Comparative Performance Analysis of Various Techniques for Psychiatric Evaluation of Social Media

Alwinder Singh[1], Jodhwinder Singh[2], Rashmi Chaudhary[3], Prashant Singh Rana[1]

Computer Science and Engineering Department, Thapar Institute of Engineering and Technology, Patiala, Punjab, India[1]

Conestoga College, Doon Valley, Kitchener, Ontario, Canada[2]

Computer Science and Engineering Department, Netaji Subhas University of Technology, New Delhi, India[3]

Corresponding author: Prashant Singh Rana, Email: prashant.singh@thapar.edu

This research study aimed to enhance the precision and efficacy of identifying the underlying reasons for mental health issues by analyzing social media content. The study was conducted on the CAMS dataset, comprising 3155 Reddit posts, and a reannotated SDCNL dataset with 1896 instances to align with the domain of interest. The study employed advanced machine learning techniques and popular deep learning architectures to determine the optimal model for the problem statement. The performance of several models was evaluated, and it was found that the combination of CNN and LSTM (ConvLSTM) architecture yielded the most accurate results, with a precision of 53%. ConvLSTM incorporates spatial and temporal information by using convolutional layers to extract features from the input data, followed by LSTM layers to capture the temporal dependencies within the data. This makes ConvLSTM particularly suitable for processing sequential data, such as natural language text and has been shown to achieve better performance than using CNN or LSTM alone. These findings indicate the potential of leveraging machine learning techniques and deep learning architectures to improve the detection and understanding of mental health issues through social media analysis. The study demonstrates the potential of machine learning techniques and deep learning architectures to identify the underlying reasons for mental health issues by analyzing social media content. The findings highlight the importance of using these techniques for mental health research and the potential for future research. The study's findings significantly impact the development of interventions and treatments for mental health issues.

Keywords: Psychiatric Evaluation, Social Media, LSTM, CNN

2023. In Deepak Bhatia & Himanshu Arora (eds.), *SCRS Proceedings of International Conference of Undergraduate Students 2023*, 7–24. Computing & Intelligent Systems, SCRS, India. https://doi.org/10.56155/978-81-955020-4-2-2

1 Introduction

The mounting mental pressure that individuals experience in today's society regarding any illness has made it incredibly challenging for many even to broach the subject of mental health. However, social media has emerged as a powerful tool enabling people to share their emotions without worrying about individuality, anonymity, or other ambiguous concerns. By expressing their feelings on social media, individuals can provide insights into their mental health and offer a window into their personal lives. This makes understanding their challenges and difficulties easier and provides targeted support and interventions as needed. In essence, social media has become a valuable tool for promoting mental health and well-being and helping individuals navigate the complex landscape of mental health issues in today's society.

Recent technological advancements have opened up new scientific avenues to understand human behavior and emotions better. In particular, research has shown that it is possible to gain insights into a person's mental state through the posts they upload on social media [1]. This has spurred many researchers to develop more accurate and reliable methods for analyzing social media data, intending to reduce the rate of mental illness in society.

By analyzing social media posts, researchers can understand how individuals feel and what challenges they may face in their personal lives. This can help in the early detection of mental health issues, allowing for timely intervention and support to be provided to those in need. In essence, social media has become a valuable tool for reducing the rate of mental illness, as it enables researchers to monitor individuals' emotional states and provide targeted support as needed.

As the field of social media analytics continues to evolve, we will likely see even more sophisticated methods for analyzing social media data and gaining insights into human behavior and emotions [2]. This can revolutionize how we approach mental health care and help us better understand and address the mental health challenges that many individuals face in today's society, leading to a happier overall environment to thrive.

1.1 Motivation

There has been an increased focus on mental health as a research topic recently, with many efforts being made to analyze the mental well-being of the masses. The tremendous growth in the field of Natural Language Processing (NLP) has made it possible to use advanced Artificial Intelligence (AI) techniques in the domain to understand textual data better and interpret the mental state of the speaker/author accordingly. Some examples of these techniques are sentiment analysis and emotion detection [3].

Although many efforts have been made, there is still a lot of scope to understand the actual cause of emotions at a deeper level. This requires knowledge of psychology along with AI. We can understand where emotions originate from if we delve deeper into the subject.

