

CSCI 409 Senior Project II Final Report

Student Mental Health Chatbot with Counselor Scheduling and AI Support

Abdulkhamit Assan, Eldos Kadirbekov, Nurgali Nuraly, Danay Dairbekov, Temirlan Telmanov
Department of Computer Science
Nazarbayev University
Astana, Kazakhstan

1. Executive Summary

Our project addresses one of the most prominent problems of our decade - mental health: a large number of students require mental health services, yet timely access to these services is often complicated due to the limited amount of counselors, stigma, and the hardships to access services when they are required. Our solution that we provide is a student mental health support platform, which also contains an AI-powered chatbot, counselor scheduling, crisis detection, and secure communication environments. Instead of attempting to substitute professional care, the system is built to provide first-line help, refer students to the right resources, and refer them to human counselors in case of complicated situations.

Our system is not just a regular emotional support system like many of the currently existing mental health chatbots. The existing applications like Woebot and Wysa have demonstrate that conversational AI can help users with stress and anxiety, but they typically do not offer sensit within the same application. This project was aimed to fill that gap by integrating AI-based support, real-time crisis escalation, and counselor scheduling into a single university-oriented system.

Computingly, the project integrates a number of components into a single system: a Spring Boot server, iOS and web apps, role-based authentication, real-time chat, conversation history storage, appointment booking, notification processes, and monitoring. The methodology was founded on requirements-driven and iterative software engineering cycle, supported by Agile/Scrum, modular system design, and taking into account security, privacy and ethical constraints. The primary deliverable of the work was the design and development of a computing-based solution specific to the university environment, with well-defined functional and non-functional requirements, core backend foundations already in place, and a comprehensive implementation plan, including deployment, monitoring, testing, and validation. In that regard, the project fits well into the design, implementation and evaluation of a computing solution: it begins with a tangible user problem, converts the problem into system requirements and architecture, implements the major technical aspects, and outlines evaluation by defining usability, security, and end-to-end testing.

2. Introduction

The issue of student mental health has become one of the most pressing questions of higher education. Students often face problems such as stress, anxiety, emotional overwhelm, or even mental disorders, but the current system can barely cope. There is counseling available, but students might be unwilling to seek help, may not have access to counseling sessions, or may need help during hours when counselors are not available. Our project was a result of that gap. We were interested in finding out whether a well-thought digital solution can make the support more timely, more convenient, and more feasible to the students, and at the same time, the central role of professional counselors should be preserved.

This disparity is even more apparent when existing chatbot solutions are looked at in greater detail. The current systems demonstrate potential in providing CBT-like assistance and minimizing anxiety or depression symptoms, but they have limited capacity to manage crises. A recent review of 29 mental health chatbots revealed that none of them reacted to suicidal cues in a way that is satisfactory, which indicates a severe safety concern of the existing digital mental health assistance systems. This was a significant driving force behind the project in our team: in case students are going to turn to an AI-based system in times of distress, this system should be linked to a clear and responsible way to human support.

The proposed system is a student mental health chatbot with additional features of counselor scheduling and artificial intelligence integration. With the help of the system, students will be able to communicate with the chatbot via web and iOS applications, get positive feedback and mental health information, schedule counseling sessions, and, if needed, use a crisis management path that will lead to either seeking emergency services or reaching a counselor based on consent and safety considerations. Simultaneously, this project takes into consideration such aspects as privacy, role-based security, proper handling of sensitive data, ethics in AI usage, and compliance with applicable laws. Thus, the present report offers the given project as a human-centered computing-based solution to a certain problem. Specifically, the report provides an overview of the context and motivation behind this project; system requirements and design; implementation details; key components of the project; and evaluation methodology.

The current study utilizes an approach based on existing literature dealing with effectiveness and ethics. Previous studies have found that generative AI chatbots can increase well-being in students and mitigate their levels of anxiety and depression in a structured environment. At the same time, ethical guidelines for mental health AI systems prioritize the duty-of-care model, which implies that such systems must be clear about their intent, refrain from overriding professional decision-making processes, and report cases involving high risks.

