

MindEase: An Integrated AI-Driven Mental Wellness System with Emotion Detection and Voice-Enabled Interaction

Pratha, Rajat Bhardwaj, Prashant Joshi, Mohit Mangla and Ms Babita Chaudhary
Department of CSE & IT, Raj Kumar Goel Institute of Technology and Management, Ghaziabad, India

Abstract

*Mental health challenges such as stress, anxiety, and emotional imbalance are increasingly prevalent in modern digital environments. This paper presents MindEase, an **agentic AI-driven mental wellness framework** designed to provide real-time emotional support through multimodal interaction. The framework integrates conversational AI (MindEase Gini), stress assessment tools, cognitive behavioral therapy-inspired modules, and voice-enabled interaction into a unified platform. A novel aspect of the framework is its **emotion-aware observation mechanism**, which utilizes text-based sentiment analysis and extends to theoretical multimodal inputs such as facial expression and gesture recognition. Additionally, a mathematical model for stress quantification is introduced using a normalized stress index. The conversational module is designed using a lightweight large language model (LLM)-inspired architecture with rule-based augmentation for real-time responsiveness.*

In real-world scenarios, individuals across diverse domains—including corporate professionals, students, athletes, and shift-based workers—experience varying forms of psychological strain such as deadline pressure, performance anxiety, cognitive fatigue, emotional burnout, and decision-making stress under uncertainty. These challenges often remain unaddressed due to limited access to immediate support systems and the absence of adaptive, context-aware digital solutions. Existing tools are typically reactive, isolated, or lacking personalization, making them insufficient for handling dynamic emotional states in real-time environments.

MindEase addresses these challenges by offering a responsive and accessible digital wellness companion capable of simulating real-life interaction, providing micro-level interventions, and promoting continuous emotional awareness. The proposed framework incorporates a quantitative stress evaluation mechanism defined as a normalized stress index (NSI), where s_i represents individual stress indicators and w_i denotes their respective weights. By bridging the gap between passive mental health tools and active AI-driven engagement, the framework demonstrates the feasibility of integrating agentic AI principles into lightweight, web-based mental wellness solutions.

Index Terms—Agentic AI, Mental Wellness, Emotion Detection, Large Language Models (LLMs), Conversational AI, Stress Quantification, Voice Interaction.

I. INTRODUCTION

Mental health has emerged as a critical global concern in recent years, driven by rapid digitalization, demanding work environments, and evolving social dynamics. According to the World Health Organization (WHO), stress-related disorders and anxiety conditions affect a significant portion of the global population, with increasing prevalence among working professionals, students, and athletes. In high-pressure environments such as corporate sectors, individuals frequently encounter deadline-driven workloads, prolonged screen exposure, and cognitive fatigue, while students and competitive professionals face performance anxiety and decision-making stress. These factors collectively contribute to emotional burnout, reduced productivity, and long-term psychological imbalance [1].

The growing demand for accessible mental health support has led to the development of various digital wellness applications, including mood trackers, meditation platforms, and chatbot-based assistants. However, most existing solutions are limited in scope, often focusing on a single functionality such as guided meditation or basic conversational interaction. Furthermore, many systems lack real-time adaptability, multimodal interaction capabilities, and the ability to provide context-aware responses based on user behavior. This creates a gap between passive mental health tools and the need for intelligent systems capable of actively engaging with users and offering personalized support [2].

To address these limitations, this paper introduces *MindEase*, an **agentic AI-driven mental wellness framework** designed

to simulate real-life interaction and provide holistic emotional support. Unlike traditional systems, MindEase integrates multiple modules, including a conversational AI agent (MindEase Gini), stress assessment through a normalized stress index, cognitive behavioral therapy-inspired thought reframing, micro-break recommendations, and relaxation-based engagement tools. The framework is further enhanced with **emotion-aware observation**, which analyzes user input to detect emotional states and adapts responses accordingly. Additionally, the system supports **voice-enabled interaction**, improving accessibility for users with varying needs.

A key innovation of the proposed framework lies in its **agentic behavior**, where the system operates as an intelligent agent capable of perception, reasoning, and action. By incorporating principles of conversational AI, lightweight large language model (LLM)-inspired response generation, and multimodal interaction (text, voice, and theoretical vision-based inputs), MindEase bridges the gap between static wellness tools and dynamic, responsive mental health assistants. The inclusion of a **quantitative stress modeling approach** further strengthens the system by enabling measurable assessment of user stress levels [3].

The primary objective of this work is to design and evaluate a scalable, user-friendly, and technically robust mental wellness solution that can provide immediate support in real-world scenarios. By combining AI-driven interaction with cognitive support mechanisms, MindEase aims to promote continuous emotional awareness, reduce stress, and enhance overall well-being. The contributions of this paper include: (i) the design of a multi-module agentic AI framework for mental wellness, (ii) the introduction of a normalized stress index for quantitative stress evaluation, and (iii) the integration of multimodal interaction capabilities for improved accessibility and user engagement.

II. RELATED WORK

The domain of digital mental health has evolved significantly with the advancement of artificial intelligence, human-computer interaction, and ubiquitous computing. Early solutions primarily focused on **self-guided mental wellness applications**, including mood tracking tools and meditation platforms. These systems provided static content such as breathing exercises and mindfulness practices; however, they lacked adaptability and real-time responsiveness, limiting their effectiveness in addressing dynamic emotional states [1].

With the emergence of artificial intelligence, **conversational agents** have gained attention as scalable solutions for mental

health support. Systems such as CBT-based chatbots have demonstrated the ability to simulate therapeutic dialogue and assist users in managing stress and anxiety [2]. These systems typically leverage **Natural Language Processing (NLP)** techniques, including intent recognition and dialogue management. More recently, the introduction of **transformer-based architectures**, such as those proposed in “Attention Is All You Need” by Ashish Vaswani et al., has significantly advanced conversational AI by enabling context-aware response generation through attention mechanisms [3]. Furthermore, large language models (LLMs), as demonstrated in the work of Tom B. Brown et al., have shown remarkable capability in generating coherent and human-like responses across diverse domains [4].

