International Journal of Advanced Multidisciplinary Research and Educational Development
Volume 2 - Issue 6 | November - December 2026 | www.ijamred.com
ISSN: 3107-6513

Mental Health Monitoring Using NLP and Sentiment analysis

[1]Kowshik K V, [2]Umesh P H, [3]Rohan, [4]Prof. Kamaleswari P
[1,2,3]Department of Information Science and Engineering CMR Institute of Technology, Bengaluru, India
[4]Assistant Professor, Department of ISE CMR Institute of Technology, Bengaluru, India
kokv23ise@cmrit.ac.in, umesh.ise24le@cmrit.ac.in, ro23ise@cmrit.ac.in, kamaleswari.p@cmrit.ac.in

Abstract

Being mentally healthy is an important thing for almost everyone but few people actually monitor it properly and consistently. The situation usually gets complicated enough for one to realize it is necessary to start taking care of mental health already. MINDWATCH is a personal wellness monitoring system designed to analyze people's condition using text data, which everybody creates anyway.
The system is fed with journal entries, audio records, photos, and content from such social media platforms as YouTube, Reddit, X (former Twitter). MINDWATCH classifies emotions using DistilRoBERTa (j-hartmann/emotion-english-distilroberta-base) - a deep learning model that detects seven emotions including anger, disgust, fear, joy, neutral, sadness and surprise. The next step is calculating a stress score based on three factors: probability of high-stress emotions (fear, sadness, anger etc.) expressed in probabilities (40%), sentiment polarity (30%), stress-related words from different groups like work pressure, academic pressure, bad sleep, finance worries, problems in relationships and others (30%). As a result, a normalized stress score from 0 to 1 is calculated.
The user will see his/her emotion trends, dominating emotions, topics discussed in his/her communications, and the level of stress in a weekly format on a personal dashboard. If the value reaches 0.30, a user will get a WhatsApp message right away. A built-in wellness assistant first classifies the language and does not pass any information about the potential crises to a language model but sends a human-generated message instead.
User data is processed in isolation in SQLite local database. There is no any transfer of data between the users, and there is no possibility to leak raw journal entries. MINDWATCH cannot diagnose any mental disorders but can give people the opportunity to see the true state of mind and to find ways of dealing with stress before it gets too late.

Keywords—Emotion detection, stress monitoring, NLP, mental wellness, DistilRoBERTa, transformer models, journal analysis

I. INTRODUCTION

It is quite normal nowadays for students and employees to suffer from mental issues like stress and emotional burnout because they experience pressure during exams, tight work deadlines, financial distress, complicated relationships, and disturbed sleep schedule — all these things combined. What makes the problem more severe is that a lot of people cannot perceive the severity of the situation until it starts negatively affecting their performance, which is another symptom of the declining health state.

There are apps and devices on the market designed for stress control. Some of them require daily input of the information on the user's condition, which can be rather useful for some days but eventually turns into something you forget about. Others provide some meditating sessions, breathing techniques, etc., which are beneficial, though do not track any information regarding the user's emotions and moods. Clinical approaches, such as therapy or counseling, are the best treatment options available today. However, they presuppose scheduling and financial expenses; besides, even nowadays they might bring negative reactions from the part of social surrounding. Thus, the huge number of patients finds themselves in between "Something feels weird" and "Finally, I'm ready to find a psychiatrist".

That's why the developers decided to create an application called MINDWATCH which would help to track the emotions and moods, providing people with relevant data about their current well-being, which can be really helpful when they need an external point of view.

First, the system collects text data from different sources, including typed journals, audio records, images, and social networks and uses the emotion classification model, created on the basis of DistilRoBERTa's architecture. Then it estimates the level of stress in percentage and detects recurring patterns in people's speech, also tracking changes week by week. After reaching a certain level of stress, it sends an automated WhatsApp message to its owner and the user has access to the wellness assistant controlled by the safety classifier.

