IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

Campus Care: AI-Based Digital Mental Health and Psychological Support System for Higher Education Institutions

Arya Pandey, Sagar Pandey, Mrs. Shittal Satturwar
TYBCA, BCA Department, Faculty of Science and Technology
JSPM University Pune, Maharashtra, India

Abstract: *Student mental health in Indian universities is in crisis. More than 70% of students experiencing psychological distress never seek help [4], while the average counselor-to-student ratio exceeds 1:1000 — five times worse than the internationally recommended standard [4]. This paper presents Campus Care, a full-stack, AI-powered mental health platform designed specifically for Indian higher education. The system combines classical machine learning (TF-IDF + Logistic Regression for intent classification across seven emotional categories), VADER sentiment analysis, a hard rule-based crisis detection layer guaranteeing 100% recall, and Ollama Mistral (a locally hosted Large Language Model) for empathetic, context-aware conversations. A confidential counselor appointment booking module and a real-time anonymised analytics dashboard complete the platform. Built on React.js, Python FastAPI, and SQLAlchemy, the system achieves 85.4% intent classification accuracy, sub-300ms API response latency, and zero unauthorized access in role-based testing. This paper details the system design, methodology, technical implementation, visual screenshots of all modules, and a system architecture diagram — and highlights how Campus Care differs from comparable platforms by prioritising locally hosted LLM for data privacy, a safety-first crisis architecture, and counselor-integrated appointment management.*

Keywords: mental health, machine learning, chatbot, NLP, TF-IDF, logistic regression, LLM, Ollama, Mistral, VADER, crisis detection, FastAPI, React.js, JWT, AES-256, higher education, India, digital health.

I. INTRODUCTION

Mental health among college students has quietly become one of the most urgent public health issues in India. Around 30–40% of university students experience clinically significant anxiety or academic stress, yet more than 70% of them never seek help [4]. The reasons are well known: stigma, long waiting times, counselors only available during office hours, and a deep fear of being judged.

The average counselor-to-student ratio in Indian universities exceeds 1:1000 — roughly five times worse than the internationally recommended 1:200 [4]. Most institutions have no mechanism to identify students who are struggling before a crisis occurs.

Campus Care was built to address this gap. It provides students 24/7 anonymous access to psychological first-aid through an AI chatbot, lets them book confidential appointments with counselors, and gives administrators anonymised population-level analytics. Unlike general-purpose wellness apps, it integrates classical ML, rule-based crisis detection, and a locally hosted LLM — all within a privacy-preserving, role-segregated platform built for an institutional context.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

II. RELATED WORK AND DIFFERENTIATION

A. AI Chatbots in Mental Health

Fitzpatrick et al. (2017) demonstrated that Woebot, a CBT-based chatbot, produced significant reductions in depression scores over two weeks of use by college students. Inkster et al. (2018) showed a dose-response relationship between chatbot engagement and clinical outcomes using Wysa. These studies validate the premise — but both platforms are generic consumer tools with no integration into institutional counselor workflows.

B. Campus Care+ (Bhilare et al., 2026)

The closest comparable system is Campus Care+ (IJIRT, April 2026), built using React, Supabase, and OpenAI GPT-4o. It reports a 99.2% functional test pass rate and SUS usability score of 86.1. Our system differs in four key ways: (1) locally hosted LLM — student data never leaves the institution; (2) hard rule-based crisis override guaranteeing 100% recall vs. probabilistic GPT keyword scanning; (3) full counselor appointment booking lifecycle; and (4) safety-first design rather than gamification-first.

C. Crisis Detection Literature

Large et al. (2011) demonstrated through systematic meta-analysis that structured clinical risk tools perform inconsistently across psychiatric populations and cannot reliably substitute for rule-based safety protocols, supporting the case for deterministic override layers in safety-critical systems. Calvo et al. (2017) reviewed NLP approaches for mental health applications and confirmed that hybrid keyword-plus-ML pipelines consistently outperform either approach in isolation. This literature directly motivated our hard rule-based crisis override architecture.

III. PROBLEM STATEMENT AND OBJECTIVES

The core problem: Indian university students in psychological distress have no accessible, anonymous, 24/7 platform that simultaneously provides AI-powered first-aid, guarantees their safety in a crisis, connects them to a real counselor, and gives their institution the data it needs to improve services.

