CARMA 2026

*8th Int. Conf. on Advanced Research Methods and Analytics*
1-3 July 2026, Valencia, Spain
DOI: https://doi.org/10.4995/CARMA2026.2026.22003

# Data-Driven Crisis Detection and Support Prioritization for Student Mental Health

**Pritom Das**[1], **Afsana Islam**[2]

[1]Department of Artificial Intelligence, Ajou University, South Korea, [2]Department of Public Policy, KDI School of Public Policy and Management, South Korea.

How to cite: Das, P.; Islam, A. (2026). Data-Driven Crisis Detection and Support Prioritization for Student Mental Health. In: 8th International Conference on Advanced Research Methods and Analytics (CARMA 2026). Valencia, 1-3 July 2026. https://doi.org/10.4995/CARMA2026.2026.22003

***Abstract***

*Rising student psychological distress and chronically under-resourced counselling services create urgent demand for scalable triage tools. We present a data-driven decision-support framework that transforms mental-health survey responses into an operational prioritization system for institutional intervention planning. Using validated instruments measuring depression (PHQ-9), anxiety (GAD-7), and stress (PSS-10), the framework implements a two-stage architecture that separates crisis detection from conditional support-level estimation among non-crisis students. A policy-aware fusion mechanism integrates these outputs into a continuous intervention priority score, ranking students by support urgency under resource constraints. Evaluated on 1,977 students across 15 universities in Bangladesh, the framework is assessed using both conventional classification metrics and a proposed allocation-based evaluation that measures how effectively high-risk students are identified under limited counselling budgets. Results emphasize the importance of aligning predictive models with operational decision objectives rather than optimizing classification performance alone.*

***Keywords:*** *Crisis Detection; AI for Social Impact; Decision-support System; Resource Prioritization.*

Article in Press

## 1. Introduction

Mental health challenges among university students have emerged as a significant global public health concern. Academic pressure, financial uncertainty, social expectations, and major life transitions often converge during university years, placing students at heightened risk of psychological distress. A substantial proportion of university students experience symptoms of depression, anxiety, or stress during their academic careers (Auerbach et al., 2018). In low- and middle-income countries, these challenges are exacerbated by limited access to mental-health services (Patel et al., 2018). Therefore, universities face increasing pressure to develop scalable

This work is licensed under a Creative Commons License CC BY-NC-SA 4.0
Editorial Universitat Politècnica de València

tools for identifying, prioritizing, and supporting students who require mental-health intervention (Al Azdi et al., 2021). Poor mental health among students has been linked to reduced academic performance, higher dropout rates, impaired social functioning, and long-term socioeconomic consequences (Bruffaerts et al., 2018).

Standardized psychological instruments such as the Patient Health Questionnaire (PHQ-9), the Generalized Anxiety Disorder scale (GAD-7), and the Perceived Stress Scale (PSS-10) provide validated measures of depressive symptoms, anxiety, and perceived stress (Spitzer et al., 2006). However, these assessments typically produce static scores that require interpretation by trained professionals. Manual review and prioritization of such assessments become difficult to scale when thousands of students must be screened.

Several studies, driven by machine learning (ML) techniques, have applied methods such as decision trees, support vector machines, random forests, and neural networks to predict stress, anxiety, or depression among university students using survey or behavioral data (Shatte et al., 2019). Rahman et al. (2024) proposed convolutional neural network–support vector machine architectures to predict perceived stress and anxiety. Other work integrates feature selection or optimization to improve transparency and predictive performance (Raian et al., 2025). Specifically, precision mental health aims to use data-driven methods to enable personalized and proactive interventions (Delgadillo & Lutz, 2020), with ML models increasingly applied to psychiatric data to identify risk patterns and support early intervention (Shatte et al., 2019).

The primary contribution of this work is a decision-support architecture tailored to the operational realities of university counselling services. The framework includes (i) a two-stage triage pipeline that decomposes a multi-class problem into a clinically motivated two-step process; (ii) a policy-aware fusion mechanism that produces a continuous intervention priority score for resource-constrained allocation; and (iii) an allocation-based evaluation methodology that measures practical utility under budget constraints, complementing standard classification metrics. Unlike prior studies that emphasize predictive accuracy alone, we emphasize on structuring the decision process to align model outputs with resource-constrained intervention planning.

## 2. Methodology

### 2.1 Data

The framework is evaluated using survey data from 1,977 students across 15 universities in Bangladesh (Syeeda et al., 2024), including validated measures of depression (PHQ-9), anxiety (GAD-7), and stress (PSS-10). To reflect real-world counselling constraints, the study also introduces allocation-based evaluation, assessing how effectively the framework identifies high-risk students under limited intervention budgets.

