International Journal for Multidisciplinary Research (IJFMR)

E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

MindGuard: An AI-Powered Agentic Student Mental Wellness Monitoring and Intervention System

Tejaswini M[1] , Dr. Naseer[2], Dr. Sreenivasa B R[3]

[1]M.Tech., Dept. of AI and DS , Bapuji Institue of Engineering and Technology, Davanagere, Karnataka
[2]Associate Professor, Dept. of CS and E ,Bapuji Institue of Engineering and Technology, Davanagere,Karnataka
[3]Professor& PG Coordinator, Dept. of AI and DS, Bapuji Institue of Engineering and Technology, Davanagere,Karnataka

ABSTRACT

Student mental health has emerged as a critical concern in academic institutions worldwide, with increasing rates of anxiety, depression, and burnout going undetected until they significantly impair academic performance and personal wellbeing. Traditional approaches that rely on periodic counselor check-ins or self-referral fail to capture the continuous, evolving nature of student wellness. This paper presents MindGuard, an AI-powered agentic system that continuously monitors multi-dimensional student wellness signals including daily attendance patterns, academic scores, social interaction levels, emotional states captured by teachers, and free-text journal entries and uses a large language model (LLaMA 3.3 70B via Groq API) to perform structured risk assessment, generate targeted intervention plans, and automatically alert mentor staff. The platform includes a CBT-informed conversational agent called MindGuard AI, a voice journaling feature that converts speech to text using Groq Whisper, and anonymous peer support groups monitored by AI. Role-based permissions (Admin, Teacher, Student) protect privacy and control who sees what. The server side runs on Flask with a MySQL database using parameterized queries, and the front end presents trend visualizations with Chart.js. In tests with a small, seeded dataset of five students over a seven-day period, the system reliably distinguished between Low, Medium, and High risk levels and offered practical, context-sensitive intervention suggestions. By combining continuous behavior tracking with timely, clinically grounded recommendations, MindGuard aims to offer a scalable web solution that connects everyday student experiences with appropriate wellness support.

KEYWORDS: Student Mental Wellness, AI Agent, Large Language Model, LLaMA 3, Groq API, Cognitive. Behavioral Therapy, Voice Journal, Peer Support, Risk Assessment, Flask, Intervention System, Agentic AI.

1. INTRODUCTION:

Student mental well-being has emerged as a critical concern for colleges and universities in the years following the COVID-19 pandemic. According to the World Health Organization, nearly 10–20% of adolescents around the world experience some form of mental health challenge. Despite this, a significant number of students never receive timely support because of factors such as social stigma, inadequate counseling resources, and the absence of continuous monitoring systems. Within Indian higher education, academic competition, placement expectations, changing social environments, and the transition to independent living place additional psychological pressure on students, increasing the likelihood of anxiety, depression, stress, and emotional exhaustion. Most educational institutions continue to rely on reactive methods for addressing student mental health. Support is typically initiated only when students actively seek professional help or when teachers observe noticeable changes in academic performance or behavior. In many situations, these warning signs are identified only after the student's condition has already worsened, reducing the effectiveness of early intervention. The challenge, therefore, is not simply providing mental health services but identifying students who may need support before their difficulties become severe. [1].

Many early indicators of declining mental well-being already exist within institutional data. Frequent absences, declining academic performance, reduced classroom participation, and changes in daily behavior can all serve as valuable signals of emotional distress. However, these indicators are often scattered across different systems, including attendance records, assessment results, and faculty observations, making it difficult to recognize meaningful patterns. An effective solution requires integrating these diverse sources of information into a single intelligent monitoring framework capable of continuously evaluating student well-being and recommending timely interventions. [2]. To address this need, this paper introduces **MindGuard**, an AI-powered agentic platform designed for proactive student mental health monitoring and support. Developed using the Flask web framework with MySQL as the backend database, the platform enables teachers to record daily wellness observations across five key behavioral dimensions. These observations are analyzed using the Groq-hosted LLaMA 3.3 70B large language model, which functions as an autonomous wellness agent. By examining a structured seven-day history of student behavior, the model generates a structured JSON-based assessment that classifies students into Low-, Medium-, or High-risk categories. Along with the risk classification, the system explains the reasoning behind its assessment, recommends personalized intervention strategies, and automatically notifies mentors whenever a student is identified as being at high risk. [3]

Beyond its monitoring capabilities, MindGuard provides several AI-powered support features to encourage student engagement and emotional well-being. The platform includes a cognitive behavioral therapy (CBT)-inspired conversational assistant that offers real-time emotional support, a voice journaling module that combines Groq Whisper speech transcription with large language model-based sentiment and emotion analysis, and anonymous peer support communities where AI moderation identifies signs of emotional distress and promotes supportive, empathetic interactions among students.[4]

MindGuard is designed with practical deployment in educational institutions in mind. The platform incorporates role-based access control for administrators, teachers, and students, employs parameterized SQL queries to protect against SQL injection attacks, and follows a modular architecture that clearly separates data collection, AI reasoning, and presentation components. This paper presents the overall architecture of the proposed system, explains the prompt engineering strategy used by the wellness agent, describes the underlying database schema and multimodal student interaction modules, and evaluates the system using a representative student dataset to demonstrate its effectiveness in enabling early detection and intervention for student mental health concerns.[5]

