IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Solace: A Mental Well-being Web-Application

[1]Shrimann Vyas, [2]Dewendra Bharambe, [3]Aarohi Pashankar,[4] Kaustubh Wagh,[5]Prasanna Bhujbal

[1]Student, [2]Assistant Professor, [3]Student, [4]Student, [5]Student

[1]Computer Engineering,

[1]ABMSP's Anantrao Pawar College Of Engineering & Research, Pune, India

Abstract: We present Solace, a novel web-based platform for mental well-being support that integrates artificial intelligence (AI) for automated screening and personalized assistance. Solace employs a random forest classifier to detect the presence and severity of stress, anxiety, and depression, a rule-based recommender to suggest targeted self-assessment questions, and an NLTK-driven rule-based chatbot using keyword recognition and predefined responses. We report multi-label (presence/absence of each condition) and multi-class (severity levels) performance metrics (accuracy, precision, recall, F1-score) for our model, which demonstrates robust detection capabilities, achieving ≈68–72% accuracy for multi-class severity classification (Table I). We discuss privacy safeguards, explainable design (using feature importance analysis), and ethical considerations in AI-driven mental health applications. Our results, along with user feedback, underscore Solace's novelty in combining predictive screening, personalized recommendations, and conversational support. The practical impact lies in scalable, accessible mental health assistance that complements clinical care**.**

***Index Terms* -** Mental well-being, artificial intelligence, ma- chine learning, stress, anxiety, depression, chatbot, recommender systems.

I. Introduction

Mental health disorders such as stress, anxiety, and depression are prevalent global health issues that significantly impair quality of life. The World Health Organization estimates that over 300 million people suffer from depression worldwide, often co-occurring with anxiety or chronic stress. Despite this burden, stigma and limited access to care deter many from seeking help. Digital interventions (e.g., smartphone and web apps) have emerged to supplement traditional therapy, offering convenient, private support. However, many existing solutions provide static content or counseling without integrated AI for personalized assessment. We introduce *Solace*, an AI-integrated web platform designed to address this gap by combining automatic screening, personalized guidance, and conversational assistance in one tool.

Key contributions of this work include: (1) Presence Detec- tion and Severity Classification of stress, anxiety, and depression using a random forest model; (2) a rule-based question recommender that suggests targeted self-assessment questions based on detected needs; (3) a rule-based chatbot, using NLTK for tokenization, that provides interactive support based on keyword matching; and (4) a comprehensive evaluation of model performance (accuracy, precision, recall, F1) with up-to- date metrics for multi-label and multi-class classification. We emphasize ethical deployment, including privacy safeguards and model explainability, to engender user trust. Solace's novelty lies in its holistic integration of multiple AI components – from screening algorithms to conversational agents – within a single platform, enabling early detection and personalized support outside of clinical settings

II. Related Work

Digital mental health has become a major area of academic investigation and commercial development. Several mobile and web programs like Companion, Wysa, Ginger, and Woebot have also tried various methods to assist users' emotional well- being. These systems generally employ formal self-assessment questionnaires, content library repositories, or mindfulness aids. Although such systems provide convenience and accessibility, most lack personalization and decision support integration.

Over the past few years, machine learning (ML) methods have been widely used to identify and assess mental illnesses like depression, anxiety, and stress from structured data, such as self-reported symptoms, digital traces of behavior, and electronic health records. Simple models like logistic regression and decision trees work well; however, more complex ensemble techniques like Random Forests have proved to be better because they can identify non-linear relationships and can work effectively with high-dimensional data.

Conversational AI or chatbots have also picked up pace in mental health treatments as a non-judgmental, always-on companion with whom users feel at ease conversing. Research has shown that these systems can promote self-disclosure and introspection, particularly when users believe that the responses are empathetic and context-sensitive. Rule-based and artificial intelligence-driven chatbots (e.g., Wysa, Woebot) have been found to be quite promising in studies for reducing symptoms of depression and anxiety. Nonetheless, there are difficulties in maintaining contextual relevance, handling crises, and constructing personalized replies over time.

Content recommendation within mental health platforms has also been promising, especially when it comes to personalizing resources like self-help articles, videos, or exercises based on user needs. Rule-based recommenders provide transparency and expert-compliant decision-making, which is desirable in clinical or sensitive situations. However, most platforms to date actually address either screening or guidance infrequently combining the two seamlessly in an easy-to-use workflow.