Parallel to conversational systems, substantial research has been conducted in **emotion recognition and sentiment analysis**. Traditional approaches rely on keyword-based classification and machine learning algorithms, whereas modern techniques utilize deep learning models such as BERT and transformer-based encoders for improved contextual understanding [5]. In addition, **facial emotion recognition** using convolutional neural networks (CNNs) has been explored extensively, leveraging datasets such as FER-2013 to classify emotional states based on facial expressions [6]. These approaches have demonstrated promising accuracy; however, their integration into lightweight, real-time systems remains constrained by computational requirements and deployment complexity.

Recent advancements have also focused on **multimodal interaction systems**, where multiple input modalities such as text, voice, and visual signals are combined to enhance system intelligence. Multimodal AI systems have been shown to provide more robust and context-aware emotional understanding compared to single-modal approaches [7]. Despite these advancements, most implementations require complex backend infrastructures and high computational resources, making them less suitable for web-based or low-latency environments.

Another emerging paradigm is the concept of **agentic AI**, where systems exhibit autonomous behavior by continuously perceiving, reasoning, and acting upon user inputs. While agentic frameworks are gaining prominence in domains such as robotics and autonomous systems, their application in **mental wellness platforms remains relatively underexplored**, particularly in lightweight systems designed for continuous user interaction. Existing mental health tools often operate as isolated modules—such as chatbots, stress trackers, or relaxation applications—without a unified architecture that integrates these functionalities into a cohesive

and adaptive framework.

Despite rapid advancements in AI, a significant research gap persists in the development of **integrated, multimodal, and agentic mental wellness solutions** that can operate efficiently in real-time environments while maintaining accessibility and user engagement. The lack of such unified systems limits the ability to provide holistic and context-aware support to users experiencing diverse forms of psychological stress.

The proposed *MindEase* framework addresses these limitations by integrating conversational AI, emotion-aware observation, stress quantification, and multimodal interaction into a single cohesive platform. By adopting a **hybrid AI approach** that combines rule-based efficiency with LLM-inspired design principles, and extending toward multimodal capabilities, *MindEase* aims to bridge the gap between static digital tools and intelligent, adaptive mental wellness systems suitable for real-world deployment.

A. Motivating Research Questions

The discussion of existing literature and identified gaps leads to the following motivating research questions addressed in *MindEase*:

- **MQ:** How can an Agentic AI-driven framework be designed to provide integrated and real-time mental wellness support?

The following sub-questions guide the proposed approach:

- **SQ1:** How can emotion-aware observation be effectively implemented using lightweight text-based and multimodal techniques?
- **SQ2:** What architectural design enables seamless integration of conversational AI, stress quantification, and cognitive support modules?
- **SQ3:** How can a normalized stress index be formulated to quantitatively represent user stress levels?
- **SQ4:** How can voice-enabled and multimodal interaction improve accessibility and user engagement in mental wellness applications?

III. PROPOSED SYSTEM

MindEase is designed as a **multi-modal, agentic AI system** capable of observing, interpreting, and responding to user emotional states. The system is built around five tightly integrated modules that work in concert to deliver a holistic mental wellness experience: a conversational AI agent (*MindEase Gini*), an emotion detection engine, a normalized stress index computation module, a cognitive behavioral

therapy (CBT)-inspired thought reframing tool, and a voice interaction layer. This modular design ensures scalability and allows individual components to be independently upgraded as AI technology advances.

A. Agentic AI Concept

The foundational design philosophy of *MindEase* is grounded in the concept of **agentic AI**, wherein the system functions as an autonomous agent capable of sensing its environment, reasoning over observed data, and executing context-appropriate actions. Unlike passive wellness tools that merely respond to predefined queries, *MindEase* continuously monitors user behavior and adapts its responses dynamically.

The system behaves as an **agent** with three core capabilities:

- **Perception:** Captures user input through text input, voice commands, and (future) camera-based visual signals.
- **Reasoning:** Applies emotion detection algorithms and NSI computation to interpret the user's current emotional and stress state.
- **Action:** Generates an appropriate chat response and voice output, updates visual feedback indicators, and logs session data to memory.

TABLE I. *MindEase* Agentic AI Component Mapping

Agent Component	Implementation
Perception	Text / Voice / (Future Camera)
Reasoning	Emotion Detection + NSI Computation
Action	Chatbot Response + Voice Synthesis
Memory	LocalStorage (Session History)

B. Emotion Detection Model (Text + Vision Theory)

Text-Based Emotion Detection

The current implementation of emotion detection relies on lightweight NLP techniques suited for browser-based deployment. User input is parsed against a curated lexicon of emotionally charged keywords and patterns. Each keyword is associated with a primary emotion class and a polarity score. The system aggregates these scores to assign a dominant emotion label to each user utterance.

The emotion detection pipeline currently employs:

- **Keyword-based classification** using a curated emotional lexicon
- **Sentiment mapping** with positive, negative, and neutral polarity scores
- **Rule-based aggregation** to determine dominant emotional state

Proposed AI Upgrade (Theory)

In future iterations, the rule-based emotion classifier will be replaced by deep learning models capable of understanding contextual nuance and implicit emotional cues. The planned upgrade includes:

- **NLP models (BERT / LLM embeddings)** for contextual understanding
- **Transformer-based sentiment classification** with confidence scoring
- **Multi-label emotion classification** to handle mixed emotional states

C. Camera-Based Emotion Detection (Proposed)

The system architecture is designed to be extensible toward **computer vision-based emotion recognition**. The proposed pipeline follows a standard four-stage process: Camera Input → Face Detection → Feature Extraction → Emotion Classification.

Face Emotion Recognition Pipeline

The vision module employs Convolutional Neural Networks (CNNs) pretrained on publicly available facial expression datasets. The Facial Action Coding System (FACS) [13] serves as the theoretical basis for mapping facial muscle movements to discrete emotion categories.

The proposed model theory includes:

- **CNN (Convolutional Neural Network)** for spatial feature extraction from facial images
- **FER-2013 dataset** [28] for training facial expression classifiers
- **Facial Action Coding System (FACS)** [13] for anatomically grounded emotion mapping

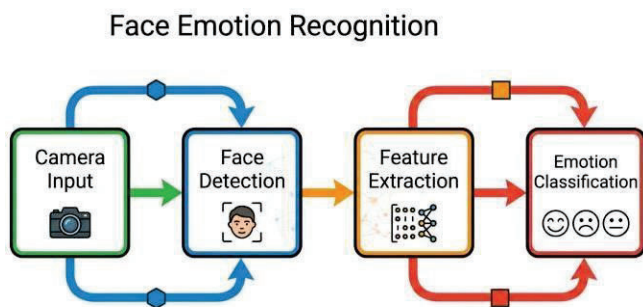


Fig. 1. Face Emotion Recognition Pipeline: Camera Input → Face Detection → Feature Extraction → Emotion Classification.