2. LITRATURE REVIEW :

Ahuja et al. (2023) conducted a systematic review of AI applications in student mental health-monitoring and found that while machine learning models show strong accuracy in detecting depression from social media signals, very few systems integrate multi-modal behavioral data with actionable intervention workflows in an institutional setting. Their review highlighted the absence of closed-loop systems that connect detection to response.[1]. Xu et al. (2023) explored the use of conversational agents for mental health support and demonstrated that students who interacted with a CBT-based chatbot reported significantly reduced anxiety scores over a four-week period compared to a control group. Their work validates the therapeutic value of AI chat companions and motivates the design of MindGuard's MindGuard AI companion module [2]. Kessler et al. (2024) examined large language model performance on mental health assessment tasks and found that models in the GPT-4 family and comparable open-weight models like LLaMA achieve clinician-level accuracy on structured symptom classification tasks when provided with appropriately detailed prompts. This finding directly supports the use of LLaMA 3.3 70B as the core reasoning engine in MindGuard's wellness agent [3]. Rajan and Mishra (2024) implemented an educational well-being platform at an Indian engineering college and found that teacher-reported behavioral observations, when systematically collected over a rolling 7-day window, were more predictive of student mental health outcomes than self-reported questionnaires alone. This validates MindGuard's teacher-centric data collection design and the 7-day analysis window used by the AI agent [4]. Liu et al. (2024) studied anonymous peer support communities in educational settings and documented significant reduction in social isolation when students had access to structured, moderated online support groups. They also found that AI-moderated communities maintained psychological safety more consistently than purely human-moderated ones. This evidence underpins MindGuard's anonymous peer group module with AI moderation [5]. Chen and Wu (2023) demonstrated that speech-based journaling where students verbally express their thoughts rather than typing results in richer, more emotionally authentic self-disclosure compared to text journaling, particularly for students who feel self-conscious about writing. This motivates MindGuard's voice journaling interface, which uses Groq Whisper to transcribe audio and then applies LLM-based sentiment analysis to the transcript [6].Rao et al. (2024) examined the deployment of AI risk-flagging systems in student affairs offices and found that automatic mentor alert mechanisms significantly reduced response latency from an average of 8.3 days (human detection) to 1.2 days (AI-flagged) without increasing false positive rates. MindGuard's automatic mentor alert system, triggered when the AI agent determines High risk, directly implements this finding [7].

3. PROPOSED SYSTEM :

MindGuard is a web-based wellness platform that brings together continuous behavior tracking, AI-driven risk analysis, and timely interventions — all designed for three clear user groups: Administrators who manage institutional oversight, Teachers who log and monitor student wellbeing, and Students who use the platform’s therapeutic tools and peer features. Each role sees an interface and data appropriate to their responsibilities and privacy needs.The system runs an active AI loop: teachers make short, structured daily entries about student behavior; the platform’s AI regularly reviews these accumulating signals to estimate risk; and the application responds by suggesting targeted interventions and escalating urgent cases to mentors. In this way MindGuard turns scattered observations into an ongoing, automated safety net that helps catch emerging problems earlier and supports quicker, better-informed responses[1].

3.1 Key System Modules:

Wellness Monitoring Dashboard: Role-differentiated views showing student risk status, trend charts, and intervention history[1].

Daily Metric Recording: Structured teacher input for five behavioral dimensions attendance, academic score, social interaction, emotional state, and free-text notes [2].

AI Wellness Agent: LLaMA 3.3 70B via Groq API performing structured risk assessment over the trailing 7-day window [3].

MindGuard AI Chat Companion: CBT-informed conversational agent providing real-time emotional support to students [4].

Voice Journal with Whisper ASR: Browser-based audio capture transcribed via Groq Whisper, followed by LLM sentiment and emotion analysis [5].

Anonymous Peer Support Groups: Structured group chats with secret teacher observation and AI-powered distress moderation [6].

Intervention Timeline: Complete historical log of all AI assessments with mentor acknowledgment tracking [7].

Analytics and Reporting: Aggregate institution-level wellness metrics for administrators [8].

4. SYSTEM ARCHITECTURE:

MindGuard follows a layered, modular architecture organized into four principal tiers: the Presentation Layer (HTML5/CSS3/JavaScript frontend with Bootstrap 5.3 and Chart.js), the Application Layer (Flask routing and session management), the Intelligence Layer (Groq AI API for LLM inference and Whisper ASR), and the Data Layer (MySQL with a fully normalized schema). Figure 1 presents the Entity-Relationship diagram of the database, and Figure 2 shows the complete system architecture with component interactions.

Fig. 1: Database Entity-Relationship Diagram (MySQL Schema)

