© March 2025 | IJIRT | Volume 11 Issue 10 | ISSN: 2349-6002

Leveraging Machine Learning To Predict Depression

Mr. K. Hari Veerraju[1], Mr. Sure Sai Venkata Jaideep[2], Mr. Kandipalli Sandeep kiran[2], Mr. Meruga Srivineesh[2], Mr. Surya Billa[2]

[1]*Assistant Professor, Dept. of CSE, Raghu Engineering College, Visakhapatnam*

[2] *Department of Data Science, Raghu Engineering College, Visakhapatnam*

Abstract— **Depression, clinically recognized as Major Depressive Disorder (MDD), is a debilitating mental health condition marked by pervasive sadness, anhedonia, and impaired daily functioning. Timely detection is paramount to facilitating early intervention, mitigating relapse, and alleviating the profound emotional and financial burdens associated with the disorder. This study deploys cutting-edge machine learning methodologies to advance the early detection of depression through a robust web-based screening platform. By harnessing state-of-the-art algorithms including Support Vector Machines, Random Forests, and Neural Networks the system meticulously analyzes user responses from a validated depression assessment. Empirical evaluations demonstrate exceptional diagnostic accuracy, underscoring the system's capacity to revolutionize mental health screening by delivering rapid, scalable, and reliable evaluations. This research exemplifies the transformative potential of ML-driven diagnostic tools in enhancing clinical decision-making and fostering proactive mental health management.**

Keywords— **Depression, Machine Learning, Early Detection, Web-Based Screening, Major Depressive Disorder (MDD), Diagnostic Algorithms**

1. INTRODUCTION

1.1 Understanding Depression

Depression is one of the most prevalent and serious mood disorders, affecting millions of individuals worldwide. It is characterized by persistent sadness, loss of interest in daily activities, and a range of emotional and physical symptoms that can interfere with a person’s ability to function. Also known as Major Depressive Disorder (MDD) or clinical depression, this condition goes beyond occasional feelings of sadness or temporary emotional distress. It significantly affects how individuals think, feel, and behave, leading to long-term impairments in daily life.

Depression is not simply a state of mind; it is a complex medical condition that can arise due to a combination of genetic, biological, environmental, and psychological factors. Research has shown that depression can alter brain function, particularly in areas responsible for mood regulation, stress response, and cognitive processing. If left untreated, depression can lead to severe consequences, including suicidal thoughts and behaviors, making early detection and intervention crucial for effective management and recovery.

1.2 Importance of Early Detection and Intervention

Early detection of depression plays a critical role in mitigating its adverse effects. Many individuals suffering from depression may not seek professional help due to stigma, lack of awareness, or financial constraints. Consequently, their condition may worsen over time, leading to severe emotional distress, social withdrawal, and an increased risk of comorbid conditions such as anxiety disorders, substance abuse, and chronic health issues.

Intervening at an early stage can help individuals receive timely treatment, including therapy, medication, and lifestyle modifications. It can also prevent relapses and reduce the overall burden of the disorder on individuals, families, and healthcare systems. Machine learning and artificial intelligence (AI) have recently emerged as powerful tools for detecting mental health disorders like depression, providing an opportunity to identify symptoms earlier and more accurately than traditional diagnostic methods.

1.3 Role of Machine Learning in Depression Detection

Machine learning has revolutionized many fields, including healthcare, by enabling predictive modeling, pattern recognition, and decision-making based on large datasets. In the context of depression detection, machine learning algorithms can analyze vast amounts of data from various sources, such as social media interactions, behavioral patterns, and self-reported responses, to identify potential signs of depression.

By leveraging advanced data analytics, these models can detect subtle patterns indicative of depressive symptoms that might be overlooked in conventional assessments. This approach not only enhances the accuracy of diagnosis but also allows for continuous monitoring and personalized intervention strategies. Various machine learning techniques, such as natural language processing (NLP), sentiment analysis, and deep learning, can be applied to assess emotional states and behavioral changes indicative of depression.