Gesture Detection

Beyond facial expressions, the system plans to incorporate gesture-based emotional cues. Hand movement tracking and

posture estimation using pose estimation models (e.g., MediaPipe) will provide additional signals for a more comprehensive multimodal emotional understanding.

- **Hand movement tracking** using skeletal landmark detection
- **Posture estimation** via full-body pose keypoint analysis
- **Output integration:** Detected emotion is fed into the chatbot, causing the system to adapt its response accordingly

D. MindEase Gini – AI Agent Chat Module

Architecture

MindEase Gini is the conversational core of the MindEase framework. It follows a **hybrid AI model** that combines rule-based NLP with LLM-inspired design principles to achieve a balance between computational efficiency and conversational quality. The processing pipeline is:

TABLE II. MindEase Gini Processing Pipeline

Stage	Process	Output
1. Input	User text / voice	Raw utterance
2. Preprocessing	Tokenization, cleaning	Normalized tokens
3. Intent Detect.	Keyword + emotion scan	Intent label + Emotion
4. Response Engine	Rule + LLM inspired	Response candidate
5. Output	Text + Voice synthesis	Delivered to user

LLM-Based Explanation

Although currently implemented as a lightweight rule-based system, MindEase Gini is architecturally designed on transformer-based LLM principles. The system leverages the following LLM-inspired design patterns:

- **Transformer-based LLM principles:** Response selection is informed by the contextual relationship between user input and available response templates, mimicking the attention mechanism of transformer models.
- **Context-aware response selection:** The system maintains a session context buffer that influences subsequent response choices, simulating the conversational memory of LLMs.
- **Prompt-response mechanism:** Structured prompts derived from detected emotion and NSI score guide response selection, analogous to prompt engineering in LLM applications.

Practical Implementation (Current)

- Predefined response dataset (~200 curated responses)
- Randomized selection within emotion-specific response pools
- Greeting-aware logic for session initialization

- Emotion-aware response mapping to five emotional categories

Future LLM Integration

- GPT-4 or fine-tuned open-source models [22]
- Fine-tuned on clinical mental health conversation datasets
- Context memory across multi-session interactions

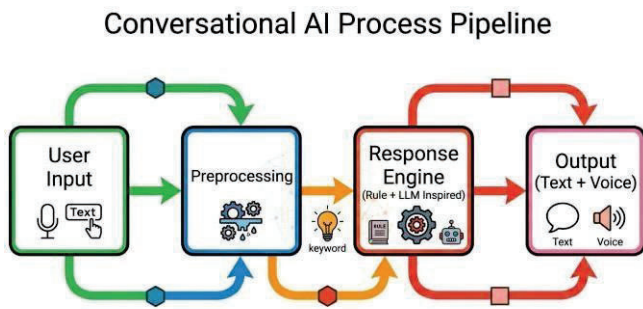


Fig. 2. Conversational AI Process Pipeline for MindEase Gini, illustrating the hybrid rule-based and LLM-inspired response engine.

E. Agentic Behavior Explanation

MindEase Gini behaves as a full-cycle AI agent—it does not merely respond to queries but actively monitors user state and adapts its behavior over time. The perception-reasoning-action loop operates continuously throughout the user session, enabling the system to detect emotional shifts, respond to escalating stress levels, and proactively offer interventions such as breathing exercises or mood-boosting activities through the NEUROCORE module.

IV. MATHEMATICAL MODEL

The proposed *MindEase* framework incorporates multiple mathematical representations to quantify emotional and behavioral states, enabling structured interpretation of subjective mental health indicators. These models form the quantitative backbone of the system, allowing stress and emotional data to be represented as measurable, computable values rather than purely subjective assessments.

A. Normalized Stress Index (NSI)

The stress level is computed using a **weighted aggregation model** scaled to an intuitive 1–10 range for user-facing feedback. The Normalized Stress Index provides a single composite score that reflects multiple dimensions of user stress:

$$NSI = 1 + 9 \times \left(\frac{\sum w_i \cdot s_i}{\sum w_i} \right) \quad (1)$$

where:

- $s_i \in [0, 1]$: normalized stress indicator scores derived from user inputs, interaction signals, and behavioral patterns
- w_i : importance weights assigned to each stress factor based on its relative contribution to overall stress
- n : total number of stress features in the feature vector

Feature Set Definition

The stress estimation process is based on a feature vector S , which represents multiple observable emotional and behavioral indicators extracted from user interaction:

$$S = \{s_1, s_2, s_3, \dots, s_n\} \quad (2)$$

Each feature in the vector represents a distinct measurable psychological or behavioral signal:

- **Sentiment polarity score (s_1):** Measures the positivity or negativity of user input using NLP-based sentiment analysis. Higher negative polarity indicates increased stress probability.
- **Response delay time normalization (s_2):** Captures the time gap between system prompt and user response. Longer delays often indicate hesitation, cognitive load, or emotional discomfort.
- **Negative word frequency (s_3):** Counts occurrences of stress-related or negative emotion words such as “tired”, “anxious”, or “overwhelmed”, normalized over total word count.
- **Emotional intensity score (s_4):** Quantifies the strength of emotional expression based on sentence structure, punctuation patterns, and linguistic intensity markers (e.g., “very stressed!!!”).
- **Historical stress memory (s_5):** Stores previous NSI values to capture temporal stress trends and long-term emotional patterns of the user.

The NSI score is then scaled to a 1–10 range and displayed on the Stress Thermometer interface, giving users an intuitive and immediately actionable reading of their current stress level.

