MindSight - Mental health and well-being surveillance, assessment and tracking solution

Susheeth G[1], Ritish N[2], Vidyashree B N[3], Mithali S Anand[4], Tejashwini B A[5]
Presidency School of Computer Science and Information Science, Presidency University, Itgalpura, Rajanukunte, Bengaluru – 560064

Abstract*— *MindSight is a web-based mental health assessment platform designed to help users evaluate their psychological well- being by self-assessing levels of stress, anxiety, and depression. Developed as part of the final year project ISR-G03 at Presidency University, the application leverages the scientifically validated DASS-21 (Depression Anxiety Stress Scales) questionnaire to collect user input across 21 questions—seven each for stress, anxiety, and depression. The system uses multiple machine learning models trained on the publicly available DASS-21 dataset to predict the severity levels of each condition, categorized as Normal, Mild, Moderate, Severe, or Extremely Severe. Among the models tested—including Logistic Regression, Decision Tree, Random Forest, XGBoost, and others—the Support Vector Machine (SVM) model achieved the highest accuracy and was selected for final deployment. The platform features a user-friendly interface built using HTML and CSS, with a Python Flask backend and scikit-learn for ML integration. Predictions are visually communicated through intuitive, color-coded indicators, ensuring accessible and engaging user experience. MindSight aims to support early mental health intervention, especially among children, by providing a quick, anonymous, and reliable self-assessment tool.

Index Terms*— *Mental Health Assessment, DASS-21, Depression, Anxiety, Stress, Machine Learning, Support Vector Ma- chine (SVM), Web Application, Flask, Python, scikit-learn, Self- Assessment Tool, Child Mental Health, Psychological Screening, Mental Well-being, SVM Classifier.

I. INTRODUCTION

Mental health has become an increasingly important concern in today's world, especially among children and young adults. Many people face challenges like stress, anxiety, and depression but do not always recognize the symptoms or seek help. Early detection and awareness can make a big difference in managing mental health problems before they grow more serious. That's where technology can play a helpful role by making mental health support more accessible and user-friendly.MindSight is a web-based mental health assessment tool developed as part of a final year project at Presidency University. It is designed to help users self-check their mental well-being by answering a set of 21 short questions. These questions are based on the DASS-21 (Depression, Anxiety, and Stress Scales), a scientifically accepted psychological tool that measures the levels of these three emotional states. The app is especially aimed at children and young adults, providing a safe, private, and easy way to assess their mental state. The app uses machine learning to analyze the responses and predict the severity level for each condition. Various machine learning models were trained and tested using the DASS-21 dataset. After comparing the performance of models like Logistic Regression, Random Forest, and XGBoost, the Support Vector Machine (SVM) was found to be the most accurate and reliable. This model was selected for deployment in the final version of the web application.

MindSight offers a simple and interactive user interface, built using HTML and CSS, while the backend is powered by Python Flask. The final results are displayed with personalized, color-coded indicators to make them easy to understand. By combining machine learning with a clean user interface, MindSight aims to make mental health self-assessment more accessible, especially for young users who may find it difficult to talk openly about their mental health