Article in Press

### 2.2. Feature Representation

Let $x_i$ represent the feature vector for student $i$, derived from psychological assessment responses and contextual variables, and $y_i$ represents the corresponding support label. The objective of this work is to construct a decision-support mechanism that prioritizes intervention under limited counselling resources. To reflect the operational workflow of university support services, we model the problem as a two-stage decision process: 1) Crisis Detection, and 2) Support-Level Estimation for Non-Crisis Students. The final output is an intervention priority score used to rank students for counselling support. Each student record consists of psychological assessment scores derived from three validated instruments: PHQ-9 for depression severity, GAD-7 for anxiety severity, and PSS-10 for perceived stress.

Let $X_i = [\mathrm{GAD}_i, \mathrm{PSS10}_i, \mathrm{context}_i]$, where $\mathrm{context}_i$ denotes contextual variables which may include demographic or academic attributes when available. Importantly, PHQ-9 values are used for target construction rather than prediction, preventing direct label leakage.

### 2.3. Two-stage architecture

Students are categorized into three tiers: Tier-0 (low support need), Tier-1 (moderate need, active monitoring), and Tier-2 (crisis-level risk, immediate intervention). For Stage 1, Tier-2 cases form the crisis class as well as Tier-0 and Tier-1 together form the non-crisis class. The architecture is motivated by both clinical and statistical considerations. The decision to escalate a student to crisis intervention is distinct from assessing which support level a non-crisis student needs, and conflating them into a single multi-class problem comes at the cost of recall on the most critical cases. Since Tier-2 cases constitute the dataset majority, a single-stage classifier tends to collapse the intermediate class; isolating Stage 1 as a high-recall binary screening phase ensures severe cases are not missed. Stage 2 operates on a more balanced non-crisis subset.

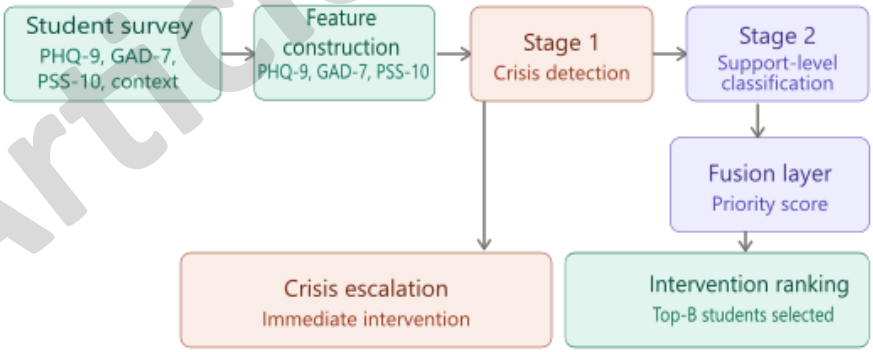



*Figure 1. System architecture*

Article in Press

### 2.4. Stage 1: Crisis Detection

The first stage identifies students who may require immediate intervention. Let $c_i \in \{0,1\}$ denotes the crisis label for student $i$, where $c_i = 1$ indicates potential crisis, and $c_i = 0$ indicates non-crisis. A classifier $f_1(x_i) \rightarrow p_i^{crisis}$ produces the probability that student belongs to the crisis category, and is defined as

$$p_i^{crisis} = \boldsymbol{P}(c_i = 1|x_i). \tag{2}$$

Students exceeding a predefined threshold $p_i^{crisis} \geq \tau_c$ are flagged for crisis escalation. This stage is designed with a high-recall objective, ensuring that potentially severe cases are not missed.

### 2.5. Stage 2: Conditional Support-Level Classification and Triage Tiers

Students not flagged as crisis cases are passed to the second stage. This stage predicts the level of support required among non-crisis students. Let $C = \{0,1,2\}$ represent the support levels, with 0, 1, and 2 corresponds to routine support, monitoring and crisis escalation, respectively. A second classifier, $f_2(x_i) \rightarrow p_{i,k}$ produces class probabilities, and is defined as

$$p_{i,k} = \boldsymbol{P}(s_i = k|x_i) \text{ for } k \in C_i, \tag{3}$$

$$\text{with } s_i = arg\ max_k p_{i,k}. \tag{4}$$

However, these probabilities are not used directly for intervention decisions. Instead, they are combined with crisis risk in the fusion stage. The final decision category is represented by a triage tier variable $t_i \in \{0,1,2\}$. If $p_i^{crisis} \geq \tau_c$, then the final triage tier assignment $C_r = 2$ . Otherwise $C_r = C_i$. Thus, students flagged as crisis cases in Stage 1 are directly assigned to Tier 2, while non-crisis students retain the support class predicted by Stage 2.