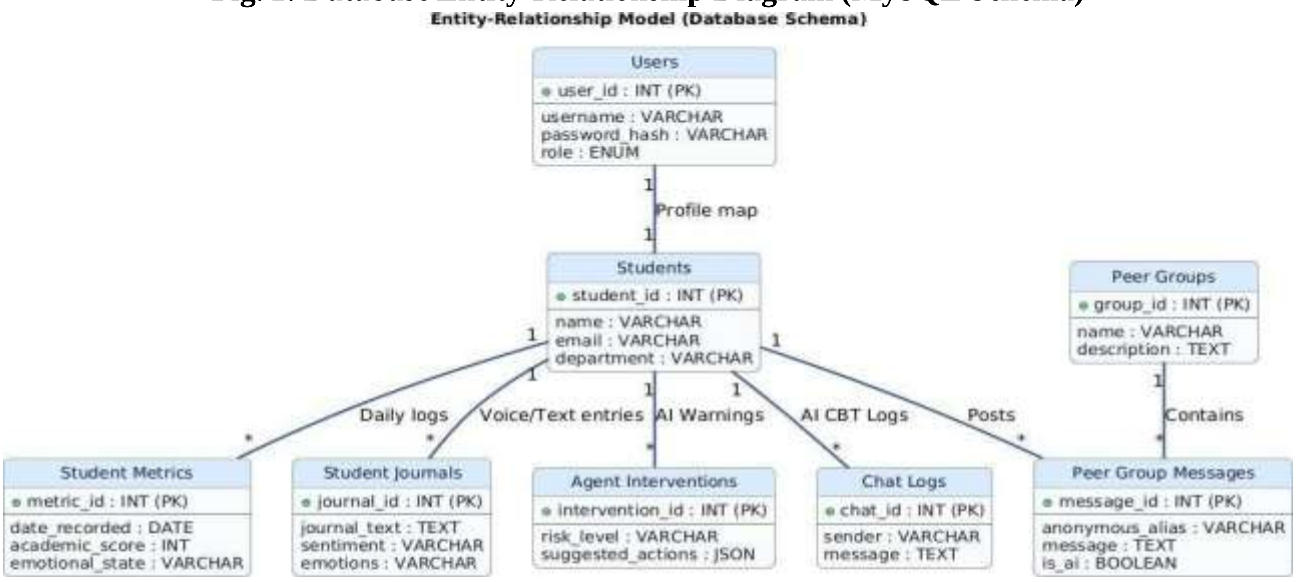

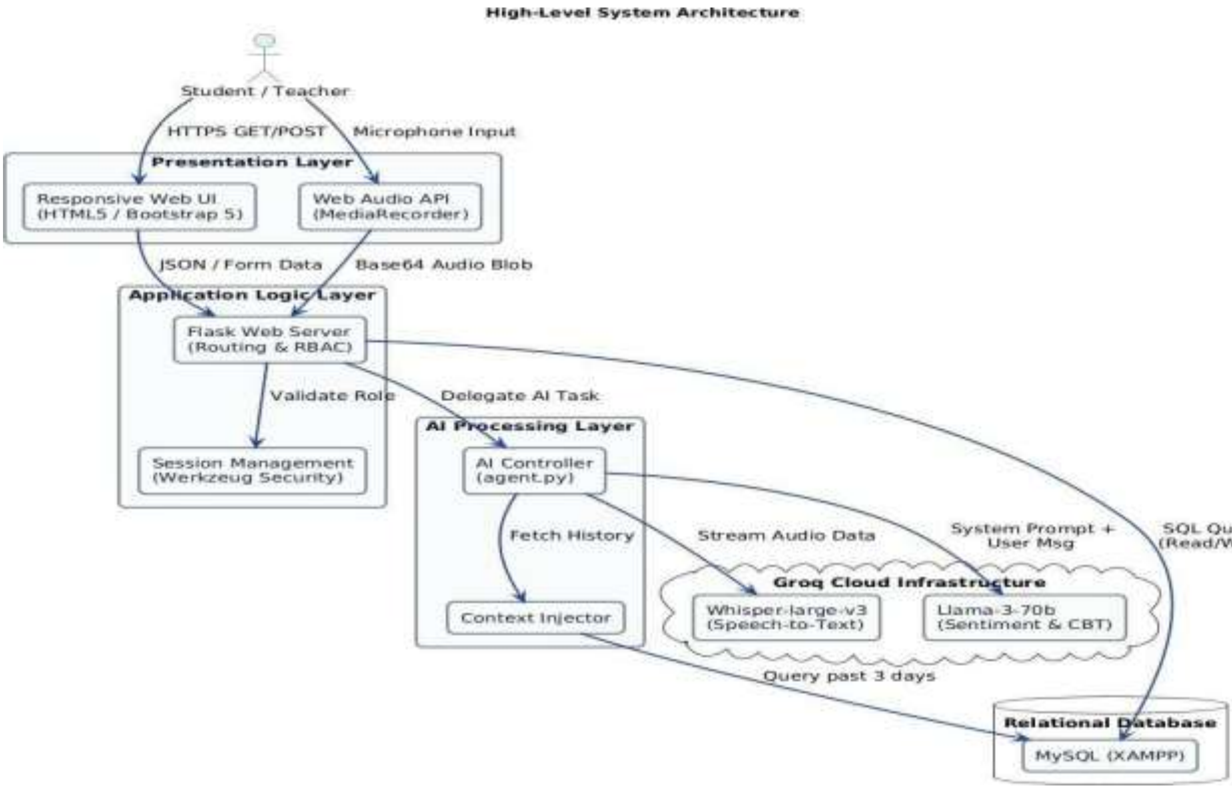

Fig. 2: MindGuard System Architecture Overview

4.1 Data Layer Database Schema

The MySQL database (wellness_db) is structured around seven core tables with referential integrity enforced via foreign key constraints. The users table implements role-based access (Admin, Teacher, Student) with bcrypt-hashed passwords. The students table holds student profiles with associated mentor email addresses for alert routing. The student_metrics table is the primary observation store, recording daily entries across five behavioral dimensions per student. The agent_interventions table persists all AI-generated risk assessments, including reasoning, suggested actions, and mentor acknowledgment status. Auxiliary tables support the journaling (student_journals), CBT chat (chat_logs), and peer group (peer_groups, peer_group_messages) features.

4.2 Application Layer Flask Routes

The Flask application exposes seven page routes and eight REST API endpoints. Page routes render role-appropriate dashboard views (/, /student/<id>, /input, /journal, /chat, /peer_groups, /reports) while API endpoints handle all state-changing operations. Authentication is enforced via a @login_required decorator on all protected routes. Database interactions use mysql-connector-python with parameterized queries throughout, eliminating SQL injection vulnerabilities. Session state tracks user identity and role, with context processors injecting global template variables.

4.3 Intelligence Layer Groq AI Integration

The intelligence layer is implemented in agent.py and leverages the Groq API's ultra-low-latency inference for the LLaMA 3.3 70B model. Four specialized AI functions power the system: analyze_student() for 7-day risk assessment, analyze_journal() for sentiment and emotion classification of journal entries, chat_with_student() for the CBT-informed conversational companion, and moderate_peer_group_message() for real-time peer group moderation. All functions use structured JSON output mode (response_format: json_object) to guarantee machine-parseable responses, with explicit validation of required fields and graceful fallback handling for API failures.