II. LITERATURE REVIEW

The application of machine learning (ML) and natural language processing (NLP) in mental health prediction has gained significant momentum in recent years. With rising concerns around stress, anxiety, and depression, particularly among younger populations, researchers have been exploring intelligent systems that can assist in early detection and intervention through digital means. Several studies have demonstrated the feasibility of AI-driven mental health assessments while also revealing critical gaps in real-world applicability, scalability, and ethical considerations. One such study, ”Prediction of Public Mental Health by using Machine Learning Algorithms,” utilizes ML models to identify early signs of mental health issues and offers valuable insights for public policy-making. It highlights how data science can be leveraged to detect trends and support mental health awareness at a broader societal level. However, it also points out challenges such as the complexity of implementation, the need for strong stakeholder collaboration, and the difficulty of integrating diverse real-world datasets due to privacy and formatting issues. In a different approach, the study titled ”Mental Health Prediction using Natural Language Processing” investigates how NLP can be applied to analyze text data, such as written responses or online content, to identify psychological distress. This technique supports scalable mental health assessments and enables automated evaluations. Nevertheless, its reliability is dependent on the effectiveness of underlying models and training data quality. Additionally, NLP systems may not capture the full spectrum of mental well-being, as they often exclude external or behavioral factors. The paper ”Application of Machine Learning to Predict Mental Health Disorders and Interpret Feature Importance” demonstrates how interpretable ML models can enhance dis- order detection using survey-based input. This approach not only improves prediction accuracy but also helps clinicians understand which features most significantly influence mental health. While beneficial, the study reports limited accuracy (79.78%) and highlights the computational demands and challenges of applying such systems in dynamic real-world environments. Beyond algorithmic methods, recent research has focused on incorporating behavioral and social data for mental health prediction. The paper ”Predicting Mental Health Problems with Personality, Behavior, and Social Networks” emphasizes the role of personality traits and social connectivity in early mental health detection. Behavioral pattern analysis can offer deeper emotional insights, particularly within educational institutions or workplace environments. However, the collection of social and behavioral data raises significant ethical and privacy concerns, and the models may not generalize well beyond specific populations like college students. Similarly, ”From Personalized Medicine to Population Health: A Survey of mHealth Sensing Techniques” provides a taxonomy of mobile health (mHealth) technologies that sup- port large-scale mental health monitoring. The paper explores how data from wearable devices and mobile applications can be combined with crowdsourced inputs to enable population- level analysis. Despite the innovation, real-world validation is limited, and the integration of personal data raises privacy and interoperability challenges. Another notable study, ”Mental Health Surveillance Among Children in the United States (2013–2019)”, presents an ex- tensive overview of mental health trends across age, gender, and socio-economic groups. The study offers valuable support for public policy development. However, it relies heavily on parent-reported data, which can introduce bias, and lacks a unified system for data collection and analysis, affecting the consistency and comprehensiveness of findings. Further advancing this line of research, the paper ”Mental Health and Well-being Surveillance, Assessment, and Tracking Solution Among Children” proposes an AI-driven framework integrating both national and international data sources to support targeted interventions. The system is positioned to assist policymakers but lacks concrete real-world validation. Additionally, the study does not adequately address ethical concerns related to data security and consent. Complementing this work, ”Mental Health and Well-being Surveillance System” explores the use of AI/ML for real- time monitoring of mental health. It supports quick response to mental health trends by integrating multiple data sources. However, the system faces implementation hurdles, such as inconsistency in data formats, challenges in real-time feasibility, and bias from imbalanced training datasets.

III. PROCESS OF EXECUTION

The MindSight project follows a step-by-step execution process to ensure accurate mental health assessment. The process begins with data collection using the DASS-21 dataset, which contains self-reported responses to 21 psychological questions. These responses are categorized into three groups: stress, anxiety, and depression. After collecting the data, it is preprocessed to remove any inconsistencies or missing values. The cleaned data is then used to train various machine learning models. After training, the best-performing model is integrated into a user-friendly web application. Users can access the app, take the assessment, and receive instant feedback based on their answers.

Data Preprocessing Module

Before training any model, the dataset must be prepared for accurate analysis. In this module, the data is cleaned by handling any missing values and normalizing the input scores. Each question in the DASS-21 dataset has responses rated from 0 to 3. These values are grouped and summed into three categories: stress, anxiety, and depression. The total scores for each category are then labeled according to their severity levels (Normal, Mild, Moderate, Severe, or Extremely Severe). This processed data is then used to create input-output pairs suitable for machine learning models.

Model Development

Multiple machine learning models were developed and tested to find the one that works best for predicting mental health conditions. The models used include Logistic Regression, Decision Tree, Random Forest, K- Nearest Neighbors (KNN), Naive Bayes, Gradient Boosting, AdaBoost, XGBoost, Stochastic Gradient Descent (SGD), and Support Vector Machine (SVM). Each model was trained using the processed DASS-21 dataset and evaluated using accuracy scores and performance metrics. Among all, the SVM model performed the best, offering the highest prediction accuracy and consistency. This model was chosen for final integration into the MindSight web application.

