




International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 14, Issue 10, October 2025

Mental Health Prediction using Machine Learning Techniques

A Data-Driven Approach to Early Detection and Risk Assessment

Siva Tejaswi Jonna, Prasanna Kumar Vemula, Venkata Naga Sai Varun Panchakarla, Hariveer

Suran Singh Kshtri, Gadde Naga Vishnu Teja

Asst. Professor, Kallam Haranadha Reddy Institute of Technology, Guntur, Andhra Pradesh, India

U.G. Students, Department of CSE (Artificial intelligence & Data Science), Kallam Haranadhareddy Institute of Technology, Guntur, Andhra Pradesh, India

ABSTRACT: This project introduces an intuitive web platform with animated features, making early depression screening simple and accessible. Users can choose to fill out a straightforward ten-question form or share their thoughts and feelings in their own words through a free text entry. The system uses a blend of machine learning models including K-Nearest Neighbors, Logistic Regression, Decision Tree, and Random Forest—to analyze the collected information. Notably, when users submit text, the platform’s model predicts the emotions behind their words, giving deeper insight into their current mental state. After extensive testing, Random Forest emerged as the most accurate model, classifying depression levels with 88% accuracy.

Built with user-friendliness and inclusivity in mind, the platform connects the animated, interactive interface to a powerful backend that analyzes a wide range of clinical, demographic, behavioral, and emotional data. It delivers timely, relevant feedback, helping individuals understand their mental health and encouraging early support when needed. Special care was taken to ensure clear communication and usability for people from diverse backgrounds. Overall, this project demonstrates how modern machine learning combined with thoughtful design can provide meaningful, real-world tools for supporting mental health.

KEYWORDS: Depression screening, Random Forest, structured assessment, behavioral data, user engagement, accessible technology, Free Text.

I. INTRODUCTION

In a world increasingly shaped by rapid technological and social change, mental health has emerged as a central pillar of human well-being, transcending geographical boundaries, cultures, and economic backgrounds. Despite its significance, timely detection and intervention for mental health issues remain elusive for many, hampered by social stigma, limited access to professional resources, and the complexity of human psychology. Against this backdrop, artificial intelligence particularly machine learning offers transformative possibilities.

This journal delves into the intersection of mental health and intelligent systems, presenting a comprehensive exploration of how machine learning can revolutionize mental health assessment and prediction. By leveraging vast datasets from surveys, behavioral records, and self-reported symptoms, machine learning models uncover nuanced patterns and risk factors that often go unnoticed in traditional clinical settings. The approach amplifies the reach of early screening, allowing individuals to receive rapid, personalized feedback through accessible web-based platforms effectively democratizing mental health care.

Our project employs a robust Random Forest classifier[9] trained on thousands of data points, integrated seamlessly into a dynamic Flask web application. The platform not only empowers users to complete structured assessments and narrative evaluations but also delivers actionable insights into their mental health risk level “Not Depressed,” “Mildly Depressed,” “Depressed,” or “Critically Depressed” with high precision. Through modular system architecture, careful

feature selection, and rigorous data preprocessing, the platform supports real-time interactivity, scalability, and continuous improvement.

Moreover, this journal illustrates the broader impact of deploying such intelligent systems in real-world contexts: from healthcare facilities and educational institutions to workplace wellness initiatives and public health analytics. By bridging the gap between data science and clinical psychology, our system supports early diagnosis, promotes mental health awareness, and enables timely interventions to prevent escalation of distress. The work further explores challenges in privacy, model interpretability, and ethical deployment, setting the stage for continuous advancement in mental health technology.

II. LITERATURE SURVEY

In recent years, there has been a significant rise in the use of machine learning for mental health assessment and prediction. As more data including psychological, behavioral, and demographic information becomes available, researchers have been able to create systems that proactively identify people at risk for conditions like anxiety, depression, and stress. For example, Wang et al [2]. (2019) showed the effectiveness of logistic regression and random forest models in predicting depression risk, while Lin et al [3]. (2020) achieved strong results using natural language processing to spot distress in social media posts. Studies often blend survey responses with machine learning algorithms like Decision Trees, Support Vector Machines (SVM), and Random Forests to create reliable and easy to understand models. In practical cases, Konda Vaishnavi et al [4]. found that the Stacking technique outperformed other models with an accuracy of 81.75%, and SRIVIDYA et al [5]. used a combination of questionnaire data and advanced algorithms, like Random Forest with hyperparameter tuning, to successfully detect issues including depression, anxiety, PTSD, and insomnia. Building on these foundations, our approach combines a symptom-based questionnaire and emotion-driven text analysis in a seamless, real-time web application, introducing an original architecture capable of classifying user text into emotional categories for immediate and personalized feedback.