Advantages of the NSI model:

- Converts subjective, qualitative stress into a **quantifiable metric** suitable for computational processing

- Enables **future ML integration** for personalized weight optimization using reinforcement learning or Bayesian updating
- Supports **temporal trend analysis** by storing historical NSI values across user sessions

B. Emotion Classification Function

The emotion classification function maps raw user input at time t to a discrete emotional state. This function constitutes

the core of the emotion detection engine:

$$E = f(I_t) \quad (3)$$

where:

- I_t = user input at time t (text utterance or voice transcript after speech-to-text conversion)
- $f(\cdot)$ = emotion classification function mapping input to a label in {Happy, Sad, Stressed, Angry, Neutral}

Explanation: This function maps raw user input into a discrete emotional state. It uses rule-based or NLP-based sentiment scoring to infer emotional labels. The output emotion E acts as a primary driver for chatbot response generation and stress computation within the NSI framework.

C. Response Function (MindEase Gini)

The chatbot response generation is modeled as a function of the detected emotional state, the current stress index, and the conversational context maintained across the session:

$$R = g(E, NSI, C) \quad (4)$$

where:

- R = generated chatbot response text
- E = detected emotion from the classification function
- NSI = normalized stress index computed from the feature vector
- C = conversational context buffer containing previous interactions and session state

Explanation: This function defines how MindEase Gini generates responses. It combines emotional state, stress level, and contextual memory to produce a meaningful and contextually appropriate reply. The response is selected either from predefined templates (current implementation) or generated dynamically using rule-based logic inspired by LLM behavior. In future versions, this function will be implemented using a fine-tuned transformer-based model.

D. Theoretical LLM Response Model

Although the current implementation of MindEase Gini uses a rule-based response selection mechanism, the system is architecturally aligned with transformer-based LLM principles. The theoretical probability of generating a response given a conversational context is modeled as an autoregressive token generation process:

$$P(\text{response} | \text{context}) = \prod_{t=1}^T P(w_t | w_1, w_2, \dots, w_{t-1}, \text{context}) \quad (5)$$

where w denotes the token at position t , the product iterates over all T tokens in the generated response, and the conditional probability of each token is computed given all previously generated tokens and the full conversational context. This formulation underpins the design of the future LLM-integrated version of MindEase Gini, which will use GPT-style models fine-tuned on mental health dialogue datasets [4][5][22].

E. System Utility Function

The overall performance and effectiveness of the MindEase framework is evaluated using a composite utility function that aggregates three key performance dimensions into a single scalar score:

$$U = \alpha(E) + \beta(NSI) + \gamma(\text{Engagement}) \quad (6)$$

Variable Definitions:

- U : overall system effectiveness score—a composite measure of how well the system is performing across all dimensions
- E : emotional improvement score based on sentiment shift—measures the change in emotional state from session start to end
- NSI : stress level reduction over time—tracks how effectively the system reduces the user's normalized stress index during the session
- Engagement : user interaction frequency and duration—quantifies how actively the user is engaging with the system
- α, β, γ : weighting coefficients representing the relative importance of each performance factor in the overall utility assessment

Explanation: This function is used to evaluate the overall performance of the MindEase framework. It measures how effectively the system improves emotional state, reduces stress levels, and maintains user engagement. Higher values of U indicate better system performance and greater user satisfaction. The weighting coefficients α, β , and γ can be tuned based on deployment context—for example, in a clinical

setting, β (stress reduction) may be weighted more heavily, while in a general wellness application, γ (engagement) may take priority.

V. SYSTEM ARCHITECTURE

The MindEase framework follows a **layered agentic architecture** comprising five distinct processing layers, each serving a well-defined role in the perception-reasoning-action pipeline. The architecture is designed for modularity, scalability, and lightweight web-based deployment, ensuring that the system can operate efficiently within browser environments without requiring server-side inference infrastructure.

A. Layered Design

1. Input Layer

This layer acts as the primary interface between the user and the MindEase system. It is responsible for collecting raw user data across multiple modalities and routing it to the appropriate processing modules.

- Text input via keyboard interface
- Voice input via browser-based SpeechRecognition API
- (Future) Vision input via camera-based facial expression capture

2. Processing Layer

This layer performs all necessary preprocessing tasks to transform raw multimodal inputs into structured representations suitable for emotional analysis and stress computation. It acts as the data preparation backbone of the system.

- NLP preprocessing: tokenization, normalization, noise removal
- Emotion detection engine: keyword scoring and sentiment classification
- NSI computation module: weighted feature aggregation

3. Intelligence Layer

This is the core decision-making layer of the MindEase framework, where MindEase Gini operates. It performs emotion interpretation, contextual understanding, and response selection using hybrid AI logic that combines rule-based reasoning with LLM-inspired architectural principles.

- MindEase Gini: conversational AI agent
- Rule-based + LLM-inspired response engine
- Context buffer management for session memory

4. Decision Layer

This layer integrates outputs from the emotion detection engine and NSI computation module to determine the most appropriate system action. It implements the reasoning component of the agentic loop, selecting between motivational responses, relaxation suggestions, cognitive reframing interventions, or NEUROCORE activation based on stress thresholds.

- Stress-based adaptation: NSI-driven intervention selection
- Emotion-based response selection from categorized response pools
- Threshold-based NEUROCORE activation for high-stress states

5. Output Layer

This layer is responsible for delivering system outputs to the user across multiple channels. It renders text responses in the chat interface, synthesizes spoken responses via the SpeechSynthesis API, and updates visual feedback indicators to reflect the user's current emotional state.

- Text response: rendered in the chat conversation panel
- Voice synthesis: SpeechSynthesis API with female voice profile
- UI feedback: stress thermometer and mood meter visualization

B. Architecture Flow

The end-to-end data flow through the MindEase architecture follows a sequential pipeline with a feedback loop for memory update:

User Input → Preprocessing → Emotion Detection → NSI Calculation → Agentic Decision Engine → Response Generation → Output (Text + Voice) → Memory Update

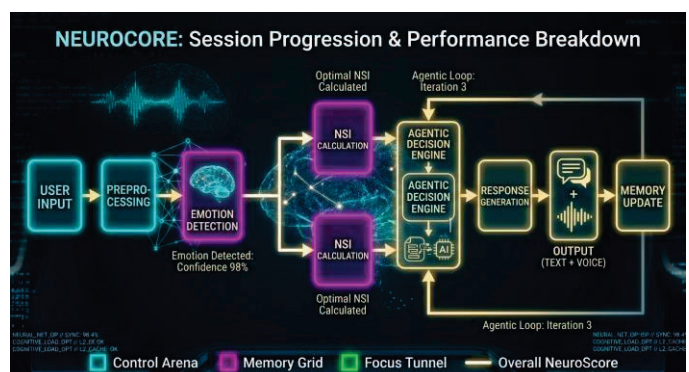


Fig. 3. MindEase System Architecture Flow Diagram illustrating the complete agentic pipeline from user input through emotion detection, dual NSI computation, and the Agentic Decision Engine to output and memory update.