Assessment of the Model

Evaluating the performance of machine learning models is a critical step in ensuring the reliability and accuracy of any predictive system, especially in sensitive areas like mental health. In MindSight, a thorough assessment of each trained model was conducted using standard evaluation metrics such as accuracy, precision, recall, F1-score, and confusion matrix analysis. These metrics were applied to each classification task—predicting the severity level of stress, anxiety, and depression.

The assessment revealed that the Support Vector Machine (SVM) model significantly outperformed other models across all three categories. SVM was particularly effective in handling the multiclass classification

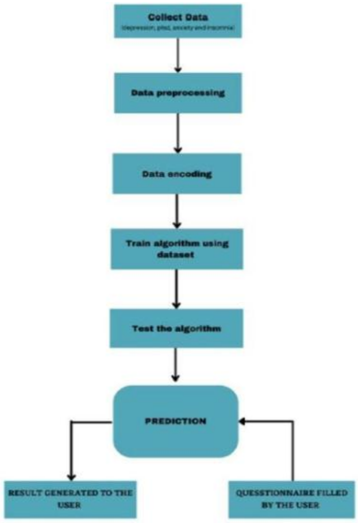

Fig. 1: Workflow Diagram

of severity levels—ranging from *Normal* to *Extremely Severe*. Its ability to create clear decision boundaries in complex data spaces made it ideal for this type of psychological assessment. The model demonstrated high generalization ability, maintaining strong performance even during cross-validation, which indicates it can make reliable predictions for new users outside the training data.This rigorous evaluation ensured that only the most robust and accurate model was integrated into the final web application, providing users with reliable insights about their mental health status. By prioritizing model precision and consistency, MindSight helps build user trust while ensuring responsible handling of sensitive psychological data.

Dataset Description

The DASS-21 dataset forms the foundation of the Mind- Sight system and plays a crucial role in its accuracy and effectiveness. This dataset is widely recognized in psychological research and includes anonymized, self- reported responses to 21 well-structured questions. These questions are grouped into three core categories:7 for Depression, 7 for Anxiety, and 7 for Stress. Each response is rated on a 4-point Likert scale (0–3), reflecting how frequently the user has experienced a specific symptom over the past week.The rich and balanced structure of the DASS-21 dataset makes it particularly suitable for training machine learning models. After collecting the responses, scores are summed within each category, and the totals are mapped to five severity levels—Normal, Mild, Moderate, Severe, and Extremely Severe—based on established psychological thresholds. This clear mapping allows for precise classification and provides a strong supervised learning framework.The dataset's credibility, completeness, and psychological validity make it a powerful tool for building mentalhealth assessment models. Its widespread acceptance in both academic and clinical communities adds value and trustworthiness to the predictions made by the MindSight platform. By leveraging this dataset, the project bridges the gap between machine learning technology and mental health awareness, ensuring that results are not only accurate but also meaningful and clinically relevant.

ML Algorithms For Classification

In this we have used, ten machine learning models—Logistic Regression, Support Vector Machine (SVM), Decision Tree, Random Forest, Gradient Boosting, AdaBoost, K-Nearest Neighbors (KNN), Naive Bayes, Stochastic Gradient Descent (SGD), and XG- Boost—are trained using DASS-21 questionnaire responses. Each model was individually evaluated for its ability to classify the severity of each condition.

Logistic Regression

Logistic Regression was used as a baseline model to classify the severity of stress, anxiety, and depression by modeling the relationship between cumulative DASS-21 scores and their corresponding labels. It helped determine if a simple linear boundary could separate classes like mild and severe depression. In development, it quickly showed how well linearly separable classes could be predicted, though it struggled with overlapping ranges. The model worked best when the score thresholds clearly defined the class boundaries.

