression. However, the study's reliance on Facebook data may limit the representation and understanding of individuals beyond the scope of the platform itself.

The authors in (Tadesse et al., 2019) address the prevalence of suicide and the growing influence of social media platforms in shaping public discourse. The paper emphasizes the significance of automated systems in monitoring and identifying individuals who may be at risk. Their research un-

Table 1: Data statistic used to validate proposed model

Class	Train	Dev	Test
Moderate Depression	3678	2169	275
No Depression	2755	848	135
Severe Depression	768	228	89
Total	7201	3245	499

derscores the importance of proactive measures to detect and support individuals in need within the context of social media platforms. The authors employ a comprehensive dataset consisting of posts from social media forums, which is annotated by mental health professionals. They extract textual features using pre-trained word embeddings and apply a deep learning model, specifically a CNN, for classification. However, the information reagrding the characteristics of the dataset, such as size, diversity, or representativeness is missing. Further, it does not explicitly address how the proposed deep learning model accounts for the potential challenges of adapting the model to new language patterns and emerging trends.

3 Methodology

The overall flow diagram of the proposed work is illustrated in Figure 1. In addition to the fine-tuned RoBERTa model, a total of eight different models were developed, namely: (i) Fine-tuned RoBERTa, (ii) Sentence-BERT + Support Vector Machine (S-BERT + SVM), (iii) Sentence-BERT + Random Forest (S-BERT + RF), (iv) Sentence-BERT + Logistic Regression (S-BERT + LR), (v) Sentence-BERT + K-Nearest Neighbors (S-BERT + KNN), (vi) Sentence-BERT + Naive Bayes (S-BERT + NB), (vii) Sentence-BERT + Gradient Boosting (S-BERT + GB), (viii) Sentence-BERT + Decision Tree (S-BERT + DT), and (ix) Sentence-BERT + AdaBoost (S-BERT + AB). The LT-EDI-2023 workshop dataset[2] was utilized to validate the proposed models. Table 1 (S et al., 2022) provides an overview of the data samples in the training, testing, and development sets (Sampath et al.), while the default pre-processing steps defined in the Ktrain library[3] were applied.

3.1 RoBERTa

RoBERTa (Robustly Optimized BERT approach) is a state-of-the-art natural language processing (NLP) model that has made significant contributions to various NLP tasks. Developed by Facebook AI in 2019, RoBERTa is built upon the architecture of BERT (Bidirectional Encoder Representations from Transformers) and leverages the power of unsupervised pretraining on large amounts of text data. Its advanced training techniques, such as using larger datasets, removing the next sentence prediction objective, and increasing the batch size, enables it to achieve superior performance on a wide range of NLP benchmarks. RoBERTa has demonstrated remarkable success in tasks like text classification, named entity recognition, sentiment analysis, question answering, and machine translation, showcasing its versatility and effectiveness. With its robustness, RoBERTa has become an invaluable tool for researchers, developers, and practitioners seeking cutting-edge solutions in the field of natural language processing.

Given the widespread use of RoBERTa in various NLP tasks, this study employed RoBERTa to detect depression from social media posts. To fine-tune RoBERTa, a maximum input size of 250 words was set for each data sample. The RoBERTa model was subsequently trained for 100 epochs, utilizing a learning rate of $2e^{-5}$ and a batch size of 32.

3.2 Sentence-BERT Representation

BERT is a state-of-the-art neural network architecture designed for pretraining language representations by leveraging the transformer architecture. The transformer model in BERT allows for capturing contextual information from both left and right contexts of a given word, enabling a deeper understanding of the meaning of words and sentences. The term “bert-base-nli-mean-tokens”[4] refers to a specific model architecture based on BERT (Bidirectional Encoder Representations from Transformers) that is used for natural language inference (NLI) tasks. The “bert-base-nli-mean-tokens” model specifically utilizes the ”mean pooling” strategy to generate fixed-length sentence representations. In this strategy, each word in a sentence is encoded using BERT, and the resulting