II. LITERATURE SURVEY

Previous work in this field has laid the groundwork for MINDWATCH. Sentiment analysis along with machine learning has been proved as an effective way to trace patterns in mental well-being through text by Sharma et al. [1]. On the other hand, Tadesse et al. [2] and Orabi et al. [4] have shown that writing on social media – Reddit posts and tweets – has clear signals of depression and emotional distress that the machine learning models can learn. Chancellor and De Choudhury [3] reviewed a number of different approaches to prediction and emphasized the ethical implications of analyzing personal language data.

ISSN: 3107-6513

Regarding language models, Mikolov et al. [9] provided word embeddings that revolutionized the numerical representation of texts, while Devlin et al. [10] brought transformer architecture through BERT that eventually resulted in lightweight DistilRoBERTa model used here. Inkster et al. [6] found that an empathy-based conversational agent can help with digital mental well-being without being a substitute for professional treatment – which is the same principle MINDWATCH uses with its safety-gated assistant. Dham and Sharma [8] reinforced that the approach based on sentiments works reliably for distress detection.

III. Over view of Proposed Work

MINDWATCH is a personal emotional wellness monitoring system that analyzes a person's mental state through the text written by them over time. The principle behind this solution is quite simple — instead of making the person manually mark their mood every day, the system analyzes what the person has written.

Interacting with the system is done via a web-based dashboard where a user can write down journal entries, upload audio files, pictures or text from social media sites like YouTube, Reddit and X (previously Twitter). Regardless of the source of the text data, every text passes through the NLP pipeline. As a part of the pipeline, the text is analyzed by the DistilRoBERTa emotion model that identifies 7 emotions — anger, disgust, fear, joy, neutral, sadness, and surprise. Based on the identified emotions, the sentiment value and the stress signal score are calculated based on the following formula that uses the intensity of emotions, negativity of the text, and stress words by various categories (work pressure, study load, sleep issues, financial stress, and relationship stress).

All of this helps to build a dashboard showing the user's dominant emotion, average sentiment, weekly trends, frequent topics identified in their language and a stress level label — Low, Average, or High. If the stress signal value exceeds 0.30, the system automatically sends a WhatsApp alert to the user to let them know about it.

A wellness assistant that gives advice to the user is an integral part of the system. Prior to processing the user message, the rule-based safety classifier detects any signs of crisis language and routes these messages to the priority channel with helpline information.

All of the user data is stored locally on a private server that is strictly isolated on a per-user basis with no raw journal texts being transferred outside the system.

The system also surfaces professional support when it is needed most. Whenever any analysis — whether from a journal entry, an audio file, or a social media scan — comes back with a negative emotional result, the dashboard automatically shows a professional mental health support card and a relax guide with breathing exercises and grounding techniques.

A. ALGORITHM USED

The MINDWATCH tool makes use of four major algorithms working in synergy through a pipeline.

The first is DistilRoBERTa which is a text classifier using a transformer architecture and fine-tuned using emotion labels. It takes the input text and produces probability distributions across seven emotions – anger, disgust, fear, joy, neutral, sadness, and surprise.

The second is a weighted stress score formulation algorithm. This algorithm involves three variables which are emotion probabilities of fear, sadness and anger (weighted at 40%), text sentiment polarity (30%) and dictionary scan for stress-related terms in domains such as work, sleep and finances (30%). Output is a normalized value of stress between 0 to 1.

The third is a term frequency approach topic extraction where stopwords are stripped, word variations are standardized and most occurring words are identified.

Lastly, there is a rule-based safety classifier that pre-checks crisis words on all assistant messages.