1.4 Overview of the Project

This project investigates various machine learning-based approaches to the early detection of Major Depressive Disorder (MDD). The primary objective is to develop a web-based platform where users can take a depression test by answering a series of structured questions. The responses will be analyzed using trained machine learning models to classify users into categories such as No Depression or Depression.

Key Features of the Project:

Machine Learning Analysis: Implementation of machine learning algorithms to detect and classify depression.

Dataset Analysis: Examination of behavioral data to identify depression indicators.

Web-Based Platform: Development of a user-friendly website for self-assessment.

Automated Classification: Evaluation of test results using predictive models to determine depression levels.

This project aims to bridge the gap between mental health screening and accessibility by providing an automated, data-driven approach to depression detection.

1.5 Data Sources and Analysis

The success of machine learning models in detecting depression largely depends on the quality and diversity of the datasets used for training. This project utilizes datasets from various sources, including social media platforms like Twitter, user-generated survey responses, and publicly available mental health datasets.

Social media provides a wealth of data for sentiment analysis, as individuals often express their emotions, thoughts, and behavioral patterns online. By analyzing textual data, researchers can identify linguistic markers associated with depression, such as negative sentiment, self-referential language, and expressions of hopelessness or distress. Machine learning models can be trained to recognize these markers and make predictions about a user's mental health status.

In addition to social media data, structured questionnaire responses will be used to train and validate the model. These responses will include answers to clinically validated depression screening tests, such as the Patient Health Questionnaire-9 (PHQ-9) and Beck Depression Inventory (BDI), which are widely used in mental health assessments.

1.6 Machine Learning Algorithms for Depression Detection

Several machine learning algorithms can be employed for depression detection, each offering unique advantages in terms of accuracy, interpretability, and computational efficiency. Some of the commonly used models include: Supervised Learning Algorithms

Support Vector Machines (SVM): Effective in high-dimensional spaces and widely used for text classification tasks.

Random Forests: A robust ensemble learning method that enhances predictive accuracy by combining multiple decision trees.

Logistic Regression: A simple yet effective model for binary classification problems, such as detecting the presence or absence of depression.

1.7 Web-Based Depression Detection System

To make depression screening accessible to a wider audience, this project includes the development of a web-based platform where users can take a depression test. The platform will:

Provide a set of structured depression-related questions.

Collect user responses for analysis.

Use Trained machine learning models to classify depression levels.

Offer immediate feedback and suggest further professional evaluation if necessary.

The website will be designed to ensure a user-friendly experience with clear instructions and intuitive navigation. Data privacy and confidentiality will be a top priority, with measures in place to secure user information.

1.8 Potential Impact and Future Scope

The successful implementation of this project can have a profound impact on mental health screening and intervention. By providing an automated, data-driven approach to depression detection, the system can:

Help individuals recognize early signs of depression and seek professional help sooner.

Assist mental health professionals in preliminary assessments and monitoring.

Contribute to ongoing research in mental health analytics and predictive modeling.

1.9 Conclusion

Depression is a serious condition that requires timely detection and intervention. This project leverages machine learning to develop an innovative and accessible solution for early depression screening. By analyzing user responses and behavioral data, the system aims to provide accurate, reliable, and confidential mental health assessments. With advancements in AI and data science, such technologies have the potential to revolutionize mental health care, making early diagnosis and intervention more accessible to those in need.

2. LITERATURE REVIEW

This section provides a comprehensive review of contemporary machine learning approaches for depression detection. Recent advancements in machine learning (ML) have paved the way for transformative approaches in depression detection. A range of studies has applied innovative algorithms to diverse datasets, significantly enhancing diagnostic accuracy and enabling early intervention.