Solace is unique in that it fills in these gaps. It provides a single web-based platform that integrates an ML-assisted mental health assessment module, a rule-based question recommender that responds to user issues, and a slim NLP-based chatbot for emotional support and follow-up conversation. By combining assessment, personalization, and real-time interaction in one system, Solace offers a holistic and integrated solution for early detection and support of mental health.

III. System Design

Solace has been designed as a modular, web-based platform that offers mental health assessments and guidance in a user- friendly environment. The system comprises two main layers: the user-facing interface and the AI-powered back-end, both working in collaboration to deliver meaningful and safe mental health support.

A. Classification Engine: Personalized Mental Health Detection

At the heart of Solace lies a machine learning-based classification engine responsible for analyzing user responses to determine whether the user may be experiencing stress, anxiety, or depression, and if so, how severe it might be.

To accomplish this, we use Random Forest classifiers, which are well-suited for handling large feature sets and capturing complex relationships within data. Instead of relying on a single decision tree, the Random Forest technique builds an ensemble of trees and aggregates their predictions. This reduces overfitting and provides more reliable results across varied user inputs.

- **Multi-label Detection:** The model first predicts which of the three conditions (stress, anxiety, depression) are likely present. Since users can suffer from multiple conditions simultaneously, this task is handled as a multi-label classification problem.
- **Severity Classification:** For each condition detected, the model then determines the severity level, typically categorized as None, Mild, Moderate, or Severe. This part of the sys- tem functions as a multi-class classifier for each individual condition.
- **Interpretability:** Although Random Forests are often seen as ”black box” models, Solace extracts feature importance scores to enhance explainability. This helps in identifying which questions influenced a prediction the most. For example, the system may highlight that “Persistent restlessness” was a key factor in predicting moderate anxiety.

This transparency allows users to better understand their results and enables clinicians (if involved) to interpret model decisions more confidently.

B.Question Recommender: Tailored Exploration of Mental States

After the initial assessment, Solace doesn't stop at just reporting scores. It engages users with a customized set of follow-up questions based on their mental health profile. This is managed by a rule-based recommendation system designed with help from psychological literature and expert consultation. How it Works: Based on the detected conditions and their severity levels, the system triggers specific question sets. For example, users showing high stress may be asked about sleep quality, recent workload, or coping habits.

Why Rule-Based: Unlike machine learning models that require large datasets to learn patterns, rule-based logic allows for safe, clear, and medically aligned branching paths, making it ideal for sensitive domains like mental health. This structured questioning ensures users feel heard and allows deeper self-reflection without overwhelming them with irrelevant queries.

C. Chatbot Engine: Empathetic Interaction and Early Guidance

Solace integrates a conversational chatbot that acts as a friendly companion during a user's interaction. While not a substitute for therapy, the chatbot offers emotionally intelligent responses, psychoeducation, and redirection to appropriate resources when needed.

- **Functionality:** The chatbot uses a keyword-driven pattern matcher built on Python's NLTK (Natural Language Toolkit). It tokenizes user inputs and looks for relevant mental health terms like "anxious," "tired," or "overwhelmed." Matched phrases are linked to pre-written supportive messages.
- **Tone and Content:** Responses are written in an affirming, human tone including mindfulness tips, motivational prompts, or links to professional help. Importantly, it has built-in flags for crisis-level keywords like "suicide" or "hopeless," which prompt it to immediately display helpline information and encourage the user to seek urgent care.
- **Safety Net:** This chatbot is not designed to diagnose or offer deep therapeutic conversations. Instead, it offers gentle engagement and encourages professional help when appropriate.

D. System Architecture Overview

Solace uses a modern web development stack built for scalability, responsiveness, and security:

- Frontend: The user interface is crafted using Django templates with Bootstrap and SCSS for styling. This ensures that users, regardless of technical literacy, can navigate the platform easily.
- Backend: The core AI functionalities are built as microservices in Python. These include:
  - Scikit-learn for the Random Forest classification models.
  - A custom-built rule engine for personalized question recommendations.
  - NLTK for chatbot response generation based on keyword matching.
- Security & Privacy: The system adheres to best practices for user privacy, ensuring data is only used for its intended purpose providing mental health insights to the individual user.