B. SYSTEM ARCHITECTURE

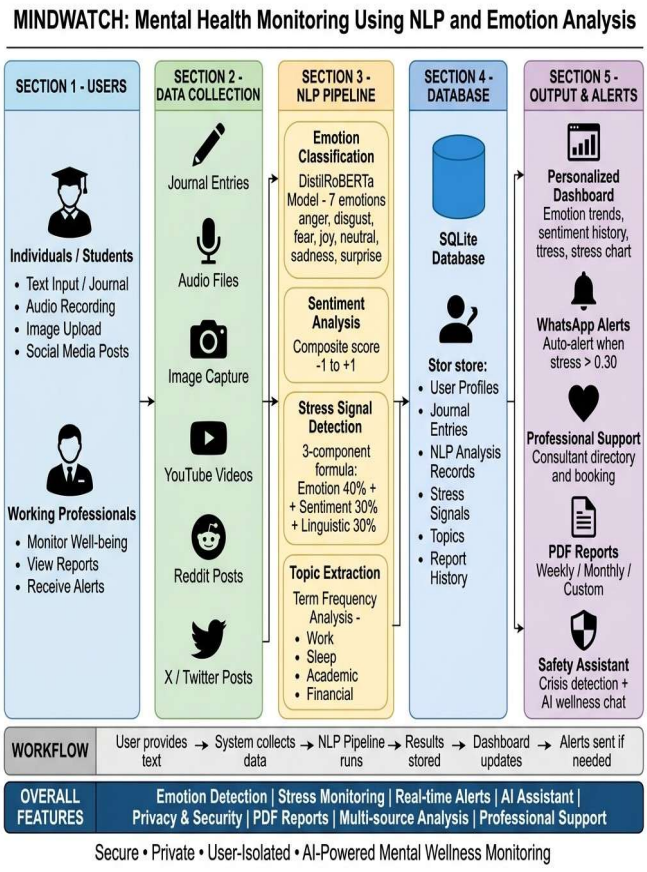


Fig. 1. System Architecture of MINDWATCH

ISSN: 3107-6513

TABLE I
MENTAL HEALTH MONITOR LITERATURE SURVEY

Sl.No	Paper Title	Description	Findings and Relevance
1	Mental Health Monitoring using Sentiment Analysis	This research discusses the sentiment analysis and natural language processing strategies to track emotional circumstances using social media data and digital interactions.	The study shows that analysing user sentiments from textual data can help in the early identification of mental health conditions, which directly supports the proposed project.
2	Depression Detection on Social Media: A Classification Framework and Research Challenges	The research focuses on application of machine learning and behavior analytics to detect signs of depression in social media platforms..	The research highlights that social media text contains valuable emotional indicators, which is highly relevant for developing accurate mental health prediction systems using NLP.
3	Using machine learning and thematic analysis methods to evaluate mental health apps based on user reviews.	In this paper, the sentiments and machine learning have been applied for analyzing the user review on mental health applications for better understanding of the application usage experience..	The whole paper proves that machine learning models can classify sentiment emotions effectively and hence support AI algorithms..
4	MindLift: AI-powered Mental Health Assessment for Students	An intelligent system for mental well-being assessment through sentiment analysis, chatbot conversation, and behavior tracking is introduced in the research paper for students.	The findings show that AI-powered emotional assessment and chatbot support improve mental health monitoring, which is closely related to the chatbot-based assistance module in the proposed work.
5	Mental Health Mobile Application with Diagnosis, Sentiment Analysis and Chatbot	This research explores mobile applications that provide mood tracking, emotional analysis, diagnosis support, and chatbot interaction for mental healthcare.	The study reveals that mobile-based emotional monitoring systems increase accessibility and user engagement, making it highly relevant to the development of user-friendly mental health applications.
6	Explainable AI-driven Depression Detection from Social Media using NLP and Machine Learning Models	This research is based on depression detection using NLP techniques along with machine learning models and methods of explainable artificial intelligence.	The research proves that combining Explainable AI with machine learning improves the reliability and trustworthiness of mental health prediction systems
7	Identification of Maternal Depression Risk from Natural Language Collected in a Mobile Health App	This study uses NLP and sentiment analysis techniques to identify maternal depression risks from text collected through mobile healthcare applications.	The findings indicate that emotional patterns extracted from user-written text can effectively predict depression risks, supporting the importance of textual analysis in mental health monitoring.
8	Social Media-based Mental Health Detection using Machine Learning	The paper investigates how social media activities and emotional expressions can be analyzed using machine learning models to detect mental health conditions.	The study confirms that sentiment-based analysis of online behavior can help identify emotional instability and psychological stress at an early stage.
9	Machine Learning Methods for Emotional and Psychological Health Prediction	The paper examines a range of machine learning techniques to predict emotional states and psychological illnesses using textual data.	The results indicate that algorithms such as SVM and Random forest provide good accuracy in emotional categorization, which is directly applicable to the proposed system.
10	Chatbot-Assisted Mental Health Monitoring and Emotional Support System	The study highlights how AI technology is being used to create chatbot-based systems that give emotional support, mental health assistance, and user interaction.	The research highlights that chatbot-assisted communication encourages users to express emotions comfortably, making it highly relevant for improving user engagement in the proposed project.