[2] https://sites.google.com/view/lt-edi-2023/home

[3] https://amaiya.github.io/ktrain/text/preprocessor.html

[4] https://huggingface.co/sentence-transformers/bert-base-nli-mean-tokens

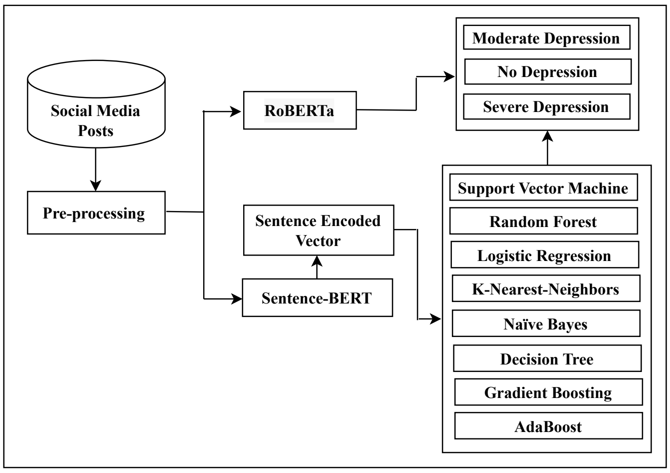

Figure 1: Flow diagram of the proposed work

contextualized word embeddings are averaged to create a single vector representation of the entire sentence. This mean pooling approach is a simple yet effective way to condense variable-length sentences into fixed-length vectors that can be fed into downstream tasks.

In this study, the text data samples underwent an initial encoding process, resulting in a 768-dimensional vector representation. Subsequently, this encoded vector, consisting of 768 dimensions, was employed in multiple well-known classical machine learning classifiers, as depicted in Figure 1.

4 Result

The performance evaluation of the proposed model encompasses various metrics, including precision (P), recall (R), F_1-score (F_1-score), accuracy (Acc), weighted precision, weighted recall, weighted F_1-score, macro-precision, macro-recall, macro-F_1-score, confusion matrix, and AUC-ROC curve. The results of validating different deep learning models on both the testing and validation datasets are presented in Table 2. Notably, among all the implemented models, the proposed fine-tuned RoBERTa model exhibited the highest macro F_1-score of 0.55 for the development dataset, as shown in Table 2. The corresponding confusion matrix and AUC-ROC curve for the proposed fine-tuned RoBERTa model are depicted in Figures 2 and 3, respectively.

In the case of Sentence-BERT combined with classical machine learning models, S-BERT + LR outperformed the classical machine learning approach, achieving a macro F_1-score of 0.51 for the development dataset. The confusion matrix and AUC-ROC curve for the S-BERT + LR model on the development dataset can be observed in Figures 4 and 5, respectively.

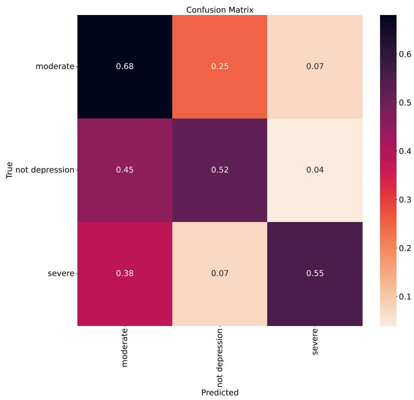

Figure 2: Confusion matrix for the proposed fine-tuned RoBERTa model (Validation dataset)

Likewise, when considering the testing dataset, the proposed fine-tuned RoBERTa model again demonstrated remarkable performance, attaining a notable macro F_1-score of 0.41. The corresponding confusion matrix and AUC-ROC curve for the fine-tuned RoBERTa model on the testing dataset are presented in Figures 6 and 7, respectively. Regard-

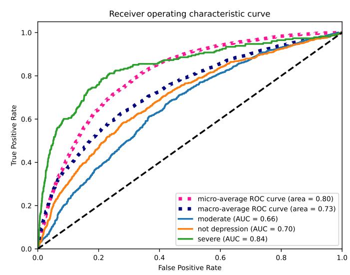

Figure 3: ROC Curve for the proposed fine-tuned RoBERTa model (Validation dataset)