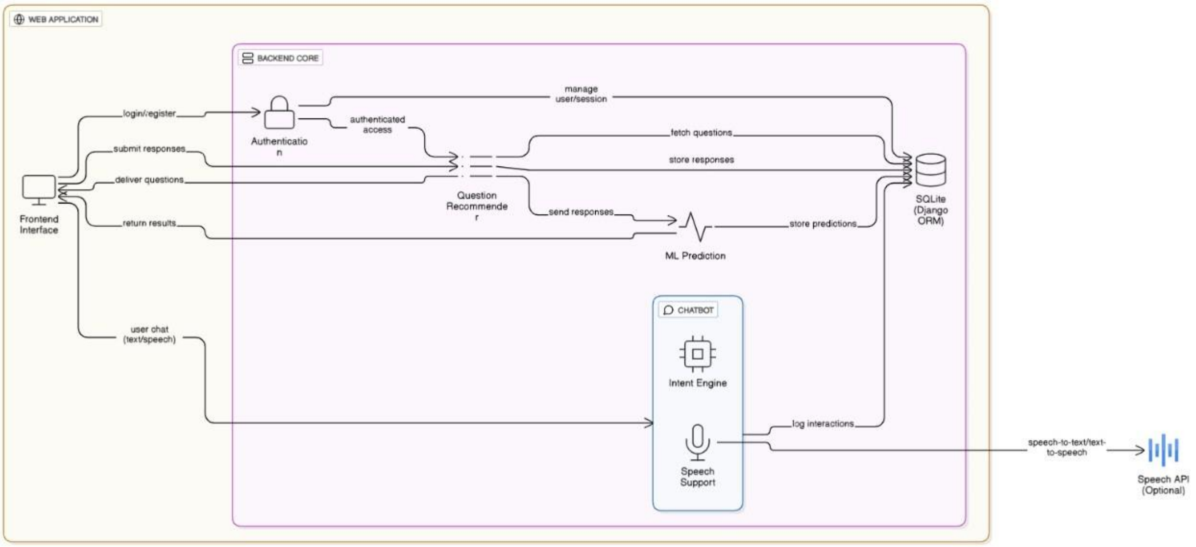

Fig. 1. system architecture of solace, illustrating the front-end, back-end ai modules (classification engine, question recommender, chatbot engine), and supporting components like data storage and privacy management.

IV. Methodology

A. Data Generation and Preparation

To design a safe and scalable mental health prediction system without compromising user privacy, we chose to simulate a custom dataset instead of using real-world clinical data. This synthetic dataset, consisting of 5,000 virtual participants, was generated using a structured approach that mirrors patterns typically observed in actual mental health assessments.

We first defined a pool of mental health indicators and assigned each simulated user a specific profile for example, high stress but no depression, or moderate anxiety combined with mild depressive symptoms. Based on the assigned profile, responses to 20-question mental health assessments were algorithmically simulated. These responses were not randomly generated but were crafted using statistical probability distributions. For instance, a user with severe anxiety would have a higher chance of responding" Always" to questions related to restlessness, insomnia, or excessive worry.

Each participant's overall responses were then translated into two sets of labels:

- Presence Labels (binary): Indicating whether the user shows any signs of stress, anxiety, or depression.
- Severity Labels (multi-class): Indicating the severity level (None, Mild, Moderate, or Severe) for each condition.

To reflect real-world diversity, we deliberately introduced some imbalance in the dataset. For example, fewer users were modeled to have severe depression compared to mild stress, to replicate the natural frequency distribution of these conditions in actual populations. Before the data was used for model training, it was cleaned and standardized. Missing values were handled gracefully, and features were normalized. Stratified sampling was applied during training so that each severity class was appropriately represented across splits.

B.Model Development and Training

The core predictive engine of Solace is composed of three independently trained machine learning models one for each condition: Stress, Anxiety, and Depression. All three models are based on the Random Forest Classifier, chosen for its resilience to overfitting, ability to handle high-dimensional data, and robust performance across imbalanced datasets.

Each model was designed to handle two tasks:

- Condition Detection (Binary Classification): Determining whether the user is likely experiencing the condition.
- Severity Prediction (Multi-Class Classification): Determining the severity level for users with a positive screening result.

Hyperparameters such as the number of decision trees, maximum depth of each tree, and class weights were fine- tuned for optimal performance. To ensure the robustness of our models, we used 5-fold cross-validation. This method allowed us to assess model consistency by rotating training and validation sets across five different data splits.

It's worth noting that although Logistic Regression was considered in early experimentation, the final deployed version of Solace exclusively relies on Random Forests due to their superior accuracy and ability to capture non-linear patterns in mental health data.

C. Evaluation Strategy

Given that users may exhibit any combination of the three mental health conditions or none at all our evaluation focused on per-condition performance.