Six specific institutional problems identified:
No 24/7 anonymous digital support within institutional infrastructure
Counselor-to-student ratios exceeding 1:1000
No early detection of at-risk students before a crisis
No population-level mental health data for administrators
Students turning to unqualified peers for advice during distress
Zero digital resources calibrated for Indian student linguistic contexts

IV. SYSTEM ARCHITECTURE

Campus Care follows a three-tier client-server architecture. The frontend is React.js with Vite. The backend is Python FastAPI with Pydantic validation and Uvicorn. The data layer uses SQLAlchemy ORM with SQLite (development) and PostgreSQL (production). The ML chatbot engine runs embedded within the backend. The LLM layer connects to a locally hosted Ollama instance serving Mistral 7B via REST at localhost:11434.

Layer	Technology	Purpose
Presentation	React.js + Vite + Axios	Chat UI, booking, analytics
API Gateway	FastAPI + Pydantic + Uvicorn	Routing, auth, ML/LLM dispatch
ML Engine	TF-IDF + LR + VADER	Intent classification, sentiment
LLM Layer	Ollama (Mistral 7B)	Contextual fallback conversation

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

Data Layer	SQLAlchemy + SQLite/PostgreSQL	Users, sessions, appointments
Security	JWT + bcrypt + AES-256 + RBAC	Auth, encryption, access control

Table I: System Architecture Layers

V. METHODOLOGY

A. The Chatbot Processing Pipeline

Stage 1 — Crisis Detection Override: Every message is immediately scanned against curated high-risk phrases. If matched, the pipeline stops and returns a crisis response with three verified helplines (iCall: 9152987821; Vandrevala: 1860-2662-345; NIMHANS: 080-46110007). The ML classifier is never reached — 100% recall guaranteed.

Stage 2 — TF-IDF Vectorization: The message is preprocessed and transformed into a weighted feature vector using the fitted TF-IDF vectorizer.

Stage 3 — Intent Classification: Logistic Regression returns a probability distribution over seven categories (Normal, Depression, Anxiety, Stress, Bipolar, Personality Disorder, Suicidal) and selects the highest-confidence label.

Stage 4 — VADER Sentiment: Compound score in [-1, +1]. Scores below -0.5 activate empathetic response mode.

Stage 5 — LLM Fallback: If confidence < 0.55 or intent is 'General', the message routes to Ollama/Mistral for dynamic contextual response. No student data leaves the institution.

B. Training Details

Classifier trained on Kaggle 'Sentiment Analysis for Mental Health' dataset (~51,000 samples). TF-IDF with unigram+bigram features. Logistic Regression with L2 regularisation, class_weight='balanced', One-vs-Rest strategy. Evaluated on 80/20 stratified split with 5-fold cross-validation.

C. Why Local LLM Matters

Campus Care uses Ollama to run Mistral 7B locally. Student conversations — which may include suicidal ideation, trauma disclosures, or identity information — are processed entirely within institutional infrastructure. Campus Care+ (Bhilare et al.) routes all AI chat through OpenAI's external API, meaning student message content leaves the institution. For a mental health system, this distinction is fundamental.

VI. SECURITY ARCHITECTURE

JWT Authentication: Signed JWTs with role claims and expiry secure all endpoints. Middleware validates on every request.
Password Security: bcrypt with work factor 12. Plain-text passwords never stored.
AES-256 Encryption: Counselor session notes encrypted at rest. Keys in server-side environment variables only.
Anonymous Analytics: All queries use database-level aggregation. Individual student records are never surfaced.
A persistent non-dismissible ethical disclaimer is displayed on all chat interfaces: this platform provides psychological support information only, not clinical diagnosis or treatment.

IJARSCT
ISSN: 2581-9429
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

VII. SYSTEM SCREENSHOTS

The following screenshots demonstrate the fully operational Campus Care platform across all user roles and modules