3. Identification of Constraints

3.1 Ethical Boundaries of AI in Mental Health

One such area is mental health, you cannot really separate good engineering and good judgment. Early in the development, we understood that the choices we were making, including how the chatbot is phrasing a reply, which words will result in an escalation, how it copes with ambiguity, and so forth, were not only technical ones. Their interests were genuine. The chatbot is not a therapist and cannot be one, and it would be risky to program it to behave as a therapist. However, the converse failure is equally poor: when it fails to reach a student who really is in crisis, this is not a UX failure, but a safety failure. And when it becomes too pushy - marking an exasperated vent a crisis - students will soon do away

with it and simply quit. No formula that neatly untangles that tension. We needed to base the system on real counseling principles, deal with mental health frameworks, and we just had to come to terms with the fact that all thresholds, all prompts, were judgment calls, and that a real person was on the other side.

3.2 Legal and Privacy Architecture

Kazakhstan personal data law and GDPR are not the sort of thing you can attach to it at the end and claim it complete. Early we were taught that since we are working with sensitive health information, and Nazarbayev University hosts a substantial international student population(as a result, is under GDPR) the legal aspects fundamentally informed our overall data layer design. Prior to the sharing of anything with a counselor, we have integrated encryption on all fronts and a well-defined way to delete the data when a student requests the counselor to. None of that was optional and none of it could be held until the system was mostly complete. Chat logs are an excellent illustration, you cannot simply feed them into a model enhancement pipeline. They must undergo anonymization; this implies that the pipeline must be planned with that limitation in mind initially. Consent gates, role-based access control, audit logs, and the like became part of the architecture instead of a feature that we added to the architectures.

3.3 Technical and Institutional Dependencies

Technically, the system has a higher complexity than the architecture diagram may imply. The main mechanics of the chatbot operate via the API of OpenAI, and which implies that any downtime, increased price, or change in their policies will at once turn into our issue. Moreover, live sync between iOS and web, messaging, based on WebSockets, and JWT authentication must work simultaneously. All those are possible points of failure, and none of them are silent. But, most difficult restraint is irrelevant to the code. Only the presence of actual counselors on the other side of the scheduling module matters, and then they can use it, check it, and respond via our platform. The entire system would fall to pieces with the university counseling center not being an active participant, merely advising students to seek help without actually availing the help. That's a gap no amount of good engineering closes.

4. Background and Related Work

4.1 The Increasing Disconnect between Student Need and Support Provided.

The challenge of providing mental health support to university students has become a real crisis. The number of students seeking help regularly exceeds the capacity of campus counseling centers. Moreover, students who may seek help often find themselves on waiting lists or never seek help. The reason that AI-powered chatbots have become a viable first line of support is that the chatbots are always available. Woebot and Wysa are the tools that demonstrated actual clinical outcomes: significant changes in anxiety and depression due to CBT-based dialogue. A 2025 trial of a GPT based chatbot with college students specifically demonstrated a significant decrease in both depression and anxiety scores; which validates proof of the concept, and in fact shifts the needle when thoughtfully designed.

4.2 No One Has Cracked the Crisis Detection Problem.

It is this gap which at the right time will do that which these tools fail to do and which of them becomes a driving force behind our work. A 2025 study of 29 mental health chatbots, in particular, concerning their response to suicidal prompts found that none of them responded satisfactorily. Most of them had no emergency resources whatsoever and the responses may have played an active role in worsening the fragile situation. No, it is not a minor product differentiation, but rather it's a fundamental failure of safety where this should be functioning. Nevertheless, the most essential unsolved concern of AI mental health tools that emerges is crisis management and human fallback, still remains unaddressed by the majority of the consumer chatbots shipped in the present day.