Web-based apps built with Flask and Django make it easy for both users and professionals to access mental health prediction tools whenever they need them. By allowing people to enter their details and even free-form text about how they're feeling, these platforms give instant, helpful feedback and support early action. In this project, we use several machine learning models Random Forest, K-Nearest Neighbours, Logistic Regression, and Decision Tree to analyze survey responses and free text. Our goal is to move beyond basic questionnaires and provide real-time, personalized results that are easy to access and scale, helping bridge the gap between technology and real mental health support.

FLASK:

Flask is a lightweight and powerful web framework for Python. It's often called a "micro-framework" because it provides the essentials for web development without unnecessary complexity. Unlike Django, which comes with built-in features like authentication and an admin panel, Flask keeps things minimal and lets us add only what we need.

Werkzeug and Jinja2: Flask is built on top of two powerful libraries:

Werkzeug handles the core web server functions, taking care of everything related to processing incoming requests and sending responses from your application. This makes the server more reliable and efficient, so developers can focus on their own features. Jinja2, meanwhile, is a flexible templating tool that lets you design web pages with dynamic content—meaning you can easily use variables, loops, and other constructs to build HTML pages that adapt to each user or situation. Together, they form the backbone of powerful and interactive Python web applications.[7]

JAVA SCRIPT:

JavaScript is a versatile, dynamically typed programming language that brings life to web pages by making them interactive. It is used for building interactive web applications, supports both client-side and server-side development, and integrates seamlessly with HTML, CSS, and a rich standard library.

JavaScript is designed to run one task at a time, making it a single-threaded language. As an interpreted language, it processes and executes code line by line rather than compiling it all at once. Additionally, the type of data a variable

holds in JavaScript is determined while the program is running, not beforehand, which is why JavaScript is considered dynamically typed.[8]

SCIKIT-LEARN:

The sklearn.ensemble module offers two main techniques Random Forest and Extra-Trees that use randomness to build multiple decision trees. Instead of relying on a single tree, these methods create a group of diverse classifiers by adding variation during tree construction. When making predictions, they average the outcomes from all these trees, resulting in a more reliable and accurate overall result.[9]

PICKLE:

The pickle module provides binary methods for converting Python object hierarchies into byte streams a process called "pickling" and restoring them from byte streams back into objects, known as "unpickling." This mechanism is also referred to as serialization, marshalling, or flattening, but the preferred terms here are "pickling" and "unpickling." [10]




Fig 1:ML Prediction System Workflow

III. PROPOSED METHODOLOGY

The proposed mental health screening system offers two user-friendly ways for individuals to assess their depression risk: a structured questionnaire and a free-text narrative input. Both approaches are designed to provide flexibility and capture diverse user experiences.

Data Collection:

User data is collected via an animated web interface. For the questionnaire path, users respond to ten specifically designed questions about depressive symptoms such as mood, energy, sleep, and concentration on a four point scale (0 to 3) reflecting increasing severity. For the free-text path, users are invited to describe their emotional state in their own words, like a journal entry. The dataset comprises both these structured responses and open-ended text, collected with user consent and anonymized for privacy.[6]

Data Preprocessing:

Prior to model training, each type of input undergoes targeted preprocessing. Questionnaire responses are numerically encoded, and missing or inconsistent data is handled appropriately. TF-IDF (Term Frequency-Inverse Document Frequency) feature extraction is applied to convert questionnaire responses into feature vectors that highlight depression-related symptoms. For free-text inputs, text is cleaned, tokenized, and similarly transformed with a TF-IDF vectorizer but this time tuned to extract emotion-specific keywords that reflect user sentiment.

Classification Logic:

Questionnaire-Based:

User choices on the four-point scale are directly linked to depression risk. If a user consistently selects the first (least severe) option for all questions, the system classifies them as “Not Depressed.” Consistently picking the highest (most severe) option results in a “Depressed” label. Answers that fall in between (mixed selections of 2s and 3s, or occasional severe responses) can classify the user as “Mildly Depressed” or “Depressed,” depending on overall pattern and thresholding logic designed to prioritize safety and early intervention.

Free-Text-Based:

The model analyzes the user’s narrative for dominant emotions using features extracted by TF-IDF from a pre-trained emotion classifier. The emotion categories are joy, sadness, fear, anger, and surprise. The system then uses a simple mapping: Joy or Surprise is labeled “Not Depressed,” Sadness signals “Depressed,” and Fear or Anger typically results in “Mildly Depressed.” This provides a nuanced bridge between raw text, emotional tone, and mental health risk.