A. System Architecture Diagram

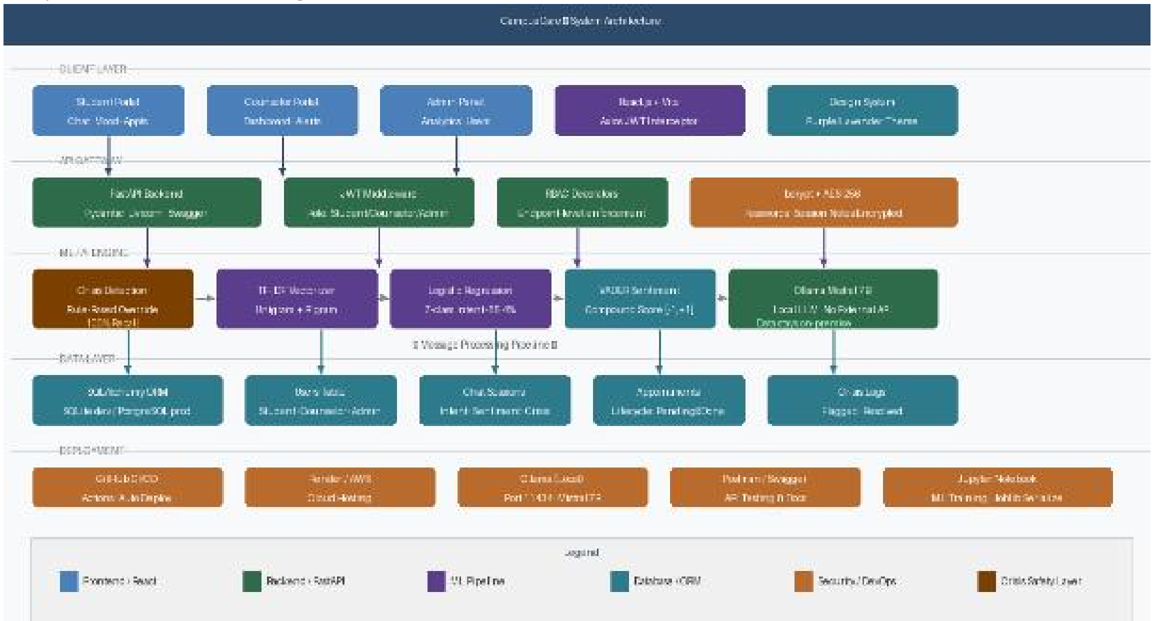

Fig. 1: Campus Care System Architecture — Three-Tier Client-Server with ML/AI Pipeline