Support Vector Machine (SVM)

SVM was implemented with an RBF kernel to effectively handle the non-linear patterns present in the questionnaire data. It created a hyperplane that maximized the margin between classes, making it ideal for distinguishing closely spaced severity levels such as moderate and severe stress. The model performed exceptionally well in identify- ing subtle differences between overlapping DASS score ranges and was selected as the final deployed model for both stress and depression due to its superior accuracy.

Decision Tree

Decision tree models were trained using raw DASS score inputs, making decisions by splitting the dataset based on the most significant features at each node. In our system, they helped visualize how individual questions influenced predictions, such as how a high score.

Random Forest

Random Forest was used to improve on the Decision Tree by training an ensemble of trees and combining their outputs for a more robust prediction. It was particularly useful for reducing variance and generalizing well across diverse score patterns. This model also provided insights into feature importance, revealing which DASS-21 items contributed most to classifying each mental health condition. It handled score fluctuations across classes better due to its averaging nature.

Gradient Boosting

Gradient Boosting was applied to iteratively learn from the mistakes of previous weak learners, refining predic tions for borderline cases. It proved effective in scenarios where score ranges overlapped slightly or where rare classes (like extremely severe anxiety) were harder to detect. This model adjusted well to complex relationships in score progression and was particularly useful during model comparison due to its high performance on mildly imbalanced data.

AdaBoost

AdaBoost focused on misclassified DASS entries during training, increasing their weight in subsequent iterations to enhance detection of difficult cases. This approach wasbeneficial for conditions with skewed class distributions, such as when extremely severe stress cases were limited. Although not as powerful as XGBoost, it served as an effective boosting technique to increase classification focus on edge cases within close score thresholds.

K-Nearest Neighbors (KNN)

KNN was used to explore instance-based learning by classifying a user based on the most similar questionnaire patterns among the training data. It predicted the severity level based on the majority class of the closest 'k' responses using Euclidean distance. This model worked well when there were clear score clusters but became less efficient as the dataset grew or when classes were not distinctly separable by raw distances.

Naive Bayes

Naive Bayes was incorporated for its speed and probabilistic interpretation of the DASS responses. Despite assuming independence between features, which is not strictly true for questionnaire data, it produced reasonable classification results. It was especially useful for quick predictions in early testing stages and helped determine the probability of a user belonging to each mental health class based on their score pattern.

Stochastic Gradient Descent (SGD)

SGD was used as a scalable optimization technique for training linear models on the DASS data, especially when experimenting with larger datasets. It updated model weights iteratively for each response, enabling faster model development cycles. While it required feature scaling and careful tuning, it helped refine other models and was effective in testing combinations of loss functions like hinge and log loss.

XGBoost (Extreme Gradient Boosting)

XGBoost emerged as the top performer for anxiety classification due to its regularized boosting framework, which controlled overfitting and handled imbalanced score ranges well. It efficiently managed overlapping GAD-7 values and produced consistent results even with subtle changes in question responses. XGBoost was integrated into the final backend pipeline for its speed, reliability, and high classification performance.

Evaluation Measures

To evaluate the performance of each machine learning model, four primary classification metrics were used: Accuracy, Precision, Recall, and F1 Score. These metrics were computed separately for stress, anxiety, and depression prediction tasks to ensure balanced evaluation across all classes.Accuracy is defined as the ratio of the number of correct predictions to the total number of predictions. While it is a useful overall measure, accuracy can be misleading when class distributions are imbalanced.

$$\text{Accuracy}=\frac{Number\ Of\ Correct\ Predictions}{Total\ Number\ of\ Predictions} \quad (1)$$

Precision is particularly important in imbalanced classification problems. It refers to the proportion of true positive predictions among all predicted positives. A high precision score indicates that the model is not labeling negative samples as positive.

$$Precision = \frac{\text{TruePositives}}{TruePositives+FalsePositives} \quad (2)$$