### 2.6. Fusion, Priority Scoring, and Intervention Allocation

The outputs of the two stages are integrated to generate a continuous intervention priority score. With the crisis probability from Stage 1 $p_i^{crisis}$, the probability of monitoring-level support $p_{i,1}$, and the probability of crisis-level support $p_{i,2}$, the priority score is defined as follows

$$r_i = \alpha p_i^{crisis} + \beta p_{i,2} + \gamma p_{i,1}. \tag{6}$$

where $\alpha > \beta > \gamma$ reflect the relative urgency of crisis versus monitoring cases. The score $r_i$ represents the relative priority of student $i$ for intervention. Students can then be prioritized according to $s_i$, to determine which individuals should receive counselling when resources are limited. The weights $\alpha, \beta$, and $\gamma$ are treated as policy parameters reflecting the relative urgency of crisis and non-crisis cases. In our experiments ($\alpha = 0.6, \beta = 0.3, \gamma = 0.1$), higher values

Article in Press

assigned to crisis probability ensure that high-risk students are prioritized. Alternative configurations can be chosen to adjust prioritization according to institutional policies.

Given a fixed counselling capacity $B$, the system selects the top-ranked students $\{i_1, i_2, \ldots . i_B\}$ such that $r_{i_1} \geq r_{i_2} \ldots . \geq r_{i_B}$. The ranking provided an actionable ranking that can guide practical intervention planning.

### 2.6. Model Training

Two types of classifiers are evaluated for both stages: 1) Multinomial Logistic Regression, and 2) Random Forest. These models provide probabilistic outputs, robustness to moderate dataset sizes, and interpretability relative to deep neural networks. The models are trained under two evaluation settings. First, an independent and identically distributed (IID) split was used, where training and test samples are randomly drawn from the same pool. Secondly, a grouped split was also used, where universities are separated across training and testing sets to evaluate cross-institution generalization. Figure 1 depicts the proposed system architecture.

**Table 1. End-to-end fused system performance across evaluation settings.**

| Evaluation | Feature set | Stage-1 | Stage-2 | Accuracy | Balanced Acc. | Macro-F1 | Tier-1 Recall | Tier-2 Recall |
|---|---|---|---|---|---|---|---|---|
| IID | enriched | Logistic | Multinomial | 0.663 | 0.465 | 0.465 | 0.267 | 0.818 |
| Grouped | totals | Logistic | Multinomial | 0.736 | 0.483 | 0.497 | 0.256 | 0.877 |

## 3. Experimental Results

### 3.1. Dataset composition and evaluation setup

The dataset contains 386 Tier-0 (low support), 195 Tier-1 (monitoring), and 1,396 Tier-2 (high support or crisis). For Stage-1 screening, this corresponds to 1,261 crisis and 716 non-crisis cases. The IID evaluation uses a 70–15–15 split with 1,383 training, 297 validation, and 297 test observations. To assess cross-institution robustness, a grouped evaluation separates entire universities across partitions, yielding 919 training, 565 validation, and 493 test observations. This grouped setting provides generalization because the model must operate on universities not observed during training, reflecting realistic deployment where models may transfer across institutions. The dataset is heavily skewed toward Tier-2 cases (70.6%), explaining the high Tier-2 recall and substantially lower Tier-1 recall observed throughout the experiments.

Article in Press

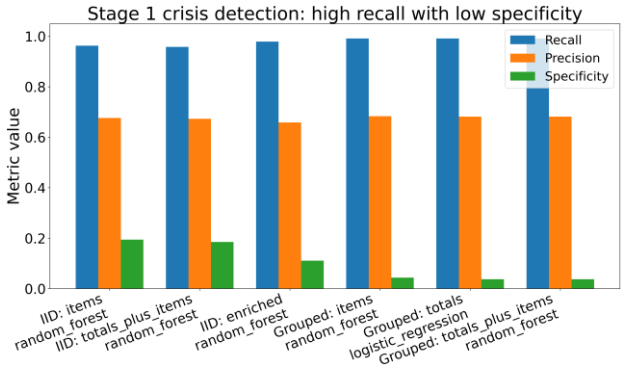



*Figure 2. Stage-1 crisis detection performance.*

### 3.2. Stage-1 crisis detection

Stage 1 screens students by predicting whether they belong to the crisis category. Figure 2 reports recall, precision, and specificity for the strongest configurations under IID and grouped evaluation. Across models, recall remains very high—above 0.95 in IID and close to 0.99 in grouped evaluation—indicating that crisis cases are rarely missed. Specificity is lower, particularly in the grouped setting, meaning some lower-tier students are flagged as crisis candidates. Consequently, Stage 1 functions as a high-recall screening filter rather than a precise classifier. This trade-off is intentional: prioritizing recall ensures potentially high-risk students proceed to Stage 2 and the fusion mechanism, where finer distinctions are made.