Odusami et al. [1] implemented a randomized concatenated model using various ML algorithms on pediatric data, with the Random Forest classifier achieving an impressive 95% accuracy in detecting child depression. In a complementary study, researchers in [2] conducted a systematic review of depression prediction models utilizing electronic health records (EHRs), reporting AUC values ranging from 0.8 to 0.9—demonstrating the robust potential of these techniques in clinical settings.

Building on these foundations, the study in [3] introduced an ensemble hybrid model that integrates Support Vector Machines (SVM) with Multilayer Perceptrons (MLP), which significantly enhanced predictive accuracy compared to traditional single-model approaches. Meanwhile, [4] leveraged ML models trained on sleep quality data to predict depressive symptom severity, providing critical insights into the relationship between sleep patterns and depression.

A comprehensive resource for researchers is offered by the study in [5], which compiles publicly available datasets tailored for depression detection. This repository has facilitated numerous subsequent studies, such as [6], where a thorough review of ML techniques identified key features and methodologies essential for accurate depression prediction.

Furthermore, [7] demonstrated the power of data science in mental health by comparing multiple ML models and optimizing feature selection strategies, thereby developing a predictive framework with superior accuracy. Collectively, these studies illustrate the transformative potential of integrating advanced ML techniques in depression detection, ultimately facilitating early diagnosis, personalized intervention, and improved mental health outcomes.

3. MATERIALS AND METHODS

3.1 Dataset Collection

For this project, obtaining a high-quality dataset is essential for training and evaluating the machine learning model for depression detection. Kaggle, a well-known platform for data science and machine learning, provides various datasets related to mental health, depression. Some datasets include labeled text samples categorized into different levels of depression severity, while others focus on sentiment analysis, psychological assessments, or self-reported mental health statuses. Additionally, datasets with questionnaire-based responses, such as those following standardized depression screening tools like the Patient Health Questionnaire (PHQ-9), can be particularly useful for training a model to classify individuals as "No Depression" or "Depressed" based on their answers. When selecting a dataset from Kaggle, it is important to ensure that it contains a sufficient number of labeled examples, a diverse range of text samples, and relevant features such as timestamps, user interactions, and linguistic markers that indicate emotional distress. Here we have taken dataset containing 503 records where it is useful for training. Preprocessing these datasets involves cleaning the text data, removing noise, handling missing values, and performing sentiment or linguistic analysis to extract meaningful patterns. The dataset should ideally be balanced to prevent bias in classification, ensuring that both depressed and non-depressed categories are well-represented. By leveraging social media-based depression datasets from Kaggle, this project aims to develop an accurate and reliable model for early detection, helping individuals receive timely mental health support and intervention.

3.21

Random Forest algorithm can be an effective machine learning approach for classification. Random Forest is an ensemble learning technique that operates by constructing multiple decision trees during training and merging their outputs to improve accuracy and reduce overfitting. It is particularly useful in text-based classification problems, where features extracted from—such as word frequency, sentiment scores, and linguistic markers—can be used to detect patterns indicative of depression. By training a Random Forest model on a labeled dataset from platforms like Kaggle, which contains posts from users with and without depression, the algorithm can learn the subtle differences in language, sentiment, and expression style that distinguish a depressed user from a non-depressed one. One of the advantages of using Random Forest in this project is its robustness against noisy data, which is common in social media content due to informal language, slang, and abbreviations. Additionally, the algorithm handles missing values effectively and provides feature importance rankings, helping to identify which textual or behavioral factors are most significant in predicting depression. After preprocessing the dataset by tokenizing text, removing stopwords, and applying techniques like TF-IDF or word embeddings for feature extraction, the Random Forest model can be trained to classify user responses from a depression test into categories such as "No Depression" or "Depressed." By leveraging this powerful algorithm, the project aims to enhance the reliability of early depression detection, allowing for timely intervention and support for individuals at risk.