This understanding can lead to the development of more effective interventions and treatments for individuals struggling with mental health issues. With mental health being a critical concern globally, research in this field can significantly improve the lives of millions of people.

Moreover, as we become more reliant on digital communication, analyzing the mental state of individuals through text data becomes increasingly important. This can aid in identifying individuals who may be at risk for mental health issues or provide support for those already struggling.

Thus, a research paper on mental health analysis through Natural Language Processing techniques can contribute to the academic discourse and have real-world implications. It can serve as a step towards bridging the gap between psychology and AI and ultimately improving the mental well-being of individuals worldwide.

Overall, the potential impact of this research on various aspects of society makes it a highly relevant and important topic to explore.

2 Literature Review

Mental health analysis has been a topic of interest in natural language processing (NLP) and artificial intelligence (AI) for decades. Researchers have attempted to use NLP techniques to analyze language patterns and identify mental health conditions such as depression, anxiety, and schizophrenia.

One of the earliest attempts in this field was made by G. P. Holman in 1974, who used a computer program to analyze the language of patients with schizophrenia. The program analyzed the frequency and type of words used by the patients and compared them to a control group. The study found that patients with schizophrenia used fewer concrete words and more abstract words than the control group, concluding that language patterns could be used to identify mental health conditions.

In the 1990s, researchers began using machine learning techniques to analyze language patterns and identify mental health conditions. For example, researchers used a decision tree algorithm to analyze the language of patients with depression. They found that specific language patterns, such as the word ”I” and negative emotion words, could accurately predict the presence of depression.

In the 2000s, researchers began exploring deep learning techniques for mental health analysis. For example, researchers used a convolutional neural network (CNN) to analyze the language of patients with depression. They found that the model could accurately predict depression with a high degree of accuracy.

In 2014, a study used a machine learning approach to predict depression in Twitter users by analyzing their tweets. The study found that certain linguistic features, such as using first-person pronouns and negative emotion words, were associated with depression.

Researchers have recently begun using GPT to analyze mental health language. For example, a study published in 2020 used BERT to analyze social media posts from individuals with depression and found that the model could accurately predict depression with a high degree of accuracy.

Overall, using NLP, AI, ML, and DL techniques for mental health analysis has a long history of research and development. As these techniques evolve, they may become increasingly useful for identifying and treating mental health conditions.

Previous research has focused on identifying mental health conditions such as depression, anxiety, and schizophrenia using NLP and AI techniques. However, a Causal analysis of mental health issues in social media posts (CAMS) takes this a step further by attempting to identify the underlying causes of these conditions.

The approach used here combines CNN and LSTM networks which gave remarkable results.

Let's assume we have a sequence of words represented by a matrix X of size T × d, where T is the length of the sequence and d is the dimension of each word embedding. We will use a CNN to extract features from the sequence and an LSTM to model the temporal dependencies.

- **Convolutional Layer**
  The convolutional layer filters the input sequence to extract relevant features [4]. The output of the convolutional layer is a feature map of size T' x h, where T' is the sequence length after applying padding and h is the number of filters.

  For each filter k, the output of the convolutional layer is computed as follows:

$$h_k = f(W_k * X + b_k) \tag{2.1}$$

where f is a non-linear activation function such as ReLU, W_k is the size k × d filter, b_k is the bias term, and * represents the convolution operation. We can then apply max pooling over the feature map to get the most important feature for each filter:

$$v_k = max(h_k) \tag{2.2}$$

The output of the convolutional layer is a vector of size h, representing the features extracted from the sequence:

$$v = [v_1, v_2, ..., v_h] \tag{2.3}$$

- **LSTM Layer**
  The LSTM layer takes the output of the convolutional layer as input and models the temporal dependencies in the sequence [5]. The output of the LSTM layer is a vector of size m, representing the final hidden state. At each time step t, the LSTM updates its hidden state h_t and cell state ctextsubscriptt as follows:

$$i_t = sigmoid(W_i * [h_{t-1}, v] + b_i) \tag{2.4}$$

$$f_t = sigmoid(W_f * [h_{t-1}, v] + b_f) \tag{2.5}$$

$$o_t = sigmoid(W_o * [h_{t-1}, v] + b_o) \tag{2.6}$$

Table 1: The structure of CAMS dataset

Text	Cause	Inference
Today would have been my best friend's 18th birthday, we'd be going out together for the first time, and we'd be sitting here making resolutions for the new year that we both know we'd never keep. None of that is even possible, though, because he's gone. It was never supposed to be like this.	4	Best friend's birthday, he's gone, never supposed to be like this
I live in a country I don't want to live in, and there is nothing I can do. I know what I want from life but can't achieve it without moving to a different country. Instead of dealing with reality and building my life in my country, I just lie on my bed doing nothing for years/months with no job, education, or even goals. When will I snap out of it? Is it immature?	2	No job, no education or even goals
It's been enough years of this already. I remember all of the pain of this last year. Imagining another year of this makes me sick. Maybe I'll drink myself to death before I wind up homeless.	4	Homeless

$$g_t = tanh(W_g * [h_{t-1}, v] + b_g) \tag{2.7}$$

$$c_t = f_t * c_{t-1} + i_t * g_t \tag{2.8}$$

$$h_t = o_t * tanh(c_t) \tag{2.9}$$

where $[h_{t-1}, v]$ is the concatenation of the previous hidden state and the output of the convolutional layer, and W_i, W_f, W_o, W_g, b_i, b_f, b_o, b_g are the learnable parameters of the LSTM.

The final output of the model can be obtained by applying a fully connected layer to the hidden state:

$$y = softmax(W_y * h_T + b_y) \tag{2.10}$$

Where W_y and b_y are the weights and bias of the fully connected layer, and softmax is the activation function used to obtain the class probabilities.

Combining CNN and LSTM models for text classification involves applying a convolutional layer to extract features from the sequence, followed by an LSTM layer to model the temporal dependencies and produce the final hidden state, which is then used to make predictions.

3 Methodology

The aim of this research is to conduct a comparative analysis of various techniques for psychiatric evaluation of social media. In recent years, social media platforms have become increasingly popular, and their use has been linked to various mental health issues such as depression, anxiety, and suicidal ideation. Therefore, identifying effective techniques to evaluate the mental health status of social media users is crucial in providing early interventions and preventing potential harm.

The problem has been simplified into a multi-class classification problem, details regarding which can be found in the following sections.

3.1 Dataset

To conduct this research, we used the CAMS dataset, which includes two parts. The first part consists of 3155 Reddit posts collected from various subreddits related to mental health [6]. These were then annotated using the expert knowledge of psychiatrists. The annotations were then used to evaluate the effectiveness of different techniques for psychiatric evaluation, such as natural language processing (NLP) and machine learning algorithms.

The second part consists of a publicly available SDCNL dataset that consists of 1896 instances for mental health causal analysis.

A combination of both of these two parts sums up the CAMS dataset. Overall, the CAMS dataset provides a rich source of social media data annotated for mental health symptoms, enabling us to conduct a comprehensive comparative analysis of various techniques for psychiatric evaluation [7].

The CAMS dataset's structure can be seen in Table 1.

The column 'Text' contains an actual post by the user, which they may write on their social media handle. The column 'Cause' denotes the category to which the text has been assigned. Lastly, the 'Inference' column summarizes the key takeaways from the text column.

The CAMS data consists of 5052 rows, each classified into different categories. Figure 1 shows data distribution among various classes.

For experimentation purposes, only a subset of data has been used, which includes the complete IntentSDCNL (SDCNL data with new annotations) data as well as some added instances (80+80+80) for classes 'Bias or Abuse', 'Jobs and Career', 'Medication' for the purpose of balancing the class while still retaining some properties of class distribution.

Furthermore, the class has been labeled with numbers, and the mapping can be observed in Table 2.

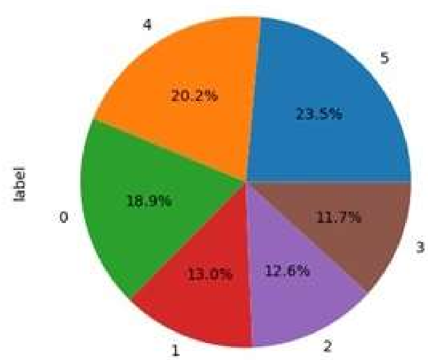

Figure 1: Distribution of Dataset among different classes

3.2 Dataset Preparation and Understanding

The dataset has been passed through various cleaning and preparation processes to convert it to a more helpful form for feeding to the model and increasing its value. The subsequent steps were employed to pre-process the dataset and transform it into a more practical format. The complete data preparation pipeline can be viewed in Figure 2.