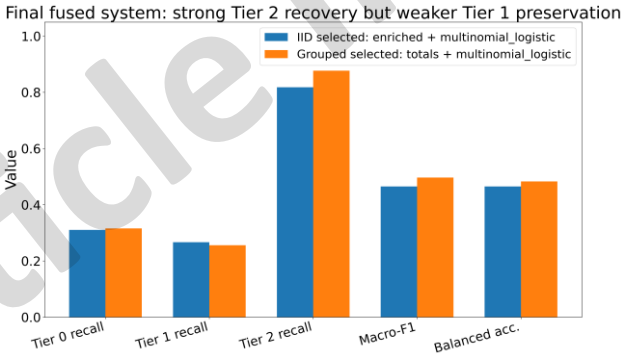



*Figure 3. End-to-end fused system performance.*

### 3.3. End-to-end fused system performance

The final system combines Stage-1 crisis probabilities with Stage-2 predictions using the policy-aware fusion rule. Table 1 summarizes the selected models. In the IID setting, the enriched feature set with logistic regression (Stage 1) and multinomial logistic regression (Stage 2) achieves accuracy 0.663, balanced accuracy 0.465, macro-F1 0.465, Tier-1 recall 0.267, and

Article in Press

Tier-2 recall 0.818. In the grouped setting, logistic regression in both stages with totals features achieves accuracy 0.736, balanced accuracy 0.483, macro-F1 0.497, Tier-1 recall 0.256, and Tier-2 recall 0.877. Figure 3 shows consistently high Tier-2 recall but lower Tier-1 recall, indicating that the fusion rule tends to push uncertain cases toward higher support levels. The system therefore operates as a risk-sensitive prioritization mechanism focused on high-risk detection rather than balanced multi-class prediction; the modest macro-F1 and low Tier-1 recall reflect difficulty in distinguishing moderate-risk (Tier-1) students from high-risk cases, with the model tending to assign uncertain cases to higher-risk categories.

### 3.4. Allocation performance under intervention constraints

Beyond classification accuracy, the framework aims to support resource allocation. Figures 4 and 5 evaluate performance under budget-constrained intervention scenarios. In the IID setting (Figure 4), the prioritization score outperforms random allocation and remains competitive with simple baselines based on anxiety or stress scores as intervention budgets increase. These baselines are intentionally chosen to reflect operational heuristics (e.g., symptom-based ranking) rather than optimized predictive models, as the goal is to evaluate decision-support utility under realistic deployment conditions. In the grouped setting (Figure 5), similar trends appear: even for unseen universities, the model identifies crisis cases more effectively than random selection. These results indicate that the framework's primary value lies in ranking students by intervention urgency rather than solely improving classification metrics.