Recall (also known as sensitivity) measures the model's ability to identify all relevant cases within a class. It is the proportion of actual positives correctly predicted by the model. F1 Score is the harmonic mean of precision and recall, providing a single metric that balances both. It is especially useful when there is an uneven class distribution or when both false positives and false negatives carry significant costs.

$$\text{F1 Score}=2 * \frac{Precision * Recall}{Precision + Recall} \quad (3)$$

Model	Precision	Recall	F1 Score
SVM	97	97	97
Decision Tree	91	91	91
Random Forest	95	95	95
Gradient Boosting	95	94	94
AdaBoost	91	91	91
KNN	90	90	90
Naive Bayes	88	89	88
SGD	89	89	89
XGBoost	96	95	95

TABLE I: Model Precision, Recall, F1 Score Average

IV. RESULTS

After training and evaluating multiple machine learning models on the DASS-21 dataset, the Support Vector Machine (SVM) emerged as the most accurate and reliable classifier. It achieved the highest accuracy across all three categories—depression, anxiety, and stress with 98.07 , 97.25, 98.69 respectively—outperforming other models such as Random Forest, XGBoost, Logistic Regression, and K-Nearest Neighbors. The SVM model was able to clearly distinguish between different severity levels ranging from Normal to Extremely Severe, making it ideal for a mental health screening tool. The predictions from the SVM model were integrated into the MindSight web application, where users receive instant and color-coded feedback based on their self- assessment. The system provided consistent results during cross-validation and testing, confirming its generalizabilty to new, unseen data. The intuitive front-end design ensured that users could easily interpret their results and understand the severity level associated with their mental state.

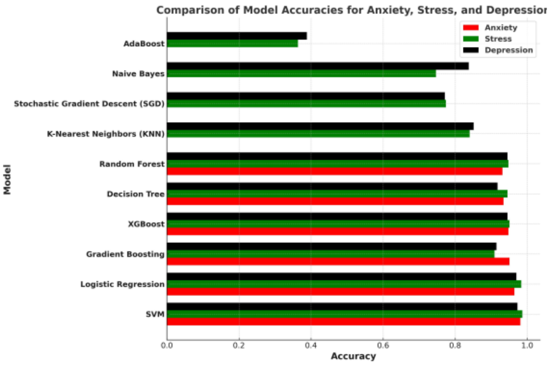

Fig. 2: Model Comparison Graph

Overall, the results demonstrate that machine learning can effectively support early detection and monitoring of mental health conditions, especially when backed by scientifically validated data like DASS-21.

Model	Stress	Anxiety	Depression
Logistic Regression	94.23	93.1	94.5
SVM	98.07	96.72	97.25
Decision Tree	92.3	93.65	92.6
Random Forest	95.38	95.92	96.2
Gradient Boosting	95.38	94.83	95.1
AdaBoost	93.84	92.5	93.84
KNN	92.88	92.88	92.7
Naive Bayes	90.76	91.2	90.8
SGD	91.53	91.53	91.4
XGBoost	96.15	95.3	96.92

TABLE II: Model Performance Accuracy Table

V. CONCLUSION AND FUTURE SCOPE

MindSight successfully demonstrates the power of combining psychological screening tools with machine learning to create an effective and user-friendly mental health assessment platform. The project proves that by using publicly available and clinically validated datasets, it is possible to train accurate models capable of providing meaningful predictions about a person's emotional well- being. Through rigorous model training and evaluation, the SVM model was chosen for deployment due to its superior performance. The final web app offers a private, anonymous, and accessible way for users—especially children and young adults—to assess their mental health status. The platform not only supports early detection but also promotes awareness and encourages proactive mental well-being. In conclusion, MindSight represents a step forward in the digital transformation of mental healthcare, providing a scalable and affordable solution for initial psychological assessment. With further development, such tools could be expanded to include more features, such as personalized tips, mental health resources, or even integration with counseling services.