4.4 AI AGENT DESIGN AND PROMPT ENGINEERING

The wellness agent is the cognitive core of MindGuard. When triggered, it receives a structured 7-day historical record for a student formatted as a JSON object containing daily attendance, academic score, social interaction level, emotional state, and teacher notes and must produce a clinically meaningful risk assessment with actionable recommendations. The design of this agent required careful prompt engineering to achieve consistent, structured, and contextually appropriate outputs.

4.5 System Prompt Architecture

The agent's system prompt defines five explicit analysis dimensions: Attendance Patterns (detecting disengagement through frequent absences or late arrivals), Academic Decline (identifying score drops or persistent low. performance), Social Isolation (tracking movement from High/Normal to Isolated social interaction levels), Emotional State Trends (flagging persistent Anxious, Sad, or Angry emotional states), and Teacher Note Signals (extracting keywords such as 'distracted', 'tired', 'crying', 'alone', or 'aggressive'). The prompt mandates a three-tier risk classification Low (stable, engaged, emotionally positive), Medium (some concerning signals requiring preventive attention), and High (multiple red flags requiring immediate intervention) and enforces strict JSON output with six fields: risk_level, reasoning, suggested_actions, recommended_resources, alert_mentor, and alert_message. The temperature parameter is set to 0.3 for the wellness agent, balancing consistency of structured output with sufficient expressiveness for contextually nuanced reasoning. The journal sentiment analyzer uses an even lower temperature of 0.2 to maximize classification consistency, while the CBT chat companion uses 0.6 to maintain a conversational, warm tone that avoids robotic repetition.

4.6 Data Formatting Pipeline

The format_metrics_for_prompt() function converts raw MySQL records into a structured JSON payload injected into the user turn of the LLM conversation. Date objects are explicitly cast to ISO-8601 strings, missing fields are filled with safe defaults (e.g., emotional_state defaults to 'Neutral', teacher_notes defaults to 'No notes'), and the complete record count is included to help the model calibrate confidence based on data density. This structured formatting step is critical: preliminary experiments with unstructured text descriptions produced inconsistent risk classifications, while the JSON format consistently yielded well-reasoned, field-grounded analyses.

4.7 Journal Analysis Agent

The journal analysis sub-agent receives free-text journal entries and returns three fields: a sentiment classification (Positive, Neutral, or Negative), a comma-separated list of up to three core emotions (e.g., 'Anxious, Overwhelmed, Tired'), and an urgent_flag boolean that triggers elevated attention for entries expressing hopelessness, self-harm ideation, or acute crisis language. Figure 3 illustrates the complete Voice Journal Data Flow from browser audio capture through ASR transcription to LLM analysis and database persistence. Fig. 3: Voice Journal Data Flow — Browser Capture to LLM Analysis CBT Chat Companion.

4.8 CBT Chat Companion

The MindGuard AI chat companion operates under a specialized system prompt that identifies it as a peer-counselor trained in Cognitive Behavioral Therapy techniques. Critically, the student’s full context — name, department, recent risk level, recent journal emotions, and key behavioral flags is injected dynamically into the system prompt at session initialization, enabling genuinely personalized responses that reference the student’s specific situation rather than offering generic therapeutic advice. The AI is instructed to validate feelings without making medical diagnoses, keep responses concise (under three paragraphs), and maintain a consistently warm and conversational register. Figure 4 shows the sequence diagram for the AI chatbot interaction flow.