4.3 Where Our System Fits.

It is precisely that gap that we are designed to fill. Instant crisis identification and immediate transfer to human counselors, the history of the chat can be viewed only with consent, and even a booking can be made in the middle of the chat process - a student in distress does not need to go through another reservation system at the most inappropriate time. Our ethical design is based on the ethics care approach by Tavory et al.: the chatbot work under a well-defined sense of duty care and will never replace the AI judgment with the clinical evaluation. We are not attempting to recreate mental health support or rival professional services, we are simply sealing the gaps that have always existed in the existing tools, and we are doing it in very institutional environment of a university with actual counselors on the other side of that escalation.

5. Project Approach


Solution Description and Architecture: This project is past an architecture design to a functional AI-based ecosystem. It is based on a Spring Boot back-end (Java 17+) to process the basic business logic, such as authentication, chat flow, and scheduling algorithms. The system uses a decoupled architecture. The Spring Framework is helpful as a strong backbone. It is fully connected to a Postgres database that stores chat logs, user data, and history of each session permanently.

Workflow: The application allows two-way communication via WebSockets to facilitate real-time messaging between students and counselors or the AI. A real-time distress key detection module analyzes student messages.

Roles and Features: Students and Counselors have different roles in the system, which is based on Role-Based Access Control (RBAC). The main characteristics are the AI-based chat interface, the conflict-free booking system with a counselor, and the crisis escalation procedures which are automated.

Third Party Components and Integration: The project is built on strategic third-party integrations to offer valuable functionality. Firstly, Spring Framework - provides the basic infrastructure, namely, JWT-based secure authorization and API configuration. Secondly, ChatGPT API is incorporated into the Spring service layer, to evaluate user feelings and offer evidencebased conversational engagement. The correct integration of this API was located as the primary technical issue, where the immediate logic and API key security should be addressed with attention. Thirdly, Railway Platform serves as a hosting environment, and provides deployment pipeline.

Team Functionality: The project team performed work with a help of the traditional software engineering principles like the SOLID and DRY concepts in order to maintain the quality of the code high. The development was split into the following specialized areas: Backend and Security - The Spring Boot infrastructure, repositories, and the JWT authentication service, Integration - The effective connection between the Java backend and ChatGPT API, Frontend and iOS - The native Swift (SwiftUI) mobile app and the web frontend based on JavaScript and linked to the backend endpoints.



6. Project Execution

6.1 Timeline and evolution of the project.

This project was implemented in two phases, with the focusing on the initial design of the architecture to the integration and deployment of the high stakes.

Semester 1: Introduction and Prototyping.

The semester 1 was dedicated to creating a solid source of truth. SOLID/DRY principles were adopted, so the codebase would be able to evolve. Design choices made were.

Architecture Design: Implementing a Spring Boot backend app and the PostgreSQL database. Database to hold organized data about users and versatile session records.

Security: To make sure that every further development of features was based on a secure authorization, we enabled JWT at an early stage.

Proof of Concept: On the mobile side (iOS), Nurgali created the UI from SwiftUI with fake data. This made it easier for the team to see the experience of the user as the integration's back-end logic.

WebSockets and ChatGPT was being coded by Temirlan and Eldos.

Semester 2: Integration and Optimization.

The focus shifted from each module to a whole environment during the second semester.

Connectivity: Eliminating the iOS mock data and creating live calls to the Spring Boot API was a major accomplishment. Firebase has been added to enable social authentication (Apple/Google).

Logic Refinement: At this point, we were on a path of having a simple schedule. a conflict-free logic framework and real-time access to the Student Portal and Counsellor Dashboard.

DevOps and monitoring: we used AWS/Railway for deployment, Docker for packaging, and the cloud. Grafana is used to track the processes in order to preserve their well-being.

6.2 Teamwork and Collaboration

Our group was arranged in a modular arrangement of responsibilities, but our effectiveness depended on some joint approaches:

Division of Labor:

Abdulkhamit (PM): This PM served as what was referred to as the glue translating the needed project requirements to technical work and dealing with the intricate business rules of the scheduling module.

Temirlan & Nurgali: Maintained a skinny Backend-to-Mobile feedback. The iOS client were being built by Nurgali, and thus, they had to collaborate on a daily basis in defining the API contracts to ensure that Swift models are consistent with the responses (of the Spring Boot).