Future Scope: In the future, it can be enhanced by incorporating more diverse datasets from various demographics to improve generalization. Integrating additional features such as social behavior patterns, physiological signals (e.g., heart rate), or natural language inputs could increase prediction accuracy. The chatbot component could be made more intelligent with NLP-based conversation flows, enabling better user interaction and mental health guidance. Additionally, partnerships with mental health professionals and institutions could help in validating and deploying this system in clinical or academic settings for real-world impact.

REFERENCES

[1] C. M. H. Saibaba, K. V. K. Alekhya, K. Yeshwanth and P. Tumuluru, ”Prediction of Public Mental Health by using Machine Learning Algorithms,” 2022 Second International Conference on Artificial Intelligence and Smart En- ergy (ICAIS), Coimbatore, India, 2022, pp. 308-311,doi: 10.1109/ICAIS53314.2022.9742837.

[2] R. Ranjana, T. Subha, K. L. Sreenidhi and K. Akshaya, ”Mental Health Prediction using Natural Language Process- ing,” 2022 International Conference on Applied Artificial Intelligence and Computing (ICAAIC), Salem, India,2022,pp.824828,doi:10.1109/ICAAIC53929.2022.9792893

[3] Y. Li, ”Application of Machine Learning to Predict Mental Health Disorders and Interpret Feature Importance,” 2023 3rd International Symposium on Computer Technology and Information Science (ISCTIS), Chengdu, China, 2023,pp.257-261, doi:10.1109/ISCTIS58954.2023.10213032

[4] D. Zhang, T. Guo, S. Han, S. Vahabli, M. Naseriparsa and F. Xia, ”Predicting Mental Health Problems with Personal- ity, Behavior, and Social Networks,” 2021 IEEE International Conference on Big Data (Big Data), Orlando, FL, USA, 2021,pp.4537-4546, doi:10.1109/BigData52589.2021.9671987.

[5] C. E. R. Booc, C. M. D. San Diego, M. L. Tee and J. D. L. Caro, ”A mobile application for campus- based psychosocial wellness program,” 2016 7th International Conference on Information, Intelligence, Systems & Applications (IISA), Chalkidiki, Greece, 2016, pp. 1-4, doi: 10.1109/IISA.2016.7785426.

[6] M. K. I. Zim, M. A. Hanif and H. Kaur, ”Prediction of personality for mental health detection using hybrid deep learning model,” 2024 IEEE International Conference on Inter-disciplinary Approaches in Technology and Management for social innovation, Gwalior, India, doi: 10.1109/IATMSI60426.2024.10503423

[7] S. BH et al., ”Mental Health Analysis of Employees using Machine Learning Techniques,” 2022 14th International Con-ference on COMmunication Systems & NETworkS (COM- SNETS), Bangalore, India, 2022, pp. 1-6, doi: 10.1109/COM-SNETS53615.2022.9668526.

[8] C. M. H. Saibaba, K. V. K. Alekhya, K. Yeshwanth and P. Tumuluru, ”Prediction of Public Mental Health by using Machine Learning Algorithms,” 2022 Second Interna-tional Conference on Artificial Intelligence and Smart En- ergy (ICAIS), Coimbatore, India, 2022, pp. 308-311, doi: 10.1109/ICAIS53314.2022.9742837.

[9] N. Jayakumar and R. N, ”Modeling Mental Health: Advances in Predictive Science towards Proactive Health Care,” 2024 IEEE International Conference for Women in Innovation, Technology & Entrepreneurship (ICWITE), Bangalore, India, 2024, pp. 202-206, doi: 10.1109/ICWITE59797.2024.10503111.

[10] X. Jin et al., ”Real-Time Scheduling of Massive Data in Time Sensitive Networks With a Limited Number of Schedule Entries,” in IEEE Access, vol. 8, pp. 6751-6767, 2020, doi: 10.1109/ACCESS.2020.2964690.

[11] S. Yang and S. N. Jat, ”Genetic Algorithms With Guided and Local Search Strategies for University Course Timetabling,” in IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), vol. 41, no. 1, pp. 93-106, Jan. 2011, doi: 10.1109/TSMCC.2010