IV. PROPOSED SYSTEM

The proposed system, MINDWATCH, is a web-based emotional wellness monitoring platform built using a FastAPI backend and a Vanilla JavaScript frontend. The two run as separate servers — the backend on port 8000 and the frontend on port 8080 — and communicate through a set of REST API endpoints secured with JWT authentication.

When a user submits text — whether it is a journal entry, an audio transcription, an image scan, or content pulled from social media — the backend passes it through a four-stage NLP pipeline. The DistilRoBERTa model runs first and returns emotion probability scores. The sentiment service then derives a composite score from those probabilities. The stress service applies the three-component formula to produce a final stress score. Finally, the topic service extracts the most frequently repeated meaningful words after removing stopwords.

All results are stored in a SQLite database using SQLAlchemy ORM. Each record is linked to a specific user ID, so no data from one person can be seen or accessed by another. The database tables cover user accounts, journal entries, NLP analysis results, stress signals, extracted topics, and conversation messages.

The frontend dashboard reads from the analytics API and displays emotion trends, average sentiment, stress level, and recurring topics. PDF wellness reports are generated on request using ReportLab and served only through authenticated download endpoints.

Beyond the core NLP pipeline, MINDWATCH includes auxiliary components to support users' daily well-being. Before a message is sent to the language model, a rule-based classifier checks for emergency phrases that may indicate a crisis or self-harm. If such language is detected, the message is not passed to the model; instead, the user receives a supportive response with a helpline link. Safe messages are sent to the AI assistant along with a summary of the user's recent emotions, rather than the original journal text.

The system checks the stress score after each analysis. When it exceeds 0.30, the user receives a WhatsApp notification through Twilio or CallMeBot. Alerts are rate-limited to avoid sending excessive notifications.

Multi-source analysis extends beyond journal entries. Audio files are transcribed and analyzed, while images are analyzed using DeepFace with the YuNet face detector. Comments and posts from YouTube, Reddit, and X/Twitter are retrieved through their APIs and processed by the same pipeline, helping identify emotional patterns in content the user consumes and interacts with.