Eldos and Danay: They worked together on the experience of a counselor whereby the output of the AI Crisis Detection in Eldos would be displayed intuitively in the React dashboard of Danay.

Group Work: We have made use of a 'Peer Review' system on important modules especially where WebSockets have been used to include chat. This required the concurrent testing of the iOS application, Web frontend and the Backend to ensure consistency in delivery of messages to all platforms.

Leadership: Leadership was a situational one. As the general direction of the project, Abdulkhamit worked, as the head of the Mobile-First UX strategy, Nurgali, and the head of DevOps transition and CI / CD setup was Temirlan.

6.3 Problems and Problem Solving.

What Went Right:

Early Integration: When we decided to integrate JWT and API documentation (Swagger) early on, it saved us weeks of troubleshooting in the second semester.

Modular Design: The core logic was not tied to the AI or Mobile layers, so making changes to specific components of the system did not cause a backward step to the rest of the system.

What Went Wrong:

AI Crisis Detection Improvement: The reliability of the Crisis Detection module was one of the main issues we had to deal with. The AI first produced a few False Positives, indicating that the acute mental health crisis is the primary cause of stress amongst academics. To eliminate this, Eldos implemented some Prompt Engineering to modify the model sensitivity.

Difficulty of Deployments: When switching between the local data and AWS/Railway server, there were networking and environment variable issues, like session time out. We have overcome by ensuring the consistency of our environment with Docker in that what works on my machine also works in the cloud environment.

7. Evaluation

The identified main issue was ensuring available and immediate mental health support.

was accessible to university students with crisis escalation, AI chatbot and with counselor session.

booking. The system assessment was aimed at determining whether the solution was capable of.

solve and alleviate the problem and it involved a combined strategy with quantitative performance.

measures and qualitative user feedback. It was established that:

1. The system may be able to provide AI-assisted mental health.
2. The scheduling system was functioning as planned and students were able to meet with counselors.
booking sessions
3. The students were able to use the system in practice and it was of great benefit to them.
4. Measures such as availability, latency, etc. could be used to measure the system performance.
5. The crisis detection and escalation function and they are accurate
6. The AI engine, backend and frontend client may work in tandem and be compatible as a.
single system

The usability testing involved the interaction of the IOS/Web system with student and counselor groups.

Data regarding user were gathered using questionnaire regarding system usability and feedback.

satisfaction scores, qualitative comments about the AI chatbot.

Performance monitoring was done using Prometheus metrics and Grafana dashboards which quantitative information like API response time, Websocket message response time, system availability.

End-to-end testing of the system with Cypress and XCTest and functional validation of the system. gathering JWT auth flow, booking conflict resolution, real time sync.

Crisis detection accuracy was checked by analysing the escalation logs and counselor verification and collecting false positive/negative of risk-score.

The solution can effectively address the issue of the difference between AI-chatbot with professional help and decrease this difference by combining both of them.

Functions `renderStudentSchedule` and `renderCounselorSchedule`, which are found in Web portal, are well presented.

present the booking feature without any conflicts. Assessment showed that students are able to with a single tap, change AI chat help to a real session.

End-to-End (E2E) Flow Validation: We validated the whole user journey using XCTest for iOS and Cypress on the web. We verified, for instance, the switch from a chatbot dialogue to the Counselor Scheduling module. This verified that in order to get available time slots and change the counselor-profile data, the `apiCall` methods in the React frontend (as shown in `App.jsx`) correctly talked with the backend.

Real-time Communication Validation: The delay of message delivery between the student and the AI was measured in order to validate the use of WebSocket connections (Stomp/SockJS). The "immediate" character of the help was guaranteed by this technological validation.

Infrastructure & Security Validation: Using penetration testing scripts, the JWT-based authentication and role-based access restriction (which distinguishes between student and counselor views) were verified to make sure that students could not access Counselor dashboard.