Model Training:

Both approaches utilize a Random Forest model (and additional models like KNN, Logistic Regression, and Decision Tree for comparison), trained on a comprehensive Kaggle dataset for questionnaire responses and the Emotions Dataset for NLP for emotion recognition in text. Feature engineering and balanced splits ensure robust, generalizable learning.

Evaluation and Confusion Matrix :

To assess performance, the model’s predictions are compared to true labels using a confusion matrix an essential diagnostic tool in multiclass classification. The provided confusion matrix summarizes how well the emotion classifier distinguishes between “anger,” “fear,” “joy,” “love,” “sadness,” and “surprise.” Strong diagonal values (for example, 236 true “anger” out of 275, 643 true “joy” out of 695, and 527 true “sadness” out of 581) indicate that the model can accurately identify these emotions in user inputs. Some off-diagonal entries (such as misclassification between “love” and “joy,” or “joy” and “surprise”) are to be expected given the subtleties of human emotion, yet overall accuracy remains high.

Figure 2 summarizes the main decision flow in our mental health prediction system, designed to deliver feedback that matches each user's preferred way of sharing how they feel. Individuals can either answer a structured symptom questionnaire or write freely about their emotions. If they choose the questionnaire, their answers are transformed into key features using TF-IDF extraction, focusing on signs related to depression. A Random Forest model then analyzes these features to predict whether they fall into categories like “Not Depressed” , “Mildly Depressed” or “Depressed.”

For those who prefer open-ended text, the system looks for emotion-specific keywords using TF-IDF, and sentiment analysis identifies the user's dominant feeling such as joy, fear, surprise, anger, or sadness. Each emotion is translated into a matching mental health status. Both input routes meet at a final mapping stage, where the system provides a holistic assessment tailored to the user. By combining both structured and expressive data, this approach makes the platform more versatile and welcoming, supporting a broader range of users and contributing to personalized, effective screening in real-world mental health scenarios.




Fig 2: Mental Health Prediction Workflow

IV. EXPERIMENTAL RESULTS

The presented mental health prediction platform was rigorously evaluated using both questionnaire and free-text inputs, leveraging a labelled dataset of over 15,000 samples. The evaluation primarily focused on four machine learning classifiers: K-Nearest Neighbors (KNN), Decision Tree, Logistic Regression, and Random Forest. Each algorithm was assessed for its predictive accuracy, with the Random Forest ultimately selected as the final model owing to its superior performance on our dataset .

1. Performance:

Notably, the Random Forest [9] model achieved the highest testing accuracy at 88.1%, considerably outperforming the next-best model. This result is even more impressive when compared with earlier studies on similar datasets, such as Konda Vaishnavi et al. [4] , where Random Forest achieved 81.75%, Logistic Regression 79.63%, KNN 80.42%, and Decision Tree 80.69%. The enhancement in accuracy can be attributed to the comprehensive feature extraction, robust preprocessing, and classification pipeline implemented in this project.

The accuracy of each classifier is summarized below:

Table 1: Accuracy of different classifiers in this study

Model	Accuracy (%)
K-Nearest Neighbors	73.2
Decision Tree	83.56
Logistic Regression	87.10
Random Forest	88.10

2. Training, Validation, and Testing Metrics:

These performance scores confirm the model's robustness and its ability to generalize to unseen inputs.

Table 2: Training, validation, and test set performance metrics.

Metric	Value
Training Accuracy	99.8%
Validation Accuracy	88.3%
Testing Accuracy	87.2%
Average Log Loss (Train)	0.11
Average Log Loss (Validation)	0.41
Average Log Loss (Test)	0.41

3. Classification Report by Emotion Labels:

The model shows especially high recall and F1-scores for critical classes like "joy" and "sadness," underscoring its effectiveness in real-world scenarios.

Table 3: Label-wise precision, recall, F1-score, and support

Label	Precision	Recall	F1-score	Support
Anger	0.90	0.83	0.86	275
Fear	0.88	0.78	0.83	224
Joy	0.84	0.96	0.90	695

Label	Precision	Recall	F1-score	Support
Love	0.83	0.61	0.70	159
Sadness	0.90	0.93	0.92	581
Surprise	0.85	0.53	0.65	66
Accuracy			0.87	2000
Macro Avg	0.87	0.77	0.81	2000
Weighted Avg	0.87	0.87	0.87	2000


4.Modalities:

In questionnaire mode, the system effectively classified a range of user response patterns, from mild to moderate and severe symptoms, and reliably triggered the override logic to output "Depressed" whenever all answers indicated maximum severity. In free-text mode, the model accurately identified the dominant emotions expressed in user narratives, such as joy, fear, anger, surprise, and sadness, and correctly translated these emotions into the corresponding depression risk category. This dual approach ensured that both structured and unstructured user input could be mapped to clinically meaningful outcomes.