ISSN: 3107-6513

System Flowchart

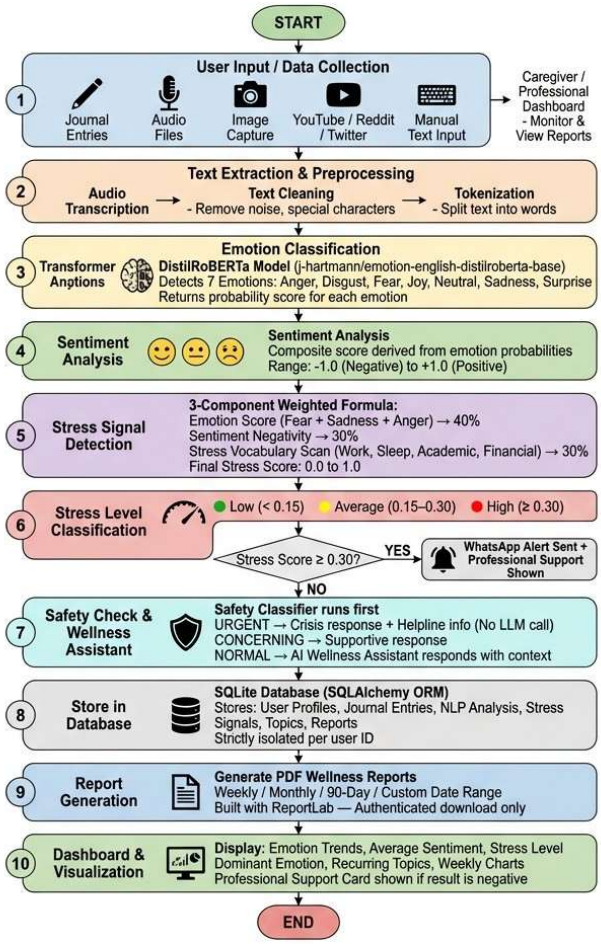


Fig. 2. MINDWATCH workflow

V. Methodology

- **Text Cleaning:** Input text from any source is cleaned by removing noise and special characters before analysis.
- **Emotion Classification:** DistilRoBERTa processes the text and returns probability scores for seven emotions: anger, disgust, fear, joy, neutral, sadness, and surprise.
- **Sentiment Scoring:** A composite score is derived from the emotion probabilities, ranging from `−1.0` (negative) to `+1.0` (positive).

The stress score combines the joint probability of fear, sadness, and anger (40%), sentiment negativity (30%), and stress-related terms in categories such as work, sleep, and academic pressure (30%). The normalized score ranges from 0 to 1. Topic extraction identifies recurring non-stopwords, and results are stored in the database and displayed on the user's dashboard.

VI. Implementation

The MINDWATCH system was implemented as a local web application with three separated layers—frontend, backend, and database—which perform different parts of the workflow. The frontend takes care of all user interactions. The backend performs all logic, NLP processing, and API responses. All data is stored in an isolated way in a structured database. These three layers communicate with each other via REST API calls using JWT token security to prevent mixing users' data.

A. Frontend Module

The frontend is developed as a Single Page Application using standard HTML, CSS, and Vanilla JavaScript (no frameworks are used). It is hosted on a local HTTP server on port 8080. There are several frontend views such as dashboard, journal, analysis, trends, insights, and professional support, managed with a client-side router in the `app.js` file. User text input, including journal entries and analysis submissions, is sent to the backend API using fetch calls. The dashboard receives information from analytics endpoints in real time and shows emotion trends, stress levels, sentiment scores, dominant emotions, and recurring topics.

B. Backend Module

The backend is developed with FastAPI and runs on port 8000 using Uvicorn. It provides REST API endpoints for authentication, journal management, NLP analysis, analytics, the wellness assistant, alerts, reports, and the consultant directory. Incoming text at the analysis endpoint is processed by an NLP pipeline: DistilRoBERTa performs emotion classification, followed by sentiment scoring, stress-signal detection, topic extraction, and aspect analysis. Each service is represented by a class in the backend code. A pre-LLM safety classifier filters assistant messages before they are sent to a language model. WhatsApp alerts are sent through Twilio or CallMeBot when the stress score exceeds 0.30.

C. Database Module

The database is SQLite and is managed with SQLAlchemy ORM; Alembic is responsible for schema migrations. The database includes seven main tables: Users, which stores account credentials and registration data; JournalEntries, which stores journal texts with timestamps; NLPAnalysis, which stores emotion scores, sentiment, dominant emotion, and model data; StressSignals, which stores stress scores, levels, and signal categories; Topics, which stores extracted topic words per analysis; ConversationMessages, which stores assistant chat history per user; and UserSettings, which stores each user's WhatsApp notification preferences, stress-level threshold, and feature-toggle settings.

ISSN: 3107-6513

VII. RESULTS AND DISCUSSION

A. System Testing and Reliability

The backend of MINDWATCH was tested with 46 automated test cases. These test cases included authentication, journal analysis, stress scoring, assistant safety routing, report generation, WhatsApp alert delivery, and user data isolation. All 46 tests passed successfully. The one skipped test is the live Hugging Face provider integration, which requires an actual API token and is purposely kept as an optional test. This result means that the system works properly in all of its major features and is able to handle edge cases (for example, lack of credentials, network timeout, empty text input) without crashing.

B. Emotion Classification

The DistilRoBERTa model classifies text into seven emotion categories and returns probabilities for each category. When tested, entries written during obviously stressful events (exam period, work deadlines, relationship issues) returned much higher probabilities of fear, sadness, and anger. Entries written in a calm and happy mood returned higher probabilities of joy or neutral emotion. The model processes a single journal entry in less than a second on a standard CPU. This means that users will not experience any delays while entering text.