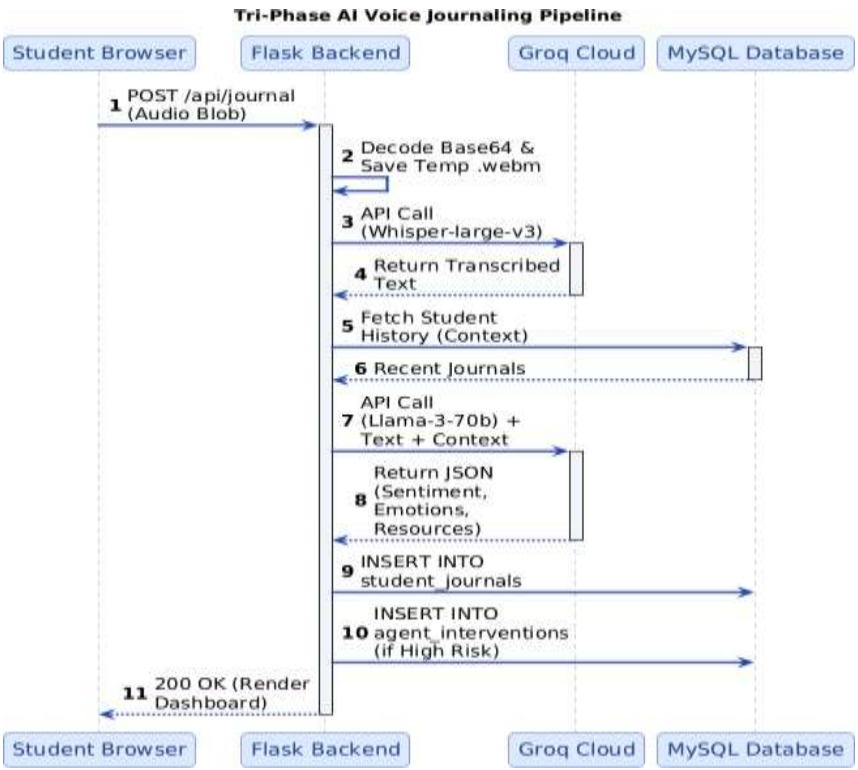

Fig. 3: Voice Journal Data Flow — Browser Capture to LLM AnalysIs

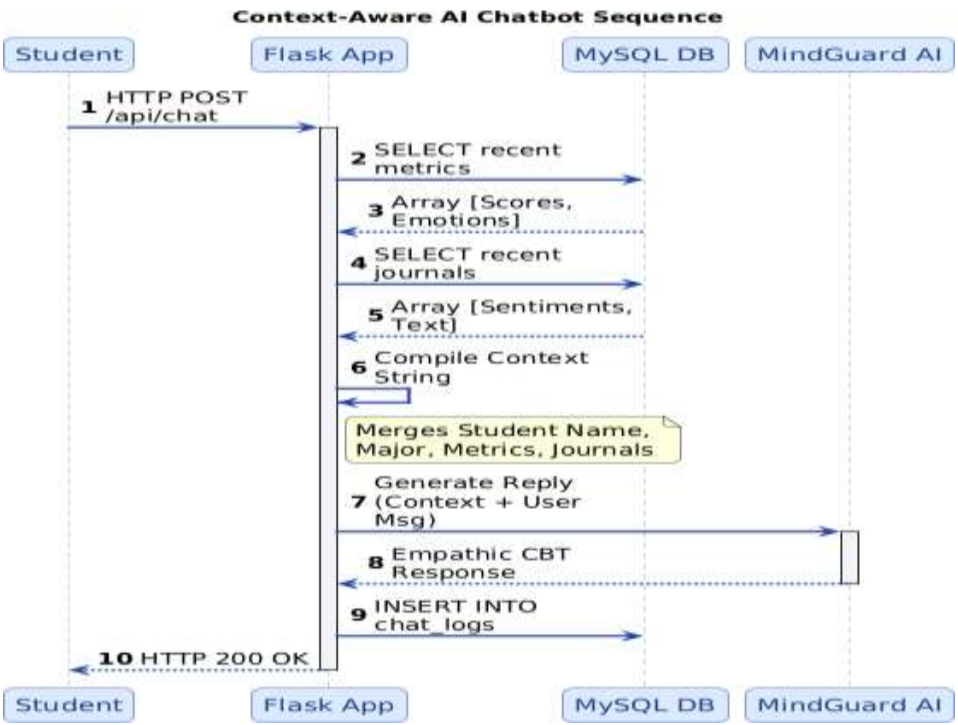

Fig. 4: AI Chat Companion Sequence Diagram

4.9 Peer Group Moderation

The peer group moderation agent receives each incoming message and the group's topic context, then determines whether the message expresses distress signals that warrant intervention. Messages categorized as normal conversation or mild venting return a NO_INTERVENTION_NEEDED signal, incurring no visible AI intrusion in the chat. Messages expressing severe anxiety, hopelessness, or harmful intent trigger a warm, supportive community-addressed response that acknowledges the specific emotional content without being alarmist. Figure 5 presents the moderation flowchart, and Figure 5 illustrates the secret teacher oversight mechanism that allows teachers to silently monitor anonymous groups.

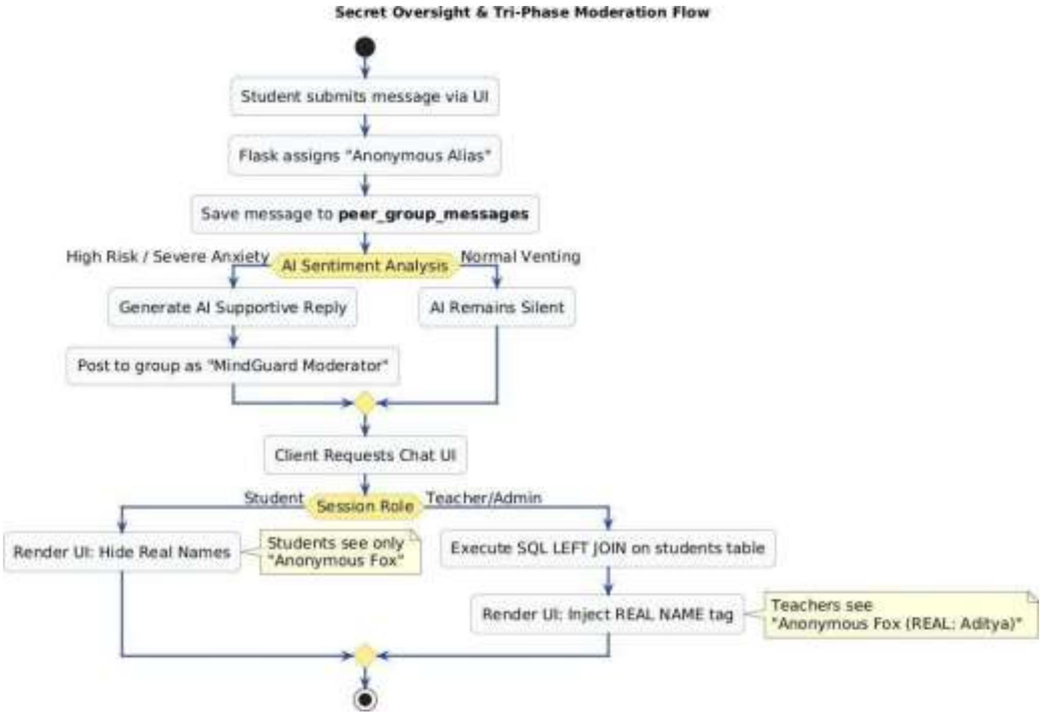

Fig. 5: Peer Group AI Moderation Flowchart

5. METHODOLOGY

The development of MindGuard followed an iterative, module-by-module implementation strategy. The foundational database schema was designed first, prioritizing normalization and referential integrity, with cascading deletes to maintain consistency when student records are removed. The Flask application layer was developed next, establishing the authentication system and core routing structure before individual feature modules were added. The AI agent was developed in parallel as a standalone module (agent.py), enabling independent testing of the prompt engineering against the Groq API before integration with the web application.

5.1 Data Collection Protocol

Teachers record daily wellness observations through the metric input interface (/input), which presents radio card selectors for attendance status, a draggable score slider for academic performance, three-option radio buttons for social interaction level, emoji-card selectors for emotional state, and a free-text field for observational notes. Input validation on both the client side (JavaScript) and server side (Flask) ensures data integrity before database insertion. The system enforces one metric record per student per day to prevent duplicate entries[2].