Table 2: Class Labels

Cause	Label
No reason	0
Bias or Abuse	1
Jobs and Career	2
Medication	3
Relationship	4
Alienation	5

3.2.1 Removal of Special Symbols

The first step in cleaning the data was removing various special characters that might occur in the social media posts due to various reasons like encoding issues during dataset collection or deliberate use by the writer etc. This also includes unnecessary information like emails, punctuation marks, and Uniform Resource Locator (URL) data.

Also, numbers have been found to be less important in this classification task, and hence numbers have also been removed.

3.2.2 Removal of Empty/Null Values

This step involves the removal of all those entries that have some missing data, as they can interfere with the working of the classifier and hence hinder its performance. This includes posts with an ”empty post” tag while collecting the data and rows with missing columns.

3.2.3 Removal of stopwords

This step includes removing some words with an abnormally high frequency in the dataset but contributing little to the value. These words hinder the performance of frequency-based classification methods as they do not carry meaning independently and hence must be removed.

Also, words with a length of less than two are being removed as they are not expected to have much importance.

3.2.4 Conversion to lowercase

The classifiers are designed to work with the frequency of occurrence of various words in the dataset. Hence, the frequency of words is a very important factor in the classification process. Case mismatch is one of the major factors that can affect the word frequencies and hence the model's overall performance. So, to prevent this from happening, all the words have been transformed into lowercase.

3.2.5 Stemming

Stemming is a technique used in natural language processing to normalize or standardize text for analysis. It involves reducing inflected or derived words to their base or root form, which helps to improve the efficiency and accuracy of text analysis tasks like classification, retrieval, and sentiment analysis. Different algorithms like the Porter or Snowball stemming algorithm can identify the root form of words, but here, Porter Stemmer has been used.

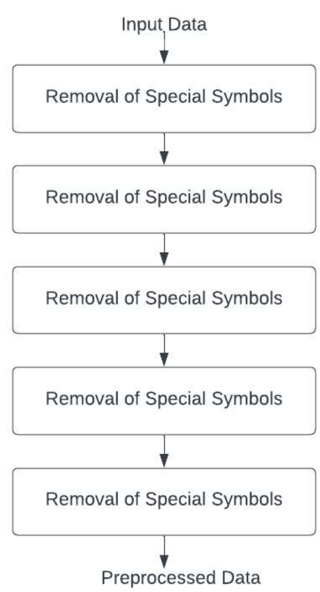

Figure 2: Data Preparation Pipeline

4 Proposed Models for Comparison

Various Machine learning and Deep learning models have been analyzed. Forthcoming models gave comparatively better performance and have been chosen for further analysis.

4.1 Deep Learning models

4.1.1 LSTM Networks

Long Short Term Memory Network includes the input to be fed to an Embedding layer followed by an LSTM layer which includes feedback. Further, half of the nodes are dropped out to make the model more robust and prevent overfitting, which is then connected to a Dense Layer, which provides probabilities of the instance being a part of the available classes. The optimizer is Adam, and the loss function used is Categorical Cross Entropy as it performs well with multiclass classification.

Table 3: Stacking Classifiers along with the Estimators and Final Estimators

Model	Estimators	Final Estimator
Stacking Classifier 1	NuSVC, GaussianNB, XGBClassifier, AdaBoostClassifier, RandomForestClassifier, ExtraTreesClassifier	Logistic Regression
Stacking Classifier 2	RandomForestClassifier, LinearSVC in a pipeline, AdaBoostClassifier, and ExtraTreesClassifier	Logistic Regression
Stacking Classifier 3	NuSVC, GaussianNB, XGBClassifier, AdaBoostClassifier, and ExtraTreesClassifier	Logistic Regression
Stacking Classifier 4	RandomForestClassifier, LinearSVC in a pipeline, AdaBoostClassifier, and ExtraTreesClassifier	Support Vector Machine