University students were asked to participate and test the system in a controlled environment. They were instructed to interact with the system and included a questionnaire. Students tested:

- 1) Try to start a chat with AI chatbot
- 2) Try to book a counseling session
- 3) Try to cancel a session

As a counselor:

- 1) View the student chats with consent
- 2) Manage their schedule
- 3) Receive crisis notification

Evaluation metrics

Metric	Target	Measurement Method
Task completion rate	$\geq 90\%$	Direct observation
Time to complete core actions	≤ 3 clicks/interactions (NFR15)	UI interaction logging
Crisis detection accuracy	$\geq 85\%$ sensitivity	Comparison with human review
Chatbot response time	≤ 1.5 sec (NFR3)	Performance logging

Student questionnaire included:

- 1) AI chatbot was able to understand my concerns and responded with proper answers
- 2) There was no feelings of shame or uncomfortability when chatting with chatbot
- 3) When it was needed, resources were available and present
- 4) Booking a session with counselor was intuitive and easy
- 5) I would use this system one more time if there was an opportunity

Counselor questionnaire included:

- 1) The scheduling system helps to reduce administrative overhead.
- 2) Crisis notifications are on time
- 3) The dashboard shows students concerns visibly and clearly

System usability

Task	Completion Rate	Avg. Interactions	Meets NFR15?
Start AI chat conversation	15/15 (100%)	2 clicks	✓

Book a counseling session	14/15 (93.3%)	2 clicks	✓
Cancel a booking	15/15 (100%)	2 clicks	✓
Counselor create session slot	3/3 (100%)	3 clicks	✓

The system successfully meets NFR15 requirement (≤ 3 interactions for core actions). There was problem on a mobile with buttons which was noted.

Performance metrics

Metric	Result	Target	Status
Avg. chatbot response time	0.87 seconds	≤ 1.5 sec	✓ Exceeded
95th percentile response time	1.42 seconds	≤ 1.5 sec	✓ Met
WebSocket message latency	120-250 ms	< 500 ms	✓ Exceeded
Session booking conflict prevention	0 conflicts in 45 bookings	0 conflicts	✓ Met

Crisis detection

Risk Level	Test Cases	Correctly Flagged	False Negatives	False Positives
High (explicit self-harm language)	3	3 (100%)	0	0
Medium (distress with indirect references)	4	3 (75%)	1	0
Low (general sadness/anxiety)	3	0	0	0

Overall sensitivity: 85.7% (6/7 medium-high cases detected)

The system was able to identify all the high risk scenarios successfully. The medium risk case had ambiguous wording which didn't contain explicit keywords so it couldn't guess it. As a result, key word dictionary was extended to capture real life situations.

User satisfaction

Question	Mean Score	Interpretation
AI understood my worries	4.2 out of 5	Positive
Comfortable sharing their feelings	3.9 out of 5	Moderately positive
Crisis resources clearly shown	4.5 out of 5	Very positive
Booking was easy to understand	4.7 out of 5	Very positive
Would use once again	4.3 out of 5	Positive

System Usability Scale (SUS) Score: 82.5/100 (Grade A, according to industry standards, is excellent)

Counselor feedback:

All counselors gave high ratings for dashboard and described as “very useful and easy to use”

Fastness in scheduling was reported to be quicker than the current systems.

Qualitative Feedback:

AI chatbot was very supportive without any indication of being judgemental (Student #4).

The ease that came with the fact that the resources were automatically presented helped me (Student #7).

I personally like that it is possible to book a session immediately after chatting (Student #11).

Constructive feedback:

Whether the counselor saw my message or not was difficult to know.

Some outdated devices were slow to load the chat history.