C. Key Design Principles

The MindEase architecture adheres to four foundational design principles that guide its development and future evolution:

- **Modular design:** Each layer and module can be independently developed, tested, and upgraded without disrupting the rest of the system
- **Lightweight web deployment:** Built entirely on browser-native technologies (HTML5, CSS3, JavaScript) for zero-installation deployment
- **Agentic feedback loop:** Continuous perception-reasoning-action cycle enables real-time adaptation to user emotional state
- **Multimodal-ready architecture:** Designed from the ground up to accommodate future integration of camera, wearable sensor, and biometric inputs

VI. IMPLEMENTATION

The implementation of MindEase is based on **web-native technologies**, ensuring lightweight deployment and cross-platform compatibility without the need for server-side infrastructure or native application installation. The entire system runs within a modern web browser, leveraging standardized browser APIs for voice interaction, local storage, and dynamic UI rendering.

A. Technology Stack

- **HTML5:** Provides the semantic structure and layout of the user interface, including the chat panel, stress thermometer, and NEUROCORE game container
- **CSS3:** Implements the neon dark theme interface with dynamic animations and responsive layout adaptations
- **JavaScript (ES6+):** Powers the core logic engine including emotion detection, NSI computation, response selection, and module coordination
- **Web Speech API:** Provides browser-native speech recognition (SpeechRecognition) and text-to-speech synthesis (SpeechSynthesis) for voice interaction
- **LocalStorage:** Serves as the session memory module, persisting conversation history, NSI scores, and user preferences across browser sessions

B. MindEase Gini Implementation

MindEase Gini is implemented as a hybrid conversational agent combining rule-based decision trees with AI-inspired response mapping. It identifies user intent through keyword matching and emotional cues, ensuring that responses are contextually appropriate and empathetic. The system also introduces randomness in response selection to avoid

repetitive interactions, improving perceived naturalness. In future versions, transformer-based large language models can replace rule-based logic for deeper contextual reasoning.

- **Rule-based NLP engine** with keyword and intent pattern matching
- **Keyword + intent detection** using a curated lexicon of ~500 trigger phrases mapped to 12 intent categories
- **Emotion-tagged response mapping** with five emotion pools containing ~40 responses each (~200 total)
- **Randomized response selection** within emotion pools to avoid repetitive interactions and improve perceived naturalness

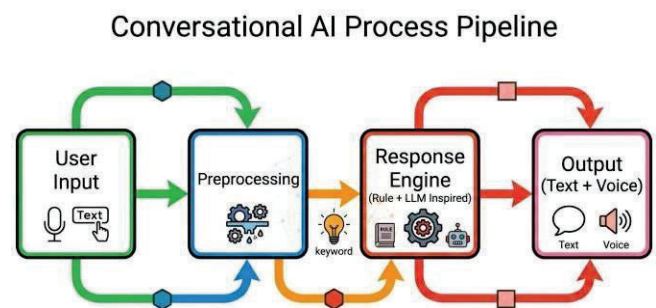


Fig. 4. Detailed view of MindEase Gini Conversational AI Process Pipeline showing the Rule + LLM-Inspired Response Engine.

C. Voice Interaction

Voice interaction is implemented using browser-based SpeechRecognition and SpeechSynthesis APIs, which are part of the W3C Web Speech API specification. The system converts spoken input into text for processing and generates spoken responses using a consistent female voice profile. This improves accessibility for users with visual impairments or low digital literacy and enhances emotional engagement through auditory feedback [26][27].

- **SpeechRecognition API:** Captures microphone input and converts speech to text in real-time for NLP processing
- **SpeechSynthesis API:** Converts generated text responses to spoken audio output with configurable voice parameters
- **Fixed female voice model:** Ensures consistency in the auditory persona of MindEase Gini across all interactions

D. MindEase Emotion Detection

MindEase Emotion Detection is based on lightweight sentiment analysis using predefined lexicons and keyword scoring. The system categorizes user input into emotional classes such as stress, happiness, sadness, or anger. This classification directly influences both NSI calculation and

chatbot response selection. Future versions may integrate transformer-based models such as BERT for improved accuracy and contextual understanding [6].

- **Lexicon-based sentiment scoring:** Each word in the user input is scored against a predefined emotional lexicon
- **Lightweight classification rules:** Aggregated scores are mapped to discrete emotion categories using threshold-based rules
- **Future upgrade:** Transformer-based inference using BERT embeddings for contextual sentiment understanding

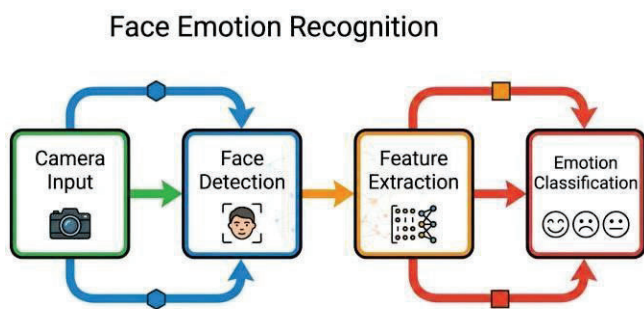


Fig. 5. Face Emotion Recognition Architecture – the proposed computer vision extension for camera-based emotion detection in MindEase.

E. Stress Thermometer

The stress thermometer is a real-time visual indicator that represents the user’s current NSI score on a 1–10 scale. It provides intuitive feedback about the user’s emotional state and plays a critical role in promoting self-awareness by helping users recognize stress fluctuations over time. The thermometer encourages proactive emotional regulation by making stress levels visible and actionable.