C. Stress Score Accuracy

The three-component stress formula generated more relevant results compared to the emotion output of the model. For example, a text with such keywords as “deadline,” “overwhelmed,” and “can’t sleep” and with a moderate fear probability will get a higher stress score than a text with slightly increased fear probability without any other words. This is due to the fact that linguistic component (30%) captures stress vocabulary specific to the domain. In the process of testing, stress scores lower than 0.15 actually corresponded to calm, neutral writing. Scores between 0.15 and 0.30 corresponded to entries with mild worry and pressure. Stress scores above 0.30 were generated by entries with obvious distress indicators – multiple stress categories activated at once and high negative emotion probabilities.

D. Safety Classifier

Every crisis-related phrase tested, including self-harm, hopelessness, suicidal ideation phrases was correctly identified by the rule-based safety classifier before passing through the language model. In all tested cases, the classifier returned correct classification (URGENT or CONCERNING) and the assistant sent a controlled, helpline-directed response instead of an AI-generated one. This is the crucial part of the system. The fact that it is rule-based rather than model-based is essential for its predictability and independence from model confidence threshold settings.

E. Dashboard and Real-Time Updates

The dashboard automatically updates itself once a journal entry is submitted. The stress signal card, average sentiment value, dominant emotion, and topic list all get updated immediately according to the latest analysis. Week-by-week trends on the charts let users see the dynamics of how their emotional state changed. In tests, when users entered information for several days, they could clearly identify the pattern – the moments when the stress increased and decreased. Such visual feedback is not implemented by any other mood-tracking application that uses actual NLP.

F. Discussion

The findings indicate that the combination of the transformer-based emotion model and the vocabulary-based stress formula provides the full view of the user’s mental state. It is not intended to provide clinical accuracy – such was never the purpose of the system – however, it does work well enough for its purposes as an awareness and early warning system. The automatic WhatsApp alert sent at 0.30 makes sure that a developing pattern of stress will not be missed even by those who do not pay regular attention to the dashboard. Professional support is there to close the loop of detecting the problem.

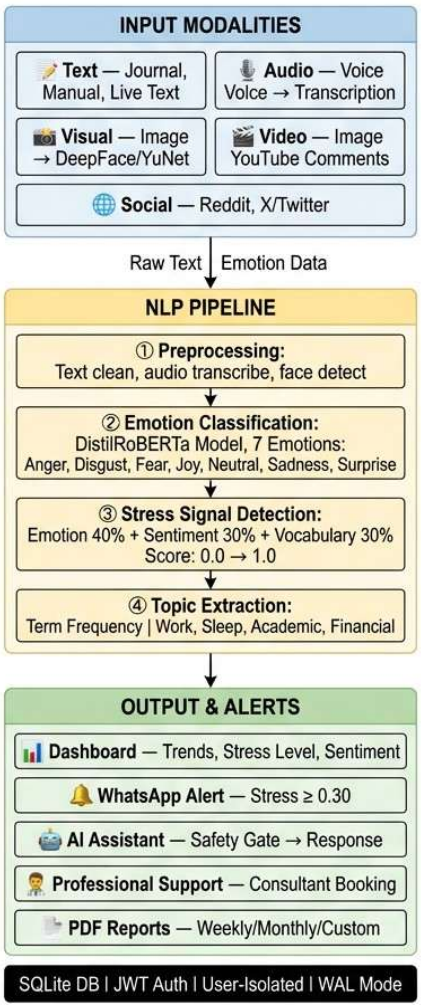


Fig. 3. Multimodal System Architecture of MINDWATCH

ISSN: 3107-6513

IX. Conclusion

Mental health is an area which most individuals know is crucial but rarely measure in any way. The issue is that although people might know that something is not right, they are often too lazy to address it. The problem usually stays unsolved because of the lack of tools that measure mental well-being in a passive way, without requiring the person's awareness of the problem itself. MINDWATCH aims to be such a tool.