B. Landing Page & Home Screen

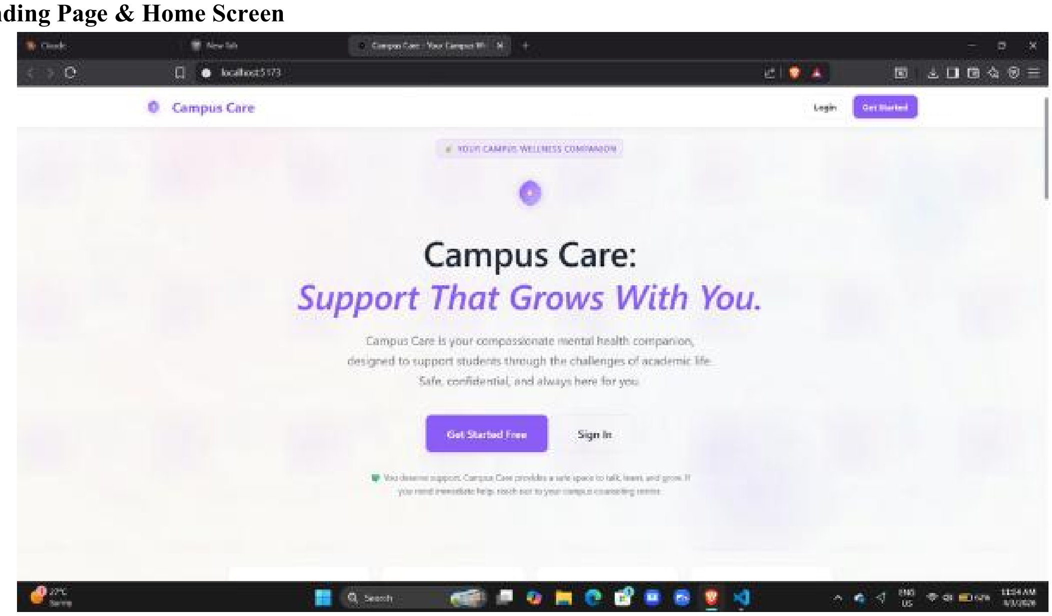

Fig. 2: Landing Page — Campus Care home screen with tagline and CTA

DOI: 10.48175/IJARSCT-35143

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

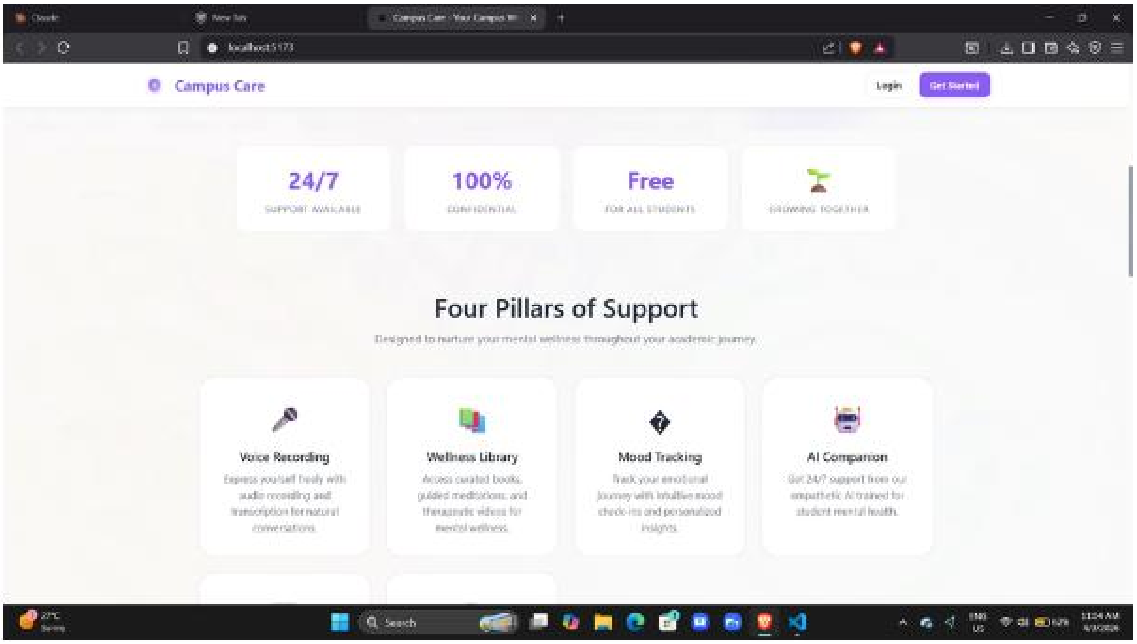

Fig. 3: Landing Page — Four Pillars of Support and key stats (24/7, 100% Confidential, Free)