Table 4: Example Output Scenarios for Mental Health Prediction

Input Mode	Example Input	Model Output
Questionnaire	All low or mild symptoms	Not Depressed
Questionnaire	Several answers at “2” or “3” severity	Mildly Depressed / Depressed
Questionnaire	All “3” (max)	Depressed (override)
Free-Text	“I feel tired, hopeless, anxious”	Emotion: Sadness – Depressed
Free-Text	“I’m excited and enjoying my semester”	Emotion: Joy – Not Depressed
Free-Text	“Feeling worried about exams”	Emotion: Fear – Mildly Depressed

User Interface Evaluation: Screenshots of both the questionnaire and text input interface (Figure 2a & 2b), alongside a confusion matrix from test results (Figure 3), demonstrate both usability and model accuracy.



The image shows a digital interface for a 'Mental Health Questionnaire'. The title is centered at the top. Below it, the first question asks about emotional feelings over the past two weeks. There are four response buttons: 'Generally positive', 'Sometimes low', 'Often sad', and 'Very low or hopeless'. A progress bar is visible, and a 'Next' button is at the bottom right.

Mental Health Questionnaire


1. How have you been feeling emotionally over the past two weeks?

Generally positive Sometimes low

Often sad Very low or hopeless

Next

Figure (2a): Questionnaire interface



The image shows a web interface titled 'Human Emotion Classifier'. It features a 'Home' button in the top right corner. The main area has a text input field with the placeholder 'Type something to analyze...'. Below the input field are two buttons: 'Predict' and 'Clear'.

Human Emotion Classifier Home

Enter text

Type something to analyze...

Predict Clear

Fig 2b: Text Input Interface




Fig 3: Confusion Matrix

V. DISCUSSION

The absolute results of our experimental setup indicate that the Random Forest classifier thanks to feature engineering, balanced preprocessing, and tailored hyperparameters offers a powerful and reliable tool for mental health screening. Its accuracy and precision, as showcased above, both advance the field and surpass benchmarks established in previous studies.

VI. CONCLUSION

This research developed a dual-mode web platform for early depression detection, using both structured questionnaires and free-text inputs. Random Forest emerged as the most reliable model, achieving 88% test accuracy and outperforming comparable approaches. The system excelled in both classifying symptom severity and interpreting user emotions from open-ended text. Its interactive design, fast assessment time, and robust machine learning foundation make the tool practical for real-world mental health applications. Overall, the platform demonstrates how AI can support accessible, efficient, and personalized mental health screening.

VII. ACKNOWLEDGMENT

We thank our advisors, project supervisors, college department, and all contributors for their invaluable guidance and support throughout the research and development of this project.

REFERENCES

1. Anshuman2305, "Mental-Health-CodeUtsava: A Machine Learning Web App using Flask on Mental Health," GitHub. [Online]. Available: <https://github.com/Anshuman2305/Mental-Health-CodeUtsava>.
2. Wang et al. (2019). Used logistic regression and random forest models to successfully predict depression risk from behavioral and demographic factors.

3. Lin et al. (2020). Demonstrated high accuracy when using natural language processing to analyze social media for signs of distress.
4. Konda Vaishnavi, U Nikhitha Kamath, B Ashwath Rao, N V Subba Reddy. "Predicting Mental Health Illness using Machine Learning Algorithms." Journal of Physics: Conference Series, vol. 2161, 2022, 012021. DOI:10.1088/1742-6596/2161/1/012021.
5. Satvik Gurjar, Chetna Patil, Ritesh Suryawanshi, Madhura Adadande, Ashwin Khore, Noshir Tarapore. "Mental Health Prediction Using Machine Learning." International Research Journal of Engineering and Technology (IRJET), vol. 9, no. 12, Dec 2022, pp. 1117-1121.
6. Praveen, "Emotions Dataset for NLP," Kaggle, 2019. Available: <https://www.kaggle.com/datasets/praveengovi/emotions-dataset-for-nlp>.
7. Flask Documentation. Available at: <https://flask.palletsprojects.com/> [Accessed October 2025].
8. GeeksforGeeks, "Introduction to JavaScript," Available: <https://www.geeksforgeeks.org/javascript/introduction-to-javascript/>. Accessed: October 2025.
9. scikit-learn, "1.11. Ensembles: Gradient boosting, random forests, bagging, voting, stacking," Available: <https://scikit-learn.org/stable/modules/ensemble.html#random-forests-and-other-randomized-tree-ensembles>. Accessed: October 2025.
10. Python Software Foundation, "pickle — Python object serialization," Available: <https://docs.python.org/3/library/pickle.html>. Accessed: October 2025.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details