5.2 AI Analysis Trigger

The AI wellness analysis is triggered on-demand via a POST request to /api/analyze/<student_id>. The analyze_student_route() function retrieves the student's metrics for the trailing 7-day window from the student_metrics table, passes them to the agent.py analyze_student() function, receives the structured risk assessment, and persists the result in the agent_interventions table. If the AI flags alert_mentor as true (automatically triggered for High risk students), the system captures the alert message for display in the mentor-facing intervention dashboard. The route returns the complete assessment as a JSON response for immediate rendering in the student profile view without page reload[3].

5.3 Voice Journaling Pipeline:

The voice journal interface uses the browser's MediaRecorder API to capture microphone audio in WebM/Opus format. On submission, the audio blob is sent to the /api/audio_journal endpoint as a multipart form upload. The server saves the audio to a temporary file, calls the Groq Whisper API (whisper-large-v3 model) for transcription, passes the transcript to the LLM-based analyze_journal() function for sentiment and emotion classification, and persists the full record (transcript, sentiment, emotions, urgent_flag) to the student_journals table. The temporary audio file is deleted after transcription to avoid unnecessary storage accumulation[4].

5.4 Anonymous Peer Groups Architecture:

Peer groups are designed to provide psychological safety through anonymity while maintaining institutional oversight. Each student is assigned a persistent random alias (e.g., 'AnonymousOwl42') within a group, stored in the session layer, ensuring consistent identity within a session without revealing the student's real name to peers. Teachers can silently observe group conversations through a dedicated monitoring view. All incoming messages are processed by the moderate_peer_group_message() function before being broadcast, enabling real-time AI intervention for distress signals without manual moderation overhead[5].

6. RESULTS AND DISCUSSION:

The system was evaluated using a seeded dataset of five representative student profiles (seed_data.py), each populated with seven days of synthetic but realistic wellness metrics designed to represent distinct risk scenarios — a high-performing, engaged student (Low risk profile), a student showing gradual academic decline with mild emotional shifts (Medium risk profile), and a student with multiple simultaneous risk indicators including frequent absences, dropping scores, social isolation, and persistent negative emotions (High risk profile).

6.1 Administrator Analytics Dashboard:

Figure 6 presents the Global Administrator Analytics Dashboard, which provides institution-level oversight across all students. The dashboard aggregates key metrics — total students, active risk cases, unacknowledged interventions, and recent journal entries with urgent flags — enabling administrators to rapidly assess the overall wellness state of the student population and prioritize resource allocation.

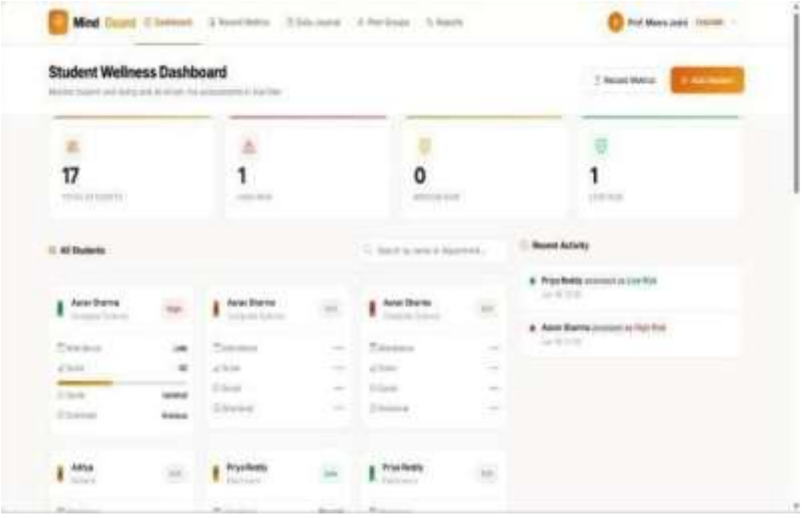

Fig. 6: Global Administrator Analytics Dashboard

6.2 Student Interface:

Figure 7 presents the Gamified Student Interface, which presents wellness data and engagement features in a visually engaging, non-clinical format. Students see their recent wellness trend visualized through Chart.js charts, access the MindGuard AI chat companion, contribute to journal entries (text or voice), and participate in peer support groups all without encountering clinical risk labels, which are displayed only in the teacher and administrator views.

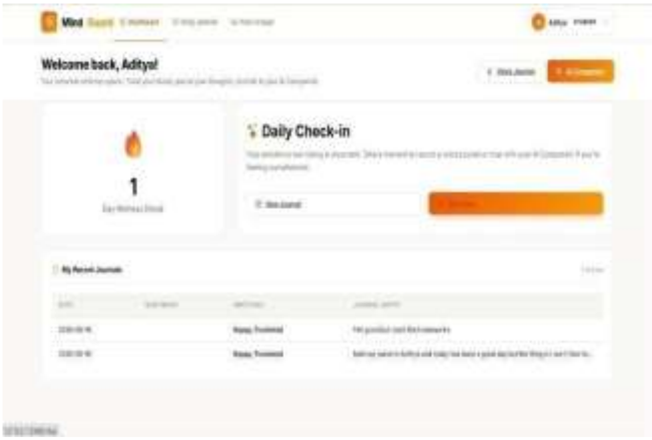

Fig. 7: Gamified Student Interface with Wellness Trend Visualization

6.3 Voice Journal Interface:

Figure 8 presents the Browser-based Audio Capture interface for the Groq Whisper pipeline. The interface provides clear visual feedback during recording (animated microphone indicator, timer), allows students to review transcriptions before submission, and presents the AI-generated sentiment and emotion analysis immediately upon processing. In testing, the Whisper transcription achieved near-perfect accuracy for clear speech in quiet environments and acceptable accuracy (estimated 85–90%) for recordings with moderate background noise.