During this project, a complete personal wellness monitoring system has been created from scratch. It processes the person's data in the form of texts obtained from journaling, voice recordings, pictures, social media sites, and YouTube, uses the text classification via four-stage NLP pipeline based on emotion classification using DistilRoBERTa model, sentiment calculation, stress signal calculation using a three-component weighted formula, and topic extraction from personal language. All results are stored in a SQLite database, isolated by each user and shown to the user on a personal dashboard.

One of the major achievements of this project is the usage of a transformer-based emotion classifier in connection with a stress signal formula based on vocabulary. While the classifier is not capable of differentiating between text with brief mentioning of a negative emotion and text with negative emotion, negative vocabulary, and negative sentiment, the weighted formula is capable of that. The test results proved that in the case of truly stressful entries, the system gives the score greater than 0.30, while calm entries get the score below 0.15.

The safety-first philosophy of the wellness assistant has proven to be crucial for the design of the system. Since all user messages are processed with the rule-based classifier, and then only messages identified as non-crisis messages are given to the model, it is guaranteed that no messages containing the crisis language will ever be handled in an automated way. This is a very powerful feature since all messages containing any signs of crisis will be automatically answered by human moderators with useful resources.

An additional practical benefit of this system is the existence of an automatic WhatsApp alert system. People rarely check their dashboards. A message arriving on their phone when stress crosses a threshold brings the system into their actual life, not just their browser.

That said, this project has clear limitations. The emotion model was not fine-tuned on mental health-specific data — it was trained on general emotion-labeled text. The stress vocabulary dictionaries, while covering seven categories, are limited to English and may miss cultural or regional expressions of distress. The system has also not been clinically validated against professional assessments, which means the scores it produces are useful for personal awareness but should not be treated as diagnostic results.

Future work could address these limitations by fine-tuning the model on mental health datasets, expanding language support, integrating physiological signals like heart rate or typing pace, and conducting a proper clinical evaluation with participants. A mobile application that runs the analysis passively in the background would also make the system far more accessible to people who need it most.

What MINDWATCH proves is that it is entirely possible to build a practical, privacy-respecting, and technically sound mental wellness monitoring system without access to cloud infrastructure or clinical resources. The tools exist. The approach works. The remaining challenge is getting it in front of the people who would benefit from it.

References

[1] A. Sharma, Gupta and Verma, "Mental health monitoring using sentiment analysis and machine learning techniques", International Journal of Advanced Research in Computer Science, vol. 12, no. 4, pp. 120–128, 2022.
[2] Deep S. Tadesse, H. Lin, B. Xu, and L. Yang, "Detection of Depression-related Posts in the Reddit Social Media Forum," IEEE Access, vol. 7, pp. 44883–44893, 2019.
[3] E. Chancellor M. De Choudhury, "Methods in Predictive Techniques for Mental Health Status on Social Media," ACM Computing Surveys, vol. 51, no. 6, 2018: 1–36.
[4] A. Orabi, P. Buddhitha, M. Husseini, and K. Orabi, "Deep Learning for Depression Detection of Twitter Users," in Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology, 2018, pp. 88–97.
[5] S. Ji, C. Pan, X. Cambria, P. Marttinen, and S. Yu, "A Survey of Knowledge Graphs: Representation, Acquisition, and Applications," IEEE Transactions on Neural Networks and Learning Systems, vol. 33, no. 2, pp. 494–514, 2022.
[6] M. Inkster, S. Sarda, and V. Subramanian, "An Empathy-driven, Conversational Artificial Intelligence Agent for Digital Mental Health," Frontiers in Digital Health, volume 1, issue 1, pages 1–9, 2018.
[7] K. Jagannatha and H. Yu, "Bidirectional RNN for Medical Event Detection in Electronic Health Records," Proceedings of NAACL-HLT, pp. 473–482, 2016.
[8] S. Dham and R. Sharma, "Sentiment Analysis for Depression Detection on Social Media Platforms using Machine Learning," International Journal of Computer Applications, vol. 176, no. 40, pp. 20–27, 2020.
[9] T. Mikolov et al., "Distributed Representations of Words and Phrases and their Compositionality," Advances in Neural Information Processing Systems (NeurIPS), vol. 26, pp. 3111–3119, 2013.
[10] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," Proceedings of NAACL-HLT, pp. 4171–4186, 2019.