Table 2: Performance of different models for the identification of depression

Model	Class	Development data				Test data			
		P	R	F_1	Acc	P	R	F_1	Acc
RoBERTa	Moderate Depression	0.76	0.68	0.72	0.63	0.58	0.66	0.62	0.52
	No Depression	0.44	0.52	0.48		0.42	0.53	0.47	
	Severe Depression	0.40	0.55	0.47		0.53	0.09	0.15	
	Macro Avg.	0.54	0.58	0.55		0.51	0.43	0.41	
	Weighted Avg.	0.65	0.63	0.64		0.53	0.52	0.49	
S-BERT + SVM	Moderate Depression	0.71	0.83	0.77	0.66	0.57	0.72	0.63	0.52
	No Depression	0.45	0.35	0.39		0.41	0.45	0.43	
	Severe Depression	0.53	0.20	0.29		0.67	0.02	0.04	
	Macro Avg.	0.57	0.46	0.48		0.55	0.40	0.37	
	Weighted Avg.	0.63	0.66	0.63		0.54	0.52	0.47	
S-BERT + RF	Moderate Depression	0.71	0.76	0.74	0.63	0.56	0.62	0.59	0.49
	No Depression	0.40	0.40	0.40		0.38	0.53	0.45	
	Severe Depression	0.47	0.14	0.21		0.40	0.02	0.04	
	Macro Avg.	0.53	0.43	0.45		0.45	0.39	0.36	
	Weighted Avg.	0.61	0.63	0.61		0.48	0.49	0.45	
S-BERT + LR	Moderate Depression	0.73	0.72	0.73	0.63	0.54	0.60	0.57	0.48
	No Depression	0.42	0.45	0.44		0.37	0.50	0.43	
	Severe Depression	0.41	0.35	0.38		0.50	0.07	0.12	
	Macro Avg.	0.52	0.51	0.51		0.47	0.39	0.37	
	Weighted Avg.	0.63	0.63	0.63		0.49	0.48	0.45	
S-BERT + KNN	Moderate Depression	0.71	0.79	0.75	0.63	0.54	0.72	0.62	0.48
	No Depression	0.41	0.31	0.35		0.32	0.27	0.29	
	Severe Depression	0.31	0.25	0.28		0.28	0.06	0.09	
	Macro Avg.	0.48	0.45	0.46		0.38	0.35	0.33	
	Weighted Avg.	0.60	0.63	0.61		0.43	0.48	0.43	
S-BERT + NB	Moderate Depression	0.74	0.47	0.57	0.47	0.56	0.41	0.47	0.44
	No Depression	0.39	0.40	0.39		0.36	0.53	0.43	
	Severe Depression	0.18	0.77	0.29		0.33	0.37	0.35	
	Macro Avg.	0.44	0.55	0.42		0.42	0.44	0.42	
	Weighted Avg.	0.61	0.47	0.51		0.47	0.44	0.44	
S-BERT + GB	Moderate Depression	0.72	0.76	0.74	0.63	0.56	0.64	0.59	0.49
	No Depression	0.43	0.40	0.41		0.38	0.49	0.43	
	Severe Depression	0.39	0.27	0.32		0.50	0.04	0.08	
	Macro Avg.	0.51	0.48	0.49		0.48	0.39	0.37	
	Weighted Avg.	0.62	0.63	0.63		0.50	0.49	0.46	
S-BERT + DT	Moderate Depression	0.70	0.54	0.61	0.50	0.57	0.51	0.54	0.43
	No Depression	0.31	0.44	0.36		0.31	0.47	0.37	
	Severe Depression	0.19	0.31	0.24		0.28	0.13	0.18	
	Macro Avg.	0.40	0.43	0.40		0.38	0.37	0.36	
	Weighted Avg.	0.56	0.50	0.52		0.44	0.43	0.43	
S-BERT + AdaBoost	Moderate Depression	0.72	0.72	0.72	0.61	0.53	0.58	0.55	0.45
	No Depression	0.42	0.42	0.42		0.34	0.44	0.38	
	Severe Depression	0.30	0.30	0.30		0.24	0.04	0.08	
	Macro Avg.	0.48	0.48	0.48		0.37	0.36	0.34	
	Weighted Avg.	0.61	0.61	0.61		0.42	0.45	0.42	