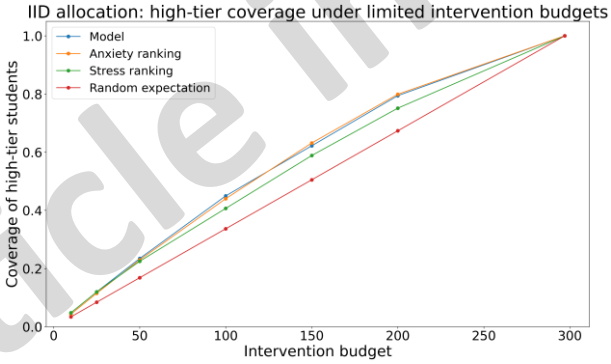



*Figure 4. Allocation performance under IID evaluation.*

Article in Press

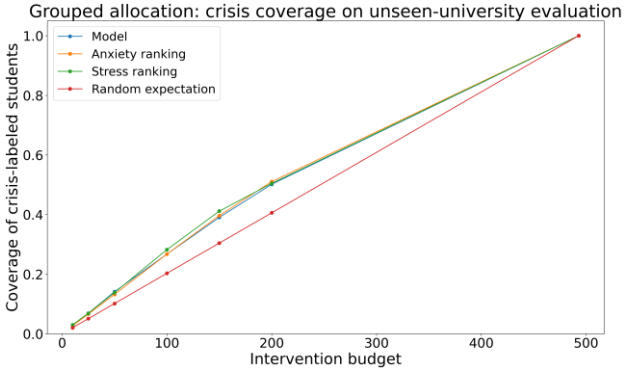



*Figure 5. Allocation performance under grouped evaluation.*

### 3.5. Cross-institution robustness

A key question for deployment is whether the model generalizes across institutions. Grouped evaluation shows that performance remains stable under this stricter setting: crisis detection stays highly sensitive and allocation performance remains competitive. However, results also indicate that the intermediate monitoring class is the hardest category to preserve. This pattern appears in both IID and grouped settings, suggesting that the difficulty arises from the structure of the decision pipeline rather than overfitting to specific universities. However, as the dataset is drawn from a single-country context, external validation across diverse institutional and cultural settings is required before broader generalization.'

## 4. Conclusion

This study proposes a two-stage decision-support framework for student mental-health screening and prioritization under human oversight. The framework separates crisis detection from conditional support-level estimation and integrates them through a policy-aware fusion rule to generate an intervention priority score. Results show three main findings. First, crisis screening achieves very high recall, ensuring most high-risk students are identified. Second, conditional classification performs moderately well, suggesting that task separation improves learnability of lower-intensity cases. Third, the fused system functions as a risk-sensitive prioritization mechanism to assist ranking under limited counselling capacity. However, the intermediate monitoring class (Tier-1) remains difficult to preserve, reflecting class imbalance and structural compression in the two-stage pipeline. As the dataset is drawn from a single-country context, broader generalization requires validation across diverse institutional and cultural settings. Future work will explore cost-sensitive learning, threshold calibration, and improved fusion strategies to enhance classification balance and robustness, while accounting for misclassification risks and fairness considerations.

Article in Press

## References

Al Azdi, Z., Islam, K., Khan, M. A., Khan, N., Ejaz, A., Khan, M. A., Warraitch, A., Jahan, I.,

& Huque, R. (2021). Effectiveness of an integrated care package for refugee mothers and children: Protocol for a cluster randomized controlled trial. *JMIR Research Protocols*, 10(5), e25047.

Auerbach, R. P., Mortier, P., Bruffaerts, R., Alonso, J., Benjet, C., Cuijpers, P., Kessler, R. C.

(2018). WHO world mental health surveys international college student project: Prevalence and distribution of mental disorders. *Journal of Abnormal Psychology, 127 (7),* 623–638.

Bruffaerts, R., Mortier, P., Kiekens, G., Auerbach, R. P., Cuijpers, P., Demyttenaere, K., … Kessler, R. C. (2018). Mental health problems in college freshmen: Prevalence and academic functioning. *Journal of Affective Disorders, 225*, 97–103.

Delgadillo, J., & Lutz, W. (2020). A development pathway towards precision mental health care. *JAMA Psychiatry, 77 (9),* 889–890.

Patel, V., Saxena, S., Lund, C., Thornicroft, G., Baingana, F., Bolton, P.,…Unützer, J. (2018). The Lancet Commission on global mental health and sustainable development. *The Lancet, 392* (10157), 1553–1598.

Raian, A., Siddique, S. A., Rajbongshi, A., Arafin, B. M. M., Islam, F., & Sheikh, R. (2025). Stress prediction among university students using swarm intelligence and explainable AI. *IEEE International Conference on Computing Applications and Systems*.

Rahman, A., Syeed, M. M., Fatema, K., Khan, R. H., Hossain, M. S., & Uddin, M. F. (2024). An interpretable hybrid CNN–SVM model for predicting stress among university students. *IEEE/ACS International Conference on Computer Systems and Applications*.

Shatte, A. B., Hutchinson, D. M., & Teague, S. J. (2019). Machine learning in mental health: A scoping review of methods and applications. *Psychological Medicine, 49* (9), 1426–1448.

Spitzer, R. L., Kroenke, K., Williams, J. B., & Löwe, B. (2006). A brief measure for assessing generalized anxiety disorder: The GAD-7. *Archives of Internal Medicine, 166* (10), 1092–1097.

Syeeda, M. M., Rahman, A., Akter, L., Fatema, K., Khan, R. H., Karim, M. R., Hossain, M. S., & Uddin, M. F. (2024). A comprehensive standardized dataset on mental health problems (MHPs) of university students. *Measurement*, 8(9).

Article in Press