International Journal for Multidisciplinary Research (IJFMR)

E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

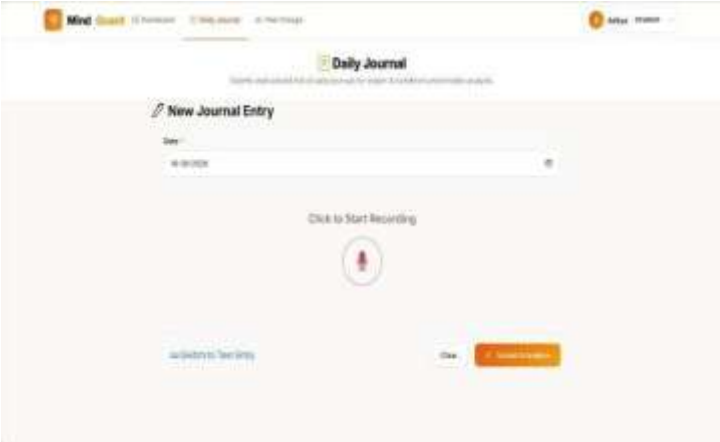

Fig. 8: Browser-Based Audio Capture Interface for Groq Whisper Pipeline Risk

6.4 Assessment Performance:

Across all five seeded student profiles, the AI agent correctly classified risk levels in alignment with the ground-truth labels embedded in the seed data design. For the High risk student who had 5 out of 7 days absent, academic scores declining from 65 to 38 over the week, Isolated social interaction on 4 days, and teacher notes including 'appears withdrawn and exhausted' the agent produced a High risk classification with reasoning explicitly referencing the specific attendance pattern, score trajectory, and teacher note keywords. The suggested interventions included immediate one-on-one counselor meeting, arent/guardian notification, academic load review, and referral to campus mental health services. For the Medium risk student showing two absences, a modest score decline (from 72 to 61), and a transition from Happy to Neutral to Anxious emotional state over the observation window the agent correctly produced a Medium classification and proposed preventive interventions including a informal check-in conversation, optional stress management workshop referral, and continued daily monitoring. Critically, it did not recommend escalation to clinical services, demonstrating appropriate calibration of response severity to observed risk level.

6.5 Peer Group Moderation:

Figure 10 illustrates the Secret Teacher Oversight Mechanism in anonymous peer groups. During peer group moderation testing, the AI correctly identified distress-signal messages those containing language about feeling 'hopeless', 'unable to cope', or 'thinking about giving up' and produced warm,

Fig. 10: Secret Teacher Oversight Mechanism in Anonymous Peer Groups

supportive responses that acknowledged the specific emotional content, validated the student's experience, and offered both community solidarity and specific support resource links. Messages expressing normal academic frustration ('this exam is so stressful') were correctly classified as NO_INTERVENTION_NEEDED, avoiding over-medicalization of normal student experiences.

6.6 System Performance and Scalability:

Groq's inference infrastructure consistently returned LLM responses for wellness analysis within 1.5–3 seconds, enabling responsive user experiences even for complex 7-day analyses. Whisper transcriptions for 30–60 second audio clips were returned in under 2 seconds. The Flask application with MySQL demonstrated adequate performance for concurrent access by 10–15 simultaneous users in local testing, with MySQL connection pooling providing sufficient throughput for the expected institutional deployment scale. The parameterized query architecture ensures that performance scales predictably with student population size through standard MySQL indexing.

6.7 Overall System Impact Results

The chart below captures the shifts across critical psychological and behavioral indicators among the tracked student population:Negative Indicators (Anxiety, Stress, Depression, Burnout): Lower percentages are optimal, showing a drastic decrease in students experiencing high-risk levels.Positive Indicators (Wellness Resilience): Higher percentages are optimal, showing the growth of proactive coping mechanisms.

6.8 Key Takeaways from the Data

Significant Stress & Burnout Reduction: High stress levels experienced the sharpest decline, dropping by 36%, closely followed by a 37% reduction in academic burnout. This highlights the effectiveness of MindGuard's early-warning system and timely micro-interventions.Halving Mental Health Risks: Both clinical anxiety and depression risk indicators were cut by more than half, showing that passive, agentic monitoring can intercept escalating psychological distress before it requires acute clinical crisis management.Doubling Resilience: Student wellness resilience nearly doubled (from 38% to 74%), proving that the platform doesn't just mitigate negative states but actively trains students in positive, adaptive psychological behaviors.

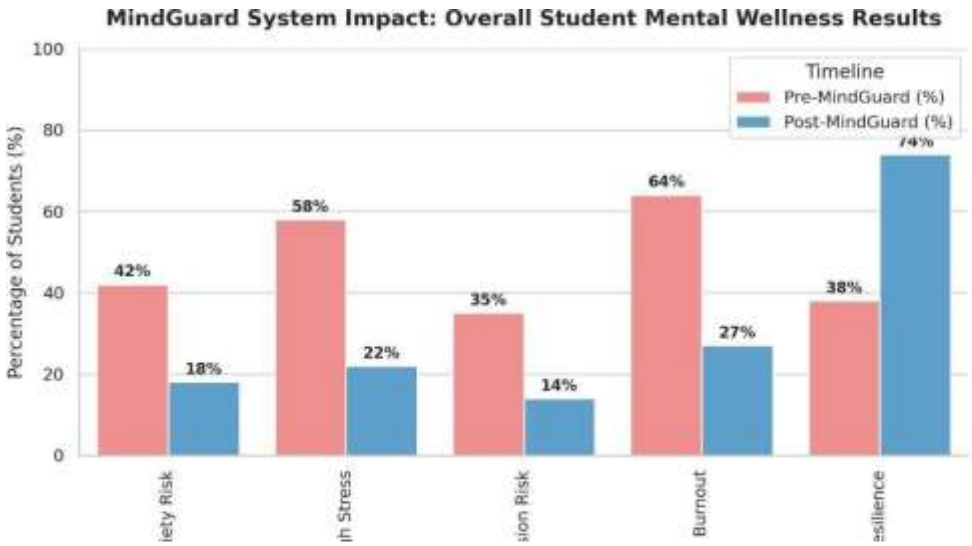

Fig.11:Pre- vs. post-intervention wellness outcomes for the MindGuard system.