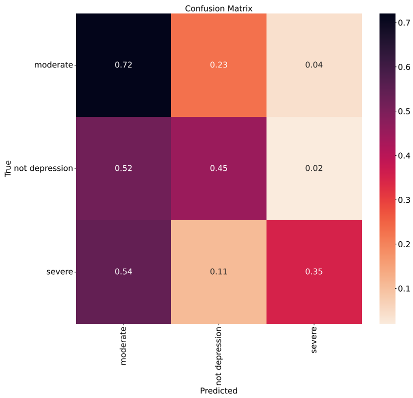

Figure 4: Confusion matrix for the proposed fine-tuned S-BERT + LR model (Validation dataset)

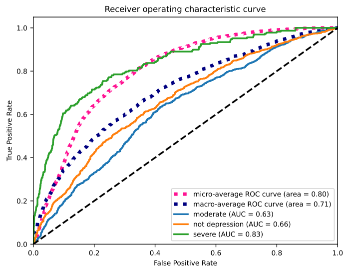

Figure 5: ROC Curve for the proposed fine-tuned S-BERT + LR model (Validation dataset)

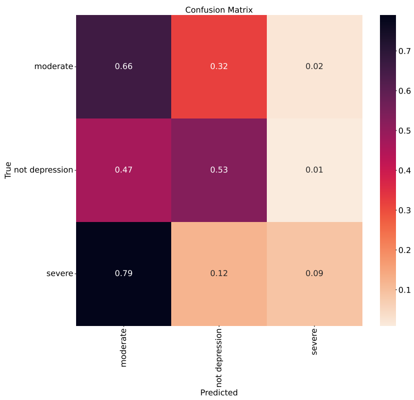

Figure 6: Confusion matrix for the proposed fine-tuned RoBERTa model (Test dataset)

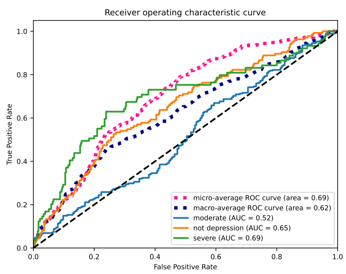

Figure 7: ROC Curve for the proposed fine-tuned RoBERTa model (Test dataset)

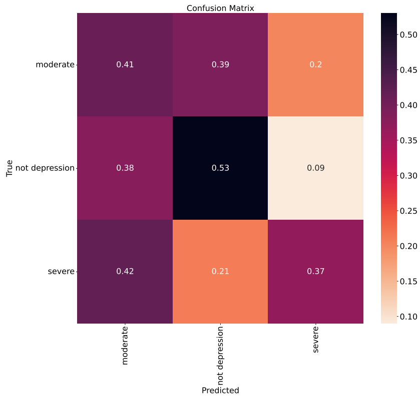

Figure 8: Confusion matrix for the proposed fine-tuned S-BERT + NB model (Test dataset)

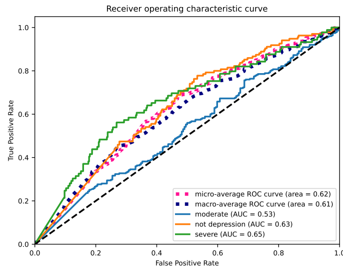

Figure 9: ROC Curve for the proposed fine-tuned S-BERT + NB model (Test dataset)

ing Sentence-BERT encoding combined with classical machine learning models, the S-BERT + NB model showcased the best performance among all implemented models for the testing dataset, achieving a macro F_1-score of 0.42. The confusion matrix and AUC-ROC curve for the S-BERT + NB model on the testing dataset can be observed in Figures 8 and 9, respectively.

5 Conclusion

Detecting depression from social media is essential for early intervention and providing timely support to those in need. Through the analysis of social media posts, we can identify indicators of depression and offer appropriate resources and assistance. Moreover, monitoring depression through social media provides valuable insights into its prevalence, distribution, and impact on different populations, aiding in public health planning, resource allocation, and targeted interventions. In this study, we explore the effectiveness of fine-tuned RoBERTa and Sentence-BERT with classical machine learning classifiers for depression identification in social media. Our findings demonstrate that RoBERTa and Sentence-BERT with Naive Bayes classifiers perform well in detecting depression. However, the overall performance of the models in this task is still limited, highlighting the need for robust systems in the future. To address this, a more comprehensive ensemble-based approach, coupled with proper pre-processing techniques to handle grammatical errors, non-standard abbreviations, and linguistic variations in social media posts, can be developed to enhance the accuracy and reliability of depression detection.