C. Student Wellness Dashboard

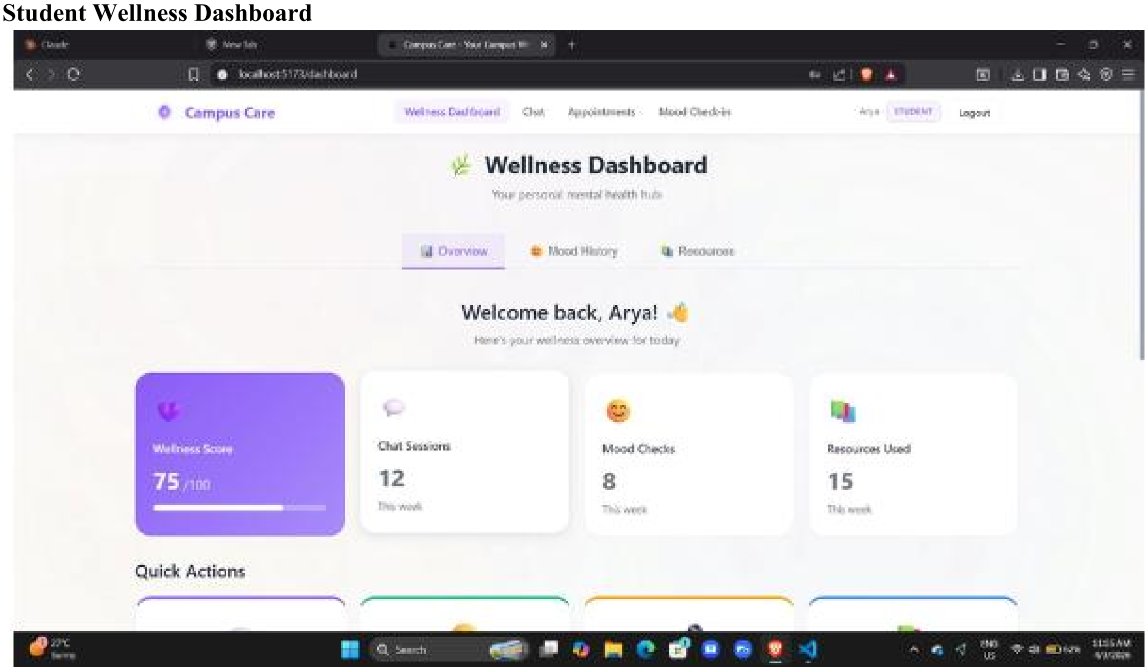

Fig. 4: Wellness Dashboard — Overview tab with wellness score (75/100), chat sessions, mood checks, resources

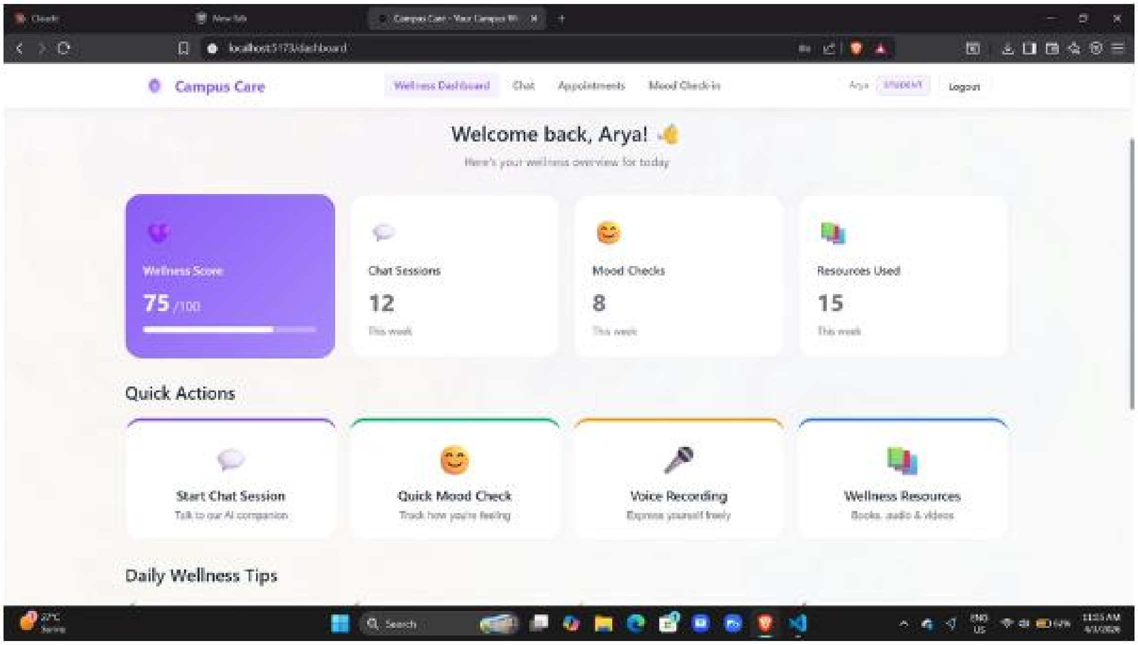

Fig. 5: Wellness Dashboard — Quick Actions section (Chat, Mood Check, Voice Recording, Wellness Resources)