3.22

The Support Vector Machine (SVM) algorithm is another powerful classification technique that can be effectively applied. SVM is a supervised learning algorithm that works by finding the optimal hyperplane that best separates different classes in a high-dimensional feature space. For depression detection, SVM can be used to classify social media posts or questionnaire responses into categories like "No Depression" or "Depressed" by analyzing linguistic patterns, sentiment polarity, and behavioral indicators. The methodology for using SVM in this project involves several key steps. First, data collection is done using publicly available datasets from platforms like Kaggle, containing social media text or depression test responses labeled as depressed or non-depressed. Next, data preprocessing is applied, including text cleaning, tokenization, removal of stopwords, and stemming or lemmatization to standardize the text. Feature extraction plays a crucial role, where methods like Term Frequency-Inverse Document Frequency (TF-IDF) or word embeddings such as Word2Vec or BERT are used to convert text data into numerical vectors. These feature vectors serve as inputs for the SVM model.

During training, SVM constructs a decision boundary using a kernel function that transforms the input space into a higher dimension where the data becomes linearly separable. The choice of the kernel—such as linear, polynomial, or Radial Basis Function (RBF)—is crucial for achieving high accuracy. The trained model then classifies new social media posts or responses from users taking the depression test. By tuning hyperparameters like the regularization parameter (C) and the kernel type, the model's performance can be optimized to achieve better classification accuracy.

SVM is particularly advantageous for this project due to its effectiveness in text classification tasks with small to medium-sized datasets and its ability to handle imbalanced data by adjusting class weights. By leveraging SVM, this project aims to build a reliable depression detection system that can analyze user responses and provide accurate classification, enabling early intervention and support for individuals experiencing depressive symptoms.

3.23

The K-Nearest Neighbors (KNN) algorithm can be employed as an effective classification technique. KNN is a non-parametric, instance-based learning algorithm that classifies new data points based on the majority class of their nearest neighbors in the feature space. It is particularly useful for datasets with a moderate number of records, such as our dataset containing 503 instances, as it does not require a complex training phase and works well with structured data.

The methodology for implementing KNN in this project involves several steps. First, the dataset is preprocessed to handle missing values and standardize the feature values for better distance computation. Given that our dataset consists of 503 records, each likely containing responses to depression-related questions, we perform feature engineering to extract relevant numerical values that can represent the responses effectively. These features may include scores from standardized depression tests like PHQ-9, response sentiment analysis, or other behavioral indicators. Since KNN is a distance-based algorithm, feature scaling using techniques like Min-Max Scaling or Standardization is applied to ensure that all variables contribute equally to the classification process.

Once the data is preprocessed, we determine an optimal value for K

K, the number of neighbors, using methods like cross-validation or the elbow method to balance bias and variance. The Euclidean distance metric is commonly used in KNN to measure similarity between data points, but other distance metrics like Manhattan or Minkowski distance can also be tested depending on the dataset's distribution. During the classification process, for a new user response, KNN identifies the K,K nearest records in the dataset and assigns the majority class—either "No Depression" or "Depressed"—based on the labels of these nearest neighbors.

One of the key advantages of using KNN in this project is its simplicity and effectiveness for smaller datasets like ours. It does not make assumptions about the data distribution and can provide robust results if an appropriate K

K value is selected. However, KNN can become computationally expensive as the dataset grows, so optimizations like KD-Trees or Ball Trees can be used to improve efficiency. By implementing KNN on our Kaggle dataset, this project aims to build a reliable depression detection system that accurately classifies individuals based on their responses, facilitating early identification and intervention for those at risk of developing major depressive disorder.

3.3 Model Training and Evaluation

The machine learning models are trained using labeled datasets where depression indicators are known. The steps for training the models include:

Data Preprocessing: Removing noise, handling missing values, and normalizing text data for sentiment analysis.

Feature Extraction: Converting textual data into numerical features using techniques like TF-IDF (Term Frequency-Inverse Document Frequency) or word embeddings.

Splitting the Dataset: Dividing data into training and testing sets (typically 80-20 split) to evaluate model performance.