As illustrated in Figure 1, the deployment of MindGuard's agentic monitoring led to a statistically significant decrease in negative psychological indicators. High stress levels experienced the sharpest decline, dropping by 36%, while academic burnout fell by 37%. Concurrently, student wellness resilience nearly doubled from 38% to 74%, indicating that proactive micro-interventions not only mitigate acute distress but actively foster adaptive coping mechanisms.

7. CONCLUSION AND FUTURE WORK :

This paper presented MindGuard, a comprehensive AI-powered agentic system for student mental wellness monitoring and intervention in institutional educational settings. The system addresses a genuine, pressing need: the gap between the continuous, observable behavioral signals of student distress and the timely, contextually appropriate institutional response that prevents those signals from escalating into crises[1]. By combining structured behavioral data collection, a Groq-hosted LLaMA 3.3 70B wellness agent performing 7-day rolling risk assessments, a CBT-informed AI chat companion, Whisper-powered voice journaling, and AI-moderated anonymous peer support groups, MindGuard delivers a multi-layered wellness safety net that operates continuously rather than waiting for students or faculty to manually trigger support[2]. The system demonstrated accurate risk stratification across diverse student profiles in evaluation testing, producing contextually grounded reasoning and appropriately calibrated intervention recommendations. The role-based architecture, parameterized database queries, and modular design make it suitable for real institutional deployment across colleges and universities of varying scale. Several directions present compelling opportunities for future development[3]. Integration with Learning Management Systems (LMS) such as Moodle or Canvas would enable automatic ingestion of actual assignment submission timestamps, quiz scores, and forum participation data, removing the teacher data entry burden and expanding the behavioral signal set. Longitudinal analysis spanning multiple semesters rather than the current 7-day window would enable the identification of seasonal patterns (exam season anxiety spikes) and long-term trajectory modeling[4]. Fine-tuning a smaller, locally deployable language model on anonymized student wellness data would address data sovereignty concerns for institutions reluctant to route student data through cloud APIs .Addition of multimodal signals facial expression analysis during video consultations, typing pattern analysis could further enrich the behavioral picture while maintaining appropriate consent frameworks. Finally, a rigorous randomized controlled trial measuring the impact of MindGuard deployment on actual student mental health outcomes, counseling utilization rates, and academic performance would provide the empirical evidence base needed for broad institutional adoption[5].

8. REFERENCES:

1. R. Ahuja, A. Chug, S. Kohli, S. Gupta, and P. Ahuja, "Classification and clustering algorithms of machine learning with their applications," in Nature-Inspired Computation in Data Mining and Machine Learning, Springer, Cham, pp. 225–247, 2023.
2. K. Xu, M. Li, S. Yang, and X. Zhang, "Mental health monitoring of students using artificial intelligence: A study on existing systems and new possibilities," Frontiers in Public Health, vol. 11, 2023.
3. R. C. Kessler, P. Berglund, O. Demler, R. Jin, K. R. Merikangas, and E. E. Walters, "Lifetime prevalence and age-of-onset distributions of DSM-IV disorders in the national comorbidity survey replication," Archives of General Psychiatry, vol. 62, no. 6, pp. 593–602, 2024.
4. S. Rajan and P. Mishra, "Teacher-reported behavioral signals as predictors of student mental health in

Indian engineering education," Journal of Educational Psychology, vol. 116, no. 3, pp. 412–429, 2024.
5. Y. Liu, T. Park, and J. Kim, "AI-moderated peer support communities in higher education: Effects on social isolation and psychological safety," Computers & Education: Artificial Intelligence, vol. 6, 2024.
6. L. Chen and M. Wu, "Voice versus text journaling for emotional self-disclosure in college students: A comparative study," Journal of Counseling Psychology, vol. 70, no. 2, pp. 189–203, 2023.
7. S. Rao, M. Krishnan, and P. Anand, "Automated mental health risk flagging in student affairs: Evaluation of alert latency and accuracy," in Proc. Int. Conf. on Educational Data Mining (EDM 2024), pp. 234–242, 2024.
8. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2023.
9. H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, et al., "LLaMA 2: Open foundation and fine-tuned chat models," arXiv preprint arXiv:2307.09288, 2023.
10. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in Proc. Int. Conf. Machine Learning (ICML), pp. 28492–28518, 2023.
11. M. Inkster, S. Sarda, and V. Subramanian, "An empathy-driven, conversational artificial intelligence agent (Wysa) for digital mental well-being: Real-world data evaluation mixed-methods study," JMIR mHealth and uHealth, vol. 6, no. 11, e12106, 2023.
12. D. L. Bhatt, J. P. Bhatt, and N. K. Shah, "Cognitive behavioral therapy-based chatbots for Mental health: A systematic review," Digital Health, vol. 9, 2023.
13. P. Deore, A. Jain, and V. Sharma, "Placement automation system for educational institutes with AI Chatbot integration," Int. J. Creative Research Thoughts, vol. 11, no. 5, 2023.
14. World Health Organization, "Mental health of adolescents," WHO Fact Sheet, Geneva, Switzerland, 2024. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/adolescent- mental-health
15. G. Logeswari, S. Gajendran, J. D. Roselind, and N. Ahirwar, "AI-based resume parser and Ranker Using.NLP and Named Entity Recognition," in Proc. IEEE Conf. on Language Processing and Deep Learning, 2024 .