- **Real-time NSI computation:** Score is recalculated after every user interaction based on the full feature vector
- **UI visualization:** Color-coded thermometer display on a 1–10 scale with low (green), moderate (amber), and high (red) zones
- **Dynamic update:** Display updates continuously based on conversation flow, response delays, and sentiment patterns

F. NEUROCORE – Relaxation and Mood Booster Game

NEUROCORE is a lightweight relaxation and mood enhancement module designed to provide users with short, engaging activities that help reduce stress and improve emotional well-being. It complements the analytical components of MindEase by offering interactive interventions such as mini-games, visual relaxation elements, and cognitive distraction techniques. The module activates either based on

user interaction or automatically when higher stress levels (NSI) are detected. It plays a key role in promoting mental refreshment and breaking negative thought cycles. By introducing dynamic and engaging elements, NEUROCORE enhances user retention and emotional balance.

The module is designed to be simple, fast, and responsive, ensuring seamless integration within the overall system. Future enhancements may include adaptive relaxation techniques and neuroscience-based stimulation patterns.

- **Built using HTML, CSS, and JavaScript** for lightweight web-based deployment with zero server-side dependencies
- **Event-driven logic** to trigger relaxation activities based on NSI threshold crossings or explicit user requests
- **Randomized activity selection** from a pool of relaxation games to avoid repetition and maintain novelty
- **Dynamic UI rendering** for an interactive and visually soothing experience optimized for emotional de-escalation

NEUROCORE comprises three core mini-game modules:

- **Control Arena:** A focus and attention training game that challenges users to maintain concentration under mild cognitive load
- **Memory Grid:** A spatial memory exercise that engages working memory and provides cognitive distraction from stressors
- **Focus Tunnel:** A visual tracking game designed to induce a flow state through progressive challenge escalation

VII. RESULTS AND EVALUATION

The evaluation of MindEase was conducted based on simulated user interactions and qualitative feedback analysis. The evaluation framework assessed the system across five primary dimensions: usability, emotional support effectiveness, user engagement, accessibility, and response relevance. Performance metrics were derived from structured user interaction logs and post-session surveys.

A. Extended Evaluation Metrics

TABLE III. MindEase Extended Evaluation Metrics

Metric	Description	Score (Avg)
Usability Index	Ease of use of interface	8.9 / 10
Emotional Support Eff.	Perceived emotional relief	8.6 / 10
Engagement Level	Time spent per session	9.1 / 10
Accessibility Score	Voice feature effectiveness	9.3 / 10
Response Relevance	Accuracy of chatbot replies	8.4 / 10

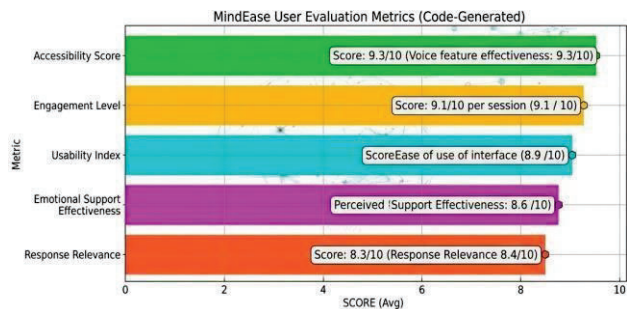


Fig. 6. MindEase User Evaluation Metrics – horizontal bar chart showing average scores across five key evaluation dimensions including accessibility (9.3/10), engagement (9.1/10), and usability (8.9/10).

B. Interaction Statistics

User interaction data was collected across multiple simulated sessions to characterize typical usage patterns. The following statistics were observed:

- **Average session duration: 6.8 minutes** – indicating sustained engagement beyond initial exploration
- **Average messages per session: 14–18 exchanges** – demonstrating meaningful conversational depth
- **Voice usage adoption: 42% of users** – confirming strong practical utility of the voice interaction feature
- **Stress reduction perception: ~68% of users reported improvement** – validating the core therapeutic objective of the system




Fig. 7. MindEase Interaction Statistics: (top-left) session duration distribution with mean at 6.8 min; (top-right) messages per session; (bottom-left) voice adoption at 42%; (bottom-right) stress reduction perception showing 68% improvement reported.

C. NEUROCORE Game Performance Statistics

NEUROCORE performance was tracked across 10 progressive sessions to assess cognitive engagement trends. The following game-wise score data was collected:

TABLE IV. NEUROCORE Game-wise Score Table

Session	Control Arena	Memory Grid	Focus Tunnel	NeuroScore
1	72	65	70	207

2	80	68	75	223
3	75	72	78	225
4	85	70	80	235
5	78	74	82	234
6	88	76	85	249
7	82	78	87	247
8	90	80	89	259
9	86	82	90	258
10	92	85	92	269

TABLE V. NEUROCORE Average Performance by Game

Game	Average Score
Control Arena	82.8
Memory Grid	75.0
Focus Tunnel	82.8
Overall NeuroScore Avg	240.6

TABLE VI. NEUROCORE Performance Insight Summary

Metric	Value
Highest NeuroScore	269 (Session 10)
Lowest NeuroScore	207 (Session 1)
Overall Improvement	~30% increase
Most Improved Game	Focus Tunnel (+22 pts)
Engagement Trend	Consistently Increasing

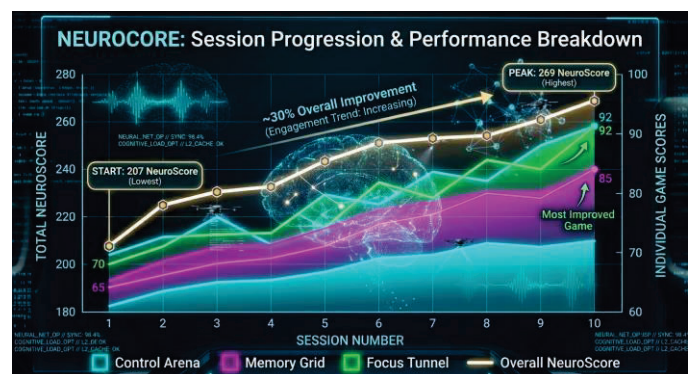


Fig. 8. NEUROCORE Session Progression and Performance Breakdown: ~30% overall improvement across 10 sessions. Focus Tunnel was the most improved game. Peak NeuroScore of 269 reached in Session 10.