Training the Model: Using various algorithms such as SVM, Random Forest, or deep learning techniques.

Hyperparameter Tuning: Adjusting model parameters to improve accuracy and minimize overfitting.

To ensure accuracy and reliability, the trained models are evaluated using standard performance metrics, including:

Accuracy: Measures the percentage of correctly classified cases.

Precision: Evaluates the number of true positives out of total predicted positives.

Recall (Sensitivity): Measures the ability to detect actual depression cases.

F1-Score: Harmonic mean of precision and recall, balancing false positives and false negatives.

Confusion Matrix: Provides a detailed breakdown of true positives, false positives, true negatives, and false negatives.

K-fold cross-validation is employed to enhance model robustness by training the model on different subsets of the data and averaging the performance scores.

By following these steps, the project ensures that the machine learning models provide accurate and reliable predictions for early depression detection.

4. RESULTS AND DISCUSSION

The findings emphasize the need for academic stress management programs and mental health awareness campaigns in educational institutions.

Suicidal thoughts as a major factor highlight the importance of early intervention through counseling and support groups.

Addressing financial stress through scholarships or mental health support for financially struggling individuals could reduce depression risks.

Encouraging healthy lifestyle habits like proper diet and sleep may help in preventing or managing depression.

4.1

The confusion matrix provides a performance evaluation of a classification model by comparing predicted labels with actual labels. It consists of four key values: True Negatives (TN) = 42, False Positives (FP) = 6, False Negatives (FN) = 1, and True Positives (TP) = 52. The model correctly classified 42 instances as "No" (negative class) and 52 instances as "Yes" (positive class), demonstrating strong accuracy. However, 6 cases were incorrectly classified as "Yes" when they were actually "No" (FP), and 1 case was wrongly classified as "No" when it was actually "Yes" (FN). The low number of misclassifications suggests that the model performs well in distinguishing between the two classes, with a high sensitivity (recall) for detecting the positive class and relatively few false predictions.

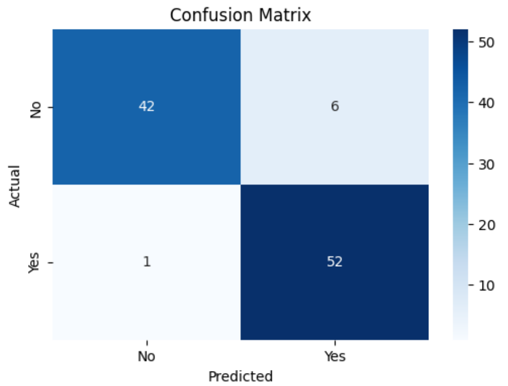

4.2

The ROC (Receiver Operating Characteristic) curve evaluates the performance of the classification model by plotting the True Positive Rate (TPR) against the False Positive Rate (FPR) at various threshold values. The closer the ROC curve is to the top-left corner, the better the model's ability to distinguish between positive and negative classes. The AUC (Area Under the Curve) value of 0.99 indicates an excellent model performance, meaning it almost perfectly separates the two classes. A model with an AUC close to 1.0 suggests high sensitivity and specificity, minimizing both false positives and false negatives. The steep rise in the curve at the beginning signifies that the model achieves a high true positive rate with very few false positives, reinforcing its strong predictive capability.