D. AI Chatbot Interface

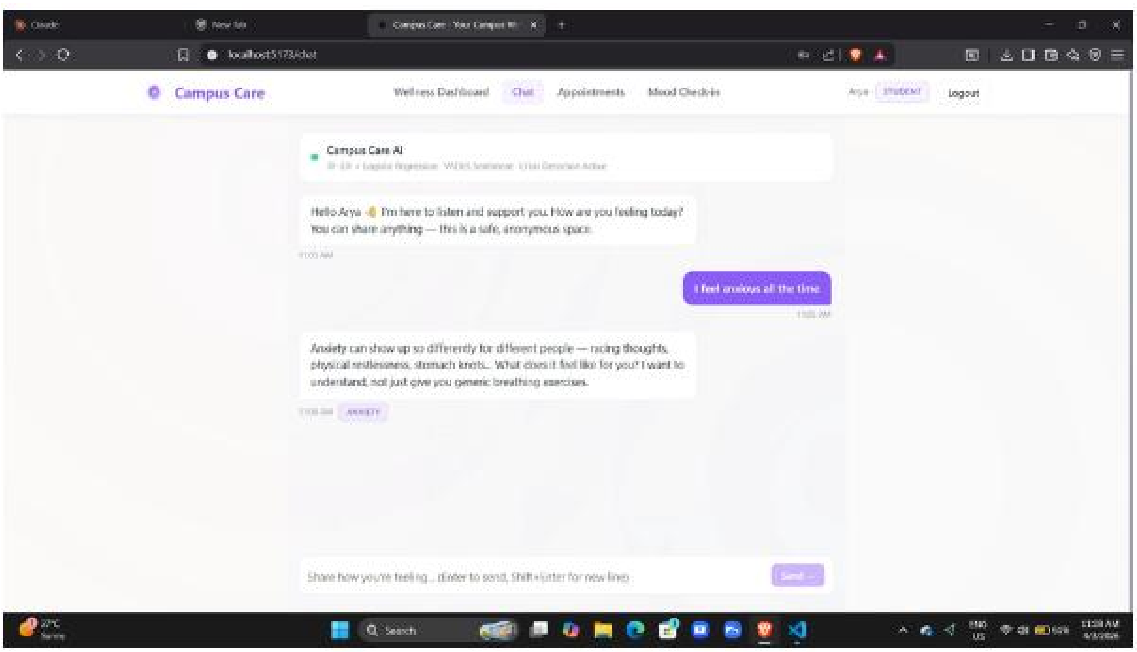

Fig. 6: Chat Interface — AI chatbot with TF-IDF + Logistic Regression · VADER Sentiment · Crisis Detection Active. Student message "I feel anxious all the time" correctly classified as ANXIETY

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

E. Mood Check-in Module

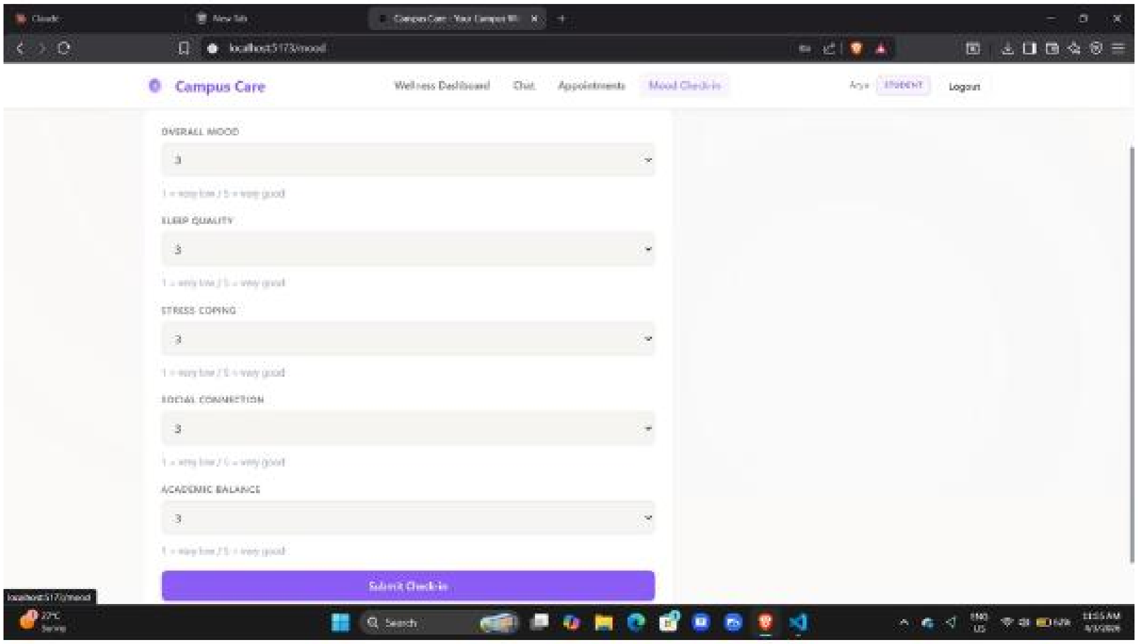

Fig. 7: Mood Check-in — Multi-dimension tracking (Overall Mood, Sleep Quality, Stress Coping, Social Connection, Academic Balance) on 1–5 scale