For binary classification tasks (presence detection), we computed the following metrics separately for each condition:

- Accuracy - the percentage of correct predictions.
- Precision -the proportion of true positives among all predicted positives.
- Recall - the proportion of true positives among all actual positives.
- F1-Score -the harmonic mean of precision and recall, capturing a balance between them.

For multi-class severity classification, we calculated macro- averaged metrics. This means we computed each metric for Mild, Moderate, and Severe categories independently, and then averaged the results treating all classes with equal weight regardless of frequency.

This granular evaluation framework helped us identify specific strengths and weaknesses in the model's ability to recognize each condition and severity level.

D. Evaluation of Chatbot and Recommender Logic

While the predictive models operate under statistical learning principles, other components of Solace specifically the chatbot and the recommendation engine follow a rule-based architecture.

The chatbot interacts with users via keyword recognition. We evaluated its performance by crafting a variety of typical user inputs and analyzing the chatbot's responses. This evaluation focused on:

- Accuracy of keyword recognition
- Relevance of selected responses
- Tone and emotional appropriateness

Special emphasis was placed on the chatbot's ability to flag high-risk inputs, such as expressions of suicidal ideation or distress. In such cases, the chatbot is programmed to bypass scripted responses and immediately present mental health hotline links or emergency resources.

The recommendation engine was similarly tested. Domain experts assessed its ability to select follow-up questions based on previous model outputs. For example, if a user showed signs of moderate depression, the engine might ask about social withdrawal or sleep disruptions. These rules were checked to ensure logical consistency and ethical appropriateness, particularly in avoiding distressing or triggering content.

TABLE I

TABLE I: classification performance (random forest, macro-averaged)

Condition	Accuracy	Precision	Recall	F1-score
Stress	0.70	0.72	0.68	0.70
Anxiety	0.72	0.70	0.73	0.71
Depression	0.68	0.65	0.70	0.67

V. Evaluations And Results

We now present the performance of Solace's classification model. Table I shows the classification results from the random forest model for stress, anxiety, and depression. All values are averaged over cross-validation folds.

The random forest classifiers achieved ≈68-72% accuracy, with macro-F1 scores around 0.67-0.71. Severity classification is an inherently difficult multi-class problem, but the RF models demonstrate reasonable discriminative ability, particularly for anxiety (72% accuracy). The metrics suggest the models can broadly differentiate between severity levels, enabling Solace to tailor its recommendations effectively. By analyzing feature importance, we can identify which questionnaire items (e.g., sleep disturbance, appetite change) were most predictive, providing a degree of explainability for the model's decisions. These performance metrics demonstrate that Solace's ML components are effective for automated mental health screening. Note: Solace outputs separate severity levels for each condition Stress, Anxiety, and Depression rather than a combined score.

VI. Discussions

A. Model Performance and Novelty

Solace's integrated approach yields promising screening capabilities. The use of a random forest classifier for both multi- label (presence) and multi-class (severity) assessment allows for nuanced user evaluation. The achieved accuracies (Table I) are encouraging for a complex task and reflect the quality of our synthetic data and feature design. We emphasize that these metrics represent automated predictions; final recommendations are always accompanied by human-readable justifications, addressing calls for transparency in AI diagnostics.

B.Limitations and Future Work

While encouraging, our evaluation has limitations. Our results are based on a synthetic dataset; while large and designed for realism, it may not fully capture the nuances of real-world clinical data. Future work will involve validating the models on real-world patient data to confirm generalizability. The chatbot, currently based on keyword matching, will be extended with deeper NLP (e.g., transformer models for true intent classification) to handle a wider range of user expressions. We also plan formal user studies and clinical trials to assess Solace's impact on outcomes.

C. Privacy and Ethical Considerations

Solace was designed with privacy and ethics as primary concerns. All personal inputs are treated as protected health information. We avoid "black box" AI by using feature importance analysis for model auditing. We also address bias: the training data were checked for demographic skew, and our models employ stratified sampling to mitigate biased outcomes. Nonetheless, algorithmic bias remains an open issue that requires ongoing monitoring.

VII. Conclusion

We have launched Solace, an integrated AI-driven web platform that aims to improve mental wellbeing with intelligent testing, tailored advice, and user interaction. The platform uses a strong Random Forest-based machine learning structure to identify the presence and measure the severity of the primary psychological conditions, namely stress, anxiety, and depression. Differently from the usual scoring system, Solace offers condition-specific severity levels to provide more detailed and actionable mental health information.