Table 4: Voting Classifiers along with the Estimators and Voting type

Model	Estimators	Voting
Voting Classifier 1	NuSVC, GaussianNB, XGBClassifier, AdaBoostClassifier, RandomForestClassifier, ExtraTreesClassifier	Hard
Voting Classifier 2	RandomForestClassifier, AdaBoostClassifier, and ExtraTreesClassifier	Hard
Voting Classifier 3	NuSVC, GaussianNB, XGBClassifier, AdaBoostClassifier, and ExtraTreesClassifier	Hard
Voting Classifier 4	NuSVC, GaussianNB, XGBClassifier, AdaBoostClassifier, RandomForestClassifier, ExtraTreesClassifier	Soft
Voting Classifier 5	RandomForestClassifier, AdaBoostClassifier, and ExtraTreesClassifier	Soft
Voting Classifier 6	NuSVC, GaussianNB, XGBClassifier, AdaBoostClassifier, and ExtraTreesClassifier	Soft

4.1.2 Convolutional Neural Network

This network first includes an Embedding layer that maps input information from a high-dimensional to a lower-dimensional space. This is followed by a 1-Dimensional

Convolution layer with 128 filters and ReLU activation, along with a MaxPooling Layer to aggregate the results of the Convolution Layer. This is repeated again and the outputs are then fed to a Dense layer which provides probabilities of the instance being a part of the available classes. The optimizer is Adam, and the loss function used is Categorical Cross Entropy as it performs well with multiclass classification.

4.1.3 Gated Recurrent Unit Network

This network also starts with an Embedding layer to decrease the input size. This is followed by a GRU layer which includes feedback [8]. Further, half of the nodes are dropped out to make the model more robust and prevent overfitting which is then connected to a Dense Layer, which provides probabilities of the instance being a part of the available classes. The optimizer is Adam, and the loss function used is Categorical Cross Entropy as it performs well with multiclass classification.

4.1.4 Combination Network

Two distinct types of combinational networks have been employed in this research, namely:

- The Convolutional Neural Network (CNN) combined with the Gated Recurrent Unit (GRU) [9], and
- The CNN combined with the Long Short-Term Memory (LSTM) network [10].

These networks are a combination of the CNN and GRU/LSTM models. They first map the inputs to a lower dimension space using Embedding Layer. It is then followed by a 1-Dimensional Convolution layer with 128 filters and ReLU activation, along with a MaxPooling Layer to aggregate the results of the Convolution Layer. Then the output is fed to two GRU/LSTM (comprising 128 units) layers, which communicate using feed-forward and feed-back communication. Finally, the results are mapped to a Dense layer with six units.

4.1.5 Bidirectional Networks

Two distinct types of Bidirectional networks have been employed in this research, namely:

- Bidirectional Long Short-Term Memory (BiLSTM) Network [11]
- Bidirectional Gated Recurrent Unit (BiGRU) Network [12]

These are bidirectional networks where two hidden layers of opposite directions are connected to the same output. These networks also start with an Embedding layer to decrease the input size. Then two LSTM/GRU layers are added in opposite directions where one LSTM/GRU layer moves forward, beginning from the start of the data sequence, and the other moves backward, beginning from the end of the data sequence. Further, half of the nodes are dropped out to make the model more robust and prevent overfitting, which is then connected to a Dense Layer, which provides probabilities of the instance being a part of the available classes. The optimizer being used is Adam, and the loss function used is Categorical Cross Entropy.

4.2 Ensemble Models

Ensemble models combine multiple models to improve the overall performance and robustness of the system [13]. Two distinct types of Ensemble models have been employed in this research, namely:

- Stacking Classifier
- Voting Classifier
  - Hard
  - Soft

4.2.1 Stacking Classifier

Stacking Classifiers [14] is an ensemble model that involves training multiple base models and combining their predictions using a meta-model. The meta-model is trained on the outputs of the base models, allowing it to learn how to combine best their predictions (Figure 3). In this research, the Stacking Classifiers have been used, which are composed of many base models:

- RandomForestClassifier,
- LinearSVC (scaled using StandardScaler in a pipeline),
- AdaBoostClassifier, and
- ExtraTreesClassifier, etc

The Stacking classifiers and the final estimators have been shown in Table 3.

4.2.2 Voting Classifier

Voting Classifiers [15] is another ensemble model type that combines the predictions of multiple base models using a voting scheme. There are two types of voting in the Voting Classifier:

- Hard Voting (Figure 4) and
- Soft Voting (Figure 5)

Hard voting involves taking the majority vote of the base models, while soft voting involves taking the weighted average of the predicted probabilities of the base models. In this research, hard and soft voting has been used in the Voting Classifier.