References

Hatoon S AlSagri and Mourad Ykhlef. 2020. Machine learning-based approach for depression detection in twitter using content and activity features. *IEICE Transactions on Information and Systems*, 103(8):1825–1832.

Jitimon Angskun, Suda Tipprasert, and Thara Angskun. 2022. Big data analytics on social networks for real-time depression detection. *Journal of Big Data*, 9(1):69.

Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. Quantifying mental health signals in twitter. In *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*, pages 51–60.

Munmun De Choudhury, Scott Counts, Eric J Horvitz, and Aaron Hoff. 2014. Characterizing and predicting postpartum depression from shared facebook data. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pages 626–638.

George Gkotsis, Anika Oellrich, Sumithra Velupillai, Maria Liakata, Tim JP Hubbard, Richard JB Dobson, and Rina Dutta. 2017. Characterisation of mental health conditions in social media using informed deep learning. *Scientific reports*, 7(1):45141.

Sharath Chandra Guntuku, J Russell Ramsay, Raina M Merchant, and Lyle H Ungar. 2019. Language of adhd in adults on social media. *Journal of attention disorders*, 23(12):1475–1485.

Sharath Chandra Guntuku, David B Yaden, Margaret L Kern, Lyle H Ungar, and Johannes C Eichstaedt. 2017. Detecting depression and mental illness on social media: an integrative review. *Current Opinion in Behavioral Sciences*, 18:43–49.

Lang He, Mingyue Niu, Prayag Tiwari, Pekka Marttinen, Rui Su, Jiewei Jiang, Chenguang Guo, Hongyu Wang, Songtao Ding, Zhongmin Wang, et al. 2022. Deep learning for depression recognition with audiovisual cues: A review. *Information Fusion*, 80:56–86.

Abhinav Kumar and Jyoti Kumari. 2021. A machine learning approach for fake news detection from urdu social media posts. In *Working Notes of FIRE 2021-Forum for Information Retrieval Evaluation (Online). CEUR*.

Jyoti Kumari and Abhinav Kumar. 2021a. A deep neural network-based model for the sentiment analysis of dravidian code-mixed social media posts. In *Working Notes of FIRE 2021-Forum for Information Retrieval Evaluation (Online). CEUR*.

Jyoti Kumari and Abhinav Kumar. 2021b. Offensive language identification on multilingual code mixing text. In *Working Notes of FIRE 2021-Forum for Information Retrieval Evaluation (Online). CEUR*.

Minsu Park, Chiyoung Cha, and Meeyoung Cha. 2012. Depressive moods of users portrayed in twitter. In *Proceedings of the 18th ACM International Conference on Knowledge Discovery and Data Mining, SIGKDD 2012*, pages 1–8.

Andrew G Reece and Christopher M Danforth. 2017. Instagram photos reveal predictive markers of depression. *EPJ Data Science*, 6(1):15.

Kayalvizhi S, Thenmozhi Durairaj, Bharathi Raja Chakravarthi, and Jerin Mahibha C. 2022. Findings of the shared task on detecting signs of depression from social media. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 331–338, Dublin, Ireland. Association for Computational Linguistics.

Kayalvizhi Sampath, Thenmozhi Durairaj, Bharathi Raja Chakravarthi, Jerin Mahibha C, Kogilavani Shanmugavadivel, and Pratik Anil booktitle = Rahood. Overview of the second shared task on detecting signs of depression from social media text.

Sunil Saumya, Abhinav Kumar, and Jyoti Prakash Singh. 2021. Offensive language identification in dravidian code mixed social media text. In *Proceedings of the first workshop on speech and language technologies for Dravidian languages*, pages 36–45.

H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*, 8(9):e73791.

Michael Mesfin Tadesse, Hongfei Lin, Bo Xu, and Liang Yang. 2019. Detection of suicide ideation in social media forums using deep learning. *Algorithms*, 13(1):7.

M Johnson Vioules, Bilel Moulahi, Jérôme Azé, and Sandra Bringay. 2018. Detection of suicide-related posts in twitter data streams. *IBM Journal of Research and Development*, 62(1):7–1.