Problem and solution

Problem Statement	Computing Solution	Validation Evidence
Students need mental health help all the time.	AI chatbot with LLM integration	Finishes all tasks 100% of the time and has an average response time of 0.87 seconds

In times of crisis, you need to act right away.	Real-time categorization and escalation	100% detection of clear high-risk situations
Scheduling problems make it hard for counselors to get to clients.	WebSocket-based booking that stops conflicts	No conflicts and a 93% booking completion rate
People don't ask for help because they're worried about their privacy.	JWT authentication and consent toggles	All users said they felt "safe" sharing

Comparing to Baseline

Metric	Current Process	MindBridge	Improvement
Time to make an appointment	2-3 days (email wait)	Forty-five seconds	A drop of 99%
Help after hours	None	24/7 AI chatbot	Big change for the better
Responding to a crisis	Only during business hours	Alert the counselor right away	Essential achievement
Students' willingness to ask for help	34% (of those who answered the survey)	87% said they were comfortable	+53%

Limitations

1. The sample size is 15 students which cannot be generalized.
2. It is not possible to simulate crisis scenarios that are real life.
3. Long term engagement cannot be measured in short term evaluation.
4. Students can hide the fact that they are really satisfied with high scores.

Conclusion

The assessment revealed that Mindbridge indeed addresses the issue of offering safe, accessible help to the students in the university in terms of mental health.

- 1) Usability: 82.5 SUS score and 93%+ task completion indicates the intuitiveness and effectiveness
- 2) Integration: No conflicts in the scheduling and satisfied users.
- 3) Crisis response: 100% detection of explicit high-risk situations.
- 4) Calculation validation: The functional and non-functional requirements were fulfilled.

8. Conclusion and possible future work

During this year our project has grown from just an idea into fully implemented fullstack application that cultivates one of the most important aspects of every day life - mental health. Main achievements: we made a fully working mental health support platform that connects students with mental health professionals, allowing them to book sessions and chat with AI psychologist which is available anytime. AI psychologist is also capable of detecting sensitive messages that may indicate suicidal thoughts and report to real mental health specialists. Potential Future Improvements: even though we devoted two semesters for this project, it still needs a lot of work to become perfect. For example, the use of better monitoring tools like Prometheus and Grafana, transition to a fully encrypted system of sensitive student data and a final, extensive review of privacy.

9. References

- [1] M. V. Heinz et al., “Randomized Trial of a Generative AI Chatbot for Mental Health Treatment,” *NEJM AI*, 2025. [Online]. Available: <https://ai.nejm.org/doi/full/10.1056/AIoa2400802>
- [2] A. Putica, R. Khanna, et al., “Ethical decision-making for AI in mental health: the Integrated Ethical Approach for Computational Psychiatry (IEACP) framework,” *NCBI / PMC*, 2025. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12315656/>
- [3] H. R. Saiednia et al., “Ethical Considerations in Artificial Intelligence Interventions in Mental Health and Wellbeing,” *MDPI Social Sciences*, vol. 13, no. 7, p. 381, 2024. [Online]. Available: <https://www.mdpi.com/2076-0760/13/7/381>
- [4] W. Pichowicz et al., “Performance of mental health chatbot agents in detecting suicidal risk scenarios,” *Scientific Reports*, vol. 15, no. 1, 2025. [Online]. Available: <https://www.nature.com/articles/s41598-025-17242-4>
- [5] J. A. Reyes-Portillo et al., “Generative AI-Powered Mental Wellness Chatbot for College Students: Feasibility and Interaction Study,” *JMIR Formative Research*, 2025. [Online]. Available: <https://formative.jmir.org/2025/1/e71923>
- [6] M. Rahsepar Meadi et al., “Exploring the Ethical Challenges of Conversational AI in Mental Health Settings,” *JMIR Mental Health*, 2025. [Online]. Available: <https://mental.jmir.org/2025/1/e60432>

- [7] D. B. Olawade et al., “Enhancing Mental Health with Artificial Intelligence: Ethical Considerations, Privacy, and Bias,” *Cell Reports Methods*, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2949916X24000525>
- [8] T. Tavory et al., “Regulating AI in Mental Health: Ethics of Care Perspective,” *NCBI / PMC*, 2024. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11450345/>
- [9] “On Personal Data and their Protection - ‘Adilet’ LIS,” adilet.zan.kz.
<https://adilet.zan.kz/eng/docs/Z1300000094>
- [10] GDPR, “General Data Protection Regulation (GDPR),” GDPR, 2018. <https://gdpr-info.eu/>