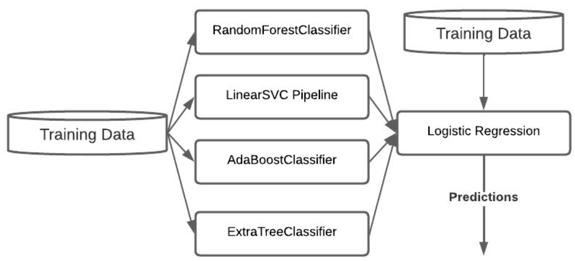

Figure 3: Working of Stacking Classifier

The various Voting Classifiers being used are shown in Table 4.

The models used while creating these ensemble models are chosen very carefully, considering each model's individual performances.

Table 5: Experimental results. F1, Precision, and Accuracy are computed for all six causal classes: 'No reason' (C0), 'Bias or abuse' (C1), 'Jobs and careers' (C2), 'Medication' (C3), 'Relationship' (C4), 'Alienation' (C5)

Model	F1-Score						Precision						Accuracy
	C0	C1	C2	C3	C4	C5	C0	C1	C2	C3	C4	C5	
Stacking Classifier	0.55	0.1	0.4	0.12	0.49	0.45	0.74	0.06	0.33	0.07	0.51	0.53	0.44
Voting Classifier (Hard)	0.51	0.1	0.52	0.34	0.3	0.49	0.71	0.08	0.4	0.09	0.6	0.51	0.4
Voting Classifier (Soft)	0.56	0.11	0.53	0.34	0.33	0.5	0.73	0.12	0.3	0.11	0.6	0.53	0.45
LSTM	0.57	0.22	0.59	0.49	0.52	0.43	0.49	0.18	0.63	0.45	0.6	0.52	0.48
CNN	0.58	0.15	0.63	0.49	0.53	0.52	0.51	0.14	0.6	0.44	0.62	0.62	0.51
GRU	0.58	0.21	0.61	0.44	0.57	0.45	0.57	0.19	0.66	0.42	0.56	0.49	0.49
CGRU	0.6	0.21	0.47	0.43	0.51	0.43	0.49	0.19	0.53	0.64	0.46	0.55	0.47
BiLSTM	0.53	0.08	0.51	0.45	0.61	0.45	0.4	0.13	0.55	0.57	0.55	0.5	0.5
BiGRU	0.56	0.23	0.51	0.51	0.45	0.43	0.53	0.24	0.45	0.59	0.49	0.43	0.46
ConvLSTM	0.6	0.19	0.52	0.47	0.57	0.55	0.51	0.3	0.57	0.53	0.57	0.51	**0.53**

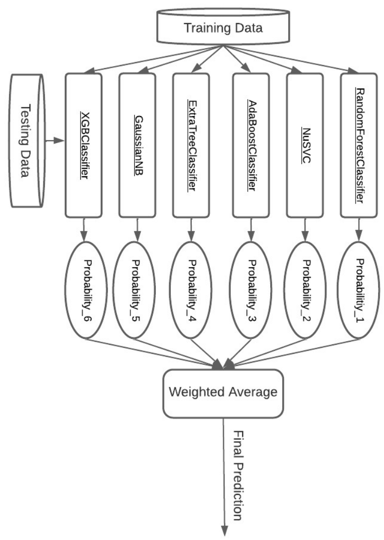

Figure 4: Working of Voting Classifier (Hard)

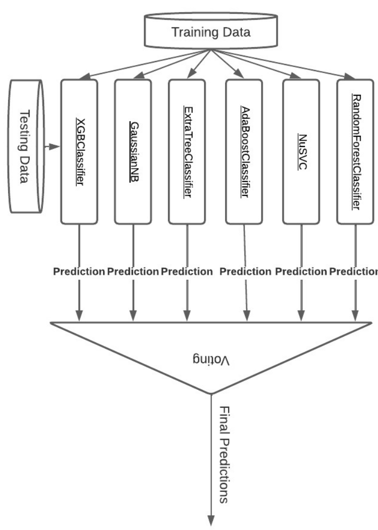

Figure 5: Working of Voting Classifier(Soft)