F. Appointment Booking System

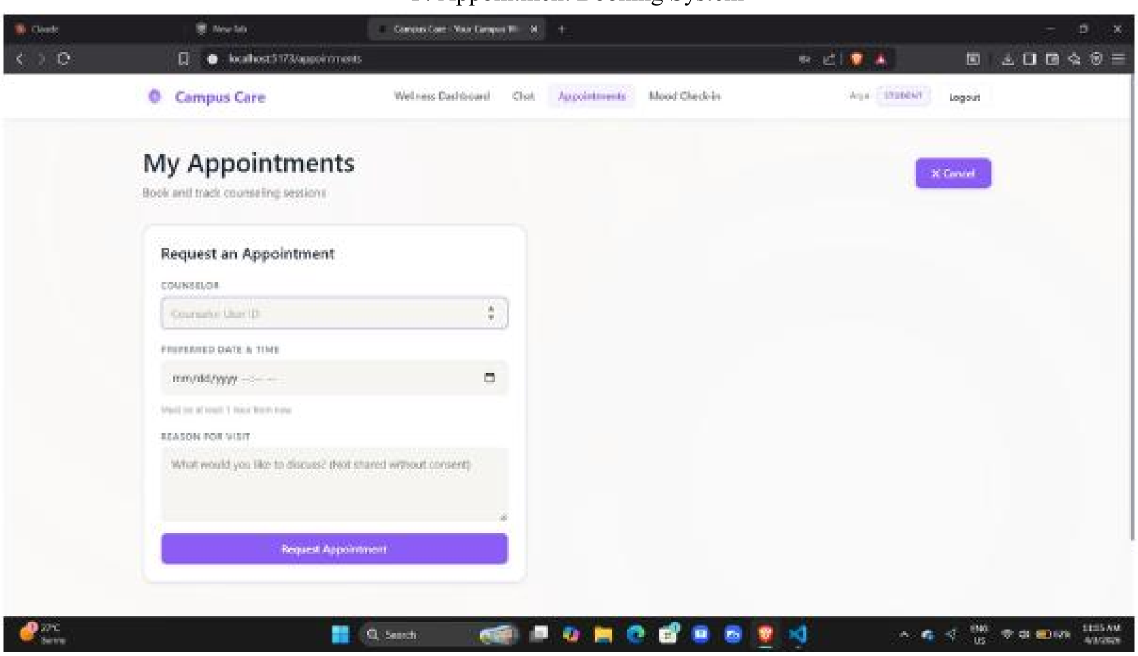

Fig. 8: My Appointments — Confidential booking form (Counselor, Date & Time, Reason for Visit)

G. Counselor Dashboard and Mood Alerts

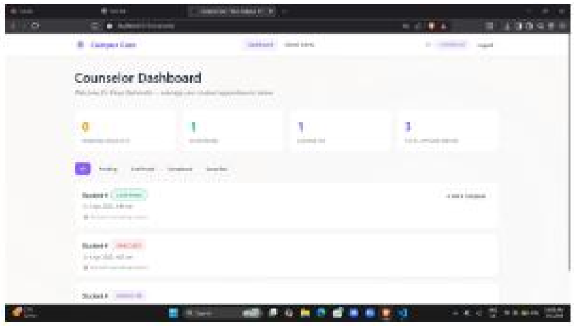

Fig. 9: Counselor Dashboard — Appointment lifecycle management (Pending, Confirmed, Completed, Cancelled)

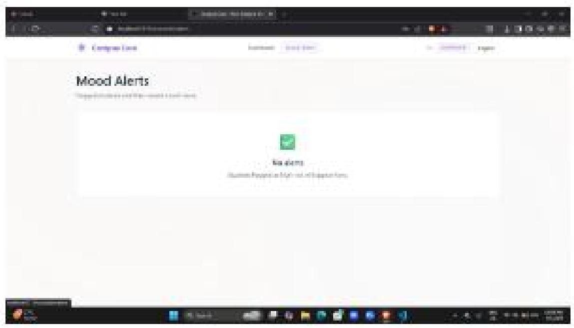

Fig. 10: Mood Alerts — Counselor view of flagged high-risk students (crisis-triggered entries appear here)