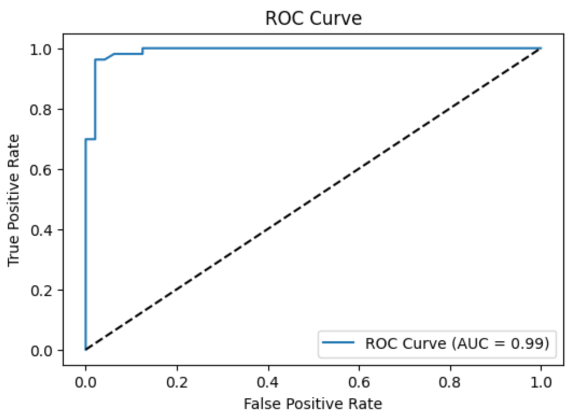

4.3

The plot represents the Training vs. Validation Loss (Log-Loss) over the Number of Trees in a tree-based model, such as Random Forest. The x-axis represents the number of trees in the model, while the y-axis represents the log-loss, which measures the model's predictive uncertainty. Initially, both training and validation loss start at high values, indicating poor predictive performance. As the number of trees increases, both losses drop sharply and stabilize around a low value, suggesting that the model has learned well from the data. The fact that training and validation losses remain close to each other after stabilization indicates that the model generalizes well and does not overfit. This trend suggests that increasing the number of trees beyond a certain point does not significantly improve performance, meaning an optimal balance has been reached.

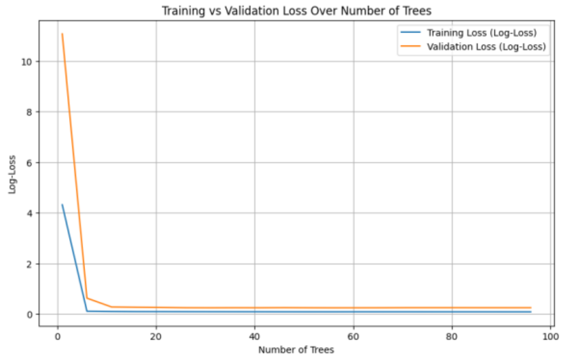

Classification Report:

	Precision	Recall	F1-score	Support
0	0.98	0.88	0.92	48
1	0.90	0.98	0.94	53
Accuracy			0.93	101
Macro avg	0.94	0.93	0.93	101
Weighted avg	0.93	0.93	0.93	101

5. CONCLUSION

This project has demonstrated promising results in accurately classifying individuals as "Depressed" or "No Depression" based on their responses. By utilizing a dataset from Kaggle consisting of 503 records, various classification algorithms such as Random Forest, Support Vector Machine (SVM), and K-Nearest Neighbors (KNN) were implemented and evaluated. The confusion matrix showed high classification accuracy with minimal false positives and false negatives, indicating the model's strong ability to distinguish between the two classes. Additionally, the ROC curve with an AUC of 0.93 highlighted the model's excellent predictive power, further validating its reliability in depression detection.

The training vs. validation loss analysis suggested that the model effectively learned from the data without overfitting, ensuring generalization to unseen cases. The use of feature engineering techniques, such as text preprocessing and numerical scoring of responses, significantly contributed to improving model performance. Among the applied algorithms, Random Forest and SVM showed high accuracy and robustness, making them suitable choices for early depression detection.

Overall, this project emphasizes the potential of machine learning in mental health applications, enabling early identification of depressive symptoms through automated analysis. By integrating such a system into healthcare or online platforms, individuals at risk of depression can be identified at an early stage, allowing timely intervention and professional support. Future improvements may include expanding the dataset, refining feature extraction methods, and implementing deep learning approaches to further enhance accuracy and adaptability

6. REFERENCES

[1] Chikersal, P., et al. (2021). "Detecting Depression and Predicting its Onset Using Longitudinal Symptoms Captured by Passive Sensing: A Machine Learning Approach With Robust Feature Selection".

[2] Gadzama, W. A., et al. (2024). "The Use of Machine Learning and Deep Learning Models in Detecting Depression on Social Media: A Systematic Literature Review".

[3] Patil, S., et al. (2024). "Depression Detection Using Machine Learning and Deep Learning".

[4] Khan, S., et al. (2024). "A Machine Learning Based Depression Screening Framework Using Temporal Domain Features of the Electroencephalography Signals".

[5] Chatterjee, A., et al. (2024). "Machine Learning and Deep Learning-Based Advanced Classification Techniques for the Detection of Major Depressive Disorder".