5 Experiment Results

The objective of this study was to compare the performance of various deep learning models and Machine Learning Models for the task of Comparative analysis of various

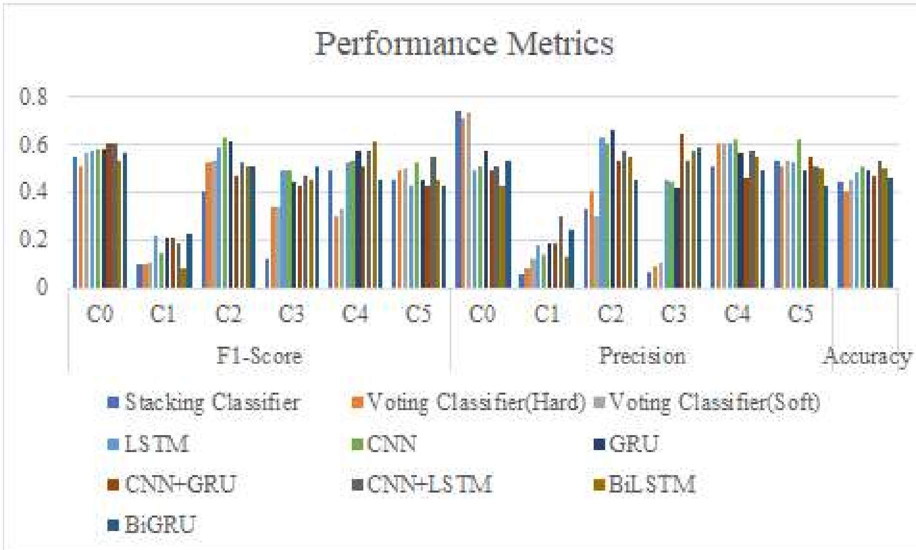

Figure 6: Performance Metrics among difference Classes

techniques for the Psychiatric evaluation of Social Media. The models evaluated in this study are the Stacking Classifier, Voting Classifier, LSTM Model, CNN Model, GRU Model, Combination of CNN and GRU (CGRU) Model, ConvLSTM Model, BiLSTM, and BiGRU models.

The Stacking Classifier is an ensemble learning technique that combines the predictions of multiple classifiers to improve the overall performance. The classifier uses six estimators: Nu-Support Vector Classification (NuSVC), GaussianNB, XGBClassifier, AdaBoostClassifier, RandomForestClassifier, and ExtraTreeClassifier. The final estimator is a logistic regression model with a maximum iteration of 1500 and the sag solver. The Stacking Classifier achieved an accuracy of 0.44 with a weighted F1 score of 0.37. Whereas the Voting classifier with a voting type hard gave an accuracy of 0.40 and in voting type soft gave an accuracy of 0.45.

The LSTM Model is a recurrent neural network (RNN) capable of processing sequential data. The model uses an embedding matrix and an embedding dimension of 300. The LSTM Model achieved an accuracy of 0.48 with a weighted F1-score of 0.47.

The CNN Model is a neural network typically used for image classification tasks but can also be applied to sequential data. The model uses the same embedding matrix and embedding dimension as the LSTM Model. The CNN Model achieved an accuracy of 0.51 with a weighted F1-score of 0.48.

The GRU Model is another RNN type similar to LSTM but has fewer parameters. The model uses the same embedding matrix and embedding dimension as the LSTM and CNN Models. The GRU Model achieved an accuracy of 0.49 with a weighted F1-score of 0.48.

The CGRU Model is a hybrid model that combines the convolutional and recurrent layers of the CNN and GRU Models, respectively. The model achieved an accuracy of 0.470 with a weighted F1-score of 0.44. The ConvLSTM Model is another hybrid model that combines the convolutional and recurrent layers of the CNN and LSTM Models, respectively. The model achieved an accuracy of 0.53 with a weighted F1-score of 0.45. The various performance metrics have been shown in Figure 6.

The model has been passed through K-fold cross-validations and gave similar results, where K=5.

Overall, the combinational model comprising of CNN and LSTM Model performed the best in accuracy and weighted F1-score, followed closely by the LSTM Model. The Stacking Classifier performed the worst among all the models evaluated in this study.

These results suggest that deep learning models, specifically CNN and LSTM Models, are effective for multi-class sentiment analysis.