Aside from its predictive functionality, Solace also has a rule- based question recommender function that selects assessment questions by symptom domains and user requirements. This guarantees a rich and clinically pertinent question set, enhanc- ing screening quality. In addition to this, the platform also has a light rule-based chatbot that utilizes simple natural language processing (NLP) methods like tokenization and stemming (through NLTK) to mimic empathetic conversation and provide relevant coping strategies. The chatbot is a judgment-free interactive interface that is available 24/7, providing a more exciting and friendly experience for users.

VIII. Acknowledgement

We acknowledge the mental health professionals who provided input to the design of Solace's questionnaires and content. We take this opportunity to express our heartfelt thanks to our project guide, Prof. Devendra Bharambe, for their consistent support and valuable guidance during the development of Solace. We are grateful to the Head of the Computer Engineering Department, Prof. Rama Gaikwad, and to all the teaching and non-teaching faculties of Anantrao Pawar College of Engineering and Research for their support and cooperation.

We also thank our Principal, Dr. Sunil B. Thakre, for providing essential infrastructure and academic support. Our appreciation goes to our team members for their collaboration, and to all individuals who contributed indirectly through discussions and feedback. Their input greatly enriched the quality of our work.

References

[1] R. Razavi et al., "Machine Learning, Deep Learning, and Data Preprocess- ing Techniques for Detecting, Predicting, and Monitoring Stress and Stress- Related Mental Disorders: Scoping Review," JMIR Mental Health, vol. 11, no. 1, 2024. [Online]. Available: https://mental.jmir.org/2024/1/e49069

[2] M. D. R. Haque *et al.*, "An Overview of Chatbot-Based Mobile Mental Health Apps: Insights From App Description and User Reviews," *JMIR mHealth uHealth*, May 2023. [Online]. Available: https://mhealth.jmir.org/ 2023/5/e45989

[3] A. Chaturvedi *et al.*, "Content Recommendation Systems in Web-Based Mental Health Care: Real-world Application and Formative Evalua- tion," *JMIR Formative Res.*, vol. 7, Jan. 2023. [Online]. Available: https://formative.jmir.org/2023/1/e42609

[4] W. Lim Ku and H. Min, "Evaluating Machine Learning Stability in Predicting Depression and Anxiety Amidst Subjective Response Errors," *Int. J. Environ. Res. Public Health*, vol. 21, no. 3, Mar. 2024. [Online]. Available: https://www.mdpi.com/1660-4601/21/3/yourarticlenumber

[5] U. Warrier *et al.*, "Ethical considerations in the use of artificial intelligence in mental health," *Egypt J. Neurol. Psychiatr. Neurosurg.*, vol. 59, Article 139, 2023. [Online]. Available: https://ejnpn.springeropen.com/articles/10. 1186/s41983-023-00745-y

[6] P. Kumar *et al.*, "A Machine Learning Implementation for Mental Health Care: Smart Watch for Depression Detection," in *Proc. 11th IEEE Int. Conf. Cloud Comput., Data Sci. Eng. (Confluence)*, 2021, pp. 568–574.

[7] R. Kovacevic, "Mental Health: Lessons Learned in 2020 for 2021 and Forward," *World Bank Blogs*, 2021. [Online]. Available: https://blogs.worldbank.org/health/ mental-health-lessons-learned-2020-2021-and-forward

[8] O. O. Lim *et al.*, "Companion: Mental Health Mobile Applications for Students," in *2022 1st*

International Conference on Technology Innovation and Its Applications (ICTIIA), 2022, pp. 1-6, doi: 10.1109/IC- TIIA54654.2022.9935882.

[9] G. Parimala, R. Kayalvizhi, and S. Nithiya, "Mental Health: Detection & Diagnosis," in *2022 International Conference on Computer Com- munication and Informatics (ICCCI)*, 2022, pp. 1-6, doi: 10.1109/IC- CCI54379.2022.9740834.

[10] O. Higgins, B. L. Short, S. K. Chalup, and R. L. Wilson, "Artificial intelligence (AI) and machine learning (ML) based decision support systems in mental health: An integrative review," *Int. J. Ment. Health Nurs.*, vol. 32, pp. 966-978, 2023, doi: 10.1111/inm.13114.

[11] J. Torous, J. Luo, and S. R. Chan, "Mental health apps: What to tell patients," *Current Psychiatry*, vol. 17, no. 3, pp. 21-25, Mar. 2018.