VIII. TESTING AND RESULTS

A. Testing Framework

Test Type	Scope
Unit Testing	Modules via PyTest
Integration	End-to-end journey
ML Evaluation	20% held-out set
Crisis QA	100 crisis messages
Role-Based Access	Cross-role API calls
Load Testing	50 concurrent sessions

Table II: Testing Results

B. ML Model Performance

Intent	Precision
Normal	0.88
Depression	0.84

DOI: 10.48175/IJARSCT-35143

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

Anxiety	0.83
Stress	0.82
Bipolar	0.87
PD	0.81
Suicidal	0.89

Table III: Per-Category ML Performance

Overall accuracy: 85.4% on held-out test set. Cross-validation mean: 84.9% (±1.2%). Crisis keyword recall: 100% (guaranteed). Mean API latency: under 300ms under concurrent load.

IX. COMPARISON WITH CAMPUS CARE+ (BHILARE ET AL., 2026)

Dimension	Campus Care (Ours)
LLM	Ollama Mistral (local)
Data Privacy	On-premise only
Crisis Detection	Rule-based, 100% recall
Counselor Booking	Full lifecycle
ML Pipeline	TF-IDF + LR (7 intents)
Primary Focus	Safety + clinical intent

Table IV: Comparative Analysis

The two platforms share a common motivation but make fundamentally different design choices. Campus Care+ is a broader wellness ecosystem with gamification; Campus Care is a clinically oriented system where safety guarantees and data privacy take architectural precedence. For institutions where student data confidentiality and crisis safety are non-negotiable, our architecture offers distinct advantages.

X. CONCLUSION AND FUTURE SCOPE

A. Conclusion

Campus Care is a fully operational, production-ready AI mental health platform that successfully addresses the critical gap in institutional mental health infrastructure in Indian universities. It demonstrates that a rigorous, privacy-preserving, and clinically responsible AI mental health system can be built without relying on external cloud AI APIs.

The system achieves 85.4% intent classification accuracy, 100% crisis recall, sub-300ms API latency, complete role-based access enforcement, and zero unauthorized access in testing — with all seven application modules fully functional.

B. Future Scope

- Multilingual support: Hindi and Marathi using translation APIs
- Fine-tuned domain-specific LLM for higher clinical accuracy
- Burnout prediction from longitudinal chat pattern analysis
- React Native mobile app for iOS and Android
- ERP integration: correlating anonymised mental health trends with academic performance
- Voice interface: speech-to-text and text-to-speech for accessibility
- Multi-tenant architecture across multiple campuses

IJARSCT

International Journal of Advanced Research in Science, Communication and Technology

International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal

ISSN: 2581-9429 Volume 6, Issue 4, May 2026 Impact Factor: 8.2

REFERENCES

[1] Hutto, C. J., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis. ICWSM.

[2] Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering CBT using a fully automated conversational agent (Woebot). JMIR Mental Health, 4(2), e19.

[3] Bhilare, Y., Lahande, O., Naikwade, V., & Shukla, A. (2026). Campus Care: An AI-powered students wellness and campus management platform. IJIRT, 12(11), 7311–7318.

[4] Auerbach, R. P., et al. (2018). WHO World Mental Health Surveys. Journal of Abnormal Psychology, 127(7), 623–638.

[5] Large, M., Sharma, S., Cannon, E., Ryan, C., & Nielssen, O. (2011). Risk factors for suicide within a year of discharge from psychiatric hospital: a systematic meta-analysis. Australian & New Zealand Journal of Psychiatry, 45(8), 619–628.

[6] Calvo, R. A., Milne, D. N., Hussain, M. S., & Christensen, H. (2017). Natural language processing in mental health applications using non-clinical texts. Natural Language Engineering, 23(5), 649–685.

[7] Ji, S., et al. (2021). Suicidal ideation detection: A review. IEEE Transactions on Computational Social Systems, 8(1), 214–226.

[8] Ollama. (2024). Run large language models locally. https://ollama.com

[9] suchintikasarkar. (2023). Sentiment Analysis for Mental Health [Dataset]. Kaggle.

[10] Pedregosa, F., et al. (2011). Scikit-learn: Machine learning in Python. JMLR, 12, 2825–2830

DOI: 10.48175/IJARSCT-35143