Table 5 provides the F1- Score, Precision, and Accuracy for different models among classes.

6 Conclusion and Future Work

In future work, it would be valuable to investigate additional approaches to improve the performance of the models on this dataset. One possible approach could be experimenting with different hyperparameters, such as changing the learning rate or number of epochs for the neural network models. Another possible improvement avenue could be exploring different pre-processing techniques, such as different tokenization strategies or word embeddings.

Another potential area for future work could be to investigate ensemble methods, such as model averaging or stacking, to combine the predictions of multiple models and potentially improve performance. It would also be interesting to explore the use of other types of neural network architectures, such as transformers, for this task.

Finally, it would be valuable to evaluate the generalizability of these models to other datasets and domains. This could involve training the models on different datasets, evaluating their performance on the original dataset used in this study, or testing the models on datasets from different domains entirely.

References

1. Cao, L., Zhang, H., Feng, L., Wei, Z., Wang, X., Li, N., & He, X. (2019). Latent suicide risk detection on microblog via suicide-oriented word embeddings and layered attention. arXiv preprint arXiv:1910.12038. Zachary Dair, Muhammad Muneeb Saad, Urja Pawar, Samantha Dockray, & Ruairi O'Reilly. (2022). Classification of Stress via Ambulatory ECG and GSR Data.

2. Mak, W. W., Poon, C. Y., Pun, L. Y., & Cheung, S. F. (2007). Meta-analysis of stigma and mental health. Social science & medicine, 65(2), 245-261.

3. Nandwani, P., & Verma, R. (2021). A review on sentiment analysis and emotion detection from text. Social Network Analysis and Mining, 11(1), 81.

4. Madasu, A., & Rao, V. A. (2023, February). Effectiveness of self normalizing neural networks for text classification. In Computational Linguistics and Intelligent Text Processing: 20th International Conference, CICLing 2019, La Rochelle, France, April 7–13, 2019, Revised Selected Papers, Part II (pp. 412-423). Cham: Springer Nature Switzerland.

5. Zhou, C., Sun, C., Liu, Z., & Lau, F. (2015). A C-LSTM neural network for text classification. arXiv preprint arXiv:1511.08630.

6. Garg, M., Saxena, C., Krishnan, V., Joshi, R., Saha, S., Mago, V., & Dorr, B. J. (2022). CAMS: an annotated corpus for causal analysis of mental health issues in social media posts. arXiv preprint arXiv:2207.04674.

7. Saxena, C., Garg, M., & Saxena, G. (2022). Explainable causal analysis of mental health on social media data. arXiv preprint arXiv:2210.08430.

8. Zulqarnain, M., Ghazali, R., Ghouse, M. G., & Mushtaq, M. F. (2019). Efficient processing of GRU based on word embedding for text classification. JOIV: International Journal on Informatics Visualization, 3(4), 377-383.

9. Luo, L. X. (2019). Network text sentiment analysis method combining LDA text representation and GRU-CNN. Personal and Ubiquitous Computing, 23(3-4), 405-412.

10. Xu, J., Chen, D., Qiu, X., & Huang, X. (2016). Cached long short-term memory neural networks for document-level sentiment classification. arXiv preprint arXiv:1610.0498

11. Liu, G., & Guo, J. (2019). Bidirectional LSTM with attention mechanism and convolutional layer for text classification. Neurocomputing, 337, 325-338.

12. Ali, W., Yang, Y., Qiu, X., Ke, Y., & Wang, Y. (2021). Aspect-level sentiment analysis based on bidirectional-GRU in SIoT. IEEE Access, 9, 69938-69950.

13. Phan, H. T., Tran, V. C., Nguyen, N. T., & Hwang, D. (2020). Improving the performance of sentiment analysis of tweets containing fuzzy sentiment using the feature ensemble model. IEEE Access, 8, 14630-14641.

14. Subba, B., & Kumari, S. (2022). A heterogeneous stacking ensemble based sentiment analysis framework using multiple word embeddings. Computational Intelligence, 38(2), 530-559.

15. Abbas, A. K., Salih, A. K., Hussein, H. A., Hussein, Q. M., & Abdulwahhab, S. A. (2020). Twitter sentiment analysis using an ensemble majority vote classifier. Journal of Southwest Jiaotong University, 55(1).