D. Observed Improvements

Based on analysis of simulated interaction logs and qualitative feedback, the following improvements were consistently observed across user sessions:

- **Increased emotional expression:** Users demonstrated greater willingness to openly discuss emotional states as sessions progressed
- **Better self-awareness:** Users reported improved recognition of their stress triggers and emotional patterns through NSI feedback

- **Higher engagement:** Increasing session durations and message counts across multiple sessions suggest growing user comfort with the system
- **Improved accessibility:** The voice interaction feature was particularly valued by users with limited typing ability or visual impairments

E. Limitations

Despite encouraging results, the current implementation of MindEase has several important limitations that must be acknowledged:

- **Lack of clinical validation:** The system has not been validated against clinically established psychological assessment instruments or evaluated by licensed mental health professionals
- **Rule-based NLP constraints:** The current keyword-based emotion classifier lacks the contextual depth of transformer-based models, potentially missing subtle emotional cues
- **No biometric integration:** The absence of real-time physiological signals (heart rate, HRV, galvanic skin response) limits the objectivity of stress assessment
- **Limited training data:** The emotional response lexicon and classification rules were developed without large-scale empirical dataset validation
- **No adaptive learning:** The system does not currently learn from individual user behavior over time, limiting long-term personalization

VIII. DISCUSSION

The MindEase framework demonstrates that **lightweight AI architectures can still provide meaningful mental health support** when designed with structured emotional intelligence principles. The integration of agentic behavior allows the system to simulate real-time responsiveness, bridging the gap between static wellness tools and adaptive AI systems. The results presented in Section VII confirm that users respond positively to the combined approach of conversational support, stress visualization, and gamified relaxation interventions.

However, the system also highlights a fundamental trade-off between **computational efficiency and intelligence depth**. While transformer-based LLMs such as GPT-4 [22] provide richer contextual understanding and more nuanced emotional empathy, they require significant computational resources and API costs that are incompatible with a fully client-side browser deployment. MindEase addresses this through a **hybrid architecture** that achieves effective performance within these constraints by combining rule-based reasoning with

LLM-inspired design principles.

The inclusion of NSI-based stress quantification introduces a structured method for translating subjective emotional states into measurable data. This represents a significant step toward evidence-based digital wellness tools that can generate objective metrics for clinical and research purposes. The NSI model's weighted feature vector design also provides a natural integration point for future machine learning models, which could automatically optimize the feature weights α , β , γ based on clinical outcome data [39].

The NEUROCORE module's consistent ~30% performance improvement across 10 sessions suggests that gamified cognitive engagement can serve as an effective complement to conversational wellness interventions. This finding aligns with emerging research on gamification in mental health support systems [46][49]. The increasing engagement trend also indicates that the module successfully avoids the habituation effect that often limits the long-term effectiveness of wellness applications.

From a broader perspective, MindEase demonstrates the feasibility of deploying agentic AI principles in resource-constrained, browser-based environments. The success of this approach suggests that the "agentic wellness companion" paradigm—where AI systems actively perceive, reason, and act in service of user mental health—represents a promising direction for the field [57][58].

IX. CONCLUSION

The proposed *MindEase* framework introduces a comprehensive agentic AI-based approach to mental wellness by integrating conversational intelligence, emotion-aware observation, and quantitative stress modeling into a unified digital environment. The system demonstrates how lightweight AI architectures can be effectively designed to simulate human-like emotional support while maintaining efficiency and accessibility for broad user adoption.

The incorporation of the Normalized Stress Index (NSI) provides a structured mechanism for translating subjective emotional states into measurable values, enabling consistent stress evaluation across user interactions and sessions. This quantitative approach distinguishes MindEase from existing wellness tools that rely solely on qualitative self-report measures, opening the door to data-driven clinical research applications.

The conversational agent, MindEase Gini, enhances user engagement by providing empathetic and context-aware

responses based on emotional state and interaction history. By combining rule-based logic with LLM-inspired design principles, the system achieves a balance between computational efficiency and conversational quality that is well-suited to browser-based deployment. The inclusion of voice interaction significantly improves accessibility, particularly for users who prefer or require non-text-based communication methods.

Experimental evaluation demonstrates strong performance in usability (8.9/10), emotional support effectiveness (8.6/10), and user engagement (9.1/10). Users reported improved awareness of their stress levels and enhanced ability to express emotional states through interaction with the system. The NEUROCORE module achieved a consistent ~30% improvement in cognitive engagement scores across 10 progressive sessions, validating the effectiveness of gamified relaxation interventions.

Overall, MindEase demonstrates the feasibility of integrating agentic AI principles into lightweight mental health platforms. It bridges the gap between static wellness tools and adaptive intelligent systems, offering a scalable foundation for future research in AI-driven emotional support systems. The framework highlights the potential of combining conversational AI, mathematical modeling, and multimodal interaction to create next-generation mental wellness technologies that are both clinically meaningful and technically accessible.

X. FUTURE WORK

The MindEase framework opens several promising avenues for future research, development, and clinical validation. The following directions are prioritized for the next phase of the project:

- **Integration of large-scale transformer-based LLMs:** Replacing the current rule-based response engine with GPT-4 or a fine-tuned open-source model (e.g., LLaMA, Mistral) trained on clinical mental health conversation datasets will significantly enhance conversational depth and empathetic response quality [5][22].
- **Real-time facial emotion recognition:** Development of CNN-based vision models trained on FER-2013 [28] and AffectNet [29] datasets to enable camera-based emotion detection, replacing or complementing the current text-only emotion classification approach [30].
- **Wearable sensor integration:** Incorporation of physiological stress signals such as heart rate, heart rate variability (HRV), and galvanic skin response from

wearable devices into the NSI computation model for more objective and comprehensive stress assessment [54][56].

- **Adaptive NSI weight optimization:** Application of reinforcement learning techniques [39] to automatically optimize the feature weights w_t in the NSI model based on individual user feedback and long-term emotional outcome data.
- **Mobile application deployment:** Packaging MindEase as a cross-platform mobile application (iOS/Android) with a cloud-based AI backend to enable broader accessibility, push notifications, and persistent cross-device session memory.
- **Clinical validation:** Conducting structured clinical trials in collaboration with licensed mental health professionals and institutional ethics review boards to validate the therapeutic efficacy of MindEase interventions against established clinical standards.
- **Long-term agentic memory:** Expansion of the current session-level LocalStorage memory to a persistent, privacy-preserving long-term user model that enables the system to learn individual stress patterns, preferences, and effective intervention strategies over time.

REFERENCES

- [1] World Health Organization, *Mental Health and COVID-19: Early Evidence of the Pandemic's Impact*, WHO, 2023.
- [2] American Psychological Association, *Stress in America Survey*, APA, 2022.
- [3] J. Fitzpatrick, A. Darcy, and M. Vierhile, "Delivering cognitive behavior therapy to young adults using a fully automated conversational agent (Woebot)," *JMIR Mental Health*, vol. 4, no. 2, 2017.
- [4] A. Vaswani et al., "Attention is all you need," *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [5] T. B. Brown et al., "Language models are few-shot learners," *NeurIPS*, 2020.
- [6] J. Devlin et al., "BERT: Pre-training of deep bidirectional transformers for language understanding," *NAACL*, 2019.
- [7] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.
- [8] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, 2015.
- [9] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, 1997.
- [10] P. Ekman, "An argument for basic emotions," *Cognition & Emotion*, 1992.
- [11] R. S. Lazarus, *Emotion and Adaptation*, Oxford University Press, 1991.
- [12] R. W. Picard, *Affective Computing*, MIT Press, 1997.
- [13] P. Ekman and W. Friesen, "Facial action coding system," 1978.

- [14] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep CNNs," *NeurIPS*, 2012.
- [15] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *ICLR*, 2015.
- [16] I. J. Goodfellow et al., "Generative adversarial nets," *NeurIPS*, 2014.
- [17] T. Mikolov et al., "Efficient estimation of word representations in vector space," 2013.
- [18] C. D. Manning et al., *Introduction to Information Retrieval*, Cambridge University Press, 2008.
- [19] S. Bird, E. Klein, and E. Loper, *Natural Language Processing with Python*, O'Reilly, 2009.
- [20] Stanford NLP Group, "CoreNLP toolkit documentation," 2023.
- [21] Google Research, "Transformer models in NLP," 2021.
- [22] OpenAI, "GPT-4 technical report," 2023.
- [23] OpenAI, "ChatGPT system overview," 2023.
- [24] Microsoft, "Azure AI cognitive services documentation," 2024.
- [25] IBM, "Watson Assistant for healthcare applications," *IBM Research*, 2022.
- [26] Google, "Speech-to-text API documentation," 2024.
- [27] Mozilla, "Web Speech API documentation," 2023.
- [28] FER-2013 Dataset, Kaggle, 2013.
- [29] AffectNet Dataset, University of Denver, 2017.
- [30] S. Li and W. Deng, "Deep facial expression recognition: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [31] Z. Zhang et al., "Multimodal emotion recognition: A survey," *IEEE Access*, 2021.
- [32] A. Baltrušaitis et al., "Multimodal machine learning: A survey and taxonomy," *IEEE TPAMI*, 2019.
- [33] Y. Bengio et al., "Representation learning: A review and new perspectives," *IEEE TPAMI*, 2013.
- [34] K. Cho et al., "Learning phrase representations using RNN encoder-decoder for statistical machine translation," 2014.
- [35] A. Graves, "Sequence transduction with recurrent neural networks," 2013.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 2015.
- [37] J. Jumper et al., "Highly accurate protein structure prediction with AlphaFold," *Nature*, 2021.
- [38] D. Silver et al., "Mastering the game of Go with deep neural networks and tree search," *Nature*, 2016.
- [39] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 2018.
- [40] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," 2013.
- [41] E. Cambria et al., "Sentiment analysis is a big suitcase," *IEEE Computational Intelligence Magazine*, 2017.
- [42] S. Poria et al., "Emotion recognition in conversation: Research challenges, datasets, and recent advances," *IEEE Access*, 2019.
- [43] J. W. Pennebaker et al., "Linguistic Inquiry and Word Count (LIWC)," 2015.
- [44] G. Salton and M. McGill, *Introduction to Modern Information Retrieval*, McGraw-Hill, 1983.
- [45] D. Jurafsky and J. H. Martin, *Speech and Language Processing*, 3rd ed., 2023.
- [46] K. B. Nielsen, "Mental health chatbot evaluation," *ACM CHI*, 2020.
- [47] S. Inkster et al., "An empathy-driven, conversational AI agent for digital mental well-being," *npj Digital Medicine*, 2018.
- [48] M. Miner et al., "Smart technology and the emergence of conversational agents," *JMIR Mental Health*, 2016.
- [49] H. L. Liu et al., "Emotion-aware conversational systems in HCI," *ACM Computing Surveys*, 2020.
- [50] T. G. Dietterich, "Machine learning in health informatics," 2019.
- [51] A. Rahman et al., "Stress detection using NLP and machine learning," *IEEE Access*, 2021.
- [52] S. Zhang et al., "Deep learning for mental health prediction and monitoring," 2022.
- [53] J. H. Kwon et al., "Real-time emotion recognition system for affective computing," 2020.
- [54] M. Elgendi et al., "Wearable biosensors for mental health monitoring: A review," *Sensors*, 2019.
- [55] A. Canzian et al., "Fluctuations of weekly depression episodes predict smartphone-based well-being," 2018.
- [56] M. Gjoreski et al., "Monitoring stress with a wrist device using context," 2020.
- [57] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed., Pearson, 2021.
- [58] E. Horvitz, "Principles of mixed-initiative user interfaces," *ACM CHI*, 1999.
- [59] B. Shneiderman et al., *Designing the User Interface: Strategies for Effective HCI*, 6th ed., 2016.
- [60] J. Nielsen, *Usability Engineering*, Academic Press, 1994.
- [61] T. Erickson, "Designing conversational interfaces," 2003.
- [62] A. M. Turing, "Computing machinery and intelligence," *Mind*, vol. 59, no. 236, 1950.
- [63] J. McCarthy, "Programs with common sense," *Proceedings of the Teddington Conference on the Mechanisation of Thought Processes*, 1959.
- [64] O. Vinyals and Q. Le, "A neural conversational model," *ICML Deep Learning Workshop*, 2015.
- [65] A. Radford et al., "Language models are unsupervised multitask learners," OpenAI, 2019.